Title

Mind Captioning: Evolving descriptive text of mental content from human brain activity

Short title

Generating text of mental content from the brain

Author

Tomoyasu Horikawa[1]*

Affiliation

[1] NTT Communication Science Laboratories, Kanagawa, Japan.

* Correspondence: horikawa.t@gmail.com

Abstract

A central challenge in neuroscience is decoding brain activity to uncover mental content comprising multiple components and their interactions. Despite progress in decoding language-related information from human brain activity, generating comprehensive descriptions of complex mental content associated with structured visual semantics remains challenging. We present a method that generates descriptive text mirroring brain representations via semantic features computed by a deep language model. Constructing linear decoding models to translate brain activity induced by videos into semantic features of corresponding captions, we optimized candidate descriptions by aligning their features with brain-decoded features through word replacement and interpolation. This process yielded well-structured descriptions that accurately capture viewed content, even without relying on the canonical language network. The method also generalized to verbalize recalled content, functioning as an interpretive interface between mental representations and text, and simultaneously demonstrating the potential for non-verbal thought-based brain-to-text communication, which could provide an alternative communication pathway for individuals with language expression difficulties, such as aphasia.

Teaser

Non-verbal thoughts can be translated into verbal descriptions by aligning semantic representations between text and the brain.

Introduction

Humans can recognize and recall intricate visual content comprising multiple semantic components, including objects, places, actions, and events, along with their interactions and relationships. These elaborate and structured mental representations form the foundation for translating thoughts into language and communicating experiences with others. Recently, substantial progress has been made in brain decoding of language-related information, enabling the direct production of linguistic outputs, such as text, from the human brain (*1–4*). However, decoding the perceptual—and not to mention mental—content associated with visual semantics to generate comprehensive descriptions of subjective experiences remains challenging. Translating brain activity linked to non-linguistic semantic information, or thoughts, into verbal descriptions could greatly enhance our ability to interpret diverse mental states, opening up numerous possibilities for applications, particularly with text-prompt-based systems [e.g., ChatGPT (*5*) and Gemini (*6*)], as well as for scientific research.

Prior research on decoding visual semantics using human functional magnetic resonance imaging (fMRI) has focused on individual components or static images. This narrow focus has hindered the decoding of complex content involving interactions between multiple elements, thus obscuring our understanding of how the brain represents rich and structured visual semantics. While studies have successfully decoded individual components in viewed (*7–9*), imagined (*7*), and dreamed (*10*) content using object- or word-level features, they have fallen short of capturing neural representations of interactions and relationships that are not predicted by their individual components (*11*) and are crucial for recognizing actions and social interactions (*12–16*).

Some researchers have incorporated caption databases (*17*) or deep neural network (DNN)-based modules, such as non-linear image captioning models (*18–20*), to produce sentence-level decoding predictions that appear to have linguistic structure. However, predictions based on database search methods are limited to existing, often deliberately structured, descriptions that may not capture the full complexity of diverse visual content. In addition, non-linear methods can introduce spurious information not "explicitly represented" in the brain (*21–24*). Specifically, non-linear captioning models can construct sentence-like structured outputs even from object-level visual features (*25*), which inherently lack relational information. Thus, successful decoding via such models might reflect the complexities of the model architecture, rather than the properties of the underlying brain representations. Therefore, these approaches, designed to produce linguistically structured outputs, are not ideal for examining whether structured visual semantics, essential for representing relational information, are genuinely encoded in the brain through decoder outputs.

To overcome these limitations, we introduce a new generative decoding method called *mind captioning*, which generates descriptive text mirroring semantic information represented in the brain (Fig. 1). Our method combines linear feature decoding analysis (*7, 9, 10*), using semantic features computed by a deep language model (LM), with a novel optimization method that generates text based on these features. Semantic features serve as intermediate representations for decoding (translating) semantic information from the brain into text. They can act as a bridge for decoding both perceptual and mental content, as shared representations exist between visual perception and mental imagery, particularly for high-level information (*7, 10, 26–28*). Additionally, deep LMs offer the advantage of effectively capturing contextual meanings, which are crucial for delineating intricate interrelationships (*29–34*).

The challenge lies in linguistically interpreting the information in semantic features decoded from brain activity. Although the ideal approach would involve examining all possible word sequences to identify the description whose semantic features best match

the decoded features, this is not feasible because the number of candidates is infinite. We thus developed an iterative optimization method that generates descriptive text from scratch by progressively aligning the semantic features of candidate descriptions with target brain-decoded features through word replacement and interpolation in a search for the best description (Fig. 1B). Crucially, we leveraged an LM pre-trained for masked language modeling (MLM) (*35*) to constrain the search space during optimization. By directly optimizing word sequences to match brain-decoded features, our method minimizes dependence on external resources such as caption databases or non-linear captioning models, thereby ensuring the generation of descriptions more closely aligned with brain representations while maintaining the interpretability of structured visual semantics in the brain.

To demonstrate the effectiveness of our method, we first validated it for perceptual content by constructing decoding models (decoders) from stimulus-induced brain activity and then tested their generalizability to activity during recall-based mental imagery. Specifically, we measured brain activity in six subjects—all Japanese individuals who were non-native English speakers—using fMRI while they viewed or recalled video clips (*36*) (fig. S1) and created a data sample by averaging fMRI volumes measured during viewing and recalling each video. To enhance the quality of the fMRI data, we averaged data samples over five repetitions for each stimulus or imagery item in the test phase. Using the samples from stimulus-induced brain activity, we trained decoders to predict semantic features, which were computed from corresponding video captions using an LM [DeBERTa-large (*37*)]. We then used these decoders to translate brain activity associated with viewed and recalled content into semantic features for novel test videos not used during training. Finally, we used the decoded features to optimize the text by employing an MLM model [RoBERTa-large (*38*)]. Through these analyses, we aimed to validate our method's capability to generate comprehensive descriptions of both viewed and recalled content from brain activity, establishing a framework for decoding non-verbal mental content and exploring the neural basis of structured visual semantics.

