Published in final edited form as:

Neuropsychologia. 2015 September ; 76: 17–26. doi:10.1016/j.neuropsychologia.2015.04.009.

Predicting brain activation patterns associated with individual lexical concepts based on five sensory-motor attributes

Leonardo Fernandino¹, Colin J. Humphries¹, Mark S. Seidenberg², William L. Gross³, Lisa L. Conant¹, and Jeffrey R. Binder¹

¹Department of Neurology, Medical College of Wisconsin, Milwaukee, WI, USA

²Department of Psychology, University of Wisconsin, Madison, WI, USA

³Department of Anesthesiology, Medical College of Wisconsin, Milwaukee, WI, USA

Abstract

While major advances have been made in uncovering the neural processes underlying perceptual representations, our grasp of how the brain gives rise to conceptual knowledge remains relatively poor. Recent work has provided strong evidence that concepts rely, at least in part, on the same sensory and motor neural systems through which they were acquired, but it is still unclear whether the neural code for concept representation uses information about sensory-motor features to discriminate between concepts. In the present study, we investigate this question by asking whether an encoding model based on five semantic attributes directly related to sensory-motor experience – sound, color, visual motion, shape, and manipulation – can successfully predict patterns of brain activation elicited by individual lexical concepts. We collected ratings on the relevance of these five attributes to the meaning of 820 words, and used these ratings as predictors in a multiple regression model of the fMRI signal associated with the words in a separate group of participants. The five resulting activation maps were then combined by linear summation to predict the distributed activation pattern elicited by a novel set of 80 test words. The encoding model predicted the activation patterns elicited by the test words significantly better than chance. As expected, prediction was successful for concrete but not for abstract concepts. Comparisons between encoding models based on different combinations of attributes indicate that all five attributes contribute to the representation of concrete concepts. Consistent with embodied theories of semantics, these results show, for the first time, that the distributed activation pattern associated with a concept combines information about different sensory-motor attributes according to their respective relevance. Future research should investigate how additional features of phenomenal experience contribute to the neural representation of conceptual knowledge.

Correspondence: Leonardo Fernandino, Department of Neurology, Medical College of Wisconsin, 8701 Watertown Plank Road, MEB 4671, Milwaukee, WI, 53226, USA. lfernandino@mcw.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Keywords

lexical semantics; fMRI; semantic memory; embodiment; concepts; multimodal processing

1. Introduction

Brain-based theories of concept representation generally begin with the premise that concepts are learned through a process of generalization or abstraction from individual experiences. Experiences are comprised of identifiable sensory, motor, spatial, temporal, affective, and cognitive components, and the relative contribution of each of these experiential components to concept formation depends on the particular concept being acquired. For example, the concept of “thunder” is learned almost entirely through auditory experiences, whereas the concept of “octopus” has no connection with audition. Concepts for manipulable objects such as “pencil” and “spoon” are learned in part through motor actions performed with those objects, whereas concepts such as “moon” and “giraffe” have very little (if any) basis in motor experience. However, most experiences are inherently multimodal in the sense that they involve simultaneous, covarying information from multiple sensory-motor modalities. Using a pair of scissors, for example, typically involves concurrent sensory-motor experiences related to action planning, proprioceptive and tactile feedback, and the perception of characteristic shape, motion, and sound features. Thus, the concept of “scissors” is likely to include representations of all of these attributes.

Sensory information is processed according to perceptual attributes that can be strictly unimodal (e.g., color) or result from the combination of two or more sensory modalities (e.g., perception of object shape often combines visual and tactile information). Embodied or “simulation” theories of concept representation propose that the information abstracted from experience during concept learning is represented to some degree in the same modality-specific and multimodal neural systems through which the learning occurred, and that concept retrieval involves some degree of activation of this sensory-motor information (Damasio, 1989; Barsalou, 2008; Glenberg and Gallese, 2012; Hoenig et al., 2011; Kiefer et al., 2007;).

Although embodiment theories have garnered extensive empirical support (for reviews, see Fischer and Zwaan, 2008; Meteyard and Vigliocco, 2008; Binder and Desai, 2011; Kiefer and Pulvermüller, 2012; Meteyard et al., 2012), most studies have aimed at demonstrating the existence of modality-specific representational systems (e.g., Warrington and Shallice, 1984; McCarthy and Warrington, 1988; Farah and McClelland, 1991), or at implicating a particular modality-specific cortical area in the representation of concrete concepts (e.g., Martin et al., 1996; Chao et al., 1999; Hauk et al., 2004; Aziz-Zadeh et al., 2006; Kiefer et al., 2008; Hsu et al., 2012). One question that remains unanswered is whether the distributed neural representation of a concept combines information originating from various sensory-motor attributes, as predicted by embodiment theories. In the present study, we explore this issue by adopting a novel approach to brain activation data relevant to this question. Instead of searching for localized activations associated with particular semantic attributes, we combine information about five different attributes to predict the activation patterns

associated with individual concepts. In other words, we ask whether information about five sensory-motor attributes of lexical concepts is sufficient to predict patterns of brain activation elicited by isolated words. The five attributes selected for study – color, shape, visual motion, sound, and manipulation – are each associated with well-studied brain networks (for a review, see Fernandino et al., 2015). In a prior study, we obtained salience ratings on each of these attributes for a set of 900 English nouns, and showed that these attribute ratings parametrically modulate word-related brain activity in distinct cortical networks containing both unimodal and multimodal nodes (Fernandino et al., 2015). In the present study, attribute-specific activation maps, derived from a set of 820 words (the "modelling set"), were combined into an "encoding model" (Haxby et al., 2014; Naselaris et al., 2011), which was used to predict word-specific, whole-brain activation patterns for the 80 remaining words (the "test set"). The encoding model consisted of a linear combination of the five attribute maps from the training set, weighted by the test word's attribute rating values. The predicted activation patterns were then compared to the observed activation patterns for each test word. Successful prediction of the test word patterns would provide direct evidence that the overall, distributed neural representation of a concrete concept encodes information about the relative relevance of specific aspects of sensory-motor experience, as rated by an independent group of participants.

