

Identifying thematic roles from neural representations measured by functional magnetic resonance imaging

Jing Wang^a, Vladimir L. Cherkassky^a, Ying Yang^a, Kai-min Kevin Chang^b, Robert Vargas^a, Nicholas Diana^a and Marcel Adam Just^a

^aCenter for Cognitive Brain Imaging, Department of Psychology, Carnegie Mellon University, Pittsburgh, USA; ^bLanguage Technologies Institute, School of Computer Science, Carnegie Mellon University, Pittsburgh, USA

ABSTRACT

The generativity and complexity of human thought stem in large part from the ability to represent relations among concepts and form propositions. The current study reveals how a given object such as *rabbit* is neurally encoded differently and identifiably depending on whether it is an agent (“the *rabbit* punches the monkey”) or a patient (“the monkey punches the *rabbit*”). Machine-learning classifiers were trained on functional magnetic resonance imaging (fMRI) data evoked by a set of short videos that conveyed *agent-verb-patient* propositions. When tested on a held-out video, the classifiers were able to reliably identify the thematic role of an object from its associated fMRI activation pattern. Moreover, when trained on one subset of the study participants, classifiers reliably identified the thematic roles in the data of a left-out participant (mean accuracy = .66), indicating that the neural representations of thematic roles were common across individuals.

ARTICLE HISTORY

Received 9 September 2015
Revised 22 March 2016
Accepted 19 April 2016

KEYWORDS

Functional magnetic resonance imaging; thematic roles; propositional representation; multivariate pattern analysis

Introduction

One of the mechanisms that contribute to the unbounded representational capacity of human language is the ability to construct multiple propositional configurations using the same conceptual elements, by assigning the concepts to different thematic roles in different propositions (Fodor & Pylyshyn, 1988). A thematic relation indicates the role that a noun or noun phrase plays in a proposition. Agent (the action doer) and patient (the action receiver) roles characterize the most fundamental thematic relations with respect to the verb or action. The question of how thematic roles are neurally coded is central to understanding the neural nature of proposition-level semantics. Some authors suggest that the thematic role representation can be viewed as an instance of relational knowledge (Wu, Waller, & Chatterjee, 2007), whereby the encoding of thematic relations relies on the temporary binding of a concept to a thematic role within an event or its propositional representation.

Most of the sparse knowledge about the neural basis of representing or processing thematic relations

has been drawn from patients with brain damage. Such patients, most of whom present with Broca’s aphasia, have more difficulty in comprehending the thematic relations when both nouns are animate (e.g., *The dog chased the cat.*), and hence both are capable of being an agent than when one of the nouns is inanimate (*The dog chased the ball.* See Berndt, Mitchum, & Haendiges, 1996, for a quantitative review; Saffran, Schwartz, & Linebarger, 1998). The fact that aphasic patients seem to rely on the animacy of individual concepts, rather than syntactic structure of the sentence, to assign thematic roles suggests that there exists a dissociable neural mechanism underlying thematic role assignment.

The deficits in thematic role assignment in aphasic patients have sometimes been attributed to the additional processing demands that thematic processing makes, above and beyond the comprehension of the individual elements of the proposition (Friederici & Frazier, 1992; Miyake, Carpenter, & Just, 1994). That is, the deficit has been attributed to the quantity of the processing rather than to its distinctive quality. Lesions in the lateral temporal cortex are found to be associated with poor performance in a sentence–

CONTACT Jing Wang ✉ jiwang1@andrew.cmu.edu

Supplemental data for this article can be accessed at <http://dx.doi.org/10.1080/02643294.2016.1182480>

© 2016 Informa UK Limited, trading as Taylor & Francis Group

picture matching task that uses reversible propositions (i.e., either noun could be the agent; Wu et al., 2007). In addition, previous functional neuroimaging studies on healthy subjects has shown that the left temporal-parietal junction (Richardson, Thomas, & Price, 2010) and other perisylvian regions (Meltzer, McArdle, Schafer, & Braun, 2010) respond more strongly to the reversible than to the non-reversible sentences. These findings have suggested the specialized and dissociable neural coding for thematic role information.