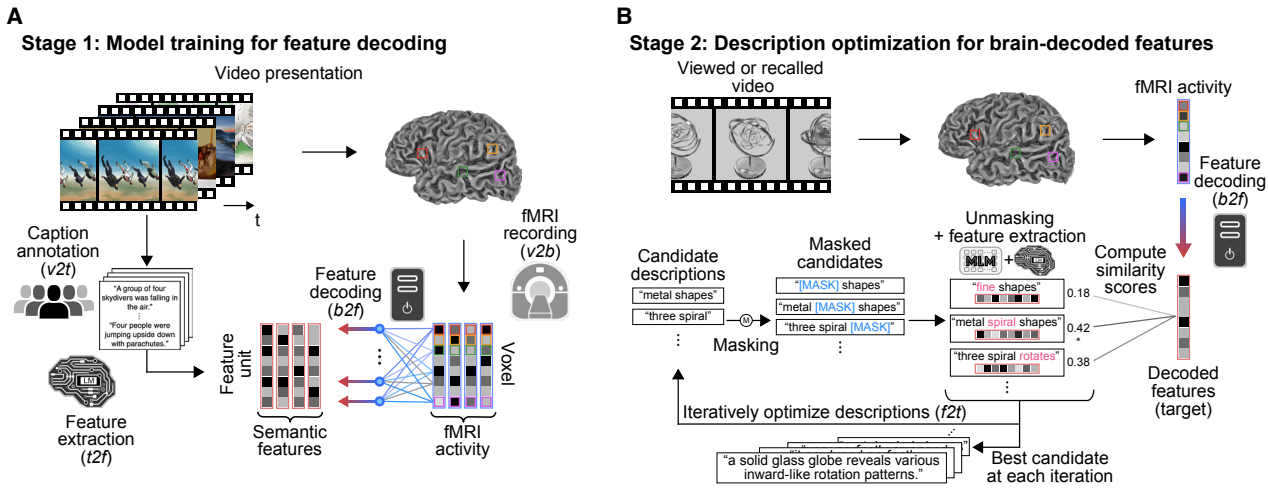

Fig. 1. Mind captioning. Our method consisted of two stages. (**A**) We first trained linear decoding models to decode whole-brain fMRI activity, measured while each subject viewed videos, into semantic features from the captions of the videos using an LM (frozen). (**B**) We then used those models to decode brain activity induced by novel video stimuli or by recall-based mental imagery of those videos, and optimized candidate descriptions iteratively by aligning their features with brain-decoded features through word replacement and interpolation, leveraging another LM pre-trained for MLM (frozen). The optimization consisted of three stages: masking, unmasking, and candidate selection. During masking, we randomly applied masks by replacing words with a mask or interpolating masks into the candidate word sequences. During unmasking, the MLM model created new candidates by filling in the masks in the masked candidates based on the context of surrounding words. During candidate selection, we computed semantic features of all new and original candidates using the LM for feature extraction. We then evaluated the similarity between those candidate features and the target brain-decoded features to select the top candidates for further optimization. The optimization process was initiated from a non-informative word ("<unk>") to avoid incorporating any prior assumptions for description generation and was repeated 100 times. See fig. S2 for details on the model and parameter validations. Each transformation process (e.g., v2t, b2f) corresponds to a Transformation ID summarized in Table 1. Because of copyright restrictions, the actual video frames used in the experiments have been replaced with schematic illustrations throughout this paper. see Materials and Methods for details of how these illustrations were generated.

Table 1. Summary of information modalities and their transformations

Transformation ID	*Input modality*	*Output modality*	*Transformation method*
v2t	Videos	Text captions	Crowdsourced manual annotation of text captions
t2f	Text captions	Semantic features	Feature extraction using a pre-trained LM (e.g., DeBERTa-large)
v2b	Videos	Brain activity	fMRI recording
b2f	Brain activity	Semantic features	L2-regularized linear regression
f2t	Semantic features	Text captions	Iterative text generation using an MLM (e.g., RoBERTa-large)

Note: Transformation IDs (e.g., v2t: video to text, b2f: brain to feature, etc.) are used consistently in Fig. 1 to denote each transformation process.

Results

Generating viewed content descriptions

The optimization of text, based on decoded features from stimulus-induced brain activity, resulted in a progressive evolution of descriptive texts (Fig. 2A). Initially, the descriptions were fragmented and lacked clear meaning. However, through iterative optimization, these descriptions naturally evolved to have a coherent structure and effectively capture the key aspects of the viewed videos. Notably, the resultant descriptions accurately reflected the content, including the dynamic changes in the viewed events (Fig. 2B). Furthermore, even when specific objects were not correctly identified, the descriptions still successfully conveyed the presence of interactions among multiple objects (e.g., Fig. 2B bottom left).

Throughout the optimization process, the features of the evolved descriptions exhibited increasingly stronger correlations with the target features and, consequently, with the features of the reference captions annotated on the viewed videos (Fig. 2C). In contrast, the same optimization performed with an untrained MLM model, which randomly suggests candidate words during unmasking—a process analogous to a genetic algorithm—did not yield comparable results (fig. S3B). These results emphasize the significance of contextual information from a pre-trained MLM model to efficiently explore descriptions closely aligned with brain representations and to enhance description quality.

To assess decoding performance, we used multiple similarity metrics to evaluate the similarity between the generated descriptions and the reference captions. We calculated similarity scores for both the captions of the viewed (correct) and irrelevant (incorrect) videos and defined discriminability as the difference between them. The generated descriptions exhibited significantly high discriminability across all metrics and subjects (Fig. 2D; Wilcoxon signed-rank test, one-tailed, $P < 0.01$, FDR corrected across metrics and subjects), indicating that these descriptions were accurate enough to differentiate video content.

To gain an intuitive understanding of the performance, we evaluated the accuracy of video identification by comparing the generated descriptions to both correct and incorrect captions across various numbers of candidates. Performance consistently exceeded chance levels for all set sizes, with approximately 50% accuracy among 100 distinct possibilities for all subjects when using feature correlation (chance = 1%; Fig. 2E), demonstrating the effectiveness of our method in translating detailed information from the brain into text through semantic features.

To further highlight the effectiveness of our method, it is worth noting that it surpasses existing approaches and holds promise for continued improvement. Specifically, our method captured subjective experiences more accurately and flexibly (fig. S5), as its generated descriptions aligned more closely with captions that individual subjects rated as highly consistent with their perceptions (figs. S1A and S5E). Additionally, our method demonstrated superior discriminability compared to approaches relying on caption databases (*17*) or non-linear captioning models (*18*) across both video-induced and image-induced fMRI data (*39*) (figs. S6 and S7). Furthermore, our method robustly generated descriptions that accurately reflected the viewed content regardless of the LMs employed, and brain encoding performance—a metric assessing alignment between the brain and models (*31, 40*)—correlated with the text generation performance of the LMs (fig. S8). These results suggest that employing LMs more closely aligned with the brain [e.g., GPT3 (*5*), OPT (*41*), or LLaMA (*42*)] may further improve the effectiveness of our method.

