Identifying the neural dynamics of category decisions with computational model-based fMRI

Emily M. Heffernan⁺, Juliana D. Adema⁺, Michael L. Mack^{*}

Department of Psychology, University of Toronto, Toronto, ON, Canada

⁺equal contribution

***Corresponding author:**

Michael L. Mack
100 St. George Street
Toronto, ON, CA M5S 3G3
michael.mack@utoronto.ca

Manuscript type: Brief report

Abstract: 126 words

Main text (less Methods): 2888 words

Acknowledgments: The authors thank Meg Schlichting for helpful discussions, Jason Chow and Ellie Croce for helping with preliminary analyses, and Alison Preston and Bradley Love for providing access to the dataset through OSF. This work was supported by the Natural Sciences and Engineering Research Council (NSERC; Discovery Grant RGPIN-2017-06753 to M.L.M), the Canada Foundation for Innovation and Ontario Research Funds (JELF 36601 to M.L.M.), a Brain Canada Future Leaders in Canadian Brain Research Grant (M.L.M), and NSERC CGS-M Fellowships (E.M.H. and J.D.A.).

Open Practices Statement: Data from this work is available as an open-access dataset: <https://osf.io/62rgs/>. This study was not preregistered.

Abstract

Successful categorization requires a careful coordination of attention, representation, and decision making. Comprehensive theories that span levels of analysis are key to understanding the computational and neural dynamics of categorization. Here, we build on recent work linking neural representations of category learning to computational models to investigate how category decision making is driven by neural signals across the brain. We uniquely combine functional magnetic resonance imaging with drift diffusion and exemplar-based categorization models to show that trial-by-trial fluctuations in neural activation from regions of occipital, cingulate, and lateral prefrontal cortices are linked to category decisions. Notably, only lateral prefrontal cortex activation was associated with exemplar-based model predictions of trial-by-trial category evidence. We propose that these brain regions underlie distinct functions that contribute to successful category learning.

Keywords: category learning, categorization, fMRI, computational modelling, drift diffusion modelling

Category learning, our ability to organize our experiences into meaningful concepts that can be leveraged in novel situations, is fundamental to the human experience. Not only are we able to group objects based on basic features such as colour and shape, but we are also capable of learning highly abstract and multivariate categories with relatively little practice. Besides commonplace categories such as “edible” and “friendly” and their antagonistic equivalents, objects can be assembled based on one or several complex perceptual features. Influential category learning models posit that novel objects are categorized according to their relative positions in a multidimensional psychological space populated with known category members in which the distance between objects determines their degree of similarity (Shepard, 1957). While this perceptual space can be composed of an unlimited number of dimensions, it is often the case that only some inform categorization decisions, and the weights of those few dimensions can vary (Nosofsky, 1986; Seger & Miller, 2010).

Although category learning has been studied for decades (Shepard, 1957; Young & Householder, 1938), more recent work has converged on a comprehensive account that formalizes the computational and neural mechanisms underlying successful learning (Zeithamova et al., 2019). While previous accounts of category learning debated the nature of concept representations (e.g., exemplar versus prototype), the multifactorial nature of concept learning has been found to implicate myriad neural systems. For instance, while exemplar-based representations recruit lateral prefrontal cortex (LPFC) and parietal areas (Mack et al., 2013), representations based on category prototypes recruit hippocampus and ventromedial prefrontal cortex (Bowman & Zeithamova, 2018). In addition to similarity-based comparisons of representations, successful learning requires higher-order inferential processes, for example, when faced with novelty or uncertainty invoked by a given stimulus (Paniukov & Davis, 2018).

A comprehensive account must incorporate each of these processes; how these multiple brain systems interact remains to be defined.

Much of this category learning work has focused on the link between predictions of representations formed during learning and how task strategies can influence what is learned by combining sophisticated neural analyses with formal model predictions (Bowman & Zeithamova, 2018; Braunlich & Seger, 2016; Mack, Love, & Preston, 2016; Mack et al., 2013). There is, however, a missing component: the decision-making process itself. Limited work has explored the neural mechanisms that govern how category knowledge, once learned, is applied to novel situations to yield measurable behavioural changes in decision making.

Investigating the neural processes of category decision making requires a computational theory of how such decisions unfold. The Exemplar-Based Random Walk (EBRW) model (Nosofsky & Palmeri, 1997) formalizes category decisions as an evidence accumulation process. Evidence in support of different categories is sampled over time through similarity-based retrieval of category exemplars motivated by the seminal Generalized Context Model (GCM) (Nosofsky, 1986). The key innovation of EBRW is its ability to predict both response probabilities and speed, thereby providing a comprehensive formal account of the behaviour underlying categorization decision making.

