

Asymmetric Compression of Representational Space for Object Animacy Categorization under Degraded Viewing Conditions

Tijl Grootswagers^{1,2,3}, J. Brendan Ritchie⁴, Susan G. Wardle^{1,2}, Andrew Heathcote^{5,6}, and Thomas A. Carlson^{2,3}

Abstract

■ Animacy is a robust organizing principle among object category representations in the human brain. Using multivariate pattern analysis methods, it has been shown that distance to the decision boundary of a classifier trained to discriminate neural activation patterns for animate and inanimate objects correlates with observer RTs for the same animacy categorization task [Ritchie, J. B., Tovar, D. A., & Carlson, T. A. Emerging object representations in the visual system predict reaction times for categorization. *PLoS Computational Biology*, *11*, e1004316, 2015; Carlson, T. A., Ritchie, J. B., Kriegeskorte, N., Durvasula, S., & Ma, J. Reaction time for object categorization is predicted by representational distance. *Journal of Cognitive Neuroscience*, *26*, 132–142, 2014]. Using MEG decoding, we tested if the same relationship holds when a stimulus manipulation (degradation) increases task difficulty, which we predicted would systematically decrease the distance of activation patterns

from the decision boundary and increase RTs. In addition, we tested whether distance to the classifier boundary correlates with drift rates in the linear ballistic accumulator [Brown, S. D., & Heathcote, A. The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, *57*, 153–178, 2008]. We found that distance to the classifier boundary correlated with RT, accuracy, and drift rates in an animacy categorization task. Split by animacy, the correlations between brain and behavior were sustained longer over the time course for animate than for inanimate stimuli. Interestingly, when examining the distance to the classifier boundary during the peak correlation between brain and behavior, we found that only degraded versions of animate, but not inanimate, objects had systematically shifted toward the classifier decision boundary as predicted. Our results support an asymmetry in the representation of animate and inanimate object categories in the human brain. ■

INTRODUCTION

Object recognition is a fast, reliable, and effortless process for humans. Early visual areas in the brain respond to simple visual features (e.g., edges, luminance contrast, or orientation), and further along the ventral stream, sensitivity to objects and object categories (e.g., faces, animals, or tools) emerges (Grill-Spector & Weiner, 2014; DiCarlo & Cox, 2007). Using multivariate pattern analysis (MVPA), several studies have analyzed pattern similarities between the neural representation of objects in inferior temporal cortex (ITC) to study its representational structure, with both fMRI (Haxby, Connolly, & Guntupalli, 2014; Kriegeskorte & Kievit, 2013; Kriegeskorte et al., 2008; Edelman, Grill-Spector, Kushnir, & Malach, 1998) and MEG (Contini, Wardle, & Carlson, in press; Cichy, Pantazis, & Oliva, 2014, 2016; Carlson, Tovar, Alink, & Kriegeskorte, 2013). These studies have provided evi-

dence of a categorical organization in ITC, after the observation that objects belonging to the same category tend to evoke similar patterns of neural activation (Sha et al., 2015; Connolly et al., 2012; Kriegeskorte et al., 2008; Haxby et al., 2001). Time series decoding studies have shown that such abstract object representations emerge later than simple visual features in the time course of object processing (Cichy et al., 2014, 2016; Kaneshiro, Guimaraes, Kim, Norcia, & Suppes, 2015; Carlson et al., 2013), consistent with a hierarchical organization of the visual stream. One robust categorical structure is the animate/inanimate distinction in human and primate ITC (Mahon & Caramazza, 2011; Kriegeskorte et al., 2008; Kiani, Esteky, Mirpour, & Tanaka, 2007; Caramazza & Shelton, 1998; Spelke, Phillips, & Woodward, 1995). Animate stimuli (e.g., humans or animals) evoke brain activation patterns that are more similar to other animate exemplars than to inanimate stimuli (e.g., plants, tools, vehicles; Cichy et al., 2014; Carlson et al., 2013; Kriegeskorte et al., 2008; Kiani et al., 2007; Downing, Jiang, Shuman, & Kanwisher, 2001).

A current topic of debate is whether categorical information in the brain, as revealed using MVPA decoding, is read out in behavior (cf. de-Wit, Alexander, Ekroll,

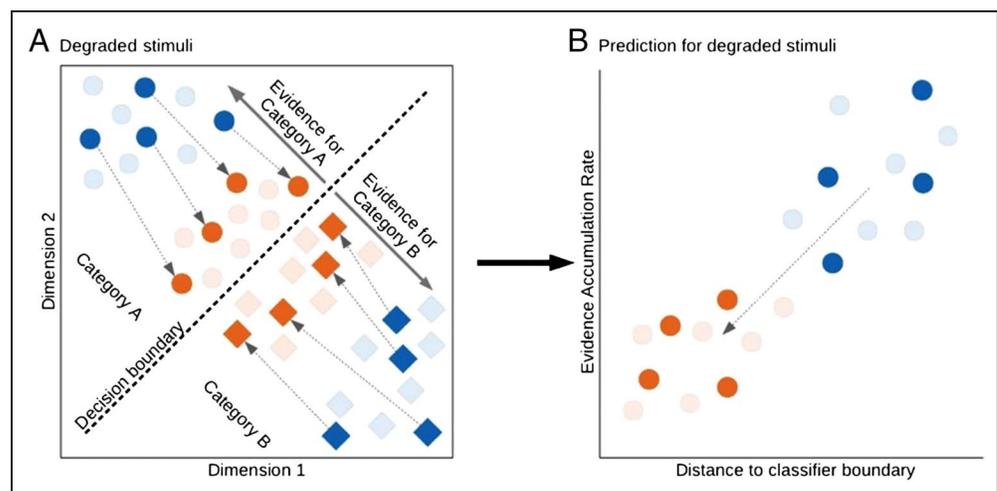
¹Macquarie University, Australia, ²ARC Centre of Excellence in Cognition and Its Disorders, Australia, ³University of Sydney, ⁴KU Leuven, ⁵University of Tasmania, ⁶University of Newcastle, Australia

& Wagemans, 2016; Williams, Dang, & Kanwisher, 2007). The presence of decodable information in neuroimaging activation patterns related to a stimulus or task does not necessarily entail that this information underlies related behavior. One recent approach to linking activation spaces to behavior (Ritchie & Carlson, 2016) is inspired by distance-to-bound models of RT (Ashby & Maddox, 1994; Pike, 1973). According to distance-to-bound models, evidence close to a decision boundary is more ambiguous, reflecting greater difficulty in categorization, whereas evidence far from the decision boundary is less ambiguous with regard to category membership. Assuming that RT is a function of stimulus discriminability and that a classifier decision boundary (which is used in MVPA decoding) reflects an observer's decision boundary, then RTs should negatively correlate with distance from the boundary; for example, stimuli that are faster to categorize should be neurally represented as further from the classifier decision boundary (Ritchie & Carlson, 2016).

In a previous application of the RT–distance approach to MEG (Ritchie, Tovar, & Carlson, 2015), a linear discriminant classifier was trained to discriminate the MEG channel activation for animate stimuli from that for inanimate stimuli. Next, the distance of each stimulus pattern to the classifier boundary in high-dimensional space was rank-order correlated to human RTs for categorizing the same stimuli as animate/inanimate. As predicted, the distance to boundary negatively correlated with RT. Moreover, the correlation over time tracked the MEG decoding time series (Ritchie et al., 2015). Interestingly, when analyzing the animate and inanimate stimuli separately, an asymmetry was observed: The correlation between distance to boundary and RT was driven by animate stimuli (Ritchie et al., 2015; Carlson, Ritchie, Kriegeskorte, Durvasula, & Ma, 2014).

A useful test of the RT–distance hypothesis would be to manipulate task difficulty experimentally and observe the effects of this behavioral manipulation on representational space. Note that, here, we use the term “representational space” under the assumption that neural representations are to some extent reflected in neuroimaging activation patterns (see, e.g., Kriegeskorte & Bandettini, 2007; Carlson, Schrater, & He, 2003; Haxby et al., 2001). The RT–distance hypothesis has so far been tested on the differences between objects. For example, in Carlson et al. (2014), an ostrich was closer to the animacy decision boundary than a human face, and participants were slower to categorize the ostrich as “animate.” The effect of increasing categorization difficulty of single-object images on the distance to the classifier boundary in activation space has not yet been tested. According to the RT–distance hypothesis, a change in behavior resulting from manipulating categorization task difficulty should be matched by a corresponding shift of the stimulus set in representational space. Numerous studies have shown that degrading object stimuli reduces categorization performance, such as by scrambling the image phase (e.g., Philastides, Ratcliff, & Sajda, 2006; Philastides & Sajda, 2006; Wichmann, Braun, & Gegenfurtner, 2006), scrambling image amplitude (e.g., Gaspar & Rousselet, 2009), reducing luminance contrast (e.g., Macé, Delorme, Richard, & Fabre-Thorpe, 2009; Macé, Thorpe, & Fabre-Thorpe, 2005), or blurring the image (e.g., Párraga, Troscianko, & Tolhurst, 2000, 2005; Bruner & Potter, 1964; Wyatt & Campbell, 1951). For degraded stimuli with longer categorization RTs, the RT–distance hypothesis predicts that these stimuli will be located closer to the animacy decision boundary, producing a “compression” of representational space compared with that for the original versions of the stimuli (Figure 1A). This predicts a correlation

Figure 1. The predicted effect of degrading stimuli on their location in representational space. (A) Stimuli of two categories are illustrated as circles (animate objects) and squares (inanimate objects) in representational space (only two dimensions are plotted here for visualization). The blue shapes represent the stimuli in a clear state, and the orange are their degraded counterparts. If distance to a classifier boundary is taken as representing evidence for a decision, this predicts that the degraded (orange) versions of the stimuli will be located closer to the classifier decision boundary



(dashed line) that separates the stimulus categories than the clear (blue) versions. (B) Stimuli from one category (animate objects) and their distance to the classifier boundary versus their rate of evidence accumulation for an animacy categorization decision. The RT–distance hypothesis predicts that the degraded stimuli, which have moved closer to the decision boundary, also have slower evidence accumulation rates (and therefore longer RTs). Note that the same prediction holds for the other category (inanimate objects).

between the shorter distance to boundary and slower accumulation rates (Figure 1B). Note that, even in the case of inanimate stimuli where there was no correlation (Ritchie et al., 2015), one would still expect a correlation when including both clear and degraded versions because of the compression of the general representational space (Figure 1B).

