
The Compositional Nature of Verb and Argument Representations in the Human Brain

Andrei Barbu* N. Siddharth* Caiming Xiong† Jason J. Corso†
Christiane D. Fellbaum‡ Catherine Hanson§ Stephen José Hanson§ Sébastien Hélie¶
Evguenia Malaia|| Barak A. Pearlmuter** Jeffrey Mark Siskind*
Thomas Michael Talavage* Ronnie B. Wilbur††

Abstract

How does the human brain represent simple compositions of objects, actors, and actions? We had subjects view action sequence videos during neuroimaging (fMRI) sessions and identified lexical descriptions of those videos by decoding (SVM) the brain representations based only on their fMRI activation patterns. As a precursor to this result, we had demonstrated that we could reliably and with high probability decode action labels corresponding to one of six action videos (*dig, walk, etc.*), again while subjects viewed the action sequence during scanning (fMRI). This result was replicated at two different brain imaging sites with common protocols but different subjects, showing common brain areas, including areas known for episodic memory (PHG, MTL, high level visual pathways, *etc.*, *i.e.*, the ‘what’ and ‘where’ systems, and TPJ, *i.e.*, ‘theory of mind’). Given these results, we were also able to successfully show a key aspect of language compositionality based on simultaneous decoding of object class and actor identity. Finally, combining these novel steps in ‘brain reading’ allowed us to accurately estimate brain representations supporting compositional decoding of a complex event composed of an actor, a verb, a direction, and an object.

1 Introduction

The compositional nature of thought is taken for granted by many in the cognitive-science and artificial-intelligence communities. For example, in computer vision, representations for nouns, such as those used for object detection, are independent of representations for verbs, such as those used for event recognition. Humans need not employ compositional representations; indeed, many argue that such representations may be doomed to failure in AI systems (Brooks, 1991). This is because concepts like *verb* or even *object* are human constructs; there is debate as to how they arise from percepts (Smith, 1996). Recent advances in brain-imaging techniques enable exploration of the compositional nature of thought. To that end, subjects underwent functional magnetic resonance imaging (fMRI) during which they were exposed to stimuli which evoke complex brain activity which was decoded, piece by piece. The video stimuli depicted events described by entire sentences composed of a *verb*, an *object*, an *actor* and a *location* or *direction* of motion. By decoding complex

*School of Electrical and Computer Engineering, Purdue University, West Lafayette IN 47907-2035

†Department of Computer Science and Engineering, SUNY Buffalo, Buffalo NY 14260-2500

‡Department of Computer Science, Princeton University, Princeton NJ 08540-5233

§Department of Psychology and Rutgers Brain Imaging Center, Rutgers University, Newark NJ 07102

¶Psychological Sciences, Purdue University, West Lafayette IN 47907

||Southwest Center for Mind, Brain, and Education, University of Texas at Arlington, Arlington TX 76019

**Hamilton Institute & Dept Computer Sci, National University of Ireland Maynooth, Co. Kildare, Ireland

††Department of Speech, Language, and Hearing Sciences and Linguistics Program, Purdue University, West Lafayette IN 47907

brain activity into its constituent parts, we show evidence for the neural basis of the compositionality of verb and argument representations.

Recent work on decoding brain activity corresponding to nouns has recovered object identity from nouns presented as image and orthographic stimuli. Hanson and Halchenko (2009) perform classification on still images of two object classes: faces and houses, and achieve an accuracy above 93% on a one-out-of-two classification task. Connolly et al. (2012) perform classification on still images of objects, two instances of each of three classes: bugs, birds, and primates, and achieve an accuracy between 60% and 98% on a one-out-of-two within-class classification task and an accuracy between 90% and 98% on a one-out-of-three between-class classification task. Just et al. (2010) perform classification on orthographically presented nouns, 5 exemplars from each of 12 classes, achieving a mean rank accuracy of 72.4% on a one-out-of-60 classification task, both within and between subjects. Pereira et al. (2012) incorporate semantic priors and achieve a mean accuracy of 13.2% on a one-out-of-12 classification task and 1.94% on a one-out-of-60 classification task when attempting to recover the object being observed. Miyawaki et al. (2008) recover the position of an object in the field of view by recovering low resolution images from the visual cortex. Object classification from video stimuli has not been previously demonstrated.

