

UC Berkeley

UC Berkeley Previously Published Works

Title

Natural speech reveals the semantic maps that tile human cerebral cortex

Permalink

<https://escholarship.org/uc/item/4vt8r82n>

Journal

Nature, 532(7600)

ISSN

0028-0836

Authors

Huth, Alexander G
de Heer, Wendy A
Griffiths, Thomas L
[et al.](#)

Publication Date

2016-04-01

DOI

10.1038/nature17637

Peer reviewed



Published in final edited form as:

Nature. 2016 April 28; 532(7600): 453–458. doi:10.1038/nature17637.

Natural speech reveals the semantic maps that tile human cerebral cortex

Alexander G. Huth^a, Wendy A. de Heer^b, Thomas L. Griffiths^{a,b}, Frédéric E. Theunissen^{a,b}, and Jack L. Gallant^{a,b}

^aHelen Wills Neuroscience Institute, University of California, Berkeley, CA 94720, USA

^bDepartment of Psychology, University of California, Berkeley, CA 94720, USA

Abstract

The meaning of language is represented in regions of the cerebral cortex collectively known as the “semantic system”. However, little of the semantic system has been mapped comprehensively, and the semantic selectivity of most regions is unknown. Here we systematically map semantic selectivity across the cortex using voxel-wise modeling of fMRI data collected while subjects listened to hours of narrative stories. We show that the semantic system is organized into intricate patterns that appear consistent across individuals. We then use a novel generative model to create a detailed semantic atlas. Our results suggest that most areas within the semantic system represent information about specific semantic domains, or groups of related concepts, and our atlas shows which domains are represented in each area. This study demonstrates that data-driven methods—commonplace in studies of human neuroanatomy and functional connectivity—provide a powerful and efficient means for mapping functional representations in the brain.

Introduction

Previous neuroimaging studies have identified a group of regions that seem to represent information about the meaning of language. These regions, collectively known as the “semantic system”, respond more to words than non-words¹, more to semantic tasks than phonological tasks¹, and more to natural speech than temporally scrambled speech². Studies that have investigated specific types of representation in the semantic system have found areas selective for concrete or abstract words^{3–5}, action verbs⁶, social narratives⁷ or other semantic features. Others have found areas selective for specific semantic domains—groups of related concepts such as living things, tools, food, or shelter^{8–13}. However, all previous studies tested only a handful of stimulus conditions, so no study has yet produced a

Users may view, print, copy, and download text and data-mine the content in such documents, for the purposes of academic research, subject always to the full Conditions of use: http://www.nature.com/authors/editorial_policies/license.html#terms

Correspondence and requests for materials should be addressed to J.L.G. (gallant@berkeley.edu).

Supplementary Information is linked to the online version of the paper.

Author Contributions All authors helped conceive of and design the experiment. W.A.dH. and A.G.H. selected and annotated stimuli and collected fMRI data. A.G.H. analyzed the data. A.G.H. and T.L.G. designed the PrAGMATiC generative model. A.G.H. and J.L.G. wrote the paper. J.L.G. contributed to all aspects of the project.

The authors declare no competing financial interests.

comprehensive survey of how semantic information is represented across the entire semantic system.

Here we addressed this problem by using a data-driven approach¹⁴ to model brain responses elicited by naturally spoken narrative stories that contain many different semantic domains¹⁵. Seven subjects listened to more than two hours of stories from *The Moth Radio Hour*² while whole-brain blood-oxygen level dependent (BOLD) responses were recorded by functional magnetic resonance imaging (fMRI). We then used voxel-wise modeling (VM), a highly effective approach for modeling responses to complex natural stimuli^{14–17}, to estimate the semantic selectivity of each voxel (Figure 1A).

Estimation and validation of semantic voxel-wise models

In VM, features of interest are first extracted from the stimuli and then regression is used to determine how each feature modulates BOLD responses in each voxel. Here we used a word embedding space to identify semantic features of each word in the stories^{12,15,18–20}. The embedding space was constructed by computing the normalized co-occurrence between each word and a set of 985 common English words (such as *above*, *worry*, and *mother*) across a large corpus of English text. Words related to the same semantic domain tend to occur in similar contexts, and so have similar co-occurrence values. For example, the words “month” and “week” are very similar (the correlation between the two is 0.74), while the words “month” and “tall” are not (correlation –0.22).

Next we used regularized linear regression to estimate how the 985 semantic features influenced BOLD responses in every cortical voxel and in each individual subject (Figure 1A). To account for responses caused by low-level properties of the stimulus such as word rate and phonemic content, additional regressors were included during VM estimation and then discarded before further analysis. We also included additional regressors to account for physiological and emotional factors, but these had no effect on the estimated semantic models (Supplemental Results 3).

One advantage of VM over conventional neuroimaging approaches is that the fit models can be validated by predicting BOLD responses to new natural stimuli that were not used during model estimation. This makes it possible to compute effect size by finding the fraction of response variance explained by the models. Here we tested how well the voxel-wise models predicted BOLD responses elicited by a new 10-minute *Moth* story (Figure 1B) that had not been used for model estimation. We found good prediction performance for voxels located throughout the semantic system, including in the lateral and ventral temporal cortex, lateral and medial parietal cortex, and medial, superior, and inferior prefrontal cortex (Figure 1C and Extended Data Fig. 1). This suggests that much of the semantic system is domain-selective.

Mapping semantic representation across cortex

By inspecting the fit models we can determine which specific semantic domains are represented in each voxel. In theory this could be done by examining each voxel separately. However, our data consist of tens of thousands of voxels per subject, rendering this approach

infeasible. A practical alternative is to project the models into a low-dimensional subspace that retains as much information as possible about the semantic tuning of the voxels^{10,14}. We found such a space by applying principal components analysis to the estimated models aggregated across subjects, producing 985 orthogonal semantic dimensions that are ordered by how much variance each explained across the voxels. It is likely that only some of these dimensions capture shared aspects of semantic tuning across the subjects; the rest reflect individual differences, fMRI noise, or the statistical properties of the stories. To identify the shared dimensions we tested whether each explained more variance across the models than expected by chance, which was defined by the principal components of the stimulus matrix used for model estimation¹⁴. At least four dimensions explained a significant amount of variance ($p < 0.001$, Bonferroni corrected bootstrap test) in all but one subject; in the last subject only three dimensions were significant (Extended Data Fig. 2). This suggests that our fMRI data contain about four statistically significant semantic dimensions that are shared across subjects.

The four shared semantic dimensions provide a way to succinctly summarize the semantic selectivity of a voxel. However, to interpret projections of the models onto these dimensions we need to understand how semantic information is encoded in this four-dimensional space. To visualize the semantic space we projected the 10,470 words in the stories from the word embedding space onto each dimension. We then used k-means clustering to identify twelve distinct categories (see Supplemental Methods for details). Each category was inspected and labeled by hand. The labels assigned to the twelve categories were *tactile* (a cluster containing words such as “fingers”), *visual* (words such as “yellow”), *numeric* (“four”), *locational* (“stadium”), *abstract* (“natural”), *temporal* (“minute”), *professional* (“meetings”), *violent* (“lethal”), *communal* (“schools”), *mental* (“asleep”), *emotional* (“despised”), and *social* (“child”). (See Supplementary Table 2 and Supplementary Results for more detailed evaluations of each category.)

