Decoding the Neural Representation of Story Meanings across Languages

Morteza Dehghani, Reihane Boghrati, Kingson Man, Joe Hoover, Sarah I. Gimbel,
Ashish Vaswani, Jason D. Zevin, Mary Helen Immordino-Yang, Andrew S. Gordon,

Antonio Damasio, Jonas T. Kaplan

University of Southern California

Author Note

This work has been funded in part by Institute for New Economic Thinking and DARPA. None of the ideas discussed and explored in this paper have been previously published or publically disclosed. Part of the fMRI data used in this paper has been previously published in Kaplan et al. (2016). The research focus and data analysis of the two papers, however, are mutually exclusive. Correspondence regarding this article should be address to Morteza Dehghani, mdehghan@usc.edu, 3620 S. McClintock Ave, Los Angeles, CA 90089-1061.

Abstract

Drawing from a common lexicon of semantic units, humans fashion narratives whose meaning transcends that of their individual utterances. However, while brain regions that represent lower-level semantic units, such as words and sentences, have been identified, questions remain about the neural representation of narrative comprehension, which involves inferring cumulative meaning. To address these questions, we exposed English, Mandarin and Farsi native speakers to native language translations of the same stories during fMRI scanning. Using a new technique in natural language processing, we calculated the distributed representations of these stories (capturing the meaning of the stories in high-dimensional semantic space), and demonstrate that using these representations we can identify the specific story a participant was reading from the neural data. Notably, this was possible even when the distributed representations were calculated using stories in a different language than the participant was reading. Relying on over 44 billion classifications, our results reveal that identification relied on a collection of brain regions most prominently located in the default mode network. These results demonstrate that neuro-semantic encoding of narratives happens at levels higher than individual semantic units and that this encoding is systematic across both individuals and languages.

Decoding the Neural Representation of Story Meanings across Languages

One of the defining characteristics of human language is its capacity for semantic extensibility. Drawing from a common lexicon of morphemes and words, humans generate and comprehend sophisticated, higher-level utterances that transcend the sum of their individual units. This is perhaps best exemplified in stories, in which sequences of events invite inferences about the intentions and motivations of characters, about cause and effect, and about theme and message. The kind of meaning that emerges over time as one listens to a story is not easily captured by analysis at the word level alone. Further, a necessary condition for generating higher-level semantic constructs is that speakers of the same language infer similar meanings from expressions of both lower and higher level semantic units. For example, it can be assumed that when speakers of the same language listen to stories, the perceived meanings of these stories have much in common. Does this imply that individuals' brain activity shows systematic patterns related to the meaning of the stories? If so, can these neural representations of meaning be associated post-hoc with particular stories? Also, when a story is translated from one language to another, the particulars of sounds and words are left behind, but a core of semantic content survives the translation. Does it then follow that the neural representations related to this semantic content are also comparable across the brains of speakers of different languages?

Understanding how conceptual knowledge is represented and organized in the human brain is one of the central problems of cognitive science and many studies have aimed at exploring and understanding the neural representations of concepts (Damasio, Tranel, Grabowski, Adolphs, & Damasio, 2004). Relying on multiple approaches and methods, this body of work has identified a collection of separable brain regions, constituting a large network, within which different conceptual categories and the corresponding words are represented (e.g. Binder, Desai, Graves, & Conant, 2009; Damasio et al., 2004; Grabowski et al., 2001; Tranel, Damasio, & Damasio, 1997). More recently a word-level *semantic map* has been proposed (Huth, de Heer, Griffiths, Theunissen, & Gallant, 2016).

In this work, our aim is to move beyond word-level semantics to investigate neuro-semantic representations at the story-level across three different languages. Specifically, we set out to determine if there are systematic patterns in the neuro-semantic representations of stories beyond those corresponding to word-level stimuli. Our aim is motivated by the long-standing understanding that discourse representations are different from the sum of all of their lexical or clausal parts. Most psycholinguistic models of discourse processing are concerned with the construction and representation of structures such as propositions and arguments (Kintsch, 1974), and situation models (Kintsch, 1988) that are, of course, derived from specific words and sentences, but are more abstract. Indeed, many interesting effects of discourse representations can be demonstrated by comparing the comprehension of the same sentence under conditions that differ in the plausibility of the most common interpretation of the sentence in isolation (Sanford & Garrod, 1998) or, conversely, comprehension of the same passage (Johnson-Laird, 1983) under conditions that differ in the availability of a situation model to begin with¹. Hence, we aim to use neuro-semantic representations of stories, containing the meaning of the stories as represented in high-dimensional semantic space, to decode neural activations (i.e. map back to the semantic space) during story reading. This will allow us to investigate the networks involved in representing higher-level semantics, and to explore whether activation patterns related to higher-level meaning are common across people and languages.

One approach that has been proven to be fruitful in exploring neuro-semantic representations is to investigate the relationship between various co-occurrence patterns of words within large textual corpora, and to relate these patterns to neural activity recorded during exposure to those words. Examining the representations of concrete nouns, Mitchell et al. (2008) proposed a computational model that predicts voxel activation using a weighted sum of various semantic features of nouns, calculated based

¹There is also emerging evidence from imaging studies that even low-level features, such as word-length, are processed differently when the words are analyzed context-free as individual words compared to when analyzed in longer passages (Schuster, Hawelka, Hutzler, Kronbichler, & Richlan, 2016).

on word co-occurrence in a large English corpus. Using fMRI data from 9 participants who viewed 60 word-picture pairs, they demonstrated that their model captures aspects of neuro-semantic representations of these concepts; it could correctly predict which word/picture participants were viewing with a higher-than-chance accuracy. The most predictive voxels were distributed around the cerebral cortex, and included sensory regions in the temporal and occipital lobes, motor related regions in the frontal and parietal lobes, and other regions of the orbitofrontal and inferior frontal cortices.

