Semantic centrality captures key aspects of knowledge construction: Behavioral and neural evidence of learning from a STEM video lecture

Yeongji Lee (Yeongji.Lee.GR@dartmouth.edu)[1] & David J. M. Kraemer (David.J.M.Kraemer@dartmouth.edu)[1]
[1]Department of Psychological and Brain Sciences, Dartmouth College

Abstract

When learning from educational content, we must simultaneously remember individual pieces of information and integrate them into a coherent conceptual framework. Previous work has shown that semantic network structure predicts what people remember from narratives, but whether these same principles extend to academic learning remains unknown. In this study, participants watched a 15-minute physics lecture and completed free recall during fMRI scanning. Using Sentence-BERT embeddings to quantify semantic centrality, we found that central information was better recalled and, critically, participants who recalled more units of central information demonstrated better conceptual understanding as rated by human raters (*r* = .73). In a voxel-wise encoding analysis, we modeled brain responses during lecture viewing using the same SBERT features, and this brain-model mapping predicted behavioral outcomes in selective cortical regions. These findings suggest that the embedding space of large language models reflects relevant aspects of how our minds organize newly learned conceptual information.

Keywords: semantic centrality; conceptual learning; memory; Sentence-BERT; voxel-wise encoding; fMRI

Introduction

How do learners transform a continuous stream of information from lectures or textbooks into structured conceptual knowledge? Despite substantial progress in understanding the neural mechanisms of learning (Cetron et al., 2020; Meshulam et al., 2021, Nguyen et al., 2022), it remains unclear how the mind encodes individual pieces of information and selectively organizes them into a coherent conceptual framework.

Recent work on naturalistic narrative offers a promising approach. Lee and Chen (2022) demonstrated that the semantic structure of a narrative plays a central role in episodic memory: events with higher semantic centrality — more connections to other events in the semantic network — are better remembered and elicit stronger hippocampal encoding responses and greater convergence across individuals. Complementing this, Caucheteux et al. (2022) showed that semantic embeddings from a large language model predict brain responses to spoken stories, and that the degree of this brain-model mapping correlates with individual comprehension scores. Together, these findings suggest that semantic structure shapes not only what is remembered, but how incoming information is encoded during narrative comprehension.

However, whether these principles extend to educational context remains an open question. Academic learning differs fundamentally from narrative memory: it lacks clear event structure, conveys abstract concepts through verbal explanation and demonstration rather than rich sensory experience, and demands that learners construct principled understanding rather than recall what happened. These differences raise the question of whether the semantic structure governs encoding and knowledge construction even in the absence of narrative structure.

In this study, we examined how semantic structure shapes the encoding and conceptual understanding of educational content at both behavioral and neural levels. Participants watched a 15-minute physics lecture and completed free recall tasks during fMRI scanning. We used Sentence-BERT embeddings to quantify the semantic centrality of each information unit in the lecture and to model neural responses during learning via voxel-wise encoding. Our behavioral results demonstrate that central information is better remembered, and that recall of central information predicts conceptual understanding. Our neural findings reveal that the same semantic features explain neural responses across distributed cortical regions, and that the degree of this brain-model mapping predicts individual learning outcomes. Together, these findings suggest that the representational space of large language models captures fundamental aspects of how knowledge is organized during learning—both in memory and in the brain.

Method

Participant

54 Dartmouth College students participated in this study (32 females; mean age = 22.30 years, SD = 7.76). Participants reported moderate prior knowledge of gravitational force (M = 3.50, SD = 1.09) but lower familiarity with both aerodynamic drag (M = 2.54, SD = 0.86) and terminal velocity (M = 2.60, SD = 0.98) on a 5-point scale.