Next, we visualized where each of the twelve categories appeared in the shared semantic space (Figure 2A). Each category label was also assigned an RGB color, where the red channel was determined by the first dimension, the green channel by the second, and the blue channel by the third. The first dimension is that which captured the most semantic variance across the voxel-wise models of all seven subjects. One end of this dimension favors categories related to humans and social interaction, including *social*, *emotional*, *violent*, and *communal*. The other end favors categories related to perceptual descriptions, quantitative descriptions and setting, including *tactile*, *locational*, *numeric* and *visual*. This is consistent with previous suggestions that humans comprise a particularly salient and strongly represented semantic domain^{16,21}. Subsequent dimensions of the semantic space captured less variance than the first and were also more difficult to interpret. The second dimension seems to distinguish between perceptual categories, including *visual* and *tactile*, and non-perceptual categories, including *mental*, *professional*, and *temporal*. The third and fourth dimensions are less clear.

Earlier studies identified the cortical regions comprising the semantic system^{1,2}, but could not comprehensively characterize their semantic selectivity. Here we were able to visualize the pattern of semantic domain-selectivity across the entire cortex by projecting voxel-wise

models onto the shared semantic dimensions. Figure 2B shows projections onto the first three dimensions for one subject, plotted together using the same RGB color scheme as in Figure 2A (Extended Data Fig. 3A shows each dimension separately). Thus, for example, a green voxel produces greater BOLD responses to categories that are colored green in the semantic space, such as *visual* and *numeric*. This visualization suggests that semantic information is represented in intricate patterns that cover the semantic system, including broad regions of prefrontal cortex, lateral and ventral temporal cortex, and lateral and medial parietal cortex. Furthermore, these patterns appear to be relatively consistent across individuals (Figure 2C; see also Extended Data Figure 3B).

Using PrAGMATiC to construct a semantic atlas

Given the apparent consistency in the patterns of semantic selectivity across individuals, we sought to create a single atlas that describes the distribution of semantically selective functional areas in human cerebral cortex. To accomplish this we developed a new Bayesian algorithm, PrAGMATiC, that produces a Probabilistic And Generative Model of Areas Tiling the Cortex²². This algorithm models patterns of functional tuning recovered by VM as a dense, tiled map of functionally homogeneous brain areas (Fig. 3A), while respecting individual differences in anatomical and functional anatomy^{23,24}. The arrangement and selectivity of these areas are determined by parameters learned from the fMRI data through a maximum likelihood estimation technique similar to contrastive divergence²⁵. Some parameters are shared; these describe properties of the cortical map that are common across the group. Other parameters are unique to each subject; these capture individual differences. Learning both shared and unique parameters simultaneously eliminates the usual requirement to perform anatomical or functional alignment of data across subjects.

The PrAGMATiC algorithm has two components: an arrangement model that determines where functional areas appear on the cortical sheet, and an emission model that determines how the cortical map is produced from an arrangement of areas. The arrangement model simulates a physical spring network that joins the centroid of each functional area to its neighbors. Equilibrium spring lengths are shared across subjects, but each spring can be stretched or compressed in any individual subject. Arrangements are also constrained by several functional landmarks, which are known regions-of-interest identified in every subject using separate functional data. These constraints ensure that the maps will be similar across subjects, but allow for substantial individual variability in the precise arrangement and size of the areas. Using the arrangement model, the emission model creates homogeneous functional areas by assigning each vertex on the cortical surface to the nearest area centroid. The functional value at each vertex is then drawn from a multivariate normal distribution. The mean functional value for each area is learned by the algorithm and is shared across subjects. Here we define the functional value as a four-dimensional vector that reflects the projection of the estimated model for each voxel onto the four shared semantic dimensions.

One important hyperparameter is the total number of areas that PrAGMATiC uses to tile the cortex. We used a cross-validation procedure to choose the total number of areas tiling each hemisphere and then tested whether each area is semantically selective. PrAGMATiC models were estimated on data from six subjects and then used to predict the semantic map

in the seventh subject using only cortical anatomy and the locations of functional landmarks in that subject. Predicted BOLD responses based on this map were compared to actual responses to determine how well the PrAGMATiC model generalizes across subjects. Prediction performance climbed quickly as the total number of areas rose from 8 to 128 and improved more gradually thereafter (Fig. 3B). In the left hemisphere, prediction performance did not improve significantly for models with 192 or more total areas ($q(\text{FDR}) > 0.01$, Tukey post hoc test with subject-wise random effects). In the right hemisphere, prediction performance did not improve significantly for models with 128 or more total areas. However, because PrAGMATiC tiles the entire cerebral cortex, these numbers include both semantically selective and nonselective areas. To identify the semantically selective areas and eliminate those that are nonselective, we tested whether the average voxel-wise semantic model in each area predicted responses significantly better than the average model for low-level features such as word rate, phoneme rate, and phonemes. This excluded areas that were not selective for either semantic or low-level features, such as motor and visual cortex. It also excluded areas that were not uniquely selective for semantic features, such as Broca's area, which was desirable because of the increased uncertainty of semantic model weights in those areas.

Figure 3C shows the semantic atlas projected onto the cortical surface of one subject (see also Extended Data Figures 4 & 5). The left hemisphere contains 77 semantic areas ($q(\text{FDR}) < 1/192$, bootstrap test) and the right contains 63 semantic areas ($q(\text{FDR}) < 1/128$, bootstrap test). A diverse tiling of areas that represent different semantic domains appear in lateral parietal cortex (LPC, Extended Data Fig. 6), medial parietal cortex (MPC, Extended Data Fig. 7), and superior prefrontal cortex (SPFC, Extended Data Fig. 8). In LPC and MPC, central areas (near the angular gyrus and subparietal sulci, respectively) are selective for *social* concepts, while surrounding areas are selective for *numeric*, *visual*, or *tactile* concepts. In SPFC, medial areas are mainly selective for *social* concepts, while dorsolateral areas are more diverse. The LPC, MPC, and SPFC also all belong to the default mode network (DMN), which is thought to be involved in introspection, rumination, and conscious thought²⁶. One interesting possibility is that semantic areas identified here represent the same semantic domains during conscious thought. This suggests that the contents of thought, or internal speech, might be decoded using these voxel-wise models¹⁷. In the lateral temporal cortex (LTC, Extended Data Fig. 9) our atlas identifies fewer distinct semantic areas than in LPC, MPC, or SPFC. This is surprising because LTC plays a key role in language comprehension^{1,27} and also belongs to the DMN. However, the quality of fMRI signals recorded in the anterior temporal lobe is poor, so LTC likely contains other semantic areas that could not be recovered using our current approach. Detailed analyses of semantic representations in LPC, MPC, SPFC, and LTC, as well as ventral temporal cortex (Extended Data Fig. 10), inferior prefrontal cortex (Extended Data Fig. 11), and opercular and insular cortex (Extended Data Fig. 12) can be found in Supplemental Materials, along with discussion and comparisons to earlier neuroimaging and lesion results.

Discussion

One striking aspect of our atlas is that the distribution of semantically selective areas is relatively symmetric across the two cerebral hemispheres. This finding is inconsistent with

human lesion studies that argue that semantic representation is lateralized to the left hemisphere¹³. However, many fMRI studies of semantic representation find only modest lateralization¹ and one study that used narrative stories found highly bilateral results similar to ours². This suggests that right hemisphere areas may respond more strongly to narrative stimuli than to the words and short phrases used in most studies. Still, more research will be needed to determine what roles these left- and right-hemisphere semantic areas play in language comprehension.

