A pluralistic approach to distributed cognition: tasks, mechanisms, and practices

by Alexander James Gillett

BA(Hons) Philosophy, MA Philosophy

Macquarie University Philosophy Department

Submitted on the 17/08/2017

Contents

Abstract v

Declaration of originality vi

Acknowledgements vii

General Introduction 1

Naturalism 5

Pluralism 6

Pragmatism 7

Overview of the chapters 10

1: The problems of distributed cognition 12

1.1. Introduction 12

1.2. Chaotic and Bloated 14

1.3. A limited approach to distributed cognition 20

1.4. The problem of cognition distributed in time 26

1.5. The flexible unit of analysis and the problem of methodological bloat 30

1.6. A pluralistic framework 37

1.6.1 Task-specificity (multilevel analysis) 38

1.6.3 Normative Patterned Practices 40

1.7. Summary 43

2: The navigation team and information bottlenecks 45

2.1. Introduction 45

2.2. Why I have chosen the navigation example 47

2.3. Preliminary details in the navigation team case study 50

2.4. The particular roles of the navigation team 55

2.5. The fix cycle and the propagation of representational states 61

2.6. Cognitive ecology and the blind man's stick (BMS) 65

2.7. Critiquing the notion of information Bottlenecks as a criterion for demarcating the flexible unit of analysis 72

2.8. Summary 80

3: Task-specificity and Multilevel Analysis 81
3.1. Introduction 81
3.2. Scaling up and scaling down 83
3.3. The problems of multi-level analysis 86
3.3.1. Cognitive artefacts and three kinds of enhancement 88
3.3.2. Two ways of transposing inappropriate conceptual baggage onto the wrong level of analysis 98
3.4. A task-specific definition of levels 105
3.4.1. Mechanistic levels 105
3.4.2. Core configurations, emergence, and task-taxonomies 111
3.5. Summary 122

4: The mutual manipulability criterion 124
4.1. Introduction 124
4.2. A mechanist approach to enumerating the components of distributed cognitive systems 125
4.3. Preliminary remarks on mutual manipulability and distributed cognition 132
4.4. Decomposing the navigation team 141
4.5. The problem of time 151
4.6. Summary 156

5: Normative Patterned Practices 158
5.1. Introduction 158
5.2. Two differing cultural-cognitive niches 160
5.3. The challenges of operationalising culture 170
5.4. Normative patterned practices 174
5.4.1. Learning, transformation, and norms 175
5.4.2. Cognitive practices 182
5.4.3. General positive attributes 190
5.5. Tackling the problems of time and methodological bloat 196
5.6. Summary 203

6: Re-examining the navigation case study 205
6.1. Introduction 205
6.2. The navigation case study revisited 206

6.2.1. Communication, the importance of buffering, and the division of labour 210

6.2.2. Parallel activity and the importance of cognitive practices 214

6.3. Developmental trajectories and resilience engineering through the degeneracy of cognitive practices 221

6.4. Cognitive-historical factors 232

6.5. Summary 240

General Conclusion 242

Summary 242

Further issues and future directions 249

Bibliography 253

Abstract

This thesis is a conceptual re-analysis of distributed cognition. I defend the position against numerous challenges and propose a novel pluralistic approach. Distributed cognition is a research framework used in many differing fields within the sciences and humanities. It proposes that the unit of analysis for exploring human cognition is flexible, and that cognition is distributed across both time and space in multiple ways (Hutchins 2001). But the stunning variety of putative cases in the literature raises a multitude of questions about whether this concept is being applied coherently. The heterogeneity of size and type of distributed cognitive systems also exacerbates the problem of cognitive bloat. I argue that standard simpler approaches to these problems are inadequate because they fail to account for the importance of cognition distributed in time.

As such, I propose a pluralistic approach that combines a number of naturalistic criteria from the literature: task-specificity (Davies & Michaelian 2016), the mutual manipulability criterion (Kaplan 2012), and normative patterned practices (Menary 2007a, 2016). These principles work in conjunction with each other in a consilient fashion to tackle the myriad problems facing a proponent of distributed cognition – as well as mitigating a further problem I call 'methodological bloat'. My pluralistic framework also provides a robust and useful framework that is suitable for both theoretical and practical purposes: i.e. it not only provides us with a principled means of designating what distributed cognition *is*, it also shows how exploring distributed cognition in specific case studies can lead to insights about human cognition "in the wild". I demonstrate these points with reference to the specific details of case studies – particularly Hutchins' (1995a) seminal navigation team. This allows me to show that my pluralistic approach is not only superior to simpler approaches to distributed cognition but also provides insights that are of note to methodological individualists.

Declaration of originality

I declare that this thesis, “A pluralistic approach to distributed cognition: tasks, mechanisms, and practices”, has been written by myself, the candidate, and has not been submitted for examination at any other institution. This thesis is a product of my own efforts, to the best of my knowledge; and where I have drawn on other sources, these have been cited extensively in the main text using APA guidelines, and are listed in full in the bibliography at the end of this thesis.

Alexander James Gillett

14/08/2017

Acknowledgements

I would like to thank my supervisor Richard Menary for greatly refining my critical faculties and writing over the last three years, for providing me with opportunities, and keeping me going through the inevitable hard times. I would also like to thank my secondary supervisors Albert Atkin and Richard Heersmink for their help in developing and discussing my ideas, and offering invaluable advice. I would also like to thank other staff and students of the Macquarie University Philosophy Department for helping me to clarify my thoughts on a number of areas in this project – these ideas have benefitted greatly from being critically examined. And I greatly appreciate their time. In particular, I would like to thank the members of the Macquarie 'Culture and Cognition' reading group; especially, Hoda Mostafavi, Graham Thomas, McArthur Mingon, Stephen Gadsby, Yves Saint-James Aquino, Kelly Hamilton, and William Hebblewhite.

I would also like to thank a number of staff from other departments at Macquarie University for discussing certain areas of this project which were crucial given its interdisciplinary nature; especially, Greg Downey, John Sutton, and Dave Kaplan. I attended several conferences and workshops during my candidature where I presented various ideas from my thesis; and the critical questions of many audience members helped to improve my arguments. It goes without saying that although this thesis owes a great debt to the above-mentioned people, that any mistakes are mine alone.

Lastly, I would like to thank my parents for allowing me the space to think expansively when I was growing up – and my brother and friends for tolerating this state of affairs. And I would like to thank my wife Emma Redfearn and daughter Freya for everything else; especially for making the last three years even more exciting than previously expected.

α: General Introduction

This is the task of natural science: to show that the wonderful is not incomprehensible.

Herbert A. Simon (1996, p. 1)

Understanding human cognition is a project complicated by the extreme sociocultural nature of our species. There is much debate about how best to understand these circumstances; and a number of different approaches have been proposed. Distributed cognition is a framework for exploring, what Hutchins' (1995a) evocatively calls, "cognition in the wild" and is utilised in many different fields across the sciences and humanities. It proposes that the appropriate unit of analysis for investigating human cognition is flexibly distributed across both space and time and involves supra-individual cognitive properties (Hutchins 2001). Distributed cognition offers novel insights into ecologically salient human cognition that is not apparent in other theoretical approaches. However, the position faces numerous challenges in order to be coherent. My aim in this project is to outline these problems in detail and examine the conceptual underpinnings of distributed cognition in order to formulate a robust and useful framework that is suitable for both practical and theoretical purposes. I shall argue for and defend a pluralistic approach.

A proponent of distributed cognition faces multiple interrelated challenges. Firstly, there is a dizzying array of case studies in the literature: studies of distributed cognition range from the very small (e.g. long-term married couples, and tool-use, etc.) to the incredibly large (e.g. institutions, and radically collaborative scientific experiments, etc.). This raises questions about whether we can use the concept

coherently across such a wide range of cases; and whether we can compare systems of such differing sizes in an appropriate fashion. Comparing systems of radically different configurations and sizes raises a series of conceptual problems that threaten to invalidate distributed cognition as a coherent project. This necessitates a having a means of negotiating multiple levels of analysis in a principled fashion. In turn, this exacerbates further issues, such as the problem of cognitive bloat: viz. how do we demarcate the boundaries of cognitive systems in a principled fashion? If we cannot provide principled criteria then this blurs the distinction between a cognitive system and its causal milieu – rendering cognitive science intractable and entailing a form of uncontrollable panpsychism. This question becomes more complicated given the differing types of systems discussed in the literature because it increases the constraints on the hunt for demarcation principles.

I will argue that simpler limited approaches to these dilemmas are insufficient for two main reasons: firstly, for practical reasons it is ill-advised to use an exclusionary definition that would dismiss much important empirical research. Instead, I will argue for a comprehensive position. A second and more problematic issue is that simpler approaches tend to overlook the importance of how cognition is distributed in time. Although ignoring these *diachronic* factors leads to a simpler position; I demonstrate that this solely *synchronic* view is actually untenable because it sacrifices important explanatory factors that are crucial for properly understanding what distributed cognition is on its own terms. As such, a more complex position is required. However, the prima facie range of diverse theoretical posits necessitated in the flexible unit of analysis in turn entails a further problem related to the feasibility of such a framework. I label this the problem of methodological bloat.

To meet these various challenges, I propose a pluralistic approach to distributed cognition that relies on three key principles drawn from naturalistic approaches in the literature: *task-specificity* (Davies & Michaelian 2016), *the mutual manipulability criterion* (Kaplan 2012), and *normative patterned*

practices (Menary 2007a, 2016). I shall show that although each of these principles cannot tackle all of these problems on their own; that a consilient combination allows us to provide a robust definition of distributed cognition that meets both naturalistic and pragmatic credentials. I will now briefly summarise how the pluralistic approach operates to make the central claims of distributed cognition – the flexible unit of analysis, and the various ways in which cognition is distributed across time and space – principled and tractable.

An investigation of a putative distributed cognitive system involves examining multiple levels of analysis; and these systems also vary greatly in terms of their size and complexity. These various levels must be elucidated on their own terms. Approaching this issue from an agent-centric perspective leads to a number of errors in which inappropriate conceptual baggage is transposed onto differing levels of analysis (which I will discuss in detail in chapter three). *Task-specificity* provides us with a criterion for establishing the centre of the flexible unit of analysis that avoids these issues by focusing on a specific cognitive task that the system is engaged in (e.g. memory, perception, problem-solving, etc.). But it does not, by itself, provide us with an adequate boundary criterion. This is a major issue for externalist positions that challenge the notion that cognitive processes are purely within the skull of isolated individuals (Adams & Aizawa 2001). The *mutual manipulability criterion* is a neutral means of tackling the problem of cognitive bloat. It uses two theoretical interventions on the target system – top-down and bottom-up – in order to differentiate between genuine components of the system (coordination dynamics) and the mere causal background (Kaplan 2012). I discuss this in chapter four.

But although the mutual manipulability criterion enables us to identify and investigate the composition of the coordination dynamics of distributed cognitive systems, it only does so over a limited timeframe: a snapshot of the system in what I shall refer to as the synchronic time scale. This fails to take into account the importance of diachronic timescales – both developmental and historical – for

understanding the cognitive behaviour of the system. The various ways in which cognition can be distributed over time are crucial for understanding how cognition is distributed in general. To omit them is to fail to provide a properly robust definition of distributed cognition. Diachronic factors can be made tractable by using normative patterned practices – these are sets of cultural practices which both shape the coordination dynamics of distributed cognitive systems and are partially constitutive of this activity (Menary 2007a; 2016). I discuss these in chapter five.

Normative patterned practices are acquired through developmental trajectories and transmitted intergenerationally; and both of these processes have a significant and multifaceted impact on cognitive behaviour (Menary 2013b, 2014, 2015; Roepstorff et al 2010). As such, normative patterned practices are involved across all the spatial and temporal scales which are entailed by a distributed cognition investigation. By operating in tandem with the other two criteria this reduces the prima facie excess complexity of my pluralistic approach, thus nullifying the problem of methodological bloat.

To reiterate: the central claim of this thesis is that these three criteria operate together to provide us with a principled and tractable way of demarcating a flexible unit of analysis and articulating how cognition is distributed in both time and space over multiple levels of analysis. To anchor this discussion, I shall repeatedly demonstrate how each of these features of the pluralistic framework relates back to a central case study: Hutchins' (1995a, 2006, 2010b) seminal example of the navigation team on board a US navy aircraft carrier in the late 1980s. This will also allow me to provide a concrete demonstration of how my pluralistic approach to distributed cognition is superior to the standard simpler approaches in the literature. But it also provides an opportunity to show how my pluralistic approach offers important insights into real world cognitive activity that would not be apparent to methodological individualism. Having practical consequences and paying attention to the ethnographic

details of case studies are important aspects of my approach. I will now briefly outline these desiderata, alongside two other important framing issues, before providing an overview of the chapters.

Naturalism

Traditional approaches to cognitive science can be described as focusing on systems that perform cognitive tasks. It is common within the literature on distributed cognition to approach the topic from a naturalistic stance that examines clearly specified particular functions, e.g. memory, decision-making, reasoning but also emotion, perception and motor-skills (see Hutchins 2001, p. 2068; Kaplan 2012, p. 546; Michaelian & Sutton 2013, p. 10; Sutton 2006, 2015, p. 430; Sutton et al 2010; Theiner et al 2010, p. 378-379; Theiner & O'Connor 2010, pp. 81-84; Wilson 2001, pp. 266-267, 2004, pp. 288-289). By naturalism I mean simply that we should pay attention to our best current available sciences (Ladyman & Ross 2007); and that this need not involve a scientistic stance. I follow Stanford in understanding naturalism as the claim "that there is only a single, interconnected project of inquiry into the world and our place within it and [the] insistence that this single integrated inquiry is responsible for simultaneously making good sense out of all of the reliable evidence we have to use at any given time" (2016, p. 92). As such, this investigation of distributed cognition is premised on naturalistic assumptions held by many other proponents within the literature. The important feature this has for my approach is that a conceptual analysis of distributed cognition must pay attention to the particular ethnographic details of case studies. This also has theoretical consequences (see below).

A critic might respond that distributed cognition and naturalism are antithetical given that the consensus view within the cognitive science community is that cognitive properties are limited to the nervous systems – in particular brains – of individuals (see Huebner 2014, p. 13; Wilson 2005, p. 227 for similar concerns). But I think that this is a misnomer for several reasons. Firstly, it fails to take into

account the wide range of active research in many different scientific fields that examines distributed cognition (as I will outline in chapter one). Allen (2017) has recently argued that ignoring these projects is itself anti-naturalistic. Secondly, Kaplan (2012) argues that these case studies, in conjunction with theoretical debates (see Menary 2010d and Robbins & Aydede 2009 for an overview), prevent traditionalists from using the definition of cognition as a skull-bound phenomenon since this simply begs the question (Clark 2008; Clark & Chalmers 1998).

Pluralism

Given the complexities of the problems facing distributed cognition I propose a pluralistic approach to the method of tackling the question of how cognition can be distributed. The problematic terrain facing a proponent of distributed cognition is incredibly complex and presents a serious set of challenges that must *all* be negotiated in order to make the position viable. To summarise, the major issues are: the miscellaneous nature of case studies in the literature; cognitive bloat; cognition distributed in time; and methodological bloat. As I will show in discussing these problems, there is no straightforward or simple singular solution that would resolve all of them.

This suggests that in considering the unit of analysis for distributed cognition that we should not look for a single top-down silver bullet, and should instead proceed in a piecemeal fashion through a series of empirically tractable principles that pick out certain features of real-world human cognitive systems. For instance, task-specificity and mutual manipulability guide research to focus on the particular relations of the components of a system and how they are integrated and related. However, although these principles lead to interesting insights in regards to the coordination dynamics of distributed cognitive systems they do not help us to understand how cognition is distributed in time, nor how they are related to the wider cultural-cognitive niche in which they are embedded.

A supplementary account of normative patterned practices satisfies this issue in a consilient manner and provides us with explanatory tractability with regards to how these coordination dynamics are sculpted by both shallow and deep timescales. As such, rather than pluralism heralding the bane of methodological bloat, it is instead a boon which properly articulates the whole system under investigation in a more profound way.

Pragmatism

Many proponents of distributed cognition emphasise its pragmatic aspects and its focus on practical concerns. For instance, in the philosophy of science literature Giere and others argue that distributed cognition can help bridge and overcome the reciprocal animosity between philosophy and sociology of science (see 2002a, pp. 638, 641; 2002b, pp. 285, 295-296; 2006, p. 114; 2007, p. 319; Giere & Moffattt 2003, pp. 301-302. Also see Brown 2011, p. 17; Cheon 2014, pp. 23-24; Magnus 2007, p. 297; Magnus & McClamrock 2015, p. 1114; Nersessian 2005, pp. 17-19; 2006, p. 702; Vaesen 2011, p. 381; cf. Toon 2014b). In another case Huebner has suggested that exploring distributed cognition is an *open empirical question* that requires adopting what he calls the "pragmatist gambit" in which we accept a moderate reconceptualization of the foundations of the cognitive sciences that moves away from methodological individualism (2014, pp. 164-165).

I think this state of affairs makes it legitimate to judge the use of distributed cognition on whether it has made these practical contributions; and I make it a central desideratum of my own view that it must be able to make definitive statements in these regards. *But the notion of practical concerns needs to be made more precise.*

There are several interrelated points here that need to be distinguished[1]. The approach I am pursuing here – based upon naturalistic credentials – follows the advice for modesty through paying attention to the case studies of distributed cognition as well as the relevant empirical studies and theoretical materials for studying systems of these kinds. There are three key factors for paying attention to other disciplines in this formulation of distributed cognition: [1] naturalistic principles dictate paying attention to the best current science relevant to the topic (Ladyman & Ross 2007; Stanford 2016); [2] distributed cognition is, by its very nature, interdisciplinary since it cuts across multiple fields in design studies, sociology, history, anthropology, philosophy, and cognitive science (Hollan et al 2000; Nersessian 2005); and lastly [3], when engaging in an interdisciplinary topic one should aim for conceptual clarity given the confusions that can arise from conceptual language barriers (Downey 2014b; Galison 1997).

As Huebner (2014) has argued, this attention to ethnographic details is not superfluous but is philosophically important. Because the very act of paying attention to the practical consequences of taking a distributed cognition stance towards a particular case study actually entails a better understanding of *what distributed cognition actually is* than if we limited ourselves to a more abstract discussion. I shall provide a performative demonstrate of this through a detailed discussion of Hutchins' (1995a) navigation case study.

In addition to specifying the need to pay attention to the particular details of case studies is the further methodological imperative that the application of distributed cognition *should* have some practical consequences for our investigation of the phenomena in question. Huebner (2014) has argued that postulations of distributed cognition cannot be explanatorily superfluous. The proponent of distributed cognition must demonstrate the practical consequences of adopting this approach. There

[1] The following argument in this section is Peircean in spirit (see Atkin 2016, ch2): my thanks to conversations with Richard Menary and Albert Atkin for help with clarifying the following distinctions.

are two different main ways of understanding practical consequences: *firstly*, as re-conceptualisation that changes both how we research and our methods of analysis; and *secondly*, as the imperative to make claims that are testable. My aim in this thesis is to provide a framework that mostly focuses on the former of these: the principles I outline in chapters three, four, and five provide a systematic framework for approaching both theoretical problems related to distributed cognition, and also investigating particular case studies in a novel fashion that generates new insights. But a secondary aim of my project is also – following Kaplan (2012) – to attempt to bring questions about distributed cognition more into the empirical domain (as I discuss in more detail in chapter four).

The consequence is that in applying the pluralistic framework for exploring how cognition is distributed across various spatiotemporal scales with a flexible unit of analysis, the major principles – task-specificity, mutual manipulability, and normative patterned practices – must aim to clarify both [1] what it is that we mean in claiming that a particular phenomenon is distributed cognition, and [2] why this claim is important for improving our understanding of the particular phenomena in question (as opposed to adopting some other theoretical stance; as well as being testable). By using the navigation case study as a central fulcrum, I will demonstrate each of these claims in a concrete fashion with particular regard to the three principles of my pluralistic framework. Viz. how paying attention to the specifics: [i] draws out important practical consequences that are often overlooked by proponents of simpler formulations of distributed cognition and methodological individualism; and furthermore [ii], how it allows us to define both Hutchins' (2001) key principles of distributed cognition – a flexible unit of analysis and supra-personal cognitive processes – more clearly.

I now provide a brief summary of the chapters.

Overview of the chapters

Chapter one begins by outlining the multiple problems facing a proponent of distributed cognition: the heterogeneous diversity of cases in the literature; the related problems of size scaling and cognitive bloat; and the problem of cognition distributed in time. The combination of these issues prevents us from putting forward a simple solution and the subsequent theoretical complexity this entails creates the further problem of methodological bloat. The chapter concludes with a précis of how my pluralistic framework meets these challenges.

Chapter two presents the most famous case study of distributed cognition – Hutchins' (1995a, 2006, 2010b) naval navigation team. I shall argue that the standard limited approaches to this case study – based on information bandwidth criterions – are insufficient for properly demarcating it as a genuine case of distributed cognition. Additionally, the navigation team example acts as a lynchpin in the thesis through which I can demonstrate how each of the principles of my pluralistic approach tackle not only the philosophical problems outlined in chapter one but also provides practical insights into the case study itself (thus, facilitating the pragmatic stance I have advocated). In turn, this shows how my position is superior to both limited approaches to distributed cognition but also raises points that are of interest to methodological individualists.

Having outlined the problem space and the shortcomings of other approaches in the previous chapters, *chapter three* begins the positive defence of my pluralistic approach. Building on work by Davies and Michaelian (2016), I shall argue that task-specificity is an alternative method for fixing the flexible unit of analysis across multiple spatial scales in a discreet and well-defined manner. Additionally, task-specificity avoids a number of conceptual problems that effect agent-centric approaches. However, it still leaves the problem of cognitive bloat. This problem is tackled in *chapter four* where I outline how a mechanistic approach to the composition of cognitive systems, using the mutual manipulability

criterion (Kaplan 2012), provides conditions by which to clearly distinguish between genuine components of a system and mere causal background conditions.

Although the mutual manipulability criterion operates well over shorter synchronic time scales it fails over longer diachronic timescales. To nullify both this as well as the problem of methodological bloat I introduce the notion of normative patterned practices (Menary 2007a, 2016) in *chapter five*. Normative patterned practices are the sets of cultural practices that govern the coordination dynamics of agents and artefacts in distributed cognitive systems. They are acquired and mastered by agents over developmental trajectories (resulting in the transformation of the neurocognitive profile of the agent and the wider system), and they are accumulatively transmitted and refined over historical time scales (altering the task space). As such normative patterned practices bind together the wide range of spatial and temporal scales that distributed cognition investigates into a manageable package – thus negating both the problems.

Chapter six concludes the project by providing a detailed re-examination of the navigation team case study and showing how the pluralistic framework operates in tandem. I pick out key examples from the synchronic and two diachronic time scales and how this not only deepens our understanding of the target cognitive behaviour but also draws out important theoretical insights.

1: The problems of distributed cognition

The very idea of a suprapersonal cognitive system stirs a deep sense of outrage.

Mary Douglas (1986, p. x)

There is a real sense in which the study of distributed cognition remains in its infancy […] In part, I believe that this is because the nature of distributed cognition is poorly understood, or at least more poorly understood than has commonly been acknowledged.

Bryce Huebner (2014, p. vii-viii)

1.1. Introduction

Provisionally, we can define distributed cognition as the claim that cognition is not solely inside the head of individuals – the intricacies of how to understand this are the topic of this first chapter; in which I outline the central interrelated problems facing a proponent of distributed cognition. There is a sheer abundance of heterogeneity in the purported cases of distributed cognition in the literature, which varies wildly in both type and size. Such a state of affairs raises concerns about whether distributed cognition is a coherent notion. Additionally, there are a range of problems concerning the differing scales of size involved in these cases and whether they can appreciably be understood using the same concept in any meaningful way. Lastly, and perhaps most

problematically, the sheer variety of cases exacerbates the problem of cognitive bloat because it complicates the search for principled boundary conditions.

Confronted with these multifaceted problems one solution would be to offer a strict and limited notion of what distributed cognition is. However, I shall argue that this is unpalatable for two reasons: firstly, from a practical point of view, distributed cognition is a widely used research paradigm across many scientific fields (e.g. design studies, cognitive ethnography, human-computer interface research, social and developmental psychology, etc.). A strict and limited notion in terms of scope fails to account for this variety, inadvertently dismisses much useful empirical work, and is therefore likely to be ignored. Secondly, many previous limited accounts of distributed cognition tend to focus solely on what I shall refer to as the synchronic scale – activity as it takes place over a short period of time. Such an approach ignores the importance of cognition distributed in time. This omission undermines the coherence of a solely synchronic approach since it leaves itself insufficient in regards to accounting for two important diachronic features – developmental trajectories and cognitive historical factors – which are crucial for understanding how synchronically distributed cognition is possible. I shall show that these diachronic factors are often crucial for understanding the synchronic coordination dynamics of distributed cognitive systems. I refer to this as the problem of cognition distributed in time.

Given the problems with a limited approach, I suggest a pluralistic method that is both applicable to this wide range of phenomena and also provides conceptual clarity for how distributed cognitive systems can be demarcated in a principled manner. My account builds from Hutchins' (2001) claim that a distributed cognition analysis begins with a flexible unit of analysis. However, I argue that Hutchins' original formulation creates an additional problem which I shall refer to as "methodological bloat" (that the diverse range of theoretical and empirical issues considered under

the ambit of distributed cognition are so broad that it defies feasible application). After outlining these problems, I conclude the chapter by providing a précis of my pluralistic framework – composed of the core principles of task-specificity, the mutual manipulability criterion, and normative patterned practices – and how it tackles these problems. I then go on to argue for these in detail in the remainder of the thesis.

The chapter is structured as follows: in the next section (1.2) I briefly survey how distributed cognition has been used in a startling number of different ways across many different fields. This relates to problems about extreme diversity of type and size that entails questioning the credibility and coherency of what distributed cognition is. In particular, this heterogeneity exacerbates the problem of cognitive bloat. In section 1.3 I explore the simpler limited option to resolving these problems and in section 1.4 introduce the problem of cognition distributed in time, which renders this approach untenable. In the remainder of the chapter I begin to outline the grounds for my pluralistic framework and how this tackles the problems of distributed cognition. Section 1.5 provides the basis by examining Hutchins' (2001) two principles of distributed cognition but shows how this entails a further problem: methodological bloat. The final section (1.6) sketches a précis of my pluralistic framework.

1.2. Chaotic and Bloated

The aim of this investigation is to make sense of distributed cognition. I agree with Huebner (2014) that distributed cognition is more poorly understood than is commonly acknowledged. An examination of the literature across a wide range of different fields testifies to this confusion. My point here in providing an extensive overview of the literature is twofold. Firstly, it shows that there is a broad discussion of distributed cognition within the literature, demonstrating that a

philosophical analysis of this problem is worthwhile and important. Secondly, the vivid diversity in these cases raises a number of concerns about how the idea is being applied. Viz. the miscellaneous nature of this assortment threatens to render attributions of distributed cognition meaningless or confused. As Heersmink has rightly observed: it is not clear what the principled criteria are for designating these cases as distributed cognition (2016, pp. 523-524). This should be of concern to philosophers.

Distributed cognition is used in an incredibly eclectic range of research fields ranging from the analysis of medical practices (Cole & Engeström 1993; Perry 2003) – such as how blood glucose meters are used (Furniss et al 2015), how heart surgeons work together in cardiac operations (Hazelhurst et al 2007), and in assessing the activity and communication within a radiology department (Rogers & Ellis 1994) – to how flight teams work together and with their equipment on commercial airliners (Hutchins 1995b; Norman 1993; Zhang 1997a). Distributed cognition has also been taken up as a research approach for fields such as educational studies (Cole & Engeström 1993; Moore & Rocklin 1998; Xu & Clarke 2012); studies of problem-solving (Theiner et al 2010; Theiner & O'Connor 2010); ecologically salient teamwork, such as emergency service dispatch centres (Blandford & Furniss 2005), café staff (Kirsh 2006), CSI teams (Baber 2010; Baber et al 2006), and sports teams (Williamson & Cox 2014); Computer Supported Cooperative Work and Human Computer Interface design (Halverson 2002; Hollan et al 2000; Nilsson et al 2012; Perry 2003; Rogers & Ellis 1994); research into external and visual representations (Heersmink 2013; Hutchins 2005; Liu et al 2008; Norman 1991; Zhang 1997b; Zhang & Norman 1994, 1995); various kinds of navigational practices (Hutchins 1995a, 2005, 2006, 2008, 2010b); and memory research in ecologically salient settings (Barnier et al 2008; Bietti & Sutton 2015; Harris et al 2011; Sutton et al 2010).

In the arts, distributed cognition has also been used as a framework for understanding the practice of drafting and revising in creative and academic writing (Frieman 2015; Menary 2007b), pedagogy in dance (Kirsh 2010), and how groups of Shakespearean actors could put on large numbers of plays without an excessive memory load (Sutton 2010; Tribble 2005). Kendal and colleagues have recently discussed distributed cognition in relation to niche construction as a perspective towards understanding human development (Flynn et al 2013; Kendal 2011). Lastly, an examination of the philosophy of science literature demonstrates the variety and incredibly diverse range of sizes of the systems that have been designated as distributed cognition. Ranging from the small: e.g. mathematical formalisms and scientific notation such as chemical formulae (Giere & Moffatt 2003); diagrams (Giere 2002b, 2006); instruments such as microscopes (Toon 2014a) and lab equipment (Heersmink 2016); and theoretical and physical models (Charbonneau 2013; Giere 2006; Knuuttila 2011). To the medium: e.g. team work enterprises such as the use and interpretation of fMRI data (Alač & Hutchins 2004); interdisciplinary fieldwork (Giere & Moffatt 2003); collaboration in chemistry laboratories (Becvar et al 2005, 2008); collaborative model-based reasoning in biomedical laboratories (Nersessian 2005, 2006, 2009; Nersessian et al 2004); and double-blind medical trials (Magnus 2007). To even more esoteric units of analysis: e.g. high energy physics experiments such as the Large Hadron Collider at CERN (Giere 2002a, 2006; Huebner 2014; Knorr Cetina 1999); and the Hubble Telescope (Giere 2006, 2007, 2012).[2]

[2] It is noteworthy that distributed cognition has also been used to look at a wide range of species. For instance, Huebner (2014) claims that honeybees can be thought of as minimally collectively minded in their foraging and hive relocation activities, but that the construction of nests by termites should not be. Theiner (2014) discusses how fire ants make rafts to survive floods and Poirier and Chicoisne (2006) discuss wolf-packs as examples of distributed cognition. Johnson has argued that distributed cognition has significant advantages for studying other animals, such as dolphins (2015) and apes (2001). Japyassú and Laland have recently argued that spider cognition is distributed by web threads (2017).

However, some of the discussions of distributed cognition in the biology literature would present additional unnecessary confusions if they were included here. For instance, O'Donnell and colleagues (2015) use distributed cognition to predict that eusocial insects will have smaller mushroom body structures in their brains due to a reduction in cognitive load. But a study of two species of ants by Kamhi and colleagues (2016)

Adams & Aizawa raise the plausible worry that the cases listed above are "an unscientific motley" because it is both too loose and over-inclusive (2001, p. 62). Sutton has responded that although he is "no enemy of motley"; the proponent of distributed cognition cannot merely offer endless examples and has to try and offer some systematic approach to this state of affairs (2006, pp. 235-236). Over and above this call to arms for some kind of organisation, the sheer miscellaneous diversity of these cases raises a number of concerns which are worth enumerating in detail.

The first point of note is the diversity in type and size of these purported cases, which range from simple dyadic relations (like those of married couples or a biochemist manipulating a mechanical model) to gigantic cases (such as the Large Hadron Collider at CERN or an entire medical department). This has raised doubts about whether distributed cognition is being used in a coherent manner or whether there are competing and contradictory formulations of the position (e.g. see Moore & Rocklin 1998). I shall return to this issue shortly when I discuss limited solutions to this problem, but a more immediate concern is how these differences are often overlooked by some proponents of distributed cognition.

Putative distributed cognitive systems have been identified as ranging from one agent interacting with a single tool (e.g. a biochemist using a mechanical model to explore molecular structure); to teams of several agents (e.g. sports teams or a navigation team on board an aircraft carrier); and even up to what has been referred to as "radical collaboration" in which hundreds of agents work on collective projects (Kukla 2012). For example, the recent scientific paper produced

---

found that they actually had slightly larger mushroom body structures in their brains. As such, I shall limit myself to only analysing the stunning array of examples in human cognition – which is already a difficult enough task without taking on these other cases. Additionally, a key factor in my account will be an emphasis on cultural practices – this has obvious limitations in being applied to animals. A notable and interesting potential exception is Keil's (2015) recent study of sheepdogs and shepherd interactions as a distributed cognitive system.

by the scientific teams working on just one detector of the Large Hadron Collider (ATLAS) that accompanied the discovery of the Higgs Boson has over two and a half thousand co-authors and ran to nine pages just naming them all (ATLAS Collaboration 2012). Such size discrepancies cannot be simply overlooked as Giere does when he "suggests" that the "…navigation aboard a traditional US Navy ship, presented in *Cognition in the Wild* [Hutchins 1995a], provides a prototype for constructing a cognitive account of HEP [High Energy Physics] experiments" (2002a, p. 639). Here Giere has compared two putative distributed cognitive systems of massively differing size: the first is composed of four individuals and a similar number of cognitive artefacts[3], and this is then compared to a second system which is an institution with well over two thousand individuals and numerous tools and devices.

One argument against such an inference can be drawn from Maynard Smith's (1968) observation that one cannot simply 'scale up' an organism's size for multiple reasons[4]. Firstly, as size increases the body of the organism will become subjected to different arrangements of forces (e.g. at very small scales the Van der Waals force is stronger than gravity) and this will have an impact on how physiological processes operate. Secondly, as size increases one has to take into account that mass and surface-area do not alter at the same rate and this can cause problems for homeostatic processes. Analogously, it is highly likely that similar alterations in size scaling will influence the types of group-level socio-cognitive processes that are present at distinct levels of analysis. As such, we should expect there to be differences in the types of cognitive properties that are present in differently sized and organised cognitive systems. To be clear, the problem is one of scale – indeed, reviews of studies on collaborative group work have showed that there is no straightforward relation between size and output (Collins & Guetzkow 1964; Forsyth 2009; Steiner 1972). But there

[3] I define these in the next section.
[4] The terminology "scale up" is taken from Davies and Michaelian (2016). I explain their work in greater detail later in this chapter and in chapter three.

are also issues with inappropriately treating cognitive properties at certain levels of analysis in terms of another level – and misattributing properties by both "scaling up" and "scaling down" (Davies & Michaelian 2016; also see Drayson 2012). Hutchins (1995a, 2008) refers to this as the problem "over-attribution" – e.g. from the agent onto the wider system or vice versa. In chapter three, building on work by Davies and Michaelian (2016), I outline a task-specific approach to tackling these problems.

To reiterate, the diversity of putative distributed cognitive systems raises questions about the legitimacy of the term and about how to understand the relations or categories of these differently sized systems. These issues are also related to, and worsen, another dilemma for which task-specificity does not help: viz. the problem of cognitive bloat. Cognitive bloat is a slippery slope argument that arises once one challenges the notion that cognition is a process that takes place solely inside the head of a cognitive agent (Rowlands 2009, p. 2). If we do not accept the skin and skull of the agent as a legitimate boundary, then what is the boundary of the cognitive system and how can this be decided in a principled fashion? Adams and Aizawa were the first to note that once one goes beyond the skull then there is a very real danger of the cognitive system expanding outwards in an uncontrollable fashion (2001, p. 57). As Rowlands notes, the issue is that once one allows in some external processes it becomes difficult to stop (2009, p. 6). Without a principled means of demarcating the boundaries of a cognitive system, the proponent of distributed cognition is in danger of including both genuine components and merely causally connected processes and features into hideously bloated systems (Kaplan 2012, p. 547). Not only is this philosophically undesirable but it also smears away important explanatory divisions and undermines our understanding of the systems under investigation (Eliasmith 2009, p. 150).

Cognitive bloat is a general problem because if one does not have a clear notion of what distributed cognition is, this can lead to the idea being applied in an unprincipled manner with no real grounds for being able to explain what *is* and *is not* a distributed cognitive system. However, the problem of cognitive bloat is exacerbated by the heterogeneity of types and sizes in the literature since this state of affairs renders the search for principled boundary conditions more fraught but also arguably provides the critic with an existence proof of the problem of cognitive bloat. I return to the problem of cognitive bloat in chapters two and four and argue, following Kaplan (2012), that we can adopt the mutual manipulability criterion from biology and neuroscience as a means of solving the boundary criterion issue in the synchronic plane in a principled manner (but not the issues size scaling or heterogeneity).

In the next section I consider a potential solution to these intertwined problems that aims to limit distributed cognition to a small subset of activity. However, I shall show this position to be wanting for both practical reasons and because it is unable to handle the problem of cognition distributed in time. The failure of this simple approach offers a motivation to make sense of a more general notion of distributed cognition that is more comprehensive insofar that it can account for the great diversity outlined above whilst also simultaneously avoiding becoming unprincipled and bloated.

1.3. A limited approach to distributed cognition

Many definitions of distributed cognition have been put forward in the philosophical literature. For instance, Cheon (2014) describes a distributed cognitive system as an information processing system; Magnus (2007) in terms of the shared task in which the system is engaged in; Poirier and Chicoisne (2006) as an emergent property; and Sutton and colleagues (2010) in terms of a set of dimensions. A

common theme amongst these positions is that distributed cognition is often considered to be *solely* a claim about how cognition can be extended or distributed beyond the skull. This often involves placing it in a hierarchy with extended cognition. Carter and colleagues define distributed cognition as the most extreme form of externalism in which "cognitive processing may not just be extended beyond the agent's head or organism [as in extended cognition] but even distributed amongst several individuals along with their epistemic artifacts" (2014, p. 72).

In addition to this hierarchical relation to extended cognition, distributed cognition is also sometimes compared to what is called "collective cognition"[5]. Giere, who popularised the use of distributed cognition in philosophy of science (see Brown 2011; Cheon 2014; Magnus 2007; cf. Heersmink 2016; Vaesen 2011), defines collective cognition as being the collective cognitive activity of multiple agents without cognitive artefacts: "'collective cognition', which I would characterize as cognition distributed among persons only, without bringing in artifacts" (2006, p. 111; also see 2002b, p. 287; 2006, p. 99; 2007, p. 314).

For our immediate purposes, we can understand a cognitive artefact as a created feature of an agent's environment that contributes to the completion of a cognitive goal by 'improving' the agent's performance (see Heersmink 2013 and Hutchins 1999 for overviews)[6]. The key point here is that distributed cognition is introduced in relation to collective cognition – thus defining it as a special case of collective cognition which is not just the collective actions of a group of agents but one which also includes specialised epistemic resources as genuine components of the system[7]. A

[5] To be clear, the following is merely a description of how others have articulated the position in what I am referring to as the limited view. This is not the position that I hold; I outline my view in greater detail in the last two sections of this chapter before elaborating this in the rest of the thesis.

[6] I use the terms epistemic tool, mental tool, thinking tool, cognitive artefact, (etc.) interchangeably. I return to discussing these in more detail in chapter two.

[7] It must be noted that there is a conflict here in how Giere presents the relationship of distributed cognition and collective cognition. In many places he seems to suggest that collective cognition is a lesser term (2002b, p. 287; 2006, pp. 99, 111; 2007, p. 314) and that distributed cognition is a more specialised term that "goes

second key feature identified by Giere, but also recognised by others (e.g. Huebner 2014), is that distributed cognitive systems, such as a group of agents and cognitive artefacts, are involved in the completion of a shared task that could not practically be completed by either an individual acting alone or a group that was unaided by cognitive artefacts (2002a, p. 641; 2002b, p. 287; 2006, p. 98; 2007, p. 314; also see Magnus 2007, p. 300).

Giere importantly states that he does not think that distributed cognition "…can usefully be defined by strictly necessary and sufficient conditions" (2002a, p. 641). But I think there are good reasons for taking these two factors to be *his* attempted set of criteria for what distributed cognition is. Although one might grant that distributed cognition cannot be defined with necessary and sufficient conditions, one can still reasonably request by what criteria it is being defined (see Heersmink 2016, pp. 523-524 for similar points[8]). I think the repetition of these two key claims in all of Giere's discussions of distributed cognition – [1] configurations of agents and cognitive artefacts and [2] the nature of the task being beyond the capacities of the unaided individual – justifies the notion that these exemplify what *he* thinks distributed cognition is.

As one would expect, these points are not universally accepted amongst philosophers. For instance, some would include systems with shared tasks that can be completed without aides, and others would count collective cognition as a form of distributed cognition: e.g. Sutton and

---

beyond" this. This is how I have presented it in the main text. But in other places this relationship is reversed and it seems as if collective cognition is a special case of distributed cognition (2002a, p. 641). This relationship is complicated even further in later works where he uses them in conjunction as "collective distributed cognition" which he distinguishes from "individual distributed cognition" (2012, p. 200). The latter term refers to one individual and cognitive artefact(s); whereas the latter refers to a group of agents and cognitive artefacts. But this distinction is confusing because it also includes a distinction between individual or collective cognition and "fully" distributed cognition (p. 199). Since Giere is inconsistent here I will choose the description he uses most often – which I have presented in the main text.

[8] Even though Heersmink thinks that distributed cognition cannot be usefully defined using necessary and sufficient conditions, it is important to note that he offers a set of dimensional criteria by which to judge situated cognitive systems in terms of degrees of integration (2016, p. 534; 2015. Also see Sutton et al 2010 for similar views about dimensions of integration).

colleagues (2010) on long-term married couples and distributed remembering. However, this is a good starting point for articulating a restricted notion of what distributed cognition is (a position I will ultimately reject). The key here is that distributed cognition is placed in a hierarchy in which the system under investigation has to be of a certain size and must contain both cognitive agents and artefacts engaged in a shared task.

In reference to the list of case studies in section 1.2 it would seem that this tacitly rejects a number of examples that are too small or which can be categorised as only collective cognition. Furthermore, the clause relating to otherwise impossible shared cognitive tasks also questions the legitimacy of a number of other cases. In particular, I suggest that this potentially threatens a number of the larger case studies in regards to what extent it makes sense to articulate how the entire system is engaged in a shared cognitive task: e.g. to what extent is a whole medical department or CERN engaged in a shared cognitive task? (see Magnus 2007, pp. 302-308 for similar concerns)[9].

In these more esoteric cases the answers to these questions are unlikely to be precise and the pragmatic constraints I outlined in the general introduction to the thesis militate against accepting vagueness here. Building from Magnus (2007, p. 301), and Vaesen (2011, p. 379) we can identify a range of possible different types of system:

1. One agent and one cognitive artefact (dyad)
2. One agent and multiple cognitive artefacts (one-to-many)
3. Multiple agents (either a dyad or many)
4. Multiple agents and one cognitive artefact (many-to-one)
5. Multiple agents and multiple artefacts (many-to-many)

[9] In chapter three I return to the complexities of cognitive tasks and multiple levels of analysis.

Even if we dismiss points [1] and [2] as only being "extended cognition"[10] and [3] as only "collective cognition", the last two options are still quite a motley assortment: i.e. how many is many before these systems count as qualitatively different from each other (e.g. groups of three components versus groups of twenty, etc.)? The problems outlined in the previous section – of size scaling and cognitive bloat – still remain because nothing here so far offers a reasonable boundary criterion[11].

The most obvious solution at this juncture might be to offer some limiting factors, so that distributed cognition is constrained to a narrower definition that applies only to a smaller subset of these cases and to regard other cases as inappropriate or fallacious. For instance, Heersmink (2015) has proposed a multidimensional framework to organise and explore dyadic distributed cognitive systems. He deems group sized systems to be composed of combinations of these dyads and argues that much larger systems are methodologically intractable because it is not feasible to properly trace out all the cognitive partnerships in these massive systems (2016, pp. 524-525).

However, I do not think that this debate can be resolved in such a straightforward manner for a number of reasons. Firstly, I think that limited approaches to the problems of distributed cognition that exclude large swathes of empirical research is an undesirable outcome. This is not to endorse an "unscientific motley" – these cases do need to be organised adequately – but it is to recognise that many of these studies examine ecologically salient human cognitive systems that

[10] In general, I will prefer the term distributed cognition for the reason that the term 'extended' is misleading. As Sutton notes, "...the word 'extended' can easily be misread as assuming a more basic inner cognitive system which only spreads later in development" (2015, p. 429). Arguably this contributes to some of the problems I will return to in chapter three.

[11] I acknowledge that one could argue that some of these can be seen as smaller systems nested within larger configurations – I return to this topic in chapters three and four where I provide a methodology for approaching this in a principled manner. However, even accepting this state of affairs does not significantly reduce the challenges facing the proponent of distributed cognition. Additionally, one might observe that from the list of examples I began the chapter with, a majority of the cases deal with small groups. One could then question whether there are any identifiable features across these, and whether these are replicated in larger groups. In chapter three I deal with the topic of scaling up and down, and in chapters five and six I shall show how one can look at comparable systems in terms of cultural practices.

would otherwise be unexamined and unexplained. Hence, I will aim to balance preserving this great variety with my pluralistic approach whilst also providing principled criteria for how to designate distributed cognition and organise these differing kinds of systems.

Secondly, many limited approaches to distributed cognition either explicitly or tacitly endorse what Kaplan (2012) has labelled an information "bottleneck" or "bandwidth" boundary demarcation criterion – in which the interior of the cognitive system is differentiated from its exterior according to the relativistic flow of information (e.g. see Cheon 2014; Davies & Michaelian 2016; Huebner 2014; Hutchins 1995a, 2001, 2010a; Poirier & Chicoisne 2006; Sutton 2006). In chapter two, with reference to Hutchins' (1995a) famous navigation team case study, I will return to this issue and show that it is inadequate for a number of reasons. Most notable of which, for our immediate discussion, is its inability to properly nullify the problem of cognitive bloat because it cannot properly differentiate between the genuine components of the system and mere causal background features.

A third issue that raises far greater concerns – which I think merits us taking an alternative approach – is the problem of cognition distributed in time. My main contention, that I return to throughout this thesis, is that simpler limited approaches to distributed cognition fail to adequately account for what distributed cognition is on their own terms because they only focus on a single or narrow temporal slice of the system. I refer to this as the synchronic temporal level – processes occurring over only a short period of time related to a specific activity. But by solely focusing on a synchronic level of analysis this theoretical position leaves itself unable to properly explain how these relations are actually distributed. I now explain this in more detail by returning to the work of one of the original exponents of distributed cognition, Ed Hutchins.

1.4. The problem of cognition distributed in time

Hutchins observes that when one examines human cognitive activity in the wild there are three "interesting" ways in which human cognitive processes can be distributed:

> …cognitive processes may be distributed *across the members of a social group*, cognitive processes may be distributed in the sense that the operation of the cognitive system *involves coordination between internal and external (material or environmental) structure*, and processes may be *distributed through time* in such a way that the products of earlier events can transform the nature of later events. The effects of these kinds of distribution of process are extremely important to an understanding of human cognition. (2001, p. 2068. My emphasis. Also see 2006, pp. 277-280; Hollan et al 2000, pp. 176-178)

For clarity, the three ways (in order of complexity) are: [1] cognition distributed as the coordination dynamics between external and internal cognitive resources (i.e. the dyadic interactions between an agent and a cognitive resource); [2] socially distributed cognition; and [3] cognition distributed across time. Hutchins argues that these kinds of distribution are vital for properly understanding human cognitive activity[12].

Most discussions in the philosophical literature have tended to focus almost exclusively on socially distributed cognition (and*, a fortiori*, coordination dynamics to the extent that one can understand larger spatial configurations as composed of smaller nested dyadic relations of agents and artefacts) and *have entirely overlooked how cognition can be distributed in time*. This latter issue is part of a more widespread criticism of the extended mind and distributed cognition literature: that these discussions have not taken temporal factors into account adequately; and that such an oversight has led to a variety of distortions and inaccuracies in our explanations of human cognition in the wild (Hutchins; 1995a, 2006, 2008, 2011; Johannsen 2010; Kirchhoff 2012; Menary 2013a;

[12] This is a point that I will support through a practical demonstration in the remainder of this thesis (see esp. chapter six).

Sutton 2008). I think that if we do not take diachronic temporal factors into account properly then we cannot give a correct account of spatially distributed cognition. Viz. although the limited view, which only focuses on how cognition is distributed spatially and synchronically, appears to be the more appealing option at first glance; in actual fact, by only focusing on this overly-confined set of factors, this position cannot properly explicate even the limited phenomena it attends to because these factors are highly constrained and shaped by longer term temporal factors which have been excluded from analysis. I demonstrate this below with a simple example.

Additionally, although Hutchins does emphasise that temporal factors are vital for understanding human cognition, I think that he has been insufficiently precise in how we think about the various ways in which cognition can be distributed in time. There are many different ways in which we can consider cognition to be distributed across time and so we need a taxonomy or classificatory system that prevents confusion here. Following Bietti and Sutton (2015) and Nersessian (2005), I divide these into a synchronic timescale and two major diachronic timescales: developmental and historical. The synchronic timescale can be considered to be occurring over micro-periods (e.g. the milliseconds involved in molecular interactions) to slightly longer periods of minutes and even – depending on the task – hours (corresponding to behavioural timescales). These shorter timescales, which we can regard as snapshots of a particular distributed cognitive system, are embedded within two differing diachronic timescales.

Firstly, the *developmental trajectories* of the system itself – these are the ontogenetic timescales over which a particular system evolves and, in the case of humans, becomes adept at some particular task in which they are repeatedly engaged. Menary (2014, 2015) refers to this as the enculturation thesis. Secondly, there are also longer-term timescales that go beyond the ontogenetic lifetimes of individuals – *cognitive historical factors* – and involve transgenerational

transmissions and refinements of epistemic tools and cognitive practices (Dutilh Novaes 2013; Hutchins 1995a; Menary & Gillett 2017; Nersessian 2005; Sterelny 2003, 2012; Tomasello 1999). These two diachronic timescales can be considered over years, decades, and even centuries in the case of the latter. I shall argue that both diachronic scales are crucial for understanding human cognition.

The destabilising impact that omitting these temporal factors can have on an analysis of distributed cognition can be made clear through considering a brief example. Before doing so it is important to note that the imprecision in Hutchins' descriptions of temporal scales is also present, to a lesser extent, in how he describes cognition distributed spatially. As will be clear from the discussion in section 1.2, the two notions of dyadic coordination dynamics and socially distributed cognition are woefully insufficient to account for the stunning diversity in the literature. In chapter three I will devise a framework for tackling and mitigating this problem and the issues that come along with attempting to cover such a wide range of phenomena.

Since much of the conceptual armature for arguing for the importance of time will be developed throughout the thesis, this first toy example is overly simple for the sake of clarity. Later examples, particularly in the concluding chapter, will be based on far more ethnographic detail – a feature that I agree with Huebner (2014) is essential for debates about distributed cognition. Norman (1991) details the example of a checklist as an everyday cognitive artefact that helps an agent perform a memory task. In chapter three I shall go into more details of how the manipulation of these devices alters both the problem space and also the neurocognitive profile of the agent (see Menary 2007a, 2014). Here my concern is to demonstrate that properly understanding how a checklist is used for a synchronic activity – in Norman's example the task is to perform pre-flight safety checks – requires paying attention to the diachronic factors. For instance, in order for the

agent to appropriately use the checklist so that they can achieve the specific task they must have learnt the "cognitive norms" (Menary 2007a) that govern how the device is physically manipulated[13]. For instance, these govern what order one reads the list (e.g. top to bottom, left to right), and they also dictate the relationship of symbols to affordances in the task space, etc. This learning process involves a developmental trajectory by which the agent becomes proficient and "enskilled" in its operation (Ingold 2000, p. 416).

The other major diachronic factors involved here are the longer term historical processes by which the checklist's format and structure have been created, maintained and refined. Hutchins (1995a, pp. 293-302) notes that the complexity of checklists and written procedures are often overlooked. In terms of cognitive historical factors this involves the extent to which the relevant task space is altered over multiple generations – what Sterelny (2003) refers to as "epistemic engineering". If we want to understand the cognitive behaviour of an agent in a particular synchronic time-slice and we fail to take into account how these diachronic factors have radically altered the task space and the agent's profile then we will fail to include vital explanatory processes in our unit of analysis. Hutchins (1995a, 2008) argues that this can lead to an "over-attribution" of the cognitive work to the skull of an individual for what is better understood as a multiply distributed process – across both space and time.

The core point here is that the limited version fails on its own terms. I.e. if we want to understand how cognition is distributed spatially over the synchronic scale, then we must also understand the diachronic temporal factors – both developmental trajectories and cognitive historical elements – that are related to this. The question is now how we can do this in a tractable manner. Once one introduces new factors this raises the spectre of cognitive bloat: viz. at what

[13] I discuss these in more detail in chapter five (section 5.4.1).

point in time do we stop our analysis? What rubric or measure is there for accounting for both synchronic and diachronic cognitive processes? Again, I suggest a good starting point to exploring these issued involves returning to Hutchins' work. From here, with modification, I will build my pluralistic framework.

1.5. The flexible unit of analysis and the problem of methodological bloat

When outlining what I referred to as the limited approach to distributed cognition in section 1.3 I noted that a common theme amongst various positions is that distributed cognition refers to the claim that cognitive processes are not limited to the skulls of individuals. I think that a crucial point of the debate here is that this is *only one aspect* of how Hutchins originally presented the idea. Hutchins defined distributed cognition as a "subfield" (2006) differentiated from mainstream cognitive science according to two principles (this passage is worth citing in full because there are many issues here):

> What distinguishes distributed cognition from other approaches is the commitment to two related theoretical principles. *The first concerns the boundaries of the unit of analysis for cognition*. While boundaries are often a matter of tradition in a field, there are some general rules one can follow. Bateson (1972) says *one should bound the unit so that things are not left inexplicable*. This usually means putting boundaries on units where the traffic is low. *The second principle concerns the range of mechanisms that may be assumed to participate in cognitive processes*. While mainstream cognitive science looks for cognitive events in the manipulation of symbols (Newell et al 1989), or more recently, patterns of activation across arrays of processing units (Rumelhart et al 1986, McClelland et al 1986) inside individual actors, *distributed cognition looks for a broader class of cognitive events and does not expect all such events to be encompassed by the skin or skull of an individual*. (2001, p. 2068 emphasis added)

Let us make the two principles clear: [1] the study of cognition should be based on a flexible unit of analysis; and [2] there are supraindividual cognitive processes. There are several important points related to this. Firstly, the relation of these two principles is important. Secondly, they differentiate distributed cognition from traditional approaches in cognitive science. Thirdly, there are concerns here about instrumentalism. Fourthly, there is an overly cursory attempt to define the boundaries of the system here that is inadequate for tackling the problem of cognitive bloat but which also, in conjunction with other going concerns, introduces a new issue: methodological bloat. I now consider each of these in turn.

Firstly, the relation of the two principles is important. Often, as was discussed in section 1.3, proponents of distributed cognition begin with the second of the two claims: supra-individual cognitive processing. But here Hutchins sets up his definition so that we begin with a claim about a flexible unit of analysis, and then *it is on this basis* that we can postulate the *possibility* of cognitive "events" distributed over mechanisms that are supra-personal. To reiterate, Hutchins' way of defining distributed cognition is with a reasonable methodological claim that naturalists should accept: when investigating a phenomenon, we should not decide our unit of analysis *a priori* but should instead negotiate the boundaries of the investigation through the investigative process itself. As Hutchins puts it: "A central claim of the distributed cognition framework is that the proper unit of analysis for cognition should not be set a priori, but should be responsive to the nature of the phenomena under study" (2010b, p. 426). It is on the basis that the unit of analysis is not predefined that we can then accept the possibility that it is potentially not always bounded by the skull of an individual. Therefore, the possibility of supra-individual cognitive processing mechanisms is a corollary *possible* ontological consequence of the reasonable methodological position.

The second issue here is how Hutchins' definition based on a flexible unit of analysis explicitly differentiates distributed cognition from traditional positions in cognitive science. The latter treat the isolated individual as the de facto unit of analysis. As Wilson (2005, p. 227) puts it: "For most of not only its short history but its long past, the science of the mind has been the science of the individual mind, with human beings serving as our paradigm example of agents with minds, cognitive agents". This is a position often referred to as "methodological individualism" (see Watkins 1952). The proponent of distributed cognition must accept that the current mainstream position in cognitive science is, mostly, focused on representations that are 'located' inside the skull of the individuals. As such, it is beholden upon a legitimate project on distributed cognition to provide both reasons and evidence to the methodological individualist which cannot be simply dismissed as explanatorily superfluous (Huebner 2014; Theiner & O'Connor 2010). This is one of the aims of remainder of this thesis.

It is important to clarify that distributed cognition should not be understood as diametrically opposed to mainstream views; i.e. as something like a Kuhnian (1970) paradigm shift or "break" – as Brown seems to hold (2011, p. 18). Here I am in agreement with others – such as Button (2008)[14], Davies and Michaelian (2016), Hollan and colleagues (2000), Huebner (2014), Sutton (2015), and Sutton and colleagues (2010) – that distributed cognition is consilient in some ways with mainstream views.

Again, key to understanding this is the flexible unit of analysis. As noted above, most philosophers and theorists have focused on the possibilities of supra-individual cognitive processing. But what is often overlooked in the literature, with the notable exception of Davies and Michaelian (2016), is that this also involves the possibility that the appropriate unit of analysis for the

[14] However, it is important to note that whilst Button sees distributed cognition and mainstream views in cognitive science as consistent he rejects both of them.

investigation of some cognitive phenomenon might equally well be the isolated individual or some sub-personal cortical region. Therefore, when investigating a cognitive phenomenon using the distributed cognition stance one must accept that in some cases the skin and skull will be an appropriate boundary; and that in other cases taking the whole individual into account will be too complicated. But importantly – and "interestingly" (Hutchins 2001, 2006) – in other cases the prearranged convention of setting the boundary at the skin will abstract and cleave off important difference makers:

> For some sorts of phenomena, the skin or skull of an individual is exactly the correct boundary. For some phenomena, the whole person is just too big and including the whole organism would involve too many interactions. For other phenomena, setting the boundary of the unit of analysis at the skin will cut lines of interaction in ways that leave key aspects of the phenomena unexplained or unexplainable. (Hutchins 2010b, p. 426)

As noted above, it is beholden upon the proponent of distributed cognition to clearly specify: what phenomena; how the flexible boundaries are demarcated in a principled fashion; and which properties are inexplicable for the methodological individualist. In the final section of this chapter I outline how my pluralistic approach meets these demands. Additionally, it is important to note that this is not just a question about boundary alignment (although this is certainly a major concern) but is rather an exploration of how multiple levels of analysis and interactions – from the neural to the niche – can be traversed and understood using a distributed cognition framework.

Thirdly, although I think that Hutchins' work is a good basis for approaching some of the problems I have outlined in previous sections (heterogeneity and size scaling), it also introduces new problems. A minor issue related to instrumentalism that can be resolved pretty straightforwardly, but which also leads to a more complex issue relating to the practical tractability of the flexible unit

of analysis. The ire of critics has been raised by a number of philosophically loose statements that Hutchins has made in regards to the methodological angle. For example,

> …to take the distributed perspective is not to make any claim about the nature of the world. Rather, it is to choose a way of looking at the world […and therefore…] there is no series of experiments or set of observations that could prove or disprove the distributed cognition perspective, because the perspective itself makes no empirical claims. (2014, pp. 36-37; see also 2010a, pp. 706-707)

Understandably, this has led some critics such as Hutto and colleagues (2014) to state that this view of distributed cognition is too 'instrumentalist' – i.e. if the position makes no claims about nature it is hard to see what taking the distributed cognition position amounts to. I agree that such comments are problematic because if there is no ontological bite in how we are to understand distributed cognition, then it is hard to understand on what basis we are arguing for a view of cognition as distributed in the first place. This issue could be seen as a conceptual cousin of cognitive bloat since the issue here is how we define the bounds of the system. So, I also agree that we should not understand the flexible unit of analysis in such a loose manner that it does not say anything meaningful about the world. But given that Hutchins also makes a number of contradictory statements on this topic that make it clear that he does think that there are specific ways in which cognition is distributed beyond the skull (e.g. 2001, p. 2068; 2006, pp. 277-280; Hutchins 2008; 2011; see also Hollan et al 2000, pp. 176-178), I agree with Sutton's (2015) counter-argument against Hutto and colleagues that distributed cognition is an interesting and viable project to pursue. Sutton argues that far from being toothless, Hutchins does "make many specific claims about the nature of cognition in particular settings" and rather than being a "blanket retreat from ontological

commitment", this is indicative of a pragmatic attitude towards metaphysical claims that is driven by scientific work (2015, pp. 429-430; also see Sutton et al 2010)[15].

The fourth and final issue in the passage above: the problem still remains of how to define the flexible unit of analysis in a coherent and precise manner. Hutchins does offer the tentative phrase "bound the system where traffic is low". As noted above in section 1.3 this is the information bottleneck approach and I shall return to discuss its limitations in the next chapter. But the other aspect here is that the flexible unit of analysis doubles the problem of heterogeneity because now the suggestion is that each particular system under investigation requires multilevel analysis. As Hutchins puts it: "…something special might be happening in systems of distributed processing, whether the processors are neurons, connectionist nodes, areas of a brain, whole persons, groups of persons, or groups of groups of persons" (2001, p. 2068). The critic can argue that trying to account for too much becomes methodologically over-excessive – i.e. rather than *enabling* researchers out in the wild, critics claim that it actually results in hamstringing them. Indeed, Hollan, Hutchins and Kirsh acknowledge that this is a serious issue (2000, p. 193). I shall term this problem

[15] As stated in the introduction, I am here conducting a conceptual reanalysis of distributed cognition with the aim that this work has a practical import – both conceptually but also, if possible, empirically as well. As such, my aim in this thesis will be to make the notion of "interesting" ways in which the unit of analysis of cognitive systems is demarcated both empirically tractable and conceptually robust. But in order for my pluralistic framework to be able do important theoretical and practical work, my position must have a non-ad hoc means of handling the issue here about how one draws the boundaries of the cognitive system. Viz. so that my account does not slip into instrumentalism; i.e. that my claims about distributed cognition are entirely dependent on the observer, and thus, we are not really making a claim about the world. To render distributed cognition this arbitrary would raise questions about why adopt the position at all in the first place.

Here I ally myself with Glennan's (2017, p. 93) recent observation that there is a family of philosophers from differing backgrounds who have attempted to be faithful to two propositions: [1] that one should adopt a perspectival/stance aspect towards the complex sciences; but also [2] maintain the objectivity of the target objects of enquiry. For instance, despite big differences, Wimsatt (2007), Ladyman & Ross (2007), and Giere (2006) all hold this kind of view – what Glennan refers to as a pluralistic-realist approach towards ontology. The key aspect of these views that is relevant for the discussion of distributed cognition – specifically the notion that the unit of analysis is flexible – is the claim that "there are multiple correct ways to parse the world" (Glennan 2017, p. 93). And that the process of determining these is dependent both on the structure of the world but also on the interests of the observer. In the remainder of this thesis – particularly chapters three, four and five – I will outline what I think are the main ingredients for negotiating this tricky terrain of boundary drawing in the exploration of ecologically salient human cognition.

“methodological bloat” and, prima facie, it seems like a reasonable objection. These kinds of objections are often levied at projects that attempt interdisciplinarity in regards to culture and cognition. Viz. many critics have claimed that making psychology dependent on the other sciences is not empirically tractable and is wrong-headed (e.g. Fodor 1980). Tooby and Cosmides capture this complaint rather pithily in a humorous analogy wherein projects such as distributed cognition are described as being in “…a situation resembling some nightmarish story Borges might have written, where scientists are condemned by their unexamined assumptions to study the nature of mirrors only by cataloguing and investigating everything that mirrors can reflect” (cited in Sutton 2010, p. 214). These worries motivate methodological individualism.