Very few functional neuroimaging studies have studied how thematic roles are represented or computed in the brain (but see Frankland & Greene, 2015). One obstacle is that the thematic roles of concepts are determined by the contexts and are computed dynamically as needed (Ferretti, McRae, & Hatherell, 2001; Kuperberg, Kreher, Sitnikova, Caplan, & Holcomb, 2007; Wassenaar & Hagoort, 2007). This makes it challenging to dissociate the neural signature of a thematic role from the neural representation of the individual concept per se and the proposition. Another obstacle is that the identification of thematic role is confounded with the serial position of words, in simple structured English sentences.

The current study examined the possibility of decoding the thematic role binding of an agent or patient from the neural activation patterns evoked by semantically reversible events depicted in an animation (e.g., a monkey patting a rabbit) in healthy individuals. This type of event offers an opportunity for studying the neural representation of thematic

roles, because the correct assignment of thematic relation is critical to the correct understanding and representation of the event. The input modality of video allows the examination of thematic role representation from a perspective of event, without the restrictions of syntactic structure imposed by text stimuli. We used multivoxel decoding on functional magnetic resonance imaging (fMRI) data evoked by animation videos to investigate the neural basis of thematic relation representation. We trained the machine-learning classifiers to distinguish the activation patterns associated with two different thematic roles: agent and patient. If the neural representations of these two thematic roles are different as we hypothesized, the trained classifier should be generalizable to a new event including the same character, being able to identify with which thematic role the character is associated in this new event. Moreover, we hypothesized that there is cross-individual commonality in the neural representations of thematic roles. This hypothesis was tested by training a classifier on one set of participants and testing it on other participants.

Method

Experimental paradigm and procedure

Functional images were acquired while participants encoded the thematic relations in an animation video depicting a simple event. Sixteen video clips (programmed using Alice 3, alice.org) approximately 5 s long depicted one animal character performing a physical action on another animal character. There were three possible animals (monkey, rabbit, and wolf), four actions (kick, pat, punch, and push), and two possible agent–patient thematic role assignments. Two of the animals performed all four actions on each other. The third animal performed and received two actions from the other animals, resulting in 16 unique video clips, as shown in Table 1. The third animal was included in the stimulus set to force participants to encode the thematic roles of both animals in each event, rather than encoding one animal with its thematic role and inferring the other animal and its role. Each of the 16 events was depicted in two videos, alternating the left–right positions of the animals. Each video was presented twice in two scans, resulting in 64 trials that entailed four presentations of each event.

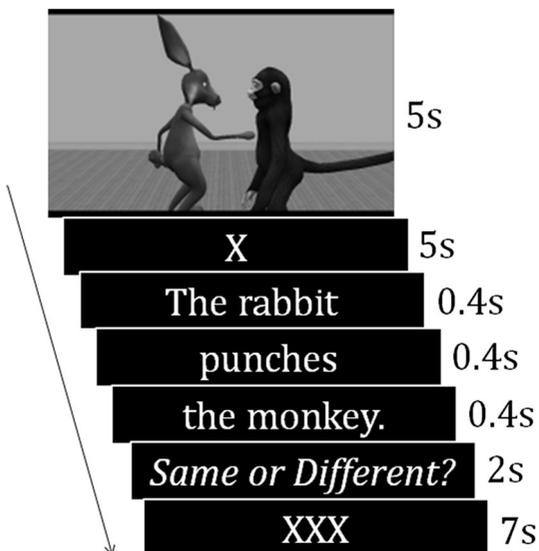


Figure 1. Illustration of stimulus presentation.

Table 1. Sentence summary of the events presented in the animation video clips.

