Narratives as Networks: Predicting Memory from the Structure of Naturalistic Events

Hongmi Lee (hongmi.lee@jhu.edu)
Department of Psychological and Brain Sciences, Johns Hopkins University,
3400 N Charles Street, Baltimore, MD 21218, United States

Janice Chen (janice@jhu.edu)
Department of Psychological and Brain Sciences, Johns Hopkins University,
3400 N Charles Street, Baltimore, MD 21218, United States

Abstract:

A naturalistic spoken recall paradigm with fMRI was used to investigate how the structure of a complex realistic experience affects memory. Subjects watched short movie clips and then verbally recalled the movie details aloud while being scanned. To quantify the structure of the movies, we transformed each movie plot into a network, where nodes are individual movie events and the connections between them are determined by content similarity. Inter-event similarity was computed by correlating high-dimensional sentence embeddings derived from human-generated text descriptions of the movie events. Behavioral results showed that the centrality of events within the network (i.e., the overall number and strength of connections with other events) positively predicted recall performance. Higher centrality also predicted stronger univariate activation and the reactivation of event-specific multi-voxel patterns during recall in the posterior medial cortex, a brain area thought to represent abstract 'situation models.' Representational similarity analysis revealed that the neural pattern similarity structure of default network areas during recall reflected the text-based narrative network structure. Our study introduces a novel approach to quantify the structure of complex narratives and demonstrates that inter-event structure predicts behavioral and neural markers of memory under naturalistic conditions.

Keywords: memory; recall; fMRI; narrative; network

Introduction

For several decades, the dominant approach to studying human memory recall has depended on trial-based learning paradigms where subjects are presented with lists of random isolated words or pictures (e.g., Murdock, 1962). However, in daily life we are more likely to encounter structured material than random lists of items. A more typical real-world use of memory is describing a recent experience in one's own words; for example, relating the events of one's day to a friend over dinner. This form of extended narration-like recall does not strictly follow the principles of isolated-item list learning and memory (Chen et al., 2017; Heusser, Fitzpatrick, & Manning, 2018; Kintsch, 1982).

Unlike a single item in a list learning paradigm, complex real-life experiences consist of multiple inter-related components or events, which often form coherent narratives. Earlier work on narrative processing showed that the canonical structure of stories (i.e., conflict and resolution) and causal relationships between events affect comprehension and memory for text passages (e.g., Trabasso & van den Broek, 1985). While the narratives used in these studies were certainly more complex than random word lists, these experiments were still limited in that they relied either on relatively short and carefully designed episodes, or on extensive manual analysis of existing text documents.

Here, we introduce an approach for quantifying and exploring narrative structure which is easily generalizable to different types of narratives. In this method, each narrative (movie plot) is automatically transformed into a network of interconnected events based on semantic similarity measured from sentence embedding distances; embeddings were generated from a deep learning based natural language model. We show that narrative structure can be used to predict memory behavior and brain responses during unguided spoken recall, as people recount the complex movie events they have seen freely from memory.

Methods

Fifteen subjects watched a series of 10 short movie clips and then verbally recalled the movie plots in their own words irrespective of the order of presentation. Both tasks were performed while the subjects were being scanned. The movies were on average 4.62 minutes long and had narratives which varied in content and structure.

To quantify and assess the inter-event structure of the movie plots, we employed an approach wherein we transformed a narrative into a graph/network (Figure 1). In this narrative network, the events that constitute a movie plot (nodes) form connections with each other (edges), and the connection strength between a pair of

