

# THE TEMPORAL DYNAMICS OF VISUAL OBJECT RECOGNITION

---

ERIKA W. CONTINI

Bachelor of Liberal Studies (Psychology Honours)

Perception in Action Research Centre & ARC Centre of Excellence in  
Cognition and its Disorders, Department of Cognitive Science

&

Department of Psychology

Faculty of Human Sciences

Macquarie University, Sydney, Australia

Presented for the degree of

*Combined Doctor of Philosophy / Master of Clinical Neuropsychology*

March 2018





# TABLE OF CONTENTS

<b>TABLE OF CONTENTS.....</b>	<b>I</b>
<b>SUMMARY .....</b>	<b>V</b>
<b>STATEMENT OF AUTHENTICATION.....</b>	<b>VII</b>
<b>ACKNOWLEDGEMENTS.....</b>	<b>IX</b>
<b>AUTHOR NOTE.....</b>	<b>XI</b>
<b>CHAPTER ONE: GENERAL INTRODUCTION .....</b>	<b>1</b>
I. INTRODUCTION TO THESIS TOPIC .....	3
II. REFERENCES .....	6
<b>1. LITERATURE REVIEW: DECODING THE TIME-COURSE OF OBJECT RECOGNITION IN THE HUMAN BRAIN: FROM VISUAL FEATURES TO CATEGORICAL DECISIONS .....</b>	<b>11</b>
1.1. ABSTRACT .....	13
1.2. INTRODUCTION.....	15
1.3. DECODING THE OBJECT RECOGNITION TIME-COURSE.....	18
1.3.1. <i>Initial Insights into Time-Series Decoding</i> .....	18
1.3.2. <i>Temporal Decoding of High-Dimensional Distributed Category Representations</i> .....	20
1.3.3. <i>Uncovering the Hierarchy of Object Category Representations</i> .....	29
1.3.4. <i>Decoding the Time-Course of Face Processing</i> .....	34
1.3.5. <i>Is a Category a Category?</i> .....	36
1.4. LINKING DYNAMIC OBJECT CATEGORY REPRESENTATIONS TO HUMAN BEHAVIOUR.....	42
1.4.1. <i>Categorisation Reaction-Times</i> .....	42
1.4.2. <i>Attention</i> .....	44
1.5. SUMMARY AND CONCLUSIONS .....	46
1.6. REFERENCES.....	48
III. INTERIM SUMMARY AND CHAPTER OUTLINE.....	59
IV. REFERENCES .....	63

<b>CHAPTER TWO: EVALUATING THE TEMPORAL DYNAMICS OF OBJECT REPRESENTATIONS AS A FUNCTION OF ANIMACY AND REAL-WORLD SIZE .....</b>	<b>69</b>
2.1. ABSTRACT .....	71
2.2. INTRODUCTION .....	73
2.3. METHOD .....	77
2.3.1. <i>Participants</i> .....	77
2.3.2. <i>Stimuli</i> .....	78
2.3.3. <i>Experimental Design and Attention Task</i> .....	79
2.3.4. <i>Display Apparatus</i> .....	80
2.3.5. <i>MEG Data Acquisition and Analysis</i> .....	81
2.4. RESULTS .....	84
2.4.1. <i>Model Testing Across Time</i> .....	87
2.5. DISCUSSION .....	93
2.6. REFERENCES .....	99

<b>CHAPTER THREE: NEURAL CODING OF VISUAL OBJECTS: NEW INSIGHTS INTO CATEGORICAL REPRESENTATIONS.....</b>	<b>107</b>
3.1. ABSTRACT .....	109
3.2. INTRODUCTION .....	111
3.3. METHOD .....	113
3.3.1. <i>Participants</i> .....	113
3.3.2. <i>Stimuli</i> .....	113
3.3.3. <i>MEG Experimental Procedure</i> .....	114
3.3.4. <i>Attention Task</i> .....	115
3.3.5. <i>Display Apparatus</i> .....	115
3.3.6. <i>MEG Data Acquisition and Analysis</i> .....	116
3.4. RESULTS .....	126
3.4.1. <i>Model Testing Across Time</i> .....	131
3.4.2. <i>Model Testing: Early Versus Late Time Windows</i> .....	133
3.5. DISCUSSION .....	137

3.6. REFERENCES.....	143
<b>CHAPTER FOUR: REACTION TIMES PREDICT DYNAMIC BRAIN</b>	
<b>REPRESENTATIONS MEASURED WITH MEG FOR ONLY SOME OBJECT</b>	
<b>CATEGORISATION TASKS .....</b>	<b>151</b>
4.1. ABSTRACT .....	153
4.2. INTRODUCTION.....	155
4.3. METHOD.....	159
4.3.1. <i>Stimuli</i> .....	159
4.3.2. <i>Behavioural Experiment and Reaction-Time Data</i> .....	159
4.3.3. <i>MEG Experiment</i> .....	161
4.3.4. <i>MEG Data Acquisition and Processing</i> .....	162
4.3.5. <i>Classification Analysis</i> .....	163
4.3.6. <i>Correlating Distances in Activation Space with Reaction-Times</i> .....	165
4.4. RESULTS .....	166
4.4.1. <i>Behavioural Task Performance</i> .....	166
4.4.2. <i>Exemplar Decoding from MEG Time-Series Data</i> .....	170
4.4.3. <i>Overall Representational Distance and Reaction-Time Correlations</i> .....	172
4.4.4. <i>Evaluating Representational Distance as a Predictor of Categorisation RTs</i> <i>Throughout the Decoding Time-Course</i> .....	175
4.5. DISCUSSION.....	183
4.6. REFERENCES.....	191
4.7. APPENDIX.....	197
<b>CHAPTER FIVE GENERAL DISCUSSION .....</b>	<b>199</b>
5. GENERAL DISCUSSION .....	201
5.1. <i>Results summary</i> .....	201
5.2. <i>Insights into the temporal dynamics of object representations</i> .....	203
5.3. <i>Is a Category Truly a Category?</i> .....	207
5.4. <i>Addressing the Connection Between Brain and Behaviour</i> .....	213
5.5. CONCLUSION .....	216

5.6. REFERENCES .....	217
<b>APPENDIX ONE: ETHICAL APPROVALS.....</b>	<b>225</b>
<b>APPENDIX TWO: PUBLISHED MANUSCRIPT .....</b>	<b>254</b>

## SUMMARY

---

Visual object recognition is a complex problem, with much still to be discovered about *how* the visual system achieves this task. Several studies have examined the emergence of object category structure, focusing particularly on animacy as an overarching principle of the neural organisation of object representations. Results from fMRI studies have highlighted additional organisational principles for category structure, such as real-world size, and biological class, however the temporal dynamics of these category organisations are yet to be established. The aim of this thesis is to build upon our understanding of visual object recognition, with a specific focus on evaluating the temporal dynamics of object category structure as measured with MEG. Using representational similarity analysis applied to MEG data, the first empirical chapter compares the temporal dynamics of animacy and real-world size dimensions of object representations. The results replicate previous findings for the animacy time-course, however there was no evidence for a distinct time-course associated with real-world size. The second empirical chapter examines alternatives to the animacy category organisation of object representations, using a novel stimulus set that includes objects which do not clearly belong to the typically evaluated ‘animate’ or ‘inanimate’ categories (e.g., robots and human-/animal-like toys). This study evaluates a range of models based on current theories of object categorisation including animacy, and the biological classes based ‘animacy continuum’, as well as novel behaviourally-generated models related to human-similarity and experience. Results show that the model of human-similarity is the best predictor of object representations late in the time-course of visual object processing. The aim of the third empirical chapter is to link these human-similarity results from the MEG data to behaviour. This study shows that object categorisation reaction times predict representational distance not only for object animacy (as shown in previous studies), but also when objects are grouped according to human-similarity. In contrast, other plausible object category organisations for the same stimulus set (i.e., living/non-living; has movement/no

movement) do not show the same relationship between brain activation patterns and behaviour. To conclude, the findings from these three studies are discussed within the broader context of the current literature related to object representations in the human brain. This thesis highlights the efficacy of a new human-similarity model of object category representations and critically evaluates what aspects of decodable neural representations are informative for understanding the link between brain and behaviour.

## STATEMENT OF AUTHENTICATION

---

This thesis is submitted in partial fulfilment of the requirements for the degree of Combined Doctor of Philosophy / Master of Clinical Neuropsychology.

I declare that the research presented in this thesis is my original work. This work has not previously been submitted for a degree or diploma in any university. I have appropriately acknowledged help or assistance that I have received during the preparation of this thesis, as well as any sources of information that I have used.

All research presented in this thesis was approved by Macquarie University Ethics Review Committee (Human Research), reference number: 5201300804. Documentation of this approval is given in Appendix One.

Signed:

Date: 29<sup>th</sup> March, 2018

Erika W. Contini

(Student Number: 41994930)





## ACKNOWLEDGEMENTS

---

Firstly, I would like to express my sincere thanks to my supervisors who assisted me throughout this process. I feel honoured to have had the pleasure of working with such wonderful scientists. To Mark Williams, for his ongoing guidance and expertise throughout this degree. Your kind support and encouragement has been greatly appreciated. To Thomas Carlson, for introducing me to the wonders of MEG, and for his enthusiasm and knowledge that assisted me with the formulation of my studies. To Erin Goddard and Susan Wardle, for their endless guidance and support regardless of the time of day, as well as their friendship and humour, which truly made all the difference. Sincere thanks to Erin for her MATLAB prowess and continued patience with helping me navigate the intricate world of coding. Special thanks to Susan, whose advice and contribution to the completion of this thesis has been invaluable. Susan, thank you for the time and energy you have so generously given me over the past few years. I will remember our meetings/life debriefs with great fondness.

I would like to thank the many members of the Department of Cognitive Science. To the Perception in Action lab, for their enthusiasm for research and support for students. To the kind people who participated in my MEG studies, and Liz Stylianou for her diligent management of the MEG lab. To Tijl Grootswagers for his assistance with data collection, hallway brainstorming and helpful advice. To the operations team, who keep the department running like a well-oiled machine. Special mention to Lesley McKnight, Marcus Ockenden, Craig Richardson, Katie Webb and Lisa Yen for their dedication to making life that little bit easier when navigating the various administrative and technical aspects of completing a PhD. To the many friendly colleagues and fellow students: Kiley, Mariia, Tijl, Lina, Simmy, Nate, Ann, Kate, Ben, Valerie, Regine, Emily, Sarah, Rebecca, Marina, Josh, Carmen, my work colleagues from the MQ Reading Clinic, - the list goes on, it has been a pleasure to see your friendly faces around the halls, and I treasure the chats and cheers that we have exchanged over the years.

To my fabulous friends. Thank you for bringing joy to my days in the most wonderful and heart-warming ways. To the Neuro and Uni crew - Nikki, Jody, Megan, Louise, Flick, Monica, Jilly, Gen, Heather, Soheil, Thushara, Huachen, and their partners, to my fabulous SASSers – Amanda and Jenna, and the mountains crew, thank you all for brightening my days, for understanding my meltdowns, for your generosity of coffee and food, and for the many hugs and kind words that have meant so much to me. I don't know if I can bake enough cake to sufficiently thank you all but never say never.

Last but by no means least, I am immensely grateful to my family. To my sister Alexis, and her husband Robert – thank you for your continued support and light-hearted encouragement that regularly put a smile on my face. I am so lucky to have you both in my life. To my wonderful and supportive parents, Tessa and Peter – I am forever in your debt (possibly literally) for everything you have done for me to see me through this degree. For the many cooked meals, the assistance with moving houses, for listening to me rave and rant, through the good times and the more trying. Thank you for your continued encouragement, faith, and positivity throughout this process. I am eternally grateful to you both.

Thank you.

## AUTHOR NOTE

---

This thesis has been prepared in the format of 'Thesis by Publication'. The reference style and formatting largely conform to standards laid out by the APA Publication Manual (6<sup>th</sup> edition). Due to the 'Thesis by Publication' format, there is some degree of repetition between the chapters. I have tried to avoid repetition as much as possible while still allowing each chapter to stand in isolation. My role in this body of research included the design of all three empirical studies, collection of all MEG and behavioural data, analysis and interpretation of the data, and writing of all five chapters of this thesis. Assistance from others with aspects of this research has been acknowledged where appropriate at the end of each manuscript. I am the first author of the published review paper, and three empirical studies presented in this thesis.



## CHAPTER ONE

---

---

### GENERAL INTRODUCTION

---

A section of this chapter has been published in *Neuropsychologia* as:

Contini E.W., Wardle S.G., Carlson T.A. (2017). Decoding the time-course of object recognition in the human brain: From visual features to categorical decisions. *Neuropsychologia*, 105, 165–176.



## **i. Introduction to Thesis Topic**

The aim of this thesis is to investigate the temporal dynamics of object recognition, with a particular focus on object category representations. Recognising the numerous objects we encounter in our daily existence is a fundamental part of human visual processing. Decades of research have been dedicated to better understanding this seemingly effortless process, with much still to be discovered about *how* the visual system achieves this complex task. While artificial intelligence systems have recently advanced to such a degree that they are able to recognise objects at a level similar to human performance (He, Zhang, Ren, & Sun, 2015), establishing the correspondence with human visual processing is ongoing (e.g., see Cadieu et al., 2014; Cichy, Khosla, Pantazis, Torralba, & Oliva, 2016).

Early studies of object recognition deficits in human patients with localised lesions laid the groundwork for much of the neuroimaging research in the field of object categorisation to date (Capitani, Laiacona, Mahon, & Caramazza, 2003; Warrington & McCarthy, 1983, 1987, 1994; Warrington & Shallice, 1984). Evaluation of patients' recognition abilities revealed an interesting pattern of deficits related to specific categories of objects. Deficits with recognising animate/living (Sheridan & Humphreys, 1993; Silveri & Gainotti, 1988; Warrington & Shallice, 1984) or inanimate/non-living objects (Sacchett & Humphreys, 1992; Warrington & McCarthy, 1983, 1987) suggested a dichotomous organisation of object representations in the brain. These findings shaped our current understanding of how the brain processes objects.

A prominent goal in visual neuroscience is to use brain activity to evaluate the categorical structure and processes underlying object representations. Neuroimaging allows us to ask these questions about object recognition noninvasively in healthy adults. Researchers in the field of object recognition have employed multivariate pattern analysis (MVPA) methods to study the brain's representation of objects, identifying patterns that are consistent across human subjects (Haxby et al., 2001) and species (Kiani, Esteky, Mirpour,

& Tanaka, 2007; Kriegeskorte et al., 2008). These consistencies suggest a common underlying organisation to object representations (Op de Beeck, Torfs, & Wagemans, 2008). However, much of the literature to date has focused on spatially localising brain regions which respond preferentially to particular object categories (e.g., Beauchamp, Lee, Haxby, & Martin, 2002; Chao, Haxby, & Martin, 1999; Downing, Jiang, Shuman, & Kanwisher, 2001; Epstein & Kanwisher, 1998; Ishai, Ungerleider, Martin, Schouten, & Haxby, 1999; Kanwisher, McDermott, & Chun, 1997; Martin, Wiggs, Ungerleider, & Haxby, 1996), with more to discover about how these representations unfold over time. There is a growing body of evidence from recent studies using time-series decoding methods, showing that object recognition processes evolve dynamically following exposure to a stimulus (for a review, see Contini, Wardle, & Carlson, 2017 in this chapter below). These recently established temporal decoding methods reveal details about object processing that may be lost in spatially focused neuroimaging methods (i.e., fMRI) which have low temporal resolution. Evaluating the temporal dynamics of object representations is thus an important complementary direction for research into visual object processing.

Here I begin by reviewing the literature related to time-series decoding of object representations in the brain. This review, published in *Neuropsychologia* (Contini et al., 2017), provides an overview of how investigating the time-course of object representations has added to our understanding of visual object processing. I summarise the current findings related to the hierarchical category structure of object representations, including a focused section for the special category of faces, which are shown to be of particular importance to the human visual system. Next, I evaluate the evidence surrounding the integrity of object category structure in the context of research that highlights the potentially confounding influence of low-level image properties on measured brain activity. Following this review, I describe recent studies that use categorisation reaction times and attention modulations to link brain-derived object representations to behavioural processes, with the goal of



advancing our understanding of how measured activity from neuroimaging techniques relates functionally to human behaviour.

While the scope of the review covers the broad field of object decoding related to time-series data, my thesis focuses on the time-course of object category structure. Following this review, I provide a brief summary of our current understanding of the temporal dynamics of object category representations and outline the contributions of this thesis.

## ii. References

- Beauchamp, M. S., Lee, K. E., Haxby, J. V., & Martin, A. (2002). Parallel visual motion processing streams for manipulable objects and human movements. *Neuron*, 34(1), 149–159.
- Cadieu, C. F., Hong, H., Yamins, D. L. K., Pinto, N., Ardila, D., Solomon, E. A., ... DiCarlo, J. J. (2014). Deep Neural Networks Rival the Representation of Primate IT Cortex for Core Visual Object Recognition. *PLOS Computational Biology*, 10(12), e1003963. <https://doi.org/10.1371/journal.pcbi.1003963>
- Capitani, E., Laiacona, M., Mahon, B., & Caramazza, A. (2003). What Are the Facts of Semantic Category-Specific Deficits? A Critical Review of the Clinical Evidence. *Cognitive Neuropsychology*, 20(3–6), 213–261. <https://doi.org/10.1080/02643290244000266>
- Chao, L. L., Haxby, J. V., & Martin, A. (1999). Attribute-based neural substrates in temporal cortex for perceiving and knowing about objects. *Nature Neuroscience*, 2(10), 913–919.
- Cichy, R. M., Khosla, A., Pantazis, D., Torralba, A., & Oliva, A. (2016). Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence. *Scientific Reports*, 6. <https://doi.org/10.1038/srep27755>
- Contini, E. W., Wardle, S. G., & Carlson, T. A. (2017). Decoding the time-course of object recognition in the human brain: From visual features to categorical decisions. *Neuropsychologia*, 105(Supplement C), 165–176. <https://doi.org/10.1016/j.neuropsychologia.2017.02.013>

- Downing, P., Jiang, Y., Shuman, M., & Kanwisher, N. (2001). A cortical area selective for visual processing of the human body. *Science*, 293(5539), 2470–2473.
- Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature*, 392(6676), 598–601.
- Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and Overlapping Representations of Faces and Objects in Ventral Temporal Cortex. *Science*, 293(5539), 2425–2430.  
<https://doi.org/10.1126/science.1063736>
- He, K., Zhang, X., Ren, S., & Sun, J. (2015). Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification. In *Computer Vision (ICCV), 2015 IEEE International Conference on* (pp. 1026–1034). IEEE. Retrieved from <http://ieeexplore.ieee.org/abstract/document/7410480/>
- Ishai, A., Ungerleider, L. G., Martin, A., Schouten, J. L., & Haxby, J. V. (1999). Distributed representation of objects in the human ventral visual pathway. *Proceedings of the National Academy of Sciences*, 96(16), 9379–9384.  
<https://doi.org/10.1073/pnas.96.16.9379>
- Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The Fusiform Face Area: A Module in Human Extrastriate Cortex Specialized for Face Perception. *Journal of Neuroscience*, 17(11), 4302–4311.
- Kiani, R., Esteky, H., Mirpour, K., & Tanaka, K. (2007). Object Category Structure in Response Patterns of Neuronal Population in Monkey Inferior Temporal Cortex. *Journal of Neurophysiology*, 97(6), 4296–4309.  
<https://doi.org/10.1152/jn.00024.2007>

- Kriegeskorte, N., Mur, M., Ruff, D. A., Kiani, R., Bodurka, J., Esteky, H., ... Bandettini, P. A. (2008). Matching Categorical Object Representations in Inferior Temporal Cortex of Man and Monkey. *Neuron*, 60(6), 1126–1141.  
<https://doi.org/10.1016/j.neuron.2008.10.043>
- Martin, A., Wiggs, C. L., Ungerleider, L. G., & Haxby, J. V. (1996). Neural correlates of category-specific knowledge. Retrieved from  
<http://www.nature.com/nature/journal/v379/n6566/abs/379649a0.html>
- Op de Beeck, H. P., Torfs, K., & Wagemans, J. (2008). Perceived Shape Similarity among Unfamiliar Objects and the Organization of the Human Object Vision Pathway. *Journal of Neuroscience*, 28(40), 10111–10123.  
<https://doi.org/10.1523/JNEUROSCI.2511-08.2008>
- Sacchett, C., & Humphreys, G. W. (1992). Calling a squirrel a squirrel but a canoe a wigwam: a category-specific deficit for artefactual objects and body parts. *Cognitive Neuropsychology*, 9(1), 73–86.  
<https://doi.org/10.1080/02643299208252053>
- Sheridan, J., & Humphreys, G. W. (1993). A verbal-semantic category-specific recognition impairment. *Cognitive Neuropsychology*, 10(2), 143–184.
- Silveri, M. C., & Gainotti, G. (1988). Interaction between vision and language in category-specific semantic impairment. *Cognitive Neuropsychology*, 5(6), 677–709.
- Warrington, E. K., & McCarthy, R. (1983). Category specific access dysphasia. *Brain*, 106(4), 859–878.
- Warrington, E. K., & McCarthy, R. A. (1987). Categories of knowledge: Further fractionations and an attempted integration. *Brain*, 110(5), 1273–1296.

Warrington, E. K., & McCarthy, R. A. (1994). Multiple meaning systems in the brain: A case for visual semantics. *Neuropsychologia*, 32(12), 1465–1473.

[https://doi.org/10.1016/0028-3932\(94\)90118-X](https://doi.org/10.1016/0028-3932(94)90118-X)

Warrington, E. K., & Shallice, T. (1984). Category specific semantic impairments. *Brain*, 107(3), 829–853.



# 1. LITERATURE REVIEW

---

## DECODING THE TIME-COURSE OF OBJECT RECOGNITION IN THE HUMAN BRAIN: FROM VISUAL FEATURES TO CATEGORICAL DECISIONS

---

Erika W. Contini<sup>1,2</sup>, Susan G. Wardle<sup>1,2\*</sup> & Thomas A. Carlson<sup>1,2,3\*</sup>

<sup>1</sup>Department of Cognitive Science, Macquarie University, Sydney, Australia

<sup>2</sup>ARC Centre of Excellence in Cognition and its Disorders and Perception in Action Research Centre,  
Macquarie University, Australia

<sup>3</sup>School of Psychology, University of Sydney, Australia

\* shared senior authorship

Corresponding author: Erika Contini [erika.contini@students.mq.edu.au](mailto:erika.contini@students.mq.edu.au)

This section has been published as:

Contini E.W., Wardle S.G., Carlson T.A. (2017). Decoding the time-course of object recognition in the human brain: From visual features to categorical decisions. *Neuropsychologia*, 105, 165–176. <https://doi.org/10.1016/j.neuropsychologia.2017.02.013>





## **1.1. Abstract**

Visual object recognition is a complex, dynamic process. Multivariate pattern analysis methods, such as decoding, have begun to reveal how the brain processes complex visual information. Recently, temporal decoding methods for EEG and MEG have offered the potential to evaluate the temporal dynamics of object recognition. Here we review the contribution of M/EEG time-series decoding methods to understanding visual object recognition in the human brain. Consistent with the current understanding of the visual processing hierarchy, low-level visual features dominate decodable object representations early in the time-course, with more abstract representations related to object category emerging later. A key finding is that the time-course of object processing is highly dynamic and rapidly evolving, with limited temporal generalisation of decodable information. Several studies have examined the emergence of object category structure, and we consider to what degree category decoding can be explained by sensitivity to low-level visual features. Finally, we evaluate recent work attempting to link human behaviour to the neural time-course of object processing.



## 1.2. Introduction

Visual object recognition is a complex problem. In everyday life, we experience an overwhelming number of objects that the brain needs to rapidly differentiate and identify. How is it that we are able to identify a chair, for example, despite large variability in lighting, colour, design, materials and viewpoint? The human brain does a remarkable job of efficiently solving this problem, and has inspired decades of behavioural, neuroscience, and computer science research. The ventral visual processing stream, which involves a number of regions throughout the occipito-temporal cortex, is well-established as the neural pathway for object recognition (e.g., Grill-Spector, Kourtzi, & Kanwisher, 2001; Haxby et al., 2001; Ishai, Ungerleider, Martin, Schouten, & Haxby, 1999). Functional neuroimaging methods have identified a number of object-selective areas in human visual cortex that are preferentially activated by specific object categories including animals (Chao, Haxby, & Martin, 1999; Martin, Wiggs, Ungerleider, & Haxby, 1996), faces (Ishai et al., 1999; Kanwisher, McDermott, & Chun, 1997), bodies (Downing, Jiang, Shuman, & Kanwisher, 2001), places (Epstein & Kanwisher, 1998), and tools (Beauchamp, Lee, Haxby, & Martin, 2002; Chao et al., 1999; Martin et al., 1996). However, the number of object-selective regions identified to date accounts for only a handful of the numerous object categories that exist (Biederman, 1987). This highlights that while identification of these regions has advanced our understanding of the neural pathways involved in higher-level visual processing, we are yet to fully understand *how* the brain solves the many challenges associated with object recognition.

Recent studies in visual object recognition have been influenced by advances in neuroimaging analysis methods, which allow for a fundamental change in the type of information that can be extracted from neuroimaging data. Early neuroimaging studies employed univariate analysis techniques that focus on identifying differences in the average activation of individual voxels or sensors for different experimental conditions. In contrast,

more recent multivariate analysis techniques (e.g. multivariate pattern analysis; MVPA) analyse patterns of activation associated with experimental conditions from multiple voxels/sensors simultaneously. In the field of neuroimaging, multivariate methods have the potential to detect differences in activation which are lost when averaging data for univariate analyses, making them more sensitive (Carlson, Schrater, & He, 2003; Cox & Savoy, 2003; Grootswagers, Wardle, & Carlson, 2016; Haxby et al., 2001; Haynes & Rees, 2006). While univariate and multivariate analyses are complementary in the information they provide, there is a fundamental difference in the types of experimental questions these methods address. Univariate magneto/electro-encephalography (M/EEG) analyses evaluate differences in *activation*, quantifying relative differences in average activity between experimental conditions, while multivariate methods have the potential to examine differences in *information*, for example by comparing differences in distributed patterns of brain activation between experimental conditions (Grootswagers et al., 2016; but see also de-Wit, Alexander, Ekroll, & Wagemans, 2016).

Decoding methods are a form of MVPA that have provided important insights into how the brain processes information (Grootswagers et al., 2016; Haxby, Connolly, & Guntupalli, 2014; Haynes, 2015; Pereira, Mitchell, & Botvinick, 2009). In the field of object recognition, decoding methods aim to map differences in complex neural activity patterns associated with perceiving objects. Early studies adopting these methods marked an important first step in our understanding of *how* objects are processed in the brain (Carlson et al., 2003; Carlson, Hogendoorn, Kanai, Mesik, & Turret, 2011; Alex Clarke & Tyler, 2014; Cox & Savoy, 2003; Haxby et al., 2001; Kriegeskorte, Mur, Ruff, et al., 2008; Liu, Agam, Madsen, & Kreiman, 2009; O'Toole, Jiang, Abdi, & Haxby, 2005). A particular strength of applying decoding methods to M/EEG data with the aim of understanding visual object recognition is that it can reveal how visual object representations change over time with high temporal resolution. The focus of this review is on what has been learned about

visual object processing in the human brain with the application of these recent, powerful temporal decoding methods.

To date, decoding methods have been used in conjunction with a variety of neuroimaging and neurophysiology approaches to investigate the neural mechanisms underlying object recognition. For example, brain decoding using single cell recordings has revealed category structure within monkey inferior temporal cortex (IT) (Hung, Kreiman, Poggio, & DiCarlo, 2005), with differentiable neural patterns associated with animate and inanimate objects, as well as more specific animate subcategories, including human and animal faces and bodies (Kiani, Esteky, Mirpour, & Tanaka, 2007). Decoding of fMRI data has shown commonalities between object response patterns in both human and monkey brains, with similar differentiation of animate/inanimate categories and face/body subcategories identified in both species (Kriegeskorte, Mur, Ruff, et al., 2008). Categorical representations are not limited to localised brain regions, as activity evoked by objects and faces often overlaps between the categories and is distributed throughout ventral temporal cortex (Haxby et al., 2001). Moreover, there has been a recent move towards looking beyond representing object categories in terms of dichotomies, such as the animate/inanimate distinction. Instead, a continuum has been proposed that spans from inanimate objects to humans, with objects categorised in terms of their biological similarity to humans (Connolly et al., 2012; Sha et al., 2015). Together, these studies emphasise a role for distributed patterns of activity in the neural representation of object category, building on the original observation of overlapping activation patterns in the human ventral stream (Haxby et al., 2001).

Decoding of time-series data using M/EEG has offered the potential to examine the time-course of object representations in the human brain, revealing a dynamic evolution of object category structure over time (e.g., Barragan-Jason, Cauchoux, & Barbeau, 2015; Carlson et al., 2011; Carlson, Tovar, Alink, & Kriegeskorte, 2013; Cauchoux, Barragan-

Jason, Serre, & Barbeau, 2014; Cichy, Pantazis, & Oliva, 2014; Clarke, Devereux, Randall, & Tyler, 2014; Goddard, Carlson, Dermody, & Woolgar, 2016; Kaiser, Azzalini, & Peelen, 2016; Simanova, Gerven, Oostenveld, & Hagoort, 2010). As visual information moves through the ventral pathway, the content of visual representations changes rapidly between brain regions. fMRI decoding studies have revealed much about the representation of objects in the human brain, however the coarse temporal resolution of fMRI limits the examination of dynamic visual processes. As a complement to the static snapshot of representational structure revealed with fMRI, neuroimaging techniques with higher temporal resolution such as M/EEG facilitate investigation of the dynamic processes of visual object recognition. The aim of this review is to discuss how time-series decoding studies have advanced our understanding of the complexities of visual object recognition by focusing on the dynamic processes involved. Time-series decoding is a relatively new approach to studying object recognition, and here we highlight the potential of this new direction to inform the field. Note that we focus on what has been learned about object processing by applying time-series decoding methods, for a more thorough discussion of the technical details of time-series decoding analyses see Grootswagers et al. (2016).

### **1.3. Decoding the Object Recognition Time-Course**

#### **1.3.1. Initial Insights into Time-Series Decoding**

It is well-established that the process of visual object recognition requires a number of hierarchically organised stages that progress through the occipito-temporal pathway (Grill-Spector & Malach, 2004; Malach, Levy, & Hasson, 2002). Early retinotopic visual areas are more sensitive to changes in low-level stimulus properties, while higher cortical areas within the ventral temporal lobe produce more complex responses to whole objects, and appear to account for more abstract properties such as object category (Altmann, Bühlhoff, & Kourtzi, 2003; Grill-Spector & Malach, 2004; Van Essen, Anderson, & Felleman, 1992). Single-unit recordings in macaques have shown that these different

processing stages progress successively in time as information passes through the occipito-temporal pathway (Schmolsky et al., 1998). With the development of M/EEG techniques, we are able to evaluate the temporal dynamics of visual object recognition in the human brain with millisecond resolution, allowing us to delve into more specific and fine-grained processes occurring in the various stages of visual object processing.

Differences in early versus late stages in object processing have been examined by comparing MEG and fMRI data for the same stimulus set. By linking both temporal and spatial neuroimaging data, Cichy et al. (2014) showed that activity early in the MEG time-course correlated more strongly with fMRI activity in V1, while later MEG activity was more strongly associated with activity in IT. The stimulus set were 96 colour images of animate and inanimate objects used in previous studies (Kiani et al., 2007; Kriegeskorte, Mur, Ruff, et al., 2008). The activity patterns associated with viewing each individual object were first compared using representational similarity analysis (RSA; Kriegeskorte, 2008), where a matrix is created based on the difference in brain activation patterns for every pairwise comparison of object images. This was done separately for the fMRI and MEG data. These 'dissimilarity matrices' provide an index of the difference in the brain response between object representations. The dissimilarity matrices were then compared across imaging modalities by examining when the relative similarity between the activation patterns for each object pair in the fMRI data most closely resembled that in the MEG data. The finding that the representational structure early in the MEG data more closely resembles V1 activity while later MEG data is closer to the structure observed in IT with fMRI is consistent with the known features of the visual processing hierarchy, and thus provides a source of validation for MEG decoding methods.

Time-series decoding methods have also expanded our understanding of the temporal intricacies associated with processing low-level stimulus properties. For example, Goddard et al. (2016) investigated temporal differences in the contribution of low and high spatial

frequencies to the representation of object identity. Stimuli were greyscale images of objects that were matched in their amplitude spectrum to control for low-level visual properties. Phase randomisation was applied selectively to spatial frequency bands such that object identity information was confined to a restricted spatial frequency band in each stimulus. Decoding analysis revealed that the activation pattern of the whole-brain MEG recordings contained information related to object identity, but there were critical differences in the processing of low versus high spatial frequencies. They found that low spatial frequencies provided object identity information earlier and in more occipitally located regions than high spatial frequency information, which supported decoding of object identity later in the time-course. Further, by applying Granger causality analysis, they found evidence for both an early feedforward and later feedback flow of information related to object identity (Goddard et al., 2016).

Most of the existing temporal decoding literature on object recognition has focused on decoding object categories and investigating the representational structure of object representations. These results are discussed in detail in the following sections.

### **1.3.2. Temporal Decoding of High-Dimensional Distributed Category Representations**


Understanding population coding is one of the overarching aims of neuroscience (Averbeck, Latham, & Pouget, 2006). A significant and influential recent development in understanding the neural mechanisms underlying object representation has been the theoretical shift away from identifying specific localised regions associated with particular object categories towards focusing on more distributed activity patterns that extend throughout higher-level visual cortex (Haxby et al., 2001). A current focus in object recognition is on interpreting the highly multidimensional activation patterns elicited by object stimuli. Specifically, one of the main themes centres on understanding the structure of object category representations. In this section we outline what temporal decoding methods have so far revealed about the neural architecture of object categories.



A useful and intuitive way to think about the highly multidimensional activation patterns evoked by viewing objects is to conceptualise them in terms of an abstract representational space (DiCarlo & Cox, 2007; Kriegeskorte & Kievit, 2013). The complex activation pattern across M/EEG sensors elicited by viewing a particular object exemplar (e.g., shoe, horse, face) can be considered as a single point in an abstract representational space. Thus in this abstract representation, the proximity between data points indicates the degree of similarity, such that a greater distance between object exemplars represents more disparate neural activation patterns, while exemplars with closer points have more similar activation patterns. Consequently, object exemplars further away from each other in representational space are easier to "decode" from each other using machine learning classification than objects which are closer together, as they share more similar activation patterns and are less separable in higher-dimensional space.

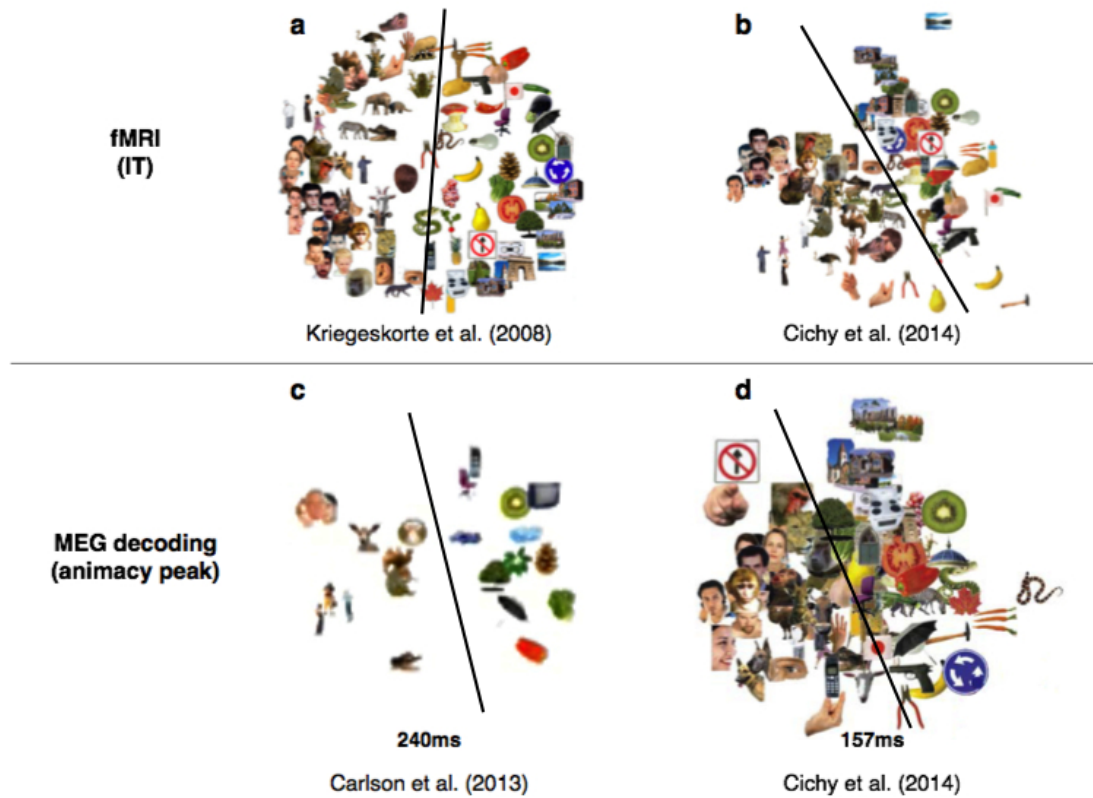
Multidimensional scaling (MDS) is a technique that can be used as a tool for visualising the representational space of objects in the brain. An example of an MDS plot for object representations is shown in Figure 1. Carlson et al. (2013) studied the first 1000 ms of the visual object recognition time course using MEG decoding. MEG recordings were acquired as participants viewed a series of single object images while performing an unrelated attention task (reporting whether a letter superimposed on top of the object image was a vowel or a consonant). Using MDS to visualise the differences in object representations as a function of time, we can appreciate the emergence of category structure throughout the time course (Figure 1). Early in the time-course (~60 ms post-stimulus onset), the representations of individual exemplars are almost wholly overlapping in the representational space, reflecting poor decodability of the individual objects based on the MEG data. This is expected given the time it takes for an image on the retina to transition to a cortical representation that is accessible to MEG (see Nowak & Bullier, 1997 for a review). By ~120 ms, differences between individual exemplars have emerged such that the exemplars are spread out in the representational space. This is the time of peak decoding for

this image set; the point at which individual exemplars are most easily distinguished from one another. As we progress through the time course, category structure begins to emerge. From 120 ms, some subcategories appear to start to cluster: note the grouping of faces and animals in the centre. By 240 ms a clear categorical distinction between animate and inanimate objects is apparent (diagonal black lines in Figure 1 indicate the category boundary). Interestingly, this animacy category distinction is maintained at 360 ms, despite much less differentiation of individual object exemplars, which are clustered tightly together in the later stages of the time-course. This is an example of the advantage of time-series decoding, as the emergence of categorical clustering over time is not captured by other neuroimaging methods with lower temporal resolution such as fMRI.



*Figure 1.* MDS plots demonstrating the evolution of object representations over time with MEG. Individual MDS panels show the representational geometry of a set of objects at a particular time-point (shown as ms post-stimulus onset). Distances between objects represent the level of similarity between the neural patterns measured with MEG such that larger distances indicate greater dissimilarity (i.e., more distinct neural patterns). The representational geometry evolves from initially entirely overlapping (poorly discriminated) objects, to maximal differentiation at 120 ms (peak decoding for this sample). The animate/inanimate category distinction (black dividing line represents the boundary) peaks at 240 ms and is maintained even at 360 ms, although differences between individual object exemplars become less defined. Images adapted with permission of the Association for Research in Vision and Ophthalmology, from Carlson et al. (2013); permission conveyed through Copyright Clearance Center, Inc.

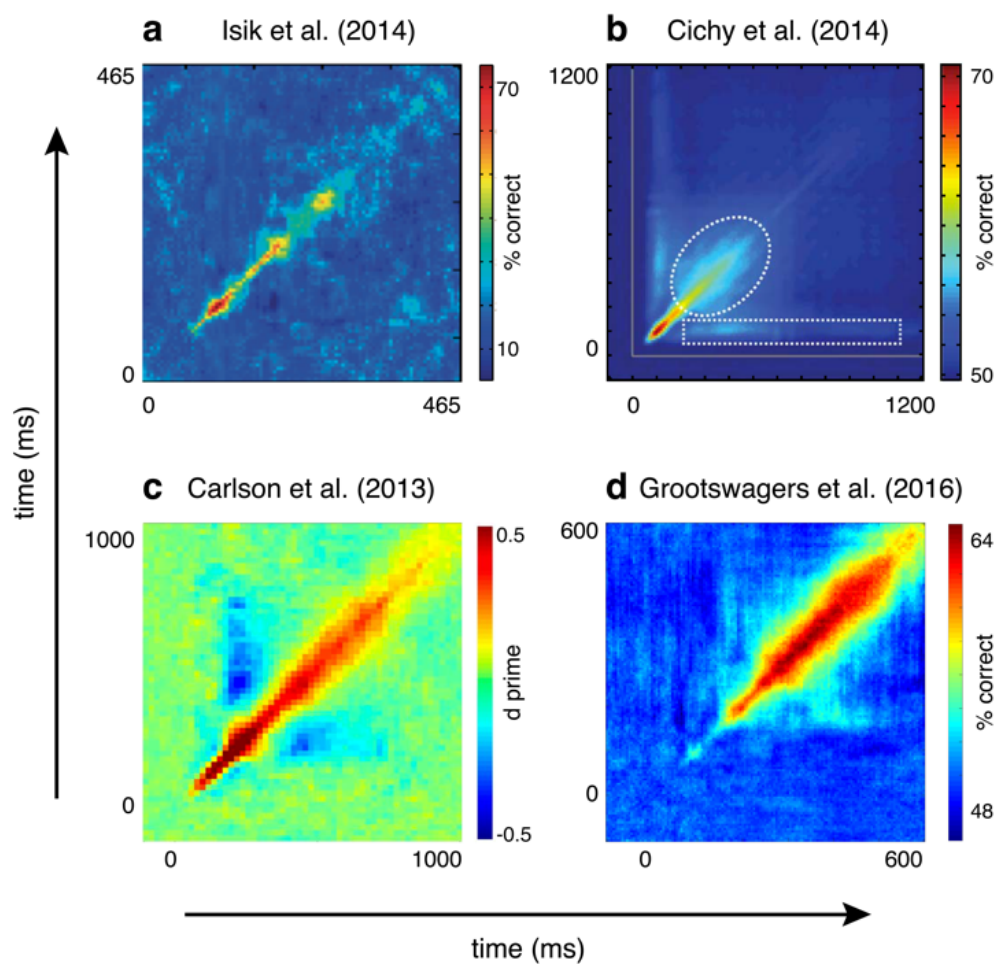
As MDS has been used in several fMRI and MEG studies with the same object stimuli (Carlson et al., 2013; Cichy et al., 2014; Kriegeskorte, Mur, Ruff, et al., 2008) it is possible to visualise the reliability and consistency of object representations across neuroimaging modalities by comparing across studies - a rare opportunity in neuroimaging (Figure 2). All three studies used the object set (or subset of the set) originally used by Kiani et al. (2007) in monkey IT. Kriegeskorte et al. (2008) and Cichy et al. (2014) both used fMRI to evaluate the representational geometry of objects in human IT (Figure 2a and b). The animate/inanimate divide is clearly visible in IT, and additional subcategory groupings such as animals and humans also cluster together. These fMRI MDS plots can be compared to those from MEG time-series decoding studies (Figure 2 c and d; Carlson et al., 2013; Cichy et al., 2014). MEG data is shown at the time of peak animacy decoding in each study: 240 ms (Figure 2c) and 157 ms (Figure 2d). Note that the time of peak decoding is dependent on the stimulus set being evaluated as well as the noise and variability in a particular data set, complicating the comparison of specific decoding onset times across studies. The emergence of feature information or category structure is better discussed in terms of relative timing within individual studies. A critical difference between the MDS plots for fMRI versus MEG is that the MEG plots represent similarity in whole-brain activation patterns, whereas the fMRI analysis is localised to IT. Considering this substantial difference, the degree of similarity in the categorical representation is remarkable (compare top and bottom rows of Figure 2).



*Figure 2.* Comparison of MDS plots from fMRI and MEG studies. Top MDS plots (a and b) show the representational geometry in human IT using fMRI; bottom MDS plots (c and d) show the representational geometry using MEG at the time of peak decoding for the animate/inanimate (animacy) category distinction. Animacy peak decoding times for the individual studies are listed below the MEG MDS plots. Note that (a) has been normalised and rigidly aligned (Procrustes alignment) for visualisation, producing a more regular spacing between object exemplars in the MDS plot compared to panels b-d. Black lines in each plot mark the (approximate) animacy boundary. Black lines in each plot mark the (approximate) animacy boundary. Panel (a) image from Kriegeskorte et al. (2008), reprinted with permission from Elsevier. Panel (c) image from Carlson et al. (2013) reprinted with permission of the Association for Research in Vision and Ophthalmology; permission conveyed through Copyright Clearance Center, Inc. Panels (b) and (d) images from Cichy et al. (2014) reprinted with permission from Macmillan Publishers Ltd: Nature Neuroscience, copyright (2014).

Another advantage of using time-series decoding to investigate object processing is that we can examine to what degree information related to object category is maintained in the brain activation patterns over time (see King & Dehaene, 2014 for a review of temporal generalisation methods). This is achieved by performing time-point by time-point comparisons across the entire time-course, whereby a classifier is trained on data from one time-point and then tested at every other time point. If information is maintained within the signal for some period of time, then the classifier is expected to generalise and successfully classify the test data, regardless of the time difference between the training and test data sets. The results of temporal generalisation are visualised in two-dimensional heat maps indicating the magnitude of classification performance for each pair of time points (Figure 3).

A number of MEG studies have examined the temporal generalisation of decoding to evaluate the dynamics of object representations (Figure 3). For each plot in Figure 3, the diagonal represents standard decoding analysis when the classifier is trained and tested on data from the same time-point, thus classification accuracy is expected to be highest along the diagonal. If decoding performance is significant off the diagonal, it is indicative of some degree of temporal generalisation of the decodable signal related to object classification. Note that the decoded category differs across studies: Isik et al. (2014) and Cichy et al. (2014) decoded object exemplars, while Carlson et al. (2013) and Grootswagers et al. (2016) decoded object animacy.



*Figure 3.* Examples of temporal cross-decoding of MEG data. Temporal generalisation of MEG decoding from (a) Isik et al. (2014), (b) Cichy et al. (2014) (c) Carlson et al. (2013), and (d) Grootswagers et al. (2016). Generalisation of decoding performance across time is assessed by training and testing the classifier on MEG data from different time points. Points on the diagonal represent training and testing on data from the same time point (i.e., regular classification). Colour indicates classifier performance as a function of time for each study; hotter colours indicate higher levels of classifier performance (measured using decoding accuracy percent (a, b, d) and d-prime in (c). Panel (a) image from Isik et al. (2014) reprinted with permission from ©The American Physiological Society. Panel (b) image from Cichy et al. (2014) reprinted with permission from Macmillan Publishers Ltd: Nature Neuroscience, copyright (2014). Panel (c) image from Carlson et al. (2013) reprinted with permission of the Association for Research in Vision and Ophthalmology; permission conveyed through Copyright Clearance Center, Inc. Panel (d) image from Grootswagers et al. (2016) reprinted with permission from the Massachusetts Institute of Technology © 2016, published by the MIT Press.

The most striking similarity between the temporal generalisation plots from these different studies is that the MEG signal underlying object decoding evolves relatively quickly, with little information generalising across time. Where there is generalisation, it tends to cluster closely around the diagonal, suggesting that the structure of object representations accessible in the whole-brain MEG signal evolves rapidly, and follows a specific neural trajectory. Another interesting feature is that the period where there is the greatest generalisation occurs relatively late after stimulus onset (see particularly Figures 3c-d). This suggests maintenance of object representations at later stages of processing (Carlson et al., 2013). Alternatively, it may reflect the gradual accumulation of multiple related information processing stages of variable duration throughout the visual processing hierarchy, leading to a greater temporal spread of information over time. Paradoxically, there are also some periods of below chance decoding, whereby a classifier systematically categorises an object incorrectly as the opposite category (e.g. see blue regions in Figure 3c). Carlson et al. (2013) suggest this may represent adaptation or inhibition following a period of excitation, resulting in later inversion of the neural representation. Consistent with this explanation, anticorrelated MEG signals have also been identified around the time of stimulus offset for both visual (Carlson et al., 2011) and auditory (Chait, Poeppel, Cheveigné, & Simon, 2007) stimuli.

An exception to the general lack of generalisation is seen in the results of Cichy et al. (2014), where there is a period of significant generalisation of early information across most of the time-course (Figure 3b; grey dotted rectangle shows the period of significant temporal generalisation). This could be because this analysis used pair-wise object decoding as opposed to leave-one-exemplar-out category classification (see Carlson et al., 2013; Grootswagers et al., 2016). With pairwise object decoding, low-level information diagnostic of object is contained in both the training and test sets for the classifier, thus maintenance of early visual information is consistent with low-level properties being useful for classification. There is a similar pattern of generalisation in Figure 3c, and this analysis also



included exemplars in both training and test sets (Carlson et al., 2013). However, in leave-one-out exemplar decoding (Figure 3d) the test exemplar is not included in the training set so early processing reflecting sensitivity to low level image properties is less likely to be used by the classifier. Kaiser et al. (2016) took this idea further by using temporal generalisation to directly examine decoding of low-level shape versus object category in a matched stimulus set (see Section 1.3.5, Figure 6).

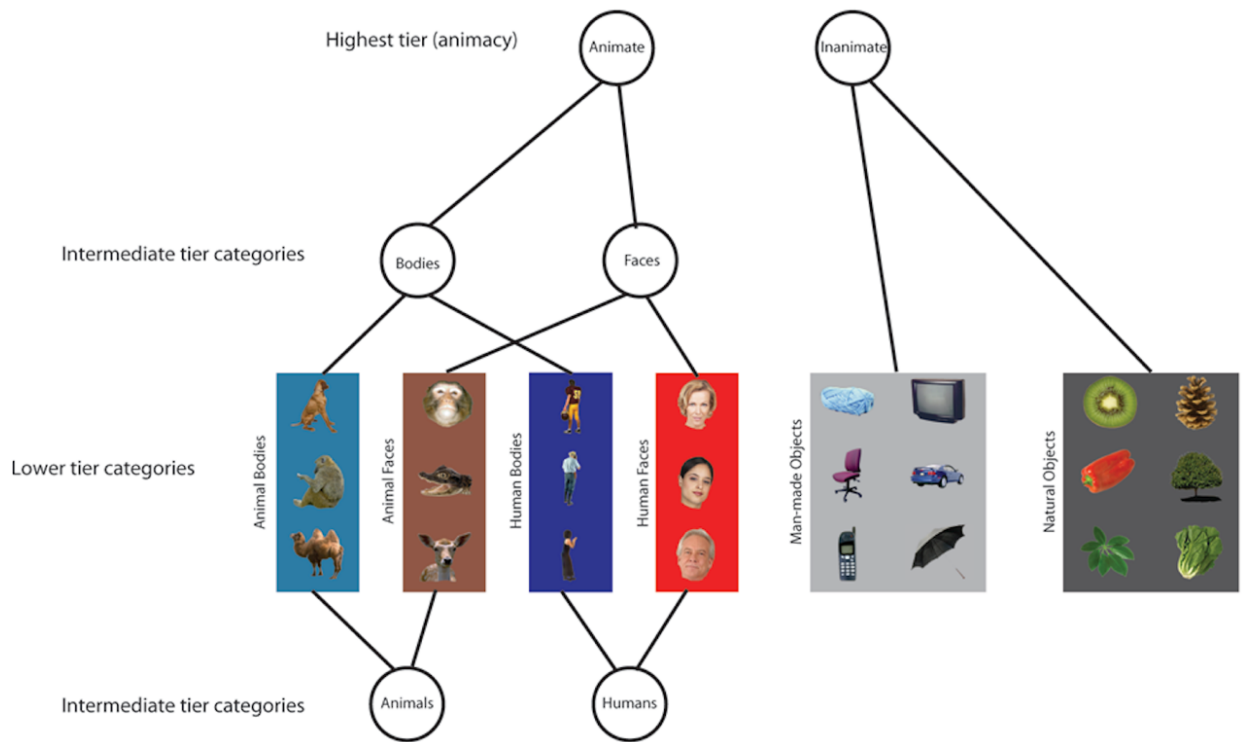
The studies in this section demonstrate the utility of examining the change in representational structure of object representations over time. Studies to date have shown that the brain's response to visual objects is highly dynamic, evolving rapidly from sensitivity to low-level visual properties to more category-like representations. There is also evidence for a robust categorisation distinction between animate and inanimate objects that is sustained. In the following section we examine what has been learnt from time-series decoding about the hierarchical structure of object category representations.

### **1.3.3. Uncovering the Hierarchy of Object Category Representations**

Visual object categorisation can be understood as a dynamic process of evidence accumulation over time (Mack & Palmeri, 2011; Nosofsky & Palmeri, 1997; Philiastides & Sajda, 2006). Functionally, the accumulation of evidence for object category membership is likely to exploit the complex network of feedback and feedforward connections within the object-selective ventral pathway, rather than operating as a linear progression of representation from low-level features through to semantic concepts of increasing abstraction (Kravitz, Saleem, Baker, Ungerleider, & Mishkin, 2013). This is a recent development and extension of the earlier idea of sequential processing stages, where objects are first categorised at an intermediate/basic level (Mervis & Rosch, 1981) (e.g., cat), with superordinate (e.g., animal) and subordinate (e.g., Siamese cat) categorisation occurring later in the visual processing hierarchy (Grill-Spector & Kanwisher, 2005).

MEG decoding methods have been applied to understanding the timing of category abstraction by selecting stimuli with a planned hierarchical category structure and assessing when these category level distinctions emerge. Results from two MEG decoding studies (Carlson et al., 2013; Cichy et al., 2014) provide insights into the category processing hierarchy by evaluating the time course of hierarchically organised object categories. The stimulus sets across both studies followed the same planned hierarchical structure (see Figure 4): the highest (most general/abstract) category tier was the commonly evaluated animate/inanimate distinction. The animate domain was then further subdivided into intermediate categories of faces/bodies as well as humans/animals, with the lowest (most specific) category level being that of human faces/human bodies, animal faces/animal bodies. For the inanimate domain, these objects could be further subcategorised into man-made and natural objects.

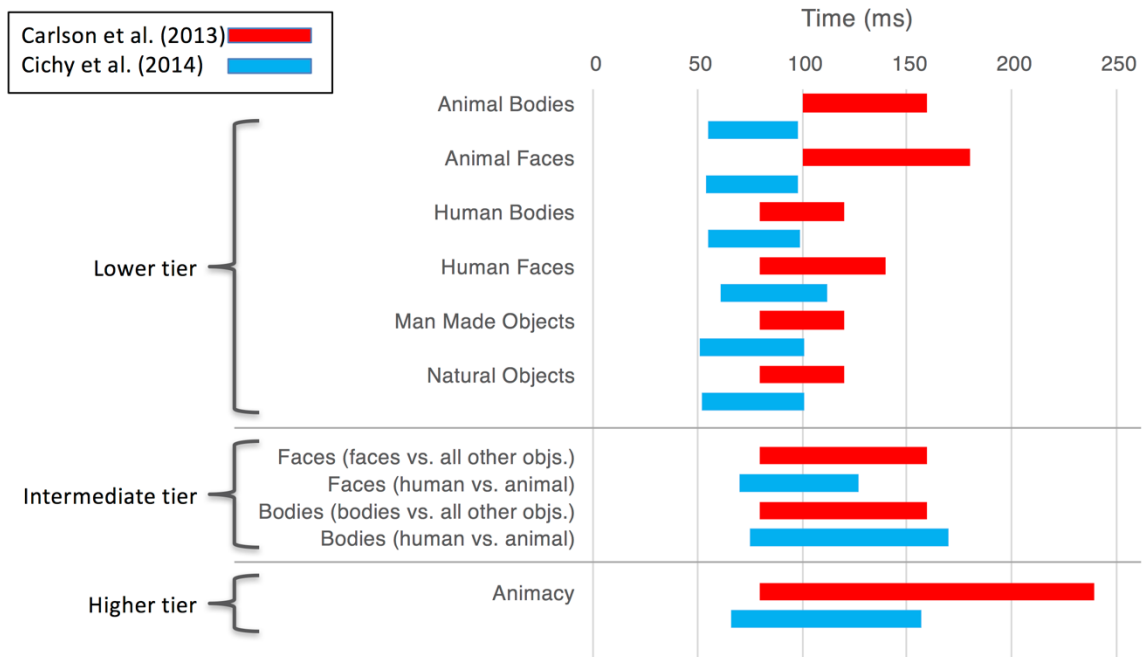
To investigate the dynamics of hierarchical object category representations, these studies looked at two critical moments in the MEG object decoding time-course: the time at which decoding first reaches significance (decoding onset) and the time that the category distinction is maximally differentiated (peak decoding time). A visual comparison of the decoding time course in both Carlson et al. (2013) and Cichy et al. (2014) for object categories at different levels of the category hierarchy is shown in Figure 5. Carlson et al. (2013) found the onset of significant decoding for individual object exemplars occurred around 80 ms after stimulus onset, which was also similar to the decoding onsets of individual categories; all falling between 80 – 100 ms. Cichy et al. (2014) similarly found little variability between overall exemplar decoding onset (~48 ms) and individual category onsets (~51 - 61 ms post-stimulus onset).



*Figure 4.* Hierarchical category structure of the object stimuli used in both Carlson et al. (2013) and Cichy et al. (2014). Diagram shows the three levels of category abstraction: lower tier (animal and human faces and bodies), intermediate tier (bodies and faces), and highest tier (animacy). Note, number of stimuli differed between the two studies: Carlson et al. (2013) used 24 images (as shown), while Cichy et al. (2014) employed 92. Image reprinted from Carlson et al. (2013). Image reprinted with permission of the Association for Research in Vision and Ophthalmology, from Carlson et al. (2013); permission conveyed through Copyright Clearance Center, Inc.

Evaluation of peak decoding times for the category hierarchy in Carlson et al. (2013) revealed a largely linear organisation of the category tiers (see Figure 5, red bars), such that lower tier categories (e.g., human faces, animal bodies) showed peak decoding times from 120 - 180 ms, followed by intermediate tier categories (faces, bodies, human, animal) which ranged from 160 - 240 ms, with the superordinate tier animacy category (animate vs. inanimate) peaking at 240 ms. The results suggest evolving hierarchical representations of object categories that progress from specific, individual exemplar individuation through to more abstract semantic groupings. Results from Cichy et al. (2014) largely support these findings (see Figure 5, blue bars), with peak decoding between intermediate tier categories (natural vs. artificial and faces vs. bodies, 122 and 136 ms respectively) occurring earlier than the superordinate animacy category peak (157 ms).

Although there is similar evidence in the results of both Carlson et al. (2013) and Cichy et al. (2014) for a roughly linear emergence of object categories from specific to more abstract object category representations, the data are not entirely straightforward. For example, for the subordinate tier categories in Cichy et al. (2014), this linear hierarchy does not hold, with subcategory peaks occurring later in the time-course than some of the intermediate tier category peaks: human versus animal faces peaked at 127 ms, while human versus animal bodies peaked at 170 ms. Furthermore, in Carlson et al. (2013), the intermediate tier categories of faces and bodies are best discriminated at a similar time point to the more specific lower tier categories (e.g., human faces, animal bodies), and similarly, Cichy et al. (2014) found no significant difference between the peak latency at which individual images were discriminated early in the time course compared to a higher tier category grouping of human versus non-human bodies.



*Figure 5.* Onset and peak decoding times for category tiers within the object category hierarchy, as evaluated by Carlson et al. (2013) and Cichy et al. (2014). Each coloured bar represents the time from the onset of significant category decoding to the time of peak decoding (relative to post-stimulus onset) for Carlson et al. (2013) (red) and Cichy et al. (2014) (blue). Note, the intermediate tier categories are evaluated using different stimulus comparisons by Carlson et al. (2013) and Cichy et al. (2014) (specific comparisons are noted in brackets).

It is plausible that these anomalies in the emergence of hierarchical category structure reflect that category representations are not a strict linear hierarchy, consistent with recent ideas on the importance of recurrent processing such as lateral connections and feedback on feedforward visual object processing (Kravitz et al., 2013). Carlson et al. (2013) suggest these discrepancies may be partially explained by the inclusion of face and body stimuli in the design, as there is evidence for special processing systems specifically tuned for the detection of faces and bodies in both human and primate ventral temporal cortex (Downing et al., 2001; Kanwisher et al., 1997). If some categories are processed differently than others, a straightforward linear emergence of category structure is not expected. A further caveat is

that these results are specific to the selected stimulus set and the imposed category structure of the experimental design.

#### **1.3.4. Decoding the Time-Course of Face Processing**

The object category of faces stands out as having a particularly strong contribution to decoded brain representations measured with EEG (Kaneshiro, Guimaraes, Kim, Norcia, & Suppes, 2015) and MEG (Van de Nieuwenhuijzen et al., 2013). For example, in the MDS results for large object sets, human faces tend to cluster prominently (Figure 2). Two EEG decoding studies have examined the temporal dynamics of face representations in detail (Barragan-Jason et al., 2015; Cauchoix et al., 2014). A strong feature of these studies is that they used large natural image stimulus sets containing pictures of faces with natural backgrounds, and faces were not repeated in the experimental design. This is in contrast to the majority of temporal object decoding studies to date, which tend to repeat the same object exemplars multiple times within the experiment to increase signal-to-noise. Both studies had participants complete a go/no-go task and trained a classifier to detect the target versus non-target faces across the EEG time-course (human vs. animal in Cauchoix et al., 2014, and famous vs. unfamiliar faces in Barragan-Jason et al., 2015). In both studies, the time-course of face decoding follows a similar pattern: there is an initial sharp rise in performance of the classifier early in the time-course, shortly followed by a relatively brief plateau or drop in decoding, before a second, more gradual rise in decoding accuracy, with peak decoding occurring relatively late in the time-course (350 ms and 600 ms for humans/animals and familiarity respectively). They suggest that the two (early vs. late) peaks in decoding reflect different levels of processing within the object recognition hierarchy. An initial fast feedforward signal driven by low-level stimulus properties is later modified to incorporate more detailed feedback from higher cortical areas (Barragan-Jason et al., 2015; Cauchoix et al., 2014).

A recent study which combined MEG decoding with fMRI localization provides further insight into different levels of face processing. Vida, Nestor, Plaut and Behrmann (2017) evaluated time-series decoding performance of face identity in specific regions of the brain linked to face processing. They used an fMRI localiser to identify two regions that were preferentially activated by faces, over and above activity associated with objects more broadly: right lateral occipital cortex and the right fusiform gyrus. At these sites, and a control site (left V1) they related the time-course of face decoding to three models of face processing: one that represented low-level, V1-like responses (“image-based” model), a higher-level “identity-based” face model, and a behavioural-rating model where participants rated a subset of the images as to how similar the face identities were on a scale from 1 to 8. Their results showed that firstly, while face identity was decodable at all three regions from ~50 - 400ms, the image-based model was a better predictor of the neural data than the identity-based model until ~200 ms, after which its performance dropped below that of the identity-based model at a number of time-points. However, this shift occurred in the face-selective regions only, with the image-based model outperforming others in the V1 control region throughout the time-course.

These results are in line with previous findings (see Section 1.3.1) that suggest low-level visual features are associated with activity in early visual areas, while high-level conceptual information is represented in cortical regions further down the ventral processing stream. The behavioural-rating model similarly correlated with the neural data significantly in all three regions from ~50-400 ms, however, after controlling for representations in the V1 control region, this significant time-window reduced to 100 - 250 ms in the face-selective regions. Notably, behavioural ratings were significantly more similar to the identity-based model than the image-based model. Significantly, by using time-series decoding methods, Vida et al. were able to show that face-selective regions process information related to *both* low-level image properties as well as higher-level face identity-based representations, highlighting the potential shortcomings of associating localised regions with single

functions. The challenges involved in empirically separating out the contribution of low-level properties from more abstract object representations in object decoding studies is discussed in detail in the following section.

### **1.3.5. Is a Category a Category?**

Above we reviewed several studies focusing on uncovering the dynamic category structure of object representations. However, when evaluating object representations, it is important to separate effects related to the abstract concept of an object "category" from other potentially confounding features that co-vary with category membership such as low-level visual properties including shape, colour, and luminance contrast (Wardle & Ritchie, 2014). For example, items within the category of tools tend to have handles, resulting in a characteristic long and often slender feature to their appearance (e.g., hammer, scissors, screwdriver, toothbrush). This relatively consistent information within this category could result in decoding based on low-level feature similarities related to shape, such that instead of decoding the category of "tools" compared to "faces", we may instead be erroneously decoding "long and thin" versus "round". Low-level stimulus properties such as orientation (Cichy, Ramirez, & Pantazis, 2015; Ramkumar, Jas, Pannasch, Hari, & Parkkonen, 2013) are readily decodable from whole-brain MEG signals, and perceptual similarity is a strong determinant of the decodability of abstract visual stimuli (Wardle, Kriegeskorte, Grootswagers, Khaligh-Razavi, & Carlson, 2016).