The dominant models of RT appeal to some form of evidence accumulation process (Ashby, 2000; Ratcliff, 1985); that is, evidence for a decision (e.g., animate or inanimate) accumulates over time, and the response is made when the amount of evidence reaches a certain threshold (Brown & Heathcote, 2008; Gold & Shadlen, 2007; Ratcliff & Rouder, 1998). Distance to the classifier boundary can be linked to evidence accumulation. Carlson et al. (2014) simulated evidence accumulation with a sequential analysis model, using the distance to the boundary for each object exemplar as a proxy for evidence strength. They found that accumulation rate correlated with categorization RTs, providing support for the theoretical link between distance to boundary and evidence accumulation (Ritchie et al., 2015; Carlson et al., 2014). One way to build on this would be to more closely relate distance to evidence accumulation; beyond correlating distance with median RTs, an existing model of evidence accumulation can be fit to participants' behavioral data, yielding independent model parameter estimates that can be correlated with distance to the boundary. As an accumulator model provides a more complete characterization of categorization behavior than average RTs, it may provide a better measure to correlate to representational distances.

The aim of this study was twofold. First, using MEG decoding, we sought to test the prediction that degrading object exemplar would compress the representational space (which we assume is reflected in the high-dimensional MEG activation space). This would in turn correlate with slower animacy RTs compared with undegraded or clear versions of the same stimuli (Figure 1). Second, we aimed to test the RT–distance hypothesis in the context of an existing model of RT distributions and choice accuracy, the linear ballistic accumulator (LBA; Brown & Heathcote, 2008), to evaluate whether distance to boundary can be related more directly to evidence accumulation model parameters.

METHODS

Participants

All participants gave informed consent in writing before the experiment. The study was conducted with the approval of the Macquarie University human research ethics committee. One hundred participants were recruited on Amazon's Mechanical Turk (MTurk) to determine the level of stimulus degradation needed for equal object recognition performance across all stimuli. Another group

of 40 new MTurk participants were asked to rate the stimuli on a 7-point scale from typical animate to typical inanimate. For the second part of the study, 20 healthy volunteers (four men; mean age = 29.3 years) with normal or corrected-to-normal vision participated in the MEG experiment. Participants in both experiments were financially compensated for their time. All analysis procedures were performed in MATLAB (The MathWorks, Natick, MA) using the statistics and machine learning toolbox. For reference purposes, the code is freely available at github.com/Tijl/Grootswagers_etal_degraded_objects.

Stimuli

We constructed a set of 48 visual object stimuli including 24 animals and 24 inanimate objects (natural and man-made) on a phase-scrambled natural image background in both a clear condition and a degraded condition. First, high-resolution images ($>512 \times 512$ pixels) of various objects were collected via an Internet search. We only selected images that included prototypical viewpoints of objects. To test the RT–distance hypothesis more generally, no human or human face images were included in the stimulus set as they are generally outliers both behaviorally and in the brain's response compared with other object exemplars. That is, face stimuli tend to have fast categorization RTs and produce pronounced responses in neuroimaging data. It is therefore possible that the inclusion of face stimuli could disproportionately explain any observed correlations between RT and neural distance. All exemplar images used in the study are shown in Figure 2A. As color can be a salient cue for image recognition (Joseph & Proffitt, 1996; Wurm, Legge, Isenberg, & Luebker, 1993; Biederman & Ju, 1988; Ostergaard & Davidoff, 1985) and would make the degrading process (described in the next section) less effective, gray-scaled versions of the object images were used in the experiment. A different random noise background was created for each exemplar by phase scrambling a natural image of a forest scene (see Figure 2B). The object images were overlaid on the noise background, producing the stimuli for the clear condition (Figure 2C).

We created a degraded condition by blurring the same set of images (including their background). As different objects require different levels of blur to equate recognition performance, we first measured the amount of blur required to impair recognition for each image. Our aim was for participants to be able to perform the task given unlimited exposure time (i.e., correctly recognize the object) but to reduce their categorization performance under brief presentation duration (i.e., by reducing speed and/or accuracy). We simulated defocused blur using an image filter (Figure 3) that convolved the amplitude spectrum of the image in the Fourier domain (Figure 3A and B) with a Fourier-transformed cylinder function (a sombrero function; Figure 3C and D), where increasing the radius of the

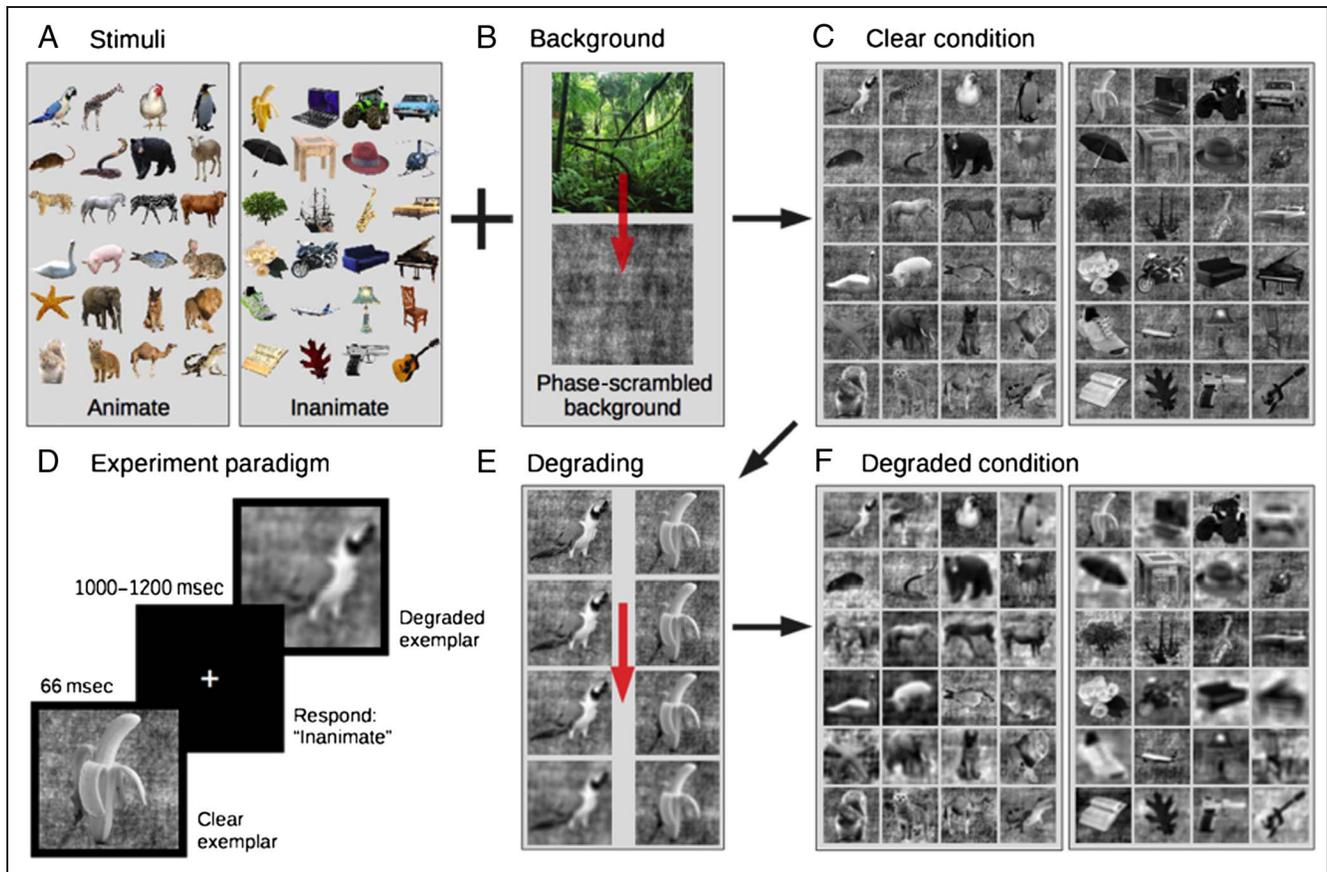


Figure 2. Stimuli and experiment design. (A) Stimuli consisted of 24 animate and 24 inanimate objects. (B) Stimuli were placed on a phase-scrambled natural image background. (C) All stimuli in the clear condition. (D) In the MEG experiment, participants saw clear and degraded stimuli for 66 msec in randomized order with a varying ISI and were asked to report the stimulus category (animate or inanimate) with a button press. (E) To create the degraded condition, stimuli were gradually blurred by simulating defocus to a level where they were equally recognizable. (F) All stimuli in the degraded condition.

cylinder function results in a greater magnitude (Figure 3E) of defocus blur (Sonka, Hlaváč, & Boyle, 1993). Images were then gradually degraded by increasing the radius in steps of 2 pixels, from a radius of 1 pixel (no degrading) to 59 pixels (very degraded). The sequence of the image coming into focus was presented to the MTurk participants. Each participant saw all 48 stimuli from both animate and inanimate categories once starting from the most degraded state, while its level of focus was gradually increased (in steps of 2-pixel radius). Participants were instructed to press the spacebar as soon as they recognized the object in the picture. The stimulus was then removed from the screen, and participants entered a name for the stimulus. To check for correct recognition, the responses on the naming task were assessed manually for validity, to allow for variations in spelling or for synonymous names. For each exemplar, the amount of focus needed for 25% of the MTurk participants to correctly recognize (i.e., name) the object (see Figure 2F) was used as the blur filter parameters for that exemplar in the degraded condition. On average, a radius of 17 pixels ($SD = 4.68$ pixels) was used for the animate exemplars and 20.5 pixels ($SD = 10.12$ pixels) for the inanimate exemplars.

MEG Experiment Design

Before the MEG experiment, we confirmed that each participant could recognize all exemplars given unlimited presentation time, even in the degraded state, as the intention of the design was to decrease categorization speed and accuracy without making the stimuli unrecognizable. Stimuli in the degraded state were shown, and participants were asked to name the object in each picture. If the participant failed to correctly recognize the stimulus, it was shown in the clear state to ensure that all objects were correctly recognized. Next, the participant was trained on the task (outside the MEG), as the brief presentation duration and fast pace of the categorization task required practice to master. If performance (animacy categorization accuracy) in the first block was lower than 80%, the participant was shown all the stimuli again in both states to identify degraded exemplars they were unable to recognize and then practiced again on a second training block. All 20 participants performed above 80% correct after the familiarization step.