The dangers of intractability are a serious concern (and notably, are the inverse cousins of Hutchins’ worries about having an overly restrictive unit of analysis that cleaves off important explanatory difference makers). Because of these worries about how much complexity environmental factors add to the problem space, Anderson (2010, pp. 296-297) states that it is beholden upon those who argue for the importance of environmental context to go beyond gestures and provide adequate details that lead to more illumination about cognition than mystification. However, merely acknowledging that something is difficult or ambitious is not to despair; instead Hollan and colleagues see it as a challenge to roll our sleeves up and explore how to tackle this complexity in a viable way (2000, pp. 193-194). I will resolve this issue in chapter five by emphasising the importance of normative patterned practices – adopted from the work of Menary (2007a, 2016) – as a factor that operates at multiple levels of analysis (I provide more details on this below in the précis).

Given that there are so many issues facing the proponent of distributed cognition one could be forgiven for seeing the position as untenable. However, in the following section I provide a précis

of how I intend to solve the many challenges I have discussed in this chapter – heterogeneity, size scaling/multilevel analysis, cognitive bloat, the problem of time, and methodological bloat – using a pluralistic framework.

1.6. A pluralistic framework

Having surveyed the many problems facing a proponent of distributed cognition we are now in a good position to understand the challenges and why a simpler limited approach will not be adequate. In its stead I propose a pluralistic framework that builds from Hutchins' starting point that distributed cognition is premised on a flexible unit of analysis that is not pre-given but is instead discerned through the process of inquiry. However, Hutchins' work still lacks principled criteria by which the flexible unit of analysis and the many ways in which cognition can be distributed across time and space can be designated. I propose the following framework – drawing together several naturalistic and pragmatically motivated criteria from the literature – which provides a comprehensive means of demarcating the flexible unit of analysis in a principled fashion.

As noted above, my aim here is to avoid a parochial approach. As such, these criteria have been constructed with the aim that they are useful to practitioners in the wide range of fields that distributed cognition analyses are being conducted. A core component of this project requires paying particular attention to the details of specific case studies and being responsive to the wide range of phenomena that are evident in human cognitive activity. Attention to ethnography is crucial because when discussions about distributed cognition become untethered from this requirement the debates quickly descend into pseudo-problems that do not have much relevance to practical research. By paying specific attention to particular case studies I am advocating a naturalistic approach to these questions that is in line with many others in the field who also argue

about the importance of paying attention to the empirical details – for instance, Hurley writes: "The issues between internalism and externalism should be resolved bottom up by […] scientific practice, not by advance metaphysics: by seeing whether any good psychological explanations are externalist, not by deciding on a criterion of the mental and using it to sort explanations as constitutive or not" (2010, p. 107; also see Huebner 2014; Kaplan 2012; Sutton et al 2010; Theiner et al 2010; Wilson 2001, 2004)[16].

The multifarious problems posed against distributed cognition necessitates taking a pluralistic approach that uses a number of theoretical principles: task-specificity (Davies & Michaelian 2016), mutual manipulability (Kaplan 2012), and normative patterned practices (Menary 2007a, 2016). I have selected these principles through a review of the philosophical literature. A pluralistic approach is needed because no single principle is sufficient to tackle the many problems I have outlined in this chapter. Instead, they are consilient with each other, and operate together – with refinements – to provide us with a complex and sophisticated understanding of what distributed cognition is.

1.6.1 Task-specificity (multilevel analysis)

Both [1] the boundaries of the target system under investigation and [2] how cognition is distributed in time and space can be identified by focusing on a particular well-defined cognitive task. The importance of task-specificity is commonly recognised in the literature (e.g. Davies & Michaelian 2016; Huebner 2014, p. 14; Hutchins 1995a, pp. 157-158; Magnus 2007; Magnus & McClamrock 2015). However, my account here goes beyond the work of these previous thinkers in two ways.

[16] To be fair it must be noted that Hutchins' work is exemplary in regards to paying attention to the particular details of case studies. However, I would contend that he does not leverage his ethnographic work appropriately towards the tackling of philosophical questions.

Firstly, previous accounts have sometimes overlooked how the type of task engaged in by the distributed cognitive system can have a dramatic impact on how agents and artefacts are coordinated (Norman 1991). As such, I adopt Steiner's (1972) taxonomy of task-types. Secondly, task-specificity avoids having an agent-centred approach (see Hutchins 2011). As Davies and Michaelian (2016) make clear, many other approaches to distributed cognition struggle to consider multiple levels of analysis because they are inappropriately "scaling-up" or "scaling-down" an agent-centric view about cognition. This entails carrying over inappropriate conceptual baggage and not adequately differentiating between different levels of analysis. I also clarify the notion of supra-individual cognitive levels with reference to Caporael's (2014) notion of "core configurations" and Wimsatt's (1986, 2007) definition of emergence as the failure of aggregativity (also see Poirier & Chicoisne 2006; Theiner & O'Connor 2010; Theiner et al 2010). A task-orientated approach allows us to have a more encompassing framework that can adequately discuss the various levels of analysis in which cognition is purportedly distributed over spatial and temporal scales *on their own terms*.

1.6.2 The Mutual Manipulability Criterion

The unit of analysis for distributed cognition is focused on a shared cognitive task. This involves appropriately demarcated levels of explanation across multiple levels – because otherwise analysis will overlook how the use of certain cognitive tools alters both the task space and the agent's neurocognitive profile (Hutchins 1995a, 1999, 2006; Norman 1991; Menary & Gillett 2017). As such, the goal of the investigation is to properly enumerate the components of the system that are involved in this task (Hollan et al 2000; Hutchins 2014). Clarifying what these components of the system are, their relations, and the overall organisation of the system can be delineated using both mechanistic and functional decomposition (Bechtel 2008, 2009; also see Clark 2008, pp. 13-15).

These issues are sometimes overlooked in the literature on distributed cognition – hence why I have emphasised the importance of paying attention to the details of particular case studies. The consequences of this are an inability to properly differentiate between genuine components of a distributed cognitive system and features which are merely necessary causal background features. Kaplan (2012) raises this as a general complaint against externalist positions in the philosophy of cognition. I shall follow Kaplan in adopting the mutual manipulability criterion from philosophy of biology/neuroscience (also see Craver 2007). This criterion operates by postulating two forms of intervention on a system (top-down and bottom-up) in order to identify what is a genuine component of the system as opposed to a mere background condition. Interestingly, whilst the mutual manipulability criterion is a useful theoretical tool for investigating synchronic timescale phenomena, it breaks down once one considers longer term, diachronic, timescales. Therefore, it needs supplementing.

1.6.3 Normative Patterned Practices

The above points may strike some as being very heterogeneous and without a common core principle linking them in a more parsimonious fashion – i.e. methodological bloat. The objection being that the investigation of distributed cognition would be too taxing since it tries to cover too much material. Now the issue has become even worse – since even the very definition of distributed cognition is appealing to a huge range of theoretical material and empirical case studies before the investigation has even properly begun. My answer to both varieties of the same issue draws on the notion of normative patterned practices (Menary 2016). These are patterns of physical interactions spread across a population which govern and mediate the manipulation of cognitive resources by agents towards a particular cognitive task (Menary 2007a, 2010a, 2013b, 2015).

Normative patterned practices can be considered as a 'heuristic' operating at all levels and scales of the purported ways in which cognition can possibly (but not necessarily in each case) be distributed across both time and space. For instance, Roepstorff and colleagues (2010) claim that they are the medium by which the macro- and micro- spatial scales interact. This also has the added benefit of bringing both diachronic temporal dimensions into play because the relationship of the micro and macro takes place through the enculturation of normative patterned practices from the niche-level being correlated with plastic changes at the level of the brain, and having an impact on the coordination dynamics of how agents manipulate cognitive resources and interact with each other (Menary 2007a, 2012, 2013a, 2013b, 2014, 2015; Roepstorff et al 2010). I argue that this claim is supported by a wide range of empirical evidence related to the ways in which cultural practices affect neural plasticity (see Downey & Lende 2012a; Draganski & May 2008 for reviews). As such, normative patterned practices provide us with a means of tackling the problem of cognition distributed in time by facilitating the exploration of how various temporal scales interact.

The key point about normative patterned practices is there is no division here between the rules of the interaction and the mediation of the interaction itself (see Kendal 2011, p. 245; Moore 2012, p. 299[17]; Ortner 1984; Roepstorff & Bubandt 2003, pp. 14-15; Roepstorff et al 2010). This not only allows us an *empirically tractable* vehicle that captures the more esoteric aspects of culture such as rules, norms and knowledge (Moore 2012, pp. 304-305; Roepstorff & Bubandt 2003, pp. 14-15; Roepstorff et al 2010, pp. 1051-1052); but also makes this more tractable by not dividing it from the physical interactions that constitute the specially engineered environments within which the

[17] There are questions here about having an all-encompassing term (see Odling-Smee 2007 and Kendal 2011, p. 245 about various distinctions between mental entities, behaviour, etc.) since there is the danger of holism that occludes important distinctions. This is one of the reasons I will give in chapter two for rejecting Hutchins' Batesonian inspired and heavy-handed new metaphysics of mind based on "cognitive ecosystems" (2010a, 2014). The rebuttal to this worry is that normative patterned practices are a more tractable and useful notion than the more nebulous notions of 'culture' (see Moore 2012; Roepstorff et al 2010; Vogeley & Roepstorff 2009).

human distributed cognitive activity is embedded – what I shall call, following Sterelny (2003), Menary (2014), and others (e.g. Boyd et al 2011), a *cultural-cognitive niche*.

Additionally, a focus on the interactional level – by noting that this is what we are concerned with in both synchronic and diachronic factors – enables us to give a more nuanced way of understanding what, and how, cognition can be distributed over larger spatial and temporal scales. For instance, Menary (2007a, 2014, 2015) emphasises that normative patterned practices are central to understanding the mastery and learning of cognitive capacities through development; but also, that they are what is transmitted inter-generationally – vertically, obliquely, and horizontally within cultural-cognitive niches. Thus, normative patterned practices are involved in thinking about the temporal organisation of action – and must be seen within the context of a longer term dynamic unit of analysis beyond the immediate activity. I.e. there is an emphasis on the diachronic and not just the synchronic (Ortner 1984, pp. 150, 158). Furthermore, a practice-based approach takes history and development seriously (p. 159). This again relates to the medium by which the cultural-cognitive niche interacts with cognition. Roepstorff and colleagues (2010, p. 1052) specifically claim that "[p]atterns of practice at the level of social interaction correlate in relevant ways with neural and psychophysical patterns". As such, they are the vehicle by which we can consider how the macro and micro scales across which cognition is distributed, interact – and therefore provide us with a means of tackling both the problem of cognition distributed in time and the problem of methodological bloat.

I now defend these claims in the rest of the thesis.

1.7. Summary

In this chapter I have outlined the general problems facing the proponent of distributed cognition. The term is used widely in many different fields and has been applied to a stunning array of miscellaneous kinds of systems. Much of this work is incredibly useful but suffers from the deficit of lacking principled conditions under which the term is being applied. The sheer variety of types and sizes of systems involved also raises concerns about how we compare and apply theoretical terminology without leading to inappropriate conceptualisations. Furthermore, and perhaps most damagingly, the heterogeneity in the literature exacerbates the problem of cognitive bloat by confusing the search for principled boundary criteria. However, contrary to a simple solution in which distributed cognition is limited to some smaller subset of these cases I have sought to devise a more parsimonious framework with the intention that it will have wide applicability in the sciences and humanities. I have argued in this manner for two reasons: firstly, a parochial approach which omits much good empirical work is likely to be simply ignored in many of the fields where distributed cognition is being used. But more problematically, the issue of cognition distributed in time demonstrates that the limited view is not able to properly account for distributed cognition properly because many synchronic processes are shaped by diachronic factors.

My more complex view devised to meet all these challenges builds from and revises Hutchins' two principles: [1] a flexible unit of analysis; and how this entails the possibility of [2] cognition distributed in both space and time beyond the skull of individual agents. I concluded the chapter with a précis of my pluralistic framework – composed of task-specificity, the mutual manipulability criterion, and normative patterned practices. The précis also outlined how these criteria provide a principled means of both demarcating the flexible unit of analysis and designating how cognition is distributed both spatially and temporally. This framework must avoid the problem of methodological bloat (i.e. be tractable for theoretical and practical researchers) and also make

claims about specific case studies that have a practical import (thus justifying the adoption of a distributed cognitive stance as opposed to the more standard methodological individualism).

In the next chapter I outline the most famous case study of distributed cognition and how it is usually explicated in terms of information flow by the limited approach. After showing the shortcomings of this I will then begin the positive defence of my position.

2: The navigation team and information bottlenecks

2.1. Introduction

In the first chapter I outlined the general problems facing a proponent of distributed cognition. There is an unwieldly diversity of purported cases in the literature and this has the potential to make the term vacuous – as well as creating problems related to size-scaling and exacerbating the problem of cognitive bloat. Additionally, there is the problem of cognition distributed in time. To tackle these I proposed a pluralistic, as opposed to parochial, approach. In this chapter I outline Hutchins' classic case study of the navigation team on board an aircraft carrier (1995a) as it is often portrayed in the philosophical literature. My key intention in this exposition is to show the limitations of the standard portrayals of this case study in terms of "information flow"; this exposition then presents an opportunity to demonstrate the superior aspects of my pluralistic framework as I return to this case study throughout the rest of the thesis.

I begin this chapter by justifying my discussion of this case study (section 2.2), before providing a concise but detailed account of the case study itself – whilst also clearing up some errors present in the philosophical commentaries (sections 2.2-2.5). The latter half of this chapter examines why the standard "information flow" approach is insufficient for multiple reasons. The focus on the synchronic group activity of agents and artefacts towards a task that could not feasibly be achieved by a solitary agent is often given as the justification for why this is a case of distributed cognition (e.g. Giere 2002a; Huebner 2014; Magnus 2007). But this still leaves the question of how to demarcate the boundaries of the flexible unit of analysis. In this chapter I shall discuss Hutchins' proposal that we can use an information bandwidth criterion – "putting boundaries where traffic is low" (2001, also see 2010a, 2014) – as the criterion for discerning this. For our purposes we can

define information flow as the cognitive processing that occurs in a cognitive system (the relationship of inputs and outputs). As such, the claim here is what Clark (2008) refers to as an "information bottleneck" which demarcates what is 'inside' the system as having higher informational bandwidth than features which are 'outside' the system.

This is not an esoteric view. Kaplan (2012) notes that this idea has also been explored by other theorists from very different perspectives such as Haugeland (1998) and Simon (1996). A similar position has also been defended recently by Huebner (2014), who emphasises the importance of the propagation of representational states across various media with a shared "trading language"[18]. In general, it is common for accounts of distributed cognition to be cashed out in terms of information flow (e.g. Cheon 2014; Davies & Michaelian 2016; Heersmink 2015; Poirier & Chicoisne 2006; Sutton 2006; Sutton et al 2010). However, whilst this idea might seem appealing at first glance, several philosophers have raised issues. For instance, Clark (2008) and Kaplan (2012) have pointed out that we have examples of high information interfaces between systems that are not integrated and are separable (e.g. distributed computing, prosthetics). Furthermore, Eliasmith (2009) points out that the information bandwidth criterion is actually an argument *against* distributed cognition because there is a sizeable difference in the informational bandwidth inter-neuronally in contrast to neuron-environment interactions. Most damagingly, I shall argue that it fails to adequately tackle the problem of cognitive bloat because it cannot provide us with an adequate means of differentiating between genuine components of a cognitive system and mere causal background conditions. As such, my argument against information bottlenecks as an adequate demarcation criterion is two pronged: on the one hand it cannot stave off the problem of

[18] He takes this term from Galison's (1997) work on radically interdisciplinary scientific research projects – e.g. high energy physics experiments which require extensive collaboration between engineers and theoretical physicists who both have different technical vocabularies. Galison adopts the term from linguistic anthropology.

cognitive bloat; and on the other hand, it fails to provide us with an argument in favour of distributed cognition and is in fact inadvertently in favour of internalism.

To further complicate matters, I contend that Hutchins' ideas about information bottlenecks are also related to some unnecessary and overly grandiose metaphysics that he inherits from the cyberneticist Gregory Bateson (1972). This leads Hutchins to inadvertently postulate a new and quite radical a priori unit of analysis – the cognitive ecosystem – consequently contradicting his emphasis on the flexible unit of analysis that I have argued is a principal principle of distributed cognition. I discuss these issues in sections 2.6-2.7, paving the way in the rest of the thesis to building a positive replacement. Viz. the failings of Hutchins' more simplistic account acts as a *preliminary* justification for the complexity of my position (and thus an opening rebuttal against methodological bloat).

2.2. Why I have chosen the navigation example

The shortcomings of the simple view and the superiority of my pluralistic framework can be clearly demonstrated by using a detailed re-analysis of Hutchins' (1995a) seminal case study of a navigation team. Hutchins' ethnographic project was conducted in the 1980s on board a US naval vessel prior to the widespread use of GPS devices – he acknowledged that many of the practices described in his study would become obsolete and cease to exist as satellite technology improved (p. 33). For the sake of clarity this re-analysis will have to be conducted over two distinct stages: a first succinct pass conducted in this chapter describes the basics of the case study in its entirety – this is slightly more embellished but consistent with the limited presentations often found in the philosophical literature. My main criticism here is that philosophers often give overly concise descriptions of case studies deeming these other details as merely of ethnographic interest. However, I agree with Huebner's

(2014) recent contention that these ethnographic details are in fact vital for properly understanding distributed cognition.

I then return repeatedly in subsequent chapters to the navigation team to show how the three core principles – task-specificity, the mutual manipulability criterion, and normative patterned practices – not only tackle conceptual problems but also make practical insights and contributions to real world cases. In summary, I have chosen the navigation case study as the most pertinent site at which to compare the two views for three notable reasons: [1] it is seminal; [2] a re-examination reveals the inadequacies of the simpler approaches to distributed cognition; and conversely [3] allows me to demonstrate the benefits of my pluralistic approach in a concrete manner. I now discuss each of these in more detail.

Firstly, Hutchins' work is acknowledged by many – both critics and allies – as inspiring an increasing interest in naturalistic research into group cognition. For instance, Rupert (2011, p. 631) states that empirical work "greatly increased" after the publication of *Cognition in the Wild* and according to Google Scholar, this study has been cited over twelve thousand times[19]. Given this prominence in the literature, one could not provide a very good overview of distributed cognition without discussing it. For instance, Brown (2011, p. 18) recently described *Cognition in the Wild* as the "foundational text for research in distributed cognition"; Huebner (2014, pp. 153-156, 240) has stated that the navigation case study is one of the best and most plausible examples of distributed cognition; and Clark (1997, p. 76) once described it as "the most successful and sustained investigation of the cognitive properties of human groups to date". Even Rupert, a vocal critic of externalist positions, has stated that "Hutchins' work provides the most balanced and powerful example" of distributed cognition (2011, p. 634). Therefore, despite the fact that much has already

[19] As of 03/08/2017 it has been cited 12,020 times: https://scholar.google.com.au/citations?user=T5WQ2EcAAAAJ&hl=en

been written on this example, it is an excellent place to examine the conceptual foundations of distributed cognition.

Secondly, whilst this first reason could be seen as mere historical exegesis, there are other factors that provide good grounds for focusing our attention here. Namely, a reanalysis of this case study shows the limitations of portrayals of distributed cognition that only focus on information flow in synchronic activity. The navigation example is often a core feature of how the simpler view is presented and popularised (e.g. Giere 2002a, 2002b, 2007). So, demonstrating shortcomings here is especially damaging. As I mentioned in chapter one (section 1.2), the navigation example is sometimes presented as a salient example of a distributed cognitive system before an analogy is drawn to another phenomenon (e.g. Giere 2002a, p. 639). My point here is not to attack this argumentative style (although I think it is a dubious strategy if there is a large size discrepancy – as I discuss in more detail in the next chapter). Instead, by showing the limited view to be an *inadequate* account of the navigation example in general I am critiquing this position at a point of relative strength rather than attacking a straw-man. Viz. the navigation case study is the most famous example of distributed cognition and the simple account is the most popular account of distributed cognition. By showing that the latter is inadequate to explaining the richness of the former this provides us with very good reasons for searching for an alternative interpretation.

The third reason follows on from the second: a re-examination of the Hutchins case study allows me to demonstrate what my pluralistic framework is capable of. For instance, a central feature of my position is an emphasis on the importance of paying attention to the particular details of case studies and the implications this has for understanding distributed cognition. Hutchins' case study is apt for these purposes because, as Nardi (1996, p. 70) puts it: "*Cognition in the Wild* is a tour de force [of] ethnography". As such, it is an excellent model for the degree of detail needed for

a proper study of distributed cognition. Additionally, I acknowledged in the first chapter that my pluralistic framework seems at first to be overly baroque and complex (the problem of methodological bloat). This might lead one to favour the simple view on the basis of parsimony. However, I contend that this is *sacrifice of accuracy for simplicity*. Repeated and refined engagements with the navigation example in subsequent chapters will demonstrate how the various principles operate in a salient real-world case study. For instance, by examining the normative patterned practices present in this case study we are able to elucidate diachronic temporal scales which relate to the resilience of the task performance of the system and why the task is performed in the way it is (chapters five and six). These features, crucial for understanding the cognitive activity as it is synchronically, would not be apparent without this criterion.

Therefore, examining this case study not only demonstrates how features of the pluralistic framework are superior to the limited view, but also provides a performative demonstration of how these principles operate in a coherent manner – thus, allaying concerns raised at the end of chapter one about the tractability of this interpretation (viz. the problem of methodological bloat). Furthermore, it is also noteworthy that my opponent is not just the more prevalent limited interpretation of distributed cognition, but also the methodological individualist. Although I tackle these issues much more in later chapters, the re-analysis of this case study does also provide us with an early opportunity to demonstrate in a concrete manner how my approach to distributed cognition has features of which the methodological individualist should take note.

2.3. Preliminary details in the navigation team case study

Hutchins conducted an ethnographic field study of navigational practices in the US navy during the 1980s observing several crews. His case study focuses on an aircraft carrier named the *Palau*, an

extremely large ship: over 180 meters long, 30 metres wide and weighing roughly 15,500 metric tonnes un-laden (1995a, p. 7). Steering and navigating such a large vessel is a challenging cognitive task and requires a dedicated team – as Hutchins puts it: "A ship is a massive object; its inertia makes it slow to respond to changes in propeller speed or rudder position" (p. 48). I.e. there is a delayed response time between the actions of the crew and the behaviour of the vessel. Because of this time-lag in response it is incredibly important that the crew of the ship plan and anticipate their future positions and courses of action – especially for entering and exiting harbours and in confined waters.

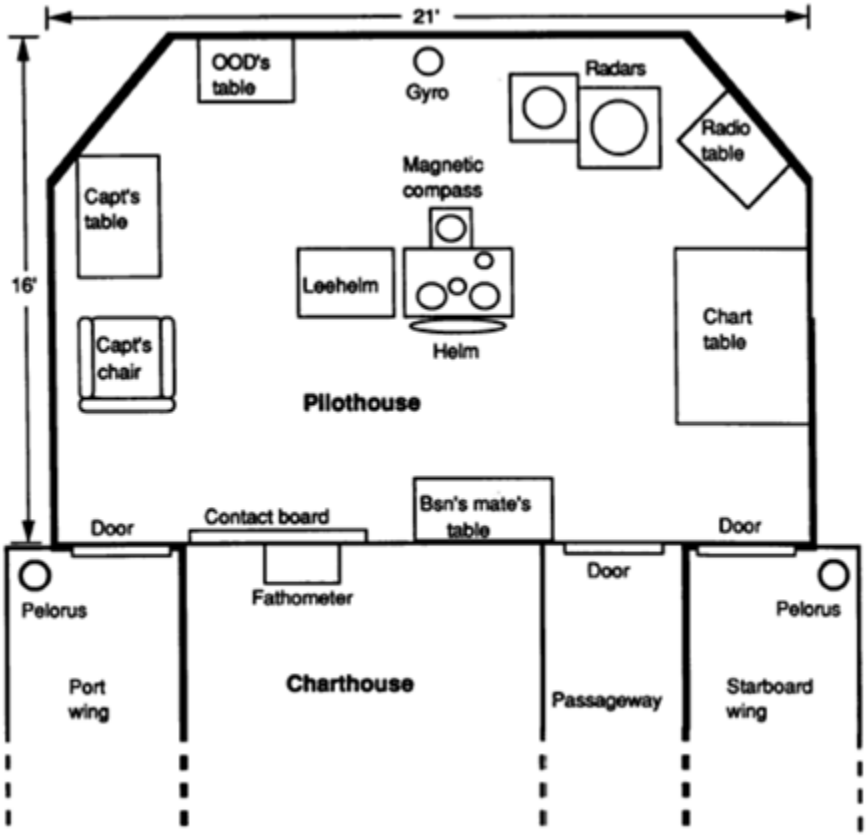

Fig 2.1: the layout of the bridge. The major hub of activity of the team is located at the chart table where the plotter and bearing recorder are located. During SEA AND ANCHOR DETAIL two pelorus operators are stationed outside – one each on the starboard and port wings to use alidades for the observation of landmarks. [Diagram from Hutchins 1995a, p. 27]

The navigation team are situated on the bridge of the ship (see figure 2.1 above). We can divide the team into two interconnected but roughly separable units: one being the Officer on Deck and the pilot who are more attentive to the actual steering of the vessel; and a second unit which is the focus of a majority of Hutchins' observations. This second unit performs what Hutchins identifies as a "computational ritual called the *fix cycle*" (p. 26 original emphasis). Henceforth, for the sake of simplicity, this is what shall be referred to as the navigation team. When the ship is not in port, this 'ritual' is carried out at all times. Therefore, unlike many of the other departments on board a naval vessel, the navigation team is effectively constantly on duty – which contributes towards making this cognitive task more taxing.

The fix cycle can be defined briefly as the activity of the navigation team in ascertaining the current position of the vessel and its future position given its current course. Hutchins identifies these as the central questions of navigation: 'where am I?', and 'given that I am at point X how do I get to point Y?' (pp. 12-13, 52). For our purposes the fix cycle is *primarily* performed by four agents using a range of different cognitive artefacts. However, it is important to note that this is not strictly always the case for two reasons which are often overlooked.

Firstly, the normative instructions for what are termed "Sea and Anchor Piloting Detail" (p. 42), which prescribes how the team should conduct itself, stipulates that *ten* personnel are required. When Vaesen (2011, p. 382) describes the navigation case study, he erroneously claims that Hutchins' states that there "is no central blueprint for successful navigation". This is wrong since Hutchins explicitly starts his analysis of the fix cycle by beginning with its normative description in the official document: "Navigation Department Watch Standing Procedure". During his fieldwork, a navigational officer handed it to him saying: "It's all in here […] You can read this and save yourself the trouble of standing watch" (1995a, pp. 26-27).

Charitably, one can interpret Vaesen's comments as indicating that not every possible natural contingency of the task is contained in this document. But equally one can argue that no blueprint has ever contained every physically possible outcome for any task – therefore making this an unreasonable request. Indeed, Hutchins notes that not every single possible aspect of the task is in the set of instructions because the construction of such a document would not be feasible and would also be unusable given both its impenetrable length and the high turnover of personnel in the team (1995a, pp. 26, 263). Additionally, the guidelines also state that where possible each position on the team, with the exception of the Navigator who is an officer, should be filled by a Quartermaster (something that is also not accomplished in practice). This document lists the entire composition of the team as follows (pp. 42, 178):

Table 2.2: Prescriptive guidelines for entire Navigation Team

i.	The Navigator as overseer (conning officer who additionally checks radar for other vessels in the vicinity)	vi.	**Port Pelorus Operator** with alidade
ii.	Assistant to the Navigator	vii.	Helmsman
iii.	**Plotter** on chart table, with hoey and other calculation devices	viii.	Quartermaster of the Watch
iv.	**Recording Bearer** with phone communication to pelorus operators, bearing record log, and watch (radar bearings if they are needed)	ix.	Emergency helmsman in the aft section with direct access to the rudder in case of loss of power
v.	**Starboard Pelorus Operator** with alidade	x.	Fathometer operator

Hutchins observes that the prescriptive organisation is not often met, and that the team were often undermanned (pp. 26-27, 222). Deranty notes that this is a common feature of real life working conditions, where there is often a gap between the prescription of work and the actualities of what must be achieved in order to make sure the task is completed (2009, pp. 79-80). As I have stated above, our main focus is only going to be on four of the members of this team: the plotter, the bearing recorder, and the two pelorus operators [iii, iv, v and vi]. With the exception of the

officer on deck and pilot, this is not because the others are not important but rather because the team is understaffed.

A second overlooked reason for why it is not always the case that four individuals are carrying out the fix cycle is more problematic for standard portrayals of distributed cognition. When the task is carried out in open waters (any stretch of water more than eight kilometres from land) the fix cycle can be done by a solitary agent because the time restrictions on the task are much reduced. For instance, Brown (2011), Cheon (2013, p. 24), Clark (1997, pp. 76-77), Huebner (2014, p. 153), Magnus (2007, p. 298) all overlook this important ethnographic detail. It is also noteworthy that Brown equivocates between describing it as a group activity taking place near a port (2011, p. 24) and as a necessary group activity at all times (pp. 19-20). I contend that attention to the particular details of case studies is important for understanding human cognition. As I will repeatedly show in later chapters; it is this attention to ethnographic details that has significant philosophical import. As such, Brown's ambivalence is detrimental.

One could argue in defence of this general misreading that the most philosophically appealing aspect of Hutchins' case study are the claims about supra-individual cognitive properties at the group level (e.g. see Rupert 2004, p. 391). However, I think that a narrow focus here on this one aspect is ultimately self-defeating. As I began to argue in chapter one, and will argue for at greater length in the coming chapters, a narrow focus on synchronic group activity overlooks a range of other ways in which cognition can be distributed – particularly diachronic factors. And it is the omission of these factors that leaves us unable to properly explain how cognition is distributed at all.

In confined waters the task must be completed within three minutes since more position fixes are required in order to manoeuvre more regularly, whereas in open waters only one position

fix is required per hour – although this becomes one every fifteen minutes when within sight or radar contact with land (Hutchins 1995a, p. 28). In the latter situation ("Normal Steaming Conditions"), the entire task of plotting a position fix and calculating the trajectory of the vessel is achieved by a solitary agent. As such, a deflationary critic could argue that it is only the context, specifically time constraints, which render the task a necessary group activity. As I will argue in chapter six, the differences between how the task is performed by an individual and by a team have significant impacts for how we understand this case study. Another deflationary point of attack is to articulate the relationship of the agent and artefacts so that the latter are mere inputs; subsequently reducing cognitive activity to the agent. These are important points to bear in mind when considering the navigation example and the various deflationary strategies that a methodological individualist might propose. In the next section I outline the case study before returning to these philosophical issues.

2.4. The particular roles of the navigation team

With the above caveats in mind, I now turn to the fix cycle in more detail. As I stated in section 2.2, the following presentation is in line with how it is often presented in the philosophical literature. This then provides us with a basis to add extra details that are made salient by my pluralistic framework.

The navigation team is composed of four key individuals – the plotter, bearing recorder, and two pelorus operators – who are engaged in this procedural activity. The task of ascertaining a position fix and calculating the ship's future trajectory is subdivided amongst this team who also use a range of cognitive artefacts. The major emphasis of the limited views is to focus on the team as a whole, but this also includes the "coordination dynamics" between the agents and the respective

cognitive artefacts they are manipulating[20]. For our purposes, the major subsets of the team that interest us are presented in the table 2.3 above.

Table 2.3: summary of the three key components of the navigation team and the various subtasks of the overall fix cycle that they perform

No.	Subsystem components	Tasks and processes
A.	{Plotter, hoey, chart, bearing record}	Computation and decision making
B.	{Bearing recorder, bearing record, phone, watch}	Communication, time-keeping, working memory, buffering, recording[21]
C.	{Pelorus operator, alidade, gyrocompass}	Observation

Here I have listed the four key individuals (there are two pelorus observers) with the cognitive artefacts that they use for their specific subcomponent of the overall task of the fix cycle. The whole task can be summarised as follows: The core aim of the team in the fix cycle is to attain a good position fix which is achieved by identifying three landmarks from the ship's projected location. These landmarks are then measured as having bearings relative to the ship which allows for the construction of "lines of position" on the chart. When three of these converge on the chart they form a triangle – the smaller the triangle, the better accuracy of the position fix. Once this is achieved the team can then project where the ship will be, *ceteris paribus* (e.g. if velocity and direction stay constant), in three minutes time and chose another set of three landmarks and begin the cycle again. I have presented the layout of this system in the figure 2.4 below.

[20] Here is an example of the nesting I mentioned briefly in the first chapter and to which I will return in the next chapters.

[21] In the literature this use of an external representation as a memory aide, following Donald (1993), is referred to as an "exogram" – i.e. an external memory resource in contradistinction to internal memories which are defined as "engrams".

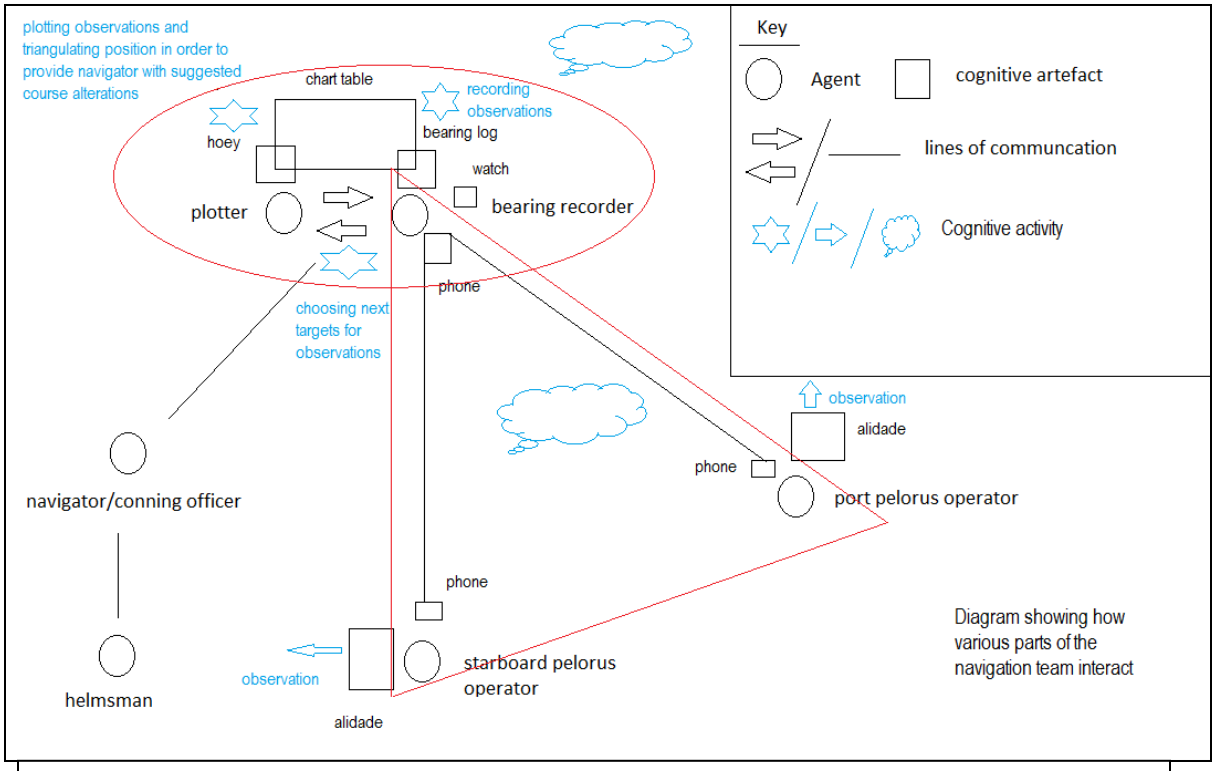

Figure 2.4: the interaction of the components in the navigation team. The plotter inside the pilothouse (decision making, computation) and the two pelorus operators outside with the alidades (observing) are connected through the bearing recorder (communication, memory).

I have listed the entire cycle without dissecting it too much; the table above catalogues the various subcomponents of the team and the specific subtasks they are engaged in during the fix cycle. I shall now discuss each of these in more detail in order of the complexity of the particular subtasks. Each of the agents in the team uses an array of cognitive artefacts to achieve their relevant sub-goal. Perhaps the simplest task is that of the two pelorus operators [C]. They use specialised observational devices called “alidades” (Hutchins 1995a, pp. 119-123). These can be understood as telescopes that have two external representational features built into the lens: [i] an aligned numbered card-insert overlaps the lens using a prism and is calibrated to a gyrocompass on the bridge of the ship; and [ii] a black line in the visual field which makes the combination more salient by superimposing this abstract representational frame of reference onto the physical image

in the telescope. This is what Hutchins (2005) calls a "conceptual blend", with the added notion that it is the stability of this external representation – what he calls a "material anchor" that enables this sophisticated and complicated representation to occur easily for the agent[22].

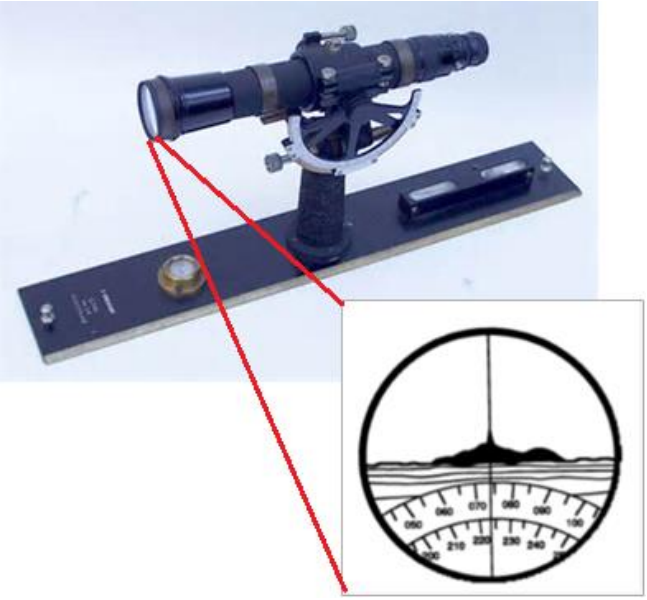

Figure 2.5: alidades (specialised telescopes) are operated by the pelorus operators to attain landmark bearings. The viewing lens is bisected by a line which aligns the visual field of the landmark with the conceptual schema of compass points.

[Diagram adapted from Hutchins 1995a, p. 120; and a stock image taken from: http://www.antiqbuyer.com/All_Archives/SURVEY/level-archive.htm accessed on 04/10/16.]

Many other theorists point to the stability of external representations as being a functionally different aspect to internal representations that offers alternative ways of doing cognitive work (see Charbonneau 2013; Donald 1993, 2010; Menary 2007b; Sutton 2006; Vorms 2012; Zhang 1997b)[23]. The gyrocompass is orientated towards true north as opposed to magnetic north – this is more useful for navigational purposes and also avoids the problems associated with the Earth's fluctuating magnetic field and issues arising from other disturbances in magnetic fields (e.g. the ship is made from a lot of metal). Having an inbuilt representation connected to the gyrocompass radically simplifies the task of the pelorus operator (and other members of the team). Their task is simply to find the appropriate target landmark and to report the reading of the instrument (as a three-digit numerical representation) at the correct time. Huebner notes that the

[22] Conceptual blends are a crucial for understanding how external representations operate in distributed cognitive systems. I discuss this in more detail in chapters five and six.
[23] I return to this topic in more detail in the next chapter.

digital representation produced in this process then acts as a “trading language” between the various subsections of the team (2014, p. 154).

The correct timing of the bearing report is orchestrated by the bearing recorder [B]. The bearing recorder plays an important role in connecting the pelorus operators who are outside the pilothouse with the plotter who is inside. He communicates to them over a phone-line (Hutchins 1995a, pp. 123-124). This has many advantages which are sometimes overlooked in the brief presentations given by some philosophers – indeed Brown (2011, p. 19-20) erroneously states that the pelorus operators communicate directly with the plotter. This misses out the crucial role that the bearing taker plays. If the system was set up as Brown describes this would lead to confusion as information would arrive from the pelorus operators when the plotter was busy with other tasks (such as another position fix or related calculation) and would lead to an increase in errors. The bridge of a naval vessel is an incredibly noisy place with much cross-chatter from the various departments engaged in the many tasks that are crucial for the vessel. Much of this cross-chatter involves discussion of numbers which increases the likelihood for errors given that they are similar to the three-digit bearing reports of the pelorus operators (see Hutchins 1995a, pp. 237-238 for an example of this). Instead the bearing taker performs the important function of “buffering” so that there is an asymmetric communication channel and separation between these other parts of the system so that they do not interact in adverse ways (pp. 194-195). This shall be spelled out in much greater detail in chapter six (section 6.2.1). For now, it is important to note that the bearing recorder is the member of the team responsible for relaying what landmarks the pelorus operators should be looking for and also when to report them. He does this by calling “mark” over the phone line. Timing is crucial for an accurate position fix. The bearing recorder notes down the coordinates that he receives into the bearing log.

The final member of the team is perhaps the most important: the plotter [A]. He takes the bearing reports from the bearing log and marks them on the chart as “lines of position” – these are lines drawn from the landmark to the hypothetical position of the ship based on the three-digit bearing reports. Where the three lines of position intersect is a representation of the location of the ship at that time. Hutchins describes the chart as an analogue computational device (1995a, p. 61). Vaesen has claimed that there is no “central unit orchestrating the cognitive endeavour” and that there is “no crew member who has internalized all relevant knowledge and skills” (2011, pp. 380-383). Likewise, Magnus has claimed: “That it would be wrong to think that the whole process is understood by someone” (2007, p. 298). Arguably both of these points are wrong given that the fix cycle is carried out in open waters by a solitary agent[24]. When I return to consider the vital importance of multiple timescales for understanding the navigation case study in chapters five and six it will become even clearer how the special set up of developmental trajectories make both of these spurious descriptions.

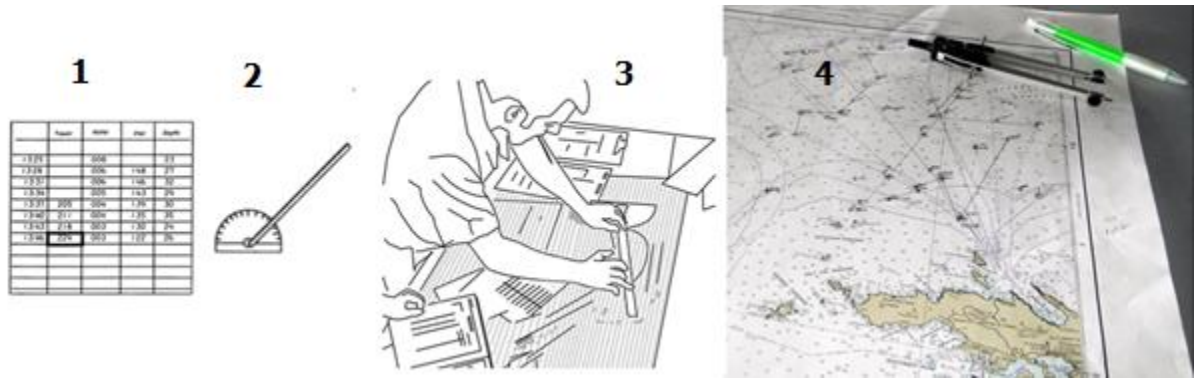

Figure 2.6: the propagation of representational states achieved by the plotter: [1] from the bearing record log to [2] the movable position of the hoey; this [3] is then aligned with the chart, and [4] finally creates the line of position. [Diagram adapted from Hutchins 1995a, p. 120; 2010b]

[24] As I noted in section 2.3; this is a difference between the individual and team levels, and how this impacts on how the task is performed, will be key point that I return to in chapter six (sections 6.2.1 and 6.2.2).

I would suggest that we can regard the interactions of the plotter and the chart (using other mediational tools) as a special subsystem that acts as a unit of control within the wider system. The plotter uses a range of computational devices to transfer the information in the bearing log into a representation of the ship's position (this involves altering the representational format from digital to analogue); and also, to aid in the decision-making process of where the ship will be for the next fix cycle in three minutes time – and thus what three landmarks will next be appropriate. A particularly useful tool for the plotter is the hoey (see figure 2.6 above) – this is a semi-circular protractor with a moveable arm that is used to calculate angles (bearing reports are in terms of 360-degrees) and also helps to reduce the amount and length of lines drawn on the chart to prevent it from becoming illegible (Hutchins 1995a, pp. 124-127).

2.5. The fix cycle and the propagation of representational states

We can now consider a more detailed exposition of the entire process of the fix cycle. I have portrayed this diagrammatically in figure 2.7 below. The plotter chooses three landmarks in order to triangulate the position of the vessel from the resultant lines of position. The bearing recorder relays these landmarks to the pelorus operators verbally across the phone network describing them *both* in terms of visual cues *and* their estimated direction as three-digit headings. The pelorus operators acquire their targets in the viewfinders of the alidades with their built-in gyroscopes. After three minutes the recorder calls "mark" and the three-digit headings are relayed and recorded in the bearing record log[25]. These values are then plotted on the chart by the plotter using the hoey to geometrically construct a position fix in the form of a triangle where the three lines of position intersect. In doing so the plotter transforms what is a digital representation (three-digit number)

[25] I will discuss why three minutes are chosen in more detail in chapter six when we consider the diachronic factors of the cultural-cognitive niche (section 6.4).

into an analogue representation (a physical position on the hoey which is then combined with the chart). This complex representation is then used by the plotter to make course suggestions to the conning officer (navigator), and to then begin the fix cycle again by choosing at least another three landmarks for the pelorus operators to target in the next three-minute operation cycle.

The core point of the fix cycle that philosophers (e.g. Giere 2006, p. 98; Huebner 2014, p. 154) focus on is that the overall task could not be completed by a solitary individual because of the time constraints in restricted manoeuvring in confined waters and because the ship is incredibly hard to manoeuvre: due to a bad turning circle and slow reaction times to control, etc. (Hutchins 1995a, p. 48). To be clear, the argument here against the methodological individualist is the claim that the fix cycle could not be done "in the head" of an individual. As Wilson puts it: for "individualism to be defensible as a view of cognition in this case, there must be an in-the-head account of all of the relevant cognitive processes" (2004, p. 178). The overall task is regarded as cognitive insofar as it involves the concerted efforts of a group of cognitive agents and how they use cognitive artefacts in working together to complete this task. As I outlined in chapter one, this description matches with the standard way that distributed cognition is discussed in the philosophical literature (section 1.3).

Of particular importance here is what Hutchins repeatedly refers to as the *propagation of representational states across various different states* (1995, pp. xvi, 65, 117-118, 131, 135, 154, 170, 190, 227, 290, 308-310, 373-374). Huebner identifies this latter feature as the core point for why this is a case of distributed cognition because "…the representation of [the] ship's location can *only* be produced through the interaction and coordination of a number of these distinct processes" (2014, p. 154, my emphasis). I think this is a crucial point. Although Huebner fails to recognise that the fix cycle can in fact be performed by a solitary agent, it is the propagation of representational

states across various representational media that forms the basis for the simpler view's claim that the methodological individualist cannot claim that all cognitive activity takes place solely within the confines of one agent's skull. And as such, that this is a legitimate case of distributed cognition.

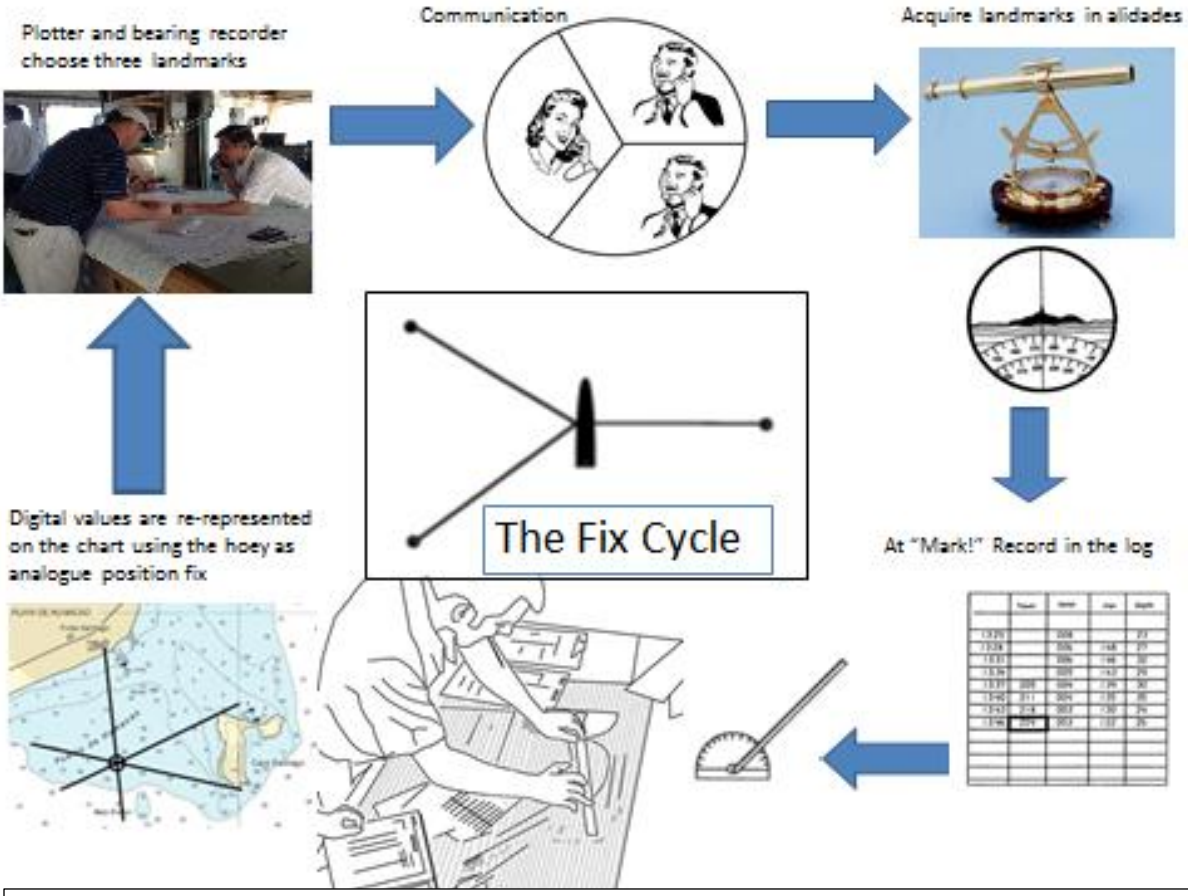

Figure 2.7: the fix cycle. Clockwise from top left panel: The plotter chooses three target landmarks. These are relayed to the two pelorus observers by the bearing recorder. The pelorus observers acquire the landmarks and these are relayed as three-digit numbers which are recorded in the bearing log book when the bearing recorder calls 'Mark'. These are reformatted by the plotter as analogue representations on the movable arm of the hoey, and blended with the chart so that the three measurements triangulate the position of the vessel. From here Future positions of the vessel are projected using dead reckoning begins the next cycle.

However, it is questionable whether this propagation of representational states, which is often conceived of in terms of information flow[26], can adequately demarcate what the distributed cognitive system is in this case. The crux here is whether there is an "information bottleneck" that distinguishes between processes 'inside' the putative cognitive system that are deemed to be of a higher bandwidth than processes connected to 'outside' the system. So, for example, in the team performance of the fix cycle one could claim that there is high informational transfer between the plotter and bearing recorder, who in turn acts as a buffer between the high information transfers from the two pelorus operators. In turn, these four individuals, and their relevant cognitive artefacts, can be seen as a distributed cognitive system because they have lower informational interfaces with the rest of the environment and other members of the crew on board the ship[27]. The issue here is whether this presents us with a sufficiently robust demarcation principle that can stave off the problems I outlined in the first chapter: particularly the problem of cognitive bloat. I shall argue, following Kaplan (2012), that information flow, by itself, cannot sufficiently bound the system and differentiate between genuine components and mere causal background conditions.

Firstly, I think that this approach is at fault for only focusing on synchronic cognitive processes whilst omitting how cognition can be distributed in time. In chapters five and six I show that this is an egregious mistake. But placing this issue to one side for the present purposes, I think there are multiple reasons for thinking that the information bottleneck is insufficient even on a solely synchronic time-scale. In the remainder of this chapter I explore these problems. But before I do so there is an important point of note: although I have gone into more detail than many of the

[26] I return to the question of 'what is information?' in the final section of this chapter. Advocates of this position often do not clarify what kind of information they are referring to. A notable exception is Cheon (2014, p. 30) who specifies that it is the "…carrying or propagating [of] natural or non-natural semantic information".

[27] Although Hutchins is not explicit on this point, Rupert (2004, p. 394) takes the following passage from *Cognition in the Wild* to be indicative of this view: "…the normally assumed boundaries of the individual are not the boundaries of the unit described by *steep gradients in the density of interaction among media*" (Hutchins 1995a, p. 157, my emphasis; also see Rupert 2011, pp. 636-637).

expositions in the philosophical literature, there is still far more to this example than has been covered here. Furthermore, these omissions which I shall discuss in subsequent chapters are actually crucial for properly cashing out why this is a case of distributed cognition. Insofar as these factors are omitted we will have an inadequate account of what distributed cognition is because, as I will show in the following chapters, the particular details of case studies are *not* mere ethnographic superfluities with no philosophical (or cognitive) bearing. They are in fact crucial for providing us with a properly spelt out understanding of what a distributed cognitive system *is* in this particular case and in general (also see Huebner 2014 for similar remarks on the importance of ethnographic details).

2.6. Cognitive ecology and the blind man's stick (BMS)

In discussing how we should demarcate a distributed cognitive system, Hutchins often states that we should do so where connectivity is low: e.g. following Bateson (1972) he states "…one should bound the unit so that things are not left inexplicable. This usually means putting boundaries on units where the traffic is low" (Hutchins 2001, p. 2068). However, he does not go into great detail on the specifics of how this achieved except to state that this is an empirical question linked to the context of the particular investigation (2010a, pp. 705-706). For instance:

> For distributed cognition, the existence of boundaries and centers are empirical questions. *Centers and boundaries are features that are determined by the relative density of information flow across a system*. Some systems have a clear center while other systems have multiple centers or no center at all. (2014, p. 37, emphasis added)

Although this certainly fits with Hutchins' broader commitments to a flexible unit of analysis, I think that his more extended discussions of information bottlenecks create difficulties for his position. It is important to disentangle the notion of information flow from these issues before assessing how it

fails to adequately demarcate the boundaries of a flexible unit of analysis. The basic problem is that information flow is connected to what Hutchins calls "cognitive ecology" (1995a, 2010a, 2010c). These issues relate to Hutchins' philosophical heritage in the cybernetic theorist Gregory Bateson's notion of "ecology of mind" or an "ecology of ideas" (1972, pp. xv). Hutchins' use of the term bears an obvious resemblance to this (but as we shall see below, there is an ambiguity in how Hutchins uses the term). It is also from Bateson that Hutchins gets his notion that cognitive processes are not bound by the skin (1972, pp. 319, 339, 454, 460, 483). But problematically, Hutchins also gets some more excessive metaphysical flourishes from Bateson, which I argue lead to unnecessary complications. It is useful to discuss these since it allows me to contrast my own more modest views on how we should understand the two key principles of distributed cognition (viz. the flexible unit of analysis, and the possibility of supra-individual cognitive processing).

In the context of information bottlenecks, following Bateson, Hutchins (2010a, pp. 705-706) claims that his approach to bounding units of analysis where connectivity is relatively low is to give a "cybernetic twist" to Plato's call that we "carve nature at the joints" (Phaedrus 265d–266a). This is a multifaceted claim that makes trying to understand Hutchins' position confusing. On the one hand, claims about carving nature at the joints tacitly implies that one can make viable claims about nature. Thus, again seemingly contradicting his claims elsewhere (e.g. Hutchins 2014), which deny that distributed cognition makes claims about nature. On the other hand, Hutchins also acknowledges that what appears to be an area of low connectivity according to one theoretical approach, may appear as a region of high connectivity from another theoretical stance (2010a, p. 706). Additionally, in briefly considering the question about which position to adopt, Hutchins states that the issue is not about which one is true but rather which approach "gives us the best scientific leverage" (p. 707). One could argue that this again exposes Hutchins to the charge of 'over-instrumentalism' levied by Hutto and colleagues (2014). This latter issue is something that I will seek

to avoid in my own approach to these questions by providing precise conditions for how the unit of analysis should be demarcated by providing more objective grounds by which to determine them.

In attempting to outline what the distributed unit of cognitive analysis is in more detail Hutchins repeatedly references an extended passage of Bateson's analogy of the 'Blind Man's Stick' (BMS) (see Hutchins 1995a, pp. 291-293; 2010a, p. 706):

> Suppose I am a blind man, and I use a stick. I go tap, tap, tap. Where do *I* start? Is my mental system bounded at the handle of the stick? Is it bounded by my skin? Does it start halfway up the stick? But these are nonsense questions. The stick is a pathway along which transforms of difference are being transmitted. *The way to delineate the system is to draw the limiting line in such a way that you do not cut any of these pathways in ways which leave things inexplicable.* If what you are trying to explain is a given piece of behavior, such as the locomotion of the blind man, then for this purpose, you will need the street, the stick, the man, the street, the stick, and so on, round and round. (Bateson 1972, p. 459. First emphasis in the original, second added)[28]

Firstly, it is interesting to note that in Hutchins' (1995a, pp. 291-293) earlier citation of this passage he leaves off the last sentence – and I think this has implications for how one understands what the point is here. Without the last sentence the claim is simply that we should not exclude important explanatory factors in explaining the behaviour of the cognitive system. This is a claim that needs defending, but one that fits with the earlier points about the flexible unit of analysis – now partially fleshed out a little more as involving "pathways" (i.e. information channels). However, the last

[28] This is a recurring analogy for Bateson who discusses it on many occasions (1972, pp. 251, 318, 458). It is also a common example in the cognitive science literature. For instance, Clark (2008, p. 31) links this example to the seminal phenomenologist of embodiment Merleau-Ponty, as well as a range of more contemporary theorists with regards to how tool-use alters peripersonal embodied space (e.g. see Iriki & Taoka 2012). Additionally, the BMS has also been taken up by the cognitive archaeologist Lambros Malafouris – from whom I take the acronym BMS. Malafouris (2008, 2010, 2013) uses the analogy to emphasise that cognitive tools have not received the credit they are due in human cognitive behaviour – he refers to this as "epistemic neglect" (2010, p. 15) – both today and in the past. The notion that we should consider distributed cognitive processing as not just a modern phenomenon but something with a long history is an important point that I shall refer to again when discussing the importance of time in later chapters (also see Jeffares 2010; Johannsen 2010; Sterelny 2003, 2012; Sutton 2008).

sentence gives us a prescription for what these flexible boundaries *should* include in this seemingly simple example – and they quickly spiral outwards quite dramatically.

Here we see the problem of cognitive bloat in action, because trying to define the boundaries of the BMS in terms of information flow fails to differentiate between features that are potentially genuine components (e.g. the stick) from mere causal background features (e.g. the texture of the walking surface). Thus, there is a blending of the target system and its general environment that is unpalatable from both a theoretical and practical point of view (see Eliasmith 2009, p. 150 for related points). The practical infeasibility of this unit of analysis would seem to make the fears about methodological bloat, alluded to in the first chapter by Tooby and Cosmides in their pithy story about mirrors, to be coming true after all (see section 1.5). Since, in order to understand the movement of the blind man and his use of the stick we must also take into account the materiality of the whole street. One could perhaps defend Bateson (and a fortiori Hutchins) by noting that it is ambiguous whether the BMS case involves a claim about what constitutes the cognitive system. However, *it is this very ambiguity itself* (about what is a part of the system and what is simply a causal background feature), which invites the dangers of cognitive bloating. Hence the necessity for adequate demarcation criteria – in chapter four I return to the BMS example and show how Kaplan's (2012) work on the mutual manipulability criterion can resolve this issue.

I think these are only the beginnings of the worries here with how Hutchins, influenced by Bateson, is characterising information bottlenecking. Because although Hutchins states that a key principle of distributed cognition is a flexible unit of analysis (which we are currently trying to make more concrete through information bottlenecking), when it is approached from a 'Batesonian' perspective it leads to a position Hutchins refers to as "cognitive ecology". Hutchins uses this term throughout his writings but it seems to take on a different form here. In his earlier writings it refers

to the intertwining construction of the social and epistemic environment of an agent – e.g. the tools and techniques, and their mutual interdependencies – and how these constrain and shape the cognitive behaviour of the agent:

> In attempting to understand the history of navigation from a cognitive perspective, it is important to consider the whole suite of instruments that are used together in doing the task. The tools of navigation share with one another a rich network of mutual computational and representational dependencies. Each plays a role in the computational environments of the others, providing the raw materials of computation or consuming the products of it. In the ecology of tools, based on the flow of computational products, each tool creates the environment for others. (1995a, pp. 113-114. Also see Malafouris 2013, Preston 1988, pp. 515-528; Sutton 2008, 2010).

Although this idea seems to survive on into later writings – e.g. "By cognitive ecology I mean that all the elements and relations potentially interact with one another and that each is part of the environment for all the others [which can] give rise to powerful cognitive processes" (Hutchins 2010c, p. 99) – it is also contradicted by other statements that imply a much more metaphysically-loaded and radical position which cannot be made consistent with a flexible unit of analysis. I think that this more radical view inadvertently results in Hutchins postulating a new a priori unit of analysis – the "cognitive ecosystem" – and thus a new radical new theory of mind, viz. "cognitive ecology" (2010a, 2014). In concluding his paper on cognitive ecology Hutchins writes:

> My hopeful prediction is that the reality of the rich interconnectivity of the brain, body, and world will draw the many strands described here together into a coherent approach to mind as a property of cognitive ecosystems. (2010a, p. 211)[29]

I agree with Haugeland who states that a core goal of investigations into cognitive systems is to be able to identify the principles by which we can "divide systems into distinct subsystems along

[29] A subsequent paper attempts to offer some "conjectures" as to the features of cognitive ecosystems (2014).

nonarbitrary lines" (1998, p. 211. My emphasis; also see Clark 2008, pp. 31-32). As such, I contend that the cognitive ecosystem approach smothers and smears distinctions that are important: viz. between genuine components of the cognition system and the causal milieu. This blending of boundaries with what Hutchins refers to as the cognitive ecosystem goes too far in eradicating useful distinctions and additionally, as shown in the BMS example above, fails to distinguish between components of a system from necessary background features. This is particularly peculiar because an exemplary feature of Hutchins' cognitive ethnographic studies is their attention to these details (viz. components and their interrelations) whereas these are often what are missing in some of the overly brisk philosophical accounts of distributed cognition. For instance, Giere's accounts of the Hubble telescope (2006, 2007, 2012) and Magnus' account of double-blind trials (2007) have notably very little ethnographic details[30].

Even if we were tempted to agree with Hutchins' 'Batesonian' inspired desire to have a larger unit of analysis in which all human cognition is *always* understood in terms of its cognitive ecosystem (what Moore & Rocklin 1998 call the 'social only' view), then this would entail reasserting an a priori unit of analysis. And whereas methodological individualism proposes the isolatable individual or brain, and the majority of the philosophical discussion of distributed cognition proposes the specialised synchronic group; Hutchins' view inadvertently results in proposing the cognitive ecosystem. Since this contradicts the primary key principle of distributed cognition – a flexible unit of analysis determined by context – I shall reject Hutchins' move here as going too far.

Arguably, given the rest of his oeuvre (especially his practical work), this is not Hutchins' intention. However, it is the unfortunate and inescapable outcome of his lack of conceptual clarity and philosophical rigour. So, whilst agreeing that understanding how a particular cognitive system is

[30] I reconsider these examples and issues attendant in them in chapter four (section 4.3).

embedded in its environment is crucial (in some, possibly many, cases), I shall not follow Hutchins' 'Batesonian' call for a new theory of mind centred on the cognitive ecosystem (2010a, p. 712). Instead, I shall stick to searching for a principled way in which to articulate how cognitive processing can be distributed across both temporal and spatial domains. In doing so I shall agree with Haugeland and Clark that we should be able to dissect these systems into subsystems in a non-arbitrary manner.

