

The Animacy Continuum in the Human Ventral Vision Pathway

Long Sha^{1,2}, James V. Haxby^{1,3}, Herve Abdi⁴, J. Swaroop Guntupalli¹,
Nikolaas N. Oosterhof^{1,3}, Yaroslav O. Halchenko¹,
and Andrew C. Connolly¹

Abstract

■ Major theories for explaining the organization of semantic memory in the human brain are premised on the often-observed dichotomous dissociation between living and nonliving objects. Evidence from neuroimaging has been interpreted to suggest that this distinction is reflected in the functional topography of the ventral vision pathway as lateral-to-medial activation gradients. Recently, we observed that similar activation gradients also reflect differences among living stimuli consistent with the semantic dimension of graded animacy. Here, we address whether the salient dichotomous distinction between living and nonliving objects is actually reflected in observable measured brain activity or whether previous observations of a dichotomous dissociation were the illusory result of stimulus sampling biases. Using fMRI, we measured neural responses

while participants viewed 10 animal species with high to low animacy and two inanimate categories. Representational similarity analysis of the activity in ventral vision cortex revealed a main axis of variation with high-animacy species maximally different from artifacts and with the least animate species closest to artifacts. Although the associated functional topography mirrored activation gradients observed for animate–inanimate contrasts, we found no evidence for a dichotomous dissociation. We conclude that a central organizing principle of human object vision corresponds to the graded psychological property of animacy with no clear distinction between living and nonliving stimuli. The lack of evidence for a dichotomous dissociation in the measured brain activity challenges theories based on this premise. ■

INTRODUCTION

Evidence for the existence of an animate–inanimate division in the human ventral vision pathway has been well documented in neuropsychology (Warrington & Shallice, 1984), electrophysiology (Kiani, Esteky, Mirpour, & Tanaka, 2007), and neuroimaging (Kriegeskorte, Mur, Ruff, et al., 2008; O’Toole, Jiang, Abdi, & Haxby, 2005; Hanson, Matsuka, & Haxby, 2004). This division is associated with distinctive brain activity topographies in ventral temporal (VT) cortex, lateral occipital (LO) cortex, and STS (Mahon, Anzellotti, Schwarzbach, Zampini, & Caramazza, 2009; Chao, Haxby, & Martin, 1999), with lateral VT, superior LO, and STS showing more activation for animate objects and medial VT and inferior LO showing more activation for inanimate objects.

Several theories have been proposed to account for this organization. Modality-specific accounts, such as the sensory-functional hypothesis (SFH; Allport, 1985; Warrington & Shallice, 1984), propose that it arises from an interaction between (1) the uneven distribution of sensory and motor properties central to different cate-

gory domains and (2) neural representations that are distributed across property-specific cortical fields. In contrast, the domain specificity hypothesis (Caramazza & Mahon, 2003; Caramazza & Shelton, 1998) proposes that evolutionary forces requiring efficient processing of important category domains led to the development of mental modules that are dedicated for representing specific category domains (see also Kanwisher & Dilks, 2013). Some (Mahon & Caramazza, 2009; Martin, 2007) have proposed a hybrid model that incorporates both domain-specific and property-based representations. And others have proposed that this organization arises from differences in retinal eccentricity biases across categories (Levy, Hasson, Avidan, Hendler, & Malach, 2001).

Although the animate–inanimate distinction is a major feature of the representational space in the ventral vision pathway, it is not the only one. Multivariate pattern (MVP) analyses of functional neuroimaging data have indicated that a broad range of object categories evoke specific distributed patterns of neural activity (Connolly, Guntupalli, et al., 2012; Eger, Ashburner, Haynes, Dolan, & Rees, 2008; Kriegeskorte, Mur, Ruff, et al., 2008; O’Toole et al., 2005; Hanson et al., 2004; Haxby et al., 2001) and that these patterns have a common basis across individuals (Haxby et al., 2011). Within the animate domain,

¹Dartmouth College, ²New York University, ³University of Trento, ⁴The University of Texas

our previous work (Connolly, Guntupalli, et al., 2012) showed that patterns of response in LO and VT cortex could reliably differentiate animal species, including coarse distinctions, such as primates versus birds, and fine distinctions, such as monkeys versus lemurs. More surprisingly, using multidimensional scaling (MDS), we found that the dominant dimension was a continuum, which we interpreted as spanning the abstract psychological dimension of animacy from “most animate” (primates) to “least animate” (insects). This continuum was associated with a cortical topography similar to the topography described for the animate versus inanimate distinction, with patterns of response for “less animate” species similar to patterns for inanimate categories from other studies (Mahon et al., 2009; Martin, 2007; Chao et al., 1999). This pattern suggests that the animate side of the animate–inanimate division is really a continuous dimension and that the “least animate” animals may be close to the representation of bona fide inanimates.

Here, we test the hypothesis that the neural representational space in the ventral pathway is dominated by a continuous dimension from most animate to least animate to inanimate categories. Our question is whether the inanimate stimuli have distinct representations beyond the end of this continuum or do they have representations similar to the animal species with “low animacy.” We define level of animacy as the degree of similarity to the animate prototype, which we assume to be humans. Thus, animals that are phylogenetically close to humans, such as monkeys, rank higher on the animacy scale than more distant relatives such as fish. We do not expect an exact match based on the amount of shared genetic material, however, because the relevant properties of animate entities such as perceived intelligence and sociability do not always predict genetic similarity. For example, creatures like dogs may rank higher than some less familiar primate cousins—an effect we observed in a previous study in which dog faces were closer to human faces than were monkey faces in the neural representational space (Connolly, Gobbini, & Haxby, 2012).

The notion of graded animacy is not new. The classical Greek philosopher Aristotle described an animacy hierarchy in his treatise *On the Soul* (Latin: *De Anima*; Aristotle, 1986) in which he placed plants at the very bottom, followed by the simple animals like worms, and proceeding through the higher animals with humans near the top—exceeded only by the gods. Ranking of animals and objects according to the abstract notion of animacy is a well-known phenomenon in linguistics in which animacy hierarchies constrain the rules of grammar (Young & Morgan, 1987). The operationalization of animacy in this study is based on our own intuitive rankings of animals along this continuum, which we presume to be in line with both common intuition and cross-cultural folk taxonomies (Atran, 1998).

Our goals for this study were (1) to replicate our previous findings an animacy continuum in object vision cortex (Connolly, Gobbini, et al., 2012) and (2) to determine whether a sharp boundary exists between the “least

animate” living stimuli and the inanimate objects. Our stimuli were images of animal species with varying levels of animacy—primates, quadruped mammals, birds, fish, invertebrates—and of two inanimate categories of tools (Figure 1A). We measured the patterns of neural response to these stimuli with fMRI and analyze the results using a variety of MVPA techniques to characterize the high-dimensional neural representational space.

METHODS

Stimuli

All image stimuli were collected online from the following 12 categories: (1) humans, (2) chimpanzees, (3) cats, (4) giraffes, (5) pelicans, (6) warblers, (7) clownfish, (8) stingrays, (9) ladybugs, (10) lobsters, (11) hammers, and (12) keys (see Figure 1A). These categories could be ranked and grouped based on a priori level of animacy from highest to lowest as primates, quadruped mammals, birds, fish, invertebrates, and tools. Within each level of animacy, we also tried to control for real-world size, so that a typically large object is paired with a small object in each level. This was done to eliminate a potential confound between levels of animacy and real-world size (Konkle & Oliva, 2012). To minimize the effect of low-level image statistics on our early visual area signals, we first tiled the images then randomized the size of the individual images for each row (Figure 1A). There were 20 different images within each category. One tiled image was created based on the same image, and one was based on the left-right mirrored version. There were 40 images for each category, resulting in 480 tiled images in total for the experiment. Each presentation was a randomly selected image from the image pool. Each visual stimulus was presented at about $13^\circ \times 13^\circ$ visual angles to participants in the fMRI scanner.