Recent work on decoding brain activity corresponding to verbs has primarily been concerned with identifying active brain regions. Kable and Chatterjee (2006) present the brain regions which attempt to distinguish between the different agents of actions and between the different kinds of actions they perform. Kemmerer et al. (2008) analyze the regions of interest (ROI) of brain activity associated with orthographic presentation of twenty different verbs in each of five different verb classes. Kemmerer and Gonzalez Castillo (2010) analyze the brain activity associated with verbs in terms of the motor components of event structure and attempt to localize the ROIs of such motor components. While prior work analyzes regions which are activated when subjects are presented verbs as stimuli, we recover the content of the resulting brain activity by classifying the verb from brain scans.

Recent work demonstrates the ability to decode the actor of an event using personality traits. Hassabis et al. (2013) demonstrate the ability to recover the identity of an imagined actor from that actor's personality. Subjects are informed of the two distinguishing binary personality traits of four actors. During fMRI, they are presented sentences orthographically which describe an actor performing an action. The subjects are asked to imagine this scenario with this actor and to rate whether the actions of the actor accurately reflect the personality of that actor. The resulting brain activation corresponding to these two binary personality traits is used to recover the identity of the actor. No prior work has recovered the identity of an actor without relying on that actor's personality. In the work presented here, the personality of the actor has no bearing on the actions being performed.

In this paper, two new experiments are presented. In Experiment 1, subjects are shown videos and asked to think of verbs that characterize those videos. Their brains are imaged via fMRI and measured neural activation is decoded to recover the verb that the subjects are thinking about. Decoding is done by means of a support vector machine (SVM) trained on brain scans of those same verbs. We know of no other work that decodes brain activity corresponding to verbs. We show early evidence that the regions identified by this decoding process are not intimately tied to a particular subject *via* an additional analysis that trains on one subject and tests on another. In Experiment 2, subjects are shown videos and asked to think of complex sentences composed of multiple components that characterize those videos. We show a novel ability to decode brain activity corresponding to multiple objects: the identity of an actor and the identity of an object. We decode the identity of an actor without relying on the personality traits of that actor. We know of no other work which recovers an entire sentence composed of multiple constituents. We find evidence that suggests underlying neural representations of mental states are independent and compose into sentences largely without modifying one another.

2 Compositionality

We discuss a particular kind of compositionality as it applies to sentence structure: objects fill argument positions in predicates that combine to form the meaning of a sentence. Pylkkänen et al. (2011) reviews work which attempts to show this kind of compositionality using a task called *complement coercion*. Subjects in this task are presented with sentences whose meaning is richer than their syntax. For example, the sentence *The boy finished the pizza* is understood as meaning that the pizza was eaten, even though the verb *eat* does not appear anywhere in the sentence (Pustejovsky, 1995). The presence of *pizza*, belonging to the category *food*, coerces the interpretation of *finish* as *finish*

eating. By contrast, *He finished the newspaper* induces the interpretation *finish reading*. Because the syntactic complexity in this prior experiment was held constant, the assumption is that coercion is a purely semantic meaning-adding function application, with little consequence for the syntax. The participants completed this task, and brain activity was measured using magnetoencephalography (MEG). The results show activity related to coercion in the anterior midline field. This result suggests an initial localization for at least some function application, but it is difficult to use MEG to distinguish whether this activity is read from the ventromedial prefrontal cortex or the anterior cingulate cortex. Earlier work on the representation of objects and actions in the brain also indicates that these representations may be independent.

