Title

Cortical transformation of stimulus-space in order to linearize a linearly inseparable task

Meng-Huan Wu^{1*}, David Kleinschmidt^{2*}, Lauren Emberson³, Donias Doko⁴, Shimon Edelman⁵,
Robert Jacobs¹, Rajeev Raizada¹

¹University of Rochester, ²Rutgers University, ³Princeton University, ⁴Frank H. Netter School of
Medicine, Quinnipiac University, ⁵Cornell University

*These authors contributed equally to this work

Abstract

The human brain is able to learn difficult categorization tasks, even ones that have linearly inseparable boundaries; however, it is currently unknown how it achieves this computational feat. We investigated this by training participants on an animal categorization task with a linearly inseparable prototype structure in a morph shape space. Participants underwent fMRI scans before and after four days of behavioral training. Widespread representational changes were found throughout the brain, including an untangling of the categories' neural patterns that made them more linearly separable after behavioral training. These neural changes were task-dependent, as they were only observed while participants were performing the categorization task, not during passive viewing. Moreover, they were found to occur in frontal and parietal areas, rather than ventral temporal cortices, suggesting that they reflected attentional and decisional reweighting, rather than changes in object recognition templates. These results illustrate how the brain can flexibly transform neural representational space in order to solve computationally challenging tasks.

Introduction

Just by a quick glance at a photograph of an animal, people can say with reasonably good accuracy whether the photo was of a cat or a dog. The apparent ease with which this is accomplished belies the computational complexity of this process. If we think of each image as a point in a “pixel space” (where each dimension is the brightness of one pixel), categories such as “dog” and “cat” correspond to parts of this image space that are highly tangled, like two pieces of paper crumpled together (DiCarlo & Cox, 2007). Somehow, the brain transforms these tangled, linearly inseparable representations to linearly separable representations that are computationally simple to read out.

How does the brain achieve this feat? In order to study how the brain transforms its representation of a linearly inseparable task, three key ingredients must be combined together. First, we need to study representational change. To observe a change, at least two fMRI scanning sessions are required, so that they can be directly compared: one before participants have learned the task, and the other after they have been trained on it. Second, we need to use a linearly inseparable task, i.e. one that cannot be solved using a linear category boundary. Although linearly separable tasks are simpler and easier for participants to learn, the real world is rarely so cooperative. Very few real-world tasks are linearly separable. We therefore chose to study an inseparable one. The third key ingredient is to use a task that allows lower-level pixel-space aspects to be clearly distinguishable from higher-level shape-space properties. Human beings can recognize a cat as being a cat rather than a dog, regardless of whether it is seen from the front, side, above or below. Each of those different viewpoints produces a vastly different visual projection on the retina (the eye's version of pixel-space), but in all cases the same 3D cat-shaped body is giving rise to them. In the present study, we therefore chose to present 3D shapes in a variety of different viewpoints, and to require participants to categorize the 3D shapes themselves, discounting irrelevant viewpoint information.

The most important difference between the present study and much of the earlier work is the fact that our study scanned the participants twice: once before training, and then a second fMRI scan after several days of behavioral training. Without pre- and post-scans of this sort, it is not possible to measure representational change. One might argue that the pre-training scan is unnecessary, as one can plausibly assume that an untrained brain will not contain any pre-existing representations of the task that is about to be learned. Several studies using only post-training scans have been carried out, and have provided very valuable insights (e.g. Reber et al., 1998; Folstein et al., 2013; Seger et al., 2015). Despite the plausibility of this assumption, we wish to argue that actually collecting pre-training fMRI data is necessary in order to truly study representational change. Even a “naive” participant will already have seen and categorized literally millions of visual stimuli, so the experimental stimulus materials may very possibly trigger some sort of indirect visual recognition memory, even if they are just patterns of dots.

In many fMRI studies, the question of interest is not only whether some neural representation exists in the brain (if a person can perform a task, then logically their brain must contain some sort of information about it, in some form), but more specifically whether we are able to detect that representation and measure its properties. Thus, even though one can logically deduce that training has changed the representational content of the participants' brains, it is scientifically of great interest to test whether information about those representational changes can actually be measured using our existing experimental techniques.

The second crucial aspect of the present study is our use of a linearly inseparable task. Again, this is not unprecedented in the literature. However, existing studies of linearly

inseparable tasks lacked one or more of our set of three crucial ingredients, thereby allowing our study to ask questions that have not been previously addressed. Specifically, using a prototype distortion paradigm, Braunlich et al. (2017) demonstrated that support vector regression can be used to predict each stimuli's distance from the category boundary and from the prototype. Using an XOR stimulus space, Li et al. (2007) demonstrated that a support vector machine can be used to decode which category each stimulus belonged to. A key difference between those studies and the one presented here is that they used flat two-dimensional stimuli, whereas we presented 3D shapes from multiple viewing angles. This allowed us to show that category learning mechanisms operated beyond recognizing stimuli differences in the pixel space, and that such mechanisms can be generalized to more naturalistic category learning settings in shape-space. Moreover, our prototype category structure is arguably more natural than the XOR task. The XOR task has an elegant logical form, but tasks with that structure rarely arise in everyday life. In contrast, many tasks have the prototype structure, for example: is this person in my tribe, or a stranger? Am I close to home, or far away?

The third crucial aspect of our present study is that we used a category structure that was defined in shape-space, not pixel-space. Critically, each exemplar in our study was presented from multiple viewing angles. In previous experiments, each exemplar was presented only in one canonical angle, meaning that two exemplars could be perfectly discriminated based solely on differences in pixel values (Braunlich et al., 2017; Folstein, Gauthier, & Palmeri, 2012; Mack et al., 2016). Such discriminations can presumably be accomplished by recruiting low-level visual cortices and frontal regions (Reber et al., 1998) without involving shape-selective regions such as lateral occipital complex (LOC). In contrast, because we presented each exemplar from multiple viewing angles, our experimental task forced participants to map stimuli that are vastly different in pixel-space to the same exemplar in shape-space. This manipulation made it less likely that low-level visual regions would underlie category learning, and much more likely that higher-level regions sensitive to object shape would play an important role in the acquisition of category structures.

In sum, to address our question of interest, we needed pre- and post-training scans, a linearly inseparable task, and 3D shapes presented from different viewpoints. All three of these aspects of the study needed to be combined together at once, in order to be able to attack this goal. While some experiments have already been performed that individually include some subset of these three necessary ingredients, to the best of our knowledge no existing study has combined all three at once, until the one presented here. Therefore, our study is able to investigate how the brain transforms its representation of a linearly inseparable task in a new way. Our experimental design allows us, for the first time, to test the following three claims at once: that observed representational changes are indeed changes, that these changes are in response to trying to solve a linearly inseparable task, and that these changes reflect shape-level rather than pixel-level stimulus properties.