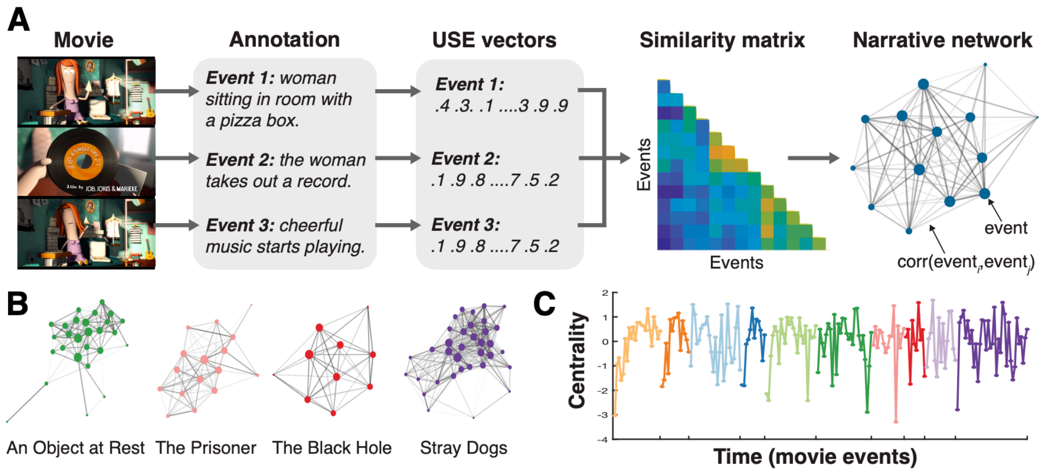

Figure 1. (A) Schematic of the narrative network approach. (B) Example narrative networks of four movies used in the experiment. The strength of connections between events was thresholded at r = .6 for visualization purposes. (C) Time course of the event centrality (PageRank) for all 10 movies used in the experiment. Each data point represents a single movie event. Different colors indicate different movies.

events is determined by their content similarity (edge weight). To build a narrative network for each movie, human annotators first segmented the movies into shorter events (20 events per movie on average, 202 events in total) and provided written descriptions for each event. The text descriptions were in turn encoded into high-dimensional vectors with Google's Universal Sentence Encoder (USE; Cer et al., 2018). The connection strength between event node pairs was computed by correlating their corresponding USE vectors. The 'centrality' or importance of each individual event was defined as the PageRank of each node in the network, such that events with stronger and greater number of connections with other events have higher centrality.

Each fMRI subject's recall speech was manually transcribed, segmented into discrete utterances, and matched to the movie events to identify which specific movie events were recalled. The recall text segments were also transformed into USE vectors to be compared with the vectors based on movie annotations.

Functional MRI data were first motion and distortion corrected using FSL (https://fsl.fmrib.ox.ac.uk/fsl) and then projected onto a standard template surface (fsaverage6) and smoothed using FreeSurfer (https://surfer.nmr.mgh.harvard.edu/). The functional data were additionally high-pass filtered, z-scored within each scanning run, and shifted by 3 TRs (TR = 1.5 s) to account for the delay in hemodynamic response.

Behavioral Results

Consistent with prior studies (Chen et al., 2017; Heusser, Fitzpatrick, & Manning, 2018), naturalistic narrative recall behavior showed characteristics distinct from free recall behavior in list-learning experiments. Subjects on average recalled 9 of the 10 movies and spent 3.27 minutes describing each movie. The recall order of movie events very strictly followed their chronological order within each movie (ρ = .97; Figure 2A), whereas the recall at the movie level was occasionally out of order (ρ = .52). The recall of a movie was graded rather than all-or-none, as the proportion of scenes recalled per movie varied across movies in each subject (M = 76.9 %). Recall probability was not higher for the events at the beginning and end of each movie (Figure 2B), contrary to strong serial position effects observed in traditional list-learning experiments (Murdock, 1962). Although there was no clear effect of temporal order, inter-event connections in the narrative network significantly affected recall performance: high centrality events were more likely to be recalled than low centrality events (Figure 2C; t_{14} = 5.14, p < .001). Event centrality was also positively correlated with recall accuracy (Figure 2D; r = .17, p = .014), measured as the similarity between USE vectors from the recall speech transcription and the movie scene description of the corresponding events.

fMRI Results

We next examined the relationship between event centrality and neural responses during recall of the