Time-varying decoding studies investigating object recognition have attempted to account for potential low-level confounds using a variety of different approaches. In fMRI studies, V1 is often used as a control region to show that unsupervised categorical clustering of object representations (e.g., by animacy) do not emerge in V1, but are present in IT (Cichy et al., 2014; Kriegeskorte, Mur, Ruff, et al., 2008). As this degree of spatial localisation is not possible with MEG, time-varying decoding studies often exploit the time-course to demonstrate that models of early visual processing may perform well early on in the time-



course but later the contribution from categorical models such as animacy emerges. For example, Carlson et al. (2013) used RSA to evaluate the performance of low-level feature-based models of vision on object representations measured with MEG. The models of early visual processing included a shape-based model that compared the image silhouettes (Jaccard, 1901), a colour-based model (CIE), and a hierarchical visual processing model (HMAX) containing layers tuned to process varying levels of stimulus complexity (Riesenhuber & Poggio, 1999; Serre et al., 2007). Notably, these models were able to successfully differentiate human faces from other objects (silhouette model), humans from human bodies (CIE), and man-made objects from other objects (HMAX). While the models were only able to distinguish a maximum of two out of a possible 10 object categories in that particular data set, these results highlight the importance of accounting for the potential influence of low-level stimulus properties on object category decoding.

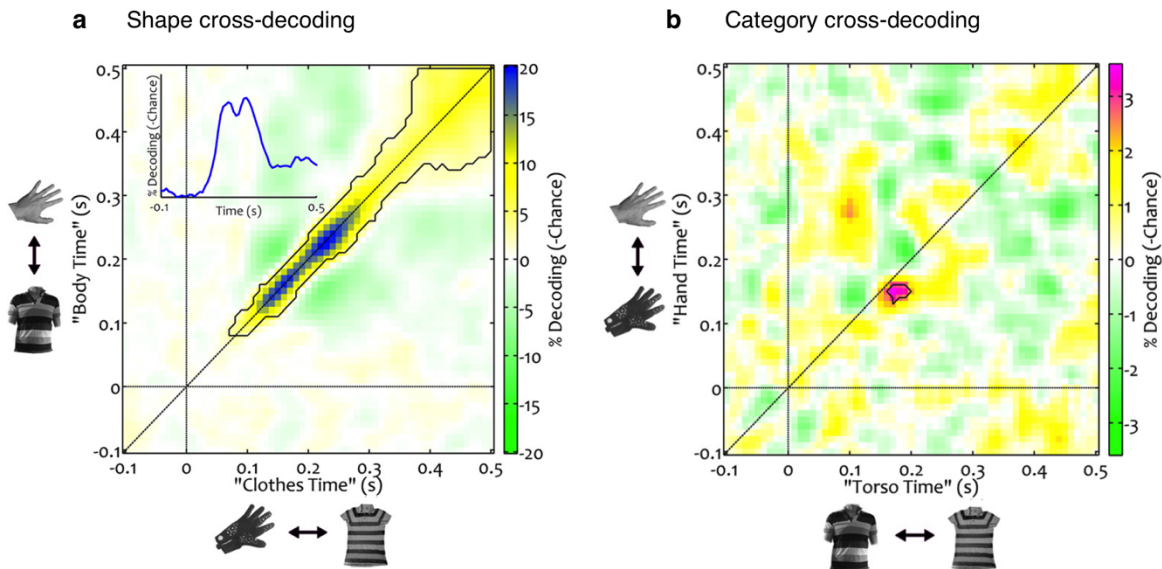
Although there is evidence that visual similarity accounts for at least some of the representational similarity shared by objects within the same category, it is unlikely to account for all observed category effects. As a control for low-level stimulus properties when decoding object category, Carlson et al. (2011) evaluated decoding accuracy of cars and faces from MEG activation patterns compared to artificially generated car and face textures which preserved local image statistics but removed recognisable form. If classification of object category is heavily based on low-level image statistics, it is expected that the classifier would find it difficult to discriminate between objects and their matched texture images, which share low-level image properties (e.g., classifying 'car' vs. 'car-texture'). However, object categories were able to be accurately decoded from their texture counterparts in the whole-brain MEG activation patterns, and this distinction emerged earlier in the time course than information differentiating the two object categories of cars and faces (Carlson et al., 2011). These results are consistent with a transition from decoding based on V1-like image properties to IT-like object category over the MEG decoding time course (Cichy et al., 2014).

The contribution of semantic meaning to category representations was explored in a study by Clarke et al. (2014), which compared the performance of both low-level visual processing and feature-based semantic models for their 302 object stimuli from 11 categories. They found that a model which combined both category-specific semantic information and low-level visual features best accounted for variability in neural object representations later in the MEG time course, over and above what could be achieved by using the HMAX model alone (A. Clarke et al., 2014). These results suggest there is a contribution of abstract category membership (as indexed by semantic similarity) to the brain representation. Similarly, an EEG study evaluating event-related potentials (ERPs) associated with object representations aimed to provide evidence for the involvement of conceptual category membership by examining decoding of the categories 'animals' versus 'tools' across three modalities of visual pictures, spoken words and written words (Simanova et al., 2010). However, classifier performance was much higher for visual pictures than for spoken or written words, complicating the critical cross-classification analysis across modality.

Similar to the frequently-used approach of assessing the potential contribution of low-level properties to object representations with the HMAX model (Carlson et al., 2013; A. Clarke et al., 2014; Isik et al., 2014; Kriegeskorte, Mur, Ruff, et al., 2008), recently Cichy, Khosla, Pantazis, Torralba and Oliva (2016) compared MEG object decoding performance with the performance of deep neural networks (DDNs). DNNs are a powerful form of computer vision model, comprised of multiple non-linear processing layers with the ability to learn tasks such as object categorization (see Kriegeskorte, 2015). The recent advancement of these models has propelled computer modelling of object recognition to a level that is beginning to rival human object categorisation performance (He, Zhang, Ren, & Sun, 2015). By comparing the decoding performance of the DNN model to MEG and fMRI brain data, Cichy et al. (2016) showed a hierarchy of both spatially and temporally ordered processes that correlated with the DDN processing layers in an ordered fashion;

deeper layers were associated with later brain activity in higher cortical processing regions. Furthermore, by analysing activity throughout the entire brain, they identified involvement of not only the ventral, but also the dorsal pathway in object recognition. A key issue for the future is to understand to what degree deep-neural networks mimic the object recognition processes in the human brain, and in what respects the two systems are different.

The above approaches are data-driven, and aim to show that categorical structure that emerges in the neuroimaging data cannot be completely explained by sensitivity to lower level visual properties confounded with object category. A more compelling line of evidence against reducing category representations entirely to low-level properties involves experimental manipulations designed to separate the two factors. A recent MEG/fMRI study (Kaiser, Azzalini, et al., 2016) sought to specifically identify object category representations independent of shared visual properties by purposefully selecting visually similar stimuli that belonged to the semantic categories of body parts and clothing (e.g., gloves vs. hands, shirts vs. torsos). With this stimulus set, visual similarity is balanced across category membership. While their fMRI results showed overlapping spatial representations for both shape and category information, the MEG data showed a specific, comparatively late time window within which category-selective information was present (Figure 6b) and in contrast to this, shape dependent responses (Figure 6a) were decoded relatively early on, from 90 ms post-stimulus onset, and was sustained throughout the time-course. Interestingly, a recent MEG study (Proklova, Kaiser, & Peelen, 2018) employing human behavioural judgements of perceptual similarity, did not show a distinct time-course for object category, but rather, object shape and texture drove object decoding performance. Further investigation is necessary to determine the circumstances that lead to dissociable category representations across neuroimaging modalities, and how these relate to each other.



*Figure 6.* Temporal cross decoding of object category (body parts vs. clothing). The black outlined areas in both plots indicate when cross-decoding of (a) shape or (b) object category was significantly above chance. Images reprinted from Kaiser et al. (2016a) with permission from ©The American Physiological Society.

Object invariance is both an interesting theoretical question for investigation and an experimental manipulation that minimises the contribution of low-level features. Object invariance is the ability to successfully recognise objects despite high variability in their appearance, for instance, due to differences in viewpoint and size. Carlson et al. (2011) showed that object category information for faces and cars could be decoded from the neural data despite changes in retinal location of the image; evidence for a position-invariant object representation. They tested a classifier on information from a novel retinal location that was not included in the original training set, and were able to successfully recover information about the object category as early as 135 ms, suggesting that position-invariant category information is present early in the neural signal. Furthermore, this was not simply due to low-level image statistics, as category information could not be decoded from abstract

textures generated to preserve the local low-level image statistics of the object images (Carlson et al., 2011).

Invariant object decoding was further investigated by Isik et al. (2014). Isik et al. showed position invariant decoding across three stimulus locations for a set of six object exemplars, while also evaluating size invariance. Using three different stimulus sizes for each exemplar, they showed that size-invariant information was present in the signal from around 125 ms post-stimulus onset, earlier in the time-course than position-invariant information which emerged around 150 ms for their stimulus set. Furthermore, by comparing the decoding latencies for the different train-test conditions, they found that comparisons which required a greater transformation of the data from the train to the test condition (e.g., for small, middle and large sized stimuli, train small/test large requires a bigger transformation than train middle/test large) resulted in later decoding onsets.

Successful object decoding across changes in retinal size or position suggests access to more abstract object representations, which is likely to reflect higher-level visual processing. Isik et al. (2014) compared their results to the different levels of the computational HMAX model (Riesenhuber & Poggio, 1999; Serre et al., 2007), which includes multiple stages representing early V1 processing (layer C1), mid-ventral processing (layer C2), and higher level IT-like processing (layer C3). Comparison of the experimental results with the response of the HMAX model for their six object exemplars revealed that only the IT layer of HMAX could successfully decode object exemplars across all transformations of size and position. Isik et al. suggest their observed relationship between the degree of image transformation and decoding time (i.e., that decoding occurs later for larger visual transformations across object size and position) is consistent with a feedforward processing hierarchy supporting invariant object representations. In the next section we consider the link between object category representations and human behaviour.

## **1.4. Linking Dynamic Object Category Representations to Human Behaviour**

One of the critical limitations of decoding approaches in neuroscience is that successful decoding of a stimulus property or experimental condition does not entail that this information is explicitly used by the brain for behaviour (cf. de-Wit et al., 2016; Ritchie, Kaplan, & Klein, in press). An emerging goal is to find ways to link measured brain activation patterns to behaviour (Philiastides & Sajda, 2006). In this section we review two recent approaches to linking the decoded dynamic object representations from time-series neuroimaging data with human behaviour. In the first approach, the goal is to link reaction times (RTs) for object categorisation to the brain's evolving representation of the objects in multidimensional space (Section 1.4.1). The second involves uncovering the effects of attention on decodable object representations (Section 1.4.2).

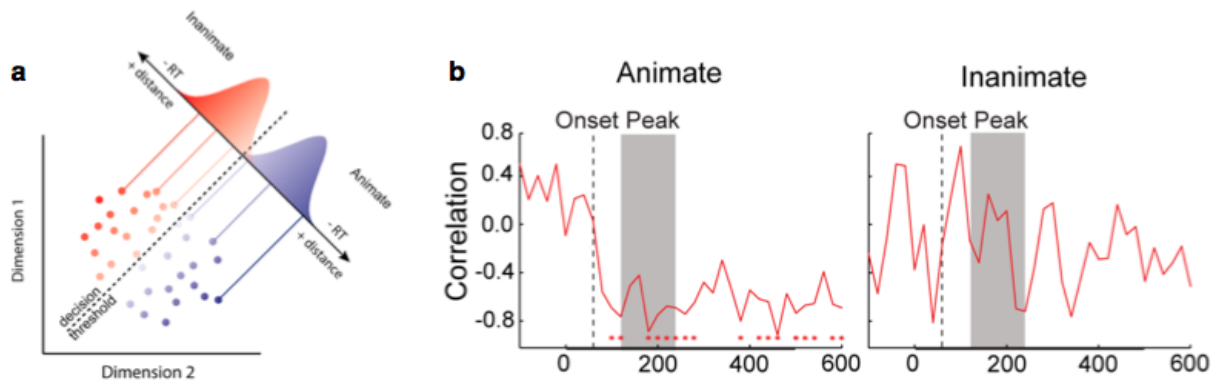
### **1.4.1. Categorisation Reaction-Times**

One fruitful approach to evaluate the relationship between brain and behaviour has been to link human RTs for categorisation to the multidimensional representation of object exemplars in neuroimaging data (Ritchie & Carlson, 2016). Human categorisation behaviour can be conceptualised as a process of evidence accumulation leading to a decision. Behavioural RTs, such as those measured by a simple button press to categorise a stimulus (e.g., fish as 'animate', boot as 'inanimate'), can be thought of as a proxy for the point at which a decision has been made. Carlson, Ritchie, Kriegeskorte, Durvasula, and Ma (2014) tested this idea by combining previously collected fMRI data from human IT (Kriegeskorte, Mur, Ruff, et al., 2008) with separately collected behavioural RTs for categorising the same 92 object images as either 'animate' or 'inanimate'. Carlson et al. (2014) linked the decision boundary from the fMRI data (boundary is that used by a classifier to categorise stimuli by animacy) to RTs for human observers to complete the same categorisation task. They showed that the distance to the classifier decision boundary predicted reaction times for animacy categorisation. Specifically, objects represented further from the animacy decision

boundary were associated with faster behavioural categorisation RTs (i.e., more easily categorised by human observers) than objects represented closer to the boundary (Figure 7a).

Ritchie, Tovar, and Carlson (2015) expanded on this idea by studying the dynamic coupling between emerging object representations and behaviour using MEG decoding. Their findings replicated the relationship between representational distance and RTs found by Carlson et al. (2014) for fMRI, as a larger distance from the classifier boundary was associated with a faster reaction time for animacy categorisation. The key advantage of examining this relationship with MEG is that its emergence following stimulus onset can be tracked over time, which cannot be determined from the fMRI data. Ritchie et al. (2015) found that as decoding accuracy increased following stimulus onset (indicating greater separability of individual object representations with more processing time), the correlation between representational distance and RTs similarly increased. Importantly, this suggests that the strength of the association between representations and behaviour appears to follow decoding performance over time (Ritchie et al., 2015). The significance of this result is that it implies that the brain "reads out" information at the optimal time for making a categorical decision. This highlights the utility of time-series decoding methods, as this result could not be accessed by examining time-averaged activation patterns as in fMRI.

Interestingly, in both the fMRI and MEG studies, the relationship between representational distance and RTs appeared to be driven by animate rather than inanimate objects (Carlson et al. 2014, Ritchie et al., 2015). This is shown in Figure 7b by comparing the left (animate) and right (inanimate) RT correlation panels. For animate object decoding, a negative correlation is observed with RTs around the time of peak decoding (indicating that faster RTs were associated with greater distances from the decision boundary for animacy). However, this relationship is not found with inanimate objects. The authors suggest this asymmetry may be because inanimacy is a negatively defined category (i.e., 'not animate').



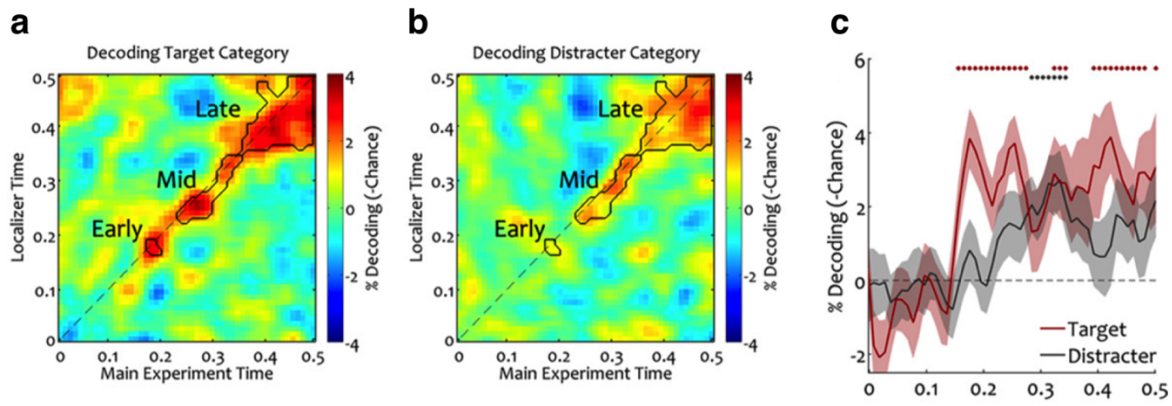
*Figure 7.* Linking behavioural RTs for object animacy categorisation to dynamic object representations measured with MEG. (a) Diagram illustrating how representational distance from the classifier category boundary relates to RTs. (b) Left panel shows the correlation between RTs and the MEG decoding time-course for animate objects, while the right panel shows the same for inanimate objects. Grey bar in both plots indicates the window of peak decoding. Red stars along the x-axis indicate time-points at which the correlation between representational distance and reaction time is significant. Images from Ritchie et al. (2015) reprinted with permission from the author.

#### 1.4.2. Attention

Attention has strong modulatory effects on neural processing (Gandhi, Heeger, & Boynton, 1999; Kastner & Ungerleider, 2000; Moran & Desimone, 1985). A recent MEG study by Kaiser, Oosterhof, & Peelen (2016) evaluated the effects of top-down attention on object category representations. Participants performed a category search task by detecting the presence of either a car or a person embedded in cluttered natural scenes. This design allows comparison of the neural signal associated with the attended versus unattended object category. The classifier was trained on data from viewing car and people exemplars that were separate to the experimental data set and presented as segmented objects on a plain background. The classifier was then tested on the data for when the objects were embedded in complex natural scenes as either a target or a distractor.



The results demonstrate a substantial effect of attention on object representations. Under attended conditions, object category could be decoded rapidly in under 200 ms even though the objects were in highly cluttered scenes (Figure 8c). However, unattended object categories embedded in complex scenes were decoded much later (Figure 8c), demonstrating a strong top-down modulation of attention on visual processing. Attended objects were decoded more easily than unattended objects across all early (180 – 220 ms), mid (230 – 340 ms), and late (350 – 500 ms) time windows (Figure 8a, b). The greater temporal generalisation of cross-decoding during the late time window indicates that the representational similarity shared by isolated objects and objects in scenes is sustained for longer in the late processing stages, which is particularly evident for attended objects (Kaiser, Oosterhof, et al., 2016). These results suggest that top-down attention rapidly modulates object category representations, facilitating rapid target detection in natural scenes (Thorpe, Fize, & Marlot, 1996). By using temporal decoding, this study was able to reveal the time course of category-specific attentional enhancement, extending previous EEG findings which instead focused on comparing evoked potentials to target presence versus absence (Codispoti, Ferrari, Junghöfer, & Schupp, 2006; Thorpe et al., 1996).



*Figure 8.* MEG decoding of attended versus unattended targets. Temporal cross-decoding heat-maps of (a) target category and (b) unattended distractor category (training on isolated objects, testing on objects embedded in cluttered natural scenes in both instances); (c) Time-course of decoding performance for the target and distractor conditions separately. Dots indicate above-chance decoding performance. Images reprinted with permission of Society for Neuroscience, from Kaiser et al. (2016b); permission conveyed through Copyright Clearance Center, Inc.

## 1.5. Summary and Conclusions

Above we have reviewed how time-series decoding methods have expanded our understanding of the neural processes underlying human object recognition. Object recognition remains a significant area of interest within the field of vision research and the recent implementation of time-series decoding methods has allowed researchers to unpack some of the neural processes that underlie *how* the brain achieves this highly complex and multifaceted task. Complementing other neuroimaging techniques, such as fMRI, which provide static snapshots of the visual object processing hierarchy, time-series decoding has unveiled a system of dynamically changing processes that extend throughout the brain. This new perspective brings a host of interesting avenues for further exploration. For example, at what point in time do high-level constructs of attention and consciousness modulate the dynamics of object processing, and what is the effect of this modulation on perception? How do space- and feature-based attention processes differentially influence the processing

hierarchy? There remain unanswered questions about the processes involved in visual working memory and how these representations in the mind's eye compare to the perception of a physical stimulus; are they fundamentally different or is visual working memory a process of revisiting points within these dynamic representations? Finally, it would be of particular interest to know how these various modulations of object representations differentially influence behaviour, further strengthening our understanding of the connection between brain and behaviour. Time-series decoding methods have greatly enriched our knowledge of neural object recognition processes, and show great potential for further advances in this and many other domains.

#### Author contributions:

E.C, S.W, T.C reviewed the literature and developed the ideas

E.C wrote the paper

SW and TC provided critical revisions on the paper

Acknowledgements: This research was supported by an Australian Research Council (ARC) Discovery project (DP160101300). T.A.C is supported by an ARC Future Fellowship (FT120100816). S.G.W. is supported by an Australian NHMRC Early Career Fellowship (APP1072245).

## 1.6. References

- Altmann, C. F., Bühlhoff, H. H., & Kourtzi, Z. (2003). Perceptual Organization of Local Elements into Global Shapes in the Human Visual Cortex. *Current Biology*, 13(4), 342–349. [https://doi.org/10.1016/S0960-9822\(03\)00052-6](https://doi.org/10.1016/S0960-9822(03)00052-6)
- Averbeck, B. B., Latham, P. E., & Pouget, A. (2006). Neural correlations, population coding and computation. *Nature Reviews Neuroscience*, 7(5), 358–366.
- Barragan-Jason, G., Cauchoix, M., & Barbeau, E. J. (2015). The neural speed of familiar face recognition. *Neuropsychologia*, 75, 390–401. <https://doi.org/10.1016/j.neuropsychologia.2015.06.017>
- Beauchamp, M. S., Lee, K. E., Haxby, J. V., & Martin, A. (2002). Parallel visual motion processing streams for manipulable objects and human movements. *Neuron*, 34(1), 149–159.
- Biederman, I. (1987). Recognition-by-components: a theory of human image understanding. *Psychological Review*, 94(2), 115.
- Carlson, Schrater, P., & He, S. (2003). Patterns of Activity in the Categorical Representations of Objects. *Journal of Cognitive Neuroscience*, 15(5), 704–717. <https://doi.org/10.1162/jocn.2003.15.5.704>
- Carlson, T. A., Hogendoorn, H., Kanai, R., Mesik, J., & Turret, J. (2011). High temporal resolution decoding of object position and category. *Journal of Vision*, 11(10), 9. <https://doi.org/10.1167/11.10.9>
- Carlson, T. A., Ritchie, J. B., Kriegeskorte, N., Durvasula, S., & Ma, J. (2014). Reaction Time for Object Categorization Is Predicted by Representational Distance. *Journal of Cognitive Neuroscience*, 26(1), 132–142. [https://doi.org/10.1162/jocn\\_a\\_00476](https://doi.org/10.1162/jocn_a_00476)

- Carlson, Tovar, D. A., Alink, A., & Kriegeskorte, N. (2013). Representational dynamics of object vision: The first 1000 ms. *Journal of Vision*, 13(10), 1–1.  
<https://doi.org/10.1167/13.10.1>
- Cauchoix, M., Barragan-Jason, G., Serre, T., & Barbeau, E. J. (2014). The Neural Dynamics of Face Detection in the Wild Revealed by MVPA. *The Journal of Neuroscience*, 34(3), 846–854. <https://doi.org/10.1523/JNEUROSCI.3030-13.2014>
- Chait, M., Poeppel, D., Cheveigné, A. de, & Simon, J. Z. (2007). Processing Asymmetry of Transitions between Order and Disorder in Human Auditory Cortex. *Journal of Neuroscience*, 27(19), 5207–5214. <https://doi.org/10.1523/JNEUROSCI.0318-07.2007>
- Chao, L. L., Haxby, J. V., & Martin, A. (1999). Attribute-based neural substrates in temporal cortex for perceiving and knowing about objects. *Nature Neuroscience*, 2(10), 913–919.
- Cichy, R. M., Khosla, A., Pantazis, D., Torralba, A., & Oliva, A. (2016). Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence. *Scientific Reports*, 6.  
<https://doi.org/10.1038/srep27755>
- Cichy, R. M., Pantazis, D., & Oliva, A. (2014). Resolving human object recognition in space and time. *Nature Neuroscience*, 17(3), 455–462.  
<https://doi.org/10.1038/nn.3635>
- Cichy, R. M., Ramirez, F. M., & Pantazis, D. (2015). Can visual information encoded in cortical columns be decoded from magnetoencephalography data in humans? *NeuroImage*, 121, 193–204. <https://doi.org/10.1016/j.neuroimage.2015.07.011>

- Clarke, A., Devereux, B. J., Randall, B., & Tyler, L. K. (2014). Predicting the Time Course of Individual Objects with MEG. *Cerebral Cortex*.  
<https://doi.org/10.1093/cercor/bhu203>
- Clarke, A., & Tyler, L. K. (2014). Object-Specific Semantic Coding in Human Perirhinal Cortex. *Journal of Neuroscience*, 34(14), 4766–4775.  
<https://doi.org/10.1523/JNEUROSCI.2828-13.2014>
- Codispoti, M., Ferrari, V., Junghöfer, M., & Schupp, H. T. (2006). The categorization of natural scenes: Brain attention networks revealed by dense sensor ERPs. *NeuroImage*, 32(2), 583–591. <https://doi.org/10.1016/j.neuroimage.2006.04.180>
- Connolly, A. C., Guntupalli, J. S., Gors, J., Hanke, M., Halchenko, Y. O., Wu, Y.-C., ... Haxby, J. V. (2012). The Representation of Biological Classes in the Human Brain. *Journal of Neuroscience*, 32(8), 2608–2618.  
<https://doi.org/10.1523/JNEUROSCI.5547-11.2012>
- Cox, D. D., & Savoy, R. L. (2003). Functional magnetic resonance imaging (fMRI) ‘brain reading’: detecting and classifying distributed patterns of fMRI activity in human visual cortex. *NeuroImage*, 19(2 Pt 1), 261–270.
- de-Wit, L., Alexander, D., Ekroll, V., & Wagemans, J. (2016). Is neuroimaging measuring information in the brain? *Psychonomic Bulletin & Review*, 23(5), 1415–1428.  
<https://doi.org/10.3758/s13423-016-1002-0>
- DiCarlo, J. J., & Cox, D. D. (2007). Untangling invariant object recognition. *Trends in Cognitive Sciences*, 11(8), 333–341.
- Downing, P., Jiang, Y., Shuman, M., & Kanwisher, N. (2001). A cortical area selective for visual processing of the human body. *Science*, 293(5539), 2470–2473.

- Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature*, 392(6676), 598–601.
- Gandhi, S. P., Heeger, D. J., & Boynton, G. M. (1999). Spatial attention affects brain activity in human primary visual cortex. *Proceedings of the National Academy of Sciences*, 96(6), 3314–3319. <https://doi.org/10.1073/pnas.96.6.3314>
- Goddard, E., Carlson, T. A., Dermody, N., & Woolgar, A. (2016). Representational dynamics of object recognition: Feedforward and feedback information flows. *NeuroImage*, 128, 385–397. <https://doi.org/10.1016/j.neuroimage.2016.01.006>
- Grill-Spector, K., Kourtzi, Z., & Kanwisher, N. (2001). The lateral occipital complex and its role in object recognition. *Vision Research*, 41(10–11), 1409–1422.
- Grill-Spector, K., & Malach, R. (2004). THE HUMAN VISUAL CORTEX. *Annual Review of Neuroscience*, 27(1), 649–677.  
<https://doi.org/10.1146/annurev.neuro.27.070203.144220>
- Grill-Spector, K., & Kanwisher, N. (2005). Visual recognition as soon as you know it is there, you know what it is. *Psychological Science*, 16(2), 152–160.
- Grootswagers, T., Wardle, S. G., & Carlson, T. A. (2016). Decoding dynamic brain patterns from evoked responses: A tutorial on multivariate pattern analysis applied to time-series neuroimaging data. *Journal of Cognitive Neuroscience*.  
[https://doi.org/10.1162/jocn\\_a\\_01068](https://doi.org/10.1162/jocn_a_01068)
- Haxby, J. V., Connolly, A. C., & Guntupalli, J. S. (2014). Decoding Neural Representational Spaces Using Multivariate Pattern Analysis. *Annual Review of Neuroscience*, 37(1), 435–456. <https://doi.org/10.1146/annurev-neuro-062012-170325>

- Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and Overlapping Representations of Faces and Objects in Ventral Temporal Cortex. *Science*, 293(5539), 2425–2430.  
<https://doi.org/10.1126/science.1063736>
- Haynes, J.-D. (2015). A primer on pattern-based approaches to fMRI: principles, pitfalls, and perspectives. *Neuron*, 87(2), 257–270.
- Haynes, J.-D., & Rees, G. (2006). Decoding mental states from brain activity in humans. *Nature Reviews Neuroscience*, 7(7), 523–534. <https://doi.org/10.1038/nrn1931>
- He, K., Zhang, X., Ren, S., & Sun, J. (2015). Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification. In *Computer Vision (ICCV), 2015 IEEE International Conference on* (pp. 1026–1034). IEEE. Retrieved from <http://ieeexplore.ieee.org/abstract/document/7410480/>
- Hung, C. P., Kreiman, G., Poggio, T., & DiCarlo, J. J. (2005). Fast readout of object identity from macaque inferior temporal cortex. *Science*, 310(5749), 863–866.
- Ishai, A., Ungerleider, L. G., Martin, A., Schouten, J. L., & Haxby, J. V. (1999). Distributed representation of objects in the human ventral visual pathway. *Proceedings of the National Academy of Sciences*, 96(16), 9379–9384.  
<https://doi.org/10.1073/pnas.96.16.9379>
- Isik, L., Meyers, E. M., Leibo, J. Z., & Poggio, T. (2014). The dynamics of invariant object recognition in the human visual system. *Journal of Neurophysiology*, 111(1), 91–102. <https://doi.org/10.1152/jn.00394.2013>
- Jaccard, P. (1901). *Etude comparative de la distribution florale dans une portion des Alpes et du Jura*. Impr. Corbaz.



- Kaiser, D., Azzalini, D. C., & Peelen, M. V. (2016). Shape-independent object category responses revealed by MEG and fMRI decoding. *Journal of Neurophysiology*, 115(4), 2246–2250. <https://doi.org/10.1152/jn.01074.2015>
- Kaiser, D., Oosterhof, N. N., & Peelen, M. V. (2016). The Neural Dynamics of Attentional Selection in Natural Scenes. *The Journal of Neuroscience*, 36(41), 10522–10528. <https://doi.org/10.1523/JNEUROSCI.1385-16.2016>
- Kaneshiro, B., Guimaraes, M. P., Kim, H.-S., Norcia, A. M., & Suppes, P. (2015). A Representational Similarity Analysis of the Dynamics of Object Processing Using Single-Trial EEG Classification. *PLOS ONE*, 10(8), e0135697. <https://doi.org/10.1371/journal.pone.0135697>
- Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The Fusiform Face Area: A Module in Human Extrastriate Cortex Specialized for Face Perception. *Journal of Neuroscience*, 17, 4302–4311.
- Kastner, S., & Ungerleider, L. G. (2000). Mechanisms of Visual Attention in the Human Cortex. *Annual Review of Neuroscience*, 23(1), 315–341. <https://doi.org/10.1146/annurev.neuro.23.1.315>
- Kiani, R., Esteky, H., Mirpour, K., & Tanaka, K. (2007). Object Category Structure in Response Patterns of Neuronal Population in Monkey Inferior Temporal Cortex. *Journal of Neurophysiology*, 97(6), 4296–4309. <https://doi.org/10.1152/jn.00024.2007>
- King, J.-R., & Dehaene, S. (2014). Characterizing the dynamics of mental representations: the temporal generalization method. *Trends in Cognitive Sciences*, 18(4), 203–210. <https://doi.org/10.1016/j.tics.2014.01.002>

- Kravitz, D. J., Saleem, K. S., Baker, C. I., Ungerleider, L. G., & Mishkin, M. (2013). The ventral visual pathway: An expanded neural framework for the processing of object quality. *Trends in Cognitive Sciences*, 17(1), 26–49.  
<https://doi.org/10.1016/j.tics.2012.10.011>
- Kriegeskorte, N. (2015). Deep neural networks: a new framework for modelling biological vision and brain information processing. *BioRxiv*, 029876.  
<https://doi.org/10.1101/029876>
- Kriegeskorte, N., & Kievit, R. A. (2013). Representational geometry: integrating cognition, computation, and the brain. *Trends in Cognitive Sciences*, 17(8), 401–412.  
<https://doi.org/10.1016/j.tics.2013.06.007>
- Kriegeskorte, N., Mur, M., & Bandettini, P. A. (2008). Representational similarity analysis - connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, 2(4). <https://doi.org/10.3389/neuro.06.004.2008>
- Kriegeskorte, N., Mur, M., Ruff, D. A., Kiani, R., Bodurka, J., Esteky, H., ... Bandettini, P. A. (2008). Matching Categorical Object Representations in Inferior Temporal Cortex of Man and Monkey. *Neuron*, 60(6), 1126–1141.  
<https://doi.org/10.1016/j.neuron.2008.10.043>
- Liu, H., Agam, Y., Madsen, J. R., & Kreiman, G. (2009). Timing, Timing, Timing: Fast Decoding of Object Information from Intracranial Field Potentials in Human Visual Cortex. *Neuron*, 62(2), 281–290. <https://doi.org/10.1016/j.neuron.2009.02.025>
- Mack, M. L., & Palmeri, T. J. (2011). The Timing of Visual Object Categorization. *Frontiers in Psychology*, 2. <https://doi.org/10.3389/fpsyg.2011.00165>
- Malach, R., Levy, I., & Hasson, U. (2002). The topography of high-order human object areas. *Trends in Cognitive Sciences*, 6(4), 176–184.

- Martin, A., Wiggs, C. L., Ungerleider, L. G., & Haxby, J. V. (1996). Neural correlates of category-specific knowledge. Retrieved from <http://www.nature.com/nature/journal/v379/n6566/abs/379649a0.html>
- Mervis, C. B., & Rosch, E. (1981). Categorization of natural objects. *Annual Review of Psychology*, 32(1), 89–115.
- Moran, J., & Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Frontiers in Cognitive Neuroscience*, 229, 342–345.
- Nosofsky, R. M., & Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review*, 104(2), 266.
- Nowak, L. G., & Bullier, J. (1997). The Timing of Information Transfer in the Visual System. In K. S. Rockland, J. H. Kaas, & A. Peters (Eds.), *Extrastriate Cortex in Primates* (pp. 205–241). Springer US. [https://doi.org/10.1007/978-1-4757-9625-4\\_5](https://doi.org/10.1007/978-1-4757-9625-4_5)
- O'Toole, A. J., Jiang, F., Abdi, H., & Haxby, J. V. (2005). Partially distributed representations of objects and faces in ventral temporal cortex. *Journal of Cognitive Neuroscience*, 17(4), 580–590.
- Pereira, F., Mitchell, T., & Botvinick, M. (2009). Machine learning classifiers and fMRI: a tutorial overview. *Neuroimage*, 45(1), S199–S209.
- Philiastides, M. G., & Sajda, P. (2006). Temporal characterization of the neural correlates of perceptual decision making in the human brain. *Cerebral Cortex*, 16(4), 509–518.

- Proklova, D., Kaiser, D., & Peelen, M. V. (2018). MEG sensor patterns reflect perceptual but not categorical similarity of animate and inanimate objects. *BioRxiv*, 238584. <https://doi.org/10.1101/238584>
- Ramkumar, P., Jas, M., Pannasch, S., Hari, R., & Parkkonen, L. (2013). Feature-Specific Information Processing Precedes Concerted Activation in Human Visual Cortex. *The Journal of Neuroscience*, 33(18), 7691–7699. <https://doi.org/10.1523/JNEUROSCI.3905-12.2013>
- Riesenhuber, M., & Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nature Neuroscience*, 2(11), 1019–1025. <https://doi.org/10.1038/14819>
- Ritchie, J. B., & Carlson, T. A. (2016). Neural decoding and “inner” psychophysics: a distance-to-bound approach for linking mind, brain, and behavior. *Frontiers in Neuroscience*, 10.
- Ritchie, J. B., Kaplan, D. M., & Klein, C. (in press). Decoding the brain: Neural representation and the limits of multivariate pattern analysis in cognitive neuroscience. *British Journal for the Philosophy of Science*.
- Ritchie, J. B., Tovar, D. A., & Carlson, T. A. (2015). Emerging Object Representations in the Visual System Predict Reaction Times for Categorization. *PLoS Comput Biol*, 11(6), e1004316. <https://doi.org/10.1371/journal.pcbi.1004316>
- Schmolesky, M. T., Wang, Y., Hanes, D. P., Thompson, K. G., Leutgeb, S., Schall, J. D., & Leventhal, A. G. (1998). Signal Timing Across the Macaque Visual System. *Journal of Neurophysiology*, 79(6), 3272–3278.
- Serre, T., Kreiman, G., Kouh, M., Cadieu, C., Knoblich, U., & Poggio, T. (2007). A quantitative theory of immediate visual recognition. *Progress in Brain Research*, 165, 33–56.

- Sha, L., Haxby, J. V., Abdi, H., Guntupalli, J. S., Oosterhof, N. N., Halchenko, Y. O., & Connolly, A. C. (2015). The Animacy Continuum in the Human Ventral Vision Pathway. *Journal of Cognitive Neuroscience*, 27(4), 665–678.  
[https://doi.org/10.1162/jocn\\_a\\_00733](https://doi.org/10.1162/jocn_a_00733)
- Simanova, I., Gerven, M. van, Oostenveld, R., & Hagoort, P. (2010). Identifying Object Categories from Event-Related EEG: Toward Decoding of Conceptual Representations. *PLOS ONE*, 5(12), e14465.  
<https://doi.org/10.1371/journal.pone.0014465>
- Thorpe, S., Fize, D., & Marlot, C. (1996). Speed of processing in the human visual system. *Nature*, 381(6582), 520–522.
- Van de Nieuwenhuijzen, M. E., Backus, A. R., Bahramisharif, A., Doeller, C. F., Jensen, O., & van Gerven, M. A. (2013). MEG-based decoding of the spatiotemporal dynamics of visual category perception. *NeuroImage*, 83, 1063–1073.
- Van Essen, D. C., Anderson, C. H., & Felleman, D. J. (1992). Information Processing in the Primate Visual System: An Integrated Systems Perspective. *Science*, 255(5043), 419.
- Vida, M. D., Nestor, A., Plaut, D. C., & Behrmann, M. (2017). Spatiotemporal dynamics of similarity-based neural representations of facial identity. *Proceedings of the National Academy of Sciences*, 114(2), 388–393.
- Wardle, S. G., Kriegeskorte, N., Grootswagers, T., Khaligh-Razavi, S.-M., & Carlson, T. A. (2016). Perceptual similarity of visual patterns predicts dynamic neural activation patterns measured with MEG. *NeuroImage*, 132, 59–70.  
<https://doi.org/10.1016/j.neuroimage.2016.02.019>

Wardle, S. G., & Ritchie, J. B. (2014). Can Object Category-Selectivity in the Ventral Visual Pathway Be Explained by Sensitivity to Low-Level Image Properties? *Journal of Neuroscience*, 34(45), 14817–14819.  
<https://doi.org/10.1523/JNEUROSCI.3566-14.2014>

### iii. Interim Summary and Chapter Outline

With the development of M/EEG methods, the temporal dynamics of object category processing is being established. Much of the research into object category structure has pointed to object animacy as an overarching organisational principle (e.g., Caramazza & Shelton, 1998; Chao et al., 1999; Cichy, Pantazis, & Oliva, 2014; Kiani et al., 2007; Kriegeskorte et al., 2008; Lu, Li, & Meng, 2016; Mahon & Caramazza, 2009; Proklova, Kaiser, & Peelen, 2016). More recently, other principles that may govern the wide-ranging patterns of activation related to object processing have been proposed (such as object size: Julian et al., 2017; Konkle & Caramazza, 2013; and a biological class-based animacy continuum: Connolly et al., 2012; Sha et al., 2015). However, the temporal dynamics of these are still being explored.

The primary goal of this thesis is to investigate the temporal dynamics of object category representations, with a focus on testing the limits of the animacy dichotomy. In pursuit of this goal, I use MEG to evaluate the time-course of a range of established and hypothesised organisation principles related to object categorisation. I use brain decoding methods and the representational similarity analysis framework (Kriegeskorte & Kievit, 2013; Nili et al., 2014) to investigate how well these proposed principles account for the representational geometry throughout the time-course of object processing and to what degree these findings link to behaviour.

It is important to note here that this thesis does not cover the related field of speeded categorisation, which is concerned with how the timing of behavioural categorisation judgements relates to different levels of category abstraction (i.e., comparing the speed at which participants are able to categorise an object as a *kookaburra* versus the more abstract category of a *bird* (see Grill-Spector & Kanwisher, 2005; Mack, Alan, Gauthier, Tanaka, & Palmeri, 2009; Palmeri & Blalock, 2000; Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976; Thorpe, Fize, & Marlot, 1996; Wu, Crouzet, Thorpe, &

Fabre-Thorpe, 2014). While a number of theories regarding a processing hierarchy exist, there is ongoing debate as to what level of object category abstraction emerges first, under what specific conditions this occurs, and indeed whether it is beneficial to relate levels of categorisation to particular stages of visual processing (for discussion of the different theoretical and computational accounts related to this topic see Mack & Palmeri, 2011, 2015). This thesis however, focuses on how the brain represents object categories by comparing their differing temporal signatures, and examines whether the emergent geometry of these representations is also useful for behavioural categorisation. Relatedly, while acknowledging that there is considerable research into localising where in the brain object processing occurs (see Grill-Spector & Weiner, 2014; Haxby et al., 2001, 2011; Ishai, Ungerleider, Martin, Schouten, & Haxby, 1999), with a particular focus on the ventral-visual processing stream and involvement of the inferior temporal cortex, MEG as a neuroimaging technique is better utilised for asking questions related to time, rather than space. Other studies that have combined both MEG and fMRI neuroimaging techniques (e.g., Cichy, Pantazis, & Oliva, 2014, 2016; Hebart, Bankson, Harel, Baker, & Cichy, 2018; Kaiser, Azzalini, & Peelen, 2016; Khaligh-Razavi, Cichy, Pantazis, & Oliva, 2018) provide a more thorough discussion of questions related to the spatial evolution of these processes over time. The following sections provide an overview of how each empirical chapter contributes to the goals of this thesis.

In the field of object categorisation, few wide-ranging organizational principles have been successfully identified. The concept of object animacy has been evaluated using both fMRI and MEG neuroimaging techniques, however, the more recently proposed organisation related to an object's real-world size has yet to be established with MEG. In **Chapter 2**, I sought to investigate the temporal dynamics of real-world size as an organisational principle in brain object representations, and evaluate how these temporal patterns compare to the time-course of object animacy.



As I was unable to detect time-varying representations related to real-world size in the MEG data, I shifted my focus to the animacy dichotomy. Recent fMRI research has extended upon the animacy dichotomy, leading to the suggestion that objects are represented along a continuum from humans to inanimate objects, with the position of a category along this continuum dependent on its biological class (Connolly et al., 2012; Sha et al., 2015). In **Chapter 3**, I evaluated the temporal dynamics of this, and other plausible models of the organisation of object category representations. In particular, I sought to test how well models related to human experiences and similarity accounted for variability in object representations. My stimulus set was specifically designed to test the limits of the contemporary animacy model by incorporating objects that would be difficult for categorisation according to standard animacy criteria.

The focus of **Chapter 4** is to link findings from neuroimaging results to human behaviour, with the important goal of determining whether the patterns I observe in the MEG data in Chapter 3 are useful for behavioural processes (Carlson, Goddard, Kaplan, Klein, & Ritchie, 2017; de-Wit, Alexander, Ekroll, & Wagemans, 2016; Krakauer, Ghazanfar, Gomez-Marin, MacIver, & Poeppel, 2017; Ritchie & Carlson, 2016). Uncovering the relationship between brain and behaviour processes remains a challenge in cognitive neuroscience, but is a vital step in clarifying the utility of measured brain activity. Linking reaction times (RTs) to brain representations is one approach to uncovering this relationship, and this method has been implemented successfully in studies using fMRI (Carlson et al., 2014) and MEG (Grootswagers, Ritchie, Wardle, Heathcote, & Carlson, 2017; Ritchie et al., 2015). Following on from my results in Chapter 3, which showed that a model based on human-similarity best accounted for object representations throughout the decoding time-course, I investigated whether this, and other plausible category organisations, could be used to reveal a correlation between brain and behaviour processes. This method of linking behavioural categorisation RTs to brain representations had previously only been implemented for categorisation based on animacy. As such, it is important to also see

whether my novel neuroimaging result related to human-similarity also relates to behaviour in this way.

Finally, **Chapter 5** provides a general discussion of the findings throughout this thesis, synthesising the common themes that emerge across my empirical studies. My findings related to the temporal dynamics of object category representations are discussed in the context of recent advancements in the field, and suggestions for future research are highlighted.

#### iv. References

- Caramazza, A., & Shelton, J. R. (1998). Domain-Specific Knowledge Systems in the Brain: The Animate-Inanimate Distinction. *Journal of Cognitive Neuroscience*, 10(1), 1–34. <https://doi.org/10.1162/089892998563752>
- Carlson, T. A., Goddard, E., Kaplan, D. M., Klein, C., & Ritchie, J. B. (2017). Ghosts in machine learning for cognitive neuroscience: Moving from data to theory. *NeuroImage*. <https://doi.org/10.1016/j.neuroimage.2017.08.019>
- Carlson, T. A., Ritchie, J. B., Kriegeskorte, N., Durvasula, S., & Ma, J. (2014). Reaction Time for Object Categorization Is Predicted by Representational Distance. *Journal of Cognitive Neuroscience*, 26(1), 132–142. [https://doi.org/10.1162/jocn\\_a\\_00476](https://doi.org/10.1162/jocn_a_00476)
- Chao, L. L., Haxby, J. V., & Martin, A. (1999). Attribute-based neural substrates in temporal cortex for perceiving and knowing about objects. *Nature Neuroscience*, 2(10), 913–919.
- Cichy, R. M., Pantazis, D., & Oliva, A. (2014). Resolving human object recognition in space and time. *Nature Neuroscience*, 17(3), 455–462. <https://doi.org/10.1038/nn.3635>
- Connolly, A. C., Guntupalli, J. S., Gors, J., Hanke, M., Halchenko, Y. O., Wu, Y.-C., ... Haxby, J. V. (2012). The Representation of Biological Classes in the Human Brain. *Journal of Neuroscience*, 32(8), 2608–2618. <https://doi.org/10.1523/JNEUROSCI.5547-11.2012>
- de-Wit, L., Alexander, D., Ekroll, V., & Wagemans, J. (2016). Is neuroimaging measuring information in the brain? *Psychonomic Bulletin & Review*, 23(5), 1415–1428. <https://doi.org/10.3758/s13423-016-1002-0>

- Grill-Spector, K., & Kanwisher, N. (2005). Visual recognition as soon as you know it is there, you know what it is. *Psychological Science*, *16*(2), 152–160.
- Grill-Spector, K., & Weiner, K. S. (2014). The functional architecture of the ventral temporal cortex and its role in categorization. *Nature Reviews. Neuroscience*, *15*(8), 536–548. <https://doi.org/10.1038/nrn3747>
- Grootswagers, T., Ritchie, J. B., Wardle, S. G., Heathcote, A., & Carlson, T. A. (2017). Asymmetric Compression of Representational Space for Object Animacy Categorization under Degraded Viewing Conditions. *Journal of Cognitive Neuroscience*, 1–16. [https://doi.org/10.1162/jocn\\_a\\_01177](https://doi.org/10.1162/jocn_a_01177)
- Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and Overlapping Representations of Faces and Objects in Ventral Temporal Cortex. *Science*, *293*(5539), 2425–2430. <https://doi.org/10.1126/science.1063736>
- Haxby, J. V., Guntupalli, J. S., Connolly, A. C., Halchenko, Y. O., Conroy, B. R., Gobbini, M. I., ... Ramadge, P. J. (2011). A Common, High-Dimensional Model of the Representational Space in Human Ventral Temporal Cortex. *Neuron*, *72*(2), 404–416. <https://doi.org/10.1016/j.neuron.2011.08.026>
- Ishai, A., Ungerleider, L. G., Martin, A., Schouten, J. L., & Haxby, J. V. (1999). Distributed representation of objects in the human ventral visual pathway. *Proceedings of the National Academy of Sciences*, *96*(16), 9379–9384. <https://doi.org/10.1073/pnas.96.16.9379>

- Julian, J. B., Ryan, J., & Epstein, R. A. (2017). Coding of Object Size and Object Category in Human Visual Cortex. *Cerebral Cortex*, 27(6), 3095–3109.  
<https://doi.org/10.1093/cercor/bhw150>
- Kiani, R., Esteky, H., Mirpour, K., & Tanaka, K. (2007). Object Category Structure in Response Patterns of Neuronal Population in Monkey Inferior Temporal Cortex. *Journal of Neurophysiology*, 97(6), 4296–4309.  
<https://doi.org/10.1152/jn.00024.2007>
- Konkle, T., & Caramazza, A. (2013). Tripartite Organization of the Ventral Stream by Animacy and Object Size. *Journal of Neuroscience*, 33(25), 10235–10242.  
<https://doi.org/10.1523/JNEUROSCI.0983-13.2013>
- Krakauer, J. W., Ghazanfar, A. A., Gomez-Marin, A., MacIver, M. A., & Poeppel, D. (2017). Neuroscience Needs Behavior: Correcting a Reductionist Bias. *Neuron*, 93(3), 480–490. <https://doi.org/10.1016/j.neuron.2016.12.041>
- Kriegeskorte, N., & Kievit, R. A. (2013). Representational geometry: integrating cognition, computation, and the brain. *Trends in Cognitive Sciences*, 17(8), 401–412.  
<https://doi.org/10.1016/j.tics.2013.06.007>
- Kriegeskorte, N., Mur, M., Ruff, D. A., Kiani, R., Bodurka, J., Esteky, H., ... Bandettini, P. A. (2008). Matching Categorical Object Representations in Inferior Temporal Cortex of Man and Monkey. *Neuron*, 60(6), 1126–1141.  
<https://doi.org/10.1016/j.neuron.2008.10.043>
- Lu, Z., Li, X., & Meng, M. (2016). Encodings of implied motion for animate and inanimate object categories in the two visual pathways. *NeuroImage*, 125, 668–680.  
<https://doi.org/10.1016/j.neuroimage.2015.10.059>

- Mack, M. L., Alan, C.-N. W., Gauthier, I., Tanaka, J. W., & Palmeri, T. J. (2009). Time course of visual object categorization: Fastest does not necessarily mean first. *Vision Research*, 49, 1961–1968.
- Mack, M. L., & Palmeri, T. J. (2011). The Timing of Visual Object Categorization. *Frontiers in Psychology*, 2. <https://doi.org/10.3389/fpsyg.2011.00165>
- Mack, M. L., & Palmeri, T. J. (2015). The dynamics of categorization: Unraveling rapid categorization. *Journal of Experimental Psychology: General*, 144(3), 551.
- Mahon, B. Z., & Caramazza, A. (2009). Concepts and Categories: A Cognitive Neuropsychological Perspective. *Annual Review of Psychology*, 60(1), 27–51. <https://doi.org/10.1146/annurev.psych.60.110707.163532>
- Nili, H., Wingfield, C., Walther, A., Su, L., Marslen-Wilson, W., & Kriegeskorte, N. (2014). A Toolbox for Representational Similarity Analysis. *PLoS Computational Biology*, 10(4), e1003553. <https://doi.org/10.1371/journal.pcbi.1003553>
- Palmeri, T. J., & Blalock, C. L. (2000). The role of background knowledge in speeded perceptual categorization. *Cognition*, 77(2), B45 – B57.
- Proklova, D., Kaiser, D., & Peelen, M. V. (2016). Disentangling Representations of Object Shape and Object Category in Human Visual Cortex: The Animate–Inanimate Distinction. *Journal of Cognitive Neuroscience*, 1–13. [https://doi.org/10.1162/jocn\\_a\\_00924](https://doi.org/10.1162/jocn_a_00924)
- Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, 8(3), 382–439.
- Ritchie, J. B., & Carlson, T. A. (2016). Neural decoding and “inner” psychophysics: a distance-to-bound approach for linking mind, brain, and behavior. *Frontiers in Neuroscience*, 10.

- Ritchie, J. B., Tovar, D. A., & Carlson, T. A. (2015). Emerging Object Representations in the Visual System Predict Reaction Times for Categorization. *PLoS Comput Biol*, *11*(6), e1004316. <https://doi.org/10.1371/journal.pcbi.1004316>
- Sha, L., Haxby, J. V., Abdi, H., Guntupalli, J. S., Oosterhof, N. N., Halchenko, Y. O., & Connolly, A. C. (2015). The Animacy Continuum in the Human Ventral Vision Pathway. *Journal of Cognitive Neuroscience*, *27*(4), 665–678. [https://doi.org/10.1162/jocn\\_a\\_00733](https://doi.org/10.1162/jocn_a_00733)
- Thorpe, S., Fize, D., & Marlot, C. (1996). Speed of processing in the human visual system. *Nature*, *381*(6582), 520–522.
- Wu, C.-T., Crouzet, S. M., Thorpe, S. J., & Fabre-Thorpe, M. (2014). At 120 msec You Can Spot the Animal but You Don't Yet Know It's a Dog. *Journal of Cognitive Neuroscience* *27*(1), 141–149.





## CHAPTER TWO

---

---

# EVALUATING THE TEMPORAL DYNAMICS OF OBJECT REPRESENTATIONS AS A FUNCTION OF ANIMACY AND REAL-WORLD SIZE

---

Erika W. Contini<sup>1,2</sup>, Erin Goddard<sup>3</sup>, Susan G. Wardle<sup>2,4</sup>, Mark Williams<sup>1,2</sup>,  
& Thomas A. Carlson<sup>2,5</sup>

<sup>1</sup>Department of Cognitive Science, Macquarie University, Sydney, Australia

<sup>2</sup>ARC Centre of Excellence in Cognition and its Disorders and Perception in Action Research Centre,  
Macquarie University, Australia

<sup>3</sup>McGill Vision Research Group, Department of Ophthalmology, McGill University, Montreal, Canada

<sup>4</sup>Laboratory of Brain and Cognition, National Institutes of Health, Bethesda, MD, USA

<sup>5</sup>School of Psychology, University of Sydney, Australia



## **2.1. Abstract**

In the field of human object recognition, object animacy has become a widely accepted principle in the organization of object representations. More recently, evidence from fMRI studies supporting a real-world size distinction in object representations has also emerged, with size-selective regions identified adjacent to those that show a preference for object animacy properties. While research has focused on the spatial patterns of activation related to object size properties, the temporal dynamics of real-world size representations are yet to be determined. Using representational similarity analysis of magnetoencephalography data, we sought to establish the temporal evolution of real-world size dimensions, and evaluate how this relates to the time-course of object animacy. Using a stimulus set that was balanced for object interactivity, we replicated previous findings related to the time-course of object animacy. Surprisingly, we were unable to detect modulations related to real-world size in the brain activation patterns. Our results are interpreted in the context of recent spatial and temporal findings showing a real-world size dimension for object category representations. Our findings highlight the importance of carefully considered stimulus sets, and the difficulty in separating out the unique contributions of individual object properties that contribute to multifaceted object representations.



## 2.2. Introduction

Understanding the neural processes involved in object recognition is an ongoing challenge in visual neuroscience. A key goal is to unveil the intricacies of higher-level object representations which afford us the ability to rapidly differentiate, with remarkable accuracy, the vast array of objects we encounter on a day-to-day basis. Several fMRI studies have evaluated the brain activity patterns in response to specific object categories, with a particular focus on brain regions within ventral occipital-temporal cortex - a region widely accepted as a central hub for object recognition processes (Caramazza & Shelton, 1998; Grill-Spector & Malach, 2004; Haxby et al., 2001; Ishai, Ungerleider, Martin, Schouten, & Haxby, 1999). These studies identified several specialised category-selective regions which respond preferentially to particular objects including faces (Ishai et al., 1999; Kanwisher, McDermott, & Chun, 1997), bodies (Downing, Jiang, Shuman, & Kanwisher, 2001), animals (Chao, Haxby, & Martin, 1999; Martin, Wiggs, Ungerleider, & Haxby, 1996), places (Epstein, Harris, Stanley, & Kanwisher, 1999; Epstein & Kanwisher, 1998) and tools (Beauchamp, Lee, Haxby, & Martin, 2002; Chao et al., 1999; Chao & Martin, 2000). While a number of category-selective regions have been identified, these account for only a subset of object categories encountered in everyday life (Biederman, 1987), and by attributing particular brain regions to relatively broad category definitions we limit the extent to which we can account for fine-grained variations at the subcategory level (Haxby et al., 2001; for discussion, see also Vida, Nestor, Plaut, & Behrmann, 2017).

In addition to clustered category-selective regions, findings from multivariate pattern analysis (MVPA) studies have evaluated coarse-scale patterns of category-selective activity distributed throughout occipital-temporal visual cortex (Haxby et al., 2001), with a particular focus on the organisational principle of animacy, which suggests that object representations are at least partly structured according to whether an object is animate or inanimate (Caramazza & Shelton, 1998; Chao et al., 1999; Huth, Nishimoto, Vu, & Gallant, 2012; Kriegeskorte, Mur, Ruff, et al., 2008; Mahon & Caramazza, 2011; Martin, 2007). More

recently, real-world size has been identified as another potential principle underlying the large-scale organization of object representations in the ventral visual stream. Using fMRI to measure patterns of BOLD activation while participants viewed a range of big (e.g., car, dishwasher) and small (e.g., coin, mug) inanimate object images, Konkle and Oliva (2012) found large-scale regions in ventral temporal cortex which showed preferential activation depending on an object's real-world size. These size-coding regions overlapped with other known category-selective regions in a logical manner: regions that responded preferentially to big objects overlapped with regions that have been shown to respond to landmarks and buildings (i.e., the parahippocampal place area; PPA), while localised activity associated with small objects was adjacent to regions known for their involvement with processing faces, bodies and tools (Konkle & Oliva, 2012). This pattern of differential activation across cortical regions as a function of real-world size was robust to changes in retinal size, and was also evident during mental imagery of big and small inanimate objects, which the authors conclude indicates higher-level conceptual processing of size independent of physical stimulus properties.

Further support for coding of size within the object representations is shown in an fMRI study by Konkle and Caramazza (2013), who evaluated the relationship between the real-world size of objects and object animacy in the large-scale organization of object representations in the human brain. Their results similarly show a size distinction within inanimate objects, however the same was not apparent for animate objects. They concluded that the structure of object representations is best encompassed by a "tripartite" organization, where object representations are primarily coded as either animate, big-inanimate, or small-inanimate objects. Further, the order of preferential activation in these regions for each of the three categorical distinctions displayed an interesting pattern. The cortical region that showed a preference primarily for big-inanimate objects had the next highest preferential activation in response to small-inanimate objects, and the lowest activity for animate objects, thus preserving the established animate-inanimate distinction. However, this pattern was

markedly different for regions with a preference for small-inanimate objects, which instead showed the next highest activation in response to animate objects, and the lowest response to big-inanimate objects (Konkle & Caramazza, 2013). The authors highlight that this region, encompassing the lateral temporal cortex, has historically been linked to activity in response to manipulable objects, and interactions between object and agent (Beauchamp et al., 2002; Bracci, Cavina-Pratesi, Ietswaart, Caramazza, & Peelen, 2012; Bracci & Peelen, 2013), which may explain why the dissociation between animate and inanimate objects is somewhat blurred in this region.

While these studies show activity patterns that differentiate large from small objects, they do not evaluate whether these size-coding regions are also able to differentiate objects of different semantic subcategories. For example, patterns of activity in a particular brain region may be able to differentiate a wardrobe (large object) from a shoe (small object), but can it also differentiate the wardrobe (large object) from a sofa (another large object)? To answer this question, a recent fMRI study by Julian et al. (2017) extended upon previous univariate size findings (e.g., Konkle & Caramazza, 2013; Konkle & Oliva, 2012) using multivariate analyses to evaluate whether real-world object size is important for object category representations. To do so, they first created representational similarity matrices (RSMs), where individual cells of the matrix corresponded to the correlation value for each pair of object activity patterns, evaluating 10 big (e.g., arm chair, stove) and 10 small (e.g., hat, mobile phone) inanimate object categories. They then performed a split-half classification analysis, where they evaluated how well the pattern information in one half of the data could classify objects in the remaining half according to object size or category. In line with previous results (Konkle & Caramazza, 2013; Konkle & Oliva, 2011, 2012), they found differential regional activation corresponding to big and small objects, with the additional novel finding of greater similarity between objects of similar sizes, despite differences in object subcategory. They argued that this provided evidence for coding of object size within object categories and suggested that the distinction between big and small

objects may reflect a fundamental difference related to object stability. This is in line with previous findings which suggest that stable, landmark-qualities of large objects may drive activity in the typically scene-selective regions of cortex that seem to also show preferential activation to large objects (Troiani, Stigliani, Smith, & Epstein, 2014).

To date, the majority of studies evaluating the neural mechanisms underlying object processing have focused on spatially localising the corresponding activity in specific brain regions. Studies using fMRI typically localise patterns of BOLD activation in object-centred regions that span the ventral temporal lobe and occipital cortex, providing a temporally blurred snapshot of the response of these regions. Consequently, they are unable to account for distinct variations in the timing of patterns of activity associated with processing different object categories (Carlson, Tovar, Alink, & Kriegeskorte, 2013). Recently, studies using magnetoencephalography (MEG) have examined the temporal dynamics of object representations and found differences in the timing of activations associated with object categories. Information that differentiates objects at an exemplar level can be decoded early on in the time course of object processing, while more complex conceptual information related to semantic category, including the animacy dimension, is decodable later on in the neural signal from around 150 ms post-stimulus onset (Carlson et al., 2013; Cichy, Pantazis, & Oliva, 2014; Grootswagers, Wardle, & Carlson, 2016; for a review, see Contini, Wardle, & Carlson, 2017). Together with findings from fMRI studies showing category sensitive regions throughout the ventral stream (for a review, Grill-Spector & Weiner, 2014), MEG studies provide support for highly dynamic representations of object categories. While previous studies have evaluated the decoding time-course of animacy (Carlson et al., 2013; Chan, Halgren, Marinkovic, & Cash, 2011; Cichy et al., 2014), the temporal dynamics of real-world size and its relation to object animacy are yet to be established. If real-world size is also a primary organisation principle for object representations then it is plausible that it may have its own unique set of temporal dynamics, which emerge independently of object animacy in the decoding time-course.



Building on the proposed “tripartite” organisation of object representations (Konkle & Caramazza, 2013), the current study used MEG decoding to investigate the temporal dynamics of object representations in relation to the dimensions of real-world size and object animacy. We created a novel stimulus set which included 36 objects, divided evenly into animate, big-inanimate, and small-inanimate objects. Previous studies have noted that big-inanimate objects tend to be more spatially stable in the natural environment (e.g., a sofa or stove) and activate brain regions associated with processing non-manipulable, landmark-type objects (He et al., 2013; Julian et al., 2017; Troiani et al., 2014). Differences in manipulability of objects have also been shown to influence behavioural reaction times (Filliter, McMullen, & Westwood, 2005). We controlled for this apparent difference between big and small object categories by including objects designed for human interaction (and are thus, manipulable) in both the big- (e.g., car), and small- (e.g., scissors) inanimate object categories. By testing models related to these different dimensions of object processing using the representational similarity analysis (RSA) framework (Kriegeskorte, Mur, & Bandettini, 2008), we evaluated whether differences in multivariate patterns could be explained by a large-scale organisation related to object size in addition to object animacy.

## **2.3. Method**

### **2.3.1. Participants**

A total of 22 English-speaking healthy adults (13 female; mean age = 26.23 years;  $SD = 8.07$ ; range = 18 - 45) were recruited from the Macquarie University community. Informed written consent was obtained prior to participation, and all brain recordings were conducted in the KIT-Macquarie Brain Research (MEG) Laboratory. Participants were paid \$20/h for their time. All participants self-reported normal or corrected-to-normal vision, were free from significant medical conditions, and were not currently taking any neuroactive medications. This study was approved by the Macquarie University Human Research Ethics Committee.

### 2.3.2. Stimuli

Stimuli consisted of 36 naturalistic images of objects with a planned categorical structure (see Figure 1). The category structure had three tiers: the highest tier in the hierarchy contained the three main categories of interest in line with the findings of Konkle et al. (2013) (specifically animals, big-inanimate objects, and small-inanimate objects), with equal numbers of exemplars in each (12 exemplars per category). Below this tier, objects were grouped according to whether they were *interactive* (six exemplars) or *non-interactive* (six exemplars). The lowest tier category was only within the interactive object categories, whereby the interactive objects were divided into equal numbers of instruments (3 exemplars) and other objects (3 exemplars) which differed depending on size as follows: in the small-inanimate object category, these were hand-held tool-type objects, while in the big-inanimate object category, these items were vehicles of various forms. All stimuli were segmented object images, presented on a uniform grey background.