Another interesting aspect of these results is that the organization of semantically selective brain areas appears highly consistent across individuals. This might suggest that innate anatomical connectivity or cortical cytoarchitecture constrains the organization of high-level semantic representations^{28,29}. It is also possible that this is due to common life experiences of the subjects, all of whom were raised and educated in western industrial societies. Future studies that include subjects from more diverse backgrounds will be needed to determine how much of this organizational consistency reflects innate brain structure versus experience.

One limitation of PRAGMATiC as used here is that each area is assumed to be functionally homogeneous. This is a common assumption in the design and analysis of many neuroimaging studies³⁰. However, many cortical maps, including semantic maps in visual cortex¹⁴, seem to contain smoothly-changing gradients of representation. It should be possible to modify the PRAGMATiC algorithm to model functional gradients explicitly. This will provide an objective tool for determining whether the semantic maps found here are best described as homogeneous areas or as gradients.

Data-driven approaches are commonplace in studies of human neuroanatomy³¹ and resting state networks^{26,32}, but are only beginning to be used in functional imaging^{14,15}. Our study demonstrates the power and efficiency of data-driven approaches for functional mapping of the human brain. Although our experiment used a simple design in which subjects only listened to stories, the data were rich enough to produce a comprehensive atlas of semantically selective areas. Furthermore, our data-driven framework is quite general. Other properties of language can be mapped (even in this same dataset) by using feature spaces that reflect phonemes, syntax and so on. Complex semantic models that incorporate information beyond word co-occurrence can be tested and compared quantitatively. The generalizability of these models can also be tested by using stimuli beyond autobiographical stories. It is sometimes difficult to synthesize the results of data-driven experiments with those from hypothesis-driven experiments, but future methodological and theoretical developments should help to bridge this divide. We expect that the semantic atlas presented here will be useful for many researchers investigating the neurobiological basis of language. We also expect that this atlas can be refined and expanded by incorporating results from future studies. To facilitate this, we have created a detailed interactive version of the semantic atlas that can be explored online at <http://gallantlab.org/huth2016>.

Methods

MRI data collection

MRI data were collected on a 3T Siemens TIM Trio scanner at the UC Berkeley Brain Imaging Center using a 32-channel Siemens volume coil. Functional scans were collected using gradient echo EPI with TR = 2.0045s, TE = 31ms, flip angle = 70 degrees, voxel size = $2.24 \times 2.24 \times 4.1$ mm (slice thickness = 3.5 mm with 18% slice gap), matrix size = 100×100 , and field of view = 224×224 mm. 30 axial slices were prescribed to cover the entire cortex and were scanned in interleaved order. A custom-modified bipolar water excitation radiofrequency (RF) pulse was used to avoid signal from fat. Anatomical data were collected using a T1-weighted multi-echo MP-RAGE sequence on the same 3T scanner.

Subjects

Functional data were collected from five male subjects and two female subjects: S1 (male, age 26), S2 (male, age 32), S3 (female, age 31), S4 (male, age 31), S5 (male, age 26), S6 (female, age 25), and S7 (male, age 30). Two of the subjects were authors on the paper (S1: AGH and S3: WAdH). All subjects were healthy and had normal hearing. The experimental protocol was approved by the Committee for the Protection of Human Subjects at University of California, Berkeley. Written informed consent was obtained from all subjects.

Natural Story Stimuli

The model estimation dataset consisted of ten 10- to 15-minute stories taken from *The Moth Radio Hour*. In each story, a single speaker tells an autobiographical story in front of a live audience. The ten selected stories cover a wide range of topics and are highly engaging. Each story was played during a separate fMRI scan. The length of each scan was tailored to the story, and included 10 seconds of silence both before and after the story. These data were collected during two 2-hour scanning sessions that were performed on different days. The model validation dataset consisted of one 10-minute story, also taken from *The Moth Radio Hour*. This story was played twice for each subject (once during each scanning session), and then the two responses were averaged. For story synopses and details of story transcription and preprocessing procedures see Supplemental Methods.

Stories were played over Sensimetrics S14 in-ear piezoelectric headphones (Sensimetrics, Malden, MA, USA). A Behringer Ultra-Curve Pro hardware parametric equalizer was used to flatten the frequency response of the headphones based on calibration data provided by Sensimetrics. All stimuli were played at 44.1 kHz using the pygame library in Python. All stimuli were normalized to have peak loudness of -1 dB relative to max. However, the stories were performed by different speakers and were not uniformly mastered, so some differences in total loudness remain.

Story transcription and preprocessing

Each story was manually transcribed by one listener, and then the transcript was checked by a second listener. Certain sounds (e.g. laughter, lip-smacking and breathing) were also marked in order to improve the accuracy of the automated alignment. The audio of each story was downsampled to 11 kHz and the Penn Phonetics Lab Forced Aligner (P2FA³³)

was used to automatically align the audio to the transcript. The forced aligner uses a phonetic hidden Markov model to find the temporal onset and offset of each word and phoneme. The CMU pronouncing dictionary was used to guess the pronunciation of each word. When necessary, words and word fragments that appeared in the transcript but not in the dictionary were manually added. After automatic alignment was complete, Praat³⁴ was used to check and correct each aligned transcript manually. The corrected aligned transcript was then spot-checked for accuracy by a different listener.

Finally the aligned transcripts were converted into separate word and phoneme representations. The phoneme representation of each story is a list of pairs (p, t) , where p is a phoneme and t is the time from the beginning of the story to the middle of the phoneme (i.e. halfway between the start and end of the phoneme) in seconds. Similarly the word representation of each story is a list of pairs (w, t) , where w is a word.

Semantic model construction

To account for response variance caused by the semantic content of the stories we constructed a 985-dimensional semantic feature space based on word co-occurrence statistics in a large corpus of text^{12,18,19}. First, we constructed a 10,470 word lexicon from the union of the set of all words appearing in the stories and the 10,000 most common words in the large text corpus. We then selected 985 basis words from Wikipedia's *List of 1000 basic words* (contrary to the title, this list contained only 985 unique words at the time it was accessed). This basis set was selected because it consists of common words that span a very broad range of topics. The text corpus used to construct this feature space includes the transcripts of 13 *Moth* stories (including the 10 used as stimuli in this experiment), 604 popular books, 2,405,569 Wikipedia pages, and 36,333,459 user comments scraped from reddit.com. In total the 10,470 words in our lexicon appeared 1,548,774,960 times in this corpus.

Next, we constructed a word co-occurrence matrix, M , with 985 rows and 10,470 columns. Iterating through the text corpus, we added 1 to $M_{i,j}$ each time word j appeared within 15 words of basis word i . A window size of 15 was selected to be large enough to suppress syntactic effects (i.e. word order) but no larger. Once the word co-occurrence matrix was complete we log-transformed the counts, replacing $M_{i,j}$ with $\log(1 + M_{i,j})$. Next, each row of M was z-scored to correct for differences in basis word frequency, and then each column of M was z-scored to correct for word frequency. Each column of M is now a 985-dimensional semantic vector representing one word in the lexicon.