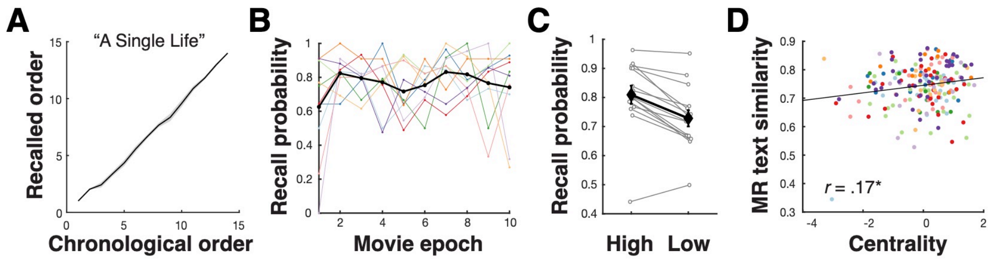

Figure 2. Behavioral results. (A) Recall order rank of movie events (1 = recalled first ~ 14 = recalled last, regardless of how many events a subject recalled) in an example movie ("A Single Life") as a function of their chronological order. Shaded areas indicate SEM across subjects. (B) Recall probability (proportion of events recalled) as a function of epochs in the movie, averaged across subjects. Colored lines represent individual movies. The black line represents the average across all 10 movies. (C) Recall probability for high vs. low centrality events based on the median split of movie events' PageRank centrality. Recall probability was computed separately for each movie and then averaged across movies for each subject. Gray lines show individual subjects' data. The black line shows the average across subjects. Error bars represent SEM across subjects. (D) Relationship between the event centrality and the Movie-Recall text similarity. Each dot represents a movie event. Different colors of the dots represent different movies.

events. A whole-brain univariate analysis was performed to identify brain regions whose activation scaled with the event centrality. For each point on the brain surface of each subject, we computed the mean activation of each movie event recalled by the subject. Event centrality was next regressed onto the event-by-event activation. A one-sample *t*-test against 0 was applied to the resulting parameter estimate maps from all subjects, producing a group-level *t* statistic map (Figure 3A). The regions whose activation showed the strongest positive relationship with event centrality were posterior medial cortex (PMC) and angular gyrus, previously known to show increased activation during memory recollection (the general recollection network; Rugg & Vilberg, 2013).

We next tested whether event centrality predicts the strength of event-specific memory reactivation using multivariate pattern analysis. Memory reactivation was measured as between-subjects similarity across event-specific brain activation patterns during movie watching and recall (Chen et al., 2017). We found that the event centrality was positively correlated with the event-specific memory reactivation in PMC (Figure 3B; t_{14} = 3.16, $p < .01$), an area indicated in the univariate analysis above and also thought to represent abstract 'situation models' of events (Ranganath & Ritchey, 2012). This positive relationship was not observed in low-level sensory areas such as auditory and motor-somatosensory cortices (p's > .7).

Finally, we performed whole-brain representational similarity analysis (Kriegeskorte, Mur, & Bandettini, 2008) to test whether the representational similarity structure based on brain activation patterns during recall reflects the narrative network structure based on the movie annotation text. We first divided each subject's cortical surface into 1000 parcels following Schaefer et al. (2018). For each parcel, we computed pairwise correlations between the activation patterns of recalled movie events in a between-subjects manner, separately for each movie. The correlation coefficients were concatenated across all movies and in turn correlated with the correlations between movie annotation USE vectors. One-sample *t*-tests were performed on the resulting brain-text correlations from all subjects. We found that parcels within and around PMC showed the highest brain-text representational similarity during recall (Figure 3C).

Conclusions and Future Directions

The current study demonstrated that unguided recall of a complex naturalistic experience is predicted by inter-event structure and has different properties from those observed in a traditional list learning paradigm. Using an approach of transforming narratives into networks, we showed that high event centrality, or rich interconnections with other events based on sentence embedding similarity, predicted better subsequent recall. In addition, high centrality was associated with higher activation and stronger event-specific pattern reactivation in recollection-related brain areas during recall. Neural patterns in high-level associative areas also exhibited a representational structure similar to the narrative network structure.