Events
The rabbit punches the monkey.
The rabbit pats the monkey.
The rabbit pushes the monkey.
The rabbit kicks the monkey.
The rabbit punches the wolf.
The rabbit pats the wolf.
The monkey pushes the wolf.
The monkey kicks the wolf.
The monkey punches the rabbit.
The monkey pats the rabbit.
The monkey pushes the rabbit.
The monkey kicks the rabbit.
The wolf punches the rabbit.
The wolf pats the rabbit.
The wolf pushes the monkey.
The wolf kicks the monkey.

The participants were asked to mentally generate an *Agent-verb-patient* sentence to describe the event in the video. To ensure that participants followed this instruction, they were asked to judge whether or not an *Agent-verb-patient* sentence that followed the video contained the same thematic bindings and animals as those in the video (Figure 1). A 5-s blank interval separated the video and the sentence. Sixteen of the 64 comprehension test sentences were inconsistent with the video, eight of which mis-stated the agent-patient relation and eight that used incorrect verbs. Participants indicated their response by pressing a *same* or *different* button on a response apparatus. The mean comprehension test accuracy in the 2-s judgement window was 98%. A 7-s blank interval separated the video and the sentence, so that the reading and comprehension of the sentence would not contaminate the fMRI signal associated with the processing of the video.

Each judgement period was followed by a 7-s fixation, consisting of an "X" in the centre of the screen, during which participants were asked to clear their minds and fixate on the "X". In addition, six 17-s fixation trials were evenly distributed throughout the scans. These trials provided the baseline measure for assessing percentage signal change in functional images.

Participants

Five healthy, right-handed native English-speaking adults (2 males) from the Carnegie Mellon community participated and gave written informed consent

approved by the Carnegie Mellon Institutional Review Boards.

fMRI imaging protocol

Data were acquired using a Siemens Verio 3T scanner and 32-channel phased-array head coil (Siemens, Erlangen, Germany) at the Scientific Imaging and Brain Research (SIBR) Center at Carnegie Mellon University. Functional images were acquired using a gradient echo planar imaging (EPI) pulse sequence with time to repetition (TR) = 1000 ms, echo time (TE) = 25 ms, and a 60° flip angle. Twenty 5-mm thick, anterior commissure (AC)-posterior commissure (PC) aligned slices were imaged with a gap of 1 mm between slices. The acquisition matrix was 64 × 64 with 3.125 × 3.125 × 5-mm³ voxels.

fMRI data pre-processing

The fMRI data were corrected for slice timing, head motion, and linear trend and were normalized into Montreal Neurological Institute (MNI) space using SPM8 (Wellcome Department of Cognitive Neurology, London). Further analyses were performed using in-house scripts on Matlab7 (Mathworks, MA, USA). The percentage signal change (PSC) in signal intensity during each presentation trial of a video was computed at each voxel relative to a baseline activation level measured during all 17-s fixation intervals. The first 4 s of each fixation period were excluded to account for the hemodynamic response delay. The fMRI data from the processing of the animation video consisted of the mean of 4 s of brain images (4 images with a TR of 1 s), the first starting at the 5th second from the video onset. The PSC was then normalized to a mean of 0 and variance of 1 across voxels for each image (Mitchell et al., 2004; Pereira, Mitchell, & Botvinick, 2009), to equate the overall intensities of the mean PSC images.

Decoding thematic relations within participants

Given an fMRI activation pattern that corresponds to a single event (e.g., *The monkey pats the rabbit*), we tested the possibility of decoding the thematic role each character was associated with. To ensure that the decodability was not due to low-level visual properties in the video, activation data in the occipital lobe,

defined by Anatomical Automatic Labeling (Tzourio-Mazoyer et al., 2002), were excluded from consideration. The fMRI data from the four presentations of each event were averaged, thus balancing the left-right position of the agent and the patient.