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

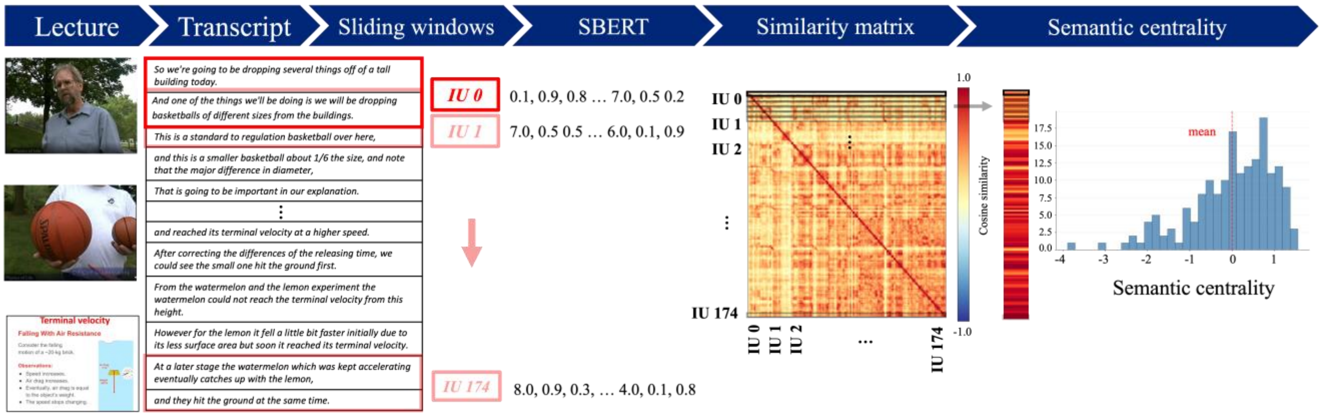

Figure 1. The procedure of calculating semantic centrality of information units in the lecture

Procedure

Prior to scanning, all participants completed a short quiz assessing their prior knowledge of physics concepts related to the lecture. Participants then completed five functional runs inside the fMRI scanner. The first run consisted of five minutes of eyes-open rest. During the following three runs, participants watched a physics video teaching terminal velocity and related Newtonian concepts. Upon completion of the video, participants completed an additional five minutes of eyes-open rest. In the final run, participants completed two free recall tasks. In both recall tasks, participants were instructed to recall everything they learned and remembered from the video in as much detail as possible, in any order. The first task was silent recall, in which participants recalled silently for five minutes without any verbal response. In the following verbal recall, participants described what they learned and remembered for a minimum of five minutes and up to ten minutes. After scanning, participants completed the same quiz administered before the scan to assess knowledge change. Resting-state and silent recall data were not included in the current analysis.

fMRI Data Acquisition

Brain images were acquired at the Dartmouth Brain Imaging Center using a 3 Tesla Siemens PRISMA scanner with a 32-channel head coil. A high-resolution T1-weighted anatomical scan was collected for each participant, followed by five functional runs ranging from 5 to 16 minutes. Functional images were acquired using a T2*-weighted echo-planar imaging (EPI) sequence (TR = 2000 ms; TE = 32 ms; flip angle = 79°; voxel size = 3 × 3 × 3 mm^3). Functional data were preprocessed using fMRIPrep version 21.0.1 (Esteban et al., 2019).

Stimuli

The stimuli (total duration: 15 min 45 sec) were adapted from a YouTube video (ESFTV, 2010) produced by Christopher Baycura and colleagues at the State University of New York College of Environmental Science and Forestry (SUNY-ESF). The original video showed a professor dropping six pairs of objects varying in size and mass to demonstrate the physics of falling objects and terminal velocity. We edited this material to create a structured learning sequence consisting of four demonstrations and the recap explanation, which was created by the researchers to summarize the definition of terminal velocity and review the key concepts from the lecture.