Representing objects in the brain Objects are static entities that can be represented by a (mostly) static neural representation. For example, the 3D representation of a soda can will look the same in many different contexts, and the appearance of the soda can is not unfolding in time. It is generally believed that the lexicon of object concepts is represented in the medial temporal lobe while different areas of the temporal lobe may be combinatoric in constructing object types (Hanson et al., 2004) although there may be modal areas associated with different representational functions. For example, lesion data suggests that the temporal pole is associated with naming people, the inferior temporal cortex is associated with naming animals, and the anterior lateral occipital regions are associated with naming tools. In addition, some regions involved in object representation are modality specific. For example, spoken-word processing involves the superior temporal lobe (part of the auditory associative cortex; Binder et al., 2000) while reading words representing objects activates occipito-temporal regions because of the visual processing (Puce et al., 1996). Specifically, auditory word processing involves a stream of information starting in Heschl’s gyri that is transferred to the superior temporal gyrus. Once the superior temporal gyrus has been reached, the modality of stimulus presentation is no longer relevant. In contrast, the initial processing for written words starts in the occipital lobe (V1 and V2), and moves on to occipito-temporal regions specialized in identifying orthographic units. The information then moves rostrally to the temporal lobe proper, where modality of presentation is no longer relevant (Binder et al., 2000).

Representing actions in the brain Unlike objects, verbs are dynamic entities that unfold in time. For instance, observing someone pick up a ball takes time as the person’s movement unfolds. Evidence reviewed in Coello and Bidet-Ildei (2012) suggests that action verbs activate both semantic units in the temporal cortex and a motor network. The motor network includes the premotor areas (including the supplementary motor area), the primary motor cortex, and the posterior parietal cortex. Some researchers went as far as suggesting that the well-known ventral/dorsal distinction in the visual pathways corresponds to a semantic (ventral) and action (dorsal) distinction. Representation of action may involve ‘mirror neurons’ that have been shown in macaque to respond jointly in perception/action tasks, where the similarity of the self action is to the perceived action of an observed individual.

3 Approach

All experiments reported follow the same procedure and are analyzed using the same methods and classifiers. Videos are shown to subjects who are asked to think about some aspect(s) of the video while whole-brain fMRI scans are acquired every two seconds. Because fMRI acquisition times are slow, roughly equal to the length of the video stimuli, a single brain volume that corresponds to the brain activation induced by that video stimulus is classified to recover the features that the subjects were asked to think about. Multiple runs separated by several minutes of rest, where no data is acquired, are performed per subject.

Imaging performed at Purdue University used a 3T GE Signa HDx scanner (Waukesha, Wisconsin) with a Nova Medical (Wilmington, Massachusetts) 16 channel brain array to collect whole-brain volumes via a gradient-echo EPI sequence with 2000ms TR, 22ms TE, 200mm×200mm FOV, and 77° flip angle. We acquired 35 axial slices with a 3.000mm slice thickness using a 64×64 acquisition matrix resulting in 3.125mm×3.125mm×3.000mm voxels.

Imaging performed at St. James Hospital in Dublin, Ireland, used a 3T Phillips Achieva scanner (Best, The Netherlands) using a gradient-echo EPI sequence with 2000ms TR and 240mm×240mm FOV. We acquired 37 axial slices with a 3.550mm slice thickness using an 80×80 acquisition matrix resulting in 3.000mm×3.000mm×3.550mm voxels.

Data was acquired in runs, with between three and eight runs per subject per experiment, and each axis of variation of each experiment was counterbalanced within each run. fMRI scans were pro-

cessed using AFNI (Cox et al., 1996) to skull-strip each volume, motion correct and detrend each run, and align each subject’s runs to each other. Voxels within a run were z-scored, subtracting the mean value of that voxel for the run and dividing by its variance. Because each brain volume has very high dimension, between 143,360 and 236,800 voxels, we eliminate voxels by computing a per-voxel Fisher score on our training set and keeping the 5,000 highest-scoring voxels. The Fisher score of a voxel v for a classification task with C classes where each class c has n_c examples is computed as

$$\frac{\sum_{c=1}^C n_c (\mu_{c,v} - \mu)^2}{\sum_{c=1}^C n_c \sigma_{c,v}^2} \quad (1)$$

where $\mu_{c,v}$ and $\sigma_{c,v}$ are the per-class per-voxel means and variances and μ is the mean for the entire brain volume. A linear SVM classifies the selected voxels.