Here, motivated by the EBRW framework and recent advances in interrogating brain function with computational theory (Kragel, Morton, & Polyn, 2015), we leverage a combination of computational models in an exploratory approach to identifying neural processes linked to categorization decision making. Our approach significantly extends prior work that has interrogated brain function with categorization models (e.g., Bowman & Zeithamova, 2018; Davis, Goldwater, & Giron, 2017; Mack et al., 2013) by 1) targeting neural response that fully

characterizes decision responses and speeds, and 2) focusing on participant-specific predictions through hierarchical model analyses. Although the EBRW model provides a strong theoretical framework for interpreting potentially informative neural processes, there are currently no analytic approaches for applying EBRW in a hierarchical manner that captures individual differences in neural function in category decision making. As such, we split the primary elements of EBRW—that is, exemplar-based category representations that drive an accumulation of noisy decision evidence—into a two-stage analytic approach to best approximate the formal structure of EBRW: First, we interrogate neural signals related to decision making with a hierarchical variant of the drift diffusion model (DDM; Ratcliff, 1978), a computational model that approximates EBRW’s decision making mechanism. Second, we evaluate the correspondence between the identified brain measures and individually tailored predictions of the GCM, the model that EBRW’s exemplar-based category representations are based on. Since EBRW formalizes that category representations impact decision making through changes to the rate of evidence accumulation (Nosofsky & Palmeri, 1997) and that this sort of category evidence has been correlated with neural function in prefrontal and parietal cortices and striatum (Davis et al., 2017), we focus on the drift rate parameter in the DDM. In essence, we test whether or not constraining drift rate to vary across trials in the same manner as fluctuations in neural signal better captures category decisions, and if so, whether or not that neural signal is related to exemplar-based predictions of category evidence. Although this approach is a departure from a pure EBRW analysis, it does provide a pragmatic solution to the challenge of linking brain and model in a category decision making framework that accounts for response times.

We test the hypothesis that brain activation during a classic categorization task (Mack et al., 2013; Medin & Schaffer, 1978) corresponds to the rate of category evidence accumulation on

a trial-by-trial basis. Given that new experiences modulate activity in a network of brain regions (e.g., Kafkas & Montaldi, 2018), we also test the hypothesis that neural decision signals may vary when category knowledge is generalized to novel relative to previously encountered stimuli. Given the exploratory nature of this approach (Thompson, Wright, & Bissett, 2020), we take a purposefully uninformed view of brain function. Specifically, we interrogate neural response from a parcellation of distinct regions across the whole brain independently identified in a large-scale, data-driven analysis of resting state and task-based interregional connectivity (Schaefer et al., 2018). This approach provides a key first approximation of how to identify neural function underlying complex categorization behaviour through the lens of a formal computational theory.

Methods

The current study leverages a previously published open-access dataset (Mack et al., 2013). This dataset, which includes behavioural and structural and functional magnetic resonance imaging (fMRI) data, was downloaded from OSF (<https://osf.io/62rgs/>). We include most methodological details here, but a full description can be found in the original paper.

Participants

Data from 20 participants were included in the primary analyses (age range of 19–33 years; mean age of 23.5 years; 14 female).

Stimuli

The stimulus set was composed of 16 objects consisting of simple shapes enclosed in a grey, horizontally oriented rectangle (Figure 1). The simple shape varied based on four salient binary-valued features (colour: red or green, shape: circle or triangle, size: large or small, and position: right or left). For each participant, the four features were randomly assigned to the four dimensions defined by the 5/4 category structure (Medin & Schaffer, 1978). This structure is

divided into two categories with the prototype member of category A corresponding to $[0,0,0,0]$ and the prototype member of category B corresponding to $[1,1,1,1]$. Nine objects served as the training items with five for category A and four for category B. The remaining seven objects served as a transfer set.

Procedures

After consent in accordance with the University of Texas Institutional Review Board, participants were instructed that they would be shown simple objects composed of different features and that the task was to learn which object belonged to one of two categories through corrective feedback. Participants performed the training phase of the experiment in a behavioural testing room on a laptop computer. On each training trial, one of the nine training stimuli was displayed for 3.5s and participants made a response to the stimulus' category by pressing one of two labelled keys on the keyboard. Then, a fixation cross was presented for 0.5s, followed by a feedback display that presented the stimulus, the correct category, and whether the participant's response was correct or incorrect for 3.5s. The nine training stimuli were presented 20 times in randomized order during the initial training outside the scanner. Participants also completed an additional four training repetitions inside the MRI scanner during an anatomical scan as a refresher of the training items' category membership.

After training, participants performed the testing phase during fMRI scanning. On each test trial, one of sixteen stimuli (consisting of the nine training stimuli and seven novel transfer stimuli) was displayed for 3.5s and participants made a category response by pressing one of two buttons on an MRI-compatible button box. A fixation cross was then presented for 6.5s. No feedback was provided during the testing phase. The 16 stimuli were presented three times in randomized order during six functional runs for 18 total repetitions per stimulus.

Figure 1. Stimuli and test performance. Stimuli composed of four binary dimensions were split into training (five A and four B items) and transfer sets (example stimulus set shown in top row). Test performance showed typical accuracy (middle) and median reaction time (bottom) performance as in previous reports (e.g., Medin & Schaffer, 1978). Lighter dots depict participant-specific averages, darker dots depict group averages, and errors bars depict 95% confidence intervals.

fMRI Data Acquisition and Preprocessing

Whole-brain imaging data were acquired on a 3.0T GE Signa MRI system (GE Medical Systems). Structural images were acquired using a T2-weighted flow-compensated spin-echo pulse sequence (TR=3s; TE=68ms, 256x256 matrix, 1x1mm in-plane resolution) with thirty-three 3mm-thick oblique axial slices (0.6mm gap), approximately 20° off the AC-PC line. Functional images were acquired with an echo planar imaging sequence using the same slice prescription as the structural images (TR=2s, TE=30.5ms, flip angle=73°, 64x64 matrix, 3.75x3.75 in-plane resolution, bottom-up interleaved acquisition, 0.6mm gap). An additional high-resolution T1-weighted 3D SPGR structural volume (256x256x172 matrix, 1x1x1.3mm voxels) was acquired for registration and brain parcellation.