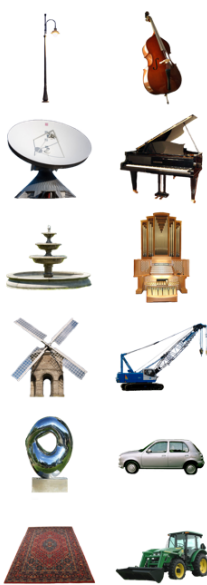

Animals	Big-inanimate objects	Small-inanimate objects
	<div> <div>  </div> <div> <div>Non-interactive</div> <div>Interactive</div> <div>Vehicles</div> </div> <div>Instruments</div> </div>	<div> <div>  </div> <div> <div>Non-interactive</div> <div>Interactive</div> <div>Instruments</div> <div>Tools</div> </div> </div>

Figure 1. Stimuli design and categorical structure. Stimuli are shown divided into the tripartite categorical structure of interest: animals, big-inanimate objects, and small-inanimate objects. Brackets show additional embedded subcategory groupings.

### **2.3.3. Experimental Design and Attention Task**


Prior to the experimental task, participants viewed a simple PowerPoint presentation of each of the 36 objects in their real-world context (i.e., participants were shown the object image prior to segmenting the objects from their backgrounds to create the stimuli for the experiment) to provide participants with information about the relative real-world size of the objects. This presentation was also used to ensure that participants successfully recognised all the objects prior to the experiment, with all participants demonstrating adequate knowledge of the objects in the stimulus set by naming and/or providing semantic information.

For the experimental task, participants completed six blocks of 216 trials, with the block order pseudorandomised for each participant. Within a single block, each of the 36 exemplars was presented six times in pseudorandom order, resulting in a total of 1296 trials per participant (36 exemplars x 6 presentations x 6 blocks). In an individual trial, participants viewed a single exemplar presented centrally on a grey background for 50 ms, with a random inter-trial interval ranging between 900 and 1200 ms, totaling approximately one hour of MEG recording time (see Figure 2). Participants were instructed to maintain fixation on central dot (diameter of 0.1 DVA), which remained on the screen at all times (with stimuli presented behind this dot). During the stimulus presentation, participants completed a simple attention task, where they were required to press a button whenever they saw the same image repeated consecutively. After a repeat trial, there was a minimum of three non-repeat trials, while the fourth and subsequent trials had a 30% chance of being a repeat, with this process resetting once a repeat had occurred. There were on average 242 repeats per participant (range 239 – 247) which remained in the analysed trials. Feedback was provided to the participant after stimulus offset (during the ISI) by a change in colour of the central fixation dot, which briefly (250 ms) turned green for correctly identified repetitions, and red for false-alarms or missed repetitions. Participants were also provided with summary statistics of their overall performance at the end of each block. The overall mean accuracy across participants

was 96% ( $SD = 5.1\%$ ), with an average participant reaction time of 514 ms ( $SD = 107$  ms; for correct trials only).

#### 2.3.4. Display Apparatus

Participants lay supine in the magnetically sealed recording chamber. Stimuli were projected onto a mirror using an InFocus IN5108 projector situated outside the chamber. This mirror reflected the image onto the ceiling of the room, located 113 cm above the participant (stimulus visual angle = 9 degrees). The experiment was run on a Dell PC desktop computer using MATLAB software (Natick, MA) and the Psychophysics Toolbox extensions (Brainard, 1997; Kleiner et al., 2007; Pelli, 1997).



*Figure 2.* Experiment design. Example showing two trials of the experimental design. ISI = interstimulus interval.

### 2.3.5. MEG Data Acquisition and Analysis

Using a 160-channel whole-head axial gradiometer (KIT, Kanazawa, Japan), continuous neuromagnetic data were acquired at a sampling rate of 1000 Hz, and filtered online from 0.03 to 200 Hz. MATLAB (Natick, MA) was used for all processing and statistical analyses of the data. Offline, we down-sampled the data to 200 Hz for the analysis, and created a single trial epoch with a time window from -100 ms to 600 ms relative to stimulus onset. We then applied Principal Components Analysis (PCA) to reduce the dimensionality of the epoched data, retaining the  $n$  components that accounted for 99% of the variance in the data. For each participant, the number of components retained ranged from 20 to 94 ( $M = 62.14$ ,  $SD = 21.49$ ).

**2.3.5.1. Classification analysis.** For each participant, we trained a classifier (linear discriminant analysis (LDA); Duda, Hart, & Stork, 2001) to discriminate MEG data (the PCA component weightings) for one exemplar from data for a second exemplar, then we repeated this process for each pair of exemplars. We employed a 10-fold cross validation procedure, where we trained the classifier on 90% of the data, and then measured classification performance using the remaining 10%, such that the classifier was never trained and tested on the same data. We repeated this process 10 times ensuring that all trials were used once as the test data. This results in a measure of how well the classifier can discriminate each exemplar identity from the PCA components for every possible pairwise comparison of the 36 stimuli. Classification accuracy is reported using d-prime ( $d'$ ).

**2.3.5.2. Representational similarity analysis (RSA).** We then used representational similarity analysis (RSA; Kriegeskorte, Mur, & Bandettini, 2008) to compare the differences in classification performance across exemplars, providing an indication of the structure of object representations as evidenced by similarities between object activity patterns. To do so, we firstly created a Representational Dissimilarity Matrix (RDM; Kriegeskorte, Mur, & Bandettini, 2008), where each cell contains the average classifier accuracy ( $d'$ ) for one pair of object exemplars, generating a separate RDM for each time point. This creates a matrix

of decoding performance that can then be compared to our model predictions (described below). The observed RDM for this experiment is a 36 x 36 matrix (see example frames in Figure 3B) that is symmetric along the negative diagonal. At each time point, we compared this observed RDM from each participant's MEG data with our model RDMs (models described in detail below, and see also Figure 6).

**2.3.5.3. Model Matrices.** To account for the potential impact of low-level visual similarity between exemplars on decoding performance, we first computed the Jaccard silhouette model RDM (Jaccard, 1901), which compares the abstract shape of each object by identifying the pixels the object occupies within the image. The model RDM is subsequently generated by computing the difference between the overlapping silhouette regions of two object images at a time, with each cell of the RDM corresponding to 'one minus the correlation value' for each exemplar pair, such that higher cell values indicate greater dissimilarity between the image silhouettes. We used Spearman's rank correlation coefficient to test the relationship between object shape and object decodability (see Results, Figure 4). To exclude the contribution of object shape, we partialled out this correlation between the Jaccard model and the observed RDM (partial Spearman's  $\rho$ ) when testing the relationship between our theoretical models and the observed data. We then used a nonparametric Wilcoxon signed rank test to assess significance at each time point, averaged across all participants. To control for false-positive results due to multiple comparisons we employed a false-discovery rate (FDR) threshold of 0.01.

In light of recent findings suggesting a tripartite organisation of object representations (Konkle & Caramazza, 2013), where real world size is reflected in inanimate object representations (Julian et al., 2017; Konkle & Oliva, 2011, 2012), we created the following eight theoretical models (shown in Figure 6) to evaluate the time course of object real-world size and how this relates to the animacy time course:

#### **2.3.5.3.1. Main models**

Model 1) *Animacy model*: Dichotomous model based on organizational principles proposed in previous studies (Caramazza & Shelton, 1998; Carlson et al., 2013; Cichy et al., 2014; Kriegeskorte, Mur, & Bandettini, 2008) comparing animals to all inanimate objects (big- and small-inanimate objects combined).

Model 2) *Tripartite model*: Depicting three groups of object representations, whereby objects are not only coded according to animacy, but inanimate object representations are also organized according to real-world size: big or small (Julian et al., 2017; Konkle & Caramazza, 2013; Konkle & Oliva, 2011, 2012). As such, this model depicts representations grouped according to whether they are animate, big-inanimate, or small-inanimate objects.

#### **2.3.5.3.2. Binary category models**

The binary category models evaluated a subset of the data, by including only two of the three major categorical distinctions in the stimulus set. Previous studies have identified a size distinction when inanimate objects are evaluated independently of animate objects ((Julian et al., 2017; Konkle & Oliva, 2012), therefore, we expected to see this represented in the decoding time-course.

Model 3) *Animacy versus big-inanimate model*: This binary category model compares only the animals and big-inanimate objects in the data set, with small-inanimate objects removed.

Model 4) *Animacy versus small-inanimate model*: This binary category model compares only the animals and small-inanimate objects in the data set, with big-inanimate objects removed.

Model 5) *Big- versus small-inanimate model*: This binary category, inanimate size model compares only the big-inanimate and small-inanimate objects in the data set, with animal objects removed.

#### **2.3.5.3.3. Cluster models**

The three cluster models evaluated whether representations within a single category show more similar patterns of activity compared to all other objects, which will not show

clustering of their representations. Studies have shown clustering of object representations related to object animacy in ventral temporal cortex (Clarke & Tyler, 2014; Kriegeskorte, Mur, Ruff, et al., 2008), so if real-world size is a similarly large-scale organisation principle, it is plausible that a similar clustering of representations may be seen for each of the size dimensions.

Model 6) *Animal cluster model*: This single category model predicts that animal object representations would be more similar to each other while being dissimilar to all other inanimate objects, which are predicted to not cluster together into groups.

Model 7) *Big-inanimate cluster model*: This single category model predicts that big-inanimate object representations will be more similar to each other while being dissimilar to all other animal and small-inanimate objects. Animal and small-inanimate object are predicted to show no group clustering.

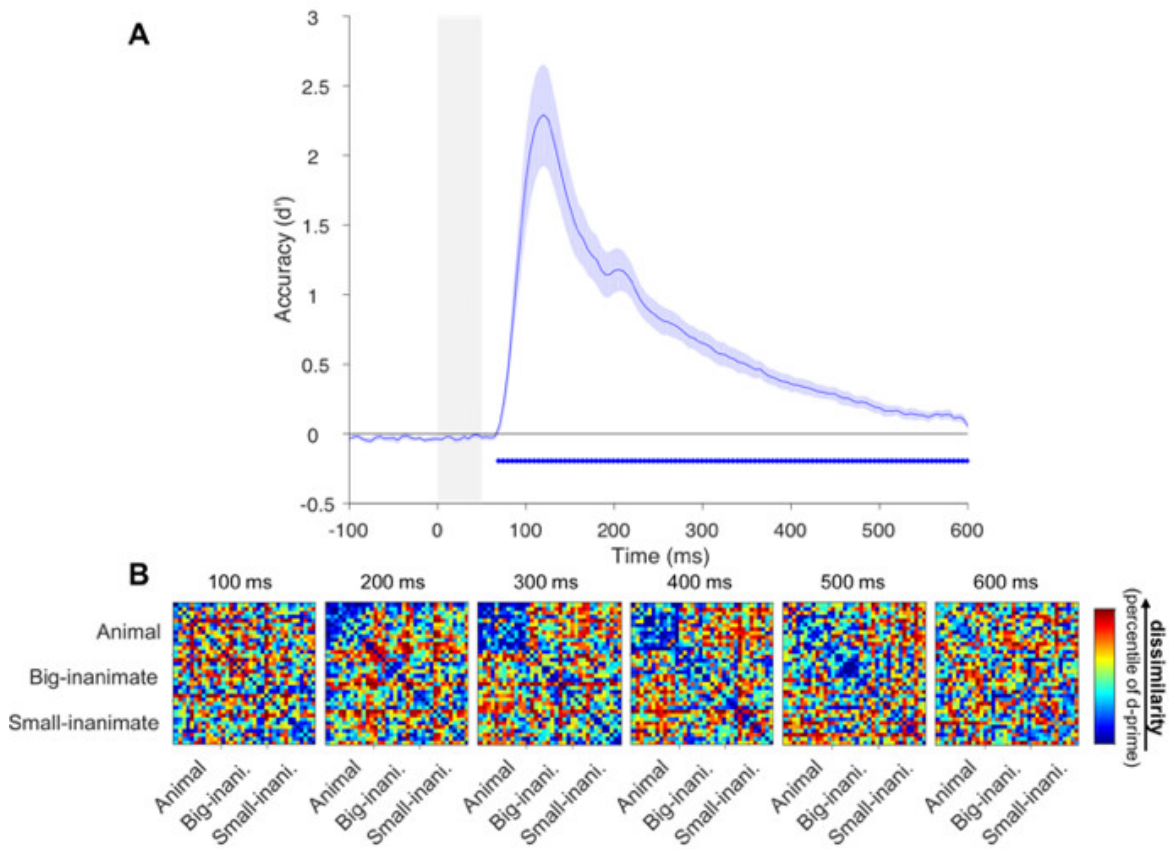
Model 8) *Small-inanimate cluster model*: This single category model predicts that small-inanimate object representations will be more similar to each other while being dissimilar to all other animal and big-inanimate objects. Animal and big-inanimate object are predicted to show no group clustering.

## 2.4. Results

We recorded participants' neuromagnetic signals using MEG while they viewed a series of briefly presented object stimuli to test predictions about the categorical organisation of neural object representations. Firstly, we evaluated whether we could decode the individual object images from participants' MEG recordings throughout the time course of visual processing (Figure 3A). To do so, we trained and tested classifiers on their ability to differentiate each pair of the 36 objects in our stimulus set from the MEG activation patterns, creating an RDM at each time point (see Figure 3B for a sample of these RDMs throughout the time course). Consistent with image decoding results from previous MEG studies (Carlson et al., 2013; Cichy et al., 2014; Redcay & Carlson, 2014; Ritchie, Tovar, & Carlson, 2015), we found significant decoding of object exemplars from 70 ms, which is sustained




throughout the time-course evaluated (i.e., until the final time-point of the trial epoch at 600 ms post-stimulus onset; one-tailed Wilcoxon signed-rank test, corrected for multiple comparisons using an FDR threshold of 0.01). Exemplar decoding performance peaks at 120 ms post-stimulus onset. When inspecting the RDMs across time (Figure 3B), coherent structure begins to appear in the RDMs late in the time course (as seen from 200 ms onwards in Figure 3B), with a cool-coloured square emerging in the top left of the plot, indicating greater similarity in activity patterns for the animal objects. This structure is sustained for a considerable period throughout the time course (clearly apparent in the 200 – 400 ms RDMs shown in Figure 3B).



*Figure 3.* Overall decoding time course and average decoding RDMs over time for all exemplars. A) Exemplar decoding performance (measured in  $d'$ ) as a function of time, averaged across decoding performance for all exemplar pairs. Grey bar indicates the period the image was on the screen. Error bars indicate 95% between-subject confidence intervals. Blue dots along the x-axis indicate time points at which decoding performance was significantly above chance (Wilcoxon signed-rank test, corrected across time points for multiple comparisons with an FDR threshold of 0.01). B) Average rank-ordered decodability of exemplar pairs across participants. RDMs shown at 100 ms intervals throughout the decoding time course. Individual cells in the RDMs represent the degree of dissimilarity between pair-wise comparisons of response patterns, reported as the percentile of  $d'$ -prime. Warm colours (red to yellow) indicate greater dissimilarity (more distinct neural representations), while cool colours (green to blue) indicate greater similarity (overlapping neural representations).

### 2.4.1. Model Testing Across Time

Once we established that the individual objects could be decoded from the MEG signal, we first evaluated the shape-dependent Jaccard model to test how well low-level shape properties correlated with the neural representations. It is important to account for differences in low-level features inherent to the different object images as it is possible for the classifier to simply rely on these conflated perceptual differences to successfully decode object category membership (Baldassi et al., 2013; Kay, Naselaris, Prenger, & Gallant, 2008; but see also Kaiser, Azzalini, & Peelen, 2016; Proklova, Kaiser, & Peelen, 2016; Wardle & Ritchie, 2014). The Jaccard RDM and time-series correlation with the MEG data are shown in Figure 4, with the onset of significant decoding from 75 ms post-stimulus onset, and two correlation peaks at 90 ms and 120 ms (peak Spearman's  $\rho$  of 0.25 at both time points).



*Figure 4.* Jaccard model correlation over time. Left panel shows the Jaccard silhouette model generated by comparing the overlapping pixels for each exemplar pair. Cells in the RDM represent ‘1 minus the proportion of overlapping pixels’, with warm colours indicating greater dissimilarity, and cool colours indicating greater similarity between the object shapes. The right panel shows the Spearman correlation between the Jaccard model and the time-series MEG data. Shaded error bars indicate the 95% confidence interval of the between-subject means. Coloured dots along the x-axis indicate time points at which the model correlation is significantly above chance (Wilcoxon signed-rank test, corrected for multiple comparisons using an FDR threshold of 0.01).

As there was a significant correlation between the Jaccard silhouette model and the MEG data throughout the time course, this suggests that low-level shape related information accounts for a considerable degree of variation in the MEG signal that is being picked up by the classifier to decode object identity. Therefore, when evaluating other models, we considered their partial correlation (given the correlation of the data with the Jaccard RDM) as a partial correlation to control for object shape when evaluating each of the theoretical models of interest.

Results from the time-series model testing are shown in Figure 5, with significant onset and peak-decoding times listed in Table 1 (note: these should not be interpreted as absolute timings<sup>1</sup>). We first sought to replicate the established animacy dichotomy identified in previous MEG object decoding studies (Carlson et al., 2013; Cichy et al., 2014; Grootswagers et al., 2016). Our results similarly show that object animacy provides a good account of the pattern of decoding performance (Figure 5A). In addition, when evaluating the tripartite model (Figure 5B), based on the results of Konkle et al. (2013), we also see a significant relationship between the model and the MEG data, with this relationship appearing more sustained throughout the time-course than that seen for animacy. If this model is truly capturing the brain representations better than a model based on an animacy dichotomy, then we should also see a significant relationship with dichotomous models (using a subset of the data), evaluating every possible dichotomous comparison between the

---

<sup>1</sup> After the completion of this study, a delay in the stimulus projector timing was identified such that images were being displayed to participants after the stimulus onset trigger was being sent to the MEG recording. Timing delays of up to 60 ms have been identified, however due to changes in the experimental equipment, we are unable to determine the precise delay for the present study. Consequently, absolute timings are uncertain, with results indicative of relative timings only. This delay was accounted for in future experiments using a photodiode on the stimulus presentation screen.

animal, big-inanimate, and small-inanimate object categories (Figure 5, C-E). These dichotomous models also account for the potential influence of unbalanced category numbers, as they have equal stimulus numbers in each category (cf. animacy model, which has only 12 animal objects compared to 24 inanimate). Upon evaluation of these three dichotomous comparisons, both the animal vs. big-inanimate (Figure 5C), and animal vs. small-inanimate object (Figure 5D) models show a significant correlation throughout the time-course, however the model comparing only big- and small-inanimate objects (Figure 5E) does not reach significance at any time-point. This result suggests that the animate objects are likely driving the results seen in the animacy (A) and tripartite (B) model results.

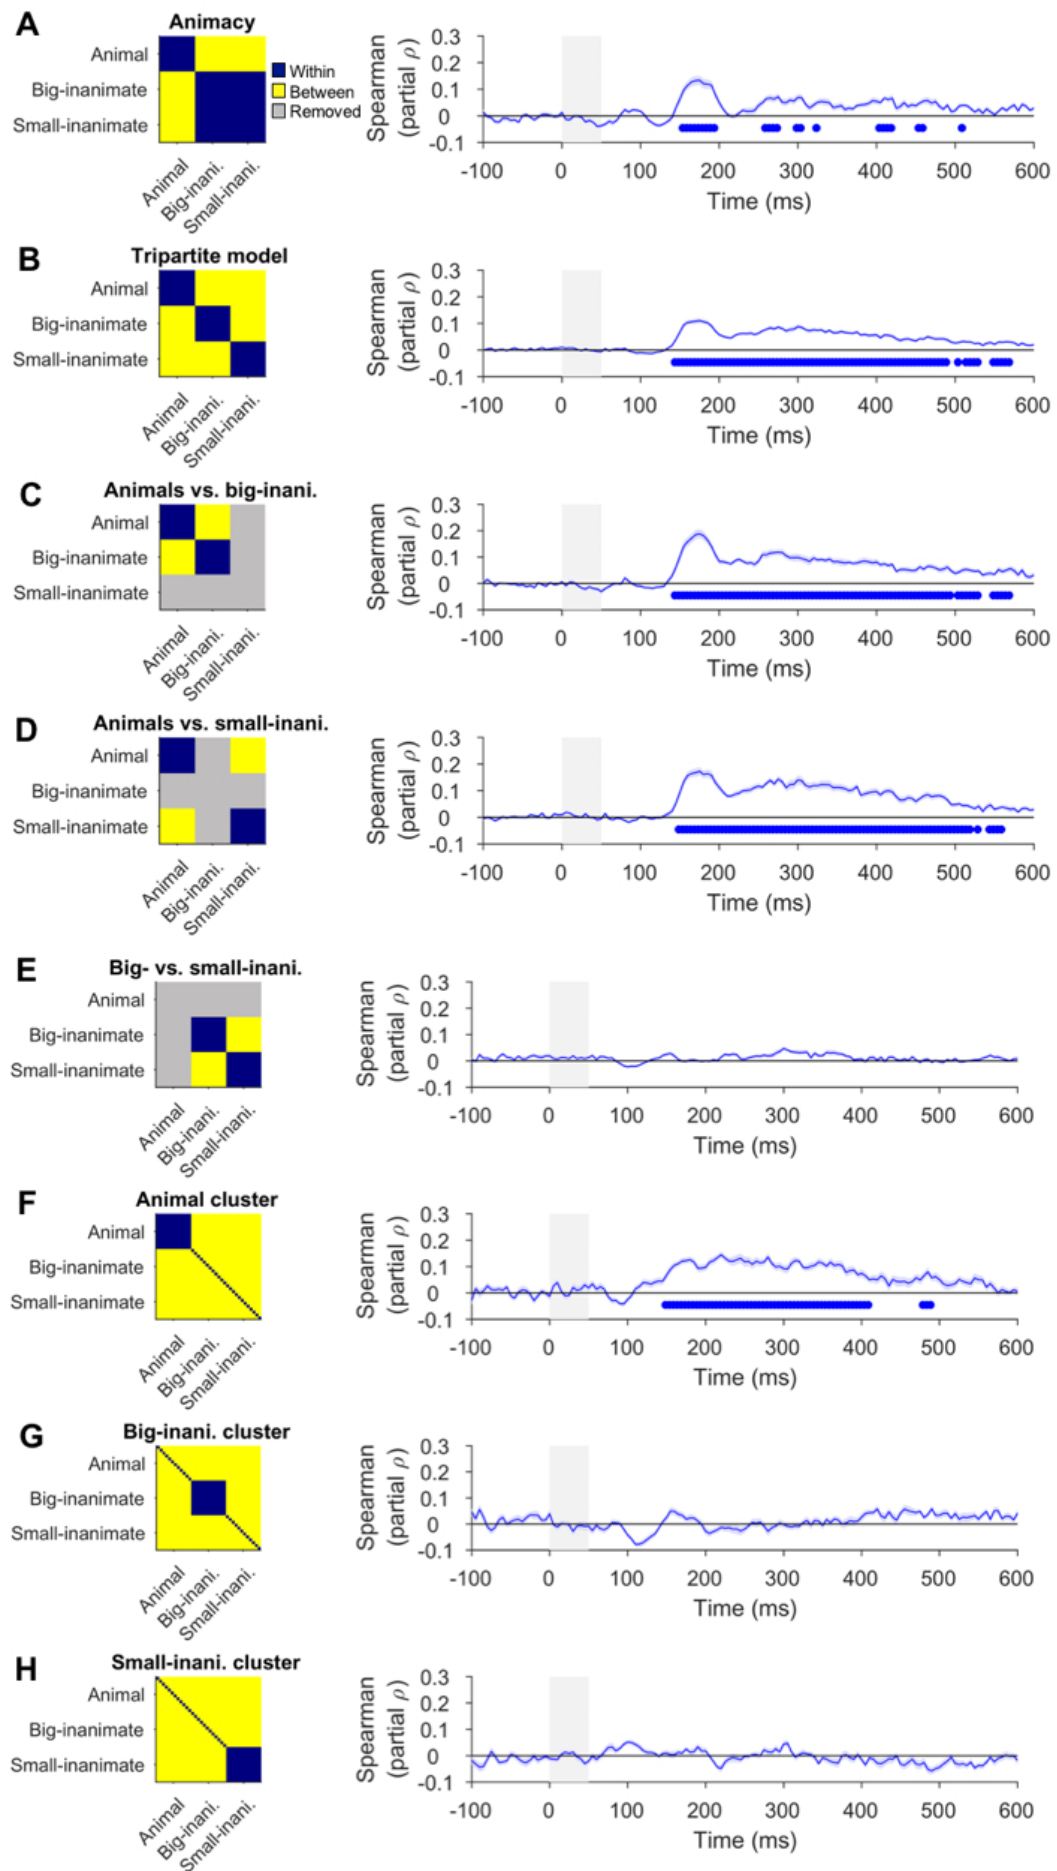
Table 1

*Onsets and peaks of significant correlation time points for each of the theoretical models*

Model	Onset (ms)	Peak (ms)
A) Animacy	155	175
B) Tripartite model	145	175
C) Animals vs. big-inani.	145	175
D) Animals vs. small-inani.	150	175
E) Big-inani. vs. small-inani.	-	-
F) Animal cluster	150	220
G) Big cluster	-	-
H) Small cluster	-	-

*Note:* inani. = inanimate objects. All reported values are significant, correcting for multiple comparisons using an FDR threshold of 0.01. Where no value is present, results did not reach significance at the desired threshold.

*Figure 5.* Representational similarity analysis with model testing. *Left panels:* The theoretical RDM models under investigation. Inani. = inanimate objects. Model colours indicate the predicted degree of similarity between object pairs, with blue indicating more similar object representations and yellow indicating greater dissimilarity. Comparisons coded in grey were removed for the purposes of that model prediction. Each model was correlated with the MEG classifier data at each time point (right panel), and included the Jaccard model RDM as a partial correlation to control for the effects of object shape. *Right panels:* Shaded error bars indicate the 95% confidence interval of the between-subject means. Coloured dots along the x-axis indicate time points at which the correlation for each model is significantly above chance (Wilcoxon signed-rank test, corrected for multiple comparisons using an FDR threshold of 0.01).



However, it is interesting to note that the tripartite model (Figure 5B) appears to provide a better, more sustained account of the neural signal than the animacy model (Figure 5A). The only difference between these two models is that the animacy model groups both the big and small objects together into one ‘inanimate object’ category; predicting that there will be greater similarity between these objects (irrespective of size) and that these objects will show greater dissimilarity compared to animate objects. The tripartite model on the other hand, predicts that for the three categories of animate, big-inanimate, and small-inanimate objects, each individual category will show greater within-category similarity, and greater dissimilarity when compared to the other two object categories. As a result, the tripartite model predicts more variability in activation patterns within the ‘inanimate objects’ category, and this provides a better account of the brain activity patterns observed, even though we do not see an effect for size when evaluated as a separate subset.

To further investigate how much each of the three object categories in the tripartite model drive the observed effects, we evaluated a series of cluster models (Figure 5, F – H). These models predict that objects within a category cluster show more similar patterns of activity to other same-category members than to all other objects, and that all other objects show more disparate object representations (i.e., show no clustering). Results from these analyses only showed clustering of animal object representations (Figure 5F), with big- (Figure 5G) and small-inanimate objects (Figure 5H) showing no clustering of object representations throughout the time-course. This suggests that there is greater cohesion of object representations within the animate object category, whereas object representations for inanimate objects may be more disparate. This may reflect a difference in how each of these categories are defined, since the definition for what constitutes an animate object is relatively constrained, with greater conceptual and visual similarity between individual category members than what may be expected for objects in both the big and small inanimate object categories.



## 2.5. Discussion

The focus of this study was to examine the time-course of object representations related to real-world size, and specifically how the time-course of object size relates to that of object animacy. Previous fMRI studies have established activation related to big- and small-inanimate objects in cortical regions adjacent to those which respond to animate objects (Julian et al., 2017; Konkle & Caramazza, 2013; Konkle & Oliva, 2012), however, the temporal evolution of size-based object representations has yet to be established. As object size is often related to object stability and manipulability (with large objects typically more stable and non-manipulable compared to small objects; Mullally & Maguire, 2011; Troiani et al., 2014), we designed our stimulus set to include objects designed for human interaction in both the big- and small-inanimate object categories. Controlling for low-level properties related to object shape, we replicated previous MEG findings which show that animacy representations emerge relatively late and remain sustained throughout the time course of object processing (Carlson et al., 2013; Cichy et al., 2014), however we observed no evidence for representations related to object size in the patterns of MEG activation.

That we identified patterns of activation reflecting object animacy yet did not find evidence for category representations organized by real-world size raises the question of why there is a discrepancy between the current MEG study and the fMRI findings reported by others (see Julian et al., 2017; Konkle & Caramazza, 2013; Konkle & Oliva, 2012). Here we discuss several possible reasons for this result.

One possibility is that the object size organisation is less accessible using MEG than fMRI. As studies have identified activation consistent with object size in regions adjacent to those associated with object animacy (Julian et al., 2017; Konkle & Caramazza, 2013; Konkle & Oliva, 2012) we hypothesized that these would be similarly detectable using MEG. Since we were unable to detect any size-related modulations in the signal, it may be that representations associated with object animacy are comparatively more widespread throughout the brain (Haxby et al., 2001), allowing for better detection when using whole-

brain analyses as in the current study, while potentially missing more localised activation, as may be the case for real-world size representations. Relatedly, the activation signal associated with object size may be comparatively weaker than that for object animacy. As a result, MEG may not have the power to detect these subtle modulations, particularly given the potential for competing information from multiple sensors which can affect the signal to noise ratio (for an example of this, see Hebart, Bankson, Harel, Baker, & Cichy, 2017).

On the contrary, a recent study (Khaligh-Razavi, Cichy, Pantazis, & Oliva, 2018) conducted after we had obtained our data, was able to detect real-world size modulations in the MEG signal, suggesting potential effect size issues related to the MEG signal are less likely to explain our inability to detect these representations. This study simultaneously compared size- and animacy-induced spatiotemporal activations obtained from fMRI and MEG recordings using a similarity-based fusion approach (Cichy, Pantazis, & Oliva, 2016). This study showed significant decoding of object animacy and real-world size throughout the time-course of object processing, however upon qualitative inspection of their results (this was not tested statistically), the effect for size appears to be somewhat weaker than that of animacy. This difference in the strength of the relationship between object size and animacy was also reflected in the tripartite results of Konkle et al. (2013; additional results not included in paper; T. Konkle, personal communication, January 21, 2015), with a considerable discrepancy between the beta weights associated with animacy and size dimensions in their fMRI results, suggesting that response modulations associated with object size are comparatively weaker than those identified for object animacy. In light of these findings, it is likely that this difference in the strength of size-based modulations similarly impacted our ability to detect these subtle differences, as we also used a comparatively limited stimulus set (36 stimuli compared to 240 in the study by Konkle & Caramazza, 2013, and 118 in Khaligh-Razavi et al., 2018. Interestingly, Khaligh-Razavi et al. (2018) showed a real-world size effect in both inanimate objects and animate objects, unlike the tripartite model (Konkle & Carramazza, 2013) which suggested the size effect

was limited to inanimate objects. Further investigation of the representation of real-world size is needed to establish how this factor is differentially coded within animate and inanimate object categories.

Another important point for consideration is whether the effects that we see across previous studies, in fact, reflect a size dimension, or alternatively, some other covarying factor that may similarly explain these apparent size-based modulations in brain representations. When trying to account for real-world size effects, researchers often highlight the problem of differentiating the individual contributions of additional factors that are often conflated with large inanimate object properties, such as *object fixedness* (i.e., how fixed an object is within its environment) or *contextual stability* (i.e., how strongly associated an object is to its contextual setting) (Julian et al., 2017; Troiani et al., 2014). If the coding of object size is mostly related to object stability and landmark presence, then it makes sense that a size distinction would not appear when comparing big and small animals (Konkle & Caramazza, 2013), as animals are inherently mobile and consequently should not activate regions associated with stable-, landmark-type feature preferences. Similarly, stimulus feature properties may also pose a potential confound when evaluating real-world size representations, as a recent study (Long, Konkle, Cohen, & Alvarez, 2016) showed that visual search performance was enhanced when a target differed in real-world size from the distractor objects, with this effect eliminated when mid-level perceptual features of the objects (e.g., texture and contour information) were removed. This suggests that real-world size may dictate certain form properties of objects in such a way that distinguished big from small objects at a perceptual level. These potential semantic and physical-feature confounds have implications for the interpretation of size effects, and warrant further investigation to tease apart any unique contributions of these differing object properties.

The idea of *manipulability* has also been posited as a potential driver of the size distinction, with activity in the PPA (which shows selectivity for large objects) showing preferential activation in response to large non-manipulable objects compared to both tools

and animals (He et al., 2013). Konkle et al. (2013) also note a fundamental difference in how humans interact with animate versus inanimate objects, with inanimate object size dictating whether we interact with objects using our hands (small inanimate objects) or our whole body (big inanimate objects), with this clear size/function association not typically evident for animate objects. This supports the idea of a tripartite organisation for object representations that does not necessarily rely on the coding of object size. Here we specifically tried to control for this by including objects in our stimulus set that are designed for human interaction in both the big and small object categories. Therefore, if the activity associated with object size identified in previous studies is, in part, related to an imbalance in object interactivity across big and small object categories, then this may provide an alternative explanation as to why we do not see a similar effect in the present study. It would be of interest for future studies to investigate whether properties related to object interactivity or object size better explain the size-related brain representations identified in previous studies.

While a tripartite organization of object representations makes theoretical sense, and with emerging evidence suggesting that factors related to size or object stability/manipulability may contribute to such an organisation, evidence from a patient case study by Masullo et al. (2012) suggests this distinction may not be quite so well-defined. By assessing the semantic memory of a patient (MR) who demonstrated severe apperceptive agnosia, they discovered that MR made significantly more errors when asked questions pertaining to musical instruments compared to other inanimate objects, and the number of errors was similar to that made for animate objects. This result does not neatly fit into a tripartite organization of object processing, as it appears that the breakdown in MR's object processing similarly affects the instrument (i.e., an inanimate object subcategory) and animate representations, which are predicted to be somewhat well differentiated by the tripartite model. The authors suggest their findings support models of object processing that differentially weight the sensory and action-related features of objects for categorisation

(Warrington & McCarthy, 1983, 1987; Warrington & Shallice, 1984; for reviews of studies supporting this theory see Gainotti, 2006, and Martin, 2007). Similarly, Bracci and Peelen (2013) used fMRI to identify overlapping responses in the lateral occipitotemporal cortex to parts of the body related to the execution of actions (e.g., hands, limbs), and inanimate objects that are used to extend upon the body (e.g., tools, sporting equipment). They conclude that these object representations appear to also contain information about the relationship between object and body. Together, the results from these studies highlight some potential difficulties in defining dimensions related to object animacy and size when attempting to account for variations in object representations. Further studies need to employ carefully designed stimulus sets in order to tease apart the contributions of these various object properties, and test models that evaluate a range of object-related features in order to establish what underlying concepts are coded in object representations.

In sum, using a stimulus set where we aimed to balance across exemplars for aspects of object interactivity, we replicated previous findings which identify activity patterns consistent with an animacy organisation throughout the time course of object processing, however we were unable to detect object real-world size modulations in our MEG data. Establishing what aspects of neural object representations are uniquely related to real-world size remains a challenge for this field, as a number of object properties, such as stability and manipulability, are conflated with object size, posing a challenge for the creation of carefully controlled stimulus sets.

In the following chapter, the focus shifts towards investigating the intricacies of object animacy, with the aim of testing alternative explanations that may provide a better account of the organisation of object representations. Recent studies have identified a biological classes continuum within the animate subcategory (Connolly et al., 2012; Sha et al., 2015), whereby object representations exist along a spectrum from humans to inanimate objects. We investigate this further by comparing the efficacy of a biological classes

continuum model to a range of models representing factors related to the human condition, as well as one specifically related to the broad idea of human-similarity.

Author contributions:

E.C, M.W, T.C designed the study

E.C collected the data

E.C, E.G, T.C conducted the analysis

E.C, E.G, S.W, M.W, T.C interpreted the results

E.C wrote the manuscript

E.C, E.G, S.W, M.W provided critical revisions on the manuscript

## 2.6. References

- Baldassi, C., Alemi-Neissi, A., Pagan, M., DiCarlo, J. J., Zecchina, R., & Zoccolan, D. (2013). Shape Similarity, Better than Semantic Membership, Accounts for the Structure of Visual Object Representations in a Population of Monkey Inferotemporal Neurons. *PLoS Computational Biology*, 9(8), e1003167. <https://doi.org/10.1371/journal.pcbi.1003167>
- Beauchamp, M. S., Lee, K. E., Haxby, J. V., & Martin, A. (2002). Parallel visual motion processing streams for manipulable objects and human movements. *Neuron*, 34(1), 149–159.
- Biederman, I. (1987). Recognition-by-components: a theory of human image understanding. *Psychological Review*, 94(2), 115.
- Bracci, S., Cavina-Pratesi, C., Ietswaart, M., Caramazza, A., & Peelen, M. V. (2012). Closely overlapping responses to tools and hands in left lateral occipitotemporal cortex. *Journal of Neurophysiology*, 107(5), 1443–1456. <https://doi.org/10.1152/jn.00619.2011>
- Bracci, S., & Peelen, M. V. (2013). Body and Object Effectors: The Organization of Object Representations in High-Level Visual Cortex Reflects Body–Object Interactions. *Journal of Neuroscience*, 33(46), 18247–18258. <https://doi.org/10.1523/JNEUROSCI.1322-13.2013>
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 433–436.
- Caramazza, A., & Shelton, J. R. (1998). Domain-Specific Knowledge Systems in the Brain: The Animate-Inanimate Distinction. *Journal of Cognitive Neuroscience*, 10(1), 1–34. <https://doi.org/10.1162/089892998563752>
- Carlson, Tovar, D. A., Alink, A., & Kriegeskorte, N. (2013). Representational dynamics of object vision: The first 1000 ms. *Journal of Vision*, 13(10), 1–1. <https://doi.org/10.1167/13.10.1>

- Chan, A. M., Halgren, E., Marinkovic, K., & Cash, S. S. (2011). Decoding word and category-specific spatiotemporal representations from MEG and EEG. *NeuroImage*, 54(4), 3028–3039. <https://doi.org/10.1016/j.neuroimage.2010.10.073>
- Chao, L. L., Haxby, J. V., & Martin, A. (1999). Attribute-based neural substrates in temporal cortex for perceiving and knowing about objects. *Nature Neuroscience*, 2(10), 913–919.
- Chao, L. L., & Martin, A. (2000). Representation of Manipulable Man-Made Objects in the Dorsal Stream. *NeuroImage*, 12(4), 478–484. <https://doi.org/10.1006/nimg.2000.0635>
- Cichy, R. M., Pantazis, D., & Oliva, A. (2014). Resolving human object recognition in space and time. *Nature Neuroscience*, 17(3), 455–462. <https://doi.org/10.1038/nn.3635>
- Cichy, R. M., Pantazis, D., & Oliva, A. (2016). Similarity-based fusion of MEG and fMRI reveals spatio-temporal dynamics in human cortex during visual object recognition. *Cerebral Cortex*, 26(8), 3563–3579.
- Clarke, A., & Tyler, L. K. (2014). Object-Specific Semantic Coding in Human Perirhinal Cortex. *Journal of Neuroscience*, 34(14), 4766–4775. <https://doi.org/10.1523/JNEUROSCI.2828-13.2014>
- Connolly, A. C., Guntupalli, J. S., Gors, J., Hanke, M., Halchenko, Y. O., Wu, Y.-C., ... Haxby, J. V. (2012). The Representation of Biological Classes in the Human Brain. *Journal of Neuroscience*, 32(8), 2608–2618. <https://doi.org/10.1523/JNEUROSCI.5547-11.2012>
- Contini, E. W., Wardle, S. G., & Carlson, T. A. (2017). Decoding the time-course of object recognition in the human brain: From visual features to categorical decisions. *Neuropsychologia*, 105(Supplement C), 165–176. <https://doi.org/10.1016/j.neuropsychologia.2017.02.013>



- Downing, P., Jiang, Y., Shuman, M., & Kanwisher, N. (2001). A cortical area selective for visual processing of the human body. *Science*, 293(5539), 2470–2473.
- Duda, R. O., Hart, P. E., & Stork, D. G. (2001). *Pattern Classification*. John Wiley & Sons.
- Epstein, R., Harris, A., Stanley, D., & Kanwisher, N. (1999). The parahippocampal place area: Recognition, navigation, or encoding? *Neuron*, 23(1), 115–125.
- Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature*, 392(6676), 598–601.
- Filliter, J. H., McMullen, P. A., & Westwood, D. (2005). Manipulability and living/non-living category effects on object identification. *Brain and Cognition*, 57(1), 61–65. <https://doi.org/10.1016/j.bandc.2004.08.022>
- Gainotti, G. (2006). Anatomical functional and cognitive determinants of semantic memory disorders. *Neuroscience & Biobehavioral Reviews*, 30(5), 577–594.
- Grill-Spector, K., & Malach, R. (2004). THE HUMAN VISUAL CORTEX. *Annual Review of Neuroscience*, 27(1), 649–677. <https://doi.org/10.1146/annurev.neuro.27.070203.144220>
- Grill-Spector, K., & Weiner, K. S. (2014). The functional architecture of the ventral temporal cortex and its role in categorization. *Nature Reviews. Neuroscience*, 15(8), 536–548. <https://doi.org/10.1038/nrn3747>
- Grootswagers, T., Wardle, S. G., & Carlson, T. A. (2016). Decoding Dynamic Brain Patterns from Evoked Responses: A Tutorial on Multivariate Pattern Analysis Applied to Time Series Neuroimaging Data. *Journal of Cognitive Neuroscience*, 29(4), 677–697. [https://doi.org/10.1162/jocn\\_a\\_01068](https://doi.org/10.1162/jocn_a_01068)
- Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and Overlapping Representations of Faces and Objects in Ventral Temporal Cortex. *Science*, 293(5539), 2425–2430. <https://doi.org/10.1126/science.1063736>

- He, C., Peelen, M. V., Han, Z., Lin, N., Caramazza, A., & Bi, Y. (2013). Selectivity for large nonmanipulable objects in scene-selective visual cortex does not require visual experience. *NeuroImage*, 79(Supplement C), 1–9.  
<https://doi.org/10.1016/j.neuroimage.2013.04.051>
- Hebart, M. N., Bankson, B. B., Harel, A., Baker, C. I., & Cichy, R. M. (2017). Representational dynamics of task context and its influence on visual object processing. *BioRxiv*, 153684.
- Huth, A. G., Nishimoto, S., Vu, A. T., & Gallant, J. L. (2012). A Continuous Semantic Space Describes the Representation of Thousands of Object and Action Categories across the Human Brain. *Neuron*, 76(6), 1210–1224.  
<https://doi.org/10.1016/j.neuron.2012.10.014>
- Ishai, A., Ungerleider, L. G., Martin, A., Schouten, J. L., & Haxby, J. V. (1999). Distributed representation of objects in the human ventral visual pathway. *Proceedings of the National Academy of Sciences*, 96(16), 9379–9384.  
<https://doi.org/10.1073/pnas.96.16.9379>
- Jaccard, P. (1901). *Etude comparative de la distribution florale dans une portion des Alpes et du Jura*. Impr. Corbaz.
- Julian, J. B., Ryan, J., & Epstein, R. A. (2017). Coding of Object Size and Object Category in Human Visual Cortex. *Cerebral Cortex*, 27(6), 3095–3109.  
<https://doi.org/10.1093/cercor/bhw150>
- Kaiser, D., Azzalini, D. C., & Peelen, M. V. (2016). Shape-independent object category responses revealed by MEG and fMRI decoding. *Journal of Neurophysiology*, 115(4), 2246–2250. <https://doi.org/10.1152/jn.01074.2015>
- Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The Fusiform Face Area: A Module in Human Extrastriate Cortex Specialized for Face Perception. *Journal of Neuroscience*, 17(11), 4302–4311.

- Kay, K. N., Naselaris, T., Prenger, R. J., & Gallant, J. L. (2008). Identifying natural images from human brain activity. *Nature*, 452(7185), 352–355.  
<https://doi.org/10.1038/nature06713>
- Khaligh-Razavi, S.-M., Cichy, R. M., Pantazis, D., & Oliva, A. (2018). Tracking the Spatiotemporal Neural Dynamics of Real-world Object Size and Animacy in the Human Brain. *Journal of Cognitive Neuroscience*, 1–18.
- Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., Broussard, C., & others. (2007). What's new in Psychtoolbox-3. *Perception*, 36(14), 1.
- Konkle, T., & Caramazza, A. (2013). Tripartite Organization of the Ventral Stream by Animacy and Object Size. *Journal of Neuroscience*, 33(25), 10235–10242.  
<https://doi.org/10.1523/JNEUROSCI.0983-13.2013>
- Konkle, T., & Oliva, A. (2011). Canonical visual size for real-world objects. *Journal of Experimental Psychology: Human Perception and Performance*, 37(1), 23.
- Konkle, T., & Oliva, A. (2012). A Real-World Size Organization of Object Responses in Occipitotemporal Cortex. *Neuron*, 74(6), 1114–1124.  
<https://doi.org/10.1016/j.neuron.2012.04.036>
- Kriegeskorte, N., Mur, M., & Bandettini, P. A. (2008). Representational similarity analysis - connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, 2(4). <https://doi.org/10.3389/neuro.06.004.2008>
- Kriegeskorte, N., Mur, M., Ruff, D. A., Kiani, R., Bodurka, J., Esteky, H., ... Bandettini, P. A. (2008). Matching Categorical Object Representations in Inferior Temporal Cortex of Man and Monkey. *Neuron*, 60(6), 1126–1141.  
<https://doi.org/10.1016/j.neuron.2008.10.043>
- Long, B., Konkle, T., Cohen, M. A., & Alvarez, G. A. (2016). Mid-level perceptual features distinguish objects of different real-world sizes. *Journal of Experimental Psychology: General*, 145(1), 95.

- Mahon, B. Z., & Caramazza, A. (2011). What drives the organization of object knowledge in the brain? *Trends in Cognitive Sciences*, 15(3), 97–103.  
<https://doi.org/10.1016/j.tics.2011.01.004>
- Martin, A. (2007). The Representation of Object Concepts in the Brain. *Annual Review of Psychology*, 58(1), 25–45.  
<https://doi.org/10.1146/annurev.psych.57.102904.190143>
- Martin, A., Wiggs, C. L., Ungerleider, L. G., & Haxby, J. V. (1996). Neural correlates of category-specific knowledge. Retrieved from  
<http://www.nature.com/nature/journal/v379/n6566/abs/379649a0.html>
- Masullo, C., Piccininni, C., Quaranta, D., Vita, M. G., Gaudino, S., & Gainotti, G. (2012). Selective impairment of living things and musical instruments on a verbal ‘Semantic Knowledge Questionnaire’ in a case of apperceptive visual agnosia. *Brain and Cognition*, 80(1), 155–159. <https://doi.org/10.1016/j.bandc.2012.06.002>
- Mullally, S. L., & Maguire, E. A. (2011). A New Role for the Parahippocampal Cortex in Representing Space. *Journal of Neuroscience*, 31(20), 7441–7449.  
<https://doi.org/10.1523/JNEUROSCI.0267-11.2011>
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10(4), 437–442.
- Proklova, D., Kaiser, D., & Peelen, M. V. (2016). Disentangling Representations of Object Shape and Object Category in Human Visual Cortex: The Animate–Inanimate Distinction. *Journal of Cognitive Neuroscience*, 1–13.  
[https://doi.org/10.1162/jocn\\_a\\_00924](https://doi.org/10.1162/jocn_a_00924)
- Redcay, E., & Carlson, T. A. (2014). Rapid neural discrimination of communicative gestures. *Social Cognitive and Affective Neuroscience*.  
<https://doi.org/10.1093/scan/nsu089>

- Ritchie, J. B., Tovar, D. A., & Carlson, T. A. (2015). Emerging Object Representations in the Visual System Predict Reaction Times for Categorization. *PLoS Comput Biol*, *11*(6), e1004316. <https://doi.org/10.1371/journal.pcbi.1004316>
- Sha, L., Haxby, J. V., Abdi, H., Guntupalli, J. S., Oosterhof, N. N., Halchenko, Y. O., & Connolly, A. C. (2015). The Animacy Continuum in the Human Ventral Vision Pathway. *Journal of Cognitive Neuroscience*, *27*(4), 665–678. [https://doi.org/10.1162/jocn\\_a\\_00733](https://doi.org/10.1162/jocn_a_00733)
- Troiani, V., Stigliani, A., Smith, M. E., & Epstein, R. A. (2014). Multiple Object Properties Drive Scene-Selective Regions. *Cerebral Cortex*, *24*(4), 883–897. <https://doi.org/10.1093/cercor/bhs364>
- Vida, M. D., Nestor, A., Plaut, D. C., & Behrmann, M. (2017). Spatiotemporal dynamics of similarity-based neural representations of facial identity. *Proceedings of the National Academy of Sciences*, *114*(2), 388–393.
- Wardle, S. G., & Ritchie, J. B. (2014). Can Object Category-Selectivity in the Ventral Visual Pathway Be Explained by Sensitivity to Low-Level Image Properties? *Journal of Neuroscience*, *34*(45), 14817–14819. <https://doi.org/10.1523/JNEUROSCI.3566-14.2014>
- Warrington, E. K., & McCarthy, R. (1983). Category specific access dysphasia. *Brain*, *106*(4), 859–878.
- Warrington, E. K., & McCarthy, R. A. (1987). Categories of knowledge: Further fractionations and an attempted integration. *Brain*, *110*(5), 1273–1296.
- Warrington, E. K., & Shallice, T. (1984). Category specific semantic impairments. *Brain*, *107*(3), 829–853.



---

NEURAL CODING OF VISUAL OBJECTS:  
NEW INSIGHTS INTO CATEGORICAL  
REPRESENTATIONS

---

Erika W. Contini<sup>1,2</sup>, Erin Goddard<sup>3</sup>, Tijl Grootswagers<sup>4</sup>, Mark Williams<sup>1,2</sup>,  
& Thomas Carlson<sup>2,4</sup>

<sup>1</sup>Department of Cognitive Science, Macquarie University, Australia

<sup>2</sup>ARC Centre of Excellence in Cognition and its Disorders and Perception in Action Research Centre,  
Macquarie University, Australia

<sup>3</sup>McGill Vision Research, Department of Ophthalmology, McGill University, Montreal, Canada

<sup>4</sup>School of Psychology, The University of Sydney, Australia

**Published as: Contini, E. W., Goddard, E., Grootswagers, T., Williams, M., Carlson, T. (2020) A humanness dimension to visual object coding in the brain, *NeuroImage*, Vol. 221, 117139, <https://doi.org/10.1016/j.neuroimage.2020.117139>.**





### 3.1. Abstract

Neuroimaging studies into human visual object recognition have largely focused on a relatively small number of object categories (for example, faces, bodies, scenes, and vehicles). Recent studies have taken a broader focus, investigating a hypothesised animate-inanimate distinction in object representations. However, such studies have evaluated representations of objects using stimuli that are clearly identified as animate or inanimate, neglecting to include objects that may not fit neatly into this dichotomy. For example, objects such as robots and animal- or human-like toys do not easily fit into this rubric as they encompass concepts like agency and human-like experiences, which are known to influence human perception. As such, it is unknown how the animate-inanimate dichotomy accounts for neural representations of these ambiguous objects. Here we generated a novel stimulus set including objects from ambiguous categories, along with new exemplars from standard object categories. We used time-series decoding methods applied to MEG data from 24 human subjects (18 female) to compare contemporary object models, such as the animacy dichotomy, to new models of human agency/experience. We show that early neural responses were best accounted for by low-level, visual similarity models of the objects, while performance later in time ( $>155$  ms post-stimulus onset) was dominated by the agency/experience models. In particular, a graded model of human-similarity provided the best account for the representation of objects, demonstrating the brain's "human-centric" focus in the construction of higher order object representations, and highlighting the importance of using broad and diverse stimulus sets when investigating visual object processing.



### 3.2. Introduction

Object recognition is fundamental for human vision and is accomplished remarkably quickly by the brain (Thorpe, Fize, & Marlot, 1996). However, the neural processes underlying these perceptual abilities remain unclear. The ventral temporal cortex (VTC) is widely known for its involvement in visual object perception (Caramazza & Shelton, 1998; Haxby et al., 2001; Ishai, Ungerleider, Martin, Schouten, & Haxby, 1999; Mahon et al., 2007). Within VTC, objects belonging to the same semantic category evoke more similar neural responses, while objects from separate categories evoke more dissimilar responses. In particular, an animate-inanimate dichotomy has been proposed as an overarching categorisation principal in humans and primates (Kiani, Esteky, Mirpour, & Tanaka, 2007; Kriegeskorte, Mur, Ruff, et al., 2008; Pinsk et al., 2009), with neuroimaging studies revealing a medial to lateral organization of animate and inanimate objects respectively across this region (Chao, Haxby, & Martin, 1999; Kanwisher, McDermott, & Chun, 1997; Konkle & Caramazza, 2013; Mahon et al., 2007; Taylor & Downing, 2011). Regions within VTC also respond preferentially to images from particular subcategories, including faces, animals, bodies (Downing, 2006; Downing, Jiang, Shuman, & Kanwisher, 2001; Haxby et al., 1994; Puce, Allison, Asgari, Gore, & McCarthy, 1996; Sergent, Ohta, & Macdonald, 1992), tools (Chao et al., 1999; Chao & Martin, 2000) and places (Epstein, Harris, Stanley, & Kanwisher, 1999; Epstein & Kanwisher, 1998; Taylor & Downing, 2011). However, these localisation studies focus on coarse distinctions between a relatively small number of object categories; a mere sample of the wide array of objects that we see in everyday life.

A complementary alternative to localising object categories within specific brain regions or subregions, is to characterise the representational structure of distributed patterns of brain activity (Haxby et al., 2001; Ishai et al., 1999). Using multivariate pattern analysis (MVPA) methods (Grootswagers, Wardle, & Carlson, 2016; Haynes, 2015; Pereira, Mitchell, & Botvinick, 2009), researchers can evaluate systematic variations within complex neural data and test specific hypotheses about the neural representation of object categories

without limiting analyses to a specific region of interest (Kriegeskorte & Kievit, 2013; Kriegeskorte, Mur, & Bandettini, 2008). Objects categorised collectively as *animate* can also be reliably separated into individual species, such as human, monkey, bird, and fish (Connolly et al., 2012; Kiani et al., 2007; Sha et al., 2015). Furthermore, the relative similarity/dissimilarity of object representations suggests that objects may be represented along an animacy continuum, whereby animate subcategories are grouped according to how biologically similar the species are to humans, with those most dissimilar appearing closer to inanimate objects in the representational space (Connolly et al., 2012; Sha et al., 2015). The animacy continuum, however, does not provide a clear prediction for subcategory differentiation within the inanimate domain, nor for how the brain would represent objects that blur the animate-inanimate distinction (e.g., robots and human-/animal-like toys). It is also unclear whether a continuum centred around ‘animacy’ best captures the stimulus dimension along which neural responses vary. As an alternative, Sha et al. (2015) propose that the neural representation of objects may be better characterised as varying with the object’s ability to perform goal-directed actions. Notably, similar factors related to agency and certain human-related experiences are also known to influence human perception of these types of objects (Gobbini et al., 2011; Gray, Gray, & Wegner, 2007), which raises these factors as other candidate organisational principles for the brain’s neural representation of objects.

In the present study, we used MEG to characterise the organisation of neural object representations, and to explore their temporal dynamics. We studied the brain’s emerging representation of 120 object stimuli from both animate and inanimate object categories, and tested a wide range of models that might account for the brain’s representation of these objects using the representational similarity analysis (RSA) framework (Kriegeskorte, Mur, & Bandettini, 2008). We found behaviourally-generated models of agency and other human-related experiences can account for the neural representations of these objects, and show the

model which best accounted for our data was a “human-centric” model that describes objects in terms of their similarity to humans.

### **3.3. Method**

#### **3.3.1. Participants**

Twenty-four English-speaking volunteers (18 female) with an average age of 24.7 years (SD = 5.47; range = 18-37) were recruited from the Macquarie University community. Informed written consent was obtained prior to participation, and participants were paid \$20/h for their time. All participants self-reported normal or corrected-to-normal vision (wearing of contacts was allowed), were free of medical conditions, and were not currently taking any neuroactive medications. This study was approved by the Macquarie University Human Research Ethics Committee.

#### **3.3.2. Stimuli**

Stimuli consisted of 120 naturalistic images of objects (Figure 1), which were displayed on a uniform grey background. Twelve object categories were used in the study: six animate (humans, primates, domestic animals, birds, fish, invertebrates) and six inanimate (plants, robots, machines, tools, toys, other non-moving objects). In this stimulus set, animate is defined as living animals, in line with previous research (Caramazza & Shelton, 1998; Carlson, Tovar, Alink, & Kriegeskorte, 2013; Connolly et al., 2012; Gobbini et al., 2011; Kriegeskorte, Mur, Ruff, et al., 2008; Sha et al., 2015). Categories were selected to include ones similar to those used by Sha et al. (2015), with the addition of robots and toys to address the questions about agency and experience. We also included machines, which, like robots, had moving parts, but did not have the humanistic/animalistic/agentive properties. Stationary objects were also included, which neither moved nor had humanistic/animalistic/agentive properties.



Figure 1. Stimuli from each of the 12 object categories. Animate object categories are ordered vertically according to the biological classes animacy continuum (Sha et al., 2015). Brackets show two examples of different groupings of the stimuli: living vs. non-living and animate vs. inanimate.

### 3.3.3. MEG Experimental Procedure

For the experimental task, participants completed eight blocks of 398 trials (3184 trials in total). Within each block exemplars were presented for 100 ms, with a random inter-trial interval ranging between 750 and 1000 ms. The eight blocks were collected in a single session totalling approximately one hour of MEG recording time. Stimuli were presented in a predetermined pseudo-randomised order, such that for each trial, the preceding and

following images had an equal probability of being from any one of the 12 object categories. The ordering of the 8 blocks was pseudo-randomised across participants.

Across trials, object images were manipulated in two ways to reduce the effects of low-level stimulus properties on our data. Firstly, a left-right flipped version of each image was included in the stimulus set, resulting in a total of 240 stimuli from 120 object images. Secondly, during image presentation, stimuli appeared in one of four locations while participants maintained fixation on a central marker, thus varying retinal location of the stimulus images. The four locations were defined by a shift from central presentation towards each of the four corners of the screen, where each stimulus location overlapped the central fixation point (details in *Display Apparatus* below). Each stimulus was presented three times at each location. This resulted in a total of 2880 trials (240 stimuli x 4 locations x 3 repetitions = 2880 trials). The additional trials were not included in the analysis: these included the first and last trial of each block, as well as 288 repeat trials that were added for the attention task (see below).

#### **3.3.4. Attention Task**

During the experiment, participants completed a one-back attention task, where they were required to press a button whenever an object image was repeated consecutively. Participants received feedback about their accuracy on the task at the completion of each block. The mean accuracy across participants was 87.38% ( $SD = 7.28\%$ ), with an average reaction time of 535 ms ( $SD = 51$  ms). Due to a malfunction of the response button during the experiment, accuracy and reaction times were missing for one of our 24 participants, as well as for one out of the eight blocks for each of two further participants. These participants were still instructed to perform the task and were unaware that the button was not recording their responses.

#### **3.3.5. Display Apparatus**

Participants lay supine in the magnetically sealed recording room. Using an InFocus IN5108 projector situated outside the chamber, stimuli were projected onto a mirror, which

reflected the image onto the ceiling, located approximately 113 cm above the participant. The total screen area was 20x15 degrees of visual angle (DVA). Throughout the experiment the screen background was held at a mean grey, and subjects were instructed to fixate on a black central fixation point (diameter of 0.1 DVA) that was always present. All stimulus locations were within a 6.9 DVA square, centred on the fixation point. Each stimulus consisted of a 256x256 pixel image (containing the segmented colour object) that was drawn to a 4.9x4.9 DVA square. Stimuli were presented one at a time, in one of four locations aligned with the upper left, upper right, lower left, or lower right corner of the 6.9 DVA square. A central square of 150 pixels (2.9 DVA) was common to all four stimulus locations. All stimuli were drawn as full colour segmented objects against a mean grey background (as in Figure 1): the same mean grey as the screen outside the stimulus location. Upon stimulus presentation, a 50x50 pixel (1x1 DVA) white square simultaneously appeared in the bottom right corner of the projection, which was aligned with a photodetector attached to the mirror to accurately record the stimulus presentation time in the MEG recording. The experiment was run on a Dell PC desktop computer using MATLAB software (Natick, MA) and the Psychophysics Toolbox extensions (Brainard, 1997; Kleiner et al., 2007; Pelli, 1997).

### **3.3.6. MEG Data Acquisition and Analysis**

MEG data were recorded in the KIT-Macquarie Brain Research Laboratory using a 160-channel whole-head axial gradiometer (KIT, Kanazawa, Japan). Continuous data were acquired at a sampling rate of 1000 Hz, and were band-pass-filtered online from 0.03 to 200 Hz. MATLAB (2013b, Natick, MA) was used for all processing and statistical analyses of the data. Offline, we down-sampled the data to 200Hz and epoched each trial into an event with a time window from -100 ms to 600 ms relative to stimulus onset. To reduce the dimensionality of the data, we applied Principal Components Analysis to the epoched data from the 160 gradiometers, and retained the first  $n$  components that accounted for 99% of the variance. The number of components retained for each participant ranged from 14 to 72 (Mean = 34.21, SD = 18.90).



**3.3.6.1. Classification analysis.** For each participant, we used linear discriminant analysis (Duda, Hart, & Stork, 2001) to classify object/exemplar identity at the single trial level, training and testing classifiers on their ability to discriminate every possible exemplar pair of the 120 object images. We used cross-validated classification accuracy as a measure of how dissimilar the patterns of brain activity were for one exemplar compared to another (Nili et al., 2014). We did not attempt to model the effects of spatial position or left-right flip in our classification analysis, but instead used a single data label (the object identity) for data obtained from both the standard and left-right-flipped versions of the stimuli, as well as all four stimulus presentation locations. By including data from all variations of the stimuli, we sought to force the classifier to generalise beyond lower-level visual features, (such as the presence or absence of stimulation at a given location in the visual field), and instead use any neural correlate of object identity. These modifications to the stimulus presentation would have introduced extra noise into the signal across trials, so would tend to reduce classifier performance relative to unvarying stimuli, but they allowed us to better target higher-level object representations. For each time-point, we trained and tested a separate classifier to discriminate each pair of exemplar identities from the PCA components. We used a 10-fold cross-validation procedure, where the classifier was trained on data from 90% of the trials and then its accuracy was evaluated using its performance when classifying the remaining 10% of the data, so that the classifier was never tested on data that were included in the training set. This process was repeated 10 times, so that all trials were used as test data once each. D-prime ( $d'$ ) was used as the metric for classification accuracy.

**3.3.6.2. Representational Similarity Analysis (RSA).** Classifier accuracies ( $d'$ ) were averaged across exemplar pairs to obtain the mean classifier performance for each time point. Additionally, to capture the pattern of classifier performance across exemplar pairs

and compare this pattern with model predictions, we constructed a Representational Dissimilarity Matrix (RDM) for each time point. The RDM is a 120x120 matrix, symmetric along the negative diagonal, where each cell is the classification accuracy ( $d'$ ) for that pair of exemplars.

For each time point we compared each participant's observed RDM with model RDMs, where each model RDM was a 120x120 matrix derived from theory, computational modelling, or behavioural data (as described in detail below). This analysis, known as 'Representational Similarity Analysis' (RSA) (Kriegeskorte, Mur, & Bandettini, 2008) tests the relationship between models of interest and the group data, measuring how well the model RDMs account for the observed pattern of results. At each time point we used Kendall's tau-a to compute the rank order correlation between each candidate model and the data, and used these correlation values to compare alternative models in their ability to account for the neural data. Figure 2 shows the model RDMs.