Other researchers have proposed different representational (Fyshe, Talukdar, Murphy, & Mitchell, 2014; Yogatama, Faruqui, Dyer, & Smith, 2014), computational (Huth, Nishimoto, Vu, & Gallant, 2012; Just, Cherkassky, Aryal, & Mitchell, 2010; Shinkareva, Malave, Mason, Mitchell, & Just, 2011), and methodological (Guimaraes, Wong, Uy, Grosenick, & Suppes, 2007; Sudre et al., 2012) modifications to the general Mitchell et al. (2008) approach in order to better capture and explore the neural architectures involved in representing concrete nouns. Notably, Huth et al. (2016) use word-embeddings, a quantitative representation of meaning calculated based on co-occurrence patterns of 985 common English words, along with a generative model of areas covering the cortex, to create a detailed semantic-map reported to be consistent across individuals. Generally, in word-level analysis applied in the above approaches, each word is treated as an isolated symbol, independent of any relation to context or to the overall meaning that the story conveys. Therefore, even though these approaches have been very useful in mapping the meaning of individual words, they do not reveal how the higher-level meaning that emerges through the course of a story is represented.

Discourse has typically been operationalized for fMRI analysis by reference to theory-dependent identification of key discourse features (Whitney et al., 2009), general linear model (Yarkoni, Speer, & Zacks, 2008) or cross-correlation (Ames, Honey, Chow, Todorov, & Hasson, 2015; Lerner, Honey, Silbert, & Hasson, 2011) analyses comparing texts differing in coherence. More recently researchers have started to investigate the neural architectures involved in the semantic representation of sentences and story segments. Wehbe et al. (2014) present a dynamic model for studying how different

regions of the brain encode various types of information while reading a story. Passages from a Harry Potter story were presented to 9 participants using Rapid Serial Visual Presentation (each word presented for 0.5 seconds). The story segments (16 words) were then represented using a combination of various syntactic features, semantic features of individual words, and low-level features such as number of letters, and used in a predictive model that was able to predict which one of the two novel passages was being read. They showed that different aspects of story processing were encoded in different brain networks. For example, activity in temporo-parietal regions in both hemispheres was related to processing sentence length and complexity, while semantic classification was most accurate in voxels from the middle and superior temporal gyri and the left inferior frontal gyrus. In general, brain regions on the lateral surfaces of both hemispheres in the temporal, occipital, and parietal regions appear to have been most predictive of linguistic content in the stories.

While previous methodologies have found semantic representations to be widely distributed throughout the brain, there now seems to be growing evidence for a special involvement of the default mode network (DMN) in representation of high-level meaning and language comprehension. The DMN was originally identified as a “resting state” network that shows high baseline activity when people are asked to rest without engaging in any specific externally-focused task (Raichle, 2015; Raichle & Snyder, 2007). This bilateral network includes midline cortical structures (medial prefrontal cortex, precuneus, and posterior parietal cortex) as well as lateral structures (inferior parietal lobe and anterior temporal lobes). Activity in these nodes is highly correlated during resting fMRI (Buckner, Andrews-Hanna, & Schacter, 2008) and has been thought to reflect such cognitive processes as mind-wandering (Smallwood & Schooler, 2015), thinking about one’s self (Qin & Northoff, 2011), remembering the past and imagining the future (Østby et al., 2012), and in general to support stimulus-independent thought (Smallwood et al., 2013).

Yet, the specific function of this network is not well understood. Recent work has implicated the DMN in more active forms of cognition (Spreng, 2012), and specifically

in semantic processing (Binder et al., 2009). A series of studies have shown that the DMN seems to be involved in representing the global meaning of passages, rather than meaning at the word or sentence-level (Ferstl, Neumann, Bogler, & Von Cramon, 2008; Lerner et al., 2011). Further, the activity in the DMN is consistent when a story is presented in different modalities (spoken vs. written), or in different languages (Russian vs. English to native speakers of these languages), indicating highly abstracted representations of the stimuli in this network (Chen et al., 2016; Honey, Thompson, Lerner, & Hasson, 2012; Regev, Honey, Simony, & Hasson, 2013). The word-level semantic map produced by Huth et al. (2016) also demonstrates significant overlap with the DMN. More recently, Simony et al. (2016) show that the DMN reconfigures consistently across subjects when processing narrative stimuli. Also, a recent study demonstrates that patterns of activity in the DMN when people are describing a narrative are highly consistent across individuals and specific to events in the narrative (Chen et al., 2016). Our own recent finding shows that activity in some DMN nodes increases throughout the course of a story, and is greatest when reading stories containing strong moral values (Kaplan et al., 2016).

Parallel to this line of work, multi-voxel pattern analysis (MVPA) has also been used to explore common representations across languages by predicting the neural response to a noun in speakers of one language, based on the noun’s neural representation in speakers of a different language (Buchweitz, Shinkareva, Mason, Mitchell, & Just, 2012; Correia et al., 2014; Zinszer, Anderson, Kang, Wheatley, & Raizada, 2015). Supporting evidence from lesion studies, Correia et al. (2014) find the left anterior temporal lobe to be an area with high predictive accuracy independent of language. The ATL is part of the DMN, confirming this region as part of a network in which the abstract representations of words are coded.

In this paper, we use *story-level* embeddings to examine the neuro-semantic representations of stories in people from three different linguistic-cultural backgrounds, across their three native languages. We demonstrate that story-level vector representations can capture story content across languages, and can be used to localize

the neuro-semantic content of stories to specific regions of the brain. We mapped the brain regions in which language-independent semantic story content was detected, revealing a set of regions in the DMN that have been implicated in the generation of rich internal experiences. While previous work has demonstrated that the DMN responds to abstract meaning of stories, we show that the patterns of activation in the DMN correspond to the unique encoding of the meaning of narratives, and that these patterns can be used to decode the meaning of narratives across three different languages.

Method

Stimuli

For our stimuli, we used real-world personal narratives written by people describing their experiences. In order to find such narratives, we started with a corpus of over 20 million weblog story posts that were compiled from Spinn3r.com (Sagae et al., 2013). Next, the corpus was queried on different topics (e.g. telling a lie, getting a divorce) via a text retrieval system (Apache Lucene). Forty stories were chosen and each was condensed to a paragraph of 145-155 words. Professional translators were used to translate these English stories into Mandarin Chinese and Farsi. Lastly, these translations were back-translated by native speakers into English, and the back-translated versions were checked for any inconsistencies with the original stories. Any minor inconsistencies were resolved by the translators.

Distributed Representation of Stories

Distributed word representations refer to n-dimensional numeric vector representations that capture some semantic and syntactic aspects of words. Prior to the recent widespread upsurge of deep neural networks, methods such as Latent Semantic Analysis (Deerwester, Dumais, Landauer, Furnas, & Harshman, 1990) and Latent Dirichlet Analysis (LDA; Blei, Ng, & Jordan, 2003) were used to build such distributed word representation.