After the practice task, participants completed the MEG experiment. On each trial within a block, stimuli were

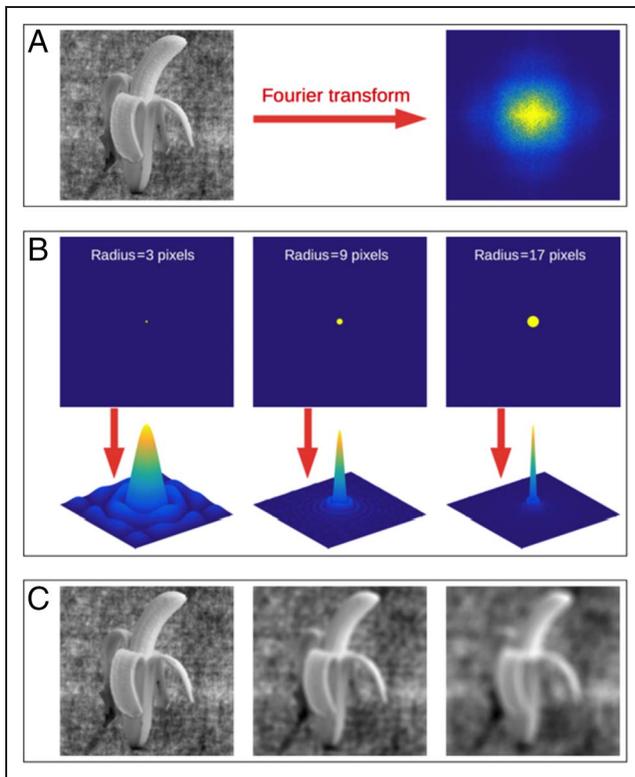


Figure 3. Image filter used to blur the experimental stimuli. (A) The original image and its amplitude spectrum, which shows the typical energy pattern for natural images (high energy at low spatial frequencies [center] and low energy at high spatial frequencies; Field, 1987). (B) Cylinder functions with radii of 3, 9, and 17 pixels and their Fourier-transformed versions (sombbrero functions), which are used as the image filters. (C) Convolving the amplitude spectrum of the original image with the sombrero functions in the Fourier domain results in images with different levels of defocus blur by removing a significant proportion of energy at high spatial frequencies.

projected (at $9^\circ \times 9^\circ$ visual angle) on a black background for 66 msec, followed by a fixation cross for a random duration between 1000 and 1200 msec. Participants were asked to categorize the stimulus as animate or inanimate as fast and accurate as possible using a button press (see also Figure 2D). To avoid possible motor response confounds, the mapping of the response buttons alternated between blocks (Grootswagers, Wardle, & Carlson, 2017; Ritchie et al., 2015), and participants received feedback for the first 10 trials of a block (red or green cross) to ensure accurate mapping of the buttons. Within each block, four repetitions of each exemplar in both conditions were presented in randomized order. Block duration was ~ 7 min. After each block, participants received feedback on performance (mean accuracy and number of missed trials). Each participant completed eight blocks, resulting in 32 trials per exemplar, 768 trials per category and condition (animate/inanimate, clear/degraded), and 3072 total trials per participant. The total time in the scanner was about 1 hr, including breaks between blocks.

MEG Acquisition and Preprocessing

Participants were fitted with a cap with five marker coils to track head movement during the session. The MEG signal was continuously sampled at 1000 Hz from 160 axial gradiometers using a whole-head MEG system (model PQ1160R-N2; KIT, Kanazawa, Japan) while participants lay in a supine position inside a magnetically shielded room (Fujihara Co. Ltd., Tokyo, Japan). Recordings were filtered online between 0.03 and 200 Hz. We examined the delay in stimulus onsets (Ramkumar, Jas, Pannasch, Hari, & Parkkonen, 2013) by comparing the photodiode responses with the stimulus onset triggers (sent by the experiment script) and found a highly consistent delay of 56.26 msec, for which we conservatively corrected by shifting the onset triggers back by 56 msec. Recordings were sliced into 700-msec epochs (-100 to 600 msec poststimulus onset). The trials were downsampled to 200 Hz (5-msec resolution) and transformed using principal component analysis, where the components that accounted for 99% of the variance were retained to reduce the dimensionality of the data (mean = 62.25 components, $SD = 12.33$ components). Finally, to increase signal-to-noise, four trials of each exemplar (with balanced response mappings) were averaged into pseudotrials (Grootswagers et al., 2017; Isik, Meyers, Leibo, & Poggio, 2014), leaving eight pseudotrials per exemplar in both conditions.

Sliding Time Window Decoding

To investigate the decoding performance of animacy over time, sliding window Naïve Bayes classifiers were used on the pseudotrials. To assess the difference in decoding performance between the clear and degraded conditions, three separate classifiers were used for decoding, one for each condition and one for both conditions combined. At each 5-msec interval t , the classifiers were trained and tested on a 25-msec window (from $t-25$ msec to t). The classifier performance was examined using leave-one-exemplar-out cross-validation (Carlson et al., 2013). In this method, the classifier is trained on the animacy of all but one exemplar and tested on trials of the left-out exemplar. This is repeated for each exemplar, and the mean decoding accuracy on the left-out exemplars is used to assess generalization accuracy. Using this cross-validation method, the classifier has to generalize the concept of animacy and cannot benefit from exploiting individual stimulus properties because the test exemplar is not in the training set (Carlson et al., 2013). We report the participant-averaged classifier accuracy over time, with significant above-chance accuracies assessed using a nonparametric Wilcoxon signed rank test. The false discovery rate (FDR) was used to control for false positives resulting from multiple comparisons.

Fitting LBA on Individual Participant Behavior

One of the aims of this study was to use a more complex model of evidence accumulation and then to correlate

distance to boundary with the drift rate parameters of the model as well as RT and accuracy. The LBA is a mathematically tractable yet simple and complete model of evidence accumulation (Brown & Heathcote, 2008). LBA simultaneously takes into account both the full distribution of RTs (i.e., the mean, variability, and positive skew of the RT distribution; see Luce, 1986) and task accuracy. LBA differs from other well-known accumulation models (e.g., the diffusion decision model [Ratcliff, 1978] or the leaky competitive accumulator [Usher & McClelland, 2001]) by modeling the response alternatives (here, animate and inanimate) with separate independent accumulators that accrue evidence accumulation in a linear and ballistic manner. In particular, the LBA model assumes that between-trial variability in the amount of evidence required for a decision, and the rate at which it accumulates, dominates over any moment-to-moment variability in evidence during a trial, whereas the latter source of noise plays a greater role in alternative models. These properties make LBA analytically simple and relatively easy to apply, which is ideal for this study, where each stimulus is treated as a separate condition, meaning that there are many parameters to estimate, each based on relatively few data points. To simplify the fitting procedure, the stimuli for both categories were separated into six bins based on accuracy. Next, a set of progressively more complex parameterizations were fit stepwise using maximum likelihood estimation (Rae, Heathcote, Donkin, Averell, & Brown, 2014; Donkin, Brown, & Heathcote, 2011). The most complex model parameterization, which was used here for further analysis, included clear/degraded, animate/inanimate, and stimulus as factors. Hence, each stimulus in both clear and degraded conditions had a separate drift rate, which is important to note, as the goal was to correlate stimulus-specific distances with drift rates. Figure 4 shows that the LBA model provided a very accurate model of the effect of stimulus degradation and of differences between animate and inanimate images, both in terms of accuracy (Figure 4A) and the entire distribution of RTs (Figure 4B), and so provides a useful characterization of the participants' behavior.

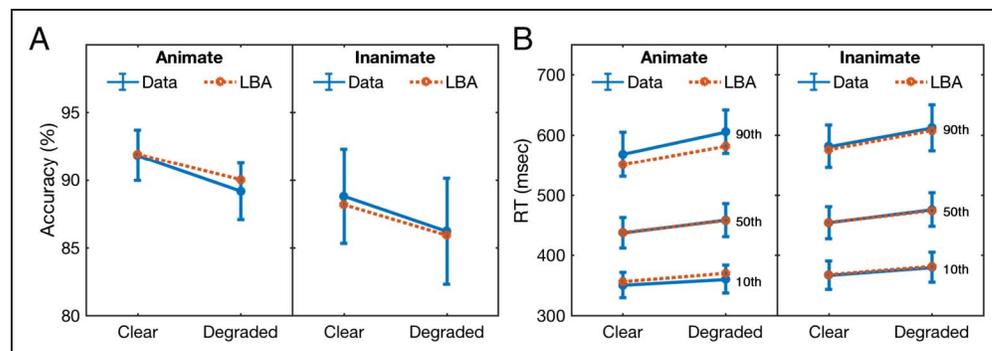
Predicting Behavior from Representational Distance

Individual exemplars can be represented as points in representational space (i.e., a multidimensional feature space). To decode animacy, we applied a discriminant classifier (Gaussian Naïve Bayes) to optimize a decision boundary in multidimensional space, separating the neural patterns for animate and inanimate exemplars. According to the RT–distance hypothesis, the neural representations of the exemplars that are close (in multidimensional space) to this boundary are predicted to be more difficult to discriminate, as there is less evidence for the decision. This forms a prediction for behavior, where less evidence for a decision would result in slower RTs or lower accuracies (Ritchie & Carlson, 2016; Ritchie et al., 2015; Carlson et al., 2014) and, in the case of the LBA decision model, slower drift rates. For the current study, we predicted that degraded stimuli would be closer to the decision boundary in representational space and therefore correlate with lower accuracy, slower RTs, and slower drift rates.

We tested this prediction by repeating the following process at each time window: First, all the trials for each exemplar were averaged to create an average representation in multidimensional space for each exemplar (one for the exemplar in clear state and one for degraded state). In this space, a decision boundary for animacy was fitted (i.e., training a Naïve Bayes classifier), and the representational distance to this boundary for each exemplar was computed. This resulted in a distance value for each exemplar in both clear and degraded states (24 animate and 24 inanimate exemplars in 2 states = 92 distance values). Next, the exemplar distances were rank-order correlated (using Spearman's ρ) to the two behavioral measures (median RT and mean accuracy) and the mean drift rate for the exemplars.