One run was taken as the test set and the remaining runs were taken as the training set. The third brain volume after the onset of each stimulus was taken along with the class of the stimulus to train an SVM. This lag of three brain volumes is required because fMRI does not measure neural activation but instead measures the flow of oxygenated blood, the blood-oxygen-level-dependent (BOLD) signal, which correlates with increased neural activation. It takes roughly five to six seconds for this signal to peak which puts the peak in the third volume after the stimulus presentation. Cross validation was performed by choosing each of the different runs as the test set.

To understand our results and to demonstrate that they are not classifying noise or irrelevant features, we perform an analysis to understand the brain regions that are relevant to each experiment. We determine these regions by two methods. First we employ a spatial searchlight (Kriegeskorte et al., 2006) which slides a small sphere across the entire brain volume and repeats the above analysis keeping only the voxels inside that sphere. We use a sphere of radius three voxels, densely place its center at every voxel, and do not perform any dimensionality reduction on the remaining voxels. We then perform an eight-fold cross validation as described above for each position of the sphere. For Experiment 1 we also back-project the SVM coefficients onto the anatomical scans—the higher the absolute value of the coefficient the more that voxel contributes to the classification performance of the SVM—and use a classifier with a different metric, $w(i)^2$, as described by Hanson and Halchenko (2009).

4 Experiment 1: Verb Representation

We conducted an experiment to evaluate the ability to identify brain activity corresponding to verbs denoting actions. Subjects are shown video clips of humans interacting with objects and are told to think of the verb being enacted, but otherwise have no task. The subjects were shown clips depicting each of these verbs prior to the experiment and were instructed about the intended meaning of each verb. One difficulty with such an experiment is that there is disagreement between human subjects as to whether a verb occurred in a video or not. To overcome this difficulty, we asked five humans to annotate the DARPA Mind’s Eye year 2 video corpus with the extent of every verb. From this corpus, we chose video clips where at least two out of the five annotators agreed on the depiction. We selected between twenty seven and thirty 2.5s video clips depicting each of six different verbs (*carry*, *dig*, *hold*, *pick up*, *put down*, and *walk*). Key frames from one clip for each of the six verbs are shown in Fig. 1. Despite multiple annotators agreeing on whether a video depicts a verb, the task of classifying each clip remains very difficult for human subjects as it is easy to confuse similar verbs such as *carry* and *hold*. We address this problem by presenting, in rapid succession, pairs of video clips which depict the same verb and asking the subjects to think about the verb that would best describe both videos.

We employed a rapid event-related design similar to that of Just et al. (2010). We presented pairs of 2.5s video clips at 12fps, depicting the same verb, separated by 0.5s blanking and followed by an average of 4.5s (minimum 2.5s) fixation. While the video clips within each pair depicted the same verb, the clips across pairs within a run depicted different verbs, randomly counterbalanced. Each run comprised 48 stimulus presentations spanning 254 captured brain volumes and ended with 24s of fixation. Eight runs for each of subjects 1 through 3 were collected at Purdue University. Three runs for subject 4 and four runs for subject 5 were collected at St. James Hospital.



Figure 1: Key frames from sample stimuli for each of the six verbs in Experiment 1. Example stimulus videos are included in the supplementary material.




Figure 2: Results for Experiment 1. (left) Per-subject classification accuracy on 1-out-of-6 verb classes averaged across class and fold. Horizontal line indicates chance performance, 16.66%. (right) Corresponding confusion matrix averaged across subject and fold is mostly diagonal, with the highest numbers of errors being made distinguishing *hold* and *carry*, two ambiguous stimuli.