Behavioral Judgments

For each stimulus category in the experiment, we collected behavioral judgments using Amazon’s Mechanical Turk (<https://www.mturk.com>). Forty anonymous raters were recruited for each test. Raters saw three images from different categories lined up on the same screen and were required to choose the odd one out based on general object category. Each triplet (of 220 possible ones) was rated by four different individuals. We scored the results by creating a stimulus category by stimulus category dissimilarity matrix (DSM), and for each image judged as odd, we added one to the two cells corresponding to the dissimilarity between the odd image and the other two images.

fMRI Participants

Eleven adult participants (seven men) from the Dartmouth College community participated in the experiment. All




Figure 1. Experimental design, main behavioral and neural findings. (A) We tested two inanimate—keys and hammers—and 10 animate categories that spanned the range of animacy from humans and chimpanzees to ladybugs and lobsters. Tiled stimulus arrays were used to mitigate any influence of category-specific biases in the retinal image. (B) Left: The DSM derived from behavioral judgments shows a distinct separation of animate and inanimate objects with a secondary axis reflecting animacy among the animates. Right: Classical MDS of behavioral DSM. (C) Left: Neural similarity structure derived from activity in LO cortex shows no evidence of a clear animate–inanimate division. Right: MDS derived using STATIS. δ = dissimilarity; λ = eigenvalue of each PC; τ = inertia for each PC expressed as percentage of variance accounted for. Ellipses represent 95% confidence intervals for the factor scores based on bootstrapped data resampling.

participants provided informed consent and were cleared for safety in MRI scanning before the experiment. Participants were paid for their time. All experimental procedures and consent materials were approved by the institutional review board of Dartmouth College, the Committee for the Protection of Human Subjects.

Procedure

For each trial, three images were shown consecutively to the participants. Each image was presented for 500 msec without gaps between images, and events were followed by a 4500-msec ISI. During image presentation, a fixation cross appeared at the center of the screen. Participants were instructed to pay attention to the images and were free to move their eyes while performing a simple 1-back task for all image triplets. If the last image in the triplets belonged to the same category as the previous two, participants were asked to report “same” by pushing the response button in the scanner, otherwise to report “different.” The mismatch trials and blank trials—where only a fixation cross appeared for 6 sec—were pseudorandomly placed in the experiment for participants to maintain attention. There were seven functional acquisition runs. Within each run, Type 1, Index 1 sequence for first-order counterbalancing of 14 trial types

(12 categories, 1 mismatch trial, and 1 blank trial) was created for run-wise presentations (Aguirre, 2007; Finney & Outhwaite, 1956). Each trial type was repeated six times. In addition, three leading dummy trials and one trailing dummy trial were added, resulting in 88 trials per run.

Image Acquisition

Brain images were acquired using a 3-T Philips Achieva Intera scanner with a 32-channel head coil. The functional imaging used gradient-echo EPI with SENSE reduction factor of 2. The MR parameters were echo time/repetition time = 35 msec/2000 msec, flip angle = 90° , resolution = 3×3 mm, matrix size of 80×80 , and field of view = 240×240 mm. There were 42 transverse slices with full-brain coverage, and the slice thickness was 3 mm with no gap. Slices were acquired in an interleaved order. Each of the seven functional acquisition runs included 264 functional acquisitions and 12 dummy acquisitions for a total time of 552 sec per run. At the end of each experiment a single, high-resolution T1-weighted (echo time/repetition time = 4.53 msec/9848 msec) anatomical scan was acquired with a 3-D turbo field echo sequence. The voxel resolution was $0.938 \times 0.938 \times 1.0$ mm with a bounding box matrix of $256 \times 256 \times 160$ (field of view = $240 \times 240 \times 160$ mm).

Image Preprocessing

To reduce noise related to participants' head movements and scanner drift, functional images were preprocessed using the following steps: First, images were corrected for slice acquisition time because of interleaved slice order within each functional acquisition. Second, participants' head movements were corrected by spatially registering all functional images to the last functional acquisition, which was closest in time to the anatomical high resolution scan. Third, data were despiked to remove signals unlikely to be physiological activities. Fourth, each functional acquisition run was detrended using up to third-order polynomials to remove any scanner drift, and motion parameters were also regressed out. Stimulus-specific BOLD signal response patterns were estimated using the general linear model for each run separately, resulting in 12 full-brain beta patterns for each of seven runs for each participant. The functional data were not spatially smoothed. All preprocessing steps were performed using AFNI (Cox, 1996).

Surface Mapping

Brain surfaces for each participant were generated using FreeSurfer's recon-all software (Fischl, 2012). Surfaces were transformed into AFNI/SUMA format and resampled to a standard surface mesh grid at various resolutions. The preparation of the surfaces in this manner was facilitated using a wrapper script (mvpa2-prep-afni-surf), which itself is freely available through PyMVPA. The surfaces were created from the T1-weighted anatomical images. The mapping of functional data and results from volume space to surface space was done using AFNI's 3dVol2Surf program. The mapping of data from surface to volume was done using AFNI's 3dSurf2Vol program.

Pattern Classification

We used linear support vector machine (SVM; Chang & Lin, 2011; Cortes & Vapnik, 1995; implemented in PyMVPA, Hanke et al., 2009) for all classification analyses with leave-one-run-out cross validation strategy.

Whole-brain Surface-based Searchlight

Conventional whole-brain searchlight analysis assigns a desired computation to a spherical local ROI centered on a given voxel and iterates the same computation through all the voxels in the brain volume (Kriegeskorte, Goebel, & Bandettini, 2006), mapping the results of some analysis back to the center voxel. Although this method is excellent for mapping informational content throughout the brain volume, it runs the risk of including voxels in the same searchlight that are not actually near

to each other on the cortical manifold, for example, voxels on either side of the sylvian fissure. Instead of a volume-based searchlight, here we use a surface-based searchlight (Oosterhof, Wiestler, Downing, & Diedrichsen, 2011; Oosterhof, Wiggert, Diedrichsen, Tipper, & Downing, 2010), which defines each local ROI centered on a surface node from a standard surface mesh (Fischl, 2012). For each surface node, the algorithm identifies the set of neighboring surface nodes and selects their corresponding voxels from the participant's native brain volume to define the searchlight's patterns. The results of the multivariate calculation are then mapped back to the corresponding central surface node. This method takes advantage of the surface anatomical alignment across participants and defines more restricted searchlight ROIs containing only voxels that are adjacent on the cortical manifold, as distances between nodes on the surfaces are computed along the cortical surface rather than using a Euclidean measure as used in volume-based approaches. Searchlight volumes were restricted to include 100 neighboring voxels per central surface node. For the definition of searchlight centers, we used a standard surface mesh with a resolution corresponding to a linear icosahedron tessellation with 32 linear divisions of each of the 20 triangles of a standard icosahedron (see AFNI's MapIcosahedron). This resolution results in 20,484 surface nodes for each full-brain surface reconstruction, which is close to the resolution provided by 3 mm isotropic voxels with about 2.5 mm between node pairs. For the definition of the surface nodes that make up a local neighborhood, we used a higher-resolution standard mesh with an icosahedron tessellation with 128 linear divisions with 327,684 nodes per full brain surface reconstruction. Searchlight analyses were implemented using PyMVPA (Hanke et al., 2009).