In recent years, neural network approaches for language processing have gained

considerable attention. One notable example is the *word2vec* (Mikolov, Chen, Corrado, & Dean, 2013) method which generates distributed word vectors from a text corpus. It efficiently trains a model that predicts a word conditioned on its surrounding context (e.g. the five words before and after the target word) by maximizing vector similarity of words that appear together and minimizing the similarity of words that do not.

Moving beyond representations at the word level, several related methods have been proposed for distributed representation of higher-level textual units (i.e. sentences, paragraphs, or documents). Paragraph vector (Le & Mikolov, 2014) is a technique that simultaneously learns vector representations for words and texts, and captures different aspects of the semantics of the text in the training process.

In this work, we used a version of paragraph vector derived from word2vec called *doc2vec* (Le & Mikolov, 2014). At a high level, doc2vec uses a technique called the *Distributed Memory* Model which aims to maximize the probability of a word w_i , given the paragraph D_i that the word is drawn from and the words adjacent to it (i.e. its context) $w_{i-j}, \dots, w_{i-1}, w_{i+1}, \dots, w_{i+j}$. This is treated as a multi-class classification problem, and paragraphs and context words are represented as n-dimensional vectors. During training, these vectors are optimized via stochastic gradient descent with backpropagation (Rumelhart, Hinton, & Williams, 1986) to maximize the probability of the target word. That is, at each iteration of gradient descent, a context for a given target word is sampled from a random paragraph. Then the representation of this paragraph D_i and the representations of the context words are used to predict the target word. The gradient from this prediction is then used to update the paragraph vector D_i and context vectors.

An important component of this model is the paragraph matrix D , which both provides context for the interpretation of individual words, and merges the semantics at the lexical level. The columns of matrix D contain ‘paragraph vectors’ that are each optimized to represent a paragraph or other units of higher level text drawn from the training corpus. These paragraph vectors are learned via a sliding window function that predicts each word in the paragraph, given the paragraph vector and vectors

representing the words contained in the sliding window. Importantly, paragraph vectors are constant across windows sampled from the same paragraph; thus, they can be thought of as representing the component of the paragraph that transcends and unifies each of the windows sampled from the paragraph. Accordingly, doc2vec represents paragraphs as unique higher-level units and not simply as aggregations of individual words. The relationship between paragraph vectors and the contexts contained in a paragraph can also be understood as hierarchical, such that words and contexts are nested in paragraphs. From this view, a portion of the variance associated with words can be attributed to their immediate context (e.g. local semantic and syntactic phenomena), but an additional portion of this variance is explained by the paragraph from which the word is drawn. Accordingly, doc2vec’s paragraph vectors represent the component of a paragraph that cannot be reduced to word-level semantics, they represent the higher-level meaning expressed by the paragraph – or, as in our case, narrative.

Supporting this view, it has been shown that doc2vec captures the overall meaning of larger pieces of text significantly better than other related techniques (Dai, Olah, & Le, 2015; Lau & Baldwin, 2016; Le & Mikolov, 2014). Through a series of computational experiments, Dai et al. (2015) show that doc2vec more effectively measures semantic similarity between long pieces of texts compared to methods that do not have an integrative component and use lexical semantics in isolation. They specifically show doc2vec results in higher classification accuracy, and in more meaningful nearest neighbors, across several different domains compared to LDA, bag of words and averaged word embeddings. Similarly, Lau and Baldwin (2016) in two different tasks (duplicate question detection in a web forum, and prediction of the similarity of a pair

of sentences) show that doc2vec outperforms word2vec and n-gram models²³.

All model training was performed using the doc2vec c library, now available through *gensim* (Řehůřek & Sojka, 2010) in python⁴. To train the paragraph vector model for English stories, we used an English corpus of weblogs of personal stories (~ 2 billion words; See Gordon & Swanson, 2009, for more details about how this corpus was compiled). Similarly, Chinese story vectors were built using a Chinese blog corpus of personal stories (~ 2 billion words; Gordon, Huangfu, Sagae, Mao, & Chen, 2013). A similar process was applied to the Farsi stories with a minor difference on the type of corpora used for training. Since Farsi blog posts are mostly written in informal Farsi, and the stories we used are in formal Farsi, we used a combination of Farsi news corpora (AleAhmad, Amiri, Darrudi, Rahgozar, & Oroumchian, 2009) along with a large weblog corpus (~ 2 billion words) of personal stories to have a better representation of stories in formal Farsi. Model training took between 20-24 hours for each language. After completion of the training process, each language model was then used to represent each of the stories as a 100-dimensional vector. This resulted in 40 story-vectors per language⁵.

²We further performed a behavioral experiment demonstrating the effectiveness of doc2vec compared to word-level operations in capturing the overall gist of stories. Participants rated the similarity between documents predicted to be similar by doc2vec and word2vec. Results indicated that doc2vec was substantially better at approximating human perceptions of document similarity. Further, we investigated how many of the first three closest neighbors of the stories match across the two representations. Knowing this detail is important as it provides insights on how much the semantic spaces constructed using the word-level technique and doc2vec are similar. Confirming our argument, there was only 21.67% overlap between the nearest neighbors of the same stories between the two representations. See Supplementary Materials for details about these experiments.

³We compared the performance of two different word-level representations of the stories to the doc2vec representations in our whole-brain analysis, and show that one of the two methods, namely LDA (Blei et al., 2003), performs at chance-level, while the other method, Linguistic Inquiry and Word Count (LIWC; Tausczik & Pennebaker, 2010), performs significantly lower than doc2vec. See Supplementary Materials for details about our whole-brain analysis.

⁴Detailed instruction on how to use the doc2vec framework in python is available at <https://radimrehurek.com/gensim/models/doc2vec.html>

⁵It should be noted that the inter-language story vectors are not linearly correlated ($r^2_{Mandarin-Farsi} =$

Participants

Ninety-five healthy participants with no history of psychological or neurological disorders were recruited from the University of Southern California community and the surrounding Los Angeles Area. We recruited participants from three groups: Americans, Chinese, and Iranians. All American participants were born in the United States and were native English speakers who grew up in exclusively English-speaking households. All Chinese and Iranian participants were born and raised in their native country, and had been in the United States for fewer than 5 years. Chinese and Iranian participants were all fluent in English in addition to their native languages.