Methods

Participants

Eighteen University of Rochester students participated in the study. They all gave written consent in accordance with University of Rochester research subjects review board. One participant was run before the paradigm was finalized, and two participants failed to learn the task (post-test categorization accuracy lower than 0.5); therefore, only fifteen participants are included for further fMRI analysis.

Stimuli and design

Stimuli were 3D animal shapes, and the generation procedure is described in <https://github.com/kleinschmidt/animorph> (Edelman, Bühlhoff, & Bühlhoff, 1999). Many aspects of the animal shapes could be parametrically altered in order to change their appearance, e.g. length and girth of the torso, ear size and position, angles of the legs, and so on. There were 55 such parameters in all. We wanted to create shape categories which were defined not just by one or two salient features, but which instead involved changes to the overall animal appearance resulting from many features all changing together. We therefore created two approximately orthogonal vectors that cut obliquely through the 55-dimensional parameter space, with those two vectors producing a two-dimensional shape space that involved changes in many different features together. These two vectors will be referred to as the x and y axes.

The resulting shape space is shown in Figure 1. Although many individual features varied, some are particularly noticeable, e.g. how fat or thin the animals' bodies were, their front knee angle, and the distances between ears and between eyes. To define the Dax and non-Dax categories in this shape space, we defined a central region as the Dax category, and the outer regions of the space as non-Dax, as can also be seen from that figure. All of the animals used in the experiment were formed by a linear combination of these two basis vectors.

In this space, we defined a linearly inseparable prototype category structure (Fig. 1). We first picked a point in the x, y space to serve as the central prototype. Animal shapes closer than a critical distance to this prototype were defined to be "Didoop Daxes", and shapes farther than the critical distance were "non-Didoop Daxes" (for brevity we refer to the categories as "Dax" and "non-Dax").

Fig. 1. 2D morph space of novel animals used in the fMRI sessions. During behavioral training, only animals in the inner and outer ring were shown to the participants; therefore, middle ring animals were not pre-assigned to either of the two categories.

This prototype stimulus design has two important properties. First, the Daxes and non-Daxes were completely tangled in the current stimulus space: no linear transformation of this parameter x, y space can make the categories linearly separable. Second, these stimuli are drawn from a continuously and parametrically varying space, which differs from previous fMRI studies of linearly-inseparable category tasks (Mack, Preston, & Love, 2013). The current methods of stimulus construction make it possible to formulate explicit models of stimuli representations and the corresponding pairwise similarity structure of the stimulus parameters that can be compared against the similarity structure of the neural representations.

We also varied the viewing angle of the stimuli, such that several stimuli vastly different in pixel space were in fact the same animal. Unlike previous studies (e.g. Reber et al., 1998), this critical experiment design feature ruled out the hypothesis that the brain merely picked up stimuli differences in the pixel space, and provided evidence that the brain did indeed learn the shape difference between animals. Stimuli could vary in seven possible pitches and five possible orientations (several extreme orientation and pitch angles are shown in Fig. 2).

Fig. 2. Each animal was displayed in various orientations and pitches. In this figure, the animals are displayed at the extremes of the ranges of orientation and pitch angles, in order to illustrate the wide range of viewpoints that were included in the categorization task. In the pitch discrimination task, the participants' task was to identify the two rightmost animals as "standing on a steep hill", and to identify the two leftmost animals as "not standing on a steep hill".

Procedure

Subjects came in for six separate sessions: a pre-training scanning session, four behavioral training sessions, and a post-training scanning session, all on separate days.

Each of the four training sessions consisted of 320 trials in four 80-trial blocks. Subjects were told they were going to learn to tell the difference between two varieties of animals called "Daxes": ones that lived on the island of Didoop, and ones that did not. On each trial, the subject saw a picture of one of the animals, and had to press one of the two buttons to report whether it was a "Didoop Dax" or a "non-Didoop Dax". Subjects were told whether they were correct or incorrect by a thumbs up or thumbs down icon, with additional positive feedback on correctly identified "Didoop Dax" trials in the form of a picture of a tropical island. The tropical island was shown for additional feedback when participants correctly identified a Didoop Dax animal, as the task context described that Dax animals lived on the island of Didoop¹.

On the first day of training, subjects saw only the canonical view of the animals until they exceeded 60% accuracy for one 80-trial block. In all subsequent blocks, variation in the orientation was introduced, and animals would be sampled from all five orientations. On the second day, after subjects again achieved 60% accuracy on one block, additional variation in the pitch angle was added until the end of all training sessions, and animals will be sampled from all viewing angles. Note that the number of trials and blocks remained identical regardless of the learning curve of each participant.

The pre- and post-training scanning sessions were identical. Subjects performed three different tasks, each on the exact same stimuli sequence presented in the same order. First, they passively viewed the stimuli. Next, the subjects performed the same classification task they were trained on (the Dax task), and a pitch angle discrimination task where they had to determine whether the animal was standing on a steep hill or not (the slant task). For all the tasks performed in the scanner, the stimulus appeared on the screen for 2 seconds from the onset of each trial, and participants had 3 seconds to respond. During the Dax task, the subject saw a picture of one of the animals, and had to press one of the two buttons to report whether it was a "Didoop Dax" or a "non-Didoop Dax". In the pitch discrimination task, subjects were asked "Is the animal standing on a steep hill or not?". The participants were informed that they should respond solely based on animals' pitch angle and not their shape. Examples of the stimuli shown at different pitch angles and orientations, along with labels indicating which ones were to be classed as on a

¹ The picture increased participants' task engagement over multiple days of training.

steep hill or not, are shown in Fig.2. Before beginning the pitch task, participants had two practice trials, one at the steepest pitch and one at the shallowest pitch, which were repeated until both were responded to correctly. The order of these two tasks was counterbalanced across subjects, while the passive task always occurred first. In the pre-training scanning session, participants were asked to categorize animals without any knowledge about the category; they were instructed to pay close attention to the animals and take their best guess.

We wished to test whether participants could learn the linearly inseparable boundary and generalize it to novel animals, instead of merely learning individual associations between each animal shape and its category label. In order to test for generalization, the stimuli that we presented during the fMRI scanning sessions consisted not only of animal shapes that had been shown during the behavioral training sessions, but also some novel stimuli drawn from previously unsampled parts of the two-dimensional shape space.

Thus, the stimuli that the participants saw in the scanning sessions were a variant of the stimuli that were used during the behavioral training, forming an “X” in the x, y space (Fig. 1). Among these stimuli, four non-Dax stimuli (referred to as outer-ring) were also shown to the participants during the behavioral training sessions, and all the other stimuli were shown only during the fMRI scans. The fMRI-only stimuli were: the prototype, four Dax stimuli (inner-ring) that were generated by rotating the four Dax stimuli seen during training by 45 degrees in the morph space, and four intermediate stimuli (middle-ring) that were equidistant from the Dax and non-Dax stimuli (on the category boundary shown in Fig. 1).