For the first searchlight analysis, we computed the correlation between the behavioral DSM and individual participant's voxel DSM within each local searchlight, and we iterated this process for the whole-brain surface. We then performed one-sample *t* test against zero for each surface node across all participants to generate a group map (Figure 2A).

For the second searchlight analysis (ROI definition), we performed 12-way classification using SVM classifiers with a leave-one-run-out cross-validation strategy within each local searchlight, and we iterated this process for the whole-brain surface. We then performed a one-sample *t* test against chance (.0833) for each surface node across all 11 participants and constructed a group mask by thresholding the *t* statistic map at $t(10) \geq 4$ ($p = .003$, two-tailed, uncorrected for multiple comparisons). We then used the SUMA clustering program (SurfClust) to isolate a continuous group mask, with threshold of $t > 4$ and a size criterion of 100 contiguous surface nodes. The ROI in each participant included the late visual areas (LO and VT), similar to our previous late visual mask (Connolly, Guntupalli, et al., 2012). The surface masks were then registered to the brain volume. Because the mapping from the standard surface

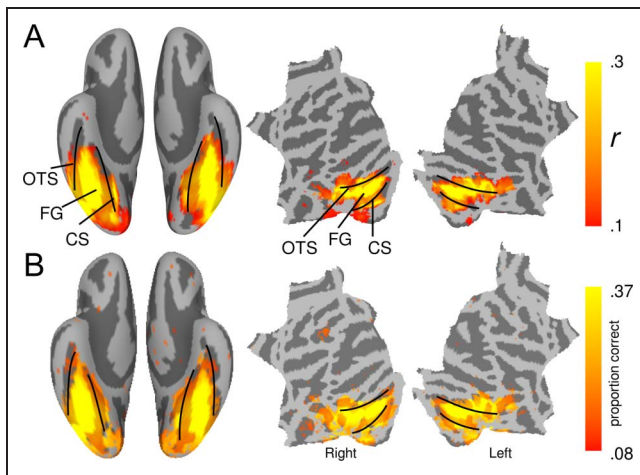


Figure 2. Searchlight results. (A) Correspondence between neural similarity and behavioral judgments was assessed as the correlation between the behavioral DSM and the neural DSM at each searchlight location. Mean correlation values were thresholded using a group t statistic for $r > 0$, $t(10) = 3$, $p < .02$. (B) Twelve-way classification accuracy was assessed at each searchlight. The mean proportion correct was thresholded at $t(10) = 4.0$, $p < .001$, and this map was clustered to generate the group ROI mask. CS = collateral sulcus; FG = fusiform gyrus; OTS = occipito-temporal sulcus.

mesh to individual volumes is idiosyncratic, each participant's volume mask contained a unique number of voxels with a mean of 2395 ± 88 .

Feature Scaling

Often MVP analyses incorporate a voxel selection step, which involves using some criterion to select the most informative voxels while discarding noisy ones. Voxel selection is an important step for reducing noise, which helps in assessing the true informational content of a data set. However, this practice has some drawbacks. For example, it is necessary when selecting voxels to choose a threshold criterion, either keeping a set percentage or a set number of features. Also, different data folds used for cross-validation will likely select different sets of features, and selected features in one participant are not guaranteed to overlap with selected features in other participants. These latter facts make it difficult to assess the topographical consistency of voxels that carry information. To avoid the arbitrary nature of threshold setting and to keep all of the voxels in our analysis, while at the same time attenuating noise, we chose voxel scaling as an alternative to voxel selection. For every data fold used for cross-validation, the training data were used to compute an F statistic based on a one-way ANOVA, which is a standard sensitivity measure often used in voxel selection. Then instead of discarding voxels with an F value below some criterion, we multiplied each voxel by its F value (after normalizing the set of F values to make their sum of squares equal to one). Conceptually, feature scaling

and feature selection are equivalent: Feature selection can be thought of as a special case of feature scaling where the weights for voxels below the criterion are set to zero, whereas the weights for those above are set to one. For classification, it is important when selecting voxels to base the selection criterion only on the training data to avoid overfitting and artificially boosting classification accuracies (Kriegeskorte, Simmons, Bellgowan, & Baker, 2009). This is also true for voxel scaling, which works the same as voxel selection for cross-validation purposes. A unique scaling F map is generated for each data fold, and this map is based only on the training data.

STATIS

STATIS is an extension of PCA used to analyze three-way data and has been used previously in neuroimaging (Abdi, Dunlop, & Williams, 2009; O'Toole et al., 2007; Shinkareva, Ombao, Sutton, Mohanty, & Miller, 2006; Kherif et al., 2003). STATIS optimally combines data tables from different sources then performs PCA on the combined data table, which is called the "compromise." STATIS provides a low-dimensional characterization of the informational structure, which entails simply plotting the factor scores for the observations. The variable (voxel) loadings describe the contributions of each variable to each principal component (PC). In addition, bootstrap data resampling provides for the estimation of confidence intervals for the factor scores of the stimuli. These confidence intervals are indicated in our plots using ellipses around the factor score centers. STATIS is well-suited for analyzing fMRI data because the only requirement for combining individual data tables is that they have the same number of rows (i.e., observations or stimuli). The columns (alternatively: variables, features, voxels) need not correspond across data tables. This is so because the individual tables may be combined horizontally—that is, a table with i rows and j columns is combined with a table with i rows and k columns to form a table with i rows and $j + k$ columns. Thus, voxels from different participants may be combined in the same compromise table without aligning brains to a standard template. Because STATIS has been thoroughly written about elsewhere, we refer the reader to Abdi, Williams, Valentin, and Bennani-Dosse (2012) for a tutorial and detailed mathematical definitions.

We used STATIS at several stages in our analyses. First, we used STATIS to optimally combine the seven runs from each individual participant into a single $12 \times n$ data table, where 12 is the number of stimulus categories and n is the size of participant's ROI. This is the first step of STATIS without performing PCA on the compromise tables. Before computing the participant-specific run-wise compromise tables, the columns (voxels) of each participant's data set were scaled using the feature scaling method described above and then centered by subtracting the column-wise means from the values in each

column. The rows were centered in a similar manner. Finally, before generating the compromise, each row was scaled by its Euclidean norm (square root of sum of squares). We chose to center and scale the rows in this manner, because this transformation is involved in calculating Pearson correlation distance between observations, which is a standard way to calculate neural DSMs with fMRI data. Because the voxels across the runs of a given participant correspond to each other, we computed the participant-specific compromise as the weighted linear sum of run-wise data tables (Abdi et al., 2012). We then computed a group-wise STATIS solution on 11 such compromise tables. The factor scores and voxel loadings for this solution are reported in the main text. Because the group compromise matrix was made up of participant-specific, horizontally stacked datasets, the PCA of this matrix yields voxel loadings for each voxel in each participant. The factor scores and voxel loadings for this solution are reported in the main text. To evaluate the shared functional topography associated with the voxel loadings across participants, the volumetric results were mapped to the standard surface mesh grid as described above.

Because STATIS provides transformation matrices from the original high-dimensional space into the lower-dimensional PC space, it is possible to project new supplemental observations into the PC space if those observations are on the same set of original variables. We exploited this feature by computing STATIS on subsets of the stimuli (i.e., just primates and tools or just the nonprimate animals) and then mapped the left-out stimuli into the PC space as supplemental observations (Figure 4). We also exploit this feature for using STATIS for classification with factor scores (Figure 5). For each data fold in cross-validation, STATIS is computed using the training data only, and then the training and testing data are mapped into the PCA space based on the training data to avoid “peeking.”