Another novel aspect of the present study is the focus on group-averaged activation patterns. Studies using multivoxel pattern analysis (MVPA) typically evaluate the performance of the encoding model (or classifier, in the case of pattern classification studies) separately for each participant, to account for individual differences in brain morphology and function-structure mapping, and perform group-level statistical tests on the resulting accuracy scores (e.g., Mitchell et al., 2008; Huth et al., 2012; Haxby et al., 2001). Here, we assess whether similarities in the neural code for concrete concepts across individuals would allow a group-level encoding model to successfully predict the group-averaged activation maps corresponding to different concepts. If prediction is successful, the encoding model can be seen as a first approximation to a subject-independent neural code for concept representation, and the predicted activation maps could be considered as rough neural signatures of the respective concepts.

2. Methods

2.1. Attribute ratings

We focused on five semantic attributes related to sensory-motor experience: Sound, Color, Manipulation, Visual Motion, and Shape. Ratings for these attributes were available for a set of 900 words (see Fernandino et al., 2015, for details). The ratings reflect the salience of each attribute to the meaning of the word on a 7-point Likert scale ranging from "not at all important" to "very important". The data set included approximately 30 ratings of each attribute for each word. Figure 1 shows mean ratings for six example words, and Table 1 lists the correlations between the five attribute ratings across all words.

2.2. fMRI data

All analyses were conducted on the data from Fernandino et al. (2015). Data collection procedures from that study are summarized below.

2.2.1. Participants—Forty-four healthy, native speakers of English (16 females; mean age 28.2, range 19 to 49) with no history of neurological or psychiatric disorders, participated in the study. All were right-handed according to the Edinburgh Handedness Inventory (Oldfield, 1971). Participants were compensated for their participation and gave informed consent in conformity with the protocol approved by the Medical College of Wisconsin Institutional Review Board.

2.2.2. Stimuli—The stimuli consisted of the 900 nouns for which attribute ratings were collected (see section 2.1. above) and 300 pseudowords. All words were relatively familiar (mean CELEX frequency = 37.4 per million, SD = 118.5), and between 3 and 9 letters in length, with a flat distribution across this length range (i.e., 126–129 words of each length). Six hundred of the words were relatively concrete and 300 were relatively abstract, as determined by either published imageability ratings, which at the time were available for 748 words (Wilson 1988; Bird et al. 2001; Clark and Paivio 2004), or consensus judgment of the authors. The pseudowords were generated by a computer program (Medler and Binder, 2005) using constrained trigram statistics (i.e., third-order approximation to English), followed by exclusion of pseudohomophones. Pseudowords were matched to the words on length, orthographic neighborhood density, and bigram and trigram metrics (Table 2).

2.2.3. Task procedure—The stimuli were back-projected in white Courier font on a black background screen that was viewed by the participant through a mirror attached to the head coil. Stimuli subtended an average horizontal visual angle of approximately 2.5°. In addition to the 1200 task trials (900 words, 300 pseudowords), 600 passive fixation events (“+”) were included to act as a baseline and provide jittering for the deconvolution analysis, resulting in a total of 1960 stimulus events distributed across 10 runs. On task trials, the stimulus string was presented for 1000 ms followed by a fixation cross for 1000 ms; on fixation events, the fixation cross appeared for 2000 ms. Thus, each stimulus was followed by a varied fixation interval of 1 s, 3 s, 5 s, etc. The probability distribution of these fixation intervals was an exponentially decaying function with a median of 1 s and maximum of 13 s.

Participants were required to decide whether the stimulus (word or pseudoword) referred to something that can be experienced through the senses (i.e., a “concreteness decision”), and to respond as quickly as possible by pressing one of two response keys. They were informed that some of the items would be pseudowords, and that the answer should be “no” in these cases. All participants responded with their right hand using the index and middle fingers. Eprime 1.0 (Psychology Software Tools) was used for stimulus presentation and response registration.

2.2.4. fMRI acquisition and preprocessing—Gradient-echo EPI images were collected in 10 runs of 196 volumes each. Twenty-three participants were scanned on a GE

1.5T Signa MRI scanner (TR = 2000 ms, TE = 40 ms, 21 axial slices, $3.75 \times 3.75 \times 6.5$ mm voxels), and the other 21 were scanned on a GE 3T Excite MRI scanner (TR = 2000 ms, TE = 25 ms, 40 axial slices, $3 \times 3 \times 3$ mm voxels). T1-weighted anatomical images were obtained using a 3D spoiled gradient-echo sequence with voxel dimensions of $1 \times 1 \times 1$ mm.

Preprocessing and statistical analysis of fMRI data were done with the AFNI software package (Cox 1996). EPI volumes were corrected for slice acquisition time and for head motion. Functional volumes were then aligned to the T1-weighted anatomical volume and transformed into Talairach standard space (Talairach and Tournoux 1988), resampled at 3 mm isotropic voxels, and smoothed with a 6 mm FWHM Gaussian kernel. Each voxel time series was then rescaled to percent of mean signal level, so that subsequent regression parameter estimates reflect percent signal change.