Jitters of 0, 2, and 4 seconds were added after the response phase ended. Each of these thirteen animal shapes was shown four times each for four runs, for 52 trials each per block and 208 trials total, 16 per unique animal shape. Animal shape, orientations, pitches, and trial order for the 208 trials were identical in all three tasks, for all subjects and sessions.

Image acquisition

Whole-brain images were acquired with a 3-T Siemens MAGNETOM Trio scanner with a 32-channel head coil located at the Rochester Center for Brain Imaging. At the start of each participant’s scanning session, a high-resolution structural T1 contrast image was acquired using a magnetization prepared rapid gradient-echo pulse sequence (repetition time = 2530, echo time = 3.44 msec, flip angle = 7°, field of view = 256 mm, matrix = 256 × 256, 1 × 1 × 1 mm sagittal left-to-right slices). An interleaved EPI pulse sequence was used for T2* contrast (repetition time = 2000 ms, echo time = 30 ms, flip angle = 90°, field of view = 256 × 256 mm, matrix = 64 × 64, 30 sagittal left-to-right slices, voxel size = 4 × 4 × 4 mm). The first six volumes of each run were discarded to allow for signal equilibration.

fMRI Data Preprocessing

Functional images were processed using SPM8. Preprocessing stages included motion correction, slice time correction and spatial normalization. Beta estimates for each voxel were obtained by fitting a general linear model (GLM) to the time series signal, where we convolved each stimulus onset with a standard hemodynamic response function. As is standard, the GLM model also included constant-term columns, one for each of the 12 runs, in order to allow it to account for run-to-run variations in baseline BOLD signal.

Candidate Models and Construction of Model RDMs

We constructed a high-level prototype distance model and a low-level pose-only model to compare their effectiveness in explaining neural representational dissimilarity matrices (RDMs). The former aimed to capture the task-relevant information between animals in the animal morph space; its RDM was a 13-by-13 matrix, where each entry was calculated as the difference in Euclidean distances between the animals and the center (e.g., two animals on the same ring will have a zero in the RDM entry). The latter aimed to capture the voxel activation responses of early visual cortex for each viewing angle, and its RDM was a 35-by-35 matrix (5 orientations and 7 pitch angles). We first constructed a set of Gabor filter banks varying in 4 spatial scales (scaling factor = 1.7 and minimum wavelength = 3) and 6 orientations (Field, 1987). For each animal in each viewing angle, a 1444-dimensional Gabor feature vector was computed by convolving the image stimulus with the Gabor filter banks. Finally, the pose-only model features were calculated by averaging Gabor feature vectors for each viewing angle across different animals, and each entry in this model was calculated as the Pearson correlation between the pose-only model features of viewing angle pairs.

The prototype distance model contained all and only information about the Dax vs. non-Dax categorization task, while the pose-only model pooled Gabor features across animal shapes and were completely irrelevant to that task. Note that the dimension of the two model RDMs are different, one is the number of different animal shapes (13x13), and the other is the number of different viewing angles for each animal shape (35x35). In theory, a 13x35=455 dimensional RDM matrix could be constructed such that each dimension represents one animal in one viewing angle. However, this number would exceed the total number of stimulus presentations, so it would be impossible to make accurate fMRI activation estimates to fill this RDM. Fortunately, our aim is not to compare the fit of the 13x13 prototype distance model against the fit of the 35x35 pose-only Gabor model, so there is no need to model both factors simultaneously. Our main hypothesis is to test how representations of animal shape-space change before and after training.

Finally, to investigate whether neural representations are encoding pitch angle information, we constructed a 7-by-7 pitch model where each entry was calculated as the pitch angle differences for each of the 7 different pitch angles.

Representational Similarity Analysis

The RDMs for the candidate models were described in the preceding subsection. Here, we describe the construction of the neural RDMs. We applied representational similarity analysis (RSA; Edelman, Grill-Spector, Kushnir, & Malach, 1998; Kriegeskorte et al., 2008) to each of the 48 bilateral Harvard-Oxford Atlas (Desikan et al., 2006) regions of interest (ROIs). To create the neural RDM for the prototype distance model, the features of each animal was calculated by averaging the beta values across trials of the same animal shape, and each entry (i, j) in this RDM was calculated as 1 minus Pearson correlation between features from animal i and animal j . The neural RDM for the pose-only Gabor model was created in a similar fashion, while features of each viewing angle were calculated by averaging beta values from trials with the same viewing angle instead. Finally, the neural RDM for the pitch model was similarly created by averaging beta values from trials with the same pitch angle together. To investigate how well each model fit the neural RDM, Kendall's tau-a correlation is calculated between the entries in upper triangular part of model and neural RDM for each ROI in each participant. Kendall's tau-a correlation is selected since the two model RDMs (described in previous section) contained

multiple ties (Nili et al., 2014). For each ROI, a Student's t-test was applied on the Fisher-Z transformed correlation values across participants. P-values across all 48 bilateral ROIs and two scanning sessions were corrected with false discovery rate (FDR; Benjamini & Hochberg, 1995).

Visualizing representation using Multidimensional Scaling

In order to visualize the representational organization of animal shapes in each ROI, we applied classical metric multidimensional scaling (MDS; Torgerson, 1952) to the group average RDM². After a double-centering operation, we decomposed the group average RDM (termed S_{avg} here) into eigenvectors and eigenvalues:

$$S_{avg} = VDV^T, \text{ where } V^T V = I \quad (1)$$

The low-dimensional coordinates of each animal shape can then be computed by:

$$P_{avg} = VD^{\frac{1}{2}} = S_{avg} * VD^{-\frac{1}{2}} \quad (2)$$

To compute confidence intervals for each animal shape, we generated 500 group average RDMs by bootstrap resampling subjects with replacement, and calculated their coordinates as

$$P_{bootstrap,avg} = VD^{\frac{1}{2}} = S_{bootstrap,avg} * VD^{-\frac{1}{2}} \quad (3)$$

We visualized the bootstrap sample coordinates by fitting a 2-dimensional ellipse that contains 95% of the points (details are described in Appendix of Abdi, Dunlop, & Williams, 2009).

Previous studies have applied the DISTATIS algorithm (Abdi et al., 2009) to the RDMs of individual subjects. Instead of calculating the group average RDM, DISTATIS calculates a weight for each participant based on the similarities between individual participant's RDM, and outputs a compromise matrix, which weights and averages each participant's RDM accordingly. We obtained qualitatively similar figures by applying this algorithm, so to save space in the present manuscript we show only the figures that were generated by applying MDS to group average RDMs.