Before returning to the discussion of information bottlenecks, it is noteworthy that whilst I have rejected the more extreme notion of cognitive ecology, the ambivalent way in which Hutchins uses this term entails that it would be imprudent to throw out everything associated with this term. For instance, in following chapters I will be discussing the importance of how social and cognitive dependencies can become entwined (Hutchins 1995a, pp. 200-203). But whereas Hutchins refers to this as part of a cognitive ecology I shall prefer the term cultural-cognitive niche (following the work of Menary 2010e, 2013a, 2014; Sterelny 2003, 2012; and influenced by Boyd et al 2011).

This term is preferred for several reasons. Firstly, ecology is too passive – human agents' interactions with their epistemic, physical, and cultural environments are incredibly active (Menary 2007a, 2010a). Secondly, ecology already has a specific meaning in the scientific lexicon with its basic functional unit being the ecosystem (Pears 1985). With the growing emphasis on the importance of interdisciplinarity given the complexity of the theoretical and practical problems faced by science and philosophy (see Andersen 2016), this kind of 'concept mining' – in which terms from one field with specific meanings are appropriated and then redeployed with an entirely different meaning – is rife for misunderstandings. As Downey has recently pointed out, when engaging in an interdisciplinary topic (such as distributed cognition) one should be wary of reinventing the wheel and generating needless neologisms without consulting the other fields since

this “the proliferation of specialised terminology” can result in isolating disciplines from each other through worsening conceptual language barriers (2014b, p. 116).

On this basis it is notable that the notion of cultural-cognitive niche construction has the advantage of being much more selective on what is designated as important: viz. aspects of the environment that cognitive agents are actively engaged in, with, and through towards cognitive tasks (Menary 2013a, pp. 27-29). Furthermore, the notion of niche construction has the additional advantage of (partially) being developed explicitly to discuss the role of cultural learning in evolution – hence this has the potential to facilitate a better dialogue between biology, the human sciences, and humanities through giving a more robust and non-eliminative account of what the cultural environment is (Gauvain 2000; Laland et al 2000; Sterelny 2003, p. 171). This point is additionally supported by recent work by Kendal (2011) and colleagues (Flynn et al 2011) who discuss the fruitful compatibility of distributed cognition and niche construction for exploring human development. I outline cultural-cognitive niches in greater detail in chapters five and six, but now return to the topic of information flow and how it is caught between two inadequacies: cognitive bloat and internalism.

2.7. Critiquing the notion of information Bottlenecks as a criterion for demarcating the flexible unit of analysis

Having unpacked the notion of an information bottleneck from its ‘Batesonian’ metaphysical baggage, we can now turn to questions surrounding its inadequacy for discerning system boundaries and inability to differentiate genuine system features from mere necessary background conditions. As I mentioned above, beyond a few sparse remarks Hutchins does not really develop what he means by low-informational-connectivity. However, there are several others in the literature who propose similar positions. For instance, Kaplan attacks a similar approach which he labels the “Simon-Haugeland bandwidth criterion” (2012, pp. 553-556). For our purposes we need only focus

on how Haugeland (1998) builds from Simon (1996). Haugeland defines a system as a ‘‘relatively independent and self-contained composites of components interacting at interfaces’’ where an interface is defined as “a point of interactive contact” between systems (or components – i.e. subsystems) in which the interactions are ‘‘well-defined, reliable, and relatively simple’’ (1998, 215). On this basis, Kaplan takes the bandwidth criterion to be

> …an arbitrary part or subsystem counts as a genuine component when the causal interactions within that part are appreciably greater than interactions between it and other parts of the system. (2012, p. 553)

Haugeland then goes on to claim that that there is a high-bandwidth connection between agent and world in acting and sensing and so there is no real interface separating agent and world. Instead, he proposes that there is an “intimate intermingling of mind, body and world” (1998, p. 224).

A second element differentiating a low-bandwidth interface from a genuine component is that the former are replaceable by functional equivalents (Haugeland 1998, p. 224). The point being that if the connections between parts of the brain and body were genuinely high-bandwidth then the substitution of motor signals from brain into an alternative – such as features of the external environment or prosthetic limbs – would be impossible in principle. Haugeland argues that there are specifics of how the brain-body interacts entailing that simple functional replacement is not possible. Since there are no real interfaces this brings the traditional boundaries of brain, body and the environment into question. In reviewing Haugeland’s argument, Kaplan summarises this nicely:

> Because no principled division of the brain-body-environment loop into components interacting at narrow-bandwidth interfaces is feasible, the brain cannot be the component exclusively or distinctively responsible for supporting cognition – since it is not a discernible component in the first place. According to the bandwidth criterion, then, cognition extends [– i.e. is distributed]. (2012, p. 554)

However, there are several problems with such an argument. Firstly, this is only an argument in favour of strongly embodied cognition. This is not to criticise embodied positions for not going far enough, but is rather to point out that there is nothing here in Haugeland's argument that necessitates a strong embedded/extended cognition thesis in which cognition is distributed beyond the body[31].

Furthermore, there are good empirical reasons for not even accepting this more modest version of argument as just advocating strong embodiment. Against Haugeland's claim that motor-signals could not in-principle be routed into a functional equivalent of a part of the body, it is extremely noteworthy that a team of neurobiologists have successfully designed a brain-machine interface that allows chimpanzees to control a multi-jointed prosthetic arm with cortical signals (Velliste et al 2008). The aim of this team, like many others, is to enable amputees and paralysed persons to regain lost functions. With an abundance of funding for these types of projects coming from the US military Defence Agency Research Projects Agency (DARPA), many teams have made great strides towards this goal in recent years (e.g. see Burck et al 2011) – to the extent that some amputees are now able to use these devices to perform such fine-grained motor-control actions as picking up a grape without crushing it, and catching a silk handkerchief as it falls[32]. Commentating on these experiments Kaplan writes: "replaceability of the limb with a functional equivalent prosthetic limb is not just in principle but in practice possible" (2012, p. 555). But importantly, Malafouris adds: that as impressive as this new research is, we must recognise it as only the most recent chapter in a much longer history which goes back as far as the manufacture of the first stone tools (2008, p. 404)

[31] The terms strong embodiment and strong embeddedness are taken from Menary's (2010b) useful hierarchical distinction of these positions. He distinguishes between weak, moderate and strong versions of each of these theses. Weak versions of embodiment and embeddedness correspond to standard materialism in the philosophy of mind (e.g. Adams & Aizawa 2001, 2010). Moderate versions correspond to situated and scaffolded views (e.g. Sterelny 2010). A strong version corresponds to enactive (e.g. Hutto et al 2014) and extended positions (e.g. Clark 2008).

[32] Videos of this are quite astonishing: e.g. see https://www.youtube.com/watch?v=suwZ5D9bk0M and https://www.youtube.com/watch?v=suwZ5D9bk0M (accessed on 09/08/2016).

– i.e. distributed cognition is not merely a modern phenomenon and has a long evolutionary history (also see 2010; Donald 1993, 2010; Jeffares 2010; Johansen 2010).

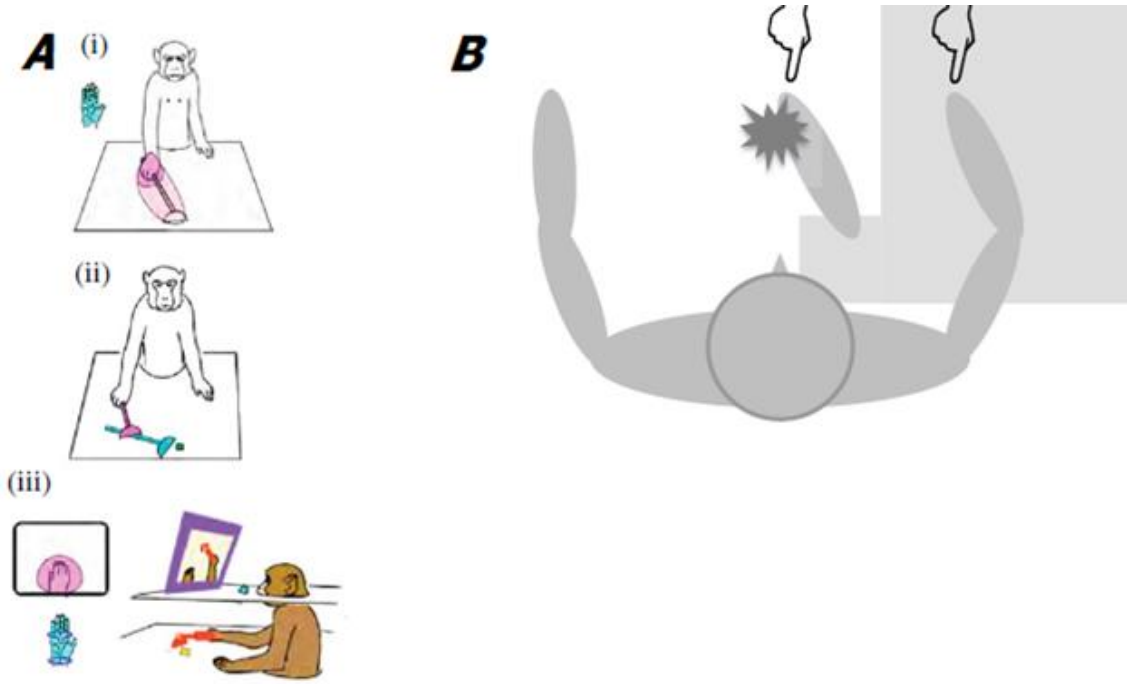

Figure 2.8: Panel (A) shows three increasingly complex ways in which Macaque monkeys are taught to use tools to obtain food that is out of reach: [i] standard use of a rake; [ii] use of one rake to reach another; and [iii] use of a rake mediated through a video system. Highlighted regions show the plasticity of the body schema to incorporate relevant features of the environment as literal extensions of the body. Panel (B) shows a human agent duped into experiencing a rubber hand as part of their body by the synchronous stimulation whilst their real hand is visually occluded. [Diagrams from Iriki & Taoka (2012, p. 12) and Hohwy (2013, p. 105)]

Another related line of evidence that objects to Haugeland's point is the work by Iriki and Taoka (2012) focusing on the functional and anatomical plastic changes induced in the brain through the acquisition of tool-use in primates (see panel A in figure 2.8 above). A core part of Iriki and Taoka's claim is that the body schema of the primate is altered through the acquisition of tool-use – as evidenced by the alteration of the receptive field for the visual system to incorporate the tool and not the hand as the limit of peripersonal space. They go on to claim that this "redeployment"

(Anderson 2010) of neural resources is only possible because of a redundancy or lassitude in which neurons are "bistable or polysemous for the hand or the tool", and that this "accidentally established equivalence between body parts and tools" is crucial for tool use (Iriki & Taoka 2012, p. 14). Importantly, the alteration of these neurons correlated with the body schema occurs only through the "….active, intentional usage of the tool, not its mere grasping by the hand" (Maravita & Iriki 2004, p. 81. Also see Malafouris 2010; Menary 2013b). This plastic redundancy is important because organisms not only grow but can also undergo injury. If the representational system was unresponsive to these possible changes it would be highly un-adaptive. As such, tool-use is the "cultural recycling" (Dehaene & Cohen 2007) of these cortical structures through the learning associated with the acquisition of cultural practices (Menary 2013b, 2014). Such a possibility is occluded by Haugeland's stance because although he allows for the high information flow between brain, body, and world; this view fails to discriminate adequately between mere interactions with the world and the functional incorporation of environmental resources through redeployment. This shows how the problem of cognitive bloat, and the failure to recognise important distinctions, diminishes our understanding of cognitive phenomena.

A similar finding about the redeployment of the equilibrium system to the tongue of agents who have lost their sense of balance through inner ear damage is yet another example of this (Doidge 2007, pp. 1-26). Of particular note is the fact that this redeployment was dependent on the plasticity of the body and nervous system, and also the use of a mediational prosthetic device in a training phase that scaffolded this extreme change (Downey 2012, p. 182). A final piece of supportive evidence for this line of thinking would be the rubber hand illusion by which an agent can be tricked into experiencing an inanimate rubber hand as part of their body through synchronous touching of a rubber hand in the visual field and the agent's real hand which is occluded from view (see Hohwy 2013, pp. 104-106 for details; and see panel B in figure 2.8 above). This again shows the

elasticity of the body schema and the flexible boundary between what is the body and what is not, which Haugeland's position is insufficiently equipped to account for.

A more pressing concern of a general nature is that interfaces need not be low-bandwidth as Haugeland proscribes. Clark (2008, pp. 32-33) highlights that distributed computing networks with multiple individual computers have extremely high bandwidth interfaces and yet we discern them as separate components of the system without conceptual or practical difficulty. Interestingly, Eliasmith relies on an information-bandwidth argument to make exactly the converse argument that Haugeland makes. Eliasmith dismisses positions which argue for a blurring of the boundary between the internal and external – something which he thinks is misguided (2009, pp. 148, 150). He points out that it is "obvious" that interactions between the brain and the environment are much slower than intra-brain interactions[33]. He criticises Clark and Chalmers (1998), amongst others, for overlooking this "huge difference" in kinds of coupling since he claims that it makes their position inconsistent.

For instance, this argument would rule out Hutchins' navigation team as a case of distributed cognition because the internal flow of information in each agent's brain seals them off from the much lower propagation of representations between each other and their tools. Furthermore, it demonstrates that there is something other than the skin and skull that is different between brain-brain and brain-world interactions. I think there might be a response that could be made here in regards to expert tool-use discussed above, insofar as expansions of peripersonal space and "perceptuomotor integration" and phenomenological "fluency" are correlated (Iriki & Taoka 2012; Nemirovsky et al 2013). I return to this issue in the following chapters in regards to enskillment and developmental trajectories. But since I do not accept information bottlenecks as

[33] In a similar fashion, Kirchhoff (2016, p. 5) points out that internalists point to the "causal density" of the brain.

adequate for demarcating the boundaries of the cognitive unit of analysis this is not an issue here for my own position.

Despite these criticisms, Eliasmith adds that he does not think that determining the appropriate system boundary is an easy task, nor that it definitely stops at the skin (just that information bottlenecking does not support distributed cognition). He concludes:

> …it is unclear whether we will be able to identify general, consistent principles for identifying system boundaries. Nevertheless, it is essential to realize that this is a task worth pursuing, and that simply blurring systems over boundaries, or suggesting that such boundaries do not really exist, is bad for both practical (i.e., trying to do science) and theoretical (i.e., appropriate conceptual application) reasons. (2009, p. 150)

This is really important for understanding how we should articulate the flexible unit of analysis so that it avoids the pitfalls of Hutchins' inadvertently radical eliminative cognitive ecology. We must take heed of these issues if distributed cognition is to be a useful framework both practically and theoretically.

According to Kaplan the central problem of the bandwidth criterion is its inability to discern between which causal relations of a system are the most explanatorily salient: "What goes wrong with the bandwidth criterion is its failure to identify, among the multitude of causal relationships in a given system, which are, and which are not, explanatorily relevant to the phenomenon to be explained" (2012, p. 556). Information bottlenecks cannot properly tackle cognitive bloat because mere causally connected features are needlessly incorporated into the system in a spurious manner under this condition. I also doubt that trying to refine the manner in which information flows – e.g. Heersmink (2015) differentiates one-way, two-way, and reciprocal in a hierarchy from shallow to deep integration (see Davies & Michaelian 2016 for a defence of one-way information flow) – can surmount the challenges I have identified here. Because *if information flow is defined too*

permissively then cognitive bloat looms, and *if it is defined in a more restrictive manner then it supports internalism*. As such, the position is trapped between two differing issues, so even if it is able to escape one prong of the argument it is unlikely to be able to avoid the other.

Despite this failure, I think that we should be wary of throwing out the whole idea. Part of the issue with the 'Batesonian' approach discussed above is that it occluded important divisions we might want to make in discerning parts and components of a system and their interrelations. Arguably, information bandwidths and information processing in general are useful theoretical tools for enquiring into the composition of distributed cognitive systems (e.g. see Cheon 2014; Davies & Michaelian 2016; Kirsh 2010; Heersmink 2015) – and are well grounded according to our naturalistic approach since they are a cornerstone of the interdisciplinary project that is cognitive science (Miller 1956; Simon 1996). Of course, this is not to deny that there are a range of philosophical issues with this – such as: 'what is information?' (Adriaans 2012; Crnkovic & Hofkirchner 2011); 'what kind of information is cognitive?' (Crnkovic & Hofkirchner 2011; de Witt et al 2016); and, most crucially, 'how do we demarcate this information that is supposedly cognitive from claims that information is everywhere?' Or, as Searle incisively puts it: the bare notion of information does not enable us to distinguish the mind from thermostats and livers (1980, p. 418). Since I am not claiming that information processing can provide us with a complete theoretical set of tools needed to demarcate distributed cognitive systems adequately, these are not an issue for the position I put forward here. Instead, they are noted merely to highlight the fact that mainstream views in cognitive science have equally deep philosophical questions at the basis of their projects, and as such this is not a specific burden solely for the proponent of distributed cognition.

2.8. Summary

In this chapter I have outlined the most famous example of distributed cognition – Hutchins' navigation team case study – as it most often appears in the philosophical literature (with the caveat that my attention to the ethnographic details is more thorough). My aim here involved two interrelated points: firstly, to highlight the shortcomings of the standard approach to this case study which focuses solely on the synchronic group. This is cashed out in terms of information flow (the propagation of representational states). However, this is neither an adequate demarcation principle (because it cannot tackle the problems of cognitive bloat and generates a host of other issues), nor a legitimate motivating claim for distributed cognition in the first place. Additionally, Hutchins' proposals to appropriately demarcate the flexible unit of analysis fail because they inadvertently lead to undermining this claim and repostulating an a priori unit of analysis (i.e. the cognitive ecosystem).

These failures lead on to my second aim: foregrounding what my position must achieve in order to be viable (in conjunction with the problems outlined in chapter one). In the rest of the thesis I now go into greater detail on how a range of important explanatory details have been overlooked – it is these omissions that prevent the limited approach from properly articulating how and why this is a case of distributed cognition. Specifically, through articulating the importance of task-specificity (chapter three), the mutual manipulability criterion (chapter four), and normative patterned practices (chapter five), I will draw out interesting insights in the navigation case study which are overlooked by both the limited view and methodological individualism. This shows the practical benefits of my position. And thus provides a performative demonstration of how my view operates in a coherent fashion with regards to the most famous case study of distributed cognition (consequently blunting the challenge of methodological bloat).

3: Task-specificity and Multilevel Analysis

…trying to understand perception by studying only neurons is like trying to understand bird flight by studying only feathers.

David Marr (1982/2010, p. 27)

3.1. Introduction

In the first chapter I outlined the range of interconnected problems facing a proponent of distributed cognition: [1] the extreme heterogeneity of purported cases in the literature which differ greatly in terms of composition, size, and complexity; [2] how this exacerbates the problem of cognitive bloat – since this state of affairs makes it even more difficult to demarcate the boundaries of what is a genuine part of a distributed cognitive system; and [3] the problem of cognition distributed in time – how the importance of various diachronic temporal scales for understanding the synchronic scale are overlooked. These issues problematize Hutchins' (2001) emphasis on a flexible unit of analysis with regards to how cognition is distributed spatially and temporally in a principled manner (an overlooked core feature of the distributed cognition position). To tackle these many issues, I proposed a novel pluralistic approach. The case for this position was further motivated in the second chapter by demonstrating the failure of the standard simpler approaches which attempt to use information flow to resolve these issues. The simple approach failed to provide a sufficiently well-motivated argument in favour of distributed cognition and to challenge methodological internalism. Viz. if one claims that there are supra-individual cognitive properties, then it is beholden upon them to specify what these are in a principled manner.

In this chapter, I begin to outline my pluralistic approach by exploring the intricacies and issues surrounding how to demarcate the flexible unit of analysis with regards to the multilevel analysis that it entails. These issues include scaling up and scaling down, clearly specifying what we mean by levels, and how to avoid the carrying across of inappropriate conceptual baggage from one level to another. Building on work by Davies and Michaelian (2016), I propose a task-specific approach combined with Caporael's (2014) notion of "core configurations" – distinct emergent organisations of agents (and artefacts) focused on a shared task – to explicate supra-individual levels of the unit of analysis in a clear manner.

The chapter is structured as follows: in the next section (3.2) I outline Davies and Michaelian's (2016) dual-problem of scaling up and scaling down. Although I disagree with part of how this is formulated I think that they are right to criticise the shortcomings of an agent-centric approach because it leads to a range of other problems: viz. confusions surrounding how agents and artefacts interact; and the inappropriate transposing of conceptual baggage from one level of analysis onto another (section 3.3). Instead, in the final section of the chapter (3.4), I supplement Davies & Michaelian's task-specific approach with an analysis of the concepts of *levels* (Craver 2007), *emergence* (Theiner et al 2010; Wimsatt 1986), *task-type* (Steiner 1972), and *core configurations* (Caporael 2014) in order to provide us with a means of treating distinct supra-individual levels on their own terms. Task-specificity then acts in the pluralistic framework as an opening gambit from which to tackle the further issues of cognitive bloat and the problem of time, to which I will return in subsequent chapters.

3.2. Scaling up and scaling down

In discussing supra-individual cognitive properties, Hutchins stresses the importance of exploring the question of whether cognitive properties are different or similar at group and individual levels (1995a, pp. 176-178). Indeed, he goes on to postulate that "…organized groups may have cognitive properties that differ from those of the individuals who constitute the group" (p. 228). This involves the notion of "emergent" cognitive properties (e.g. Poirier & Chicoisne 2006; Theiner & O'Connor 2010; Theiner et al 2010), to which I will return in the final section of this chapter. But the first point to note is that this involves the flexible unit of analysis: viz. such an investigation not only involves multiple levels of analysis; but also, a framework in which to understand these multiple levels of analysis *on their own terms*.

My goal in this chapter is to explore the topic of multilevel analysis and its attendant issues. Although this is overlooked by some in the literature, Davies and Michaelian (2016) rightly focus on the multiple levels of analysis as a hallmark of distributed cognition. I.e. the focus of an investigation of cognition is not just the individual, but can also involve processing taking place at both the supra-personal and sub-personal level. Therefore, a theoretical position must be able to attend to cognitive processing on multiple levels if it is to have adequate scope to account for cognition. Davies and Michaelian label the issues that other approaches have in negotiating these multiple levels of analysis as the dual-problem of "scaling up" and "scaling down" (pp. 310-311). They go on to note that sub-personal processing is often overlooked in debates about distributed cognition which tend instead to focus on the more controversial notion of supra-individual cognitive processing (ibid). I agree entirely with this point. However, they go on to argue that distributed cognition is preferable to both traditional internalism (methodological individualism) and extended cognition because these positions cannot scale down. Given that much of cognitive science, especially neuroscience, focuses on specific sub-personal cognitive processes this *would* render both

methodological individualism and extended cognition erroneous according to naturalistic credentials – i.e. (as discussed in the general introduction) a philosophical position should be informed by the best available current science (Ladyman & Ross 2007; Stanford 2016). However, whilst I do think that distributed cognition is superior to both these other positions because of its "ontological flexibility" (Davies & Michaelian 2016, p. 318) in being able to handle multiple levels of analysis; Davies and Michaelian have not phrased this point correctly with regards to their criticism of extended cognition and methodological individualism. As such, I will first untangle this point in order to protect the crucial insight: viz. although methodological individualism and extended cognition do discuss sub-personal cognitive processing; their agent-centric stances entail a range of problems related to carrying-over unsuitable conceptual baggage when trying to analyse supra-individual levels.

Davies and Michaelian's claim that neither methodological individualism nor extended cognition properly consider processes occurring below the level of the individual as a "unified whole"; instead they are "black boxed" (p. 311). They use the example of planning as a particular cognitive task to make their point:

> The intracranialist view [methodological individualism] can do no better than to say that planning, for example, is something done by a cognitive agent, operating on his own. The extended cognition view can do no better than to say that planning is something done by a cognitive agent and which sometimes does and sometimes does not extend into external resources. (ibid)

We can query the validity of this claim by taking an archetypal proponent of each position and seeing what they say about planning in terms of sub-personal cognitive processing. For instance, the most famous exponent of extended cognition, Clark, has recently discussed planning in terms of "predictive processing" (2016). This functional process is instantiated in cortical structures in the

frontal lobes and potentially in coordination with a range of external resources such as diaries, timetables, etc. Clark also postulates that this relationship is mediated by "prediction error minimisation"[34].

For traditional internalist positions, one possible account of planning at the sub-personal level is Fodor's (1983) famous account of mental modules. This holds that planning occurs in a domain-specific module, taking place entirely within the skull. A module need not be an anatomically localised region but they do have a range of functional properties that curtail how they are structurally instantiated in the brain – e.g. informationally encapsulated and fast, etc. Thus, the neural region cannot be too distributed (although these criteria are now also a matter of debate; see Anderson 2010; Carruthers 2006 for discussion).

My purpose in briefly outlining these two positions is not to assess their merit, but rather to make it clear that both methodological individualism and extended cognition *can and do* discuss sub-individual cognitive processing. As such, I think that this incidental aspect of Davies and Michaelian's criticism is misplaced. But there is also a further complication here. Although Hutchins certainly discusses sub-individual cognitive processing, he does so in quite an ambiguous manner which makes it difficult to understand what distributed cognition contributes to these discussions. Given that Hutchins has claimed that opposing positions such as parallel distributed processing (2001) and massive modularity (2014) are both instances of distributed cognition, it is debatable whether distributed cognition can make a meaningful contribution here[35].

---

[34] Clark has recently advocated that predictive processing is the best framework for articulating embodied, embedded, and extended cognition. This is because error minimisation acts as a unifying principle that brings action and perception together. Furthermore, Clark (2016) explains how he sees this framework as drawing together all the various interests he has over the years into a coherent account. Conversely, see Hohwy (2013) for an alternative internalist account of predictive processing.

[35] For details on how parallel distributed processing and massive modularity are opposing theoretical positions see Anderson (2010).

As such, and for the sake of simplicity, in this thesis I will for the most part be focusing on how distributed cognition "scales up" from the individual. But this does not mean I am entirely eschewing the neuronal or sub-personal level. When I come to discuss the importance of cognition distributed in time, in chapter five, the interplay of the various spatial levels – from neurons to niches – will become clearer. My present aim is to provide a framework to discuss these various spatial levels, and their differences and similarities, in a principled fashion that avoids a range of complications. Indeed, the vagueness in Hutchins' account carries over onto the supra-personal levels. I agree with Huebner (2014) that a major issue in the debates about distributed cognition is that it is unclear and imprecise what is meant by supra-individual cognitive processes. In order to make distributed cognition a viable position this must be rectified. It is beholden upon the proponent of supra-individual cognitive properties to delineate what these are in a clear fashion. Davies and Michaelian observe that agent-centric approaches seem to hold the high ground in these regards (2016, pp. 312-313). But this is not so. As I shall show in the next section, agent-centric positions are undermined by a range of problems related to scaling up and down: viz. the carrying-across of inappropriate conceptual baggage onto levels which should be treated on their own terms.

3.3. The problems of multi-level analysis

Most positions in the philosophy of cognition take an agent-centric approach to defining the centre of cognitive systems. In this section, I aim to show a number of problems that occur for the agent-centric view and in the final section of this chapter I will then outline the alternative task-centric approaches. Whilst Davies and Michaelian attacked both methodological individualism and extended cognition, here I limit my discussion to externalist positions because my aim is to

formulate a coherent account of distributed cognition[36]. Clark has popularised an agent-centric approach to externalism with the tagline: cognition is "organism centred [but] not organism bound" (2008, pp. 123, 135; e.g. see Giere 2006, 2007, 2012 for an account of distributed cognition based on this proposition). It is this type of claim that Davies and Michaelian (2016) are attacking (also see Hutchins 2011). To see why, I return to the simple case study from the first chapter, Norman's (1991) pre-flight checklist: this example cuts to heart of the issues here.

In order to complete a series of pre-flight safety checks, one can imagine an agent going about this in two ways. They could attempt to remember a laborious and complex procedure solely in one's head. Or alternatively, they could rely on a checklist that not only outlines the important safety checks that must be carried out in advance, but also allows the agent to keep track of the checks they have made through being malleable – i.e. by making marks in the appropriate space on the form, which can be visually discerned, and correspond to a specific completed check. This is another example of what Donald (1993) called an exogram – an external memory source (as seen in the navigation case study in the previous chapter).

Norman notes that many people would see the use of the checklist as 'improving' the performance of the agent in this memory task. Importantly, he goes on to point out that how the agent's performance is 'improved' is often mischaracterised because we fail to take into account the distinction between multiple levels of analysis – in particular, between the agent level and the wider system level. In turn, this can create a series of pseudo-problems. *Firstly*, the impact of the cognitive artefact on the agent's performance is often misunderstood as "amplifying" the agent's capacities (e.g. Bruner et al 1966). Instead, the interaction is far more complex. The manipulative activities of an agent on cognitive artefacts not only alter the task-space but also "transform" the neurocognitive

[36] This is not to say that methodological individualism does not have problems with carrying across inappropriate conceptual baggage when scaling down. For instance, Drayson (2012) gives a good overview of these problems, e.g. the "mereological" and "homunculi" fallacies – I return to these below.

profile of the agent – both functionally and structurally (see Menary 2007a, 2013b, 2014, 2015). *Secondly*, a related issue here is that multiple levels of analysis are either implicitly overlooked or misconceptualised. The resultant errors flow in both directions: with invalid inferences of agential features onto group level systems; and system's level properties inappropriately attributed onto the agent – what Hutchins calls "overattribution" (1995a, 2008). Both these sets of problems have a common source: the failure to properly differentiate and designate a level of analysis on its own terms. I now discuss these in more detail.

3.3.1. Cognitive artefacts and three kinds of enhancement

Norman (1991) uses the checklist example to show the importance of being clear about multiple levels of analysis in reference to how agents interact with cognitive artefacts. In the previous chapter I rejected an information theoretic approach as adequate – by itself – for cashing out the interactions of agents and cognitive artefacts. Here my primary concern is on explicating how this relationship impacts on the performance of the agent and how it is often misunderstood.

Human agents cohabit in what I am referring to, following Menary (2010e, 2014) and Sterelny (2003), as "cultural-cognitive niches". These are structured environments abundant in cognitive technologies, physical tools, sculpted environmental pressures, conspecifics, teachers, cultural knowledge, and public symbol systems. Human agents learn to manipulate these various tools and techniques to achieve cognitive goals (Menary 2007a, 2010a, 2010b, 2013b). A wide range of theorists, such as Hutchins (1999) and Luria (1928), have claimed that tool-use is "absolutely fundamental" to human cognition and what it means to be human (also see Cole & Engström 1993; Clark 2008; Dennett 1996, 2017; Donald 1993, 2010; Engelbart 1988; Heersmink 2013; Henrich 2016, Jeffares 2010; Johnson-Frey 2003; Malafouris 2008, 2010, 2013; Menary 2007a, 2017; Preston 1998;

Skagestad 1993; Sterelny 2010; Sutton 2008; Theiner & Drain 2016; Vaesen 2012; Vygotsky 1978, 1981). But there is an important nuance in how we understand what role these cognitive tools play in cognitive activity – and being imprecise on this matter misses out key features that can later lead to a range of inaccuracies in how we understand distributed cognitive systems.

For instance, Bruner and colleagues (1966) describe the effects of using cognitive artefacts as "*amplificatory*" of an agent's cognitive capacities in a similar manner to how physical tools amplify physical capacities. Bruner and colleagues identify three kinds of cultural amplifiers: of motor capacities (e.g. cutting tools, wheels, etc.); of sensory capacities (e.g. smoke signalling, magnifiers, telescopes, etc.); and of "ratiocinative capacities" – and this is of "infinite variety, ranging from language to myth and theory and explanation" (p. 56). The issue here is whether amplification is an appropriate term to describe the role each of these cultural artefacts plays in human activity (Cole & Griffin 1980; Hutchins 1999, 2006; Norman 1991; Menary & Gillett 2017). For instance, is it right to state that cognitive artefacts operate just as a spade increases one's capacity to dig, or how a loud hailer increases the capacity to project one's voice? This would entail that a calculator amplifies one's capacity to compute certain mathematical equations and a diary amplifies one's capacity to remember certain pieces of information. As I will demonstrate, this is not an accurate portrayal.

Although amplification of motor capacities is an appropriate characterisation, this terminology does not fit adequately with the latter two cases. I think this can be made particularly clear with reference to Humphreys' (2004, pp. 4-6) terminology of "epistemic enhancers". He identifies three ways in which a physical or epistemic tool can "enhance" a human agent's epistemic

capacities[37]: [1] *extrapolation* – the quantitative extension of the dimensions and domain of an existing capacity. A good example is the use of an optical telescope or microscope[38]. This fits with Bruner and colleagues' notion of amplification. [2] A second kind of epistemic enhancer is *conversion*. This involves altering a particular stream of information into a differing but already existent medium/modality. For instance, a sonar device alters sound waves into a visual medium. In the discussion below, this kind of enhancement will be one of great importance because many cognitive tools and techniques operate by transforming a task that is difficult in one domain into another which is much more suited to our anthropic profile (e.g. we are a species that is predominately vision orientated). A good example of this is Charbonneau's (2013) historical analysis of how physical models were used in the exploration of molecular structures in biochemistry. Here the conversion was from an abstract symbolic format (mathematical notation), in which the computational problem was prohibitively difficult, into the format of three-dimensional wooden models. This then enabled what Charbonneau calls "haptic reasoning" – the physical manipulation of the model's structure as a proxy for the target phenomena, which made the investigation much easier (2013, p. 592).

Lastly [3], whereas the former two kinds of epistemic enhancers lead to quantitative changes (either amplifying a pre-existing cognitive capacity, or altering a cognitive domain into another that is also already available to the agent), Humphreys identifies a form of enhancement in

[37] Although the following examples I list here are involve sensory-modalities this is not of concern for two reasons. Firstly, as stated in the general introduction, we are exploring what Theiner and O'Connor refer to as a "big tent" conception of cognition, which includes both perception and action as within the ambit of purview (also see Allen 2017; Anderson et al 2012; Kaplan 2012; Rowlands 2009, 2010). Secondly, Humphreys does go on to consider his distinctions of epistemic enhancement with regards to computation (2004, p. 5. Also see 2009 where he discusses computer-aided mathematics).

[38] This point is made with the caveat that one must be clear to specify what kind of microscope or telescope one is using since there are various kinds of these instruments involving different forms of intervention on their target to generate data. Some of these introduce novel capacities rather than straightforward extrapolation. For instance, a scanning tunnelling microscope. See Hacking (1983, pp. 183-209) for an in-depth discussion of this topic.

which the manipulation of some kinds of cognitive tools enables a genuinely qualitative change to the cognitive capacities of the agent. He refers to this as *augmentation*; and gives the example of gaining access to previously inaccessible features of the world through the use of scientific instruments such as Geiger counters and scanning tunnelling microscopes (which, respectively, allow us to investigate radiation and the atomic level – these would otherwise be infeasibly accessible).

With Humphreys' more sophisticated analysis we can return to Norman's example of the pre-flight safety checklist. At first glance, one might be tempted to claim that the pilot's ability to perform this memory task is extrapolated (and thus "amplified") by using the checklist. Norman (1991) notes that this is a mistake, because rather than 'amplifying' the agent's memory capacities this mediational device instead *converts* the task domain (also see Hutchins 1999, 2006). The structure and format of the checklist reorganises the task by stipulating a linear order to the procedure which can improve performance. There are two interrelated implications from this: [i] the transformation of the task space; and [ii] the transformation of the neurocognitive profile of the agent (Menary 2007a, 2014, 2015).

Firstly, the mediational role of the cognitive artefact and other epistemic resources demonstrates the impact that their manipulation and deployment can have on the task space (Hutchins 1995a; Menary 2007a). This shows that we must pay attention to the environmental resources involved in a specific cognitive task (Lave 1988). Tversky nicely summarises this view as follows: "Cognition is inescapably affected by the immediate who, what, where, when, and perhaps why" (2009, p. 201). Behavioural studies provide evidence that the task-context and material tools in the environment are vital for understanding the neurocognitive profile of agents. For instance, Kirsh (1995) examined the important role of how expert agents utilise space in a range of cognitive

tasks for simplifying choices, perceptions, and internal computations: e.g. in procedural tasks such as cooking a meal or reassembling a dismantled engine, parts can be placed in order of salience in easily perceptible groupings.

This cognitive ethnographic research has been complemented by lab-based experiments by Zhang (1997b) and colleagues (Zhang & Norman 1994) intervening on the material structures of tools to show that these had consequences for task performance. In particular, material tools that instantiated normative aspects of a task curtailed possible decision moves in an abstract task space to only those which are salient. In an experiment using the Tic-Tac-Toe game, it was found that one particular instantiation of four isomorphic versions of this game made it much easier for players to pick the optimal starting moves. In comparison with the other isomorphic versions, the optimal move was less visually perceptible (Zhang 1997b). Figure 3.1 below shows the four isomorphs – *Line*, *Numbers*, *Shapes*, or *Colours*. The aim of the test is to make a set of three matching units versus a computer player. Participants found the optimal solution easily in the Line isomorph but not in the others. The experiment shows that the physical format of the epistemic resources involved in a cognitive task can have a significant impact on performance and what kinds of actions are taken. These results were replicated in another test on isomorphs in the game *Tower of Hanoi* (Zhang & Norman 1994). Again, it was found that one particular isomorph improved performance, but interestingly the major difference was that the physical structure of this isomorph (the coffee cup version) made incorrect game moves a physical impossibility – i.e. it was physically impossible to break the rule of placing a bigger unit on a smaller unit – thus limiting an agent's choices and decisions to only those that were salient.

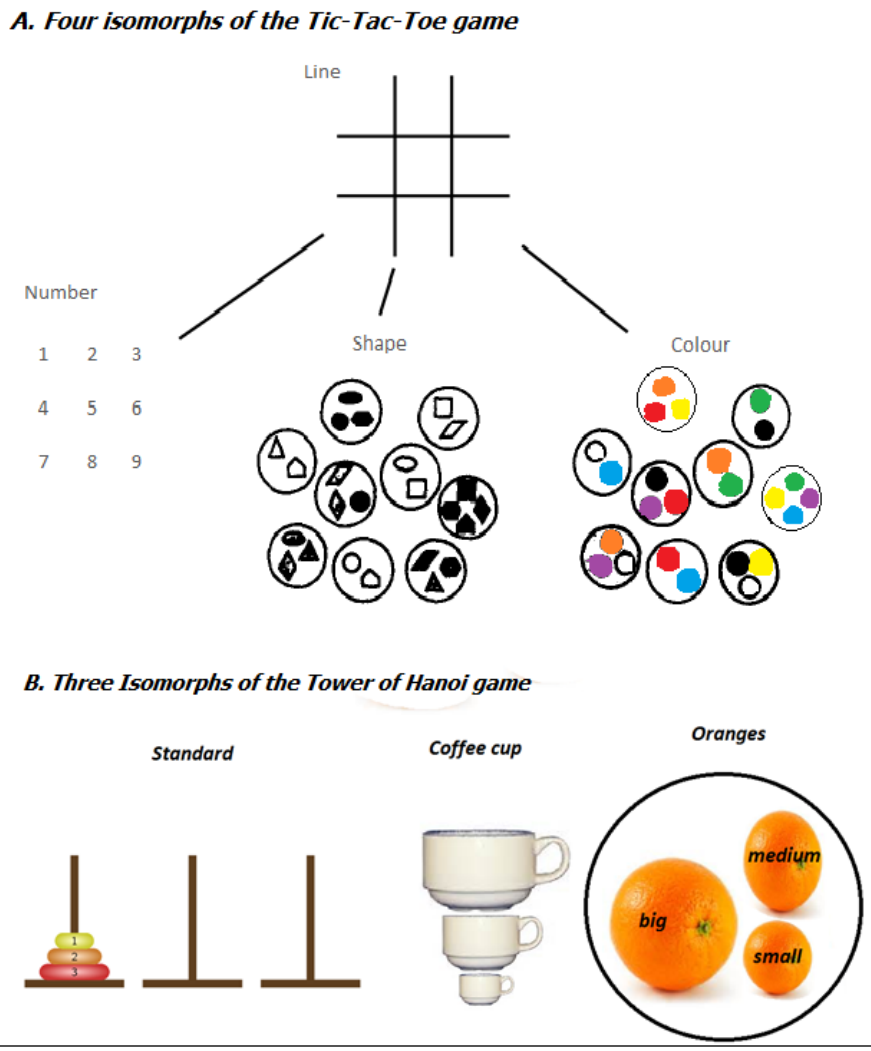

Figure 3.1: TOP PANEL [A]: Four isomorphs of the abstract problem space of the Tic-Tac-Toe game. The basic arrangement is LINE, in which the agent must get three adjacent squares in a row on the 3x3 grid. It is perceptually obvious that the best strategy here is to choose the middle square. In the other three isomorphs (NUMBER, SHAPE, and COLOUR) the best strategy is not as perceptually salient and task performance is drastically reduced. BOTTOM PANEL [B]: Three isomorphs for the abstract problem space of the Tower of Hanoi game in which one moves the arrangement of shapes (standard discs, coffee cups, or oranges) from one 'tower' to another according to two rules: [1] only move one piece at a time; [2] a smaller piece cannot be moved onto a larger one. Task performance was best with the coffee-cup isomorph because more constraints are instantiated in the representational format. This limits actions to only salient moves. For instance, one cannot physically break rule 2 because smaller cups do not fit on larger coffee cups without spilling. There are no such constraints in the other isomorphs. [Diagrams adapted from Zhang (1997b, p. 188), Zhang & Norman (1994, pp. 92, 97).

Although these experiments focused on games, this is not just some idle non-ecologically valid finding. For instance, both Charbonneau (2013) and Vorms (2012) have shown that these

factors play a role in the creation, construction, and manipulation of physical models, diagrams, and notation in scientific practices. Whilst some might be more willing to grant this as a factor in the use of physical models and diagrams where visual factors can be an important factor in an agent's reasoning processes (e.g. see Abrahamsen & Bechtel 2015; Giere 2002b, 2006; Knuuttila 2011; Nersessian 2005, 2006, 2009); it is reasonable to expect critics to be more sceptical of this in the case of mere notation (e.g. see Hofstadter 2001, p. 503). But a cognitive analysis of numerical systems by Zhang and Norman (1995) gives us good reasons to believe that Arabic-Hindu numerals have come to be the predominant "mental tools" for numerical tasks because of their format (cf. Chrisomalis 2004). This includes perceptual grouping features. Arabic-Hindu numerals have what they call a 2D structure in which the base value is clearly represented by the position of symbols and this drastically reduces cognitive load and simplifies the normative procedures involving their manipulation. In other notation systems size can often conflate with value to make their use more confusing – for instance, Roman numerals have an incongruent representational format in which larger numbers appear as visually smaller sets of symbols (e.g. 37 = XXVII, whereas 501 = DI).

Zhang and Norman's cognitive analysis is supported by a range of behavioural experiments on mathematics students conducted by Landy and colleagues (2007, 2014) in which they altered non-salient visual features around algebraic equations to create incorrect perceptual groupings that would go against the standard normative procedures – FOIL: first, outer, inner, last (see figure 3.2 below). Not only did they find that these non-salient perceptual groupings heavily influenced task behaviour, in post-experiment interviews several participants claimed that they had become aware of the misleading features *but still could not resist them* (also see De Cruz & De Smedt 2013 on the importance of notation, and its visual and physical format, for the tackling of more advanced problems in mathematics).

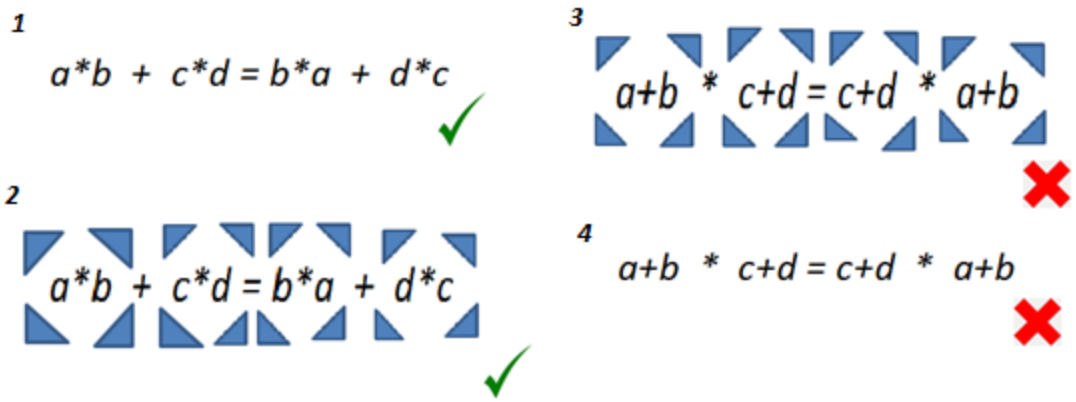

Figure 3.2: Four algebraic equations. Equations 1 and 4 are unaltered. Equation 1 is correct according to the FOIL procedures, which dictates the order in which the symbols in equation are manipulated. Equation 4 violates these rules. The other two equations (2 and 4) are equivalent to these equations (according to their respective columns), but are also encompassed by meaningless perceptual markings (blue triangles). In equation 2 they are congruent with FOIL whereas they are incongruent in equation 3. Despite having no relevance to the actual operations carried out, Landy and colleagues use this experiment as well as others to argue that visual perception plays a greater role in mathematical cognition than is normally recognised. [Diagram adapted from Landy & Goldstone (2007, p. 726)]

My point here is *not* that tool-use automatically distributes cognition beyond the skull of individuals, but rather that *how we use or manipulate epistemic resources* is vital for understanding internal processing and the structure of the task space itself. Although this might seem like a deflationary position it is actually quite far-reaching. Because although several critics have tried to argue that tool-use just involves "cognitive offloading", following Menary (2007a, 2012) I think that this misses a major factor in what is going on here – viz. that it is the normative physically embodied manipulation of these cognitive resources that enables overall improved task performance to be achieved and cognitive conversion to take place at the level of the individual perspective. Additionally, it is incredibly noteworthy that the repeated importance of norms in these empirical cases demonstrates my claim that a solitary focus on information flow is insufficient for properly understanding distributed cognition. In chapter five I will discuss this at greater length in the context

of developmental trajectories and the acquisition of normative patterned practices that mediate these interactions.

The above points concern the alteration of the task space mediated by the manipulation of the cognitive artefact and how this dramatically alters the way the agent interacts with the task domain. A second related aspect of this is how it also alters the cognitive functions being deployed by the agent – i.e. the transformation of the neurocognitive profile of the agent themselves. For instance, in Norman's checklist example, rather than the agent having to remember a procedure of checks and keeping track of them as they are performed – involving both long-term and working memory – they can keep track with the checklist. This second strategy alters the task-space into a perceptual coordination task involving the correct manipulation of the external linguistic symbols in conjunction with matching features in the environment and checking these off as the task progresses. Furthermore, there are good reasons to think that the alteration of cognitive functions is also associated with corresponding plastic changes in neural circuitry correlated with this activity (Menary 2013b, 2014).

Convergent evidence from neuroimaging studies of problem solving shows differing neural correlates for well-defined and ill-defined cognitive tasks (Jausovec & Jausovec 2000; Luo & Knoblich 2007; Qiu et al 2008). Ill-defined problem spaces have multiple solutions and no clear procedure, whereas well-defined problem spaces have a single solution with a clear procedure (Jaarsveld et al 2010). I contend that we can think of the conversion effect on task spaces (resulting from the correct manipulation of certain cognitive artefacts) in terms of shifting from an ill-defined to a well-defined problem. In the studies on external representations discussed above the common theme was the way these devices channelled behavioural repertories and rendered certain task solutions more salient. If the task space radically alters depending on whether we are considering an agent by

themselves completing it solely in their head compared to an agent and artefact coordination dynamics, then this gives us good reasons for thinking that we cannot simply or straightforwardly scale up the cognitive properties of single agent system to the agent-artefact system. And this is because both the task-space has radically changed, and also the neurocognitive properties of the agent have been altered.

To properly understand this, we must also take into account the developmental trajectories by which agents gain the expertise to manipulate these cognitive artefacts according to cognitive norms (Menary 2007a, 2010a). It is this learning procedure that transforms the agent's neurocognitive profile. Here we have an example of why we need to consider how cognition is distributed over time. I return to this topic in chapter five.

To summarise: in Norman's (1991) example of simple dyadic tool-use the crucial point is *the importance in the different levels of analysis*. Norman emphasises that if we do not adequately differentiate between what he refers to as the systems level and personal level aspects of the unit of analysis then we are more likely to make the mistake of considering the use of a cognitive artefact as amplificatory. Indeed, at the systems level the overall performance of the system is enhanced. The pilot – by manipulating the checklist – performs the memory task of pre-flight safety checks in a more optimal manner. But at the level of analysis of the agent, the checklist mediates between the agent and the task, consequently changing the nature of the task from one primarily drawing on memory to a perceptual one that involves manipulating symbols (on the list) in conjunction with their targets in task space (a series of procedures for ensuring the safety of the aircraft).

Lastly, given these alterations in the task space, this engenders a reciprocal change in the patterns of activation in the neural populations at the sub-individual level. Therefore, Norman is right to emphasise that the level of analysis is key for looking at distributed cognitive systems. What

this demonstrates is that the interactions of agents and cognitive artefacts must be dealt with carefully, particularly in reference to properly discriminating different levels of analysis. By paying attention to this relationship and how it impacts on the task space we get a clearer understanding not only of the distributed cognitive system but also its components as well. I now turn to the second set of complications that arise when one omits these important details.

3.3.2. Two ways of transposing inappropriate conceptual baggage onto the wrong level of analysis

The key difference between the approach advocated by Davies and Michaelian (2016) and Hutchins (2011, 2014), and other externalist positions which are agent-centric, is that the latter approaches always hold that the centre of organisation is the agent. However, Davies and Michaelian also recognise that the methodological individualist, whom they see as beginning from an uncontroversial starting point, can rightly respond that if the agent is not central to our investigation of human cognition, then what is? They add that this potentially exacerbates the problem of cognitive bloat because it leaves us bereft of a boundary criterion (2016, pp. 312-313). Whilst this is a partly reasonable objection – insofar that we are challenging the traditional demarcation of what is cognitive and so must offer a viable alternative – it is wrong to suggest that the traditional view is itself unblemished.

Kaplan (2012) has noted that the very debate itself about distributed cognition has rendered the traditional position of internalism question-begging, and thus not uncontroversial – i.e. one can no longer respond to the question 'what is cognitive?' by pointing to skull-bound processes because this is the very issue that has been brought into dispute. In the next section I will formulate a positive alternative to the agent-centric approach based on task-specificity. But first it is important

to detail two ways in which an agent-centric approach runs into conceptual difficulties in attempting to articulate multiple levels of analysis.

The central issue here is that rather than being conceived of on their own terms, the agent-centric approach inadvertently transposes across inappropriate conceptual baggage which generates problems for multilevel analysis. By beginning with an agent-centred approach the concern here is that levels of analysis that are sub-personal or supra-personal will be improperly conceived in agential terms. When scaling down, Drayson labels this a 'mereological fallacy" in which there is the false ascription of a property to the part that should only be ascribed to the whole (2012, p. 3). In particular, it is the false projection of personal-level descriptions onto sub-personal explanandum (also see Fabry 2017, p. 14).

Drayson notes that a related issue is the "homunculus fallacy" – the attempt to explain higher-level cognitive properties (e.g. intelligence) through the postulation of that property at a lower-level of explanation (2012, pp. 5-6). She notes that the difference between these issues is that the former relates to the correct use of concepts and the latter is an issue about explanatory practices (p. 6). Whereas Drayson is focusing on scaling down, I contend that both of these fallacies arise with scaling up if one adopts an agent-centric approach (also see Huebner 2014 for similar points). I think that this conceptual baggage then clouds how certain issues are debated and prevents other options from coming forward. Furthermore, this error occurs in both directions: viz. there is an inappropriate scaling up of agential features onto the group or system level (residual internalism); and reciprocally, there is an inappropriate and mistaken postulation of cognitive properties on to the individual that should be attributed to supra-individual level – what Hutchins calls "overattribution" (1995a, 2008). I now discuss each of these in turn.

The problem of residual internalism: The first issue here is an inverse of the mereological fallacy that Drayson identifies (from the personal level onto the sub-personal level): it is the false attribution of agential features onto a supra-individual level of analysis that should instead be treated on its own terms. Davies and Michaelian note that one issue for agent-centric approaches to distributed cognition is that whilst they can certainly recognise the existence of supra-individual cognitive processes and systems, they will struggle to individuate them appropriately (2016, p. 310). The projection of agential properties onto supra-individual levels (composed of agents and artefacts) is an invalid inference because there is no well-motivated reason for attributing agential features onto higher levels of organisation (just as these reasons are lacking in cases of scaling down). For instance, Theiner and O'Connor note that one good argument against such inferences is that it is highly unsettled what phenomenal consciousness of agency is at the individual level; therefore, to project upwards towards a notion of collective consciousness is dubious at best (2010, p. 106).

I argue that this problematic situation arises precisely because some other approaches to distributed cognition begin with an agent-centred approach and its attendant individualist bias; consequently, taking this as the blueprint for how to approach cognitive properties at differing levels of analysis. Clark and Chalmers' (1998) "parity principle" argument for the extended mind is an exemplar of this view. It holds that when considering a process in the world, we should treat it as cognitive *if,* were it in the skull, we would see it as such:

> If, as we confront some task, a part of the world functions as a process which, *were it done in the head*, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world is (so we claim) part of the cognitive process. Cognitive processes ain't (all) in the head! (p. 8 original emphasis)

There are several problems with this analogy. Firstly, Menary has noted that this argumentative approach is "fatally flawed" because "it assumes the very position it is meant to displace" (2010a, p.

234). Viz. by biasing the primacy of the skull this amounts to a "residual form of internalism" because it prejudices the discussion in favour of skull-bound processing (2007a, p. 63). Furthermore, I think this leads to erroneous discussions about agential properties at the group level (e.g. consciousness, agency, awareness, etc.) that are not properly motivated and which smuggle in concepts that lead to debates about inflated homunculi (see Huebner 2014 for a review; also see Giere 2006, 2007, 2012).

Following Theiner and colleagues (2010, pp. 378-379), I think that the proponent of distributed cognition must be very careful about what cognitive properties one postulates at supra-individual levels of analysis; and beginning from an agent-centric perspective skews this already difficult issue towards treacherous conceptual terrain. Following Wilson (2001, pp. 266-267, 2004, pp. 288-289), and building on the naturalistic platform I laid out in the general introduction, I think that we should focus on demonstrating that the systems level instantiates *at least one paradigmatic cognitive function towards a specific task*: e.g. attention, perception, problem-solving, memory, etc. Theiner and O'Connor note that the advantage of this task-specific strategy is that it focuses in on specific well-defined properties, or clusters of properties; and as such, prevents the discussion from descending into "cheap but unreliable inferences" about agential features which critics would otherwise attempt to use as knock-down arguments (2010, p. 82). However, as demonstrated in the previous sub-section, we must also be careful to distinguish between levels and differentiate between how this cognitive property is instantiated.

Secondly, the parity principle overlooks the fact that "Things in the world behave differently than things in the mind" (Kirsh 2010, p. 446). Viz. the manipulations of external representations in a putative cognitive process are not equivalent to internal processes. Although critics such as Adams & Aizawa (2001) rightly attack proponents of the parity principle for overlooking this point; it must

also be recognised that some advocates of distributed cognition have long known about this important functional distinction. For instance, Donald discusses the functional differences between internal and external representations of memory – engrams and exograms – in great detail (1993, pp. 314-316; 2010, pp. 72-73). Exograms can be frozen in time, are physically constructed in multiple substrates, and can be public. Additionally, both engrams and exograms have differing retrieval procedures which draw on different neural populations, etc. We can see these differences at play in Norman's example (and also in the bearing log book from the navigation case study), where the physical format of the checklist enables retrieval of information through the visual system rather than by accessing an engram.

As such, Sutton (2006, 2010) has noted that rather than this being a knock-down argument against distributed cognition; it instead shows that the position must be premised on a different basis. We need a different starting point for regarding cognition that does not lead to these mistakes – in the next and final section of this chapter I outline a task specific approach to this problem. But before I do so there is a reciprocal version of this problem that must also be discussed.

The problem of overattribution: this issue can be demonstrated by again returning to Norman's checklist example. Following Menary (2007a), I was careful in the proceeding sections to stipulate that it is the *physically embodied manipulation* of this mediating device (the coordination dynamics between agent and artefact) that enhances task performance at the systems level and facilitates cognitive conversion at the personal level. A failure to make this distinction results in the problem Hutchins labels "overattribution" (1995a, 2008). This is a variety of the mereological fallacy in which cognitive properties which should properly be attributed to the system level are inappropriately attributed to the agent (who is only part of the wider system).

The insufficiently clear manner in which the coordination dynamics of a distributed cognitive system are cashed out – i.e. solely in terms of information flow or, alternatively, what is sometimes referred to as "coupling" (e.g. see Clark & Chalmers 1998) – has led to criticism. Critics attack this portrayal of the distributed cognition position for making a category error of seeing a causal relation as an instance of constitution. Adams & Aizawa have labelled this the "coupling-constitution" fallacy (2010, pp. 67-68). Rupert captures the crux of the problem nicely with the following analogy:

> One wishes to understand an important historical event, say, Nazi Germany's invasion of Poland. In order fully to understand this event as an historical event, one would need to know, among many other things, a great deal about the economic conditions in Germany during the nineteen-twenties. This does not imply that the economic conditions in Germany during the nineteen-twenties are *part* of the invasion. (2004, p. 396. Original emphasis)

Here we see a repeat of the inability of information processing *by itself* to distinguish between genuine components of a cognitive system and features which are mere causal background conditions (thus, making it a conceptual cousin of the problem of cognitive bloat and highlighting the slippery slope it entails). This is a serious issue for what Sutton (2010) labels "first-wave" proponents of extended and distributed cognition arguments which rely on the parity principle[39]. But my position is unaffected by this objection because I am not adopting an agent-centric approach. Instead, we need a viable alternative which, rather than claiming that the cognitive artefact is part of the agent, takes the interaction of the agent and artefact to form a new "hybrid" at the systems level (Menary 2007a). Menary puts this formally as follows (2007a, pp. 62-63): rather than the category mistake of $X + Y \rightarrow X = Y$ (the mistaken notion that the coupling of X and Y entails that X is *a part of* Y) – the issue that bothers Adams & Aizawa (2010, p. 68). It is instead the argument that the

[39] A position which is still defended despite the criticisms I will now outline (e.g. Carter et al 2014; Magnus 2007, p. 300; Pritchard 2010, 2016).

physical manipulations of the artefact that mediate the interaction of the agent and the artefact (X and Y) form a *larger* hybrid system (Z), which is a new unit of analysis composed of the agent and certain environmental features. Importantly, this larger distributed cognitive system composed of the interactions of agent and artefact has different cognitive properties to the agent considered by itself (Menary 2007a, 2013b; Menary & Gillett 2017). In later chapters I outline Menary's key observation that the interaction of agent and artefact is not merely causal but is also normative – this entails the importance of cultural practices (especially over diachronic timescales) and how these alter the neurocognitive profiles of agents and integrate distributed cognitive systems.

Before doing so I must first outline a framework for properly articulating multiple levels of analysis that avoids the problems I have just discussed. It is noteworthy at this juncture that the numerous issues here advise against Hutchins' (2014) suggestion of a 'fractal" approach to the topic of multilevel analysis, since this would bias the discussion towards the notion of self-similarity. Whilst it is certainly possible that there are points of similarity across differing levels of analysis I would urge caution here because of the conceptual errors that lurk in the shadows of the fractal terminology. Furthermore, it overlooks the importance of seeking the differences of cognitive properties at multiple levels of analysis. Given that this is of such importance even in an example as simple as the manipulation of a checklist or log book for a memory task; it is highly likely that this will also be of great importance for complex organisations of coordination dynamics. This renders the notion of cognitive fractals of very limited use – as such, I shall not discuss them further[40].

[40] A noteworthy exception is recent work by Goldstone and Theiner (2017) on what they label "multiple, interacting levels of cognitive systems" (MILCS). Here they formulate a network systems theory account of information flow intended to be equally viable at both neural and social levels. There claim being that are potentially common principles which govern repeated patterns of structure at these differing levels of analysis because both kinds of system face similar issues: e.g. bandwidth size; noise; resilience; and rates of interconnectivity (I return to these latter three topics in chapter 6, sections 6.2.1 and 6.3). Goldstone and Theiner base their claim on a range of empirical evidence in network science that has applied formalisms to both social and neural networks.

3.4. A task-specific definition of levels

Having demonstrated the importance of multiple levels of analysis for properly exploring distributed cognition, and problems for agent-centric approaches in articulating this, I shall now outline an alternative. Davies and Michaelian have claimed that there is no need to start with the whole agent if we have a well-specified cognitive task (2016, p. 317). This is also a view shared by others, for example Huebner writes approvingly of Wilson's work that: "We both reject the assumption that individuals should be the privileged locus of explanation, and hold that the specification of tasks, and the corresponding process of reverse engineering, can be carried out on groups as well as individuals" (2014, p. 226. Also see Anderson et al 2012; Magnus 2007). To make this viable, the notions of levels and task-specificity must be clarified. I will first discuss the issue of levels since this then provides us with the kinds of distinctions that task-specificity must be able to navigate.

3.4.1. Mechanistic levels

In the previous sections of this chapter the following discreet kinds of levels have been identified: sub-personal, personal, and supra-personal. The notion of 'level' is a particularly fraught term in philosophy. In a useful "field guide", Craver provides an overview of the many ways in which this ambiguous term is used (2007, pp. 163-195). He divides these into two major groups – levels of

---

MILCS is a novel and interesting approach that potentially provides the distributed cognition proponent with a network systems theory language that can be applied across all cognitive domains and across all the wide variety of differently scaled systems. However, whilst this approach avoids the homuncular fallacy and is also less likely to commit the overattribution fallacy; I have similar reservations about this position that I do about all information flow accounts of distributed cognition that I discussed in chapters one and two. Namely, that they are insufficient for tackling the problem of cognitive bloat and the problem of cognition distributed in time. As such, I still think that the MILCS approach – insofar that it is applied to systems composed of human agents – must be supplemented by an account of normative patterned practices (see chapter five). This point will become especially in chapter six when I discuss how information channels in the navigation team are sculpted and partially constituted by diachronically distributed normative patterned practices.

nature and levels of science – which he then places in a hierarchy (see figure 3.3 below)[41]. Levels of science either refer to hierarchical branches of science (e.g. fields, paradigms, etc.) or their epistemic constructs (e.g. models, theories, etc.) (p. 171). Levels of nature imply that there are distinct units within nature – e.g. size – and that the aim of science is to attempt to carve nature at these joints. For instance, Oppenheim and Putnam (1956) famously divided reality into six levels (particles, atoms, molecules, cells, organs, organisms, and societies). There are many issues with such a coarse grain of analysis – a particularly telling example comes from Wimsatt (2007, p. 208) who notes that size cannot be a sufficient indicator of level because black holes and bacteria can be the same size but are obviously radically different. Thus, we would not want to claim that they are on the same level of nature. For instance, a black hole would not be subject to Brownian motion. Wimsatt is keen to point out that this is not just a "philosopher's silly hypothetical example" – as demonstrated by a series of letters in the journal *Nature* discussing the properties of black holes ranging in size from 10^{-2} to 10^{-4} mm in diameter (ibid). This repeats the concerns I outlined in chapter one for attempting to draw comparisons between differently sized distributed cognitive systems without careful consideration.