Data and Code

The implementation of STATIS that we used in this article is available as a part of PyMVPA (www.pymvpa.org). In addition, we are supplying the data and the Python code in the form of an IPython notebook (ipython.org/notebook.html) that we used to run the STATIS and classification analyses. The purpose of this material is to allow readers to reproduce the nonstandard analyses reported in this article and so that readers can explore first-hand the effects of changing various analysis parameters. The data included are anonymous, preprocessed, and masked to include just the ROI. To make it easier for readers to run these analyses, the IPython notebook and data are provided within a modified NeuroDebian (neuro.debian.net/) virtual machine that has been updated to include all of the Python dependencies needed to run the analyses. This material is available for download and is presently hosted at haxbylab.dartmouth.edu.

RESULTS

Behavioral Ratings of Similarity

We collected behavioral similarity ratings from 40 anonymous raters using an online survey. We computed an average behavioral DSM based on these ratings (Figure 1B). MDS of this DSM (Figure 1B, right) reveals a sharp distinction between animate and inanimate categories expressed by the first dimension. The second dimension captures variation that suggests an animacy gradient that ranges from mammals to invertebrates and fish. Thus, the behavioral ratings of the stimuli show both a clear distinction between animates and inanimates and a continuous animacy gradient within the animate domain.

Searchlight Analyses and ROI Definition

To assess the distribution of category-sensitive brain regions and also to define an ROI for further analysis, we began our analyses using whole-brain multivariate searchlights.

First, we searched for brain regions with neural response DSMs that resembled the behavioral rating DSM. We used a representational similarity analysis (Kriegeskorte, Mur, & Bandettini, 2008) surface-based searchlight (Connolly, Guntupalli, et al., 2012; Oosterhof et al., 2010, 2011; Kriegeskorte et al., 2006) to map correlations between behavioral and neural DSMs. Correlations between the behavioral and neural DSMs were highest in bilateral LO and VT cortices (Figure 2A). The behavioral ratings show two separate dimensions: (1) a strong animate–inanimate divide and (2) an animacy gradient on the second dimension. The similarity searchlight showed reliable but somewhat weak correlations between neural similarity and behavioral similarity (Figure 3A, $r \leq .3$). It is not possible to know given these correlations which aspects of behavioral and neural similarity are correlated with each other. Therefore, it is necessary to further evaluate the neural representational structure to assess the relative contributions of an animacy dichotomy versus an animacy continuum.

A second searchlight analysis used a 12-way SVM (Chang & Lin, 2011; Cortes & Vapnik, 1995; SVM) pattern classifier to map regions with distinguishing information about our 12 categories. Again bilateral LO and VT, that is, nearly all of object vision cortex, showed significantly above chance classification accuracy (Figure 2B). Next, we defined an ROI for further analysis using the surface nodes with highest average MVP classification accuracy across participants. Note that, in both the representational similarity analysis and SVM searchlight analyses, surface nodes in early vision along the calcarine sulcus are visible on the results maps (Figure 2A and B). These surface nodes were not contiguous with the larger sets of surface nodes identified throughout LO and VT, and therefore, these did not cluster with the larger set, leaving only the voxels that can be seen as colored voxels in




Figure 3. (A) The complete set of factor scores for STATIS analysis in ROI. Error bars show the *SEM*. Each bar plot is shown on the same scale for the y axis, which ranges from -0.06 to 0.06 . Note that although some lower-order PCs (e.g., #7) show separation between low-animate and inanimate objects, none of these show a systematic difference. (B) The functional topographies associated with the first two PCs, rendered as the mean voxel loadings across participants. The voxel loadings were calculated for each voxel in the volumes for each participant and then mapped to the standard surface mesh grid. The color scales are all the same across PCs, which range from -0.0018 to 0.0018 . Like for the factor scores, the absolute magnitudes of individual loadings are not meaningful because the data were normalized before analysis. Only the relative magnitudes and the sign carry meaning. (C) Scree plot shows the percentage of variance accounted for by each PC for the STATIS analysis in the ROI. PC1 and PC2 together account for 65% of the variance. CS = collateral sulcus; FG = fusiform gyrus; OTS = occipito-temporal sulcus; Occ. pole = occipital pole.

Figure 3. The above-chance 12-way classification in early visual cortex (Figure 2B) indicates that early vision carried some signal for distinguishing between our categories; however, the similarity structure in early vision was only weakly correlated with behavioral structure compared with later visual regions. Furthermore, our searchlight analysis using a dichotomous animate versus inanimate and an animacy continuum model (Figure 7) in a multiple regression to predict neural similarity shows that neither model strongly predicted structure in early vision. Because the focus of the present article was to explore structure in later vision, we do not report a detailed analysis of structure in early vision.

Structural Analysis Using STATIS

We used STATIS (Abdi et al., 2012) to analyze the structure of the category-related neural patterns within this ROI, which encompassed all of the LO complex and extended to the posterior intraparietal sulcus. Like PCA,

STATIS computes PCs, each of which has factor scores (Figure 3A) for the categories and loadings (Figure 3B) for the voxels in each participant. The reliability of factor scores was assessed using a bootstrap method to derive confidence intervals around the factor score centroids—the ellipses in Figure 1C. The first PC accounts for 55% of the variance, the second PC accounts for 10%, and the third PC accounts for $\approx 7\%$ (Figure 3C). Factor scores on the first PC (Figures 1C and 3A) show a continuous gradient of animacy: primates and quadruped mammals cluster on one end and tools on the other, with fish and invertebrates closer to tools and birds closer to mammals. We did not observe the sharp contrast between animate and inanimate categories that was evident in the behavioral ratings. Rather, the activation patterns for tools are similar to those for stingrays, ladybugs, and lobsters, placing the tools near the less animate end of the animacy gradient, whereas the activation patterns for humans are similar to those for chimpanzees and cats. These results are robust with respect to whether,

instead of voxel scaling (see Experimental Methods), we used voxel selection with 300 and 2000 voxels and when we used no voxel selection or scaling (results not shown).

To examine the relationship between the dominant first dimension, which was based on responses to both animals and tools, and the animacy dimension as defined solely within animal classes, we used STATIS to define PCs based on different subsets in our data: (1) PCs on only primates and tools (Figure 4A) and (2) PCs based on only animal categories without primates and tools (Figure 4B). For each solution, we projected the remaining categorical patterns onto the corresponding PC space and found graded levels of animacy with little change compared with the first PC defined based on all categories. These results indicate that neural activity in our ROI treats both animate and inanimate objects along a single continuous dimension that reflects the level of animacy, and the animate–inanimate distinction is one aspect of this dimension.