One American participant, two Chinese participants, and two Iranian participants were excluded either for excessive motion or for not completing the scans. This left 30 American participants (mean age: 23.83 ± 0.92 , 15 male), 30 Chinese participants (mean age: 23.47 ± 0.39 , 16 male), and 30 Iranian participants (mean age: 26.66 ± 0.60 , 16 male). Subjects were paid \$20 per hour for their participation and gave informed consent approved by the Institutional Review Board of the University of Southern California.

fMRI Experiment

Participants read study instructions in their native language and were given the opportunity to ask questions to the experimenter before entering the fMRI scanner. For each participant there were 5 story-reading scans (556 seconds), one resting state scan (360 seconds), and one additional task run not presented here (question answering scan; for more details about the experimental setup see Kaplan et al., 2016). During the story scans, each story was preceded by a context slide (2s) identifying the protagonist of the story (e.g. American Mother). Following a 1.5 second delay, the story was presented over the course of 3 slides of text, each displayed for 12 seconds.

$\overline{0.0008}, r^2_{\text{Mandarin-English}} = 0.0003, r^2_{\text{English-Farsi}} = 0.0183$), while there is much higher correlation between the story vectors in the same language (average $r^2 = 0.308$).

fMRI Parameters

Imaging was performed using a 3T Siemens MAGNETOM Trio System with a 12-channel matrix head coil at the Dana and David Dornsife Neuroscience Institute at the University of Southern California. Functional images were acquired using a gradient-echo, echo-planar, T2*-weighted pulse sequence (TR = 2000 msec, one shot per repetition, TE = 25 msec, flip angle = 90°, 64 x 64 matrix). Forty slices covering the entire brain were acquired with a voxel resolution of $3.0 \times 3.0 \times 3.0$ mm. Functional data were continuously acquired for each run, with a short break between runs. A T1-weighted high-resolution ($1 \times 1 \times 1$ mm) image was acquired using a three-dimensional magnetization-prepared rapid acquisition gradient (MPRAGE) sequence (TR = 2530 msec, TE = 3.09 msec, flip angle = 10°, 256 x 256 matrix). Two hundred and eight coronal slices covering the entire brain were acquired. Total scan time for each participant was approximately 80 minutes.

fMRI Data Preprocessing

Univariate data analysis was performed with FSL (FMRIB's Software Library <http://fsl.fmrib.ox.ac.uk/fsl/fslwiki/>). Data were preprocessed using standard steps: motion correction (Jenkinson et al. 2002), 8mm FWHM spatial smoothing, highpass temporal filtering using Gaussian-weighted least-squares straight line fitting with a sigma of 60s (corresponding to a period of 120s) and slice timing correction. Data were also corrected for magnetic field inhomogeneities using field maps acquired for each subject.

Data were then analyzed using the General Linear Model. Each component of the task (context, the combined 3 screens of the story, and the question) was modeled by convolving the task design with a double-gamma hemodynamic response function. The model for each functional scan included one regressor for all of the context periods, one regressor for all of the question periods, and a set of regressors for each 12-second story period in the scan. Also included in the model were the temporal derivative of each story regressor, and six motion correction parameters. This analysis resulted in a

statistical map of standardized z scores for each story for each participant, representing the brain activity for that story relative to resting baseline.

To register the data to a common space for cross-subjects analysis, we used FSL’s FLIRT tool in two stages (Jenkinson et al. 2002; Jenkinson and Smith 2001). First, functional images from each participant were aligned with their own T1-weighted MPRAGE using a 6 degrees of freedom rigid-body warp. Next, the MPRAGE was registered to the standard MNI atlas with a 12 degrees of freedom affine transformation, and then this transformation was refined using FNIRT nonlinear registration (Andersson et al. 2007a; 2007b).

Analysis

As discussed earlier, 90 participants (30 from each culture) were scanned while reading 40 different stories. The fMRI data were preprocessed and the z-score activation map for each story was converted to a vector of 212,018 dimensions (voxels) for each participant-story pair, resulting in a matrix of $40 \times 212,018$ for each participant. Next, as discussed previously, the text of each of the forty stories was converted to a 100 dimensional semantic vector using the paragraph vector method described earlier, resulting in a matrix of dimension 40×100 for each language.

Searchlight-based multi-voxel pattern analysis was performed in order to investigate specific brain networks that encode neuro-semantic representations of stories. Specifically, a ridge regression model was fitted on the neural activity recorded in a sphere (of 4-voxel radius), successively centered around every voxel for each story (as predictor variables), with its accompanied 100 dimensional semantic representation of that story (as observation variables). The fitted model was evaluated using k -fold cross-validation: the ridge regression model was trained on every possible pair of 38 stories and tested on the two remaining stories, resulting in $\binom{40}{2}$ analyses per voxel. In each fold, using the trained model on the 38 stories, the story vectors were predicted for the two left-out stories, $Story_A$ and $Story_B$. A similar approach to Mitchell et al. (2008) was then used to evaluate the predictions: if the predicted vectors were more

similar (as assessed by cosine similarity) to the vectors of their target stories, as compared to the vectors of their non-target stories, then the classification was counted as correct. More formally, this can be described as:

$$\cos(\vec{predicted}_A, \vec{actual}_A) + \cos(\vec{predicted}_B, \vec{actual}_B) > \cos(\vec{predicted}_A, \vec{actual}_B) + \cos(\vec{predicted}_B, \vec{actual}_A) \quad (1)$$

The accuracy of the classification for the voxel was then averaged over the 780 folds. For each participant, this analysis was then repeated at all the voxels in the brain, resulting in 780 (cross-validation) * 212,018 (voxels) * 90 (participants) classifications. For the inter-language analyses, the fMRI vectors of one cultural group were modeled with story vectors generated from a different language. To be more specific, we used the American subjects' fMRI data, which were recorded while reading stories in English, to map to Farsi and Mandarin story vectors; the Iranian subjects' fMRI data, recorded while reading stories in Farsi, to map to English and Mandarin story vectors; and the Chinese subjects' fMRI data, recorded while reading stories in Mandarin, to map to English and Farsi story vectors. The inter-language analysis added a factor of three to the number of analyses, resulting in total of 44,650,990,800 classifications. The analysis was implemented using the *scikit-learn* library in python (Pedregosa et al., 2011), and was ran on 30 Google Computing Engine machines⁶, each with 4 CPUs and 26GB of RAM, and it took about six weeks complete.