Cross-task Linear Support Vector Machine (SVM) Classifier

During the categorization task, participants pressed the same button to indicate animal categories. In order to identify brain regions that were sensitive to button responses but not category information, for each ROI in each participant we trained a linear SVM classifier (MATLAB `fitsvm` function) in one task to distinguish between button 1 and button 2 presses, and tested it on another task. A Student's t-test was performed on the accuracy for each ROI to determine whether the accuracy was higher than 0.5. Cross-task SVM classifier was applied here instead of RSA since there were only two conditions (button 1 and 2 presses), and the resulting dissimilarity matrix would only contain one unique correlation value.

² Pearson correlations were used to calculate fMRI RDMs, so we were aware that this violated classical MDS's assumption of Euclidean distance matrix. However, since all eigenvalues we obtained were positive and this was used mainly as a visualization technique, we left the development of rigorous resampling and projection techniques for future work.

Results

Participants successfully learned the linearly inseparable category boundary

We trained participants to categorize novel animals with various orientations and pitches as “Dax” or “non-Dax”. The task was fairly challenging, in large part due to the fact that the different presentations of each animal spanned a wide range of different viewing angles. In order to learn which animals were Dax or non-Dax, the subjects needed to categorize the intrinsic shape of each animal, abstracted away from the viewing angle that it happened to be seen from. The challenging nature of the task can be seen from the fact that, over the course of training, the participants’ performance did not reach plateau even after four days of training. However, the 92% mean accuracy in the final training session shows that participants did indeed succeed at learning the task.

After behavioral training, participants performed a similar Dax categorization task in the scanner without feedback. For statistical tests of whether the participants categorized particular stimuli as Dax or non-Dax, the proportion of category-label responses for each animal stimulus were calculated for each of the 15 participants, and then those 15 values were subjected to a group-level random effects t-test against the chance level of 50%. Four outer-ring animals were classified as non-Dax ($t_{14}=37$, $p<1e-14$, 95% non-Dax responses averaged across animals and participants, chance = 50%), while the novel inner-ring ($t_{14}=9.5$, $p<1e-7$) and prototype animals ($t_{14}=10.9$, $p<1e-7$) are classified as Dax (Fig. 3). Furthermore, proportion of Dax responses was significantly higher for prototype than that for inner-ring animals ($t_{14}=4.56$, $p<0.001$); this prototype effect corroborated with previous studies (Knowlton & Squire, 1993). Overall, this suggested that participants accurately learned the linearly inseparable category boundary and were able to generalize it on novel animals.

In-scanner Behavioral Performance

Fig. 3. Group average proportion of Dax responses during the category discrimination task performed in the MRI scanner. Left: pre-training scans. Right: post-training scans.

Representational distances in intraparietal sulcus and inferior frontal gyrus were ranked according to distances to prototype

If participants could indeed learn the linearly inseparable category boundary, then how did their brains manage to untangle this category information from complex and linearly inseparable low-level visual representations? According to prototype theory (Cutzu & Edelman, 1998; Posner & Keele, 1968), novel stimuli were assigned to the category with the closest prototype. We therefore constructed a prototype distance model where the dissimilarity between each animal pair was calculated as their Euclidean distance differences between the prototype animal in the morph space. We then correlated this model RDM with that of all bilateral Harvard-Oxford ROIs using RSA. 25 out of 48 Harvard-Oxford Atlas regions were significantly correlated with prototype distance model when participants were performing the Dax classification task after behavioral training ($t > 2.56$, FDR-corrected across all ROIs and pre/post sessions); no ROIs are significantly correlated with the model before training (Fig. 4). Among these regions, intraparietal sulcus (IPS)³ and Inferior Frontal Gyrus (IFG) had the highest t -values. This was consistent with previous studies demonstrating that intraparietal and frontal areas were differentially activated during visual category learning tasks (Seger et al., 2000).

Fig. 4. T-value maps for Kendall's τ_a correlation between RDMs of the prototype distance model and the activation patterns within each ROI (collected during the categorization task, in

³ Throughout the manuscript we used the term “intraparietal sulcus (IPS)” instead of the the name that happens to be used in the text files that accompany the Harvard-Oxford atlas downloads, which is “Lateral Occipital Cortex, superior division”. This is because this ROI is indeed in the parietal cortex, so in referring to it as the intraparietal sulcus (IPS) we thereby avoid creating any unnecessary confusion. Moreover, recent papers in the goal-relevant visual processing literature which discuss the IPS provide coordinates for that region that fall into this Harvard-Oxford ROI (Henderson & Serences, 2019; Swisher, Halko, Merabet, McMains, & Somers, 2007).

the post-behavioral-training MRI scan). The t-value of $t=2.56$ corresponds to an FDR-corrected p-value of 0.05.

While the subjects were performing the Dax categorization task in the scanner, they pressed one of two buttons to indicate which category they judged each stimulus to be. When interpreting the resulting fMRI activation, we must therefore be careful to distinguish between genuinely categorization-related activation and purely motor-related activation. The very design of the task can make these two types of activation difficult to pull apart, as every categorization decision is accompanied by its corresponding button press. Fortunately, our experimental design avoids this problem, as the internal state of the participants before training as opposed to after training are very different. Specifically, in the pre-training fMRI scan, the participants were pressing Dax and non-Dax buttons, using the same buttons as they would later use in the post-training fMRI scan. However, before behavioral training they do not yet have any category-structure information encoded in their brains. Thus, during this pre-training scan, their motoric and category-related activations were dissociated.

Therefore, in order to distinguish between motor-output and Dax-category neural information, we performed a cross-decoding analysis (see, e.g. Kaplan et al. (2015) for a review. Note that we used a classifier cross-decoding analysis instead of similarity-based analysis because there were only two button-press categories, and similarity matrices made from only two conditions are uninformative due to containing only one unique off-diagonal data point). Specifically, we trained a linear SVM classifier to distinguish between fMRI data elicited by the two different button-press responses during the post-training Dax task, and tested that same trained SVM on fMRI data from the pre-training Dax task. Following the logic laid out above, the rationale is as follows: if the SVM succeeds in classifying pre-behavioral-training fMRI test-data from a given brain area, then the activation in that area must be representing motor-output button-press responses rather than actual Dax-category information, because the pre-training fMRI data cannot contain any Dax-category information, due to the participants not yet knowing what the Dax-category is. In contrast, if the SVM fails to classify the pre-behavioral-training fMRI test-data, then the information that the SVM had extracted from its post-behavioral-training Dax-task fMRI data must have been about the Dax-category itself.