The Functional Topography of the Animacy Continuum

To characterize the functional topography for the animacy dimension, we examined the voxel loadings for each of our PCs (Figure 3B). For PC1, positive loadings indicate more activation for highly animate categories,

that is, those categories with positive factor scores on PC 1, and negative loadings indicate more activation for less animate or inanimate categories—those that have negative factor scores on PC1. The functional topography in bilateral VT resembles previously reported coarse-scale animate versus inanimate activation topographies in which the lateral fusiform gyrus responds more to animate objects and the medial fusiform responds more to inanimate objects (Connolly, Guntupalli, et al., 2012; Mahon et al., 2009; Chao et al., 1999). Bilateral superior LO cortex showed more activation for highly animate objects, whereas bilateral inferior and posterior LO cortices showed more activation to less animate or inanimate objects. This topography is very similar to the map of the first PC based on primates and tools (Figure 4A), as well as the map defined using just eight animal categories without primates and tools (Figure 4B). Although the voxel loading maps convey the relative contributions of activation gradients to the overall representational space, the statistical maps show that even lower-order PCs have some topological consistency across participants (data not shown beyond PC2). Lower-order PCs, however, should be interpreted with caution. PCA determines a set of orthogonal PCs, but it does not follow that the underlying neural generators—the true latent variables that shape the representational space—are orthogonal or even independent of each other. We do not know

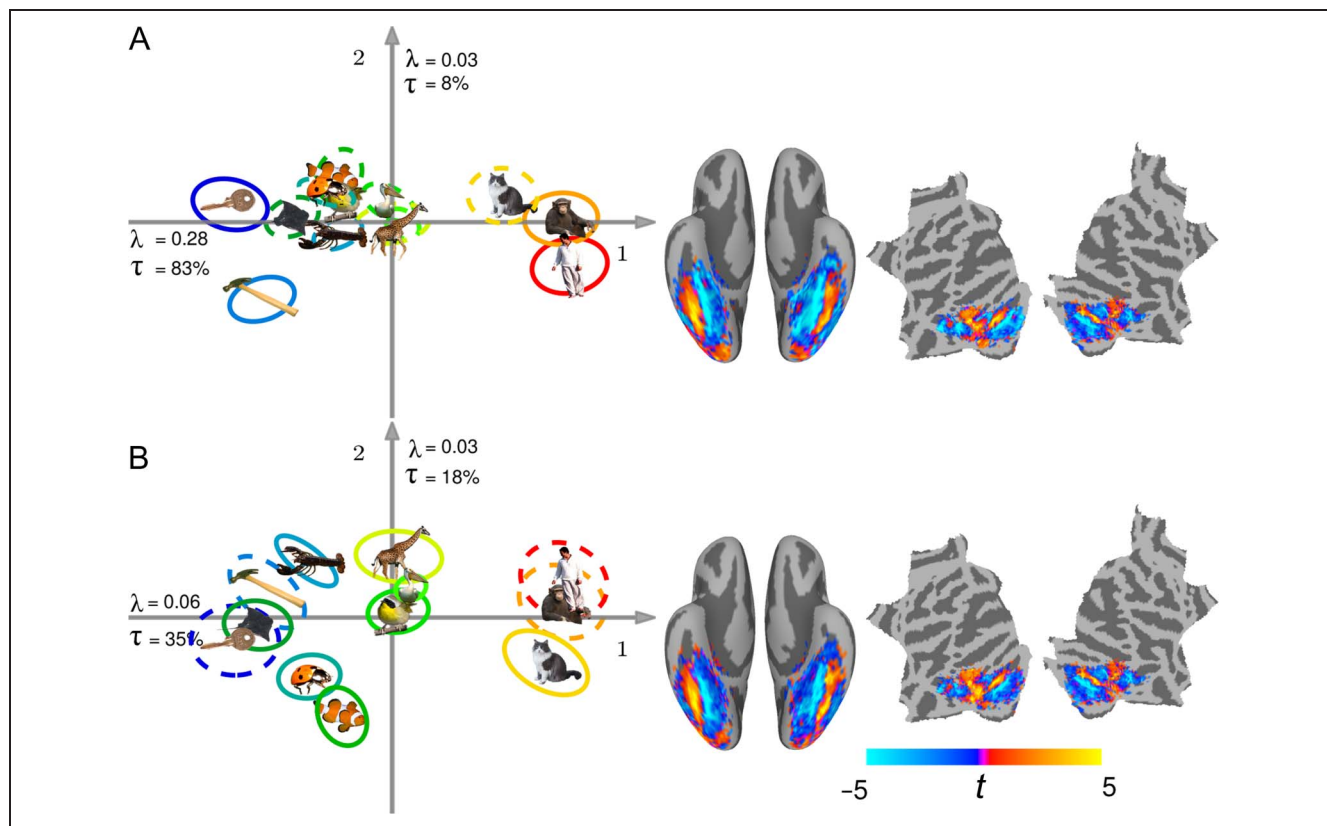


Figure 4. (A) MDS showing the factor scores for first two PCs from STATIS calculated using just primates and tools. The remaining stimulus patterns were then projected into the PC space as supplemental observations (dashed ellipses). (B) MDS from STATIS based on just the nonprimate animals with primates and tools projected into the space as supplemental observations.

the true number of underlying neural generators—and even the first PC may capture variation that is generated by two or more nonindependent signals. The results of the PCA indicate that the representational space is high-dimensional with numerous neural generators that are multiplexed in the cortex of the ventral vision pathway, each with a topography that is, at least partially, shared across participants.

The Animacy Continuum Captures Most of the between Biological Class Differentiation

To further characterize the contributions of the PCs to the representation of our stimulus categories, we performed a series of analyses to compare classification performance as a function of the inclusion and exclusion of each PC. This was done by running classification analysis using just the factor scores of the PCs. Our stimuli included one superordinate class of inanimate stimuli (tools) and five different superordinate biological classes with different levels of animacy. If the first PC captures animacy information, then this PC by itself should be a sufficient basis for classifying between-class pairings among animal classes and tools. However, within-class distinctions, such as warblers versus pelicans, may not be captured well by the animacy dimension, and classification of these pairings may depend on other dimensions. To test this hypothesis, we performed within-subject SVM classification for all 66 unique pairs of categories, and we present the results for within- and between-superordinate category pairs separately (Figure 5). Classification using all PCs resulted in mean classification of 0.87 ± 0.02 (mean \pm SEM) for between-class pairs and 0.77 ± 0.03 for within-class pairs. Classification based on just the first PC resulted in mean classification accuracy of 0.79 ± 0.01 for between-class pairs and 0.61 ± 0.02 for within-class pairs. Classification of between-class pairs based on all other PCs (0.78 ± 0.02) was equivalent to classification based on the first PC alone, $t(10) = 0.99$, *ns*, whereas classification of within-class pairs based on all other PCs (0.78 ± 0.03) was significantly better than classification based on the first PC, $t(10) = 5.92$, $p < .001$, and equivalent to classification using all PCs, $t(10) = 1.27$, *ns*. These results show that the first PC contains mostly information about between-class distinctions—whereas information about finer distinctions within classes is more distributed across PCs. For example, the distinction between pelicans and warblers was not evident on the first PC (Figures 1C and 3A) but was seen in the second, third, fifth, and sixth PCs (Figure 3A). Similarly, the distinction between humans and chimpanzees was evident in only the fourth, sixth, tenth, and eleventh PCs (Figure 3A). In addition, some between-class distinctions are carried on lower-order PCs: The fourth PC provides between-class distinction for birds versus invertebrates, although not carrying much distinguishing information for within-class pairs for either birds or invertebrates.




Figure 5. The contributions of PC1 versus all other PCs for within- and between-class pairwise distinctions measured using SVM pattern classifiers. Error bars show 95% confidence intervals.

Information Content on the High and Low Ends of the Animacy Continuum Topography

To examine whether the discriminating information about “high animacy” stimuli was localized in cortex with positive values on the animacy continuum dimension and information about “low animacy” stimuli is localized to cortex with negative values on the animacy continuum, we subsampled our ROI into the voxels that loaded the most positively and the most negatively on the first PC, keeping 300 voxels for each partition. We then tested pairwise classification for high animacy categories (humans vs. chimpanzees) and inanimate categories (hammers vs. keys) in each division. To evaluate the outcome of the classification accuracies, we used a two-way ANOVA with voxel loading as one independent variable with two levels (positive vs. negative loadings), animacy of the stimuli as the second independent variable also with two levels (primates vs. tools), and with participants as the random variable. Main effects were not observed for voxel loading (mean proportions correct: positive = 0.70 ± 0.15 ; negative = 0.66 ± 0.16 ; $F(1/10) = 0.99$, *ns*) nor for animacy (primates = 0.69 ± 0.16 ; tools = 0.67 ± 0.16 ; $F(1/10) = 0.24$, *ns*), and there was no interaction (pos-primates = 0.72 ± 0.15 ; pos-tools = 0.69 ± 0.17 ; neg-primates = 0.66 ± 0.17 ; neg-tools = 0.66 ± 0.17 ; $F(1/10) = 0.14$, *ns*). Classification accuracies for 12-way SVM classification were equivalent across all categories with one exception: Accuracies for cats were higher in positive voxels than in negative voxels, but these results did not represent any general trend (Figure 6).