The matrix used for voxel-wise model estimation was then constructed from the stories: for each word-time pair (w, t) in each story we selected the corresponding column of M , creating a new list of semantic vector-time pairs, (M_w, t) . These vectors were then resampled at times corresponding to the fMRI acquisitions using a 3-lobe Lanczos filter with the cutoff frequency set to the Nyquist frequency of the fMRI acquisition (0.249 Hz).

Voxel-wise model estimation and validation

A linearized finite impulse response (FIR) model^{14,17} consisting of four separate feature spaces was fit to every cortical voxel in each subject's brain. These four feature spaces were

word rate (1 feature), phoneme rate (1 feature), phonemes (39 features), and semantics (985 features). The word rate, phoneme rate, and phoneme features were used to account for responses to low-level properties of the stories that could contaminate the semantic model weights (see Supplemental Methods for details of how these low-level models were constructed). A separate linear temporal filter with four delays (1, 2, 3, and 4 time points) was fit for each of these 1026 features, yielding a total of 4104 features. This was accomplished by concatenating feature vectors that had been delayed by 1, 2, 3, and 4 time points (2, 4, 6, and 8 seconds). Thus, in the concatenated feature space one channel represents the word rate 2 seconds earlier, another 4 seconds earlier, and so on. Taking the dot product of this concatenated feature space with a set of linear weights is functionally equivalent to convolving the original stimulus vectors with linear temporal kernels that have non-zero entries for 1-, 2-, 3-, and 4-time point delays.

Before doing regression, we first z-scored each feature channel within each story. This was done to match the features to the fMRI responses, which were also z-scored within each story. However, this had little effect on the learned weights.

The 4104 weights for each voxel were estimated using L2-regularized linear regression (a.k.a. ridge regression). To keep the scale of the weights consistent and to prevent bias in subsequent analyses, a single value of the regularization coefficient was used for all voxels in all subjects. This regularization coefficient was found by bootstrapping the regression procedure 50 times in each subject. In each bootstrap iteration, 800 time points (20 blocks of 40 consecutive time points each) were removed from the model estimation dataset and reserved for testing. Then the model weights were estimated on the remaining 2937 time points for each of 20 possible regularization coefficients (log spaced between 10 and 1000). These weights were used to predict responses for the 800 reserved time points, and then the correlation between actual and predicted responses was found. After the bootstrapping was complete, a regularization-performance curve was obtained for each subject by averaging the bootstrap sample correlations first across the 50 samples and then across all voxels. Next, the regularization-performance curves were averaged across the seven subjects and the best overall value of the regularization parameter (183.3) was selected. The best overall regularization parameter value was also the best value in three individual subjects. For the other four subjects the best regularization parameter value was slightly higher (233.6).

To validate the voxel-wise models, estimated semantic feature weights were used to predict responses to a separate story that had not been used for weight estimation. Prediction performance was then estimated as the Pearson correlation between predicted and actual responses for each voxel over the 290 time points in the validation story. Statistical significance was computed by comparing estimated correlations to the null distribution of correlations between two independent Gaussian random vectors of the same length. Resulting *p*-values were corrected for multiple comparisons within each subject using the false discovery rate (FDR) procedure³⁵.

All model fitting and analysis was performed using custom software written in Python, making heavy use of NumPy³⁶, SciPy³⁷, and pycortex³⁸.

Semantic principal components analysis

We used principal components analysis (PCA) to recover a low-dimensional semantic space from the estimated semantic model weights. We first selected only the 10,000 best predicted voxels in each subject according to the average bootstrap correlation (for the selected regularization parameter value) obtained during model estimation. This was done to avoid including noise from poorly modeled voxels. Then we removed temporal information from the voxel-wise model weights by averaging across the four delays for each feature. The weights for the word frequency, phoneme frequency, and phoneme features were then discarded, leaving only the 985 semantic model weights for each voxel. Finally, we applied PCA to these weights, yielding 985 principal components (PCs). Partial scree plots showing the amount of variance accounted for by each PC are shown in Extended Data Figure 2. See Supplemental Methods for details.

PrAGMATiC

The PrAGMATiC generative model²² has two components: an arrangement model and an emission model. The arrangement model defines a probability distribution over possible arrangements of the functional areas. This model assumes that the location of each area is defined by a single point called the area centroid. Each centroid is modeled as being joined to nearby centroids by springs. While exact centroid locations can vary from subject to subject, the equilibrium length of each spring is assumed to be consistent across subjects. The probability distribution over possible locations of the centroids is defined using the total potential energy of the spring system. This distribution assigns a high probability to low-energy arrangements of the centroids (i.e. where the springs are not stretched much and so store little potential energy) and low probability to high-energy arrangements (where the springs are stretched a lot).


The second component is the emission model, which defines a probability distribution over semantic maps given an arrangement of functional areas. In the emission model each area centroid is assigned a particular semantic value in the 4-D common semantic space. This value determines what type of semantic information is represented in that area. To generate a semantic map from any particular arrangement, each point on the cortical surface is first assigned to the closest area centroid (creating a Voronoi diagram). Then the semantic value for each point is sampled from a spherical Gaussian distribution in semantic space, centered on the semantic value of the centroid.

A consequence of modeling semantic maps using a Voronoi diagram is that every point on the cortex must be assigned to an area, while we know that many points on the cortex are not semantically selective. We distinguished between semantically selective and non-selective areas by testing whether the mean semantic voxel-wise model in each area predicted responses significantly better on a held-out story than a baseline model that accounts for responses to phonemes and word rate.

To train the generative model we derived maximum likelihood estimation (MLE) update rules similar to the Boltzmann learning rule with contrastive divergence²⁵. We used these learning rules to iteratively update the spring lengths and semantic values, maximizing the

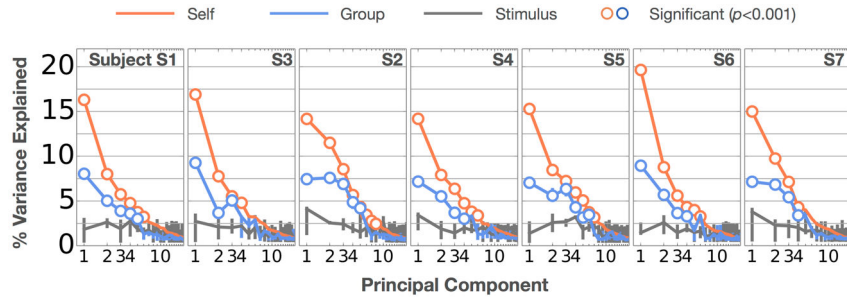
probability of the observed maps and minimizing the probability of unobserved maps. For details see Supplemental Methods.

Extended Data




Extended Data Figure 1. Voxel-wise model prediction performance

Cortical flatmaps showing prediction performance of voxel-wise semantic models for all seven subjects, formatted similarly to Figure 1C in the main text. Models were tested using one 10-minute story that was not included during model estimation. Prediction performance was then computed as the correlation between predicted and measured BOLD responses. (Left column) Raw prediction performance. Note that the colormap here is scaled 0–1 rather than 0–0.6 as in the main text in order to match the scale of the adjusted prediction performance maps. (Right column) Prediction performance corrected to account for different amounts of noise in the BOLD responses (see Supplemental Methods for details). The voxel-wise semantic models predict BOLD responses in many brain areas, including superior and inferior prefrontal cortex (SPFC, IPFC), lateral and ventral temporal cortex (LTC, VTC), and lateral and medial parietal cortex (LPC, MPC). As explained in the main text, these same regions have been previously identified as the “semantic system” in the human brain.