The result of this analysis is a map of story-prediction accuracy. In other words, the above method scores each voxel neighborhood based on how accurately it can predict the story vectors. To assess the statistical significance of these accuracy values, non-parametric one-sample *t*-tests were performed with the FSL Randomise program (Winkler, Ridgway, Webster, Smith, & Nichols, 2014). For a given classification, the accuracy maps from 30 participants were concatenated and then binarized, expressed as 30 positive or negative signs indicating above- or below-chance decoding. Under the null

⁶<https://cloud.google.com/compute/>

hypothesis, positive and negative signs are equally likely to occur and may be randomly flipped with no effect on their distribution. Random flipping occurs for the set of signs associated with each voxel, followed by the calculation of its t -statistic. We performed 5mm HWHM spatial smoothing of the variance (over subjects), a commonly included step for increasing study power. This is performed for all voxels in the map, creating a permuted t -map. The t -map then undergoes TFCE, a procedure that further improves power by taking into account the underlying extent of spatial support. TFCE enhances values within cluster-like structures, while preserving voxel-level inference (Smith & Nichols, 2009). The TFCE map ultimately contributes a single value, its maximum, to an empirical null distribution of maximum TFCE values. This distribution was populated by repeating the permutation procedure 5000 times. The preceding steps were all implemented in a relatively standard one-line command to FSL Randomise (see Supplementary Materials). The value at the 95th percentile of the distribution was used to threshold the original, unpermuted data, thereby controlling the familywise error rate at the 0.05 level. To facilitate comparison across different classification types, we conservatively restricted visualization to the set of voxels that survived thresholding across all of the classification types. For example, we created a unified intra-language decoding map, selecting only voxels present in all three of the English, Chinese, and Farsi classification maps.

Results

The searchlight analyses yielded maps of brain regions that contained information regarding the content of stories. Decoding of stories presented in the participants' native language was successful in a set of spatially-defined brain regions (Figure 1). Maps for all three within-language classifications were similar (for individual intra-language classification maps see Supplementary Materials Figure 3). The most prominent cluster was in the bilateral posteromedial cortices, including the precuneus and posterior cingulate. Also on the medial surface was a bilateral cluster including the paracingulate gyrus, superior frontal gyrus and frontal pole. On the lateral surfaces, the

junction of the temporo-parieto-occipital lobes was implicated, specifically, a cluster centered on the angular gyrus bilaterally and extending into the posterior supramarginal gyrus bilaterally and superior lateral occipital cortex. Rostrally, somewhat lateralized to the right but also present on the left, were the superior frontal sulcus and middle frontal gyrus.

The brain activations of participants that were evoked by native-language stories were also successfully decoded using the representations of the same stories, translated into different, unfamiliar languages. We mapped this semantic information to a similar set of regions as found for native-language story decoding, namely the cortical midline structures and fronto-parietal regions Figure 2 (for individual inter-language maps see Supplementary Materials Figure 4).

Across all combinations of cultural backgrounds and story languages, the overall accuracies of each searchlight map were broadly similar (Table 1). The mean accuracies of significant voxels in each map were generally above 0.55 and the maximum accuracies on each map were all above 0.6 (For whole-brain analysis accuracies, please see Supplementary Materials Figure).

Discussion

In this study, we investigated neural representations that transcend narratives' language-specific words and syntactic features, and the degree to which these representations are systematic across people. Using distributed representations of narrative texts, we were able to decode neural activity corresponding to the reading of particular stories irrespective of the language being used to convey the story. Within- and between-language narrative decoding highlighted similar patterns of neural activity, with the highest classification performance in the DMN.

The spatial distribution of classification performance implicates the DMN in these high-level semantic representations. One of the distinguishing features of narrative comprehension is that it requires the integration of information over time. To grasp the meaning of a story we must continually stitch together individual words, link events

with their causes, and remember what came before to build a holistic understanding. While some aspects of semantic integration appear to happen very quickly (Christiansen & Chater, 2016; Hagoort, Hald, Bastiaansen, & Petersson, 2004), there is evidence that the DMN is involved in a process that unfolds over longer time scales (Hasson, Chen, & Honey, 2015; Honey, Chen, Müsch, & Hasson, 2016; Lerner et al., 2011). The relatively slow fluctuations found in the DMN (Zou et al., 2008) may therefore relate to a semantic integration process that extends over long time windows such as those required to understand a story. This is consistent with fMRI data showing that shared stimulus-locked activity in the DMN increases as longer segments of a story remain intact (Lerner et al., 2011; Simony et al., 2016), and with evidence that there are shared patterns of activity in the DMN across participants when describing scenes from an audiovisual story (Chen et al., 2016). The DMN has a profile of anatomical and functional connectivity (van den Heuvel & Sporns, 2013) that situates it to participate in the process of abstracting across information from a diversity of neural systems, and it may function as a very high-level convergence and divergence zone (Fernandino et al., 2015; Kaplan et al., 2016). In this hierarchical model of neural architecture, the DMN can be activated by, and in turn is able to reactivate, multiple lower-level brain systems that represent information in more explicit, mapped formats (Meyer & Damasio, 2009). Such properties may be instrumental to the DMN’s apparent role in constructing meaning that transcends the specifics of the language in which information is presented.

An essential component of this research was the use of high-dimensional representations of entire narratives (Le & Mikolov, 2014), rather than focusing on word-level and sentence-level semantics. Recently, Wehbe et al. (2014) showed that English language passage-level content from a Harry Potter story could be predicted from fMRI data. There are several key differences between our approach and the approach used by Wehbe et al. We exposed participants to completely novel 150-word stories in a naturalistic setting, while Wehbe et al. used short passages (16 words) from the same story that all participants had read before. Further, the method that we use for representing the stories is fully automatic and does not provide room for researcher

degree of freedom. In other words, we did not choose which set of variables to include and which ones not to include in the representation of the stories. This was done automatically by training the models on millions of other stories. The doc2vec method takes into account the context in which local lexical information should be interpreted, and it integrates the information at the lexical level to form higher order representations of the stories. To generate representations of passages Wehbe et al. relied on a wide variety of features ranging from low-level constructs, such as individual word properties, to higher-level constructs, such as character mentions and behavior. In contrast, while the distributed representations of the narratives used in this study likely capture some low-level semantic and syntactic features, previous research has demonstrated that they also reliably capture document-level features that cannot be reduced to a simple concatenation of lower-level entities (Dai et al., 2015; Lau & Baldwin, 2016; Le & Mikolov, 2014)⁷. Thus, by modeling narratives with document-level embeddings that can capture story-level meaning and using these representations for cross-linguistic decoding, we were able to target the neural encoding of story-level semantics.