Repeating this process over time and participants resulted in three time-varying correlations (two for each behavioral measure and one for drift rate) for each participant. We then report the participant-averaged time-varying correlations and used a nonparametric Wilcoxon signed rank test at each time window to test for significant above-zero correlations at the group level. Note that

Figure 4. Fits of the LBA model to (A) accuracy and (B) the distribution of RT, with 95% confidence intervals. The RT distribution is illustrated by plotting the 90th percentile (upper lines, representing the slowest RTs), the median (50th percentile, middle lines), and the 10th percentile (lower lines, representing the fastest RTs).



this approach differs from Ritchie et al. (2015), where distances and RTs were averaged over participants first and the correlations were performed at the group level. FDR adjustment was used to control for false positives resulting from multiple comparisons. This process allowed us to compare the distance to boundary correlations with the three variables. We also sought to assess the possibility of modulating effects of animacy, for example, whether the animate–inanimate asymmetry reported in Carlson et al. (2014) and Ritchie et al. (2015) was replicated and whether it was affected by degrading the stimuli. For this, we computed the time-varying correlations separately for animate and inanimate exemplars.

RESULTS

Behavioral Results

Overall, participants performed well on the categorization task (mean accuracy = 87.9%, $SD = 12.7\%$) with a median RT of 457.5 msec ($SD = 64.6$ msec). Figure 5 shows the median RT (Figure 5A) and mean accuracy (Figure 5B) separate for category (animate vs. inanimate) and stimulus condition (clear vs. degraded). ANOVA showed that animate exemplars were easier to categorize than inanimate exemplars (compare left vs. right groups in each plot), with significantly faster median RTs ($F(1, 19) = 32.09, p < .0001$) and higher accuracies ($F(1, 19) = 6.27, p < .0001$). As predicted, degrading the stimuli made the animacy categorization task more difficult (compare blue and yellow lines), significantly lowering accuracy ($F(1, 19) = 41.81, p < .0001$) and increasing median RT ($F(1, 19) = 85.06, p < .0001$) for both animate and inanimate exemplars. There was no significant interaction between animacy and degrading for either median RT ($F(1, 19) = 0.04, p = .84$) or accuracy ($F(1, 19) = 0.01, p = .92$).

To test whether distance is related to evidence accumulation, we obtained drift rates for each exemplar and participant from the LBA fits. The LBA produces a separate drift rate parameter for the correct accumulator

(i.e., the animate accumulator for animate stimuli and the inanimate accumulator for inanimate stimuli) and the incorrect accumulator (i.e., the inanimate accumulator for animate stimuli and vice versa); however, only the drift rates for the correct response accumulators were included in our analysis. The resulting drift rate parameters are summarized in Figure 5C, split by animacy and degradation. Degrading the stimuli resulted in significantly slower drift rates ($F(1, 19) = 41.20, p < .0001$). There was a main effect of animacy on drift rate ($F(1, 19) = 4.44, p < .05$) and a significant interaction between animacy and stimulus clarity on drift rate ($F(1, 19) = 18.52, p < .0001$), as degrading reduced the drift rate for animate exemplars more than inanimate exemplars (see Figure 5C). In summary, we confirmed that the degraded exemplars were generally harder to categorize (in terms of longer RT, lower accuracy, and slower drift rate), demonstrating that our stimulus categorization difficulty manipulation (blur) was successful.

MEG Decoding

A prediction of degrading the stimuli is that decoding performance will be lower relative to that for clear stimuli. To evaluate the effect of blurring the stimuli on decoding performance over time, sliding time window classifiers were trained on predicting stimulus animacy at each time point. The results are presented in Figure 6A as mean cross-validated decoding accuracy over participants, separately for clear and degraded stimuli. The decoding onset in the clear condition is at 75 msec, which is consistent with previous findings showing an animacy decoding onset of approximately 60–80 msec (Ritchie et al., 2015; Carlson et al., 2013). In the degraded condition, decoding onset was 170 msec, and decoding performance was significantly lower than that in the clear condition over the entire time course (black marks above the x axis). Peak decoding performance for degraded objects was also later, 380 msec compared with 345 msec for clear stimuli. Initially, decoding performance for the

Figure 5. Behavioral results. The distributions of (A) the median RT, (B) accuracy, and (C) drift rate for the correct response accumulator, split up by exemplar category (animate/inanimate) and condition (clear/degraded, blue and orange lines). Error bars represent $\pm SEM$.

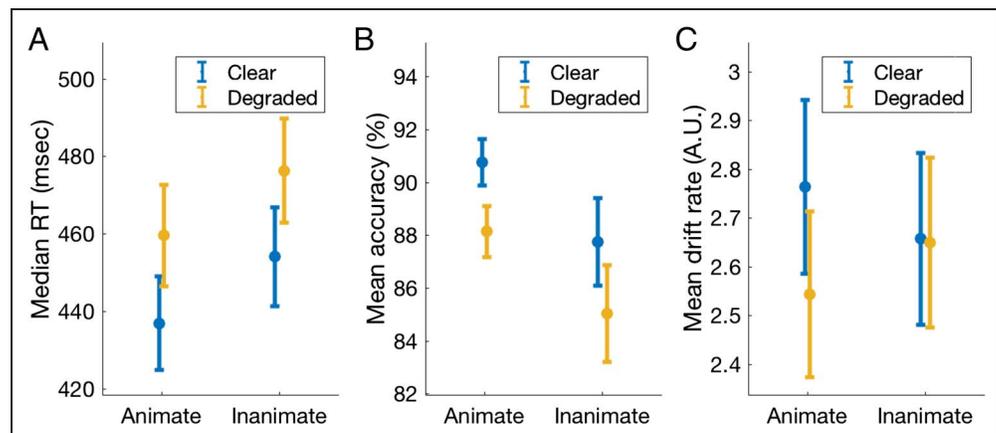
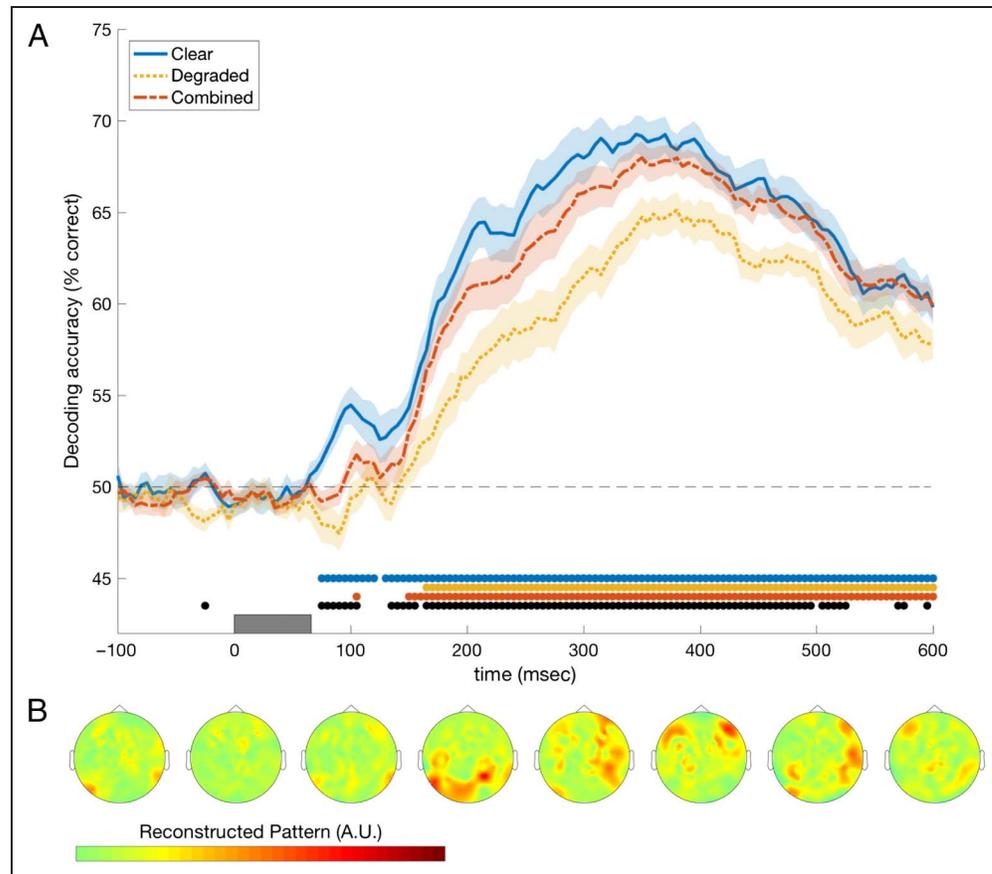


Figure 6. Decoding animacy from the MEG signal. (A) Decoding was performed using leave-one-exemplar-out cross-validation with 25-msec sliding time window classifiers. At each time point t , the graph shows the mean classifier accuracy over participants at the window $[t-25 \text{ msec}, t]$. Shaded areas show standard error between participants. Colored marks above the x axis indicate significant above-chance (50%) decoding. Black marks indicate significant (FDR-adjusted $p < .05$) differences in classifier accuracy between the clear and degraded conditions. The gray bar on the x axis indicates the time that the stimulus was on the screen (0–66 msec). (B) The mean classifier weights at 100-msec intervals over cross-validation folds and over participants for the combined condition. Weights were transformed into activation patterns before projection onto the sensors (Haufe et al., 2014).



combined data for clear and degraded conditions closely matched that for the degraded condition but then raised to a level closer to the performance in the clear condition. In summary, the neural patterns for blurred stimuli were more difficult to decode from the whole-brain MEG activation patterns than those for the clear versions of the same stimuli.