The decoding used a leave-one-out cross-validation procedure, partitioning the fMRI data into a training set and a test set in each fold. Each test event contained two animal characters, each of whose thematic role assignments was tested. Two logistic regression classifiers, one per animal character, were trained to associate fMRI activation patterns with one of the two thematic roles. The training data for a given classifier were the activation patterns evoked by events that involved the given animal character in the role of either the agent or the patient. For instance, when the event *The rabbit punches the monkey* was being tested, a classifier was trained on all the other 11 events involving *rabbit*: 5 events in which *rabbit* was the *agent* and 6 in which it was the *patient*. (Note that if the unbalanced priors of the two classes did affect the classification, it would favour the incorrect answer.) Another classifier was trained independently on *monkey* for the same test event. The trained classifier for a character was applied to the activation patterns of the left-out test event, to predict its thematic role. The accuracy of each prediction was either 1 or 0. Because the two characters in the same event were trained and tested in two independent models, knowing the thematic role of one character from one classification model did not enter into the decoding of the other character.

Three hundred (out of more than 17,000 non-occipital) voxels with the most stable profiles were used in the classification. A voxel's activation profile was its vector of activation levels evoked by the set of the 16 video animations in the training set. A voxel's stability was measured by the mean of the pairwise correlations between its profiles across the four presentations. This voxel selection approach was performed in each cross-validation fold.

Two alternative approaches to selecting voxels were also tested. In the first alternative approach, voxels were selected based on the stability over six different role-character combinations. In the second alternative approach, voxels were selected based on the difference in activation magnitudes to agent and patient. A two-sample *t*-test between the events when the character was an agent versus the events

when it was a patient role in the training set was performed at each voxel. Voxels with the lowest *p* values were selected, as each of them responded reliably differently to the two thematic roles of the same character.

Decoding thematic relations across participants

To investigate whether there is a commonality across participants in the neural representation of thematic roles, a thematic role classifier was trained on the fMRI data of four participants and was tested on the left-out participant in a cross-validated manner. To ensure that the classifier was modelling the thematic role rather than a particular scenario, when one event was being tested, data associated with the same event from all the participants in the training set were also excluded. To identify voxels with stable responses both within and across participants, 500 voxels with stable response profiles over the 15 non-testing events *within* each training participant were first selected, resulting in a pre-selected set with no more than 2000 voxels. To estimate the cross-participant consistency of these pre-selected voxels, voxel responses to each event within each training participant were averaged across the four presentations. The cross-participant consistency was the mean correlation of responses between pairs of participants. The 600 voxels that were most consistent across participants in the pre-selected set were used for training and testing.

The decoding accuracy was evaluated based on an empirically generated null distribution. To obtain the distribution of accuracies at chance level, the same classification was performed except that the label of each test item was randomly selected. Repeating this procedure for 10,000 iterations resulted in the null distribution, against which the significance level of the actual accuracy was computed.

The statistical reliability of classification outcome using the voxels at selected locations was assessed relative to a chance-level distribution. The actual consistency of each voxel was measured by the voxel's probability of being selected across the 80 folds (5 participants \times 16 events). The null distribution was generated by an 8000-iteration random permutation. In each iteration, the selected voxels were randomly relocated for every fold. The probability of a voxel being selected over all the 80 folds indicated the chance-

level consistency. The maximum probability value across all the voxels was considered as one observation of the chance-level distribution. This procedure was repeated for 8000 iterations, resulting in a null distribution. Using the maximum chance probability across all voxels controlled for the familywise error rate. The critical value at $p = .05$ was then empirically obtained based on this null distribution for thresholding the actual probability map. Additionally an arbitrary threshold of cluster size ($n > 5$ voxels) was applied.