While previous work has investigated cross-linguistic decoding, it has focused on translating semantic representations of individual words (Buchweitz et al., 2012; Correia et al., 2014; Correia, Jansma, Hausfeld, Kikkert, & Bonte, 2015; Zinszer, Anderson, Kang, Wheatley, & Raizada, 2016) rather than higher-level semantic constructs. Most recently, Zinszer et al. (2016) showed that neuro-semantic representations of words are preserved across both individuals and languages by translating English words into their Chinese counterparts using brain activation of independent groups of native English and Chinese speakers. Notably, other work has examined modal (Regev et al., 2013; Wilson, Molnar-Szakacs, & Iacoboni, 2008) and linguistic (Honey et al., 2012) representational invariance in narrative processing, however these studies have not investigated whether neural representations of story-level features can be decoded. Thus, while Honey et al. (2012) found that a set of regions showed similar activation

⁷Please refer to the Supplementary Materials for the additional experiments and analyses performed to validate the method used for calculating the distributed representations of stories.

patterns among both English and Russian-speaking participants who listened to time-matched recordings of the same narrative, the fidelity of the invariant representations has not previously been addressed.

Accordingly, while this research collectively makes a strong case for the cross-linguistic invariance of word-level semantic representation and suggests that invariance might exist for higher-level narrative constructs, the fidelity with which this higher-level representational invariance holds for higher-level semantic constructs has been less certain. In contrast, by recording participant responses to multiple narratives, we were able to explore not only cross-linguistic invariance, but also test whether fMRI data during story reading contained sufficient information for identifying which narrative the activity corresponded to. Once we learned the relationship between the neural maps and the story representations, we were able to predict the representation of the left-out stories based just on fMRI activity data.

To our knowledge, this research is the first to demonstrate within and cross-linguistic neural decoding of entire narrative sequences. By conducting cross-language decoding, we were able to functionally control for low-level, language-specific semantic constructs. If the accuracies of the within-language story predictions were contingent on the variation of specific lexical or syntactic features captured in the story representations, a substantial drop in accuracy should have been observed in the cross-language predictions. However, the accuracies of the within- and between-language predictions were comparable, which indicates that these predictions exploited systematic encodings of story-level elements that transcended the idiosyncratic lexical and syntactic features of individual languages.

Overall, our research is an attempt to demonstrate that abstracted beyond the level of independent concepts and language units, the brain seems to systematically encode high-level narrative elements. Further, despite the remarkable variety of human language, the combination of a shared cognitive architecture and overlapping socio-cultural experiences produces a remarkable cross-language consistency not only in the words and concepts that we use, but also in the narratives that we construct. Our

results indicate that this similarity is echoed in the neuro-semantic representation of narrative-level information.

Author Contribution

MD and JTK developed the concept and were involved in all phases of the study. Data analysis was performed by RB. KM, JH and JDZ were involved in interpretation of the results and writing of the manuscript. SIG contributed to the study design and data collection. AV was consulted on the computational analysis of the data. ASG contributed to the study design and assisted with NLP model training. AD was involved in the study design and contributed to the interpretation of the study. All authors approved the final version of the manuscript for submission.

References

- AleAhmad, A., Amiri, H., Darrudi, E., Rahgozar, M., & Oroumchian, F. (2009).
Hamshahri: A standard persian text collection. *Knowledge-Based Systems*, 22(5),
382–387.
- Ames, D. L., Honey, C. J., Chow, M. A., Todorov, A., & Hasson, U. (2015). Contextual
alignment of cognitive and neural dynamics. *Journal of cognitive neuroscience*.
- Binder, J. R., Desai, R. H., Graves, W. W., & Conant, L. L. (2009). Where is the
semantic system? a critical review and meta-analysis of 120 functional
neuroimaging studies. *Cerebral Cortex*, 19(12), 2767–2796.
- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. *the Journal
of machine Learning research*, 3, 993–1022.
- Buchweitz, A., Shinkareva, S. V., Mason, R. A., Mitchell, T. M., & Just, M. A. (2012).
Identifying bilingual semantic neural representations across languages. *Brain and
language*, 120(3), 282–289.
- Buckner, R. L., Andrews-Hanna, J. R., & Schacter, D. L. (2008). The brain’s default
network. *Annals of the New York Academy of Sciences*, 1124(1), 1–38.
- Chen, J., Leong, Y. C., Honey, C. J., Yong, C. H., Norman, K. A., & Hasson, U.
(2016). Shared memories reveal shared structure in neural activity across
individuals. *Nature Neuroscience*.
- Christiansen, M. H., & Chater, N. (2016). The now-or-never bottleneck: A fundamental
constraint on language. *Behavioral and Brain Sciences*, 39, e62.
- Correia, J., Formisano, E., Valente, G., Hausfeld, L., Jansma, B., & Bonte, M. (2014).
Brain-based translation: fmri decoding of spoken words in bilinguals reveals
language-independent semantic representations in anterior temporal lobe. *The
Journal of Neuroscience*, 34(1), 332–338.
- Correia, J., Jansma, B., Hausfeld, L., Kikkert, S., & Bonte, M. (2015). Eeg decoding of
spoken words in bilingual listeners: from words to language invariant
semantic-conceptual representations. *Frontiers in psychology*, 6, 71.
- Dai, A. M., Olah, C., & Le, Q. V. (2015). Document embedding with paragraph

vectors. *arXiv preprint arXiv:1507.07998*.