2.3. Prediction analyses

For the prediction analyses, the 900 word stimuli were split into two groups: a modelling set, consisting of 820 items, and a test set, consisting of 80 items (40 concrete and 40 abstract). The 80 test words were selected randomly with the constraint that the concrete and abstract subsets were matched in word frequency, number of letters, number of phonemes, number of syllables, orthographic and phonological neighborhood densities, and bigram frequency (Table 3).

The prediction analysis was conducted in four steps. First, for each participant, we generated activation maps for each of the five attributes of word meaning (Attribute Maps, or AMs) based on the words in the modelling set (Figure 2A). This was done by converting their attribute ratings into z-scores and including them as simultaneous predictor variables in a Generalized Least Squares (GLS) regression model. Word length, number of phonemes, number of syllables, word frequency, bigram frequency, orthographic and phonological neighborhood density, and the participant's z-transformed RT for each trial were included as nuisance regressors. Two binary regressors were also added to account for activity associated with early visual, orthographic, and phonological processing of the stimulus, as well as the subsequent motor response: one regressor coded for "word" events and the other for "pseudoword" events. Noise was modeled with linear, second-order, and third-order trends, as well as with the estimates of the motion parameters. For each attribute, a group-level AM was obtained by averaging the individual AMs (beta values) across participants.

Second, we computed predicted activation maps (Predicted Maps, PMs), for each of the 80 words in the test set, as linear combinations of the AMs, whereby each AM was weighted by the z-score of the word's corresponding attribute rating (Figure 2B). The PM for a given test word corresponds to the hypothetical activation pattern that would be associated with the meaning of that word if the word's meaning were completely captured by the five attribute ratings (i.e., sound, color, manipulation, visual motion, and shape).

Third, we generated activation maps for each of the 80 words in the test set, relative to a pseudoword baseline (Observed Maps, OM; Figure 2C). For each participant, a separate GLS regression was done for each of the 80 test words, with the following explanatory variables: (1) a binary regressor coding for the presentation of the selected test word; (2) a

binary regressor coding for presentation of all the non-selected words (i.e., the other 899 words in the stimulus set); (3) a binary regressor coding for presentation of the pseudowords; (4) five continuous regressors coding for each of the five semantic attribute ratings for all non-selected words; and (5) a continuous regressor coding the response time for each trial. The resulting OM for a given test word, thus, revealed the unique activation pattern elicited by that word. For each test word, a group-level OM was obtained by averaging the individual OMs (beta values) across all participants.

Finally, we assessed the accuracy of the PMs by calculating the voxel-by-voxel pairwise correlation between each PM and each of the 80 OMs (Figure 2D). For each PM, we then rank-ordered all the OMs according to correlation strength and noted the percentile rank of the OM for the corresponding target word (Figure 2E). This percentile rank, scaled to a 0 to 1 range, was assigned to the PM as its accuracy score. Thus, each PM was assigned an accuracy score corresponding to how similar it was to its respective OM relative to the other 79 OMs, with 0 corresponding to least similar and 1 corresponding to most similar. For instance, if the OM for the word "coffee" were the most highly correlated (among the 80 OMs) to the PM for the same word, that word would receive an accuracy score of $79/79 = 1$. If, instead, it were the second most highly correlated to its respective PM, its accuracy score would be $78/79 = 0.987$. We then conducted a one-tailed Wilcoxon signed rank test to verify whether the median prediction accuracy across all 80 words was significantly higher than chance (0.5).

In principle, any type of information about the stimuli that correlates with the attribute ratings could contribute to the prediction success of our encoding model. Thus, it is important to rule out other factors (e.g., word length) as possible drivers of prediction accuracy. For this purpose, the test set consisted of two matched subsets, one with 40 concrete words and the other with 40 abstract words. The two subsets were matched on all lexical attributes, except for concreteness (i.e., the extent to which they referred to sensory-motor experiences). As Table 3 shows, inter-stimulus variance for the attribute ratings was much smaller among abstract than among concrete words. In other words, abstract words contained much less sensory-motor information. If prediction accuracy were indeed driven by the sensory-motor aspects of word meaning, performance should be higher for concrete than for abstract words. Finally, we also investigated the relative contribution of each of the five attributes to prediction accuracy (for concrete words) by evaluating encoding models based on all possible combinations of these attributes.

Voxel selection is an important step in MVPA, since the inclusion of non-informative voxels can severely degrade the model's performance (Cox and Savoy, 2003; Mitchell et al., 2004). We excluded non-brain voxels and voxels with low signal-to-noise ratio (SNR) by confining the analysis to a mask where the temporal SNR was higher than an arbitrary threshold ($tSNR > 50$). The resulting mask encompassed most of the brain, except for the regions that are typically affected by signal dropout (portions of the inferior temporal and orbitofrontal cortex).

3. Results

3.1. Analyses including all five attributes

When all 80 words were analyzed together, activation patterns were predicted significantly better than chance (median = 0.63, 95% CI [0.58, 0.77], $V = 2413.5$, $p = 0.00007$), showing that information about these five aspects of sensory-motor experience is encoded in the activation pattern elicited by lexical concepts (Figure 3). Since concrete concepts contain more sensory-motor information than abstract concepts (as reflected in the larger variance of attribute ratings for the former; Table 3), we expected that prediction accuracy would be driven mainly by the concrete words in the test set. As shown in Figure 3, when analyzed separately, concrete words were predicted significantly better than chance (median = 0.67, 95% CI [0.55, 0.81], $V = 610$, $p = 0.0036$), while abstract words were not (median = 0.50, 95% CI [0.29, 0.63], $V = 360.5$, $p = 0.75$). This was further confirmed by a Wilcoxon rank sum test, which showed that prediction was significantly higher for concrete than for abstract words ($W = 1048$, $p = 0.0086$). Since concrete and abstract words were closely matched on non-semantic lexical properties of the stimuli (Table 3), this result confirms that prediction accuracy was driven by semantic information alone.