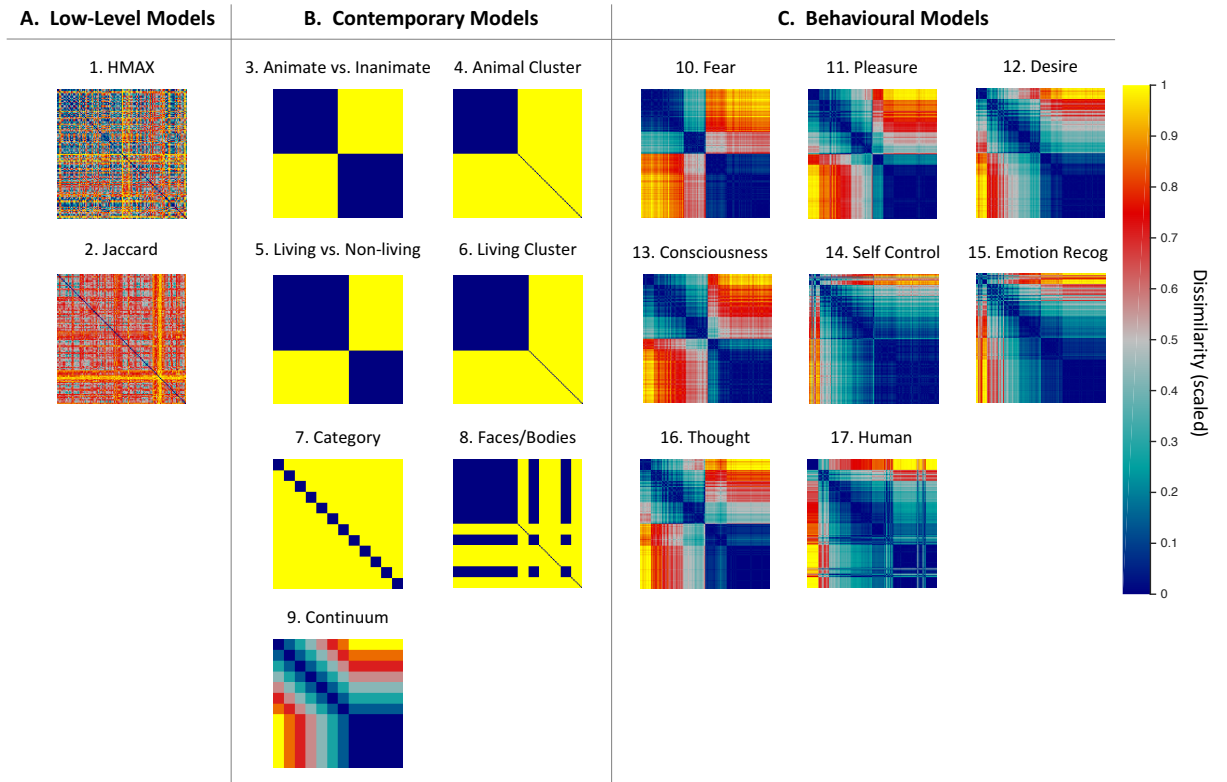


Figure 2. RDMs used for model testing. Models are grouped according to whether they are low-level models (A), contemporary models (B), or behavioural models (C). See text for descriptions of each model. Model axes refer to all 120 image exemplars (grouped by category in the same order as in the RDM based on average decoding of the MEG data in Figure 5). Colour bar indicates predicted degree of dissimilarity between exemplar pairs.

**3.3.6.2.1. Low-level models.** The HMAX and Jaccard silhouette models were included to test for the effects of low-level stimulus properties on the similarity/dissimilarity of neural responses, as measured using classifier performance.

*HMAX (1):* Computational model of low-level visual processes. We applied the HMAX model (Riesenhuber & Poggio, 1999; Serre et al., 2007) as detailed in Serre et al., (2007) to simulate the responses of low-level visual areas. HMAX was applied to images at only a single image location and based on the standard orientation of each stimulus (i.e. not left-right flipped). The responses of the final HMAX layer (C2) for every stimulus were

vectorized. We then generated the model RDM by taking the Euclidean distance between the vectorized model responses for each pair of stimuli.

*Jaccard (silhouette model; 2)*: An abstract shape model that measures the shape of each object in terms of the pixels that the image occupies (Jaccard, 1901). We generated the model RDM by comparing the overlapping silhouette regions of two images at a time and obtaining a measure of the difference. This model was generated based on the standard orientation of each stimulus (i.e., not flipped), independent of location.

**3.3.6.2.2. Contemporary models of object representations.** The ‘contemporary’ models (Figure 2, models 3-9) were created based on organisational structures proposed in previous studies, with the term ‘contemporary’ used to highlight that these reflect the current understanding of object category structure. Descriptions of each model are provided below.

*Dichotomy models (3 and 5)*: The animate vs. inanimate dichotomy model (Caramazza & Shelton, 1998; Carlson et al., 2013; Cichy, Pantazis, & Oliva, 2014; Kriegeskorte, Mur, Ruff, et al., 2008), is a multi-category model that grouped all animate and inanimate objects separately (implying that objects within these groupings were more similar to each other, and more dissimilar to objects in the other grouping). Similarly, the living vs. non-living dichotomy model (Gainotti, 2000; Huth, Nishimoto, Vu, & Gallant, 2012; Warrington & Shallice, 1984) grouped all living and non-living objects separately. The living category included the same items as the animate category but with the addition of plants.

*Cluster models (4 and 6)*: The animal cluster model (4) is a single-category model that only grouped all animate objects together, suggesting that animate objects will be more similar to each other, and more dissimilar to all other objects, but that inanimate objects will not cluster. The living cluster model (6) follows the same principles, but grouping all living objects together. The cluster models were created to determine whether the effect of the dichotomy models was driven by cohesion within the in-group alone (i.e., animate, living), with more disparate object representations in the out-group category (i.e., inanimate, non-living) (Clarke & Tyler, 2014).

*Faces/bodies model (7):* Faces and bodies stand out as special categories for object recognition (Barragan-Jason, Cauchoix, & Barbeau, 2015; Cauchoix, Barragan-Jason, Serre, & Barbeau, 2014; Gobbini et al., 2011; Haxby et al., 2001; Van de Nieuwenhuijzen et al., 2013), and so were of interest given the inclusion of toys and robots in our stimulus set. As such, the faces/bodies model is single-category model, grouping together all object categories that had faces or bodies, including all animate objects, as well as robots and toys.

*Category model (8):* The category model was included as a measure of category individuation, as it proposes that items within individual categories have distinctly related patterns due to common visual and semantic properties, and these patterns are more different to those of objects from other categories (Clarke & Tyler, 2014). This model grouped each individual category as being more similar to within-category items and more dissimilar to other categories.

*Continuum model (9):* The continuum model is a graded model based on the animacy continuum proposed by Sha et al. (2015; see also Connolly et al., 2012). The continuum included a gradient of similarity between object categories that varied along a dimension related to biological classes, such that categories more similar to humans (biologically), would have more similar activity patterns, and those more dissimilar to humans would have activity patterns more similar to inanimate objects. For this model, plants were included on the continuum as they are a biological category, and were represented on the continuum between invertebrates and inanimate objects. All non-living inanimate objects were treated as a single category, most dissimilar to the human category.

**3.3.6.2.3. Behavioural models.** The behavioural models include the agency/experience models (Figure 2, models 10 – 16) and the human model (17). These models were created by obtaining behavioural ratings of the stimuli according to a specific question (detailed below). A total of 325 Amazon’s Mechanical Turk workers residing in either the United States of America or Canada, completed one of the eight surveys online (number of participants per survey ranged from 40 – 43). Participants included 146 females

(1 other, 1 no response), and had an average age of 35.27 years ( $SD = 10.26$ , range = 18.9 – 70.8; one age value missing). In each survey we asked workers to answer a single question for each of the stimuli:

10. *Fear* – How much is it capable of feeling afraid or fearful?

11. *Pleasure* – How much is it capable of experiencing physical or emotional pleasure?

12. *Desire* - How much is it capable of longing or hoping for things?

13. *Consciousness* - How much is it capable of having experiences and being aware of things?

14. *Thought* - How much is it capable of thinking?

15. *Emotion-recognition* - How much is it capable of understanding how others are feeling?

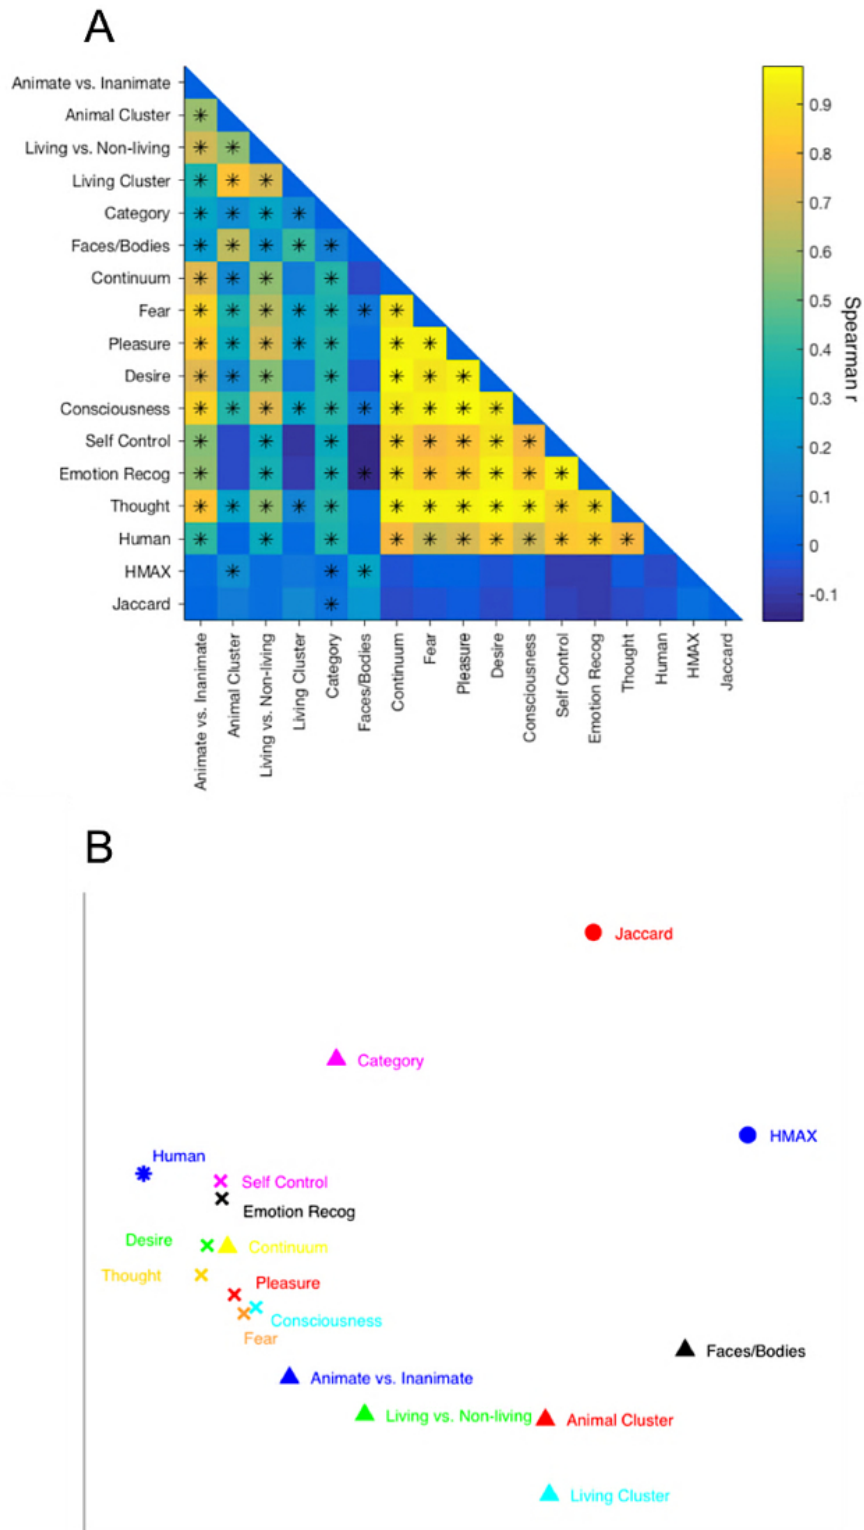
16. *Self-Control* - How much is it capable of exercising self-restraint over desires, emotions or impulses?

17. *Human* – How similar is this to a human?

Surveys 10-16 were based on a subset of the mental capacity surveys used in Gray et al. (2007), which vary as to how much they loaded onto the author's 'Experience' and 'Agency' factors that were established in their study. The seven agency/experience models were based on the results of these surveys. The 'Human' survey (17) was added to address a meta-representational idea of categorization, that of "human-ness": a complex factor which may encompass biology, agency, and visual similarity. Each survey required participants to rate all 120 images on a 7-point scale from 'Not at all' to 'Very much so' in response to the specific question. Each survey took approximately 10 minutes to complete and participants were paid \$1.50 for their time. The surveys were created and administered using the Qualtrics online survey platform. For each survey, participants provided voluntary consent and basic demographic information before completing the survey. Participants were only allowed to complete one of the eight surveys available, resulting in unique individuals for each survey. Stimulus order was randomised separately for each participant.

To construct the models based on agency and experience (shown in Figure 2), an RDM was created for each set of survey responses by obtaining the absolute difference between image ratings for each pairwise comparison of the 120 images, using the mean ratings of each image. These RDMs, based on the survey ratings, provide hypothetical models of the degree of dissimilarity between the neural responses associated with each image. For graphical purposes, we scaled these difference values between 0 and 1 for each model, such that warmer colours indicate greater dissimilarity, while cooler colours depict greater similarity between the neural representations in the pair-wise comparison.

**3.3.6.3. Model intercorrelations.** As the models we used in this study were not orthogonal, we measured the degree of overlap by performing correlations (Spearman) between each of the models (see Figure 3A). By evaluating the strength of these correlations, we obtained an estimate of how much the models overlap in terms of the hypotheses being tested. Of particular note, the mTurk behavioural models based on the agency and experience factors from Gray, et al. (2007) and the human model we created are all highly correlated (see clustering in Figure 3B MDS plot of the representational geometry): this was expected, since each of these models capture slightly different aspects of similarity to humans.



*Figure 3.* Model intercorrelations. A) Model correlation matrix. Cell colour indicates correlation strength ( $*p < .05$ , adjusted for multiple comparisons across time points using a FDR of  $q < .01$ ), with yellow cells indicating a stronger correlation between models, while blue indicates a weak/no correlation. B) MDS plot showing the representational geometry of model similarity.



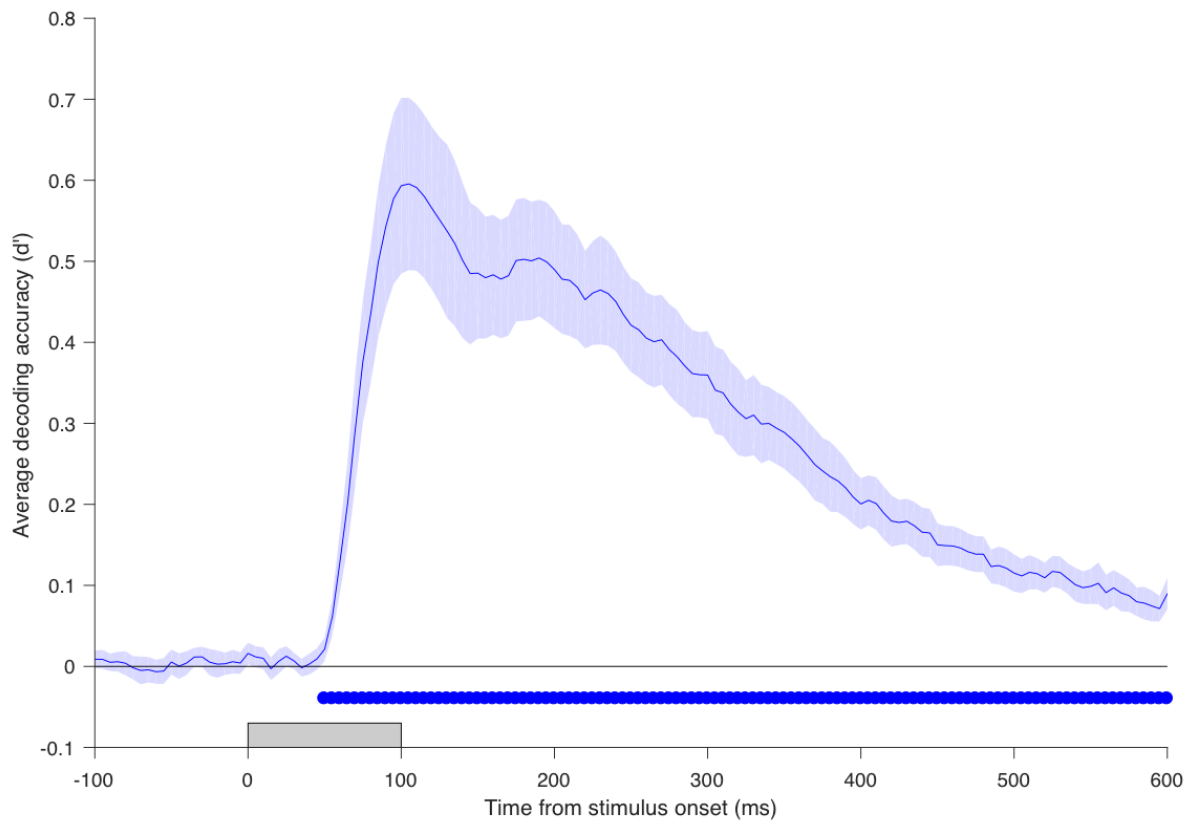
We deliberately chose stimuli that were visually diverse within each subcategory, to minimise the extent to which visual similarity would produce seemingly ‘categorical’ patterns of results. The model correlation data suggests that our stimulus set provided good separation of visual similarity and object category, since few models correlated with the visual feature confound models. Importantly, this should minimise the confounding contribution of low-level visual similarity when we evaluate our hypothesis driven models. Exceptions to this included the animal cluster, category, and faces/bodies models, which each showed a significant correlation with one, or both of the HMAX and Jaccard models. This suggests that despite our stimulus diversity, there was still greater visual homogeneity of exemplars within the category groupings in these models than between category groupings. This means that, particularly for the animal cluster, category, and faces/bodies models, any correlation between these models and the observed pattern of classifier performance may have been driven by low-level visual similarity rather than by the higher-level category structure represented by these models.

**3.3.6.4. Noise ceiling.** We computed a noise ceiling for our data following the methods of Nili et al, (2014). The noise ceiling provides an estimate of the theoretical maximum model correlation that is possible given the noise in the data. To compute the noise ceiling we generated a theoretically ‘ideal’ model by taking the between-subjects average RDM, then correlating each subject’s data with this ideal model using Kendell’s tau-a. The upper bound of the noise ceiling is the average of these correlation values. To calculate the lower bound of the noise ceiling we repeated a similar process, except that when we correlated each subject’s data with the ‘ideal’ model, we excluded that subject’s own data from the average. Model correlations that approach the noise ceiling indicate that the model provides a near-complete account of the explainable variance in the data, whereas greater

distance from the noise ceiling indicates that a better model could account for further variance in the data.

### **3.4. Results**

We scanned participants using MEG while they viewed briefly presented object stimuli. We trained and tested classifiers on their ability to discriminate each pair of the 120 objects in our stimulus set, using the MEG sensor recordings. Evaluation of the overall time-varying object decoding performance revealed significant sustained decoding of individual exemplars commencing from 50 ms post stimulus-onset (Figure 4), with peak decoding performance at approximately 105 ms post stimulus onset. Across individuals, there was a tendency for classifier performance to show a second smaller peak at a later time, which is reflected in the smaller peak in the averaged data, at around 180-195 ms post stimulus-onset. These results are consistent with results from previous MEG decoding studies (Carlson et al., 2013; Cichy et al., 2014; Goddard, Carlson, Dermody, & Woolgar, 2016) and align with what we know about the time taken for visual information to reach cortical areas (Thorpe et al., 1996).



*Figure 4.* Average decoding performance (measured in d-prime) over time for all exemplar pairs. Grey bar indicates the period the image was on the screen. Error bars indicate 95% between-subject confidence intervals. Blue dots along the x-axis indicate time points at which decoding performance was significantly above chance (one-tailed t-test, adjusted for multiple comparisons across time points using a FDR of  $q < .01$ ).

We next examined the pattern of classifier performance across different stimulus object pairs. Figure 5 shows an RDM of decoding performance averaged across participants and across all time-points with above chance classifier decoding. Upon inspection, decoding performance varies widely depending on the exemplar comparison, however the human stimuli tend to be more dissimilar to other objects from categories that are further down the “continuum” (seen in the frequency of yellow cells). An exception to this is seen with the robot and toy stimuli, which appear more similar to the human and animate objects than

other inanimate objects (seen in the bluish horizontal bands on the left side of the RDM at the level of ‘robots’ and ‘toys’).

Having established that we could decode the individual object images from the MEG data, we then tested a range of hypotheses about category representations by comparing the observed RDM with a series of model RDMs (Figure 6). We first considered the maximum correlation with the data that would be possible for any model to achieve, by calculating the noise ceiling. The noise ceiling (Figure 6-1) increases steadily from around 55 ms post stimulus onset, peaking at 100 ms (lower-bound tau-a correlation of 0.26), before gradually returning to baseline noise levels over the subsequent ~500ms. As expected, the noise ceiling varies with average decoding accuracy (see Figure 4): where decoding accuracy is higher, the noise ceiling is higher and there is also less disparity between the lower and upper limits of the noise ceiling, indicating that there is more variance in the RDM that could potentially be captured by one or more models. However, as the peak correlation for the noise ceiling is still relatively low, this suggests a high degree of between-subject noise, or ‘unexplainable variance’ in the data.


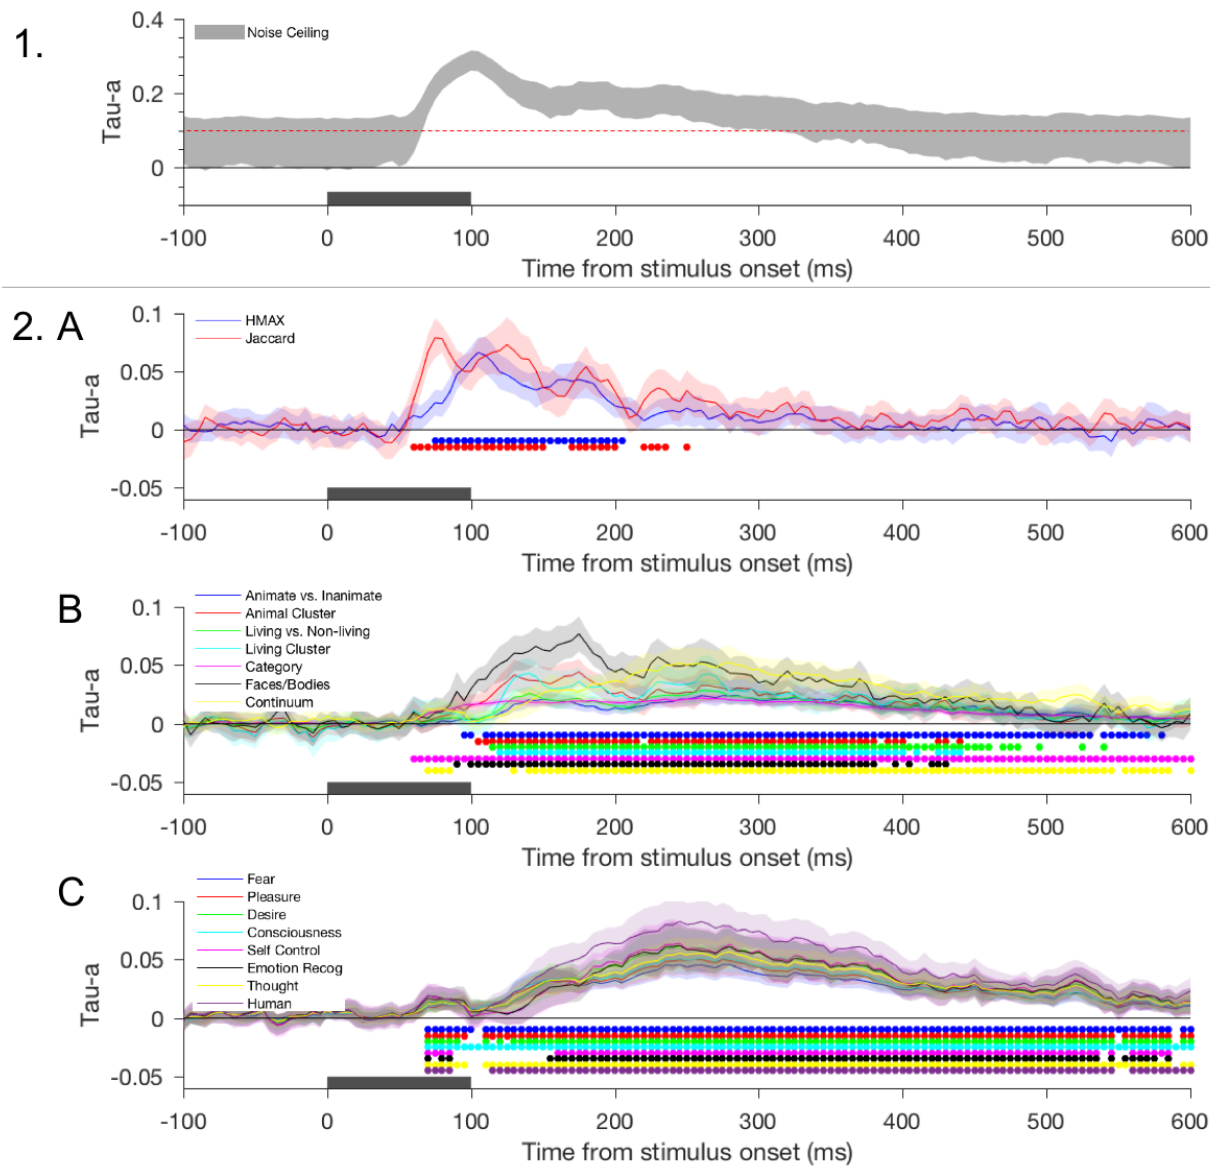


Figure 5. RDM showing average decoding performance over time (50-600 ms post stimulus onset) using scaled  $d'$  (proportion of maximum  $d'$ ). The colour of each cell indicates the degree of similarity between the neural responses to that pair of objects (as measured in classification performance). Warm colours (red to yellow) indicate greater dissimilarity (more distinct neural representations), while cool colours (green to blue) indicate greater similarity (overlapping neural representations).



*Figure 6.* Noise ceiling and RSA model testing correlations over time. 1) Noise ceiling time-course. The boundaries of the grey shaded area indicate the upper and lower limits of the noise ceiling (the maximum expected correlation of any model). 2) Correlation between the classifier data and (A) the low-level visual feature models, (B) contemporary models, and (C) behavioural-rating models. Red dotted line in (1) indicates the maximum value of the y-axis (Kendall's tau-a = 0.1) to aid comparison with the plots in (2). Shaded error bars in A, B and C indicate the 95% confidence interval of the between-subject means. Plotting conventions as in Figure 4.

### 3.4.1. Model Testing Across Time

At every time point, we rank-correlated each subject's RDM with each model RDM using Kendall's Tau-a (Figure 6-2). We then tested whether the group subjects' mean correlation was significantly above zero at each time point (one-tailed t-test  $p < .05$ , adjusted for multiple comparisons across time points using a false discovery rate (FDR) of  $q < .01$ ).

**3.4.1.1. Low-level visual feature models.** Both the Jaccard and HMAX models (Figure 6-2A) were significantly above zero only in the early stages of the time course, peaking at 75 and 105 ms respectively, with neither model reaching significance after 250 ms post stimulus onset. This is in line with the established literature about the time-course of visual object recognition, with responses related to lower-level visual stimulus properties occurring earlier on, and more complex semantic (and thus category-relevant) responses occurring later in the time-course (Carlson, Simmons, Kriegeskorte, & Slevc, 2014; Carlson et al., 2013; Cichy et al., 2014; Clarke, Devereux, Randall, & Tyler, 2015). However, it should be noted that both models show high correlations despite the fact that our design incorporated left-right flips of the images and spatial displacement of the images to control for some low-level stimulus properties; changes which would have reduced the performance of these models. The low-level models were generated using only the standard orientation of each stimulus, and yet they performed well at predicting the pattern of classifier performance when the classifier was required to generalise across left-right flips and the four spatial locations. That these models were still significantly correlated with the data affirms the importance of low-level visual similarity in neural representations, which may have contributed to categorical effects, both historically and in the present study.

**3.4.1.2. Contemporary models.** For the contemporary models (Figure 6-2B), there was variability in the onset of significant correlation with the MEG data, with onsets ranging from 60 ms (category model) to 120 ms (living-cluster model). The peak correlation times varied amongst these models, and did not follow the same order as the onsets: the living-cluster model had the earliest peak correlation at 140 ms (tau-a = .043), followed by animal-

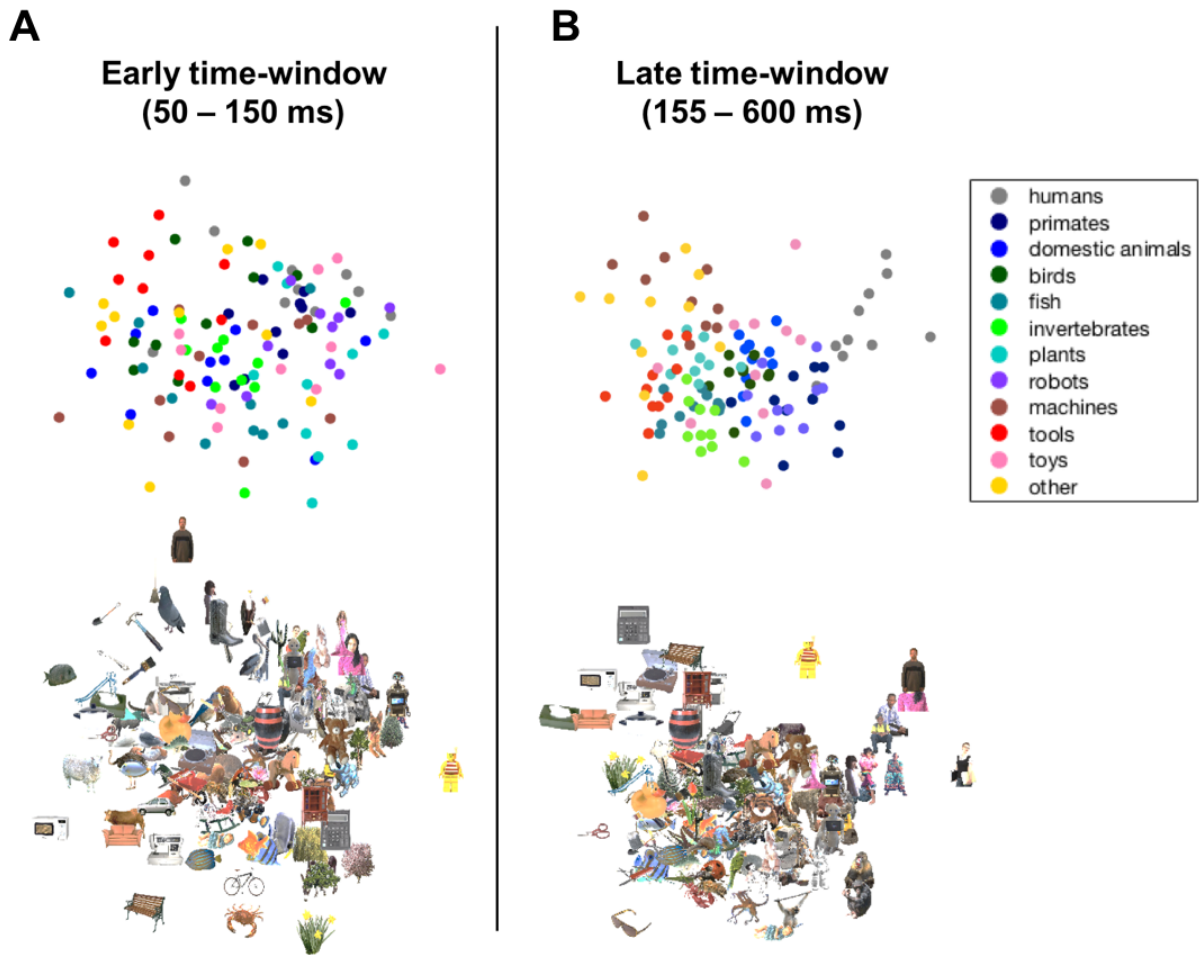
cluster and faces/bodies models at 175 ms ( $\tau_a = .044$  and  $.077$  respectively), living/non-living and category models at 260 ms ( $\tau_a = .029$  and  $.022$  respectively), continuum model at 280 ms ( $\tau_a = .052$ ), and animate/inanimate model at 285 ms ( $\tau_a = .024$ ). The models with the earliest peaks tended to have higher peak correlations, but notably two of these models (animal-cluster and faces/bodies) were among those showing significant overlap with one or both of the low-level feature models (see Figure 3), suggesting that low-level visual similarity may have contributed to the earlier onset and high peaks for these models. Interestingly, the cluster versions of the animal and living models had higher and earlier peak correlations than their non-cluster counterparts. While, in both instances, objects in these cluster categories are more likely to be visually similar, they also represent specific conceptually inclusive criteria, with the inanimate and non-living categories likely thought of as not meeting the criterion of ‘animate’ or ‘living’. Previous studies have found similar effects of an animate advantage using reaction times (RTs) (Carlson, Ritchie, Kriegeskorte, Durvasula, & Ma, 2014; Ritchie, Tovar, & Carlson, 2015), showing a significant relationship between RTs and representational distance from a categorisation boundary for objects within the animate category, which did not hold for inanimate objects. These results suggest that it is less useful to think of inanimate objects as a grouped category, and emphasise the need to re-evaluate how we conceptualise and model neural object categorisation.

**3.4.1.3. Behavioural rating models.** Figure 6-2C shows the correlations over time for the behavioural rating models. Due to the generally high level of overlap between the internal structure of these models (see Figure 3), it is not surprising that as a group they have similar correlation time courses. In each case, the onset of above chance correlation with the model was approximately 70 ms, and they tended to show a peak correlation at later times than the visual similarity and contemporary models, with peaks ranging from 245 to 280 ms. Despite their similarity, the human-factor model appears to have a slightly higher peak correlation ( $\tau_a = .083$ ) than the other behavioural models, which have peaks from  $.046$  to  $.064$ .



### 3.4.2. Model Testing: Early Versus Late Time Windows


Across all models tested, there was a broad early versus late division between the low-level stimulus property models, and the more category-related models respectively. As such, we elected to perform statistical comparisons across average model performance in separate early and late time-windows. As overall decoding was significant from 50 ms, we excluded time points before this, as this period would largely represent noise. The remaining data were separated into early and late time windows based on the two peaks evident in the overall decoding performance (Figure 4), such that the early time-window (50 – 150 ms) captures the first peak, while the late time-window (155 – 600 ms) captures the second peak. In Figure 7 we illustrate the representational geometry of the data using multidimensional scaling (MDS) for the early and late time-windows separately. For each RDM, we used MDS to provide a visualisation of these dissimilarities in a 2-dimensional space, using the (default) *Kruskal's normalized stress1* criterion. When visually comparing the MDS plots, objects in the late time-window tend to cluster into their individual object categories, whereas objects in the early time-window appear more disparate. Of particular note, in the late time-window, the animate/inanimate distinction is somewhat apparent (right/left halves of the MDS plots respectively), however, the clusters for toys and robots seem to fall closer to the animate objects, located on the right/animate side of the MDS plot.



*Figure 7.* MDS plots showing the representational geometry of the 120 object images for A) early (50 – 150 ms), and B) late (155 – 600 ms) time-windows. In each case the data are plotted in a space defined by the first two dimensions extracted using MDS (where each dimension has arbitrary units). Upper MDS dot plots show category clustering, where colours indicate the different object categories (see key for category labels) and each spot represents a different object. Lower MDS image plots show the same data, replacing the dots with the corresponding object image.

We computed the correlation between each of our models and the average classifier performance across the early and late time-windows (Figure 8). To facilitate comparison between the early and late time windows (with different noise ceilings) we plotted model performance as a proportion of the average noise ceiling (lower bound) for each time window. Due to this, we shortened the late-time window to only cover the period of 155 – 350 ms, so as not to artificially inflate the results in this time window where the noise-ceiling gradually returns to baseline levels. We also conducted a series of repeated measures t-tests to assess between-model performance in each of these time windows separately (adjusted for multiple comparisons across time points using a FDR of  $q < .01$ , Figure 8).

In the early time window, the low-level feature models and faces/bodies had the highest average correlations, with the Jaccard model achieving the highest average correlation, followed by HMAX and then faces/bodies. Each of these models performed significantly better than all lower-correlated models (with the exception of HMAX, which did not significantly differ from the face/bodies model). Few of the lower-correlated models differed significantly, though of particular note, the animal and living cluster models were significantly better than their counterparts (the animate/inanimate and living/non-living dichotomous models) suggesting that the correlation of the animate/inanimate and living/non-living dichotomous models is primarily driven by the similarity between objects within the animate and living categories rather than by homogeneity within the inanimate and non-living categories.



*Figure 8.* Average model correlations (tau-a) as a proportion of the noise ceiling across early (A) and late (B) time windows. Models are arranged in order of highest average rank correlation within each time window (highest on the left), and error bars indicate 95% confidence intervals of the between-subjects average correlation. Paired t-tests determined significant differences between model correlations (lines indicate significance at  $p < .05$ , adjusted for multiple comparisons across time points using a FDR of  $q < .01$ ). Note, the y-axis range for the late time-window is double that of the early time-window.

Conversely, the later time window is dominated by the human model, which performs significantly better than all other models in this time window. This model is also the only model that significantly outperforms the faces/bodies model, suggesting that the organisation of neural representations according to this human-centric model is not merely driven by the presence of faces/bodies, which are well-established as significant factors in object processing (Farah, 1996; Kanwisher et al., 1997; Thorpe et al., 1996). After the human model, the agency/experience models tend to have the next best fit with the data overall, along with the faces/bodies and continuum models. The faces/bodies and continuum models were the only contemporary models to outperform both low-level feature models in the late time-window. As such, the organisation of representations depicted by these models appears more relevant to the higher level conceptual organisation of object representations than the other contemporary models we evaluated. Overall, the contemporary and low-level feature models tend to have the lowest average correlations for the later time window, with the category, animate vs. inanimate, and living vs. non-living models not significantly different from either Jaccard or HMAX. As in the early time window, the animal and living cluster models have again performed significantly better than their dichotomous counterparts.

### **3.5. Discussion**

Many candidate categories, such as animate-inanimate and living/non-living, have been proposed as organisational principles in the brain's representation of visual objects. However, these models vary in the extent to which their proposed categories overlap with low-level stimulus feature differences, and the contribution of these factors is not always tested. Additionally, these models have typically not been tested with stimuli that are ambiguous in their category. Here we sought to provide an in-depth evaluation of contemporary models of visual object categorisation by directly comparing their ability to account for neural responses to a diverse range of object stimuli, as measured using MEG. In addition to contemporary models, we created new theoretical and behaviour-based models, and models based on low-level visual similarity, which we compared with

contemporary models. To test the predictive power of contemporary models for non-standard stimuli, we included novel stimuli that did not conform to the typical categories of contemporary models such as the frequently evaluated concept of animacy. Overall, our results showed that the best performing model was one based on a broad concept of human-similarity.

The evolution of the classifier performance and model correlations over time reported here are consistent with the established temporal flow of information associated with object recognition and categorisation processes (Contini, Wardle, & Carlson, 2017). These complex and multi-stage processes begin with initial cortical responses that primarily vary with the low-level visual properties of the stimulus (DiCarlo, Zoccolan, & Rust, 2012); information about the stimulus accumulates over time (from 80-800ms), with the influence of semantics (Clarke & Tyler, 2014) and category structure at intermediate (e.g. face, body), and subsequently higher (e.g. animate/inanimate) levels emerging at later times (Carlson et al., 2013; Cichy, Khosla, Pantazis, Torralba, & Oliva, 2016; Cichy et al., 2014; Contini et al., 2017). Our results are consistent with this evolution: at early times the best performing models were the low-level feature models (Jaccard and HMAX), while the more complex categorisation models (contemporary and behavioural models) had stronger correlations with the data at later times. While low-level stimulus properties contribute to the rich representation of object categories, particularly in the early stages of the object recognition time-course, our late time-window results and those from a growing number of studies have shown that low-level features are unable to wholly account for the complexities of object categorisation (Bracci & Op de Beeck, 2016; Carlson et al., 2013; Cichy et al., 2014; Coggan, Baker, & Andrews, 2016; Kaiser, Azzalini, & Peelen, 2016; Kriegeskorte, Mur, Ruff, et al., 2008; Proklova, Kaiser, & Peelen, 2016; Van den Hurk, Van Baelen, & Op de Beeck, 2017). Since the later time epoch is of greater relevance to the question of higher-level object representation, in the remainder of our discussion we consider how the model

correlations from this later time window inform our understanding of the neural representation of objects.

One of the most surprising results in the later time window is that the animate vs. inanimate model was among the lowest performing models, even though the animate/inanimate dichotomy is well-established in the visual object recognition literature (Kiani et al., 2007; Kriegeskorte, Mur, Ruff, et al., 2008; Carlson et al., 2013; Cichy et al., 2014; Proklova et al., 2016; Caramazza & Shelton, 1998). Our study, however, specifically included stimuli from subcategories that do not fall at polar ends of the animate/inanimate spectrum, while previous studies tend to use objects that have clear membership to animate or inanimate categories (Bracci & Op de Beeck, 2016; Konkle & Caramazza, 2013; Proklova et al., 2016) with some studies, in addition, utilising stimuli from the same stimulus set (Carlson et al., 2013; Cichy et al., 2014; Kiani et al., 2007; Kriegeskorte, Mur, Ruff, et al., 2008). The performance of the animate vs. inanimate model (and similarly with the living vs. non-living model) in our study was likely reduced by the inclusion of robots and toys as we grouped these ambiguous categories with the inanimate objects, reasoning that they are man-made and not animals/living. When visually inspecting the MDS plots showing the approximate representational geometry in the late time-window (Figure 7), these objects are represented closer to humans and animate objects than to the inanimate objects, suggesting that an animate/inanimate distinction may not be the best way to classify this stimulus set, and highlighting the impact of stimulus selection on defining the organisation of object categories. Indeed, a recent fMRI study by Bracci, Kalfas, & Op de Beeck (2017) similarly showed that visually confusing objects (e.g., a mug in the shape of a cow) exhibited neural activity patterns that were more similar to animate objects (i.e., an actual cow) than inanimates. Furthermore, as exemplar typicality affects the distinctiveness of category representations (Iordan, Greene, Beck, & Fei-Fei, 2016), the inclusion of these ambiguous object categories may have disproportionately affected a strict dichotomous categorisation model, like that of animate vs. inanimate. Further testing of any

category boundary effects (including the animate/inanimate boundary) should attempt to replicate the effect with new stimulus sets which include categories and stimuli that are visually diverse.

The superior performance of the human model is an interesting finding which builds on our existing understanding of the representational geometry of object categories. The model builds on the continuum idea of Connolly et al. (2012) and Sha et al. (2015) as it represents a kind of human-similarity continuum. However, unlike the animacy continuum which is based on biological classes, the human model was not limited by biology, and could incorporate factors such as faces/bodies, which are important to the system irrespective of biology (i.e., as seen with the traditionally inanimate, and non-biological categories of robots and toys) (Gobbini et al., 2011; Tong, Nakayama, Moscovitch, Weinrib, & Kanwisher, 2000). Results from an fMRI study by Gobbini et al. (2011) are consistent with a level of cross-over between animate/inanimate object categories that does not fit into a dichotomy, nor a continuum based on biological classes. The authors compared human observers' perception of human faces and robots, and found that robots evoked activation in areas associated with faces (though to a lesser extent than humans), while also activating object areas and areas associated with mechanical movements. This supports the idea of more a complex model of object categorisation that incorporates factors such as agency and human-related experiences; factors which are not given great weight in contemporary models of neural object representations. Given the relative strength of our human-centric model in accounting for the data, the idea of "humanness" as an important dimension in the neural representation of objects warrants further exploration.

Our best performing human-centric model likely encompasses a complex set of features, including both visual and conceptual factors associated with being human. In our study, we did not impose a definition or any criteria against which people should rate the objects when asked 'How similar is it to a human?' (with responses from this survey used to generate the human model). Accordingly, we do not know which features people were using



to rate object ‘humanness’, raising an interesting area for further investigation. The brain likely makes use of both visual and semantic information for representing objects (Carlson, Simmons, et al., 2014; Clarke & Tyler, 2014; Coggan, Baker, et al., 2016; Coggan, Liu, Baker, & Andrews, 2016), with our data suggesting that the semantic component encompasses information about concepts such as function, agency, and human experience. Indeed, a recent study by Connolly et al. (2016) showed an overlap between regions sensitive to the perceived threat of animals and those associated with social cognition in humans, highlighting the importance of agent-related dimensions to object processing.

Presently, we still do not have a clear understanding of how different semantic concepts relate to object representations and category structure. A recent paper by Martin (2016) puts forward a model of neural object representations that attempts to explain how concepts such as these are encompassed in multifaceted representations of object categories. Martin (2016) suggests that neural patterns are formed from complex interactive circuits based on a range of systems throughout the brain, including those associated with action, perception and emotion. This idea shifts the focus away from models based on broad categories, with a view to a more holistic approach to object representations that considers interactions between various circuits throughout the brain. The success of the humanness model, which includes complex constructs such as function, agency, and human experience, lends support to the interactive and multifaceted proposal by Martin (2016). Finally, it is noteworthy that while the human centric model was the best performing model in the present study, the model was well below the theoretical maximum correlation (i.e. the noise ceiling), which highlights that there remains considerable work to be done in understanding how objects and object categories are represented in the brain.

Author contributions:

E.C, E.G, M.W, T.C designed the study

E.C, T.G collected the data

E.C, E.G, T.G conducted the analysis

E.C, E.G, T.C interpreted the results

E.C wrote the manuscript

All authors reviewed the manuscript

### 3.6. References

- Barragan-Jason, G., Cauchoix, M., & Barbeau, E. J. (2015). The neural speed of familiar face recognition. *Neuropsychologia*, 75, 390–401.  
<https://doi.org/10.1016/j.neuropsychologia.2015.06.017>
- Bracci, S., Kalfas, I., & Op de Beeck, H. (2017). The ventral visual pathway represents animal appearance over animacy, unlike human behavior and deep neural networks. *BioRxiv*, 228932.
- Bracci, S., & Op de Beeck, H. P. (2016). Dissociations and Associations between Shape and Category Representations in the Two Visual Pathways. *The Journal of Neuroscience*, 36(2), 432–444. <https://doi.org/10.1523/JNEUROSCI.2314-15.2016>
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 433–436.
- Caramazza, A., & Shelton, J. R. (1998). Domain-Specific Knowledge Systems in the Brain: The Animate-Inanimate Distinction. *Journal of Cognitive Neuroscience*, 10(1), 1–34. <https://doi.org/10.1162/089892998563752>
- Carlson, T. A., Ritchie, J. B., Kriegeskorte, N., Durvasula, S., & Ma, J. (2014). Reaction Time for Object Categorization Is Predicted by Representational Distance. *Journal of Cognitive Neuroscience*, 26(1), 132–142. [https://doi.org/10.1162/jocn\\_a\\_00476](https://doi.org/10.1162/jocn_a_00476)
- Carlson, T. A., Simmons, R. A., Kriegeskorte, N., & Slevc, L. R. (2014). The Emergence of Semantic Meaning in the Ventral Temporal Pathway. *Journal of Cognitive Neuroscience*, 26(1), 120–131. [https://doi.org/10.1162/jocn\\_a\\_00458](https://doi.org/10.1162/jocn_a_00458)
- Carlson, Tovar, D. A., Alink, A., & Kriegeskorte, N. (2013). Representational dynamics of object vision: The first 1000 ms. *Journal of Vision*, 13(10), 1–1.  
<https://doi.org/10.1167/13.10.1>
- Cauchoix, M., Barragan-Jason, G., Serre, T., & Barbeau, E. J. (2014). The Neural Dynamics of Face Detection in the Wild Revealed by MVPA. *The Journal of Neuroscience*, 34(3), 846–854. <https://doi.org/10.1523/JNEUROSCI.3030-13.2014>

- Chao, L. L., Haxby, J. V., & Martin, A. (1999). Attribute-based neural substrates in temporal cortex for perceiving and knowing about objects. *Nature Neuroscience*, 2(10), 913–919.
- Chao, L. L., & Martin, A. (2000). Representation of Manipulable Man-Made Objects in the Dorsal Stream. *NeuroImage*, 12(4), 478–484.  
<https://doi.org/10.1006/nimg.2000.0635>
- Cichy RM, Khosla A, Pantazis D, Torralba A, Oliva A (2016) Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence. *Sci Rep* 6
- Cichy, R. M., Pantazis, D., & Oliva, A. (2014). Resolving human object recognition in space and time. *Nature Neuroscience*, 17(3), 455–462.  
<https://doi.org/10.1038/nn.3635>
- Clarke, A., Devereux, B. J., Randall, B., & Tyler, L. K. (2015). Predicting the Time Course of Individual Objects with MEG. *Cerebral Cortex*, 25(10), 3602–3612.  
<https://doi.org/10.1093/cercor/bhu203>
- Clarke, A., & Tyler, L. K. (2014). Object-Specific Semantic Coding in Human Perirhinal Cortex. *Journal of Neuroscience*, 34(14), 4766–4775.  
<https://doi.org/10.1523/JNEUROSCI.2828-13.2014>
- Coggan, D. D., Baker, D. H., & Andrews, T. J. (2016). The Role of Visual and Semantic Properties in the Emergence of Category-Specific Patterns of Neural Response in the Human Brain. *ENeuro*, 3(4), ENEURO-0158.  
<https://doi.org/10.1523/ENEURO.0158-16.2016>
- Coggan, D. D., Liu, W., Baker, D. H., & Andrews, T. J. (2016). Category-selective patterns of neural response in the ventral visual pathway in the absence of categorical information. *NeuroImage*, 135, 107–114.  
<https://doi.org/10.1016/j.neuroimage.2016.04.060>

- Connolly, A. C., Guntupalli, J. S., Gors, J., Hanke, M., Halchenko, Y. O., Wu, Y.-C., ... Haxby, J. V. (2012). The Representation of Biological Classes in the Human Brain. *Journal of Neuroscience*, 32(8), 2608–2618.  
<https://doi.org/10.1523/JNEUROSCI.5547-11.2012>
- Connolly, A. C., Sha, L., Guntupalli, J. S., Oosterhof, N., Halchenko, Y. O., Nastase, S. A., ... Haxby, J. V. (2016). How the Human Brain Represents Perceived Dangerousness or “Predacity” of Animals. *The Journal of Neuroscience*, 36(19), 5373–5384. <https://doi.org/10.1523/JNEUROSCI.3395-15.2016>
- Contini, E. W., Wardle, S. G., & Carlson, T. A. (2017). Decoding the time-course of object recognition in the human brain: From visual features to categorical decisions. *Neuropsychologia*, 105(Supplement C), 165–176.  
<https://doi.org/10.1016/j.neuropsychologia.2017.02.013>
- DiCarlo, J. J., Zoccolan, D., & Rust, N. C. (2012). How Does the Brain Solve Visual Object Recognition? *Neuron*, 73(3), 415–434.  
<https://doi.org/10.1016/j.neuron.2012.01.010>
- Downing, P. (2006). Domain Specificity in Visual Cortex. *Cerebral Cortex*, 16(10), 1453–1461. <https://doi.org/10.1093/cercor/bhj086>
- Downing, P., Jiang, Y., Shuman, M., & Kanwisher, N. (2001). A cortical area selective for visual processing of the human body. *Science*, 293(5539), 2470–2473.
- Duda, R. O., Hart, P. E., & Stork, D. G. (2001). *Pattern Classification*. John Wiley & Sons.
- Epstein, R., Harris, A., Stanley, D., & Kanwisher, N. (1999). The parahippocampal place area: Recognition, navigation, or encoding? *Neuron*, 23(1), 115–125.
- Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature*, 392(6676), 598–601.

- Farah, M. J. (1996). Is face recognition ‘special’? Evidence from neuropsychology. *Behavioural Brain Research*, 76(1–2), 181–189. [https://doi.org/10.1016/0166-4328\(95\)00198-0](https://doi.org/10.1016/0166-4328(95)00198-0)
- Gainotti, G. (2000). What the Locus of Brain Lesion Tells us About the Nature of the Cognitive Defect Underlying Category-Specific Disorders: A Review. *Cortex*, 36(4), 539–559. [https://doi.org/10.1016/S0010-9452\(08\)70537-9](https://doi.org/10.1016/S0010-9452(08)70537-9)
- Gobbini, M. I., Gentili, C., Ricciardi, E., Bellucci, C., Salvini, P., Laschi, C., ... Pietrini, P. (2011). Distinct Neural Systems Involved in Agency and Animacy Detection. *Journal of Cognitive Neuroscience*, 23(8), 1911–1920. <https://doi.org/10.1162/jocn.2010.21574>
- Goddard, E., Carlson, T. A., Dermody, N., & Woolgar, A. (2016). Representational dynamics of object recognition: Feedforward and feedback information flows. *NeuroImage*, 128, 385–397. <https://doi.org/10.1016/j.neuroimage.2016.01.006>
- Gray, H. M., Gray, K., & Wegner, D. M. (2007). Dimensions of Mind Perception. *Science*, 315(5812), 619–619. <https://doi.org/10.1126/science.1134475>
- Grootswagers, T., Wardle, S. G., & Carlson, T. A. (2016). Decoding Dynamic Brain Patterns from Evoked Responses: A Tutorial on Multivariate Pattern Analysis Applied to Time Series Neuroimaging Data. *Journal of Cognitive Neuroscience*, 29(4), 677–697. [https://doi.org/10.1162/jocn\\_a\\_01068](https://doi.org/10.1162/jocn_a_01068)
- Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and Overlapping Representations of Faces and Objects in Ventral Temporal Cortex. *Science*, 293(5539), 2425–2430. <https://doi.org/10.1126/science.1063736>
- Haxby, J. V., Horwitz, B., Ungerleider, L. G., Maisog, J. M., Pietrini, P., & Grady, C. L. (1994). The functional organization of human extrastriate cortex: a PET-rCBF study of selective attention to faces and locations. *Journal of Neuroscience*, 14(11), 6336–6353.

- Haynes, J.-D. (2015). A primer on pattern-based approaches to fMRI: principles, pitfalls, and perspectives. *Neuron*, 87(2), 257–270.
- Huth, A. G., Nishimoto, S., Vu, A. T., & Gallant, J. L. (2012). A Continuous Semantic Space Describes the Representation of Thousands of Object and Action Categories across the Human Brain. *Neuron*, 76(6), 1210–1224.  
<https://doi.org/10.1016/j.neuron.2012.10.014>
- Iordan, M. C., Greene, M. R., Beck, D. M., & Fei-Fei, L. (2016). Typicality sharpens category representations in object-selective cortex. *NeuroImage*, 134, 170–179.  
<https://doi.org/10.1016/j.neuroimage.2016.04.012>
- Ishai, A., Ungerleider, L. G., Martin, A., Schouten, J. L., & Haxby, J. V. (1999). Distributed representation of objects in the human ventral visual pathway. *Proceedings of the National Academy of Sciences*, 96(16), 9379–9384.  
<https://doi.org/10.1073/pnas.96.16.9379>
- Jaccard, P. (1901). *Etude comparative de la distribution florale dans une portion des Alpes et du Jura*. Impr. Corbaz.
- Kaiser, D., Azzalini, D. C., & Peelen, M. V. (2016). Shape-independent object category responses revealed by MEG and fMRI decoding. *Journal of Neurophysiology*, 115(4), 2246–2250. <https://doi.org/10.1152/jn.01074.2015>
- Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The Fusiform Face Area: A Module in Human Extrastriate Cortex Specialized for Face Perception. *Journal of Neuroscience*, 17, 4302–4311.
- Kiani, R., Esteky, H., Mirpour, K., & Tanaka, K. (2007). Object Category Structure in Response Patterns of Neuronal Population in Monkey Inferior Temporal Cortex. *Journal of Neurophysiology*, 97(6), 4296–4309.  
<https://doi.org/10.1152/jn.00024.2007>
- Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., Broussard, C., & others. (2007). What's new in Psychtoolbox-3. *Perception*, 36(14), 1.

- Konkle, T., & Caramazza, A. (2013). Tripartite Organization of the Ventral Stream by Animacy and Object Size. *Journal of Neuroscience*, 33(25), 10235–10242.  
<https://doi.org/10.1523/JNEUROSCI.0983-13.2013>
- Kriegeskorte, N., & Kievit, R. A. (2013). Representational geometry: integrating cognition, computation, and the brain. *Trends in Cognitive Sciences*, 17(8), 401–412.  
<https://doi.org/10.1016/j.tics.2013.06.007>
- Kriegeskorte, N., Mur, M., & Bandettini, P. A. (2008). Representational similarity analysis - connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, 2(4). <https://doi.org/10.3389/neuro.06.004.2008>
- Kriegeskorte, N., Mur, M., Ruff, D. A., Kiani, R., Bodurka, J., Esteky, H., ... Bandettini, P. A. (2008). Matching Categorical Object Representations in Inferior Temporal Cortex of Man and Monkey. *Neuron*, 60(6), 1126–1141.  
<https://doi.org/10.1016/j.neuron.2008.10.043>
- Mahon, B. Z., Milleville, S. C., Negri, G. A. L., Rumiat, R. I., Caramazza, A., & Martin, A. (2007). Action-Related Properties Shape Object Representations in the Ventral Stream. *Neuron*, 55(3), 507–520. <https://doi.org/10.1016/j.neuron.2007.07.011>
- Martin, A. (2016). GRAPES—Grounding representations in action, perception, and emotion systems: How object properties and categories are represented in the human brain. *Psychonomic Bulletin & Review*, 1–12.
- Nili, H., Wingfield, C., Walther, A., Su, L., Marslen-Wilson, W., & Kriegeskorte, N. (2014). A Toolbox for Representational Similarity Analysis. *PLoS Computational Biology*, 10(4), e1003553. <https://doi.org/10.1371/journal.pcbi.1003553>
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10(4), 437–442.
- Pereira, F., Mitchell, T., & Botvinick, M. (2009). Machine learning classifiers and fMRI: a tutorial overview. *Neuroimage*, 45(1), S199–S209.



- Pinsk, M. A., Arcaro, M., Weiner, K. S., Kalkus, J. F., Inati, S. J., Gross, C. G., & Kastner, S. (2009). Neural representations of faces and body parts in macaque and human cortex: a comparative fMRI study. *Journal of Neurophysiology*, *101*(5), 2581–2600.
- Proklova, D., Kaiser, D., & Peelen, M. V. (2016). Disentangling Representations of Object Shape and Object Category in Human Visual Cortex: The Animate–Inanimate Distinction. *Journal of Cognitive Neuroscience*, 1–13.  
[https://doi.org/10.1162/jocn\\_a\\_00924](https://doi.org/10.1162/jocn_a_00924)
- Puce, A., Allison, T., Asgari, M., Gore, J. C., & McCarthy, G. (1996). Differential Sensitivity of Human Visual Cortex to Faces, Letterstrings, and Textures: A Functional Magnetic Resonance Imaging Study. *Journal of Neuroscience*, *16*(16), 5205–5215.
- Riesenhuber, M., & Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nature Neuroscience*, *2*(11), 1019–1025. <https://doi.org/10.1038/14819>
- Ritchie, J. B., Tovar, D. A., & Carlson, T. A. (2015). Emerging Object Representations in the Visual System Predict Reaction Times for Categorization. *PLoS Comput Biol*, *11*(6), e1004316. <https://doi.org/10.1371/journal.pcbi.1004316>
- Sergent, J., Ohta, S., & Macdonald, B. (1992). Functional neuroanatomy of face and object processing: a positron emission tomography study. *Brain*, *115*(1), 15–36.
- Serre, T., Kreiman, G., Kouh, M., Cadieu, C., Knoblich, U., & Poggio, T. (2007). A quantitative theory of immediate visual recognition. *Progress in Brain Research*, *165*, 33–56.
- Sha, L., Haxby, J. V., Abdi, H., Guntupalli, J. S., Oosterhof, N. N., Halchenko, Y. O., & Connolly, A. C. (2015). The Animacy Continuum in the Human Ventral Vision Pathway. *Journal of Cognitive Neuroscience*, *27*(4), 665–678.  
[https://doi.org/10.1162/jocn\\_a\\_00733](https://doi.org/10.1162/jocn_a_00733)

- Taylor, J. C., & Downing, P. E. (2011). Division of labor between lateral and ventral extrastriate representations of faces, bodies, and objects. *Journal of Cognitive Neuroscience*, 23(12), 4122–4137.
- Thorpe, S., Fize, D., & Marlot, C. (1996). Speed of processing in the human visual system. *Nature*, 381(6582), 520–522.
- Tong, F., Nakayama, K., Moscovitch, M., Weinrib, O., & Kanwisher, N. (2000). Response properties of the human fusiform face area. *Cognitive Neuropsychology*, 17(1), 257–280. <https://doi.org/10.1080/026432900380607>
- Van de Nieuwenhuijzen, M. E., Backus, A. R., Bahramisharif, A., Doeller, C. F., Jensen, O., & van Gerven, M. A. (2013). MEG-based decoding of the spatiotemporal dynamics of visual category perception. *NeuroImage*, 83, 1063–1073.
- Van den Hurk, J., Van Baelen, M., & Op de Beeck, H. P. (2017). Development of visual category selectivity in ventral visual cortex does not require visual experience. *Proceedings of the National Academy of Sciences*, 201612862. <https://doi.org/10.1073/pnas.1612862114>
- Wardle, S. G., & Ritchie, J. B. (2014). Can Object Category-Selectivity in the Ventral Visual Pathway Be Explained by Sensitivity to Low-Level Image Properties? *Journal of Neuroscience*, 34(45), 14817–14819. <https://doi.org/10.1523/JNEUROSCI.3566-14.2014>
- Warrington, E. K., & Shallice, T. (1984). Category specific semantic impairments. *Brain*, 107(3), 829–853.

---

REACTION TIMES PREDICT DYNAMIC  
BRAIN REPRESENTATIONS MEASURED  
WITH MEG FOR ONLY SOME OBJECT  
CATEGORISATION TASKS

---

Erika W. Contini<sup>1,2</sup>, Erin Goddard<sup>3\*</sup>, & Susan G. Wardle<sup>2,4\*</sup>

<sup>1</sup>Department of Cognitive Science, Macquarie University, Sydney, Australia

<sup>2</sup>ARC Centre of Excellence in Cognition and its Disorders and Perception in Action Research Centre,  
Macquarie University, Australia

<sup>3</sup>McGill Vision Research Group, Department of Ophthalmology, McGill University, Montreal, Canada

<sup>4</sup>Laboratory of Brain and Cognition, National Institutes of Health, Bethesda, MD, USA

(\* = equal contribution)

**Published as: Contini, E. W., Goddard, E., Wardle, S. G.  
(2021) Reaction times predict dynamic brain  
representations measured with MEG for only some object  
categorisation tasks, *Neuropsychologia*, Vol. 151, 107687,  
<https://doi.org/10.1016/j.neuropsychologia.2020.107687>.**



#### 4.1. Abstract

A recent approach to linking brain and behaviour processes in visual neuroscience is to relate behavioural categorisation reaction times (RTs) to brain representations measured with neuroimaging. In this framework, objects are assumed to be represented in a multidimensional activation space, with the distances between object representations indicating their degree of similarity. Previous studies have shown that faster RTs are associated with greater distances from a classification decision boundary for categorisation according to object animacy, however no other object categorisation tasks have been evaluated using this method. Using one stimulus set (which included typically animate and inanimate objects but was also designed to include visually confusing objects for classification; i.e., robots and toys) we conducted four categorisation tasks assessing animacy, living, moving, and human-similarity concepts and linked the categorisation RTs to MEG time-series decoding data. Our results show a sustained RT-distance relationship throughout the time course of object processing for not only animacy, but also categorisation according to human-similarity. Interestingly, this sustained RT-distance relationship was not observed for the living and moving category organisations, despite comparable classification accuracy of the MEG data across all four category organisations. Our findings show for the first time that behavioural RTs predict representational distance for an organisational principle other than animacy, however further research is needed to determine why this relationship is observed only for some category organisations and not others.



## 4.2. Introduction

Despite decades of research into human visual object recognition, we still do not have a clear understanding of the neural processes that afford us this rapid and efficient ability. With the implementation of multivariate pattern analysis (MVPA), neuroimaging studies have revealed coherent structure within neural object representations related to categorical concepts, for example, based on whether objects are animate or inanimate (Carlson, Tovar, Alink, & Kriegeskorte, 2013; Cichy, Pantazis, & Oliva, 2014), and this organisation is observed in both human and non-human primates (Kriegeskorte et al., 2008). Magnetoencephalography (MEG) studies evaluating object recognition processes have revealed that early neural activity is associated with the processing of low-level properties of visual stimuli, while more complex semantic and categorical information related to object category is represented relatively late (~120 - 240 ms) after stimulus onset (Carlson et al., 2013; Cichy et al., 2014; Clarke, Devereux, Randall, & Tyler, 2015; Goddard, Carlson, Dermody, & Woolgar, 2016). This temporal hierarchy of processing stages is consistent with fMRI data showing the flow of neural activity throughout the ventral-temporal processing stream, from early visual cortex (V1) through to the inferior temporal cortex (IT); a region known for its involvement in object processing (Cichy et al., 2014; Kriegeskorte et al., 2008; Vida, Nestor, Plaut, & Behrmann, 2017). Further evaluation suggests a hierarchy of object representations which initially include information that distinguishes individual object exemplars, with more abstract conceptual category information, like the animate/inanimate distinction, emerging later (Carlson et al., 2013; Cichy et al., 2014; Contini, Wardle, & Carlson, 2017; Kriegeskorte et al., 2008).