- Damasio, H., Tranel, D., Grabowski, T., Adolphs, R., & Damasio, A. (2004). Neural systems behind word and concept retrieval. *Cognition*, *92*(1), 179–229.
- Deerwester, S. C., Dumais, S. T., Landauer, T. K., Furnas, G. W., & Harshman, R. A. (1990). Indexing by latent semantic analysis. *JAsIs*, *41*(6), 391–407.
- Fernandino, L., Binder, J. R., Desai, R. H., Pendl, S. L., Humphries, C. J., Gross, W. L., ... Seidenberg, M. S. (2015). Concept representation reflects multimodal abstraction: A framework for embodied semantics. *Cerebral Cortex*, bhv020.
- Ferstl, E. C., Neumann, J., Bogler, C., & Von Cramon, D. Y. (2008). The extended language network: a meta-analysis of neuroimaging studies on text comprehension. *Human brain mapping*, *29*(5), 581–593.
- Fyshe, A., Talukdar, P. P., Murphy, B., & Mitchell, T. M. (2014). Interpretable semantic vectors from a joint model of brain-and text-based meaning. In *Proc. of acl*.
- Gordon, A., Huangfu, L., Sagae, K., Mao, W., & Chen, W. (2013). Identifying personal narratives in chinese weblog posts. In *Intelligent narrative technologies workshop, boston, ma*.
- Gordon, A., & Swanson, R. (2009). Identifying personal stories in millions of weblog entries. In *Third international conference on weblogs and social media, data challenge workshop, san jose, ca* (Vol. 46).
- Grabowski, T. J., Damasio, H., Tranel, D., Ponto, L. L. B., Hichwa, R. D., & Damasio, A. R. (2001). A role for left temporal pole in the retrieval of words for unique entities. *Human brain mapping*, *13*(4), 199–212.
- Guimaraes, M. P., Wong, D. K., Uy, E. T., Grosenick, L., & Suppes, P. (2007). Single-trial classification of meg recordings. *Biomedical Engineering, IEEE Transactions on*, *54*(3), 436–443.
- Hagoort, P., Hald, L., Bastiaansen, M., & Petersson, K. M. (2004). Integration of word meaning and world knowledge in language comprehension. *science*, *304*(5669), 438–441.

- Hasson, U., Chen, J., & Honey, C. J. (2015). Hierarchical process memory: memory as an integral component of information processing. *Trends in cognitive sciences*, 19(6), 304–313.
- Honey, C. J., Chen, J., Müsch, K., & Hasson, U. (2016). How long is now? the multiple timescales of language processing. *Behavioral and Brain Sciences*, 39.
- Honey, C. J., Thompson, C. R., Lerner, Y., & Hasson, U. (2012). Not lost in translation: neural responses shared across languages. *The Journal of Neuroscience*, 32(44), 15277–15283.
- Huth, A. G., de Heer, W. A., Griffiths, T. L., Theunissen, F. E., & Gallant, J. L. (2016). Natural speech reveals the semantic maps that tile human cerebral cortex. *Nature*, 532(7600), 453–458.
- Huth, A. G., Nishimoto, S., Vu, A. T., & Gallant, J. L. (2012). A continuous semantic space describes the representation of thousands of object and action categories across the human brain. *Neuron*, 76(6), 1210–1224.
- Johnson-Laird, P. N. (1983). *Mental models: Towards a cognitive science of language, inference, and consciousness* (No. 6). Harvard University Press.
- Just, M. A., Cherkassky, V. L., Aryal, S., & Mitchell, T. M. (2010). A neurosemantic theory of concrete noun representation based on the underlying brain codes. *PloS one*, 5(1), e8622.
- Kaplan, J. T., Gimbel, S. I., Dehghani, M., Immordino-Yang, M. H., Sagae, K., Wong, J. D., ... Damasio, A. (2016). Processing narratives concerning protected values: A cross-cultural investigation of neural correlates. *Cerebral Cortex*, bhv325.
- Kintsch, W. (1974). *The representation of meaning in memory*. Lawrence Erlbaum.
- Kintsch, W. (1988). The role of knowledge in discourse comprehension: a construction-integration model. *Psychological review*, 95(2), 163.
- Lau, J. H., & Baldwin, T. (2016). An empirical evaluation of doc2vec with practical insights into document embedding generation. *arXiv preprint arXiv:1607.05368*.
- Le, Q. V., & Mikolov, T. (2014). Distributed representations of sentences and documents. In *Icml* (Vol. 14, pp. 1188–1196).

- Lerner, Y., Honey, C. J., Silbert, L. J., & Hasson, U. (2011). Topographic mapping of a hierarchy of temporal receptive windows using a narrated story. *The Journal of Neuroscience*, *31*(8), 2906–2915.
- Meyer, K., & Damasio, A. (2009). Convergence and divergence in a neural architecture for recognition and memory. *Trends in neurosciences*, *32*(7), 376–382.
- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Mitchell, T. M., Shinkareva, S. V., Carlson, A., Chang, K.-M., Malave, V. L., Mason, R. A., & Just, M. A. (2008). Predicting human brain activity associated with the meanings of nouns. *science*, *320*(5880), 1191–1195.
- Østby, Y., Walhovd, K. B., Tamnes, C. K., Grydeland, H., Westlye, L. T., & Fjell, A. M. (2012). Mental time travel and default-mode network functional connectivity in the developing brain. *Proceedings of the National Academy of Sciences*, *109*(42), 16800–16804.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., . . . Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, *12*, 2825–2830.
- Qin, P., & Northoff, G. (2011). How is our self related to midline regions and the default-mode network? *Neuroimage*, *57*(3), 1221–1233.
- Raichle, M. E. (2015). The brain’s default mode network. *Annual review of neuroscience*, *38*, 433–447.
- Raichle, M. E., & Snyder, A. Z. (2007). A default mode of brain function: a brief history of an evolving idea. *Neuroimage*, *37*(4), 1083–1090.
- Regev, M., Honey, C. J., Simony, E., & Hasson, U. (2013). Selective and invariant neural responses to spoken and written narratives. *The Journal of Neuroscience*, *33*(40), 15978–15988.
- Řehůřek, R., & Sojka, P. (2010, May 22). Software Framework for Topic Modelling with Large Corpora. In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks* (pp. 45–50). Valletta, Malta: ELRA.