Results

Neural signature of thematic roles can be decoded across contexts

The neural representation of thematic roles was decodable across different events. A classifier was first trained on the multivoxel activation patterns associated with the same character bound to either the agent or the patient role in a set of events and was then tested on an unlabelled activation pattern evoked by a test event that was new to the classifier. The trained classifiers were able to decode whether the character was the agent or the patient in the test event with a reliable mean accuracy of .67 ($p < .05$) across five participants, with a range of .59–.75. The accuracies of all but one individual participant were significantly above chance (the lowest accuracy = .594, at the margin of the critical value of .60). Thus, it is possible to decode the thematic role associated with a single character in an event that is new to the model.

The mean accuracy of identifying the agent and patient in the test event was .64 and .70, respectively, with no statistically significant difference between them (paired-sample t test over characters, agent vs. patient, $t = -1.00$, $p > .3$). This result indicates that the accuracy of decoding thematic roles was not dominated by a single role.

Two alternative approaches to selecting voxels revealed similar accuracies of identifying thematic roles. As described above, the set of voxels that resulted in the reliable decoding was selected based on the stability over events. The motivation of using this approach was to find voxels that tuned to the semantic content of the entire event. The underlying assumption was that the thematic role information

can be identified by examining the neural representation associated with processing the full event. An alternative hypothesis is that there might be voxels that are tuned directly to the thematic roles. To test this hypothesis, classification was performed on a set of 300 voxels selected based on either their stability over role-character combinations, or their activation level differences between agent and patient. Accuracies based on these two alternative voxel selection methods were .66 ($SD = .04$) and .64 ($SD = .08$), respectively, numerically lower but similar to the results based on voxels selected at propositional level. The brain regions of voxels selected by these different methods were generally similar (examples in Figure S1 in the supplemental data).

Cross-participant commonality of thematic role representation

There was a commonality in the neural representations of thematic roles across participants. The common neural signatures allowed the identification of thematic roles in a novel event processed by any given participant, based on a classifier having been trained exclusively on data from other participants. The accuracy of identifying the agent or patient role that a character played across participant was .66, significantly above chance level (critical value at $p = .05$ was .59, based on a 5000-iteration random permutation test that randomly shuffled the labels of the test data). Thus the thematic role information was decodable from individual participants, and such an effect was common across individuals.

The voxels that led to the reliable cross-participant decoding of thematic roles were located in multiple brain regions. The main clusters that permitted accurate classification of thematic roles were in the left anterior temporal lobe, the right posterior middle temporal gyrus, and the precuneus. These regions that have previously been associated with thematic role processing or representations, as discussed below (Figure 2; see Table 2 for the list of consistently selected voxel clusters in the current study).

Discussion

The study demonstrated the capability of distinguishing the neural signature of agent versus patient thematic role bindings. The thematic role signature

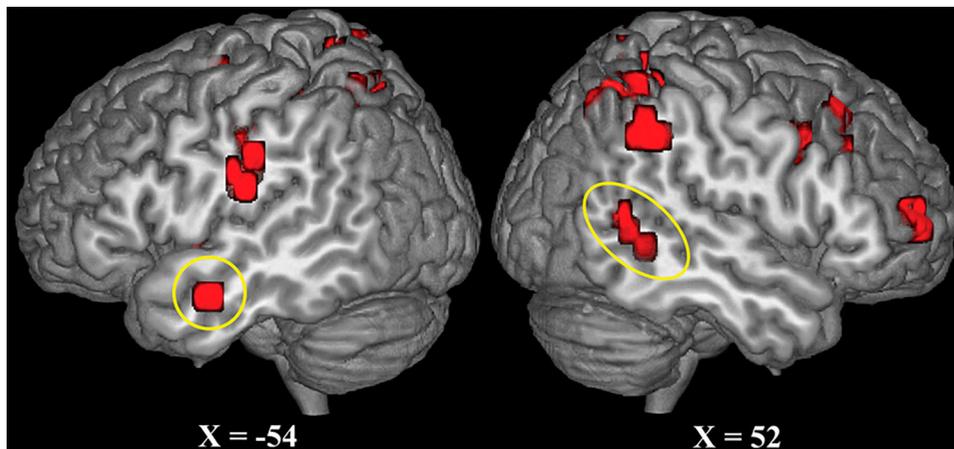


Figure 2. Consistently selected brain regions that support cross-participant thematic role decoding. The two regions highlighted with circles are regions or are adjacent to regions that have been previously identified for representing thematic role information. [To view this figure in colour, please see the online version of this Journal.]