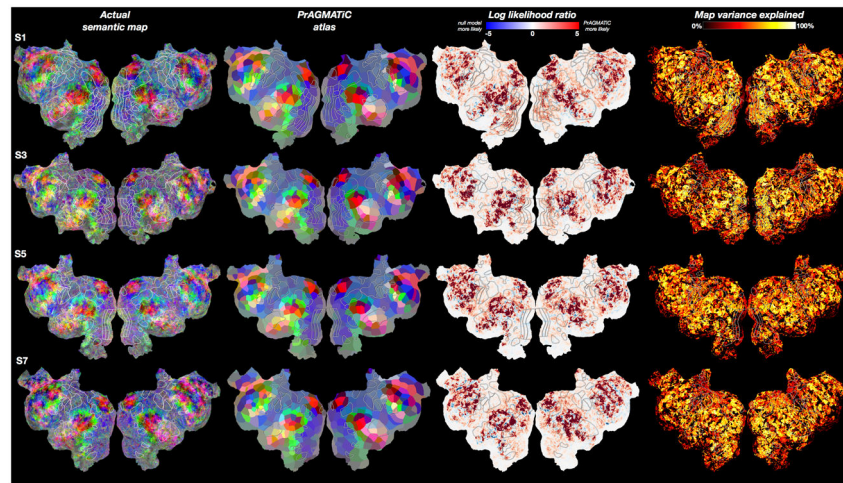
Extended Data Figure 2. Amount of variance explained by individual subject and group semantic dimensions

Principal components analysis (PCA) was used to discover the most important semantic dimensions from voxel-wise semantic model weights in each subject. To reduce noise, we used only the 10,000 best voxels in each subject, determined by cross-validation within the model estimation dataset. Here we show the amount of variance explained in the semantic model weights by each of the 20 most important PCs. Orange lines show the amount of variance explained each subject’s own PCs, blue lines show the variance explained by the PCs of combined data from the other six subjects, and gray lines show the variance explained by the PCs of the stories. (The Gale-Shapley stable marriage algorithm was used to re-order the group and stimulus PCs to maximize their correlation with the subject’s PCs.) Error bars indicate 99% confidence intervals. Confidence intervals for the subjects’ own PCs and group PCs are very small. Hollow markers indicate subject or group PCs that explain significantly more variance than the corresponding stimulus PCs ($p < 0.001$, bootstrap test). Six PCs explain significantly more variance in one out of seven subjects, five PCs in two subjects, four PCs in three subjects, and three PCs in one subject. Thus, four PCs seem to comprise a semantic space that is common across most individuals.




Extended Data Figure 3. Separate cortical projections of semantic dimensions 1-4 on subject S2 and combined cortical projections of dimensions 1-3 for subjects S1, S3, and S4

(a) Voxel-wise semantic model weights for subject S2 were projected onto each of the common semantic dimensions defined by PCs 1-4. Voxels for which model generalization performance was not significantly greater than zero ($q(\text{FDR}) > 0.05$) are shown in gray. Positive projections are shown in red, negative projections in blue and near-zero projections in white. Voxels with fMRI signal dropout due to field inhomogeneity are shaded with black hatched lines. (b) Like Figures 2B and 2C in the main text, this figure shows the result of projecting voxel-wise models onto the first three common semantic dimensions, and then coloring each voxel using an RGB colormap. The red color component corresponds to the projection on the first PC, the green component to the second, and the blue component to the third. Semantic information seems to be represented in complex patterns distributed across the semantic system and the patterns seem to be largely conserved across individuals.



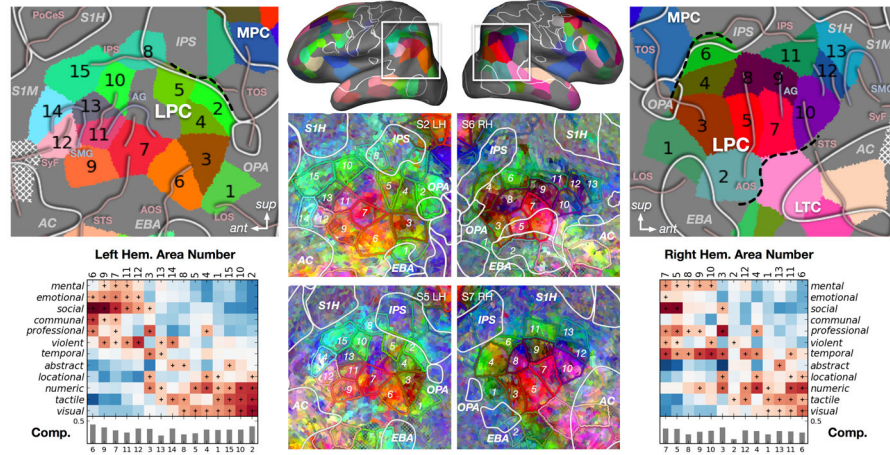
Extended Data Figure 4. PrAGMATiC atlas likelihood maps

Comparison of actual semantic maps (Figure 2, Extended Data Figure 3) to the maps generated from the PrAGMATiC atlas (Figure 3). PrAGMATiC atlases for the left and right hemispheres were fit using data from all seven subjects. The left hemisphere atlas has 192 total areas and the right hemisphere has 128 (including non-semantic areas). Here we show (first column) the actual semantic maps for four subjects, (second column) the PrAGMATiC atlas on each subject's cortical surface, (third column) the log likelihood ratio of the actual semantic map under the PrAGMATiC atlas versus a null model, and (fourth column) the fraction of variance in the semantic map that the PrAGMATiC atlas explains for each location on the cortical surface. The likelihood ratio maps show that most areas where there are large semantic model weights (i.e. the semantic system) are much better explained by PrAGMATiC than by a null model and thus appear red, while areas where the weights are small (i.e. somatomotor cortex, visual cortex, etc.) are about equally well explained by both PrAGMATiC and the null model and thus appear white. Variance explained was computed by subtracting the PrAGMATiC atlas from the actual semantic map (in the space of the four group semantic dimensions), squaring and summing the residuals and then dividing by the sum of squares in the actual map. The variance explained maps show that the PrAGMATiC atlas captures a large fraction of the variance in the semantic maps (37–47% in total).