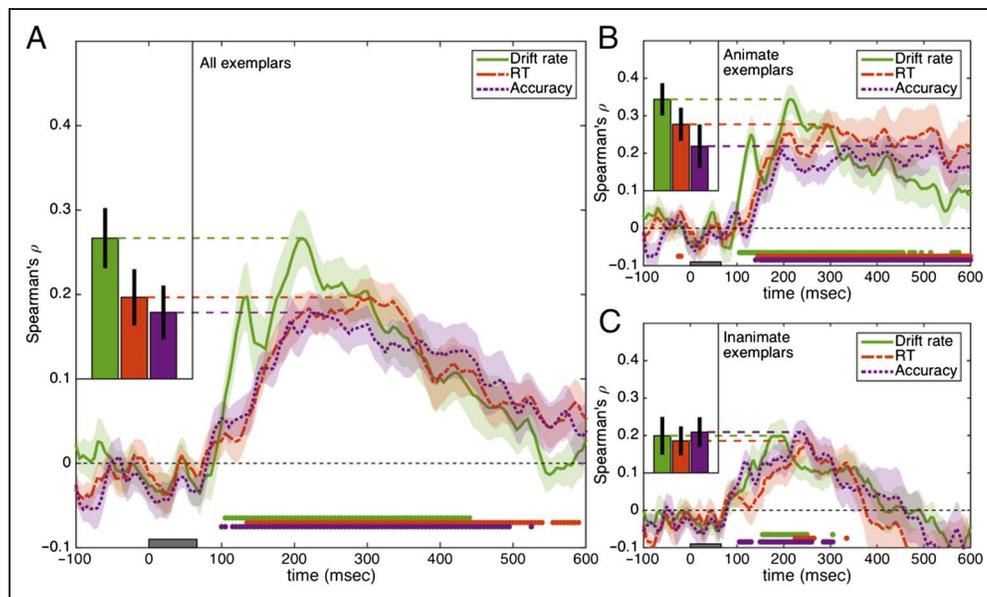
Predicting Behavior from Representational Distance

To investigate the relationship between decodability and behavior, we computed the time-varying distance to the classifier decision boundary for all the exemplars for each participant. These participant-specific distances were then rank-order correlated to the participant's behavioral measures for each exemplar: RT, accuracy, and the drift rate as fitted by LBA. Note that, for accuracy and drift rate, the RT–distance hypothesis predicts a positive correlation (closer to the boundary corresponds to lower accuracy and slower drift rate); however, a negative correlation is predicted between RT and distance (closer to the boundary corresponds to longer RTs). For ease of comparison, the sign of the correlation for RT was inverted in Figure 7.

Figure 7A shows the mean time-varying correlations over participants for drift rate (green line), RT (red line), and accuracy (purple line). The peaks of the time-varying correlations are shown in the inset bar graphs. Although drift rate appears to have an earlier and higher peak correlation than RT and accuracy, this difference was not statistically significant. The similarity between these time-varying correlations is likely due to correlations among the different behavioral measures (e.g., exemplars with fast RTs are likely to have high accuracies). These results further show that accuracy and drift rate can be predicted equally well using distance to boundary. Although the LBA model has been used before to fit nonhuman primate neural activations (Cassey, Heathcote, & Brown, 2014), this is the first time it has been related to neuroimaging data measured with MEG. Thus, our results are promising considering that drift rate more closely represents the accumulation of evidence for a decision, compared with RT or accuracy. In summary, distance to boundary correlates with the behavioral measures as well as the fitted drift rate parameters, and these follow similar trajectories.

Having established that distance to boundary predicts RT, accuracy, and drift rate for animacy categorization, we next investigated the relative contributions of animate and inanimate exemplars. In Ritchie et al. (2015), time-varying correlations for the inanimate exemplars were

Figure 7. The correlation between distance to boundary, behavioral variables, and drift rate. (A) The distance model applied to all the exemplars. (B) Results for animate exemplars only. (C) Results for inanimate exemplars. Three measures (RT, accuracy, and drift rate parameter) were correlated with distance to boundary over time. RT refers to inverted normalized RT. Accuracy is the mean accuracy for exemplars, and drift rate was estimated by fitting LBA. Shaded areas refer to the standard error over participants, and colored marks above the x axis indicate significant above-zero correlations. The peak correlations with their standard errors are compared in the inset (bars, top left). The gray bar on the x axis indicates the time the stimulus was on the screen (0–66 msec).



not significant. Thus, the effect was specific to animate exemplars, and the same asymmetry was also reported in Carlson et al. (2014). In this study, however, because of the degrading of stimuli, we also predicted a correlation for inanimate exemplars (when using both clear and degraded exemplars), as we predicted that the representations of degraded inanimate exemplars would still shift toward the classifier decision boundary (Figure 1). Figure 7 shows the result of computing the time-varying correlation separately for animates (Figure 7B) and inanimates (Figure 7C). Animate exemplars reach higher correlations on all behavioral measures and follow the same trends as seen in Figure 7A for the combined stimulus set. In contrast, inanimate exemplars have lower correlations overall and were less sustained over time. However, significant above-chance correlations between distance and RT, accuracy, and drift rate were still present for the inanimate exemplars. In addition, more sustained correlations (more significant time points) are present for drift rate and accuracy than for RT, for the inanimate exemplars (which is of interest, as previous studies only used RT).

Comparing the Decoding Time Courses with the Time Course of Predicting Behavior

Ritchie et al. (2015) found that the time-varying correlation with behavior matched the time-varying classifier decoding performance. To examine whether this relationship was present in the current study, the time-varying correlations with behavior (RT, accuracy, and drift rate) were rank-order correlated to the time-varying decoding

result. The results are presented in Figure 8 and show significant correlations between the results from Figure 7 and the decoding trajectory (Figure 6A), rising and falling at approximately the same time. When comparing the animate and inanimate trajectories separately (Figure 7B and C), only the animate time-varying correlation matches decoding performance. These results are consistent with the findings from Ritchie et al. (2015) who also found higher correlations between RT correlation and decoding trajectories for the animate exemplars. In addition, the decoding trajectory also correlates significantly with the accuracy results from Figure 7 and is significant for both animate and inanimate exemplars. Correlations with the drift rate parameter are also significant but appear lower in magnitude. This is evident when comparing Figures 6 and 7, where drift rate has a dual-peak structure, which differs from the trajectories for the RT and accuracy correlations.

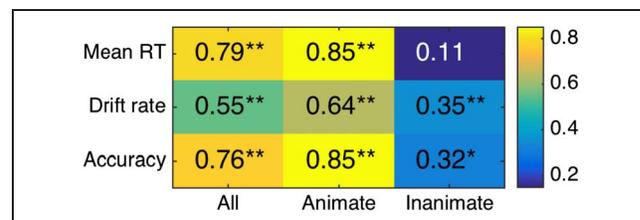


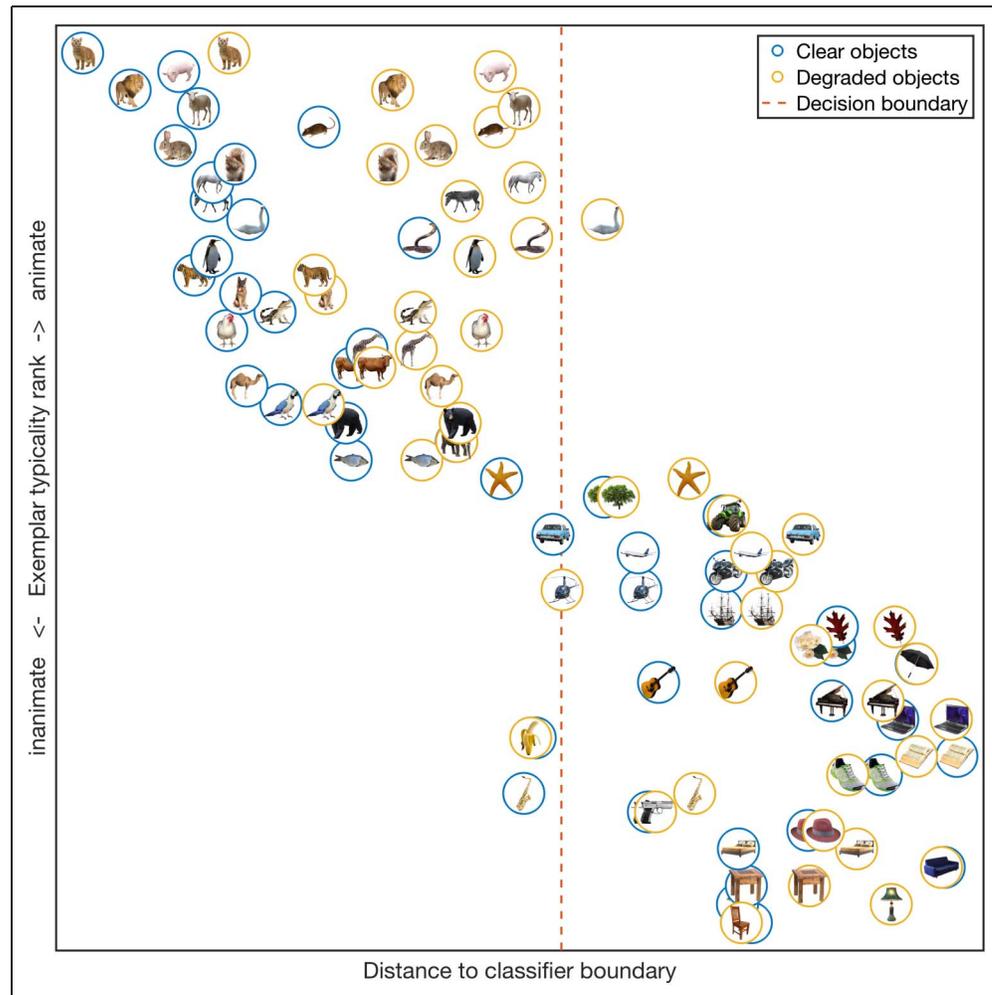
Figure 8. Similarity (Spearman's ρ) between the decoding trajectories (Figure 6A) and correlation trajectories from Figure 7. The time-varying decoding performance (for all clear and degraded stimuli combined) and time-varying correlations were rank-order correlated. High values indicate similar trajectories (e.g., matching rises, falls, and peaks). Asterisks indicate significant correlations (* $p < .01$, ** $p < .001$).

Compression of Degraded Objects in Representational Space

The RT–distance hypothesis predicts that the neural patterns for degraded objects (which have slower RTs and lower accuracies; see Figure 5) will be closer to the category boundary in representational space, as the distance to the boundary represents the degree of evidence for category membership (Figure 1). The results above show that classifier accuracy is lower for degraded objects, indicating that degraded items are harder to separate in representational space (Figure 6A). Next, we showed that distance to boundary correlates with all three behavioral measures (Figure 7A), and this correlation is mostly driven by animate exemplars where it is sustained over time (Figure 7B), with some significant correlations for inanimates across a more limited portion of the time course (Figure 7C). Together, these results seem to favor our prediction based on the RT–distance hypothesis: Degraded objects (both animate and, to some extent, inanimate) are closer to the boundary than their clear counterparts, and this shift in representational space correlates with RT, accuracy, and drift rate.