We carried out precisely this test, training the SVM with Matlab's built-in `fitsvm` function. The classifier's decoding performance was calculated for each participant in each ROI, and the average of the participants accuracies was tested against the 50% chance-level using a t-test, in the standard manner. We found that the only ROI with significant above-chance accuracy was the Postcentral Gyrus (somatosensory cortex), and even that region fails to reach significance after correcting for multiple comparisons ($t_{14} = 2.46$, $p < 0.04$ uncorrected, $p = 0.97$ corrected). Although it is not unexpected to find that somatosensory cortex contained button-press information, it is nonetheless reassuring to see that such information was present only in sensorimotor areas. In contrast, the frontal and parietal areas that our Representational Similarity Analyses found to contain information about the structure of the Dax-category did not show significant results in this button-press cross-decoding task (inferior frontal gyrus ROI: $t_{14} = -1.48$, $p > 0.9$ uncorrected; intraparietal sulcus: $t_{14} = -2.24$, $p > 0.95$ uncorrected), suggesting that the information that they encoded did indeed reflect the Dax category structure, rather than merely the button-presses that the subjects used for giving their responses while performing the categorization task. Moreover, the fact that the button-press cross-decoding performed somewhat worse than chance in those frontal and parietal regions suggests that their neural information

content had changed markedly between the pre-training and post-training scans, even though the subjects were engaging in the same sorts of button presses in both cases. This is precisely what would be expected if, as we suggest, those regions encoded Dax-category information in a manner that was distinct from merely encoding button-press responses.

A common finding in cross-decoding analyses is that the results are different, depending on which condition is used for training the classifier and which for testing. Specifically, if the two different conditions differ in how noisy they are, the cross-decoding classifiers typically perform best when trained on the less-noisy condition (Kaplan et al., 2015). Our present analyses also follow this pattern. The results described in the preceding paragraph were obtained when training the SVM on the post-behavioral-training fMRI scans, during which the subjects were successfully able to perform the Dax-category task. During the pre-behavioral-training fMRI scans, the subjects had to try to perform the Dax task even though they had not yet had an opportunity to learn which animals were Daxes and which were not. It is therefore to be expected that their neural responses during this scan would be noisy and highly variable, and that training an SVM on this pre-behavioral-training fMRI data would yield little cross-decoding transfer when tested on the post-behavioral-training scans. This is indeed what we found. When the SVM was trained on the pre-training scans and tested on the post-training scans, no regions reached statistical significance, even without any multiple comparison correction.

To visualize how the representational structure changes after behavioral training, we applied MDS on IFG. Qualitatively, the outer ring animals was clearly separated from other stimuli after behavioral training (Fig. 5, right), but not before training (Fig. 5, left).

Neural representations of categories in IFG become separated after training

Fig. 5. Left and Right: MDS plot of IFG during Dax discrimination task before (left) and after (right) the behavioral training. It can be seen that the outer-ring of animals (shown in saturated red) became markedly separated from the others after behavioral training. Middle: Representational change of each animal shape, showing the pre-vs-post training differences in MDS position. Note that the coordinates in the left figure were aligned to the right one via Procrustes analysis.

As a comparison, we also performed RSA with the pose-only model, that captured low-level visual model features like pitch and orientation (averaged across all animal shapes) to see whether there were low-level brain representations changes before and after behavioral training. Confirming that the pose-model reflects low-level visual processing, in both fMRI sessions, the

correlations with the pose-model were significant in early visual cortices only. Moreover, we find no differences in the fit of the pose-model for pre and post-training, indicating that there is no effect of categorization training on low-level visual cortex representations (Fig. 6). Overall, this suggested that widespread representational change was only observed in the task-relevant shape dimension, not the irrelevant viewing angle one.

Fig. 6. RSA correlation with the task-irrelevant Gabor filter model and the task-relevant prototype distance model. Each arrow shows an ROI's representational change elicited by behavioral training. Some named ROIs of particular interest are highlighted in black.

Previous studies (Mack, Love, & Preston, 2016) had demonstrated that hippocampus representations can adapt and reflect learned category structure of the current task. However, in our study the hippocampus RDM (extracted with Harvard-Oxford subcortical atlas) was not significantly correlated with our prototype distance model. A searchlight analysis with spherical radius of 3 voxels also revealed no such clusters around previously reported MNI coordinates. Future studies with higher voxel precision and a better hippocampus mask are needed to test whether the medial temporal lobe regions can represent category structure in a similar manner as the cortical regions reported here.

Representational change is task-dependent

Was the representational change widespread across the brain observed only when participants were performing the Dax categorization task? In other words, was attention on the task-relevant dimension necessary to separate Dax and non-Dax animals? To test this hypothesis, we applied the same RSA pipeline on beta activation images while participants were performing the passive viewing and pitch discrimination task. None of the ROIs were significantly correlated with the prototype distance model in either task (Fig. 7). This suggested that untangling of linearly inseparable category information might be an attentional or decisional effect, rather than a task-independent re-tuning of low-level visual cortex representations.

**IFG representational change
not present during passive task**

**Some object-sensitive regions do
not show category separation**

Fig. 7. Left: MDS plot of the IFG during the passive viewing task after behavioral training. Right: MDS plot of the LOC during the Dax discrimination task after behavioral training. In contrast to the category-separation that was shown in Fig.5, the plots here show no category separation

Task-dependent dynamic switching between different representations in multiple frontal and parietal regions

Recent studies have shown that dorsal pathway regions dynamically switch between different representations depending on the current task, with neural representational structures reconfiguring themselves in order to more strongly represent the information relevant to the task that is being performed at the time (Bracci, Daniels, & Op de Beeck, 2017; Vaziri-Pashkam & Xu, 2017). Although the primary question of interest in our present study was to investigate training-induced representational changes in the Dax task, the fact that our participants also performed a slant-discrimination task provides us with an opportunity to ask whether our data also shows this task-dependent representational switching effect.

In order to test this, we created a 7-by-7 pitch angle model where each entry in this RDM represents the difference between pitch angles. To construct the neural RDM for each ROI, for each animal pitch-angle we averaged the activation patterns across viewpoints and animal shapes. We then correlated this model RDM with that of all bilateral Harvard-Oxford ROIs using RSA to see how these model correlations changed depending on which task the participants were performing. Specifically, for each subject and each ROI we carried out a paired t-test, comparing the Fisher-z transformed pitch-task and Dax-task correlations across different task conditions. This comparison was performed only for the post-training scans, because the categorization task

model fit was trivially poor during the pre-training scans simply because the subjects had not yet learned how to perform that task.

Dynamic task-dependent representational switching would manifest itself as the pitch-model fitting better while the subjects were performing the pitch-task, and the prototype distance model fitting better during the Dax task. We found precisely this effect in several multiple frontal and parietal regions ($t > 3.084$, paired t-test, FDR-corrected across all ROIs). The intersection of these regions (Table 2) and those containing category-relevant representational changes (Table 1) were as follows: Middle Frontal Gyrus, Inferior Frontal Gyrus, Superior Parietal Lobe, posterior Supramarginal Gyrus and Intraparietal Sulcus.