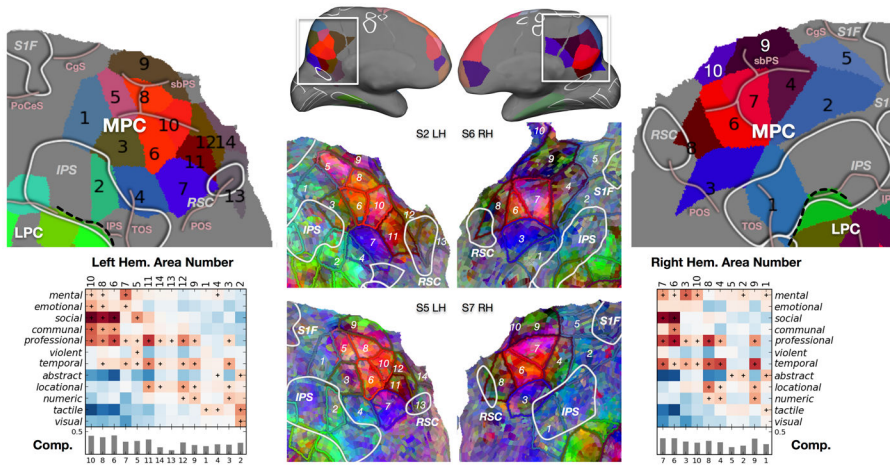
Extended Data Figure 5. Comparison of PrAGMATiC models fit with different initial conditions

As with many clustering algorithms, PrAGMATiC optimizes a non-convex objective function and so can find many potential locally optimal solutions. To reduce the effect of non-convexity on our results, we re-fit the model 10 times (each time with a different random initialization), and then selected the model fit that yielded the best likelihood (i.e. performance on the training set) as the PrAGMATiC atlas (Figure 3). Here we show (top) the PrAGMATiC atlas and (bottom) the second best model out of the 10 that were estimated. The parcellations given by these two models are very similar. However, there are a few differences, which illustrate uncertainty in the model. Some of these differences are due to statistical thresholding: a few areas that were found to be significantly semantically selective in the best model are missing in the alternate model (see left medial prefrontal cortex), and some significant areas in the alternate model are missing from the best model (left ventral occipital cortex). Other differences suggest alternative parcellations for a few regions, where, for example, the same region of cortex is parcellated into 3 areas in the best model and 4 areas in the alternate model. Yet it is clear that none of the differences between these two models are sufficient to change any of the interpretations given in the main text.



Extended Data Figure 6. Semantic atlas for lateral parietal cortex (LPC)

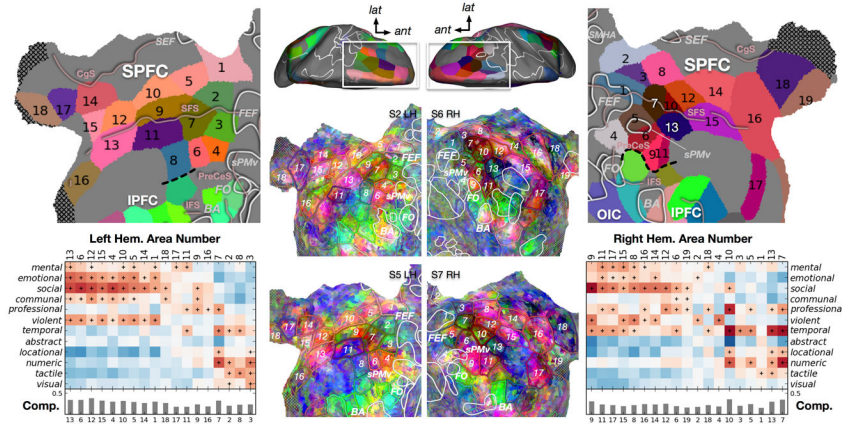
The PRAGMATiC atlas divides LPC into 15 areas in the left hemisphere and 13 areas in the right. Here we show (top left and right) the atlas for each hemisphere, (top middle) 3-D brains indicating the location of LPC, (bottom middle) individual maps for two subjects in each hemisphere, and (bottom left and right) the average predicted response of each area to the 12 semantic categories identified earlier (responses consistently greater than zero across subjects are marked with “+”). Bars show how completely this 12 category interpretation captures the average semantic model in each area. LPC appears to be organized around the angular gyrus (AG), with a core that is selective for *social*, *emotional*, and *mental* concepts (L6, 7, 9, 11; R5, 7) and a periphery that is selective for *visual*, *tactile*, and *numeric* concepts (L2, 4, 5, 8, 10, 15; R6, 11).



Extended Data Figure 7. Semantic atlas for medial parietal cortex (MPC)


The PRAGMATiC atlas divides MPC into 14 areas in the left hemisphere and 10 areas in the right. Here we show (top left and right) the atlas for each hemisphere, (top middle) 3-D brains indicating the location of MPC, (bottom middle) individual maps for two subjects in each hemisphere, and (bottom left and right) the average predicted response of each area to the 12 semantic categories identified earlier (responses consistently greater than zero across subjects are marked with “+”). Bars show how completely the 12 category interpretation

captures the average semantic model in each area. Like LPC, MPC appears to be organized around a core group of areas that are selective for *social* and *mental* concepts (L6, 8, 10; R6, 7). Dorsolateral MPC areas (L2, 4; R1) are selective for *visual* and *tactile* concepts. Anterior dorsal areas (L5, 9; R4, 9) are selective for *temporal* concepts. Ventral areas (L11, 12, 14; R8) are selective for *professional*, *temporal*, and *locational* concepts. Just above retrosplenial cortex one distinct area in each hemisphere is selective for *mental*, *professional* and *temporal* concepts (L7; R3). Overall, right MPC responds more than left MPC to *mental* concepts.




Extended Data Figure 8. Semantic atlas for superior prefrontal cortex (SPFC)

The PrAGMATiC atlas divides SPFC into 18 areas in the left hemisphere and 19 areas in the right. Here we show (top left and right) the atlas for each hemisphere, (top middle) 3-D brains indicating the location of SPFC, (bottom middle) individual maps for two subjects in each hemisphere, and (bottom left and right) the average response of each area in the atlas to the 12 semantic categories identified earlier (responses consistently greater than zero across subjects are marked with “+”). Bars show how completely the 12 category interpretation captures the average semantic model in each area. The organization in SPFC seems to follow the long rostro-caudal sulci and gyri of the dorsal frontal lobe. Posterior-lateral SPFC areas (L4, 6; R6, 9, 11) are selective for *social*, *emotional*, *communal*, and *violent* concepts. Posterior superior frontal sulcus areas (L2, 3, 7, 8; R1, 5, 7) are selective for *visual*, *tactile*, and *numeric* concepts. Superior frontal gyrus contains a long strip of areas (L1, 5, 10, 12–15; R8, 12, 14–16) selective for *social*, *emotional*, *communal*, and *violent* concepts.



Extended Data Figure 9. Semantic atlas for lateral temporal cortex (LTC)


The PrAGMATiC atlas divides LTC into 8 areas in both the left and right hemispheres. Here we show (top left and right) the atlas for each hemisphere, (top middle) 3-D brains indicating the location of LTC, (bottom middle) individual maps for two subjects in each hemisphere, and (bottom left and right) the average response of each area in the atlas to the 12 semantic categories identified earlier (responses consistently greater than zero across subjects are marked with “+”). Bars show how completely the 12 category interpretation captures the average semantic model in each area. Anterior LTC areas (L4-8; R3-8) are selective for *social*, *emotional*, *mental*, and *violent* concepts. Posterior LTC areas (L1-3; R1-2) are selective for *numeric*, *tactile*, and *visual* concepts.