learned by a classifier from the fMRI data associated with one set of propositions enabled the identification of the thematic role of a character in a new set of propositions. Even though thematic roles are computed dynamically and are temporarily bound to a word-level concept (animal characters in this case) during the course of comprehending an event, it was still possible to identify the neural correlates that distinguish agent and patient from each other from the fMRI signal. The voxels that are tuned to thematic roles provide as accurate, if not more accurate, identification of thematic roles as voxels showing maximal activation differences between agent and patient.

One previous study identified two different brain areas that respectively encoded the word identity of

agent and word identity of patient (Frankland & Greene, 2015). The classification problem was to distinguish different word concepts at the same thematic role (man-as-agent, cat-as-agent, etc.), or, in other words, role-modulated word concepts; thus the question being addressed was “which brain areas encode the semantic content of the agent, and which brain areas encode the semantic content of the patient?”. The current study took a different approach, directly decoding the thematic role of a given character (rabbit-as-agent or rabbit-as-patient), to identify the neural basis of thematic role representation per se. In addition, there was a possibility in Frankland and Greene (2015) that the identified regions reflected a mixture of the “agency” of one noun relative to another and the actual role it fulfilled in the sentences, particularly when (a) two thirds of the sentences were formed by agent and patient from the two semantic categories of human and non-human, and (b) the responses in the regions of interests were associated with affective arousal of the sentence. The current study used animal characters of high semantic homogeneity with homogenous appearances in the video presentation, to ensure that the representational differences were driven by the thematic role assignments.

Moreover, the representational differences characterized by only four participants’ data were generalizable to a new participant with above-chance accuracy. This finding indicates reliable commonality across people in the neural representations of agent and patient. Among the brain regions that were

Table 2. Brain regions that support cross-participant thematic role decoding.

Region	MNI coordinates (mm)			Number of voxels
	x	y	z	
L anterior middle temporal gyrus	-50	-9	-20	6
R posterior middle temporal gyrus	56	-56	-8	15
Precuneus	6	-65	28	10
L postcentral gyrus	-50	-21	22	83
L inferior parietal lobule	-34	-53	46	14
L inferior parietal lobule	-34	-65	40	6
L superior parietal lobule	-19	-50	64	7
R superior parietal lobule	22	-56	64	12
R inferior parietal lobule/angular gyrus	38	-68	40	58
R inferior frontal gyrus	41	10	28	13
R middle frontal gyrus	41	26	34	22
L insula	-37	1	-2	7
R rostralateral prefrontal cortex	31	54	-2	17
Medial prefrontal cortex	6	38	34	6

Note: MNI = Montreal Neurological Institute.

consistently selected across participants and events, the anterior temporal lobe and the posterior middle temporal gyrus are the two regions that have been identified by two previous studies investigating thematic role representations. The cluster in the left anterior middle temporal lobe was close to the regions identified by Frankland and Greene (2015) that decode the object identities of agent or patient.

The anterior temporal lobe has been long proposed to play a critical role in concept representation based on neuropsychology and functional neuroimaging research (Patterson, Nestor, & Rogers, 2007; Price, 2012), likely to be sensitive to both integrative semantic and syntactic processing (Rogalsky & Hickok, 2009). Recent functional imaging studies have shown that multiple features of object concepts, including colour, shape, typical location, and manner of manipulation, are jointly encoded in this region (Coutanche & Thompson-Schill, 2014; Peelen & Caramazza, 2012). The current study further suggests that the anterior temporal area is sensitive not only to single word-level semantics. The fact that this region shows stable activation profiles over the events suggests that it represents or binds a proposition-level feature, the thematic role, to a concept.