Assessing structured semantic information in generated descriptions

A key advantage of structured descriptions over simple word lists is their ability to organize words to convey contextual meaning that goes beyond a mere list of individual entities, capturing their interrelations. To capture analogous structures in visual experiences, we define structured visual semantics as a range of visually grounded relational structures—including semantic roles, attribute–object relations, and spatial configurations—illustrated by contrasts such as "a bird eats a snake" versus "a snake eats a bird" or "some grass in a mug" versus "a mug in some grass," which are linguistically describable and sensitive to word order (*43*). These distinctions are essential for building rich representations of visual structure—such as action direction and visually implied interrelations—that form the foundation for perceiving social and human–object interactions (*11–16*).

To assess whether our generated descriptions captured the relational structure of visual scenes through proper word arrangement, we tested the effect of shuffling their word order—either for all words or nouns only (up to 1,000 shuffled variants; see Materials and Methods)—under the hypothesis that if the original word order accurately conveyed visual relationships, shuffling would reduce discriminability.

While the shuffled descriptions retained reasonably high accuracy in identifying videos—indicating that word lists alone provide informative cues—shuffling all words, or even just the nouns, significantly impaired discriminability (Fig. 2F; Wilcoxon signed-rank test, one-tailed, $P < 0.01$, FDR corrected across subjects). This reduction in discriminability remained robust even when using the minimally disrupted shuffled descriptions, which had the highest fluency (or linguistic acceptability) as assessed by the pseudo-log-likelihood score from MLM scoring (*44, 45*) (fig. S4F). These results demonstrate that our method generates descriptions that capture more detailed information than simple word lists.

Interestingly, the impact of shuffling was more pronounced when using features from deeper LM layers to generate descriptions (fig. S9), underscoring the importance of the deep structure of the LMs in constructing contextual semantic representations and accurately decoding structured information.

Disentangling the origins of descriptive coherence

A crucial step in understanding the representational basis of our method is to determine whether the coherence of generated descriptions—essential for depicting visual relations—reflects structured semantic information in the target features or is instead artificially imposed by the priors from the MLM model used in the text optimization process (Fig. 1B). We reasoned that if brain-decoded features genuinely represent specific structured semantic information—uniquely conveyed by the particular word order of the generated description—then these features should exhibit greater similarity to the features of the original description than to those of shuffled variants that preserve the same words but disrupt their original order. We tested this by comparing feature correlations between the brain-decoded features and both the original and shuffled descriptions.

Supporting our reasoning, shuffling the word order significantly reduced feature correlations, especially when it substantially altered the original meanings (Fig. 2G). The original descriptions scored highest among all variants generated through all-word shuffling and ranked in the top 0.001% for noun-only shuffled variants (six subjects pooled). These results suggest that the specific word order of the generated descriptions is meaningfully guided by the semantic information encoded in the feature representations.

To further explore the possibility that the MLM model imposed coherent structure on the generated descriptions, we examined the outputs of decoders trained with semantic features from word-order-shuffled captions—a method intended to prevent the decoders from learning structured semantics. As a result, we confirmed that incoherent word sequences were generated, though they still contained words that semantically matched individual components of the viewed videos (fig. S3C).

We also confirmed that our text optimization method accurately reconstructs original descriptions from model-derived (i.e., not brain-decoded) features after sufficient iterations (fig. S10). Importantly, reconstructions that matched the target captions yielded feature correlations of 1.0, whereas even minor deviations—such as word omissions or reordering—led to notable reductions in correlation, highlighting the fidelity of the semantic feature space. Although the method struggled with long or shuffled captions—likely due to the incremental, token-based optimization strategy and residual priors in the MLM model—these priors alone were insufficient to enforce coherent word sequences. These results indicate that the semantic feature space preserves fine-grained information sufficient for accurate text reconstruction, with the model priors playing a supportive—but not determinative—role in shaping the linguistic form of the output.

Together, these results reinforce our finding that the text optimization process does not impose restrictive constraints on forming linguistically structured outputs when generating coherent and structured descriptions from brain activity. Instead, the observed coherence is likely guided by the structured information present in the brain-decoded features, indicating that these features—not the optimization algorithm—are the primary source of the generated structure.