We performed an eight-fold cross validation (fewer for subjects 4 and 5) for a six-way classification task, where runs constituted folds. The results are presented in Fig. 2. The per-subject accuracies, averaged across class and fold, were: 80.73%, 87.24%, 78.91%, 35.94%, and 43.75% (chance 16.66%). Note that the last two were trained on fewer runs than the first three. This demonstrates the ability to recover the verb that the subjects were thinking about. The robustness of this result is enhanced by the fact that it was replicated on two different fMRI scanners at different locations run by different experimenters.

To evaluate whether the brain regions used for classification generalize across subjects, we performed an additional analysis on the data for subjects 1 and 2. One run out of the eight was selected as the test set and the data for one of the two subjects was classified. The training set consisted of all seven other runs for the subject whose data does not appear in the test set. The test was performed on the run omitted from the training set, even though it was gathered from a different subject, to preclude the possibility that the same stimulus sequence appeared in both the training and test sets. We performed cross validation by varying which subject contributes the test data and which subject contributes the training data, and within each of these folds we varied which of the eight runs is the test set. These two cross validations yielded accuracies of 33.59% (subject 1→subject 2) and 41.41% (subject 2→subject 1), averaged across class and fold, where chance again is 16.66%.

To locate regions of the brain used in the previous analysis, we used a spatial-searchlight linear-SVM method on subject 1. We use the accuracy to determine the sensitivity of each voxel and threshold upward to less than 5% of the cross-validation measures. These measures are overlaid and (2-stage) registered to MNI152 2mm anatomicals shown in Fig. 3(left). Notable are visual-pathway areas (lateral occipital-LO, lingual gyrus-LG, and fusiform gyrus) as well as prefrontal areas (inferior frontal gyrus, middle frontal gyrus, and cingulate) and areas consistent with the ‘mirror system’ (Arbib, 2006) and the so-called ‘theory of mind’ (pre-central gyrus, angular gyrus-AG, and superior parietal lobule-SPL) areas (Dronkers et al., 2004; Turken and Dronkers, 2011). Fig. 3(right) shows the




Figure 3: (left) Searchlight analysis for Experiment 1 indicating the classification accuracy of different brain regions on the anatomical scans from subject 1, averaged across stimulus, class, and run. (right) A similar analysis using a $w(i)^2$ metric.

decoded ROIs from a similar SVM classifier with a different metric, $w(i)^2$ (Hanson and Halchenko, 2009), showing similar brain areas but, due to higher sensitivity, also indicates sub-cortical regions (hippocampal) associated with encoding processes not seen with the cross-validation accuracy metric. As argued in Section 2, lateral-occipital areas are involved in visual processing specifically related to language, and the fusiform gyrus is a hetero-modal area that could hold abstract representations of the elements contained in the videos (*e.g.*, semantics). This data brings initial support for the hypothesis that concepts have both modality-specific and abstract representations. Hence, the elements used by the SVM to classify the videos are also neuroscientifically meaningful.

5 Experiment 2: Argument Representation

We conducted a further experiment to evaluate the ability to recover compositional semantics for entire sentences. Subjects were shown videos that depict sentences of the form: the *actor verb the object direction/location*. They were asked to think about the sentence depicted in each video and otherwise had no task. Videos depicting three verbs (*carry*, *fold*, and *leave*), each performed with three objects (*chair*, *shirt*, and *tortilla*), each performed by four human actors, and each performed on either side of the field of view were filmed for this task. The verbs were chosen to be discriminable based on features described by Kemmerer et al. (2008):

<i>leave</i>	–state-change	–contact
<i>fold</i>	+state-change	+contact
<i>carry</i>	–state-change	+contact

Nouns were chosen based on categories found to be easily discriminable by Just et al. (2010): *chair* (furniture), *shirt* (clothing), and *tortilla* (food) and also selected to allow each verb to be performed with each noun. Because these stimuli are not as ambiguous as the ones from Experiment 1, they were not shown in pairs. All stimuli enactments were filmed against the same nonvarying background, which contained no other objects except for a table (Fig 4).