The categorical distinction of animate versus inanimate objects has been proposed as one of the fundamental organisational principles in the neural representation of complex objects. In recent years, a number of studies have extended this organisational principle, to incorporate factors such as real-world size (Julian, Ryan, & Epstein, 2017; Konkle & Caramazza, 2013; Konkle & Oliva, 2012) and biological class (Connolly et al., 2012; Sha et

al., 2015), the latter of which depicts object categories along a continuum based on biological similarity of species, spanning from humans at one end of the continuum, to invertebrates, and lastly inanimate objects at the other. Our previous work also provided support for the idea of a continuum, as opposed to a dichotomous representation of object categories, with a graded model of human-similarity (see Chapter 3). This model, generated from human behavioural ratings, outperformed a biological classes continuum model as well as the more established animacy dichotomy model when ambiguous object categories (toys and robots) were included in the stimulus set.

Despite the aforementioned findings, it remains unclear whether any of these neural representations is predictive of perception and behavior. While these studies aimed to characterise brain activity involved in object processing, detecting a pattern of activity using neuroimaging methods does not necessitate functional involvement of that signal in behavioural processes. For example, Williams et al. (2007) showed that patterns of activity in the lateral occipital cortex (LOC), as measured using fMRI, were related to correct object categorisation behaviour, while activity in retinotopic regions was not, despite both of these regions containing information about object category. This highlights the importance of looking beyond measured brain activation patterns to determine how this information is used functionally, and whether information contained in these patterns of activity is predictive of our daily experiences and behaviour - a fundamental goal of neuroscience research (Carlson, Goddard, Kaplan, Klein, & Ritchie, 2017; de-Wit, Alexander, Ekroll, & Wagemans, 2016; Krakauer, Ghazanfar, Gomez-Marin, MacIver, & Poeppel, 2017; Ritchie & Carlson, 2016).

One approach to link brain processes to behaviour is to evaluate the relationship between reaction times (RTs) and neural representations (Carlson, Ritchie, Kriegeskorte, Durvasula, & Ma, 2014; Grootswagers, Ritchie, Wardle, Heathcote, & Carlson, 2017; Philastides & Sajda, 2006; Ritchie, Tovar, & Carlson, 2015). A seminal study evaluating this relationship was conducted by Carlson et al. (2014), who evaluated the relationship between object representations in human inferior temporal cortex (IT), as measured with



fMRI and behavioural categorisation RTs. To delineate the object representations from the fMRI data into categories (such as animate and inanimate), the authors employed a linear classifier to categorise the object representations. This classifier creates a categorisation decision boundary, whereby objects positioned far from this boundary have clearer evidence of belonging to a particular category while those close to the decision boundary are less certain. The authors used decision-making theory to hypothesise that a correlation between the distance to classifier boundary and category decision RT would suggest that the representational space is meaningful for behavioural categorisation processes.

As predicted, Carlson et al. (2014) found that representational distance from a decision boundary can be used as a proxy for the decision-making process. They found a negative correlation between representational distance from the object classifier boundary and behavioural RTs, indicating that participants had faster RTs for objects that were represented further from the decision boundary, with slower RTs for less distinctly represented objects that appeared closer to the classifier boundary. One caveat to this result was noted when the relationships between representational distance and RT for animate and inanimate objects were evaluated separately. This analysis revealed that the animate objects alone were driving the correlation between RT and representational distance, with the authors suggesting that the inanimate object category may lack the cohesion that the animate category possesses (Carlson et al., 2014).

In light of this evidence for a potential link between brain and behaviour data for object categorisation, the RT-distance approach has more recently been applied to the brain activation patterns measured with MEG to evaluate the time-course of this relationship. Using MEG, Ritchie et al. (2015) replicated the animacy results from Carlson et al. (2014), showing a negative relationship between representational distance and RTs that was similarly driven by the animate objects. The authors additionally built on these findings by tracking this relationship throughout the time course of object processing. Ritchie et al. (2015) highlighted that the correlation between RT and representational distance became

significant for the animate objects around the peak decoding time, suggesting this is the optimal time for the brain to “read-out” object category information for use in behaviour. Grootswagers et al. (2017) also showed an asymmetrical relationship between MEG-derived brain representations and behaviour for animate compared to inanimate objects, and additionally evaluated this relationship using degraded stimuli, thus increasing task difficulty. They showed that degrading a stimulus resulted in the expected compression of the representational space for animate but not inanimate objects, further supporting an asymmetry in the representations of these two object categories (Grootswagers et al., 2017). These studies highlight the utility of RT data as a means to gain insight into which neural processes or organisational principles are likely relevant for behaviour.

Research using the RT-distance method to evaluate object categorisation processes has focussed on the animate versus inanimate distinction. However, in light of our previous findings, which show that a model based on human-similarity best accounted for the organisation of neural object representations (beyond that achieved by animacy), we sought to test whether a judgement of human-similarity also shows a relationship between representational distance and behavioural categorisation RTs. In addition to the established animacy dimension, and our recently proposed dimension of human-similarity, we also evaluated categorisation based on living and moving concepts. A number of both neuroimaging (Clarke, Taylor, & Tyler, 2010; Huth, Nishimoto, Vu, & Gallant, 2012) and behavioural (Filliter, McMullen, & Westwood, 2005) studies have evaluated the contribution of living and moving factors to object representations and categorisation processes, including patient studies that highlight category-specific deficits consistent with these factors (Gainotti, 2000; Masullo et al., 2012). As we are yet to fully understand how the brain processes object representations and object category structure, it is important that we continue to test alternate models and theories, such as these, against the widely researched animacy distinction, with the aim of finding better accounts of neural object recognition processes.

Here, we used a single stimulus set across all four of these plausible category organisations (animacy, living, moving, human-similarity), with the stimulus set designed such that it could be carved up in different ways depending on the categorisation task at hand. For example, a tree stimulus would be categorised as inanimate for an animacy categorisation task, but for a living/non-living categorisation, it would be categorised as living. Thus, the tree would be grouped together in a category with animals for the living task, but not for the animacy task. To provide a more stringent test of the established animacy dimension as well as the plausible related dimensions of living, moving and human-similarity, we included stimuli that did not have clear category membership based on the established animacy criteria, namely those of robots and toys. These stimuli have physical features that resemble animate objects (i.e., faces/bodies) but are man-made and thus technically inanimate objects, so if the animacy category is the true organisational principle for these neural representations then these inanimate objects should be categorised as such.

### **4.3. Method**

Data for the behavioural and MEG components of the study were collected separately, with two different sets of participants. Details for each component are provided below.

#### **4.3.1. Stimuli**

The same set of object stimuli was utilised for both the behavioural and MEG components of the study. Stimuli were 120 naturalistic object images (see Chapter 3, Figure 1), including a range of objects from twelve object categories including six animate categories (humans, primates, domestic animals, birds, fish, invertebrates), and six inanimate categories (plants, robots, machines, tool, toys, and other stationary objects).

#### **4.3.2. Behavioural Experiment and Reaction-Time Data**

A total of 400 Amazon's Mechanical Turk workers, residing in either Canada or the United States of America, each completed one of four behavioural RT categorisation tasks online ( $n = 100$  per task). Each task followed a similar format, whereby a trial consisted of

a single object exemplar presented centrally on a grey background. Using a two-alternative forced-choice paradigm, participants were required to categorise each of the 120 objects as quickly as possible, using a standard keyboard button press. Images were each presented for 500ms, with the next trial triggered when the participant pressed a button to make their categorisation choice. No feedback was given to participants. Participants' categorisation choice and reaction time data were collected for each of the 120 object exemplars. Experimental tasks and associated categorisation<sup>2</sup> options were as follows:

Task 1) Animacy: '*animals (living human/animal/insect)*' or '*not an animal*';

Task 2) Living: '*living*' or '*non-living*';

Task 3) Moving: '*has movement*' or '*no movement*';

Task 4) Human: '*more like a human*' or '*less like a human*'.

Participants completed a pre-test practice task prior to commencing the experimental task to familiarise themselves with the task requirements and to also serve as a screening task to ensure participants were responding in line with the task instructions. The stimuli for the pre-test practice task consisted of eight trials using novel exemplars (these were not included in the experimental stimulus set) which participants were required to categorise, with a minimum accuracy rate of 75% to continue to the experimental task (however, no feedback was given to participants). Pre-test exemplars were the same for each of the four tasks, and included four animals (giraffe, parrot, bear, penguin) and four inanimate objects (laptop computer, palm tree, lamp, arm chair). Instructions differed slightly for Task 4 (Human), where participants were provided with a diagram showing four of the practice set exemplars along a continuum from *more like a human* to *less like a human* (i.e., bear, parrot,

---

<sup>2</sup> For clarity, category organisations used for the tasks and to group the stimuli in the current study are identified with a capital letter (e.g., Animacy, Living). Instances where no capital letter is used (e.g., animacy) refer to the concept itself, rather than the experimental task or design.

tree, lamp), so that participants understood that an object did not have to be a human to be categorised as *more like a human*.

### 4.3.3. MEG Experiment

The MEG experimental data were reused from our previous experiment (see Chapter 3). An overview of the methods for the MEG experiment is provided below.

**4.3.3.1. Participants.** Participants for the MEG component included 24 healthy adults from the Macquarie University community (18 female; mean age = 24.7, range = 18 - 37). All participants reported normal or corrected-to-normal vision, provided informed consent prior to participation, and were paid \$20/h for their time. The research in this study was approved by the Macquarie University Human Research Ethics Committee.

**4.3.3.2. MEG experimental task procedure.** The MEG experimental task was run on a Dell PC desktop computer using Matlab (Natick, MA), with the Psychophysics Toolbox extensions (Brainard, 1997; Kleiner et al., 2007; Pelli, 1997). Stimuli were projected onto a mirror located inside the MEG magnetically sealed room (MSR), using an InFocus IN508 projector. The mirror reflected this image onto the ceiling of the MSR, approximately 113 cm above the participant, who lay supine in the MEG chamber. Object images were segmented and presented on a grey background. The stimuli were approximately 4.9 x 4.9 degrees of visual angle (DVA), and were presented in one of four centrally overlapping screen quadrants located around a black central fixation cross (diameter of 0.1 DVA; see Chapter 3 for a detailed description). A left-right flipped version of each of the 120 object images was also included in the stimulus set, resulting in a total stimulus set of 240 images. we included the four stimulus locations and flipped version of the stimuli to reduce the effects of low-level stimulus properties on the data. Each stimulus was presented at each location three times, resulting in a total of 2880 trials for data analysis (240 stimuli x 4 locations x 3 repetitions). Additional trials, which were not analysed, included the first and

last trials of each block, and the 288 repeat trials that were included for the attention task (see below).

Subjects completed a one-back attention task during the MEG experiment, where they pressed a button whenever a stimulus was repeated consecutively. Participants received feedback about the accuracy of their performance on this task at the end of each block. In a single testing session of approximately one hour duration, participants completed eight blocks (pseudo-randomised across participants) consisting of 398 trials (3184 total trials). On each trial, the stimulus was presented for 100 ms, followed by a random inter-stimulus interval between 750 and 1000 ms where just the fixation cross was present on the screen. Stimulus presentation times were recorded using a photodetector attached to the mirror. The photodetector was aligned with a white square (1x1 deg) that appeared simultaneously on the screen with each image presentation. Within each block, the stimuli appeared in a predetermined pseudo-randomised order, such that there was equal probability that the preceding and following images for each trial would belong to any one of the 12 object categories.

#### **4.3.4. MEG Data Acquisition and Processing**

The neuromagnetic signal was recorded using a 160-channel whole-head axial gradiometer (KIT, Kanazawa, Japan) in the KIT-Macquarie Brain Research Laboratory. Data were acquired continuously, sampled at a rate of 1000 Hz, and were filtered online from 0.03 to 200 Hz. we then down-sampled the data offline, to 200 Hz, and epoched the individual trials into single events with a time window from 100 ms pre-stimulus to 600 ms post-stimulus. For the MEG decoding analyses, we subsequently reduced the dimensionality of the epoched data from the 160 gradiometers using Principal Components Analysis (PCA). A threshold was set to retain the first  $n$  components that accounted for 99% of the variance in the data, resulting in an average of 34 retained components (range = 14 – 72).

#### 4.3.5. Classification Analysis

For each MEG participant's data, we trained classifiers to perform a series of classifications on the components resulting from PCA of the MEG sensor recordings. In each case, we trained the classifier using linear discriminant analysis (LDA), which we implemented using the Matlab function `fittediscr.m` (from the Statistics and Machine Learning Toolbox) with 'DiscrimType' set to 'diaglinear'. For each classification, we divided the set of stimuli into two groups which differed according to the behavioural categorisation task in question (see *Behavioural experiment and reaction time data* section above for details of the behavioural tasks). For Task 1 (Animacy), the stimuli were grouped for the classifier according to the contemporary animacy definition used in previous studies (Carlson et al., 2013; Cichy et al., 2014; Sha et al., 2015) with the data from each trial labelled as 'animate' or 'inanimate' based on whether the stimulus for that trial was animate (human or animal) or inanimate (plant, robot, machine, tool, toy or other inanimate object). Similarly, for Task 2 (Living), the classification groupings were the same as for animacy, except that plants were grouped with the other animate objects to form the 'living' group, while the remaining objects comprised the 'non-living' group. The procedure for determining classification grouping differed for Tasks 3 and 4 (Moving and Human), as these categorisations do not specify clear category membership definitions for all stimuli within the data set. As such, the data were divided into two groups for the classifier according to the mean behavioural categorisation response for each stimulus (see Figure 2, plots C and D). For all four tasks, we balanced the number of stimuli in each of the two groups for the classifier by removing stimuli from the larger grouping until both groups were even. Stimuli to be removed were randomly selected such that, where possible, an even number were removed from each subcategory (e.g. humans, primates, tools etc) so as to maintain an even distribution of subcategory exemplars. This process resulted in slightly different stimulus numbers for each of the four tasks. To control for differences in power due to variations in set size, we repeated the analysis with a subset of the data, in which the number of stimuli for each of the four

tasks was reduced to 100 (50 per category subdivision for the classifier). Furthermore, we ran a second control analysis for the Animacy and Living category organisations where, instead of using predetermined category groupings, we divided the stimuli into groups for the classifier based on the mean behavioural categorisation performance (in line with the procedure for the Moving and Human category organisations) to ensure that any differences observed between the tasks in the results were not due to methodological differences in how the categories were defined.

Across all four category organisations, we performed two classifications for data from each time bin. First, we measured classifier performance by performing classification with four-fold cross-validation, in each case training on data from 75% of the trials and testing the classifier's accuracy in predicting the category (e.g. animate or inanimate) for the remaining 25% of trials in order to avoid circularity. Across the four instances, where we trained on 75% of the data and tested on the remainder, data from each trial were included in the test data set once. We averaged accuracy across these four cross-validation folds to obtain a single accuracy value for each MEG participant at each time bin.

Second, for each of the four category organisations, we trained a classifier on the same categorisation where data from all trials were included in the training set, in order to also obtain a single decision boundary for the classifier at each time bin, for each MEG participant. This decision boundary is a linear hyperplane in an  $n$ -dimensional space, where  $n$  is the number of components that were retained following the PCA reduction of the data. Data from each trial defined a point in this  $n$ -dimensional space, and linear discriminant analysis (LDA) seeks to find the linear hyperplane that best divides these points according to the labels given to each trial's data, so that data from one category (e.g. 'animals' trials) fall on one side of the decision boundary and data from the other category (e.g. 'not an animal') fall on the other side of the decision boundary. For each trial, we computed the Euclidean distance between the data point corresponding to that trial, and the single decision boundary. Since LDA often does not perfectly separate data according to their category



labels, some data points will be on the incorrect side of the decision boundary. To preserve this information (discussed in greater detail below), we multiplied the distance by negative one where the point was on the incorrect side of the decision boundary (i.e. if it would be labelled incorrectly by the classifier). Once we had the signed distance values for each trial, we averaged data across trials of the same exemplar to obtain a single signed distance value for each exemplar (for an example, see Figure 4). We then tested whether these signed distance values from the decision boundary calculated from the MEG data were correlated with reaction times for each categorization task, as described below.

#### **4.3.6. Correlating Distances in Activation Space with Reaction-Times**

In order to test whether decodable information about category membership in the neural signal is also predictive of behaviour, we correlated the representational distances obtained from the classification analysis (see above) with behavioural reaction times following methods used previously (Carlson et al., 2014; Grootswagers et al., 2017; Ritchie et al., 2015) with minor modifications as described below. For each of the four different behavioural categorisation tasks, individual RTs for the 120 object exemplars were normalised within subjects. We then calculated Spearman's rank-order correlation between the median normalised RT (from the behavioural categorisation task) and the mean exemplar representational distance from the decision boundary. We computed this correlation for each participant, at every time-point from -100 to 600 ms relative to stimulus onset, and then obtained a participant-averaged correlation providing a single overall time-course of the relationship between representational distance and RTs for each of the different categorisation tasks, as well as individual subcategory division time-course correlations. This follows the method performed by Grootswagers et al. (2017), which differs slightly from that of Ritchie et al. (2015) who computed the time-series correlation at the group level (rather than individual level, as done here) after averaging distances and RTs across all participants. Using bootstrapping to test for significance, we generated a null distribution of correlations by shuffling the 120 exemplar labels for each participant's MEG data and

correlated these with the median RTs for the object exemplars, whose labels had also been shuffled. We repeated this procedure 10000 times to obtain the null distribution of correlation values. We then compared the experimental correlation between representational distance and RTs with this null distribution (two-tailed t-test, adjusted for multiple comparisons across time points using a false discovery rate (FDR) of  $q < .01$ ).

## **4.4. Results**

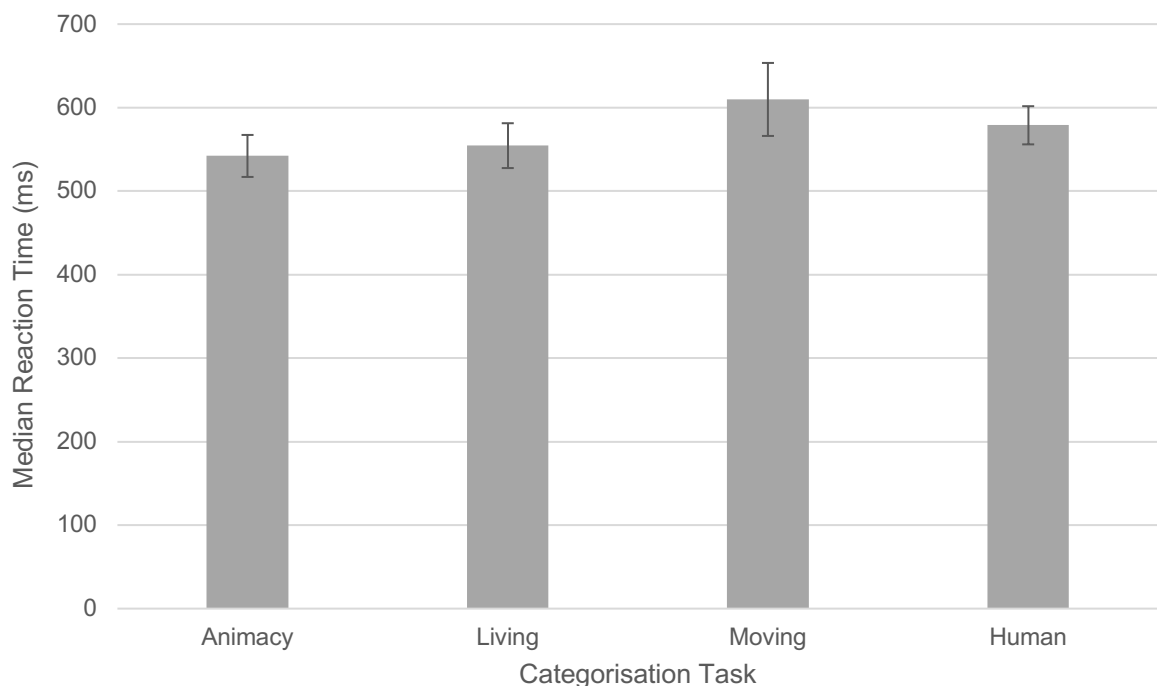
### **4.4.1. Behavioural Task Performance**

Using the same set of 120 object exemplars as the MEG study, we ran four separate behavioural categorisation tasks (Animacy, Living, Moving, and Human; see Methods for task descriptions) with a two-alternative forced choice design. Task accuracy was measured for the Animacy and Living tasks, as these two tasks have clearly defined category membership. Participants demonstrated good categorisation performance overall, obtaining 89% accuracy on average for both the Animacy and Living tasks. It was not possible to compute a meaningful accuracy score for the Moving and Human tasks as they do not have clearly defined binary category membership for each of the object exemplars. As such, category membership (for the classifier analysis) was based on participant categorisation, creating circularity if we were to compute measurements of accuracy. Furthermore, we did not specify a definition for the categorisation of “has movement” or “no movement”, with participants able to define their own criteria for categorisation. This, for example, could have centred around concepts such as voluntary movement, or moving parts, and as such is not clearly delineated. Similarly, the “more like a human” and “less like a human” categorisation is more of a continuum than a binary distinction, as we have shown previously (see Chapter 3).

Upon inspection of the median reaction times for each exemplar, we noted that the stingray exemplar was a consistent outlier, with a slower categorisation reaction time across each of the four tasks. This consistently slow response time may indicate that participants had difficulty identifying the object depicted in this exemplar. As such, ambiguity in

exemplar identity may be conflated with categorisation RT, causing the RT to no longer be a meaningful representation of categorisation decision time. Therefore, the RT and MEG data for this object were removed from further analyses. Overall median reactions times are shown in Figure 1 (excluding the stingray exemplar), and were similar across all four categorisation tasks (range 542 – 609 ms).

We conducted a series of correlations between the median exemplar RTs for each of the four categorisation tasks in order to assess how similar the behavioural processes were across tasks (see Table 1). Results showed significant correlations between every pair-wise comparison of the four sets of categorisation RTs (Pearson correlations, corrected for multiple comparisons using a false discovery rate (FDR) of  $q < .05$ ). Given the same stimulus set was used for each of these tasks, with only a subset of exemplars expected to shift category grouping across tasks, it is perhaps unsurprising that we see this level of overlap in the RTs.



*Figure 1.* Median categorisation reaction times for each of the four categorisation tasks. Error bars indicate one standard deviation.

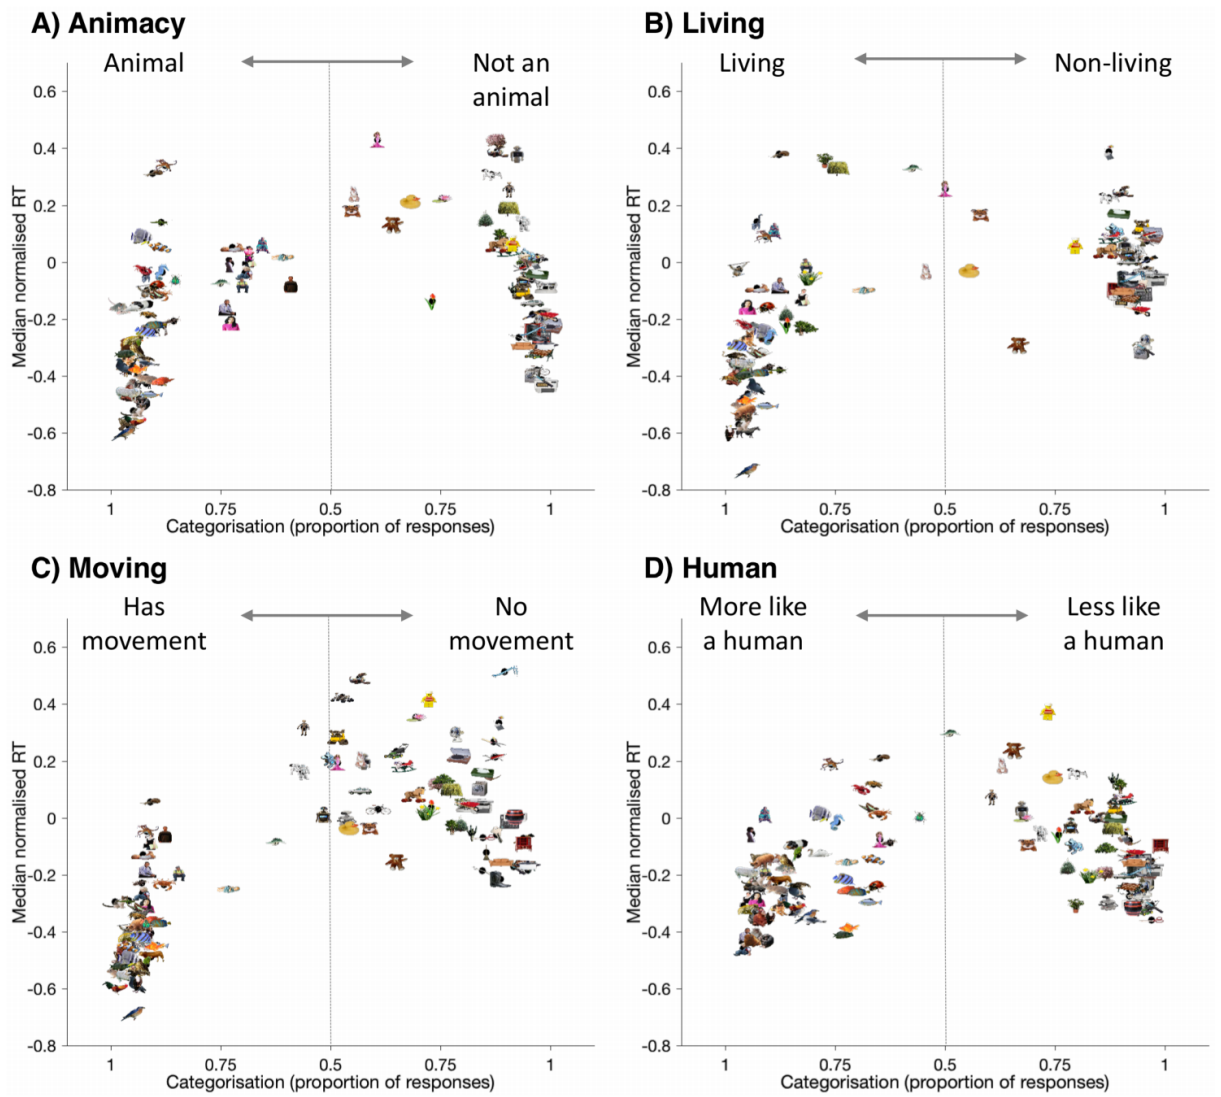
Table 1

*Correlations Between Categorisation Task Reaction-Times*

Task	Animacy	Living	Moving	Human
Animacy	1			
Living	.705	1		
Moving	.630	.729	1	
Human	.668	.453	.398	1

*Note:* Values represent Pearson correlations. All values are significant, corrected for multiple comparisons using a false discovery rate (FDR) of  $q < .05$ .

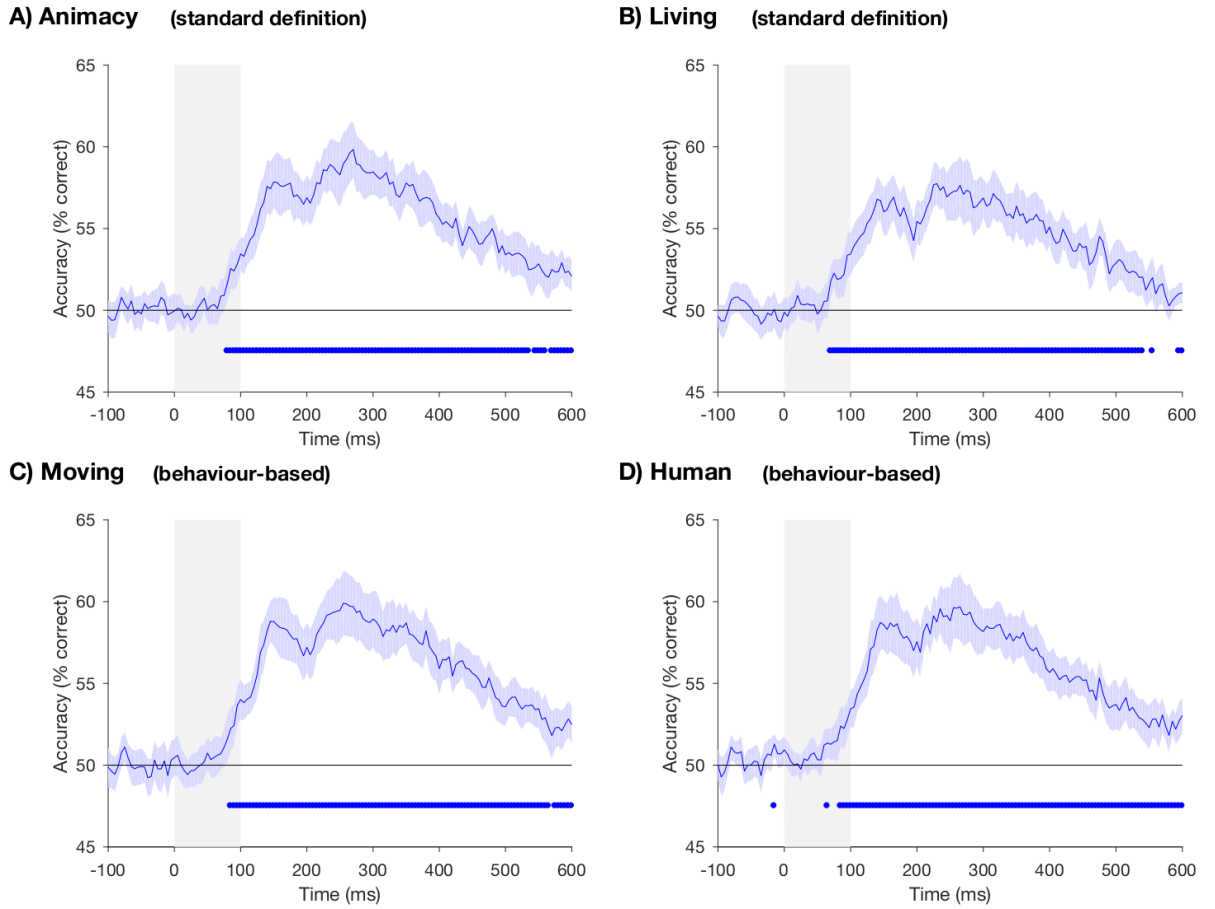
The relationship between mean categorisation performance and normalised RT for the exemplars is plotted in Figure 2 for each of the four categorisation tasks. These plots only include exemplars used for the classification analysis for each task (where some exemplars have been removed to balance the stimulus sets; see Section 4.3.5). As expected, for each task the fastest RTs were for those exemplars which were most consistently categorised, while those exemplars that were less consistently categorised tended to have slower RTs.



*Figure 2.* Categorisation performance and reaction time for the different behavioural tasks. Objects were categorised as per the following binary categorisation tasks: A) Animacy: living human/animal/insect vs. not an animal; B) Living: living vs. non-living; C) Moving: has movement vs. no movement; and D) Human: more like a human vs. less like a human. RTs were normalised to have a mean of zero. The x-axis shows the behavioural categorisation for each exemplar as a proportion of all categorisations, with objects represented towards the extreme ends of the axis (closer to a categorisation value of 1) indicating greater consistency in their categorisation as that category (a score of 1 indicates *all* participants consistently categorised the exemplar as that category). Objects represented towards the centre of the x-axis (closer to a proportion of 0.5) had greater variability in categorisation across participants. Note, the stingray data were removed from the data set as RTs to this exemplar were consistent outliers.

#### 4.4.2. Exemplar Decoding from MEG Time-Series Data

We first sought to determine whether the four different category organisations (Animacy, Living, Moving, and Human) of the stimulus set could be decoded from the MEG time-series data at the individual trial level. Results from the overall time-varying decoding performance (Figure 3) revealed significant and sustained decoding of object category membership for all four category organisations, consistent with results from previous time-series decoding studies (Carlson et al., 2013; Cichy et al., 2014; Goddard et al., 2016; Grootswagers et al., 2017; Kaiser, Azzalini, & Peelen, 2016; Kaiser, Oosterhof, & Peelen, 2016; Ritchie et al., 2015; Simanova, Gerven, Oostenveld, & Hagoort, 2010). The time of decoding onset (i.e. first time-point of significant decoding) was similar for all four category organisations, with onsets ranging from 70 – 85 ms. There was slightly more variation in the peak decoding times (i.e. the time-point at which decoding performance is highest) across the four category organisations, with the Living organisation showing the earliest peak decoding time, at 230 ms, followed by Moving (255 ms), Human (265 ms), and lastly Animacy, with a peak at 270 ms. Overall decoding performance across the four category organisations was very similar, ranging from 58 – 59% accuracy at their respective peak decoding times. It is unsurprising that the decoding time-courses are so similar, as there is substantial overlap in the way the objects were categorised for each of the four different tasks (i.e., only a subset of the exemplars are swapping category label across the four decoding plots).



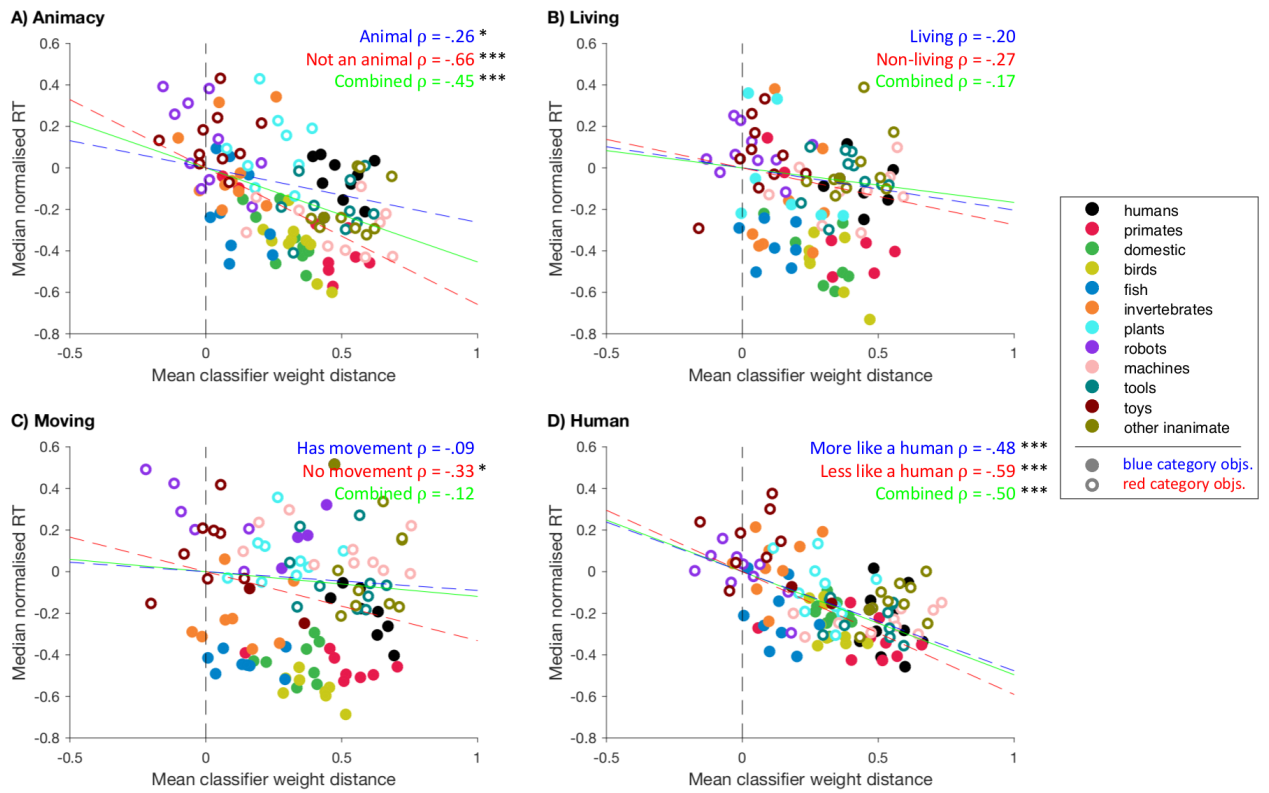
*Figure 3.* Overall MEG decoding performance for all stimuli according to the different category organisations for the classifier. Object exemplars in the upper plots were divided into two groups for the classifier based on the established standard category definitions of A) Animacy (animate/inanimate) or B) Living (living/non-living), while for the lower plots, categories for the classifier were determined based on mean behavioural task categorisation of objects from the Moving (has movement/no movement) and Human (more like a human/less like a human) categorisation tasks. Grey bar indicates the period the image appeared on the screen. Error bars indicate 95% between-subject confidence intervals. Blue dots along the x-axis indicate time points at which decoding performance was significantly above chance (two-tailed t-test, corrected for multiple comparisons using a false discovery rate (FDR) of  $q < .01$ ).

#### 4.4.3. Overall Representational Distance and Reaction-Time Correlations

The scatter plots in Figure 4 show the signed representational distance (i.e., mean classifier weight distance) as a function of median normalised RT for each object exemplar in each of the four category organisations (Animacy, Living, Moving, and Human). For each individual category organisation, we computed Spearman's  $\rho$  using the group means for a 50 ms time window that centred on the peak decoding time for each of the four category organisations. For each, we computed the overall correlation (Figure 4, green solid line), with all object exemplars included, as well as separate correlations (Figure 4, red and blue dashed lines) for each of the two category subdivisions provided for the classifier (e.g., correlations for 'animal' and 'not an animal' exemplars separately). When evaluating the correlations with all objects included, results showed a significant negative correlation between representational distance and RT for Animacy (Spearman's  $\rho = -.45, p < .001$ ) and Human (Spearman's  $\rho = -.50, p < .001$ ), indicating that greater representational distance from the category decision boundary is associated with faster behavioural categorisation RTs around the time of peak decoding for these two category organisations. We did not find the same relationship for Living (Spearman's  $\rho = -.17$ ) or Moving (Spearman's  $\rho = -.12$ ), with no significant overall correlation for either of these category organisations ( $p > .05$  for both).

Interestingly, both subdivisions within the Animacy category organisation show a similar relationship, with sustained significant negative correlations between representational distance and RT. This differs from previous studies which only found a significant relationship for animate objects (Carlson et al., 2014; Ritchie et al., 2015), or otherwise a comparatively shorter period of significance for inanimate objects (Grootswagers et al., 2017). Below we discuss differences between our stimulus set and those used by others that may account for this discrepancy. Like the Animacy category organisation, the Human category organisation also shows a negative correlation that is significant and sustained for each of its subdivisions.



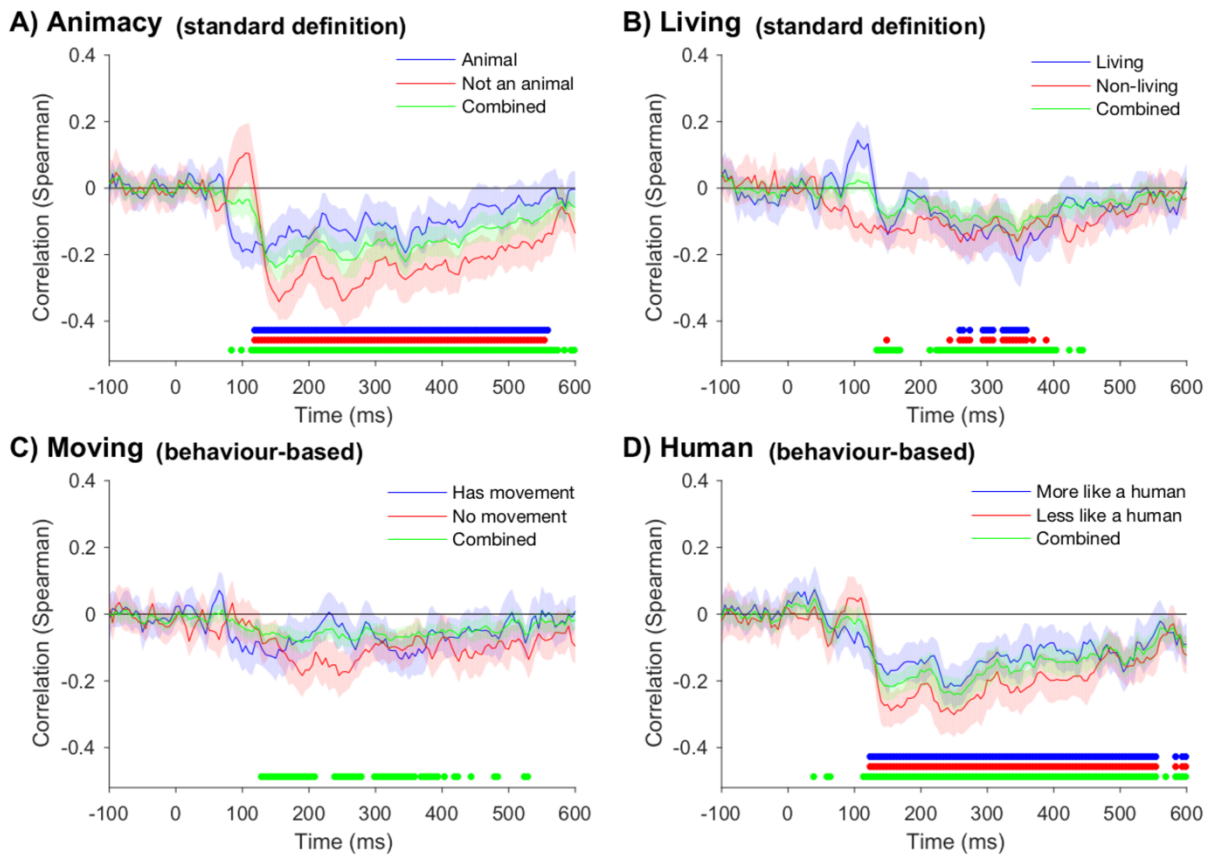


**Figure 4.** Scatter plots for the rank-order correlation between exemplar representational distance and reaction time for each of the four category organisations. Plots are based on the mean results over a 50 ms period centred around the MEG peak decoding time for each of the four tasks as follows: A) Animacy peak at 270 ms; B) Living peak at 230 ms; C) Moving peak at 255 ms; D) Human peak at 265 ms. Individual objects are shown as either filled dots or hollow dots to indicate category membership according to the relevant classifier, while colours indicate subcategory membership (see key for subcategory descriptions). The dashed line at zero on the x-axis shows the decision boundary for the classifier. Objects categorised correctly by the classifier have a positive x-value, while those incorrectly categorised (according to individual category organisations) have a negative x-value. Coloured lines show Spearman's  $\rho$  for all objects combined (green line), as well as correlations for each category subdivision (red and blue lines) ( $*p < .05$ ;  $**p < .01$   $***p < .001$ ).

Our calculation of the representational distance differed to that used in previous studies (Carlson et al., 2014; Grootswagers et al., 2017; Ritchie et al., 2015) as we used signed distance values (as opposed to taking the absolute value of the distance) to preserve incorrect classifier performance (i.e., when exemplars are classified on the wrong side of the decision boundary according to the task-dependent category membership). The distance to the classifier boundary includes a sign: exemplars of the first predicted category have a positive sign and those of the second predicted category have a negative sign. Taking the absolute value of this distance means that a linear correlation can be used to test the prediction that RT will be longest for stimuli that are closest to this boundary. Instead of taking the absolute distance, we achieved this transformation by multiplying the signed distances of exemplars of the second category by -1. This subtle difference in analysis only makes a difference where there are exemplars that the classifier consistently misclassifies. In previous studies, where the stimulus set was comprised of exemplars that clearly belonged to one category or the other, it would likely not have made a difference to take the absolute value of the distance. However, when there are exemplars that the classifier consistently misclassifies (as for our stimulus set, which deliberately included ambiguous exemplars such as robots and toys), the absolute value is inappropriate, since it treats an object for which the classifier is significantly below chance in the same way as one for which the classifier is significantly above chance. This is equivalent to predicting that those exemplars that are consistently misclassified will have shorter RTs than those for which the classifier is at chance. By using the signed distance, we instead predicted that exemplars which the classifier consistently misclassifies would have the longest RTs, since it would suggest that the subject would have to actively suppress the information carried on this neural representation in order to perform the task successfully. For comparison to previous studies, we also verified these results using the absolute value distance method (see Section 4.7 Appendix), and found that the results were overall consistent with the current analysis.

#### **4.4.4. Evaluating Representational Distance as a Predictor of Categorisation RTs Throughout the Decoding Time-Course**

Having established that there is a relationship between representational distance and RT for at least two of our tasks around the time of peak decoding, we then evaluated the time-course of this relationship. Categorisation reaction times were correlated with the representational distance calculated from the MEG data from 100ms prior to stimulus onset, to 600ms after stimulus onset. The time-series correlations for each of the four category organisations are shown in Figure 5. Results for each category organisation are presented two-fold: one version showing the correlation between representational distance and RTs throughout the time-course for individual subcategories (i.e., depicting ‘animal’ and ‘not an animal’ category subdivisions separately), and a second version showing the overall correlation for both categories combined. This was done for comparison with previous studies, which showed the overall correlation of RTs with representational distance was primarily driven by one subcategory correlation (i.e., animate objects; see Carlson et al., 2014; Grootswagers et al., 2017; Ritchie et al., 2015). For the Animacy category organisation (Figure 5A), RTs correlated with the representational distance from the classifier boundary, with a significant and sustained negative correlation from 115 ms post stimulus-onset for the overall combined correlation, and from 120 ms for both category subdivisions individually. This result replicates previous findings relating behavioural categorisation RTs for animacy to representational distances (Carlson et al., 2014; Grootswagers et al., 2017; Ritchie et al., 2015) while using a new stimulus set.



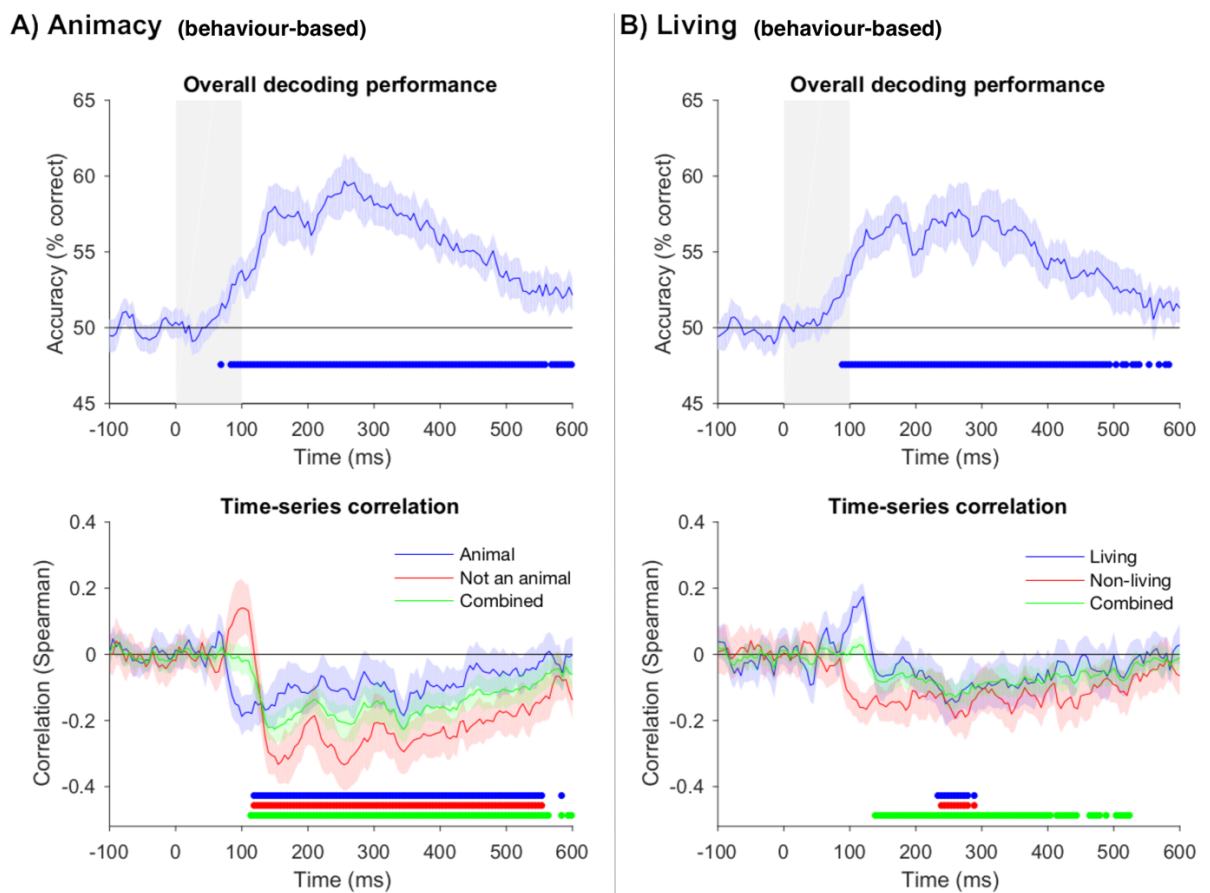
*Figure 5.* Time-series correlations between representational distance from the decision boundary set by the classifier and categorisation reaction times for each of the different category organisations. Time-varying rank-order correlation (Spearman's  $\rho$ ) between the mean object representational distance and the median normalised behavioural categorisation RT for the four category organisations: A) Animacy, B) Living, C) Moving, and D) Human. In each plot, overall correlations are plotted for all stimuli combined, as well as separately for each category subdivision (see figure legends). Shaded error bars indicate the 95% confidence interval of the between-subject means. Coloured dots along the x-axis indicate time points at which the correlation is significantly above chance (two-tailed t-test, corrected for multiple comparisons using a false discovery rate (FDR) of  $q < .01$ ).

Next, we extended this approach by evaluating the relationship between categorisation reaction time and MEG representational distance for three new tasks. Previous studies have only considered animacy categorisation, which does not allow for comparison of the animacy model with alternatives. RTs for the Human category organisation (Figure 5D) were also correlated with representational distance throughout the time-course, with a significant and sustained overall combined correlation from 115 ms, and significant sustained individual category subdivision correlations from 125 ms. Notably, the same sustained relationship was not exhibited for the Living and Moving categorisation tasks (Figure 5B and C), despite similarly successful decoding of the individual exemplars (Figure 3B and C; discussed above). The overall combined correlation between RT and representational distance for the Living category organisation reached significance from 135 – 170 ms, and was sustained again for a later period, from 225 – 405 ms, with only patchy periods of significance between 260 – 360 ms for the individual category subdivisions. Similarly for the Moving task, the overall correlation between representational distance and RT was significant for patchy periods of the time-course from 130 – 395 ms, with no significant time-points for the category subdivisions individually. Furthermore, the peak times for the overall combined RT-distance correlation did not consistently coincide with peak classifier decoding times (cf. Figure 3). Peak RT-distance times occurred before the classifier decoding peak for Animacy (RT-distance peak at 150 ms; decoding peak at 270 ms), and Moving (RT-distance peak at 190 ms; decoding peak at 255 ms) category organisations. In contrast to this, the timing peaks for the Human category organisation were similar (RT-distance peak at 250 ms; decoding peak at 265 ms), while the Living category organisation had a peak RT-distance correlation that occurred after the classifier peak decoding time (RT-distance peak at 354 ms; decoding peak at 230 ms). The variability in results across the four category organisations provides evidence for differing relationships between RTs and representational distances depending on task and classification criteria, indicative of a more complex relationship than suggested by previous studies.

For the main analyses above, we used predefined category organisations for the Animate and Living tasks when splitting the data for the classifier (i.e., the data were split into animate/inanimate and living/non-living subcategories using standard definitions of these category labels). However, this process differed from that used for the Moving and Human classifier divisions, which were based on the mean behavioural categorisation choice (as these tasks do not have clearly defined subcategory groupings). Furthermore, by using predefined groupings for Animacy and Living, some stimuli were categorised by participants on average in the behavioural task as belonging to the category contrary to this standard definition. If neural category representations are similar to the categorisation used by participants in the behavioural task, this may affect the reported link between RTs and the structure of the object representations measured with MEG.

Given this methodological difference between the definition of the four categorisation tasks, we ran an additional set of control analyses for the Animacy and Living tasks whereby we split the data for the classifier based on categorisation in the behavioural tasks, in order to be consistent with the method used for the Moving and Human tasks. Overall decoding performance and time-series correlations between RT and representational distance are shown in Figure 6 for these control analyses. The results from using behavioural performance to define the Animate and Living category boundaries are consistent with the results obtained using the predefined category boundaries in the original analysis (cf. Figures 3 and 5, panels A and B). The onset of significant decoding for the Animacy (behaviour-based) category organisation was at 85 ms (peak at 255 ms), while decoding onset for the Living (behaviour-based) organisation was slightly later than in the original analysis, at 90 ms (peak at 265 ms). Classifier decoding accuracy was 60% and 58% for the behaviour-based Animacy and Living organisations respectively. There is a significant and sustained relationship between representational distance and RT for the Animacy task, and as observed previously, there is also a significant correlation between representational distance and RT for the Living task, though this is across considerably fewer time-points than for Animacy

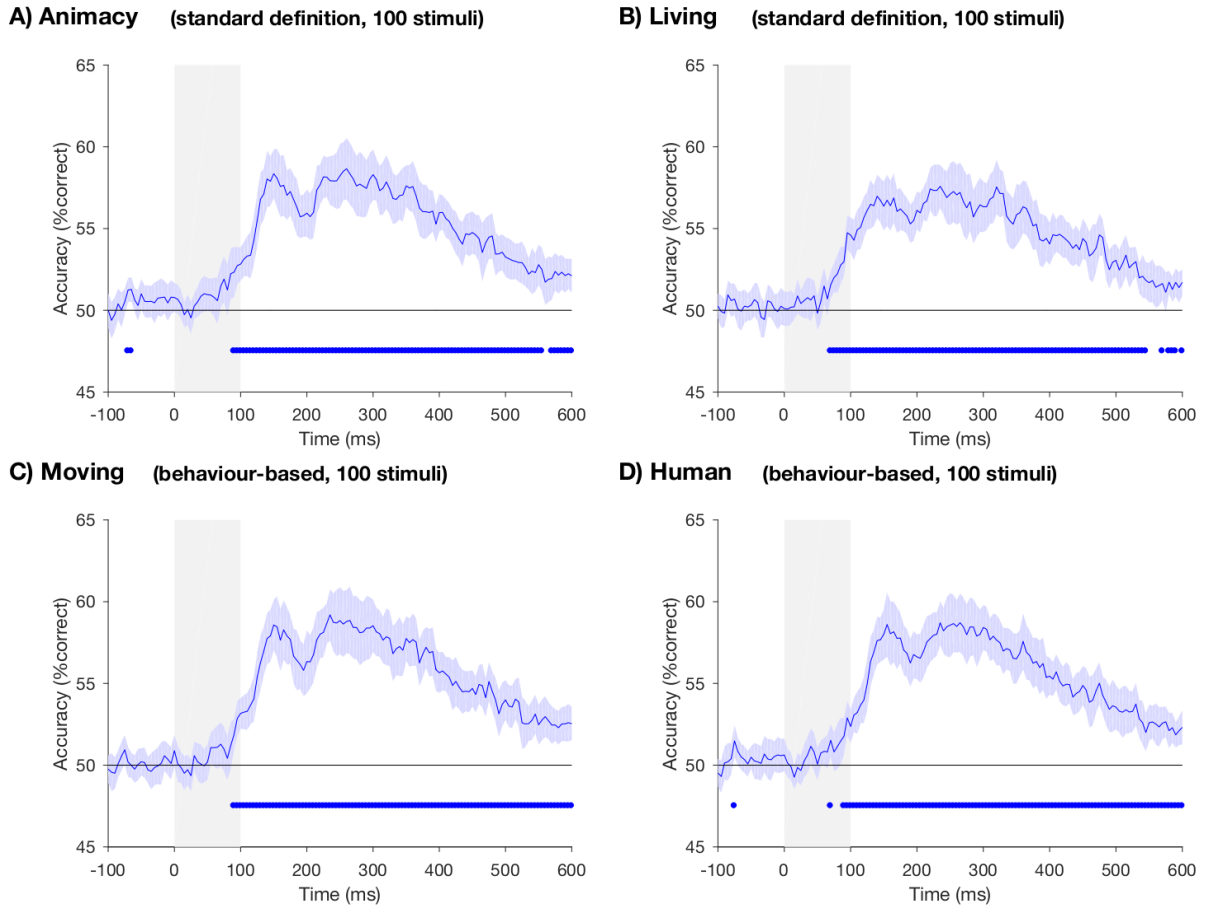
when evaluating the category subdivision correlations individually. Overall, the control analysis reveals that the results of the original analysis are robust, with little difference in the relationship between representational distance and RT when using standard definitions of Animacy and Living category groupings, compared to when objects are grouped based on participant categorisation behaviour.



*Figure 6.* Overall MEG decoding performance and time-series correlations of representational distance and categorisation reaction times for the Animacy and Living tasks using participant categorisation behaviour to define category membership for the classifier. Upper panels: Overall decoding performance for the control analysis, using behaviour-based category organisations for A) Animacy and B) Living tasks. Grey bar indicates the period the stimulus was visible on the screen. Lower panels: Time-series correlation between representational distance and RTs for the control analysis, using behaviour-based category organisations for A) Animacy and B) Living tasks. Plotting conventions as in Figure 5.

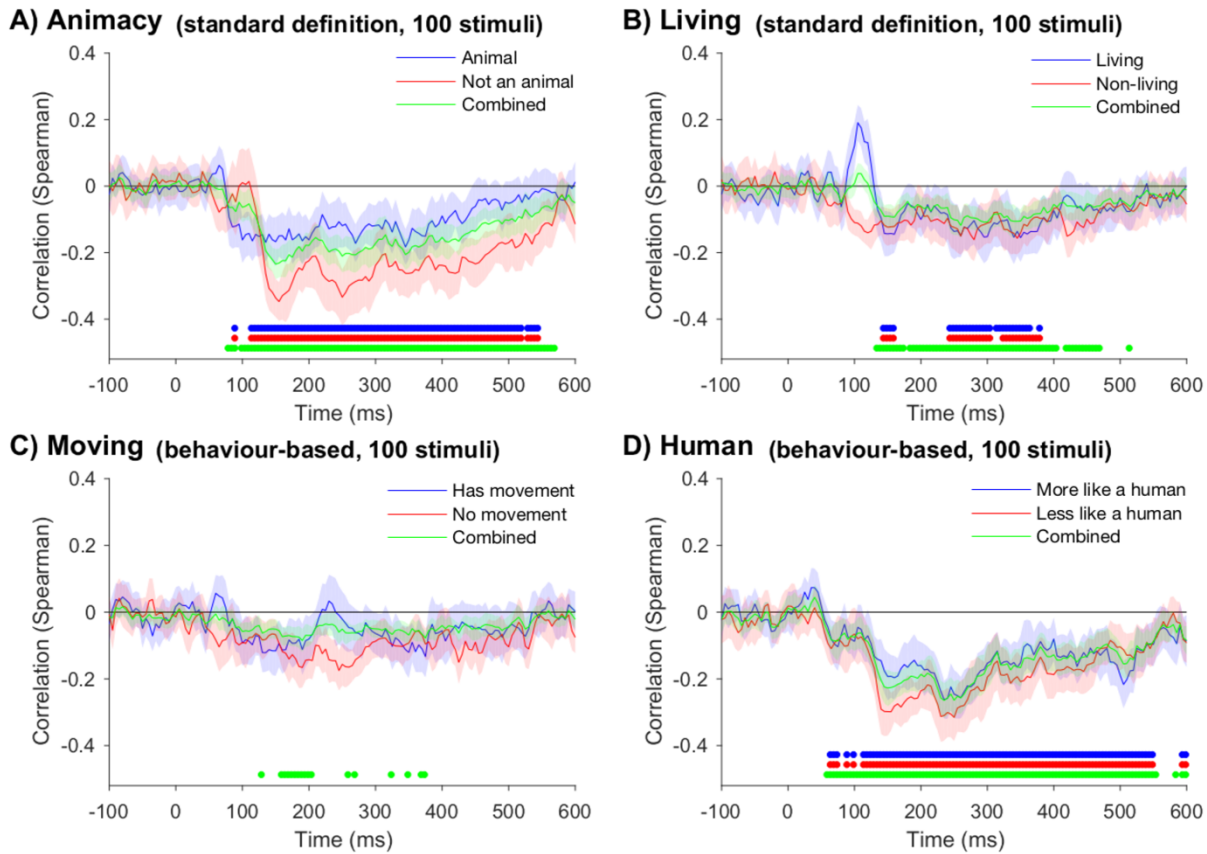
There were slight differences in the number of stimuli included in the analysis for each of the four category organisations as a product of balancing the category subdivisions for the classifier (described in the Methods). Stimulus set numbers for the original analysis ranged from 100 to 118 across the four category organisations, split evenly into two subgroups for training the classifier. To account for any effects of these differences in stimulus numbers on the observed differences between the results for the four category organisations, we conducted a second control analysis where we adjusted the number of stimuli in each set so that all correlations between RTs and representational distances were based on 100 stimuli (50 per category subdivision for the classifier). While every effort was made to include the same stimuli for each of the four category organisations, the different inclusion criteria for each necessitated minor differences in individual exemplars between stimulus sets. Using matched stimulus set numbers, the decoding onsets and peaks remained mostly similar to the main analysis, with the Living category organisation showing the earliest onset at 70 ms (peak at 235 ms), and all three other category organisations showing decoding onsets at 90 ms, though with slightly different peaks: the Moving organisation had a peak at 235 ms; Animacy at 260 ms, and Human at 265 ms. The overall classifier decoding performance was again similar for each of the four category organisations (classifier decoding accuracy range = 58 – 59%; see Figure 7), and remains consistent with the main analysis which included some variation in stimulus set numbers across tasks (cf. Figure 3).





*Figure 7.* Overall MEG decoding performance for the different categorisation tasks, using size-matched stimulus sets. Object exemplars in the upper plots were divided into two groups for the classifier based on the standard category definitions of A) Animacy (animate/inanimate) or B) Living (living/non-living), while for the lower plots, categories for the classifier were determined based on mean behavioural task categorisation of objects from the Moving (has movement/no movement) and Human (more like a human/less like a human) categorisation tasks. Stimulus sets for each of the four category organisations were matched in number, to each contain 100 stimuli (50 per subdivision for the classifier). Grey bar indicates the period the image appeared on the screen. Error bars indicate 95% between-subject confidence intervals. Blue dots along the x-axis indicate time points at which decoding performance was significantly above chance (two-tailed t-test, corrected for multiple comparisons using a false discovery rate (FDR) of  $q < .01$ ).

Similarly, when evaluating the time-series correlations using matched stimulus set numbers (Figure 8), we see the same pattern of results as the original analysis (cf. Figure 4). Focusing on the category subdivision correlations (red and blue lines in Figure 8), significant and sustained correlations are again observed throughout the time course for both Animacy and Human category subdivisions, with a comparatively shorter period where the subdivision correlations are significant for Living, and no significant correlations observed for the Moving organisation.



*Figure 8.* Time-series correlations between representational distance from the classifier boundary and categorisation reaction times for the size-matched stimulus sets. Time-varying rank-order correlation (Spearman's  $\rho$ ) between the mean object representational distance and the median normalised behavioural categorisation RT using matched number sets of stimuli for the four category organisations: A) Animacy, B) Living, C) Moving, and D) Human. Plotting conventions as in Figure 5.

In sum, the results from both sets of control analyses support the original analysis, consistently showing a significant and sustained relationship between representational distance and RT throughout the time-course for the individual subdivisions of the Animacy and Human category organisations. In comparison, for the Living category organisation, this relationship is significant for a much shorter time period, and neither subdivision reaches significance at any time for the Moving category organisation.

#### **4.5. Discussion**

In this study, we aimed to evaluate how categorisation task type influences the relationship between object representational distance and behavioural categorisation RTs; a relationship which has previously been established for categorisation based only on animacy (Carlson et al., 2014; Grootswagers et al., 2017; Ritchie et al., 2015). In addition to categorising objects as either animate or inanimate (Animacy task), we evaluated separate behavioural categorisation RT data according to whether or not an object was living (Living task), whether or not it had movement (Moving task), and whether it was more or less like a human (Human task). This study also employed a new stimulus set specifically designed to include objects that test the limits of the animacy dichotomy. Our results showed that the binary object category membership for all four task-dependent category organisations of the stimulus set could be successfully decoded from the MEG data (see Figure 3). However, representational distances from the classification decision boundary predicted behavioural categorisation RTs for *some*, but not all, of the four categorisation tasks. Specifically, using a novel stimulus set, we found a sustained relationship between representational distance from the category classification boundary and behavioural RTs throughout the time-course for categorisation based on animacy, consistent with previous work (Carlson et al., 2014; Grootswagers et al., 2017; Ritchie et al., 2015). Importantly, we did not find the same clear relationship between representational distance and RTs for the living and moving categorisations. We did, however, find a consistent relationship between representational distance and behavioural RTs based on categorisation according to human-similarity, which

showed a sustained relationship throughout the time-course, at a level similar to that seen with animacy. This is the first time that representational distance has been shown to correlate with behavioural categorisation RTs for any task other than animacy.