Extended Data Figure 10. Semantic atlas for ventral temporal cortex (VTC)


The PrAGMATiC atlas divides VTC into 6 areas in the left hemisphere and 1 area in the right. Here we show (top left and right) the atlas for each hemisphere, (top middle) 3-D brains indicating the location of VTC, (bottom middle) individual maps for two subjects in each hemisphere, and (bottom left and right) the average response of each area in the atlas to the 12 semantic categories identified earlier (responses consistently greater than zero across subjects are marked with “+”). Bars show how completely the 12 category interpretation captures the average semantic model in each area. VTC is relatively homogeneous: all areas

are selective for *numeric*, *tactile*, and *visual* concepts. In left VTC areas close to the parahippocampal place area (PPA) are also selective for *locational* concepts (L5-6).



Extended Data Figure 11. Semantic atlas for inferior prefrontal cortex (IPFC)

The PrAGMATiC atlas divides IPFC into 12 areas in the left hemisphere and 9 areas in the right. Here we show (top left and right) the atlas for each hemisphere, (top middle) 3-D brains indicating the location of IPFC, (bottom middle) individual maps for two subjects in each hemisphere, and (bottom left and right) the average response of each area in the atlas to the 12 semantic categories identified earlier (responses consistently greater than zero across subjects are marked with “+”). Bars show how completely the 12 category interpretation captures the average semantic model in each area. Posterior IPFC areas in the precentral sulcus (L1-3; R1, 2) are selective for *visual*, *tactile*, and *numeric* concepts. Areas on the inferior frontal gyrus (L8; R4, 7) are selective for *social* and *violent* concepts. Areas in the inferior frontal sulcus and anterior middle frontal gyrus (L4-7; R5-6) are selective for *visual*, *tactile*, and *numeric* concepts. Areas in the orbitofrontal sulci (L10; R9) are also selective for *visual*, *tactile*, *numeric*, and *locational* concepts.



Extended Data Figure 12. Semantic atlas for opercular and insular cortex (OIC)

The PRAGMATiC atlas divides OIC into 4 areas in the left hemisphere and 3 areas in the right. Here we show (top left and right) the atlas for each hemisphere, (top middle) 3-D brains indicating the location of OIC, (bottom middle) individual maps for two subjects in each hemisphere, and (bottom left and right) the average response of each area in the atlas to the 12 semantic categories identified earlier (responses consistently greater than zero across subjects are marked with “+”). Bars show how completely the 12 category interpretation captures the average semantic model in each area. These areas are homogeneously selective for *abstract* concepts, with more posterior and superior areas also responding to *emotional*, *communal*, and *mental* concepts.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work was supported by grants from the National Science Foundation (IIS1208203), the National Eye Institute (EY019684), and from the Center for Science of Information (CSoI), an NSF Science and Technology Center, under grant agreement CCF-0939370. A.G.H. was also supported by the William Orr Dingwall Neurolinguistics Fellowship. We thank J. Sohl-Dickstein and K. Crane for technical discussions about PRAGMATiC, J. Nguyen for assistance transcribing and aligning stimuli, B. Griffin for segmenting and flattening cortical surfaces, and N. Bilenko, J. Gao, M. Lescroart, and A. Nunez-Elizalde for general comments and discussions.

References

1. Binder JR, Desai RH, Graves WW, Conant LL. Where is the semantic system? A critical review and meta-analysis of 120 functional neuroimaging studies. *Cereb Cortex*. 2009; 19:2767–96. [PubMed: 19329570]
2. Lerner Y, Honey CJ, Silbert LJ, Hasson U. Topographic mapping of a hierarchy of temporal receptive windows using a narrated story. *J Neurosci*. 2011; 31:2906–15. [PubMed: 21414912]
3. Friederici AD, Opitz B, von Cramon DY. Segregating semantic and syntactic aspects of processing in the human brain: an fMRI investigation of different word types. *Cereb Cortex*. 2000; 10:698–705. [PubMed: 10906316]
4. Noppeney U, Price CJ. Retrieval of abstract semantics. *Neuroimage*. 2004; 22:164–70. [PubMed: 15110006]
5. Binder JR, Westbury CF, McKiernan KA, Possing ET, Medler DA. Distinct brain systems for processing concrete and abstract concepts. *J Cogn Neurosci*. 2005; 17:905–917. [PubMed: 16021798]
6. Bedny M, Caramazza A, Grossman E, Pascual-Leone A, Saxe R. Concepts are more than percepts: the case of action verbs. *J Neurosci*. 2008; 28:11347–53. [PubMed: 18971476]
7. Saxe R, Kanwisher N. People thinking about thinking people The role of the temporo-parietal junction in ‘theory of mind’. *Neuroimage*. 2003; 19:1835–1842. [PubMed: 12948738]
8. Caramazza A, Shelton JR. Domain-specific knowledge systems in the brain the animate-inanimate distinction. *J Cogn Neurosci*. 1998; 10:1–34. [PubMed: 9526080]
9. Mummery CJ, Patterson K, Hodges JR, Price CJ. Functional neuroanatomy of the semantic system: divisible by what? *J Cogn Neurosci*. 1998; 10:766–77. [PubMed: 9831743]
10. Just MA, Cherkassky VL, Aryal S, Mitchell TM. A neurosemantic theory of concrete noun representation based on the underlying brain codes. *PLoS One*. 2010; 5:e8622. [PubMed: 20084104]
11. Warrington EK. The selective impairment of semantic memory. *Q J Exp Psychol*. 1975; 27:635–657. [PubMed: 1197619]