Anatomical and functional MRI data for each participant were preprocessed using the fMRIPrep automated MRI workflow (version 1.0.15; Esteban et al., 2019), which included brain

extraction, motion correction, coregistration between functional and T1 volumes, and normalization to the MNI 2009c asymmetric brain template. AROMA-identified noise components were regressed out of the functional timeseries. For each run, trial-level beta parameters were estimated from the functional timeseries across the whole brain using the LS-S approach (Mumford, Turner, Ashby, & Poldrack, 2012). This approach provides whole brain voxel-level estimates of BOLD response for each trial across the entire experiment (i.e., the degree of brain activation on each trial for each participant). These beta estimates across trials were averaged within 100 regions of interest (ROIs) as defined by a resting-state and task-based brain parcellation (Schaefer et al., 2018) and an additional eight ROIs from subcortical regions including right and left hippocampus, caudate, putamen, and thalamus.

Figure 2. Brain parcellation and analysis schematic. A) Neural activation from 108 ROIs derived from a data-driven resting-state functional parcellation and anatomically defined subcortical regions were investigated. B) Mean beta estimates within each ROI were extracted and leveraged to predict trial-by-trial changes in drift rate in DDM simulations of categorization decisions.

Brain-informed Drift Diffusion Modelling

DDM analyses were conducted by first averaging beta estimates (i.e., trial-specific estimates of neural activation) within each ROI for each trial resulting in an ROI-specific timeseries (Figure 2B) that characterized how BOLD response within the ROI varied across

trials. Separate DDM simulations were conducted for each ROI wherein trial-by-trial changes in drift rate were linked to the ROI timeseries. Additionally, the relationship between drift rate (v) and ROI activation was allowed to vary by stimulus type (training vs. transfer items). Thus, drift rate was modelled as a linear regression of ROI neural activation, stimulus type, and their interaction ($v \sim ROI + type + ROI:type$). This relationship was evaluated with DDM simulations implemented with the Hierarchical Drift Diffusion Model (HDDM) library (Wiecki, Sofer, & Frank, 2013), which performs hierarchical Bayesian parameter estimation to predict response choices and times. Markov chain Monte Carlo sampling was conducted for 20,000 samples with 10,000 burn-in and thinning set to 2. The deviance information criterion (DIC) for each ROI-based model was compared to a baseline model which included no link between drift rate and neural activation (i.e., drift rate was a free parameter within the constraints of the hierarchical parameter estimation). ROI-based models with smaller DIC values differing from the baseline model by at least 10 were considered a significantly better fit (Spiegelhalter, Best, Carlin, & Van Der Linde, 2002).

Parameter estimates from ROI-based models meeting the DIC criterion were further explored by analyzing the posterior distributions of effects due to neural activation and stimulus type (training vs. transfer items) on drift rate. Specifically, the existence of an effect was quantified with the probability of direction, pd , which is defined as proportion of posterior samples in the most probable direction (i.e., pd ranges from 0.5 to 1 with values closer to 1 for more likely effects). pd is akin to the frequentist p -value and is best interpreted as the degree of evidence against a null effect (Makowski, Ben-Shachar, Chen, & Lüdtke, 2019). Given that we performed 108 independent model fits, the issue of multiple comparisons inflating the familywise error rate had to be considered. As such, we performed a permutation test on each of

the regions that demonstrated significantly lower DIC values. The permutation test consisted of shuffling the ROI beta values across trials within each participant, performing hierarchical Bayesian parameter estimation, and saving the resulting DIC value. This process was repeated 1,000 times, each time with a random shuffling of beta values in order to define a null distribution of DIC values for the DDM model fit to that specific ROI. An empirical significance level was then calculated as the proportion of the null distribution that was smaller than the observed DIC value and compared to a Bonferroni-corrected significance level of 4.63×10^{-4} ($\alpha = 0.05$ and 108 tests). To aid in interpretation, significance values based on these permutation tests are reported as Bonferroni-adjusted p -values (p_{adj}) such that they can be compared to $\alpha = 0.05$. Although this overall statistical approach mixes frequentists and Bayesian rationales for interpreting statistical significance, it provides an appropriately conservative view for the type of data-driven exploratory analysis central to the current study.

Linking Category Evidence and Neural Response

Leveraging the HDDM provides a quantitative means for interrogating the link between neural activation and categorization decisions. In particular, finding that trial-by-trial fluctuations in the neural activation of certain brain regions is related to model-based predictions of behaviour suggests these brain regions are performing an important role in mapping sensory information onto category knowledge. However, the DDM is agnostic to specific mechanisms underlying computations of category evidence, thus it represents only one component of the motivating EBRW framework.

To test the hypothesis that the degree of evidence for one category over another is reflected in the neural dynamics of the identified ROIs, we performed an additional analysis that

linked participant-specific cognitive model predictions of categorization to neural activation in the ROIs identified by the brain-informed DDM analysis. Specifically, we generated predictions of category evidence with the GCM (Nosofsky, 1986). A key mechanism of GCM is selective attention, whereby diagnostic feature dimensions for a given task are weighted to varying degrees in calculating similarity to stored category exemplars and thus modulate the evidence for each category. It has been previously shown that individual differences in both categorization performance and neural representations (e.g., Mack et al., 2013; Braunlich & Love, 2019) are related to attention weighting in the GCM.