The link between high event centrality and behavioral and neural markers of memory may have several different but not mutually exclusive explanations. High

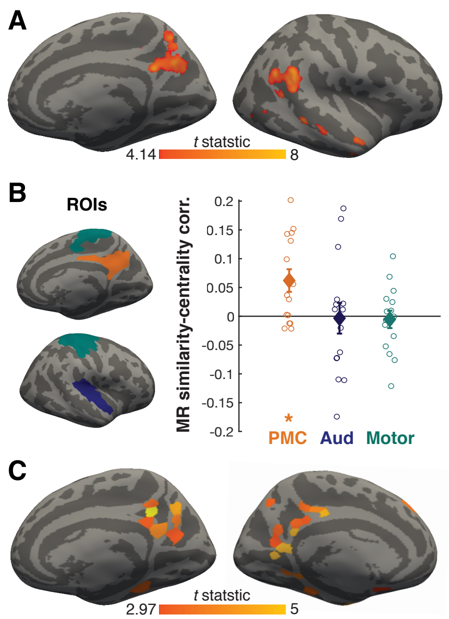

Figure 3. fMRI results. (A) The whole-brain surface map of brain areas where univariate activation during recall had a positive relationship with the centrality of the recalled events. $p < .001$ (uncorrected). (B) Three regions of interest (Left; orange = posterior-medial cortex, blue = auditory cortex, green = motor-somatosensory cortex) and the mean Pearson correlation between the event-specific centrality and movie-recall between-subjects pattern similarity (Right). Hollow circles represent individual subjects. Filled diamonds represent the average across subjects. Error bars indicate SEM across subjects. (C) The whole-brain surface map (medial view) of cortical parcels that show positive correlations between the event-wise movie annotation text similarity matrices and the recall fMRI pattern similarity matrices. $p < .01$ (uncorrected).

centrality events might be encoded more robustly during movie watching through repeated presentation and reactivation of the common semantic features that connect the events to other events. It is also possible that high centrality events are more easily accessed during recall as they are more likely to be cued by other related events. Further analyses will examine these and other potential mechanisms behind the centrality effects.

In the current study, connections between events in a narrative network were defined by the semantic similarity between the movie annotations of the events. However, narrative networks can be built based on other types of inter-event relationships, for example causality (Trabasso & van den Broek, 1985). Indeed, our preliminary analyses (not included here) suggest that if centrality is calculated from human ratings of causality between event pairs, this also predicts behavioral recall performance and memory reactivation measured as movie-recall fMRI pattern similarity. Further analyses will explore the relationship between different types of narrative structure and their relative contribution to brain responses during naturalistic recall.

Acknowledgments

We acknowledge the support of the Alfred P. Sloan Foundation (Fellowship to JC). We acknowledge the use of fMRI data collected under Princeton University's IRB #5516.

References

Cer, D., Yang, Y., Kong, S., Hua, N., Limtiaco, N., St. John, R., Constant, N., Guajardo-Cespedes, M., Yuan, S., Tar, C., Sung, Y-H., Strope, B., & Kurzweil, R. (2018). Universal sentence encoder. *arXiv,* Available at: http://arxiv.org/abs/1803.11175.

Chen, J., Leong, Y. C., Honey, C. J., Yong, C. H., Norman, K. A., & Hasson, U. (2017). Shared memories reveal shared structure in neural activity across individuals. *Nature Neuroscience, 20*, 115–125.

Heusser, A. C., Fitzpatrick, P. C., & Manning, J. R. (2018). How is experience transformed into memory? *bioRxiv,* Available at: https://doi.org/10.1101/409987.

Kintsch, W. (1982). Memory for text. In A. Flammer & W. Kintsch (Eds.), *Discourse Processing.* Amsterdam, Netherlands: North-Holland.

Kriegeskorte, N., Mur, M., & Bandettini, P. (2008). Representational similarity analysis — connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience, 2*, 4.

Murdock, B., B., Jr. (1962). The serial position effect in free recall. *Journal of Experimental Psychology, 64,* 482–488.

Ranganath, C., & Ritchey, M. (2012). Two cortical systems for memory-guided behaviour. *Nature Reviews Neuroscience, 13*, 713–726.

Rugg, M. D. & Vilberg, K. L. (2013). Brain networks underlying episodic memory retrieval. *Current Opinion in Neurobiology, 23,* 255–260.

Schaefer, A., Kong, R., Gordon, E. M., Laumann, T. O., Zuo, X.-N., Holmes, A. J., Eickhoff, S. B., & Yeo, B. T. T. (2018). Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity MRI. *Cerebral Cortex, 28*, 3095–3114.

Trabasso, T. & van den Broek, P. (1985). Causal thinking and the representation of narrative events. *Journal of Memory and Language, 24,* 612–630.