Semantic Centrality

To quantify and assess how the information was organized and interconnected within the lecture video, we calculated the semantic centrality of each information unit. Following Lee and Chen (2022), who used semantic embeddings to quantify inter-event structure in narratives, we adapted this approach to measure the interconnectedness of an educational video. The analysis consisted of four steps (Figure 1). First, we segmented the lecture transcript into 175 clauses using a rule-based algorithm. Next, we created overlapping sliding windows of two clauses with a step size of one clause, which we refer to as an "information unit" in this study. Third, we converted each information unit into a 768-dimensional numerical vector using Sentence-BERT (Reimers & Gurevych, 2019), a modified BERT model (Devlin et al., 2019) optimized for semantic comparison. Lastly, we computed pairwise cosine similarity between all information units, creating a 174 by 174 similarity matrix. For each information unit, we calculated degree centrality by summing its cosine similarity values to all other units. This measure quantifies how semantically connected each unit is to the rest of the lecture content. We normalized these values by dividing each by the total sum of all degree centrality values and multiplying by 100, then z-scored the resulting distribution. These standardized values constitute the semantic centrality measure used in all subsequent analyses.

Measuring Item Level Memory

To measure how much information was retrieved from each participant's verbal recall, we used a cosine similarity to determine each information unit's recall in their recall transcription. We transcribed their verbal recall, then created sliding windows of clauses using the same procedure as the lecture video. Then these sliding windows were converted into semantic embeddings using the Sentence-BERT model. To identify the recall unit that corresponds to the units of the lecture, every unit of a recall transcription was compared with every information unit of the lecture. In this study, we used the threshold of cosine similarity of 0.72 to classify the recall units that are similar enough to each lecture unit. To validate this threshold, a human rater examined recall-lecture pairs across a range of cosine similarity values (0.7 – 0.8) and judged whether each pair was semantically equivalent; 0.72 best aligned with these judgments. For this validation, recall transcriptions were segmented into sliding windows of two sentences rather than two clauses.

Measuring Conceptual Understanding

In order to measure individuals' conceptual understanding, two raters quantified their level of knowledge reflected in participants' free recall transcriptions. For consistent assessment, we developed a rubric that characterized the different levels of understanding for each concept that was taught in the lecture. The rubric comprised four domains that were mainly derived from the core concepts and their mechanistic relationships to the demonstrations. Critically, the rubric was designed to assess the quality of reasoning rather than the quantity of information recalled—each domain evaluated whether participants demonstrated clear, logical reasoning about underlying mechanisms, independent of how many information units they reproduced. Two raters evaluated each recall transcript on a 15-point scale following the rubric, and the ratings were averaged across the raters for analyses. The raters showed strong agreement over the 54 ratings, with a Krippendorff's alpha of .789.

Voxel-wise Encoding Model

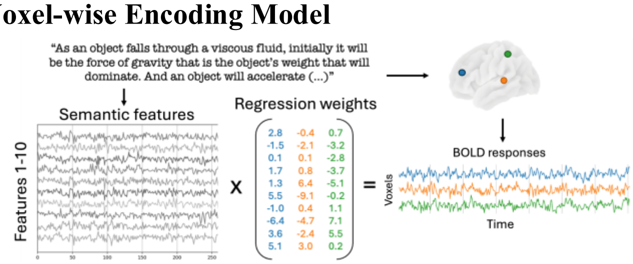

Figure 2. Voxel-wise modeling. For each participant, we fit regularized linear regression to estimate each voxel response.

A subset of 26 subjects were included in the fMRI analysis. For the semantic features, we z-scored each of the 768 SBERT contextualized semantic features within each run. For fMRI data, we applied a cerebral cortex mask and performed signal cleaning using Nilearn's signal.clean function (Abraham et al., 2014). Specifically, we applied linear detrending, high-pass filtering (cosine basis, cutoff = 0.01 Hz), and regressed out nuisance covariates (global signal, cerebrospinal fluid signal, six motion parameters, and spike regressors identified using nltools with a 3 SD threshold; Chang et al., 2018). The clean signal was z-scored within each run, and 5 time points were trimmed from the beginning and end of each run to account for temporal edge effects.