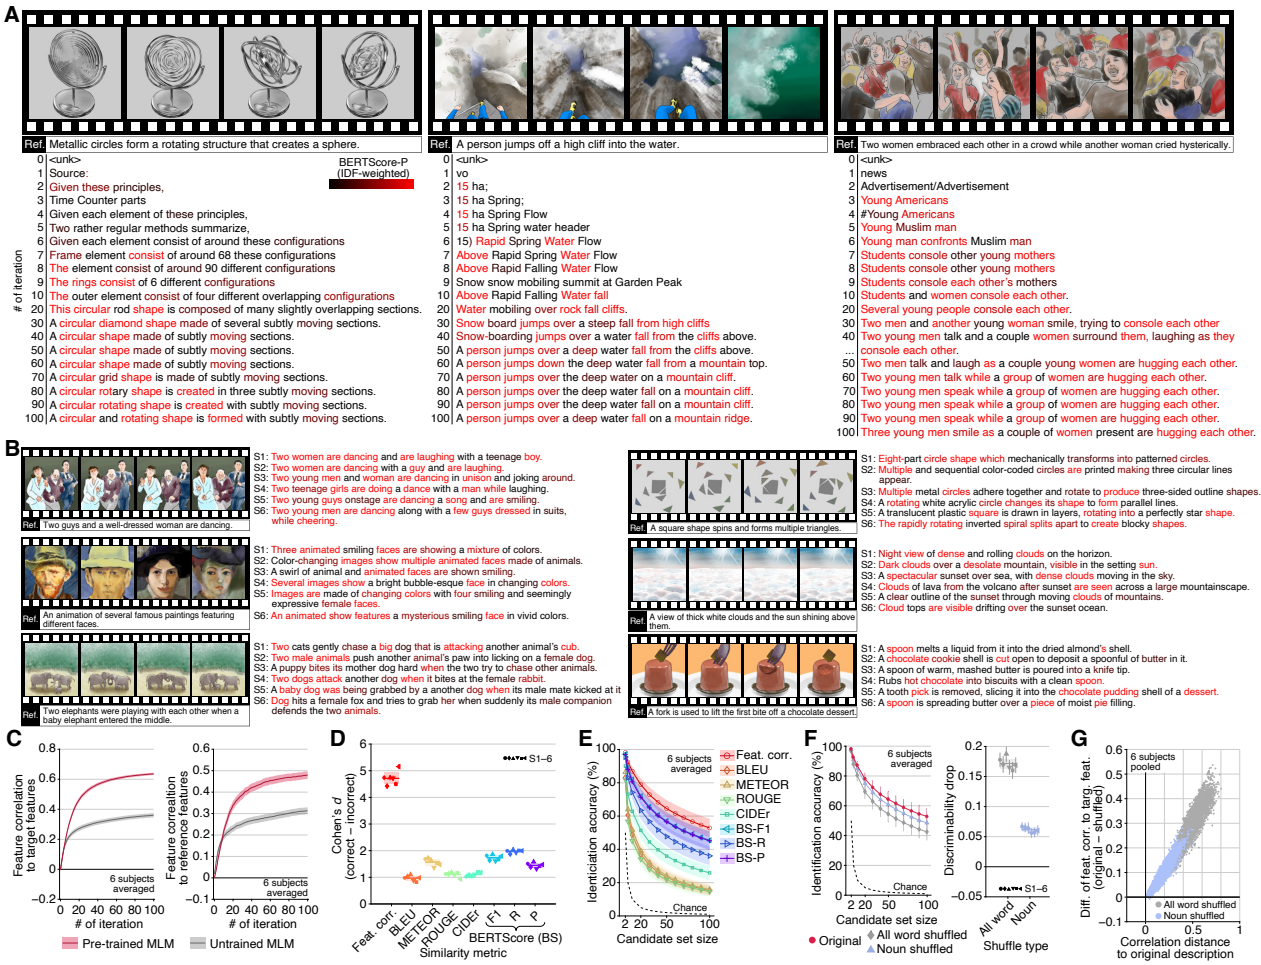

Fig. 2. Generating viewed content descriptions. Descriptions were generated using features from all LM layers decoded from whole-brain activity. (**A**) Evolved descriptions during the optimization (see https://horikawa-t.github.io/MindCaptioningProject/ for more results with the original videos). (**B**) Descriptions after 100 iterations for all subjects (see fig. S3A for more example). In (A), (B), the color indicates accuracy [inverse document frequency (IDF)-weighted BERTScore-P]. A reference caption of the video is shown below frames. (**C**) Feature correlations between features of generated descriptions and those decoded from the brain, as well as those computed from correct references. (**D**) Cohen's *d* of discriminability (see fig. S4B for raw scores). (**E**) Video identification accuracy with varying numbers of candidates. (**F**) Effects of word-order shuffling on video identification accuracy and discriminability. (**G**) Scatterplot of the correlation distances (one minus feature correlation) between the original and shuffled descriptions against the difference in feature correlations to target features between original and shuffled descriptions. Each dot indicates a shuffled description. Shades in (C), (E), and error bars in (D), (F) indicate 95% confidence intervals (C.I.) across samples ($N = 72$). Shades in (D), (F) indicate 95% C.I. across subjects ($N = 6$). See fig. S4 for individual results.

Contributions from different brain areas

Having confirmed the ability to generate accurate and well-structured descriptions from whole-brain activity, we next examined the contributions of specific brain regions to this decoding, focusing on whether complex structured semantics can be derived independently of the language network. Although numerous studies have shown that a broadly distributed semantic network encodes the meanings of both visual and linguistic information (*46–49*), most have focused on category- or word-level representations. Research on structured semantics has primarily concentrated on language processing, often associated with the fronto-temporal language network (*50–52*), whereas the neural substrates for structured visual semantics remain comparatively underexplored. Emerging evidence suggests that, while the language network is also recruited in processing meaningful visual scenes (*53, 54*), other regions—particularly the lateral occipital temporal cortex—encode certain relational aspects of visual information, including interactions and the directedness of actions among persons and objects (*11–16*). These findings suggest the intriguing potential to enable the direct decoding and communication of rich structured semantic information from brain activity while bypassing the linguistic processing typically required to translate thoughts into words. Pursuing this possibility could broaden the scope of brain–machine interfaces (BMIs) for converting non-verbal thoughts into text, opening new avenues for semantic decoding that do not rely on language.

To establish a foundation for evaluating decoding performance, we began by analyzing how the brain encodes structured visual semantic information in videos. We constructed two encoding models: one based on the semantic features used in our decoding analysis and the other on visual features from a video recognition DNN pre-trained to classify object and action categories [TimeSformer (*55*)]. The semantic encoding model effectively predicted brain activity in the language network and in regions involved in recognizing objects, actions, and interactions (*11–16*), spanning the higher visual cortex (HVC) and extended areas of the parietal and frontal cortices (Fig. 3A). While the visual model performed better in the lower visual cortex (LVC, V1–V3), the semantic model progressively outperformed it in the HVC and language network (Fig. 3B–E). The shift in relative superiority between the visual and semantic models occurs at the midpoint of the category-selective regions, between their posterior and anterior halves, suggesting a functional boundary (Fig. 3B). Additionally, voxels in the HVC and language network were better predicted by features from deeper layers of the visual DNN and the LM (Fig. 3F). Notably, the LM layers yielding the highest performance for language network voxels were much deeper than those for the HVC (Wilcoxon rank-sum test, one-tailed, $P < 0.01$, FDR corrected across subjects), highlighting an indicative link between the language network and contextualized semantic information. These results demonstrate that the language network, along with other regions, is involved in encoding contextual semantics, consistent with previous studies on its activation by non-verbal visual semantics (*53, 54*).

We then assessed decoding performance using video identification analysis based on descriptions generated from these brain regions with varying selectivity (Fig. 3G). Focusing on voxels better predicted by the semantic model yielded higher performance, approaching whole-brain activity results for some subjects (fig. S4H). In contrast, decoding from voxels better predicted by the visual model resulted in weaker performance, indicating limited contributions from these voxels even though they are widely distributed across the posterior side of visual category-selective areas. Using only the language network, despite its involvement in encoding contextual semantics, did not produce high performance, suggesting that its contribution may be more supportive than essential. Notably, ablating the language network did not profoundly

impact performance, achieving almost 50% accuracy from 100 possibilities (chance = 1%; six subjects averaged), showing decreased accuracy and discriminability by word order shuffling (Fig. 3H; Wilcoxon signed-rank test, one-tailed, $P < 0.01$, FDR corrected across subjects), and even generating intelligible descriptions (fig. S3D).

These results suggest that accurate descriptions capturing structured semantics can be generated without relying on the language network, indicating that structured visual semantic information is represented across regions extending from the anterior portions of the occipital visual cortex into parietal and frontal cortices outside the language network. Such representations may underlie the comprehension of complex visual events in individuals with global aphasia (*53*). These findings also provide further support for the distinction between language and non-verbal thought (*56*).