3.2. Analyses on subsets of attributes

Since the predictive power of the model was specific to concrete words, we focused on this subset in the ensuing analyses. Prediction accuracies for models based on all possible combinations of the 5 attributes are shown in Figure 4. When prediction maps were based on only one attribute, prediction accuracy across the 40 concrete words was indistinguishable from chance for all attributes, with no significant differences between them ($p = 0.84$, one-way Kruskal-Wallis rank sum test), and with a mean (across attributes) of 0.53. When prediction was based on two attributes at a time, prediction accuracies were above chance level for most attribute dyads, except for Sound & Motion and for Manipulation & Motion, but with no significant differences between attribute pairs ($p = 0.66$). Mean accuracy across pairs was 0.56. With three attributes included at a time, activation patterns were predicted better than chance for 5 out of the 10 triads, again with no significant differences between them ($p = 0.40$), and with a mean accuracy of 0.60. Finally, when predictions were based on four attributes at a time, prediction was successful for all but one combination of attributes (the one in which Motion was left out). Mean prediction accuracy was 0.65, with no significant differences between the models ($p = 0.49$). These results show that prediction accuracy increased as a function of the number of attributes included in the model (Figure 5), but revealed no significant differences in the relative contribution of individual attributes.

4. Discussion

An encoding model based on the relative relevance of five sensory-motor attributes – Sound, Color, Visual Motion, Manipulation and Shape – to conceptual content was able to predict the distributed fMRI activation pattern elicited by concrete words. Prediction accuracy for a matched set of abstract words, on the other hand, was at chance level, confirming that prediction success for the concrete words was based solely on semantic information. These results show that the neural code for concept representation includes information about

sensory-motor attributes. Finally, while prediction accuracy increased linearly with the number of sensory-motor attributes included in the model, we found no significant differences between the relative contributions of the attributes, suggesting that all of them contributed to prediction success.

These results make three main contributions to the existing literature on concept representation. First, while many previous studies focused on the important issue of whether activation in modal sensory or motor systems is modulated by sensory or motor conceptual content, they did not focus on the multimodal nature of the concept representation system. The present study offers the first evidence that the neural code for concept representation includes information about the relative relevance of specific aspects of sensory-motor experience, involving different modalities. Second, the present study establishes a novel method that can be used to assess the contribution of other attributes (sensory, motor, affective, temporal, spatial, social, etc.) to the neural representation of conceptual knowledge. Finally, our results show that the neural code for concepts is consistent enough across participants to allow successful prediction of activation patterns at the group level. This finding is important in the context of previous studies that have shown that the neural representation of a lexical concept is measurably affected by the individual's unique sensory-motor experiences during concept acquisition (Hoenig et al., 2011; Kiefer et al., 2007; Weisberg et al., 2007). A corollary of this finding is that a lexical concept is not identically represented across different brains, but reflects each participant's unique life history. The relative extent of this inter-subject variability, however, remains an open issue. Is it large enough that inter-subject averaging would lead to complete loss of information? Our results indicate that, despite individual differences, the neural activation pattern associated with a lexical concept displays a certain degree of invariance across participants, making it possible to discriminate between concepts in group-averaged data. The generalizability of our results is further strengthened by the fact that the attribute ratings were derived from a separate group of participants.

The present approach to relating neural activation patterns to word meanings shares some similarities with previous studies on this topic but also differs in important ways. Mitchell et al. (2008) introduced the approach of predicting neural activation patterns for concepts using a linear combination of basis images weighted by semantic feature values. They showed that a set of 25 semantic features related to sensory-motor experience could predict neural activation associated with concrete concepts, in individual participants, with a mean accuracy of 0.77 (chance level = 0.5). One major difference from the present study is that the semantic features of the test concepts were determined by text co-occurrence rather than explicit attribute judgments. For example, one semantic feature was the degree to which a test concept co-occurs with the word "eat" in a large text corpus. The features were thus defined by measuring co-occurrence with 25 words, primarily verbs: see, hear, listen, taste, smell, eat, touch, rub, lift, manipulate, run, push, fill, move, ride, say, fear, open, approach, near, enter, drive, wear, break, clean. Although some of these have strong relations to modality-specific neural processing systems (e.g., see, hear, listen, taste, smell, touch, manipulate, run, fear), many involve multiple sensory-motor attributes from various modalities (eat, rub, lift, fill, move, push, say, open, approach, enter) or are ambiguous as to experiential content (ride, drive, wear, break, clean). Thus the relationships between these

features and particular neural processing systems are somewhat unclear. The co-occurrence method adds to this uncertainty in two ways. First, it is debatable whether co-occurrence with a set of verbs adequately captures the targeted content of a concept, since other verbs are available in most cases that could be used to convey the same general meaning (e.g., "observe," "watch," "look," and "view" instead of "see"). Second, many of the words selected as co-occurrence targets have multiple meanings, some of which may be unrelated to the targeted neural modality. The target verb "run," for example, was likely selected to identify content related to lower limb motor actions, but usages such as "the car is running" ("car" was one of the test concepts) are unrelated to this targeted content.