Thus, to quantify participant-specific predictions of category evidence, we fit to each participant’s behaviour by optimizing GCM parameters of dimensional attention weights (w), sensitivity (c), and decision gain (γ) with a genetic algorithm approach (differential evolution in *scipy* version 1.3.0) to maximize likelihood of response probabilities during the last four repetitions of training (Mack et al., 2013). We focused on end-of-learning behaviour to estimate model parameters with data independent from the data collected during the testing period. Following prior reports, r was set to 1 such that similarity values were computed according to city-block distance (Mack et al., 2013; Zaki, Nosofsky, Stanton, & Cohen, 2003). Participant-specific optimized parameters were then leveraged to predict for each stimulus the degree of discriminatory evidence (ev) for the most probable category. This measure of category evidence for stimulus x is the absolute value of the difference between the summed similarity for category A and B:

$$ev_x = \left| \sum_{y \in A} e^{-c \sum_{i \in d} w |x_i - y_i|} - \sum_{y \in B} e^{-c \sum_{i \in d} w |x_i - y_i|} \right|,$$

where d is the set of feature dimensions, A and B are the set of training items in the two categories, w is the set of optimized dimension weights, and c is sensitivity. We then evaluated

the relationship between category evidence (ev) and trial-by-trial neural activation with a mixed effects linear regression. Specifically, category evidence was modelled as the response, neural activation as the predictor, and random intercepts were included for participants. Given that category evidence is exponentially distributed, a log link function was included in the regression model. Regression analyses were conducted with Bayesian estimation (*rstanarm* R package version 2.19.3, R version 4.0.0). Neural activation models were compared to a baseline model that only included random intercepts.

Results

Testing phase performance (Figure 1) showed typical results consistent with previous reports (e.g., Medin & Schaffer, 1978). Responses to A and B training items demonstrated clear learning that was retained during test. Responses to novel transfer items varied according to match to category exemplars as determined by each participants' learning performance (Mack et al., 2013). Median reaction times (RTs) at the group level did not vary across items; however, there was variability across participants and trials. The brain-based DDM analysis offers a means for accounting for this trial-by-trial variability in response choices and times with neural function.

Figure 3. Brain regions linked to category decisions. DDM models informed by neural activation from occipital (red), mid cingulate (blue), and lateral PFC (green) each accounted for category decisions significantly better than the baseline model not informed by brain signals.

Brain Regions Related to Category Decisions

Across the brain, only six of the tested ROIs showed brain-informed HDDM predictions with significantly better accounts of category decisions relative to the baseline model.

Interestingly, these six ROIs were composed of three pairs of adjacent regions, with each pair showing similar effects. To simplify presentation of the results and best reflect the nature of their similar effects, we combined these pairs into three ROIs (Figure 3): an occipital region including extrastriate cortex, a region of midcingulate cortex, and a region of left lateral prefrontal cortex (PFC) extending into insula. HDDM simulations of these combined ROIs demonstrated significantly lower DICs (occipital: 6,712.9, mid cingulate: 6,713.5, lateral PFC: 6,711.5) relative to baseline (6,725.3). Of these three regions, permutation tests of occipital ($p_{adj} = 0.0019$) and lateral PFC ($p_{adj} = 0.0016$) survived multiple comparison correction whereas mid cingulate did not ($p_{adj} = 0.095$).

Although these three regions demonstrated similar overall fits to behaviour, the nature of the relationship between ROI activation and drift rate was unique across ROIs (Figure 4). In the occipital ROI, neural activation was positively related to drift rate but only for transfer items ($pd_{train} = 0.8, pd_{transfer} = 0.998$). In mid cingulate, neural activation was negatively related to drift rate but only for trained items ($pd_{train} > 0.999, pd_{transfer} = 0.676$). Finally, the lateral PFC region showed a negative relationship between neural activation and drift rate for both stimulus types ($pd_{train} = 0.997, pd_{transfer} = 0.985$).

Figure 4: Effects of neural activation on drift rate. Posterior distributions of the ROI activation effects on drift rate for training (purple) and transfer (orange) stimuli relative to 0 (dotted line) are shown separate for occipital (top), mid cingulate (middle), and lateral PFC (bottom) ROIs. Shaded regions represent 95% prediction intervals and asterisks denote effects with less than 2.5% of the posterior distribution overlapping with 0.

Category Evidence in Neural Activation

The key prediction of EBRW indicates that category decision making is driven by similarity-based comparisons to category exemplars. Extending this hypothesis to neural function, it follows that brain regions key for category decision making will exhibit activation profiles that track category evidence. To test this prediction, we evaluated the association between category evidence, as derived from GCM-based model fits of learning behaviour, and neural activation in the DDM-identified brain regions. The GCM captured the transfer behaviour well (mean RMSE=0.042) and the participant-specific optimized parameters (mean values: $c = 6.41$, $w_1 = 0.325$, $w_2 = 0.134$, $w_3 = 0.416$, $w_4 = 0.124$, $\gamma = 4.48$) were consistent with prior reports (Mack et al., 2013; Zaki et al., 2003). These parameter sets were used to calculate category evidence as defined in the methods.

Of the three ROIs, only lateral PFC showed a significant relationship with model-based predictions of category evidence (Figure 5; $R^2 = 0.145$) with higher lateral PFC activation associated with less discriminatory category evidence ($\beta = -0.004$, $CI = [-0.007, -0.001]$, $pd = 0.989$). These findings support the hypothesis that lateral PFC activation fluctuates as a function of category evidence (Paniukov & Davis, 2018) and that this category evidence plays a role in the accumulation of evidence in category decisions (Nosofsky & Palmeri, 1997).