These results indicate that information that discriminates among high animacy categories is not found predominantly in the cortex that responds most strongly

to those stimuli and, similarly, information that discriminates low animacy and inanimate stimuli is not found predominantly in the cortex that responds most strongly to those stimuli. Rather the information about high-animacy, low-animacy, and inanimate categories is distributed across the ventral vision pathway that contains the animacy continuum.

No Evidence for a Sharp Animate–Inanimate Distinction in Cortex

The dichotomous distinction between animate and inanimate objects is psychologically salient and prominent in the similarity space defined by behavioral judgments (Figure 1B). By contrast, the similarity structure defined by neural signal in the ventral vision pathway is dominated by the animacy continuum with overlap between the lowest animacy species and inanimate tools. To test whether the dichotomous distinction between animate and inanimate objects may be represented in some other brain region, we built two model-based DSMs: (1) a dichotomous model in which minimal distance is posited within living and nonliving categories and maximum distance between these two categories and (2) a continuum model that posits minimal dissimilarity between objects at the same animacy level (primates, mammals, birds, etc.) and increasing dissimilarity further along the animacy continuum (Figure 7A). We then used these models as predictors of




Figure 6. Proportion of times each stimulus was classified correctly in 12-way SVM classification using all voxels, only voxels that loaded positively on the first PC, and only voxels that loaded negatively on the first PC. The error bars show 95% confidence intervals based on the SEM across participants.




Figure 7. Multiple regression searchlight using model DSMs shows no evidence for a sharp animate–inanimate distinction. Model DSMs (top) for the animacy continuum model and the dichotomous living versus nonliving model (stimuli arranged as in Figure 1) were used as regressors in multiple regression to predict neural similarity at each searchlight. The map shows the difference in beta-weights for the continuum model minus the dichotomous model. The map is thresholded using the R^2 fit for the regression model with $R^2 > .2$.

neural DSMs in a multivariate linear least-squares regression searchlight analysis.

The results of this analysis revealed a good fit between the model and neural activity throughout bilateral LO cortex. The model accounted for 51% of the variance in searchlight DSM centered on the peak voxel in the middle of the right fusiform. Comparison of the β coefficients for the two model DSMs show that the fit of the model throughout LO cortex is attributable to the high correlation between the continuum model and neural representational structure in this region. β values for the dichotomous model were only higher in early visual cortex. No region outside of early visual cortex showed any evidence for representing the dichotomous distinction between animate and inanimate objects. Moreover, the dichotomous model was not a good fit for the neural structure in early vision, as it accounted for only a small percentage of the variance. Furthermore, structural analysis using STATIS in early visual cortex revealed overlapping representations for living and nonliving objects (results not shown). Thus, the neural basis for the psychologically predominant, dichotomous animate–inanimate distinction (first dimension in Figure 1B) was not revealed in our experimental data.

DISCUSSION

The results of this study suggest that the view of ventral vision pathway as encoding a dichotomous living–nonliving distinction is incorrect; furthermore, a graded dimension of animacy describes a major part of the representational space in the human ventral vision pathway. Representations of inanimate objects share the same location on

this dimension as animals with “low animacy,” such as invertebrates and fish. The animacy continuum is carried by the same topography that has been associated with animate–inanimate domain specificity. We propose that the representation of the animate–inanimate distinction in the ventral vision pathway is better understood as one facet of an animacy continuum. Our previous results showed that a continuous dimension from “least animate” to “most animate” captured most of the variance in neural representation among biological classes (Connolly, Guntupalli, et al., 2012). The current study shows that bona fide inanimate categories are represented along this continuum in a manner similar to animals with low animacy, a pattern that blurs the distinction between the animate and inanimate domains in neural representational space.

The similarity structure for neural responses in the ventral vision pathway differed significantly from the similarity structure based on behavioral ratings. In particular, the dominant dimension in the neural representational space is for the animacy continuum, and the representation of tools is embedded in that continuum on the low animacy end. By contrast, the primary distinction in the behavioral ratings was between all living animals and the inanimate tools (first dimension; Figure 1A), with equivalent values for animals, except humans, suggesting invariance across stimuli depicting different species. The second dimension reflected a continuum of the living animals with an ordering that closely matched the animacy continuum in the neural representational space. The distinction between humans and other animals, especially chimpanzees, was not evident on the animacy continuum but was only seen on lower-order PCs. These discrepancies between behavioral and neural representational spaces suggest that the categorical distinction between all living, animate entities and nonliving entities in the ventral vision pathway is actually just one facet of a more general organization principle.

The neural representation of an animacy continuum may be shared with other primates. The between-category correlations for population responses in monkey IT cortex (Kiani et al., 2007) show that insects and butterflies are more highly correlated with fish and reptiles than with the categories of mammalian species, with birds in an intermediate position. The data in Kiani et al. (2007), however, appear to show a larger distinction between their inanimate and low animacy categories (their Figure 10) than we find in human VT cortex.

The contrast between responses to primates and tools replicates the topography for the distinction between animate and inanimate stimuli seen in many previous studies. We show, however, that the animacy continuum, which involves no nonliving stimuli, has essentially the same topography (Figure 4B) and that the responses to “low animacy” animals are very similar to the responses to tools (Figure 1C). This topography, therefore, does not reflect the binary animate–inanimate distinction with invariance across living animals but, rather, reflects a con-

tinuum in which the inanimate stimuli are embedded among low animacy animals.

Information about high animacy stimuli is not limited to areas that have positive values on this dimension, and likewise, information about low animacy stimuli is not limited to areas with negative values on the animacy continuum. Rather, in general, the distinctions among high animacy stimuli are as strong in the cortex with weak responses to these stimuli as in the cortex with strong responses and likewise for distinctions among low animacy animals and inanimates—consistent with previous reports by us and others (Huth, Nishimoto, Vu, & Gallant, 2012; Brants, Baeck, Wagemans, & Op de Beeck, 2011; Haxby et al., 2001, 2011; Op de Beeck, Brants, Baeck, & Wagemans, 2010). Thus, both high and low animacy categories are represented in patterns of brain activity in both lateral and medial VT cortex.

The topographies of PCs (Figure 3B) are pattern basis functions that capture the differential patterns of response that carry stimulus distinctions. For example, although the animacy continuum does not capture the distinction between responses to warblers and pelicans, adding in the second and third PCs does capture the differential topographies for these responses. In general, the first PC—that is, the animacy continuum—captures coarser distinctions among animal classes, whereas the information about finer within-class distinctions is more distributed across PCs. The patterns that capture a distinction between two categories are not restricted to the cortex that has positive values for those categories on the dominant animacy continuum topography, indicating that information about animal species and objects is broadly distributed.