The present approach attempts to remove such ambiguity by deriving the semantic representation of concepts directly from human judgments about attribute salience. Rather than inferring semantic content indirectly, the rating method asks human observers to make explicit links between concepts and experimenter-defined attributes. The limitation of this approach is that human judgments are inevitably subjective and open to error, although averaging responses from a sample of observers after removal of intra-class outliers mitigates this problem to some degree. Hoffman and Lambon Ralph (2013) provided evidence supporting the validity of such an approach. They obtained salience ratings on 8 attributes – visual form, observed motion, color, touch, sound, taste, smell, and performed actions – for 160 object concepts. These attribute representations predicted lexical processing time, and did so more accurately than traditional feature-based representations. It remains an open question, however, whether attribute-based representations derived from human judgments are more or less accurate than text-based co-occurrence measures at predicting neural activity patterns (for a review of the differences between approaches to word semantics based on explicit sensory-motor features and those based on statistical distributions of word co-occurrences – and how they can be integrated – see Andrews et al., 2009).

In a related study, Huth et al. (2012) investigated the cortical representation of semantic categories based on BOLD responses to natural movies. In that study, participants watched movie clips displaying a variety of objects, people, animals and actions, which were coded in the fMRI regression model using 1705 taxonomic category labels from the WordNet lexicon. The authors then identified four significant factors underlying the model fit using principal components analysis. The maps based on these four components showed a degree of correspondence between categorical structure and cortical representation (i.e., similar categories tended to be represented by similar sets of voxels) in individual participants. However, the underlying factors in such an analysis must be interpreted *post hoc*, and the resulting maps offer little insight into how those factors are related to other known functions of the corresponding cortical areas. Our approach focused instead on five sensory-motor dimensions involved in concept acquisition (at least for concrete concepts), and investigated the extent to which these dimensions capture the neural representation of lexical concepts. This approach offers a representational space for the description of concepts in which the dimensions can be related directly to known functions of the brain.

Perhaps the most important difference between our study and those of Mitchell et al. (2008) and Huth et al. (2013) is that concepts in our study were cued solely by word stimuli, with

no pictorial representations of the concepts themselves. Mitchell et al., on the other hand, presented word-picture pairs as stimuli (e.g., the word "airplane" accompanied by a line drawing of an airplane), while Huth et al. presented movie clips. The use of pictorial stimuli in these studies leaves open the possibility that the elicited activation patterns in part reflected perceptual rather than purely semantic aspects of the objects (Haxby et al. 2001; Spiridon and Kanwisher 2002; Cox and Savoy 2003; Kriegeskorte et al. 2008; Haxby et al. 2011). Since similar categories of objects tend to share similar visual properties, similarity in cortical representation could have been driven to a large extent by perceptual similarity. The present study avoided this confound by relying exclusively on word stimuli, thus ensuring that any sensory-related information contributing to prediction accuracy originated from the concept representation itself, rather than the physical stimulus used to cue it. This was further confirmed by the contrasting results obtained for concrete and abstract words.

A final difference between studies is that our analysis was done on group-averaged activation maps, and thus relied entirely on concept-related neural activation patterns that are common across individuals. This approach offers the potential of predicting a participant's activation pattern for a concept based solely on attribute maps generated from other participants. However, the averaging procedure certainly results in loss of information about inter-subject variability, and this may explain in part why the accuracy levels obtained in the present study are generally lower than those obtained by Mitchell et al. (2008). Another factor that probably contributed to the difference in accuracy was the number of semantic features used to generate the predictions (5 in the present study, 25 in Mitchell et al.). Presumably, increasing the number of attributes in our analysis would lead to higher prediction accuracies, provided the extra attributes capture significant aspects of the concept's neural representation.

One potential criticism of the present study is that the task may have induced participants to engage in mental imagery, and that the activation patterns may thus reflect perceptual, rather than conceptual, representations. This argument is based on the assumption that mental imagery and concept retrieval constitute categorically distinct kinds of phenomena, such that imagery engages perceptual representations while conceptual processing does not (e.g., Machery, 2007). In light of this criticism, previous studies have taken steps to minimize conscious, deliberate conceptual elaboration and imagery by using more implicit tasks such as lexical decision and masked word presentation, thus focusing on the automatic, unconscious aspects of conceptual processing (e.g., Kiefer et al., 2012; Trumpp et al., 2013; Willems et al., 2010).

While these studies have been important in demonstrating that early automatic processes in concept retrieval also involve sensory-motor representations, we see the categorical distinction between concepts and imagery as an artificial dichotomy. It seems more likely that concept retrieval can involve a variable amount of sensory-motor reenactment, where the extent of the reenactment depends on the specific demands of the task. We know, for instance, that word imageability (a direct measure of how readily mental images come to mind) affects behavior even on "shallow" tasks such as lexical decision (e.g., Evans et al., 2012), suggesting that "images" are activated to some degree even during such tasks. When more extensive or detailed sensory-motor information is required (e.g., during narrative

comprehension), retrieval of this information may be experienced as conscious imagery. In all cases, however, sensory-motor information about a concept is being retrieved. Thus, we see automatic, unconscious activation of conceptual features, on the one hand, and deliberate, vivid imagery, on the other, as two ends of a continuous range of conceptual knowledge retrieval. Studies that have focused on implicit concept retrieval have made a fundamental contribution to our understanding of conceptual processing; however, implicit activation of conceptual features is only one aspect of a larger, more complex phenomenon.