Figure 5. Relationship between trial-by-trial lateral PFC activation and GCM-predicted category evidence (ev). Both group-level (thick black line) and participant-level (thin lines) effects are depicted with participant lines coloured according to the direction of the effect (blue = negative, yellow = positive). The shaded ribbon and shaded region of the distribution depict 95% prediction intervals of the group effect of LPFC activation.

Discussion

By integrating a formal category decision making model, EBRW, with whole-brain neural measures, we demonstrate that activation in specific brain regions relates to the trial-by-trial dynamics of category decisions. Specifically, we found that activation in LPFC was associated with category decisions and the speed of those decisions through the lens of the computational model. Notably, trial-by-trial activation in this region also related to exemplar-based predictions of category evidence. In both cases, the link to LPFC activation was an inverse relationship: higher activation was accompanied by less category evidence and slower decisions. These findings support an account of LPFC engagement tied to the difficulty of the current category

decision, such that LPFC is recruited to resolve conflicts and help drive decision making in ambiguous circumstances (Davis et al., 2017; O’Bryan, Worthy, Livesey, & Davis, 2018).

The LPFC has been previously reported to be involved in a variety of tasks consistent with this interpretation. Monkey studies suggest that neurons in LPFC are sensitive to multidimensional feature representations (Mendoza-Halliday & Martinez-Trujillo, 2017) and code for category boundaries (Seger & Miller, 2010). Additionally, human work points to a specific role for LPFC in the control of memory retrieval, particularly in the face of ambiguity and competing information (e.g., Badre & Nee, 2018; Thompson-Schill et al., 1998). Indeed, recent fMRI work focused on category learning has specifically pointed to LPFC playing a role in representing evidence for different categories (Davis et al., 2017; O’Bryan, Walden, Serra, & Davis, 2018; O’Bryan, Worthy, et al., 2018; Paniukov & Davis, 2018). Here, we extend these findings by directly linking trial-by-trial neural signals from LPFC to category decision making *and* the expression of exemplar-based category knowledge to uniquely support the notion that LPFC is tracking category evidence in a behaviourally relevant manner (Nosofsky & Palmeri, 1997; Paniukov & Davis, 2018).

The current findings are consistent with prior studies of LPFC’s role in category decisions, but important differences are noteworthy. In particular, whereas we defined the exemplar-based decision signal as the difference between category summed similarities, others have looked for neural engagement reflecting the summed similarity for the most likely category or a combined signal reflecting both exemplar similarities and dissimilarities (O’Bryan, Walden, et al., 2018; O’Bryan, Worthy, et al., 2018). These variants of decision signals are, in most tasks, correlated; however, they are intuitive and informative model quantities and may represent distinct variables computed in different regions. Future work employing the methods we propose

here with different candidate model-based decision signals may help dissociate distinct functions in these brain regions (O'Bryan, Worthy, et al., 2018; Paniukov & Davis, 2018; Zeithamova et al., 2019).

Although the occipital region was not associated with model-based predictions of category evidence, activation in this region did exhibit distinct relationships with evidence accumulation as formalized in the DDM. Decision-related activation was restricted to novel transfer items, such that higher activation was accompanied by quicker and more accurate responses. Prior studies support the notion of concept representation in perceptual cortices. Specifically, it is thought that recurrent feedforward/feedback loops with medial temporal lobe and PFC allow visual regions to make inferences about stimulus features (Hindy, Ng, & Turk-Browne, 2016; Lee & Mumford, 2003). Moreover, occipital regions respond to the similarity between particular stimuli and category representations (Braunlich & Love, 2019), as well as demonstrating an increase in activation in category-relevant visual areas (Folstein et al., 2013). Thus, the link between occipital activation and decision making we observe may be due to the engagement of neural representations tuned to diagnostic visual dimensions. It follows that this neural tuning may be most helpful and, therefore, recruited when generalizing to new stimuli.

Activation within mid cingulate cortex (MCC) also improved DDM predictions; however, its potential role in category learning is less clear. The relationship between MCC activation and drift rate was restricted to trained items, such that lower activation was associated with higher drift rates. While this relationship resembles that of the one observed in the LPFC, the lack of a similar relationship for transfer items is puzzling. Activation in MCC is often observed across a broad spectrum of cognitive tasks which calls into question its specificity for any given domain (Yarkoni, Poldrack, Nichols, Van Essen, & Wager, 2011). That this region did

not survive the conservative multiple comparison threshold further warrants caution in interpreting a distinct role of mid cingulate in category learning; future work will undoubtedly shed light on this potential node in the category learning neural network.