This experiment, like Experiment 1, also used a rapid event-related design. We collected multiple videos, between 4 and 7, for each cross product of the verb, object and human actor. Variation along the side of field of view and direction of motion was accomplished by mirroring the videos about the vertical axis. Such mirroring induces variation in direction of motion (leftward *vs.* rightward) for the verbs *carry* and *leave* and induces variation in the location in the field of view where the verb *fold* occurs (left half *vs.* right half). We presented 2s video clips at 10fps followed by an average of 4s (minimum 2s) fixation. Each run comprised 72 stimulus presentations spanning 244 captured brain volumes, with eight runs per subject, and ended with 24s of fixation. Each run was individually counterbalanced for each of the four conditions (verb, object, actor, and mirroring). We collected data for three subjects at Purdue University but discarded the data for one of the three due to subject motion. One subject did eight runs without exiting the scanner. One subject exited the scanner between runs six and seven, which required cross-session registration. All subjects were aware of the experiment design, were informed of the intended depiction of each stimulus prior to the scan, and were instructed to think of the intended depiction after each presentation.

This experimental design supports the following classification analyses:

- event** one-out-of-9 verb&noun (*carry*, *fold*, and *leave*, each performed on *chair*, *shirt*, and *tortilla*)
- verb** one-out-of-3 verb (*carry*, *fold*, and *leave*)
- object** one-out-of-3 noun (*chair*, *shirt*, and *tortilla*)
- actor** one-out-of-4 actor identity
- direction** one-out-of-2 motion direction for *carry* and *leave* (leftward *vs.* rightward)
- location** one-out-of-2 location in the field of view for *fold* (right *vs.* left)



Figure 4: Key frames from sample stimuli in Experiment 2. Example stimulus videos are included in the supplementary material.

The analysis performed was exactly the same as that for Experiment 1, including eight-fold cross validation for each of our analyses, where runs constituted folds. Fig. 5 presents an overview of the results along with per-subject classification accuracies and aggregate confusion matrices for each of the above analyses. Note that we achieve significantly above-chance performance on all six analyses with only a single fold for a single subject across all six analyses performing below chance.

Verb performance is well above chance (76.22%, chance 11.11%). This replicates Experiment 1 with different videos and a new verb and adds to the evidence that brain activity corresponding to verbs can reliably be decoded from fMRI scans. **Object** performance was significant as well (60.42%, chance 33.33%). Given neural activation, we can decode which object the subjects are thinking about. We know of no other work that decodes brain activity corresponding to objects from videos. The fact that the verb and object can be decoded independently already provides evidence of argument compositionality. Were the neural representations not compositional at this level, decoding would not be possible. For example, if the representation of *carry* was neurally encoded as a combination of *walk* and a particular object, verb performance would not exceed chance, because our experiment is counterbalanced with respect to the object with which the action is being performed. While this indicates that the representations for verbs and objects are independent of each other to some degree, we also seek to quantify the level of independence. If the representation of *carry* is somewhat different depending on which object is being carried, we expect that performance would increase when we jointly classify the object and the verb. This seems to not be the case. The accuracy of **event** is almost identical to the joint independent accuracy of **verb** and **object**: $0.5538 \approx 0.5289 = 0.8212 \times 0.6441$ (subject 1) and $0.4097 \approx 0.3967 = 0.7031 \times 0.5642$ (subject 2), indicating that the representation of these verbs is independent of the objects that the verbs are being performed with. This is also confirmed by the confusion matrix for **event** in Fig. 5(c) which remains diagonal.