Craver opts for a much more *localised* definition of level which he ties to the notion of mechanistic explanation. By 'localised' this simply means that the levels of analysis in a particular phenomenon do not extend globally across onto other systems. Mechanistic levels are levels of nature but they are levels of composition which are neither spatial nor material (2007, pp. 188-189). They are levels of the interactions between the relata (components) that form the organisation of the system. The notion here is that a particular system is composed of sets of inter-level relations between the whole of the organisation of the system and the parts or components which, in concert, produce the behaviour of the higher level of explanation or analysis. This is useful for our purposes

[41] For reasons of space I will not discuss most of these as they are superfluous for our purposes.

here for several reasons. *Firstly*, it nullifies the problem of size since we are now focused on interrelations rather than spatial comparisons. But this does not mean that we cannot compare distributed cognitive systems. Indeed, a *second* benefit of this approach is that it makes us look at composition in a way that enables distributed cognitive systems to be formed of nested hierarchies that are themselves other distributed cognitive systems. *Thirdly*, Craver notes that a mechanistic conception of levels corresponds to components rather than objects so that in "...many cases, the components picked out in a mechanistic decomposition fail to correspond to paradigmatic entities with clear spatial boundaries" (p. 190). As such, mechanistic levels are not monolithic divisions in the structure of the world; instead they are much more local (pp. 190-191).

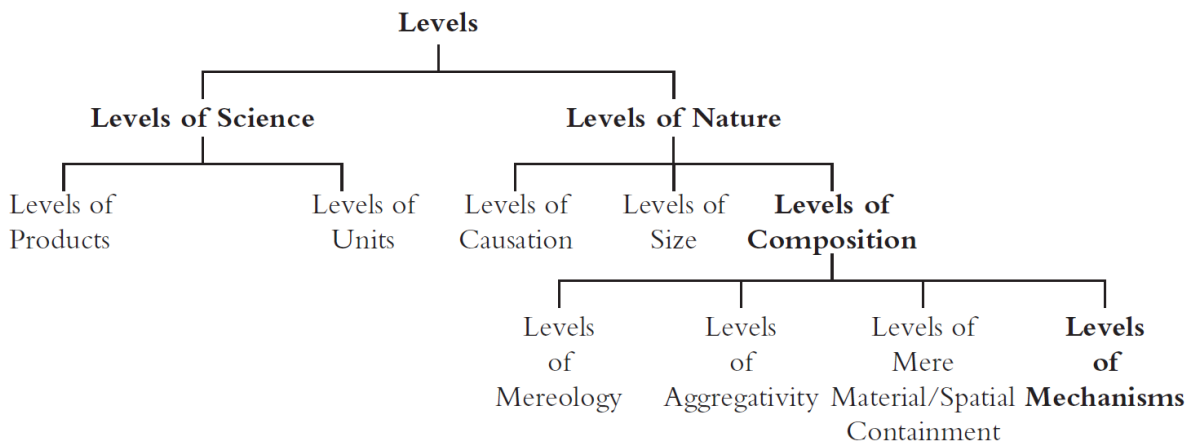

Figure 3.3: Taxonomy of various notions of levels in the philosophical literature can. For our purposes, I shall only be discussing levels of analysis (see main text). [Diagram from Craver (2007, p. 171)]

I think this makes the mechanistic approach very amenable to my pluralistic framework because it places a larger emphasis on the particularities of case studies and the compositional hierarchy of the system under investigation (p. 191). In the next chapter I outline the importance of mechanistic composition for my pluralistic view in greater detail. Our primary concern here is to make the notions of sub-personal, personal, and supra-personal more concrete with regards to task-

specificity. Careful arguments must be put forward in favour of multilevel investigations and explanations because they face charges of explanatory superfluity (Huebner 2014; Theiner & O'Connor 2010) – i.e. what role do the different levels actually play in an explanation if they are 'non-fundamental' (Craver 2007, p. 163). My suggestion is that we can borrow from three main sources in these regards – Caporael's (2014) core configurations, Wimsatt's (1986) definition of emergence; and Steiner's (1972) task taxonomy – and that the notion of task-specificity allows us to do this in a non-ad hoc manner. Before doing so, I must first show how my pluralistic approach is consilient with both the sub-personal and personal levels of analysis in order for my position to be a properly multilevel account.

Firstly, the foremost level of focus for cognitive science and philosophy of mind is the personal level. Because of the degree of focus by others here I shall not dwell on it at great length, nor do I think it is a particularly problematic level of explanation for distributed cognition compared to other positions. I.e. topics of consciousness, agency, etc. are perennial problems for philosophy and science. A point of concern here is that if one adopts Davies and Michaelian's alternative to an agent-centric approach to individuating cognitive systems this can result in the balkanisation of the mind.

We can make this clear with the Norman example. When the agent is engaged in this memory task – which is converted into a perceptual task and one involving symbol manipulation and reading through the mediation of the checklist – one could argue that for this cognitive task only those regions which are functionally involved in this task are part of the cognitive system. This entails the seemingly counterintuitive conclusion that only parts of the agent are cognitive on the task-specific view. Take the blind man's stick example discussed in the previous chapter, based on a

task-specific approach the stick is only part of the blind man's cognitive system whilst he is in locomotion.

Although this is a position that Bateson agrees with (1972, p. 459), it does open up a task-specific approach to Rupert's (2010) "fleeting" cognitive systems challenge. I.e. cognitive systems become so short-lived that they are potentially coming into and out of existence over incredibly short-periods of time at a frenetic rate (also see Menary 2010c, pp. 13-15 for a discussion). By a principle of parsimony this *could* be seen as undesirable. Davies and Michaelian bite the bullet on this point and counter-argue that "It may not be a problem to posit extremely short-lived cognitive systems, but it is a problem to posit too many systems" (2016, p. 310). Indeed, Anderson and colleagues (2012) have also proposed task-specific synergies of neural, bodily, and environmental components can be "softly-assembled" in which the components are not fixed. Additionally, Allen (2017) has recently pointed out that a naturalist should be comfortable with short-lived cognitive systems given that there is research into the cognitive processing of very short-lived organisms. As such, this shall not concern us here any further given my adherence to a naturalist position (as I stated in the general introduction).

Another issue that a critic might make against a task-specific approach is that it negates or overlooks the importance of the human agent. In response, it is noteworthy that the key criticism against agent-centric approaches is that they have failed to consider other potential organising elements in distributed cognitive systems (Davies & Michaelian 2016; Hutchins 2011; Kirchhoff 2012). But the reciprocal is not true: there is no reason for a task-specific approach to ignore the agent or personal level. As Hutchins clearly states: just because distributed cognition considers larger units of analysis in which agents are just components, or in which agents might not be central; this is not to deny that they are integral components in these larger systems: "Any attempt to

explain the cognitive properties of such a larger system without reference to the properties of its most active integral parts would be deficient" (1995a, p. 287). But he goes on to note that where methodological individualists have gone wrong is in not recognising the importance of supra-individual levels of analysis for properly understanding the neurocognitive profiles of the agents themselves (pp. 287-288; also see Sutton 2010; Menary 2007a, 2013b). Again, the notion of nested hierarchies is important here.

The second main level of analysis in the previous discussions was sub-personal levels of description. Again, I shall take it as given that this is not a matter of concern – not because the details of this position have been exhaustively and definitively outlined, but rather because there is much philosophy of mind and philosophy of neuroscience that explores this topic (e.g. Bechtel 2008; Craver 2007; Krakauer et al 2017). For the purposes of the multilevel analysis necessitated by distributed cognition, it is only necessary to note that the mechanism by which higher-levels of organisation (e.g. acquiring specific sets of cultural practices) can engender plastic changes on the nervous system involves a form of neuronal redeployment whereby both the function of neural populations and their structure are altered by habitual learning (Anderson 2010; Dehaene & Cohen 2007; Downey & Lende 2012b; Menary 2013b, 2014). The enculturation of human neurocognitive profiles involves both the micro and macro scales (Roepstorff et al 2010). Therefore, sub-personal processes are an important part of the overall multilevel analysis in which distributed cognition is engaged (see Downey & Lende 2012b, p. 40; Fabry 2017; and Li 2003 for similar points).

I return to this in greater detail in chapter five. The only other stipulation is that, following Davies and Michaelian, any consideration of sub-personal levels is centred on the cognitive task in which the system is engaged in. On this point, it is noteworthy that Craver notes that just such an approach is the standard picture when discussing the various sub-personal levels of neural

processing involved in spatial memory: from the molecular level and cellular-electrophysiological level, to the level of spatial map formation and level of spatial memory (2007, pp. 165-170). He adds that one can go to higher levels of organisation – what I refer to as supra-individual levels. Of the three main kinds of levels identified in the previous discussions, this is the one that we must now adequately individuate.

Table 3.4: There are various repeated groupings of agents (and artefacts) involved in shared tasks that are a common feature of human social groups that are hypothesised to have been stable over long periods of evolutionary time. [Table adapted from Caporael (2014, p. 62)]

Core configurations	Group size	Modal Tasks	Affords/Scaffolds
Dyad	2	"Up-close" interactions; sex, artefacts, infant-caregiver interaction	Micro-coordination
Task-group	5	Foraging, hunting, gathering, direct interface with habitat	"Distributed cognition"[42]
Deme (band)	30	Movement from place to place, general processing and maintenance, work group coordination	Shared construction of reality (includes indigenous psychologies), relational group identity
Macrodeme (Macroband)	300	Seasonal gathering, exchange of individuals, resources, and information	Stabilising and standardising language; ontologies, collective identities

3.4.2. Core configurations, emergence, and task-taxonomies

I suggest that above the level of the individual, Caporael's (2014, pp. 60-63) notion of "core configurations" delineates a set of supra-individual conglomerations of agents and environmental resources that are stable, repeatable and *task-dependent*. She identifies several discrete arrangements of agents in task-specified groupings ranging from dyads (two agents), task groups

[42] It is noteworthy that in this table Caporael exemplifies the typical narrow portrayal of distributed cognition I discussed in the first chapter: namely, that distributed cognition only occurs at the discrete level of small groups. I reject this in favour of the notion that multiple discrete configurations are potentially distributed cognitive systems.

(five), bands (thirty), and macrobands (three-hundred) – and which she claims are "repeated assemblages" found in human niches (2014, p. 62. See table 3.4 above). For my purposes the important factor here is that we can readily identify discrete arrangements of agents (and cognitive artefacts) above the level of the individual. Furthermore, we can treat these as various levels in a unit of analysis that are composed of smaller, nested, cognitive systems, and the relations between these component parts and how they are organised (also see Hutchins 2014).

Core configurations, as task-specific groupings, meet the criteria for what Wimsatt refers to as "robustness" – this is the notion that a particular identifiable phenomenon can be said to be a genuine feature of reality (as opposed to a mere construct or artifice) if there are multiple convergent lines of evidence supporting its claim to be part of one's ontology (2007, ch4; also see Huebner 2014, pp. 14-15; Theiner & O'Connor 2010, pp. 103-105 for a similar discussion). Robustness is important because it offers prima facie defeasible reasons to reject the explanatory superfluity of system level analyses of cognitive properties. Indeed, there are a wide range of examples in which groups of humans and their interactions are studied by many different fields and methods – group psychology, sociology, anthropology, economics, etc. Across these diverse fields stable and repeated patterns of activity have been observed. For instance, [1] *Social loafing*: the "reduction of individual effort exerted when people work in groups compared to when they work alone" (Forsyth 2009, pp. 294-295). An example of this is the phenomenon of "collaborative inhibition" in collective memory research, in which a nominal group (composed of aggregate scores of separate individuals) outperforms a collaborative group in a memory task (see Barnier et al 2008 for a discussion).

[2] *Groupthink*: the pressure to reach a consensus can lead to the suppression of alternative views (Solomon 2006). For group deliberation to be effective there are a range of conditions that

need to be satisfied: [a] the encouragement of the expression of true opinions; and [b] a diversity of deliberators. Levy notes that these conditions are quite weak and there are many ecologically salient ways in which they can be subverted – such as charismatic leaders and "informational cascades" (runaway feedback loops and deliberating mechanisms) – that can lead to groupthink (2017, p. 5). Minson & Mueller add that in collaborative work, when several members triangulate on a solution or decision this can lead to overconfidence and a "myopic disregard of alternative viewpoints" (2012, p. 219). Sunstein argues that social pressures can lead to group members being silent even if they do not think that they are wrong because they fear reprisals such as ridicule (2005, pp. 985-986). This leads to a suppression of minority reports on topics and contributes to groupthink and conformity bias. Furthermore, Thagard (2006) argues that, in general, emotional factors are a characteristic feature of group decision-making; and that this has a role at both the individual and group level (again demonstrating the importance of multilevel analysis). Fricker (2007) argues that social status can have an impact on whether the views of certain individuals are given appropriate credibility. Fricker labels this a form of "epistemic injustice" where the speaker is denigrated *qua* provider of knowledge. But there is also a flip side here insofar that the group is deprived of knowledge because they have inappropriately undermined the legitimacy of the speaker and forsaken the message without properly assessing its actual epistemic credibility; and this results in an "epistemic dysfunction" for the whole system[43].

[3] *Group polarization* (Isenberg 1986; Sunstein 2002): groups can be driven to more extreme conclusions, through the amplification of shared biases, than if individuals assess the same information separately. This can occur because of social influences on behaviour (e.g. individuals

[43] A pertinent and particularly egregious example that Fricker describes involves a woman in a male dominated work place who has to rely on an ally to put forward suggestions in problem solving and decision-making debates (2007, pp. 46-47). She writes down her suggestions and passes them surreptitiously to a male colleague who then puts forward the idea. If she does not adopt this strategy then the idea is inevitably shot down no matter its actual epistemic warrant.

want to be perceived favourably by others in the group); and because of limited “argument pools”. A particularly virulent form of this in recent political debates is “the echo chamber effect”, which is caused through social media algorithms feeding users only messages that they are interested in or agree with – this exacerbates and isolates differing opinions rather than facilitating debate (Iyengar & Hahn 2009).

[4] *Conformity biases*: a significant proportion of individuals are prepared to go along with majority decisions even if they are obviously wrong. Asch (1951) demonstrated this in a series of highly intriguing experiments in which groups of participants were asked to individually make judgements about which two lines of a set were equal in length (see figure 3.5 below). The actual perception task itself had an obvious answer. But unknown to the participant everyone else in the group were actors who occasionally collectively gave deliberately wrong answers to test whether the participant could be swayed to go against their own visual experience in favour of the group majority. Interestingly, more than a third of participants could be swayed in such a manner.

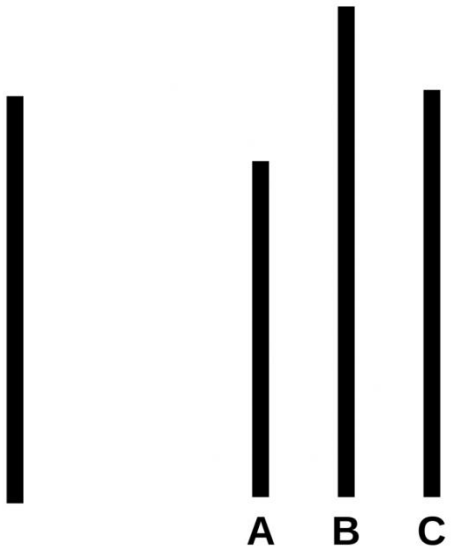

Figure 3.5: The Asch conformity test. Participants are asked to match the lines of the same length in what they think is a visual perception task. They do this task individually but in the presence of confederates who are in league with the experimenter. For test trials the confederates all give the same obviously wrong answer. Asch (1951) found that this caused a significant proportion (over a third) of participants to give the wrong answer.

Lastly [5], there are issues related to *the communication of information in group work*. Information in ecologically salient groups is often not spread evenly so that all members have access to all the necessary information for the task. Stasser & Titus (1985) have showed that the pooling of information in these contexts can lead to instances where key unshared information can be overlooked in group decision-making. Conversely however, Thagard (1993) has argued that this disparity of communication can actually facilitate diversity of opinion within communities thus preventing there from being an absolute consensus. His contentions are supported by computational models (see Hutchins 1995a). Relatedly, von Hippel (1994) has labelled information that is hard to acquire and transmit as "sticky" and this prevents it from being shared and properly utilised in group interactions.

This wide range of cases (see D'Agostino 2008 and Sutton 2013 for overviews) demonstrates that "process losses" are a robust and stable finding. Forsyth defines process losses as reductions in the effectiveness of performance due to "actions, operations, or dynamics that prevent the group from reaching its full potential" (2009, p. 293). For instance, if one asks individuals to pull a rope attached to pressure gauge which measures this force, one will find that a nominal group that is comprised of an additive score of these individual efforts far exceeds a collaborative effort of the same collective. For instance, two men pulled approximately only 1.9 times more than a solitary individual, three men 2.5 times more, and an eight-man team only pulled 4 times more than a solitary individual (Steiner 1966, p. 276; 1972, pp. 32-33). As we shall see in chapter six when we return to the navigation case study, this situation becomes more interesting once we take more complex forms of social organisation into account. This will also enable us to demonstrate that there can also be benefits for well-organised core configurations – expert teams (Williamson & Cox 2014). Collins and Guetzkow labelled these "assembly bonus effects": in which the group is able to achieve a cognitive task in a manner with a productivity that "exceeds the potential of the most capable

member" (1964, p. 58). Mercier and Sperber (2011) note that individuals are particularly bad at judging this group-level benefit (also see Levy 2017).

One might argue at this juncture that these are just social phenomena. Against this deflationist proposal we can respond that some cognitive properties are genuinely present at the wider system or supra-individual level as well (Hutchins 1995a, 2001). A core part of the distributed cognition position is that the social and the cognitive are inextricably intertwined, and that this interaction generates supra-individual cognitive properties that are *emergent* from the individual level (see Hutchins 1995a; Vygotsky 1978, p. 58; Wertsch 1985, pp. 59-61). It is standard within the distributed cognition literature to understand emergence using Wimsatt's (1986, 2007) 'conditions of the failure of aggregativity' (see Poirier & Chicoisne 2006; Theiner & O'Connor 2010; Theiner et al 2010). Bechtel (2008, p. 129) argues that Wimsatt's account provides us with a non-"spooky" definition of emergence – an important feature given the worries that such a term raises for some naturalistic philosophers (e.g. see Ladyman & Ross 2007). Wimsatt identified four conditions (1986, pp. 260-269; also see 2007, pp. 280-281) which he sees as separately necessary and jointly sufficient for aggregativity. As such, emergence is understood as the failure of a systems level property to meet these conditions:

IS: (Intersubstitution) a condition on the intersubstitution or rearrangement of parts. I.e. does the systems level property remain invariant under the rearranging of the parts?

QS: (Size Scaling) a condition on size scaling (primarily, though not exclusively, for quantitative properties) with addition or subtraction of parts. I.e. does the systems level property remain qualitatively similar through a quantitative change in value?

RA: (Decomposition and Re-aggregation) a condition on invariance under the decomposition and re-aggregation of parts. I.e. does the systems level property remain unaltered through the dismantling and reassembling of the system?

CI: (Linearity) a linearity condition that there be no cooperative or inhibitory interactions among parts in the production or realization of the system property. I.e. are there no relations in the system which are mutually inhibitive or cooperative?

Cheon has argued that the notion of emergent clusters of properties is not very useful and that the link between emergence and distribution is not definitive (2014, p. 25). However, I agree with Kirchhoff (2014) that a distributed cognitive system must be organised in a non-aggregative fashion. If this were not the case then the supra-individual level of explanation would be superfluous (also see Huebner 2014, pp. 256-257). Furthermore, I think the notion of emergent properties is useful insofar that it forces the investigation to focus on the differences of the cognitive properties of the system at various levels – which is vital for avoiding the conceptual pitfalls I have outlined in this chapter.

Based on Wimsatt's framework, Theiner and O'Connor propose that supra-individual cognitive properties are dependent on the social organization and interactions among individuals – insofar that they are not reducible to the mere aggregation of their components (2010, pp. 81, 84-85, 105). In light of the above examples, and in response to the critic, the core point is not solely about whether performance at the systems level is enhanced or diminished, but rather that the cognitive property is an *organisation-dependent property that is not strictly an aggregation of its component parts* (Michaelian & Sutton 2013, p. 15). As I outlined in the introduction, a naturalistic approach to cognition focuses on specific cognitive functions – e.g. memory, decision-making, reasoning, but also perception, emotion and motor skills, etc. (see Allen 2017; Hutchins 2001, p.

2068; Kaplan 2012, p. 546; Rowlands 2009, p. 7; 2010, p. 109; Theiner et al 2010, p. 378-379; Wilson 2001, pp. S266-267; 2004, pp. 288-289).

On this basis we can argue that a naturalistic approach to exploring distributed cognition – as advocated and defended by Kaplan (2012), Michaelian & Sutton (2013, p. 10), Sutton (2006; 2015, p. 430), Sutton and colleagues (2010), Theiner and colleagues (2010, pp. 378-379), and Wilson (2001, p. S266, 2004, p. 288) – entails a task-specific approach, which focuses on how these supra-individual cognitive properties are not merely aggregations of their components. In this chapter we have already seen how the simple dyadic interaction of one agent and artefact can lead to a range of complicated arrangements: e.g. how the act of remembering, perception, or decision-making is distributed at the systems level in a manner that renders the cognitive function different (in a variety of ways) in comparison to the individual level. In the following chapters I consider larger core configurations.

In addition to adopting a naturalistic approach to tasks, we must add a refinement that the task-structure has an impact on the organisation of the coordination dynamics that compose the cognitive system. We can use Steiner's (1966, 1972) taxonomy of various kinds of task types and demands to explore how these alter the way processes can be enacted. Steiner distinguished between various forms of tasks in his analysis of productivity, which he saw as determined by three variables which he labelled: task demands; resources; and processes (1972, pp. 6-9). Task demands are the requirements imposed on the agents by the task itself and the normative context within which the task takes place. Task demands determine what resources are relevant and how they are combined and utilised for optimal outcomes. Steiner identifies knowledge, abilities, skills, and tools as resources. These categories can be understood as falling into the following three sections according to various questions we can ask about the task structure (Forsyth 2009, pp. 298-304):

- *Component*: can the task be subdivided?
  - *Divisible*: task has components that can be assigned as specific sub-goals (e.g. playing soccer)
  - *Unitary*: task cannot be divided (e.g. pulling a rope)
- *Quantity versus Quality*: what is more important for the task outcome – quantity or quality?
  - *Maximising*: the more produced by the system the better (e.g. scoring the most goals in soccer)
  - *Optimising*: a correct or optimal product is the goal of the task (e.g. problem solving to find the correct solution)
- *Interdependence/permitted processes*: how are the inputs from the various components of the system combined?
  - *Additive*: all individual components of the system have their inputs added together for the product. (e.g. pulling a rope)
  - *Compensatory*: the product of the system is derived from an averaging of the individual efforts. (e.g. attempting to make a prediction about an event and then taking the average from the various proposals by team members)
  - *Disjunctive*: the system selects one effort from all others as the best. In this type of task, the system has the potential to be as good as its best member. (e.g. brainstorming a solution to a maths problem and picking the best answer)
  - *Conjunctive*: all components of the system contribute to the final task outcome. Depending on whether the task is divisible or unitary the outcome is, respectively, either better than the worst member or equal to the worst. (e.g. making a meal in a commercial kitchen; climbing a mountain as a team)
  - *Discretionary*: efforts of the system are combined according to procedures that are decided by the social group itself. Performance in these kinds of tasks is variable because the group itself decides the norms by which coordination dynamics are dictated. (e.g. a jury)

Although Steiner's taxonomy provides us with a much more nuanced framework for specifying how the components of a cognitive system interact towards a particular task, it does not adequately

provide us, by itself, with a means for distinguishing amongst these various features what is genuine component of these task-centric systems, and what are mere causal background conditions. This is the problem of cognitive bloat.

Davies and Michaelian (2016) claim that task-specificity can tackle the problem of cognitive bloat. Although I agree that task-specificity does provide us with an "ontological flexibility" (pp. 317-318) that can outline multiple levels of analysis whilst avoiding the pitfalls attendant on agent-centric approaches; task-specificity, by itself, is unable to stave off this further challenge. We can see the issue here by turning to Davies and Michaelian's diagrammatic portrayal of the ontological flexibility of task-specificity (see figure 3.6 below). Taking this set up one can consider how various arrangements of components are related to each other at differing levels of analysis. For instance, one could consider how [A] is a cognitive process engaging with an external environment for a specific cognitive task. At a higher level of analysis, one could see both [A] and [B] as the components of a hybrid cognitive system [C] and how it relates to its environment [D]. In turn, these could form a distributed cognitive system [E] at another level of analysis in relation to another task environment [F], etc. (2016, p. 317).

We can make this clearer with a couple of concrete examples. Firstly, returning to Norman's checklist example (1991): in which [A] and [B] are two brain regions, [C] is the brain, [D] is a cognitive artefact and [E] is the integrated distributed system of agent and artefact which can then interact with a wider environment featuring other cognitive systems [F]. Secondly, Hutchins' navigation team case study (1995a): in which [A] is an agent (Plotter), [B] is a cognitive artefact (chart), and [C] is the hybrid system of these integrated components. [D] is another cognitive system in the team (bearing recorder), and [E] is the higher level distributed cognitive system formed of these cognitive systems coming together. [F] is yet another cognitive system that they are

interacting with (pelorus operator). In this set up the navigation team would be a larger conglomeration [G].

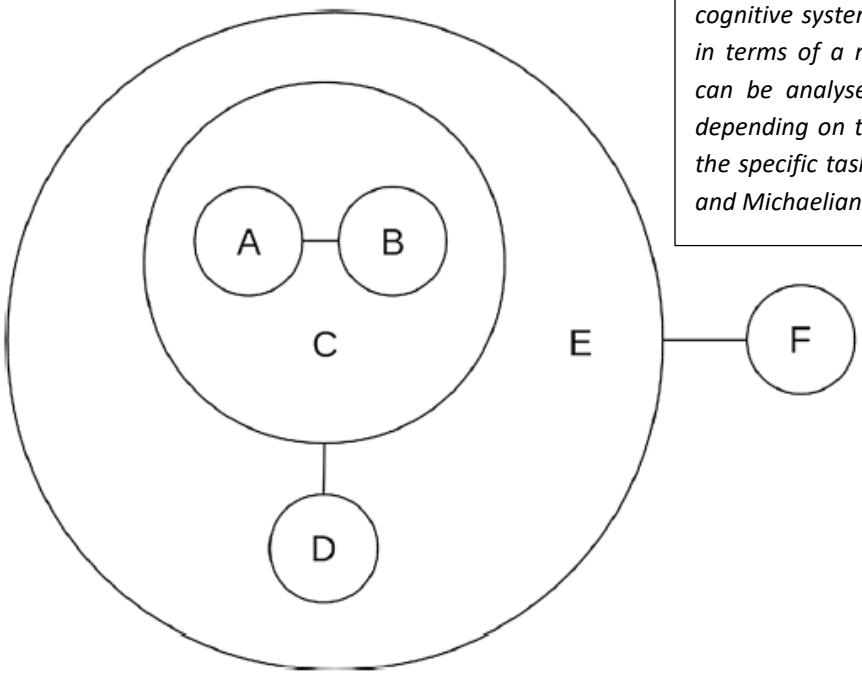

Figure 3.6: Various components of a cognitive system are shown as related in terms of a nested hierarchy. These can be analysed in a range of ways depending on the level of enquiry and the specific task domain. [From Davies and Michaelian (2016, p. 317)]

The advantage of the task-specific approach on display in both these examples is that the agent is not necessarily the centre of the system and nor do we have to understand these higher and lower level systems solely in terms of their agential features. Additionally, it highlights how distributed cognitive systems can be composed of other cognitive systems to form nested hierarchies. However, this diagram does reveal another point of tension here. How are we supposed to understand the interactions or relations between the various components? Like many other proponents of distributed cognition, Davies and Michaelian discuss these solely in terms of information flow (2016, p. 313). As I argued in the previous chapter, this approach is insufficient for tackling the problem of cognitive bloat because it cannot differentiate between genuine components and mere background conditions. This point is especially clear in the nested hierarchy

on display in this diagram. For instance, one can ask why the system stops at the interaction of [E] and [F]. Surely, if it is only an interaction in terms of information flow, then this forms the larger system [G]. And when this larger system interacts with another it will form the larger system [H] and so on; thus, entailing the escalating bloating of the system that Adams & Aizawa (2001) warned about. Given that task-specificity cannot tackle this issue it is in need of supplementation. In the next chapter I turn to the mutual manipulability criterion to stave off this challenge.

3.5. Summary

In this chapter I have begun a positive defence of my pluralistic approach to distributed cognition by outlining the multifaceted importance of the first of the three core principles: *task-specificity*. A proper analysis of distributed cognition requires examining multiple levels of analysis. Agent-centric approaches are prone to carrying-across inappropriate conceptual baggage when they scale both up and down. These mereological fallacies occur in both directions: with agential features erroneously ascribed to supra-individual cognitive properties; and cognitive properties which are genuinely supra-individual over-attributed to the agent. Because these issues rule out an agent-centric approach as appropriate for exploring distributed cognition we must provide an alternative. In these regards, I built upon Davies and Michaelian's (2016) task-specific approach, which avoids succumbing to these mereological fallacies by providing a neutral basis upon which the flexible unit of analysis can be centred.

Furthermore, I supplemented this position by making the notion of supra-individual levels of analysis more concrete. Here I turned to Caporael's (2014) notion of core configurations in conjunction with Steiner's (1972) taxonomy of task types. The conjunction of these concepts provides us with a firm way of centring supra-individual cognitive properties to specific cognitive

tasks in a manner that is robust (Wimsatt 2007) and is emergent from the organisation of the composition of the system (Poirier & Chicoisne 2006; Theiner & O'Connor 2010; Theiner et al 2010). However, this still leaves problems which task-specificity (*by itself*) is insufficient for tackling: viz. how to adequately demarcate the genuine components of the system from mere causal background features. As such, in the next chapter I build on the preliminary discussion here of mechanistic explanations and show how the mutual manipulability criterion can tackle the problem of cognitive bloat.

4: The mutual manipulability criterion

4.1. Introduction

Chapter one outlined the main problems facing a proponent of distributed cognition: viz. the extreme heterogeneity of cases in the wider literature, how this creates problems related to size scaling, and exacerbates the problem of cognitive bloat. Then, in chapter two, I demonstrated the shortcomings of the information-flow view that is standard in the literature for dealing with these problems and thus paved the way for my pluralistic framework. Chapter three began the positive defence of this position by showing how a task-centric approach can negotiate the need for a multilevel unit of analysis whilst avoiding a series of conceptual problems which plague agent-centric accounts. In doing so I went beyond Hutchins in beginning to outline how the flexible unit of analysis – a crucial feature of distributed cognition – could be defined in a more principled manner.

However, I concluded the previous chapter with the recognition that task-specificity was insufficient to tackle the problem of cognitive bloat. In this chapter I return to the central case study outlined in chapter two – Hutchins' (1995a) naval navigation team – to demonstrate an alternative approach to tackling the demarcation problem that is at heart of cognitive bloat: *the mutual manipulability criterion* (Craver 2007; Kaplan 2012; van Eck & de Jong 2016). The mutual manipulability criterion demarcates the boundaries of the unit of analysis in a manner that makes it clear what is considered to be a genuine part of the system as opposed to a mere necessary causal background feature. This principle not only allows us to tackle the theoretical problems associated with bounding a distributed cognitive system and thereby avoiding cognitive bloat, but also provides us with practical insights into actual real-world case studies. As outlined in the general introduction, my aim here is to investigate distributed cognition from a naturalistic and pragmatically motivated

stance. This entails paying close attention to the details of particular case studies to extract not just an analytic definition of what distributed cognition *is* but also what it *does*. I.e. it tells us something of practical interest about the particular target phenomena. However, although it allows us to nullify the dangers of cognitive bloat, the problem of time still remains at large; and I shall conclude this chapter by outlining these issues and why it motivates the need for us to take cultural practices into account.

The chapter is structured as follows: the next section (4.2) outlines the mutual manipulability criterion in general and demonstrates, with reference to some examples, how it nullifies the problem of cognitive bloat. Section 4.3 discusses some important meta-theoretical aspects of this principle in respects to the rest of the framework. I then demonstrate how the mutual manipulability criterion operates in regards to the central case study of Hutchins' navigation team – showing how this principle draws out important practical insights into the phenomenon in question that are otherwise overlooked (section 4.4). The chapter concludes by discussing the limitations of the mutual manipulability criterion in regards to temporal factors (section 4.5). Whilst this has been previously acknowledged, in differing ways, in the literature (e.g. Kaplan 2012; Kirchhoff 2016); I use this discussion to set up the requirements and issues for what the final principle of my pluralistic framework – normative patterned practices – must resolve in the next chapter.

4.2. A mechanist approach to enumerating the components of distributed cognitive systems

Huebner (2014, p. 160-161) notes that there are multiple constant background activities ongoing in human brains that are not involved in the production of personal-level representations for specific

cognitive tasks. For instance, homeostatic processes dedicated to regulating heart rate and respiration, digestion, and proprioception[44]. But whilst Huebner is probably right on this claim (although he notes that it is an empirically open question), we need some method or set of principles for distinguishing between background conditions, which are necessary for the overall functioning of the system, from features which are more specifically tied to constituting a task-specific cognitive process. As I have shown in previous chapters, this failure leads to a slippery slope in which any causal connection to the system is fallaciously designated as a constituent – entailing the problem of cognitive bloat. In order to tackle this question, I turn to the recent work on mechanistic approaches towards explanation in philosophy of biology and neuroscience (Craver 2007; Kaplan 2012).

Bechtel (2008, p. 13) defines a mechanism as a structure that performs a function in virtue of its component parts, operations, and their organisation. Mechanistic explanations are orientated towards explaining how the components of a mechanism operate and are organised so as to generate a particular behaviour (Craver 2007, p. 122). With regards to distributed cognition, the pertinent issue here is the ability to identify the genuine component parts, operations, and the organisation of a putative distributed cognitive system. Using the conceptual apparatus of the previous chapter, these are arranged into nested task-centric hierarchies (I will demonstrate this in reference to the navigation case study below in section 4.4). This is a crucial aspect of any properly formulated and coherent distributed cognition framework that wants to properly account for what human cognition is in the real world.

[44] It is noteworthy that this situation is analogous to the bridge of the naval vessel in which the navigation team is only one department – I discussed the issue of cross-chatter communication in section 2.4 and return to this topic in greater detail in chapter six (section 6.2.1). It also demonstrates that the balkanisation of the mind entailed by a task-specific approach, which I acknowledged at the end of the previous chapter, fits well with the phenomena under investigation. Namely, these basic homeostatic processes entail that even a standard methodological individualist position will have to engage in a selection of only some internal processes from the total set for their unit of analysis. As such, balkanisation at larger core configurations is not a special issue just for proponents of distributed cognition.

A common worry for some critics (and proponents) is that smearing or blurring the boundaries of cognitive systems can lead to the occluding of important distinctions that are vital for properly explicating a target system or phenomenon (see remarks by Clark 1997, p. 163, 2008, pp. 31-32; Craver 2007, p. 189; Eliasmith 2009, pp. 149-150; Haugeland 1998, p. 211). This leads to concerns of critics such as Adams and Aizawa who fear that without sufficient demarcation criterion that if "...anything that is causally connected to a cognitive process is part of the cognitive process, then there is the threat of cognition bleeding into everything" (2001, p. 57). The danger here is that *if* there is not a strict or reasonable criterion for discerning what is and is not a genuine part of a cognitive system, then this threatens some form of "pancognitivism, where everything is cognitive" (ibid). As such, even in distributed cognitive systems – where we are postulating a dense interplay of brain, body, and world – it is crucial to identify different components of the target system under investigation. As Kaplan insists, we must be able to properly specify what the genuine components of a putative distributed cognitive system are in comparison to mere necessary background causal conditions: "the boundaries of a given mechanism are to be drawn around all and only the entities, activities, and organizational features explanatorily relevant to the target phenomenon they produce" (2012, p. 556). Kaplan proposes that the notion of mutual manipulability can resolve this issue.

Mutual manipulability is a naturalistic method for assessing constitutive relations based on mechanistic explanatory strategies drawn from practices in the biological sciences (see Craver 2007). It can be summarised as follows: to see if a component is a genuine part of a system one would ideally be able to make an intervention on the component and see a change in the system as a whole, and reciprocally make an intervention on the system as a whole and see a change in component. Craver states this more formally as follows:

> (i) *x* is part of *S*; (ii) in the conditions relevant to the request of explanation there is some change to *X*'s *Φ*-ing that changes *S*'s *Ψ*-ing; and (iii) in the conditions relevant to the request for explanation there is some change in *S*'s *Ψ*-ing that changes *X*'s *Φ*-ing. (2007, p. 153)

I.e. A component, *x*, is taken to be part of a system, *S*, iff both: an alteration of the component part causes a change in the system, *and* an alteration of the system causes a change in the part. It is important to note that this definition says nothing about localisation – which would preclude a distributed account. Indeed, Craver and Tabery note that mechanisms are not necessarily localisable: "Components of mechanisms might be widely distributed (as are many brain mechanisms) and might violate our intuitive or tutored sense of the boundaries of objects (as an action potential violates the cell boundary). The assumption of localization is often an important heuristic in the search for mechanisms; however, this heuristic often must be abandoned as the mechanism's organization reveals itself" (2015). Craver then goes on to elaborate this simple formulation (see 2007, pp. 155–160). Kaplan (2012, p. 557) summarises this as follows[45]:

> (M1) When ϕ is set to the value ϕ_1 in an (ideal) intervention, then Ψ takes on the value $f(\phi_1)$ [or some probability distribution of values $f(\phi_1)$].
>
> (M2) When Ψ is set to the value Ψ_1 in an (ideal) intervention, then ϕ takes on the value $f(\Psi_1)$ [or some probability distribution of values $f(\Psi_1)$]

Where Ψ is a variable for a higher-level phenomenon present at the systems level to be explained, and ϕ is a variable standing for a lower-level component of the mechanism that putatively underlies the higher-level phenomenon (Ψ).

[45] Although the original formulation of the mutual manipulability criterion was given by Craver, I have here preferred Kaplan's more concise version (but there is very little to distinguish them). Much of the following discussion draws heavily on both Craver (2007, pp. 152-160) and Kaplan (2012, pp. 557-562).

We can demonstrate how the mutual manipulability criterion works by considering some examples. Kaplan (2012, pp. 565-568) reassess two of Clark's classic thought experiments from the extended mind literature: Bluefin tuna swimming behaviour (1997, pp. 219-220), and Otto and his notebook (Clark & Chalmers 1998, pp. 12-18). Since much has been written on the latter case, I will only consider the former here. Bluefin tuna are known to exploit local properties of sea currents to achieve speeds that exceed what is theoretically possible with just their musculature. In this case the higher-level phenomenon (Ψ) to be explained is "[a distributed] tuna-local-ocean-environment propulsion mechanism" (Kaplan 2012, p. 565). The putative component (ϕ) under investigation (about whether it is a genuine part of this system or just a mere causal background feature) is the local environmental features: vortices and eddies. Kaplan's suggestion is that the mutual manipulability criterion offers us a way of coming to a concrete answer in this matter by using two "ideal interventions" that experimentally test how effects propagate in the system. A bottom-up manipulation (M1) alters properties of the component to see if it engenders a change in the overall system behaviour. And a top-down manipulation (M2) engages the system behaviour to see if this elicits a change in the component. So, in the case of the Bluefin tuna, a bottom-up experiment influences the local environmental features to see if they have an impact on the "distributed tuna-local-ocean-environment propulsion mechanism", and the top-down experiment engages the system behaviour to see if it causes a reciprocal change in the component properties. *Both manipulations are required* because this evaluates the constitutive inter-level relations between the part and the whole[46]. A successful experiment in only one direction is not sufficient for ascertaining

[46] Although not central to our argument it is interesting to note that a mechanistic approach to defining the constitution of a cognitive system avoids Ross and Ladyman's (2010) recent critique. They claim that composition is an ill-founded metaphysical posit with which to consider cognition because it is not a well-motivated according to our best available science. However, mechanistic accounts avoid this critique because whereas Ross and Ladyman are attacking constitution as an abstract universal category; as I outlined in the previous chapter, mechanistic composition is not a global designation and only apples to specific systems under investigation (Craver 2007).

whether a component is part of the system. It is only when both M1 and M2 are satisfied that we can be clearer in what we are claiming. For example, in the case of the Bluefin tuna, the criterion allows us to identify the specific aspects of the environment that can be picked out against the general background causal milieu.

In chapter two, I outlined the case of the blind man's stick (BMS) (Bateson 1972; Hutchins 1995a, 2010a; Malafouris 2008, 2010, 2013). The issue with this case was that an information bottleneck was not sufficient for discerning genuine components of the system from the mere causal background. As such, the putative distributed cognitive system ballooned out to include the road. Whilst it must be acknowledged that the sounds produced from probes by a blind person on different materials provide information (Downey 2016), if we are to ambiguously include both these information channels and their sources in this distributed cognitive system then it escalates outwards excessively to include almost anything the agent is interacting with.

Such a state of affairs leaves us unable to properly demarcate genuine components of the system from features that are just causal background conditions (e.g. gravity, homeostatic processes, atmosphere, etc.). The mutual manipulation criterion allows us to interrogate the putative composition of the BMS system and identify the genuine components. The higher-level phenomena (Ψ) or task we wish to explain is the locomotion of the blind person. It is claimed that this involves the agent (specifically their auditory[47] and motor control system and relevant muscles, etc.), the stick, and the street.

We can now apply the two ideal interventions to this system: firstly, a series of bottom-up manipulations (M1) on each of these putative components to see if they engender a change in the overall behaviour of the system. Arguably, interventions on the agent or the stick will drastically

[47] Interestingly, the neural regions correlated with visual inputs (e.g. calcarine cortex (V1)) are active for auditory sensory stimulus in blind agents (see Downey 2016 for an overview).

alter the performance of the system. Likewise, if one were to inhibit the agent's interaction with the street or change its materiality this would also alter the overall system behaviour. However, as with the above example, this is not sufficient to claim that they are genuine components of the system; we must also conduct a top-down manipulation (M2) and engage the system in its task behaviour to see if this produces reciprocal changes in the putative components. The act of walking in the BMS case will obviously alter the state of the agent (in regards to their location and states of musculature; and auditory, motor-control, and spatial-reasoning neural correlates of the nervous system, etc.), but also involves the operation and adjustment of the stick (via the plasticity of the body schema as discussed in chapter two (see Iriki & Taoka 2012; Malafouris 2008, 2010; Menary 2013b)).

In contrast, there is no reciprocal change generated in the street per se. But we must be clear here: the way that the stick interacts with the street and probes the environment to generate different auditory stimulus and tactile feedback is not altering the general state of the street (except in the unusual cases where the blind agent uses excessive force on an extremely weak surface). Instead, one can argue that these probing actions create a looping interaction of brain, body and *specific features* of the world[48]: these are specific sounds that can be used for echolocation and also tactile sensations. Both are mediated and generated by the stick in conjunction with the environment. As with the tuna motile system above, the mutual manipulability criterion allows us to identify specific fine-grained aspects of the environment that can be reasonably identified as genuine components of a distributed cognitive system. By clearly differentiating between these and mere causal background conditions, the mutual manipulability criterion provides us with a robust means of staving off the challenge of cognitive bloat. It provides the researcher with a means of drawing a clear line and avoiding the slippery slope that critics think is inherent in externalist views.

[48] I take the notion of looping interactions from Menary (2012, 2013a).

Nullifying the problem of cognitive bloat is not the only appealing aspect of what this criterion can do for distributed cognition – as I will now discuss in the next section.

4.3. Preliminary remarks on mutual manipulability and distributed cognition

Before turning to how this principle operates in regards to the navigation case study, it is worthwhile enumerating the many reasons that make mutual manipulability a thoroughly attractive criterion and clarifying some conceptual issues. There are numerous ways in which the mutual manipulability criterion is an excellent supplement to task-specificity for rendering the dual notions of distributed cognition (a flexible unit of analysis and the possibility of supra-individual cognitive events) in a more thorough and principled fashion. *Firstly*, because it is drawn from practices within science it is a well-motivated principle to adopt from a naturalistic standpoint (Craver 2007, pp. 2-3) and therefore fits the desiderata I outlined in the general introduction to the thesis. *Secondly*, as noted in the previous chapter, Craver points out that mechanistic explanations – including cognitive mechanisms – often transgress compartmental and physical boundaries such as the skin and skull (2007, p. 141; also see Clark 1997; Haugeland 1998; Wilson 2004). Thus, because mechanistic explanations are not necessarily localisable (Craver & Tabery 2015), the mutual manipulability criterion is amenable to distributed cognition[49].

Importantly, a *third* feature is that the mutual manipulability criterion is neutral insofar that it favours neither internalism nor externalism, and is non-question begging because it requires "no special assumptions about the nature of cognition" (Kaplan 2012, p. 557). Given that van Eck and de Jong (2016, p. 12) have rightly criticised both sides in the debate for having a priori assumptions that

[49] A point of clarification: the use of the term 'local' here does not clash with the localised notion of levels that were discussed in the previous chapter (section 3.4) because the latter usage means simply that mechanistic levels are 'not-globalisable' whereas the current discussion is about the real possibility that particular mechanisms need not necessarily be physically situated in a particular (small) location.

neither will accept, this is especially welcome. *Fourthly*, Kirchhoff (2016, p. 8) points out another related aspect: because the mutual manipulability criterion is both neutral and lacks any special assumptions about cognition it therefore also avoids issues about the 'Mark of the Cognitive'. Adams & Aizawa (2001, 2010) attempted to solve the problem of cognitive bloat and stem the challenge of externalism in defence of internalism by formulating a criterion in which cognition is identified as representations with underived content. However, many have criticised this approach for numerous reasons – e.g., Allen (2017) has recently argued that the Mark of the Cognitive is not naturalistic because it relies on a distinction (between derived and underived content) not found within the cognitive sciences and also fails to attend to a wide range of empirical work on cognition. Avoiding this debate is incredibly beneficial.

A *fifth* reason is that the mutual manipulability criterion can give a concrete answer to the question of whether a putatively extended component is part of a mechanism or just a necessary background condition (see Kaplan 2012, p. 563; Kirchhoff 2016, p. 8). I.e. it is theoretically testable – for instance, it has recently been applied to spider cognition (Japyassú & Laland 2017). Kaplan criticises other putative measures for not being able to tackle this problem satisfactorily, which consequently makes them effectively empty. By making claims about distributed cognition empirically tractable, it is Kaplan's aim to bring these debates in line with mainstream discussions in cognitive science. This bolsters the views of Sutton (2015) and others (discussed in chapter one, section 1.5) that distributed cognition is consilient, to some extent, with mainstream views in cognitive science rather than a radical opponent. A *sixth* factor in favour of adopting the mutual manipulability criterion is that it is only a sufficient condition as to whether a component is part of a mechanism (Kaplan 2012, p. 560; also see Craver 2007, p. 141). Because it is not a necessary condition, there are limitations as to what conclusions can be drawn by failures to meet it. This

limitation on its scope makes the mutual manipulability criterion amenable to operating as part of my pluralistic framework (I discuss this in greater detail in the final section of this chapter).

A *seventh* notable feature is that the mutual manipulability criterion fits the flexible unit of analysis which I argued is one of the hallmarks of distributed cognition (see section 1.5) – because its adoption entails *that the boundaries are set through inquiry rather than pre-defined*. An *eighth* and related conducive aspect here is the fact that mechanistic explanations are multilevel (Craver 2007). As Bechtel (2009) puts it: insofar as a mechanistic approach involves the functional decomposition of a mechanism into its parts and also its re-composition to see how it is organised, this necessitates an integrated multi-level perspective whereby one is not only "looking down", but also "up", and even "around" to locate the mechanism in its context. This matches nicely with the emphasis on the importance of taking care to treat multiple levels properly on their own terms. In chapter three I showed how these multiple levels of analysis can be articulated in a principled fashion that avoids a range of conceptual problems by using a task-based approach.

As such, there is much to admire about this proposal. But I must note that the use of mutual manipulability to demarcate the boundaries of distributed cognitive systems also raises the question of whether the adoption of such a strategy will move my project back towards focusing solely on philosophical questions. This definitional work in articulating how a cognitive system is bounded is vital in order to avoid the problem of cognitive bloat. But my pragmatic claim is to reject the notion that such work equates with a cessation of the interrogation. Instead, in addition to this definitional work, a higher understanding of the system under investigation is achieved through exploring the practical implications of treating the system as distributed in various *interesting* ways (Hollan et al 2000; Hutchins 2001, 2006). As I argued in the first two chapters, the issue for Hutchins' own presentation of what these interesting ways *are*, is a general lack of principled means by which to

delineate what this actually entails in a sufficiently robust manner. My alternative formulation has been to develop a set of criteria for how cognition can be distributed in interesting ways across both space and time, and how the flexible unit of analysis can be established in a principled manner.

A demonstration of how the mutual manipulability criterion has pragmatic benefits can be clarified by returning to the case studies in the philosophy of science literature that I critiqued in chapter 2 (section 2.6) for lacking a sufficient amount of ethnographic detail. I identified two prominent case studies: Magnus' (2007) account of double-blind medical trials, and Giere's (2006, 2007, 2012) account of the Hubble telescope. I will now discuss each in turn. In agreement with Huebner (2014), I think that ethnographic detail cannot be dismissed as philosophically superfluous – indeed it is crucial to adequately elucidate the genuine components of the task-specific system under investigation if one wants to properly apply both M1 and M2 in a cogent manner. Consequently, the fact that the mutual manipulability criterion forces the investigator to pay very close attention to the details of the putatively distributed cognitive system can be designated as another desideratum of this principle. Attention to detail is where a distributed cognition analysis can provide insights into the target phenomena that would otherwise be overlooked by other frameworks (e.g. methodological individualists).

A prominent case study in the philosophy of science literature is Magnus' (2007) double-blind medical trial[50]. Unfortunately, his account has a general lack of details in regards to the components of this putative distributed cognitive system. Although Magnus provides an excellent analysis of the task-structure, he provides no details as to: [1] how many agents (nurses, doctors, and patients) are involved in the trial; [2] how they are related to each other (except to note that there are instances as to how the agents have knowledge about which patients are receiving the

[50] For instance, see discussions in Brown (2011), Cheon (2014), and Kerr and Gelfert (2014).

real treatment and which are the controls receiving a placebo); and [3] what and how many epistemic tools are necessary for carrying out the task. Since there are no details provided here, it is impossible to apply a bottom-up intervention (M1) on a component of the system to see what effect it might have on the wider systems shared task performance. It is also not possible to monitor the effects of any top-down manipulation (M2) of the system behaviour since none of the components have been sufficiently specified. This example demonstrates the importance of specifying the components of the system under investigation. Huebner (2014) argues convincingly that a proper investigation of distributed cognition must engage in a functional decomposition of the target phenomenon in order to bring this field in line with mainstream cognitive science. As such, we currently cannot even judge whether this is or is not a case of distributed cognition.

Another prominent case is Giere's (2006, 2007, 2012) analysis of the Hubble telescope. Although he does provide more detail in regards to the many institutions and locations of teams working on this huge scientific project, he is still lax in regards to the details of how many agents are working together on this collective scientific project. Heersmink (2016) has raised doubts about whether it is possible to adequately outline the coordinate dynamics of these kinds of gigantic systems. I think this is an empirical question so I will not comment (except to note that Nersessian and colleagues' (2004) excellent and rich investigation of biomedical engineering laboratories itself required an interdisciplinary team). The above example has already demonstrated how the lack of details here is detrimental. But my concern in this case study is to show the practical strengths of the mutual manipulability criterion in unpicking conceptual problems.

Giere's concerns in this case study are to explore Knorr Cetina's suggestions that the individual knower is replaced in "radically collaborative" scientific projects[51]. In high energy physics

[51] I take this term from Kukla (2012).

experiments she claims that the individual knower is replaced by the communicative coordination dynamics in the radical collaboration; so, although no individual knows how the knowledge is produced in its entirety the epistemic agent becomes the experiment itself: "…the individual has been turned into an element of a much larger unit that functions as a *collective epistemic subject*" (1999, pp. 167-168, 178. My emphasis). In the previous chapter I rejected the notion of projecting agential features onto supra-individual configurations because of the conceptual problems this generates. Giere (2006, p. 109) also, understandably, sees Knorr Cetina's claim as provocative. Namely, that distributed cognition entails the "erasure of the individual as an epistemic subject" (Knorr Cetina 1999, p. 171). Leaving aside the fact that the nesting of individual human agents does not need to nullify their importance (as I demonstrated in the previous chapter); Giere takes this claim to entail that we cannot point to any one individual or even a small group as responsible for producing knowledge claims in these contexts (2006, pp. 109-110). This is the problem of distributed epistemology (to which I return in the general conclusion of the project).

Giere rightly, in my view, rejects the notion of an extended "super agent". He turns to the example of the Hubble telescope to try and demonstrate what he sees as the confusions taking place here:

> Consider again the Hubble telescope. As a distributed cognitive system it extends at least from the telescope in orbit through a series of intermediaries to the Space Telescope Science Institute in Maryland. If one adds the Abell 1689 cluster of galaxies used as a gravitational lens, the system extends 2.2 billion light-years out into space. Are we to say that its mind extends from the telescope in orbit to Maryland, or 2.2 billion light years out into space? (2006, p. 112)

Giere uses this intuitively absurd scenario to gain leverage for the following question: "Is the Hubble system as a whole epistemically (as opposed to just causally) responsible for the final conclusions?" (2007, p. 317). I would like to note that the question of whether the entire system is involved in, and

responsible for, making and producing a particular knowledge claim is entirely different in kind and degree from a claim about whether the entire system – including a galaxy billions of light years away – is all one colossal mind. I think Giere's set up of the issues in this manner involves a rhetorical sleight of hand that attempts to use the absurdity of the latter question to render the former equally insensible.

However, the first two principles of my pluralistic framework that I have so far outlined provide us with a means of approaching these questions in a sensible fashion and demystifying them. Firstly, Giere (2012) follows Clark's (2008) agent-centric approach ("cognition is organism centred if not organism bound"). In contrast, by beginning with a task-centric analysis there are no good grounds for transposing agential features beyond the personal level – to do so is a mereological fallacy (Drayson 2012). Additionally, by centring our investigation on the specific task we are able to note the various levels and parts that are involved in this system – thus, there is no need to erase the individual agents who are components in this wider system. Secondly, we can examine these components using the dual interventions of the mutual manipulability criterion to assess whether any of these are genuine components of this distributed system or just mere causal background features.

We need not go through all the components (which would admittedly be a huge task in itself) to get an interesting result. Consider only the Abell 1689 cluster of galaxies that Giere emphasises in their important role in the system (i.e. the gravitational lensing effect). If this cluster of galaxies is a genuine part of the cognitive system then it extends out 2.2 billion light years into space – an obviously absurd conclusion. However, using a bottom-up intervention (M1) we can note

that if one were to obscure or omit Abell 1689[52], then this would have a dramatic impact on the successful operation of the telescope. However, by itself this is not sufficient for claiming that this is part of the distributed epistemic production system that is the Hubble telescope. For this claim to be justified we also need to make a top-down intervention (M2). And here it will be recognised that there is no intervention or manipulation of the entire system that will engender any specific change in Abell 1689 (i.e. there is no behavioural change or implementation of the coordination dynamics of the Hubble telescope system that will ever have an impact on this cluster of galaxies). As such, we have good reasons for postulating that, although highly important, Abell 1689 is a mere causal background feature of the system. There is obviously far more to be said on this topic. But for our immediate concerns I have demonstrated that the mutual manipulability criterion has practical benefits.

Having shown the practical as well as philosophical benefits of this criterion, I now outline a few last points of order before moving on to the navigation case study. Bechtel (2008, p. 15) identifies two ways of engaging in mechanistic decomposition: decomposing a mechanism into parts (structural decomposition) and into its operations (functional decomposition) – but this is not to deny that the parts and operations of a system are not intimately interwoven. Mechanistic decomposition is focused on explicating the underlying processes that are responsible for higher-level phenomenon. Craver (2007, p. 188) states that to give a functional decomposition of a system into its working parts is to give a mechanistic explanation of the system. Functional decomposition is not simply about cutting a system up into little pieces. Craver gives the example of cutting a rat brain up into one-centimetre cubes – this is unlikely to result in a series of components for various cognitive functions such as semantic memory (2007, pp. 187-188; also see Haugeland 1998, pp. 211-

[52] This could be achieved simply by not including them in the information process itself; or by a much more complex sci-fi-esque procedure involving the use of interstellar gas to conceal them from the system.

214 who makes a similar point using the example of a TV). Instead, Clark (2008, pp. 14-15) describes what he calls "distributed functional decomposition" as the analysis of a system "in terms of the flow and transformation of energy, information, control, and where applicable, representations".

It is important to note that whilst I rejected the possibility of using information flow (*by itself*) as a means of demarcating the unit of analysis in chapter two, I did not reject information per se wholesale as part of our analysis (since to do so would entail a paradigmatic shift away from cognitive science, which I rejected (see Sutton 2015)). Instead, following Menary (2013a, p. 28; 2013b, pp. 353-354), we can cash these interactions out in two ways: [1] coordination dynamics; and [2] the ways in which the former are embedded, shaped, and transformed in the cultural-cognitive niches which agents partially create, maintain, and manipulate. For our purposes here, the key points are that these two ways of thinking about distributed cognitive systems are not only compatible, but are also both mandatory for a proper understanding of human cognition in the wild (also see Kirchhoff 2014). Both ways of understanding the relations of cognitive systems provide us with complementary pieces of the puzzle.

The significant factor here is *time* and the ways in which cognition is distributed across three different temporal scales: [1] a *synchronic* factor – the dynamic interactions in the contemporary event (Bietti & Sutton 2015; Kirchhoff 2012, 2016); and two longer term *diachronic* factors, [2] the developmental trajectories by which agents acquire the habitual embodied practices that mediate the coordination dynamics (Menary 2012, 2013b, 2014, 2015), and [3] the cognitive-historical factors whereby the contemporary epistemic landscape is shaped by previous generations (Bietti & Sutton 2015; Dutilh-Novaes 2013; Hutchins 1995a, 2008; Sterelny 2003; Tomasello 1999). At the end of this chapter (and in the next) I shall go on to argue, drawing on Menary (2007a, 2015), that if we are to correctly understand the synchronic coordination dynamics that are clarified by mechanistic

principles, then we must understand how they are shaped and "enculturated" by situating the system within the wider cognitive-cultural niche. The notion of "normative patterned practices" (Menary 2016) enables us to do so whilst also avoiding problems of methodological bloat that could potentially derail the pluralistic approach I am adopting here. This will also be the focus of the next chapter.

This concludes the preliminary remarks. Having demonstrated the many desirable features of the mutual manipulability criterion, I now show how it operates in regards to our central case study: Hutchins' navigation team.

4.4. Decomposing the navigation team

It is interesting to note that Hutchins' often used mechanistic terminology in describing the navigation team (1995a, pp. 131, 135, 138, 177, 192, 203, 225). This is suggestive – but nothing more – that he would be conducive to a task-driven functional decomposition of the navigation team to reveal the nested hierarchy of this core configuration. In chapter two (section 2.4) I broke the overall team performance of the fix cycle down into three subsystems which are focused on specific subtasks in the overall task. This again demonstrates the importance of task-specificity (chapter three) – since it allows us to properly identify putative components to be interrogated using the mutual manipulability criterion. These subsystems are as follows: [A] the plotter and how he interacts with the chart and other epistemic tools to construct lines of position and make decisions about future landmark targets; [B] the communication network with the bearing recorder at the hub; and [C] the two pelorus operators and their observation devices. I have summarised these in the table 4.1 below:

Table 4.1: summary of the three key components of the navigation team and the various subtasks of the overall fix cycle that they perform		
No.	Subsystem components	Tasks and processes
A.	{Plotter, hoey, chart, bearing record}	Computation and decision making
B.	{Bearing recorder, bearing record, phone, watch}	Communication, time-keeping, working memory (exogram), buffering, recording
C.	{Pelorus operator, alidade, gyrocompass}	Observation

Adopting a multilevel analysis, I first use both bottom-up (M1) and top-down (M2) interventions to discern whether these subsystems are themselves distributed cognitive systems before moving on to the whole team – this puts into practice the nested hierarchy I discussed at the end of chapter three. In chapter two, I queried whether the focus by Hutchins (1995a, 2001, 2010a, 2014) and others on information flow (the propagation of representational states across various media) was sufficient for demarcating the navigation team as a genuine case of distributed cognition. Applying the mutual manipulability criterion provides us with a principled means for properly demonstrating that the navigation team that performs the fix cycle is a distributed cognitive system.

Before discussing each of these components I will briefly recap how the fix cycle is performed by the navigation team (see figure 4.2 below). Firstly, the plotter chooses at least three landmarks for the next position fix. These are then conveyed verbally to the two pelorus operators by the bearing recorder. The pelorus operators acquire these landmark targets in alidades and wait for the bearing recorder to call “mark” before relaying the three-digit headings for each of these. The bearing recorder jots these down in the bearing record. The plotter takes these written numerical representations and reformats them as analogue angles on the hoey and then brings these into conjunction with the relevant features of the chart to construct lines of position. Where the lines of position intersect, a position fix is established. From here the position of the vessel is ascertained using dead reckoning and new landmarks are chosen and the cycle begins again. I now

consider the subcomponents of the fix cycle before considering the whole system as composed of these nested subsystems.

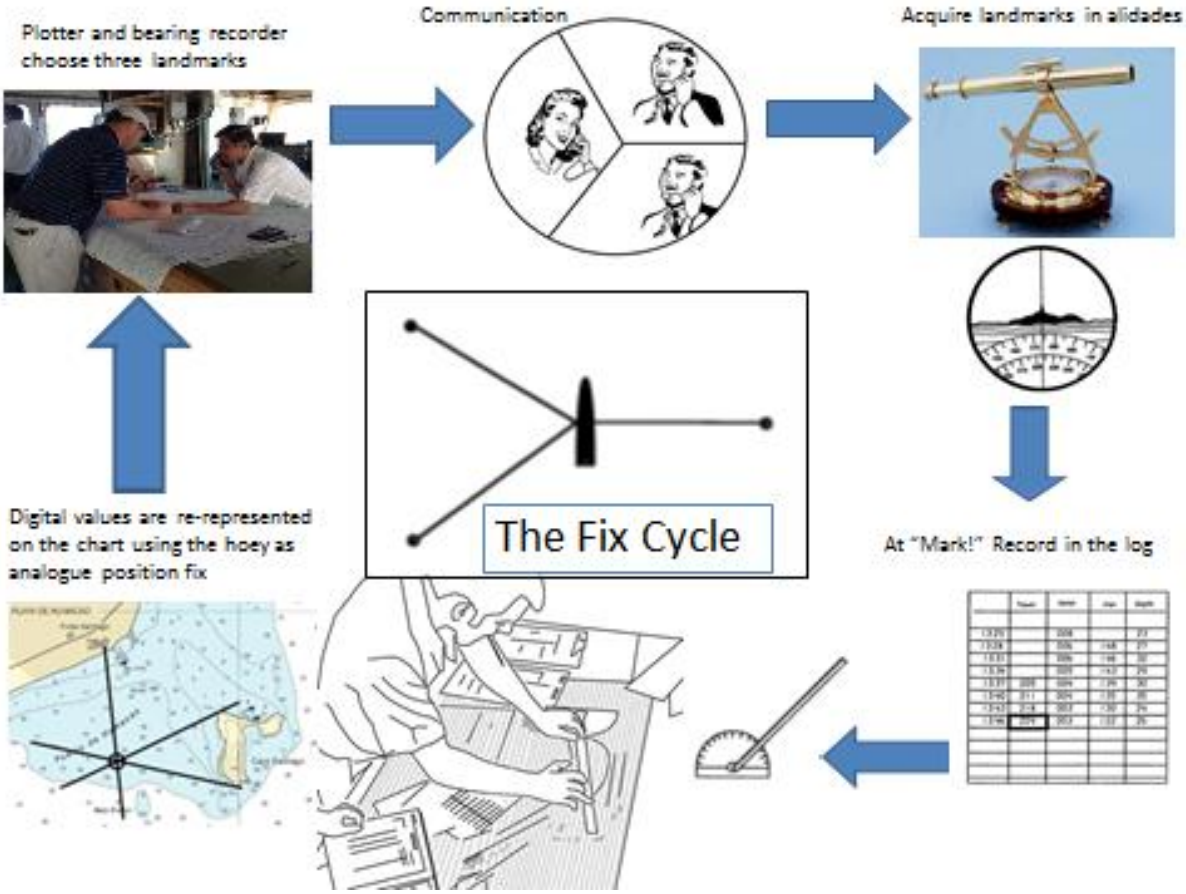

Figure 4.2: The fix cycle. Beginning in the top left panel and moving clockwise, the plotter and bearing recorder choose three target landmarks. These are then relayed to the two pelorus observers by the bearing recorder. The pelorus observers acquire the landmarks and when the bearing recorder calls 'Mark' these are relayed as three-digit numbers. The bearing recorder writes these down in the bearing log book. The plotter takes these digital representations and reformats them as analogue representations on the movable arm of the hoey. This is transposed onto the chart so that the three measurements triangulate the position of the vessel. This is then used to project the future positions of the vessel using dead reckoning. And this is also used as the source for beginning the next cycle.