Our approach offers a novel method for interrogating neural measures with psychological theory and extends the recent surge in developing innovative approaches for bridging formal theories with brain measures (Bowman & Zeithamova, 2018; Davis, Love, & Preston, 2012; Daw, O'Doherty, Dayan, Seymour, & Dolan, 2006; Forstmann, Wagenmakers, Eichele, Brown, & Serences, 2011; Mack et al., 2013; Nosofsky, Little, & James, 2012; O'Doherty, Hampton, & Kim, 2007). Much of this work focusses on localizing hypothesized representational structure of stimuli (e.g., similarity relationships between items within and between categories) to neural representations found in activation patterns (Bowman & Zeithamova, 2018; Davis, Xue, Love, Preston, & Poldrack, 2014; Mack et al., 2016; Mack, Preston, & Love, 2020). Indeed, targeting representations in both model and brain has proven to be an invaluable tool in characterizing neural mechanisms underlying cognition and in evaluating formal theories. However, these approaches necessarily average over potentially informative neural dynamics that may explain behaviour on a trial-by-trial basis. Such trial-level approaches have become standard for research in reinforcement learning wherein model predictions of how value influences choices depend on prior experiences and responses (Daw et al., 2006; Leong, Radulescu, Daniel, DeWoskin, & Niv, 2017; Radulescu, Niv, & Ballard, 2019). We extend this approach by linking neural signal to a comprehensive prediction of behaviour in terms of choice probabilities and response times (Mack & Preston, 2016). This tight integration between brain and model is afforded by the EBRW framework that makes explicit predictions for the speed of category decisions based on the structure of category knowledge (Nosofsky & Palmeri, 1997). However, our approach is not

limited to category learning; any theory that makes clear predictions about response speed is a potential candidate for the methods that we outline here.

The current findings also extend an already rich literature supporting EBRW's comprehensive behavioural predictions (Annis & Palmeri, 2019; Nosofsky & Palmeri, 1997, 2015; Palmeri, 1997) to make the first step towards characterizing this model's neural mechanisms. The central tenet of EBRW is that exemplar-based category representations drive not only what categorization responses observers make but also the speed of those responses. The hypothesized mechanism posits that the degree to which a to-be-categorized stimulus matches stored memory traces of category exemplars dictates which exemplars are retrieved from memory to inform a random walk accumulation of evidence. Given the literature's view on LPFC, as noted above, that this region was solely linked to EBRW's mechanisms of category evidence accumulation provides compelling empirical support for the theory.

Our approach relies on first estimating trial-level beta series from the fMRI data (Mumford et al., 2012). This step of the analysis is inherently noisier than alternative methods of linking model to brain dynamics that leverage parametric regressors of time-varying model predictions in a mass univariate approach (e.g., Davis et al., 2012; O'Doherty et al., 2007). However, relationships between model mechanism and neural function identified with the parametric regressor approach are ultimately explanatory in nature. In contrast, by linking neural signal to drift rate in the DDM, our approach isolates brain function that gives rise to behaviour within a generative framework that makes predictions about category decisions on a trial-by-trial basis (Kragel et al., 2015). First focusing on isolating brain regions that comprehensively account for response choice and time allows for a more targeted evaluation of how such regions are further linked to model measures like category representation.

In summary, a comprehensive account of category learning requires an understanding of the dynamics of attention, representation, and decision making both at the level of computational and neural processes. Here, we build on recent advances that link computational model predictions of category representations to neural coding (Mack, Love, & Preston, 2018; Zeithamova et al., 2019) in order to isolate the neural signals of category decision making. We also extend methods of linking brain and behaviour through the DDM (Frank et al., 2015; Mack & Preston, 2016; Roberts & Hutcherson, 2019; White, Mumford, & Poldrack, 2012) to demonstrate that trial-by-trial neural signals from occipital, mid cingulate, and lateral PFC track the accumulation of evidence in category decisions. Importantly, LPFC activation tracked participant-specific predictions of exemplar-based category evidence as formalized by EBRW (Nosofsky & Palmeri, 1997). More generally, this approach offers a novel method for quantitatively connecting behavioural data to neural processes with cognitive theory.

References

- Annis, J., & Palmeri, T. J. (2019). Modeling memory dynamics in visual expertise. *Journal of Experimental Psychology: Learning Memory and Cognition*, 45(9), 1599–1618.
<https://doi.org/10.1037/xlm0000664>
- Badre, D., & Nee, D. E. (2018). Frontal Cortex and the Hierarchical Control of Behavior. *Trends in Cognitive Sciences*, Vol. 22, pp. 170–188. <https://doi.org/10.1016/j.tics.2017.11.005>
- Bowman, C., & Zeithamova, D. (2018). Abstract memory representations in the ventromedial prefrontal cortex and hippocampus support concept generalization. *The Journal of Neuroscience*, 28(10), 2605–2614. <https://doi.org/10.1523/JNEUROSCI.2811-17.2018>
- Braunlich, K., & Love, B. C. (2019). Occipitotemporal representations reflect individual differences in conceptual knowledge. *Journal of Experimental Psychology: General*, 148(7), 1192–1203. <https://doi.org/10.1037/xge0000501>
- Braunlich, K., & Seger, C. (2016). Categorical evidence, confidence, and urgency during probabilistic categorization. *NeuroImage*, 125, 941–952.
<https://doi.org/10.1016/j.neuroimage.2015.11.011>
- Davis, T., Goldwater, M., & Giron, J. (2017). From Concrete Examples to Abstract Relations: The Rostrolateral Prefrontal Cortex Integrates Novel Examples into Relational Categories. *Cerebral Cortex*, 27, 2652–2670. <https://doi.org/10.1093/cercor/bhw099>
- Davis, T., Love, B. C., & Preston, A. R. (2012). Learning the exception to the rule: Model-based fMRI reveals specialized representations for surprising category members. *Cerebral Cortex*, 22(2), 260–273.
- Davis, T., Xue, G., Love, B. C., Preston, A. R., & Poldrack, R. A. (2014). Global neural pattern similarity as a common basis for categorization and recognition memory. *The Journal of*

Neuroscience : The Official Journal of the Society for Neuroscience, 34(22), 7472–7484.