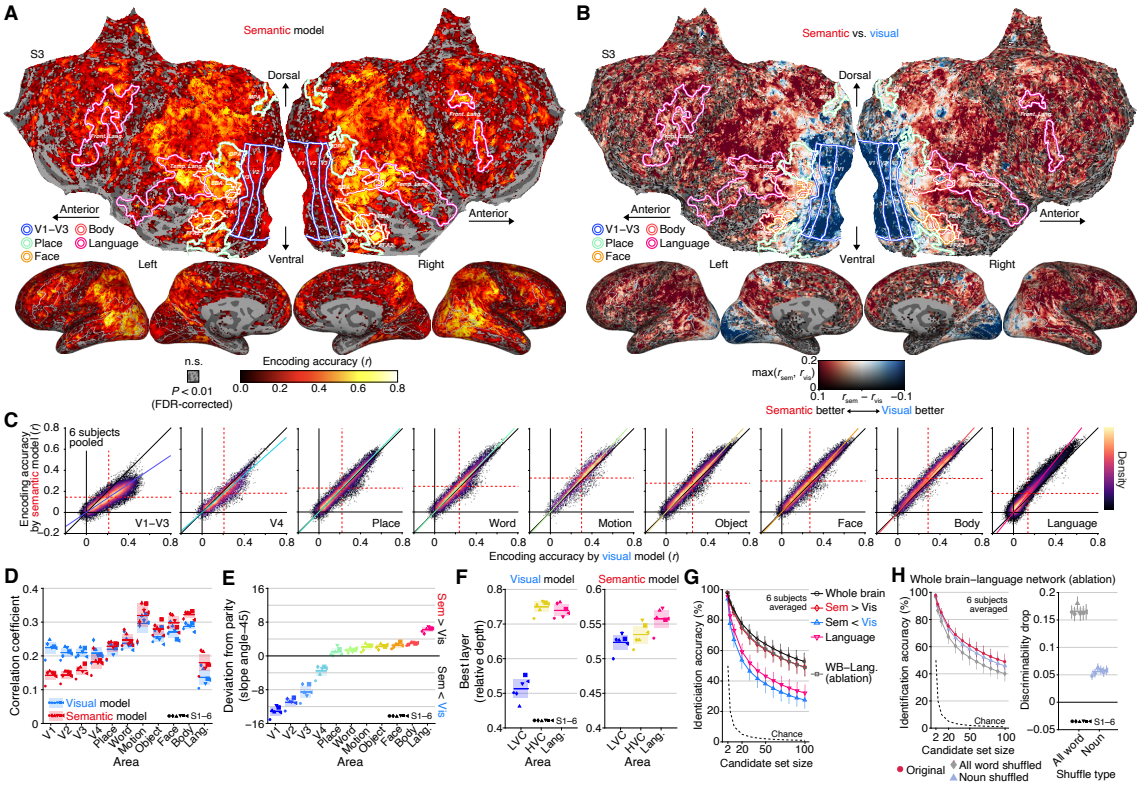

Fig. 3. Contributions from different brain areas. A cross-validation analysis was performed within training perception data. Encoding models were trained using features from each layer to generate predictions from multiple layers. The final predictions were constructed based on the best layer per voxel determined using nested cross-validation. (**A**) Encoding performance of the semantic model. **(B)** Performance difference between the semantic and visual models. (**C**) Density heatmap of accuracy (solid lines, best linear fit). (**D**) Mean encoding accuracy within each area. (**E**) Deviation from parity based on the slope angles of the best linear fit. (**F**) Mean of the best layers. The indices of layers with the highest performance were averaged across voxels in each area and then converted into relative depth. (**G**) Video identification accuracy from different brain areas. Decoders were trained using voxels with higher encoding accuracy according to the semantic or visual models, voxels within the language network, or voxels from the entire brain except for the language network. (**H**) Effects of word-order shuffling on video identification accuracy and discriminability without using the language network. In **(D–H)**, error bars and shades indicate 95% C.I. across samples and subjects, respectively. See figs. S4 and S11 for individual results.

Generating recalled content descriptions

Finally, we investigated whether the decoders trained on brain activity induced by visual stimuli—hereafter referred to as *perception-trained decoders*—could be used to generate descriptions of mental content based on brain activity induced by mental imagery of recalled videos, applying the same evaluation procedures as in the perception data analysis. The analysis successfully generated descriptions that accurately reflected the content of the recalled videos, although accuracy varied among individuals (Fig. 4A). These descriptions were more similar to the captions of the recalled videos than to irrelevant ones (fig. S12A and B), with proficient subjects achieving nearly 40% accuracy in identifying recalled videos from 100 candidates (Fig. 4B; chance = 1%). Shuffling word order resulted in a notable reduction in video identification accuracy and discriminability (Fig. 4C; Wilcoxon signed-rank test, one-tailed, $P < 0.01$, FDR corrected across subjects). Excluding the language network from the analysis slightly, but not substantially, reduced video identification accuracy (Fig. 4D).

We further confirmed that semantic features serve as effective intermediate representations, allowing perception-trained decoders to generalize across cognitive states and decode mental content (Fig. 4E). Specifically, we applied perception-trained decoders to both stimulus- and imagery-induced brain activity, evaluating feature-based pairwise video identification performance using features from individual layers of three model types—visual (TimeSformer) (*55*), visuo-semantic (CLIP) (*57*), and semantic (DeBERTa-large) (*37*). When quantifying generalizability as the angle between the parity line and a line connecting chance level (50%) with the observed accuracies, generalizability increased with layer depth in visual features, was relatively high for visuo-semantic features, and was highest for semantic features across all layers. These results demonstrate that semantic features robustly bridge neural representations shared across perception and imagery, providing a reliable foundation for translating internal experiences into coherent text.

Notably, when the same text generation analysis was applied to brain activity during the preparation period (fig. S1B), all subjects showed lower accuracy than during the imagery period, with most subjects (except S2) performing at levels closer to chance (fig. S12D). This finding suggests that volitional mental imagery, rather than text reading of verbal cues, is essential for eliciting semantic neural representations, enabling the accurate generation of target content descriptions.

Collectively, these results confirm the generalizability of the perception-trained decoders to generate descriptions of recalled content, showcasing their capability to effectively verbalize mental representations. Additionally, we were able to obtain comprehensible descriptions of the recalled content with reasonable identification accuracy from single-trial fMRI activity (Fig. 4F and fig. S14), demonstrating the potential applicability of our method to experiences that are difficult to reproduce, such as dreams (*27*).