A. The Plotter: in the first nested subsystem, the plotter takes digital numerical representations that have been recorded in the bearing log and reformats these as analogue angles on the movable arms of the hoey. These are then aligned with the matching landmarks on the chart to construct lines of position. Three lines of position intersect to triangulate the position of the ship. We can identify the cognitive task in this instance as computation – this is the higher-level phenomenon that we are aiming to describe by reference to the parts of this subsystem: viz. the agent, the chart, the hoey, and other calculating devices such as paper and pencil, and calculators. The bottom-up intervention (M1) on this subsystem targets these components to see if it elicits a functional alteration at the level of the system as a whole. So, we can imagine that one or other of these devices is missing or broken and see if the overall behaviour of the system drops. Arguably, omitting any one of these (with the possible exception of the paper and pen for disposable calculations, and calculator) will drastically reduce the performance of the system (omissions or breakages in pen, paper, and calculator will perhaps just slow down the process and might increase errors). But merely conducting a bottom-up intervention (M1) does not yet tell us whether these factors are more than just necessary background conditions for the system[53].

In order to decide this question, we must also conduct a top-down intervention (M2) in which we engage the system in its target behaviour (computation) to see what changes this causes at the level of the components. Interestingly, in this case I think the mutual manipulability criterion does pull apart a number of these epistemic tools so some are designated as just mere causal background features whereas others should be regarded as genuine constitutive parts of this subsystem. For instance, in the process of achieving a position fix the plotter manipulates and alters

[53] One could object that if we took the tools away one-by-one task performance would degrade to the point of inability. As such, although none of the tools are necessary on their own, they are jointly sufficient. However, it could be argued that even with this point in mind we can still recognise that certain tools are more important for others in that their omission would lead to much greater detriments to behaviour. For instance, the chart is much more important than any of the other artefacts. To remove this device would make a position fix nigh on impossible.

both the hoey and the chart. So, at the level of the components the plotter, hoey, and chart are all altered by a top-down intervention (M2). But the bearing log and other epistemic tools are unaltered. Instead, we can see these as just mere inputs to this particular subsystem. Thus, by using both M1 and M2 we have good and principled reasons for regarding the conglomeration of plotter, chart, and hoey as a distributed cognitive system (within the wider overall system). On the other hand, the plotter also uses the chart to make decisions as to future landmarks. In this case, it becomes apparent that only the chart is part of a distributed cognitive system for this task-relative aspect (and sometimes the bearing recorder is involved as well – in this case the task becomes a group decision-making discussion). None of the other epistemic tools are utilised in this task. The plotter consults the position fix on the chart and computes the future location using two "precomputational" cognitive artefacts (the three-minute rule and the planned route)[54]. From here he then makes decisions about which landmarks will best provide a position fix

B. The Recorder: the next system involves the communication network at the heart of the navigation team. The core part of this role in the navigation team is what Hutchins terms "buffering" (1995a, pp. 194-195). As will be discussed in more detail below (and in chapter six), the various subtasks of the fix cycle take differing lengths in an extremely noisy environment. As such, the role of the recorder is to buffer these temporal factors and to avoid errors that could be caused through cross-talk (e.g. if the pelorus operators reported directly to the plotter without meditation – as Brown (2011) mistakenly puts it – this could lead to confusion if the plotter was working on a different position fix, etc.).

Building on from this, the dual interventions of M1 and M2 suggest a split in what parts of this system can be regarded as properly constitutive as opposed to necessary background features.

[54] I return to these in chapter six since they involve diachronic temporal factors that would overly complicate the current discussion.

In bottom-up manipulations (M1) we can see that alterations of the watch, log book, or phone will have an effect on the system behaviour as a whole (as will interventions on the agents). But when we conduct a top-down intervention (M2) to engage the subsystem in its main task behaviour (communication and recording) we can see that only the agents and the log book are picked out by this. No intervention at the system level alters the watch or phone in any meaningful sense. As such, we can see these as mere background phenomena – even though it should be noted that keeping track of the precise timing is essential for properly conducting the fix cycle. But, as is noted by Kaplan (2012, pp. 561-562), one could make similar points about heart-rate, oxygen supply and other autonomic and physiological systems that are essential for human cognitive functions (also see van Eck & de Jong 2016, p. 19). Whilst M1 interventions on these subsystems and features would engender changes in the overall performance of the agent, there are no reciprocal M2 interventions by which we engage the agent in a specific task-behaviour and it elicits a specific alteration in these features. One can imagine a possible scenario in which all the crews' watches break and the recorder is tasked with trying to count out the minutes as accurately as he can – performance would be extremely depleted but not impossible. Likewise, in the story with which Hutchins begins *Cognition in the Wild* in which the ship losses power (1995a, ch1) – and thus also the loss of the steam-powered phone network – he recounts how the recorder has to physically walk to-and-from the bridge and the observation posts in order to get bearing reports (which is also similar to how the task is performed by a solitary agent as I discussed in chapter two). Walking only effects the state of the agent not the phone itself. As such, the watch and phone are just necessary background features rather than genuine components of the cognitive system[55].

[55] One possible objection could be the adjustment of time on the watch. This perhaps presents us with a fuzzy or grey area. For instance, in some cases the fix cycle is performed at a different rate – this would involve making changes to recording time on the watch.

C. The Pelorus Operator: the third nested subsystem picked out by a functional decomposition is the observation of landmarks by the pelorus operators using their specialised telescopes (alidades), which are calibrated to the gyrocompass. As with the previous two cases, the mutual manipulability criterion again identifies only some of the epistemic tools as proper components of the distributed cognitive system. Again, (M1) bottom-up interventions on all parts of the system affect overall performance – indeed, the loss of the gyrocompass is an incredibly severe real-life accident that is documented extensively by Hutchins (1995a, ch1 & 8). For instance, the pelorus operators have to make measurements of magnetic north rather than true north (and these measurements then have to be converted into true north by the plotter – see below). However, there is no activity at the level of the system as a whole that alters the gyrocompass, and as such it fails to meet the top-down intervention (M2), so it should not be considered more than a necessary background condition (in which necessary is understood as being relative to the normal functioning of the system). Instead, it picks out the alidade which the observer physically manipulates in order to gain a measurement as a proper part of this distributed cognitive system.

We have now considered all the subsystems that comprise the fix cycle. Moving up the nested hierarchy (shown in figure 4.3 below), the question now is: what happens when we apply the mutual manipulability criterion to the system as a whole? This is a crucial point because deflationary positions have suggested that one can consider the navigation team as a set of discreet but separable units. Wilson (2005, p. 234) has argued that the fix cycle is more accurately understood as an instance of what he calls "the social manifestation thesis" – that there are certain cognitive properties at the individual level which are only present in certain social organisational structures but which lack any group-level cognitive activity (also see Wilson 2001, 2004). This would make it an *aggregate* of smaller units akin to Clark and Chalmers' (1998) notion of the extended mind (e.g. [plotter and chart] plus [recorder and log book] plus [pelorus operators and alidade]). The mutual

manipulability criterion allows us to seriously and fairly assess Wilson's deflationary position by considering hypothetical M1 and M2 interventions of the overall fix cycle. Arguably, a bottom-up intervention (M1) on any of these subsystems [A, B, or C] would have dramatic impacts on the behaviour of the system as a whole to complete the fix cycle. As figure 4.3 below shows, the performance of the fix cycle is constituted by these component systems. The fact that the entire fix cycle can be performed by a solitary individual (as I noted in chapter two) does not diminish this point in this context because the time-constraints make it physically impossible that the task could be completed by a solitary individual[56].

Figure 4.3: The nested hierarchy of the navigation team's performance of the fix cycle. Several levels of core configurations are shown. These are subsystems that constitute the overall performance of the shared task. [Diagram adapted from Craver (2007, p. 189)]

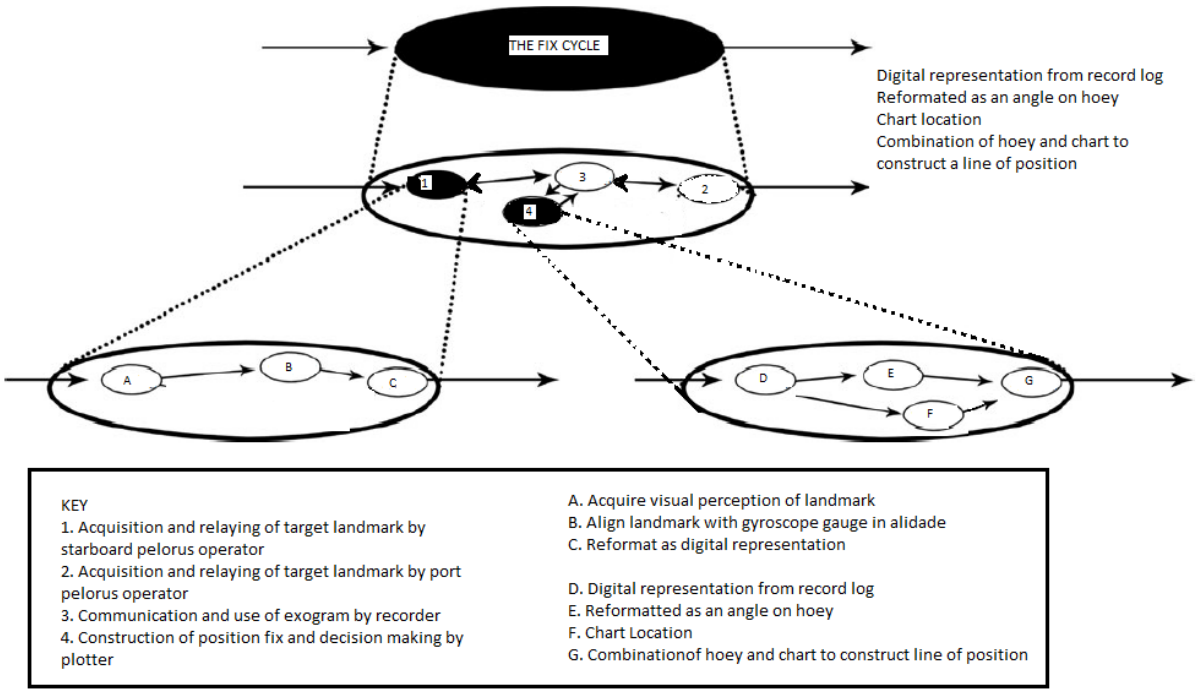

[56] I return to the difference between individual and group performance of the fix cycle in chapter six.

An example clarifies this. The failure of the gyrocompass causes a shift in the division of cognitive labour with the bearing recorder having to perform more of the computational load normally assigned to the plotter[57]. The lines of position require more computational work as each three-digit bearing has to be recalibrated to true north from a measurement of magnetic north (including local distortions). This involves an extra four steps of computation that can be remembered through the mnemonic "Can Dead Men Vote Twice" which stands for 'C+D=M, M+V=T' (compass heading plus deviation equals magnetic north; and magnetic heading plus variation equals true heading). Hutchins notes that although this computational structure is well-known it trebles the workload (1995a, p. 320). Additionally, as mentioned above, the loss of power in this scenario makes communication much more labour intensive as the bearing recorder and pelorus operators have to physically meet each other in order to relay the three-digit bearings.

However, by itself this kind of bottom-up manipulation (M1) does not decide whether these are genuine components of a distributed cognitive system. This issue is decided by a top-down intervention (M2) in which we engage the system in the task of the fix cycle and see what impact this has on component parts. Hutchins' documents this extensively in activity charts that he uses to breakdown the overall system and how various subsystems combine to perform the fix cycle once every three minutes (see figure 4.4 below). Here we can pick out how the various subsystems (*A*, *B*, and *C*) are all engaged in the overall task behaviour. A period of approximately one minute is shown of one loop of the fix cycle; beginning with the observing and reporting of landmark targets by the pelorus operators (the two red circles in the bottom left hand corner). Then the jotting down of these digital readings by the recorder in the log book (blue circle). Before these reports are then

[57] Although this is not quite right, because, as was mentioned above, actually this cognitive workload is covered by several "precomputational" cognitive artefacts – to avoid confusion I discuss this later in chapter six when I detail the diachronic temporal factors involved in the case study.

passed on to the plotter who uses them to construct lines of position on the chart using a hoey to reformat the digital representations into analogue representations (orange circle).

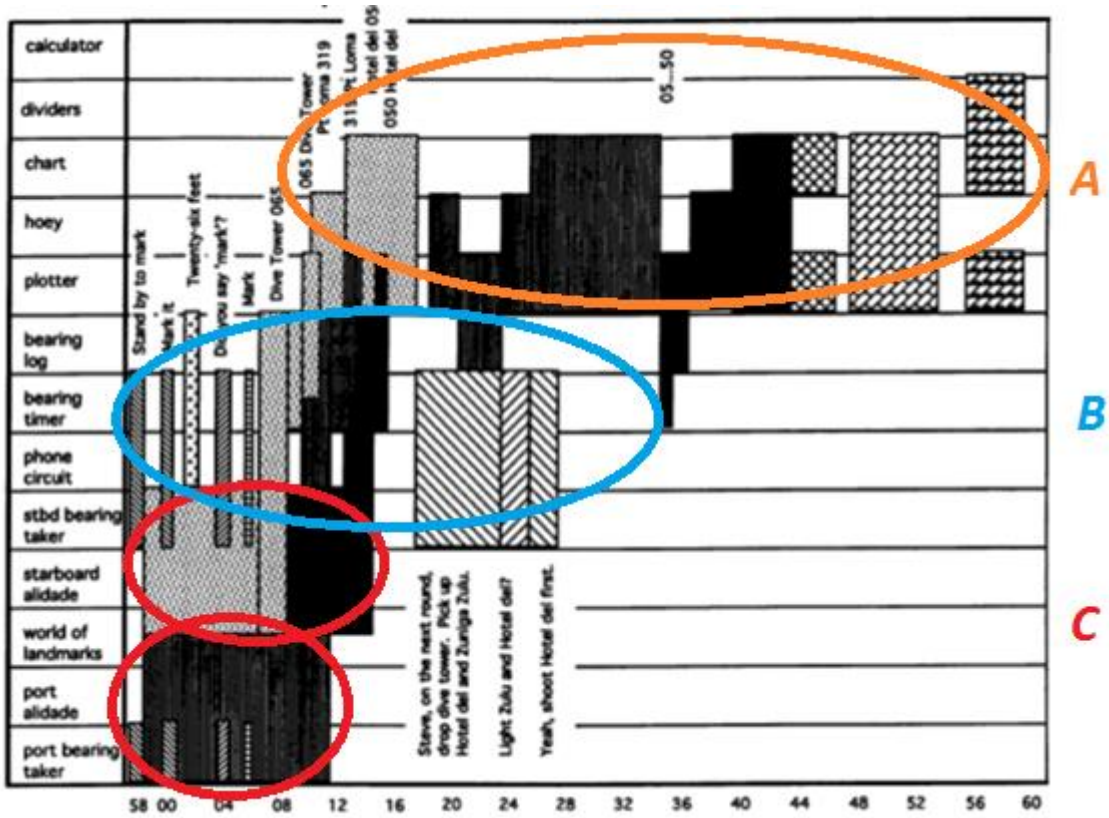

Figure 4.4: in this activity chart shows snapshot of how the activity throughout the navigation team is divided amongst the members who are interacting with various cognitive resources over a period of approximately one minute. The three main subsystem components of the overall system have been identified: [A] the plotter manipulating the chart and hoey to compute lines of position; [B] the bearing recorder as the hub of the communication network and acting as a buffer and timekeeper by using discrete representations and the log as an exogram; and [C] the two pelorus observers using the alidades to make observations of landmark targets and juxtapose them in the abstract space of coordinates. Additionally, we can begin to see here that temporal distribution is key to properly understanding a fairly synchronic or short-lived cognitive activity. For more details see main text. [Diagram adapted from Hutchins (1995a, p. 187)]

Based on the two theoretical interventions, we can therefore postulate that both aspects of the mutual manipulability criterion have been met; thus, providing us with good reasons to claim

that this is a distributed cognitive system. In contrast to the simpler accounts of distributed cognition which rely *solely* on information flow (the propagation of representational states across various media), the mutual manipulability criterion not only staves off the challenge of cognitive bloat, but also provides us with a principled reason for why and how cognition is distributed in synchronic activity. However, it was notable that my exposition had to omit several important diachronic factors (as mentioned in footnotes). I now return to this issue in the final section of this chapter before exploring this theme in even greater detail in the remaining two chapters.

4.5. The problem of time

What is interesting about the activity chart shown in figure 4.4 above is that, in addition to being a useful theoretical tool for collating the ethnographic information for the purposes of providing a fine-grained examination using task-specificity and mutual manipulability, it also demonstrates the crucial importance of time. The workload of the shared task is unevenly divided across the members of the team, and each of these is operating at different times on the differing subtasks (this is an instance of the parallel activity that I will discuss in chapter six, section 6.2.2). Although Hutchins does discuss the importance of time, his focus is more on cognitive historical factors – what he called "precomputations" (1995a, 2008) – and developmental trajectories (1995a, chs6&7; also see Alač & Hutchins 2004). As I discussed in chapter one (section 1.4), there are shortcomings in how precisely Hutchins discusses this topic – viz. he fails to distinguish adequately between these two forms of diachronic temporal distribution.

In both these diachronic cases Kaplan acknowledges that there are limitations about the timescales over which the mutual manipulability criterion operates (2012, p. 559). I think that this limitation is a serious motivating factor in favour of supplementing the mutual manipulability

criterion with the notion of normative patterned practices – which I detail in the next chapter. But there is also an issue here in regards to what the activity chart shows in relation to the synchronic temporal scale – or, more accurately, what we can reasonably mean by the term 'synchronic'. This is important because it has formed a core part of Kirchhoff's (2016) critical analysis of the use of the mutual manipulability criterion in debates about extended and distributed cognition. He argues that cases like the fix cycle are "dynamic" and involve nonlinear causal pathways that bring in to question the widespread notion that constitutive and causal relations are different in kind.

Kirchhoff (2016, p. 7) notes that constitutive relations are taken throughout philosophy – implicitly and explicitly – to be atemporal, whereas causal relations are, as Bennett (2011, p. 93) notes, "paradigmatically diachronic". This division between causal and constitutive relations according to temporality is a major issue for the proponent of distributed cognition since it places an immediate barrier towards the importance of cognition distributed in time – which, as I have repeatedly claimed, is crucial for understanding human cognitive systems as they are in the wild. My aim in this section is to disentangle what is the legitimate issue of cognition distributed in time from this confusion over whether a synchronic activity can also be understood as dynamic or nonlinear.

Kirchhoff's issue is to question whether constitutive relations can be "instantaneous" in these cases. But I think this argument rests on a confusion between what we can differentiate as 'dynamic but synchronic' and 'properly diachronic' timescales. This point can be made clearer by returning to the distinctions I made in chapter one, following Bietti & Sutton (2015), Nersessian (2005) and others, in which I identified one synchronic timescale corresponding to behavioural factors and which are task-relative (e.g. ranging from micro-periods of milliseconds up to slightly longer periods involving minutes and hours). These shorter timescales correspond to snapshots that I argue should be understood as embedded in two longer term diachronic temporal scales: *firstly*, an

ontogenetic timescale relating to the developmental trajectories of the system itself (e.g. the lifetime of an organism); and *secondly*, a historical timescale relating to the intergenerational transmission and refinement of epistemic resources and cultural practices (e.g. ranging over decades and potentially centuries). I have argued that both these diachronic temporal factors have a significant impact not only on the cognitive capacities of the human agents but also on the task-spaces within which they act (Menary 2015; Menary & Gillett 2017; also see Fabry 2017).

It is hard to see how the mutual manipulability criterion could be appropriately applied in these latter contexts – for instance, there is a built-in barrier by which one generation influences the next in the case of "downstream epistemic engineering" (Sterelny 2003). This is the process in which the epistemic niche within which the next generation develops and engages with cognitive problems has been sculpted and shaped by the work of previous generations. The agents of generation T_1 shapes that of T_2 and so on (T_{n+1}). And although members of generation T_{n+1} are reciprocally shaped and both shape their own current cognitive-cultural niche through acquiring, mastering and deploying cognitive practices (Menary 2014, 2015; Roepstorff et al 2010); they cannot retroactively influence the cultural-cognitive niche of generation T_1. This point becomes even more apparent when we consider multiple generations. This is not a critical blow for my adoption of the mutual manipulability criterion because it was always intended to operate in a limited fashion in conjunction with other complementary principles to designate how cognition is distributed in time and space in a more robust manner. As such, both cases of diachronic temporal distribution raise interesting limitations for the criterion and suggest what normative patterned practices must do to supplement and complete the pluralistic framework. This shall be my goal in the next chapter.

However, Kirchhoff (2016) claims that mutual manipulability is also complicated by dynamic systems in which there are nonlinear feedback loops. Kirchhoff states his issue concisely as follows:

> ...if we assume that some constituents perform their activity at a much faster timescale [than] the 'higher-level' process of the system as a whole, and if we further assume that the constitution relation between the whole and its component parts is synchronic (holds at an instant *t*), *how*, then, can the constitution relation hold synchronically – under the presupposition that some parts perform their role at a temporal scale faster than the higher-level process of the system as a whole? (2016, p. 10. My emphasis)

For example, the fix cycle is spread over several agents and the artefacts they are manipulating (subsystems A, B, and C). These functional units are both sequentially and nonlinearly related and spread over timescales occurring at different rates (as shown in the activity chart diagram in figures 4.3 and 4.4 above). Kirchhoff (2016) has argued that there is a problem in cases such as these because there are *dynamic processes that have temporal durations of differing lengths at differing levels of analysis*. So, for instance, the act of observing by the pelorus operators is shorter in duration than the entire fix cycle of which it is putatively a component. It is this kind of disparity which Kirchhoff argues is an issue.

Contra Kirchhoff, I contend that the problem of differing dynamic temporal lengths can be accommodated by the mutual manipulability criterion. I think this is especially clear in the Hutchins' case study because although the cognitive activity is temporally extended over three minutes at the higher-level property (the fix cycle itself) and the actual activity of components is much shorter, if one takes snapshots at time-slices none of these are causal relations, instead they are constitutive (this is especially clear in the nested hierarchy of figure 4.3). The activity of the subsystems A, B, and C *compose* the fix cycle – as is shown by the fact that they meet the dual criterion of mutual manipulability. Each of these activities is constitutive of the overall system behaviour and the fact that they operate at differing timescales does not create an issue that warrants serious concern. This becomes even more apparent when we realise that each of the subsystems of the fix cycle can themselves be decomposed into a series of components which also operate at differing timescales.

The crucial point is that if one cuts intersections through components performing the higher-level task behaviour of the fix cycle these inter-level relations should still be characterised as constitutive even if they have differing temporal durations.

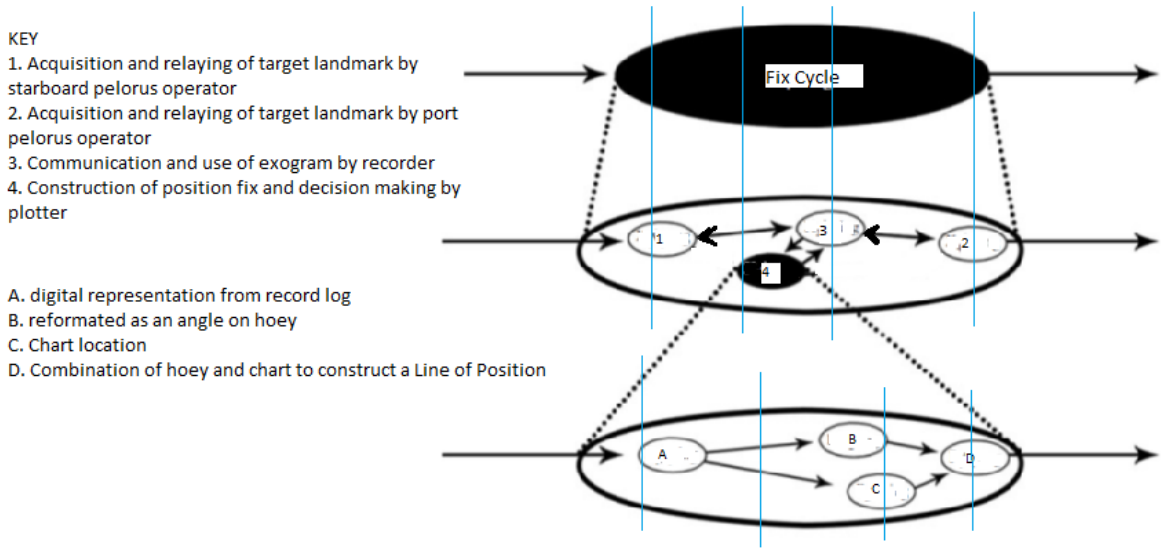

Figure 4.5: Three levels of analysis of the fix cycle: At the top is the system behaviour itself. The middle level shows the interaction between the three subsystems. The bottom level shows the sub-processes of just the plotter. Inter-level cuts (blue lines) show that even if the subcomponents of the overall system behaviour have differing temporal durations, that this activity is still constitutive rather than causal of the overall system behaviour. [Diagram adapted from Craver (2007, p. 189)]

Figure 4.5 above shows a series of inter-level cuts to demonstrate that even though the overall system level behaviour is temporally extended in time, the component behaviour at each stage is constitutive rather than causal. To regard what we mean by synchronic differently is to hold the studying of cognitive phenomena to unreasonable standards that don't match the work of practising researchers. Viz. it is reasonable to take the passage of a few milliseconds or minutes to be a snapshot of a cognitive system at a particular point in time – to go beyond this is to entertain debates about the metaphysics of time that are both unnecessary for practical work of exploring

cognition and currently undetermined by our best current science (on the latter issue see Ladyman & Ross 2007, pp. 162-167). As such, it is not within the ambit of the naturalistic and pragmatic project that I am conducting here. Therefore, Kirchhoff's (2016) notion of "reciprocal continuous causation" is an unnecessary amendment. Instead, the problem of time is the further issue of how this synchronic scale is situated within two important diachronic factors which the mutual manipulability criterion is unable to parse.

4.6. Summary

In this chapter I have outlined the second core feature of my pluralistic framework, *the mutual manipulability criterion*, and shown how it complements task-specificity. This mechanistic explanatory strategy composed of two interventions on the system (top-down and bottom-up) allows us a principled and neutral means by which to discern genuine components of a distributed cognitive system from mere causal background features. It therefore demarcates the flexible unit of analysis in a robust and principled fashion that avoids the problem of cognitive bloat. I have also shown how this operates in regards to the central example of the navigation case study. It not only demonstrates that this is a legitimate distributed cognitive system, but also leads to practical insights into the cognitive phenomena itself.

However, the mutual manipulability criterion has severe limitations on the timescales it operates over and so is insufficient by itself for considering how cognition can be distributed in time beyond the synchronic scale. As such, I turn in the next chapter to the notion of normative patterned practices, which I shall argue is the central and unifying principle of the pluralistic framework. This latter aspect is crucial, because whereas the mutual manipulability criterion gives

us a principled means of fending off cognitive bloat, the problem of methodological bloat still remains and must be nullified.

5: Normative Patterned Practices

Our brain and nervous system are our most cultural organs.

Greg Downey and Daniel Lende (2012b, p. 23)

5.1. Introduction

In the first two chapters I showed that the proponent of distributed cognition faces many differing kinds of challenges that motivate the need for a pluralistic approach. In chapters three and four I outlined how the first two principles of my pluralistic framework operate in regards to fixing the flexible unit of analysis across multiple levels of analysis. But the focus has primarily been on how cognition is distributed across space. This leaves the problem of understanding how cognition is distributed in time. Most, but not all, discussions of distributed cognition focus almost exclusively on synchronic timescales – a snapshot of the target system at a particular moment devoid of longer temporal factors. When one attempts to extend these criteria out over longer temporal scales – what I refer to as diachronic factors – the issue of how to demarcate the unit of analysis again becomes vague. Rather than ignoring these factors as unnecessary complications that can be abstracted away from or idealised; *I argue that how cognition is distributed in time diachronically is crucial for properly understanding how it is distributed at all*[58]. Specifically, my claim, following

[58] Humphreys (2004, pp. 141-147) distinguishes between abstraction and idealisation. He defines idealisation as a process whereby we take a property of the target system and alter it into a related property that is more desirable (for tractability). In contrast he identities four meanings of abstraction which are (for the most part) related to the omission of certain properties (either through ignoring them, or treating them as having a value of zero, or through controlled experiments that eliminate them) so as to focus on another specific set of properties in the target system. Humphreys' distinction is noteworthy because although I think that it is wrong to try and idealise diachronic factors as synchronic trajectories, I accept that in some cases it might be

Menary (2013a, 2013b) and others (e.g. Bietti & Sutton 2015; Hutchins 1995a, Nersessian 2005; Sutton 2008), is that in order to properly understand a human cognitive system in the wild one must understand how it is distributed in time according to two major diachronic temporal scales – what Nersessian (2005, p. 38) calls *developmental trajectories* and *cognitive-historical factors*.

My aim in this chapter is to make these notions tractable. I shall argue that the key to understanding this problem is a proper understanding of *the relationship of culture and cognition*. I will outline how this can be best understood using two concepts adopted primarily from the work of Menary (2007a, 2012, 2013a, 2013b, 2014, 2015, 2016) and others (Boyd et al 2011; Downey & Lende 2012b; Fabry 2017; Roepstorff et al 2010; Sterelny 2003, 2012; Tomasello 1999, 2009): *normative patterned practices* and *cultural-cognitive niches*. Normative patterned practices are sets of patterned activity spread across a population which *both shape and are shaped by* the coordination dynamics involved in specific cognitive tasks. Cultural-cognitive niches are the epistemic aspects of the cultural environments which distributed cognitive systems are embedded in. Furthermore, I shall use normative patterned practices to stave off the conceptual problem of methodological bloat outlined in chapter one by showing that this principle can tie together and organise the rest of this comprehensive framework in a tractable manner. This completes my pluralistic framework's rebuttal of the multiple problems outlined in chapter one; thus, offering a complete and robust defence of distributed cognition.

The chapter is structured as follows. In the next section (5.2) I outline the importance of recognising that humans are cultural animals for understanding human cognitive activity by drawing a brief comparison between the activities of the plotter from Hutchins' navigation case study and

---

justifiable to abstract away from diachronic factors on the basis of utility. However, to this I would like to add two caveats: firstly, I also contend that standard investigations of human cognition often unknowingly overlook that a solitary focus on synchronic factors is omitting these important diachronic factors. Secondly, a major aim of my approach to distributed cognition is that it provides investigators with a framework which makes diachronic temporal factors tractable so that they need not be abstracted away.

navigational practices from a differing cultural-cognitive niche. I then turn to the issues of how to operationalise 'culture' for studying human cognition (section 5.3). Following Roepstorff and colleagues (2010), Menary (2007a, 2015) and others, I shall recommend that we focus on practices – in particular normative patterned practices – and I outline the multiple key beneficial features of this approach (section 5.4). In the final section (5.5) I elaborate on this discussion to clarify how this allows me to tackle both the problem of methodological bloat and the issue of time. In the next and final chapter, I recapitulate the entire pluralistic framework and show how it elucidates the central case study of the navigation team across all three major temporal scales in a manner that is: [1] both practically and theoretically salient; [2] goes beyond standard simple approaches; and [3] offers insights that should be of note to methodological individualists.

5.2. Two differing cultural-cognitive niches

In the previous chapters we have discussed the coordination dynamics of distributed cognitive systems – how they operate over multiple levels of analysis (chapter three), and how they can be differentiated against the causal milieu (chapter four). In this chapter my aim is to show how the coordination dynamics of distributed cognitive systems are shaped, formed, and partially constituted through being embedded in cultural-cognitive niches; and that *an adequate understanding of distributed cognition necessarily involves taking these factors into account*. As I outlined briefly in chapter two (section 2.6), a cultural-cognitive niche is the sculpted social environment in which humans conduct their daily cognitive behaviour.

Henrich remarks that humans are "peculiar" and "puzzling" primates because although we are an astonishingly successful species – spreading across the planet and inhabiting almost every terrestrial habitat – this is in spite of the fact that our species is weak, slow, and requires cooked

food despite lacking innate capacities for making fire or cooking (2016, p. 1). Henrich argues that the key to understanding this is that we are a "cultural species" – i.e. it is not our immense brain size but our reliance upon cultural knowledge and intelligence that enables us to survive and thrive (pp. 2-3; also see Boyd et al 2011; Sterelny 2003, 2012; Tomasello 1999, 2009). This is not to say that big brains are not important – indeed humans can be seen (albeit not straightforwardly) as part of a more general pattern of various species that have larger brains (that are specially structured) for larger behavioural repertories and increased social learning (see Deacon 1997; Downey & Lende 2012a; Herculano-Houzel 2009; Hofman 2014; Reader et al 2011; Rilling 2006)[59]. Instead, Henrich's point is that our success is dependent on collective efforts rather than individual brilliance. Therefore, enlarged brains are not primarily for de novo behaviour, but for reliably acquiring the cognitive and cultural capital of a particular cultural-cognitive niche (also see Sterelny 2011, p. 810; 2012, p. 151; Menary 2014, p. 291). Human cognitive activity takes place within a specialised and structured environment. Sterelny (2003, p. 123) identifies our niche constructing capacities as one of the hallmarks of our species. Although many other species engage in niche constructing activities (Odling-Smee et al 2003), humans have a unique propensity for cumulative niche construction: the high-fidelity transmission of cultural resources across generations that enables the retention of innovations. Tomasello (1999) refers to this as the "ratchet effect" and he identifies it as a qualitatively unique feature of human culture (2009; also see Boyd et al 2011; Dean et al 2012; Henrich 2016; Sterelny 2003, 2012). Human niche constructing activities alter not only the physical and general cultural environment but also the epistemic environment as well (Sterelny 2003).

I refer to this with the slightly different nomenclature "cultural-cognitive niche" because the term cognitive niche has also been used by evolutionary psychologists to support their massive

[59] This is complicated by the fact that our species has also been engaged in a form of self-domestication which some have argued has led to a reduction in human brain sizes since the Pleistocene (Henrich 2016; Henneberg 1989; cf. Ruff 2002).

modularity position (e.g. Pinker 2010). Boyd and colleagues (2011) have critiqued this position because it places too much emphasis on individual intelligence and underplays the importance of culture and cooperation. They provide evidence against the modularist view through natural experiments from history (e.g. lost European explorers and lost cultural knowledge through the death of elders in oral traditions), which demonstrate the importance of culture rather than individual ingenuity.

As has long been recognised by anthropologists, despite the acknowledgment that culture is crucial for understanding our species there are serious questions about understanding how to operationalise 'culture' in order to understand its impact on human cognition (Vogeley & Roepstorff 2009). Before discussing the complexities of the issue of operationalisation, I will first demonstrate the relevance of cultural-cognitive niches to the topic of distributed cognition. In particular, I will compare the performance of the plotter engaging in a navigation task with an agent from a differing 'navigational niche' (these are displayed in figure 5.1 below). The comparative strategy is notable because, as Fleck notes, inhabitants of a cultural-cognitive niche are usually unaware of how their cognitive processing is influenced: "The individual within the collective is never, or hardly ever, conscious of the prevailing thought style which almost always exerts an absolutely compulsive force upon his thinking, and with which it is not possible to be at variance" (1979, p. 41). Importantly, Boas adds that a lot of how human cognitive behaviour is cultural habituated is unconscious and often occluded behind post-hoc rationalisation which hide how it is enculturated (1910, p. 382).

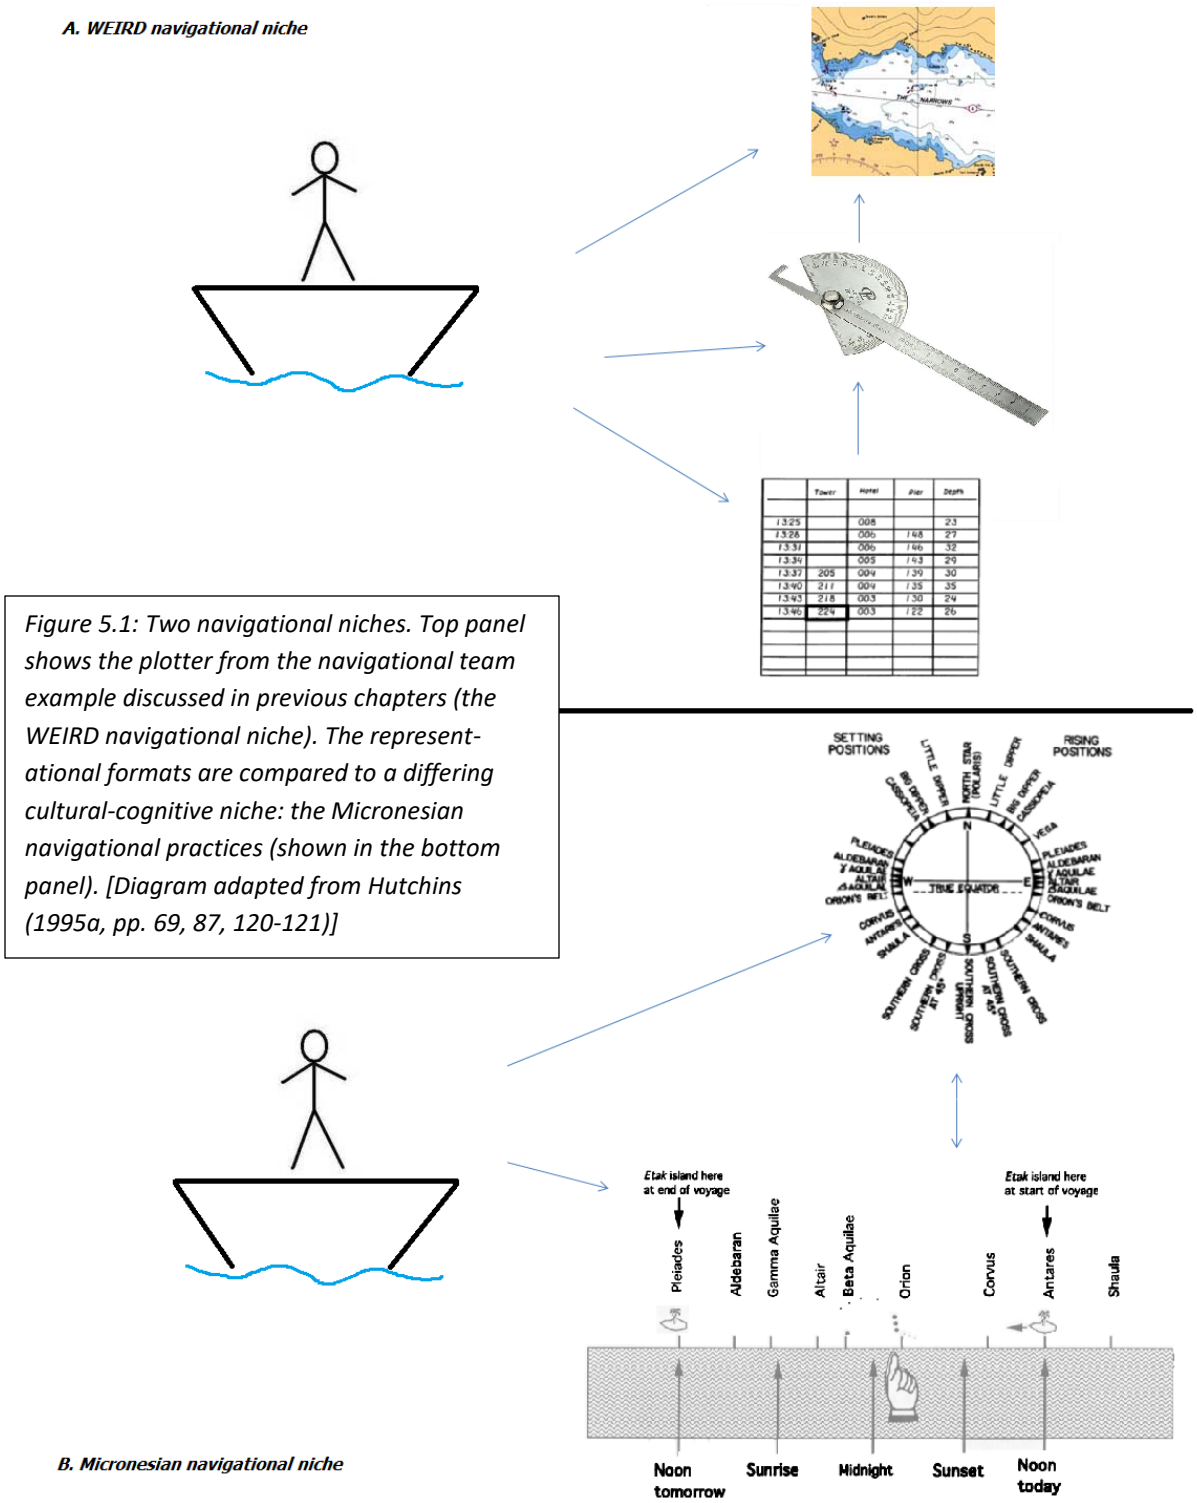

Figure 5.1: Two navigational niches. Top panel shows the plotter from the navigational team example discussed in previous chapters (the WEIRD navigational niche). The represent-ational formats are compared to a differing cultural-cognitive niche: the Micronesian navigational practices (shown in the bottom panel). [Diagram adapted from Hutchins (1995a, pp. 69, 87, 120-121)]

By examining multiple cultural-cognitive niches we can overcome these deficits. I refer to the navigational niche in which the fix cycle is embedded as 'WEIRD' (Western, Educated,

Industrialised, Rich, Democratic), following Henrich and colleagues (2010) famous paper, for reasons that I will go into in the next section. By navigational niche I mean the cultural practices, techniques, and tools that comprise and shape how they perform the task. Hutchins identifies the two central abstract questions in navigation as: "where am I?" and "given that I at location *X* how do I get to location *Y*?" (1995a, pp. 12-13, 50, 52). From this abstract starting point, one can then move onto a comparative analysis of how human agents in differing cultural-cognitive niches deploy cognitive resources to tackle this abstract task. As I previously discussed in chapter three, the interactions between agents and artefacts embedded in cultural-cognitive niches have a dramatic impact on how the task is apprehended and what representations are involved (Kirsh 1995; Zhang 1997b); in turn this impacts the neurocognitive profiles of the agents and the functional behaviour of the wider system (Menary 2007a, 2013b, 2014, 2015; Menary & Gillett 2017). In this chapter I refine how culture shapes and partially constitutes these coordination dynamics.

Hutchins identifies three main features of the WEIRD navigational niche as follows (1995a, pp. 95-112): [i] increasing use of physical artefacts; [ii] digital measurement; and [iii] the importance of the chart. We can see all three of these at play in how the plotter performs his subtasks of the overall fix cycle. There is a heavy reliance on the use of the hoey and the chart in order to tackle the central questions of navigation. These devices act as material anchors that enable the agent to reformat digital representations into analogue representations. In particular, the navigational chart combines multiple discreet measurement systems for organising space and time. As Hutchins notes, navigation is a task that humans have been engaging in for millennia (1995a, p. 21). The basic problems have been tackled by a vast array of individuals across multiple generations and this has reorganised the task space through the incremental creation and refinement of tools to help solve aspects of the task space (also see Boyd et al 2011; Henrich 2016; Sterelny 2003; Tomasello 1999).

Although this is a crucial factor, I return to the historical aspects of this in chapter 6 (section 6.4). Our primary concern here is that the modern navigational problem for a naval navigator is not one of trying to figure out how to use the information and solve the problem – as Hutchins notes "that has already been worked out" (1995a, p. 21). Rather, their task is to use the already existing tools and techniques to resolve everyday problems based on the present information and desired goal. This is an instance of conversion as I discussed in chapter three (section 3.3.1). For instance, in order to tackle the central questions of navigation the plotter uses the sophisticated multimodal representation on the chart to locate the vessel in relation to a set of coordinates (an abstract representational space) that is *blended* with physical space[60]. The plotter is then able to project the future passage of the vessel using a set of techniques called 'the three-minute rule' (I discuss this more in chapter six, section 6.4).

The significance and 'WEIRDness' of these procedures and representational formats can be made more salient by comparing them to a different cultural-cognitive niche. This comparison then shows the importance that culture has in reshaping [1] the task space, [2] the neurocognitive profile of the agents, and [3] the functional behaviour of the distributed cognitive system. Hutchins (1995a, pp. 65-93) highlights the peculiarities of the WEIRD navigational niche by contrasting it to another cultural-cognitive niche: what I shall refer to as the Micronesian navigational niche. In contrast to a reliance on physical tools and the use of discreet and digital measurement systems, the Micronesian navigational niche provides us with a comparative group that achieves nautical navigation to a highly successful degree in an entirely different way. I now summarise the key differences:

Firstly, the Central Caroline Islands – east of the Philippines – are notable for being comprised of less than 0.2 percent land (p. 67). As such, the local population often have to traverse

[60] I discuss the importance of conceptual blends below in more detail.

great distances at sea without any landmarks. *Secondly*, they do not use any external representational devices: no compasses, maps, or other navigational equipment (Hutchins 2008, p. 2013). *Thirdly*, Micronesian navigators do not use discrete measurement systems of time and space. Instead, journeys are measured utilising two sophisticated representations (see figure 5.1 above): [1] a sidereal compass of how stars rise and fall in the sky at different times of day – Heersmink (2013) refers to this as a "naturefact" (the appropriation of a physical structure as a serendipitous representation of another factor one wants to measure). These movements trace out "star paths" related to certain constellations. Hutchins notes that although seeing a star is a simple perceptual task, constellations do not exist independent of cultural groups who have arranged a set of stars into a perceptual group: "While an eye can register a pinpoint of light, seeing that pinpoint as a star is a cultural accomplishment" (2011, p. 441). He further elaborates that this is an embodied cognitive activity dependent on the agent's "…brain, of course, but also on his body and his eyes, and on a set of traditional cultural practices that orchestrate the interactions among a complex collection of elements" (2008, p. 2012). Combined, this provides expert navigators who have learnt how to see these patterns with a sidereal compass that defines thirty-two directions (Hutchins 1995a, p. 69; 2005).

In conjunction with this naturefact, navigators also use another sophisticated multifaceted representation: [2] two forms of interrelated "fictive motion" – the blending of a dynamic scenario with a static situation (Fauconnier & Turner 2002; Hutchins 2005). When the canoe is beyond the sight of land the navigators imagine that it remains stationary whilst the world moves around the canoe (1995a, p. 72; 2008, p. 2013). Agents learn to attend to the sensations of how their bodies feel whilst sailing in order to judge their speed, as well as the direction of their canoe in the dark. Additionally, an imaginary *moving* island called an "etak" is projected by the navigators to be beyond the horizon and travelling parallel to their course – and this allows them to divide their

journey into several non-discrete stages. Hutchins adds that the developmental trajectories related to learning these cultural representations "…establish and maintain the required relations among the elements of the distributed system" (2008, p. 2013).

Before moving on to the comparative analysis it is important to note the role of "conceptual blends" (Hutchins 2005, 2010c) that are occurring in this distributed cognitive system. In both the etak and the sidereal compass, internal representational architecture is brought into conjunction with a physical feature of the environment. As mentioned, this involves learning to see the world in a certain way – i.e. altering one's perception and affordances. But it is crucial to recognise that this is not just a transformation of the agent's perceptual affordances but also involves a transformation of how these structures can be manipulated as external representations. This ability enables the agents to articulate and conceptualise more sophisticated cognitive tasks – this is what Menary (2007a, 2015) refers to as the transformation thesis (which I discuss in greater detail below in section 5.4).

Whilst the Micronesian navigational system might seem overly baroque I suggest that one's prima facie reaction to this is premised on the radical alterity of this system to the unconsciously domesticated status of our own WEIRD navigational niche. These differences are important to recognise for two reasons: firstly, it provides us with an example of a distributed cognitive system that does not involve any physical cognitive artefacts. Secondly, it highlights the important differences that being enculturated into differing cultural-cognitive niches can have on the representational schemas of agents approaching the same abstract task. This should not just be understood simply as: in one niche the agent uses physical tools and in the other they use naturefacts. Instead, the differences have a profound and deeper impact on how the agents cognise about spatial relations. For instance, the WEIRD navigator primarily uses a map; and spatial

navigation involves an objective space in which bearings intersect in allocentric (or absolute) space in the location of targets. In contrast, a Micronesian navigator does not use any artificial tools but relies on a set of techniques that blend serendipitous features of their environment with highly imaginative fictive motions; and spatial navigation is subjective (or egocentric) so that bearings intersect in the observer and nowhere else (i.e. the observer is a privileged point).

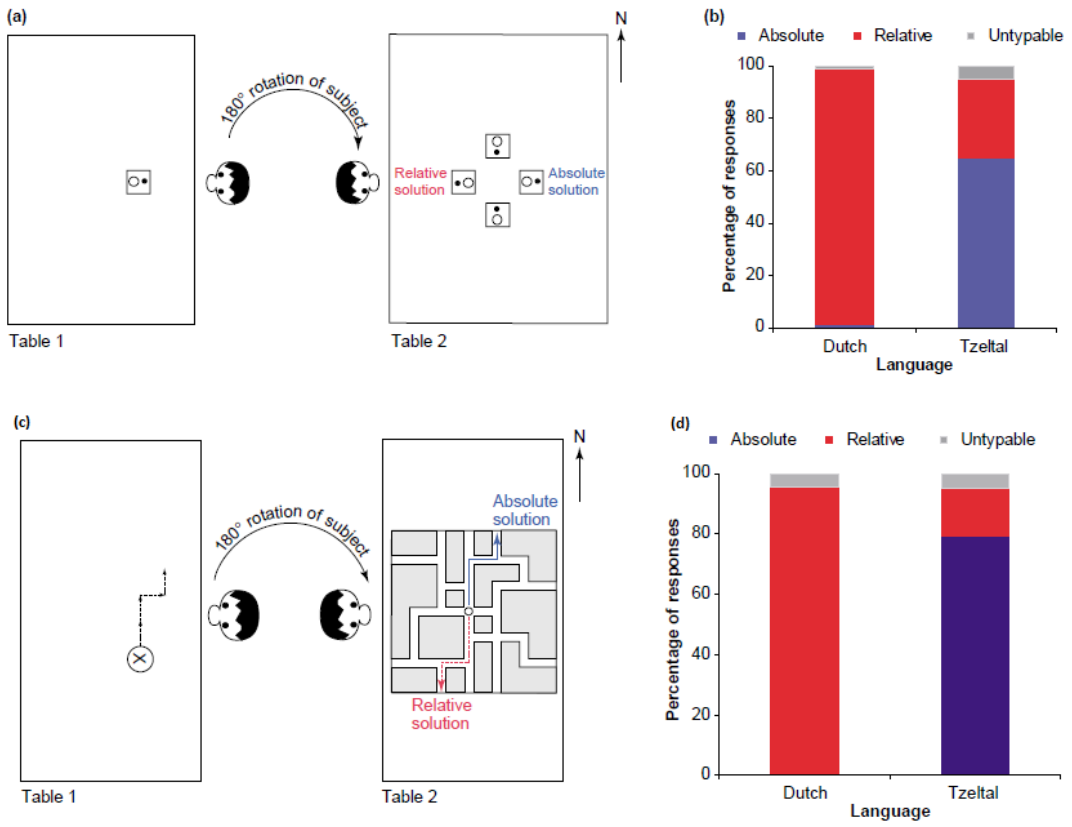

Figure 5.2: Two experiments by Majid and colleagues (2004) comparing spatial reasoning of egocentric and allocentric language speakers (respectively Dutch and Tzeltal). Both experiments involve participants looking at a task set up in one room [table 1 as shown in panels (a) and (c)] before rotating 180° to look at another task set up in another room [table 2]. There are two different solutions – relative and absolute – which are favoured by one's spatial language terms. These experiments show that agents from allocentric and egocentric frames of reference see the same task differently (results in the tables in panels (b) and (d)). [Diagram from Majid et al (2004, p. 110)]

Convergent empirical evidence is suggestive that these differences impact on the basic ways in which agents cognise about spatial relations: Firstly [1], the basic framework of allocentric or

egocentric entails differing task performance on spatial tasks (Haun et al 2006, 2011; Levinson 2003; Majid et al 2004). In these experiments, speakers of primarily allocentric or egocentric languages are presented with a task set up before rotating 180° to complete a spatial reasoning task that has two possible solutions (relative or absolute). They found that most speakers of a language favoured a particular way of tackling the problem. This shows that one's language group structures how one comprehends and thinks about space – even for something as simple as a rotation of 180° (see figure 5.2 above). This has significant real-world importance: for instance, speakers of Guugu Yimithirr – an allocentric language from Northern Australia – are able to reliably point to cardinal directions without recourse to any external cognitive device – even when blindfolded or inside a building with no reference to landmarks outside (Haviland 1998). This capacity is evident in many other allocentric language groups (e.g. Boroditsky & Gaby 2010; Levinson 2003).

Secondly [2], differing strategies for navigating in space have been repeatedly shown to have differing neural correlates (Bohbot et al 2007; 2012; Konishi & Bohbot 2013): [i] "response strategies" which involve following instructions (e.g. proceed to the end of the road then turn left and carry on until the third right, etc.), which correlates with neural activity in the caudate nucleus; and [ii] "spatial strategies" use landmarks to construct a mental map relative to a set of other points in space, which correlates with activity in the grey matter of the hippocampus (in particular the right para-hippocampal cortex which is correlated with spatial memory). Since differing cognitive tools and techniques favour differing strategies, we can infer from this that differing configurations of distributed cognitive systems are likely to have an impact on neuroplasticity. Indeed, expertise and prolonged habitual practices in regards to tackling navigational problems with spatial strategies is correlated with neuro-plastic changes to the para-hippocampal structures (Maguire et al 2000).

Spatial cognition and navigation is not an isolated case of enculturation. Downey and Lende (2012a, 2012b) summarise a wide range of cognitive domains in which culture dramatically impacts on the neurocognitive profile of humans: including the shaping of music and phoneme perception by one's experience of the acoustic environment (Boas 1889; Gaser & Schlaug 2003; Nan et al 2008; Roepstorff et al 2010); the refinement or stunting of olfactory perception through one's language group (Downey 2014a; Wnuk & Majid 2014); the shaping of decision-making capacities (Henrich et al 2005, 2010; Nisbet et al 2001); and the transformation of mathematical cognition using either abacus or pen and paper as its primary external representational tool (Cantlon & Brannon 2007; Hu et al 2011; Tanaka et al 2012; Tang et al 2006). This evidence is highly suggestive that culture – particularly the process of enculturation by which an agent becomes inculcated into a cultural-cognitive niche (Menary 2014) – is important for understanding distributed cognition. But how to understand culture is itself a vexed question, as I will now discuss in the next section.

5.3. The challenges of operationalising culture

Many theorists in philosophy and the human sciences agree that culture is crucial for understanding human cognition (e.g. Boyd et al 2011; Henrich 2016; Menary 2014; Rogoff 1990; Shore 1996; Sterelny 2003; Tomasello 1999). As Geertz puts it:

> ...man's nervous system does not merely enable him to acquire culture, it positively demands that he do so if it is going to function at all. Rather than culture acting only to supplement, develop, and extend organically based capacities logically and genetically prior to it, it would seem to be ingredient to those capacities themselves. A cultureless human being would probably turn out to be not an intrinsically talented though unfulfilled ape, but a wholly mindless and consequently unworkable monstrosity. (1973, pp. 47, 68)

However, whilst this point might be relatively widely accepted, culture is a notoriously ambiguous term. This is a point that has long been recognised in anthropology (see Boas 1910; Downey & Lende 2012b, pp. 36-41; Moore 2012; Ortner 1984, Roepstorff et al 2010; Vogeley & Roepstorff 2009 for discussion) but has more recently become a stumbling block for other fields which have been trying to operationalise the term 'culture' in order to make it more tractable for empirical research (e.g. see Laland et al 2010, p. 138).

The complexities and difficulties of this are apparent in the new fields of cultural psychology (Nisbet et al 2001) and cultural neuroscience (Chiao 2009). This work is exemplary insofar as it represents a step away from a mono-focus on Western, particularly North American, participants. Arnett points out that modern psychology is incredibly parochial in its sampling of the human population, and that claims about universality in psychology are pretty weak given that 68% of all data samples come from only 5% of the world's total population (2008, pp. 602, 605; also see Henrich et al 2010, p. 63). Following Henrich and colleagues, we can label this the problem of WEIRD participants (Western, Educated, Industrialised, Rich, and Democratic). Henrich and colleagues have showed that the situation is even worse than a simple skewing of the sample to one particular culture because WEIRDs "are among the least representative populations one could find for generalizing about humans" (2010, p. 61). This is because WEIRDs are an outlier compared to the rest of the world's population when one conducts meta-analyses (p. 79-80).

A startling example of the outlier status of WEIRDs is the well-known Muller-Lyer illusion in which a participant is presented with two equal length parallel lines which have arrows at the ends either pointing inwards or outwards and asked which one is longer (see the right panel in figure 5.3 below). Previous investigations of this illusion which focused solely on WEIRD participants have reliably found that agents pick the line with arrow pointing inwards as longer. However, when

Henrich and colleagues (2010, p. 64) took this visual illusion to participants living in a wide range of small communities they found widely differing results that showed that WEIRD people are an outlier in how strongly they respond to this illusion (this is shown in the graph on the left in figure 5.3 below in which the WEIRD populations are the two right-hand columns). As such, this example demonstrates that it is important to take enculturating factors into account when investigating human cognition. This point becomes even stronger when one considers that this is just one of many cases across a wide range of domains – including spatial reasoning, moral reasoning, visual perception, and socioeconomic reasoning – in which Henrich and colleagues have showed that WEIRD populations are not only different to many cultural groups, but are outliers[61].

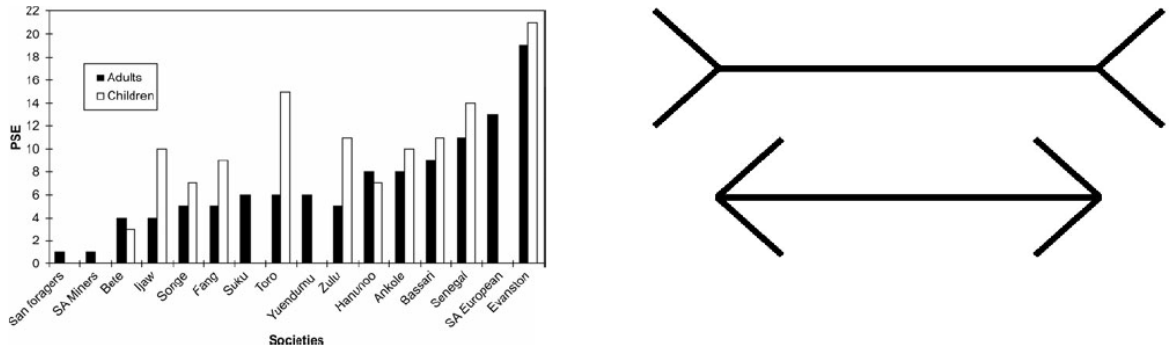

Figure 5.3: Right panel: the Muller-Lyer illusion – two equal length parallel lines are visually judged to be of different lengths depending on which the arrows are facing. However, as shown in the graph in the left panel (from Henrich et al 2010, p. 64), this illusion affects WEIRD participants far more than multiple other cultural groups. Indeed, the response of the two WEIRD populations in the right-hand columns are outliers to the general lower response of fourteen other groups studied who were from smaller-scale societies.

[61] There is a large debate about whether the capacity for reason is a human universal or whether it varies across populations. Nisbet and colleagues (2001) argued that 'Western' populations used what they referred to as "analytic" reasoning (based on the principle of non-contradiction and discretely identifiable units), whereas 'East Asian' populations instead have a "holistic" reasoning style (which tries to reconcile differences and considers the context or field much more). Nisbet and colleagues link these "qualitative differences" in reasoning styles with social hierarchical features (p. 305). Mercier characterises this as claiming that some cultures "spurn debate that threatens social harmony" (2011, pp. 87-88, 92-93). He rejects this and argues that although reasoning styles might be culturally variant to some degree this is not a qualitative difference, as Nisbet and colleagues claim, and is instead merely the cultural shaping of human universal capacity for making human communication robust (see Mercier & Sperber 2011).

The increased work in cultural psychology and cultural neuroscience is to be applauded for making strides towards countering the lopsided focus on WEIRD participants. However, several commentators have pointed out that these new fields also have a range of problems themselves. Firstly, rather than exploring the variety of potential differences on offer between the many different cultural groups across the planet most work in cultural psychology and cultural neuroscience has compared a generic WEIRD group (drawn either from Europe or North America) with what is usually labelled an "East Asian" group (from Japan, Korea, or China). These two unwieldy groups are then compared across a limited set of abstract binaries – e.g. collectivist vs individualist. Not only does this give the wrong impression that there are only a limited number of coarse-grained cultural identities and a small axis of binary differences between them (Cohen 2009, p. 194). It also fails to attend to differences within categories such as "collectivist" and "individualist" which can vary dramatically and thus cannot be straightforwardly transposed across different cultural populations (Downey 2016, pp. 51-52). Indeed, a meta-analysis of the literature has undermined the notion that 'East Asians' and 'Westerners' are homogeneous groups and shows that the intergroup differences are far less stark than is often claimed (Oyserman et al 2002; also see Roepstorff 2011).

Perhaps most damagingly, Roepstorff and colleagues (2010) argue that it is not clear what is meant by 'culture' in many of these studies (also see Downey 2016; Reynolds Losin et al 2010). They point out that these studies try to use culture as a label for a homogeneous population, but this has been shown definitively by the field of anthropology to be meaningless because it ignores intra-cultural variation (2010, p. 1051). Instead, we need a principle that is more tractable but which also pays attention to this heterogeneity. I.e. we need a way to "operationalise" culture for investigating its impact on human cognition (Downey & Lende 2012b, p. 36; Vogeley & Roepstorff 2009, p. 514).

In these regards, Roepstorff and Bubandt argue that we can appeal to the anthropological tradition to clarify how culture is not a "homogenous and static entity but a heterogeneous and dynamic set of practices" (2003, p. 15; also see Ortner 1984). Vogeley and Roepstorff define culture as a "set of competences, practices and beliefs in groups that shapes and influences the group members and that is […] in continuous and dynamic exchange with its members (rather than being a rigid body of standardisations of language, habits or belief systems)" (2009, p. 511). Building on this, Roepstorff and colleagues (2010, p. 1051) have suggested that the notion of "patterned practices" is the key because it picks out regularities in everyday activities that we can identify in the coordination dynamics of core configurations.