We used voxel-wise encoding to model how 768 SBERT features relate to BOLD responses in each cortical voxel and in each individual subject (Figure 2). A separate linear temporal filter with five delays (1, 2, 3, and 4 TRs) was fit for each of these 768 features to account for hemodynamic response, yielding 3,840 features in total. For each participant, data during video watching were split into training (80%) and test (20%) sets. The regularization parameter was selected via nested cross-validation on the training set; a single regularization coefficient was used across all voxels, determined by bootstrapping the regression procedure 40 times per participant.

To validate the voxel-wise models, estimated semantic feature weights were used to predict responses to a separate testing dataset that had not been used for weight estimation. Model performance was assessed as the Pearson correlation between predicted and actual responses on held-out test data. Statistical significance was determined by comparing correlations to a null distribution generated from independent Gaussian random vectors of the same length.

Results

The Information Unit With Stronger Semantic Connection Is Better Remembered

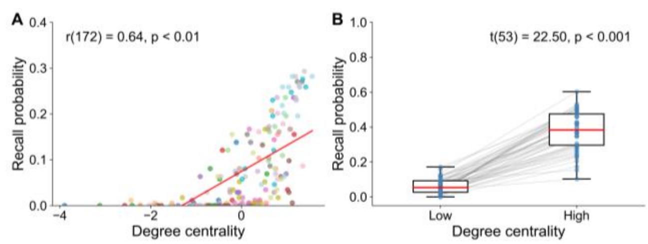

Figure 3. Semantic centrality predicts recall probability. (A). Each point represents an information unit. Recall probability is calculated by the number of participants who recalled the unit divided by the total number of participants. (B). Individual recall rates for high and low centrality information. Each line represents one participant (n = 54), showing the proportion of units recalled in each centrality group.

We asked whether the semantic centrality of each item predicts how likely participants are to remember it. For each unit, we computed recall probability as the proportion of participants who successfully recalled it during free recall.

Semantic centrality was a strong predictor of recall probability ($r(172) = .60$, $p < 0.01$, 95% confidence interval CI = [.54,.72]), consistent with the centrality-recall relationship previously observed in Lee and Chen's narrative study (r = .56) (2022). That this effect is of comparable magnitude in an educational context — despite the absence of narrative event structure — suggest that semantic organization may be a domain general principle of what gets encoded.

At the individual level, we split information units into high and low centrality groups at the mean (high: n = 98 units; low: n = 76 units) and computed each participant's recall rate within each group. High centrality units were recalled at substantially higher rates (M = 0.38; SD = 0.11) than low centrality units (M = 0.06; SD = 0.04); ($t(53) = 22.50$, p < 0.001, Cohen's $d_z = 3.06$, CI of the difference = [0.29, 0.35]; Figure 3B). Strikingly, low centrality units — which comprised nearly half the lecture content — were barely recalled by any participant, irrespective of individual conceptual understanding scores. Taken together, these findings indicate that the semantic structure of the lecture, rather than individual differences in attention or ability alone, strongly shapes what is encoded into memory.

Recall Of Central Information Units Predicts Better Conceptual Understanding

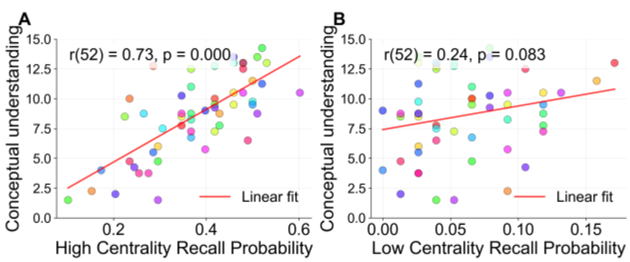

Figure 4. Recall probability of central information predicts conceptual understanding while recall probability of non-central information didn't. (A). The relationship between recall of high centrality units and conceptual understanding. (B). The relationship between recall of low centrality units and conceptual understanding.