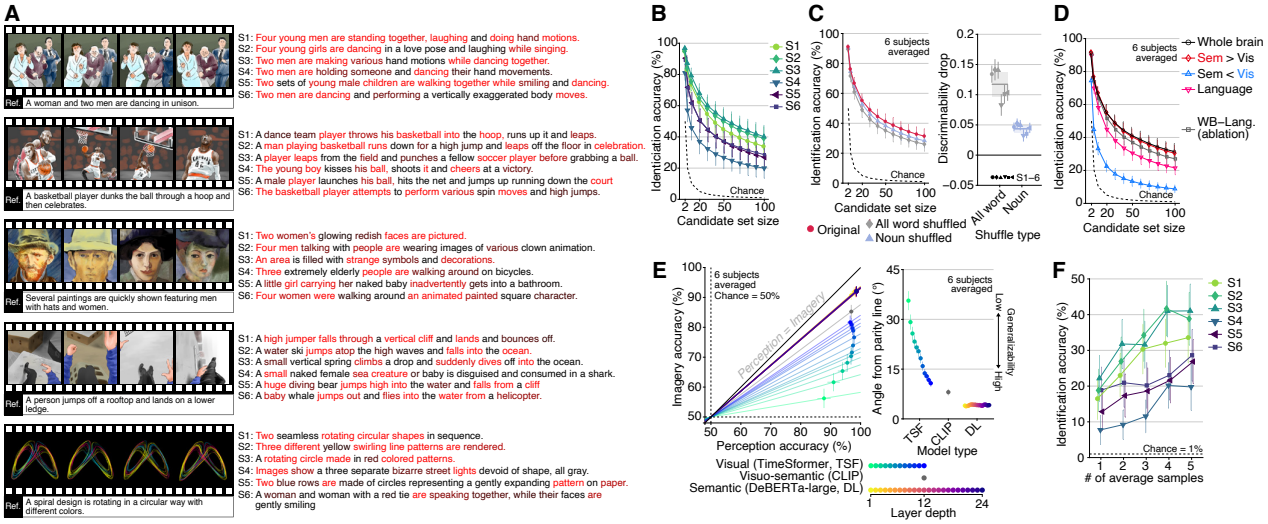

Fig. 4. Generating recalled content descriptions. We applied decoders trained on stimulus-induced brain activity (DeBERTa-large; all layers; decoded from whole-brain activity, unless otherwise stated) to analyze the brain activity of subjects engaged in recalling video scenes from memory, prompted by verbal descriptions (see fig. S1B). (**A**) Descriptions of recalled content generated after 100 iterations for all subjects (see fig. S3E for more examples). (**B**) Identification accuracy of recalled videos for individual subjects. (**C**) Effects of word-order shuffling on identification accuracy and discriminability of recalled videos. (**D**) Identification accuracy of recalled videos from different brain areas. (**E**) Comparison of feature-based pairwise video identification accuracy between perception and imagery. Generalizability was quantified as the angle between the parity line (the diagonal line indicating equal accuracy for perception and imagery) and a line connecting the chance level (50%) with the observed accuracy for each model and layer. Error bars (left and right panels) indicate 95% C.I. across samples and jackknife-estimated standard errors, respectively. (**F**) Identification accuracy of recalled videos with a varying number of averaged samples. For (A), (C), and (D), conventions are the same as for Figs. 2B, F, and 3G. See figs. S6B, D, and S12–S14 for additional results.

Discussion

We have successfully generated descriptive text representing visual content experienced during perception and mental imagery by aligning semantic features of text with those linearly decoded from human brain activity. Our success is attributed to two key factors: the advancement of deep LMs that provide contextualized semantic representations similar to those in the human brain (*29–34*) and our innovative text optimization method utilizing the MLM model for word-level optimization while efficiently constraining the search space. Together, these factors facilitate the direct translation of brain representations into text, resulting in optimally aligned descriptions of visual semantic information decoded from the brain. These descriptions were well structured, accurately capturing individual components and their interrelations without using the language network, thus highlighting the existence of fine-grained structured semantic information outside this network. Our method enables the intelligible interpretation of internal thoughts, demonstrating the feasibility of non-verbal, thought-based brain-to-text communication.

Our methodological configuration was designed to generate accurate and comprehensive descriptions of mental content from brain activity while facilitating interpretation of the neural basis of structured visual semantics. Rather than mapping brain activity directly to text (*19, 20*)—a strategy requiring large-scale brain data and introducing architectural complexity that complicates attribution—we adopted a two-stage approach: first decoding brain activity into a feature space, then generating text from these features, thereby enabling localization of potential information loss. We used stimulus-induced activity to ensure stable decoder training and employed linear mappings to probe explicit brain representations (*21–24*), prioritizing interpretability and mitigating overfitting—particularly crucial for generalizing to imagery. We aligned brain activity with semantic features to leverage high-level representations shared between perception and imagery (*7, 10, 26–28*), enhancing cross-state generalizability. Indeed, semantic features demonstrated superior generalizability when decoding imagery using perception-trained decoders, outperforming visual and visuo-semantic counterparts—commonly used in captioning models (*25, 58*) and their decoding applications (*18*) (Fig. 4E). We used iterative text optimization to generate descriptions, which accurately reconstructed original descriptions from model-derived features—thereby supporting both the fidelity of the semantic feature space and the efficacy of the optimization (fig. S10). The method generates descriptions optimally aligned with brain-decoded features, eliminating reliance on caption databases (*17*) or learned feature-to-text mapping modules (*18*)—both dependent on curated, labeled datasets that limit output diversity and blur attribution to brain-derived features—and consequently outperformed these alternatives across both video and image datasets (figs. S5–S7). Altogether, our approach balances interpretability, generalizability, and performance—establishing a transparent framework for decoding non-verbal thought into language and paving the way for systematic investigation of how structured semantics are encoded across the human brain.

A key feature distinguishing our method from previous visual-semantic decoding approaches is its ability to generate descriptions optimally aligned with brain-decoded features without relying on external caption resources (*17*) or additional learnable captioning modules (*18*). Unlike one-directional, single-pass generation with autoregressive LMs (*3, 4*), our use of a bidirectional MLM (*35*) enables full contextual integration from both directions and mitigates model-prior bias through iterative feature-guided optimization—achieving high-fidelity verbalization whenever the target representation lies within the searched space (fig. S10). Notably, this optimization process is robust enough to generate descriptions even from a non-informative initial

state (i.e., <unk> token; Fig. 2). While some limitations remain, emerging diffusion-based generators (*59–61*) may further expand the search space and enhance descriptive fidelity. Requiring curated labels only at the decoder-training stage, our framework is readily extensible beyond the visual domain to other sensory modalities (e.g., audition, touch) and cognitive domains (e.g., numerical reasoning, conceptual thought), offering a versatile framework for semantically grounded brain decoding across diverse forms of mental content.