I will now outline this in greater detail and discuss the benefits of adopting such an approach in the next section; before highlighting how this resolves the remaining challenges facing the proponent of distributed cognition (viz. the problem of cognition distributed in time and methodological bloat) in the final section of this chapter.

5.4. Normative patterned practices

In this section I outline several key features of normative patterned practices. Firstly, I will discuss the normative aspect of patterned practices and how they relate to developmental trajectories that have various transformative effects on the system. Secondly, I outline a taxonomy of cognitive practices – based on the work of Menary (2007a, 2016) – that supplements our understanding of coordination dynamics. Thirdly, and lastly, I summarise four key positive attributes that make normative patterned practices particularly suited for analysing distributed cognition.

5.4.1. Learning, transformation, and norms

Menary defines patterned practices as patterns of activity spread out across a population that are physically embodied, repeated, and stable habits that are shared by members of a cultural-cognitive niche and which are reliably involved in how human agents interact with each other and their various environmental resources (2013a, 2013b, 2014). These stable patterns of activity arise from habitual interactions but which also shape that activity (Roepstorff et al 2010), and must be learnt by agents as they are inculcated into cultural-cognitive niches.

Particular cultural-cognitive niches have developmental aspects (Menary 2014; Sterelny 2012; Stotz 2010). These structures enhance learning in a number of ways such as the use of props, toys, and tools that support trial-and-error learning; and the implicit and explicit teaching of skills, values, ideas, information, and expected modes of social interaction and behaviour – this is what Sterelny refers to as "hybrid learning" (2012, 2014). In these developmental niches the acquisition and mastery of a set of patterned practices impacts the agent's neurocognitive profile: by both initiating neuro-plastic changes to brain regions which are cannibalised towards new cultural functions (Menary 2014; also see Dehaene & Cohen 2007) and physiological changes associated with new behavioural repertories to enact the manipulation of cognitive resources; but also leads to functional changes so that the agent is able to tackle novel cognitive tasks that would otherwise be in feasible or potentially even unthinkable (Menary 2007a, 2015).

Here we see the importance of developmental trajectories associated with how the agent acquires and then masters these cognitive-cultural practices. Through hybrid learning and the repeated deployment of these patterned practices towards a specific cognitive task the agent becomes habituated. Their performance moves over time from 'clunky' or 'jerky' behaviour characteristic of the novice, to the expert's phenomenological "fluency" and "perceptuo-motor

integration" in their interactions with external cognitive resources and the task space (Nemirovsky et al 2013; Menary 2013b, p. 361). A notable feature of our species, in comparison to others, is that the hybrid learning of these developmental niches – which is often necessary for sophisticated cognitive tasks such as mathematics and scientific reasoning – involves explicit teaching in combination with a high propensity for social learning (Dean et al 2012; Keil 2011; Tomasello 2009). Tomasello and colleagues have termed this the "Vygotskian intelligence hypothesis" (Hermann et al 2007; also see van Shaik & Burkart 2011).

On this topic, it is incredibly important to note that the process by which patterned practices are mastered demonstrates two points related to the work of Vygotsky. Firstly, that some cognitive functions begin as what he referred to as "interpsychological" before they become "intrapsychological" (Vygotsky 1978, 1981; also see Wertsch 1985). Menary notes that the cultural origin of these cognitive functions entails that our biologically endowed capacities can be supplemented and transformed (2013b). This enculturation process is not only dependent on the structures of the niche but also on specialised features of the sub-individual level of analysis: humans have a high degree of plasticity (Sterelny 2003). Of particular importance is what Menary (2014) refers to as "learning driven plasticity". The plasticity of the human cortex and nervous system allows for neuronal resources to be cannibalised and exapted for new, culturally generated functions (Anderson 2010; Dehaene & Cohen 2007). It is noteworthy that this again shows the importance of a multilevel analysis for distributed cognition; and also clarifies why we cannot ignore the sub-personal level of analysis when investigating putative distributed cognitive systems.

Learning driven plasticity relates to the second point of Vygotsky's work: that the process by which an interpsychological phenomenon is mastered sufficiently by an agent to become intrapsychological results in the internalisation of the process which subsequently alters the

cognitive function: “internalisation *transforms* the process itself and changes its structure and function” (1981, p. 163. My emphasis). In regards to the coordination dynamics of a distributed cognitive system this entails that diachronic factors are crucial for understanding the cognitive properties of the system. Because the ontogenetic status of the agent dictates how they engage with the cognitive resources in their niche.

It must be noted that a critic might claim that the notion of internalisation makes distributed cognition at best transitory – because by internalising external cognitive resources they are only used as scaffolds by novices who then abandon these once they have mastered the cognitive practice. To this we can note two rebuttals: firstly, Menary notes that many sophisticated and complex cognitive tasks never become fully internalised – for example, complex mathematical operations such as imaginary numbers (2013b; 2015; also see De Cruz & De Smedt 2013). Indeed, Sutton and colleagues (2010) point out that we should recognise the natural variation by which different agents interact with the cultural-cognitive niche: some individuals utilise external tools whereas others prefer culturally acquired techniques (e.g. in a memory task one could use an exogram or a mnemonic).

Secondly, Hutchins has argued that there are some interpsychological properties that can never be internalised (1995a, pp. 283-285). In the next chapter I return to consider his examples of “parallel activity” and “communication” in relation to the navigation team case study (sections 6.2.1 and 6.2.2). Thirdly, and perhaps most decisively, rather than falling into a trap of a pseudo-internalism, Vygotsky’s axis of ‘interpsychological *prior* to intrapsychological’ offers us a viable alternative to the parity principle in defending distributed cognition since it does not prioritise the skull as the origin of cognitive properties (Menary 2007a; also see Cole & Engeström 1993; Nardi

1996; Theiner & Drain 2016). This also matches with the task-specific approach I outlined in chapter three.

The developmental trajectories by which human agents learn and master the deployment of normative patterned practices for the successful completion of cognitive tasks are transformative of the neurocognitive profiles of human agents (Menary 2015). Menary refers to this as the "transformation thesis" (2007a). As I noted in chapter three, many theorists are imprecise in regards to this proposition – often failing to carefully distinguish between various levels of analysis. Following Humphreys (2004), I identified three main ways that we can think about how the manipulation of an epistemic resource alters a human agent's neurocognitive profile: [1] *extrapolation* – the amplifying of an existing capacity (e.g. the use of a telescope to make an observation); [2] *conversion* – the altering of a cognitive task into one that makes it more amenable for our anthropic biases and thus more tractable (e.g. using a list to convert a memory task in to a perceptual one with a far reduced cognitive load on long term memory); and [3] *augmentation* – whereas the former enhancements are merely quantitative changes to the agent's neurocognitive profile, manipulations engendered through the use of this kind of epistemic tool (and the changes associated with the process of learning the normative patterned practices that guide this activity) entail a qualitative change in which the distributed system – agent-manipulating-cognitive-resource – is capable of cognitive tasks that would otherwise be impossible. I contend that Menary is usually discussing the transformation thesis in regard to this more qualitative change to the agent's neurocognitive profile because the transformatory aspects of developmental trajectories often involve the acquisition of cognitive skills for tasks that would otherwise be impossible, or distinctly novel (2015, p. 3).

We can see this clearly in regards to Hutchins' navigation case study. The plotter's capacities to interact with spatial reasoning problems are drastically altered by being inculcated into a set of normative patterned practices for interacting with an array of cognitive technologies; e.g.: [i] inscriptional notation and measuring systems provide public representational systems that allow for discrete and digital judgements about space; [ii] manipulations of the chart facilitate sophisticated conceptual blends of abstract and physical space that would be hard to keep stable if they were purely in the head (Hutchins 2005); and [iii] use of the chart also transforms the orientational framework within which navigational tasks are conceived – from a personal view to a bird's eye view. It is notable the latter is a distinct cultural achievement: an abstract perspective that would not otherwise exist and is radically different from traditional wayfinding relations to the environment (Hutchins 1995a, p. 62; Ingold 2000, p. 227).

This discussion shows how important developmental trajectories are for properly articulating what distributed cognitive systems *are*. The information flow accounts I discussed in chapter two which only focus on synchronic coordination dynamics omit a whole range of factors. Viz. the transformation and generation of novel behavioural repertories; neural plastic changes (including the reuse of cortical tissue for new functions); and alterations of the functional unit of analysis through the mastery and habituation of practices – development and internalisation also relates to differences of multilevel analysis (thus, if these are overlooked then one is likely to commit one of the many conceptual errors I outlined in chapter two). Perhaps most damagingly, a solitary focus on a synchronic temporal plane renders these positions incapable of actually establishing that these are cases of distributed cognition. Viz. simpler accounts of distributed cognition which ignore cognition distributed in time cannot explain why and how coordination dynamics and flows of information have the particular patterns that they do. Importantly, this should not just be understood as outlining the conditions of the possibility for the distributed

cognitive systems. Instead, normative patterned practices are also partially constitutive of cognitive activity insofar that they not only shape coordination dynamics but are also that which is enacted in how the wider system performs a cognitive task.

We can see this in the case of the plotter from the navigation team case study. The coordination dynamics by which he manipulates the hoey, chart, and numerical symbols from the bearing log book are orchestrated by normative patterned practices. These dictate how to combine these epistemic resources. The plotter's physically embodied movements and other computational processes are shaped by his acquisition of these practices. But these practices are also enacted through his instantiating these coordination dynamics. As such, this shows how we can see the information flow and normative patterned practices accounts as consilient and supplementary.

Considering the developmental aspects of the navigation case study also highlights another aspect of Menary's position: the importance of cognitive norms (2007a, 2010a, 2010b). I have discussed the fact that the cognitive practices that agents must learn in order to tackle cognitive tasks effectively become habitual and are repeated patterns. But it is also important to recognise that embodied interactions with environmental resources are normatively constrained (i.e. in a particular cognitive task there are actions that one can and cannot make). Tomasello (2009) contends that a qualitatively different aspect of human culture is its normative dimension. For instance, Vaesen notes that humans are unusual amongst primates in attributing particular functions to specific tools (2012, p. 206). Humans use tools for specific functions, and specific tools can be said to have specific affordances. Behavioural experiments have shown that this "functional fixedness" can be enculturated relatively quickly – this leads some theorists to posit the existence of a conceptual system that organises and stores functional information. This hypothesis is supported by neuropathological cases of apraxia – where patients either suffer from conceptual or motoric

errors and disorders in their purposeful execution of actions using tools (see Vaesen 2012, p. 206 for more details).

Henrich identifies that social norms have radically influenced a wide range of human activities: kinship relations, mating, food sharing, parenting, and reciprocity (2016, p. 5). He also notes the importance of rituals for cementing pro-sociality. This is crucial for maintaining Tomasello's ratchet – the accumulativeness of cultural-cognitive niche construction. Even more strongly, Henrich argues that for a cultural animal, sociality (especially social interconnectedness) is more important than brute individual intelligence (2016, pp. 212-214).

Menary's central point is that norms are important for understanding human cognitive behaviour – i.e. human cognitive behaviour has become evolutionarily shaped in cultural-cognitive niches to be shaped, transformed, and guided by cognitive norms (2007a, 2010a, 2010b). They are normative insofar as there are correct and wrong ways to deploy epistemic resources towards the completion of a cognitive task. These "cognitive oughts" must be learned in order for an agent to appropriately utilise a particular environmental resource to successfully complete a cognitive task (Menary 2007a, p. 139). Menary identifies five kinds of cognitive norms that are involved in human activity towards cognitive tasks (2007a, p. 137; 2010, p. 239; 2010b, pp. 570-571):

1. *Purposive*: these dictate how practices are engaged towards the goal of a particular activity (e.g. how manipulations of pen and paper or an abacus are towards the completion of mathematical tasks).

2. *Corrective*: there are norms for correcting cognitive activity towards a goal (e.g. redrafting or revising a draft paper, correcting errors in the workings of an equation, or refining the parameters of a mathematical model).[62]
3. *Manipulative*: there are norms that guide how epistemic resources are to be manipulated (e.g. when reading English, one reads left to right and top to bottom).
4. *Interpretative*: how the products of epistemic actions are to be interpreted are also normatively constrained (e.g. when moving an abacus these movements have specific numeric values).
5. *Creative*: there are norms governing the creation of epistemic resources (e.g. whether one needs to create new notation or terminology for a novel concept).

Menary does not aim for this list to be exhaustive. Instead, the key point here is that insofar as a majority of human cognitive behaviour takes place within cultural-cognitive niches – either explicitly relying on environmental resources or structured through being embedded in the particular epistemic niche – it is inherently normative. Since norms have to be learnt and mastered in order for them to be properly deployed this demonstrates that the coordination dynamics of distributed cognition are incredibly reliant on developmental trajectories. As such, if one overlooks these in an analysis then one does not have a proper understanding of the cognitive behaviour of the system. This again shows the inadequacy of a solitary focus on information flow.

5.4.2. Cognitive practices

Menary's framework also provides us with another way of nuancing our understanding of coordination dynamics that is supplementary to the mechanistic approach I discussed in the

[62] In the next chapter I will elaborate on the importance of error detection and correction in the navigation case study (see section 6.3.2).

previous chapter. Menary argues that embodied actions or manipulations play a crucial role in many cognitive tasks. These are the coordination dynamics by which internal and external resources are mediated[63]. In chapter three, when discussing the role of task-specificity in centring the flexible unit of analysis, I emphasised the importance of embodied manipulative interactions for avoiding conceptual mistakes related to the incorrect attribution of properties to the agent rather than the system level (e.g. mereological fallacies, overattribution, and residual internalism). Fabry highlights the wide scope of what should be counted as coordination dynamics in these contexts: "…any kind of embodied interaction with the [cultural-]cognitive niche that functionally contributes to the performance of a cognitive practice" (2017, p. 3). Menary (2016) identifies six kinds of cognitive practices: [1] biological coupling; [2] epistemic practices; [3] self-correcting practices; [4] tool-use; [5] use of public symbol systems; and [6] blended practices[64]. The following descriptions draw heavily on Menary (2007a, 2010a, 2010b, 2016) but have been expanded to show how they apply to a wide variety of cases, including the navigation team:

Biological coupling: these are direct sensorimotor interactions of an organism to its environment; for example, exploratory saccadic eye patterns (see Ballard et al 1995). A good example of how these can be shaped by cultural practices is provided by Downey (2011, 2012) who has explored how extensive training practices in Capoeira (a Brazilian martial art) can greatly transform sensory systems such as one's sense of balance and space. Downey (2015) has gone on to emphasise the importance of motility of the body in its direct sensorimotor interactions with the

[63] Menary (2007a, 2010a), following Rowlands (1999), terms this the "manipulation thesis". But to avoid confusions with the mutual manipulability criterion I will simply stick to discussing Menary's work under the notion of cognitive practices.

[64] In previous versions of this framework (2007a, 2010a, 2010b), Menary presented only the use of representations as involving cognitive practices. In recent work he has altered his position to designate all of these various kinds of coordination dynamics as cognitive practices (2016). This is a prudent reconfiguration because all of these kinds of manipulations have normative aspects; and as such, should be counted as being mediated by cognitive practices.

environment that arises through habitual practices. Downey's example is athletics but he goes on to point out that these factors are just as important for more sedentary practices:

> Our underlying biological nature, like all living organisms, is dynamic and subject to degeneration, improvement, modification, and adaptability. Rather than producing only "meaning" in some symbolic, discursive, or political sense, we can see in a phenomenon like sport how habitual action modifies the body, in muscular, neurological, skeletal, and other fashions. *Sport makes more apparent a process that occurs in all cultural practices, even sitting still and reading or typing while focusing our eyes less than a meter from our faces*. (2015, p. 129, emphasis added)

Implicit here is the relevance of the physical manipulation of epistemic resources and the impact this has on the basic motor-programmes and altered perceptual capacities of the agent. In the navigation team case study, we can see this in the embodied activities of the plotter who aim to "think like a compass" (Hutchins 1995a, p. 141), and in how the physical manipulations of the dividers and hoey on the chart involve sensorimotor contingencies (2010b).

Self-corrective actions: these interactions constrain future actions to goal-specific movements. These are practices whereby the agent enacts a patterned practice that keeps their attention on task-salient features. These can take the form of techniques such as verbal cueing in egocentric language use in children (Menary 2007a; Vygotsky 1978), realigning one's body in space to improve perception (Clark 2016; Menary 2010b, p. 564), or what Sutton (2007) calls "instructional nudges", and "caretaking practices" – e.g. the phrase "keep the eye on ball" for cricket batsmen acts as verbal command which helps direct attention. Hutchins (1995a, p. 313) notes that "verbal shadowing" is common in "cognition in the wild" because it is a self-regulatory function that improves task performance (Clark 1997 also refers to this as "phonological looping"). For instance, the plotter uses verbal shadowing when computing a position fix[65]. Additionally, this function can

[65] This action also has further functions related to group error detection which I discuss in the next chapter.

also be instantiated by actual features of the environment rather than physical movements of the body. Kirsh (1995) refers to this as the "intelligent use of space" whereby agents use "jigs" in a range of everyday cognitive tasks to constrain cognitive acts to only those which are task-salient (also see Zhang 1997b; Zhang & Norman 1994). An example of a self-corrective external representation in the navigation team is the viewfinder of the alidade which has a cross-hair in the visual field of the pelorus operator so as to focus attention on the target landmark.

Epistemic practices: here, Menary draws heavily on Kirsh and Maglio's (1994) identification of "epistemic actions" which they differentiate from pragmatic actions. The latter are simply moves in physical space whereas epistemic actions involve the manipulation of the environment as its own representation – i.e. they are not just physical movements in space but can also be conceived as moves in an abstract conceptual space of the cognitive task itself. Kirsh and Maglio used this notion to describe the behaviour of expert Tetris players who would make multiple rapid rotations of the playing piece (Zoid) that are akin to exploratory movements of how best to use the piece. Although the expert players made far more physical movements and had to put in more external work their performance was far better in comparison to novices who relied more heavily on internal resources (e.g. mental visualisations). Kirsh notes that "Epistemic actions are everywhere" (2006, p. 252). Other examples include the use of gestures, pen and paper, abacus, and body parts for mathematical tasks (De Cruz & De Smedt 2013); jigsaw puzzle pieces for problem solving (Kirsh 2010); and the movement of the nautical slide rule and the hoey for computational tasks in the navigation case study (Hutchins 1995a, 2005).

Kirsh and Maglio (1994, p. 514) argue that epistemic actions operate to simplify and offload the cognitive task in three main ways: by reducing memory load; reducing time complexity; and reducing error. In chapter three I noted that this is a limited way of understanding what is going on in these interactions that does not take multiple levels of analysis into account. Instead, we can

recognise that the developmental trajectories involved in learning how to properly deploy normative patterned practices entails a transformation of both the neurocognitive profile of the agent and the task-space (Menary 2014, 2015). For instance, the use of the hoey by the plotter *converts* a digital computational task into an analogue task that involves the physical manipulation of the moving arms in coordination with the visual perception of the alignment of a public symbol system – this renders the task far easier, faster and more robust.

The use of epistemic tools: in our discussion of self-corrective actions and epistemic practices it was evident that a number of these involved the manipulation of "culturally constituted objects" in the cultural-cognitive niche (Hutchins 2010c, p. 99). The use of epistemic tools plays a very important role in everyday cognition. Menary gives examples such as calculators, rulers, pens and paper, and computers. These environmental resources are often referred to as cognitive artefacts (Heersmink 2013, Hutchins 1999, Norman 1991). The navigation team is replete with numerous epistemic tools – alidades, record log book, hoey, and the chart – and the coordination dynamics of the agents with these devices performs a range of cognitive functions (e.g. perception, memory, and computation).

As noted above when discussing self-correcting actions, one of their key features is that they bring stability to cognitive tasks. This stability in part is due to what Donald (1993) refers to as an "exogram" – external memory storage structures that reduce the cognitive loads on working memory. Furthermore, both Hutchins (2010) and Donald (2010) note that this feature allows for the division of cognitive labour across a social group (i.e. enables and increases collaboration) – and this is even more the case with public representational systems. For instance, in the navigation team the physical structure of the chart acts as a material anchor that facilitates the joint attention of the plotter and recorder for decision making in the choice of target landmarks and the assessment of bearing reports by the pelorus operators.

Use of public symbol systems: another key epistemic resource in cultural-cognitive niches are public representational systems such as alphabets, numerals, diagrams, notational systems for formal logics and computer programming languages, pictographic and logographic writing systems, and maps, etc. As noted above, the capacity for a cognitive task to be distributed across groups is greatly aided by these types of cognitive artefacts. Representational systems allow agents to engage in sophisticated cognitive tasks that would otherwise be improbable (Menary 2007b). Henrich (2016) notes that learning a particular public symbol system provides an agent with new cognitive techniques for engaging with cognitive tasks (also see Menary 2007a, 2015). For instance, representational notations are arguably essential for the development of mathematics and logic (De Cruz & De Smedt 2013; Dutilh Novaes 2013; Luria 1976; Menary 2015).

Beyond the obvious augmentative effects to communication, the use of writing enables agents to record detailed and nuanced arguments and other records without placing a large cognitive load on working and long-term memory. Neuroimaging evidence shows that learning to read and use public symbol systems has an impact on cortical structures – thickening the corpus callosum which connects the two hemispheres of the brain (Castro-Caldes et al 1999) – and is also associated with behavioural differences, which *crucially* are not just related to memory tasks but also to other cognitive domains such as improving/transforming logical reasoning (Luria 1976; cf. Mercier 2011) but diminishing facial recognition capacities (Dehaene 2009). Furthermore, this "exographic" nature (Sutton 2008) also facilitates the refinement and increased complexity of multi-stage logical arguments, technical treatise, and creative thinking (Frieman 2015; Menary 2007b).

The sophistication of the WEIRD navigational niche is dependent on the historical development of an array of measuring and notational systems that go back as far as the ancient Sumerians and Babylonians (Hutchins 1995a; Ifrah 1998). This shows the importance of these systems across diachronic historical timescales – the downstream accumulation of a cultural-

cognitive niche. But notational systems also have a causal and constitutive role in synchronic cognitive activity. The common example of this given in the literature is Rumelhart and McClelland's (1986) description of long-multiplication whereby a complex calculation such as 365 x 475 is broken down using an algorithm into simpler single-digit computations which can be spatially arranged on a page so as to keep track of the calculations and offload working memory. Menary (2007a, 2013b, 2015) argues that the cognitive system is partially constituted by the normative patterned practices that guide and shape the manipulations of the symbols. For instance, the computational activities of the plotter using pen and paper to calculate lines of position are shaped and governed by the agent's internalisation of norms that dictate how to use Arabic-Indian numerals and coordinate systems of the Mercator projection.

Blended practices: these involve combinations of the above manipulations for complex cognitive tasks. In his case study of a flight simulator, Hutchins (2010c) shows how these are combined with gestures that trace imaginary trajectories that enable pedagogy and communication. Hutchins has gone on to show that these sorts of blended practices are common in everyday cognition and more sophisticated settings; for example, Alač and Hutchins' (2004) ethnographic study of how an expert and a novice examine and read the data emerging from an fMRI scanner. This output is an incredibly complex array of colours and symbols on a digital screen – a phase map. Alač and Hutchins go to great pains to stress how difficult both the expert and the novice find it to extract meaningful patterns from this data – what they call "seeing as" which is the blending of an actually instantiated material space with an imaginary one. This is another example of the importance of conceptual blends to investigations of distributed and enculturated cognition (Hutchins 2005, 2011). The expert creates and deploys an array of transitory translation tools that enable her to decipher the data and explain this to the novice. She traces imaginary trajectories across these devices and in reference to the phase map. This enables her to draw out the salient

aspects and interpret the image not only for herself but also to scaffold the novice as well. Achieving this would not have been possible without extensive training in the appropriate practices; as is evident by the novice's difficulties. In the navigation team the chart is the definitive blended practice, I will return to discussing a range of important facets of this in the last chapter.

Lastly, a general point of note: Fabry emphasises the limitations imposed by the physical attributes of human bodies on the embodied manipulative aspect of coordination dynamics – e.g. the degrees of freedom in joints and muscles in human limbs curtail possible movements (2017, p. 20). We can elaborate on this by drawing on Changizi's (2003) analysis of behavioural repertories in mammals, which shows that an increase in motor programmes is due to the recombination of pre-existing muscle groups into novel patterns. The question here is whether enculturation into differing sets of normative patterned practices enables and favours certain kinds of re-combinations. For example, Downey's (2012) analysis shows how differing sets of practices in Capoeira or gymnastics for learning how to do handstands entails physiological differences that also have consequences for the affordances of the participants. And at the neural level, both Anderson (2010) and Dehaene and Cohen (2007) have discussed whether the prior function of cortical regions constrains how they can be redeployed by cultural practices.

In addition to bodily and neural constraints, I think there is one more form of constraint that is overlooked here: this is what Engelbart (1988) refers to as a "Neo-Whorfian hypothesis" – how the materiality of the environment and its potential conduciveness to manipulation has an impact on possible cognitive capacities. For instance, he points out that in the case of writing that the availability of materials with which to make inscriptions has a serious impact on whether and how this cognitive practice develops. Engelbart sought to prove this point with a delightful experiment in which he tied a brick to a pencil and recorded his comparative performance in speed and legibility. At first glance Skagestad notes that this experiment appears absurd given its obvious outcome; but

it is intended to induce a "mental jolt" and force us to realise "…the extent to which our intellectual, social, and cultural life depends on the physical characteristics – literal weight and otherwise – of such tools as pencils" (1993, p. 161).

Skagestad goes on to add that the point can be made much clearer by considering the history of writing implements (e.g. Petroski 1989) where we can clearly see that some forms of writing materials inhibited use over long periods of time, thus acting as a constraint against tasks that required protracted use (e.g. writing lengthy essays and treatise). For instance, the Norse only had chisels and stone tablets, which, as Skagestad notes, "…made the writing of even the simplest sentence an extremely laborious and time-consuming task" (1993, p. 162). Petroski makes similar points about the use of charcoal (1989).

In conjunction with the manner in which the other principles of my pluralistic framework carve up the coordination dynamics of distributed cognitive systems, this taxonomy supplements and enriches our terminology for drawing out how to understand these relationships. In the final section of this chapter I will clarify the points of consilience between normative patterned practices, task-specificity, and the mutual manipulability criterion in more detail. But firstly, it is useful to briefly summarise some general positive attributes of normative patterned practices that make it particularly well suited for discussing distributed cognition.

5.4.3. General positive attributes

Although the preceding discussion demonstrates that the concept of normative patterned practices is useful to the proponent of distributed cognition; as I have argued, normative patterned practices must also be recognised as crucial for adequately understanding what distributed cognition actually is. Indeed, if one overlooks their role in articulating the coordination dynamics of a system then

one's conception of distributed cognition will be inadequate because it fails to understand how and why they are patterned as they are, and why the system should be regarded as distributed.

In addition to this vital role, and given that the principle is being used in multiple different ways throughout the framework, it is worthwhile clarifying four general positive attributes – viz. normative patterned practices: [1] are empirically tractable; [2] relate directly to the neurocognitive profile of human agents; [3] mediate multiple temporal and spatial levels of analysis; and lastly [4], are consilient with distributed cognition. I now discuss each of these in more detail.

Firstly, as noted in the previous chapter when discussing the mutual manipulability criterion, empirical tractability is a desirable meta-theoretical feature from a naturalistic and pragmatic philosophical perspective. Moore (2012, pp. 304-305), Roepstorff and Bubandt (2003, pp. 14-15), Roepstorff and colleagues (2010, pp. 1051-1052) all emphasise that normative patterned practices are preferable over other ways of operationalising culture in these regards – insofar as patterned practices are observable and testable (in principle). For instance, we can examine the normative patterned practices involved in a particular cognitive task, and potentially engage in the double-interventions of the mutual manipulability criterion to identify what role they are playing in the deployment of coordination dynamics – i.e. we can attempt to discern the precise manner in which cognitive activity is governed, shaped, and partially constituted by the process of enculturation. I discuss this in more detail below in the final section of this chapter. And in the next chapter I will provide numerous concrete examples of this in relation to the navigation team case study.

Additionally, practices cover a wide range of interrelated social and cultural phenomena (Roepstorff & Bubandt 2003, pp. 14-15). In this sense they can be said to incorporate capabilities, habits, collective actions, meanings, symbols, shared representations, and "other sorts of patterned, shared conditionings of the nervous system" (Downey & Lende 2012b, p. 37). Related to this

encompassing feature: there is not a division here between the rules of the interaction and the mediation of the interaction itself (see Kendal 2011, p. 245; Moore 2012, p. 299; Ortner 1984; Roepstorff & Bubandt 2003, pp. 14-15; Roepstorff et al 2010). This is useful because of the importance of cognitive norms (as I discussed above), but also because D'Andrade has proposed that having a conception of culture as a set of rules provides an opening gambit in bridging between cognitive science and anthropology (1995, p. 143).

Secondly, many theorists have noted the relationship between normative patterned practices and the neurocognitive profiles of human agents. Shore (1996, p. 5) follows Ruth Benedict in arguing that cultural patterns are sets of specific organisations of cultural artefacts and accompanying psychological patterns. Wacquant (2014) argues that habits are an empirical prompt, or point of investigation for seeing how patterns of practice alter and sculpt the psychological and embodied profile of an agent as concretely embedded in a social context. Roepstorff and colleagues claim that specific differences in normative patterned practices between specifically well-defined cultural groups are correlated with differences in neural correlates: "Patterns of practice at the level of social interaction correlate in relevant ways with neural and psychophysical patterns" (2010, pp. 1052). Hence, as discussed in section 5.3 above, they advise that cognitive scientists should use normative patterned practices rather than a more nebulous notion of 'culture' in investigating the relationship of the cultural niche to the brain and cognitive behaviour (pp. 1056, 1058).

As I have outlined previously, the notion that normative patterned practices are correlated with plastic changes in the brain is well supported in the literature on brain plasticity. Menary refers to this as "learning-driven plasticity" (2014): repeated engagements and deployment of normative patterned practices drive structural and functional stages over extended developmental periods. Multiple studies show that neural plasticity is driven by specific cultural practices over a wider range

of domains and timescales: such as juggling over just a few months (Draganski et al 2004); studying for a medical exam over just a few months (Draganski et al 2006); driving a taxi for a long period of time (Maguire et al 2000); and studying mathematics to high degree over several decades (Aydin et al 2007). Enculturation has a widespread effect on cognitive domains to differing degrees and over differential timescales (from months to decades). Indeed, Roepstorff explicitly notes that patterned practices are distributed beyond the individual in both space and time (2013, p. 62).

Thirdly, by drawing a direct link between how developmental trajectories involved in enculturation alter the neurocognitive profiles of agents this not only mediates between the micro and macro levels on the spatial scale of analysis (Roepstorff et al 2010, p. 1058); but also mediates between the various levels on the temporal scale. It does this by emphasising the importance of both 'dynamic-but-synchronic' activity in the present but also the importance of diachronic features (Menary 2013a; Ortner 1984, pp. 150, 158). As Ortner puts it: a practice-based approach takes history and development seriously (1984, p. 159). In regards to developmental trajectories this relates to how agents have to learn and master the requisite normative patterned practices associated with using epistemic resources in the cultural-cognitive niche towards cognitive tasks (Menary 2007a, 2014, 2015). In relation to cognitive-historical factors, normative patterned practices are the medium by which downstream epistemic engineering (Sterelny 2003, 2012) takes place through high fidelity inter-generational transmission. This then shapes the problem-space for successive generations and enables the tackling of cognitive tasks that would otherwise be impossible due to the temporal scales involved (Boyd et al 2011; Henrich 2016; Shea 2009). I return to discussing this in more detail in the next section to make it explicit how this can resolve the problem of cognition distributed in time; and in the next chapter I provide extensive examples of how this operates in regards to the navigation case study.

The last and fourth point of consilience for normative patterned practices and my approach is the explicit link drawn by Roepstorff and colleagues between their emphasis on patterned practices as a replacement for the concept of culture and distributed cognition (2010, p. 1056). This textural support is supplemented by the fact that Hutchins discusses practices throughout his oeuvre. In the first section of this chapter I discussed the navigational practices in WEIRD and Micronesian cultural-cognitive niches (1995a). Another key example that Hutchins discusses are trajectory-based cultural practices: the projection of imaginary trajectories onto real spaces that then organise cognitive behaviour, e.g. queuing (2014, p. 38). Hutchins argues that this is a culturally proscribed way of seeing space – i.e. an instance of patterned practices affecting perception (again highlighting the importance of conceptual blends). He claims that the trajectory-based practice of forming a line is a physical grouping that reduces the dimensionality (i.e. cognitive load) of a cognitive task in a social setting (2012). *But this reduction is not in the mind of one agent but in the shared practices of the participants* (also see 2014, p. 40): "Concepts in the wild are manifest in practices, and practices include the social and material settings in which they are situated" (2012 p. 315).

In general Hutchins argues that "Cultural practices assemble agencies into working assemblages and put the assemblages to work" (2001, p. 2069). He develops this point in more detail by critiquing how Clark's (2008) agent-centric approach underplays the role of practices in the assemblage and coordination of distributed cognitive systems: "the assembly process itself is extended and orchestrated by the cultural practices that constitute the cognitive niche" (Hutchins 2011, p. 442). Against Clark he highlights that "…*cultural practices are the things people do in interaction with one another*" and provides numerous examples of how everyday cognitive activity is constantly shaped and constituted by cultural practices; e.g. speaking, reading, writing, and

"…particular ways of seeing (or hearing, or feeling, or smelling, or tasting) the world" (pp. 440-441. Original emphasis).

The key issue here is that normative patterned practices are amenable to multilevel analysis precisely because they are patterns of activity both at the level of the individual and the group. Menary summarises many of these points nicely as follows with the claim that:

> …*our minds are extended by the practices* (real activities) of thought, which we learn in the niche (and where the niche may contain many social and cultural institutions). *This claim works on both an individual and group level. Cognitive practices are nothing but patterns of activity spread out across a population.* Individuals learn and acquire these practices during development and many of these practices will be collaborative in nature. In other words, cognitive practices are real activities that take place in a public forum; *they can be carried out individually or in collaborative groups* (where everyone is a practitioner). Cognitive practices are *genuine 'components'* of our mental and cognitive capacities, they are dynamic, active, processes by means of which we think and successfully complete cognitive tasks. (2013a, p. 27. My emphasis)

Here we see many of the key features which make normative patterned practices a good way to supplement the other principles of the pluralistic framework. Menary stresses the importance of the various spatial and temporal scales between the wider cultural-cognitive niche and the actual activity of a system as it is taking place synchronously and the developmental trajectories involved as well.

I now consider how these features of the concept of normative patterned practices help us to directly resolve the remaining challenges facing the proponent of distributed cognition: the problem of cognition distributed in time, and methodological bloat.

5.5. Tackling the problems of time and methodological bloat

In the preceding discussion I advanced the view that humans are highly culturally shaped organisms and that this has a significant impact on human cognitive activity. In order to make this relationship tractable I have proposed normative patterned practices for understanding this relationship between human agents and the cultural-cognitive niches they inhabit. I listed a range of desirable theoretical features that make normative patterned practices attractive for investigating the relationship between culture and cognition (e.g. that they are empirically tractable; that they relate micro- and macro- spatial scales, etc.). But there are three other points that are crucial for making my framework complete: [1] how they are conducive to both task-specificity and the mutual manipulability criterion; [2] how they provide a means by which to tractably analyse both diachronic temporal scales; and lastly [3], how they go beyond just being merely consilient with the two other principles and unify my pluralistic framework so as to avoid the problem of methodological bloat. Although I have touched on these issues throughout the preceding discussion, here I make these central points explicit.

Normative patterned practices are the third and last of the main principles that operate together, with task-specificity and the mutual manipulability criterion, to make a designation of distributed cognition more precise and accountable. It is important to recognise that normative patterned practices are not only conducive to both these other features of the pluralistic framework, but also supplementary. Firstly, in regards to task-specificity, following Steiner (1972) we articulated the importance of task-demands – these are the resources, context, and norms that guide a particular task. Normative patterned practices – as sets of procedures for the right and wrong ways by which agents relate to each other and to mediating tools in order to complete tasks – elaborate what these task-demands are in further detail. They also make the connection between task-demands and the wider cultural-cognitive niche more concrete and tractable.

Understanding the relationship between normative patterned practices and the mutual manipulability leads onto the second point of interest about normative patterned practices: the way in which they allow us to explore diachronic temporal scales in a principled fashion. As I noted in chapter four (section 4.3), Menary argues that the coordination dynamics in distributed cognitive systems could be understood *both* mechanistically but also in terms of the cognitive practices that have to be acquired and mastered and which govern these interactions (2013a, p. 28; 2013b, pp. 353-354). The acquisition and mastery of cognitive practices involved here in order for an agent to competently tackle a cognitive task involves a period of learning (in particular, hybrid learning which takes place in specialised structured developmental niches). Hence, understanding the synchronic deployment of specific coordination dynamics necessarily involves the importance of diachronic temporal factors which make the former possible (and as such, a solitary focus on synchronic information flow is inadequate for properly understanding distributed cognitive systems).

I have followed Kaplan (2012) in adopting the mutual manipulability criterion as an objective means by which to discriminate between the genuine components of a distributed cognitive system and the mere causal background features in a principled manner. To briefly recap: this criterion works by using two idealised interventions on the system in question. A bottom-up procedure which manipulates the component to see if it elicits a response in the system's behaviour as a whole; and a top-down manipulation of the system behaviour to see if it engenders a change in the putative component. Only if both conditions are met does it count as a genuine component of the system. Whilst this allows us to articulate the coordination dynamics in properly ethnographically detailed case studies, this only operates over synchronic timescales.

In this chapter I have demonstrated the importance of developmental trajectories for understanding the successful task performance of both the WEIRD plotter and Micronesian sailor

(section 5.2 above). I shall supplement this in the next and final chapter by providing multiple examples of how diachronic temporal factors are crucial for properly explaining the synchronic distributed cognitive activity of the navigation team. As such, the mutual manipulability criterion requires supplementing. Normative patterned practices fill this explanatory gap by mediating both diachronic temporal scales. Normative patterned practices are involved in the developmental trajectories by which agents learn and master the relevant procedures and strategies for utilising epistemic resources in their cultural-cognitive niche towards the successful completion of a cognitive task. These repeated deployments habituate the practices so that behavioural repertories are sculpted into stable patterned activity and this correlates with plastic changes at the neural level. Both the physically instantiated behavioural repertories and neural circuitry display clear evidence of alterations that correlate with the degree of expertise. Empirical evidence for these claims can be seen in longitudinal studies of expert navigators (Maguire et al 2000) and mathematicians (Aydin et al 2007) in the case of neural plasticity; and by the performance of expert bartenders in behavioural experiments which obscured task relative features that only affected expert fluency (Beach 1993).

This clearly shows that normative patterned practices facilitate discussion of developmental trajectories in a manner that is supplementary to the mutual manipulability criterion. However, if they only operated over this diachronic timescale then this would still leave us short of a full explanation since the problem of cognition distributed in time also entails that we have a medium by which to explore and understand cognitive-historical factors. In particular, the retention and refinement of cognitive resources in downstream epistemic engineering; and how this impacts on the cognitive task space. Serendipitously, normative patterned practices also allow us to conceptualise how the cultural-cognitive niche is constructed over time through the inter-generational transmission of resources and techniques. The creation, maintenance, and

manipulation of these are governed by normative patterned practices which are also transmitted across generations through hybrid learning (Menary 2012, 2014).

Although this demonstrates how normative patterned practices allow us to parse both diachronic temporal scales, *we should not just see these as mere causal background conditions*. This is a vital point. Because if this were not the case then diachronic factors would merely be enabling conditions and not constitutive of the cognitive processes of the system; and thus, not really distributed in time in the sense by which we discussed spatially distributed cognition. This disparity would not fit with the inclusive and expansive approach I have endorsed. And although I have argued that synchronic coordination dynamics are dependent upon diachronic factors, I have also claimed that normative patterned practices should be seen as partially constitutive of distributed cognitive systems. This claim now needs defending – and I shall do so by considering the conjunction of normative patterned practices and mutual manipulability.

As I have repeated throughout this chapter: normative patterned practices both shape and are shaped by the coordination dynamics of distributed cognitive systems (Vogeley & Roepstorff 2009). The interaction of the agents and epistemic resources in a putative distributed cognitive system are the *enactment* of various kinds of cognitive practices (e.g. biological coupling, epistemic actions, etc.). As such, they *are partially constitutive* of the capacities of the distributed cognitive system (Menary 2013b, p. 363), since they are both *that which is deployed in cognitive activity* as the coordination dynamics of the system, as well as the normative procedures which guide cognitive activity and alter the task-space. In this manner I argue that they meet both aspects of the mutual manipulability criterion. Because if one intervenes on the practices then one alters the cognitive behaviour of the system (bottom-up manipulation), and if one engages the system in its target behaviour then it deploys specific cognitive practices (top-down manipulation). This is why it is

appropriate to categorise a cognitive system as distributed in time – and the combination of the mutual manipulability criterion and normative patterned practices provides us with a tractable means of doing so. Furthermore, it is noteworthy that my approach here avoids a temporal version of the problem of cognitive bloat because the combination of these principles enables the investigator to select and distinguish specific sets of normative patterned practices – the medium of diachronic process – that should count as constitutive, rather than just letting in all diachronic factors indiscriminately.

Returning to an example in the navigation case study helps to clarify this matter. When the pelorus operator manoeuvres the alidade so as to acquire a landmark target in the viewfinder, the agent engages a series of motor programs to manipulate the telescope. At the same time, the agent has their visual capacities 'extrapolated' (using Humphrey's (2004) terminology) so that they can see further. This comes into conjunction with how the perceptual capacities have been transformed – what Alač & Hutchins (2004) refers to as learning "seeing as" – so that various specifics features of the environment become salient affordances. The stable anchor in the viewfinder of the alidade – a self-correcting action – not only eases the cognitive load here for the agent by keeping the target in line; what Hutchins (2005) refers to as a "material anchor". It also creates a "conceptual blend" by bringing the abstract symbol system of the gyroscope into alignment with the visual field. All of these interactions involve biological couplings that have been habituated by repeated deployment of these cognitive practices (Menary 2007a). These coordination dynamics occurring in the synchronic plane involve the actual deployment of various normative patterned practices that have been acquired through developmental trajectories by the agent themselves (thus, entailing neurological changes and alterations of behavioural repertoires). And these practices have also been devised and refined across cognitive-historical scales; this entails the epistemic engineering of the problem space itself. So, although these normative patterned practices are the medium by which we

analyse and discuss the two diachronic factors here. They are also being deployed and enacted in the synchronic plane. When we apply the mutual manipulability criterion to these various practices, a bottom-up manipulation on curtailing any of them (as shown in chapter 4, section 4.4) has a serious impact on this task-specific system core configuration (agent-artefact). And a top-down intervention would witness the agent deploying these cognitive practices in order to achieve the task. As such, this demonstrates how specific normative patterned practices, and cognition distributed in time, can be seen as partially constitutive of the activity of distributed cognitive systems.

Importantly, another point of consilience between normative patterned practices and the other two principles comes via Giere's (2002b, p. 292) declaration that understanding distributed cognitive systems requires more than simply enumerating the components. It also requires an understanding of the organization of the system. On this point it is noteworthy that van Eck and de Jong (2016, p. 19) have criticised Kirchhoff's (2014) previous attempts to combine mechanistic explanatory strategies and Menary's emphasis on cognitive practices because they contend that the latter cannot tackle the problem of cognitive bloat. I think this criticism is misplaced because Kirchhoff's claim is not about cognitive bloat but rather that distributed cognitive systems cannot be seen as just mere aggregates and must be properly organised. As my discussion of task-specificity and multilevel analysis in chapter three showed, many in the literature hold that distributed cognitive systems must have an emergent organisation (Poirier & Chicoisne 2006; Theiner & O'Connor 2010; Theiner et al 2010). If this were not the case then claims about distributed cognition would be in danger of explanatory superfluity (see Huebner 2014). Given that normative patterned practices provide us with a means for articulating how systems are assembled and organised across diachronic timescales, this again demonstrates the consilience of how the principles of my pluralistic framework interact to explain differing features of distributed cognition.

Finally, normative patterned practices provide us with a means to blunt the challenge of methodological bloat. In chapter one I acknowledged that the pluralistic framework faces a reasonable critique from the methodological individualist. Viz. the *prima facie* complexity of my position is overly demanding and attempts to include too much in the unit of analysis – so much so that it entails the paralysis of research (section 1.5). My response is to point to the wide applicability of normative patterned practices. They operate over both micro and macro spatial scales: they are cultural patterns of activity correlated with alterations to neurons in learning-driven plasticity; they sculpt the shape of behavioural repertories in the manipulation of cognitive resources, and the coordination dynamics of distributed cognitive systems; and sets of normative patterned practices partially constitute what a cultural-cognitive niche *is* (alongside techniques and tools).

They also operate across both synchronic and diachronic temporal scales: being deployed in contemporary activity; learnt and mastered in developmental trajectories; and transmitted inter-generationally in "Tomasello's ratchet" – "the accumulative improvement of innovation through social learning" (Sterelny 2003, pp. 116-117) – as the set of procedures for creating, maintaining, and manipulating cognitive resources (Menary 2007a, 2012). As such, normative patterned practices can be seen as a unifying heuristic for an investigator using the pluralistic framework. *I.e. one can use this principle as the means by which to interrogate the case study under investigation and how it operates at the various differing levels of analysis.* Therefore, despite the seemingly huge burden placed on a proponent of distributed cognition by my pluralistic framework; this excess dissipates through how normative patterned practices combine with the other two criteria.

To demonstrate this consilience, I conclude this chapter with a brief summary of how the three main principles of my pluralistic framework operate together: firstly, *task-specificity* centres the unit of analysis and provides an appropriate means to interrogate multiple levels of analysis on

their own terms (i.e. without committing mereological fallacies). Secondly, the *mutual manipulability criterion* designates genuine components of the system from mere causal background features. Finally, *normative patterned practices* provide us with a medium by which to understand how synchronic and diachronic temporal scales interact, and how neural to niche level spatial scales are related. They allow us to focus on the patterns of activity that mediate the use of culturally acquired epistemic tools – and their linkage to historical and developmental time scales. Insofar as these patterned practices are involved in the coordination dynamics that are delineated by the mutual manipulability criterion, I think this gives us grounds for asserting that these are legitimate components of this distributed cognitive system – centred on the shared task and not necessarily the agent. These diachronic elements are often crucial for understanding how a cognitive activity takes place synchronically.

I now consider this in relation to the central case study of this thesis, Hutchins' navigation team.

5.6. Summary

Many theorists accept that humans are definitively cultural animals and that this has an impact for how we understand human cognition. Enculturation is diachronic in two ways: firstly, cultural practices have a significant impact on the neurocognitive profile of the agent through the developmental trajectories by which an agent becomes inculcated into a particular cultural-cognitive niche. Secondly, the cultural-cognitive niches of humans are also special insofar that they are highly accumulative in the retention and refinement of epistemic resources across generations, which historically alters the shape of the everyday cognitive tasks facing members of the niche.

But there is an issue in how we operationalise ‘culture’ for understanding human cognition. Following Menary (and others), I have argued that we can make this tractable by using normative patterned practices: sets of cultural practices spread across a population that govern how agents and artefacts interact towards cognitive tasks. They are acquired over developmental trajectories and accumulate over generations. Elaborating on this, I also showed how the concept not only nullifies the issue of methodological bloat but also tackles the issue of how cognition is distributed in time. Normative patterned practices provide us with a mediational concept that is active and present at all spatial and temporal levels and as such simplifies the pluralistic framework and also explains the interconnection of cognition diachronically distributed in time with synchronic and spatially distributed cognition. Although I showed this to some extent in the discussion above, in the final chapter I return to the central example of the navigation case study at greater length to concretely show the importance of diachronic temporal factors for understanding distributed cognition and how normative patterned practices make this tractable.

6: Re-examining the navigation case study

…humans are fundamentally social beings. […] If humans lived in zoos, we would be classified as "obligatory gregarious".

Georg Theiner, Colin Allen & Robert L. Goldstone (2010, p. 380)

6.1. Introduction

In the previous three chapters I have argued for a pluralistic framework to tackle the many problems faced by a proponent of distributed cognition. In this final chapter I return to the central case study of the navigation team to demonstrate the practical benefits that this approach has for explicating cases of distributed cognition in greater depth. My aim here is to show how the three concepts of task-specificity, mutual manipulability, and normative patterned practices operate together to both provide deeper insights into what distributed cognition is, and shed light on the case study itself – thus demonstrating the usefulness of the framework for both practical and theoretical researchers.

The chapter is structured around considering the coordination dynamics and normative patterned practices over the three important temporal levels that I have identified. After briefly recapping the fix cycle, section 6.2 examines two important features of the navigation team in the synchronic plane: communication channels (6.2.1) and the division of cognitive labour (6.2.2). Section 6.3 builds on the second of these and shows how developmental trajectories – on multiple levels – are crucial to the organisation of the system and how the team displays *resilient* behaviour in a noisy environment. The final section (6.4) looks at the cognitive-historical factors involved in the case: in particular, our species unique penchant for spreading cognitive activity intergenerationally

in order to achieve tasks that would otherwise be impossible – what I shall refer to as "virtual collaboration" (Tomasello 1999, p. 41).

6.2. The navigation case study revisited

In chapter two (section 2.3) it was noted that although the navigation case study is sometimes portrayed in the philosophical literature as *only* being possible if it is performed by a team (e.g. Huebner 2014, p. 154; Magnus 2007, p. 298), this is actually a misnomer since the task is performed by a solitary individual when the vessel is more than eight kilometres from land. The main reason that entailed that the task is only achievable by a team was not the actual complexity of the task demands but rather the time constraint of three minutes on the fix cycle[66]. However, when the vessel is further away from shore the position fix need only be completed once per hour (and once every fifteen minutes when within site or radar contact with land) since it is not as essential to know the precise location of the vessel (Hutchins 1995a, p. 28). This is referred to as "Normal Steaming Conditions". In this situation, the solitary agent conducts the entire fix cycle – to recap (see figure 6.1 below): the agent chooses three appropriate landmarks; these are then acquired and observed as three-digit discreet numbers and recorded in the bearing log. The agent then takes these digital representations and reformats them as analogue angles on the movable arms of the hoey which are brought into coordination with the chart to create a line of position. When at least three of these are combined appropriately they triangulate the position of the vessel.

[66] I shall return to the importance of three minutes in section 6.4.

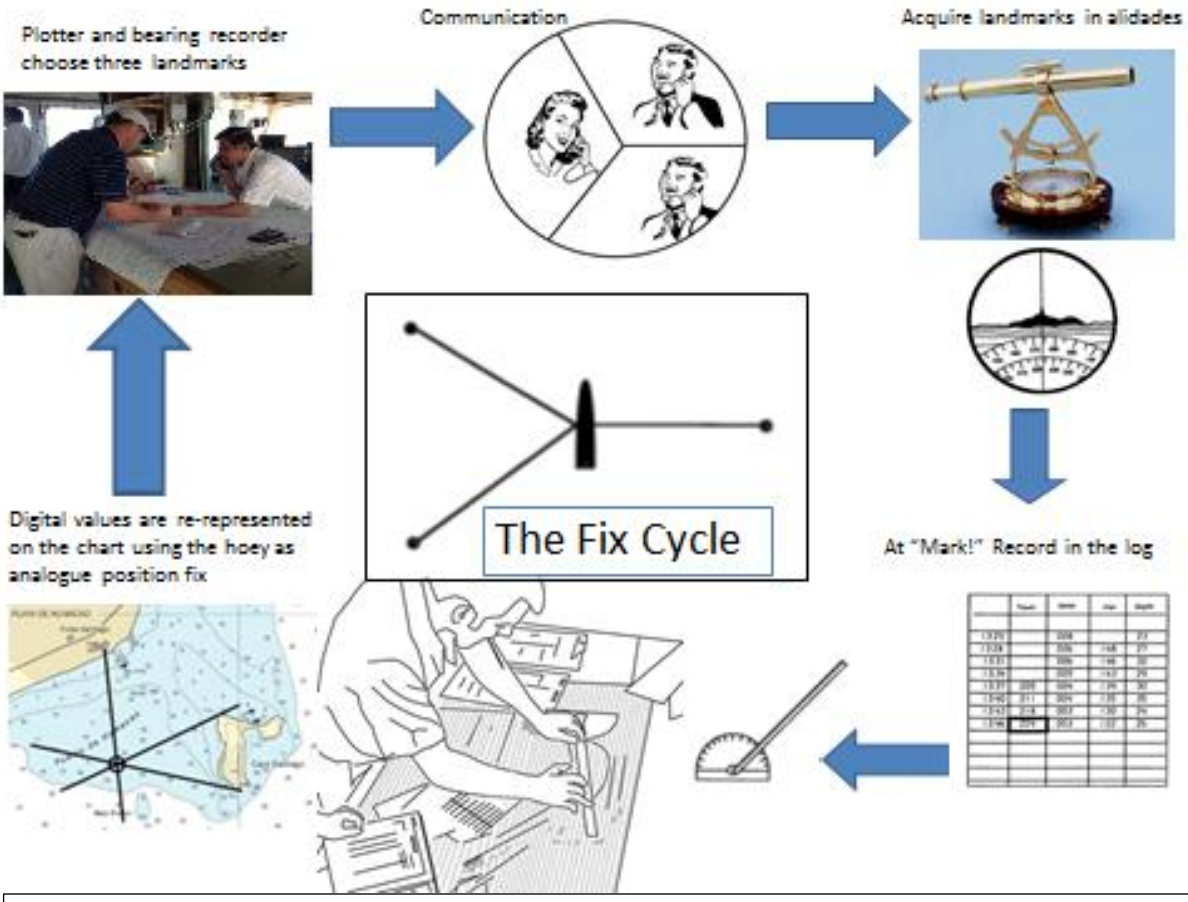

Figure 6.1: the fix cycle. Clockwise from top left panel: The plotter chooses three target landmarks. These are relayed to the two pelorus observers by the bearing recorder. The pelorus observers acquire the landmarks and these are relayed as three-digit numbers which are recorded in the bearing log book when the bearing recorder calls 'Mark'. These are reformatted by the plotter as analogue representations on the movable arm of the hoey, and blended with the chart so that the three measurements triangulate the position of the vessel. From here Future positions of the vessel are projected using dead reckoning begins

Hutchins (1995a, 2006, 2010c) argues that this latter act is much more sophisticated than is usually appreciated since what is being constructed is a complex representational schema that blends several forms of conceptual space with physical space: e.g. the physical space of the chart; the conceptual iconic representation of the surrounding environment; and the conceptual spaces of the Mercator projection (the mapping framework which includes several discrete measurement systems: compass directions; a longitudinal and latitudinal grid system imposed on the world; and

the nautical mile)[67]. He also adds that such a sophisticated cognitive achievement is made possible by the fact that the chart and the hoey act as "material anchors" (2005, 2010c) – i.e. the physical structures of the hoey and the chart act as scaffolds through which looping interactions from the plotter manipulating them makes the conceptual representations more stable (see figure 6.2 below). In part this stability is enabled because representations are "frozen" (Huebner 2014, p. 179) allowing for, and simplifying, more prolonged engagements and refinements (also see Clark 2008; De Cruz & De Smedt 2013; Kirsh 2010; Menary 2007b). It also enables the conceptual blending of several representational systems and formats in the visual field (Hutchins 2005, 2010b, 2010c). This complex representation is used to calculate not only the current position of the vessel but also project the future position of the vessel using "dead reckoning"[68] – this then guides decisions about choosing future landmark targets.

Figure 6.2: the plotter uses the flexible arms of the hoey in conjunction with the chart to reformat the digital coordinates of bearing reports into lines of position. The material features of these cognitive artefacts have a representational function that complements internal processing through being both stable in the visual field and also enacting computational processes through their manipulation.

[Figure from Hutchins (2010b, p. 441)]

For our current purposes, the key feature here is how the agent goes about deciding in which order to shoot the three landmark targets. The appropriate guiding norms for organising in what order the agent shoots the three target landmarks is a rule of thumb: "shoot beam bearings

[67] I discuss this particular coordinate system in more detail in section 6.4.

[68] This is where the navigator monitors the motion of the vessel from its previous known location by keeping track of how fast the ship is moving and which direction it is travelling in (Hutchins 1995a, p. 56).

first" (Hutchins 1995, pp. 206). Beam bearings are landmark targets which are closest to the horizontal or orthogonal angle off either side of the vessel (for example, in the central panel of the figure 6.1 above the beam bearing is on the starboard side). Agents are taught to shoot these targets first because they are the most subject to change over time as the vessel is moving. This is because a bearing – and subsequent line of position – is the measurement of an angle. The closer a landmark target is to the horizontal the more this angle is subject to change as the fulcrum of the angle (i.e. the vessel) moves. This is also a factor with other landmarks as well – hence the necessity to perform the task quickly with a team – but it is most pronounced in the case of beam bearings.

Drawing on the discussion of normative patterned practices in the previous chapter (section 5.4.1), we can label this as a "manipulative cognitive norm" (Menary 2007a) since it orchestrates how the cognitive labour is subdivided. However, there is an interesting issue here because although "shooting beam bearings first" is a good rule of thumb whilst one individual is completing the whole task; when the same task is completed by a team it leads to process losses because it *fails to take into account the differences between group and individual cognitive properties*. This failure is a perfect example of where my pluralistic approach to distributed cognition can be of use because it places an emphasis on multiple levels of analysis and the importance of paying attention to the components and the differences between the individual and system level (Hutchins 1995a, pp. 174-177; Menary & Gillett 2017; Norman 1991).

In this particular subtask, Hutchins identifies that the group system is different to the individual system in two important respects with regards to two cognitive properties: [1] communication; and [2] parallel activity (1995a, p. 284). I will now discuss each of these in turn in the following sections (6.2.1 and 6.2.2 respectively). The most obvious difference between the group performance and individual performance is that there is no stage where the decisions made in

conjunction with the chart (as a material anchor) have to be relayed to separate observers. Instead, the single agent makes the decision and then carries out the action themselves. Some accounts of this communicative act sometimes skip over the details here; but, as I have argued, it is these specific details that are of central importance for properly understanding why distributed cognition is a useful explanatory framework – as I shall now demonstrate.

6.2.1. Communication, the importance of buffering, and the division of labour

In the division of cognitive labour within the navigation team, the task of monitoring and facilitating communication is delegated to the bearing recorder; and this is their primary function alongside the recording of observations in the record log and time keeping. In particular, Hutchins picks out a specific aspect of this compartmentalisation of communication: "buffering" (1995a, p. 195). The bridge of the ship is an incredibly busy place, and the fix cycle is not the only task that is taking place. Alongside the navigation team there are a number of other crew members who are all performing tasks and making noise. Hutchins presents an especially clear example of how many points of confusion can arise in the bridge whilst the fix cycle is being performed because much of the cross-talk from other departments in the bridge also involves three-digit numbers – which exacerbates the likelihood of error (p. 233-234).

Buffering is a way of preventing higher error rates which can be caused by cross talk at the group level by controlling the speed and timing with which communicative signals move from one part of a system to another. Hutchins states that "buffering prevents the uncontrolled propagation of effects from one part of the system to another" (p. 195). In this case, the combination of the bearing recorder and bearing log act as a buffer between the pelorus observers and the plotter. The plotter is taking the observations from the two observers and reformatting them into position fixes

– as noted above, this is an incredibly complex task. If, as Brown (2011, pp. 19-20) erroneously describes, new observations were being delivered straight to the plotter whilst he was carrying out this activity this would likely lead to distractions.

For example, imagine that the plotter is examining a line of position and taking a three-digit number and converting this into an angle on the movable arm of the hoey when suddenly a new landmark report comes in. This new three-digit number could easily lead to confusion and cause the plotter to construct an incorrect line of position. Norman refers to these kinds of errors as "data-driven" "action slips" – where attention is unconsciously misdirected away from one's task goal by the current sensory input (1987, p. 105, 109; also see 1986, pp. 254-255). To prevent this, the bearing recorder acts as a break in this communication channel and thus makes the performance of the team more *resilient* by reducing error rates caused by cross communication[69]. The message between observer and plotter is mediated by the bearing log which then allows the plotter to tackle the complicated construction of the multimodal representation that is the position fix with fewer distractions, and at a less hasty pace. In these regards, the bearing log acts as an exogram (Donald 1993, 2010) – an external memory source that alters the cognitive load on working and long-term memory (as discussed in chapter three in the case of pre-flight checklists).

Hutchins argues that communication is an irreducibly supra-personal cognitive property in a socially distributed cognitive system (1995a, p. 284). To see this, consider how the agents in this core configuration interact. Any task that is performed by a group involves joint action and attention, collaboration, and a division of cognitive and physical labour which transforms the organisation of the task-structure and the coordination dynamics of the system (also see Wertsch 1985, pp. 64-65). This means that the task structure and coordination dynamics are fundamentally different at the

[69] I return to issues of resilience and flexibility and how this involves diachronic temporal factors in section 6.3 below.

social level than if the same task were performed by an individual. Not only are there the coordination dynamics that arise in which an individual manipulates environmental resources in conjunction with internal resources, but there are also a whole range of behaviours – guided by normative patterned practices – which dictate how individuals collaborate and perform joint actions. Hutchins claims that any division of labour, physical or cognitive, involves distributed cognition:

> All divisions of labour, whether the labour is physical or cognitive in nature, require distributed cognition in order to coordinate the activities of the participants. Even a simple system of two men driving a spike with hammers requires some cognition on the part of each to coordinate his own activities with those of the other. (1995a, p. 176)

This is an example of what is called a "joint action". Knoblich & Sebanz define joint action as "…social interactions wherein two or more individuals coordinate their actions in space and time to bring about a change in the environment" (2006, p. 100)[70]. Key to understanding collaboration and the organised division of physical and cognitive labour is communication.

Communication is a supra-individual cognitive property that changes the nature of the task space – and thus the coordination dynamics that form the task-centric core configuration. It is worth noting two points about communication. *Firstly*, communication is highly multimodal – involving speech, gesture, tools, context, and other embodied features such as eye gaze and facial expression, etc. (Hutchins 2010c). In chapter three I showed how differing external representational mediums can have differing impacts on cognitive behaviour (section 3.3.1). The particular cultural practices of a cognitive-cultural niche also impact on how lines of communication are structured and formatted (e.g. what is communicated and how). *Secondly*, communication is a property that varies greatly for different core configurations. Small units can be very intimate but in larger units the role of

[70] Also see Butterfill and Sebanz (2011) and Galantucci and Sebanz (2009) for overviews which explicitly draw links between this literature and distributed cognition in regards to how the primary unit of analysis for cognitive science need not always be a solitary agent but can be multiple interacting agents.

buffering and codifying the propagation of representations becomes more important for maintaining a coherent system (see Huebner 2014, pp. 165, 170, 179-180).

Hutchins identifies four parameters across which communication varies: [1] the patterns of interconnectivity between the agents; [2] what is communicated; [3] the strength of this communication; and [4] temporal factors (1995a, pp. 248-252; also see Thagard 2006, p. 83). The way these variables interrelate have significant impacts on cognitive properties at both the individual component and group system level. Importantly, in the case of the latter, they are not merely aggregative of the members of the group. Instead, they cash out the emergent organisation of the system. This influences how, for example, a memory or decision-making task is distributed and enacted. For instance, in the navigation case study the patterns of interconnectivity between the pelorus operators and the plotter are buffered through the recorder; and what is communicated differs depending on direction. From the plotter come linguistic landmark descriptions, projected digital headings, and error corrections (see section 6.3 below), and from the pelorus operator come queries and digital headings. Both are relayed by the recorder and both involve a mixture of mathematical and natural languages, and theoretical vocabulary. But it is also important to note the relative strength – persuasiveness – of particular communicative acts. This involves considering the social hierarchy and expertise (topics I cover in section 6.3). All of these factors[71]; shape how a cognitive task is distributed across the team.