To decode complex brain activity corresponding to an entire sentence, we can combine **actor**, **verb**, **object**, and **direction** or **location**. We perform significantly above chance on this one-out-of-72 ($4 \times 3 \times 3 \times 2$) classification:

$$0.3281 \times 0.8212 \times 0.6441 \times \left(\frac{1}{3} \times 0.6823 + \frac{2}{3} \times 0.8333\right) = \mathbf{0.1359} \gg \mathbf{0.0139} = \frac{1}{72} \text{ (subject 1)}$$

$$0.3281 \times 0.7031 \times 0.5642 \times \left(\frac{1}{3} \times 0.6458 + \frac{2}{3} \times 0.7170\right) = \mathbf{0.0902} \gg \mathbf{0.0139} = \frac{1}{72} \text{ (subject 2)}$$

(Since **direction** applied to *carry* and *leave* while **location** disjointly applied to *fold*, this yields a binary classification task across all verbs.) Thus we are able to classify *entire sentences* compositionally from their individual words.

To locate regions of the brain used in the previous analyses, we applied the same searchlight linear-SVM method that was performed in Experiment 1 to subject 1's data from this experiment and identified similar areas in visual-pathway, parietal, and prefrontal areas. The resulting ROIs, shown in Fig. 6, are overlaid and color coded according to the specific visual feature being decoded. In general, it is clear that the decoding is sensitive to action/category information and various visual object-and-motion features. Many of the same regions active for **verb** in Experiment 1 also show activity in this experiment. **Direction** and **location** activity is present in the visual cortex with significant **location** activity occurring in the early visual cortex. **Object** activity is present in the temporal cortex, and agrees with previous work on object-category encoding (Gazzaniga et al., 2008).




Figure 5: Results for Experiment 2. (a) Per-subject mean classification accuracies averaged across fold. Note that all six analyses perform above chance. (b) Per-subject classification accuracies showing the means and variances of performance across the different folds for each class. The horizontal line indicates chance performance. (c) Corresponding confusion matrices, averaged across subject and fold. Note that they are mostly diagonal.



Figure 6: Searchlight analysis for Experiment 2 indicating the classification accuracy of different brain regions on the anatomical scans from subject 1 averaged across stimulus, class, and run.

6 Conclusion

We have demonstrated that it is possible to read a subject’s brain activity and decode a complex action tableau corresponding to a sentence from its constituents. To do so, we showed novel work which decodes brain activity associated with verbs and simultaneously recovers lexical aspects of different parts of speech. Our results indicate that the neural representations for verbs and objects compose together to form the meaning of a sentence apparently without modifying one another. These results indicate that representations which attempt to decompose meaning into constituents may have a neural basis.

Acknowledgments

AB, NS, and JMS were supported, in part, by Army Research Laboratory (ARL) Cooperative Agreement W911NF-10-2-0060. CX and JJC were supported, in part, by ARL Cooperative Agreement W911NF-10-2-0062 and NSF CAREER grant IIS-0845282. CDF was supported, in part, by NSF grant CNS-0855157. CH and SJH were supported, in part, by the McDonnell Foundation. BAP was supported, in part, by Science Foundation Ireland grant 09/IN.1/I2637. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either express or implied, of the supporting institutions. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes, notwithstanding any copyright notation herein. Dr. Gregory G. Tamer, Jr. provided assistance with imaging and analysis.