5. Conclusions and future directions

Our results show that the distributed cortical representation of concrete lexical concepts can be predicted, to a considerable extent, by information about the relative relevance of five attributes of sensory-motor experience – sound, color, motion, shape and manipulation – to the content of those concepts. Furthermore, these results show that the activation patterns elicited by concrete concepts are consistent enough across participants to be useful in predicting group-averaged data.

This approach advances our understanding of the neural substrates of conceptual knowledge by framing individual concepts as points in a multidimensional representational space, defined by elementary features of phenomenal experience that can be mapped onto known functions of the brain. By adding new attributes to the encoding model, we hope to determine which ones capture relevant aspects of this representational space based on how much they contribute to prediction accuracy. By restricting the analysis to different brain regions, we also plan to explore how their contributions to concept representation relate to their known functions in other domains.

Acknowledgements

We would like to thank two anonymous reviewers for their helpful comments and suggestions.

This project was supported by National Institute of Neurological Diseases and Stroke grant R01 NS33576, by National Institutes of Health General Clinical Research Center grant M01 RR00058, and by National Institute of General Medical Sciences grant T32 GM89586.

References

- Andrews M, Frank S, Vigliocco G. Reconciling embodied and distributional accounts of meaning in language. *Topics in Cognitive Science*. 2014; 6(3):359–370. [PubMed: 24935903]
- Aziz-Zadeh L, Wilson SM, Rizzolatti G, Iacoboni M. Congruent embodied representations for visually presented actions and linguistic phrases describing actions. *Current Biology*. 2006; 16(18):1818–1823. [PubMed: 16979559]
- Balota DA, Yap MJ, Cortese MJ, Hutchison KA, Kessler B, Loftis B, Neely JH, Nelson DL, Simpson GB, Treiman R. The English Lexicon Project. *Behavior Research Methods*. 2007; 39:445–459. [PubMed: 17958156]
- Barsalou LW. Grounded cognition. *Annual Review of Psychology*. 2008; 59:617–645.
- Bartels A, Zeki S. The architecture of the colour centre in the human visual brain: new results and a review. *European Journal of Neuroscience*. 2000; 12(1):172–193. [PubMed: 10651872]
- Binder JR, Desai RH. The neurobiology of semantic memory. *Trends in Cognitive Sciences*. 2011; 15(11):527–536. [PubMed: 22001867]

- Binder JR, Desai RH, Graves WW, Conant LL. Where is the semantic system? A critical review and meta-analysis of 120 functional neuroimaging studies. *Cerebral Cortex*. 2009; 19(12):2767–2796. [PubMed: 19329570]
- Bird H, Franklin S, Howard D. Age of acquisition and imageability ratings for a large set of words, including verbs and function words. *Behavior Research Methods, Instruments, & Computers*. 2001; 33(1):73–79.
- Brysbaert M, Warriner AB, Kuperman V. Concreteness ratings for 40 thousand generally known English word lemmas. *Behavior Research Methods*. 2014; 46(3):904–911. [PubMed: 24142837]
- Caspers S, Zilles K, Laird AR, Eickhoff SB. ALE meta-analysis of action observation and imitation in the human brain. *NeuroImage*. 2010; 50(3):1148–1167. [PubMed: 20056149]
- Chao LL, Haxby JV, Martin A. Attribute-based neural substrates in temporal cortex for perceiving and knowing about objects. *Nature Neuroscience*. 1999; 2(10):913–919. [PubMed: 10491613]
- Clark JM, Paivio A. Extensions of the Paivio, Yuille, and Madigan (1968) norms. *Behavior Research Methods, Instruments, & Computers*. 2004; 36(3):371–383.
- Clark VP, Parasuraman R, Keil K, Kulansky R, Fannon S, Maisog JM, et al. Selective attention to face identity and color studied with fMRI. *Human Brain Mapping*. 1997; 5(4):293–297. [PubMed: 20408231]
- Corbetta M, Miezin FM, Dobmeyer S, Shulman GL, Petersen SE. Selective and divided attention during visual discriminations of shape, color, and speed: functional anatomy by positron emission tomography. *The Journal of Neuroscience*. 1991; 11(8):2383–2402. [PubMed: 1869921]
- Cox DD, Savoy RL. Functional magnetic resonance imaging (fMRI) “brain reading”: detecting and classifying distributed patterns of fMRI activity in human visual cortex. *NeuroImage*. 2003; 19(2 Pt 1):261–270. [PubMed: 12814577]
- Cox RW. AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical Research, an International Journal*. 1996; 29(3):162–173.
- Damasio AR. Time-locked multiregional retroactivation: a systems-level proposal for the neural substrates of recall and recognition. *Cognition*. 1989; 33(1–2):25–62. [PubMed: 2691184]
- Evans GAL, Lambon Ralph MA, Woollams AM. What’s in a word? A parametric study of semantic influences on visual word recognition. *Psychonomic Bulletin & Review*. 2012; 19(2):325–331. [PubMed: 22258820]
- Farah MJ, McClelland JL. A computational model of semantic memory impairment: Modality specificity and emergent category specificity. *Journal of Experimental Psychology: General*. 1991; 120(4):339–357. [PubMed: 1837294]
- Fernandino L, Binder JR, Desai RH, Pendl SL, Humphries CJ, Gross WL, Connant LL, Seidenberg MS. Concept Representation Reflects Multimodal Abstraction: A Framework for Embodied Semantics. *Cerebral Cortex*. 2015
- Fischer MH, Zwaan RA. Embodied language: A review of the role of the motor system in language comprehension. *Quarterly Journal of Experimental Psychology*. 2008; 61(6):825–850.
- Glenberg AM, Gallese V. Action-based language: A theory of language acquisition, comprehension, and production. *Cortex*. 2012; 48(7):905–922. [PubMed: 21601842]
- Grill-Spector K, Malach R. The Human Visual Cortex. *Annual Review of Neuroscience*. 2004; 27(1):649–677.
- Hauk O, Johnsrude I, Pulvermüller F. Somatotopic representation of action words in human motor and premotor cortex. *Neuron*. 2004; 41(2):301–307. [PubMed: 14741110]
- Haxby JV, Connolly AC, Guntupalli JS. Decoding Neural Representational Spaces Using Multivariate Pattern Analysis. *Annual Review of Neuroscience*. 2014; 37(1):435–456.
- Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P. Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*. 2001; 293(5539):2425–2430. [PubMed: 11577229]
- Haxby JV, Guntupalli JS, Connolly AC, Halchenko YO, Conroy BR, Gobbini MI, et al. A common, high-dimensional model of the representational space in human ventral temporal cortex. *Neuron*. 2011; 72(2):404–416. [PubMed: 22017997]