To visually compare our results with the predictions as shown in Figure 1A, we plotted all exemplars in both clear and degraded states relative to the decision boundary in representational space (at the time of peak correlation between drift rate and distance to boundary [210 msec]; Figure 7A) in Figure 9. On the y axis, the stimuli are ordered by their typicality ratings. Comparing Figures 1A and 9 suggests that the prediction that degraded stimuli are located closer to the category boundary than clear stimuli holds only for animate, but not inanimate, exemplars. This result is consistent with previous results from Ritchie et al. (2015) and Carlson et al. (2014), which showed no clear relationship between distance to boundary and RT for inanimate stimuli. Although even if the basic relationship between behavior and distance to the classifier boundary does not hold for inanimate exemplars, it is still surprising that there is also no relationship between degrading inanimate stimuli and where they are located relative to the category boundary. This asymmetric compression of representational space is however consistent with the interaction between animacy and degrading on the fitted drift rates (Figure 5C), which suggests that there is little or no

Figure 9. Reconstruction of the representational space for animacy decoding. The locations of exemplars in their clear state are plotted in blue circles; and their degraded counterparts, in yellow circles. The x axis represents distance from an object to the boundary and is the mean distance to the boundary over all participants at 210 msec. The exemplars are ordered on the y axis according to their typicality rating obtained from 40 participants on MTurk. Note that some objects (e.g., the degraded starfish) have a negative mean distance and are thus placed on the opposite side of the boundary. Note that degraded animate exemplar representations are in general closer to the linear decision boundary than the clear versions, demonstrating compression of representational space for animate but not inanimate objects.



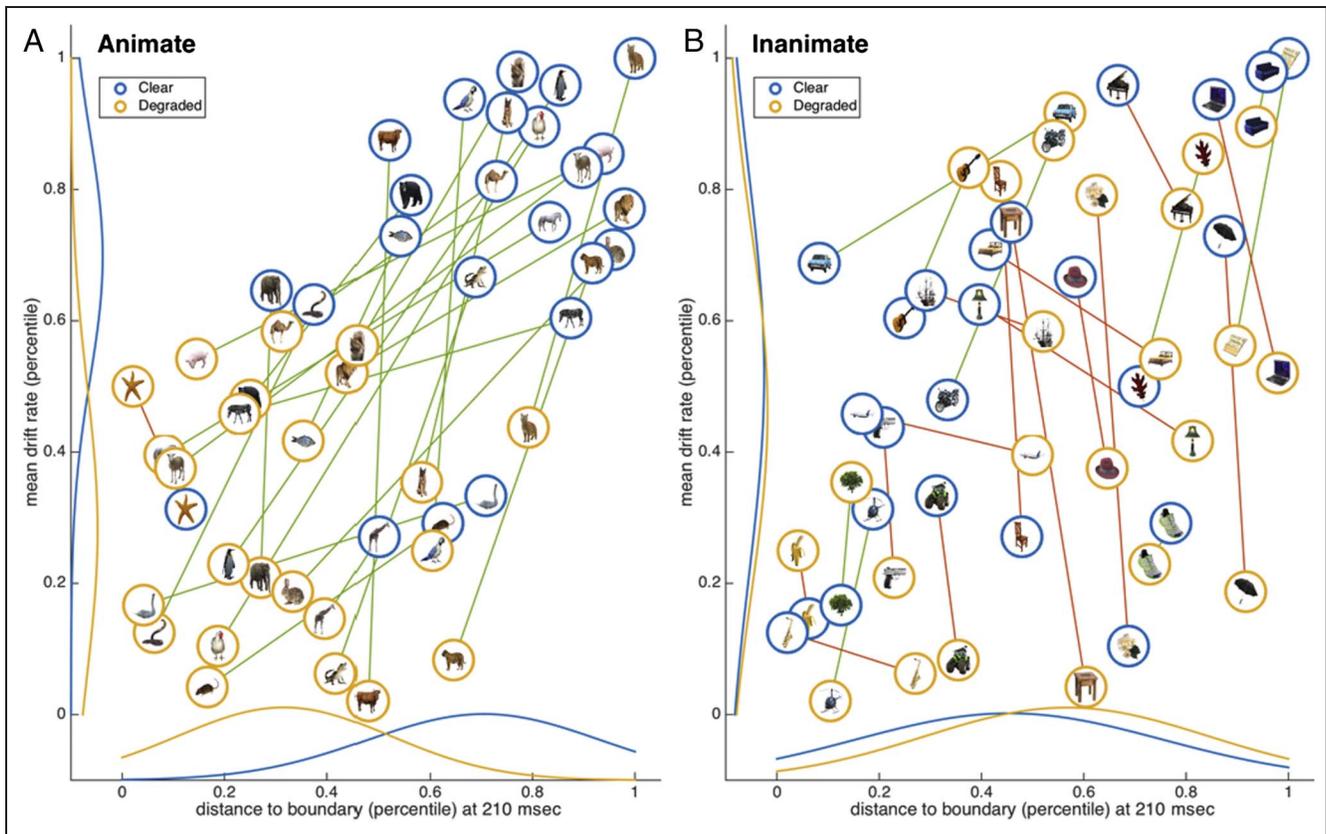


Figure 10. The effect of degrading on the relationship between drift rate and distance to boundary for animate (A) and inanimate (B) exemplars. Items were rank-ordered for mean drift rate over participants and mean distance to boundary at the peak LBA prediction time (210 msec; see Figure 7A). Here, the *y* axis represents the decision boundary. Blue circles are the exemplars in their clear versions; and yellow circles, degraded versions. Lines connect the two versions of each exemplar and are colored green if the RT–distance hypothesis correctly predicts the direction of the relationship (i.e., the degraded versions of the objects are closer to the boundary and have a slower drift rate) and red if not. The distributions of the exemplars on the variables are shown on the axes.

effect of degrading on the representation of inanimate exemplars.

The compression of exemplar distances toward the animacy boundary is also predicted to match slower drift rates (as shown in Figure 1B). To compare our results with this second prediction, we plotted the mean drift rates over participants against the mean distance to boundary over participants for each exemplar in both states in Figure 10 separately for animate (Figure 10A) and inanimate (Figure 10B) exemplars. The result is shown at the time of peak correlation between drift rate and distance to boundary (210 msec; Figure 7A). The exemplars are plotted as a function of their mean distance to boundary (*x* axis) and mean drift rate (*y* axis). The RT–distance hypothesis predicts that exemplars with a higher drift rate are located further away from the boundary. Lines in Figure 10 connect the two states of each exemplar and are colored green if this prediction is true for each exemplar. This is the case for most of the animate exemplars (Figure 10A), but (not surprisingly, considering the asymmetry in Figure 9) for only a few of the inanimate exemplars (Figure 10B).

DISCUSSION

Our aim was to test the RT–distance hypothesis in the context of a specific prediction, namely, that degrading objects will result in a compression of representational space and that the shift in representational space for degraded objects will match impaired categorization behavior. The results showed that degraded object images produced slower RTs, lower accuracies, and slower evidence accumulation rates in an animacy categorization task (Figure 5). When examining the distances of individual objects, we observed an asymmetric compression: Only the degraded animate objects had moved closer toward the classifier boundary. Unexpectedly, there was not a corresponding consistent shift toward the boundary for degraded inanimate exemplars. In addition, we found that distance to boundary correlates with drift rate, which is a measure that is more closely related to the decision process than descriptive statistics (mean RT or accuracy).

In this study, all MEG channels were included for training and testing the classifier, as is the standard for many

MEG decoding studies (e.g., Grootswagers et al., 2017; Kaiser, Azzalini, & Peelen, 2016; Cichy et al., 2014; Carlson et al., 2013). A limitation of this approach is that the spatial source of the decodable signal is unknown. Because participants were performing an animacy task, it is not possible to entirely exclude a contribution from the decision process (e.g., originating in frontal executive areas) as a possible source of the decodable signal. However, previous research has found no difference in decoding performance between active animacy categorization versus a distractor task (Ritchie et al., 2015). In addition, the representational structure in the MEG response to visual objects has been found to correspond best to fMRI representations in the ventral visual stream, showing that the early and late MEG responses after stimulus onset match the fMRI responses in V1 and IT, respectively (Cichy et al., 2014, 2016). Moreover, examining the classifier weights (Figure 6B) suggests that the decodable signal originates from occipital sensors at 200 msec and more temporal at 300 msec. This suggests that the most likely sources of the decodable signal in our study originate from areas in the ventral visual stream.

Asymmetric Effects of Stimulus Degradation on the Neural Representation of Animacy

We found asymmetrical effects for animate and inanimate objects. Correlations between distance and behavior (RT/accuracy/drift rate) in this study were driven by animate stimuli. Inanimate stimuli had smaller and less sustained correlations between distance and behavior and did not show compression toward the classifier boundary in representational space for degraded versions of the stimuli. Previous studies also found that correlations between distance and RT were almost exclusively driven by animate stimuli (Ritchie et al., 2015; Carlson et al., 2014), which is consistent with our results. Our study differed in many aspects from earlier studies, for example, by using grayscale stimuli on a controlled background and excluding human or human face stimuli. Human faces/bodies have faster RTs (Crouzet, Kirchner, & Thorpe, 2010) and highly decodable neural responses (Kaneshiro et al., 2015; Carlson et al., 2013; Kriegeskorte et al., 2008) and could therefore potentially significantly drive the RT–distance correlations in previous studies. However, we have shown that these findings are robust when human faces/bodies are omitted (also note that correlations were calculated within participants, rather than on the pooled group means as in previous work). Carlson et al. (2014) argue for a conceptual difference between the two animacy categories, suggesting that “inanimate” is not an equivalent category to animate, for example, because inanimate is negatively defined (i.e., as “not animate”) and is less restricted than the animate category. Furthermore, a clear hierarchical subdivision (e.g., vertebrate–invertebrate) has been reported only within the animate category (Kiani et al., 2007).

We observed that degraded versions of animate objects had compressed toward the boundary, consistent with the RT–distance hypothesis. The lack of compression for the inanimate side of the boundary, and the absence of strong correlations between inanimate categorization behavior and distance to boundary, suggests that, although the animate–inanimate distinction is highly decodable, it does not sufficiently capture the structure of brain representations linked to responding “inanimate.” Future work could explore this issue using a different categorization task where both categories are similarly constrained (e.g., faces vs. tools) or use a different model of the animacy task (e.g., model it as an animal detection task, instead of animacy categorization). However, it is still possible that responses to other category dichotomies will be based on one category versus “not” that category. For example, a processing bias for faces means that they are easier to recognize than tools (cf. Wu, Crouzet, Thorpe, & Fabre-Thorpe, 2015) and an effective face/tool categorization strategy would be to simply try to detect whether or not a stimulus is a face. In addition, outside the laboratory, objects are categorized effectively without having to exhaustively test all two-way categorization combinations. Thus, instead of using dichotomous categorization, applying a different task to examine the RT–distance relationship might yield new insights, for example, by using go/no-go tasks (Crouzet et al., 2010; Kirchner & Thorpe, 2006; Thorpe, Fize, & Marlot, 1996).