The finding that representational distance was related to only certain categorisation behavioural tasks is particularly interesting. The same stimulus set was used for all the RT categorisation tasks and to obtain the neural representation data, with passive viewing of the stimuli in the MEG (i.e., so that the measured brain representations were independent of categorisation behaviour; see Bracci, Daniels, & Op de Beeck, 2017; Harel, Kravitz, & Baker, 2014; Hebart, Bankson, Harel, Baker, & Cichy, 2018). The classification of the neural data was similar for all four category organisations (range 58 – 59% classifier decoding accuracy), yet there was substantial variability in how well the representational distance from the classifier decision boundary correlated with behavioural categorisation RTs. This highlights an important caveat to decoding analyses: the fact that a classifier can learn a given category boundary does not necessarily indicate that this classification is based on information that is important or useful for behaviour (de-Wit et al., 2016; Ritchie & Carlson, 2016; Ritchie et al., 2015).

Given the same stimulus set was used for all four categorisation tasks, any differences in RTs should reflect the different categorisation strategies used by participants in each task. The behavioural categorisation tasks used in the current study were specifically employed such that participants would have to adopt different decision criteria to complete each of the four tasks, requiring them to effectively carve up the same stimulus set in different ways depending on task requirements. When correlating the RTs, significant correlations are observed between exemplar RTs across all four categorisation tasks. However, there appears to be some variability in the strength of the correlations between tasks: while Animacy and Human categorisation RTs are strongly correlated, Animacy RTs are also strongly correlated with both the Living and Moving RTs, whereas the Human RTs are only moderately correlated with the Living and Moving RTs. If all four tasks were

equally correlated with each other, there may be concerns that the different categorisation tasks are ultimately tapping into the same concept. However, perhaps these different relationships between task RTs indicate the presence of unique, task-dependent variance that is being captured in the Animacy and Human RTs. Even so, the question remains: what aspects of these categorisation tasks does the RT-distance relationship actually reflect? We discuss this question further in later paragraphs.

One interesting finding from the current experiment is that there was a comparatively strong relationship between representational distance and behavioural RTs for both animate and inanimate categories individually (and similarly for the individual categories of ‘more like a human’ and ‘less like a human’). This differs from findings in the previous literature which have either shown no clear relationship for inanimate objects, with the overall correlation driven by objects in the animate category (Carlson et al., 2014; Ritchie et al., 2015), or otherwise a comparatively weaker effect for the inanimate objects (Grootswagers et al., 2017). The stimulus set for the current study included categories of objects that were designed to be more difficult to categorise, namely by the inclusion of robots and toys, which represent typically animate objects but would traditionally be categorised as inanimate due to being man-made. Across all four behavioural categorisation tasks, participants appear to be less certain of how to categorise these objects, with comparatively slower median RTs, and with mean behavioural categorisation values closer to the mid-point between the two category options, indicating greater inconsistency in how participants categorised these objects (see Figure 2). As a result, we observe a reasonable degree of variation in RTs and categorisation responses to (traditionally defined) “inanimate” objects using these subcategories. It is possible that the weaker effect for inanimacy observed in previous studies is driven by the use of inanimate stimuli that are more homogenous and easier to categorise than the set used in the current study.

In addition to the relative difficulty observed when participants categorised robots and toys compared to less ambiguous subcategories, evaluating the representational

distances for the individual exemplars reveals that the classifier also appears to have difficulty categorising these objects. Exemplars from both the robot and toy subcategories are positioned closer to the classifier decision boundary, and in many cases, categorised on the incorrect side of the boundary, across all four categorisation groupings (see Figure 4). The classifier appears to have difficulty categorising these objects and, in many cases, is opting to incorrectly classify these items to obtain optimal classification of the data set, grouping robots and toys in the ‘animate’/‘living’/‘has movement’/‘more like a human’ category subdivision despite being trained with these objects belonging to the opposite category. This may indicate that the classifier is using brain responses based on low or mid-level visual features in order to categorise these objects. Similarly, participants may be taking time to override the uncharacteristic visual features of these items in order to categorise the objects in the opposing category, thus explaining the slower categorisation RTs seen across tasks for these objects.

The classifier’s grouping of the robot and toy exemplars with more animate categories is perhaps not surprising given that these object exemplars include features such as faces and bodies, which are visually more similar to features typical of animate objects than inanimate objects. Separating out the contribution of visual properties to the underlying object representations remains an ongoing challenge in the field of object recognition (Bracci & Op de Beeck, 2016; Carlson et al., 2013; Kaiser, Azzalini, et al., 2016; Proklova, Kaiser, & Peelen, 2016), with a recent study showing that *animal appearance*, as opposed to animacy as a semantic concept, dominates object representations in the ventral visual cortex (Bracci, Kalfas, & de Beeck, 2017). Faces and bodies are readily decoded from brain activation patterns (Carlson et al., 2013; Kaneshiro, Guimaraes, Kim, Norcia, & Suppes, 2015; Kriegeskorte et al., 2008; Van de Nieuwenhuijzen et al., 2013) and typically are processed very efficiently by the visual system (Fabre-Thorpe, 2011; Farah, Wilson, Drain, & Tanaka, 1998; Rossion, Joyce, Cottrell, & Tarr, 2003). Specific regions of cortex, notably regions along the right fusiform gyrus, have been shown to exhibit preferential activation in

response to faces and bodies (Downing, Jiang, Shuman, & Kanwisher, 2001; Ishai, Ungerleider, Martin, Schouten, & Haxby, 1999; Kanwisher, McDermott, & Chun, 1997), and these regions overlap with other cortical regions associated with animate object processing (Downing, 2006; Downing, Jiang, Shuman, & Kanwisher, 2001; Haxby et al., 2001; Kanwisher, McDermott, & Chun, 1997; Peelen & Downing, 2017; Proklova et al., 2016). Objects with faces likely create a difference in the overall activation level in face areas compared to when there is no face present, and consequently, it is plausible that this localised difference in activation magnitude may be enough to drive differences in the whole brain MEG signal. The advantage of including difficult categories such as robots and toys in our stimulus set is that participants are unable to adopt a “has face/body” criterion for categorisation, as these features do not reliably predict category membership. The consistently slow RTs and short representational distances for the robot and toy exemplars highlight the difficulty these categories pose for both the visual system and human categorisation behaviour.

In light of our results, what can be said about the RT-distance method as a means of linking brain and behaviour processes? The relationship we observe between representational distance and RTs for object animacy replicates work from previous studies (Carlson et al., 2014; Grootswagers et al., 2017; Ritchie et al., 2015). This effect generalises across stimulus sets, including the current one, which was designed to incorporate objects that would be challenging for both the classifier, and human behavioural categorisation. Additionally, we have the novel finding of showing a similar relationship between representational distance and RTs for human-similarity. This is consistent with the results from our previous study showing the concept of human-similarity to be an important principle in the organisation of object representations (see Chapter 3). It could be argued that these consistent results across studies validate the RT-distance approach as a reliable method of linking brain and behaviour processes. Consequently, as we did not observe an RT-distance relationship for the Living and Moving category organisations, one possibility is

that the underlying large-scale brain representations are naturally organised along dimensions related to both animacy and human-similarity. Or alternatively, given that a model based on human-similarity outperformed the animacy model in our previous experiment (Chapter 3), perhaps what we see in previous animacy studies is actually a reflection of this human-similarity principle, and a product of using similar stimulus sets with object categories that have clear animate or inanimate category membership. Evaluating more varied stimulus sets that are specifically designed to push the boundaries of the categories being tested is central to tackling questions about the neural representations of semantic concepts.

Linking brain processes to behaviour is one of the major goals in neuroscience, and a vital step in validating neuroimaging findings, however identifying adequate and meaningful ways of doing so remains a difficult task. While we acquired the MEG and behavioural data on separate occasions, and from two separate participant populations, this method was also used by Carlson, et al. (2014), who were the first to show a relationship between representational distance and behavioural categorisation RTs for animacy. As the RT-distance relationship is observed despite using separate data sets, this result likely represents something coarse about the relationship between RTs and brain representations overall, rather than a tight link between individual categorisation behaviour and brain activity measured with MEG. Nevertheless, correlating representational distances with RTs does not imply a causal process, as it is not possible to determine from the correlations whether this is a direct relationship or is otherwise mediated by another variable (or multiple variables) that have yet to be identified. It may be the case that an RT-distance relationship for living and moving category organisations is being masked by some other confounding variable that we have not accounted for. Consequently, that we observe the RT-distance relationship with only half of our category organisations could be a consequence of these issues in the method rather than reflecting anything informative about the underlying representational structure.



The RT-distance method used here is only one approach to linking brain and behaviour processes (Ritchie & Carlson, 2016). Other researchers have taken alternative approaches by evaluating how changes in task requirements modulate brain representations. For example, studies have manipulated participants' attention to particular stimulus features during neurophysiological recordings, in which participants make different conceptual (i.e., semantically related) or physical (i.e., stimulus related) decisions about the same stimuli, showing how task-based attention modulates the brain representations measured with fMRI or MEG (Harel et al., 2014; Hebart et al., 2018; Nastase et al., 2016). At a more global level, Kaiser, Oosterhof, and Peelen (2016) evaluated the decodability of attended versus unattended objects as a whole within a scene, linking shifts in covert attention to the strength of the brain representations. These studies (as well as ours) employ linear classification methods to construct the representational spaces, and as such, may miss important information about representations that deviates from a strictly linear organisation (Ritchie & Carlson, 2016). Finding more detailed and sophisticated methods of linking behaviour to brain representations remains a necessary goal for future studies.

Our study is the first to use the RT-distance method for categorisation tasks other than animacy. Here, we report a novel relationship between RTs and representational distance from a classification decision boundary for categorisation based on human-similarity. However, this relationship was not consistently found for all categorisation tasks evaluated, and given the correlational nature of these findings, conclusions regarding causation are difficult. An important task that remains for future studies will be to determine *why* a relationship is observed between representational distance and behavioural RTs for some category organisations but not others. Despite these caveats, identifying a novel human-similarity category organisation highlights the need to evaluate a range of theories and models that extend beyond animacy. Studies should endeavour to evaluate theories and models with both strongly grounded predictions, as well as those for which predictions may be less well defined, to extend research beyond the current state, and ensure we are not

limiting our exploration due to any underlying confirmation biases (Kriegeskorte & Kievit, 2013). While it is far from a complete explanation of the neural mechanisms underlying object categorisation, the link between representational distance and behavioural RTs for a novel category organisation provides a necessary step in the right direction towards linking brain and behavioural processes in order to understand the complexities of object recognition (Ritchie & Carlson, 2016).

Author contributions:

E.C, E.G, S.W designed the study

E.C collected the data

E.C, E.G conducted the analysis

E.C, E.G, S.W interpreted the results

E.C wrote the manuscript

S.W and E.G provided critical revisions on the manuscript

#### 4.6. References

- Bracci, S., Daniels, N., & Op de Beeck, H. (2017). Task Context Overrides Object- and Category-Related Representational Content in the Human Parietal Cortex. *Cerebral Cortex*, 27(1), 310–321. <https://doi.org/10.1093/cercor/bhw419>
- Bracci, S., Kalfas, I., & de Beeck, H. O. (2017). The ventral visual pathway represents animal appearance over animacy, unlike human behavior and deep neural networks. *BioRxiv*, 228932.
- Bracci, S., & Op de Beeck, H. P. (2016). Dissociations and Associations between Shape and Category Representations in the Two Visual Pathways. *The Journal of Neuroscience*, 36(2), 432–444. <https://doi.org/10.1523/JNEUROSCI.2314-15.2016>
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 433–436.
- Carlson, T. A., Goddard, E., Kaplan, D. M., Klein, C., & Ritchie, J. B. (2017). Ghosts in machine learning for cognitive neuroscience: Moving from data to theory. *NeuroImage*. <https://doi.org/10.1016/j.neuroimage.2017.08.019>
- Carlson, T. A., Ritchie, J. B., Kriegeskorte, N., Durvasula, S., & Ma, J. (2014). Reaction Time for Object Categorization Is Predicted by Representational Distance. *Journal of Cognitive Neuroscience*, 26(1), 132–142. [https://doi.org/10.1162/jocn\\_a\\_00476](https://doi.org/10.1162/jocn_a_00476)
- Carlson, Tovar, D. A., Alink, A., & Kriegeskorte, N. (2013). Representational dynamics of object vision: The first 1000 ms. *Journal of Vision*, 13(10), 1–1. <https://doi.org/10.1167/13.10.1>
- Cichy, R. M., Pantazis, D., & Oliva, A. (2014). Resolving human object recognition in space and time. *Nature Neuroscience*, 17(3), 455–462. <https://doi.org/10.1038/nn.3635>
- Clarke, A., Devereux, B. J., Randall, B., & Tyler, L. K. (2015). Predicting the Time Course of Individual Objects with MEG. *Cerebral Cortex*, 25(10), 3602–3612. <https://doi.org/10.1093/cercor/bhu203>

- Clarke, A., Taylor, K. I., & Tyler, L. K. (2010). The Evolution of Meaning: Spatio-temporal Dynamics of Visual Object Recognition. *Journal of Cognitive Neuroscience*, 23(8), 1887–1899. <https://doi.org/10.1162/jocn.2010.21544>
- Connolly, A. C., Guntupalli, J. S., Gors, J., Hanke, M., Halchenko, Y. O., Wu, Y.-C., ... Haxby, J. V. (2012). The Representation of Biological Classes in the Human Brain. *Journal of Neuroscience*, 32(8), 2608–2618. <https://doi.org/10.1523/JNEUROSCI.5547-11.2012>
- Contini, E. W., Wardle, S. G., & Carlson, T. A. (2017). Decoding the time-course of object recognition in the human brain: From visual features to categorical decisions. *Neuropsychologia*, 105(Supplement C), 165–176. <https://doi.org/10.1016/j.neuropsychologia.2017.02.013>
- de-Wit, L., Alexander, D., Ekroll, V., & Wagemans, J. (2016). Is neuroimaging measuring information in the brain? *Psychonomic Bulletin & Review*, 23(5), 1415–1428. <https://doi.org/10.3758/s13423-016-1002-0>
- Downing, P. (2006). Domain Specificity in Visual Cortex. *Cerebral Cortex*, 16(10), 1453–1461. <https://doi.org/10.1093/cercor/bhj086>
- Downing, P., Jiang, Y., Shuman, M., & Kanwisher, N. (2001). A cortical area selective for visual processing of the human body. *Science*, 293(5539), 2470–2473.
- Fabre-Thorpe, M. (2011). The characteristics and limits of rapid visual categorization. *Frontiers in Psychology*, 2, 243. <https://doi.org/10.3389/fpsyg.2011.00243>
- Farah, M. J., Wilson, K. D., Drain, M., & Tanaka, J. N. (1998). What is "special" about face perception? *Psychological Review*, 105(3), 482.
- Filliter, J. H., McMullen, P. A., & Westwood, D. (2005). Manipulability and living/non-living category effects on object identification. *Brain and Cognition*, 57(1), 61–65. <https://doi.org/10.1016/j.bandc.2004.08.022>

- Gainotti, G. (2000). What the Locus of Brain Lesion Tells us About the Nature of the Cognitive Defect Underlying Category-Specific Disorders: A Review. *Cortex*, 36(4), 539–559. [https://doi.org/10.1016/S0010-9452\(08\)70537-9](https://doi.org/10.1016/S0010-9452(08)70537-9)
- Goddard, E., Carlson, T. A., Dermody, N., & Woolgar, A. (2016). Representational dynamics of object recognition: Feedforward and feedback information flows. *NeuroImage*, 128, 385–397. <https://doi.org/10.1016/j.neuroimage.2016.01.006>
- Grootswagers, T., Ritchie, J. B., Wardle, S. G., Heathcote, A., & Carlson, T. A. (2017). Asymmetric Compression of Representational Space for Object Animacy Categorization under Degraded Viewing Conditions. *Journal of Cognitive Neuroscience*, 1–16. [https://doi.org/10.1162/jocn\\_a\\_01177](https://doi.org/10.1162/jocn_a_01177)
- Harel, A., Kravitz, D. J., & Baker, C. I. (2014). Task context impacts visual object processing differentially across the cortex. *Proceedings of the National Academy of Sciences*, 111(10), E962–E971. <https://doi.org/10.1073/pnas.1312567111>
- Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and Overlapping Representations of Faces and Objects in Ventral Temporal Cortex. *Science*, 293(5539), 2425–2430. <https://doi.org/10.1126/science.1063736>
- Hebart, M. N., Bankson, B. B., Harel, A., Baker, C. I., & Cichy, R. M. (2018). The representational dynamics of task and object processing in humans. *ELife*, 7. <https://doi.org/10.7554/eLife.32816>
- Huth, A. G., Nishimoto, S., Vu, A. T., & Gallant, J. L. (2012). A Continuous Semantic Space Describes the Representation of Thousands of Object and Action Categories across the Human Brain. *Neuron*, 76(6), 1210–1224. <https://doi.org/10.1016/j.neuron.2012.10.014>
- Julian, J. B., Ryan, J., & Epstein, R. A. (2017). Coding of Object Size and Object Category in Human Visual Cortex. *Cerebral Cortex*, 27(6), 3095–3109. <https://doi.org/10.1093/cercor/bhw150>

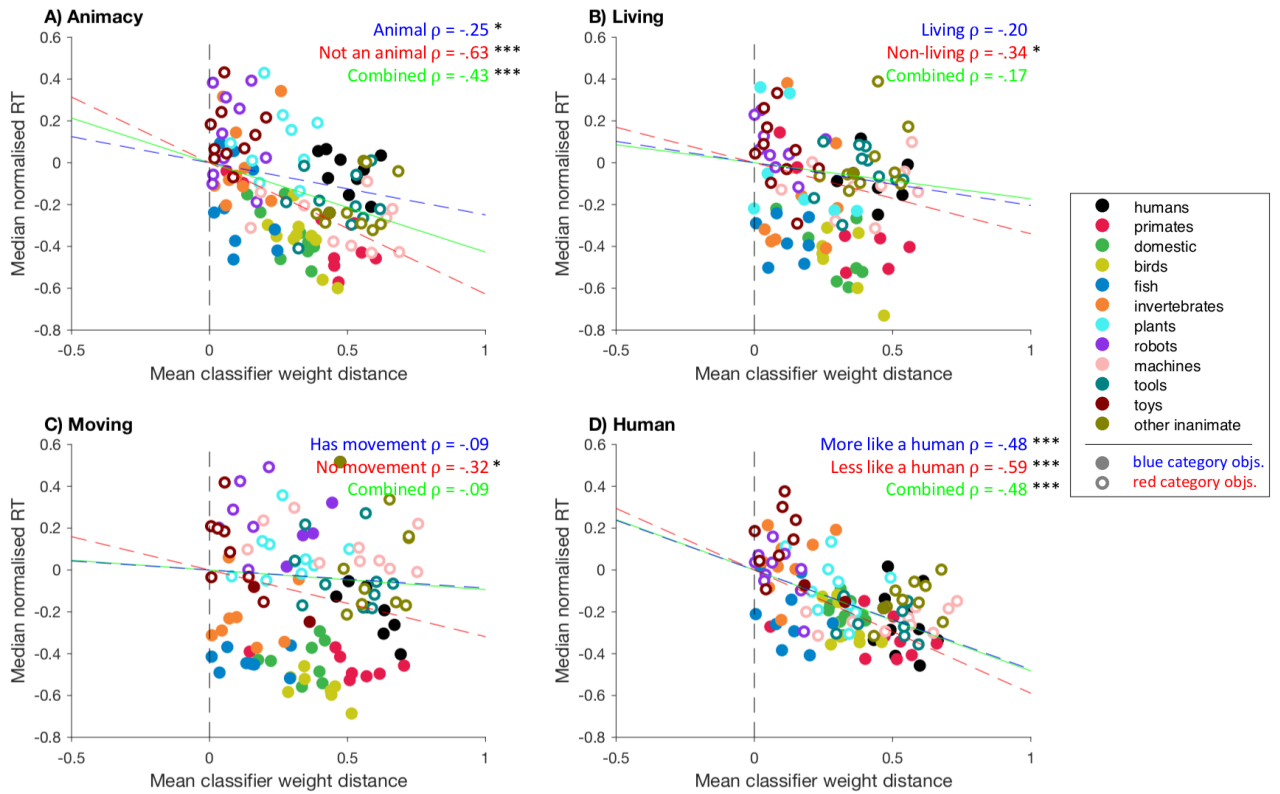
- Kaiser, D., Azzalini, D. C., & Peelen, M. V. (2016). Shape-independent object category responses revealed by MEG and fMRI decoding. *Journal of Neurophysiology*, *115*(4), 2246–2250. <https://doi.org/10.1152/jn.01074.2015>
- Kaiser, D., Oosterhof, N. N., & Peelen, M. V. (2016). The Neural Dynamics of Attentional Selection in Natural Scenes. *The Journal of Neuroscience*, *36*(41), 10522–10528. <https://doi.org/10.1523/JNEUROSCI.1385-16.2016>
- Kaneshiro, B., Guimaraes, M. P., Kim, H.-S., Norcia, A. M., & Suppes, P. (2015). A Representational Similarity Analysis of the Dynamics of Object Processing Using Single-Trial EEG Classification. *PLOS ONE*, *10*(8), e0135697. <https://doi.org/10.1371/journal.pone.0135697>
- Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The Fusiform Face Area: A Module in Human Extrastriate Cortex Specialized for Face Perception. *Journal of Neuroscience*, *17*(11), 4302–4311.
- Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., Broussard, C., & others. (2007). What's new in Psychtoolbox-3. *Perception*, *36*(14), 1.
- Konkle, T., & Caramazza, A. (2013). Tripartite Organization of the Ventral Stream by Animacy and Object Size. *Journal of Neuroscience*, *33*(25), 10235–10242. <https://doi.org/10.1523/JNEUROSCI.0983-13.2013>
- Konkle, T., & Oliva, A. (2012). A Real-World Size Organization of Object Responses in Occipitotemporal Cortex. *Neuron*, *74*(6), 1114–1124. <https://doi.org/10.1016/j.neuron.2012.04.036>
- Krakauer, J. W., Ghazanfar, A. A., Gomez-Marin, A., MacIver, M. A., & Poeppel, D. (2017). Neuroscience Needs Behavior: Correcting a Reductionist Bias. *Neuron*, *93*(3), 480–490. <https://doi.org/10.1016/j.neuron.2016.12.041>
- Kriegeskorte, N., & Kievit, R. A. (2013). Representational geometry: integrating cognition, computation, and the brain. *Trends in Cognitive Sciences*, *17*(8), 401–412. <https://doi.org/10.1016/j.tics.2013.06.007>

- Kriegeskorte, N., Mur, M., Ruff, D. A., Kiani, R., Bodurka, J., Esteky, H., ... Bandettini, P. A. (2008). Matching Categorical Object Representations in Inferior Temporal Cortex of Man and Monkey. *Neuron*, 60(6), 1126–1141.  
<https://doi.org/10.1016/j.neuron.2008.10.043>
- Masullo, C., Piccininni, C., Quaranta, D., Vita, M. G., Gaudino, S., & Gainotti, G. (2012). Selective impairment of living things and musical instruments on a verbal ‘Semantic Knowledge Questionnaire’ in a case of apperceptive visual agnosia. *Brain and Cognition*, 80(1), 155–159. <https://doi.org/10.1016/j.bandc.2012.06.002>
- Nastase, S. A., Connolly, A. C., Oosterhof, N. N., Halchenko, Y. O., Guntupalli, J. S., Castello, M. V. di O., ... Haxby, J. V. (2016). Attention selectively reshapes the geometry of distributed semantic representation. *BioRxiv*, 045252.  
<https://doi.org/10.1101/045252>
- Peelen, M. V., & Downing, P. E. (2017). Category selectivity in human visual cortex: Beyond visual object recognition. *Neuropsychologia*. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0028393217301215>
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10(4), 437–442.
- Philiastides, M. G., & Sajda, P. (2006). Temporal characterization of the neural correlates of perceptual decision making in the human brain. *Cerebral Cortex*, 16(4), 509–518.
- Proklova, D., Kaiser, D., & Peelen, M. V. (2016). Disentangling Representations of Object Shape and Object Category in Human Visual Cortex: The Animate–Inanimate Distinction. *Journal of Cognitive Neuroscience*, 1–13.  
[https://doi.org/10.1162/jocn\\_a\\_00924](https://doi.org/10.1162/jocn_a_00924)
- Ritchie, J. B., & Carlson, T. A. (2016). Neural decoding and “inner” psychophysics: a distance-to-bound approach for linking mind, brain, and behavior. *Frontiers in Neuroscience*, 10.

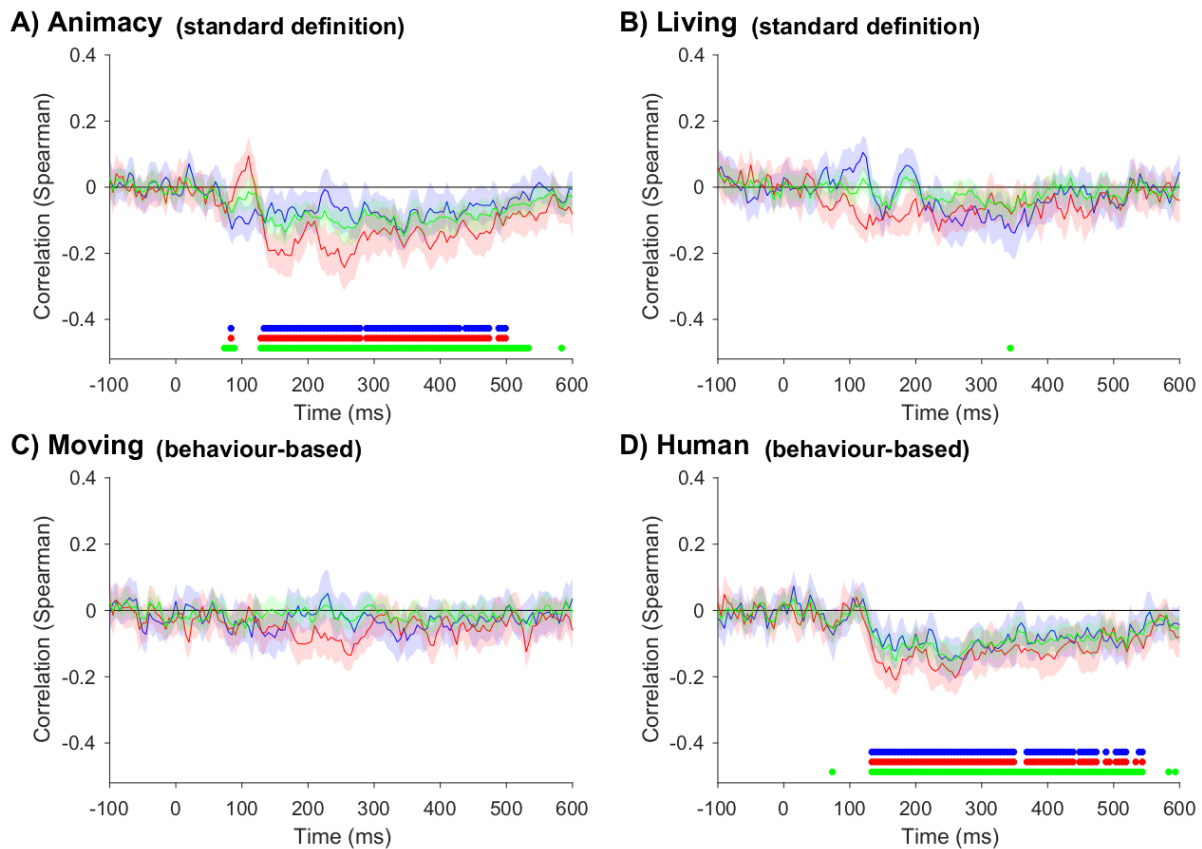
- Ritchie, J. B., Tovar, D. A., & Carlson, T. A. (2015). Emerging Object Representations in the Visual System Predict Reaction Times for Categorization. *PLoS Comput Biol*, 11(6), e1004316. <https://doi.org/10.1371/journal.pcbi.1004316>
- Rossion, B., Joyce, C. A., Cottrell, G. W., & Tarr, M. J. (2003). Early lateralization and orientation tuning for face, word, and object processing in the visual cortex. *NeuroImage*, 20(3), 1609–1624. <https://doi.org/10.1016/j.neuroimage.2003.07.010>
- Sha, L., Haxby, J. V., Abdi, H., Guntupalli, J. S., Oosterhof, N. N., Halchenko, Y. O., & Connolly, A. C. (2015). The Animacy Continuum in the Human Ventral Vision Pathway. *Journal of Cognitive Neuroscience*, 27(4), 665–678. [https://doi.org/10.1162/jocn\\_a\\_00733](https://doi.org/10.1162/jocn_a_00733)
- Simanova, I., Gerven, M. van, Oostenveld, R., & Hagoort, P. (2010). Identifying Object Categories from Event-Related EEG: Toward Decoding of Conceptual Representations. *PLOS ONE*, 5(12), e14465. <https://doi.org/10.1371/journal.pone.0014465>
- Van de Nieuwenhuijzen, M. E., Backus, A. R., Bahramisharif, A., Doeller, C. F., Jensen, O., & van Gerven, M. A. (2013). MEG-based decoding of the spatiotemporal dynamics of visual category perception. *NeuroImage*, 83, 1063–1073.
- Vida, M. D., Nestor, A., Plaut, D. C., & Behrmann, M. (2017). Spatiotemporal dynamics of similarity-based neural representations of facial identity. *Proceedings of the National Academy of Sciences*, 114(2), 388–393.
- Williams, M. A., Dang, S., & Kanwisher, N. G. (2007). Only some spatial patterns of fMRI response are read out in task performance. *Nat Neurosci*, 10(6), 685–686.



## 4.7. Appendix



*Supplementary Figure 1.* Correlations between reaction times and classifier weight distances using the absolute value of the distance from the classifier decision boundary. Plots are based on the mean results over a 50 ms period centred around the MEG peak decoding time for each of the four tasks as follows: A) Animacy peak at 270 ms; B) Living peak at 230 ms; C) Moving peak at 255 ms; D) Human peak at 265 ms. Individual objects are shown as either filled dots or hollow dots to indicate category organisation grouping for the classifier, while colours indicate subcategory membership (see key for subcategory descriptions). The dashed line at zero on the x-axis shows the decision boundary for the classifier. Objects categorised correctly by the classifier have a positive x-value, while those incorrectly categorised (according to individual category organisations) have a negative x-value. Coloured lines show Spearman's  $\rho$  for all objects combined (green line), as well as correlations for each category subdivision (red and blue lines) ( $*p < .05$ ;  $**p < .01$   $***p < .001$ ).



*Supplementary Figure 2.* Time-series correlations between categorisation reaction times and representational distance from the classifier boundary when using the absolute value for the distances. Time-varying rank-order correlation (Spearman's  $\rho$ ) between the mean object absolute value representational distance and the median normalised behavioural categorisation RT for the four category organisations: A) Animacy, B) Living, C) Moving, and D) Human. In each plot, overall correlations are plotted for all stimuli combined, as well as separately for each category subdivision (see figure legends). Shaded error bars indicate the 95% confidence interval of the between-subject means. Coloured dots along the x-axis indicate time points at which the correlation is significantly above chance (two-tailed t-test, corrected for multiple comparisons using a false discovery rate (FDR) of  $q < .01$ ).

## CHAPTER FIVE

---

---

### GENERAL DISCUSSION

---



## 5. General Discussion

This thesis presented three MEG decoding studies which investigated the temporal dynamics of object category representations. In the following sections, I firstly provide an overview of the findings from each of my empirical chapters (Section 5.1), before discussing the overarching themes. In Section 5.2, I discuss the utility of temporal decoding studies in advancing our understanding of object recognition processes. Section 5.3 discusses the veracity of object category representations, and highlights some of the methodological issues that make studying these concepts difficult. The importance of linking brain data to behavioural processes is discussed in Section 5.4, followed by the overall concluding statement for this thesis.

### 5.1. Results summary

In **Chapter 2**, using a stimulus set which controlled for object interactivity, I was unable to detect variations in brain representations related to the real-world size of an observed object. This finding was surprising given the growing number of studies which have identified a size-based dimension to object representations (Julian, Ryan, & Epstein, 2017; Khaligh-Razavi, Cichy, Pantazis, & Oliva, 2018; Konkle & Caramazza, 2013; Konkle & Oliva, 2012). I did, however, replicate previous findings of an organisation according to animacy, with these representations emerging relatively late, and shown to be sustained throughout the time course of object processing. These findings emphasise an ongoing challenge in the interpretation of object representations from neurophysiological data: the multifaceted and interrelated nature of object properties, such as real-world size and landmark stability, as well as low-level image properties, make measurement of the unique contributions of these aspects difficult.

Given that the results observed in Chapter 2 remained consistent with a representational organisation based on animacy, together with recent findings of a biological classes continuum within the animate domain (Connolly et al., 2012; Sha et al., 2015), I shifted my focus in **Chapter 3** to test possible alternative explanations for the observed

animacy result. To do so I firstly created a stimulus set that was designed to test the limits of the animacy dichotomy, by including objects that visually resembled animate objects, yet are inanimate (e.g. robots, toys). Secondly, I tested a range of plausible conceptual and behavioural models of object category representations in addition to animacy. These included a model based on the biological classes continuum (Connolly et al., 2012; Sha et al., 2015), a faces/bodies model (which grouped objects irrespective of whether they were animate or inanimate, resulting in the robot and toy stimuli being grouped with all the animate objects), and a range of behaviourally generated models based on categorisation along a range of human-related agency/experience dimensions (e.g., ability to experience fear, desire, conscious thought). In addition, I tested a behaviour-based model measuring the broad concept of human-similarity. The latter model provided the best account of object decoding late in the time-course of object processing ( $>155$  ms after stimulus onset), while the low-level models dominated earlier on (50 – 150 ms after stimulus onset). Interestingly, despite being outperformed by the human-similarity model, many of the agency/experience models exceeded the performance of models evaluated in previous studies, such as animacy, living/non-living, and biological classes continuum models. These results highlight the multifaceted nature of object representations, which likely encompass concepts related to human-similarity, including agency and human experience dimensions.

In **Chapter 4** I linked this result obtained from the MEG data to behavioural categorisation reaction times, to investigate whether the information decoded from the neurophysiological data is meaningful for categorisation behaviour. I showed that categorisation based on human-similarity parallels the animacy categorisation performance: in both cases, there was a sustained relationship between categorisation RTs and representational distance from a classification decision boundary. This is the first time that a categorical organisation other than animacy has shown this link between brain and behaviour using the RT-distance method. Interestingly, this was the only other organisation which showed an RT-distance relationship that rivalled object animacy performance, despite

adequate decoding of other object category organisations from the brain data. The results from this study highlight three important points: firstly, that categorisation tasks other than animacy also predict representational distances in MEG, secondly, that a human-similarity based dimension provides a comparable account of time-varying object representations and warrants further investigation, and finally, that further research is required to determine why we see a relationship between brain and behaviour for some categorisation tasks and not others.

In sum, while I was unable to detect real-world size representations in the time-varying brain representations using MEG, we now have evidence for additional categorical factors present in the organisation of neural object representations. Most prominently, object representations appear to also be characterised by a dimension related to human-similarity, and this dimension also relates to categorisation behaviour. I discuss the implications of these findings below.

## **5.2. Insights into the temporal dynamics of object representations**

Our ability to visually recognise an object is so rapid and automatic (Fabre-Thorpe, 2011; Mack & Palmeri, 2011, 2015; Thorpe, Fize, & Marlot, 1996; VanRullen & Thorpe, 2001) that differentiating the underlying neural processes is not an easy feat. At any moment, we can efficiently integrate multiple facets of an object's properties, including information such as its orientation, colour, contrast, and shape into a global structure to which we can assign meaning, and associate with other object concepts stored in our memories (Kiefer & Pulvermüller, 2012; Mahon & Caramazza, 2009; Martin, 2007, 2016; Nosofsky, 1986). This occurs within as little as 150 ms after first seeing an object image (Thorpe et al., 1996). When one considers the overwhelming number of possible objects that humans manage to differentiate, and the ease with which this transpires, it is remarkable that this is achieved in such a short period of time. Given the clearly complex nature of object recognition, evaluating the temporal intricacies of this process provides valuable insights into what information is important, and how this is utilised during the short period of time required to

achieve this complex task. These temporal details are necessary to develop a full understanding of how the brain processes visual object properties.

Upon seeing an object, cortical activity systematically varies from stimulus onset through to the point that a conscious behavioural decision is made. Object processing begins with an initial feedforward sweep of activity, predominantly related to low-level feature processing, with top-down feedback related to higher-level conceptual information of an object occurring later in time (for reviews, see Grill-Spector & Malach, 2004; Grill-Spector & Weiner, 2014; Logothetis & Sheinberg, 1996). Differences in minimum latencies required for rapid object detection (a mostly feature-based process) versus the comparatively slower recognition (involving higher conceptual analysis) highlight the temporal disparity between feedforward and feedback processes required to make these decisions (Bar, 2003; Goddard, Carlson, Dermody, & Woolgar, 2016). This important variability in the object processing signal is blurred in more spatially sensitive neuroimaging modalities such as fMRI; changes in the neural activity are measured indirectly via the temporally sluggish BOLD response ( $\sim 5$  s; see Logothetis, 2002; Logothetis, Pauls, Augath, Trinath, & Oeltermann, 2001), at a temporal resolution that is usually  $>1$  Hz, compared to the millisecond resolution of MEG, where the measured perturbations of magnetic fields are directly related to neural activity. Consequently, spatially sensitive modalities such as fMRI are more appropriate for addressing questions about *where* activity is occurring rather than the processes that led to its occurrence. Trying to understand the processes involved in object recognition by focusing only on localised activity in the brain, is akin to gazing into a bakery store window to figure out how the cakes were made. While localisation of activity has a firm place in the neuroscience of object recognition, research in the temporal domain allows us to further unpack the evolving processes required to achieve this complex task, allowing us to ask new and exciting questions that extend our understanding beyond simply what is represented where.



For my analyses, I took advantage of recent advancements in the analysis of time-series data by employing multivariate decoding methods, including representational similarity analysis (RSA; Kriegeskorte, Mur, & Bandettini, 2008), and multidimensional scaling. These methods have aided a new and exciting direction of research within the field of object recognition (e.g., Carlson, Tovar, Alink, & Kriegeskorte, 2013; Cichy, Pantazis, & Oliva, 2014; Grootswagers, Wardle, & Carlson, 2016; Hebart, Bankson, Harel, Baker, & Cichy, 2018; Isik, Meyers, Leibo, & Poggio, 2014; Kaiser, Azzalini, & Peelen, 2016; Kaiser, Oosterhof, & Peelen, 2016; Kriegeskorte et al., 2008). By using these methods across multiple experiments in this thesis, I demonstrated how a range of computational and conceptual models of object representations show temporal variability in the extent to which they account for object decoding performance. In Chapter 3, the computational low-level visual models provide the best account of the data early in the time-course of object processing, while the conceptual models have peaks comparatively later. These findings are consistent with those from previous studies showing processing of low-level feature information progressing to more concept based, categorical distinctions in the evolving object representations (Carlson, Tovar, Alink, & Kriegeskorte, 2013; Cichy, Pantazis, & Oliva, 2014; Clarke, Devereux, Randall, & Tyler, 2015; Hebart, Bankson, Harel, Baker, & Cichy, 2018; for a review, see Contini, Wardle, & Carlson, 2017).

While this important distinction between early and late representations is evident in the temporal dynamics observed from the MEG data, we are still limited as to how we can interpret relative differences in the observed timings of the different models. For example, in Chapter 3, most of the conceptual models of object representations follow a similar time-course, including similar peak decoding times (see Chapter 3, Figure 6). Does this similarity indicate some common underlying representation that is being captured despite the conceptual differences that the models were designed to assess? Given the millisecond temporal resolution of MEG, why don't we see greater variation between these models? For my experiments, I utilised whole-brain MEG data to evaluate object representations. While

activity related to object processing extends throughout the brain (Grill-Spector & Malach, 2004; Haxby et al., 2001), comparatively stronger activation of particular regions, such as the fusiform face area, may dominate this whole-brain signal and thus hamper our ability to obtain more fine-grained distinctions of object representations. Furthermore, as the MEG recordings were limited to 160 MEG sensors, we are already averaging the signal over multiple brain regions at each sensor. Studies that link MEG with fMRI data are beginning to tackle this problem, allowing for better temporal resolution of activity in particular regions of interest (Cichy, Kriegeskorte, Jozwik, Bosch, & Charest, 2017; Cichy et al., 2014; Hebart et al., 2018). Indeed, results from a recent study that linked fMRI and MEG neuroimaging data suggested that neither category membership nor low-level visual feature information could sufficiently explain the relationship between the object stimuli used, the associated neural representations, and categorisation behaviour (Cichy et al., 2017). These results emphasise the complex interplay between the different processes involved in object recognition, and the importance of behavioural relevance in shaping these multifaceted representations.

There is also the peculiar finding of a sustained relationship with MEG exemplar decoding performance long after the image is removed from view. This pattern is observed across many of the concept-based models tested in Chapter 3, as well as the animacy and tripartite models in Chapter 2. This is even present for models which show comparatively weaker performance, such as the animacy model in Chapter 3: despite having one of the lowest average correlations across the late time-window (see Chapter 3, Figure 8B) it remains significant until just prior to the end of the 600 ms trial epoch. Similarly, the RTs in Chapter 4 correlated with classifier decoding performance for both animacy and human-similarity categorisation throughout most of the time-course. This sustained relationship is not unique to my experiments, appearing in several previous MEG decoding studies (Carlson, Ritchie, Kriegeskorte, Durvasula, & Ma, 2014; Carlson et al., 2013; Cichy et al., 2014; Grootswagers, Ritchie, Wardle, Heathcote, & Carlson, 2017; Kaiser, Azzalini, et al.,

2016). Given that we know object representations are dynamic, with little generalisation of activity patterns across time (Carlson et al., 2013; Cichy et al., 2014; Grootswagers et al., 2016; Isik et al., 2014), better methods are needed to determine what information remains in this dynamic signal at late latencies and how this relates to behaviour.

### **5.3. Is a Category Truly a Category?**

**5.3.1. Human-similarity versus animacy.** A key goal in visual neuroscience is to identify the semantic and conceptual features that characterise how objects are represented throughout the brain. The animacy dichotomy is arguably the most widely replicated large-scale organisation that has been identified to date (e.g., Caramazza & Shelton, 1998; Chao, Haxby, & Martin, 1999; Cichy et al., 2014; Kiani, Esteky, Mirpour, & Tanaka, 2007; Kriegeskorte, Mur, Ruff, et al., 2008; Lu, Li, & Meng, 2016; Mahon & Caramazza, 2009; Proklova, Kaiser, & Peelen, 2016). More recently, researchers have identified organisational factors related to real-world size (Julian, Ryan, & Epstein, 2017; Khaligh-Razavi, Cichy, Pantazis, & Oliva, 2018; Konkle & Caramazza, 2013; Konkle & Oliva, 2012) and a biological classes animacy continuum (Connolly et al., 2012; Sha et al., 2015). In light of these recent advancements in our understanding of object representations, one of the primary aims of this thesis was to evaluate the temporal dynamics of alternative categorical organisation principles to that of animacy, and to examine how these fit with our current understanding of object processing. The results from both Chapters 3 and 4 provide consistent evidence for a dimension related to human-similarity in object representations which emerges relatively late in the time-course of object processing. As discussed in Chapter 3, this result is consistent with emerging ideas about categorical representations existing along a continuum rather than a strict dichotomy (Connolly et al., 2012; Sha et al., 2015), and highlights the potential contributions of more functional, goal-related concepts

linked to behaviour in the formation of these representations (for discussion, see Martin, 2016; Peelen & Downing, 2017).

While the results from Chapter 4, using the RT-distance approach, show a similar relationship between brain and behaviour processes for both animacy and human-similarity, there is greater discrepancy between the efficacy of models based on these same concepts in the RSA analysis of Chapter 3. Here, I observed considerably better performance of the human-similarity model compared to an animacy model, with animacy one of the lowest performing models evaluated (see Chapter 3, Figures 6 and 8). It is important to note here that the human-similarity model in Chapter 3 was generated based on human behavioural ratings, which resulted in considerably more variability in how category structure was specified in the model compared to the animacy model for that experiment, which employed a standard binary categorical definition of animate and inanimate objects. When evaluating the time-series results for these models, there was a tendency for models with greater variability to better account for the neural representations in the late time-window (see Chapter 3, Figure 8). Hence, one might argue that evaluating the performance of a graded model versus a binary model is not a fair comparison when accounting for highly complex data. Consequently, the use of more variable RT data (in Chapter 4) may explain the similar correlations for both the Animacy and Human category organisations seen in my RT-distance results.

However, the difference in model construction for the RSA model testing in Chapter 3 cannot entirely account for the pattern of results across these two studies. If it were merely a matter of model gradation, then in Chapter 4 we might expect the Living category organisation to show a similar relationship between representational distance and RTs to that seen with Animacy, since binary category models based on these organisations showed comparable performance in Chapter 3. This, however, was not the case in Chapter 4, where the Animacy category organisation clearly outperformed the Living organisation in the RT-distance analysis. Given that model variability alone does not predict its performance, and

that the human-similarity factor provided the best account of the data across the two studies (Chapters 3 and 4), it seems reasonable to conclude that a human-similarity dimension accounts for some unique aspects of object representations and warrants further investigation. This is consistent with the work of Connelly et al. (2012) and Sha et al. (2015), who highlighted the idea of agency and human-related factors as having importance for how objects are represented in the brain.

The identification of category-specific deficits from studies of patients with brain lesions has inspired much of the research into object category representations to date (Forde, 1999; Kiefer & Pulvermüller, 2012; Mahon & Caramazza, 2009; Warrington & Shallice, 1984). A review of 79 patient case studies highlights that animate objects appear to be overrepresented in category-specific deficits, with fewer cases showing a selective inanimate deficit by comparison (Capitani, Laiacona, Mahon, & Caramazza, 2003). In the context of this, studying semantic categories centred on animate object representations seems justified. However, results from a recent patient case study by Panis, Torfs, Gillebert, Wagemans, and Humphreys (2017) suggest that semantics play a weaker role in explaining the deficits in object processing. The authors instead propose that problems with the cycling of feedforward and feedback information between semantic and feature-based processes may be at play. The documentation of object representations using healthy adults is essential to providing a full account of how objects are represented in the brain, however these results are often not related back to deficits observed in patients. Relating findings from neuroimaging studies using healthy adults to those seen in patient populations is essential for advancing theoretical and computational models of object recognition.

**5.3.2. The role of experimental design in investigating object category representations.** Evaluating the structure of object category representations is inherently difficult, with the reality that stimulus selection ultimately shapes what we are able to discover with neuroimaging tools (Carlson, Goddard, Kaplan, Klein, & Ritchie, 2017). As discussed in Chapter 2, categorical organisations that are proposed by researchers may

reflect categorical dimensions that are *correlated* with the concepts under investigation, rather than the concepts themselves. For example, in Chapter 2 I discussed the potential overlap between the object properties of landmark stability and the *big* real-world size dimension, which has been raised in multiple studies (see Julian et al., 2017; Mullally & Maguire, 2011; Troiani, Stigliani, Smith, & Epstein, 2014). Therefore, despite large objects being readily dissociable from small objects in certain temporal or spatial brain activity patterns, this dissociation may not be purely related to object size per se, but some other correlated dimension. Further discussion of what we can conclude from neuroimaging results continues below in Section 5.4.

Separating out the contribution of low-level stimulus properties from higher-level conceptual information is crucial to developing our understanding of object categorisation processes, and remains a challenge in this field. A number of results suggest that low-level visual properties on their own cannot completely account for representations related to animacy throughout the time-course of object processing (Bankson, Hebart, Groen, & Baker, 2017; Carlson, Hogendoorn, Kanai, Mesik, & Turret, 2011; Carlson et al., 2013; Cichy et al., 2014; Clarke & Tyler, 2014; Isik et al., 2014; Kaiser, Azzalini, et al., 2016; Kriegeskorte, Mur, Ruff, et al., 2008). In Chapter 3, I compare multiple higher-level conceptual models of object categorisation to low-level visual models, highlighting that while visual features dominate early in the time-course, their relationship with the decoding data later on is significantly weaker than most of the conceptual models tested (see Chapter 3, Figure 8), in particular that of human-similarity.

Despite these findings, there are examples in which visual properties appear to contribute to object category representations. Schmidt, Hegele, and Fleming (2017) recently showed that mid-level perceptual shape features of objects (e.g., symmetry, repetitive features, curvature) can be used in visual processing to determine where an object sits along an animacy-based continuum, with these features accurately separating the categories of animals, plants, and minerals. While I did not specifically measure mid-level properties for

each of the stimuli used in my second RSA and model testing experiment (Chapter 3), I did include multiple screen locations across trials and a left-right flipped version of each stimulus to minimise the influence of low-level features on classifier performance. Additionally, I evaluated two visual feature models (HMAX and Jaccard silhouette), with neither model providing a good account of the MEG data late in the decoding time-course, where the other conceptual models performed well. Therefore, I interpret these findings to be indicative of a dimension, related to or correlated with human-similarity, that cannot be fully explained by low-level features.

In Chapter 4, I raised the point that inanimate objects in my stimulus set which had faces (e.g., robots and toys) were challenging for the classifier, with the classifier opting to incorrectly classify many of these stimuli as animate objects in order to optimise the positioning of the decision boundary. A similar result has recently been shown in an fMRI study by Bracci, Kalfas, and Op de Beeck (2017), who compared human behavioural categorisation performance to that of deep neural networks, and activity in the ventral occipitotemporal cortex (VTC). Results showed that while DNNs categorised objects similarly to human behavioural performance, activity in VTC differed, representing object appearance rather than animacy by depicting inanimate objects that *looked* like animate ones (e.g., a cow-shaped mug), as being closer to animate objects. My results, combined with those of Bracci et al. (2017), draw attention to how our ability to uncover the intricacies of object category representations is heavily dependent on the kinds of stimuli we employ, and emphasise the importance of designing stimulus sets in ways that challenge theories of categorisation, to test the limits and generalisability of results.

Even with carefully designed stimulus sets, the limitation remains that results for any given study are, to some extent, specific to the stimulus set employed. Many neuroimaging studies that have evaluated aspects of object category representations have used the same stimulus set (or a subset thereof) from Kriegeskorte, Mur, Ruff, et al. (2008) (e.g., Carlson et al., 2013; Cichy et al., 2014; Goddard et al., 2016; Mur et al., 2013). A problem inherent

in decoding analyses applied to neuroimaging data is that it is not possible to know exactly what aspects of the neuroimaging data the classifier is employing to make its decision (Carlson & Wardle, 2015; de-Wit, Alexander, Ekroll, & Wagemans, 2016; Naselaris & Kay, 2015). Additionally, the inclusion or exclusion of certain object categories could plausibly result in a marked shift in boundary location. Consequently, even though the stimulus set employed in Chapters 3 and 4 was designed to test the limits of object animacy as well as other plausible category organisations, the results are still dependent on the exemplars and subcategories included in this stimulus set. Future studies should endeavour to replicate the present findings of the dimension of human-similarity using new stimulus sets in order to test for the generalisability of these results.

A similar generalisation issue applies to how models of categorisation are created to compare neuroimaging data to behavioural judgements using RSA. For example, in the RSA experiment in Chapter 3, I had participants categorise each of the objects individually, along a 7-point scale which represented a continuum between two category distinctions (e.g., more like a human versus less like a human). Other studies have focused on similarity judgements between items within a stimulus set, employing methods such as the inverse MDS approach (Kriegeskorte & Mur, 2012) as used in a study by Mur et al. (2013), where participants arranged all objects from a stimulus set in a prescribed two-dimensional space based on how similar participants felt the objects were. Other studies have employed pair-wise dissimilarity judgements (e.g., Connolly et al., 2012; Cooke, Jäkel, Wallraven, & Bühlhoff, 2007; Cortese & Dyre, 1996), or an odd-one-out procedure (e.g., Connolly et al., 2012; Sha et al., 2015) to create models of object similarity for testing using the RSA framework. Future studies should compare how these different approaches to obtaining models of behavioural categorisation differentially account for object representations throughout the time-course. This is particularly important to ascertain whether we are measuring behaviour in the most appropriate way for the investigation of brain representations.



#### **5.4. Addressing the Connection Between Brain and Behaviour**

A fundamental goal in cognitive neuroscience is finding adequate methods to link brain activation patterns observed using sophisticated neuroimaging techniques to human behaviour processes. In Chapter 4, I used the RT-distance method to evaluate how neural representations from the MEG data related to behavioural categorisation reaction times. Despite being able to similarly decode the object exemplars for each of the four category organisations from the MEG data, successful decoding performance did not predict a relationship with behavioural RTs, with only some of the category organisations showing a consistent relationship between representational distances and RTs throughout the decoding time-course. This pattern of results highlights an important issue in the field of cognitive neuroscience: information that may be readily decoded from neurophysiological and neuroimaging data does not necessarily entail its functional use in human behaviour. Furthermore, this result speaks to concerns that have been raised about interpreting decoding results as behaviourally relevant “information” in the brain activation patterns (de-Wit et al., 2016). While the classifier may be able to determine a reasonable classification boundary to differentiate object categories from the corresponding whole-brain MEG activation patterns, it is difficult to determine what features of the underlying neural representations are being used by the classifier to decode one category from another in the MEG data. The classifier could utilise features in the neural signal that do not underlie categorisation behaviour in a human observer, resulting in representational distances that are not functionally relevant for behaviour. Defining what features are pertinent to neural and behavioural object categorisation processes remains an ongoing challenge for cognitive neuroscience, and one that is not easily solved.

Addressing the influence of task-related effects on measured brain activity is also an important consideration when trying to define the processes underlying object recognition. I obtained data from separate people for the brain and behaviour components of my experiment in Chapter 4 to avoid potential confounds related to the neural processes for

active categorisation. Previous research has evaluated the effect of active categorisation versus passive viewing on MEG-derived neural representations, showing little difference between decoding performance of these two data sets (Ritchie, Tovar, & Carlson, 2015). To show this, Ritchie et al. (2015) had participants view object images in two task conditions during an MEG recording: via a *categorisation* task, where participants categorised each stimulus according to whether the object was animate or inanimate; and a *distracted viewing* task, where participants categorised a letter superimposed over the object images as either a consonant or a vowel. This study found similar object exemplar decoding performance from the MEG data for both the active and distracted viewing tasks, with the authors suggesting that this was indicative of a ‘core representation’ of object category structure that is independent of task-related processes.

The independence of object category representations using MEG data is an interesting finding, one which is somewhat at odds with research from other neuroimaging modalities which do show a role of task requirements on object representations (Bracci, Daniels, & Op de Beeck, 2017; Harel, Kravitz, & Baker, 2014; Hebart et al., 2018; Nastase et al., 2016). Studies using fMRI have shown that the geometry of object representations shifts depending on task requirements, with attention enhancing features relevant to categorisation task requirements and minimising irrelevant information in the signal (Dobs, Schultz, Bühlhoff, & Gardner, 2018; Kaiser, Oosterhof, et al., 2016; Nastase et al., 2016). A similar effect of task is seen in a study by Harel, Kravitz and Baker (2014), who used fMRI to measure multivariate pattern responses to objects while participants performed a range of physical (e.g. colour) versus conceptual (e.g. size) categorisation tasks for the same set of objects. Results showed that object information in early visual areas was independent of task, with decoding performance maintained across both categorisation tasks in these areas. However, in object processing regions further along the ventral visual pathway, information related to specific object identities was task-dependent, such that the ability to decode object identity information was reduced when attempting to decode across different conceptual

tasks (Harel et al., 2014). Importantly, they note that while there was a reduction in decoding performance across tasks in some object-related regions, it was still possible to decode object identity in these regions despite differences in task requirements, suggesting the presence of some stable object identity information that is independent of task-related modulations. In light of these findings, it could be the case that my results using passive viewing of objects which show a relationship between brain representations and behaviour (Carlson et al., 2014; Ritchie et al., 2015), provide an indication of these stable aspects of object representations that are independent of task requirements.

Alternatively, there is recent research to suggest that the absence of task-related effects may be due to the neuroimaging modality utilised. In a follow-up study linking the fMRI data from Harel et al. (2014) to time-series MEG data, Hebart, Bankson, Harel, Baker, and Cichy (2018) evaluated object identity decoding performance between tasks in the MEG data alone. While their results showed differences in object identity decoding performance for conceptual versus physical tasks emerged from 530 ms post stimulus onset, they found no difference in classification accuracy between tasks. The authors suggest that the effect of task seen in the fMRI data (Harel et al., 2014) may be hidden by information originating from multiple brain sources in the MEG signal. That task dependent modulations of information appear somewhat undetectable in the MEG data in previous studies (Hebart et al., 2018; Ritchie et al., 2015) suggests that this may represent a signal-to-noise issue in whole-brain MEG data. In light of these findings, the RT-distance approach to linking MEG brain processes to behaviour may lack the sensitivity to identify subtle nuances of categorisation behaviour that are detectable via other neuroimaging modalities. Given the equivocal evidence for both stable and task-dependent components of object representations in the literature, delineating the contributions of these components of object representations remains an important task for future research (Hebart et al., 2018).

## 5.5. Conclusion

The work in this thesis has contributed to our understanding of the temporal dynamics of object representations. While I was unable to detect a real-world size dimension in the MEG object representations, this result highlights the difficulties that exist in defining and subsequently measuring specific conceptual properties of object categories. In my subsequent experiments, I employed a stimulus set that was designed to pose challenges for categorisation based on the current dominating principle in this field: the animacy dichotomy. I have shown that a large-scale organisation principle related to human-similarity provided the most consistent concept-based account of the time-varying object representations obtained using MEG. In addition, I made the important contribution of linking these human-similarity based brain representations to behaviour, showing for the first time, that representational distance relates to categorisation behaviour for categorisation other than that based on animacy. I have discussed the utility of evaluating the temporal aspects of object processing, and the challenges that exist in separating the contribution of low-level visual feature information from higher-level concepts which vary throughout the object processing time-course. I raised important issues related to stimulus set, and task design, and discussed the current evidence for task-related effects on neurophysiological data. Research into the brain processes that underlie visual object categorisation must develop better methods of linking findings to behaviour, to ensure that what we measure using these sophisticated neuroimaging techniques is ultimately meaningful for human behaviour.