12. Mitchell TM, et al. Predicting human brain activity associated with the meanings of nouns. *Science*. 2008; 320:1191–5. [PubMed: 18511683]
13. Damasio H, Grabowski TJ, Tranel D, Hichwa RD, Damasio AR. A neural basis for lexical retrieval. *Nature*. 1996; 380:499–505. [PubMed: 8606767]
14. 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; 76:1210–24. [PubMed: 23259955]
15. Wehbe L, et al. Simultaneously Uncovering the Patterns of Brain Regions Involved in Different Story Reading Subprocesses. *PLoS One*. 2014; 9:e112575. [PubMed: 25426840]
16. Naselaris T, Prenger RJ, Kay KN, Oliver M, Gallant JL. Bayesian reconstruction of natural images from human brain activity. *Neuron*. 2009; 63:902–15. [PubMed: 19778517]
17. Nishimoto S, et al. Reconstructing visual experiences from brain activity evoked by natural movies. *Curr Biol*. 2011; 21:1641–6. [PubMed: 21945275]
18. Deerwester S, Dumais ST, Furnas GW, Landauer TK, Harshman R. Indexing by Latent Semantic Analysis. *J Am Soc Inf Sci*. 1990; 41:391–407.
19. Lund K, Burgess C. Producing high-dimensional semantic spaces from lexical co-occurrence. *Behav Res Methods, Instruments, Comput*. 1996; 28:203–208.
20. Turney PD, Pantel P. From frequency to meaning: Vector space models of semantics. *J Artif Intell Res*. 2010; 37:141–188.
21. Caramazza A, Mahon BZ. The organisation of conceptual knowledge in the brain: The future’s past and some future directions. *Cogn Neuropsychol*. 2006; 23:13–38. [PubMed: 21049320]
22. Huth, AG.; Griffiths, TL.; Theunissen, FE.; Gallant, JL. PrAGMATiC: a Probabilistic and Generative Model of Areas Tiling the Cortex. 2015. p. 1-15.at <<http://arxiv.org/abs/1504.03622>>
23. Amunts K, Malikovic a, Mohlberg H, Schormann T, Zilles K. Brodmann’s areas 17 and 18 brought into stereotaxic space-where and how variable? *Neuroimage*. 2000; 11:66–84. [PubMed: 10686118]
24. Fedorenko E, Hsieh PJ, Nieto-Castañón A, Whitfield-Gabrieli S, Kanwisher N. New method for fMRI investigations of language: defining ROIs functionally in individual subjects. *J Neurophysiol*. 2010; 104:1177–94. [PubMed: 20410363]
25. Hinton GE. Training Products of Experts by Minimizing Contrastive Divergence. *Neural Comput*. 2002; 14:1771–1800. [PubMed: 12180402]
26. Buckner RL, Andrews-Hanna JR, Schacter DL. The brain’s default network: anatomy, function, and relevance to disease. *Ann N Y Acad Sci*. 2008; 1124:1–38. [PubMed: 18400922]
27. DeWitt I, Rauschecker JP. Phoneme and word recognition in the auditory ventral stream. *Proc Natl Acad Sci*. 2012; 109:E505–E514. [PubMed: 22308358]
28. Riesenhuber M. Appearance isn’t everything: news on object representation in cortex. *Neuron*. 2007; 55:341–4. [PubMed: 17678848]
29. Dehaene S, Cohen L, Sigman M, Vinckier F. The neural code for written words: A proposal. *Trends Cogn Sci*. 2005; 9:335–341. [PubMed: 15951224]
30. Op de Beeck HP, Haushofer J, Kanwisher NG. Interpreting fMRI data: maps, modules and dimensions. *Nat Rev Neurosci*. 2008; 9:123–35. [PubMed: 18200027]
31. Caspers S, et al. Organization of the Human Inferior Parietal Lobule Based on Receptor Architectonics. *Cereb Cortex*. 2013; 23:615–628. [PubMed: 22375016]
32. Cohen AL, et al. Defining functional areas in individual human brains using resting functional connectivity MRI. *Neuroimage*. 2008; 41:45–57. [PubMed: 18367410]
33. Yuan J, Liberman M. Speaker identification on the SCOTUS corpus. *Proc Acoust*. 2008
34. Boersma P, Weenink D. Praat: doing phonetics by computer. 2014
35. Benjamini, Y.; Hochberg, Y. Controlling the False Discovery Rate: A practical and powerful approach to multiple testing. *J R Stat Soc Ser B* 1995. at <<http://www.jstor.org/stable/2346101>>
36. Oliphant, TE. Guide to NumPy. Brigham Young University; 2006.
37. Jones, E.; Oliphant, TE.; Peterson, P. SciPy: Open source scientific tools for Python. 2001.

38. Gao JS, Huth AG, Lescroart MD, Gallant JL. Pycortex: an interactive surface visualizer for fMRI. *Front Neuroinform.* 2015; 9:1–12. [PubMed: 25750622]

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript




Figure 1. Voxel-wise modeling

(a) Seven subjects listened to over two hours of naturally spoken narrative stories while BOLD responses were measured using fMRI. Each word in the stories was projected into a 985-dimensional word embedding space constructed using word co-occurrence statistics from a large corpus of text. A finite impulse response (FIR) regression model was estimated individually for every voxel. The voxel-wise model weights describe how words appearing in the stories influence BOLD signals. (b) Models were tested using one 10-minute story that was not included during model estimation. Model prediction performance was computed as the correlation between predicted responses to this story and actual BOLD responses. (c) Prediction performance of voxel-wise models for one subject. Semantic models accurately predict BOLD responses in many brain areas, including the lateral and ventral temporal cortex (LTC, VTC), lateral and medial parietal cortex (LPC, MPC), and superior and inferior prefrontal cortex (SPFC, IPFC). These regions have previously been identified as the “semantic system” in the human brain.




Figure 2. Principal components of voxel-wise semantic models

Principal components analysis (PCA) of voxel-wise model weights reveals four important semantic dimensions in the brain (Extended Data Fig. 2). **(a)** An RGB colormap was used to color both words and voxels based on the first three dimensions of the semantic space. Words that best match the four semantic dimensions were found and then collapsed into 12 categories using *k*-means clustering. Each category (Supplementary Table 2) was manually assigned a label. The 12 category labels (large words) and a selection of the 458 best words (small words) are plotted here along four pairs of semantic dimensions. The largest axis of variation lies roughly along the first dimension, and separates perceptual and physical categories (*tactile*, *locational*) from human-related categories (*social*, *emotional*, *violent*). **(b)** Voxel-wise model weights were projected onto the semantic dimensions and then colored using the same RGB colormap (see Extended Data Fig. 3 for separate dimensions). Projections for one subject (S2) are shown on that subject's cortical surface. Semantic information seems to be represented in intricate patterns across much of the semantic system. **(c)** Semantic PC flatmaps for three other subjects. Comparing these flatmaps, many patterns appear to be shared across individuals. (See Extended Data Fig. 3 for other subjects.)

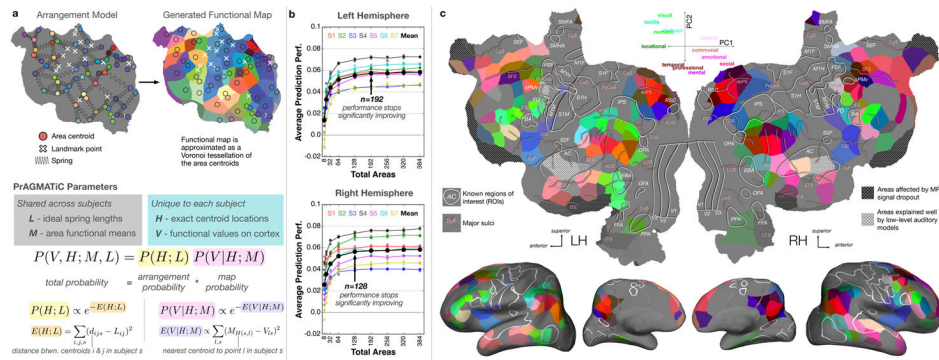


Figure 3. PrAGMATiC: a generative model for cortical maps

To create an atlas that describes the distribution of semantically selective functional areas in the human cerebral cortex we developed PrAGMATiC, a probabilistic and generative model of areas tiling the cortex. **(a)** PrAGMATiC has two parts: an arrangement model and an emission model. The arrangement model is analogous to a physical system of springs joining neighboring area centroids. To enforce similarity across subjects, springs also join areas to 19 regions-of-interest that were localized separately. The emission model assigns the functional mean of the closest area centroid to each point on the cortex, forming a Voronoi tessellation. Spring lengths and area means are shared across subjects while exact area locations are unique to each subject. These parameters are fit using maximum likelihood estimation. **(b)** A leave-one-out procedure was used to choose the number of areas in each hemisphere. PrAGMATiC models were estimated on six subjects and then used to predict BOLD responses for the seventh. Prediction performance improved significantly up to 192 total areas in the left hemisphere and 128 areas in the right. **(c)** A semantic atlas was estimated using data from all seven subjects. Areas where the semantic model did not predict better than models based on low-level features (i.e. word rate, phonemes) were removed. The remaining areas were plotted on one subject's cortical surface using the same RGB colormap as Figure 2. Areas dominated by signal dropout are shown in black hatching, and areas where the low-level models performed well are shown in white hatching. This atlas shows the functional organization of the semantic system that is common across subjects.