Next, we asked whether this semantic centrality effect shapes learning outcomes in individuals. Participants first showed reliable knowledge gains: post-lecture quiz scores (M = 15.00; SD = 2.46) were significantly higher than pre-lecture scores (M = 12.61; SD = 2.85); ($t(53) = 7.97$, $p < 0.001$, Cohen's d_z = 1.09, CI of the difference = [1.79, 2.99]), confirming successful knowledge acquisition.

We then tested whether individual differences in how much participants remembered high centrality units predicted conceptual understanding as evaluated by human raters. Recall rate for high centrality units correlated strongly with conceptual understanding scores ($r(52) = .73$, $p < .001$, 95% CI [.57, .83]; (Figure 4A)) indicating that participants who retained more semantically central items demonstrated better conceptual understanding. By contrast, recall of low centrality units showed no reliable relationship with conceptual understanding scores ($r(52) = .24$, $p = .083$, 95% CI [-.03, .48]; (Figure 4B). The confidence interval crossing zero confirms this null result. This dissociation suggests that indiscriminate memory does not drive conceptual learning — it is specifically memory for semantically central, highly interconnected information that reflects and supports the construction of abstract knowledge.

Neural Score

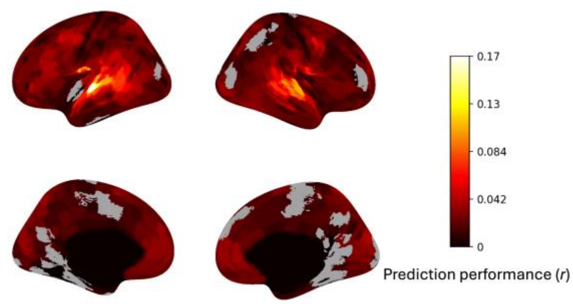

Figure 5. Average Prediction Performance across 26 subjects ($p < 0.001$, uncorrected). For each subject, voxel-wise model performance was calculated by correlating the predicted fMRI response and the actual fMRI response.

Using the subset of participants (N = 26) in behavioral analysis, we modeled fMRI responses while they were watching the lecture video using voxel-wise modeling. For each subject, we fitted a voxel-wise encoding model to estimate their BOLD response and tested the model performance on the held-out data which was not used for training. The model performance at each voxel was defined as "neural score" to represent the degree of how well the voxel response is predicted by the model.

At the group level, the results showed that the actual fMRI responses are significantly correlated with predicted fMRI responses at 444 out of 500 parcels (Schaefer et al., 2018) (Figure 5). Consistent with prior work (Caucheteux et al., 2022; Huth et al., 2016), significant neural scores were distributed across cortical networks, peaking in Heschl's gyrus, middle temporal gyrus (MTG), and superior temporal gyrus (STG). The strong alignment in these auditory and language regions likely reflects tracking of lexical and short-range semantic information — a process common to all participants hearing the lecture and therefore does not co-vary with learning outcomes.

Predicting Behavioral Outcomes From Neural Score

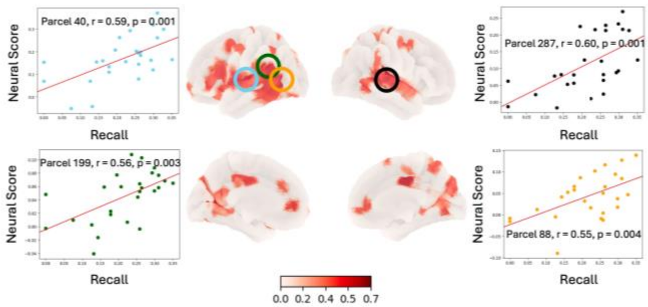

Figure 6. Correlation between neural score and recall (*r*), $p < 0.05$, uncorrected