References

- M. A. Arbib. *Action to language via the mirror neuron system*. Cambridge University Press, 2006.
- J. R. Binder, J. A. Frost, T. A. Hammeke, P. S. F. Bellgowan, J. A. Springer, J. N. Kaufman, and E. T. Possing. Human temporal lobe activation by speech and nonspeech sounds. *Cerebral Cortex*, 10(5):512–28, 2000.
- R. A. Brooks. Intelligence without representation. *Artificial intelligence*, 47(1):139–59, 1991.
- Y. Coello and C. Bidet-Ildes. Motor representation and language in space, object and movement perception. In Y. Coello and A. Bartolo, editors, *Language and Action in Cognitive Neuroscience*, chapter 4, pages 83–110. Psychology Press, 2012.
- A. C. Connolly, J. S. Guntupalli, J. Gors, M. Hanke, Y. O. Halchenko, Y.-C. Wu, H. Abdi, and J. V. Haxby. The representation of biological classes in the human brain. *The Journal of Neuroscience*, 32(8):2608–18, 2012.
- R. W. Cox et al. AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical Research*, 29(3):162–73, 1996.
- N. F. Dronkers, D. P. Wilkins, R. D. Van Valin, Jr., B. B. Redfern, J. J. Jaeger, et al. Lesion analysis of the brain areas involved in language comprehension. *Cognition*, 92(1-2):145–77, 2004.
- M. S. Gazzaniga, R. B. Ivry, and G. R. Mangun. *Cognitive Neuroscience: The Biology of the Mind*. W. W. Norton & Company, New York, third edition, 2008.
- S. J. Hanson and Y. O. Halchenko. Brain reading using full brain support vector machines for object recognition: There is no “face” identification area. *Neural Computation*, 20(2):486–503, 2009.
- S. J. Hanson, T. Matsuka, and J. V. Haxby. Combinatorial codes in ventral temporal lobe for object recognition: Haxby (2001) revisited: Is there a “face” area? *Neuroimage*, 23(1):156–66, 2004.
- D. Hassabis, R. N. Spreng, A. A. Rusu, C. A. Robbins, R. A. Mar, and D. L. Schacter. Imagine all the people: How the brain creates and uses personality models to predict behavior. *Cerebral Cortex*, 23, 2013.
- M. A. Just, V. L. Cherkassky, S. Aryal, and T. M. Mitchell. A neurosemantic theory of concrete noun representation based on the underlying brain codes. *PLoS One*, 5(1):e8622, 2010.
- J. W. Kable and A. Chatterjee. Specificity of action representations in the lateral occipitotemporal cortex. *Journal of Cognitive Neuroscience*, 18(9):1498–517, 2006.
- D. Kemmerer and J. Gonzalez Castillo. The two-level theory of verb meaning: An approach to integrating the semantics of action with the mirror neuron system. *Brain and Language*, 112(1):54–76, 2010.
- D. Kemmerer, J. Gonzalez Castillo, T. Talavage, S. Patterson, and C. Wiley. Neuroanatomical distribution of five semantic components of verbs: Evidence from fMRI. *Brain and Language*, 107(1):16–43, 2008.
- N. Kriegeskorte, R. Goebel, and P. Bandettini. Information-based functional brain mapping. *Proceedings of the National Academy of Sciences of the United States of America*, 103(10):3863–8, 2006.
- Y. Miyawaki, H. Uchida, O. Yamashita, M. Sato, Y. Morito, H. C. Tanabe, N. Sadato, and Y. Kamitani. Visual image reconstruction from human brain activity using a combination of multiscale local image decoders. *Neuron*, 60(5):915–29, 2008.
- F. Pereira, M. Botvinick, and G. Detre. Using Wikipedia to learn semantic feature representations of concrete concepts in neuroimaging experiments. *Artificial Intelligence*, 194:240–52, 2012.
- A. Puce, T. Allison, M. Asgari, J. C. Gore, and G. McCarthy. Differential sensitivity of human visual cortex to faces, letterstrings, and textures: a functional magnetic resonance imaging study. *The Journal of Neuroscience*, 16(16):5205–15, 1996.
- J. Pustejovsky. *Generative Semantics*. MIT Press, 1995.
- L. Pykkänen, J. Brennan, and D. K. Bemis. Grounding the cognitive neuroscience of semantics in linguistic theory. *Language and Cognitive Processes*, 26(9):1317–37, 2011.
- B. C. Smith. *On the origin of objects*. MIT Press Cambridge, MA, 1996.
- U. Turken and N. F. Dronkers. The neural architecture of the language comprehension network: Converging evidence from lesion and connectivity analyses. *Frontiers in systems neuroscience*, 5(1), 2011.