- Hoenig K, Müller C, Herrnberger B, Sim E-J, Spitzer M, Ehret G, Kiefer M. Neuroplasticity of semantic representations for musical instruments in professional musicians. *NeuroImage*. 2011; 56(3):1714–1725. [PubMed: 21356317]
- Hoffman P, Lambon Ralph MA. Shapes, scents and sounds: quantifying the full multi-sensory basis of conceptual knowledge. *Neuropsychologia*. 2013; 51(1):14–25. [PubMed: 23159700]
- Hsu NS, Frankland SM, Thompson-Schill SL. Chromaticity of color perception and object color knowledge. *Neuropsychologia*. 2012; 50(2):327–333. [PubMed: 22192637]
- Huth AG, Nishimoto S, Vu AT, Gallant JL. A Continuous Semantic Space Describes the Representation of Thousands of Object and Action Categories across the Human Brain. *Neuron*. 2012; 76(6):1210–1224. [PubMed: 23259955]
- Kiefer M, Pulvermüller F. Conceptual representations in mind and brain: Theoretical developments, current evidence and future directions. *Cortex*. 2012; 48(7):805–825. [PubMed: 21621764]
- Kiefer M, Sim E-J, Herrnberger B, Grothe J, Hoenig K. The sound of concepts: four markers for a link between auditory and conceptual brain systems. *The Journal of Neuroscience*. 2008; 28(47):12224–12230. [PubMed: 19020016]
- Kiefer M, Sim E-J, Liebich S, Hauk O, Tanaka J. Experience-dependent plasticity of conceptual representations in human sensory-motor areas. *Journal of Cognitive Neuroscience*. 2007; 19(3):525–542. [PubMed: 17335399]
- Kiefer M, Trumpp N, Herrnberger B, Sim E-J, Hoenig K, Pulvermüller F. Dissociating the representation of action- and sound-related concepts in middle temporal cortex. *Brain and Language*. 2012; 122(2):120–125. [PubMed: 22726721]
- Kriegeskorte N, Mur M, Ruff DA, Kiani R, Bodurka J, Esteky H, et al. Matching Categorical Object Representations in Inferior Temporal Cortex of Man and Monkey. *Neuron*. 2008; 60(6):1126–1141. [PubMed: 19109916]
- Lewis JW. Cortical networks related to human use of tools. *The Neuroscientist*. 2006; 12(3):211–231. [PubMed: 16684967]
- Lewis JW, Wightman FL, Brefczynski JA, Phinney RE, Binder JR, DeYoe EA. Human brain regions involved in recognizing environmental sounds. *Cerebral Cortex*. 2004; 14(9):1008–1021. [PubMed: 15166097]
- Machery E. Concept empiricism: a methodological critique. *Cognition*. 2007; 104(1):19–46. [PubMed: 16814274]
- Martin A, Wiggs CL, Ungerleider LG, Haxby JV. Neural correlates of category-specific knowledge. *Nature*. 1996; 379(6566):649–652. [PubMed: 8628399]
- McCarthy RA, Warrington EK. Evidence for modality-specific meaning systems in the brain. *Nature*. 1988; 334(6181):428–430. [PubMed: 3261389]
- McKeefry DJ, Zeki S. The position and topography of the human colour centre as revealed by functional magnetic resonance imaging. *Brain*. 1997; 120(12):2229–2242. [PubMed: 9448578]
- Meteyard, L.; Vigliocco, G. The role of sensory and motor information in semantic representation. In: Calvo, P.; Gomila, T., editors. *Handbook of Cognitive Science: An Embodied Approach*. Elsevier; 2008.
- Meteyard L, Cuadrado SR, Bahrami B, Vigliocco G. Coming of age: A review of embodiment and the neuroscience of semantics. *Cortex*. 2012; 48(7):788–804. [PubMed: 21163473]
- Naselaris T, Kay KN, Nishimoto S, Gallant JL. Encoding and decoding in fMRI. *NeuroImage*. 2011; 56(2):400–410. [PubMed: 20691790]
- Oldfield RC. The assessment and analysis of handedness: the Edinburgh inventory. *Neuropsychologia*. 1971; 9(1):97–113. [PubMed: 5146491]
- Patterson K, Nestor PJ, Rogers TT. Where do you know what you know? The representation of semantic knowledge in the human brain. *Nature Reviews Neuroscience*. 2007; 8(12):976–987. [PubMed: 18026167]
- Rauschecker JP, Scott SK. Maps and streams in the auditory cortex: nonhuman primates illuminate human speech processing. *Nature Neuroscience*. 2009; 12(6):718–724. [PubMed: 19471271]
- Spiridon M, Kanwisher N. How distributed is visual category information in human occipito-temporal cortex? An fMRI study. *Neuron*. 2002; 35(6):1157–1165. [PubMed: 12354404]