The blurring between tools and low animacy animals is further corroborated by studies of clearly inanimate objects that behave in ways that resemble the behavior of animate entities. Animations of simple geometric shapes that move in ways to suggest social interactions consistently evoke strong activity in the lateral parts of VT cortex (Gobbini, Koralek, Bryan, Montgomery, & Haxby, 2007; Beauchamp, Lee, Haxby, & Martin, 2003; Schultz et al., 2003; Castelli, Frith, Happé, & Frith, 2002). Robots that have features of human faces but are clearly mechanical and thus are never mistaken for living faces evoke strong activity in lateral fusiform cortex when they make recognizable facial expressions (Gobbini et al., 2011). Industrial robots that have minimal structural similarity to animals evoke stronger activity in the right fusiform cortex when they perform goal-directed actions, as compared with non-goal-directed actions (Shultz & McCarthy, 2012). Gobbini et al. (2011) proposed that the occipital face area, fusiform face area, and face-responsive posterior STS “participate in the representation of agentic forms and actions but do not distinguish between animate and inanimate agents” (p. 1915). Together these studies show that the responses to inanimate agents in the ventral vision pathway are on the higher animacy side of the animacy continuum. The common denominator seems to be

the performance or potential to perform self-initiated, complex, goal-directed actions. The word *anima* means soul in Latin. Clearly, robots do not have souls, but they can perform actions that are the result of complex internal states that must be inferred. Thus, the animacy continuum appears to be related to the perception of things, be they animate or mechanical, that involves inference of internal states of varying levels of complexity. Perception of these internal states affords prediction of behavior and can condition how we interact with these things. Representation of level of animacy, as broadly defined here, thus is a necessary precondition for recruitment of the appropriate cognitive resources for understanding and interacting effectively with animals and objects.

Where Does the Animacy Continuum Come from?

Previous theories of the animate–inanimate distinction were premised on observations of a dichotomous distinction. The strong form of the domain specificity hypothesis proposed that evolution provided for discrete modules for processing conspecifics, animals, and tools. Our observations that the “least animate” animals evoke activity similar to tools and that the “most animate” animals evoke activity similar to conspecifics challenge the notion of discrete modules. However, the idea that evolution played a role in determining neural organization of categories is consistent with the animacy continuum hypothesis: It can be argued that the abstract property of animacy has ecological and evolutionary significance. Alternatively, the strong form of the SFH proposed that visual features were more important for identifying animals and that functional motor features were more important for tools. Our observation that the “least animate” animals share representational space with tools challenges this strong version of SFH. There are no obvious functional motor properties relevant to spiders, ladybugs, and stingrays that can explain their similarity to tools. It is also not obvious why visual properties are less important for recognizing a ladybug compared with a chimpanzee.

We interpret the animacy continuum as corresponding to an abstract psychological dimension of perceived animacy that is measurable in human behavioral experiments, such as that reported in Figure 1B where the animacy continuum is reflected on the second dimension of the MDS space. Evidence of animacy as a graded psychological dimension is found in the writings of Aristotle, in linguistics, and in folk taxonomies. The notion of graded animacy has ecological validity: It is a fact about the world that animals vary with respect to intelligence, agency, and the degree to which animals share characteristics with the animate prototype—humans. Leading theories of brain evolution propose that the large size and computational power of the human brain arose from the pressures of living in large complex social groups (Dunbar, 1998). The animacy continuum hypothesis is consistent with this view in that understanding the actions and intentions of other

sentient beings is a central component. By extension, the similarity of brain activity for perceiving animals and perceiving humans will be proportional to the degree to which the animals display agentic properties.

In conclusion, we address a likely objection to our interpretation of the results: that animacy makes a poor candidate for an organizing principle of neural representation because it is not a simple, primary, or basic property. In the tradition of David Hume and John Locke, the modern empiricist notion of “nothing in the brain not first in the senses” pervades the cognitive sciences (Machery, 2006). Neoempiricists will likely judge animacy as too abstract to play a central role in representation. We disagree. As a perceptual act, perceiving the property of animacy is necessarily tied to sensory input, but we find that accounts for animacy in terms of sensory properties are convoluted and not parsimonious. No simple sensory feature, such as color, texture, or motion energy, has a clear relationship to the animacy continuum. More complex features—the presence of faces and bodies; gaze, expression, and other intentional movements; complex vocalizations—are related to the continuum, but no single feature is invariably present or necessary by itself to perceive level of animacy or agency. The level of animacy can be perceived based on action without form (Gobbini et al., 2007; Castelli et al., 2002) and on form without action (still images), and both are associated with the lateral-to-medial gradient for animacy in the ventral pathway. The congenitally blind can perceive agents and show a normal lateral-to-medial gradient for animacy in the ventral pathway despite the lack of visual sensory input or experience (Mahon et al., 2009). The common denominator, therefore, is more the abstract property of animacy or agency than a simple correlation with sensory features.

Acknowledgments

Funding was provided by National Institutes of Mental Health grants F32MH085433-01A1 (Connolly) and 5R01MH075706 (Haxby) and National Science Foundation grant NSF1129764 (Haxby). We also thank those who have helped in providing expert advice and other assistance, especially M. Ida Gobbini, Michael Hanke, Scott Pauls, and the members of the Haxby Lab.

Reprint requests should be sent to Andrew C. Connolly, Dartmouth College, 6207 Moore Hall, Hanover, NH 03755, or via e-mail: andrew.c.connolly@dartmouth.edu.

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*, 89–95.
- Abdi, H., Williams, L. J., Valentin, D., & Bannani-Dosse, M. (2012). Stas and distas: Optimum multitable principal component analysis and three way metric multidimensional scaling. *Wiley Interdisciplinary Reviews: Computational Statistics*, *4*, 124–167.