The two clusters centered at the posterior middle temporal gyrus and the precuneus were previously found to show repetition suppression to semantically equivalent events, some of which were described by syntactically different sentences (Menenti, Gierhan, Segaert, & Hagoort, 2011). However, these areas were not identified as encoding thematic roles by Frankland and Greene (2015). One possible explanation regarding this discrepancy is the difference in the stimulus modality. The current study and Menenti et al. (2011) used animation videos or picture stimuli along with sentences, whereas the Frankland and Greene (2015) presented written sentences only. The posterior superior temporal sulcus is known for processing biological motion (Giese & Poggio, 2003; Grossman & Blake, 2002). In the current study, the animation stimuli may have evoked the processing of the biological motions entailed in the actions performed by the animal characters. It is possible that the representation of an action itself leads to the formation of a mental structure, an event, or proposition representation that includes various thematic roles associated with the event. The precuneus region that is identified by two

studies may reflect the mentalization of the characters or attribution of observed behaviours (Spunt & Lieberman, 2012) evoked by the pictorial or video presentation.

The question of how thematic roles are neurally represented is related to the fundamental issues of representation and processing of propositions. Although we have shown that thematic roles can be decoded from their neural representations, the precise nature of the neural coding remains unknown. More specifically, it is unknown whether the thematic role of a concept is encoded by modulating the context-independent neural representation of that concept itself or, alternatively, the thematic role assignment is bound directly to the representation of the proposition. It may be possible to distinguish between such alternatives by using imaging modalities with higher temporal resolution (such as magnetoencephalography, MEG) to determine whether the representation of a concept changes when its thematic role is assigned.

Several facets of the current findings suggest that the differential encoding of agent and patient occurs at the level above and beyond neural representation of single concept. The two thematic roles of agent and patient were decodable from the same temporal window with similar accuracy. Moreover, the anterior temporal lobe, often associated with higher level binding or integration of semantic information, was involved in thematic role identification here and in other studies. Other regions such as the rostralateral prefrontal cortex are also involved in the dynamic computation of relations (Wendelken, Nakhachenko, Donohue, Carter, & Bunge, 2007). But it seems reasonable that propositional representations are more than the sum of their parts, even if this sum is an emergent structure.

In summary, the current study reveals how proposition-level information, specifically the thematic roles of agent and patient, are neurally encoded. The neural signatures of thematic role can be identified in concept binding areas of the brain.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding information

This work was supported by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research

Projects Activity (IARPA), via Air Force Research Laboratory (AFRL) [contract number FA8650-13-C-7360].