While our method generates linguistic outputs through brain decoding, it differs from previous language decoding attempts as it does not rely on brain activity associated with language production (*2, 4*) and perception (*1, 3*). Instead, we trained decoders on brain activity induced by non-linguistic visual stimuli to predict semantic features linked to the visual content of viewed videos. This approach enabled the generation of descriptions for both viewed and recalled content without involving the language network (Figs. 3 and 4). Notably, all subjects in our study were non-native English speakers—specifically, native Japanese speakers—nevertheless, our method proficiently generated text outputs in English. These results illustrate that our method can directly translate non-linguistic brain semantics into linguistic descriptions, even when the subject is neither native nor proficient in the output language. Consequently, it can be applied to decode brain activity in non-linguistic subjects, including infants and animals, providing insights into how they develop the neural basis for processing complex visual semantics.

Moreover, by enabling the translation of non-verbal visual semantics in the brain into text, our method opens new communication channels in BMI applications, extending possibilities beyond traditional approaches for individuals with language or motor impairments. For instance, this approach could serve as an effective means of communication for individuals with aphasia, who struggle with language expression due to damage in language areas. Furthermore, our method complements visual-based BMI systems (*62*) by providing an alternative communication pathway for individuals with conditions like amyotrophic lateral sclerosis (ALS), where degeneration of motor-related activity limits the effectiveness of motor-based BMIs. Thus, our method holds the potential to enhance communication and interaction in clinical and assistive settings.

The present study sheds light on the neural bases of structured visual semantics by examining the effects of ablating the language network and the posterior parts of category-selective regions during text generation analysis (Fig. 3). Excluding these regions did not substantially impact the quality of the generated descriptions, despite the broad distributions of semantic-feature-predictive voxels across the cortex, including within these regions—consistent with previous studies (*46–49*). This outcome underscores the significance of brain areas outside these regions in representing structured visual semantics and aligns with research on visual representations involved in recognizing interactions and actions (*11–16*). Our findings build on these studies by demonstrating that such representations are sufficient for constructing fine-grained, cohesive descriptions of both viewed and recalled content.

Furthermore, our encoding analysis, which contrasts models based on LM-derived contextual semantic features with visual DNN features related to object and action categories, revealed a functional boundary within the category-selective areas that separates posterior from anterior regions (Fig. 3B). Although this analysis specifically focused on neural representations linked to visual stimulus perception, this boundary coincides with known distinctions between visual and linguistic semantics (videos versus audio stories) (*63*) and between perceptual and mnemonic systems (*64*). Consequently, our findings suggest an alternative perspective on this boundary:

posterior regions may primarily support isolated semantic (or categorical) representations, while anterior regions integrate these into contextualized representations. The proximity of these anterior regions to language areas suggests they may play a bridging role, transforming non-verbal information into verbal expressions and connecting non-verbal and verbal semantics. Further research is needed to clarify how these anterior regions interact with the language network to achieve this integration.

While our method has shown the ability to generate descriptions that resemble captions rated by the subjects as highly consistent with their perceptions (fig. S5E), there remains potential for improvement in capturing the full spectrum of subjective experiences, particularly by refining the alignment and depth of captions annotated to video stimuli. Specifically, we relied on captions provided by independent annotators, which may not fully align with each subject's unique perceptions, potentially affecting decoding performance. Although our use of rich annotations for each video (20 captions per video) likely mitigated some variability, training decoders on subjects' own reports might yield even closer alignment. Furthermore, because we instructed annotators to focus on visual content rather than subjective aspects like emotional reactions (*65*), the generated descriptions were predominantly concrete and rarely reflected abstract dimensions such as impressions and emotions (*66*). With annotations that more accurately reflect and encompass various dimensions of subjective experience, our method may capture the content of a subject's mind more comprehensively.

A limitation of our study is the use of natural videos sourced from the web (*36*). While this approach enhances ecological validity, it constrains our ability to precisely identify the relational structures captured by our method and to assess its generalizability to atypical scenes (e.g., "a man bites a dog"). Notably, the word orders of the generated descriptions were successfully optimized to align with brain-decoded features (Fig. 2G), yielding descriptions with higher discriminability than their shuffled variants (Fig. 2F). Although these results suggest that the generated descriptions reflect some relational structure, the lack of experimental control makes it difficult to determine which relational structures are being captured. Critically, it remains unclear whether the success of our method reflects true generalizability beyond common relational patterns, or instead relies on implicit biases toward typical scene structures. This potential bias could be introduced at any stage of the pipeline, such as through model priors, training data distribution, or the experimental design—including stimulus selection. Future studies should address these issues by incorporating strong out-of-distribution probes using controlled stimuli that systematically manipulate distinct relational structures—including contrasting or atypical configurations (*43, 67*). We consider such experiments a critical next step for a more rigorous assessment of both the type of relational information decodable from brain activity and the generalization capacity of our method across a broader range of relational contexts.

A potential concern with our demonstration of generating mental content descriptions is that the verbal prompts used to cue the target videos may have influenced brain activity during the imagery period due to the slow hemodynamic response. During the preparation period, as subjects might have started to recall the videos while reading these prompts, it is difficult to fully differentiate brain activity associated with text reading from that related to mental imagery. However, because our decoders were specifically trained on brain activity induced by non-linguistic visual stimuli, they prioritize semantic information directly linked to visual content over linguistic cues. Furthermore, our analysis showed that descriptions generated during the imagery period were of higher quality than those from the preparation period (fig. S12C and D), suggesting that volitional mental imagery effectively recruited the neural representations necessary for accurate descriptions. Nonetheless, future investigations applying our method to spontaneous mental imagery (e.g., mind-wandering or

dreaming) using subjects' verbal reports as a reference would be necessary to clarify its ability to generate descriptions of mental content free from the influence of external stimuli.

To accurately characterize our primary contribution, it is essential to frame our method as an interpretive interface rather than a literal reconstruction of mental content. While our approach enables the linguistic interpretation of non-verbal mental representations, the resulting outputs inevitably reflect not only brain-derived information but also prior knowledge of the world—often implicitly embedded in the experimental and modeling framework. These framework-dependent priors manifest in several aspects discussed above—such as the choice of language models, the language and type of annotations (e.g., English versus Japanese, visual versus emotional), and the properties of experimental materials—all of which collectively shape how brain representations are ultimately verbalized, particularly under supervised decoder training. For example, our finding that decoders trained on coherent and incoherent captions yielded correspondingly structured or unstructured outputs (Fig. 2 and fig. S3C) illustrates how the generated text is shaped by the world knowledge provided during training. Thus, the decoded content should be viewed not as a pure readout or reconstruction of brain states, but as a translation filtered through the lens of a specific interpretive framework. Our key contribution, therefore, should be understood not as the faithful recovery of the brain's intrinsic "language," but as the construction of a versatile and expressive pathway for interpreting non-linguistic mental representations—by leveraging the universality of natural language as instantiated through the semantic expressiveness and generative power of LMs. This interpretive bridge advances the field by enabling more flexible and nuanced renderings than previously possible, while simultaneously raising fundamental questions: How much of the decoded output truly originates in the brain, and how much reflects the constraints of our tools? Would decoding fidelity improve with more brain-aligned semantic spaces? How might an individual's linguistic profile—their native language, vocabulary, syntax, and expressive habits—systematically shape the decoded output? Addressing these questions requires not only improving decoding performance, but also critically examining the epistemological assumptions inherent in the frameworks we use to interpret the mind.