To returning to the notion of buffering in general as it is in the wild, it is important to note that it can have both a positive and a negative role. For instance, in the context of scientific communities and the division of cognitive labour, both Thagard (1993) and Hutchins (1995a, pp. 239-261) have separately identified that asymmetric communication channels at the niche or

[71] I discussed the temporal aspect above in regards to avoiding action slips through the use of buffering.

community level is essential for mitigating factors of conformity bias that leads to groupthink (a common problem – as I discussed in chapter three, section 3.4.2). Buffers prevent information from propagating too quickly and too far through a system and this enables the possibility of alternative views to coexist and be generated whilst different parts of the overall system have access to different sets of information. On the other hand, sometimes asymmetric communication systems, or types of information which are hard to transmit – what von Hippel (1994) calls "sticky information" – can prevent the building of necessary consensus or can prevent an important piece of information from reaching the right location in a community (also see Zollman 2010). This can exacerbate group polarisation and echo chamber effects (Iyengar & Hahn 2009).

The balance between allowing too much or too little information flow is a general problem for all complex systems – too much and the system will descend into chaos with the slightest naturally occurring perturbation, too little and the system won't be flexible enough to handle errors and the vagaries of nature (Kauffman 1995). Divisions of labour take on different forms depending on the level of analysis. Discussions of the division of cognitive labour in scientific communities have often focused on the diversity of social types of agents that comprise the system (e.g. Kitcher 1993; Kuhn 1977; Weisberg & Muldoon 2009). But I submit that normative patterned practices play a strong role in shaping the coordination dynamics involved. I shall now discuss this in the context of the collaborative performance of the fix cycle (whilst acknowledging that the change in core configuration limits the analogical inferences one can draw from this point).

6.2.2. Parallel activity and the importance of cognitive practices

Another irreducible cognitive property that Hutchins identifies in the navigation team is related to communication; and the possibilities of the division of cognitive and physical labour that are not

present when there is only a solitary agent to perform the task. As stated above (section 6.2), when deciding in what order to perform the task of measuring the angles of the landmarks, the solitary individual follows the rule of thumb: 'shoot beam bearings first'. When the team perform the same task of shooting the landmarks they also follow the same rule – for the same reason (viz. that beam bearings are the measurements most susceptible to error if they are delayed).

However, Hutchins (1995a, p. 268) points out that this is a mistaken or suboptimal division of labour leading to process losses[72] because it fails to utilise the capacity of the system for "parallel activity". When the task is performed by a single agent all the actions must be conducted in a linear fashion since the agent can only focus properly on one feature of the task at a time, whereas "…socially distributed cognition can have a degree of parallelism of activity that is not possible in individuals" (ibid). But in the team, there are two observers and three targets; as such, *the structure of the task has changed* because the resources have been altered. Steiner notes that there is no straightforward relationship between changes in group size and productivity (1972, p. 67; also see Collins & Guetzkow 1964). Arguably this is yet another reason why any approach to distributed cognition must pay attention to the particular core configuration of the system under investigation, and why it is also unwise to draw analogies between systems of vastly varying size (like comparing the navigation team and CERN *a la* Giere (2002a, p. 639)).

[72] As defined in chapter three, process losses are the reduction in the effectiveness of the behaviour of the system through inefficient coordination dynamics (Steiner 1972).

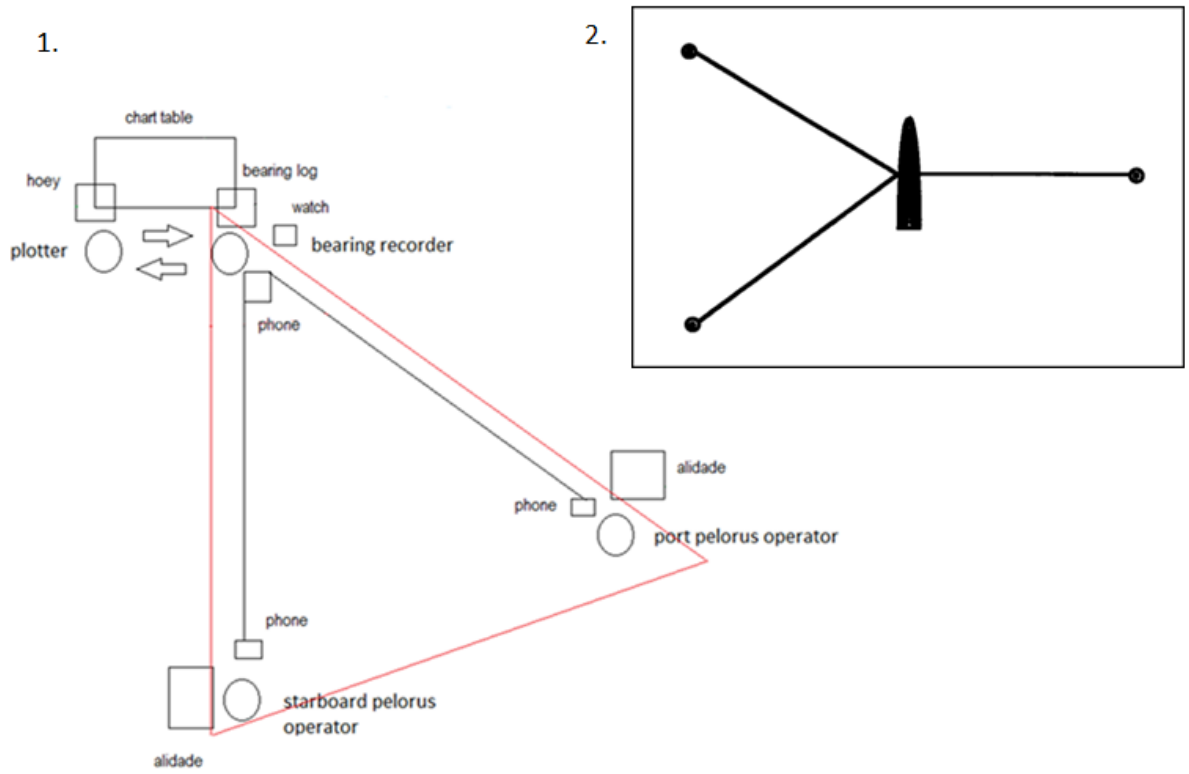

Figure 6.3: The communication network subsystem operating in the navigation team (highlighted in the red in panel [1]): between the two pelorus observers and the bearing recorder with the phone network linking them. In panel [2] three possible landmark headings are shown around a bird's eye view of a vessel. In this case, there is a beam bearing on the starboard side of the vessel. [Panel 2 adapted from Hutchins (1995a, p. 208)]

When the rule of thumb 'shoot beam bearings first' is used to organise the coordination dynamics of the how the pelorus operators jointly perform the observation task it increases the likelihood of process losses. For instance, in figure 6.3 (above), I have stripped the navigation team down to the essential features that we are discussing here: the two pelorus observers and their alidades; the bearing recorder with the bearing log and watch; and the phone network linking them. Hutchins analyses how this subsystem interacts at great length in order to make the point that the capacity for parallel activity is occluded because the team has wrongly transposed a normative

strategy that works well at the individual level which subsequently fails when it scales up to the group level (1995a, chapter 4).

We can make this point clear by considering a hypothetical but highly possible position fix in which the port pelorus operator has to shoot two targets and the starboard pelorus operator only one beam bearing target (see panel 2 in figure 6.3 above). We can organise the behaviour of the team in two ways. *Firstly*, if organised using the standard rule of thumb 'shoot beam bearings first', then when the bearing recorder calls "Mark!" the starboard operator acquires his target first and reports it whilst the port operator is inactive and waits. It is only after the starboard operator has finished that he then acquires and shoots both his targets. Let us call this strategy *A*. In contrast, we can imagine an alternative arrangement in which the division of labour is orchestrated by an alternative normative strategy (*B*): 'whichever pelorus operator has two targets goes first'[73]. Strategy *B* is designed to make optimal use of the capacity of the system for parallel activity. So, when the bearing recorder calls "Mark!" both operators acquire targets but the port operator reports his bearing first and then, while the port operator acquires his second target, the starboard operator reports his target. After which the port operator reports his second target. The speed in which this can be done mitigates any worries that might arise with the beam bearing not being shot first. This can be made especially clear by comparing the two strategies using activity charts:

[73] Hutchins formulates this rule of thumb as follows: "The operator who has two bearings should report before the operator who has only one, and he who has two bearings should always shoot the beamier of the bearings before the other" (1995a, p. 218). Additionally, he is keen to stress that it is a much more complex rule of thumb (with multiple clauses) than what is needed at the individual level.

Strategy A:						
Port operator	Wait	Wait	Shoot	Report	Shoot	Report
Starboard operator	Shoot	Report				

Strategy B:				
Port operator	Shoot	Report	Shoot	Report
Starboard operator	Shoot	Wait	Report	

In these charts, it is clear that strategy *A* takes a much longer period of time to be completed and involves multiple points at which one of the agents is left idle; and therefore, entails process losses. As such, by using strategy *A* the team is performing well-below the optimum; leading to a less accurate position fix. As Hutchins puts it strategy *A* "…fails to take advantage of the parallelism of activity that is possible with two pelorus operators" (1995a, p. 207).

A key point here for explaining the cognitive behaviour of this system is that the normative patterned practices – the repeated sets of motor-patterns that normatively constrain and mediate (as in right and wrong ways of being enacted) the coordination dynamics of internal and external resources – have a significant impact on how the cognitive task is achieved. In particular, it is *the incongruence of adopting a useful strategy at the individual level and how it can lead to negative results at the group level.* Work by Harris and colleagues (2011) suggests that the issues of differing strategies at different levels of organisation is a common phenomenon in ecologically salient human cognition (also see Lave 1988). This highlights the importance of multilevel analysis for examining cases of distributed cognition. Harris and colleagues studied the performance of long-term married couples at tackling memory tasks – both individually and collaboratively. A key reason for doing this

research is that much work on group performances in memory tasks has looked at non-ecologically salient collections of strangers who perform worse than nominal groups – composed of the members performing in isolation and having their scores combined afterwards (see Barnier et al 2008 for a review). This surprising but stable finding of process losses is referred to as "collaborative inhibition". The aim of research by Harris and colleagues (2011; also see Barnier et al 2008; Sutton et al 2010) was to see if collaborative inhibition effects could be overcome by "expert teams" (Eccles & Tenenbaum 2004; Williamson & Cox 2014). Expert teams are not merely aggregates of individual experts, they are emergent systems that have fluid coordination dynamics and consilient shared normative patterned practices that guide these interactions (also see Theiner et al 2010).

A critic might respond that one can explain these factors from a methodological individualist stance. To this there are two responses: firstly, we can respond to this challenge by pointing to the manner in which the mutual manipulability criterion carves off the coordination dynamics of the system from its causal background. Insofar as these coordination dynamics are constitutive of a cognitive task then they are supra-individual cognitive properties. In the case of expert teams this can entail an "assembly bonus effect" in which collaborative performance towards a shared task outstrips the capacities of the individuals who comprise the group (Collins & Guetzkow 1964; e.g. see Meade et al 2009). However, my point here is not that socially distributed cognition requires an assembly bonus effect (since this would nullify the navigation team who display process losses); rather, what is important is how the organisation of core configurations structured by normative patterned practices is *emergent* and must be analysed in order to properly understand the cognitive behaviour of the system and its components (also see Michaelian & Sutton 2013; Theiner & O'Connor 2010).

For instance, although Harris and colleagues' analysis of long-term married couples showed variable performance at a memory task when performing as a dyad it also showed that collaborative facilitation correlated with dyads that had shared and complementary strategies for performing the task at the group level and which were also not incongruent with individual level strategies. As they put it: "Our data suggest that what is more important than the similarity or differences between individual strategies is the presence or absence of a group level strategy that coordinates recall and capitalizes on the relative knowledge and skills of individual group members" (2011, p. 297). When a couple had conflicting or differing strategies at the individual level this led to collaborative inhibition.

This again shows how normative patterned practices are involved in understanding cognitive behaviour in ecologically salient task settings. But more importantly, it also demonstrates the need to properly distinguish between levels of organisation in analysis because these impact on how a system performs a cognitive task. If we only focus on the group level, then our analysis will fail to take this into account; thus, leaving our account of distributed cognition unable to properly explicate this key difference between the individual level and group level. By having a task-specific approach my pluralistic approach is well-suited to accommodating this important multilevel factor. Conversely, methodological individualism by not even accepting higher levels of cognitive organisation overlooks this altogether and therefore cannot properly explain the cognitive accomplishments going on in this system nor how they could be refined where there are process losses.

A second response to the deflationary critic is to point out that this discussion has only considered the task phenomenon in the synchronic time frame. As I have consistently argued, this is insufficient for understanding distributed cognition properly. For instance, key to understanding how expert teams are established and maintained is the "shared history" of these units of analysis

(Bietti & Sutton 2015). I will now outline these diachronic scales extensively in the remaining two sections of this chapter.

6.3. Developmental trajectories and resilience engineering through the degeneracy of cognitive practices

Having outlined two supra-individual cognitive properties present in the navigation team at the synchronic level I now turn to diachronic factors. In this section I discuss how developmental trajectories are essential for understanding how and why the navigation team case study is a distributed cognitive system. Against the notion that focusing on information flow is sufficient for adequately demarcating distributed cognition, I argue that synchronic activity (e.g. information flow, the propagation of representational states, coordination dynamics, etc.) is only possible insofar as it is embedded within a diachronic unit of analysis that is the condition for this integration. Importantly, in regards to multilevel analysis outlined in chapter three, this relates not just to learning on an individual level but also on a collective level.

Crucially, developmental trajectories are also involved in how the task is accomplished to a high degree of success despite the activity taking place in a noisy environment and the team performance being ridden with errors (due to the high turnover of personnel and relative inexperience of some of the team members). The specialised developmental pathway of the team members is a social hierarchy that fortuitously matches the computational complexities of the division of labour. This leads to capacities for error detection that I argue is linked to the notion of *degeneracy*. Degeneracy is a group level property for resilient behaviour through structurally diverse repertories for achieving the shared task. I argue that this is only possible because of both the

organisational structure of the team and the individual diachronic temporal factors involved in learning.

With some notable exceptions (e.g. Nersessian 2005; Sutton 2008), how cognition is distributed across time is a sorely underappreciated aspect in the literature on distributed cognition. As noted in chapter two (section 2.5), many proponents of distributed cognition (Cheon 2014; Giere 2006; Huebner 2014; Magnus 2007) take the key point of the navigation example to be that the task is subdivided so that each member of the team or specialist subunit of the larger system "only needs to know what to do when certain conditions are produced in the environment" (Hutchins 1995a, p. 199). Hutchins describes this using the term "daemons" from computer science which refers to an agent that monitors the world waiting for certain conditions to pertain and then taking a limited and specific action (1995a, p. 191). As Hutchins goes on: "Each participant *knows* how to coordinate his activities with the technologies and persons he interacts with" (p. 200 my emphasis). This forms an interlocking system so that "The whole [fix] cycle is something that emerges from the interactions of the individuals with one another and with the tools of the space" (ibid).

Whilst it is certainly the case that there are coordination dynamics operating between agents and artefacts, and agents with each other (as I discussed and made concrete by using the mutual manipulability criterion in chapter four); to stop our analysis at this juncture would be to support a limited view of distributed cognition which I do not think is ultimately viable because this overlooks a crucial point. Viz. we must also understand how the coordination dynamics are shaped by diachronic processes that dictate how they are enacted: as developmental trajectories in which agents have to learn how to use the tools and techniques to do their particular subtask and how to work with the team (Hutchins 1995a, pp. 21, 169, 264-268).

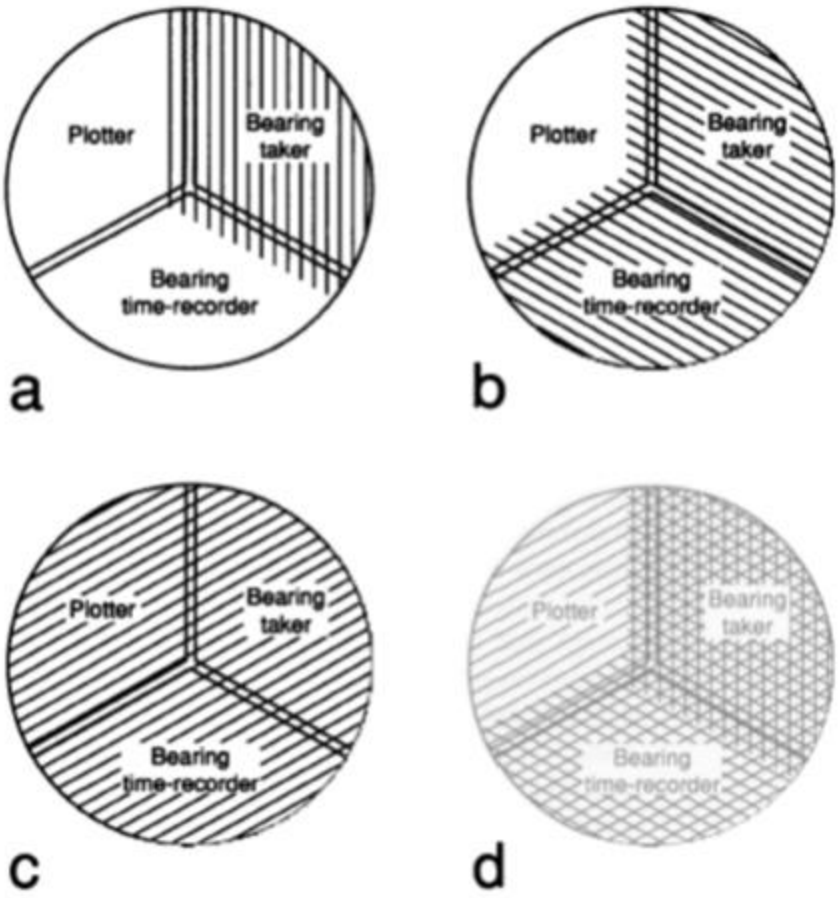

Figure 6.4: knowledge is distributed throughout the team unevenly. Panel A shows that the pelorus operator has the least amount of overlapping knowledge beyond their particular task. Panels B and C show the increasing knowledge the recorder and plotter have of the overall task, as well as their own particular duties, as we move up the social hierarchy and the matching developmental trajectory of team members. And Panel D shows the degeneracy of the system as a whole entailed by this overlapping distribution of knowledge. [Diagram from Hutchins (1995a, p. 266)]

As outlined in previous chapters, developmental trajectories are the ontogenetic processes by which an agent acquires and masters a set of normative patterned practices that are involved in a particular cognitive task (Menary 2007a, 2015). These have a significant impact not only on the agent's neurocognitive profile (altering their neural structures and behavioural repertories) but also transforms how they approach the task space (both in terms of affordances and also what cognitive resources they can utilise). But developmental trajectories also have a particularly social aspect

which the navigation team case study highlights. Novice members join the team as pelorus operators, performing the simplest task (observing landmarks) of the fix cycle before moving up the social hierarchy. Nersessian stresses that this is an initiation rite for community membership into the cultural-cognitive niche (2006, p. 708). But it is also important to recognise that this entails that higher ranked team members have experience of performing lower ranked roles of the fix cycle (see figure 6.4 above). This has several important outcomes.

Firstly, it is important to emphasise that the individual developmental trajectories of the members of team lead to overlapping and distributed knowledge structures. This is not just about what the other does but what they *should* do – this entails that it is a normative relation. Additionally, this normative aspect entails that this epistemic function is also suffused with social and ethical features: e.g. if one fails to perform one's functions appropriately this is not just a cognitive failing but a failure of one's duties in a social division of labour (see Forsyth 2009; Hutchins 1995a, pp. 190-191; Moore & Rocklin 1998; pp. 110-111). This blends the social and cognitive aspects of what agents "ought" to do towards shared cognitive tasks.

Nisbet and colleagues (2001) in their comparative longitudinal cognitive-historical analysis of human enculturated cognitive domains have commentated that there is an important link between social hierarchy and how cognitive tasks are approached in certain settings. This is particularly important with regards to the navigation team which has both a high turnover of personnel and large number of novices comprising its agential components (Hutchins 1995a, p. 263). This results in a high amount of errors in the team's performance. However, this issue is mitigated and managed because of the serendipitous social organisation of the team: viz. developmental trajectories match both the social hierarchy and the computational structure of the task (see figure

6.4 above). This entails that those with more social power also have more expertise – something which is not always the case in divisions of cognitive labour (Moore & Rocklin 1998, pp. 110-111).

Secondly, a fundamental feature of how the team is able to communicate is based on these overlapping knowledge structures. Both the plotter and the recorder have had experience as pelorus operators and so have a good understanding of what they are asking the incumbents to do (pp. 267-268). This in turns effects and constrains their decision making with regards to the choices of landmarks: viz. viewing angles that might be obscured due to oddities of ships infrastructure; whether a pelorus operator has been to this particular port before; how to describe a landmark visually so that they can find it, etc. (pp. 267-270). Rather than acting as daemons with a limited engagement outside of their particular subtask space, the recorder and especially the plotter are constantly drawing on their experience of having performed other duties in order to perform their own role. *As such, understanding these developmental trajectories are vital for properly explicating how the team actually functions synchronically as a unit*. Furthermore, it is important to note that this is an emergent system level property, which is not straightforwardly an amalgamation of individual cognitive properties: viz. the capacity of the system to act in a flexible and resilient manner. Theiner and O'Connor list the capacity of a system to adapt its behaviour and modify its environment as two exemplars of a cognitive system (2010, pp. 82-83; also see Poirier & Chicoisne 2006). Crucially, the structure of the team's developmental trajectories enables good *error detection* and *correction*. I claim that these aspects are not apparent if one adopts a methodological individualist stance.

The structure of this social hierarchy is crucial to the high performance of the team given its circumstances: [i] a high turnover of personnel; [ii] the presence of novices, and [iii] a noisy environment entail that there are a high number of errors. But the existence of errors per se is not

fatal; indeed Norman (1986, 1987) notes that it is a typical feature of human behaviour to make frequent mistakes in one's activities – even for experienced agents. Given this state of affairs, Norman has consistently argued that one should try to design a system's structure not only to try and reduce error but also to make errors easier to discover and correct when they inevitably occur (1987, p. 131). Namely, one should design systems with human error in mind. Hutchins, following Norman (1987), identifies the three goals of designing for error: [1] to eliminate, avoid, and prevent errors; [2] to facilitate detection and recovery; and [3] to facilitate learning from errors so as to reduce their likelihood in the future (1995a, pp. 272-279). I contend that in the performance of the fix cycle many of these goals are met due to a serendipitous match of the developmental trajectories of agents, the structure of the social hierarchy, and computational structure of the fix cycle.

A key concept for understanding this is what Hutchins calls "horizons of observation" – these are the lines of communication and observation between the agents who comprise the team performing the shared activity (1995a, p. 268) (these are portrayed as the hatched regions in figure 6.4 above). For instance, when the plotter constructs the lines of position that form a position fix he does so in public space which can be observed by the bearing recorder (Hutchins 2006; 2010b). These are areas in which there is an overlap of expertise and knowledge between members of the team and the particular tasks they are engaged in. In turn, this overlap facilitates a high rate error detection and correction because team members further up in the hierarchy have been trained in how to perform lower tasks (p. 279). As members of the team become enskilled and master these activities they are then promoted up the hierarchy to the next level of the task (e.g. from observations with alidade, to keeping time and noting bearings).

I think that Hutchins fails to emphasise that horizons of observation are not just synchronic organisational structures, but should also be recognised as constructed diachronically through developmental trajectories[74]. The plotter and the bearing recorder are able to monitor the behaviour of the pelorus operators precisely because they have prior experience of how to perform this task. This means that if the pelorus operators make inaccurate bearing reports, these errors are much more likely to be spotted because the plotter and bearing recorder will have some experience-based expectations as to what they should be. In regards to the above three desiderata for designing a system for human error we can understand the social hierarchy of the team as improving the detection and correction of errors (see Hutchins 1995a, pp. 276-277). Additionally, the social hierarchy also utilises the inevitable eventuation of errors as *opportunities for learning*. Not only does the novice who made the mistake learn by receiving feedback, but because this takes place in a public forum it also: [1] contributes to the learning of others in the team and improves their skills in error detection, and [2] gives experts an opportunity to practice their skills in error correction (pp. 277-279). Hutchins adds that this can occur both through the mere observation of error correcting practices or through collaborative efforts. By continuously working together agents come to have shared expectations of how their compatriots are likely to act, and collectively they "build and refine their shared knowledge" (Williamson & Sutton 2014, p. 126).

Another important area in which we can refine Hutchins' analysis is his description of the social hierarchy as having redundancy in regards to the distribution of expertise and knowledge, which in turn enables *resilient* and *flexible* behaviour: "If one human component fails for a lack of knowledge, the whole system does not grind to a halt. […] In response to a breakdown, the system adapts by changing the nominal division of labor" (1995a, p. 223). I think we can refine how to

[74] To be clear: although Hutchins does discuss cognition distributed in time, I think he fails to make it really clear for the reader that this synchronic activity is dependent on diachronic factors.

understand error detection by acknowledging the distinction that can be made between redundancy and degeneracy. Mason (2014) notes that although this distinction is often overlooked; degeneracy is a system level property that allows for a more refined understanding of how systems can respond adaptively to perturbations. Briefly, we can summarise the distinction as follows:

- Redundancy: the same kinds of structure performing the same function.
- Degeneracy: different structures performing the same function.

A simple example helps explain this difference[75]: imagine two ways of heating a room. One way is to have three radiators linked up to a gas boiler system (like a conventional central heating system in the UK). This equates to redundancy since if one of the radiators fails for some reason – e.g. a broken water pipe – this is compensated for by the other two. In contrast, another way to heat the room might be to use an electric heater running from the electrical mains, a gas heater, and a wood burner (as might be more common in an Australian home, which commonly lacks central heating). This is degeneracy, which Mason and colleagues defines as the notion that "different pathways can lead to the same output" (2015, p. 1).

The key difference between the two can be summarised as follows: with redundancy identical structures are functionally similar, whereas with degeneracy dissimilar structures are functionally similar. As such, the latter system is more resilient and flexible. For instance, in the above scenarios if there is a break down in the gas supply in the UK home, then it is does not matter whether the three radiators are hooked up to a single boiler or three boilers, the system cannot perform its function. Whereas in the Australian home, because each of the heating units relies on a different power source – making them structurally dissimilar – the system is able to operate and achieve its goal across a wider range of the contingencies and vagaries that are so common in

[75] A simpler version of this example was originally conveyed to me by Greg Downey and is used here with permission.

nature (especially in human environments which are extremely changeable because of the varied ecological and social conditions in which we live (Sterelny 2003, 2012)).

However, before moving on to consider how this informs us here, first we must disentangle the concept of degeneracy from some misconceptions. Mason and colleagues (2015) observe that the term has powerful negative connotations in popular culture, and combined with the fact that it is often mistakenly conflated with redundancy, they claim that this is why it has been overlooked. Mason (2014, p. 13) goes into great detail in how the "de-" prefix in modern culture is often attached to negative connotations where in the past the term was neutral. This is unfortunate because: "Degeneracy can give rise to *novelty* in ways that redundancy cannot" and thus gives us insights into complex systems that we might otherwise miss (Mason et al 2015, p. 1. My emphasis). My claim here is that this concept can also allow us to think about distributed cognition in an interesting and new way. I.e. we can use the concept of degeneracy to bolster Menary's (2007a) arguments about the importance of cognitive practices by providing details about the structural multiplicity of cognitive practices for identical cognitive functions: viz. how this entails both more resilient and flexible cognitive behaviour (Theiner & O'Connor 2010); and also enhances the cognitive system's capacity for novel behaviour (Menary 2015).

In the navigation case study, we can see degeneracy of practices in the deployment of strategies for overcoming errors and other difficulties. For instance, in the case of checking the report of a putative landmark heading the plotter not only draws on his personal memory of having performed this same role (which would be structurally similar in kind and therefore only redundancy) but he also utilises the exogram features of the navigation chart. The target landmarks have been chosen based on the plotter's calculations of the vessel's dead reckoning – his prediction of where the ship would be based on the last position fix given that the heading and speed do not alter. As

such, the plotter's manipulations of the chart contain expectations for what the values of these headings should be so that lines of position intersect to form a new position fix.

I argue that this cognitive activity – an epistemic action – creates a *structurally different* mode of double-checking the bearing heading values. In regards to horizons of observation this means that there are two structurally different domains in which error detection and correction takes place – audibly (utilising on-board memory) and visually (utilising the chart as an exogram). This degeneracy improves the resilience of the position fix because, as discussed in chapter three, each of the sets of normative patterned practices and structural features of the associated external representations will make some features of a task more or less salient (also see Vorms 2012; Zhang 1997b; Zhang & Norman 1994, 1995).

Degeneracy of cognitive practices provides us with a means of understanding adaptability which is an exemplar of cognitive systems (Theiner & O'Connor 2010); but it is important to reiterate that this degeneracy and the resilient performance of the task by the team in adverse conditions is only made possible based on the developmental trajectories operating at both the individual and group level that facilitate these properties. Hollnagel and colleagues define resilience "as the ability to: [1] respond to events; [2] monitor ongoing developments; [3] anticipate future threats and opportunities; and [4] learn from past failures and successes alike (2011, p. xxx). "To be resilient, a system always keeps an eye on whether its adaptive capacity, as it currently is configured and performs, is adequate to meet the demands it will or could encounter in the future" (p. 122).

Hollnagel and colleagues go on to discuss the construction of "distributed decision making environments" (pp. 48-53) which can facilitate the division of cognitive labour and identify three basic patterns for how adaptive systems fail (p. 127): *decompensation* (an inability to adapt to changes causes a cascade of failures); *working at cross-purposes* (process losses generated through

conflicting coordination dynamics); and *getting stuck in outdated behaviours* (overly focusing on past success and failing to adapt). Resilience engineering is the attempt to control a system's adaptive capacities in an optimal manner based on empirical evidence. I contend that the discussion of distributed cognition – particularly the degeneracy of normative patterned practices – in relation to the resilient behaviour of the navigation team contributes to this enterprise. It is noteworthy in these respects that Hollnagel (2001) has previously recognised the importance of distributed cognition to these discussions.

In general, the differing strategies offer different affordances – and this point is amplified by the fact that many of these strategies can be combined. For instance, co-gesture speech in combination with physical external representations (either body parts, or tools) can create sophisticated conceptual blends and complex combinatorial representations (Hutchins 2010c), and in turn these will draw out and make more or less salient differing features of the abstract task structure (see Charbonneau 2013; Vorms 2012). I discussed an example of this in the previous chapter in regards to Alač and Hutchins' (2004) cognitive ethnography of how neuroscientists interpret fMRI data. They collaborate and create intermediary graphs and diagrams as scaffolds for interpreting the data into meaningful patterns. Furthermore, this flexibility should be understood not just as redundancy but as degeneracy because they are using structurally different ways of achieving the same functional goal.

This leads on to a second point in regards to degeneracy that was implicit in the above discussion: degeneracy is suggestive of the need for diversity in regards to cognitive practices. I.e. the above example shows that having a degenerate multiplicity of practices related to solving a cognitive task is a boon insofar that it renders a cognitive system's performance more flexible and resilient. The distributed cognitive system is more capable of tackling novel problems and handling

obstacles that might block certain strategies. So, having a range of flexibility enables performance across larger swathes of "epistemic landscapes" (Weisberg & Muldoon 2009). This then leads to more resilient behaviour in tackling cognitive tasks. Arguably this is significant to a range of other cognitive domains for which there is not room to outline in detail here. But one could speculate about the connections here to divisions of labour in scientific communities and what Kuhn called the "essential tension" (1977; also see D'Agostino 2008). I return to this in the general conclusion of this project. However, for now I turn to the last temporal domain involved in the navigation case study.

6.4. Cognitive-historical factors

In the previous section I showed that several key cognitive properties of the navigation team, as it is collaboratively performed synchronically by the team, are dependent on developmental trajectories operating at both the level of the individual (as the acquisition of skills) and at the group level (through shared histories that enhance collaboration), and in the interaction of these two levels (in error detection, and enabling resilient and flexible behaviour in a noisy environment). This showed that how cognition is distributed in time is important for understanding how cognition is distributed in general. Additionally, this analysis provided practical insights into how the navigation team operates, and also provided us with a more in-depth understanding of why this is a case of distributed cognition.

In this section I continue to show how cognition distributed in time is important for understanding human cognitive activity by focusing on what, following Nersessian (2005), I refer to as cognitive-historical features. In particular, I aim to make this claim more precise by showing how it refers to the sets of epistemic resources, and the normative patterned practices that govern their usage, which comprise a cultural-cognitive niche in which agents are inculcated into various ways of

approaching cognitive tasks. With reference to the navigation case study this pertains to what I referred to in the previous chapter as the 'WEIRD navigational niche' (section 5.2).

Perhaps the most definitive aspect of the WEIRD navigation niche is the use of discreet and digital public representational systems for measuring time and space. This has huge implications for how agents – inculcated into this cultural-cognitive niche – approach and compute spatial navigation tasks. An extensive range of empirical evidence suggests that learning a discreet numerical system has transformative effects for the neurocognitive profiles of agents: initiating both neuro-plastic changes to the structure and connections of cortical regions, especially the Intra-parietal sulcus; as well as altering functional behaviour towards tasks involving mathematical reasoning (Ansari 2008; Dehaene 1997, 2007; Dehaene & Cohen 2007; Menary 2010b, 2013b, 2015; Menary & Gillett 2017; Menary & Kirchhoff 2014; Nieder & Dehaene 2009).

One particularly pertinent case for our purposes is a neuroimaging study by Tang and colleagues (2006) who compared the performance of Chinese and English speakers of equal intelligence in a set of mathematical tasks. They found that not only did the Chinese speakers outperform the English speakers but also that their behaviour was correlated with differing neural regions (motor regions in the former and perisylvian language regions in the latter). In a commentary, Cantlon & Brannon (2007) identify numerous aspects of the differing cultural-cognitive niches that could be responsible for these qualitative differences in the neurocognitive profiles. Firstly [1], abacus-users vs. pen-and-paper: with the former placing a larger emphasis on motor patterns. These physical movements themselves can act as accompanying stable representational gestures (Hutchins 2005). Expert users can even operate without the device (simply moving their fingers in the air through the sets of appropriate motor patterns) or completely internalise the

process in their imagination – this developmental trajectory is also associated with neurological changes (see Hu et al 2011; Tanaka et al 2012).

Secondly [2], another cultural source of these differences could be writing systems (logographic vs alphabetic – the former has a much higher visual complexity); and this is related to [3] differences in working memory loads due to the relative complexity of verbal referents for the base-ten counting system. To count to one-hundred in Chinese only requires eleven words as opposed to twenty-eight in English (Butterworth 1999)[76]. Another explanation for these results [4], could be differences in preferred cognitive strategies: Lave comments that these are often overlooked or poorly understood (1988, pp. 101-102; also see Landy et al 2014, pp. 7-8). Given that there are many idiosyncratic strategies to approaching the task of numerical cognition, and although these vary widely intra-group, Cantlon and Brannon hypothesised that these variations could be smaller than the inter-groups differences. The claim here being that differing cultural-cognitive niches might have shared sets of strategies that do not substantially overlap with another niche. Lastly [5], differences in the social milieu of education systems – including an increased social value given to mathematics in China as opposed to English speaking countries – could also contribute to explain these neurocognitive differences. The important point to take from this discussion is that these cultural practices altered the task space and had a corresponding alteration on the neurocognitive profile of the agent.

[76] In Chinese 1-10 each have a specific word; but beyond this larger numbers are just made from straightforward amalgamations of these simpler terms: for example, 11 is “ten-one”, and 20 is “two-ten” and 21 is “two-ten-one”, etc. In contrast, in English there is a different word for every number between 1 and 20 before moving to a special name for each decade (e.g. “twenty”) and the words for 1-9 following it (e.g. “twenty-one”).

Evidence of the extra cognitive loading of the latter system is evident in behavioural experiments where children from each language group are given money to purchase products. Chinese speakers are much more likely to use notes with values over ten whereas English speakers were far more likely to use notes below ten even when using a larger valued note was a more expedient strategy. As Butterworth notes, the teens represent no real difficulties for Chinese speakers because of the superior conceptual tools provided by their cultural-cognitive niche (1999, pp. 129-134).

This evidence suggests that differing notational systems for spatial navigation will have significant outcomes for both the shaping of the task space and the neurocognitive profiles of the agents. But in addition to these ontogenetic factors, there is a crucial cognitive-historical feature of notational system. Hutchins emphasises the base-sixty structure of the various units of measurement used by the team to complete the position fix (as opposed to the more common base-ten structure of Hindu-Arabic numerals)[77]: time is measured in terms of minutes composed of sixty seconds, and hours composed of sixty minutes; direction is measured in terms of 360° of the compass and degrees of angular measurement – latitude or longitude (each 180°); and distance is measured using the nautical mile. The last of these discrete units is the most complicated because the length of a nautical mile is based on the system of angular measurement so that it is one minute of one arc on the surface of the Earth (i.e. there are 360 x 60 nautical miles around the circumference of the Earth, which is 21,600). As such, this unit of measurement has changed in absolute length over history dependent on how measurements of the size of the Earth have been refined (Hutchins 1995a, p. 60). Base-sixty is thus the structure within which the fix cycle is performed and this has significant impacts on how the task itself is performed.

Before we consider these transformative effects, it is important to note that this public symbolic system has an extremely long history, going back to the ancient Sumerians in the third millennium BCE before being passed to the ancient Babylonians (Ifrah 1998). Whilst Hutchins (1995a, p. 166) accepts that it might seem strange at first to connect the contemporary performance of a cognitive task back to activity that occurred several thousand years ago; when we compared this to the analogue or approximate judgements of space and time in the Micronesian navigational niche (section 5.2) it became apparent how arbitrary the conditions of the WEIRD navigational niche are.

[77] An important point of note: this does not involve sixty different symbols but instead uses Hindu-Arabic numerals in sets arranged around sixty rather than one hundred. Learning sixty different notations and words is far more cognitively taxing (Ifrah 1998, p. 40; Zhang & Norman 1995).

Yet these contingent cognitive historical factors have a significant impact on how the fix cycle is performed. As such, to ignore these details is to fail to properly understand the cognitive activity of this system. This is important to note because it highlights a crucial aspect of collective human cognitive effort: the construction of these measurement systems and how they are intertwined *could not feasibly have been created by a solitary individual in a single life-time* because it involves pay-offs that transcend a single lifetime (Shea 2009, p. 2435). Even a group within a single generation would be hard pressed to devise all the facets of these sophisticated measurement systems. They are a truly intergenerational cognitive product.

Multiple theorists have pointed to this capacity – of collaborative effort across generations to complete tasks that would otherwise be beyond an individual or a single generation – as a hallmark of our species (Boyd et al 2011; Henrich 2016; Hutchins 1995a; Johansen 2010; Sterelny 2003; Tomasello 1999, 2009). Following Tomasello I shall refer to this as "virtual collaboration" (1999, p. 41). A good example of this in the navigation team is the Mercator projection: a set of discreet and digital external representations for thinking about spatial coordinates anywhere on the earth. The discreet and digital coordinate systems we use to devise maps and charts partition space into rectangles. But the earth is a sphere and the task of reconciling these two geometrical planes creates a range of complex problems (Synder 1987). Solutions to these have to meet a range of "conflicting constraints" (Kauffman 1995), such as: [i] being faithful to the true shapes of physical features; [ii] correct angular relationships amongst positions; [iii] equal area or correct relative proportions; and [iv] constant scale values for measuring distances (Hutchins 1995, p. 64). There is no optimal solution and one can only balance the set of goals against various different external representational formats (see Charbonneau 2013; Vorms 2012; and see Synder 1987 for an overview of alternative versions of the Mercator projection and other projection frameworks; and see figure 6.5 below).

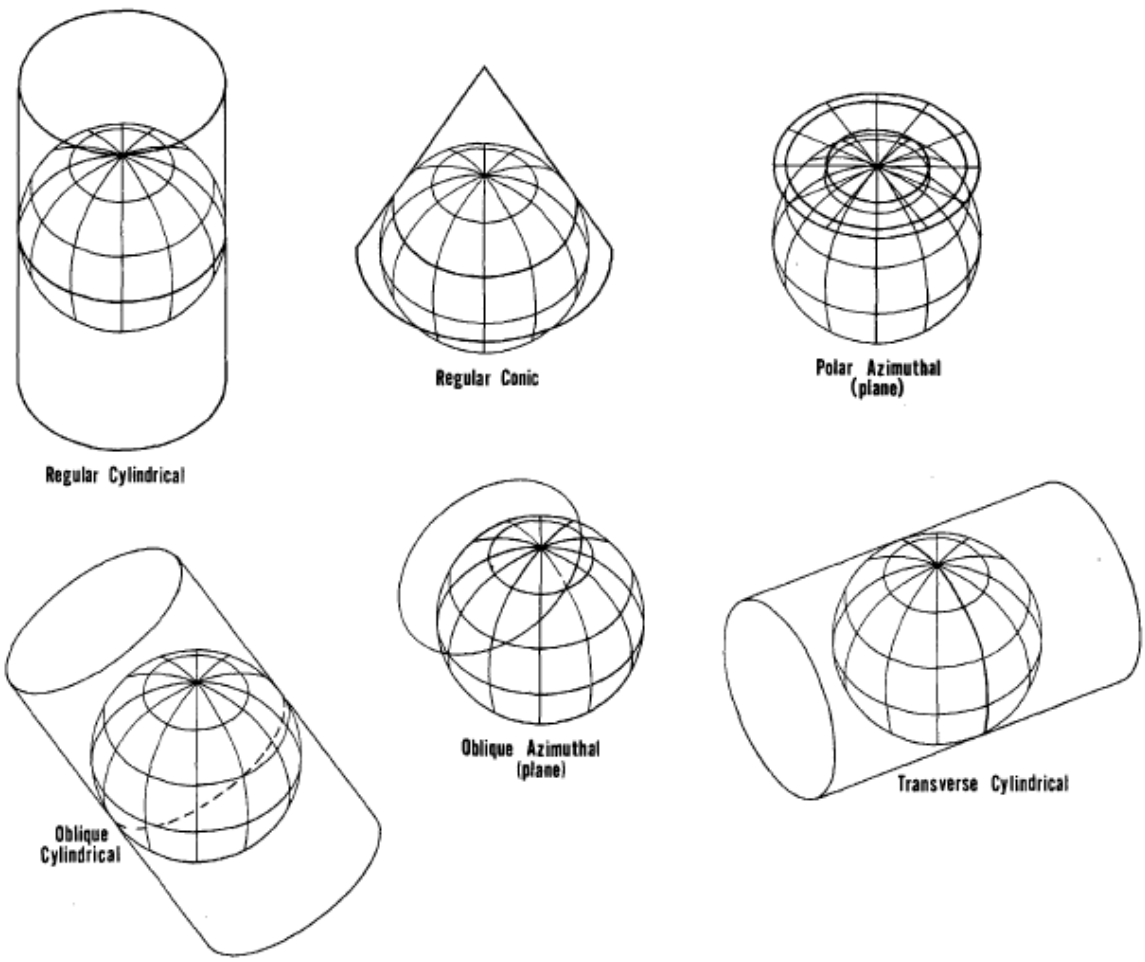

Figure 6.5: projecting the approximately spherical surface of the Earth onto a plane is a complex mathematical procedure that can be done in a number of different ways. Each format offers different affordances and constraints. [Diagram adapted from Synder (1987, p. 6)]

The Mercator projection is one of the most commonly used for most WEIRD spatial navigation. One of the main reasons for this is that this projection preserves the 90° angles and straight lines of latitude and longitude. This entails that navigators do not need to keep recalibrating bearings on long journeys (however, this saving of cognitive load is offset by the fact that the mapping projection becomes increasingly inaccurate as one moves away from the equator).

The Mercator projection was devised from previous work in the cultural-cognitive niche and it has itself been refined multiple times over since its invention in 1569 (Hutchins 1995a, pp. 63-64). Although the abstract geometrical problems were solved by a few individuals and could in principle have been solved by one individual; it has recently been argued that the construction of this mapping projection involved extensive usage of tables of rhumb lines (Leitão & Gaspar 2014). As such, it involved the necessary division of collaborative cognitive labour to create this landmark achievement. Furthermore, the task of putting this system into practice and blending the conceptual schema of the Mercator projection with physical spaces to create a systematic mapping of the surface of the Earth was a task that involved hundreds of thousands of surveyors (see Stewart 2012 on the technique of plotting triangles to map out space reliably). Indeed, the first detailed map of the entire Earth was not finished until 1972 after more than four centuries work (Synder 1987, p. 41). Nersessian (2005, pp. 20-21) argues that a cognitive-historical analysis reveals the important factors of the cultural-cognitive niche that are necessary for properly explain how cultural achievements are achieved: "These sociocultural factors, taken together with cognitive factors, help to explain the nature of the theoretical, experimental, and mathematical knowledge and the methodological practices" (also see Fleck 1979 for similar arguments).

A specific way that the Sumerian-Babylonian base-sixty system influences the performance of the fix cycle is the three-minute timing of the fix cycle itself; since this makes the calculations of speed and distance far simpler (Hutchins 1995a, pp. 151-152; 2010b, pp. 427, 429). This is referred to as the "three-minute rule" and it exploits a "serendipitous interaction"[78] between the two measurement systems: distance – a nautical mile is 2000 yards; and time – one hour; and in which 1/20 of an hour is three minutes and 1/20 of a nautical mile is 200 yards. Thus, the distance

[78] These serendipitous interactions are fairly common and are often overlooked. See Blandford and Furniss (2005), Hutchins (1995b), and Norman (1993) for other examples.

travelled in yards over three minutes is equivalent to speed in nautical miles per hour. As Hutchins puts it: “A complex computation is realized by a simple strategy of situated seeing in a carefully constructed environment” (1995a, p. 152). The deployment of this acquired normative patterned practice shapes and partially constitutes the cognitive activity of the system.

This is another example of the general and widespread phenomenon in which the task space of the cognitive task is altered through the usage of a cognitive artefact or technique (Clark 2008; Humphreys 2004; Kirsh 1995; Menary 2007a; Norman 1991; Zhang 1997b). The crucial point is that without taking into account these diachronic features we would not be able to properly understand this cognitive behaviour nor how this task was achieved. Indeed, Hutchins asserts that it is simply false to claim that the cognitive “achievement” resides solely in the head of the individual in this case; instead it is a cultural achievement involving the agent, environmental resources, and the learnt normative patterned practices which structure and assemble this system (2008; 2011). We can arguably extend this into a general point about a vast majority of human cognition (also see Boyd et al 2011, Henrich 2016; Perry 2003).

The key point here is that without taking the cognitive-historical factors into account we would be unable to properly understand the synchronic activity of the system: [i] not only how the particular coordination dynamics of the system towards the fix cycle are structured and why certain contingent coordination dynamics and representational structures are enacted; but also [ii] how the acquisition of these cultural practices and epistemic resources (that are transmitted across generations) come to partially constitute the cognitive activity of the system by shaping and being shaped by the synchronically deployed coordination dynamics. As such, cognitive-historical factors are vital for investigating distributed cognition.

6.5. Summary

In this final chapter I have showed how the three main principles of my pluralistic framework operate together to draw out significant insights into the most famous case study of distributed cognition. This discussion has not only demonstrated that my position goes further than other simpler approaches to distributed cognition; but also provided several concrete examples of how cognition is distributed that should be of interest to methodological individualists as well.

In particular, I showed how cognition distributed in time was vital for understanding the cognitive properties of the individuals and the group, and the differences between them, in the fix cycle. On the synchronic level this involved distinguishing differing sets of normative patterned practices and cognitive properties between groups and individuals involved in the task-space: viz. buffering and communication networks; and parallel activity. I then went on to show how these synchronic coordination dynamics were constrained, enabled, shaped, and partially constituted by normative patterned practices operating over two diachronic timescales. Firstly, developmental trajectories operating on multiple levels which I disentangled with a task-centric approach. These are the processes by which individuals acquire and master the appropriate normative patterned practices towards the successful completion of cognitive tasks in a resilient manner in a noisy environment; and as the shared histories by which individuals learn to collaborate and refine their shared knowledge. The interaction of developmental trajectories at the group level and individual level explains the resilient behaviour of the team in the face of a noisy and error strewn environment. In particular, this involved the degeneracy of cognitive practices to orchestrate adaptive coordination dynamics.

The second diachronic timescale were cognitive-historical factors: the processes by which virtually collaborative cognitive activity is accumulated and refined over generations in a particular

cultural-cognitive niche. This activity is passed from one generation to the next through the transmission of epistemic resources – tools and techniques – and the normative patterned practices that govern their creation, manipulation, and maintenance. This accumulative epistemic engineering heavily constrains, enables, and reshapes the task space in which synchronic cognitive activity takes place.

Lastly, in response to the dangers of methodological bloat, this chapter has acted as a performative demonstration of how normative patterned practices make a properly thorough investigation of cognition distributed across a wide range of spatial and temporal scales tractable.

Ω: General Conclusion

Summary

In this thesis I have explored the multiple problems of distributed cognition and shown how my pluralistic framework not only resolves these, but also leads to new insights into the particular case studies under investigation as well as the concept itself. I will first recap each of these problems before then summarising how the three main principles of my pluralistic framework – task-specificity, the mutual manipulability criterion, and normative patterned practices – not only handle these issues but also provide us with important insights into real world case studies. After this summary I then consider some further issues and future directions for this project.

In chapter one I catalogued the incredibly miscellaneous range of purported cases of distributed cognition that are in the literature of multiple fields in the humanities and sciences. Although this plethora demonstrates that distributed cognition is a viable and useful research framework, it also raises multiple philosophical questions. Firstly, the sheer diversity of sizes and types amongst these examples – from neural systems to niches – raises concerns about how we can compare them and whether distributed cognition is being used in a meaningfully comparable manner across this range. I demonstrated that one cannot straightforwardly make comparisons of differently sized systems without doing a disservice to the actual phenomena – thus necessitating a principled means of being able to analyse multiple levels of configuration across both space and time *on their own terms*.

Secondly, this motley assortment exacerbates the problem of cognitive bloat – viz. the question of how to appropriately demarcate the boundaries of cognitive systems and distinguish

between genuine components and mere causal background phenomena. Although I did not reject information flow entirely (because to do so would place my position in opposition to mainstream approaches in cognitive science; and I have argued, following Sutton (2015) and others, that distributed cognition should be viewed as consilient in some respects with mainstream approaches rather than as a paradigmatic combatant), I did show that it is insufficient for tackling the problem of cognitive bloat for several reasons. Most problematically, it is either too permissive because information bottlenecks cannot distinguish between genuine components and mere causal background effects (Kaplan 2012); or it inadvertently supports methodological individualism because information flow is drastically higher within the cortex than with external processes (Eliasmith 2009). As such, a further principle was required here to tackle this problem.

A simpler parochial approach was offered to resolve these tensions: to propose a set of criteria that would limit distributed cognition to a smaller subset of these cases and reject the others. However, I showed that this approach was undesirable for three reasons: firstly, for practical reasons I think that it is very unwise to postulate a set of overly narrow conceptual constraints that will ultimately be ignored by working scientists and other theorists. Much of this empirical work focuses on ecologically salient cognitive systems which are often overlooked by other approaches – and dismissing large swathes of this important research is unwise. Secondly, simpler approaches have tended to focus on *only* one of the two key features of distributed cognition – the notion of supra-individual cognitive properties distributed beyond the skull of the isolated individual in space. Following Hutchins (2001), I have argued that this ontological claim needs to be based on the naturalistic methodological claim that the unit of analysis for exploring human cognition should not be decided a priori but should be flexible. How to designate both these claims in a principled fashion introduces a new problem – particularly in relation to time.

This is the third major issue for simpler approaches (e.g. the information bandwidth criteria), which have tended to focus solely on the synchronic temporal level and have overlooked the importance of cognition distributed in time over diachronic scales. This is incredibly problematic because I showed that if we want to properly understand cognitive behaviour and the coordination dynamics of distributed cognitive systems as they occur synchronically, we often need to understand how these are shaped and constrained by both developmental trajectories and cognitive historical factors. These impact not only the capacities of agents and what available cognitive resources there are, but also how the task-space within which they operate has been shaped by the activities of previous generations in an accumulative fashion.

For instance, in Hutchins' (1995a) navigation team case study – which I have used as the lynchpin example for my thesis – the plotter uses the three-minute rule to calculate where the ship will approximately be for the next fix cycle (ceteris paribus). He does so because this length of time makes the calculations very easy. In turn this is because the WEIRD navigational cultural-cognitive niche is premised on a digital numerical measurement system that is base-sixty for contingent historical reasons that goes back to the ancient Sumerians and Babylonians. The plotter learns this cognitive practice in a developmental niche and this then dictates and partially constitutes how he thinks and behaves towards the task-space. Additionally, the task-space has itself been shaped and altered by the "accumulative downstream epistemic engineering" (Sterelny 2003) of numerous generations that have measured and recorded details of the geography of the Earth and recorded these in charts – what Hutchins refers to "precomputational" artefacts (1995a). This amounts to an intergenerational "virtual collaboration" (Tomasello 1999) that allows our species to tackle cognitive tasks that would otherwise be infeasible or even incomprehensible because it distributes the workload across time, and allows collective effort to engage with tasks that involve gigantic temporal and spatial scales (Boyd et al 2011; Henrich 2016; Shea 2009). As Hutchins notes, these

techniques and tools change the nature of the task-space (1995a, p. 21); and as I have argued, following Menary (2007a, 2013b, 2014, 2015), this also impacts on the neurocognitive profile of the agent (because different tasks are correlated with different cortical responses and behavioural repertoires).

To tackle the problems of unsystematised heterogeneity, cognitive bloat, and cognition distributed in time I formulated a pluralistic framework. However, this more complex approach introduces a new problem with regards to the practicalities of analysing such a complex state of affairs for real world case studies. I labelled this the problem of "methodological bloat". The three central principles of my novel pluralistic framework were drawn from naturalistically motivated work in the literature. I modified these previous criteria for three main reasons: [i] to make them operate in tandem with one another more fluently; [ii] to make clear their theoretical consequences for understanding what distributed cognition is; and [iii] to make practical insights into the particular case studies that they are applied to (thus meeting the pragmatic desideratum I specified in the general introduction).

In summary, the principles functioned as follows: To facilitate the multilevel analysis that is entailed by distributed cognition (given the differing size configurations at supra-individual levels), Davies and Michaelian (2016) argued that a non-agent centric approach is necessary. This is because agent-centric approaches cause a range of problems such as: [1] falsely over-attributing cognitive properties to the agents that are better understood as properties of the distributed system of which the agent is just a component (Hutchins 1995a, 2008; Norman 1991; Menary & Gillett 2017); and [2] misunderstanding sub-individual and supra-individual levels by misattributing agential features to them – thereby committing a mereological fallacy (Drayson 2012) and entailing residual internalism (Menary 2007a) – instead of considering them *on their own terms*.

The alternative is to consider the flexible unit of analysis on a task-specific basis. A cognitive system is defined as one that engages in at least one clearly specified cognitive task (Davies & Michaelian 2016; Theiner et al 2010; Wilson 2001, 2004). A cognitive task is defined in a naturalistic manner according to those explored by the various fields of cognitive sciences (Goldstone & Theiner 2017; Kaplan 2012; Miller 2003). This is refined further by Steiner's (1966, 1972) taxonomy which distinguishes how differing task structures interact with coordination dynamics of the distributed cognitive system and local resources. Lastly, the notion of level above the individual needed to be made more precise and here I turned to Caporael's (2014) notion of core configurations – task-relevant emergent organisations of agents and artefacts at distinct levels of mechanistic composition (Craver 2007; Wimsatt 1986). Task-specificity thus provides us with a principled means for negotiating and traversing the differing sizes and types of distributed cognition in the literature.

However, task-specificity still leaves the problem of cognitive bloat unchecked. Here I turned to Kaplan's (2012) use of the mutual manipulability criterion – taken from mechanistic work in biology and neuroscience (Craver 2007) – which distinguishes between genuine components of a system and mere causal background conditions through two idealised interventions. A bottom-up manipulation of the putative component to see if this engenders a response in the higher-level behaviour of the system. And a top-down manipulation of the system-level behaviour to see if this results in a reciprocal change in the component. Only by satisfying both of these conditions can it be claimed that a component is a genuine feature of the system. For instance, in the navigation team: although the bearing log, pen and paper calculations, and calculators are all used in a position fix; the plotter only reciprocally manipulates both the chart and the hoey when demarcating lines of position. As such, only the latter should be considered as genuine parts of this distributed cognitive system which is nested within the navigation team. The other features are merely background conditions for the activity of this system (even if their omission would cause great hardships for the

cognitive activity). By distinguishing between genuine components and mere causal background features the mutual manipulability criterion nullifies the challenge of cognitive bloat.

But although the mutual manipulability criterion works well over synchronic timescales it breaks down beyond this range – as acknowledged by Kaplan (2012, p. 559) who notes that it was only ever intended to operate over the short-time periods considered in cognitive psychology and neuroscience research (milliseconds to minutes). Given that cognition is distributed over several distinct temporal periods this entailed that the mutual manipulability criterion required supplementing. Here I firstly clarified and refined Hutchins' (1995a, 2001, 2006) original presentation of cognition distributed in time by building from Nersessian's (2005) more concise formulation to identify three main temporal domains: one synchronic and two diachronic – developmental trajectories and cognitive historical-features (also see Bietti & Sutton 2015). I then drew on Menary's (2007a, 2016) work on normative patterned practices: sets of cultural practices that govern the coordination dynamics at play in the distributed cognitive system. These practices are acquired by agents through developmental trajectories that alter the neurocognitive profile of the agent (Menary 2014). Through this enculturation process the cultural practices come to be partially constitutive of the agent's cognitive functioning insofar that they both shape cognitive behaviour and are that which is shaped by it (Menary 2013b). This demonstrates how normative patterned practices bind the multiple levels of analysis together.

Practices are also the medium by which cognitive-historical factors are mediated – in relation to both techniques and tools that are transmitted and refined across generations. They are the norms which guide the creation, maintenance, and manipulation of these cognitive resources (Menary 2012); thus, enabling the cranking of Tomasello's ratchet in the accumulative cultural-cognitive niche (Sterelny 2003; Tomasello 1999). As such, normative patterned practices bind

together not only the multiple temporal levels but also micro and macro spatial scales (Roepstorff et al 2010): because they have an impact on neural structures and functions (see Downey & Lende 2012b for an overview); and because sets of normative patterned practices are the medium by which we identify distinct cultural-cognitive niches. For instance, the three-minute rule and the discreet digital number system are sets of practices that are part of the WEIRD navigational cultural-cognitive niche. In addition to functionally different behaviour (e.g. differing behavioural repertoires), evidence in comparative neuroscientific studies implies that differing mathematical practices (in regards to techniques and tools) have differing neural correlates (Cantlon & Brannon 2007; Dehaene 2007; Hu et al 2011; Tanaka et al 2012; Tang et al 2006).

By operating across all temporal levels and being the medium by which spatial levels of analysis are mediated, normative patterned practices also provide us with a response to the challenge of methodological bloat. Viz. the proponent of distributed cognition knows that they must focus on the normative patterned practices associated with the task in order to map out the coordination dynamics of the system and its embedded context. From here they can then move to analysing these interrelations and discerning what the genuine composition of the system is using the mutual manipulability criterion. Additionally, this tripartite focus also places a large emphasis on paying attention to the particular details of case studies. Huebner (2014) has argued for the importance of ethnographic details for investigating distributed cognition. I demonstrated this in a concrete fashion through an in-depth re-examination of Hutchins' navigation case study in which I showed how my pluralistic framework draws out new insights both in the target phenomena but also enabled a deeper understanding of what the concept of distributed cognition is itself. Thus, meeting Huebner's desideratum through a performative demonstration.

Further issues and future directions

In this thesis I have demonstrated how my pluralistic framework draws out important insights from the most famous case of distributed cognition – Hutchins' (1995a) navigation team. But a more extensive treatment of case studies from across the diverse range that I catalogued in chapter one is needed in order to properly demonstrate the flexibility and viability of my approach.

Small steps were taken in these regards by showing how my pluralistic framework operates in relation to a several core configurations of differing sizes. In chapters two and three I discussed two smaller dyadic configurations: pre-flight checklists (Norman 1991) and the blind man's stick (Bateson 1972; Hutchins 1995a, 2010a; Malafouris 2008, 2010, 2013). In chapter four I returned to the latter of these and also considered a much larger system, the Hubble telescope (Giere 2006, 2007, 2012). Although these discussions were brief, they showed that – just as with the more prolonged discussion of the navigation team – my pluralistic framework reveals both philosophical and practical aspects of these case studies. For instance, these analyses uncovered features that were either not present or obvious in the original: e.g. in the case of the pre-flight checklist I supplemented the systems level and personal level descriptions with a corresponding account at the sub-personal level (chapter 3). In the other cases I cleared up confusions in the original: e.g. in the blind man's stick and the Hubble telescope examples I used my pluralistic framework to identify genuine components of the system and individuate cognitive properties at the appropriate levels (see chapter four). So, despite the brevity of the engagements with these case studies, they still show that my pluralistic framework does work for cases of varying sizes. But to properly demonstrate this I will need to examine other case studies in future work. Given that my pluralistic framework provided additional insights in these cases, this is suggestive that such a project would not merely be a rehashing of the field but would also be informative in and of itself; as well as placing accounts of distributed cognition under a systematic rubric.

One particular direction of development is in regards to the philosophy of science literature and discussions about the division of cognitive labour. The case studies explored in this literature range from simple dyadic cases such as the use of physical and theoretical models by a single agent (Charbonneau 2013; Giere 2006; Knuuttila 2011), to team-work in laboratories and fieldwork (Giere & Moffatt 2003; Nersessian 2005, 2006, 2009), up to much larger core configurations involving several hundred agents and artefacts, for example the Hubble telescope (Giere 2006, 2007, 2012). Questions about the legitimacy of these case studies, as well as their composition, and other conceptual matters which are more the concern of philosophy of cognition; also intersect with topics which are more directly the concern of philosophy and sociology of science. These have not been within the purview of this project. But it is interesting to note that the contention of Giere and other proponents is that distributed cognition is a positive boon to this field – particularly in regards to the debates between philosophy and sociology of science, by bridging them in a conciliatory manner (Giere 2002a, pp. 638, 641; 2002b, pp. 285, 295-296; 2006, p. 114; 2007, p. 319; Giere & Moffatt 2003, pp. 301-302; also see Brown 2011, p. 17; Cheon 2014, pp. 23-24; Magnus 2007, p. 297; Magnus & McClamrock 2015, p. 1114; Nersessian 2005, pp. 17-19; 2006, p. 702; Vaesen 2011, p. 381; cf. Toon 2014b). As such, in contrast to the other case studies outlined in chapter one, exploring these cases in the future will be of the greatest use because of this nexus of problems; viz.: [1] reassessing, reorganising, and deepening our understanding of these putative cases with my pluralistic framework; [2] the general exploration of human cognition in the wild; and [3] exploring what possible insights my pluralistic framework might have for questions in philosophy of science.

On this last point I offer some speculations to conclude this thesis, which are based on the argumentative apparatus in the preceding chapters. In chapter four I used the mutual manipulability criterion to disentangle some of the confusions surrounding what is sometimes referred to as

"distributed epistemology" in Giere's Hubble telescope example (2006, 2007, 2012)[79]. But there is still much to discuss and contemplate in cases such as these. The interaction of agents and artefacts in the scientific context is an incredibly important topic given that many fields are becoming increasingly reliant on technology and computers (Humphreys 2004, 2009) – with some theorists labelling this context a new paradigm in the progress of science (Hey et al 2009). Interrelated with the increasing use of technology is the increasing specialisation, interdisciplinarity, and collaborative nature of scientific enquiry (Andersen 2016; Stanford 2015). Both these topics raise questions about distributed epistemology: Who or what is the appropriate epistemic agent? How should we understand certain forms of scientific reasoning in which large amounts of the cognitive workload is completed by machinery? What is the impact of this state of affairs for how scientific work is conducted and for future scientific progress?