## 5.6. References

- Bankson, B. B., Hebart, M. N., Groen, I. I., & Baker, C. I. (2017). The temporal evolution of conceptual object representations revealed through models of behavior, semantics and deep neural networks. *BioRxiv*, 223990.
- Bar, M. (2003). A cortical mechanism for triggering top-down facilitation in visual object recognition. *Journal of Cognitive Neuroscience*, 15(4), 600–609.  
<https://doi.org/10.1162/089892903321662976>
- Bracci, S., Daniels, N., & Op de Beeck, H. (2017). Task Context Overrides Object- and Category-Related Representational Content in the Human Parietal Cortex. *Cerebral Cortex*, 27(1), 310–321. <https://doi.org/10.1093/cercor/bhw419>
- Bracci, S., Kalfas, I., & Op de Beeck, H. (2017). The ventral visual pathway represents animal appearance over animacy, unlike human behavior and deep neural networks. *BioRxiv*, 228932.
- Capitani, E., Laiacona, M., Mahon, B., & Caramazza, A. (2003). What Are the Facts of Semantic Category-Specific Deficits? A Critical Review of the Clinical Evidence. *Cognitive Neuropsychology*, 20(3–6), 213–261.  
<https://doi.org/10.1080/02643290244000266>
- Caramazza, A., & Shelton, J. R. (1998). Domain-Specific Knowledge Systems in the Brain: The Animate-Inanimate Distinction. *Journal of Cognitive Neuroscience*, 10(1), 1–34. <https://doi.org/10.1162/089892998563752>
- Carlson, T. A., Goddard, E., Kaplan, D. M., Klein, C., & Ritchie, J. B. (2017). Ghosts in machine learning for cognitive neuroscience: Moving from data to theory. *NeuroImage*. <https://doi.org/10.1016/j.neuroimage.2017.08.019>
- Carlson, T. A., Hogendoorn, H., Kanai, R., Mesik, J., & Turret, J. (2011). High temporal resolution decoding of object position and category. *Journal of Vision*, 11(10), 9.  
<https://doi.org/10.1167/11.10.9>

- Carlson, T. A., Ritchie, J. B., Kriegeskorte, N., Durvasula, S., & Ma, J. (2014). Reaction Time for Object Categorization Is Predicted by Representational Distance. *Journal of Cognitive Neuroscience*, 26(1), 132–142. [https://doi.org/10.1162/jocn\\_a\\_00476](https://doi.org/10.1162/jocn_a_00476)
- Carlson, T. A., & Wardle, S. G. (2015). Sensible decoding. *NeuroImage*, 110, 217–218. <https://doi.org/10.1016/j.neuroimage.2015.02.009>
- Carlson, Tovar, D. A., Alink, A., & Kriegeskorte, N. (2013). Representational dynamics of object vision: The first 1000 ms. *Journal of Vision*, 13(10), 1–1. <https://doi.org/10.1167/13.10.1>
- Chao, L. L., Haxby, J. V., & Martin, A. (1999). Attribute-based neural substrates in temporal cortex for perceiving and knowing about objects. *Nature Neuroscience*, 2(10), 913–919.
- Cichy, R. M., Kriegeskorte, N., Jozwik, K. M., Bosch, J. J. F. van den, & Charest, I. (2017). Neural dynamics of real-world object vision that guide behaviour. *BioRxiv*, 147298. <https://doi.org/10.1101/147298>
- Cichy, R. M., Pantazis, D., & Oliva, A. (2014). Resolving human object recognition in space and time. *Nature Neuroscience*, 17(3), 455–462. <https://doi.org/10.1038/nn.3635>
- Clarke, A., Devereux, B. J., Randall, B., & Tyler, L. K. (2015). Predicting the Time Course of Individual Objects with MEG. *Cerebral Cortex*, 25(10), 3602–3612. <https://doi.org/10.1093/cercor/bhu203>
- Clarke, A., & Tyler, L. K. (2014). Object-Specific Semantic Coding in Human Perirhinal Cortex. *Journal of Neuroscience*, 34(14), 4766–4775. <https://doi.org/10.1523/JNEUROSCI.2828-13.2014>
- Connolly, A. C., Guntupalli, J. S., Gors, J., Hanke, M., Halchenko, Y. O., Wu, Y.-C., ... Haxby, J. V. (2012). The Representation of Biological Classes in the Human Brain. *Journal of Neuroscience*, 32(8), 2608–2618. <https://doi.org/10.1523/JNEUROSCI.5547-11.2012>

- Contini, E. W., Wardle, S. G., & Carlson, T. A. (2017). Decoding the time-course of object recognition in the human brain: From visual features to categorical decisions. *Neuropsychologia*, 105(Supplement C), 165–176.  
<https://doi.org/10.1016/j.neuropsychologia.2017.02.013>
- Cooke, T., Jäkel, F., Wallraven, C., & Bülthoff, H. H. (2007). Multimodal similarity and categorization of novel, three-dimensional objects. *Neuropsychologia*, 45(3), 484–495.
- Cortese, J. M., & Dyre, B. P. (1996). Perceptual similarity of shapes generated from fourier descriptors. *Journal of Experimental Psychology: Human Perception and Performance*, 22(1), 133.
- de-Wit, L., Alexander, D., Ekroll, V., & Wagemans, J. (2016). Is neuroimaging measuring information in the brain? *Psychonomic Bulletin & Review*, 23(5), 1415–1428.  
<https://doi.org/10.3758/s13423-016-1002-0>
- Dobs, K., Schultz, J., Bülthoff, I., & Gardner, J. L. (2018). Task-dependent enhancement of facial expression and identity representations in human cortex. *NeuroImage*, 172, 689–702. <https://doi.org/10.1016/j.neuroimage.2018.02.013>
- Fabre-Thorpe, M. (2011). The characteristics and limits of rapid visual categorization. *Frontiers in Psychology*, 2, 243. <https://doi.org/10.3389/fpsyg.2011.00243>
- Forde, E. M. (1999). Category specific recognition impairments: A review of important case studies and influential theories. *Aphasiology*, 13(3), 169–193.
- Goddard, E., Carlson, T. A., Dermody, N., & Woolgar, A. (2016). Representational dynamics of object recognition: Feedforward and feedback information flows. *NeuroImage*, 128, 385–397. <https://doi.org/10.1016/j.neuroimage.2016.01.006>
- Grill-Spector, K., & Malach, R. (2004). THE HUMAN VISUAL CORTEX. *Annual Review of Neuroscience*, 27(1), 649–677.  
<https://doi.org/10.1146/annurev.neuro.27.070203.144220>

- Grill-Spector, K., & Weiner, K. S. (2014). The functional architecture of the ventral temporal cortex and its role in categorization. *Nature Reviews. Neuroscience*, 15(8), 536–548. <https://doi.org/10.1038/nrn3747>
- Grootswagers, T., Ritchie, J. B., Wardle, S. G., Heathcote, A., & Carlson, T. A. (2017). Asymmetric Compression of Representational Space for Object Animacy Categorization under Degraded Viewing Conditions. *Journal of Cognitive Neuroscience*, 1–16. [https://doi.org/10.1162/jocn\\_a\\_01177](https://doi.org/10.1162/jocn_a_01177)
- Grootswagers, T., Wardle, S. G., & Carlson, T. A. (2016). Decoding Dynamic Brain Patterns from Evoked Responses: A Tutorial on Multivariate Pattern Analysis Applied to Time Series Neuroimaging Data. *Journal of Cognitive Neuroscience*, 29(4), 677–697. [https://doi.org/10.1162/jocn\\_a\\_01068](https://doi.org/10.1162/jocn_a_01068)
- Harel, A., Kravitz, D. J., & Baker, C. I. (2014). Task context impacts visual object processing differentially across the cortex. *Proceedings of the National Academy of Sciences*, 111(10), E962–E971. <https://doi.org/10.1073/pnas.1312567111>
- Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and Overlapping Representations of Faces and Objects in Ventral Temporal Cortex. *Science*, 293(5539), 2425–2430. <https://doi.org/10.1126/science.1063736>
- Hebart, M. N., Bankson, B. B., Harel, A., Baker, C. I., & Cichy, R. M. (2018). The representational dynamics of task and object processing in humans. *ELife*, 7. <https://doi.org/10.7554/eLife.32816>
- Isik, L., Meyers, E. M., Leibo, J. Z., & Poggio, T. (2014). The dynamics of invariant object recognition in the human visual system. *Journal of Neurophysiology*, 111(1), 91–102. <https://doi.org/10.1152/jn.00394.2013>
- Julian, J. B., Ryan, J., & Epstein, R. A. (2017). Coding of Object Size and Object Category in Human Visual Cortex. *Cerebral Cortex*, 27(6), 3095–3109. <https://doi.org/10.1093/cercor/bhw150>



- Kaiser, D., Azzalini, D. C., & Peelen, M. V. (2016). Shape-independent object category responses revealed by MEG and fMRI decoding. *Journal of Neurophysiology*, 115(4), 2246–2250. <https://doi.org/10.1152/jn.01074.2015>
- Kaiser, D., Oosterhof, N. N., & Peelen, M. V. (2016). The Neural Dynamics of Attentional Selection in Natural Scenes. *The Journal of Neuroscience*, 36(41), 10522–10528. <https://doi.org/10.1523/JNEUROSCI.1385-16.2016>
- Khaligh-Razavi, S.-M., Cichy, R. M., Pantazis, D., & Oliva, A. (2018). Tracking the Spatiotemporal Neural Dynamics of Real-world Object Size and Animacy in the Human Brain. *Journal of Cognitive Neuroscience*, 1–18.
- Kiani, R., Esteky, H., Mirpour, K., & Tanaka, K. (2007). Object Category Structure in Response Patterns of Neuronal Population in Monkey Inferior Temporal Cortex. *Journal of Neurophysiology*, 97(6), 4296–4309. <https://doi.org/10.1152/jn.00024.2007>
- Kiefer, M., & Pulvermüller, F. (2012). Conceptual representations in mind and brain: Theoretical developments, current evidence and future directions. *Cortex*, 48(7), 805–825. <https://doi.org/10.1016/j.cortex.2011.04.006>
- Konkle, T., & Caramazza, A. (2013). Tripartite Organization of the Ventral Stream by Animacy and Object Size. *Journal of Neuroscience*, 33(25), 10235–10242. <https://doi.org/10.1523/JNEUROSCI.0983-13.2013>
- Konkle, T., & Oliva, A. (2012). A Real-World Size Organization of Object Responses in Occipitotemporal Cortex. *Neuron*, 74(6), 1114–1124. <https://doi.org/10.1016/j.neuron.2012.04.036>
- Kriegeskorte, N., & Mur, M. (2012). Inverse MDS: Inferring Dissimilarity Structure from Multiple Item Arrangements. *Frontiers in Psychology*, 3. <https://doi.org/10.3389/fpsyg.2012.00245>

- Kriegeskorte, N., Mur, M., & Bandettini, P. A. (2008). Representational similarity analysis - connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, 2(4). <https://doi.org/10.3389/neuro.06.004.2008>
- Kriegeskorte, N., Mur, M., Ruff, D. A., Kiani, R., Bodurka, J., Esteky, H., ... Bandettini, P. A. (2008). Matching Categorical Object Representations in Inferior Temporal Cortex of Man and Monkey. *Neuron*, 60(6), 1126–1141. <https://doi.org/10.1016/j.neuron.2008.10.043>
- Logothetis, N. K. (2002). The neural basis of the blood–oxygen–level–dependent functional magnetic resonance imaging signal. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 357(1424), 1003–1037. <https://doi.org/10.1098/rstb.2002.1114>
- Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., & Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature*, 412(6843), 150–157. <https://doi.org/10.1038/35084005>
- Logothetis, N. K., & Sheinberg, D. L. (1996). Visual object recognition. *Annual Review of Neuroscience*, 19(1), 577–621.
- Lu, Z., Li, X., & Meng, M. (2016). Encodings of implied motion for animate and inanimate object categories in the two visual pathways. *NeuroImage*, 125, 668–680. <https://doi.org/10.1016/j.neuroimage.2015.10.059>
- Mack, M. L., & Palmeri, T. J. (2011). The Timing of Visual Object Categorization. *Frontiers in Psychology*, 2. <https://doi.org/10.3389/fpsyg.2011.00165>
- Mack, M. L., & Palmeri, T. J. (2015). The dynamics of categorization: Unraveling rapid categorization. *Journal of Experimental Psychology: General*, 144(3), 551.
- Mahon, B. Z., & Caramazza, A. (2009). Concepts and Categories: A Cognitive Neuropsychological Perspective. *Annual Review of Psychology*, 60(1), 27–51. <https://doi.org/10.1146/annurev.psych.60.110707.163532>

- Martin, A. (2007). The Representation of Object Concepts in the Brain. *Annual Review of Psychology*, 58(1), 25–45.  
<https://doi.org/10.1146/annurev.psych.57.102904.190143>
- Martin, A. (2016). GRAPES—Grounding representations in action, perception, and emotion systems: How object properties and categories are represented in the human brain. *Psychonomic Bulletin & Review*, 1–12.
- Mur, M., Meys, M., Bodurka, J., Goebel, R., Bandettini, P. A., & Kriegeskorte, N. (2013). Human object-similarity judgments reflect and transcend the primate-IT object representation. *Frontiers in Psychology*, 4, 128.
- Naselaris, T., & Kay, K. N. (2015). Resolving ambiguities of MVPA using explicit models of representation. *Trends in Cognitive Sciences*, 19(10), 551–554.
- Nastase, S. A., Connolly, A. C., Oosterhof, N. N., Halchenko, Y. O., Guntupalli, J. S., Castello, M. V. di O., ... Haxby, J. V. (2016). Attention selectively reshapes the geometry of distributed semantic representation. *BioRxiv*, 045252.  
<https://doi.org/10.1101/045252>
- Nosofsky, R. M. (1986). Attention, similarity, and the identification–categorization relationship. *Journal of Experimental Psychology: General*, 115(1), 39.
- Panis, S., Torfs, K., Gillebert, C., Wagemans, J., & Humphreys, G. W. (2017). Neuropsychological evidence for the temporal dynamics of category-specific naming. Retrieved from <https://ora.ox.ac.uk/objects/uuid:1435e0f4-7840-48f2-9d02-b9f6aca6667a>
- Peelen, M. V., & Downing, P. E. (2017). Category selectivity in human visual cortex: Beyond visual object recognition. *Neuropsychologia*. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0028393217301215>

- Proklova, D., Kaiser, D., & Peelen, M. V. (2016). Disentangling Representations of Object Shape and Object Category in Human Visual Cortex: The Animate–Inanimate Distinction. *Journal of Cognitive Neuroscience*, 1–13.  
[https://doi.org/10.1162/jocn\\_a\\_00924](https://doi.org/10.1162/jocn_a_00924)
- Ritchie, J. B., Tovar, D. A., & Carlson, T. A. (2015). Emerging Object Representations in the Visual System Predict Reaction Times for Categorization. *PLoS Comput Biol*, 11(6), e1004316. <https://doi.org/10.1371/journal.pcbi.1004316>
- Schmidt, F., Hegele, M., & Fleming, R. W. (2017). Perceiving animacy from shape. *Journal of Vision*, 17(11), 10–10.
- Sha, L., Haxby, J. V., Abdi, H., Guntupalli, J. S., Oosterhof, N. N., Halchenko, Y. O., & Connolly, A. C. (2015). The Animacy Continuum in the Human Ventral Vision Pathway. *Journal of Cognitive Neuroscience*, 27(4), 665–678.  
[https://doi.org/10.1162/jocn\\_a\\_00733](https://doi.org/10.1162/jocn_a_00733)
- Thorpe, S., Fize, D., & Marlot, C. (1996). Speed of processing in the human visual system. *Nature*, 381(6582), 520–522.
- VanRullen, R., & Thorpe, S. J. (2001). The Time Course of Visual Processing: From Early Perception to Decision-Making. *Journal of Cognitive Neuroscience*, 13(4), 454–461. <https://doi.org/10.1162/08989290152001880>
- Warrington, E. K., & Shallice, T. (1984). Category specific semantic impairments. *Brain*, 107(3), 829–853.

## APPENDIX ONE

---

---

### ETHICAL APPROVALS

---

MACQUARIE UNIVERSITY  
RECEIVED  
11 NOV 2013  
RESEARCH OFFICE

OFFICIAL USE ONLY	
SURNAME:	Carlson
REF. NO.:	5201300804
ARCHIVE:	7 <input type="checkbox"/> 25 <input type="checkbox"/>

MACQUARIE UNIVERSITY  
HUMAN RESEARCH ETHICS COMMITTEE  
(HREC)

APPLICATION FORM

ALL RESEARCHERS MUST COMPLETE SECTIONS 1-11  
ATTACH APPENDICES UNLESS RELEVANT TO YOUR APPLICATION

Check Yes, No, or N/A by marking a box with an X to indicate your answer.

For further information on how to complete this form see *National Statement*  
<http://www.nhmrc.gov.au/files/nhmrc/publications/attachments/e72.pdf>  
and *Australian Code for the Responsible Conduct of Research (The Code)*  
<http://www.nhmrc.gov.au/files/nhmrc/file/publications/synopses/r39.pdf>

SECTION 1: PROJECT SUMMARY AND ADMINISTRATION

1.1 Full project title:

Decoding the neural representations of objects in the human brain

(a) Short name by which the project will be known (if appropriate):

n/a

(b) Name of Chief Investigator/Supervisor:

Thomas Carlson

If the project is to be undertaken by an Honours, Postgraduate or Higher Degree student, then the supervisor will be considered the Chief Investigator for this application. The student may be nominated as a co-investigator.

(c) If this research is being conducted as a teaching project, provide the unit code and name:

n/a

**1.2 Indicate the Human Research Ethics Committee (HREC) that you consider to be the primary one for this project: (see *National Statement* Chapter 5.3)**

**Note:** If the primary site for this project is affiliated with a public health or justice health organisation, please complete your application on the NEAF instead by clicking [here](#).

HREC (HUM)

- (a) Has this project been submitted for approval to another HREC? ☐ Yes ☒ No

**If you answered 'Yes'—before completing this form—please contact the ethics secretariat for further information regarding the external application approval process.**

- (b) Will this project be submitted for approval to another HREC? ☐ Yes ☒ No  
(See *National Statement* 5.3.2)

If you answered 'Yes' give the name of the HREC(s), and indicate the status of the application (i.e. submitted, approved, deferred or rejected): Please do not submit to more than one HREC concurrently:

- (c) Will this project be submitted for approval/permission to another institution? ☐ Yes ☒ No  
(for example schools, childcare centres, DEC, private companies)

If you answered 'Yes' give the name of the institution, and attach all correspondence to this application.

- (d) Has this project been previously submitted/approved by MQ HREC? ☐ Yes ☒ No

If Yes, what was the reference number and original date of approval?

- (e) Are you applying for a new clearance because the previous clearance has expired? ☐ Yes ☒ No  
(i.e. it has been five years or more since the original clearance was issued)

If Yes, please append your final report to this application and indicate below any differences between this application and the previously-approved project.

### 1.3 Investigator Details

List the details of the Chief Investigator, and any Co-Investigators, Associate Investigators, Supervisors, Research Assistants or Research Co-ordinators

#### Chief Investigator/Supervisor

**Note:** If the project is to be undertaken by an Honours/Masters/PhD student, the supervisor will be considered the Chief Investigator. The student may be named as a co-investigator.)

Name:	Thomas Carlson
Title: (Prof, Dr. etc.)	Dr
MQ Staff no. (Mandatory)	MQ20130093
Qualifications:	PhD
Positions held:	Senior Lecture
Faculty:	Faculty of Human Sciences
Department/Unit:	Department of Cognitive Science
Tel No. (W):	02 9850 4133
Mobile No:	04 1449 1771
E-mail address:	thomas.carlson@mq.edu.au

**Co-Investigator** ☐ **Associate Investigator** ☒, **Research Assistant** ☐ **Research Coordinator** ☐  
(please mark one with an X)

Name:	Lincoln Colling
Title:	Dr
Co-investigator type:	Staff : <input checked="" type="checkbox"/> <b>OR</b> Student: <input type="checkbox"/>
Staff/student no. (Mandatory)	MQ41247191
Qualifications:	PhD
Positions held: (if student, specify degree, course in which enrolled)	Senior Research Officer
Faculty:	Faculty of Human Sciences
Department:	Department of Cognitive Science
Tel No. (W):	
Mobile No:	0406515140
E-mail address: (Students: Please use your MQ student email address)	lincoln.colling@mq.edu.au

**Co-Investigator** ☒ **Associate Investigator** ☐, **Research Assistant** ☐ **Research Coordinator** ☐  
(please mark one with an X)

Name:	Erika Contini
Title:	Ms
Co-investigator type:	Staff : <input checked="" type="checkbox"/> <b>OR</b> Student: <input type="checkbox"/>
Staff/student no. (Mandatory)	41994930
Qualifications:	Bachelor of Liberal Studies (Psychology honours)
Positions held:	Combined PhD / Master of Clinical Neuropsychology



(if student, specify degree, course in which enrolled)	
Faculty:	Faculty of Human Sciences
Department:	Department of Psychology / Department of Cognitive Science
Tel No. (W):	9850 2958
Mobile No:	
E-mail address: (Students: Please use your MQ student email address)	erika.contini@mq.edu.au

**Co-Investigator** ☐ **Associate Investigator** ☒, **Research Assistant** ☐ **Research Coordinator** ☐  
(please mark one with an X)

Name:	John Brendan Welsh Ritchie
Title:	Mr
Co-investigator type:	Staff : <input checked="" type="checkbox"/> <b>OR</b> Student: <input type="checkbox"/>
Staff/student no. (Mandatory)	MQ20137981
Qualifications:	Masters Degree
Positions held: (if student, specify degree, course in which enrolled)	Research Associate
Faculty:	Faculty of Human Sciences
Department:	Department of Cognitive Science
Tel No. (W):	
Mobile No:	0452028455
E-mail address: (Students: Please use your MQ student email address)	john.ritchie@mq.edu.au

**Co-Investigator** ☐ **Associate Investigator** ☐, **Research Assistant** ☐ **Research Coordinator** ☐  
(please mark one with an X)

Name:	
Title:	
Co-investigator type:	Staff : <input type="checkbox"/> <b>OR</b> Student: <input type="checkbox"/>
Staff/student no. (Mandatory)	
Qualifications:	
Positions held: (if student, specify degree, course in which enrolled)	
Faculty:	
Department:	
Tel No. (W):	
Mobile No:	
E-mail address: (Students: Please use your MQ student email address)	

**Co-Investigator** ☐ **Associate Investigator** ☐, **Research Assistant** ☐ **Research Coordinator** ☐  
 (please mark one with an X)

Name:	
Title:	
Co-investigator type:	Staff : <input type="checkbox"/> <b>OR</b> Student: <input type="checkbox"/>
Staff/student no. (Mandatory)	
Qualifications:	
Positions held: (if student, specify degree, course in which enrolled)	
Faculty:	
Department:	
Tel No. (W):	
Mobile No:	
E-mail address: (Students: Please use your MQ student email address)	

**1.4 Nominated Contact Person (nominate one MQ staff member from Q1.3, who will be the first point of contact for the MQ HREC regarding this protocol)**

Name: Thomas Carlson
----------------------

**1.5 Type of Research** (please mark with an X the most appropriate):

- |  |                                     |
|--|-------------------------------------|
| i. This application is being submitted for staff research purposes   | <input checked="" type="checkbox"/> |
| ii. This application is being submitted by staff and/or students for teaching purposes                       | <input type="checkbox"/>            |
| iii. This application involves research which will be submitted by a student in the following degree program | <input type="checkbox"/>            |

- Type of degree undertaken:
- |                          |                          |
|--------------------------|--------------------------|
| • Doctorate              | <input type="checkbox"/> |
| • Masters                | <input type="checkbox"/> |
| • Postgraduate Diploma   | <input type="checkbox"/> |
| • Honours                | <input type="checkbox"/> |
| • Other – please specify | <input type="checkbox"/> |

--

## 1.6 Duration of study

- (a) Indicate the proposed date of commencement of the project (Researchers are reminded that projects may not commence without the written approval of the Human Research Ethics Committee (HREC)) (See *National Statement* 5.2.5-5.2.12 and *The Code* section 1.6)

DATE: [29/10/2013]

- (b) Proposed completion date of the project:

DATE: [29/10/2018]

(Note that approval is given for a five-year period, conditional on completion of a yearly Progress Report (using the Progress Report Form). After five years from the original approval date, you will need to submit a Final Report for the work and a new application for approval if you wish the project to continue. The five-year time limit enables the Committee to fully review research in an environment where legislation, guidelines and codes are always changing.)

## 1.7 Funding/Contract Details (see *National Statement* chapter 5.4)

- (a) Is this a funded project or do you intend to apply for funding? ☒ Yes ☐ No

If you answered 'No', go to section 2.

If you answered 'Yes', list the funding bodies that support this project or to which submission is planned:

ARC Future Fellowship (ARC Ref. FT120100816) (IRIS Grant Ref. 9201101354)

- (b) Status

Approved

☒ Yes ☐ No

Pending

☐ Yes ☐ No

Submitted

☐ Yes ☐ No

If not, please specify the planned date of submission: [ ]

- (c) Please state the title of the project as it appears on your grant application:

Decoding the neural representation of objects in the human brain

- (d) Will this study still be undertaken if funding is not successful? ☒ Yes ☐ No

- (e) Will the project require a research contract? ☐ Yes ☒ No

If 'Yes', please complete a contracts advice request form

[http://www.research.mq.edu.au/for/researchers/how\\_to\\_request\\_a\\_research\\_contract](http://www.research.mq.edu.au/for/researchers/how_to_request_a_research_contract)

**SECTION 2: METHODS AND NATURE OF RESEARCH***(See National Statement section 3).*

**2.1 The nature of this project is most appropriately described as involving: (please mark with an X the relevant box or boxes - more than one may apply)**

a. experiments	<input checked="" type="checkbox"/>
b. questionnaire(s) (please attach these to the application)	<input type="checkbox"/>
c. interviews (please attach questions to this application)	<input type="checkbox"/>
d. focus groups (please attach questions to this application)	<input type="checkbox"/>
e. documentary (film, audio and/or video)	<input type="checkbox"/>
f. on-line data collection (See Section 3 of the National Statement)	<input type="checkbox"/>
g. action research	<input type="checkbox"/>
h. participant observation	<input type="checkbox"/>
i. physiological investigations(s)	<p>If any of the options from (i) to (p) apply to your project, please complete your application on the NEAF instead by clicking <b>here</b>.</p>
j. clinical studies	
k. biomechanical devices(s) or invasive devices or procedures	
l. blood, tissue, other bodily fluids, recombinant DNA, gene therapy, or bio hazardous material (If you tick this item, please contact Work Health & Safety regarding the non-GMO risk assessment form. For further queries about this please email: <a href="mailto:biohazard@mq.edu.au">biohazard@mq.edu.au</a> )	
m. behavioural genetics, genetic screening and/or genetic related epidemiological studies (If you tick this item, please refer to Chapter 3.5 of the National Statement before proceeding with the completion of the application).	
n. exposure to ionizing radiation (APPENDIX F must be completed if you tick this item)	
o. the administration of illegal drugs, legal drugs (eg, alcohol, nicotine) and/or medication that is NOT part of a clinical trial (APPENDIX G must be completed if you tick this item)	
p. a clinical trial (APPENDIX C must be completed if you tick this item)	
q. other (please provide details)	

--

- 2.2 Describe the project using lay terms and include the aims of the project, the research plan and the methods. Include a description of the projected number, sex and age range of participants (including inclusion and exclusion criteria) and provide a detailed description of what will be required of participants. You must satisfy the HREC that the study has a clear rationale (include references) and conducted in accordance with the accepted ethical principles governing research involving humans.**

**Note:** Please refrain from using abbreviations and acronyms without providing the full version in the first instance (e.g. 'National Health and Medical Research Council (NHMRC)' subsequently becomes 'NHMRC').

*The description must be in at least font size 11 and no longer than two (2) pages.*

The primary aim of this project is to investigate the neural mechanisms underlying performance on choice-decision tasks. In many types of choice-decisions tasks where participants are asked to discriminate or categorise two classes of stimuli performance will vary across different exemplars of each category. For example, if participants are asked to categorise pictures of either depicting animate or inanimate objects performance on some exemplars – for example, a human face – may differ relative to performance on other exemplars – for example, a banana. These performance differences may take the form of differences in reaction time or differences in the accuracy of categorisation.

Recent work has been able to link task performance to differences in how particular stimuli are represented in the brain (Carlson et al., 2013). This research covered by this application aims to extend this earlier to in at least two ways. First, by extending the approach adopted in this earlier work to new domains with the aim of assessing the generalizability of this approach. This will involve not only using stimuli from different categories but also presenting stimuli through different sensory modalities by, for example, using auditory stimuli instead of visual stimuli. The second aim will be to further our understanding of the link between how the brain represents stimuli and performance on decision tasks. This can be done by modulating the performance on the task by increasing task difficulty by, for example, degrading stimuli or by altering stimuli in ways that are known to change how they are represented in the brain – for example, inverting a human face.

The experiments will employ participants aged 18 years and over (17 years and over if Macquarie University students). We anticipate that the majority of these participants will be Macquarie University students.

The experiments will all be variations of a theme. In the experiments the participants will be asked to view or listen to stimuli and to make simple decisions about them (e.g., “press a button when you see an inanimate object”). In some cases, participants may also be asked to view or listen to stimuli passively, without making decisions. The experiments make also differ in terms of the non-invasive measures that they will employ. In some experiments only the responses produced by the participant will be recorded – for example, the button they pressed and the time taken to press the button. These measures will be used to quantify task performance in terms of reaction time and accuracy. In other experiments, we will also record signals with MRI, EEG, and MEG. These signals are radiofrequency waves, electrical potential differences, and magnetic fields respectively. From these signals it is possible to infer how the brain represents the stimuli presented to the participants.

#### References

Carlson, T. A., Ritchie, J. B., Kriegeskorte, N., Durvasula, S., & Ma, J. (2013). RT for Object Categorization Is Predicted by Representational Distance. *Journal of Cognitive Neuroscience*, 1–11. doi:10.1162/jocn\_a\_00476

### SECTION 3: RECRUITMENT OF PARTICIPANTS

(See *National Statement* sections 1 and 2)

#### 3.1 How many participants will be involved in this study? Comment on how this number will allow the aims of the study to be achieved:

The study will involve a series of experiment, with approximately 20 participants expected for each experiment. This number way, however, vary slightly between experiments and exact number to recruit will be decided on the basis of considerations about the statistical power needed to achieve meaningful results using the proposed analysis techniques.

#### 3.2 Recruitment of participants

##### (a) What processes will be used to identify and select potential participants?

Participants will be able to sign up for research using the Department of Cognitive Science's MACCS SONA subject pool system (maccs.sona-systems.com). In addition to the MACCS SONA subject pool, the experiments that can be completed for course credit will also be placed on the Department of Psychology MQ-PSY SONA subject pool (mq-psy.sona-systems.com).

##### (b) Who will be involved in the recruitment of participants?

The investigators named above (including research assistants)

##### (c) Describe how and where the initial contact will be made with potential participants:

A study description will be placed on the MACCS / PSY SONA subject websites, and volunteers will be signup through the website.

##### (d) Describe the research data that will be collected directly from participants:

For MEG experiments we will record magnetic fields around the head of the participant. These magnetic fields reflect activity occurring in the brain of the participant.  
For EEG experiments we will record differences in electrical potential across the scalp of the participant. These electrical potential differences reflect activity occurring in the brain of the participant.  
For MRI experiments we will record radiofrequency energy responses inside a strong magnetic field. These responses can be used to infer changes in blood flow occurring in the brain and can be used to infer brain composition.  
For behavioural experiments will record button press responses (or similar: e.g., a vocal response) that reflect decisions made on the task. The accuracy of the decisions and the time taken to make the decisions can be determined on the basis of these responses

##### (e) If recruiting will be done through an organisation, specify how consent from the organisation will be obtained. Please provide copies of relevant correspondence with the organisation.

##### (f) Will participants be involved in any related studies?

☐ Yes ☒ No

If you answered 'Yes', please provide further details below:

#### 3.3 Indicate where the research will be undertaken. Please specify all location(s) and venues (including the URL for web-based studies).

Macquarie University including the Australian Hearing Hub (behavioural, MEG, EEG) and the Macquarie University Hospital (MRI)

**3.4 Does recruitment involve a direct personal approach from the researchers to the potential participants?**

☐ Yes ☒ No

If you answered 'Yes', what precautions will be taken to minimise any pressure (real or perceived) on individuals to enrol?

**3.5 Does recruitment involve the circulation/publication of an advertisement?**

☐ Yes ☒ No

If you answered 'Yes', provide a copy of the advertisement and indicate where and how often it will be published.

**3.6 Will participants receive any financial or other benefits as a result of participation?**

☒ Yes ☐ No

If you answered 'Yes', what is the amount/benefit and the justification for this? Please note it is University policy that this amount is limited to \$100 per participant. If you exceed this limit, you must provide a detailed justification.

(See *National Statement* 2.2.9 Coercion and pressure and 2.2.10 Reimbursing participants)

If the proposed research activity involves Aboriginal and Torres Strait Islander people, refer to Principle 11 of the Australian Institute of Aboriginal and Torres Strait Islander Studies (AIATSIS) *Guidelines for Ethical Research in Australian Indigenous Studies* (<http://www.aiatsis.gov.au/research/docs/GERAIS.pdf>) and Q12 of APPENDIX A of the Macquarie University ethics application form.

Participants will receive \$15 per hour for behavioural experiments and \$20 per hour for EEG/MEG/MRI experiments. This will be given to compensate them for their time and costs (e.g., travel costs etc).

Participants taking part as part of their Psychology course will receive course credit.

**3.7 Is the research targeting any particular ethnic or community group?**

☐ Yes ☒ No

(See *National Statement* 4 – Ethical consideration specific to participants, 4.7 – Aboriginal and Torres Strait Islander Peoples, 4.8 – People in other countries)

If you answered 'No', proceed to SECTION 4

If you answered 'Yes', which group is being targeted?

**3.8 Will recruitment be conducted in consultation with a representative of this group?**

☐ Yes ☐ No

(a) If you have not consulted a representative of this group, please provide your reasons for not undertaking any consultation.

- (b) If you have consulted a representative, with whom have you consulted?  
How do they represent this group?

#### SECTION 4: DESCRIPTION OF PARTICIPANTS

##### 4.1 What is the age range of participants involved in this study?

Participants will be 18 years and above. In the case of participants that are Macquarie University students some may be aged 17 years. We anticipate that many (but not all) participants will be Macquarie University Students.

- (a) Are the participants Macquarie University students? ☒ Yes ☐ No

##### 4.2 Does your research involve working with minors?

(See *National Statement* chapter 4.2 – Children and young people)

- ☐ Yes ☒ No

If you answered 'Yes' to 4.2, you will need to contact **Ms Elaine Parica** (Enquiries Officer) in Human Resources (+61 2 9850 1036) regarding the requirements under the new Working with Children Check. Please also refer to the **NSW Commission for Children and Young People**.

##### 4.3 Are the participants who are the focus of this research/teaching project:

(Please mark with an X the boxes that apply to the participants involved in your research. More than one category may be applicable)

- |   |                                     |
|---|-------------------------------------|
| a. in a teacher–student relationship with the researchers? (See <i>National Statement</i> chapter 4.3)  | <input checked="" type="checkbox"/> |
| b. a teaching unit co-ordinator in which the research takes place? (See <i>National Statement</i> chapter 4.3)                                | <input checked="" type="checkbox"/> |
| c. in an employer–employee relationship with the researchers? (See <i>National Statement</i> chapter 4.3)                                     | <input checked="" type="checkbox"/> |
| d. in any other dependent relationship with the researchers? (See <i>National Statement</i> chapter 4.3)                                      | <input type="checkbox"/>            |
| e. Macquarie University staff members?  | <input checked="" type="checkbox"/> |
| f. Aboriginal or Torres Strait Islander?<br>(if you tick this box APPENDIX A must be completed) (See <i>National Statement</i> chapter 4.7)   | <input type="checkbox"/>            |
| g. Residing in countries outside Australia<br>(if you tick this box APPENDIX B must be completed) (See <i>National Statement</i> chapter 4.8) | <input type="checkbox"/>            |
| h. wards of state? (See <i>National Statement</i> chapter 4.2)  | <input type="checkbox"/>            |
| i. prisoners or detainees? (See <i>National Statement</i> chapter 4.3)  | <input type="checkbox"/>            |



j. refugees or asylum seekers?	<input type="checkbox"/>
k. members of the armed services? (See <i>National Statement</i> chapter 4.3)	<input type="checkbox"/>
l. mentally ill? (See <i>National Statement</i> chapter 4.5)	<input type="checkbox"/>
m. intellectually impaired? (See <i>National Statement</i> chapter 4.5)	<input type="checkbox"/>
n. physically disabled?	<input type="checkbox"/>
o. unconscious or critically-ill patients? (See <i>National Statement</i> chapter 4.4)	Please complete <b>NEAF</b> instead
p. in a carer-client relationship with the researcher (e.g. medical practitioner, social worker, psychologist, legal practitioner, etc.)? (See <i>National Statement</i> chapter 4.3)	
q. in a carer-client relationship with other professional workers (e.g. medical practitioner, social worker, psychologist, legal practitioner, etc.)? (See <i>National Statement</i> chapter 4.3)	<input type="checkbox"/>

**4.4 If you marked any items in 4.3 with an X (disregarding option “o”), please give details and explain how you will mitigate real or perceived coercion.**

Some of the investigators named above will be teaching units in the departments of Psychology or Cognitive Science, and students enrolled in these units may choose to participate, either for payment or for course credit.

Some Macquarie University members may also choose to participate in this research.

For these reasons, it will be emphasised on the consent forms associated with this project that participants are free to withdraw their participation at any time without consequence. Furthermore, all signups will be managed through the SONA system rather than directly through the researchers to mitigate any perceived coercion.

**SECTION 5: RISKS AND BENEFITS (See *National Statement* chapter 2.1)**

**5.1 What expected benefits (if any) will this research have for the wider community?**

The benefits to the wider community will be primarily scientific. The wider community may benefit from an enhanced understanding of how people understand, recognise, and categorise auditory and visual material

**5.2 What expected benefits (if any) will this research have for participants?**

Participants may also benefit from the scientific knowledge gained by completion of this research. Furthermore, participants may also benefit from a greater understanding of the process of scientific research that may be gained by participating. A greater understanding of the scientific process will be of particular benefit to students enrolled in Psychology or related courses.

**5.3 Are there any risks to participants as a result of participation in this research project?** ☒ Yes ☐ No

- (a) If you answered ‘Yes’, what are those risks (e.g. physical or psychological distress)?  
If you answered ‘No’, go to Q5.4

Some of the methods used in these experiments (e.g., MRI, MEG) will require participants to lie down in a somewhat restricted space. Some participants may find this restricted space uncomfortable.

In addition to the space restrictions, MRI also makes a loud "hammering" sound when in operation. Some participants may find this noise uncomfortable.

- (b) Explain how these risks will be negated/minimised/managed.

During the experiment the participant will be able to communicate with the experimenter at all times. If participants feels any discomfort and feels the need to terminate the experiment then the participant will be able to communicate this to the experimenter, and it will be possible to immediately cease the experiment.

- (c) Indicate who will be primarily responsible for dealing with any problems and their expertise in handling these.

The experimenter, or research support staff, will be primarily responsible for dealing with any problems.

- (d) Explain how these risks will be monitored.

Participants will be in constant communication with the experimenter throughout the experiment and the experiment will immediately cease upon request from the participant.

(e) Explain how these risks, if they result in harm to participants, will be reported.

The risks are only in terms of discomfort, and we do not anticipate any harm to the participants. However, if serve adverse effects do occur then they will be reported to the Macquarie Unviersity HREC with 72 hours.

**5.4 Does the project:**

- (a) require any physically-invasive, or potentially-harmful, procedures? ☐ Yes ☒ No
- (b) involve research on patients? ☐ Yes ☒ No
- (c) involve the use of any drugs and/or devices? ☐ Yes ☒ No

If you answered 'Yes' to (a), (b) and/or (c), please complete your application on the NEAF instead by clicking [here](#).

**5.5 Will the true purpose of the research be concealed from the participant(s)?** ☐ Yes ☒ No

(See *National Statement* chapter 2.3 – Qualifying or waiving conditions for consent)

If you answered 'No' go to Q5.7.

If you answered 'Yes', give details of the deception and a justification for its use. Also provide details of how participants will be debriefed once their participation has been completed.

**5.6 Are participants given the opportunity to re-consent to the research after they have been debriefed?**

(attach debrief and consent form)

☐ Yes ☐ No ☐ N/A

If you answered 'No', explain your answer.

**5.7 Does this research involve the direct investigation of any illegal behaviour?** ☐ Yes ☒ No

**5.8 Does this research have the potential to elicit information about illegal behaviour?** ☐ Yes ☒ No

(See *National Statement* chapter 4.6 – People who may be involved in illegal activities.)

If you answered 'No' to both questions above, go to **SECTION 6**

If you answered 'Yes' to either question above, please provide the details below:

**5.9 Has this illegal behaviour already been dealt with by the criminal justice system?** ☐ Yes ☐ No ☐ Unsure

If you answered 'Yes' go to **SECTION 6**

- (a) Describe the types of illegal behaviour that are the subject of this research or might be identified in this project.

- (b) Is it possible to link the identity of participants and information about illegal activity? ☐ Yes ☐ No

If you answered 'No', what steps have been taken to ensure that participants cannot be identified?

If you answered 'Yes', please assess the risk to participants and provide a justification for exposing participants to this risk.

- (c) Are any of the participants in this study likely to be perpetrators, victims and/or witnesses of the illegal activity? ☐ Yes ☐ No

If you answered 'Yes', please provide further details.

- (d) Do any of the illegal activities involve minors? ☐ Yes ☐ No

If you answered 'Yes', please provide further details.

## SECTION 6: PRIVACY AND PUBLICATION OF RESULTS

Refer to *The Code*. For health related information refer to the Statutory Guidelines made under the *Health Records and Information Privacy (HRIP) Act 2002 (NSW) Statutory Guidelines on Research via Privacy NSW HRIP Act* and *NHMRC The Regulation of Health Information Privacy in Australia* <http://www.nhmrc.gov.au/publications/synopses/nh53syn.htm>

### 6.1 Is there a requirement for the researchers to obtain information about participants from:

- |   |                              |  |
|---|------------------------------|--|
| a. from Commonwealth departments or agencies?   | <input type="checkbox"/> Yes | <input checked="" type="checkbox"/> No |
| b. from State departments or agencies?  | <input type="checkbox"/> Yes | <input checked="" type="checkbox"/> No |
| c. from other third parties, such as universities, hospitals, schools, private practices, business etc? | <input type="checkbox"/> Yes | <input checked="" type="checkbox"/> No |

IF YOU ANSWERED 'Yes' TO ANY OF THE ABOVE ITEMS YOU MUST COMPLETE APPENDIX D. THIS IS A REQUIREMENT OF STATE AND COMMONWEALTH PRIVACY LEGISLATION.

### 6.2 Will you be recording (voices or images) any part of your study on audio-recording, film/video, or other electronic medium?

☐ Yes ☐ No

(See *National Statement 2.3* – Qualifying or waiving conditions for consent)

If you answered 'Yes', what is the medium, what will be recorded and what are the circumstances under which this recording will be undertaken? For what purposes is the recording required? Does the proposed research activity involve the secretive use of any photography, video-recording, audio- recording or other recording method?

The proposed research does not include any secretive auditory or visual recordings The only recordings that will take place are those that constitute the response measures in the proposed experiment. These include recording button press responses times and accuracy of responses. In addition, the EEG, MEG, and MRI experiments will also involve recording electrical potentials, magnetic fields, and radiofrequency waves around the participants' heads.

### 6.3 Will the recordings be sent to persons outside the research team?

☒ Yes ☐ No ☐ N/A

If you answered 'Yes', explain who would have access and describe how participant confidentiality will be ensured.

In exceptional circumstances where specific expertise in data analysis is required on a short term basis then recordings obtained with EEG, MEG, or MRI or data derived from the recording of button press responses may be shared with people outside of the research team. If this does occur the shared data will be shared in the form of de-identified data from which it would be impossible to identify the participants. This will only occur when short term assistance is required with data and in cases where it is not feasible to amend the current application to include this person because of time constraints or because of the very temporary nature of their involvement. Where the assistance is required on a longer term basis then the current application will be amended to add the additional researchers.

### 6.4 How will the results of the study be disseminated (publication and presentation of the research results)?

(See *National Statement 1.3* - Guidelines)

The results will be disseminated through scholarly journal articles, conference presentations, and via press releases organised by the Macquarie University PR Unit. Some of the data may also be presented in the form of a empirical thesis.

**6.5 Do you intend to present data relating to specific individuals (e.g. quotes, video, audio) in presentations or publications?**

☒ Yes

☐ No

If you answered 'Yes', please provide specific details.

In some cases the recorded obtained by EEG, MEG or MRI from a specific participant may be presented when their results are representative of the group. This will only be done when it is not feasible to produce a group average from the data. When individual data is presented it will be de-identified. It should also be noted that is not possible to identify participants from recordings of this time.

**6.6 How will the confidentiality of data collected/disseminated, including the identity of participants, be ensured (this includes data referred to in 6.2 above)?**

Only an allocated participant number will identify the data that is collected, and any personal information that is collected will be accessible only to the investigators named on this project.

**6.7 Will feedback or results be made available to participants?**  
(See *National Statement 1.5 – Justice*.)

☒ Yes

☐ No

Give details of how feedback or results will be made available to individual participants or, if relevant, to other groups. If no feedback is planned, a justification for this must be provided.

Where appropriate, participants will be verbally debriefed at the end of the experiment. Participants who are taking part for course credit may also receive an information sheet will references and readings about the research area to enhance the educational value of participation.

**6.8 Are you planning to retain data for the minimum period of five (5) years from the most recent publication of the research?**  
(See *The Code 2.1.1*)

☒ Yes

☐ No

If you answered 'Yes', please indicate how long you intend to retain data for.

It is usually the requirement for journal publication that the data are retained for 7 years from the date of publication.

If you answered 'No', please explain why not.

**6.9 Provide details about where the hard and electronic copies of data will be securely stored while the project is ongoing and after it is completed (e.g. will it be stored in a locked filing cabinet in the researcher's University office, in a locked cabinet in the researcher's home office?).**

Data recorded on computers in labs will be copied to researchers' password protected computers, most of which are in (lockable) offices in the University.

**6.10 Who will have access to these data (e.g. Researcher, supervisor, other researchers not stated on the application, any other third party)?**

Only the researchers listed on this application will have access to these data.

**6.11 Is there any possibility that information of a personal nature could be revealed to persons not directly connected with this project?**

☐ Yes

☒ No

If you answered 'Yes', please provide details.



**SECTION 7: PARTICIPANT INFORMATION AND CONSENT (See *National Statement* Chapter 2.2)**

**7.1 Will written consent be obtained?**

☒ Yes ☐ No

If you answered 'Yes', attach a copy of the Participant Information and Consent Form. (Note that two (2) copies of the consent form must be signed, with the participant and the investigator each retaining one (1) of the copies.)

If you answered 'No', give reasons.

(See *National Statement* Chapter 2.3 - Qualifying or waiving conditions for consent)

**7.2 Will the consent of minors participating in the research be obtained?**

☐ Yes ☐ No ☒ N/A

(See *National Statement* 4.2.7 – Children and young people)

As well as parental consent when required by law (less than 16 years), any teenager must be allowed to give their own consent/assent when taking part in research. Young children who are involved as research participants should be offered the option of assenting to the research.

In the case where a research project involves children, it is necessary to provide parental/guardian information and consent forms. A space may be included for the child to sign their consent/assent on the parental/guardian consent form.

Depending on the age of the child involved, it may be appropriate to also provide separate information and consent forms for children and for adults.

Note: For additional guidance on this matter, researchers may wish to consult Understanding Consent in Research Involving Children (Spriggs 2010).

If you answered 'Yes', describe the protocol to be used for obtaining consent and explain how minors will be able to indicate a desire to withdraw from the research.

If you answered 'No', give reasons.

**7.3 In the case of participants for whom competence in English is not adequate for informed consent, what arrangements have been made to ensure comprehension of the Participant Information and Consent Form? (See *National Statement* 1.4 - Justice)**

n/a

**7.4 Please indicate whether the following details have been provided in your Participant Information and Consent Form.**

	Yes	No	N/A
i. A short title for the project	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
ii. A brief statement of the aims of the research	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
iii. The names of the researchers, their Department affiliations, contact telephone numbers and email addresses. If you intend to provide participants with your home phone number, please indicate why this is necessary in the space provided in 7.4 (b) below.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
iv. Acknowledgment that the research is "being conducted to meet the requirements for the degree of (name of degree) under the supervision of (Supervisor's name, contact telephone number, email address and Department affiliation)"	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
v. An explanation of what each participant is expected to do and an estimate of the time commitment involved	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
vi. An acknowledgment of any recording using audio-recording, video-recording, or photographs and explanation of how this material will be used	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
vii. An objective statement of any risks or discomforts	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
viii. Any payment of money or other remuneration, e.g. course credits	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
ix. Information about how confidentiality of the data will be maintained, i.e. how privacy will be maintained, who, if anyone, will have access to the data other than the researchers/supervisor and for what purpose, the form in which the data will be published.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
x. Any plans to make the data available in non-identifiable form to other researchers (other than those listed on this application form) in the future.	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
xi. Information about how participants can obtain feedback regarding the results of the research.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
xii. An acknowledgement of any sources of funding for the research, including commercial or other sponsors.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
xiii. A statement indicating that participation is voluntary and guaranteeing participants the right to withdraw from further participation at any time without having to give a reason and without adverse consequence <i>NOTE 1. The statement about withdrawal is N/A for studies using anonymous questionnaires.</i> <i>NOTE 2. MQ students who receive course credits for their participation must be assured in writing in the consent form that they will not forfeit their course credits if they choose to withdraw from the research</i>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
xiv. A signed statement of agreement to participate in the research, e.g. "I agree to participate in this research," with the consent form signed and dated by the participant and signed and dated by the investigator or other witness.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
xv. A statement indicating that the participant has been given a signed copy of the consent form to keep.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
xvi. A footnote regarding complaint procedures as follows: <i>The ethical aspects of this study have been approved by the Macquarie University Human Research Ethics Committee. If you have any complaints or reservations about any ethical aspect of your participation in this research, you may contact the Committee through the Director, Research Ethics (telephone [02] 9850 7854, email: <a href="mailto:ethics@mq.edu.au">ethics@mq.edu.au</a>). Any complaint you make will be treated in confidence and</i>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

investigated, and you will be informed of the outcome

(a) If you answered 'No' to any of the above please provide details below:

--

(b) If you intend to provide participants with a home phone number, please provide your explanation below:

--

**SECTION 8: POSSIBLE CONFLICT OF INTEREST (See National Statement Chapter 5.4)**

**8.1 In undertaking this research do any "conflict of interest" issues arise?** ☒ Yes ☐ No

*(For instance, is the researcher in a dual role such as researcher/teacher, researcher/carer, researcher/employer? Does the researcher have access to personal files/databases as a condition of employment rather than as a researcher?).*

If you answered 'Yes', please provide details of the conflict of interest and considerations or mechanisms in place to address these issues.

As noted above, some of the investigators named above will be teaching units in the departments of Psychology or Cognitive Science, and students enrolled in these units may choose to participate, either for payment or for course credit.

Some Macquarie University members may also choose to participate in this research.

For these reasons, it will be emphasised on the consent forms associated with this project that participants are free to withdraw their participation at any time without consequence. Furthermore, all signups will be managed through the SONA system rather than directly through the researchers to mitigate any perceived coercion

**8.2 Will this research be undertaken on behalf of (or at the request of) commercial entity, or any other sponsor?** ☐ Yes ☒ No

If you answered 'Yes', who is the commercial entity or sponsor?

--

**8.3 Do the researchers have any affiliation with or financial involvement in any organisation or entity with direct or indirect interests in the subject matter or materials of this research?** ☐ Yes ☒ No

If you answered 'Yes', please provide details.

--

**8.4 Do the researchers expect to obtain any direct or indirect financial or other benefits from conducting this project?** ☐ Yes ☒ No

(Note that such benefits should be included in the Participant Information and Consent Form).

If you answered 'Yes', please provide details.

--

**8.5 Have conditions been imposed upon the use, publication or ownership of the results including the review of data, manuscript draft or scientific presentation by any other party than the listed researchers?**

☐ Yes

☒ No

(Note: The Committee is unlikely to approve arrangements that involve the censorship of *research findings in publications*.)

If you answered 'Yes', please provide details.

--

- (a) Where the project is a collaborative enterprise with any organisations or individuals including contracted research activities undertaken through one of the University's corporate entities (eg. Access Macquarie, MGSM Pty Ltd), please identify which organisation and individual will have the primary responsibility for the project:

<b>Organisation:</b>	
<b>Individual:</b>	

If this is a contracted research activity, please attach a copy of the contract to your ethics application

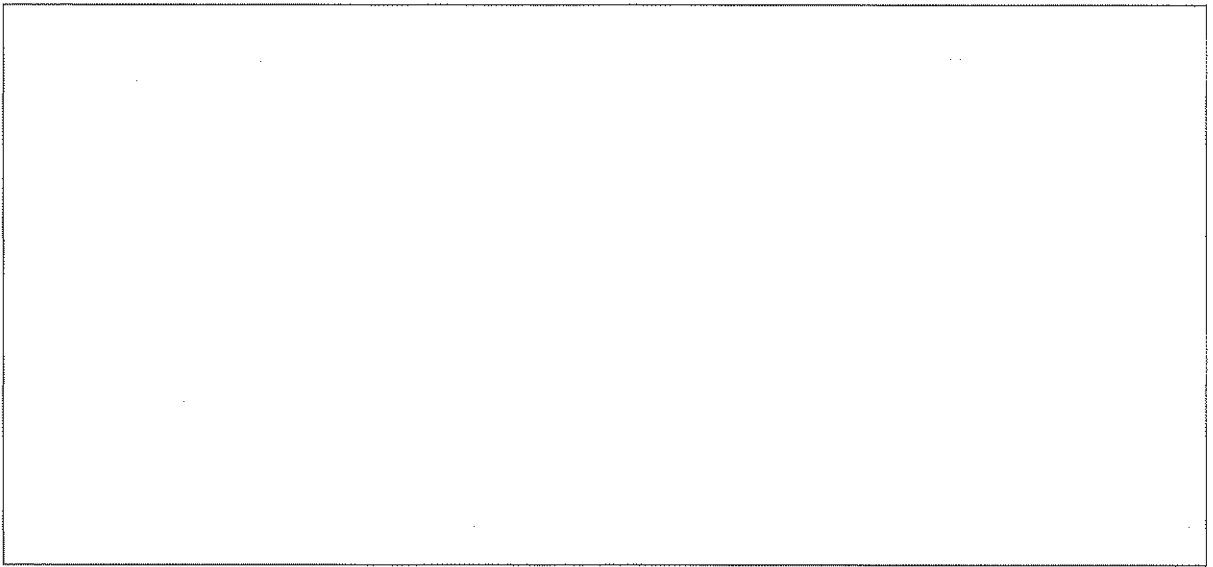
## SECTION 9: OTHER ETHICAL CONSIDERATIONS

**9.1 Are there any further ethical considerations that you wish to raise?**

☐ Yes

☒ No

If you answered 'Yes', detail what these considerations are.



**SECTION 10: SUPPORT DOCUMENTATION****10.1 Have you included the following support documents with your application?**

	Yes	No	N/A
a. copies of any correspondence with other institutions or human research ethics committees	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
b. a copy of the contract if this is contracted research (refer Q1.7 (e))	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
c. copies of relevant pages of your grant application (refer Q1.7 (a-d))	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
d. copies of Information and Consent Forms (Note that the consent form is to be printed on your respective Faculty/Department letterhead).	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
e. copies of any interviews, questionnaires, or surveys to be used	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
f. copies of any participant recruitment advertisements	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
g. Research involving Aboriginal and Torres Strait Islander people (APPENDIX A)	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
h. Research to be undertaken outside Australia (APPENDIX B)	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
i. Privacy and access to personal information (APPENDIX D)	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>

**PLEASE ONLY SUBMIT THE APPENDICES THAT ARE RELEVANT TO YOUR RESEARCH. IF IN DOUBT CONTACT THE ETHICS OFFICER, on [ethics.secretariat@mq.edu.au](mailto:ethics.secretariat@mq.edu.au) or phone (02) 9850 6848.**

<b>SECTION 11: CERTIFICATION</b>
----------------------------------

**11.1 To the best of my belief the proposed project conforms in all respects with the *National Statement on Ethical Conduct in Human Research*.**

Information about and links to legislation, guidelines and codes governing research with humans is available at:

[http://www.research.mq.edu.au/for/researchers/how to obtain ethics approval/human research ethics/policy](http://www.research.mq.edu.au/for/researchers/how_to_obtain_ethics_approval/human_research_ethics/policy)

Signature of Chief Investigator:

Name of Chief Investigator (block letters): THOMAS CARLSON

Date:



23 December 2013

Dr Thomas Carlson  
Department of Cognitive Science  
Faculty of Human Sciences  
Macquarie University

Dear Dr Carlson

RE: *Decoding the neural representations of objects in the human brain*

Thank you for the email dated 09 December 2013 responding to the issues raised by the Macquarie University Human Research Ethics Committee (HREC (Medical Sciences)).

On 16 December 2013 the Committee Executive considered your correspondence and approved your application. This research meets the requirements set out in the *National Statement on Ethical Conduct in Human Research* (2007).

**Details of this approval are as follows:**

**Reference No:** 5201300804

**Approval Date:** 16 December 2013

This letter constitutes ethical approval only.

The following documentation has been reviewed and approved by the HREC (Medical Sciences):

Documents reviewed	Version no.	Date
Macquarie University (MQ) HREC Application Form	2.3	July 2013
Correspondence from Dr Lincoln Colling addressing the HREC's feedback	No Version	09/12/2013
MQ Participant Information and Consent Form (PICF) – Behavioural Only	1.1	09/12/2013
MQ PICF – EEG	1.1	09/12/2013
MQ PICF – MRI	1.1	09/12/2013
MQ PICF – MEG	1.1	09/12/2013
MQ PICF– Combined MEG/EEG	1.1	09/12/2013
SONA Advertisements – Behavioural only, EEG, MEG /Combined MEG/EEG, MRI	1.1	09/12/2013



**Standard Conditions of Approval:**

1. Continuing compliance with the requirements of the *National Statement*, which is available at the following website:

<http://www.nhmrc.gov.au/book/national-statement-ethical-conduct-human-research>

2. Approval is for five (5) years, subject to the submission of annual reports.

**First Annual Report Due:** 1 December 2014

3. All adverse events must be reported to the HREC within 72 hours.

4. Proposed changes to the protocol must be submitted to the Committee for approval before implementation.

It is the responsibility of the Chief investigator to retain a copy of all documentation related to this project and to forward a copy of this approval letter to all personnel listed on the project.

Please do not hesitate to contact the Ethics Secretariat should you have any questions regarding your ethics application.

The HREC (Medical Sciences) wishes you every success in your research.

Yours sincerely

**Dr Karolyn White**

Director, Research Ethics

Chair, Human Research Ethics Committee (Medical Sciences)

This HREC is constituted and operates in accordance with the National Health and Medical Research Council's (NHMRC) National Statement on Ethical Conduct in Human Research (2007) (the National Statement) and the CPMP/ICH Note for Guidance on Good Clinical Practice.


## APPENDIX TWO

---

---

PUBLISHED MANUSCRIPT

---



# Decoding the time-course of object recognition in the human brain: From visual features to categorical decisions

Erika W. Contini<sup>a,b,\*</sup>, Susan G. Wardle<sup>a,b,1</sup>, Thomas A. Carlson<sup>a,b,c,1</sup>

<sup>a</sup> Department of Cognitive Science, Macquarie University, Sydney, Australia

<sup>b</sup> ARC Centre of Excellence in Cognition and its Disorders and Perception in Action Research Centre, Macquarie University, Australia

<sup>c</sup> School of Psychology, University of Sydney, Australia

## ARTICLE INFO

### Keywords:

MEG  
EEG  
MVPA  
Time-series decoding  
Object recognition  
Object categorisation

## ABSTRACT

Visual object recognition is a complex, dynamic process. Multivariate pattern analysis methods, such as decoding, have begun to reveal how the brain processes complex visual information. Recently, temporal decoding methods for EEG and MEG have offered the potential to evaluate the temporal dynamics of object recognition. Here we review the contribution of M/EEG time-series decoding methods to understanding visual object recognition in the human brain. Consistent with the current understanding of the visual processing hierarchy, low-level visual features dominate decodable object representations early in the time-course, with more abstract representations related to object category emerging later. A key finding is that the time-course of object processing is highly dynamic and rapidly evolving, with limited temporal generalisation of decodable information. Several studies have examined the emergence of object category structure, and we consider to what degree category decoding can be explained by sensitivity to low-level visual features. Finally, we evaluate recent work attempting to link human behaviour to the neural time-course of object processing.

## 1. Introduction

Visual object recognition is a complex problem. In everyday life we experience an overwhelming number of objects that the brain needs to rapidly differentiate and identify. How is it that we are able to identify a chair, for example, despite large variability in lighting, colour, design, materials and viewpoint? The human brain does a remarkable job of efficiently solving this problem, and has inspired decades of behavioural, neuroscience, and computer science research. The ventral visual processing stream, which involves a number of regions throughout the occipito-temporal cortex, is well-established as the neural pathway for object recognition (e.g., Grill-Spector et al., 2001; Haxby et al., 2001; Ishai et al., 1999). Functional neuroimaging methods have identified a number of object-selective areas in human visual cortex that are preferentially activated by specific object categories including animals (Chao et al., 1999; Martin et al., 1996), faces (Ishai et al., 1999; Kanwisher et al., 1997), bodies (Downing et al., 2001), places (Epstein and Kanwisher, 1998), and tools (Beauchamp et al., 2002; Chao et al., 1999; Martin et al., 1996). However, the number of object-selective regions identified to date accounts for only a handful of the numerous object categories that exist (Biederman, 1987). This highlights that while identification of these regions has advanced our understanding of

the neural pathways involved in higher-level visual processing, we are yet to fully understand *how* the brain solves the many challenges associated with object recognition.

Recent studies in visual object recognition have been influenced by advances in neuroimaging analysis methods, which allow for a fundamental change in the type of information that can be extracted from neuroimaging data. Early neuroimaging studies employed univariate analysis techniques that focus on identifying differences in the average activation of individual voxels or sensors for different experimental conditions. In contrast, more recent multivariate analysis techniques (e.g. multivariate pattern analysis; MVPA) analyse patterns of activation associated with experimental conditions from multiple voxels/sensors simultaneously. In the field of neuroimaging, multivariate methods have the potential to detect differences in activation which are lost when averaging data for univariate analyses, making them more sensitive (Carlson et al., 2003; Cox and Savoy, 2003; Grootswagers et al., 2016; Haxby et al., 2001; Haynes and Rees, 2006). While univariate and multivariate analyses are complementary in the information they provide, there is a fundamental difference in the types of experimental questions these methods address. Univariate magneto/electro-encephalography (M/EEG) analyses evaluate differences in *activation*, quantifying relative differences in average activity

\* Corresponding author at: Department of Cognitive Science, Macquarie University, Sydney, Australia.

E-mail address: [erika.contini@students.mq.edu.au](mailto:erika.contini@students.mq.edu.au) (E.W. Contini).

<sup>1</sup> shared senior authorship.

between experimental conditions, while multivariate methods have the potential to examine differences in *information*, for example by comparing differences in distributed patterns of brain activation between experimental conditions (Grootswagers et al., 2016; but see also de-Wit et al., 2016).

Decoding methods are a form of MVPA that have provided important insights into how the brain processes information (Grootswagers et al., 2016; Haxby et al., 2014; Haynes, 2015; Pereira et al., 2009). In the field of object recognition, decoding methods aim to map differences in complex neural activity patterns associated with perceiving objects. Early studies adopting these methods marked an important first step in our understanding of *how* objects are processed in the brain (Carlson et al., 2011; Carlson et al., 2003; Clarke and Tyler, 2014; Cox and Savoy, 2003; Haxby et al., 2001; Kriegeskorte et al., 2008; Liu et al., 2009; O'Toole et al., 2005). A particular strength of applying decoding methods to M/EEG data with the aim of understanding visual object recognition is that it can reveal how visual object representations change over time with high temporal resolution. The focus of this review is on what has been learned about visual object processing in the human brain with the application of these recent, powerful temporal decoding methods.

To date decoding methods have been used in conjunction with a variety of neuroimaging and neurophysiology approaches to investigate the neural mechanisms underlying object recognition. For example, brain decoding using single cell recordings has revealed category structure within monkey inferior temporal cortex (IT) (Hung et al., 2005), with differentiable neural patterns associated with animate and inanimate objects, as well as more specific animate subcategories, including human and animal faces and bodies (Kiani et al., 2007). Decoding of fMRI data has shown commonalities between object response patterns in both human and monkey brains, with similar differentiation of animate/inanimate categories and face/body subcategories identified in both species (Kriegeskorte et al., 2008). Categorical representations are not limited to localised brain regions, as activity evoked by objects and faces often overlaps between the categories and is distributed throughout ventral temporal cortex (Haxby et al., 2001). Moreover, there has been a recent move towards looking beyond representing object categories in terms of dichotomies, such as the animate/inanimate distinction. Instead, a continuum has been proposed that spans from inanimate objects to humans, with objects categorised in terms of their biological similarity to humans (Connolly et al., 2012; Sha et al., 2015). Together, these studies emphasise a role for distributed patterns of activity in the neural representation of object category, building on the original observation of overlapping activation patterns in the human ventral stream (Haxby et al., 2001).

Decoding of time-series data using M/EEG has offered the potential to examine the time-course of object representations in the human brain, revealing a dynamic evolution of object category structure over time (e.g., Barragan-Jason et al., 2015; Carlson et al., 2011; Carlson et al., 2013; Cauchoix et al., 2014; Cichy et al., 2014; Clarke et al., 2014; Goddard et al., 2016; Kaiser et al., 2016a, 2016b; Simanova et al., 2010). As visual information moves through the ventral pathway, the content of visual representations changes rapidly between brain regions. fMRI decoding studies have revealed much about the representation of objects in the human brain, however the coarse temporal resolution of fMRI limits the examination of dynamic visual processes. As a complement to the static snapshot of representational structure revealed with fMRI, neuroimaging techniques with higher temporal resolution such as M/EEG facilitate investigation of the dynamic processes of visual object recognition. The aim of this review is to discuss how time-series decoding studies have advanced our understanding of the complexities of visual object recognition by focusing on the dynamic processes involved. Time-series decoding is a relatively new approach to studying object recognition, and here we highlight the potential of this new direction to inform the field. Note that we focus on

what has been learned about object processing by applying time-series decoding methods, for a more thorough discussion of the technical details of time-series decoding analyses see Grootswagers et al. (2016).

## 2. Decoding the object recognition time-course

### 2.1. Initial insights into time-series decoding

It is well-established that the process of visual object recognition requires a number of hierarchically organised stages that progress through the occipito-temporal pathway (Grill-Spector and Malach, 2004; Malach et al., 2002). Early retinotopic visual areas are more sensitive to changes in low-level stimulus properties, while higher cortical areas within the ventral temporal lobe produce more complex responses to whole objects, and appear to account for more abstract properties such as object category (Altmann et al., 2003; Grill-Spector and Malach, 2004; Van Essen et al., 1992). Single-unit recordings in macaques have shown that these different processing stages progress successively in time as information passes through the occipito-temporal pathway (Schmolesky et al., 1998). With the development of M/EEG techniques, we are able to evaluate the temporal dynamics of visual object recognition in the human brain with millisecond resolution, allowing us to delve into more specific and fine-grained processes occurring in the various stages of visual object processing.

Differences in early versus late stages in object processing have been examined by comparing MEG and fMRI data for the same stimulus set. By linking both temporal and spatial neuroimaging data, Cichy et al. (2014) showed that activity early in the MEG time-course correlated more strongly with fMRI activity in V1, while later MEG activity was more strongly associated with activity in IT. The stimulus set were 96 colour images of animate and inanimate objects used in previous studies (Kiani et al., 2007; Kriegeskorte et al., 2008). The activity patterns associated with viewing each individual object were first compared using representational similarity analysis (RSA; Kriegeskorte, 2008), where a matrix is created based on the difference in brain activation patterns for every pairwise comparison of object images. This was done separately for the fMRI and MEG data. These 'dissimilarity matrices' provide an index of the difference in the brain response between object representations. The dissimilarity matrices were then compared across imaging modalities by examining when the relative similarity between the activation patterns for each object pair in the fMRI data most closely resembled that in the MEG data. The finding that the representational structure early in the MEG data more closely resembles V1 activity while later MEG data is closer to the structure observed in IT with fMRI is consistent with the known features of the visual processing hierarchy, and thus provides a source of validation for MEG decoding methods.

Time-series decoding methods have also expanded our understanding of the temporal intricacies associated with processing low-level stimulus properties. For example, Goddard et al. (2016) investigated temporal differences in the contribution of low and high spatial frequencies to the representation of object identity. Stimuli were greyscale images of objects that were matched in their amplitude spectrum to control for low-level visual properties. Phase randomisation was applied selectively to spatial frequency bands such that object identity information was confined to a restricted spatial frequency band in each stimulus. Decoding analysis revealed that the activation pattern of the whole-brain MEG recordings contained information related to object identity, but there were critical differences in the processing of low versus high spatial frequencies. They found that low spatial frequencies provided object identity information earlier and in more occipitally located regions than high spatial frequency information, which supported decoding of object identity later in the time-course. Further, by applying Granger causality analysis, they found evidence for both an early feedforward and later feedback flow of information related to object identity (Goddard et al., 2016).

Most of the existing temporal decoding literature on object recognition has focused on decoding object categories and investigating the representational structure of object representations. These results are discussed in detail in the following sections.

## 2.2. Temporal decoding of high-dimensional distributed category representations

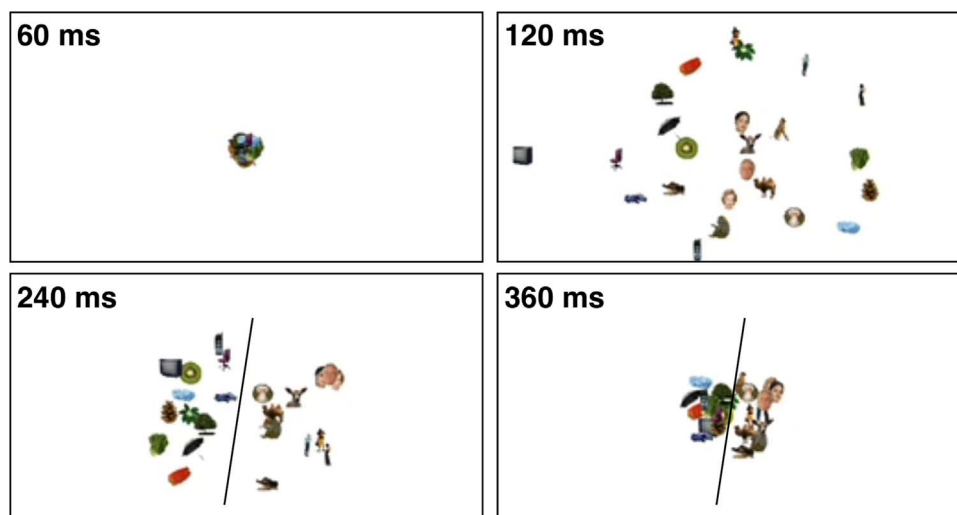
Understanding population coding is one of the overarching aims of neuroscience (Averbeck et al., 2006). A significant and influential recent development in understanding the neural mechanisms underlying object representation has been the theoretical shift away from identifying specific localised regions associated with particular object categories towards focusing on more distributed activity patterns that extend throughout higher-level visual cortex (Haxby et al., 2001). A current focus in object recognition is on interpreting the highly multidimensional activation patterns elicited by object stimuli. Specifically, one of the main themes centres on understanding the structure of object category representations. In this section we outline what temporal decoding methods have so far revealed about the neural architecture of object categories.