(<http://is.muni.cz/publication/884893/en>)

- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, *323*, 533–536.
- Sagae, K., Gordon, A. S., Dehghani, M., Metke, M., Kim, J. S., Gimbel, S. I., . . . Immordino-Yang, M. H. (2013). A data-driven approach for classification of subjectivity in personal narratives. In *Cmn* (pp. 198–213).
- Sanford, A. J., & Garrod, S. C. (1998). The role of scenario mapping in text comprehension. *Discourse processes*, *26*(2-3), 159–190.
- Schuster, S., Hawelka, S., Hutzler, F., Kronbichler, M., & Richlan, F. (2016). Words in context: The effects of length, frequency, and predictability on brain responses during natural reading. *Cerebral Cortex*.
- Shinkareva, S. V., Malave, V. L., Mason, R. A., Mitchell, T. M., & Just, M. A. (2011). Commonality of neural representations of words and pictures. *Neuroimage*, *54*(3), 2418–2425.
- Simony, E., Honey, C. J., Chen, J., Lositsky, O., Yeshurun, Y., Wiesel, A., & Hasson, U. (2016). Dynamic reconfiguration of the default mode network during narrative comprehension. *Nature Communications*, *7*.
- Smallwood, J., & Schooler, J. W. (2015). The science of mind wandering: empirically navigating the stream of consciousness. *Annual review of psychology*, *66*, 487–518.
- Smallwood, J., Tipper, C., Brown, K., Baird, B., Engen, H., Michaels, J. R., . . . Schooler, J. W. (2013). Escaping the here and now: evidence for a role of the default mode network in perceptually decoupled thought. *Neuroimage*, *69*, 120–125.
- Smith, S. M., & Nichols, T. E. (2009). Threshold-free cluster enhancement: addressing problems of smoothing, threshold dependence and localisation in cluster inference. *Neuroimage*, *44*(1), 83–98.
- Spreng, R. N. (2012). The fallacy of a “task-negative” network. *Frontiers in psychology*, *3*, 145.
- Sudre, G., Pomerleau, D., Palatucci, M., Wehbe, L., Fyshe, A., Salmelin, R., &

- Mitchell, T. (2012). Tracking neural coding of perceptual and semantic features of concrete nouns. *NeuroImage*, 62(1), 451–463.
- Tausczik, Y. R., & Pennebaker, J. W. (2010). The psychological meaning of words: Liwc and computerized text analysis methods. *Journal of language and social psychology*, 29(1), 24–54.
- Tranel, D., Damasio, H., & Damasio, A. R. (1997). A neural basis for the retrieval of conceptual knowledge. *Neuropsychologia*, 35(10), 1319–1327.
- van den Heuvel, M. P., & Sporns, O. (2013). Network hubs in the human brain. *Trends in cognitive sciences*, 17(12), 683–696.
- Wehbe, L., Murphy, B., Talukdar, P., Fyshe, A., Ramdas, A., & Mitchell, T. (2014). Simultaneously uncovering the patterns of brain regions involved in different story reading subprocesses. *PloS one*, 9(11), e112575.
- Whitney, C., Huber, W., Klann, J., Weis, S., Krach, S., & Kircher, T. (2009). Neural correlates of narrative shifts during auditory story comprehension. *Neuroimage*, 47(1), 360–366.
- Wilson, S. M., Molnar-Szakacs, I., & Iacoboni, M. (2008). Beyond superior temporal cortex: intersubject correlations in narrative speech comprehension. *Cerebral Cortex*, 18(1), 230–242.
- Winkler, A. M., Ridgway, G. R., Webster, M. A., Smith, S. M., & Nichols, T. E. (2014). Permutation inference for the general linear model. *Neuroimage*, 92, 381–397.
- Yarkoni, T., Speer, N. K., & Zacks, J. M. (2008). Neural substrates of narrative comprehension and memory. *Neuroimage*, 41(4), 1408–1425.
- Yogatama, D., Faruqui, M., Dyer, C., & Smith, N. A. (2014). Learning word representations with hierarchical sparse coding. *arXiv preprint arXiv:1406.2035*.
- Zinszer, B. D., Anderson, A. J., Kang, O., Wheatley, T., & Raizada, R. (2015). You say potato, i say tǔdòu: How speakers of different languages share the same concept. In *Proceedings of the 37th annual conference of the cognitive science society*.
- Zinszer, B. D., Anderson, A. J., Kang, O., Wheatley, T., & Raizada, R. D. (2016). Semantic structural alignment of neural representational spaces enables translation

between english and chinese words. *Journal of Cognitive Neuroscience*.

Zou, Q.-H., Zhu, C.-Z., Yang, Y., Zuo, X.-N., Long, X.-Y., Cao, Q.-J., ... Zang, Y.-F. (2008). An improved approach to detection of amplitude of low-frequency fluctuation (alff) for resting-state fmri: fractional alff. *Journal of neuroscience methods*, 172(1), 137–141.

Table 1

Decoding accuracy of story vector models for three languages, based on the fMRI activations of three cultural groups. Coordinates are in X,Y,Z MNI space

Culture	Story Language	Mean Accuracy of Significant Voxels	Maximum Accuracy among Searchlight Spheres	Coordinates of Center Voxel of Maximum Searchlight Sphere
Americans	English	0.566	0.622	48, 40, 50
	Farsi	0.575	0.624	48, 39, 50
	Mandarin	0.579	0.637	44, 40, 54
Iranians	English	0.565	0.623	68, 26, 47
	Farsi	0.549	0.605	48, 33, 48
	Mandarin	0.553	0.607	67, 26, 47
Chinese	English	0.562	0.622	71, 29, 51
	Farsi	0.55	0.608	70, 26, 55
	Mandarin	0.552	0.611	67, 22, 53

Figure 1. Intra-language searchlight maps. Colored voxels indicate regions in which English, Farsi, and Chinese stories, presented to participants in their native languages, were all successfully decoded. The most prominent cluster was found in the posteromedial cortices, bilaterally. Other clusters on the medial surface included the superior frontal gyrus, paracingulate gyrus, and frontal pole. On the lateral surface, a prominent cluster was centered on the angular gyrus.

Figure 2. Inter-language searchlight maps. Colored voxels indicate regions in which story representations, which were generated from stories in a language unfamiliar to the participant, were successfully decoded in English, Farsi, and Chinese speakers. A similar set of regions were found as for intra-language decoding, including the cortical midline structures and fronto-parietal regions.