- Talairach, J.; Tournoux, P. Co-planar stereotaxic atlas of the human brain. 3-Dimensional proportional system: an approach to cerebral imaging. New York: Thieme; 1988.
- Trumpp NM, Traub F, Kiefer M. Masked Priming of Conceptual Features Reveals Differential Brain Activation during Unconscious Access to Conceptual Action and Sound Information. *PLoS ONE*. 2013; 8(5):1–10.
- Wade A, Augath M, Logothetis N, Wandell B. fMRI measurements of color in macaque and human. *Journal of Vision*. 2008; 8(10):1–19. 6.
- Weisberg J, van Turenout M, Martin A. A neural system for learning about object function. *Cerebral Cortex*. 2007; 17(3):513–521. [PubMed: 16581980]
- Warrington EK, Shallice T. Category specific semantic impairments. *Brain*. 1984; 107(Pt 3):829–854. [PubMed: 6206910]
- Willems RM, Toni I, Hagoort P, Casasanto D. Neural dissociations between action verb understanding and motor imagery. *Journal of Cognitive Neuroscience*. 2010; 22(10):2387–2400. [PubMed: 19925195]
- Wilson M. MRC Psycholinguistic Database: Machine-usable dictionary, version 2.00. *Behavior Research Methods, Instruments, & Computers*. 1988; 20(1):6–10.

Highlights

- Does the neural code for concept representation include sensorimotor information?
- We used an encoding model to predict concept-specific neural activation patterns
- The model encoded the relative relevance of five sensorimotor attributes
- It successfully predicted whole-brain activation patterns associated with concepts
- Sensorimotor information plays a role in the neural code for concept representation

Figure 1.
Examples of words used in the study. Every word had a rating for each of the five semantic features.

Figure 2. Prediction analysis. A. Group-level attribute maps were generated from the 820-word modelling set through least squares multiple regression. B. In this example, the predicted map for “coffee” is generated by multiplying each attribute map by the z score of the corresponding attribute rating for that word and adding them together. C. The voxel-by-voxel correlations between the predicted map for “coffee” and each of the 80 group-level observed maps are computed. D. The observed maps are ranked by correlation strength, and

prediction accuracy is determined from the percentile rank of the observed map for “coffee”. Steps B-D are repeated for each word in the test set.

Figure 3. Median prediction accuracy when all 5 attributes were included in the analysis. Error bars represent the 95% confidence interval. * $p < .05$; ** $p < .005$; *** $p < .0005$.

Figure 4. Median prediction accuracy for concrete words, for different subsets of attributes. Error bars represent the 95% confidence interval.

Figure 5.
Mean prediction accuracy for concrete words by number of attributes included in the model.

Table 1

Correlations between semantic attribute ratings.

	Sound	Color	Manipulation	Visual Motion
Color	0.12	–	–	–
Manipulation	0.33	0.31	–	–
Visual Motion	0.59	0.37	0.25	–
Shape	0.28	0.76	0.58	0.53

Table 2

Lexical data [mean (standard deviation)] for words and pseudowords used in the study. Length: Number of letters. Orth: Number of orthographic neighbors. Orth_F: Averaged frequency (per million) of the orthographic neighbors. N2_F: Averaged frequency (per million) of the constrained bigrams for the wordform. N2_C: Number of wordforms that share the same constrained bigrams. N3_F: Averaged frequency (per million) of the constrained trigrams for the wordform. N3_C: Number of wordforms that share the same constrained trigrams.

	Length	Orth	Orth_F	N2_F	N2_C	N3_F	N3_C
Words	5.99 (1.99)	4.12 (5.75)	54.4 (271)	964 (1087)	68.2 (62.6)	151 (240)	12.4 (19.5)
Pseudowords	5.99 (1.99)	4.26 (5.75)	43.4 (93)	976 (759)	68.7 (49.1)	141 (176)	11.4 (10.4)
t test (p)	0.98	0.71	0.49	0.86	0.91	0.51	0.40

Table 3

Lexical and semantic attributes [mean (standard deviation)] for the two subsets of test words. Concreteness data is from Brysbaert et al., (2014). All other lexical attributes were obtained from the English Lexicon Project (Balota et al., 2007; <http://elexicon.wustl.edu>)

	Concrete	Abstract	T test (p)
Number of letters	5.75 (1.97)	5.75 (2.02)	0.70
Number of phonemes	4.47 (1.75)	4.87 (1.9)	0.33
Number of syllables	1.67 (0.8)	2 (1.01)	0.11
Log Frequency HAL	9.86 (1.31)	9.44 (1.5)	0.18
Orth. neighborhood	4.87 (5.51)	4.27 (5.81)	0.64
Phon. neighborhood	9.57 (10.22)	8.72 (11.52)	0.73
Bigram frequency	1722 (842)	1902 (933)	0.37
Concreteness	4.81 (0.19)	2.21 (0.65)	< .0001
Sound rating	2.39 (1.47)	0.96 (0.84)	< .0001
Color rating	3.32 (1.07)	0.60 (0.61)	< .0001
Manipulation rating	2.43 (1.42)	0.78 (0.53)	< .0001
Motion rating	2.42 (1.68)	0.81 (0.77)	< .0001
Shape rating	3.90 (1.27)	0.33 (0.26)	< .0001