The amount of compression toward the decision boundary was different between individual exemplars (Figure 9). Although the degraded stimuli were equated for recognizability in an object-naming task, some showed a larger displacement toward the boundary than others (e.g., compare difference between clear and degraded versions of the fish and the sheep in Figure 9). The naming task may be more difficult than the categorization task, and although they likely rely on the same underlying representation (Riesenhuber & Poggio, 2000), different amounts of evidence may be required for naming versus categorization. For example, some animals might be equally easy to name, but when making an animacy decision, some animals are more typically animate than others, which likely makes them easier to categorize than less typical animals. Typicality ratings for our stimuli indeed increased with distance to boundary, but again only for the animate stimuli (Figure 9). Animacy categorization is known to be influenced by typicality (Rosch & Mervis, 1975; Rosch, 1973; Posner & Keele, 1968). Typical exemplars have also been found to be more decodable (Jordan, Greene, Beck, & Fei-Fei, 2016). For example, mammals such as cat, tiger, and squirrel are all far from the boundary and have matching fast RTs and high drift rates. Conversely, the fish is closer to the boundary and is possibly less typically conceived of as “animate,” for example, because fish moves and behaves differently than mammals. An extreme example is the starfish, which is the animal closest to the boundary in the clear condition and is on the wrong side of the

boundary (i.e., consistently predicted “inanimate” by the classifier) in the degraded condition, suggesting that it is hard to categorize it as animate (see Figure 9). We found that participants mostly categorized the starfish as inanimate and that, when asked, some reported that they do not consider it an animate object. Note that, to observe a correlation between distance to the boundary and RTs, distances and RTs have to systematically differ between exemplars in the same category. Our results thus provide some support for the presence of a more continuous than dichotomous neural representation of animacy in the brain (cf. Sha et al., 2015; Connolly et al., 2012).

Further evidence in support of an animacy asymmetry comes from our observation that inanimate stimuli in general needed to be blurred more to equate their recognizability to animate exemplars (see Methods section). This suggests that animate objects are more homogeneous than inanimate objects; thus, there is greater variance in the “inanimate” than “animate” category (which could be caused by the lack of categorical structure in inanimates). This is consistent with animate objects in general sharing more features than inanimate objects (Garrard, Ralph, Hodges, & Patterson, 2001; McRae, de Sa, & Seidenberg, 1997). The blur filter removes details from the stimuli, but object shape is preserved. It could be the case that the general shape within the broad category of animate objects is more similar and provides more alternatives (cf. Bracci & Op de Beeck, 2016). Although some higher-level animal subdivisions might be visually dissimilar (e.g., compare fish with birds), subgroups often share similar shapes. For example, outside the laboratory, to recognize a blurry zebra (e.g., without glasses), likely alternatives that have similar shapes (e.g., horse, deer, moose) need to be eliminated. In contrast, viewing a blurry piano provides less alternatives, which share the same shape. Together, this supports the notion that inanimate and animate are not equivalent categories, which is consistent with patient studies that found selective deficits in recognition of animate or inanimate objects (Caramazza & Mahon, 2003; Caramazza & Shelton, 1998). Future research could explore whether homogeneity of the category determines the effect of degrading stimuli, by using the distance-to-bound approach with more homogeneous groups of inanimate objects, such as tools or fruits, or restricting the animate category to only domestic animals.

Drift Rate Is Predicted by Distance to Boundary

The drift rate parameter from the LBA model more directly reflects the speed of evidence accumulation and builds on previous work that used RT as a proxy for the decision process (Ritchie et al., 2015; Carlson et al., 2014). We found that correlations with LBA drift rate were on par with those for RT and that there were some qualitative differences in the trajectories. The correlation between LBA drift rate and distance peaked earlier than

RT and had an earlier onset—although this should be interpreted with caution as earlier onsets can be caused by stronger signal-to-noise rather than true underlying differences (Grootswagers et al., 2017). Moreover, drift rate followed a different trajectory than the category decoding trace (as seen in Figure 8), which suggests that it may capture a different part of the (neural) decision process. It is interesting that the drift rate trajectory does have a dual-peak structure, which resembles previous results of MEG animacy decoding (Ritchie et al., 2015; Cichy et al., 2014; Carlson et al., 2013) but was not present in our MEG decoding results. This dual-peak structure may suggest that, at some point (in the period between the peaks), distance to boundary is used to a lesser extent for forming the categorization decision. Alternatively, the decision could have already been made after the first peak, as categorization happens very fast (Crouzet et al., 2010; Kirchner & Thorpe, 2006; Thorpe et al., 1996), and the second peak reflects, for example, a feedback process that has the same representational structure (i.e., the same distances). Taken together, it is sensible to relate distance directly to evidence accumulation using drift rate as it is more closely related to decision processes and therefore may successfully relate to a wider range of behavior than RT alone. Future research could explore this further by using tasks that allow for more variance in RT and accuracy (e.g., incentivizing different speed–accuracy trade-offs). Moreover, as we showed that it is possible to link one decision model with neural distance to boundary, this could be tested with other models of decision-making, such as exemplar-based models of choice (Nosofsky & Stanton, 2005) or their LBA-based extension (Donkin & Nosofsky, 2012).

The Neural Dynamics of Visual Object Categorization

We found that significant decoding performance of clear stimuli started around 70 msec and peaked at 345 msec, occurring later for degraded stimuli (165 and 380 msec, respectively). Note that differences in decoding onsets have to be interpreted with caution, as a later onset can be a result of lower overall decoding performance or differences in variance (Grootswagers et al., 2017; Isik et al., 2014). Still, we observed that the first local maximum in the decoding trace for the clear objects was absent for the degraded objects (Figure 6A). This difference suggests that some information in the early response is predictive of animacy in the clear condition. As the timing of the early peak corresponds to early to midlevel visual areas (Cichy et al., 2016; VanRullen & Thorpe, 2001; Thorpe et al., 1996), the predictive information in the first local maximum in the clear decoding condition could reflect low-level visual information (e.g., high spatial frequencies) that is removed by degrading the stimuli (Kirchner & Thorpe, 2006).

We found that the peak correlation between distance and drift rate occurred at 210 msec and the onset of significant correlations with drift rate was at 100 msec (Figure 7). These values did not match the onset or peak of decoding, suggesting that the optimal time for readout does not necessarily correspond with the time that the information can be best decoded from the signal. In contrast, Ritchie et al. (2015) found correlations during the whole period of significant decoding. A possible explanation for this difference is that the fast-paced task and short stimulus duration (66 msec compared with 500 msec in Ritchie et al., 2015) in the current study may have promoted a faster readout of animacy by exploiting low-level visual cues (Hong, Yamins, Majaj, & DiCarlo, 2016; Kirchner & Thorpe, 2006; Thorpe et al., 1996). This would, in addition, explain why exemplars such as banana and helicopter, which have more rounded shapes than other inanimate objects, are closer to the boundary (Figure 9).

Conclusions

In this study, we tested whether representational space is compressed when degrading stimuli and whether this matches behavior in a dichotomous categorization task. We found that degrading stimuli made them harder to categorize and that this was accompanied by a compression of representational space, as predicted by the RT-distance hypothesis. This compression was only observed for animate stimuli, suggesting an asymmetry in the neural representation of animacy. Moreover, we showed that neural distance to boundary can be directly related to a current model of evidence accumulation (LBA) as the fitted drift rates from this model correlated with distance to the boundary. Connecting linear classifiers to models of the decision processes is a step toward relating brain imaging to behavior, a fundamental and complex challenge in cognitive neuroscience (de-Wit et al., 2016; Purcell & Palmeri, 2016; Forstmann & Wagenmakers, 2015).

Acknowledgments

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

Reprint requests should be sent to Tijn Grootswagers, Department of Cognitive Science, Australian Hearing Hub, 16 University Avenue, Macquarie University, NSW 2109, Australia, or via e-mail: t.grootswagers@gmail.com.

REFERENCES

- Ashby, F. G. (2000). A stochastic version of general recognition theory. *Journal of Mathematical Psychology*, *44*, 310–329.
- Ashby, F. G., & Maddox, W. T. (1994). A response time theory of separability and integrality in speeded classification. *Journal of Mathematical Psychology*, *38*, 423–466.
- Biederman, I., & Ju, G. (1988). Surface versus edge-based determinants of visual recognition. *Cognitive Psychology*, *20*, 38–64.
- Bracci, S., & Op de Beeck, H. (2016). Dissociations and associations between shape and category representations in the two visual pathways. *Journal of Neuroscience*, *36*, 432–444.
- Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, *57*, 153–178.
- Bruner, J. S., & Potter, M. C. (1964). Interference in visual recognition. *Science*, *144*, 424–425.
- Caramazza, A., & Mahon, B. Z. (2003). The organization of conceptual knowledge: The evidence from category-specific semantic deficits. *Trends in Cognitive Sciences*, *7*, 354–361.
- Caramazza, A., & Shelton, J. R. (1998). Domain-specific knowledge systems in the brain: The animate-inanimate distinction. *Journal of Cognitive Neuroscience*, *10*, 1–34.
- 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*, 132–142.
- Carlson, T. A., Schrater, P., & He, S. (2003). Patterns of activity in the categorical representations of objects. *Journal of Cognitive Neuroscience*, *15*, 704–717.
- Carlson, T. A., Tovar, D. A., Alink, A., & Kriegeskorte, N. (2013). Representational dynamics of object vision: The first 1000 ms. *Journal of Vision*, *13*, 1.
- Cassey, P., Heathcote, A., & Brown, S. D. (2014). Brain and behavior in decision-making. *PLoS Computational Biology*, *10*, e1003700.
- Cichy, R. M., Pantazis, D., & Oliva, A. (2014). Resolving human object recognition in space and time. *Nature Neuroscience*, *17*, 455–462.
- 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*, 3563–3579.
- Connolly, A. C., Guntupalli, J. S., Gors, J., Hanke, M., Halchenko, Y. O., Wu, Y.-C., et al. (2012). The representation of biological classes in the human brain. *Journal of Neuroscience*, *32*, 2608–2618.
- Contini, E. W., Wardle, S. G., & Carlson, T. A. (in press). Decoding the time-course of object recognition in the human brain: From visual features to categorical decisions. *Neuropsychologia*. <https://doi.org/10.1016/j.neuropsychologia.2017.02.013>.
- Crouzet, S. M., Kirchner, H., & Thorpe, S. J. (2010). Fast saccades toward faces: Face detection in just 100 ms. *Journal of Vision*, *10*, 16.
- de-Wit, L., Alexander, D., Ekroll, V., & Wagemans, J. (2016). Is neuroimaging measuring information in the brain? *Psychonomic Bulletin & Review*, *23*, 1415–1428.
- DiCarlo, J. J., & Cox, D. D. (2007). Untangling invariant object recognition. *Trends in Cognitive Sciences*, *11*, 333–341.
- Donkin, C., Brown, S., & Heathcote, A. (2011). Drawing conclusions from choice response time models: A tutorial using the linear ballistic accumulator. *Journal of Mathematical Psychology*, *55*, 140–151.
- Donkin, C., & Nosofsky, R. M. (2012). A power-law model of psychological memory strength in short- and long-term recognition. *Psychological Science*, *23*, 625–634.
- Downing, P. E., Jiang, Y., Shuman, M., & Kanwisher, N. (2001). A cortical area selective for visual processing of the human body. *Science*, *293*, 2470–2473.