Existing theories of how category learning affects dimensions of the stimulus space, and their relations to our current findings

In the present study, we investigated how learning a categorization task resulted in the transformation of a stimulus space. This question has previously been explored by classical theories of category learning, most notably the Generalized Context Model (Nosofsky, 1986). That model predicted that categorization training should induce an expansion of category-relevant dimensions. Augmented with the dimensional modulation theory (Goldstone, 1994), the theories also predicted that such expansion might be greater across category boundaries than within categories, leading to a classical categorical perception effect. Some compelling examples of such effects were found by Folstein and colleagues (e.g. Folstein, Palmeri & Gauthier, 2013), who found behavioral and fMRI evidence of dimensional modulation: an expansion of stimulus space across the category boundary, along the task-relevant dimension.

In contrast, the present study investigated a different hypothesis: in our prototype task design, a computationally efficient way to transform the linearly inseparable space into a separable one would be to represent all stimuli in terms only of their distance from the central prototype Dax animal. In such a transformation, which would be analogous to how kernel algorithms in machine learning can solve this sort of task, the stimulus space would not be radically warped. Recall that the animals were arranged in concentric rings around the center of the stimulus space, as shown in Figure 1. If the distance-from-the-center transformation were applied to this space, then all the animals within a given ring would become more similar to each other, even if they started off on opposite sides of the stimulus space. It would be as if the space started off like a stretched out Chinese fan, which then gets transformed by being folded back into a narrow strip. In contrast, dimensional expansion could stretch or compress the circular rings into ovals, but diametrically opposite sides of the stimulus space would always remain opposite to each other.

As Fig.6 shows, almost all brain areas showed an increased fit with this prototype distance model after training, compared to before. Some areas, most notably the Inferior Frontal Gyrus and Intraparietal Cortex, showed very marked increases in fit. Nonetheless, the prototype model was far from capturing everything in the data. As the MDS plots in Fig.5 of neural representational space in the Inferior Frontal Gyrus show, the stimuli in the outermost ring did indeed move further away from the central prototype after training, and to some extent the outer-ring stimuli may even have slightly bunched together. However, if the representations had perfectly matched the prototype model's predictions, then all of these outer ring stimuli would have collapsed together into a single point (and, moreover, the stimuli on the other rings would have collapsed together to their own separate points too). Clearly, no such collapse took place. So, although the degree of fit of the prototype model increased markedly from pre- to post-

training, it remained a very incomplete description of how the neural representational space actually behaved.

Although the IFG exhibited the nonlinear warping illustrated in Fig.5, this does not rule out the possibility that other brain regions might have shown more classical dimensional modulation effects. In our 2D stimulus space, all of the stimulus dimensions were relevant, so the prediction of the Generalized Context Model would be of an expansion in all directions. Our RSA methods would be unable to reveal a uniform expansion of this sort, as any uniform scaling leaves the relative similarities between different stimuli unchanged. To rectify this problem, we carried out an exploratory analysis, suggested by a reviewer, to examine: 1) whether categorization training resulted in an overall expansion of our stimulus space, and 2) whether categorization training differentially expanded representational distances across the category boundary more than within the boundary.

In short, we did not find statistically significant evidence for uniform expansion of this sort. That is not to say that such expansion was entirely absent: several ROIs, listed below, did indeed show weak evidence of expansion. However, as with all of the tests involving the 48 Harvard-Oxford ROIs in this study, multiple comparisons correction (using FDR) was carried out. After applying this correction, none of the ROIs survived as significantly showing the expansion effect.

The details of this additional analysis were as follows: for each ROI and each participant, we extracted distance entries from the 13-by-13 animal shape neural RDM. Entries were grouped by whether 1) both animals were inner-ring animals (inner-inner), 2) both animals were non-Daxes (outer-outer), and 3) one animal was Dax and the other wasn't (inner-outer). A two-way ANOVA was performed on the averaged entries in each of the 48 ROIs, where one factor was whether two animals were within or between the categories, and the other factor was before / after behavioral training. Even without multiple correction for the 48 ROIs, we found no main effect of animal category on representational distances in any ROI. Before multiple correction, a main effect of training on representational distances was observed in the Frontal Pole, Frontal Medial Cortex, Frontal Orbital Cortex, Frontal Operculum Cortex, anterior Middle Temporal Gyrus, and Inferior Temporal Gyrus. We found no significant interaction effect between animal category and training on representational distances, and no effects were significant in the Tukey post-hoc tests between pairs of animal category relations.

In summary, this additional analysis did not definitively rule out the hypothesis that representational space was expanded in all directions. Indeed, some regions showed a weak tendency towards this, but not, in our data at least, to a degree that reached statistical significance. Similarly, we did not find significant evidence that representational distances were expanded more across the category boundary compared to within each category. Although these results might at first sight seem inconsistent with the dimensional modulation theory, we believe that they are not so much inconsistent as simply inconclusive, for this particular question. Our stimuli and task were not designed to test for dimensional modulation of this sort, and indeed a study which seeks to be a sensitive probe of such questions would probably end up being structured very differently. The question of how dimensional modulation, i.e. expansion along task-relevant dimensions, might relate to more non-linear warping of stimulus space is an interesting one, and seems likely to be a fruitful area of investigation for future work.

Discussion

In the present study, we examined how learning a linearly inseparable category boundary affected neural representations across the brain. Our results suggested several findings. First, we found that after participants successfully learned this category boundary, the representations of Dax and non-Dax animals became linearly separated in a low-dimensional space. Second, this separation reflected a task-dependent attentional mechanism; it was only present when participants were performing the Dax categorization task and appeared most prominently in regions like the IFG and the IPS.

Previous fMRI studies suggested that learning-dependent changes during visual category learning paradigms can be observed in parietal cortex (Freedman & Assad, 2006; Hebart, Donner, & Haynes, 2012; Hebart, Schriever, Donner, & Haynes, 2016; Mack et al., 2013) and prefrontal cortex (Hebart et al., 2012, 2016; Jiang et al., 2007). Furthermore, it was hypothesized that these regions in the frontoparietal network represented abstract category or rule information independent of motor responses (Hebart et al., 2012) and stimulus types. Our results further suggest that linearly inseparable category boundaries could also be represented in similar frontoparietal network regions.

Recently, increasing evidence demonstrated that object representations were present in both the ventral and dorsal visual pathways (Bracci et al., 2017; Jeong & Xu, 2016; Konen & Kastner, 2008; Li, Mayhew, & Kourtzi, 2009; Vaziri-Pashkam & Xu, 2017). It was hypothesized that while visual representations in ventral pathway were largely task-independent, those in dorsal pathway were shaped by the current task to reflect the most diagnostic feature dimension (Bracci et al., 2017; Vaziri-Pashkam & Xu, 2017). The current study offers additional evidence that the posterior parietal cortex represents linearly inseparable category information only when this information is task-relevant.