<https://doi.org/10.1523/JNEUROSCI.3376-13.2014>

Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, 441(7095), 876–879.

Esteban, O., Markiewicz, C. J., Blair, R. W., Moodie, C. A., Isik, A. I., Erramuzpe, A., ...

Gorgolewski, K. J. (2019). fMRIPrep: a robust preprocessing pipeline for functional MRI.

Nature Methods, 16(1), 111–116. <https://doi.org/10.1038/s41592-018-0235-4>

Folstein, J. R., Palmeri, T. J., & Gauthier, I. (2013). Category learning increases discriminability of relevant object dimensions in visual cortex. *Cerebral Cortex*, 23(4), 814–823.

<https://doi.org/10.1093/cercor/bhs067>

Forstmann, B. U., Wagenmakers, E.-J., Eichele, T., Brown, S., & Serences, J. T. (2011).

Reciprocal relations between cognitive neuroscience and formal cognitive models: opposites attract? *Trends in Cognitive Sciences*, 15(6), 272–279.

Frank, M. J., Gagne, C., Nyhus, E., Masters, S., Wiecki, T. V., Cavanagh, J. F., & Badre, D.

(2015). FMRI and EEG predictors of dynamic decision parameters during human reinforcement learning. *Journal of Neuroscience*, 35(2), 485–494.

<https://doi.org/10.1523/JNEUROSCI.2036-14.2015>

Hindy, N. C., Ng, F. Y., & Turk-Browne, N. B. (2016). Linking pattern completion in the hippocampus to predictive coding in visual cortex. *Nature Neuroscience*, 19(5), 665–667.

<https://doi.org/10.1038/nn.4284>

Kafkas, A., & Montaldi, D. (2018, July 27). How do memory systems detect and respond to novelty? *Neuroscience Letters*, Vol. 680, pp. 60–68.

<https://doi.org/10.1016/j.neulet.2018.01.053>

- Kragel, J. E., Morton, N. W., & Polyn, S. M. (2015). Neural Activity in the Medial Temporal Lobe Reveals the Fidelity of Mental Time Travel. *Journal of Neuroscience*, 35(7), 2914–2926. <https://doi.org/10.1523/JNEUROSCI.3378-14.2015>
- Lee, T., & Mumford, D. (2003). Hierarchical Bayesian inference in the visual cortex. *The Journal of the Optical Society of America: A*, 20(7), 1434–1448. <https://doi.org/10.1364/josaa.20.001434>
- Leong, Y. C., Radulescu, A., Daniel, R., DeWoskin, V., & Niv, Y. (2017). Dynamic Interaction between Reinforcement Learning and Attention in Multidimensional Environments. *Neuron*, 93(2), 451–463. <https://doi.org/10.1016/j.neuron.2016.12.040>
- Mack, M. L., Love, B. C., & Preston, A. R. (2016). Dynamic updating of hippocampal object representations reflects new conceptual knowledge. *Proceedings of the National Academy of Sciences of the United States of America*, 113(46), 13203–13208. <https://doi.org/10.1073/pnas.1614048113>
- Mack, M. L., Love, B., & Preston, A. (2018). Building concepts one episode at a time: the hippocampus and concept formation. *Neuroscience Letters*, 680, 31–38. <https://doi.org/10.1016/j.neulet.2017.07.061>
- Mack, M. L., & Preston, A. R. (2016). Decisions about the past are guided by reinstatement of specific memories in the hippocampus and perirhinal cortex. *NeuroImage*, 127, 144–157. <https://doi.org/10.1016/j.neuroimage.2015.12.015>
- Mack, M. L., Preston, A. R., & Love, B. C. (2013). Decoding the brain's algorithm for categorization from its neural implementation. *Current Biology*, 23(20), 2023–2027. <https://doi.org/10.1016/j.cub.2013.08.035>
- Mack, M. L., Preston, A. R., & Love, B. C. (2020). Ventromedial prefrontal cortex compression

during concept learning. *Nature Communications*, 11(1), 1–11.

<https://doi.org/10.1038/s41467-019-13930-8>

Makowski, D., Ben-Shachar, M. S., Chen, S. H. A., & Lüdecke, D. (2019). Indices of Effect Existence and Significance in the Bayesian Framework. *Frontiers in Psychology*, 10, 2767.

<https://doi.org/10.3389/fpsyg.2019.02767>

Medin, D., & Schaffer, M. (1978). Context theory of classification learning. *Psychological Review*, 85(3), 207–238. <https://doi.org/10.1037/0033-295X.85.3.207>

Mendoza-Halliday, D., & Martinez-Trujillo, J. (2017). Neuronal population coding of perceived and memorized visual features in the lateral prefrontal cortex. *Nature Communications*, 8.