- Edelman, S., Grill-Spector, K., Kushnir, T., & Malach, R. (1998). Toward direct visualization of the internal shape representation space by fMRI. *Psychobiology*, *26*, 309–321.
- Field, D. J. (1987). Relations between the statistics of natural images and the response properties of cortical cells. *Journal of the Optical Society of America A*, *4*, 2379–2394.
- Forstmann, B. U., & Wagenmakers, E.-J. (2015). *An introduction to model-based cognitive neuroscience*. New York: Springer.
- Garrard, P., Ralph, M. A. L., Hodges, J. R., & Patterson, K. (2001). Prototypicality, distinctiveness, and intercorrelation: Analyses of the semantic attributes of living and nonliving concepts. *Cognitive Neuropsychology*, *18*, 125–174.
- Gaspar, C. M., & Rousselet, G. A. (2009). How do amplitude spectra influence rapid animal detection? *Vision Research*, *49*, 3001–3012.
- Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, *30*, 535–574.
- Grill-Spector, K., & Weiner, K. S. (2014). The functional architecture of the ventral temporal cortex and its role in categorization. *Nature Reviews Neuroscience*, *15*, 536–548.
- Grootswagers, T., Wardle, S. G., & Carlson, T. A. (2017). Decoding dynamic brain patterns from evoked responses: A tutorial on multivariate pattern analysis applied to time series neuroimaging data. *Journal of Cognitive Neuroscience*, *29*, 677–697.
- Haufe, S., Meinecke, F., Görgen, K., Dähne, S., Haynes, J.-D., Blankertz, B., et al. (2014). On the interpretation of weight vectors of linear models in multivariate neuroimaging. *Neuroimage*, *87*, 96–110.
- Haxby, J. V., Connolly, A. C., & Guntupalli, J. S. (2014). Decoding neural representational spaces using multivariate pattern analysis. *Annual Review of Neuroscience*, *37*, 435–456.
- 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*, 2425–2430.
- Hong, H., Yamins, D. L. K., Majaj, N. J., & DiCarlo, J. J. (2016). Explicit information for category-orthogonal object properties increases along the ventral stream. *Nature Neuroscience*, *19*, 613–622.
- Jordan, M. C., Greene, M. R., Beck, D. M., & Fei-Fei, L. (2016). Typicality sharpens category representations in object-selective cortex. *Neuroimage*, *134*, 170–179.
- 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*, 91–102.
- Joseph, J. E., & Proffitt, D. R. (1996). Semantic versus perceptual influences of color in object recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 407.
- Kaiser, D., Azzalini, D. C., & Peelen, M. V. (2016). Shape-independent object category responses revealed by MEG and fMRI decoding. *Journal of Neurophysiology*, *115*, 2246–2250.
- 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*, e0135697.
- 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*, 4296–4309.
- Kirchner, H., & Thorpe, S. J. (2006). Ultra-rapid object detection with saccadic eye movements: Visual processing speed revisited. *Vision Research*, *46*, 1762–1776.
- Kriegeskorte, N., & Bandettini, P. (2007). Analyzing for information, not activation, to exploit high-resolution fMRI. *Neuroimage*, *38*, 649–662.
- Kriegeskorte, N., & Kievit, R. A. (2013). Representational geometry: Integrating cognition, computation, and the brain. *Trends in Cognitive Sciences*, *17*, 401–412.
- Kriegeskorte, N., Mur, M., Ruff, D. A., Kiani, R., Bodurka, J., Esteky, H., et al. (2008). Matching categorical object representations in inferior temporal cortex of man and monkey. *Neuron*, *60*, 1126–1141.
- Luce, R. D. (1986). *Response times: Their role in inferring elementary mental organization*. New York: Oxford University Press.
- Macé, M. J.-M., Delorme, A., Richard, G., & Fabre-Thorpe, M. (2009). Spotting animals in natural scenes: Efficiency of humans and monkeys at very low contrasts. *Animal Cognition*, *13*, 405–418.
- Macé, M. J.-M., Thorpe, S. J., & Fabre-Thorpe, M. (2005). Rapid categorization of achromatic natural scenes: How robust at very low contrasts? *European Journal of Neuroscience*, *21*, 2007–2018.
- Mahon, B. Z., & Caramazza, A. (2011). What drives the organization of object knowledge in the brain? *Trends in Cognitive Sciences*, *15*, 97–103.
- McRae, K., de Sa, V. R., & Seidenberg, M. S. (1997). On the nature and scope of featural representations of word meaning. *Journal of Experimental Psychology: General*, *126*, 99.
- Nosofsky, R. M., & Stanton, R. D. (2005). Speeded classification in a probabilistic category structure: Contrasting exemplar-retrieval, decision-boundary, and prototype models. *Journal of Experimental Psychology: Human Perception and Performance*, *31*, 608.
- Ostergaard, A. L., & Davidoff, J. B. (1985). Some effects of color on naming and recognition of objects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *11*, 579.
- Párraga, C. A., Troscianko, T., & Tolhurst, D. J. (2000). The human visual system is optimised for processing the spatial information in natural visual images. *Current Biology*, *10*, 35–38.
- Párraga, C. A., Troscianko, T., & Tolhurst, D. J. (2005). The effects of amplitude-spectrum statistics on foveal and peripheral discrimination of changes in natural images, and a multi-resolution model. *Vision Research*, *45*, 3145–3168.
- Philiastides, M. G., Ratcliff, R., & Sajda, P. (2006). Neural representation of task difficulty and decision making during perceptual categorization: A timing diagram. *Journal of Neuroscience*, *26*, 8965–8975.
- Philiastides, M. G., & Sajda, P. (2006). Temporal characterization of the neural correlates of perceptual decision making in the human brain. *Cerebral Cortex*, *16*, 509–518.
- Pike, R. (1973). Response latency models for signal detection. *Psychological Review*, *80*, 53–68.
- Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, *77*, 353.
- Purcell, B. A., & Palmeri, T. J. (2016). Relating accumulator model parameters and neural dynamics. *Journal of Mathematical Psychology*, *76*, 156–171.
- Rae, B., Heathcote, A., Donkin, C., Averell, L., & Brown, S. (2014). The hare and the tortoise: Emphasizing speed can change the evidence used to make decisions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *40*, 1226–1243.
- Ramkumar, P., Jas, M., Pannasch, S., Hari, R., & Parkkonen, L. (2013). Feature-specific information processing precedes concerted activation in human visual cortex. *Journal of Neuroscience*, *33*, 7691–7699.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*, 59–108.
- Ratcliff, R. (1985). Theoretical interpretations of the speed and accuracy of positive and negative responses. *Psychological Review*, *92*, 212.

- Ratcliff, R., & Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science, 9*, 347–356.
- Riesenhuber, M., & Poggio, T. (2000). Models of object recognition. *Nature Neuroscience, 3*, 1199–1204.
- 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*, 190.
- Ritchie, J. B., Tovar, D. A., & Carlson, T. A. (2015). Emerging object representations in the visual system predict reaction times for categorization. *PLoS Computational Biology, 11*, e1004316.
- Rosch, E. (1973). On the internal structure of perceptual and semantic categories. Retrieved from psycnet.apa.org/psycinfo/1974-20611-001.
- Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology, 7*, 573–605.
- Sha, L., Haxby, J. V., Abdi, H., Guntupalli, J. S., Oosterhof, N. N., Halchenko, Y. O., et al. (2015). The animacy continuum in the human ventral vision pathway. *Journal of Cognitive Neuroscience, 27*, 665–678.
- Sonka, M., Hlavác, V., & Boyle, R. (1993). *Image processing, analysis and machine vision*. London: Chapman & Hall.
- Spelke, E. S., Phillips, A., & Woodward, A. L. (1995). Infants’ knowledge of object motion and human action. In D. Sperber, D. Premack, & A. J. Premack (Eds.), *Causal cognition: A multidisciplinary debate* (pp. 44–78). New York: Clarendon Press/Oxford University Press.
- Thorpe, S., Fize, D., & Marlot, C. (1996). Speed of processing in the human visual system. *Nature, 381*, 520–522.
- Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review, 108*, 550.
- VanRullen, R., & Thorpe, S. J. (2001). The time course of visual processing: From early perception to decision-making. *Journal of Cognitive Neuroscience, 13*, 454–461.
- Wichmann, F. A., Braun, D. I., & Gegenfurtner, K. R. (2006). Phase noise and the classification of natural images. *Vision Research, 46*, 1520–1529.
- Williams, M. A., Dang, S., & Kanwisher, N. G. (2007). Only some spatial patterns of fMRI response are read out in task performance. *Nature Neuroscience, 10*, 685–686.
- Wu, C.-T., Crouzet, S. M., Thorpe, S. J., & Fabre-Thorpe, M. (2015). At 120 msec you can spot the animal but you don’t yet know it’s a dog. *Journal of Cognitive Neuroscience, 27*, 141–149.
- Wurm, L. H., Legge, G. E., Isenberg, L. M., & Luebker, A. (1993). Color improves object recognition in normal and low vision. *Journal of Experimental Psychology: Human Perception and Performance, 19*, 899.
- Wyatt, D. F., & Campbell, D. T. (1951). On the liability of stereotype or hypothesis. *Journal of Abnormal and Social Psychology, 46*, 496–500.