Previous studies have hypothesized that whether representational changes in visual cortices were task-independent (i.e., persisted in passive viewing tasks) depended on how the stimulus morph space was formed (Folstein, Gauthier, & Palmeri, 2012; Folstein, Palmeri, Van Gulick, & Gauthier, 2015). According to that hypothesis, complex stimuli generated from a factorial space (select an origin and two orthogonal axes, and generate the stimuli by picking (x, y) values on this plane) should yield task-independent representational changes in ventral visual pathway regions. Although the stimuli used in the current study were also generated factorially, we did not observe significant representational changes in all ROIs during the passive viewing and pitch discrimination task. This difference might be due to our usage of a linearly inseparable category boundary, as opposed to the separable boundary used by Folstein and colleagues. Furthermore, these previous studies used fMRI rapid-adaptation paradigms to demonstrate these modulation effects (Folstein et al., 2013; Jiang et al., 2007), so future studies are needed to investigate whether or not experimental paradigm differences also played an interactive role in it.

Two initial motivations for the prototype design of the present study were as follows: first, this type of linearly inseparable task structure is simpler, and perhaps more ecologically valid, than the XOR task that is often thought of as the canonically inseparable task. Second, studies using non-linear kernels in machine learning (Schölkopf, Smola, & Bach, 2002) have used a similar task structure to demonstrate how a transformation of stimulus space can turn a linearly inseparable task into a separable one. The fact that we observed similar transformations in the post-training neural representational spaces does not, of course, imply that the brain actually implements this sort of kernel algorithm. However, both the machine learning work and our neural data suggest that transforming an inseparable stimulus space into a linearly

inseparable form is a computationally useful step for a system to take. The question of how the brain implements this step at the level of neural circuits remains unaddressed, and could be a fruitful area for future work. Another advantage of the prototype structure over the XOR task is that it could potentially be better suited for future studies of internal category structure, a topic that is somewhat under-investigated.

Previous studies (Nosofsky, 1986) have suggested that categorization training can expand category-relevant dimensions. Since all dimensions in our stimulus space were relevant to the categorization task, the generalized context model (GCM) would predict an overall expansion in our stimulus space. However, the prototype category structure of our stimulus space would not be learnable simply by expanding the overall space, as no such expansion would be able to move the inner “island” of Dax animals outside of its enclosing ring of non-Dax animals. Instead, some sort of non-linear transformation is required, and the kernel transformation mentioned in the previous paragraph can be a plausible mechanism.

Clearly, the present study involves the learning of only one task, so from this study alone it is impossible to know whether these findings will generalize to the learning of other computationally challenging tasks that require the transformation of stimulus space. We do not see any reason why the particular task and stimuli here should be non-representative, but only further studies can tell.

In summary, our study investigated how the neural representation of a stimulus space becomes reshaped by the learning of a task. Specifically, we investigated a task with a linearly inseparable prototype structure, going beyond the linearly separable tasks that have typically been used in previous studies. By visualizing neural representational structure with MDS, we found that some regions (notably, the inferior frontal gyrus, and the intraparietal sulcus) showed a marked untangling of the categories' neural patterns that made them more linearly separable after behavioral training. However, these neural changes did not appear to reflect permanent changes in representation, but instead were dynamically task-dependent, being observed only while participants were performing the categorization task, but not during passive viewing. The task-dependent nature of these changes, together with the fact that they were found to occur in frontal and parietal areas rather than ventral temporal cortices, suggests that they reflected attentional and decisional reweighting, rather than changes in object recognition templates. Although classical theories of category learning (Nosofsky, 1986) did a good job of explaining how linearly separable tasks might be learned, they were less able to account for linearly inseparable tasks such as the one used in the present study. Our results provide new insights into the non-linear warping of neural representational space, and how the brain uses such transformations to solve computationally challenging tasks.

References

- Abdi, H., Dunlop, J. P., & Williams, L. J. (2009). How to compute reliability estimates and display confidence and tolerance intervals for pattern classifiers using the Bootstrap and 3-way multidimensional scaling (DISTATIS). *NeuroImage*, 45(1), 89–95. <https://doi.org/10.1016/J.NEUROIMAGE.2008.11.008>
- Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*, 289–300.
- Bracci, S., Daniels, N., & Op de Beeck, H. (2017). Task Context Overrides Object- and Category-Related Representational Content in the Human Parietal Cortex. *Cerebral Cortex*, 27(1), 310–321. <https://doi.org/10.1093/cercor/bhw419>
- Braunlich, K., Liu, Z., & Seger, C. A. (2017). Occipitotemporal Category Representations Are Sensitive to Abstract Category Boundaries Defined by Generalization Demands. *The Journal of Neuroscience*, 37(32), 7631–7642. <https://doi.org/10.1523/JNEUROSCI.3825-16.2017>
- Cutzu, F., & Edelman, S. (1998). Representation of object similarity in human vision: psychophysics and a computational model. *Vision Research*, 38(15–16), 2229–2257. [https://doi.org/10.1016/S0042-6989\(97\)00186-7](https://doi.org/10.1016/S0042-6989(97)00186-7)
- Desikan, R. S., Ségonne, F., Fischl, B., Quinn, B. T., Dickerson, B. C., Blacker, D., ... Hyman, B. T. (2006). An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *Neuroimage*, 31(3), 968–980.
- DiCarlo, J. J., & Cox, D. D. (2007). Untangling invariant object recognition. *Trends in Cognitive Sciences*, 11(8), 333–341. <https://doi.org/10.1016/j.tics.2007.06.010>
- Edelman, S., Bülthoff, H. H., & Bülthoff, I. (1999). Effects of parametric manipulation of inter-stimulus similarity on 3D object recognition. *Spatial Vision*.
- Edelman, S., Grill-Spector, K., Kushnir, T., & Malach, R. (1998). Toward direct visualization of the internal shape representation space by fMRI. *Psychobiology*, 26(4), 309–321. <https://doi.org/10.3758/bf03330618>
- 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(12), 2379. <https://doi.org/10.1364/JOSAA.4.002379>
- Folstein, J. R., Gauthier, I., & Palmeri, T. J. (2012). How category learning affects object representations: Not all morphspaces stretch alike. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 38(4), 807–820. <https://doi.org/10.1037/a0025836>
- 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>
- Folstein, J. R., Palmeri, T. J., Van Gulick, A. E., & Gauthier, I. (2015). Category Learning Stretches Neural Representations in Visual Cortex. *Current Directions in Psychological Science*, 24(1), 17–23. <https://doi.org/10.1177/0963721414550707>
- Freedman, D. J., & Assad, J. A. (2006). Experience-dependent representation of visual categories in parietal cortex. *Nature* 2006 443:7107, 443(7107), 85. <https://doi.org/10.1038/nature05078>
- Hebart, M. N., Donner, T. H., & Haynes, J.-D. (2012). Human visual and parietal cortex encode visual choices independent of motor plans. *NeuroImage*, 63(3), 1393–1403. <https://doi.org/10.1016/J.NEUROIMAGE.2012.08.027>