A useful and intuitive way to think about the highly multidimensional activation patterns evoked by viewing objects is to conceptualise them in terms of an abstract representational space (DiCarlo and Cox, 2007; Kriegeskorte and Kievit, 2013). The complex activation pattern across M/EEG sensors elicited by viewing a particular object exemplar (e.g., shoe, horse, face) can be considered as a single point in an abstract representational space. Thus in this abstract representation, the proximity between data points indicates the degree of similarity, such that a greater distance between object exemplars represents more disparate neural activation patterns, while exemplars with closer points have more similar activation patterns. Consequently, object exemplars further away from each other in representational space are easier to "decode" from each other using machine learning classification than objects which are closer together, as they share more similar activation patterns and are less separable in higher-dimensional space.

Multidimensional scaling (MDS) is a technique that can be used as a tool for visualising the representational space of objects in the brain. An example of an MDS plot for object representations is shown in Fig. 1. Carlson et al. (2013) studied the first 1000 ms of the visual

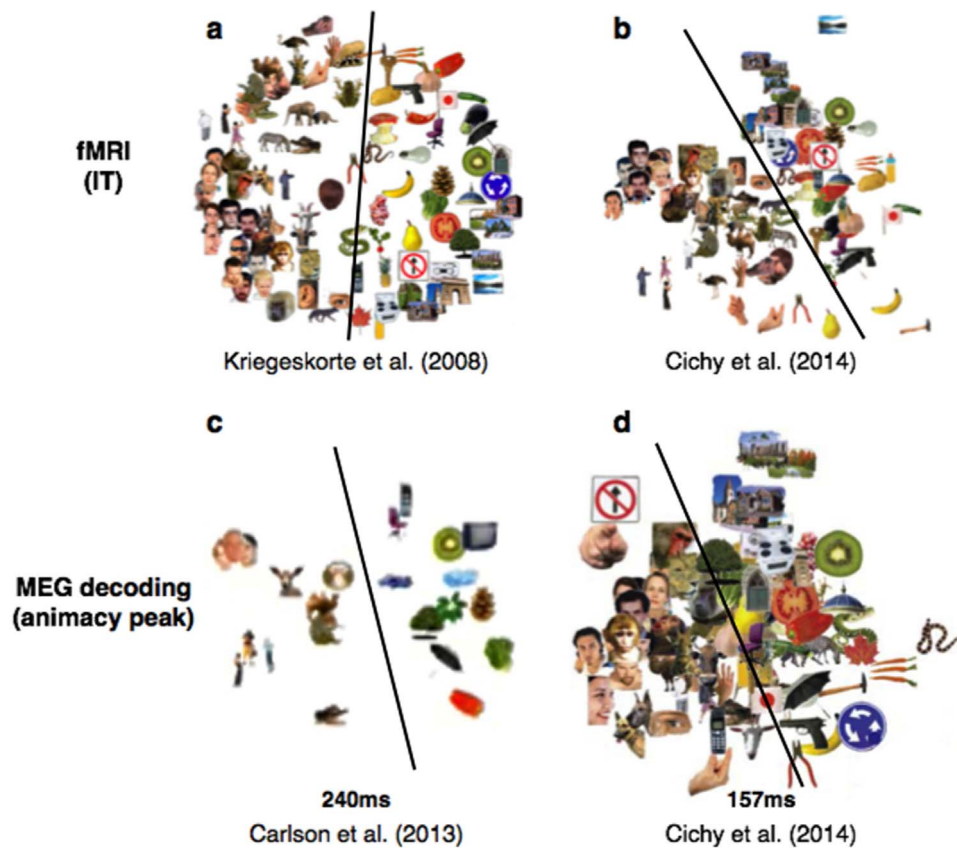
object recognition time course using MEG decoding. MEG recordings were acquired as participants viewed a series of single object images while performing an unrelated attention task (reporting whether a letter superimposed on top of the object image was a vowel or a consonant). Using MDS to visualise the differences in object representations as a function of time, we can appreciate the emergence of category structure throughout the time course (Fig. 1). Early in the time-course (~60 ms post-stimulus onset), the representations of individual exemplars are almost wholly overlapping in the representational space, reflecting poor decodability of the individual objects based on the MEG data. This is expected given the time it takes for an image on the retina to transition to a cortical representation that is accessible to MEG (see Nowak and Bullier, 1997 for a review). By ~120 ms, differences between individual exemplars have emerged such that the exemplars are spread out in the representational space. This is the time of peak decoding for this image set; the point at which individual exemplars are most easily distinguished from one another. As we progress through the time course, category structure begins to emerge. From 120 ms, some subcategories appear to start to cluster: note the grouping of faces and animals in the centre. By 240 ms a clear categorical distinction between animate and inanimate objects is apparent (diagonal black lines in Fig. 1 indicate the category boundary). Interestingly, this animacy category distinction is maintained at 360 ms, despite much less differentiation of individual object exemplars, which are clustered tightly together in the later stages of the time-course. This is an example of the advantage of time-series decoding, as the emergence of categorical clustering over time is not captured by other neuroimaging methods with lower temporal resolution such as fMRI.

As MDS has been used in several fMRI and MEG studies with the same object stimuli (Carlson et al., 2013; Cichy et al., 2014; Kriegeskorte et al., 2008) it is possible to visualise the reliability and consistency of object representations across neuroimaging modalities by comparing across studies - a rare opportunity in neuroimaging (Fig. 2). All three studies used the object set (or subset of the set) originally used by Kiani et al. (2007) in monkey IT. Kriegeskorte et al. (2008) and Cichy et al. (2014) both used fMRI to evaluate the representational geometry of objects in human IT (Fig. 2a and b). The animate/inanimate divide is clearly visible in IT, and additional subcategory groupings such as animals and humans also cluster



**Fig. 1.** MDS plots demonstrating the evolution of object representations over time with MEG. Individual MDS panels show the representational geometry of a set of objects at a particular time-point (shown as ms post-stimulus onset). Distances between objects represent the level of similarity between the neural patterns measured with MEG such that larger distances indicate greater dissimilarity (i.e., more distinct neural patterns). The representational geometry evolves from initially entirely overlapping (poorly discriminated) objects, to maximal differentiation at 120 ms (peak decoding for this sample). The animate/inanimate category distinction (black dividing line represents the boundary) peaks at 240 ms and is maintained even at 360 ms, although differences between individual object exemplars become less defined. Images adapted with permission of the Association for Research in Vision and Ophthalmology, from Carlson et al. (2013); permission conveyed through Copyright Clearance Center, Inc.





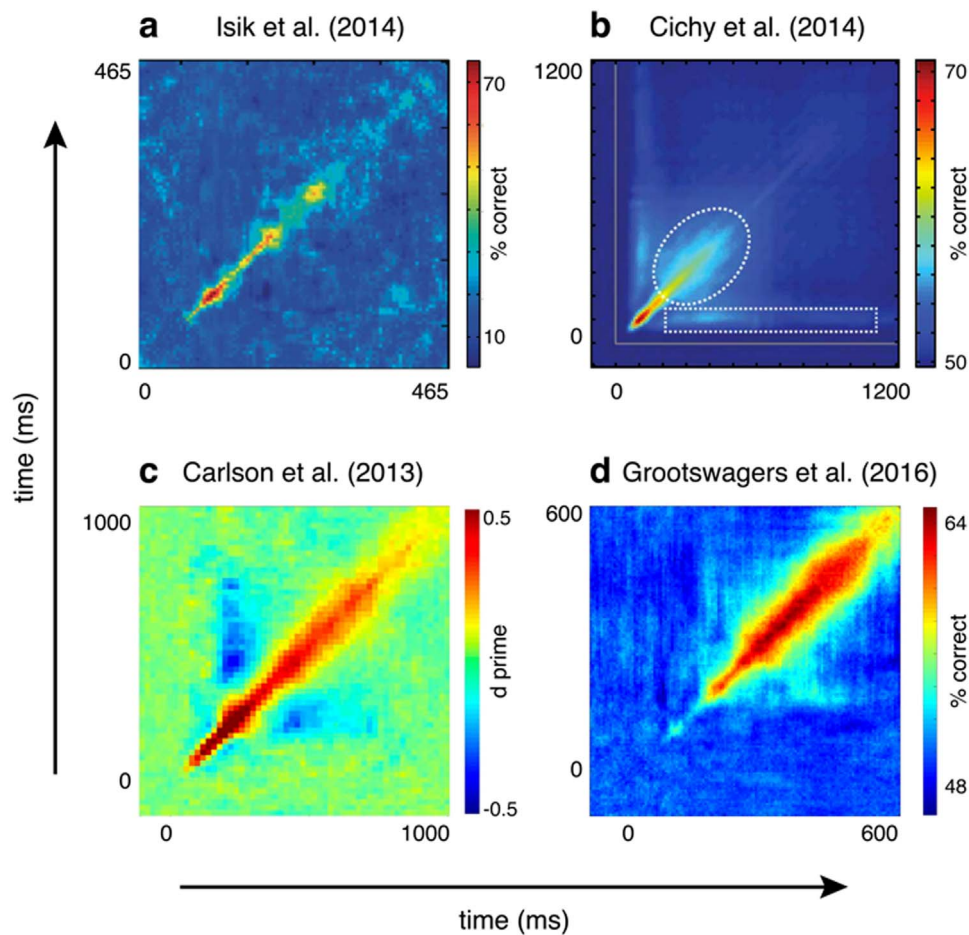
**Fig. 2.** Comparison of MDS plots from fMRI and MEG studies. Top MDS plots (a and b) show the representational geometry in human IT using fMRI; bottom MDS plots (c and d) show the representational geometry using MEG at the time of peak decoding for the animate/inanimate (animacy) category distinction. Animacy peak decoding times for the individual studies are listed below the MEG MDS plots. Note that (a) has been normalised and rigidly aligned (Procrustes alignment) for visualisation, producing a more regular spacing between object exemplars in the MDS plot compared to panels b–d. Black lines in each plot mark the (approximate) animacy boundary. Panel (a) image from [Kriegeskorte et al. \(2008\)](#), reprinted with permission from Elsevier. Panel (c) image from [Carlson et al. \(2013\)](#) reprinted with permission of the Association for Research in Vision and Ophthalmology; permission conveyed through Copyright Clearance Center, Inc. Panels (b) and (d) images from [Cichy et al. \(2014\)](#) reprinted with permission from Macmillan Publishers Ltd: [Nature Neuroscience](#), copyright (2014).

together. These fMRI MDS plots can be compared to those from MEG time-series decoding studies ([Fig. 2c](#) and [d](#); [Carlson et al., 2013](#); [Cichy et al., 2014](#)). MEG data is shown at the time of peak animacy decoding in each study: 240 ms ([Fig. 2c](#)) and 157 ms ([Fig. 2d](#)). Note that the time of peak decoding is dependent on the stimulus set being evaluated as well as the noise and variability in a particular data set, complicating the comparison of specific decoding onset times across studies. The emergence of feature information or category structure is better discussed in terms of relative timing within individual studies. A critical difference between the MDS plots for fMRI versus MEG is that the MEG plots represent similarity in whole-brain activation patterns, whereas the fMRI analysis is localised to IT. Considering this substantial difference, the degree of similarity in the categorical representation is remarkable (compare top and bottom rows of [Fig. 2](#)).

Another advantage of using time-series decoding to investigate object processing is that we can examine to what degree information related to object category is maintained in the brain activation patterns over time (see [King and Dehaene, 2014](#) for a review of temporal generalisation methods). This is achieved by performing time-point by time-point comparisons across the entire time-course, whereby a classifier is trained on data from one time-point and then tested at every other time point. If information is maintained within the signal for some period of time, then the classifier is expected to generalise and successfully classify the test data, regardless of the time difference between the training and test data sets. The results of temporal generalisation are visualised in two-dimensional heat maps indicating the magnitude of classification performance for each pair of time points ([Fig. 3](#)).

A number of MEG studies have examined the temporal generalisation of decoding to evaluate the dynamics of object representations ([Fig. 3](#)). For each plot in [Fig. 3](#), the diagonal represents standard decoding analysis when the classifier is trained and tested on data from the same time-point, thus classification accuracy is expected to be highest along the diagonal. If decoding performance is significant off the diagonal, it is indicative of some degree of temporal generalisation of the decodable signal related to object classification. Note that the decoded category differs across studies: [Isik et al. \(2014\)](#) and [Cichy et al. \(2014\)](#) decoded object exemplars, while [Carlson et al. \(2013\)](#) and [Grootswagers et al. \(2016\)](#) decoded object animacy.

The most striking similarity between the temporal generalisation plots from these different studies is that the MEG signal underlying object decoding evolves relatively quickly, with little information generalising across time. Where there is generalisation, it tends to cluster closely around the diagonal, suggesting that the structure of object representations accessible in the whole-brain MEG signal evolves rapidly, and follows a specific neural trajectory. Another interesting feature is that the period where there is the greatest generalisation occurs relatively late after stimulus onset (see particularly [Fig. 3c–d](#)). This suggests maintenance of object representations at later stages of processing ([Carlson et al., 2013](#)). Alternatively, it may reflect the gradual accumulation of multiple related information processing stages of variable duration throughout the visual processing hierarchy, leading to a greater temporal spread of information over time. Paradoxically, there are also some periods of below chance decoding, whereby a classifier systematically categorises an object incorrectly as the opposite category (e.g. see blue regions in [Fig. 3c](#)).



**Fig. 3.** Examples of temporal cross-decoding of MEG data. Temporal generalisation of MEG decoding from (a) Isik et al. (2014), (b) Cichy et al. (2014) (c) Carlson et al. (2013), and (d) Grootswagers et al. (2016). Generalisation of decoding performance across time is assessed by training and testing the classifier on MEG data from different time points. Points on the diagonal represent training and testing on data from the same time point (i.e., regular classification). Colour indicates classifier performance as a function of time for each study; hotter colours indicate higher levels of classifier performance (measured using decoding accuracy percent (a, b, d) and d-prime in (c)). Panel (a) image from Isik et al. (2014) reprinted with permission from ©The American Physiological Society. Panel (b) image from Cichy et al. (2014) reprinted with permission from Macmillan Publishers Ltd: *Nature Neuroscience*, copyright (2014). Panel (c) image from Carlson et al. (2013) reprinted with permission of the Association for Research in Vision and Ophthalmology; permission conveyed through Copyright Clearance Center, Inc. Panel (d) image from Grootswagers et al. (2016) reprinted with permission from the Massachusetts Institute of Technology © 2016, published by the MIT Press.

Carlson et al. (2013) suggest this may represent adaptation or inhibition following a period of excitation, resulting in later inversion of the neural representation. Consistent with this explanation, anticorrelated MEG signals have also been identified around the time of stimulus offset for both visual (Carlson et al., 2011) and auditory (Chait et al., 2007) stimuli.


An exception to the general lack of generalisation is seen in the results of Cichy et al. (2014), where there is a period of significant generalisation of early information across most of the time-course (Fig. 3b; grey dotted rectangle shows the period of significant temporal generalisation). This could be because this analysis used pair-wise object decoding as opposed to leave-one-exemplar-out category classification (see Carlson et al., 2013; Grootswagers et al., 2016). With pairwise object decoding, low-level information diagnostic of object is contained in both the training and test sets for the classifier, thus maintenance of early visual information is consistent with low-level properties being useful for classification. There is a similar pattern of generalisation in Fig. 3c, and this analysis also included exemplars in both training and test sets (Carlson et al., 2013). However, in leave-one-out exemplar decoding (Fig. 3d) the test exemplar is not included in the training set so early processing reflecting sensitivity to low level image properties is less likely to be used by the classifier. Kaiser et al. (2016a) took this idea further by using temporal generalisation to directly examine decoding of low-level shape versus object category in a

matched stimulus set (see Section 2.5, Fig. 6).

The studies in this section demonstrate the utility of examining the change in representational structure of object representations over time. Studies to date have shown that the brain's response to visual objects is highly dynamic, evolving rapidly from sensitivity to low-level visual properties to more category-like representations. There is also evidence for a robust categorisation distinction between animate and inanimate objects that is sustained. In the following section we examine what has been learnt from time-series decoding about the hierarchical structure of object category representations.

### 2.3. Uncovering the hierarchy of object category representations

Visual object categorisation can be understood as a dynamic process of evidence accumulation over time (Mack and Palmeri, 2011; Nosofsky and Palmeri, 1997; Philastides and Sajda, 2006). Functionally, the accumulation of evidence for object category membership is likely to exploit the complex network of feedback and feedforward connections within the object-selective ventral pathway, rather than operating as a linear progression of representation from low-level features through to semantic concepts of increasing abstraction (Kravitz et al., 2013). This is a recent development and extension of the earlier idea of sequential processing stages, where objects are first categorised at an intermediate/basic level (Mervis and Rosch,



**Fig. 4.** Hierarchical category structure of the object stimuli used in both Carlson et al. (2013) and Cichy et al. (2014). Diagram shows the three levels of category abstraction: lower tier (animal and human faces and bodies), intermediate tier (bodies and faces), and highest tier (animacy). Note, number of stimuli differed between the two studies: Carlson et al. (2013) used 24 images (as shown), while Cichy et al. (2014) employed 92. Image reprinted with permission of the Association for Research in Vision and Ophthalmology, from Carlson et al. (2013); permission conveyed through Copyright Clearance Center, Inc.

1981) (e.g., cat), with superordinate (e.g., animal) and subordinate (e.g., Siamese cat) categorisation occurring later in the visual processing hierarchy (Mack and Palmeri, 2011).

MEG decoding methods have been applied to understanding the timing of category abstraction by selecting stimuli with a planned hierarchical category structure and assessing when these category level distinctions emerge. Results from two MEG decoding studies (Carlson et al., 2013; Cichy et al., 2014) provide insights into the category processing hierarchy by evaluating the time course of hierarchically organised object categories. The stimulus sets across both studies followed the same planned hierarchical structure (see Fig. 4): the highest (most general/abstract) category tier was the commonly evaluated animate/inanimate distinction. The animate domain was then further subdivided into intermediate categories of faces/bodies as well as humans/animals, with the lowest (most specific) category level being that of human faces/human bodies, animal faces/animal bodies. For the inanimate domain, these objects could be further subcategorised into man-made and natural objects.

To investigate the dynamics of hierarchical object category representations, these studies looked at two critical moments in the MEG object decoding time-course: the time at which decoding first reaches significance (decoding onset) and the time that the category distinction is maximally differentiated (peak decoding time). A visual comparison of the decoding time course in both Carlson et al. (2013) and Cichy et al. (2014) for object categories at different levels of the category hierarchy is shown in Fig. 5. Carlson et al. (2013) found the onset of significant decoding for individual object exemplars occurred around 80 ms after stimulus onset, which was also similar to the decoding onsets of individual categories; all falling between 80–100 ms. Cichy et al. (2014) similarly found little variability between overall exemplar decoding onset (~48 ms) and individual category onsets (~51–61 ms post-stimulus onset).

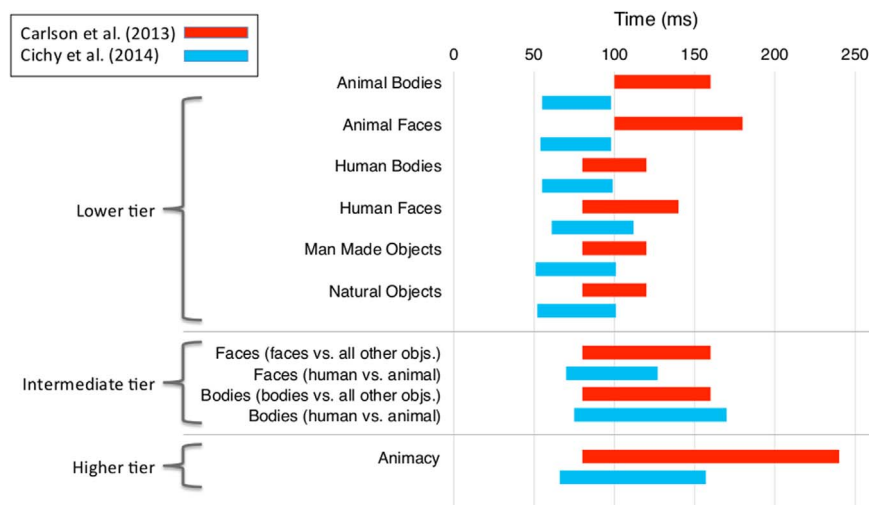
Evaluation of peak decoding times for the category hierarchy in Carlson et al. (2013) revealed a largely linear organisation of the category tiers (see Fig. 5, red bars), such that lower tier categories (e.g., human faces, animal bodies) showed peak decoding times from

120–180 ms, followed by intermediate tier categories (faces, bodies, human, animal) which ranged from 160–240 ms, with the superordinate tier animacy category (animate vs. inanimate) peaking at 240 ms. The results suggest evolving hierarchical representations of object categories that progress from specific, individual exemplar individuation through to more abstract semantic groupings. Results from Cichy et al. (2014) largely support these findings (see Fig. 5, blue bars), with peak decoding between intermediate tier categories (natural vs. artificial and faces vs. bodies, 122 and 136 ms respectively) occurring earlier than the superordinate animacy category peak (157 ms).

Although there is similar evidence in the results of both Carlson et al. (2013) and Cichy et al. (2014) for a roughly linear emergence of object categories from specific to more abstract object category representations, the data are not entirely straightforward. For example, for the subordinate tier categories in Cichy et al. (2014), this linear hierarchy does not hold, with subcategory peaks occurring later in the time-course than some of the intermediate tier category peaks: human versus animal faces peaked at 127 ms, while human versus animal bodies peaked at 170 ms. Furthermore, in Carlson et al. (2013), the intermediate tier categories of faces and bodies are best discriminated at a similar time point to the more specific lower tier categories (e.g., human faces, animal bodies), and similarly, Cichy et al. (2014) found no significant difference between the peak latency at which individual images were discriminated early in the time course compared to a higher tier category grouping of human versus non-human bodies.

It is plausible that these anomalies in the emergence of hierarchical category structure reflect that category representations are not a strict linear hierarchy, consistent with recent ideas on the importance of recurrent processing such as lateral connections and feedback on feedforward visual object processing (Kravitz et al., 2013). Carlson et al. (2013) suggest these discrepancies may be partially explained by the inclusion of face and body stimuli in the design, as there is evidence for special processing systems specifically tuned for the detection of faces and bodies in both human and primate ventral temporal cortex (Downing et al., 2001; Kanwisher et al., 1997). If some categories are





**Fig. 5.** Onset and peak decoding times for category tiers within the object category hierarchy, as evaluated by Carlson et al. (2013) and Cichy et al. (2014). Each coloured bar represents the time from the onset of significant category decoding to the time of peak decoding (relative to post-stimulus onset) for Carlson et al. (2013) (red) and Cichy et al. (2014) (blue). Note, the intermediate tier categories are evaluated using different stimulus comparisons by Carlson et al. (2013) and Cichy et al. (2014) (specific comparisons are noted in brackets).

processed differently than others, a straightforward linear emergence of category structure is not expected. A further caveat is that these results are specific to the selected stimulus set and the imposed category structure of the experimental design.

#### 2.4. Decoding the time-course of face processing

The object category of faces stands out as having a particularly strong contribution to decoded brain representations measured with EEG (Kaneshiro et al., 2015) and MEG (Van de Nieuwenhuijzen et al., 2013). For example, in the MDS results for large object sets, human faces tend to cluster prominently (Fig. 2). Two EEG decoding studies have examined the temporal dynamics of face representations in detail (Barragan-Jason et al., 2015; Cauchoix et al., 2014). A strong feature of these studies is that they used large natural image stimulus sets containing pictures of faces with natural backgrounds, and faces were not repeated in the experimental design. This is in contrast to the majority of temporal object decoding studies to date, which tend to repeat the same object exemplars multiple times within the experiment to increase signal-to-noise. Both studies had participants complete a go/no-go task and trained a classifier to detect the target versus non-target faces across the EEG time-course (human vs. animal in Cauchoix et al., 2014, and famous vs. unfamiliar faces in Barragan-Jason et al., 2015). In both studies, the time-course of face decoding follows a similar pattern: there is an initial sharp rise in performance of the classifier early in the time-course, shortly followed by a relatively brief plateau or drop in decoding, before a second, more gradual rise in decoding accuracy, with peak decoding occurring relatively late in the time-course (350 ms and 600 ms for humans/animals and familiarity respectively). They suggest that the two (early vs. late) peaks in decoding reflect different levels of processing within the object recognition hierarchy. An initial fast feedforward signal driven by low-level stimulus properties is later modified to incorporate more detailed feedback from higher cortical areas (Barragan-Jason et al., 2015; Cauchoix et al., 2014).

A recent study which combined MEG decoding with fMRI localization provides further insight into different levels of face processing. Vida et al. (2017) evaluated time-series decoding performance of face identity in specific regions of the brain linked to face processing. They used an fMRI localiser to identify two regions that were preferentially activated by faces, over and above activity associated with objects more broadly: right lateral occipital cortex and the right fusiform gyrus. At these sites, and a control site (left V1) they related the time-course of face decoding to three models of face processing: one that represented

low-level, V1-like responses (“image-based” model), a higher-level “identity-based” face model, and a behavioural-rating model where participants rated a subset of the images as to how similar the face identities were on a scale from 1 to 8. Their results showed that firstly, while face identity was decodable at all three regions from ~50–400 ms, the image-based model was a better predictor of the neural data than the identity-based model until ~200 ms, after which its performance dropped below that of the identity-based model at a number of time-points. However, this shift occurred in the face-selective regions only, with the image-based model outperforming others in the V1 control region throughout the time-course.

These results are in line with previous findings (see Section 2.1) that suggest low-level visual features are associated with activity in early visual areas, while high-level conceptual information is represented in cortical regions further down the ventral processing stream. The behavioural-rating model similarly correlated with the neural data significantly in all three regions from ~50–400 ms, however, after controlling for representations in the V1 control region, this significant time-window reduced to 100–250 ms in the face-selective regions. Notably, behavioural ratings were significantly more similar to the identity-based model than the image-based model. Significantly, by using time-series decoding methods, Vida et al. were able to show that face-selective regions process information related to *both* low-level image properties as well as higher-level face identity-based representations, highlighting the potential shortcomings of associating localised regions with single functions. The challenges involved in empirically separating out the contribution of low level properties from more abstract object representations in object decoding studies is discussed in detail in the following section.

#### 2.5. Is a category a category?

Above we reviewed several studies focusing on uncovering the dynamic category structure of object representations. However, when evaluating object representations, it is important to separate effects related to the abstract concept of an object “category” from other potentially confounding features that co-vary with category membership such as low-level visual properties including shape, colour, and luminance contrast (Wardle and Ritchie, 2014). For example, items within the category of tools tend to have handles, resulting in a characteristic long and often slender feature to their appearance (e.g., hammer, scissors, screwdriver, toothbrush). This relatively consistent information within this category could result in decoding based on low-level feature similarities related to shape, such that instead of decoding

the category of “tools” compared to “faces”, we may instead be erroneously decoding “long and thin” versus “round”. Low-level stimulus properties such as orientation (Cichy et al., 2015; Ramkumar et al., 2013) are readily decodable from whole-brain MEG signals, and perceptual similarity is a strong determinant of the decodability of abstract visual stimuli (Wardle et al., 2016).

Time-varying decoding studies investigating object recognition have attempted to account for potential low-level confounds using a variety of different approaches. In fMRI studies, V1 is often used as a control region to show that unsupervised categorical clustering of object representations (e.g., by animacy) do not emerge in V1, but are present in IT (Cichy et al., 2014; Kriegeskorte et al., 2008). As this degree of spatial localisation is not possible with MEG, time-varying decoding studies often exploit the time-course to demonstrate that models of early visual processing may perform well early on in the time-course but later the contribution from categorical models such as animacy emerges. For example, Carlson et al. (2013) used RSA to evaluate the performance of low-level feature based models of vision on object representations measured with MEG. The models of early visual processing included a shape-based model that compared the image silhouettes (Jaccard, 1901), a colour-based model (CIE), and a hierarchical visual processing model (HMAX) containing layers tuned to process varying levels of stimulus complexity (Riesenhuber and Poggio, 1999; Serre et al., 2007). Notably, these models were able to successfully differentiate human faces from other objects (silhouette model), humans from human bodies (CIE), and man-made objects from other objects (HMAX). While the models were only able to distinguish a maximum of two out of a possible 10 object categories in that particular data set, these results highlight the importance of accounting for the potential influence of low-level stimulus properties on object category decoding.

Although there is evidence that visual similarity accounts for at least some of the representational similarity shared by objects within the same category, it is unlikely to account for all observed category effects. As a control for low-level stimulus properties when decoding object category, Carlson et al. (2011) evaluated decoding accuracy of cars and faces from MEG activation patterns compared to artificially generated car and face textures which preserved local image statistics but removed recognisable form. If classification of object category is heavily based on low-level image statistics, it is expected that the classifier would find it difficult to discriminate between objects and their matched texture images, which share low-level image properties (e.g., classifying ‘car’ vs. ‘car-texture’). However, object categories were able to be accurately decoded from their texture counterparts in the whole-brain MEG activation patterns, and this distinction emerged earlier in the time course than information differentiating the two object categories of cars and faces (Carlson et al., 2011). These results are consistent with a transition from decoding based on V1-like image properties to IT-like object category over the MEG decoding time course (Cichy et al., 2014).

The contribution of semantic meaning to category representations was explored in a study by Clarke et al. (2014), which compared the performance of both low-level visual processing and feature-based semantic models for their 302 object stimuli from 11 categories. They found that a model which combined both category-specific semantic information and low-level visual features best accounted for variability in neural object representations later in the MEG time course, over and above what could be achieved by using the HMAX model alone (Clarke et al., 2014). These results suggest there is a contribution of abstract category membership (as indexed by semantic similarity) to the brain representation. Similarly, an EEG study evaluating event-related potentials (ERPs) associated with object representations aimed to provide evidence for the involvement of conceptual category membership by examining decoding of the categories ‘animals’ versus ‘tools’ across three modalities of visual pictures, spoken words and written words (Simanova et al., 2010). However, classifier performance was


much higher for visual pictures than for spoken or written words, complicating the critical cross-classification analysis across modality.

Similar to the frequently-used approach of assessing the potential contribution of low-level properties to object representations with the HMAX model (Carlson et al., 2013; Clarke et al., 2014; Isik et al., 2014; Kriegeskorte et al., 2008), recently Cichy et al. (2016) compared MEG object decoding performance with the performance of deep neural networks (DNNs). DNNs are a powerful form of computer vision model, comprised of multiple non-linear processing layers with the ability to learn tasks such as object categorization (see Kriegeskorte, 2015). The recent advancement of these models has propelled computer modelling of object recognition to a level that is beginning to rival human object categorisation performance (He et al., 2015). By comparing the decoding performance of the DNN model to MEG and fMRI brain data, Cichy et al. (2016) showed a hierarchy of both spatially and temporally ordered processes that correlated with the DDN processing layers in an ordered fashion; deeper layers were associated with later brain activity in higher cortical processing regions. Furthermore, by analysing activity throughout the entire brain, they identified involvement of not only the ventral, but also the dorsal pathway in object recognition. A key issue for the future is to understand to what degree deep-neural networks mimic the object recognition processes in the human brain, and in what respects the two systems are different.

The above approaches are data-driven, and aim to show that categorical structure that emerges in the neuroimaging data cannot be completely explained by sensitivity to lower level visual properties confounded with object category. A more compelling line of evidence against reducing category representations entirely to low-level properties involves experimental manipulations designed to separate the two factors. A recent MEG/fMRI study (Kaiser et al., 2016a) sought to specifically identify object category representations independent of shared visual properties by purposefully selecting visually similar stimuli that belonged to the semantic categories of body parts and clothing (e.g., gloves vs. hands, shirts vs. torsos). With this stimulus set, visual similarity is balanced across category membership. While their fMRI results showed overlapping spatial representations for both shape and category information, the MEG data showed a specific, comparatively late time window within which category-selective information was present (Fig. 6b) and in contrast to this, shape dependent responses (Fig. 6a) were decoded relatively early on, from 90 ms post-stimulus onset, and was sustained throughout the time-course.

Object invariance is both an interesting theoretical question for investigation and an experimental manipulation that minimises the contribution of low-level features. Object invariance is the ability to successfully recognise objects despite high variability in their appearance, for instance, due to differences in viewpoint and size. Carlson et al. (2011) showed that object category information for faces and cars could be decoded from the neural data despite changes in retinal location of the image; evidence for a position-invariant object representation. They tested a classifier on information from a novel retinal location that was not included in the original training set, and were able to successfully recover information about the object category as early as 135 ms, suggesting that position-invariant category information is present early in the neural signal. Furthermore, this was not simply due to low-level image statistics, as category information could not be decoded from abstract textures generated to preserve the local low-level image statistics of the object images (Carlson et al., 2011).

Invariant object decoding was further investigated by Isik et al. (2014). Isik et al. showed position invariant decoding across three stimulus locations for a set of six object exemplars, while also evaluating size invariance. Using three different stimulus sizes for each exemplar, they showed that size-invariant information was present in the signal from around 125 ms post-stimulus onset, earlier in the time-course than position-invariant information which emerged around 150 ms for their stimulus set. Furthermore, by comparing the decoding



**Fig. 6.** Temporal cross decoding of object category (body parts vs. clothing). The black outlined areas in both plots indicate when cross-decoding of (a) shape or (b) object category was significantly above chance. Images reprinted from Kaiser et al. (2016a) with permission from ©The American Physiological Society.

latencies for the different train-test conditions, they found that comparisons which required a greater transformation of the data from the train to the test condition (e.g., for small, middle and large sized stimuli, train small/test large requires a bigger transformation than train middle/test large) resulted in later decoding onsets.

Successful object decoding across changes in retinal size or position suggests access to more abstract object representations, which is likely to reflect higher-level visual processing. Isik et al. (2014) compared their results to the different levels of the computational HMAX model (Riesenhuber and Poggio, 1999; Serre et al., 2007), which includes multiple stages representing early V1 processing (layer C1), mid-ventral processing (later C2), and higher level IT-like processing (layer C3). Comparison of the experimental results with the response of the HMAX model for their six object exemplars revealed that only the IT layer of HMAX could successfully decode object exemplars across all transformations of size and position. Isik et al. suggest their observed relationship between the degree of image transformation and decoding time (i.e., that decoding occurs later for larger visual transformations across object size and position) is consistent with a feedforward processing hierarchy supporting invariant object representations. In the next section we consider the link between object category representations and human behaviour.

### 3. Linking dynamic object category representations to human behaviour

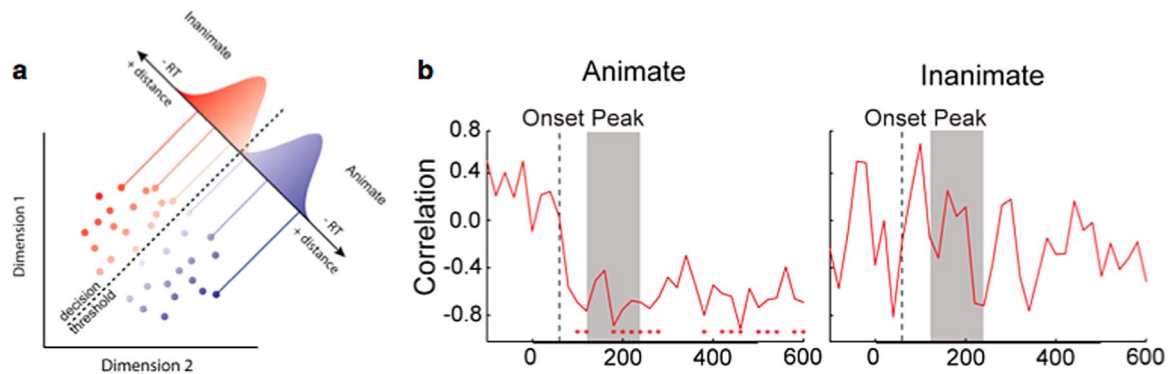
One of the critical limitations of decoding approaches in neuroscience is that successful decoding of a stimulus property or experimental condition does not entail that this information is explicitly used by the brain for behaviour (cf. de Wit et al., 2016; Ritchie et al., in press). An emerging goal is to find ways to link measured brain activation patterns to behaviour (Mack and Palmeri, 2011; Nosofsky and Palmeri, 1997; Philastides and Sajda, 2006). In this section we review two recent approaches to linking the decoded dynamic object representations from time-series neuroimaging data with human behaviour. In the first approach, the goal is to link reaction times (RTs) for object categorisation to the brain's evolving representation of the objects in multidimensional space (Section 3.1). The second involves uncovering the effects of attention on decodable object representations (Section 3.2).

#### 3.1. Categorisation reaction times

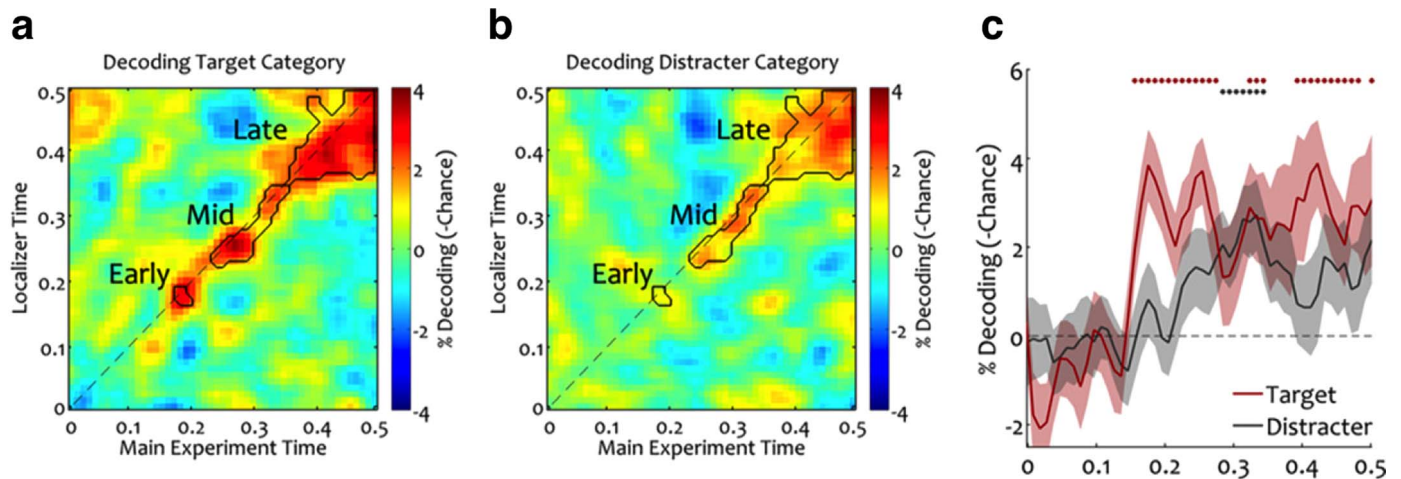
One fruitful approach to evaluate the relationship between brain and behaviour has been to link human RTs for categorisation to the multidimensional representation of object exemplars in neuroimaging data (Ritchie and Carlson, 2016). Human categorisation behaviour can be conceptualised as a process of evidence accumulation leading to a decision. Behavioural RTs, such as those measured by a simple button press to categorise a stimulus (e.g., fish as 'animate', boot as 'inanimate'), can be thought of as a proxy for the point at which a decision has been made. Carlson et al. (2014) tested this idea by combining previously collected fMRI data from human IT (Kriegeskorte et al., 2008) with separately collected behavioural RTs for categorising the same 92 object images as either 'animate' or 'inanimate'. Carlson et al. (2014) linked the decision boundary from the fMRI data (boundary is that used by a classifier to categorise stimuli by animacy) to RTs for human observers to complete the same categorisation task. They showed that the distance to the classifier decision boundary predicted reaction times for animacy categorisation. Specifically, objects represented further from the animacy decision boundary were associated with faster behavioural categorisation RTs (i.e., more easily categorised by human observers) than objects represented closer to the boundary (Fig. 7a).

Ritchie et al. (2015) expanded on this idea by studying the dynamic coupling between emerging object representations and behaviour using MEG decoding. Their findings replicated the relationship between representational distance and RTs found by Carlson et al. (2014) for fMRI, as a larger distance from the classifier boundary was associated with a faster reaction time for animacy categorisation. The key advantage of examining this relationship with MEG is that its emergence following stimulus onset can be tracked over time, which cannot be determined from the fMRI data. Ritchie et al. (2015) found that as decoding accuracy increased following stimulus onset (indicating greater separability of individual object representations with more processing time), the correlation between representational distance and RTs similarly increased. Importantly, this suggests that the strength of the association between representations and behaviour appears to follow decoding performance over time (Ritchie et al., 2015). The significance of this result is that it implies that the brain "reads out" information at the optimal time for making a categorical





**Fig. 7.** Linking behavioural RTs for object animacy categorisation to dynamic object representations measured with MEG. (a) Diagram illustrating how representational distance from the classifier category boundary relates to RTs. (b) Left panel shows the correlation between RTs and the MEG decoding time-course for animate objects, while the right panel shows the same for inanimate objects. Grey bar in both plots indicates the window of peak decoding. Red stars along the x-axis indicate time-points at which the correlation between representational distance and reaction time is significant. Images from [Ritchie et al. \(2015\)](#) reprinted with permission from the author.



**Fig. 8.** MEG decoding of attended versus unattended targets. Temporal cross-decoding heat-maps of (a) target category and (b) unattended distractor category (training on isolated objects, testing on objects embedded in cluttered natural scenes in both instances); (c) Time-course of decoding performance for the target and distractor conditions separately. Dots indicate above-chance decoding performance. Images reprinted with permission of Society for Neuroscience, from [Kaiser et al. \(2016b\)](#); permission conveyed through Copyright Clearance Center, Inc.

decision. This highlights the utility of time-series decoding methods, as this result could not be accessed by examining time-averaged activation patterns as in fMRI.

Interestingly, in both the fMRI and MEG studies, the relationship between representational distance and RTs appeared to be driven by animate rather than inanimate objects ([Carlson et al., 2014](#), [Ritchie et al., 2015](#)). This is shown in [Fig. 7b](#) by comparing the left (animate) and right (inanimate) RT correlation panels. For animate object decoding, a negative correlation is observed with RTs around the time of peak decoding (indicating that faster RTs were associated with greater distances from the decision boundary for animacy). However, this relationship is not found with inanimate objects. The authors suggest this asymmetry may be because inanimacy is a negatively defined category (i.e., 'not animate').

### 3.2. Attention

Attention has strong modulatory effects on neural processing ([Gandhi et al., 1999](#); [Kastner and Ungerleider, 2000](#); [Moran and Desimone, 1985](#)). A recent MEG study by [Kaiser et al. \(2016b\)](#) evaluated the effects of top-down attention on object category representations. Participants performed a category search task by detecting the presence of either a car or a person embedded in cluttered natural scenes. This design allows comparison of the neural signal associated with the attended versus unattended object category. The classifier was trained on data from viewing car and people exemplars that were

separate to the experimental data set and presented as segmented objects on a plain background. The classifier was then tested on the data for when the objects were embedded in complex natural scenes as either a target or a distractor.

The results demonstrate a substantial effect of attention on object representations. Under attended conditions, object category could be decoded rapidly in under 200 ms even though the objects were in highly cluttered scenes ([Fig. 8c](#)). However, unattended object categories embedded in complex scenes were decoded much later ([Fig. 8c](#)), demonstrating a strong top-down modulation of attention on visual processing. Attended objects were decoded more easily than unattended objects across all early (180 – 220 ms), mid (230 – 340 ms), and late (350 – 500 ms) time windows ([Fig. 8a, b](#)). The greater temporal generalisation of cross-decoding during the late time window indicates that the representational similarity shared by isolated objects and objects in scenes is sustained for longer in the late processing stages, which is particularly evident for attended objects ([Kaiser et al., 2016b](#)). These results suggest that top-down attention rapidly modulates object category representations, facilitating rapid target detection in natural scenes ([Thorpe et al., 1996](#)). By using temporal decoding, this study was able to reveal the time course of category-specific attentional enhancement, extending previous EEG findings which instead focused on comparing evoked potentials to target presence versus absence ([Codispoti et al., 2006](#); [Thorpe et al., 1996](#)).

## 4. Summary and conclusions

Above we have reviewed how time-series decoding methods have expanded our understanding of the neural processes underlying human object recognition. Object recognition remains a significant area of interest within the field of vision research and the recent implementation of time-series decoding methods has allowed researchers to unpack some of the neural processes that underlie *how* the brain achieves this highly complex and multifaceted task. Complementing other neuroimaging techniques, such as fMRI, which provide static snapshots of the visual object processing hierarchy, time-series decoding has unveiled a system of dynamically changing processes that extend throughout the brain. This new perspective brings a host of interesting avenues for further exploration. For example, at what point in time do high-level constructs of attention and consciousness modulate the dynamics of object processing, and what is the effect of this modulation on perception? How do space- and feature-based attention processes differentially influence the processing hierarchy? There remain unanswered questions about the processes involved in visual working memory and how these representations in the mind's eye compare to the perception of a physical stimulus; are they fundamentally different or is visual working memory a process of revisiting points within these dynamic representations? Finally, it would be of particular interest to know how these various modulations of object representations differentially influence behaviour, further strengthening our understanding of the connection between brain and behaviour. Time-series decoding methods have greatly enriched our knowledge of neural object recognition processes, and show great potential for further advances in this and many other domains.

## Acknowledgements

This research was supported by an Australian Research Council (ARC) Discovery project (DP160101300). T.A.C is supported by an ARC Future Fellowship (FT120100816). S.G.W. is supported by an Australian NHMRC Early Career Fellowship (APP1072245).

## References

- Altmann, C.F., Bühlhoff, H.H., Kourtzi, Z., 2003. Perceptual organization of local elements into global shapes in the human visual cortex. *Curr. Biol.* 13 (4), 342–349. [http://dx.doi.org/10.1016/S0960-9822\(03\)00052-6](http://dx.doi.org/10.1016/S0960-9822(03)00052-6).
- Averbeck, B.B., Latham, P.E., Pouget, A., 2006. Neural correlations, population coding and computation. *Nat. Rev. Neurosci.* 7 (5), 358–366.
- Barragan-Jason, G., Cauchoix, M., Barbeau, E.J., 2015. The neural speed of familiar face recognition. *Neuropsychologia* 75, 390–401. <http://dx.doi.org/10.1016/j.neuropsychologia.2015.06.017>.
- Beauchamp, M.S., Lee, K.E., Haxby, J.V., Martin, A., 2002. Parallel visual motion processing streams for manipulable objects and human movements. *Neuron* 34 (1), 149–159.
- Biederman, I., 1987. Recognition-by-components: a theory of human image understanding. *Psychol. Rev.* 94 (2), 115.
- Carlson, Hogendoorn, H., Kanai, R., Mesik, J., Turret, J., 2011. High temporal resolution decoding of object position and category. *J. Vis.* 11 (10), 9. <http://dx.doi.org/10.1167/11.10.9>.
- Carlson, Schrater, P., He, S., 2003. Patterns of activity in the categorical representations of objects. *J. Cogn. Neurosci.* 15 (5), 704–717. <http://dx.doi.org/10.1162/jocn.2003.15.5.704>.
- Carlson, T.A., Ritchie, J.B., Kriegeskorte, N., Durvasula, S., Ma, J., 2014. Reaction time for object categorization is predicted by representational distance. *J. Cogn. Neurosci.* 26 (1), 132–142. [http://dx.doi.org/10.1162/jocn\\_a\\_00476](http://dx.doi.org/10.1162/jocn_a_00476).
- Carlson, Tovar, Alink, A., D.A., Kriegeskorte, N., 2013. Representational dynamics of object vision: the first 1000 ms. *J. Vis.* 13 (10). <http://dx.doi.org/10.1167/13.10.1>.
- Cauchoix, M., Barragan-Jason, G., Serre, T., Barbeau, E.J., 2014. The neural dynamics of face detection in the wild revealed by MVPA. *J. Neurosci.* 34 (3), 846–854. <http://dx.doi.org/10.1523/JNEUROSCI.3030-13.2014>.
- Chait, M., Poeppel, D., Cheveigné, A. de, Simon, J.Z., 2007. Processing asymmetry of transitions between order and disorder in human auditory cortex. *J. Neurosci.* 27 (19), 5207–5214. <http://dx.doi.org/10.1523/JNEUROSCI.0318-07.2007>.
- Chao, L.L., Haxby, J.V., Martin, A., 1999. Attribute-based neural substrates in temporal cortex for perceiving and knowing about objects. *Nat. Neurosci.* 2 (10), 913–919.
- Cichy, R.M., Khosla, A., Pantazis, D., Torralba, A., Oliva, A., 2016. Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence. *Sci. Rep.* 6. <http://dx.doi.org/10.1038/srep27755>.
- Cichy, R.M., Pantazis, D., Oliva, A., 2014. Resolving human object recognition in space and time. *Nat. Neurosci.* 17 (3), 455–462. <http://dx.doi.org/10.1038/nn.3635>.
- Cichy, R.M., Ramirez, F.M., Pantazis, D., 2015. Can visual information encoded in cortical columns be decoded from magnetoencephalography data in humans? *NeuroImage* 121, 193–204. <http://dx.doi.org/10.1016/j.neuroimage.2015.07.011>.
- Clarke, A., Devereux, B.J., Randall, B., Tyler, L.K., 2014. Predicting the time course of individual objects with MEG. *Cereb. Cortex*. <http://dx.doi.org/10.1093/cercor/bhu203>.
- Clarke, A., Tyler, L.K., 2014. Object-specific semantic coding in human perirhinal cortex. *J. Neurosci.* 34 (14), 4766–4775. <http://dx.doi.org/10.1523/JNEUROSCI.2828-13.2014>.
- Codispoti, M., Ferrari, V., Junghöfer, M., Schupp, H.T., 2006. The categorization of natural scenes: brain attention networks revealed by dense sensor ERPs. *NeuroImage* 32 (2), 583–591. <http://dx.doi.org/10.1016/j.neuroimage.2006.04.180>.
- Connolly, A.C., Guntupalli, J.S., Gors, J., Hanke, M., Halchenko, Y.O., Wu, Y.-C., Haxby, J.V., 2012. The representation of biological classes in the human brain. *J. Neurosci.* 32 (8), 2608–2618. <http://dx.doi.org/10.1523/JNEUROSCI.5547-11.2012>.
- Cox, D.D., Savoy, R.L., 2003. Functional magnetic resonance imaging (fMRI) 'brain reading': detecting and classifying distributed patterns of fMRI activity in human visual cortex. *NeuroImage* 19 (2 Pt 1), 261–270.
- de Wit, L., Alexander, D., Ekroll, V., Wagemans, J., 2016. Is neuroimaging measuring information in the brain? *Psychonomic Bull. Rev.* 23 (5), 1415–1428. <http://dx.doi.org/10.3758/s13423-016-1002-0>.
- DiCarlo, J.J., Cox, D.D., 2007. Untangling invariant object recognition. *Trends Cogn. Sci.* 11 (8), 333–341.
- Downing, P., Jiang, Y., Shuman, M., Kanwisher, N., 2001. A cortical area selective for visual processing of the human body. *Science* 293 (5539), 2470–2473.
- Epstein, R., Kanwisher, N., 1998. A cortical representation of the local visual environment. *Nature* 392 (6676), 598–601.
- Gandhi, S.P., Heeger, D.J., Boynton, G.M., 1999. Spatial attention affects brain activity in human primary visual cortex. *Proc. Natl. Acad. Sci.* 96 (6), 3314–3319. <http://dx.doi.org/10.1073/pnas.96.6.3314>.
- Goddard, E., Carlson, T.A., Dermody, N., Woolgar, A., 2016. Representational dynamics of object recognition: feedforward and feedback information flows. *NeuroImage* 128, 385–397. <http://dx.doi.org/10.1016/j.neuroimage.2016.01.006>.
- Grill-Spector, K., Kourtzi, Z., Kanwisher, N., 2001. The lateral occipital complex and its role in object recognition. *Vis. Res.* 41 (10–11), 1409–1422.
- Grill-Spector, K., Malach, R., 2004. The human visual cortex. *Annu. Rev. Neurosci.* 27 (1), 649–677. <http://dx.doi.org/10.1146/annurev.neuro.27.070203.144220>.
- Grootswagers, T., Wardle, S.G., Carlson, T.A., 2016. Decoding dynamic brain patterns from evoked responses: a tutorial on multivariate pattern analysis applied to time-series neuroimaging data. *J. Cogn. Neurosci.* [http://dx.doi.org/10.1162/jocn\\_a\\_01068](http://dx.doi.org/10.1162/jocn_a_01068).
- Haxby, J.V., Connolly, A.C., Guntupalli, J.S., 2014. Decoding neural representational spaces using multivariate pattern analysis. *Annu. Rev. Neurosci.* 37 (1), 435–456. <http://dx.doi.org/10.1146/annurev-neuro-062012-170325>.
- Haxby, J.V., Gobbini, M.I., Furey, M.L., Ishai, A., Schouten, J.L., Pietrini, P., 2001. Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* 293 (5539), 2425–2430. <http://dx.doi.org/10.1126/science.1063736>.
- Haynes, J.-D., 2015. A primer on pattern-based approaches to fMRI: principles, pitfalls, and perspectives. *Neuron* 87 (2), 257–270.
- Haynes, J.-D., Rees, G., 2006. Decoding mental states from brain activity in humans. *Nat. Rev. Neurosci.* 7 (7), 523–534. <http://dx.doi.org/10.1038/nrn1931>.
- He, K., Zhang, X., Ren, S., Sun, J., 2015. Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification. In *Computer Vision (ICCV), 2015 IEEE International Conference on* pp. 1026–1034. IEEE.
- Hung, C.P., Kreiman, G., Poggio, T., DiCarlo, J.J., 2005. Fast readout of object identity from macaque inferior temporal cortex. *Science* 310 (5749), 863–866.
- Ishai, A., Ungerleider, L.G., Martin, A., Schouten, J.L., Haxby, J.V., 1999. Distributed representation of objects in the human ventral visual pathway. *Proc. Natl. Acad. Sci.* 96 (16), 9379–9384. <http://dx.doi.org/10.1073/pnas.96.16.9379>.
- Isik, L., Meyers, E.M., Leibo, J.Z., Poggio, T., 2014. The dynamics of invariant object recognition in the human visual system. *J. Neurophysiol.* 111 (1), 91–102. <http://dx.doi.org/10.1152/jn.00394.2013>.
- Jaccard, P., 1901. Etude comparative de la distribution florale dans une portion des Alpes et du Jura. *Impr. Corbaz*.
- Kaiser, D., Azzalini, D.C., Peelen, M.V., 2016a. Shape-independent object category responses revealed by MEG and fMRI decoding. *J. Neurophysiol.* 115 (4), 2246–2250. <http://dx.doi.org/10.1152/jn.01074.2015>.
- Kaiser, D., Oosterhof, N.N., Peelen, M.V., 2016b. The neural dynamics of attentional selection in natural scenes. *J. Neurosci.* 36 (41), 10522–10528. <http://dx.doi.org/10.1523/JNEUROSCI.1385-16.2016>.
- Kaneshiro, B., Guimaraes, M.P., Kim, H.-S., Norcia, A.M., Suppes, P., 2015. A representational similarity analysis of the dynamics of object processing using single-trial EEG classification. *PLOS One* 10 (8), e0135697. <http://dx.doi.org/10.1371/journal.pone.0135697>.
- Kanwisher, N., McDermott, J., Chun, M.M., 1997. The fusiform face area: a module in human extrastriate cortex specialized for face perception. *J. Neurosci.* 17, 4302–4311.
- Kastner, S., Ungerleider, L.G., 2000. Mechanisms of visual attention in the human cortex. *Annu. Rev. Neurosci.* 23 (1), 315–341. <http://dx.doi.org/10.1146/annurev.neuro.23.1.315>.
- Kiani, R., Esteky, H., Mirpour, K., Tanaka, K., 2007. Object category structure in

- response patterns of neuronal population in monkey inferior temporal cortex. *J. Neurophysiol.* 97 (6), 4296–4309. <http://dx.doi.org/10.1152/jn.00024.2007>.
- King, J.-R., Dehaene, S., 2014. Characterizing the dynamics of mental representations: the temporal generalization method. *Trends Cogn. Sci.* 18 (4), 203–210. <http://dx.doi.org/10.1016/j.tics.2014.01.002>.
- Kravitz, D.J., Saleem, K.S., Baker, C.I., Ungerleider, L.G., Mishkin, M., 2013. The ventral visual pathway: an expanded neural framework for the processing of object quality. *Trends Cogn. Sci.* 17 (1), 26–49. <http://dx.doi.org/10.1016/j.tics.2012.10.011>.
- Kriegeskorte, N., 2008. Representational similarity analysis—connecting the branches of systems neuroscience. *Front. Syst. Neurosci.* <http://dx.doi.org/10.3389/neuro.06.004.2008>.
- Kriegeskorte, N., 2015. Deep neural networks: a new framework for modelling biological vision and brain information processing. *bioRxiv*, 29876. <http://dx.doi.org/10.1101/029876>.
- Kriegeskorte, N., Kievit, R.A., 2013. Representational geometry: integrating cognition, computation, and the brain. *Trends Cogn. Sci.* 17 (8), 401–412. <http://dx.doi.org/10.1016/j.tics.2013.06.007>.
- Kriegeskorte, N., Mur, M., Ruff, D.A., Kiani, R., Bodurka, J., Esteky, H., Bandettini, P.A., 2008. Matching categorical object representations in inferior temporal cortex of man and monkey. *Neuron* 60 (6), 1126–1141. <http://dx.doi.org/10.1016/j.neuron.2008.10.043>.
- Liu, H., Agam, Y., Madsen, J.R., Kreiman, G., 2009. Timing, timing, timing: fast decoding of object information from intracranial field potentials in human visual cortex. *Neuron* 62 (2), 281–290. <http://dx.doi.org/10.1016/j.neuron.2009.02.025>.
- Mack, M.L., Palmeri, T.J., 2011. The timing of visual object Categorization. *Front. Psychol.* 2. <http://dx.doi.org/10.3389/fpsyg.2011.00165>.
- Malach, R., Levy, I., Hasson, U., 2002. The topography of high-order human object areas. *Trends Cogn. Sci.* 6 (4), 176–184.
- Martin, A., Wiggs, C. L., Ungerleider, L. G., Haxby, J. V., 1996. Neural correlates of category-specific knowledge. Retrieved from <http://www.nature.com/nature/journal/v379/n6566/abs/379649a0.html>.
- Mervis, C.B., Rosch, E., 1981. Categorization of natural objects. *Annu. Rev. Psychol.* 32 (1), 89–115.
- Moran, J., Desimone, R., 1985. Selective attention gates visual processing in the extrastriate cortex. *Front. Cogn. Neurosci.* 229, 342–345.
- Nosofsky, R.M., Palmeri, T.J., 1997. An exemplar-based random walk model of speeded classification. *Psychol. Rev.* 104 (2), 266.
- Nowak, L.G., Bullier, J., 1997. The Timing of Information Transfer in the Visual System. In: Rockland, K.S., Kaas, J.H., Peters, A. (Eds.), *Extrastriate Cortex in Primates*. Springer, US, 205–241. [http://dx.doi.org/10.1007/978-1-4757-9625-4\\_5](http://dx.doi.org/10.1007/978-1-4757-9625-4_5).
- O'Toole, A.J., Jiang, F., Abdi, H., Haxby, J.V., 2005. Partially distributed representations of objects and faces in ventral temporal cortex. *J. Cogn. Neurosci.* 17 (4), 580–590.
- Pereira, F., Mitchell, T., Botvinick, M., 2009. Machine learning classifiers and fMRI: a tutorial overview. *Neuroimage* 45 (1), S199–S209.
- Philastides, M.G., Sajda, P., 2006. Temporal characterization of the neural correlates of perceptual decision making in the human brain. *Cereb. Cortex* 16 (4), 509–518.
- Ramkumar, P., Jas, M., Pannasch, S., Hari, R., Parkkonen, L., 2013. Feature-specific information processing precedes concerted activation in human visual cortex. *J. Neurosci.* 33 (18), 7691–7699. <http://dx.doi.org/10.1523/JNEUROSCI.3905-12.2013>.
- Riesenhuber, M., Poggio, T., 1999. Hierarchical models of object recognition in cortex. *Nat. Neurosci.* 2 (11), 1019–1025. <http://dx.doi.org/10.1038/14819>.
- Ritchie, J.B., Carlson, T.A., 2016. Neural Decoding and 'Inner' Psychophysics: A Distance-to-Bound Approach for Linking Mind, Brain, and Behavior. *Front. Neurosci.* 10: 190. <http://dx.doi.org/10.3389/Fnins>.
- Ritchie, J.B., Kaplan, D.M., Klein, C., in press. Decoding the brain: Neural representation and the limits of multivariate pattern analysis in cognitive neuroscience. *British Journal for the Philosophy of Science*, <http://dx.doi.org/10.1101/127233>.
- Ritchie, J.B., Tovar, D.A., Carlson, T.A., 2015. Emerging object representations in the visual system predict reaction times for categorization. *PLoS Comput Biol.* 11 (6), e1004316. <http://dx.doi.org/10.1371/journal.pcbi.1004316>.
- Schmolekys, M.T., Wang, Y., Hanes, D.P., Thompson, K.G., Leutgeb, S., Schall, J.D., Leventhal, A.G., 1998. Signal timing across the macaque visual system. *J. Neurophysiol.* 79 (6), 3272–3278.
- Serre, T., Kreiman, G., Kouh, M., Cadieu, C., Knoblich, U., Poggio, T., 2007. A quantitative theory of immediate visual recognition. In: T.D, Paul Cisek, J.F.K. (Eds.), *Progress in Brain Research* 165. Elsevier, 33–56. [http://dx.doi.org/10.1016/S0079-6123\(06\)65004-8](http://dx.doi.org/10.1016/S0079-6123(06)65004-8).
- Sha, L., Haxby, J.V., Abdi, H., Guntupalli, J.S., Oosterhof, N.N., Halchenko, Y.O., Connolly, A.C., 2015. The animacy continuum in the human ventral vision pathway. *J. Cogn. Neurosci.* 27 (4), 665–678. [http://dx.doi.org/10.1162/jocn\\_a\\_00733](http://dx.doi.org/10.1162/jocn_a_00733).
- Simanova, I., Gerven, M., van, Oostenveld, R., Hagoort, P., 2010. Identifying object categories from event-related EEG: toward decoding of conceptual representations. *PLOS One* 5 (12), e14465. <http://dx.doi.org/10.1371/journal.pone.0014465>.
- Thorpe, S., Fize, D., Marlot, C., 1996. Speed of processing in the human visual system. *Nature* 381 (6582), 520–522.
- Van de Nieuwenhuijzen, M.E., Backus, A.R., Bahramisharif, A., Doeller, C.F., Jensen, O., van Gerven, M.A., 2013. MEG-based decoding of the spatiotemporal dynamics of visual category perception. *NeuroImage* 83, 1063–1073.
- Van Essen, D.C., Anderson, C.H., Felleman, D.J., 1992. Information processing in the primate visual system: an integrated systems perspective. *Science* 255 (5043), 419.
- Vida, M.D., Nestor, A., Plaut, D.C., Behrmann, M., 2017. Spatiotemporal dynamics of similarity-based neural representations of facial identity. *Proc. Natl. Acad. Sci.* 114 (2), 388–393.
- Wardle, S.G., Kriegeskorte, N., Grootswagers, T., Khaligh-Razavi, S.-M., Carlson, T.A., 2016. Perceptual similarity of visual patterns predicts dynamic neural activation patterns measured with MEG. *NeuroImage* 132, 59–70. <http://dx.doi.org/10.1016/j.neuroimage.2016.02.019>.
- Wardle, S.G., Ritchie, J.B., 2014. Can object category-selectivity in the ventral visual pathway be explained by sensitivity to low-level image properties? *J. Neurosci.* 34 (45), 14817–14819. <http://dx.doi.org/10.1523/JNEUROSCI.3566-14.2014>.