Mumford, J., Turner, B., Ashby, F., & Poldrack, R. (2012). Deconvolving BOLD activation in event-related designs for multivoxel pattern classification analyses. *NeuroImage*, 59(3), 2636–2643. <https://doi.org/10.1016/j.neuroimage.2011.08.076>

Nosofsky, R. M. (1986). Attention, Similarity, and the Identification-Categorization Relationship. *Journal of Experimental Psychology: General*, 115(1), 39–57. <https://doi.org/10.1037/0096-3445.115.1.39>

Nosofsky, R. M., Little, D. R., & James, T. W. (2012). Activation in the neural network responsible for categorization and recognition reflects parameter changes. *Proceedings of the National Academy of Sciences of the United States of America*, 109(1), 333–338. <https://doi.org/10.1073/pnas.1111304109>

Nosofsky, R. M., & Palmeri, T. J. (1997). An Exemplar-Based Random Walk Model of Speeded Classification. *Psychological Review*, 104(2), 266–300. <https://doi.org/10.1037/0033-295X.104.2.266>

Nosofsky, R. M., & Palmeri, T. J. (2015). *An Exemplar-Based Random-Walk Model of*

- Categorization and Recognition* (Vol. 1; J. R. Busemeyer, Z. Wang, J. T. Townsend, & A. Eidels, Eds.). <https://doi.org/10.1093/oxfordhb/9780199957996.013.7>
- O'Bryan, S. R., Walden, E., Serra, M. J., & Davis, T. (2018). Rule activation and ventromedial prefrontal engagement support accurate stopping in self-paced learning. *NeuroImage*, 172, 415–426. <https://doi.org/10.1016/j.neuroimage.2018.01.084>
- O'Bryan, S. R., Worthy, D. A., Livesey, E. J., & Davis, T. (2018). Model-based fMRI reveals dissimilarity processes underlying base rate neglect. *ELife*, 7, 1–23. <https://doi.org/10.7554/eLife.36395>
- O'Doherty, J. P., Hampton, A., & Kim, H. (2007). Model-based fMRI and its application to reward learning and decision making. *Annals of the New York Academy of Sciences*, 1104, 35–53. <https://doi.org/10.1196/annals.1390.022>
- Palmeri, T. J. (1997). Exemplar similarity and the development of automaticity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23(2), 324–354. <https://doi.org/10.1037/0278-7393.23.2.324>
- Paniukov, D., & Davis, T. (2018). The evaluative role of rostrolateral prefrontal cortex in rule-based category learning. *NeuroImage*, 166, 19–31. <https://doi.org/10.1016/j.neuroimage.2017.10.057>
- Radulescu, A., Niv, Y., & Ballard, I. (2019). Holistic Reinforcement Learning: The Role of Structure and Attention. *Trends in Cognitive Sciences*, 23(4), 278–292. <https://doi.org/10.1016/j.tics.2019.01.010>
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2), 59–108. <https://doi.org/10.1037/0033-295X.85.2.59>
- Roberts, I. D., & Hutcherson, C. A. (2019, July 1). Affect and Decision Making: Insights and

- Predictions from Computational Models. *Trends in Cognitive Sciences*, Vol. 23, pp. 602–614. <https://doi.org/10.1016/j.tics.2019.04.005>
- Schaefer, A., Kong, R., Gordon, E., Laumann, T., Zuo, X., Holmes, A., ... Yeo, B. (2018). Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity MRI. *Cerebral Cortex*, 28(9), 3095–3114. <https://doi.org/10.1093/cercor/bhx179>
- Seger, C., & Miller, E. (2010). Category learning in the brain. *Annual Review of Neuroscience*, 33, 203–219. <https://doi.org/10.1146/annurev.neuro.051508.135546>
- Shepard, R. N. (1957). Stimulus and response generalization: A stochastic model relating generalization to distance in psychological space. *Psychometrika*, 22(4), 325–245. <https://doi.org/10.1007/BF02288967>
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & Van Der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 64(4), 583–616. <https://doi.org/10.1111/1467-9868.00353>
- Thompson-Schill, S. L., Swick, D., Farah, M. J., D’Esposito, M., Kan, I. P., & Knight, R. T. (1998). Verb generation in patients with focal frontal lesions: A neuropsychological test of neuroimaging findings. *Proceedings of the National Academy of Sciences of the United States of America*, 95(26), 15855–15860. <https://doi.org/10.1073/pnas.95.26.15855>
- Thompson, W. H., Wright, J., & Bissett, P. G. (2020). Open exploration. *ELife*, 9. <https://doi.org/10.7554/eLife.52157>
- White, C. N., Mumford, J. A., & Poldrack, R. A. (2012). Perceptual criteria in the human brain. *The Journal of Neuroscience*, 32(47), 16716–16724. <https://doi.org/10.1523/JNEUROSCI.1744-12.2012>
- Wiecki, T. V., Sofer, I., & Frank, M. J. (2013). HDDM: Hierarchical Bayesian estimation of the

Drift-Diffusion Model in Python. *Frontiers in Neuroinformatics*, 7.

- Yarkoni, T., Poldrack, R. A., Nichols, T. E., Van Essen, D. C., & Wager, T. D. (2011). Large-scale automated synthesis of human functional neuroimaging data. *Nature Methods*, 8(8), 665–670. <https://doi.org/10.1038/nmeth.1635>
- Young, G., & Householder, A. (1938). Discussion of a set of points in terms of their mutual distances. *Psychometrika*, 3(1), 19–22. <https://doi.org/10.1007/BF02287916>
- Zaki, S. R., Nosofsky, R. M., Stanton, R. D., & Cohen, A. L. (2003). Prototype and exemplar accounts of category learning and attentional allocation: A reassessment. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 29(6), 1160–1173.
- Zeithamova, D., Mack, M. L., Braunlich, K., Davis, T., Seger, C. A., van Kesteren, M. T. R., & Wutz, A. (2019). Brain Mechanisms of Concept Learning. *The Journal of Neuroscience*, 39(42), 8259–8266. <https://doi.org/10.1523/JNEUROSCI.1166-19.2019>