- Hebart, M. N., Schriever, Y., Donner, T. H., & Haynes, J.-D. (2016). The Relationship between Perceptual Decision Variables and Confidence in the Human Brain. *Cerebral Cortex*, 26(1), 118–130. <https://doi.org/10.1093/cercor/bhu181>
- Henderson, M., & Serences, J. T. (2019). Human frontoparietal cortex represents behaviorally relevant target status based on abstract object features. *Journal of Neurophysiology*, 121(4), 1410–1427. <https://doi.org/10.1152/jn.00015.2019>
- Jeong, S. K., & Xu, Y. (2016). Behaviorally Relevant Abstract Object Identity Representation in the Human Parietal Cortex. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, 36(5), 1607–1619. <https://doi.org/10.1523/JNEUROSCI.1016-15.2016>
- Jiang, X., Bradley, E., Rini, R. A., Zeffiro, T., VanMeter, J., & Riesenhuber, M. (2007). Categorization Training Results in Shape- and Category-Selective Human Neural Plasticity. *Neuron*, 53(6), 891–903. <https://doi.org/10.1016/J.NEURON.2007.02.015>
- Knowlton, B. J., & Squire, L. R. (1993). The learning of categories: parallel brain systems for item memory and category knowledge. *Science (New York, N.Y.)*, 262(5140), 1747–1749. <https://doi.org/10.1126/SCIENCE.8259522>
- Konen, C. S., & Kastner, S. (2008). Two hierarchically organized neural systems for object information in human visual cortex. *Nature Neuroscience*, 11(2), 224–231. <https://doi.org/10.1038/nn2036>
- Kriegeskorte, N., Mur, M., & Bandettini, P. A. (2008). Representational similarity analysis – connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, 2, 4. <https://doi.org/10.3389/neuro.06.004.2008>
- Li, S., Mayhew, S. D., & Kourtzi, Z. (2009). Learning Shapes the Representation of Behavioral Choice in the Human Brain. *Neuron*, 62(3), 441–452. <https://doi.org/10.1016/J.NEURON.2009.03.016>
- Li, S., Ostwald, D., Giese, M., & Kourtzi, Z. (2007). Flexible coding for categorical decisions in the human brain. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, 27(45), 12321–12330. <https://doi.org/10.1523/JNEUROSCI.3795-07.2007>
- 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., 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>
- Nili, H., Wingfield, C., Walther, A., Su, L., Marslen-Wilson, W., & Kriegeskorte, N. (2014). A Toolbox for Representational Similarity Analysis. *PLoS Computational Biology*, 10(4), e1003553. <https://doi.org/10.1371/journal.pcbi.1003553>
- Nosofsky, R. M. (1986). Attention, similarity, and the identification–categorization relationship. *Journal of Experimental Psychology: General*, 115(1), 39.
- Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, 77(3p1), 353.
- Reber, P. J., Stark, C. E. L., & Squire, L. R. (1998). *Cortical areas supporting category learning identified using functional MRI (nondeclarative memory occipital cortex frontal cortex)*. *Neurobiology* (Vol. 95).
- Schölkopf, B., Smola, A. J., & Bach, F. (2002). *Learning with kernels: support vector machines*,

- regularization, optimization, and beyond*. MIT press.
- Seger, C. A., Poldrack, R. A., Prabhakaran, V., Zhao, M., Glover, G. H., & Gabrieli, J. D. . (2000). Hemispheric asymmetries and individual differences in visual concept learning as measured by functional MRI. *Neuropsychologia*, 38(9), 1316–1324.
[https://doi.org/10.1016/S0028-3932\(00\)00014-2](https://doi.org/10.1016/S0028-3932(00)00014-2)
- Swisher, J. D., Halko, M. A., Merabet, L. B., McMains, S. A., & Somers, D. C. (2007). Visual topography of human intraparietal sulcus. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, 27(20), 5326–5337.
<https://doi.org/10.1523/JNEUROSCI.0991-07.2007>
- Torgerson, W. S. (1952). Multidimensional scaling: I. Theory and method. *Psychometrika*, 17(4), 401–419. <https://doi.org/10.1007/BF02288916>
- Vaziri-Pashkam, M., & Xu, Y. (2017). Goal-Directed Visual Processing Differentially Impacts Human Ventral and Dorsal Visual Representations. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, 37(36), 8767–8782.
<https://doi.org/10.1523/JNEUROSCI.3392-16.2017>

Tables

Regions	T values (dof = 14)	Corrected P values
Frontal Pole	4.177	0.004
Insular Cortex	3.868	0.005
Superior Frontal Gyrus	4.019	0.004
Middle Frontal Gyrus	4.384	0.003
Inferior Frontal Gyrus, pars triangularis	6.008	0.001
Inferior Frontal Gyrus, pars opercularis	5.426	0.001
Precentral Gyrus	4.563	0.003
Middle Temporal Gyrus, anterior division	2.980	0.020
Middle Temporal Gyrus, posterior division	3.248	0.013
Middle Temporal Gyrus, temporooccipital part	4.581	0.003
Inferior Temporal Gyrus, temporooccipital part	3.566	0.008
Postcentral Gyrus	3.619	0.007
Superior Parietal Lobule	4.451	0.003
Supramarginal Gyrus, anterior division	3.863	0.005
Supramarginal Gyrus, posterior division	3.243	0.013
Angular Gyrus	4.231	0.004
Intraparietal Sulcus	7.031	<0.001
Lateral Occipital Cortex, inferior division	4.119	0.004
Paracingulate Gyrus	3.204	0.014
Cingulate Gyrus, posterior division	3.004	0.020
Precuneous Cortex	4.553	0.003
Frontal Orbital Cortex	5.229	0.002
Occipital Fusiform Gyrus	2.780	0.028
Frontal Operculum Cortex	4.193	0.004
Occipital Pole	3.844	0.005

Table 1. Harvard-Oxford Atlas regions significantly correlating with the prototype distance model during the categorization task in the post-training session.

Regions	T values (dof = 14)	Corrected P values
Middle Frontal Gyrus	3.808	0.015
Inferior Frontal Gyrus, pars triangularis	3.587	0.018
Precentral Gyrus	3.084	0.028
Superior Parietal Lobule	4.835	0.006
Supramarginal Gyrus, posterior division	3.320	0.023
Lateral Occipital Cortex, superior division	3.940	0.015
Supracalcarine Cortex	3.249	0.023

Table 2. Harvard-Oxford Atlas regions significantly correlating with the pitch model more during the pitch task than categorization task in the post-training session.