- Aguirre, G. K. (2007). Continuous carry-over designs for fMRI. *Neuroimage*, *35*, 1480–1494.
- Allport, D. (1985). Distributed memory, modular subsystems and dysphasia. In S. K. Newman & R. Epstein (Eds.), *Current Perspectives in Dysphasia* (pp. 32–60). Edinburgh: Churchill Livingstone.
- Aristotle, A. (1986). *De Anima: On the soul*. London: Penguin Classics.
- Atran, S. (1998). Folk biology and the anthropology of science: Cognitive universals and cultural particulars. *Behavioral and Brain Sciences*, *21*, 547–569.
- Beauchamp, M. S., Lee, K. E., Haxby, J. V., & Martin, A. (2003). fMRI responses to video and point-light displays of moving humans and manipulable objects. *Journal of Cognitive Neuroscience*, *15*, 991–1001.
- Brants, M., Baeck, A., Wagemans, J., & Op de Beeck, H. P. (2011). Multiple scales of organization for object selectivity in ventral visual cortex. *Neuroimage*, *56*, 1372–1381.
- Caramazza, A., & Mahon, B. Z. (2003). The organization of conceptual knowledge: The evidence from category-specific semantic deficits. *Trends in Cognitive Sciences*, *7*, 354–361.
- Caramazza, A., & Shelton, J. R. (1998). Domain-specific knowledge systems in the brain: The animate–inanimate distinction. *Journal of Cognitive Neuroscience*, *10*, 1–34.
- Castelli, F., Frith, C., Happé, F., & Frith, U. (2002). Autism, Asperger syndrome and brain mechanisms for the attribution of mental states to animated shapes. *Brain*, *125*, 1839–1849.
- Chang, C.-C., & Lin, C.-J. (2011). Libsvm: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology (TIST)*, *2*, 27.
- Chao, L. L., Haxby, J. V., & Martin, A. (1999). Attribute-based neural substrates in temporal cortex for perceiving and knowing about objects. *Nature Neuroscience*, *2*, 913–919.
- Connolly, A. C., Gobbini, M. I., & Haxby, J. V. (2012). Three virtues of similarity-based multivariate pattern analysis: An example from the human object vision pathway. *Visual population codes: Toward a common multivariate framework for cell recording and functional imaging*. Cambridge, MA: MIT Press.
- Connolly, A. C., Guntupalli, J. S., Gors, J., Hanke, M., Halchenko, Y. O., Wu, Y.-C., et al. (2012). The representation of biological classes in the human brain. *The Journal of Neuroscience*, *32*, 2608–2618.
- Cortes, C., & Vapnik, V. (1995). Support-vector networks. *Machine Learning*, *20*, 273–297.
- Cox, R. W. (1996). Afni: Software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical Research*, *29*, 162–173.
- Dunbar, R. I. (1998). The social brain hypothesis. *Brain*, *9*, 178–190.
- Eger, E., Ashburner, J., Haynes, J.-D., Dolan, R. J., & Rees, G. (2008). fMRI activity patterns in human loc carry information about object exemplars within category. *Journal of Cognitive Neuroscience*, *20*, 356–370.
- Finney, D., & Outhwaite, A. D. (1956). Serially balanced sequences in bioassay. *Proceedings of the Royal Society of London, Series B, Biological Sciences*, *145*, 493–507.
- Fischl, B. (2012). Freesurfer. *Neuroimage*, *62*, 774–781.
- Gobbini, M. I., Gentili, C., Ricciardi, E., Bellucci, C., Salvini, P., Laschi, C., et al. (2011). Distinct neural systems involved in agency and animacy detection. *Journal of Cognitive Neuroscience*, *23*, 1911–1920.
- Gobbini, M. I., Koralek, A. C., Bryan, R. E., Montgomery, K. J., & Haxby, J. V. (2007). Two takes on the social brain: A comparison of theory of mind tasks. *Journal of Cognitive Neuroscience*, *19*, 1803–1814.
- Hanke, M., Halchenko, Y. O., Sederberg, P. B., Hanson, S. J., Haxby, J. V., & Pollmann, S. (2009). Pymvpa: A python toolbox for multivariate pattern analysis of fMRI data. *Neuroinformatics*, *7*, 37–53.
- Hanson, S. J., Matsuka, T., & Haxby, J. V. (2004). Combinatorial codes in ventral temporal lobe for object recognition: Haxby (2001) revisited: Is there a face area? *Neuroimage*, *23*, 156–166.
- Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*, *293*, 2425–2430.
- Haxby, J. V., Guntupalli, J. S., Connolly, A. C., Halchenko, Y. O., Conroy, B. R., Gobbini, M. I., et al. (2011). A common, high-dimensional model of the representational space in human ventral temporal cortex. *Neuron*, *72*, 404–416.
- Huth, A. G., Nishimoto, S., Vu, A. T., & Gallant, J. L. (2012). A continuous semantic space describes the representation of thousands of object and action categories across the human brain. *Neuron*, *76*, 1210–1224.
- Kanwisher, N., & Dilks, D. (2013). The functional organization of the ventral visual pathway in humans. In L. Chalupa & J. Werner (Eds.), *The New Visual Neurosciences* (pp. 733–748). Cambridge, MA: MIT Press.
- Kherif, F., Poline, J.-B., Mériaux, S., Benali, H., Flandin, G., & Brett, M. (2003). Group analysis in functional neuroimaging: Selecting subjects using similarity measures. *Neuroimage*, *20*, 2197–2208.
- Kiani, R., Esteky, H., Mirpour, K., & Tanaka, K. (2007). Object category structure in response patterns of neuronal population in monkey inferior temporal cortex. *Journal of Neurophysiology*, *97*, 4296–4309.
- Konkle, T., & Oliva, A. (2012). A real-world size organization of object responses in occipitotemporal cortex. *Neuron*, *74*, 1114–1124.
- Kriegeskorte, N., Goebel, R., & Bandettini, P. (2006). Information-based functional brain mapping. *Proceedings of the National Academy of Sciences, U.S.A.*, *103*, 3863–3868.
- Kriegeskorte, N., Mur, M., & Bandettini, P. (2008). Representational similarity analysis connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, *2*, Article 4.
- Kriegeskorte, N., Mur, M., Ruff, D. A., Kiani, R., Bodurka, J., Esteky, H., et al. (2008). Matching categorical object representations in inferior temporal cortex of man and monkey. *Neuron*, *60*, 1126–1141.
- Kriegeskorte, N., Simmons, W. K., Bellgowan, P. S., & Baker, C. I. (2009). Circular analysis in systems neuroscience: The dangers of double dipping. *Nature Neuroscience*, *12*, 535–540.
- Levy, I., Hasson, U., Avidan, G., Hendler, T., & Malach, R. (2001). Center-periphery organization of human object areas. *Nature Neuroscience*, *4*, 533–539.
- Machery, E. (2006). Two dogmas of neo-empiricism. *Philosophy Compass*, *1*, 398–412.
- Mahon, B. Z., Anzellotti, S., Schwarzbach, J., Zampini, M., & Caramazza, A. (2009). Category specific organization in the human brain does not require visual experience. *Neuron*, *63*, 397–405.
- Mahon, B. Z., & Caramazza, A. (2009). Concepts & categories: A cognitive neuropsychological perspective. *Annual Review of Psychology*, *60*, 2751.
- Martin, A. (2007). The representation of object concepts in the brain. *Annual Review of Psychology*, *58*, 25–45.
- Oosterhof, N. N., Wiestler, T., Downing, P. E., & Diedrichsen, J. A. (2011). A comparison of volume-based and surface-based multi-voxel pattern analysis. *Neuroimage*, *56*, 593–600.
- Oosterhof, N. N., Wiggett, A. J., Diedrichsen, J., Tipper, S. P., & Downing, P. E. (2010). Surface-based information mapping reveals crossmodal vision-action representations in human

- parietal and occipitotemporal cortex. *Journal of Neurophysiology*, *104*, 1077–1089.
- Op de Beeck, H. P., Brants, M., Baeck, A., & Wagemans, J. (2010). Distributed subordinate specificity for bodies, faces, and buildings in human ventral visual cortex. *Neuroimage*, *49*, 3414–3425.
- O'Toole, A. J., Jiang, F., Abdi, H., & Haxby, J. V. (2005). Partially distributed representations of objects and faces in ventral temporal cortex. *Journal of Cognitive Neuroscience*, *17*, 580–590.
- O'Toole, A. J., Jiang, F., Abdi, H., Pénard, N., Dunlop, J. P., & Parent, M. A. (2007). Theoretical, statistical, and practical perspectives on pattern-based classification approaches to the analysis of functional neuroimaging data. *Journal of Cognitive Neuroscience*, *19*, 1735–1752.
- Schultz, R. T., Grelotti, D. J., Klin, A., Kleinman, J., Van der Gaag, C., Marois, R., et al. (2003). The role of the fusiform face area in social cognition: Implications for the pathobiology of autism. *Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences*, *358*, 415–427.
- Shinkareva, S. V., Ombao, H. C., Sutton, B. P., Mohanty, A., & Miller, G. A. (2006). Classification of functional brain images with a spatio-temporal dissimilarity map. *Neuroimage*, *33*, 63–71.
- Shultz, S., & McCarthy, G. (2012). Goal-directed actions activate the face-sensitive posterior superior temporal sulcus and fusiform gyrus in the absence of human-like perceptual cues. *Cerebral Cortex*, *22*, 1098–1106.
- Warrington, E. K., & Shallice, T. (1984). Category specific semantic impairments. *Brain*, *107*, 829–853.
- Young, R. W., & Morgan, W. (1987). *The Navajo language: A grammar and colloquial dictionary*. Albuquerque: University of New Mexico Press.