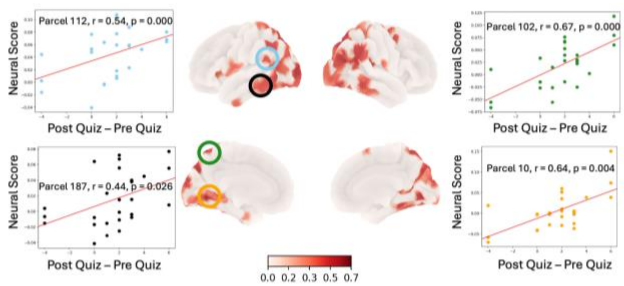

Figure 7. Correlation between neural score and [Post-Scan Quiz Score] - [Pre-Scan Quiz Score] (*r*), $p < 0.05$, uncorrected

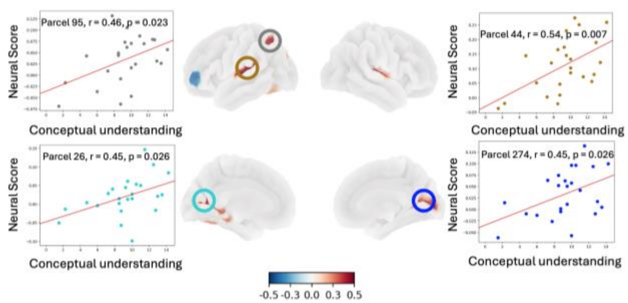

Figure 8. Correlation between neural score and conceptual understanding (*r*), $p < 0.05$, uncorrected

We next asked whether individual variation in brain-model alignment predicts learning outcomes. At each parcel, we correlated individual neural scores with three behavioral measures: number of recalled lecture units (Figure 6), pre-to-post lecture quiz score gain (Figure 7), and conceptual understanding scores (Figure 8).

Notably, the regions where SBERT best predicted fMRI responses at the group level did not consistently overlap with regions where neural scores predicted behavioral outcomes. Neural scores predictive of recall spanned auditory-lexical regions where group-level alignment peaked, but also extended to left inferior frontal gyrus (IFG), left posterior cingulate, and bilateral precuneus. Regions predictive of quiz score gain shifted toward parietal and occipital cortex, including bilateral inferior parietal and supramarginal gyri and lateral and medial occipital cortex. Regions predictive of conceptual understanding were more limited, including left intraparietal sulcus (IPS), Heschl's gyrus, and bilateral lingual gyrus. Scatter plots for a subset of regions are provided in Figures 6-8 to illustrate these relationships. Across all behavioral measures, neural scores positively predicted performance, with one exception: a single parcel in Figure 8 showed a negative relationship with conceptual understanding.

Discussion

In this study, we investigated how newly learned knowledge emerges in the mind during educational video viewing. We found that semantic features derived from a large language model capture the internal structure of abstract ideas — specifically, how interconnected they are within the broader conceptual framework — and that this structure predicts both what learners remember and how their brains respond during learning. These findings suggest that the embedding space of large language models captures meaningful mechanisms underlying how new knowledge is acquired and organized in the learner's mind.

Our behavioral results show that semantic centrality strongly predicts how likely the information is remembered by participants, and recall of semantically central content strongly predicted conceptual understanding, while recall of peripheral content did not. This pattern argues against a simple quantity-of-recall account of learning — encoding the most interconnected information matters more than indiscriminately accumulating large amounts of content. Notably, this selectivity was consistent across individuals regardless of overall comprehension level: even lower-performing participants concentrated their recall on high-centrality units, suggesting that semantic structure strongly shapes what gets encoded during lecture viewing.

Why does semantic centrality so powerfully predict memory in educational content? We suggest that educational videos represent a structurally distinct case where semantic organization becomes the primary available scaffold for encoding. Narrative memory is supported by multiple organizing frameworks — temporal sequencing, chronological ordering, and causal chaining — that learners can exploit during encoding and retrieval (Lang, 2009; Trabasso & Van Den Broek, 1985; Zwaan, 1996). In our physics lecture, multiple forces act simultaneously rather than sequentially, and demonstrations are juxtaposed without strong causal links between them. Absent these scaffolds, learners may default to semantic relatedness as the dominant organizing principle. The strong correlation between recall of central items and conceptual understanding further aligns with prior work showing that semantic centrality captures overarching themes analogous to a document summary (Erkan & Radev, 2004). Participants who successfully encoded semantically central content may therefore have extracted the lecture's core arguments, enabling them to construct more coherent conceptual frameworks.