For instance, Humphreys (2004, p. 7) details the overwhelming statistics involved in mapping the human genome: 500 million trillion base-to-base comparisons (500,000,000,000,000,000,000). This took 20,000 CPU hours on what was then the world's second largest networked supercomputer. This is equivalent to eight hundred and thirty-three days and eight hours, or over two years, of constant calculation. This involved eighty terabytes of information (80,000,000,000,000), which is equivalent to a stack of manuscripts over seven hundred and seventy kilometres high[80]. The point of all these large numbers – what he calls the "quantity of data issue" – is that they make "in principle" stances "completely infeasible […] in practice" (Humphreys 2004, 2009). Viz. it is no longer accurate to consider something knowable if a human *could,* given an endless amount of time, retrace every step of the knowledge collection process.

[79] The term "extended" is sometimes used (e.g. Carter et al 2014; Kerr & Gelfert 2014; Pritchard 2016). I have preferred to stick with the "distributed" nomenclature for reasons outlined earlier in the thesis: viz. that the term extended leads to a residual internalism and creates the erroneous impression that cognition begins in the head before leaking outwards (Menary 2007a; Sutton 2015).

[80] For reference, the international space station orbits only 330-410km above the surface of the Earth.

There are several ways one could respond to this. Firstly, one could give up on the notion of an individual knower in such a context and subdivide and distribute the checking procedure over a large group of human agents (e.g. Knorr Cetina 1999). Their collective effort could then be counted as some form of "radical collaboration" (Kukla 2012). I have advised against eliminating the individual on the basis that multiple levels of analysis and the differences between them are a vital aspect of distributed cognition (Hutchins 1995a; Menary & Gillett 2017). Either way, this still leaves the question of how to understand their collective efforts in regards to epistemic matters (Huebner 2014; Kerr & Gelfert 2014; Kukla 2012). Another way, building from the facts about enculturated cognition outlined in chapter five, would be to recognise that these cases are *not as novel* as they might at first seem. Viz. it is the case that human epistemic achievements have always been predicated on large connections of what Tomasello refers to as "virtual collaborators" (1999). Henrich (2016) emphasises that most human cognitive achievements are cultural achievements and arise as much through deliberate creativity as much as through fortuitous copying errors and mistakes (also see Boyd et al 2011; Hutchins 2008). I.e. once one takes the diachronic scale into account we can see that the epistemic credit of human cognition is often radically distributed. In one sense this complicates the matter in a perpendicular manner – by adding a temporal dimension of distribution to the spatial problem of epistemology. But given how I have shown that the diachronic temporal factors create the conditions of possibility for synchronically distributed cognition, I think is suggestive that this might in fact be a boon rather than a burden. Exploring these states of affairs requires a much greater exposition. But in this brief discussion I have shown that there is much scope for investigating these topics with my pluralistic framework in future works.

Thank you for your time.

Bibliography

Abrahamsen, A. & Bechtel, W. (2015) Diagrams as tools for scientific reasoning. *Review of Philosophy and Psychology 6*, 117-131.

Adams, A. & Aizawa (2001) The bounds of cognition. *Philosophical Psychology, 14*, 43-64.

Adams, A. & Aizawa (2010) Defending the Bounds of Cognition. (pp. 67-80) in R. Menary (Ed.) The Extended Mind. Cambridge, MA: MIT Press.

Adriaans, P. (2012) Information. In E. N. Zalta (Ed.) The Stanford Encyclopedia of Philosophy (Winter 2016 Edition). Available at: https://plato.stanford.edu/entries/information/ accessed on 02/12/2016.

Alač, M. & Hutchins, E. (2004) I See What You Are Saying: Action as Cognition in fMRI Brain Mapping Practice. *Journal of Cognition and Culture, 4*(3), 629-661.

Allen, C. (2017) On(not) defining cognition. *Synthese.* [published online], 1-17.

Andersen, H. (2016) Collaboration, interdisciplinarity, and the epistemology of contemporary science. *Studies in History and Philosophy of Science 56*, 1-10.

Anderson, M. L. (2010) Neural Reuse: A fundamental organizational principle of the brain. *Behavioral and Brain Sciences, 33*, 245-264.

Anderson, M. L., Richardson, M. J. & Chemero, A. (2012) Eroding the Boundaries of Cognition: Implications of Embodiment. *Topics in Cognitive Science 4 (4)*, 717-730.

Ansari, D. (2008) Effects of development and enculturation on number representation in the brain. *Nature Reviews Neuroscience 9*, 278-291.

Arnett, J.J. (2008) The Neglected 95%: Why American Psychology Needs to Become Less American. *American Psychologist 63*, 7, 602-614.

Asch, S. (1951) Effects of Group Pressure upon the Modification and Distortion of Judgments. (pp.177-190) in H. Guetzkow (Ed.) *Groups, leadership, and men*. Pittsburgh, PA: Carnegie.

Atkin, A. (2016) *Peirce.* New York: Routledge.

ATLAS Collaboration (2012) Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Physics Letters B,* 716, 1-29.

Aydin, K., Ucar, A., Oguz, K. K., Okur, O. O., Agayev, A., Unal, Z., Yilmaz, S. & Ozturk, C. (2007) Increased Gray Matter Density in the Parietal Cortex of Mathematicians: A Voxel-Based Morphometry Study. *American Journal of Neuroradiology 28*, 1859-1864.

Baber, C. (2010) Distributed cognition at the crime scene. *AI & Society, 25*, 423-432.

Baber, C., Smith, P., Cross, J., Hunter, J. & McMaster, R. (2006) Crime scene investigation as distributed cognition. *Pragmatics & Cognition, 14 (2)*, 357-385.

Ballard, D., Hayhoe, M. & Pelz, J. (1995) Memory representations in natural tasks. *Journal of Cognitive Neuroscience, 7*(1), 66–80.

Barnier, A. J., Sutton, J., Harris, C. B. & Wilson, R. A. (2008) A conceptual and empirical framework for the social distribution of cognition: The case of memory. *Cognitive Systems Research 9*, 33–51.

Bateson, G. (1972) *Steps to an Ecology of Mind*. New York: Ballantine Books.

Beach, K. (1993) Becoming a Bartender: The Role of External Memory Cues in a Work-directed Educational Activity. *Applied Cognitive Psychology 7*, 191-204.

Bechtel, W. (2008) *Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience*. London: Routledge.

Bechtel, W. (2009) Looking down, around, and up: Mechanistic explanation in psychology. *Philosophical Psychology*, *22* (5), 543-564.

Becvar, A., Hollan, J. & Hutchins, E. (2005) Hands as molecules: Representational gestures used for developing theory in a scientific laboratory. *Semiotica 156–1/4*, 89-112.

Becvar, A., Hollan, J. & Hutchins, E. (2008) Representational Gestures as Cognitive Artifacts for Developing Theories in a Scientific Laboratory. (pp. 117-144) in M. S. Ackerford, C. A. Halverson, T. Erickson & W. A. Kellogg (Eds.) *Resources, Co-evolution, and Artifacts: Theory in CSCW*. London: Springer-Verlag.

Bennett, K. (2011) Construction area (no hard hat required). *Philosophical Studies 154*, 79-104.

Bietti, L. M. & Sutton, J. (2015) Interacting to remember at multiple timescales: Coordination, collaboration, cooperation and culture in joint remembering. *Interaction Studies 16 (3)*, 419-450.

Blandford, A. & Furniss, D. (2005) DiCoT: A Methodology for Applying Distributed Cognition to the Design of Teamworking Systems. (pp. 26-38) in S. W. Gilroy & M. D. Harrison (Eds.) *Interactive Systems – Design, Specification, and Verification.* Springer.

Boas, F. (1889) On Alternating Sounds. *American Anthropologist 2*, 47-53.

Boas, F. (1910) Psychological Problems in Anthropology. *The American Journal of Psychology 21*(3), 371-384.

Bohbot, V. D., Lerch, J., Thorndycraft, B., Iaria, G. & Zijdenbos, A. P. (2007) Gray Matter Differences Correlate with Spontaneous Strategies in a Human Virtual Navigation Task. *The Journal of Neuroscience 27 (38)*, 10078-10083.

Bohbot, V. D. McKenzie, S., Konishi, K., Fouquet, C., Kurdi, V., Schachar, R., Boivin, M. & Robaey, P. (2012) Virtual navigation strategies from childhood to senescence: evidence for changes across the lifespan. *Frontiers in Ageing Neuroscience 4,* 28.

Boroditsky, L. & Gaby, A. (2010) Remembrances of Times East: Absolute Spatial Representations of Time in an Australian Aboriginal Community. *Psychological Science 21*(11), 1635-1639.

Boyd, R., Richerson, P. J. & Henrich, J. (2011) The cultural niche: Why social learning is essential for human adaptation. *PNAS 108,* 10918–10925.

Brown, M. J. (2011) Science as socially distributed cognition: Bridging philosophy and sociology of science. (pp. 17-30) in K. François, B. Löwe, T. Müller, & B. van Kerkhove (Eds.), *Foundations of the formal sciences VII*. London: College Publications.

Bruner, J., R., Olver, R. R. & Greenfield, P. (1966) *Studies in cognitive growth: A collaboration at the center for cognitive studies*. New York: John Wiley and Sons.

Burck, J. M., Bigelow, J. D. & Harshbarger, S. D. (2011) Revolutionizing Prosthetics: Systems Engineering Challenges and Opportunities. *Johns Hopkins APL Technical Digest, 30* (3), 186–197.

Button, G. (2008) Against 'Distributed Cognition'. *Theory, Culture & Society 25*(2), 87-104.

Butterfill, A. & Sebanz, N. (2011) Editorial: Joint Action: What Is Shared? *Review of Philosophy and Psychology 2*, 137-146.

Butterworth, B. (1999) *The Mathematical Brain.* London: Macmillan.

Cantlon, J. F., & Brannon, E. M. (2007) Adding up the effects of cultural experience on the brain. *Trends in Cognitive Sciences, 11*(1), 1-4.

Caporael, L. R. (2014) Evolution, Groups, and Scaffolded Minds. (pp. 57-76) in in L. R. Caporael, J. R. Griesemer & W. C. Wimsatt (eds.) *Developing Scaffolds in Evolution, Culture, and Cognition*. Cambridge, MA: The MIT Press.

Carter, J. A., Kallestrup, J., Palermos, S. O. & Pritchard, D. (2014) Varieties of Externalism. *Philosophical Issues 24,* 63-109.

Carruthers, P. (2006) *The Architecture of the Mind Massive Modularity and the Flexibility of Thought.* Oxford: Oxford University Press.

Castro-Caldas, A., Cavaleiro Miranda, P., Carmo, I., Reis, A., Leote, F., Ribeiro, C., & Ducla-Soares, E. (1999) Influence of learning to read and write on the morphology of the corpus callosum. *European Journal of Neurology, 6*, 23-28.

Changizi, M. A. (2003) Relationship between Number of Muscles, Behavioral Repertoire Size, and Encephalization in Mammals. *Journal of Theoretical Biology 220*, 157-168.

Charbonneau, M. (2013) The cognitive life of mechanical molecular models. *Studies in History and Philosophy of Biological and Biomedical Sciences, 44(4),* 585-594.

Cheon, H. (2014) Distributed Cognition in Scientific Contexts. *Journal for General Philosophy of Science 45*, 23-33.

Chiao, J. Y. (2009) Cultural neuroscience: a once and future discipline. *Progress in Brain Research, 178*, 287-304.

Chrisomalis, S. (2004) A Cognitive Typology for Numerical Notation. *Cambridge Archaeological Journal 14*(1), 37-52.

Clark, A. (1997) *Being There: Putting Brain, Body, and World Back Together.* Cambridge, MA: MIT Press.

Clark, A. (2008) *Supersizing the Mind: Embodiment, Action, and Cognitive Extension.* Oxford: Oxford University Press.

Clark, A. (2016) *Surfing Uncertainty: Prediction, Action, and the Embodied Mind.* New York: Oxford University Press.

Clark, A. & Chalmers, D. (1998) The extended mind. *Analysis, 58*(1), 7-19.

Cohen, A. B. (2009) Many Forms of Culture. *American Psychologist 64*(3), 194-204.

Cole, M. & Engström, Y. (1993) A cultural-historical approach to distributed cognition. (pp. 1-46) in Salomon, G. (ed.) *Distributed cognitions: Psychological and educational considerations*. Cambridge: Cambridge University Press.

Cole, M. & Griffin, M. (1980) Cultural Amplifiers Reconsidered. (pp. 343-364) in D. Olson (Ed.) *The social foundations of language and thought – Essays in honor of Jerome Bruner*. New York: Norton.

Collins, B. E. & Guetzkow, H. (1964) *A Social Psychology of Group Processes for Decision-Making*. New York: John Wiley & Sons.

Craver, C. F. (2007) *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience.* Oxford: Clarendon Press.

Craver, C. F. & Tabery, J. (2015) Mechanisms in Science. In E. N. Zalta (Ed.) *The Stanford Encyclopedia of Philosophy* (Spring 2017 Edition). Available at: https://plato.stanford.edu/cgi-bin/encyclopedia/archinfo.cgi?entry=science-mechanisms accessed on 28/07/2017.

Crnkovic, G. D. & Hofkirchner, W. (2011) Floridi's "Open Problems in Philosophy of Information", Ten Years Later. *Information 2*, 327-359.

D'Agostino (2008) Naturalizing the essential tension. *Synthese 162,* 275-308.

D'Andrade, R. (1995) *The development of cognitive anthropology*. Cambridge: Cambridge University Press.

Davies, J. & Michaelian, K. (2016) Identifying and individuating cognitive systems: a task-based distributed cognition alternative to agent-based extended cognition. *Cognitive Processing, 17 (3)*, 307-319.

De Cruz, H., & De Smedt, J. (2013) Mathematical symbols as epistemic actions. *Synthese, 190*, 3-19.

Deacon, T. (1997) *The symbolic species: The co-evolution of language and the human brain*. London: Norton.

Dean, L. G., Kendal, R. L., Schapiro, S. J., Thierry, B., & Laland, K. N. (2012) Identification of the Social and Cognitive Processes Underlying Human Cumulative Culture. *Science 335*, 1114-1118.

Dehaene, S. (1997) *The Number Sense: How the Mind Creates Mathematics*. London: Penguin.

Dehaene, S. (2007) Symbols and quantities in parietal cortex: elements of a mathematical theory of number representation and manipulation. (pp. 527-574) in P. Haggard, Y. Rossetti and M. Kawato (Eds.) *Attention & Performance XXII. Sensori-motor foundations of higher cognition*. Cambridge, MA: Harvard University Press.

Dehaene, S. (2009) *Reading in the Brain.* New York: Viking.

Dehaene, S. & Cohen L. (2007) Cultural Recycling of Cortical Maps. *Neuron, 56*, 384-398.

de Witt, L. Alexander, D. & Ekroll, V. & Wagemans, J. (2016) Is neuroimaging measuring information in the brain? *Psychonomic Bulletin & Review 23 (5)*, 1415-1428.

Dennett, D. C. (1996) *Kinds of Minds: Towards an Understanding of Consciousness.* London: Weidenfeld & Nicolson.

Dennett, D. C. (2017) *From Bacteria to Bach and Back: The Evolution of Minds.* St. Ives: Allen Lane.

Deranty, J.P. (2009) What is work? Key insights from the psychodynamics of work. *Thesis Eleven*, *98*, 69-87.

Doidge, N. (2007) *The brain that changes itself: Stories of personal triumph from the frontiers of brain science*. New York: Penguin.

Donald, M. (1993) *Origins of the Modern Mind – Three Stages in the Evolution of Culture and Cognition.* Cambridge, MA: Harvard University Press.

Donald, M. (2010) The Exographic Revolution: A Neuropsychological Sequelae. (pp. 71-80) in L. Malafouris & C. Renfrew (Eds.) *The cognitive life of things – Recasting the boundaries of the mind*. Oxford: McDonald Institute Monograph.

Douglas, M. (1986) *How Institutions Think.* New York: Syracuse University Press.

Downey, G. (2011) Learning the 'Banana-Tree': Self Modification through Movement. (pp. 77–90) in T. Ingold (Ed.) *Redrawing Anthropology: Materials, Movements, Lines.* Surrey: Ashgate.

Downey, G. (2012) Balancing between Cultures: Equilibrium in Capoeira. (pp. 169–194) in D. H. Lende & G. Downey (Eds.) *The Encultured Brain: An introduction to Neuroanthropology*. Cambridge, MA: MIT Press.

Downey, G. (2014a) Giving names to aromas in Aslian languages. *PLOS Blogs.* Available at: http://blogs.plos.org/neuroanthropology/2014/03/09/giving-names-to-aromas-in-aslian-languages/. Accessed on 22/06/2017.

Downey, G. (2014b) '*Habitus in Extremis*': From Embodied Culture to Bio-Cultural Development. *Body & Society, 20(2)*, 113-117.

Downey, G. (2015) Beneath the Horizon: The Organic Body's Role in Athletic Experience. (pp. 114-136) in K. Ram & C. Houston (Eds.) *Phenomenology in Anthropology: A sense of Perspective*. Bloomington, Indianapolis: Indiana University Press.

Downey, G. (2016) Sensory Enculturation and Neuroanthropology: The Case of Human Echolocation. (pp. 41-55) in Chiao, J. Y., Shu-Chen, L., Seligman, R. & Turner, R. (eds.) *The Oxford Handbook of Cultural Neuroscience.* Oxford: Oxford University Press.

Downey, G. & Lende, D. (2012a) Evolution and the Brain. (pp. 103-137) in D. Lende & G. Downey (Eds.) *The Encultured Brain: An Introduction to Neuroanthropology.* Cambridge, MA: MIT Press.

Downey, G. & Lende, D. (2012b) Neuroanthropology and the Encultured Brain. (pp. 23-65) in D. Lende & G. Downey (Eds.) *The Encultured Brain: An Introduction to Neuroanthropology.* Cambridge, MA: MIT Press.

Draganski, B., Gaser, C., Busch, V., Schuierer, G., Bogdahn, U. & May, A. (2004) Neuroplasticity: Changes in grey matter induced by training. *Nature 427*, 311-312.

Draganski, B., Gaser, C., Kemperman, G., Kuhn, H. G., Winkler, J., Büchel, C. & May, A. (2006) Temporal and Spatial Dynamics of Brain Structure Changes during Extensive Learning. *The Journal of Neuroscience 26*(23), 6314-6317.

Draganski, B. & May, A. (2008) Training-induced structural changes in the adult human brain. *Behavioural Brain Research 192,* 137-142.

Drayson, Z. (2012) The uses and abuses of the personal/subpersonal distinction. *Philosophical Perspectives 26* (*1*), 1-18.

Dutilh Novaes, C. (2013) Mathematical Reasoning and External Symbolic Systems. *Logique & Analyse, 221*, 45-65.

Eccles, D. W., & Tenenbaum, G. (2004) Why an Expert Team is More Than a Team of Experts: A Social-Cognitive Conceptualization of Team Coordination and Communication in Sport. *Journal of Sport & Exercise Psychology 26*, 542-560.

Eliasmith, C. (2009) Dynamics, Control, and Cognition. In pp. 134-155 Robbins, P. & Aydede, M. (eds.) *The Cambridge Handbook of Situated Cognition.* Cambridge: Cambridge University Press.

Engelbart, D. C. (1988) Augmenting Human Intellect: A Conceptual Framework. *Summary Report, SRI Project No. 3578, AFOSR-3223, Contract AF 49(638)-1024, 1962*.

Fabry, R. (2017) Betwixt and between: the enculturated predictive processing approach to cognition. *Synthese.* [published online], 1-36.

Fauconnier, G. & Turner, M. (2002) *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities.* New York: Basic Books.

Fleck, L. (1979) *Genesis and development of a scientific fact.* Tr. Bradley, F. & Trenn, T. J. (1935) *Entstehung und Entwicklung einer wissenschaftlichen Tatsache: Einführung in die Lehre von Denkstil und Denkkollektiv*. Chicago: University of Chicago Press.

Flynn, E. G., Laland, K. N., Kendal, R. L. & Kendal, J. R. (2013) Developmental niche construction. *Developmental Science 16*(2), 296-313.

Fodor, J. (1980) Methodological solipsism considered as a research strategy in cognitive psychology. *Behavioral and Brain sciences 3,*63-109.

Fodor, J. (1983) *The Modularity of Mind.* Cambridge, MA: The MIT Press.

Forsyth, D. R. (2009) *Group Dynamics – fifth edition.* Belmont, CA: Wadsworth Cenage Learning.

Fricker, M. (2007) *Epistemic Injustice: Power and the Ethics of Knowing.* Oxford: Oxford University Press.

Frieman, M. (2015) The Art of Drafting and Revision: Extended Mind in Creative Writing. *New Writing: International Journal for the Practice and Theory of Creative Writing 12*(1), 48-66.

Furniss, D., Masci, P., Curzon, P., Mayer, A., & Blandford, A. (2015) Exploring medical device design and use through layers of Distributed Cognition: How a glucometer is coupled with its context. *Journal of Biomedical Informatics 53*, 330-341.

Galantucci, B. & Sebanz, N. (2009) Joint Action: Current Perspectives. *Topics in Cognitive Science 1*, 255-259.

Galison, P. (1997) *Image and logic: a material culture of microphysics.* Chicago, Illinois: University of Chicago Press.

Gaser, C. & Schlaug, G. (2003) Brain Structures Differ between Musicians and Non-Musicians. *The Journal of Neuroscience, 23*(27), 9240-9245.

Gauvain, M. (2000) Niche construction, social co-construction, and the development of the human mind. *Behavioral and Brain Sciences 23*, 153.

Geertz, C. (1973) *The interpretation of culture*. New York: Basic Books.

Giere, R. (2002a) Discussion Note: Distributed Cognition in Epistemic Cultures. *Philosophy of Science, 69*, 637-644.

Giere, R. (2002b) Scientific cognition as distributed cognition. (pp. 285–299) in P. Carruthers, S. Stitch, & M. Siegal (Eds.) *The cognitive basis of science*. Cambridge: Cambridge University Press.

Giere, R. N. (2006) *Scientific Perspectivism*. Chicago: The University of Chicago Press.

Giere, R. (2007) Distributed Cognition without Distributed Knowing. *Social Epistemology, 21*(3), 313-320.

Giere, R. N. (2012) Scientific Cognition: human centred but not human bound. *Philosophical Explorations, 15*(2), 199-206.

Giere, R. N. & Moffatt, B. (2003) Distributed Cognition: Where the Cognitive and the Social Merge. *Social Studies of Science, 33*(2), 301-310.

Glennan, S. (2017) *The New Mechanical Philosophy*. Oxford: Oxford University Press.

Goldstone, R. L. & Theiner, G. (2017) The multiple, interacting levels of cognitive systems (MILCS) perspective on group cognition. *Philosophical Psychology 30* (3), 338-372, DOI: 10.1080/09515089.2017.1295635

Hacking, I. (1983) *Representing and Intervening: Introducing topics in the philosophy of natural science*. Cambridge: Cambridge University Press.

Halverson, C. (2002) Activity Theory and Distributed Cognition: Or What Does CSCW Need to DO with Theories? *Computer Supported Cooperative Work, 11*, 243–267.

Harris, C. B., Keil, P. G., Sutton, J., Barnier, A. J. & McIlwain, D. J. F. (2011) We Remember, We Forget: Collaborative Remembering in Older Couples. *Discourse Processes*, *48*(4), 267-303.

Haugeland, M. (1998) Mind Embodied and Embedded. (pp. 207-237) in J. Haugeland. *Having Thought: Essays in the Metaphysics of Mind.* Cambridge, MA: Harvard University Press.

Haun, D. B. M., Call, J., Janzen, G. & Levinson, S. C. (2006) Evolutionary Psychology of Spatial Representations in the Hominidae. *Current Biology 16*, 1736-1740.

Haun, D. B. M., Rapold, C. J., Janzen, G. & Levinson, S. C. (2011) Plasticity of human spatial cognition: Spatial language and cognition covary across cultures. *Cognition 119*, 70-80.

Haviland, J. B. (1998) Guugu Yimithirr Cardinal Directions. *Ethos 26 (1),* 25-47.

Hazelhurst, B., McMullen, C. K., & Gorman, P. N. (2007) Distributed cognition in the heart room: How situation awareness arises from coordinated communications during cardiac surgery. *Journal of Biomedical Informatics, 40,* 539–551.

Heersmink, R. (2013) A Taxonomy of Cognitive Artifacts: Function, Information, and Categories. *Review of Philosophy and Psychology, 4*(3), 465-481.

Heersmink, R. (2015) Dimensions of integration in embedded and extended cognitive systems. *Phenomenology & Cognitive Science*, *14*, 577–598.

Heersmink, R. (2016) The cognitive integration of scientific instruments: Information, situated cognition, and scientific practice. *Phenomenology and the Cognitive Sciences 15,* 517-537.

Henneberg, M. (1989) Decrease of Human Skull Size in the Holocene. *Human Biology 60*(3), 395-405.

Henrich, J. (2016) *The Secrets of Our Success: How culture is driving human evolution, domesticating our species, and making us smarter.* Oxford: Princeton University Press.

Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H., McElreath, R., Alvard, M., Barr, A., Ensminger, J., Henrich, N. S., Hill, K., Gil-White, F., Gurven, M., Marlowe, F. W., Patton, J. Q. & Tracer, D. (2005) "Economic man" in cross-cultural perspective: Behavioral experiments in 15 small-scale societies. *Behavioral and Brain Sciences, 28*, 795-814.

Henrich, J., Heine, S. J. & Norenzayan, A. (2010) The weirdest people in the world? *Behavioral and Brain Sciences, 33*, 61-83.

Herculano-Houzel, S. (2009) The human brain in numbers: a linearly scaled-up primate brain. *Frontiers in Human Neuroscience 3*, 31.

Hermann, E., Call, J., Hernández-Lloreda, M. V., Hare, B. & Tomasello, M. (2007) Humans Have Evolved Specialized Skills of Social Cognition: The Cultural Intelligence Hypothesis. *Science 317 (5843)*, 1360-1366.

Hey, T., Tansley, S. & Tolle, K. (eds.) (2009) *The Fourth Paradigm: Data-Intensive Scientific Discovery*. Washington: Microsoft Research.

Hofman, M. A. (2014) Evolution of the human brain: when bigger is better. *Frontiers in Human Neuroscience 8,* 15.

Hofstadter, D. R. (2001) Epilogue: Analogy as the Core of Cognition. (pp. 499-538) in D. Gentner, K. J. Holyoak & B. K. Kokinov (eds.) *The Analogical Mind: Perspectives From Cognitive Science.* Cambridge, MA: The MIT Press.

Hohwy, J. (2013) *The Predictive Brain. Oxford: Oxford University Press.*

Hollan, J., Hutchins, E., & Kirsh, D. (2000) Distributed Cognition: Toward a New Foundation for Human-Computer Interaction Research. *ACM Transactions on Computer-Human Interaction, 7* (2)*,* 174–196.

Hollnagel, E. (2001) Extended cognition and the future of ergonomics. *Theoretical Issues in Ergonomics Science 2* (3), 309-315.

Hollnagel, E., Pariès, J., Woods, D. D. & Wreathall, J. (eds.) (2011) *Resilience Engineering in Practice.* Surrey: Ashgate.

Hu, Y., Geng, F., Tao, L., Hu, N., Du, F., Fu, K. & Chen, F. (2011) Enhanced White Matter Tracts Integrity in Children With Abacus Training*. Human Brain Mapping 32,* 10-21.

Huebner, B. (2014) *Macrocognition: A Theory of Distributed Minds and Collective Intentionality*. Oxford: Oxford University Press.

Humphreys, P. (2004) *Extending Ourselves: Computational Science, Empiricism, and the Scientific Method*. Oxford: Oxford University Press.

Humphreys, P. (2009) The philosophical novelty of computer simulation methods. *Synthese 169,* 615-626.

Hurley, S. (2010) The Varieties of Externalism. (pp. 101-153) in R. Menary (Ed.) *The Extended Mind*. Cambridge, MA: MIT Press.

Hutchins, E. (1995a) *Cognition in the wild*. Cambridge, MA: MIT Press

Hutchins, E. (1995b) How a Cockpit Remembers Its Speeds. *Cognitive Science, 19,* 265-288

Hutchins, E. (1999) Cognitive artifacts. (pp. 126-128) in R. A. Wilson & F. C. Keil (Eds.) *The MIT Encyclopaedia of the Cognitive Sciences*. Cambridge, MA: MIT Press.

Hutchins, E. (2001) Cognition, Distributed. (pp. 2068-2072) in R. A. Wilson & F. C. Keil (Eds.) *The MIT Encyclopedia of the Cognitive Sciences*. Cambridge, MA: MIT Press.

Hutchins, E. (2005) Material anchors for conceptual blends. *Journal of Pragmatics, 37*, 1555-1577

Hutchins, E. (2006) The distributed cognition perspective on human interaction. (pp. 375-398) in N. J. Enfield & S. C. Levinson (Eds.) *Roots of human sociality: Culture, cognition and interaction*. Oxford: Berg.

Hutchins, E. (2008) The role of cultural practices in the emergence of modern human intelligence. *Philosophical Transactions of the Royal Society B: Biological Sciences, 363*, 2011-2019

Hutchins, E. (2010a) Cognitive Ecology. *Topics in Cognitive Science 2*, 705-715.

Hutchins, E. (2010b) Enaction, imagination, and Insight. In J. Stewart, O. Gapenne, & E. A. Di Paolo (Eds.) *Enaction: Towards a New Paradigm Cognitive Science*. Cambridge, MA: MIT Press

Hutchins, E. (2010c) Imagining the cognitive life of things. (pp. 91-102) in L. Malafouris & C. Renfrew (Eds.) *The cognitive life of things: Recasting the boundaries of the mind*. Oxford: McDonald Institute Monograph

Hutchins, E. (2011) Enculturating The Supersized Mind. *Philosophical Studies, 152,* 437-446.

Hutchins, E. (2012) Concepts in practice as a source of order.

Hutchins, E. (2014) The cultural ecosystem of human cognition. *Philosophical Psychology, 27(1)*, 34-49

Hutto, D., Kirchhoff, M. D. & Myin, E. (2014) Extensive enactivism: why keep it all in? *Frontiers in Human Neuroscience*, *8*, 706.

Ifrah, G. (1998) *Universal History of Numbers: from prehistory to the invention of the computer.* London: Harvill Press.

Ingold, T. (2000) *The Perception of the Environment – Essays on Livelihood, dwelling and skill.* London: Routledge.

Iriki, A. & Taoka, M. (2012) Triadic (ecological, neural, cognitive) niche construction: a scenario of human brain evolution extrapolating tool use and language from the control of reaching actions. *Philosophical Transactions of the Royal Society B: Biological Sciences, 367*, 10-23.

Isenberg, D. (1986) Group Polarization: A Critical Review and Meta-Analysis. *Journal of Personality and Social Psychology 50*(6), 1141-1151.

Iyengar, S. & Hahn, K. S. (2009) Red Media, Blue Media: Evidence of Ideological Selectivity in Media Use. *Journal of Communication 59*, 19-39.

Jaarsveld, S., Lachmann, T. & Hamel, R. & van Leeuwen, C. (2010) Solving and Creating Raven Progressive Matrices: Reasoning in Well- and Ill-Defined Problem Spaces. *Creativity Research Journal 22*(3), 304-319.

Japyassú, H. F. & Laland, K. N. (2017) Extended spider cognition. *Animal Cognition 20*(3), 375-395.

Jaušovec, N. & Jaušovec, J. (2000) EEG activity during the performance of complex mental problems. *International Journal of Psychophysiology 36*, 73-88.

Jeffares, B. (2010) The co-evolution of tools and minds: cognition and material culture in the hominin lineage. *Phenomenology and Cognitive science 9*, 503-520.

Johannsen, M. (2010) Technological Conceptualization: Cognition on the Shoulders of History. (pp. 59-70) in L. Malafouris & C. Renfrew (Eds.) *The cognitive life of things – Recasting the boundaries of the mind*. Oxford: McDonald Institute Monograph.

Johnson, C. M. (2001) Distributed primate cognition: A review. *Animal Cognition 3*(4), 167-183.

Johnson, C. M. (2015) The Cognitive Ecology of Dolphin Social Engagement. (pp. 229-256) in D. L. Herzing & C. M. Johnson (Eds.) *Dolphin Communication and Cognition*. Cambridge: MIT Press.

Johnson-Frey, S. H. (2003) What's so special about human tool use? *Neuron 39*, 201-204.

Kamhi, J. F., Gronnenberg, W., Robson, S. K. A. & Traniello, J. F. A. (2016) Social complexity influences brain investment and neural operation costs in ants. *Proceedings of the Royal Society B 283*, 20161949.

Kaplan, D. M. (2012) How to demarcate the boundaries of cognition. *Biology & Philosophy 27*, 545-570.

Kauffman, S. (1995) *At Home in the Universe: The Search for the Laws of Self-Organization and Complexity.* Oxford: Oxford University Press.

Keil, F. C. (2011) Science Starts Early. *Science 331*, 1022-1023.

Keil, P. G. (2015) Human-Sheepdog Distributed Cognitive Systems: An Analysis of Interspecies Cognitive Scaffolding in a Sheepdog Trial. *Journal of Cognition and Culture 15*, 508-529.

Kendal, J. R. (2011) Cultural Niche Construction and Human Learning Environments: Investigating Sociocultural Perspectives. *Biological Theory 6*, 241-250.

Kerr, E. & Gelfert, A. (2014) The 'Extendedness' Of Scientific Evidence. *Philosophical Issues 24,* 253-281.

Kirchhoff, M. D. (2012) Extended cognition and fixed properties: steps to a third-wave version of extended cognition. *Phenomenology and the Cognitive Science 11*, 287-308.

Kirchhoff, M. D. (2014) Extended cognition & constitution: Re-evaluating the constitutive claim of extended cognition. *Philosophical Psychology 27*(2), 258-283.

Kirchhoff, M. D. (2016) From mutual manipulation to cognitive extension: challenges and implications. *Phenomenology and the Cognitive Sciences,* (online) *1-16.*

Kirsh, D. (1995) The intelligent use of space. *Artificial Intelligence 73*, 31-68.

Kirsh, D. (2006) Distributed cognition – A methodological notice. *Pragmatics & Cognition 14*(2)*,* 249-262.

Kirsh, D. (2010) Thinking with external representations. *AI & Society 25,* 441-454.

Kirsh, D., & Maglio, P. (1994) On Distinguishing Epistemic from Pragmatic Actions. *Cognitive Science 18*, 513-549.

Kitcher, P. (1993) *The Advancement of Science*: *Science Without Legend, Objectivity Without Illusions*. Oxford: Oxford University Press.

Knoblich, G. & Sebanz, N. (2006) The Social Nature of Perception and Action. *Current Directions in Psychological Science 15*(3), 99-104.

Konishi, K. & Bohbot, V. D. (2013) Spatial navigational strategies correlate with gray matter in the hippocampus of healthy older adults tested in a virtual maze. *Frontiers in Aging Neuroscience 5,* 1.

Knorr Cetina, K. (1999) *Epistemic cultures: How the Sciences Make Knowledge.* Cambridge, MA: Harvard University Press.

Knuuttila, T. (2011) Modelling and representing: an artefactual approach to model-based representation. *Studies in History and Philosophy of Science 42*, 262-271.

Krakauer, J. W., Ghazanfar, A. A., Gomez-Marin, A., MacIver, M. A. & Poeppel, D. (2017) Neuroscience Needs Behavior: Correcting a Reductionist Bias. *Neuron 93*, 480-490.

Kuhn, T. (1970) *The Structure of Scientific Revolutions, Second Edition*. The University of Chicago Press.

Kuhn, T. (1977) *The Essential Tension: Selected Studies in Scientific Tradition and Change.* Chicago: The Chicago University Press.

Kukla, R. (2012) "Author TBD": Radical Collaboration in Contemporary Biomedical Research. *Philosophy of Science 79*, 845-858.

Ladyman, J. & Ross, D. with Spurrett, D. & Collier, J. (2007) *Every Thing Must Go – Metaphysics Naturalized*. Oxford: Oxford University Press.

Laland, K. N., Odling-Smee, J., & Feldman, M. W. (2000) Niche construction, biological evolution, and cultural change. *Behavioral and Brain Sciences 23*, 131-175.

Laland, K. N., Odling-Smee, J., & Myles, S. (2010). How culture shaped the human genome: bringing genetics and the human sciences together. *Nature Reviews Genetics 11*, 137-148.

Landy, D., Allen, C., & Zednik, C. (2014) A perceptual account of symbolic reasoning. *Frontiers in Psychology 5*, 1-10.

Landy, D., & Goldstone, R. L. (2007) How Abstract is Symbolic Thought? *Journal of Experimental Psychology 33*(4), 720-733.

Lave, J. (1988) *Cognition in Practice: Mind, mathematics and culture in everyday life.* Cambridge, Cambridge University Press.

Leitão, H. & Gasper, J. A. (2014) Globes, Rhumb Tables, and the Pre-History of the Mercator Projection. *Imago Mundi 66*(2), 180-195.

Levinson, S. C. (2003) *Space in Language and Cognition: Explorations in Cognitive Diversity.* Cambridge: Cambridge University Press.

Levy, N. (2017) Due Deference to Denialism: Explaining Ordinary People's Rejection of Established Scientific Findings. *Synthese.* [online]. 1-15.

Li, S-C. (2003) Biocultural Orchestration of Developmental Plasticity Across Levels: The Interplay of Biology and Culture in Shaping the Mind and Behavior Across the Life Span. *Psychological Bulletin 29*(2), 171-194.

Liu, Z., Nersessian, N. J. & Stasko, J. T. (2008) Distributed Cognition as a Theoretical Framework for Information Visualization. *IEEE Transactions on Visualization and Computer Graphics 14*(6), 1173-1180.

Luo, J. & Knoblich, G. (2007) Studying insight problem solving with neuroscientific methods. *Methods 42*, 77-86.

Luria, A. R. (1928) The Problem of the Cultural Behavior of the Child. *The Pedagogical Seminary and Journal of Genetic Psychology, 35 (4)*, 493-506.

Luria, A. R. (1976) *Cognitive Development: Its Cultural and Social Foundations.* Cambridge, MA: Harvard University Press.

Magnus, P.D. (2007) Distributed Cognition and the Task of Science. *Social Studies of Science, 37(2),* 297-310.

Magnus, P. D. & McClamrock, R. (2015) Friends with benefits! Distributed cognition hooks up cognitive and social conceptions of science. *Philosophical Psychology, 28(8),* 1114-1127.

Maguire, E. A., Gadian, D. G., Johnsrude, I. S., Good, C. D., Ashburner, J., Frackowiak, S. J., & Frith, C. D. (2000) Navigation-related structural change in the hippocampi of taxi drivers. *PNAS 97*(8), 4398-4403.

Majid, A., Bowerman, M., Kita, S., Haun, D. B. M. & Levinson, S. C. (2004) Can language restructure cognition? The case for space. *TRENDS in Cognitive Sciences 8*(3), 108-114.

Malafouris, L. (2008) Beads for a Plastic Mind: the 'Blind Man's Stick' (BMS) Hypothesis and the Active Nature of Material Culture. *Cambridge Archaeological Journal 18*(3), 401-414.

Malafouris, L. (2010) Knapping Intentions and the Marks of the Mental. (pp. 13-22) in L. Malafouris & C. Renfrew (Eds.) *The cognitive life of things: Recasting the boundaries of the mind*. Oxford: McDonald Institute Monograph.

Malafouris, L. (2013) *How Things Shape the Mind: A Theory of Material Engagement*. Cambridge, MA: MIT Press.

Maravita, A. & Iriki, A. (2004) Tools for the body (schema). *Trends in Cognitive Sciences 8*(2), 79-86.

Marr, D. (1982/2010) *Vision: A computational Approach*. London: MIT Press.

Mason, P. H. (2014) Degeneracy: Demystifying and Destigmatizing a Core Concept in Systems Biology. *Complexity 20* (3), 12-21.

Mason, P. H., Domínguez, J.F., Winter, B. & Grignoli, A. (2015) Hidden in plain view: degeneracy in complex systems. *BioSystems 128*, 1-8.

Maynard Smith, J. (1968) *Mathematical Ideas in Biology*. Cambridge: Cambridge University Press.

Meade, M. L., Nokes, T. & Morrow, D. G. (2009) Expertise promotes facilitation on a collaborative memory task. *Memory 17*(1), 39-48.

Menary, R. (2007a) *Cognitive Integration: Mind and Cognition Unbounded.* Hampshire: Palgrave Macmillan.

Menary, R. (2007b) Writing as thinking. *Language Sciences, 29*, 621-632.

Menary, R. (2010a) Cognitive Integration and the Extended Mind. (pp. 227-244) in R. Menary (Ed.) *The Extended Mind*. Cambridge, MA: MIT Press.

Menary, R. (2010b) Dimensions of mind. *Phenomenology and Cognitive Science, 9*, 561-578.

Menary, R. (2010c) Introduction. (pp. 1-25) in Menary, R. (Ed.) *The Extended Mind*. Cambridge, MA: MIT Press.

Menary, R. (Ed.) (2010d) *The Extended Mind*. Cambridge, MA: MIT Press.

Menary, R. (2010e) The holy grail of cognitivism: a response to Adams and Aizawa. *Phenomenology and Cognitive Science, 9*, 605-618.

Menary, R. (2012) Cognitive practices and cognitive character. *Philosophical Explorations, 15*(2), 147-164.

Menary, R. (2013a) Cognitive integration, enculturated cognition and the socially extended mind. *Cognitive Systems Research, 25-26*, 26-34.

Menary, R. (2013b) The Enculturated Hand. (pp. 349-367) in Radman, Z. (ed.) *The Hand, an Organ of the Mind: What the Manual Tells the Mental.* London: MIT Press.

Menary, R. (2014) Neural Plasticity, Neuronal Recycling and Niche Construction. *Mind and Language 29* (3), 286-303.

Menary, R. (2015) Mathematical Cognition: A Case of Enculturation. In T. Metzinger & J. M. Windt (Eds.), *Open MIND 25(T),* 1-20.

Menary, R. (2016) Pragmatism and the Pragmatic Turn in Cognitive Science. (pp. 217-237) in A. K. Engel, K. J. Friston and D. Kragic (Eds.), *The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science. Strüngmann Forum Reports (Vol. 18)*. Cambridge, MA: MIT Press.

Menary, R. & Gillett, A. J. (2017) Embodying Culture: Integrated Cognitive Systems and Cultural Evolution. (pp. 72-87) in J. Kiverstein (ed.) *The Routledge Handbook of Philosophy of the Social Mind.* New York: Routledge.

Menary, R. & Kirchhoff, M. (2014) Cognitive transformations and extended expertise. *Educational Philosophy and Theory, 46*(6), 610-623.

Mercier, H. (2011) On the Universality of Argumentative Reasoning. *Journal of Cognition and Culture 11,* 85-113.

Mercier, H. & Sperber, D. (2011) Why do humans reason? Arguments for an argumentative theory. *Behavioral and Brain Sciences 34*, 57-111.

Michaelian, K. & Sutton, J. (2013) Distributed Cognition and Memory Research: History and Current Directions. *Review of Philosophy and Psychology 4*(1), 1-24.

Miller, G. A. (1956) Lucky number seven, plus or minus two: some limits on our capacity for processing information. *The Psychological Review 63*(2), 81-97.

Miller, G. A. (2003) The cognitive revolution: a historical perspective. *TRENDS in Cognitive Sciences 7* (3)*,* 141-144.

Minson, J.A., & Mueller, J. S. (2012) The cost of collaboration: why joint decision making exacerbates rejection of outside information. *Psychological Sciences 23*, 219-224.

Moore, J. D. (2012) *Visions of Culture: An Introduction to Anthropological Theories and Theorists (Fourth edition).*Plymouth: Altamira Press.

Moore, J. L. & Rocklin, T. R. (1998) The Distribution of Distributed Cognition: Multiple Interpretations and Uses. *Educational Psychology Review 10*(1)*,* 97-113.

Nan, Y., Knöshe, T. R., Zysset, S. & Friederici, A. D. (2008) Cross-Cultural Music Phrase Processing: An fMRI Study. *Human Brain Mapping 29*, 312-328.

Nardi, B. A. (1996) Concepts of cognition and consciousness: four voices. *Australian Journal of Information Systems 4*, 64-79.

Nemirovsky, R., Kelton, M. L., & Rhodehamel, B. (2013) Playing Mathematical Instruments: Emerging Perceptuomotor Integration with an Interactive Mathematics Exhibit. *Journal for Research in Mathematics Education 44*(2), 372-415

Nersessian, N. J. (2005) Interpreting Scientific and Engineering Practices: Integrating the Cognitive, Social, and Cultural Dimensions. (pp. 17-56) in M. E. Gorman, R. D. Tweney, D. C. Gooding & A. P. Kincannon (Eds.) *Scientific and Technological Thinking*. London: Lawrence Erlbaum Associate.

Nersessian, N. J. (2006) Model-Based Reasoning in Distributed Cognitive Systems. *Philosophy of Science 73*(5), 699-709.

Nersessian, N. J. (2009) How Do Engineering Scientists Think? Model-Based Simulation in Biomedical Engineering Research Laboratories. *Topics in Cognitive Science 1*, 730-757.

Nersessian, N. J., Kurz-Milcke, E., Newstetter, W. C., & Davies, J. (2004) Research Laboratories as Evolving Distributed Cognitive Systems. (pp. 857-862) in A. Markman & L. Barsalou (eds.) *Proceedings of The 25th Annual Conference of the Cognitive Science Society.* Hillsdale, NJ:

Nieder, A., & Dehaene, S. (2009) Representation of number in the brain. *Annual Review of Neuroscience 32*, 185–208.

Nilsson, M., van Laere, J., Susi, T. & Ziemke, T. (2012) Information fusion in practice: A distributed cognition perspective on the active role of users. *Information Fusion 13*, 60-78.

Nisbet, R. E., Choi, I., Peng, K. & Norenzayan, A. (2001) Culture and Systems of Thought: Holistic Versus Analytic Cognition. *Psychological Review 108*(2), 291-310.

Norman, D. A. (1986) Cognitive engineering. (pp. 31-61) in D. A. Norman and S. Draper (eds.), *User Centered System Design: New Perspectives in Human-Computer Interaction*. Hillsdale, NJ: Erlbaum.

Norman, D. A. (1987) *The Psychology of Everyday Things*. Basic Books.

Norman, D. (1991) Cognitive Artifacts. (pp. 17-38) in J. M. Carroll, (Ed.) *Designing Interaction*. Cambridge: Cambridge University Press.

Norman, D. (1993) *Things That Make Us Smart: Defending Human Attributes in the Age of the Machine*. Reading, MA: Addison-Wesley.

O'Donnell, S., Bulova, S. J., DeLeon, S., Khodak, P., Miller, S. & Sulger, E. (2015) Distributed cognition and social brains: reductions in mushroom body investment accompanied the origins of sociality in wasps (Hymenoptera: Vespidae). *Proceedings of the Royal Society B 282*, 20150791.

Odling-Smee, F. J. (2007) Niche inheritance: A possible basis for classifying multiple inheritance systems in evolution. *Biological Theory 2*, 276-289.

Odling-Smee, F. J., Laland, K. N. & Feldman, M. F. (2003) *Niche Construction: The Neglected Process in Evolution*. Princeton, NJ: Princeton University Press.

Oppenheim, P. & Putnam, H. (1958) Unity of Science as a Working Hypothesis. (pp. 3-36) in H. Feigl, M. Scriven, & G. Maxwell (eds.) *Concepts, Theories, and the Mind-Body Problem, Minnesota Studies in the Philosophy of Science II*. Minneapolis: University of Minnesota Press.

Ortner, S. (1984) Theory in Anthropology since the Sixties. *Comparative Studies in Society and History 26 (1),* 126-166.

Oyserman, D., Coon, H. M. & Kemmelmeier, M. (2002) Rethinking Individualism and Collectivism: Evaluation of Theoretical Assumptions and Meta-Analyses. *Psychological Bulletin 128*(1), 3-72.

Pears, N. (1985) *Basic Biogeography (Second Edition)*. New York: John Wiley & Sons.

Perry, M. (2003) Distributed Cognition. In (pp. 193-223) J. M. Carroll (ed.) *HCI Models, Theories, and Frameworks: Toward an Interdisciplinary Science*. San Francisco: Morgan Kaufmann.

Petroski, H. (1989) *The Pencil: A History of Design and Circumstance*. New York: Knopf.

Pinker S (2010) The cognitive niche: Coevolution of intelligence, sociality, and language. *PNAS 107*(Supplement 2), 8993-8999.

Poirier, P. & Chicoisne, G. (2006) A framework for thinking about distributed cognition. *Pragmatics & Cognition 14(2)*, 215–234

Preston, B. (1988) Cognition and Tool Use. *Mind & Language 13*(4), 513-547.

Pritchard, D. (2010) Cognitive ability and the extended cognition thesis. *Synthese 175*, 133-151.

Pritchard, D. (2016) Intellectual Virtues, Extended Cognition, and the Epistemology of Education (pp. 113-127) in J. Baehr (ed.) *Intellectual Virtues and Education: Essays in Applied Virtue Epistemology*. London: Routledge.

Qiu, J., Hong, L., Yang, D., Luo, Y., Ling, L. Wu, Z. & Zhang, Q. (2008) The neural basis of insight problem solving: An event-related potential study. *Brain and Cognition 68*, 100-106.

Reader, S. M., Hager, Y. & Laland, K. N. (2011) The evolution of primate general and cultural intelligence. *Philosophical Transactions of the Royal Society B 366*, 1017-1027.

Reynolds Losin, E. A., Dapretto, M. & Iacoboni, M. (2010) Culture and neuroscience: additive or synergistic? *SCAN 5*, 148-158.

Rilling, J. K. (2006) Human and nonhuman primate brains: Are they allometrically scaled versions of the same design? *Evolutionary Anthropology 15*, 65-77.

Robbins, P. & Aydede, M. (eds.) (2009) *The Cambridge Handbook of Situated Cognition.* Cambridge: Cambridge University Press.

Roepstorff, A. (2011) Culture: A Site of *Relativist* Energy in the Cognitive Sciences. *Common Knowledge 17*(1), 37-41.

Roepstorff, A. (2013) Why Am I Not Just Lovin' Cultural Neuroscience? Toward a Slow Science of Cultural Difference. *Psychological Inquiry 24*(1), 61-63.

Roepstorff, A. & Bubandt, N. (2003) General Introduction: The critique of culture and the plurality of nature. (pp. 9-26) in Roepstorff, A., Bubandt, N. and Kull, K (eds.) *Imagining Nature: Practices of Cosmology and Identity*. Aarhus: Aarhus University Press.

Roepstorff, A., Niewöhner, J. & Beck, S. (2010) Enculturing brains through patterned practices. *Neural Networks 23*, 1051–1059.

Rogers, Y. & Ellis, J. (1994) Distributed cognition: an alternative framework for analysing and explaining collaborative working. *Journal of Information Technology 9,* 119-128.

Rogoff, B. (1990) *Apprenticeship in Thinking: Cognitive Development in Social Context.* Oxford: Oxford University Press.

Ross, D. & Ladyman, J. (2010) The Alleged Coupling-Constitution Fallacy and the Mature Sciences. (pp. 155-166) in Menary, R. (Ed.) *The Extended Mind*. Cambridge, MA: MIT Press.

Rowlands, M. (1999) *The Body in Mind: Understanding Cognitive Processes.* Cambridge: Cambridge University Press.

Rowlands, M. (2009) Extended cognition and the mark of the cognitive. *Philosophical Psychology 22*(1), 1-19.

Rowlands, M. (2010) *The New Science of the Mind: From Extended Mind to Embodied Phenomenology.* London: The MIT Press.

Ruff, C. (2002) Variation in Human Body Size and shape. *Annual Review of Anthropology 31*, 211-232.

Rumelhart, D. E., Smolensky, P. & Hinton, G. E. (1986) Schemata and sequential thought processes in PDP models. (pp. 7-57) in J. McClelland & D. E. Rumelhart (eds.) *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Vol. 2.* Cambridge, MA: MIT Press

Rupert, R. D. (2004) Challenges to the Hypothesis of Extended Cognition. *The Journal of Philosophy 101*(8), 389-428.

Rupert, R. D. (2010) Representation in Extended Cognitive Systems: Does the Scaffolding of Language Extend the Mind? (pp. 325-353) in R. Menary (Ed) *The Extended Mind.* Cambridge, MA: MIT Press.

Rupert, R. (2011) Empirical Arguments for Group Minds: A Critical Appraisal. *Philosophy Compass 6/9*, 630-639.

Searle, J. (1980) Minds, Brains, and Programs. *Behavioral and Brain Sciences 3,* 417-457.

Shea, N. (2009) Imitation as an Inheritance System. *Philosophical Transactions of the Royal Society B 364*, 2429-2443.

Shore, B. (1996) *Culture in Mind – Cognition, Culture, and the Problem of Meaning*. Oxford: Oxford University Press.

Simon, H. A. (1996) *The Sciences of the Artificial: Third edition.* London: The MIT Press.

Skagestad, P. (1993) Thinking With Machines: Intelligence Augmentation, Evolutionary Epistemology, and Semiotic. *Journal of Social and Evolutionary Systems* 16(2), 157-180.

Solomon, M. (2006) Groupthink versus The Wisdom of Crowds: The Social Epistemology of Deliberation and Dissent. *The Southern Journal of Philosophy 44*, 28-42.

Stanford, K. P. (2015) Unconceived alternatives and conservatism in science: the impact of professionalization, peer-review, and Big Science. *Synthese*, 1-18. doi:10.1007/s11229-015-0856-4 [online].

Stanford, K. P. (2016) Naturalism without Scientism. (pp. 91-108) in K. J. Clark (ed.) *The Blackwell Companion to Naturalism*. Chichester: John Wiley & Sons.

Stasser, G. & Titus, W. (1985) Pooling of unshared information in group decision making: biased information sampling during group discussion. *Journal of Personality and Social Psychology 48*, 1467-1478.

Steiner, I. (1966) Models for inferring relationships between group size and potential group productivity. *Behavioural Science, 11(4),* 273-283.

Steiner, I. (1972) *Group processes and Productivity*. New York: Academic Press.

Sterelny, K. (2003) *Thought in a Hostile World: The Evolution of Human Cognition*. Oxford: Blackwell Publishing.

Sterelny, K. (2010) Minds: extended or scaffolded? *Phenomenology and the Cognitive Sciences 9* (4), 465-481.

Sterelny, K. (2011) From hominins to humans: how *sapiens* became behaviourally modern. *Philosophical Transactions of the Royal Society B: Biological Sciences 366*, 809-822.

Sterelny, K. (2012) *The Evolved Apprentice: How evolution made humans unique*. Cambridge, MA: MIT Press.

Sterelny, K. (2014) Constructing the Cooperative Niche. (pp. 261-279) in G. Barker, E. Desjardins & T. Pearce (eds.) *Entangled Life –Organism and Environment in the Biological and Social Sciences*. London: Springer.

Stewart, I. (2012) *Seventeen equations that changed the world.* London: Profile Books.

Stich, S. (2010) Philosophy and WEIRD intuition. *Behavioral and Brain Sciences 33,* 110-111.

Stotz, K. (2010) Human nature and cognitive-developmental niche construction. *Phenomenology and the Cognitive Science, 9*, 483-501.

Sunstein, C. R. (2002) The Law of Group Polarization. *Journal of Political Philosophy 10*, 175-195.

Sunstein, C.R. (2005) Group Judgments: Statistical Means, Deliberation, and Information Markets. *New York University Law Review 80*, 962-1049.

Sutton, J. (2006) Distributed cognition – Domains and dimensions. *Pragmatics& Cognition 14(2),* 235-247.

Sutton, J (2007) Batting, Habit, and Memory: The Embodied Mind and the Nature of Skill. *Sport in Society 10*(5), 763-786.

Sutton, J. (2008) Material Agency, Skills and History: Distributed Cognition and the Archaeology of Memory. (pp. 37-55) in C. Knappett & L. Malafouris (eds.) *Material Agency: Towards a Non-Anthropocentric Approach.* London: Springer.

Sutton, J. (2010) Exograms and Interdisciplinarity: History, the Extended Mind, and the Civilizing Process. (pp. 189-226) in R. Menary (Ed) *The Extended Mind.* Cambridge, MA: MIT Press.

Sutton, J. (2013) Skill and Collaboration in the Evolution of Human Cognition. *Biological Theory 8* 28-36.

Sutton, J. (2015) Remembering as Public Practice: Wittgenstein, memory, and distributed cognitive ecologies. (pp. 409-443) in D. Moyal-Sharrock, A. Coliva & V. A. Munz (eds.) *Mind, Language, and Action: proceedings of the 36th Wittgenstein symposium*. Berlin: De Gruyter.

Sutton, J., Harris, C. B., Keil, P. G. & Barnier, A. J. (2010) The psychology of memory, extended cognition, and socially distributed remembering. *Phenomenology and the Cognitive Sciences 9*(4), 521-560.

Snyder, J. P. (1987) *Map Projections – A Working Manual. U.S. Geological Survey Professional Paper 1395*. Washington, D.C: United States Government Printing Office.

Tanaka, S., Seki, K., Hanakawa, T., Harada, M., Sugawara, S. K., Sadato, N., Watanabe, K. & Honda, M. (2012) Abacus in the brain: a longitudinal functional MRI study of a skilled abacus user with a right hemispheric lesion. *Frontiers in Psychology 315,* 1-13.

Tang, Y., Zhang, W., Kewel, C., Feng, S., Ji, Y., Shen, J., Reiman, E. M., & Llu, Y. (2006) Arithmetic processing in the brain shaped by cultures. *PNAS 103*(28), 10775-10780.

Thagard, P. (1993) Societies of Minds: Science as Distributed Computing. *Studies in History and Philosophy of Science 24*(1), 49-61.

Thagard, P. (2006) *Hot Thought: Mechanisms and Applications of Emotional Cognition*. Cambridge, MA: MIT Press.

Theiner, G. (2014) Onwards and Upwards with the Extended Mind: From Individual to Collective Epistemic Action. (pp. 191-208) in L. R. Caporael, J. R. Griesemer & W. C. Wimsatt (eds.) *Developing Scaffolds in Evolution, Culture, and Cognition*. Cambridge, MA: The MIT Press.

Theiner, G., Allen, C. & Goldstone, R. L. (2010) Recognising group cognition. *Cognitive Systems Research 11*, 378-395.

Theiner, G. & Drain, C. (2016) What's the Matter with cognition? A 'Vygotskian' perspective on material engagement theory. *Phenomenology and the Cognitive Sciences* [online].

Theiner, G. & O'Connor, T. (2010) The Emergence of Group Cognition. (pp. 78-117) in A. Corradini & T. O'Connor (eds.) *Emergence in Science and Philosophy.* New York: Routledge.

Tomasello, M. (1999) *The Cultural Origins of Human Cognition*. London: Harvard University Press.

Tomasello, M. (2009) *Why We Cooperate*. London: MIT Press.

Toon, A. (2014a) Empiricism for Cyborgs. *Philosophical Issues 24,* 409-425.

Toon, A. (2014b) Friends at last? Distributed cognition and the cognitive/social divide. *Philosophical Psychology 27* (1), 112-125.

Tribble, E. (2005) Distributing cognition in the globe. *Shakespeare Quarterly 56*, 135-155.

Tversky, B. (2009) Spatial Cognition: Embodied and Situated. (pp. 201-216) in Robbins, P. & Aydede, M. (eds.) (2009) *The Cambridge Handbook of Situated Cognition.* Cambridge: Cambridge University Press.

Vaesen, K. (2011) Giere's (In)Appropriation of Distributed Cognition. *Social Epistemology 25*(4), 379-391.

Vaesen, K. (2012) The cognitive bases of human tool use. *Behavioral and Brain Sciences 35*, 203–262.

van Eck, D. & de Jong, H. L. (2016) Mechanistic explanation, cognitive systems demarcation, and extended cognition. *Studies in History and Philosophy of Science 59*, 11-21.

van Shaik, C. P. & Burkart, J. M. (2011) Social learning and evolution: the cultural intelligence hypothesis. *Transactions of the Royal Society, Biological Sciences 366*, 1008-1016.

Velliste, M, Perel, S., Spalding, M. C., Whitford, A. S., & Schwartz, A. B. (2008) Cortical control of a prosthetic arm for self-feeding. *Nature 453,* 1098-1101.

Vogeley, K. & Roepstorff, A. (2009) Contextualising culture and social cognition. *Trends in Cognitive Sciences 13*, 511-516.

von Hippel, E. (1994) 'Sticky information' and the locus of problem solving. *Management Science 40*, 429-439.

Vorms, M. (2012) Formats of Representation in Scientific Theorizing. (pp. 250-272) in P. Humphreys & C. Imbert (eds.) *Models, Simulations, and Representations*. London: Routledge.

Vygotsky, L. S. (1978) *Mind in society: The development of higher psychological processes*. Cambridge, MA: Harvard University Press.

Vygotsky, L. S. (1981) The genesis of higher mental functions. (pp. 144-189) in J. W. Wertsch (ed.) *The Concept of Activity in Soviet Psychology*. New York: M. E. Sharpe, Inc.

Wacquant, L. (2014) Homines in Extremis: What Fighting Scholars Teach Us about Habitus. *Body & Society 20*(2), 3-17.

Watkins, J. W. N. (1952) The Principles of Methodological Individualism. *The British Journal of Philosophy 3*, 186-189.

Weisberg, M. & Muldoon, R. (2009) Epistemic landscapes and the division of cognitive labour. *Philosophy of Science 76*(2), 225-252.

Wertsch, J. (1985) *Vygotsky and the Social Formation of Mind.* Cambridge, MA: Harvard University Press.

Williamson, K. & Cox, R. (2014) Distributed Cognition in Sports Teams: Explaining successful and expert performance. *Educational Philosophy and Theory 46* (6), 640-654.

Williamson, K. & Sutton, J. (2014) Embodied Collaboration in Small Groups. (pp. 107-133) in C. T. Wolfe (ed.) *Brain Theory: Essays in Critical Neurophilosophy*. Basingstoke: Palgrave Macmillan.

Wilson, R. A. (2001) Group-level cognition. *Philosophy of Science 68 (3)* (Supplement), S262-S273.

Wilson, R. A. (2004) *Boundaries of the Mind: The Individual in the Fragile Sciences.* Cambridge: Cambridge University.

Wilson, R. A. (2005) Collective memory, group minds, and the extended mind thesis. *Cognitive Processing 6*(4), 227-236.

Wimsatt, W. C. (1986) Forms of Aggregativity. (pp. 259-291) in Donagan, A., Perovich, J. R. & Wedin, M. V. (eds.) *Human Nature and Natural Knowledge: Essays Presented to Marjorie Grene on the Occasion of her Seventy-Fifth Birthday.* Boston: D. Reidel Publishing Limited.

Wimsatt, W. (2007) *Re-Engineering Philosophy for Limited Beings: piecewise approximations to reality.* Cambridge, MA: Harvard University Press.

Wnuk, E. & Majid, A. (2014) Revisiting the limits of language: The odor lexicon of Maniq. *Cognition 131*, 125-138.

Xu, L. & Clarke, D. (2012) What Does Distributed Cognition Tell Us about Student Learning of Science? *Research in Science Education 42*(3), 491-510.

Zhang, J. (1997a) Distributed Representation as a Principle for the Analysis of Cockpit Information Displays. *The International Journal of Aviation Psychology 7*(2), 105-121.

Zhang, J. (1997b) The Nature of External Representations. *Cognitive Science 21(2)*, 179-217.

Zhang, J. & Norman (1994) Representations in Distributed Cognitive Tasks. *Cognitive Science 18,* 87-122.

Zhang, J. & Norman (1995) A representational analysis of numeration systems. *Cognition 57*, 271-295.

Zollman, K. J. S. (2010) The epistemic benefit of transient diversity. *Erkenntnis 72 (1)*, 17-35.