References

- Berndt, R. S., Mitchum, C. C., & Haendiges, A. N. (1996). Comprehension of reversible sentences in “agrammatism”: A meta-analysis. *Cognition*, 58(3), 289–308.
- Coutanche, M. N., & Thompson-Schill, S. L. (2014). Creating concepts from converging features in human cortex. *Cerebral Cortex*, 25(9), 2584–2593.
- Ferretti, T. R., McRae, K., & Hatherell, A. (2001). Integrating verbs, situation schemas, and thematic role concepts. *Journal of Memory and Language*, 44(4), 516–547.
- Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28(1-2), 3–71.
- Frankland, S. M., & Greene, J. D. (2015). An architecture for encoding sentence meaning in left mid-superior temporal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 112(37), 11732–11737.
- Friederici, A. D., & Frazier, L. (1992). Thematic analysis in agrammatic comprehension: Syntactic structures and task demands. *Brain and Language*, 42(1), 1–29.
- Giese, M. A., & Poggio, T. (2003). Cognitive neuroscience: Neural mechanisms for the recognition of biological movements. *Nature Reviews Neuroscience*, 4(3), 179–192.
- Grossman, E. D., & Blake, R. (2002). Brain areas active during visual perception of biological motion. *Neuron*, 35(6), 1167–1175.
- Kuperberg, G. R., Kreher, D. A., Sitnikova, T., Caplan, D. N., & Holcomb, P. J. (2007). The role of animacy and thematic relationships in processing active English sentences: Evidence from event-related potentials. *Brain and Language*, 100(3), 223–237.
- Meltzer, J. A., McArdle, J. J., Schafer, R. J., & Braun, A. R. (2010). Neural aspects of sentence comprehension: Syntactic complexity, reversibility, and reanalysis. *Cerebral Cortex*, 20(8), 1853–1864.
- Menenti, L., Gierhan, S. M. E., Segaert, K., & Hagoort, P. (2011). Shared language: Overlap and segregation of the neuronal infrastructure for speaking and listening revealed by functional MRI. *Psychological Science*, 22(9), 1173–1182.
- Mitchell, T. M., Hutchinson, R., Niculescu, R. S., Pereira, F., Wang, X., Just, M., & Newman, S. (2004). Learning to decode cognitive states from brain images. *Machine Learning*, 57(1–2), 145–175.
- Miyake, A., Carpenter, P. A., & Just, M. A. (1994). A capacity approach to syntactic comprehension disorders: Making normal adults perform like aphasic patients. *Cognitive Neuropsychology*, 11(6), 671–717.
- Patterson, K., Nestor, P. J., & Rogers, T. T. (2007). Where do you know what you know? The representation of semantic knowledge in the human brain. *Nature Reviews Neuroscience*, 8(12), 976–987.
- Peelen, M. V., & Caramazza, A. (2012). Conceptual object representations in human anterior temporal cortex. *Journal of Neuroscience*, 32(45), 15728–15736.
- Pereira, F., Mitchell, T., & Botvinick, M. (2009). Machine learning classifiers and fMRI: A tutorial overview. *NeuroImage*, 45 (Suppl. 1), S199–S209.
- Price, C. J. (2012). A review and synthesis of the first 20 years of PET and fMRI studies of heard speech, spoken language and reading. *NeuroImage*, 62(2), 816–847.
- Richardson, F. M., Thomas, M. S. C., & Price, C. J. (2010). Neuronal activation for semantically reversible sentences. *Journal of Cognitive Neuroscience*, 22(6), 1283–1298.
- Rogalsky, C., & Hickok, G. (2009). Selective attention to semantic and syntactic features modulates sentence processing networks in anterior temporal cortex. *Cerebral Cortex*, 19(4), 786–796.
- Saffran, E. M., Schwartz, M. F., & Linebarger, M. C. (1998). Semantic influences on thematic role assignment: Evidence from normals and aphasics. *Brain and Language*, 62(2), 255–297.
- Spunt, R. P., & Lieberman, M. D. (2012). An integrative model of the neural systems supporting the comprehension of observed emotional behavior. *NeuroImage*, 59(3), 3050–3059.
- Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., ... Joliot, M. (2002). Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *NeuroImage*, 15(1), 273–289.
- Wassenaar, M., & Hagoort, P. (2007). Thematic role assignment in patients with Broca’s aphasia: Sentence-picture matching electrified. *Neuropsychologia*, 45(4), 716–740.
- Wendelken, C., Nakhabenko, D., Donohue, S. E., Carter, C. S., & Bunge, S. A. (2007). “Brain Is to Thought as Stomach Is to ??”: Investigating the role of rostralateral prefrontal cortex in relational reasoning. *Journal of Cognitive Neuroscience*, 20(4), 682–693.
- Wu, D. H., Waller, S., & Chatterjee, A. (2007). The functional neuroanatomy of thematic role and locative relational knowledge. *Journal of Cognitive Neuroscience*, 19(9), 1542–1555.