Our fMRI findings complement these behavioral results: SBERT features explained neural activity to varying degrees across the cortex, with the strongest effects in primary auditory cortex, middle temporal cortex, and superior temporal gyrus. The regions where SBERT features best predicted neural responses were not, in general, the regions where neural score tracked behavioral outcomes. We can interpret this dissociation in two complementary ways. First, SBERT features are decomposable — any given voxel may track multiple aspects of the model's representation (lexical, short-range contextual, sentence-level), but only a subset of these aspects supports the construction of causal knowledge of physics. Auditory and superior-temporal regions are likely to encode lexical and sentence-surface-level information that is engaged across all learners regardless of their level of understanding. Higher-order regions, where the alignment with the model is weaker but their neural scores track behavior, may engage schematic information (e.g. syntactic information across sentences and paragraphs) to integrate conceptualized information and inferential elaboration. Second, the three behavioral measures vary in how far they are abstracted from sentence-level information. Recall is measured as cosine similarity between recalled clauses and lecture clauses, which is computed in the same feature space in SBERT. However, quiz scores are less aligned to sentence level feature as they require transfer of sentence level knowledge to a new context or a new format. Conceptual understanding score is far abstracted from extraction of sentence level content, as it requires extraction of causal structure or schematic relations presented throughout the entire lecture. The narrowing regional findings across our three measures may therefore reflect the degree of feature-task representational alignment, which could vary depending on how much each behavioral measurement relies on the models' representational space.

In summary, we demonstrate that semantic network structure shapes how learners preferentially encode information from educational content and reconstruct their conceptual understanding. The centrality principle provides a domain-general framework for organizing recall and plays a critical role in conceptual learning by enabling learners to extract and integrate the most interconnected ideas. Our neural findings extend this framework by showing that the same semantic features that predict behavioral outcomes also explain neural responses during learning, particularly in temporal and auditory cortices associated with language comprehension. Notably, the brain regions where SBERT features best predicted neural activity did not always correspond to regions where neural scores predicted behavioral outcomes, suggesting that some aspects of semantic processing may support perceptual encoding without directly contributing to deeper conceptual understanding. These findings advance our understanding of how the mind constructs knowledge from experience, suggesting that the semantic geometry of large language models may capture domain-general principles of how information is selectively encoded, organized, and abstracted into conceptual understanding. At a practical level, identifying semantically central information in a lecture could inform educational design, allowing strategic emphasis on the content most likely to support both retention and conceptual comprehension during learning.

References

Abraham, A., Pedregosa, F., Eickenberg, M., Gervais, P., Mueller, A., Kossaifi, J., Gramfort, A., Thirion, B., & Varoquaux, G. (2014). Machine learning for neuroimaging with scikit-learn. *Frontiers in Neuroinformatics, 8,* 14. https://doi.org/10.3389/fninf.2014.00014

Caucheteux, C., Gramfort, A., & King, J. R. (2022). Deep language algorithms predict semantic comprehension from brain activity. *Scientific reports, 12*(1), 16327. https://doi.org/10.1038/s41598-022-20460-9

Cetron, J. S., Connolly, A. C., Diamond, S. G., May, V. V., Haxby, J. V., & Kraemer, D. J. (2020). Using the force: STEM knowledge and experience construct shared neural representations of engineering concepts. *npj Science of Learning, 5*(1), 6. https://doi.org/10.1038/s41539-020-0065-x

Chang, L. J., Jolly, E., Cheong, J. H., Rapuano, K. M., Greber, N., Reddan, M. C., & Richardson, H. (2018). *cosanlab/nltools: 0.3.11.* Zenodo. https://doi.org/10.5281/zenodo.2229813

Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 4171–4186). Association for Computational Linguistics. https://doi.org/10.18653/v1/N19-1423

Erkan, G., & Radev, D. R. (2004). LexRank: Graph-based lexical centrality as salience in text summarization. *Journal of Artificial Intelligence Research, 22*, 457–479. https://doi.org/10.1613/jair.1523

ESFTV. (2010, August 11). *Physics of Life—Falling Bodies* [Video]. YouTube. https://www.youtube.com/watch?v=MfMs3odd9WE

Esteban, O., Markiewicz, C. J., Blair, R. W., Moodie, C. A., Isik, A. I., Erramuzpe, A., Kent, J. D., Goncalves, M., DuPre, E., Snyder, M., Oya, H., Ghosh, S. S., Wright, J., Durnez, J., Poldrack, R. A., & Gorgolewski, K. J. (2019). fMRIPrep: A robust preprocessing pipeline for functional MRI. *Nature Methods, 16*(1), 111–116. https://doi.org/10.1038/s41592-018-0235-4

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. https://doi.org/10.1038/nature17637

Lang, A. (2009). Effects of chronological presentation of information on processing and memory for broadcast news.

Journal of Broadcasting & Electronic Media, 33(4), 441–452. https://doi.org/10.1080/08838158909364093

Lee, H., & Chen, J. (2022). Predicting memory from the network structure of naturalistic events. *Nature Communications, 13*(1), 4235. https://doi.org/10.1038/s41467-022-31965-2

McClelland, J. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review, 102*(3), 419–457. https://doi.org/10.1037/0033-295X.102.3.419

Meshulam, M., Hasenfratz, L., Hillman, H., Liu, Y. F., Nguyen, M., Norman, K. A., & Hasson, U. (2021). Neural alignment predicts learning outcomes in students taking an introduction to computer science course. *Nature Communications, 12*(1), 1922. https://doi.org/10.1038/s41467-021-22202-3

Moore, I. L., & Long, N. M. (2024). Semantic associations restore neural encoding mechanisms. *Learning & Memory, 31*(3), a053996. https://doi.org/10.1101/lm.053996.124

Nguyen, M., Chang, A., Micciche, E., Meshulam, M., Nastase, S. A., & Hasson, U. (2022). Teacher–student neural coupling during teaching and learning. *Social Cognitive and Affective Neuroscience, 17*(4), 367–376. https://doi.org/10.1093/scan/nsab103

O'Donnell, R. E., Clement, A., & Brockmole, J. R. (2018). Semantic and functional relationships among objects increase the capacity of visual working memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 44*(7), 1151–1158. https://doi.org/10.1037/xlm0000508

Reimers, N., & Gurevych, I. (2019). Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)* (pp. 3982–3992). Association for Computational Linguistics. https://doi.org/10.18653/v1/D19-1410

Rogers, T. T., & McClelland, J. L. (2004). *Semantic cognition: A parallel distributed processing approach.* MIT Press. https://doi.org/10.7551/mitpress/6161.001.0001

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*(9), 3095–3114. https://doi.org/10.1093/cercor/bhx179

Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science, 331*(6022), 1279–1285. https://doi.org/10.1126/science.1192788

Trabasso, T., & Van Den Broek, P. (1985). Causal thinking and the representation of narrative events. *Journal of Memory and Language, 24*(5), 612–630. https://doi.org/10.1016/0749-596X(85)90049-X

Wiese, H., & Schweinberger, S. R. (2015). Getting connected: Both associative and semantic links structure semantic memory for newly learned persons. *Quarterly Journal of Experimental Psychology, 68*(11), 2131–2148. https://doi.org/10.1080/17470218.2015.1008526

Zwaan, R. A. (1996). Processing narrative time shifts. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22*(5), 1196–1207. https://doi.org/10.1037/0278-7393.22.5.1196