The generation of mental content descriptions was successful even from single-trial fMRI activity of mental imagery (Fig. 4F and fig. S14). However, this success raises ethical concerns regarding potential invasions of privacy. Key issues include the risk of unintentionally disclosing primitive thoughts before individuals have chosen to verbalize them. Additionally, unwanted biases inherent in the LMs (*68*) could distort the results within feasible optimization limits. Moreover, although requiring intensive data collection from willing participants may currently ensure consent (*3*), advances in inter-individual alignment technology could reduce this requirement (*69, 70*). Therefore, it is imperative to establish regulations that promote the ethical use of these technologies (*71*) while ensuring explicit informed consent and safeguarding subjects' mental privacy and autonomy in deciding which thoughts to disclose.

Materials and Methods

Subjects

Six healthy subjects (S1, male, aged 37–38; S2, female, aged 37–38; S3, male, aged 33–34; S4, female, aged 35–36; S5, male, aged 29–30; and S6, male, aged 22–23) with normal or corrected-to-normal vision participated in the experiments. All were native Japanese speakers and non-native English speakers; S1–S5 were proficient in English, while S6 had limited proficiency. All subjects provided written informed consent, and the study protocol was approved by the Ethics Committee of NTT (R03-004 and R03-009). The sample size was determined based on prior fMRI studies with similar protocols (*7, 66*). Data from each subject were collected over multiple scanning sessions spanning approximately six months. Experimental parameters and analytical pipelines were determined from a preliminary experiment with S1, who was exposed to the same stimuli multiple times, potentially influencing their brain responses.

Visual stimuli

Visual stimuli consisted of 2,196 short videos (all, 0.152 s to 90.1 s, mean = 6.61 s, median = 4.51 s; training, 0.152 s to 90.1 s, mean = 6.70 s, median = 4.62 s; test, 0.30 s to 20.1 s, mean = 3.94 s, median = 2.90 s) from a previous study (*36*) (https://goo.gl/forms/XErJw9sBeyuOyp5Q2). These videos covered diverse content (objects, scenes, actions, and events) and were resized to fit a 16-degree visual angle, maintaining the original aspect ratio. They were presented at the center of a gray background without sound. Sixteen duplicates were excluded to avoid redundancy, resulting in the final set of 2,180 unique videos used in our experiment.

Experimental design

We conducted two main experiments: a video presentation experiment and an imagery experiment (fig. S1). Visual stimuli were shown on an LCD monitor at the rear of the fMRI scanner, and audio stimuli were delivered via S14 earphones (Sensimetrics). Each experimental session lasted up to 2 hours. Subjects were given adequate time for rest between runs (every 8–10 minutes) and could take a break or stop the experiment at any time. The total duration for both experiments was approximately 17.1 hours.

Video presentation experiment. The video presentation experiment included training sessions (60 runs) and test sessions (10 runs). Each run had 36 or 37 stimulus blocks and 3 or 4 randomly inserted evaluation blocks, averaging 695.2 s per run. Subjects spent approximately 11.8 hours on training sessions and 1.8 hours on test sessions.

In each stimulus block, videos shorter than 10 s were repeated until the total duration exceeded 10 s. Videos longer than 10 s were presented once, followed by less than 1 s of rest, making the block duration divisible by 1 s (TR). Subjects viewed the stimuli without fixation to recognize details.

In each evaluation block, five descriptions (both in English and translated into Japanese) depicting the visual contents of videos were presented, and subjects were then asked to rate how consistent each description was with what they had perceived in the preceding stimulus block using two button boxes. The descriptions were taken from twenty captions for the preceding video, with occasional ones from other videos. Subjects rated the descriptions on a five-point scale or marked them with an “x” if they were unrelated (with higher scores indicating closer alignment with their perception; random initial score). Subjects completed the ratings at their own pace and proceeded to the next block by selecting “Proceed?” and pressing a button.

Each block was followed by a 2-s rest period, with 32-s and 8-s rest periods at the beginning and end of each run, respectively. During the rest periods, the subjects were

instructed to maintain fixation on a central spot, which consisted of a bull's eye and crosshairs, in order to keep their attention focused on the screen.

In the training session, 2,180 unique videos were each presented once in a pseudo-randomized order. This order remained consistent for all subjects. Evaluated descriptions were also consistent for all subjects. In the test session, 72 videos from the last two runs of the training session were each presented five times, divided between two runs, and shown in a pseudo-random order within each run (see Materials and Methods "MRI data preprocessing" for the data handling of the overlapping data in the training session).

Imagery experiment. The imagery experiment comprised 30 runs, each with 12 trials consisting of a preparation block, an imagery block, a video presentation block, and an evaluation block. Subjects were required to engage recall-based visually imagery of one of the 72 videos presented during the test session of the video presentation experiment. Each run averaged 434.6 s, totaling approximately 3.6 hours per subject.

During a preparation block, a verbal description (both in English and translated into Japanese) of a target video was presented to subjects to prompt them to prepare mental imagery of the visual content of the target. The descriptions were selected from the set of twenty captions collected for each video and were consistent across trials and subjects. Subjects were encouraged to imagine all details, even those not explicitly described, to mimic the video presentation experiment. The description served as a guide to vividly imagine the complete visual content of the target. Subjects pressed a button when ready, and a beep sound with less than 1 s of rest signaled the start of the imagery period, aligning with the TR.

During the imagery period, subjects recalled the visual content of the target video with their eyes closed as if actually watching it. The imagery block duration matched the stimulus block duration in the video presentation experiment (repeatedly recalling for videos shorter than 10 s, once for videos longer than 10 s), with an additional 2 s to ensure full replay. A beep signaled the end of the imagery period, prompting subjects to open their eyes.

After the imagery block, the target video was presented as in the video presentation experiment, allowing subjects to compare their mental imagery with the actual video. This was followed by two 3-s blocks where subjects rated the accuracy and vividness of their imagery on a five-point scale. Subjects adjusted the score from its random initial setting using a button box in their right hand.

Each imagery, stimulus, and evaluation block was followed by a 2-s rest period, with 32- and 8-s rest periods at the beginning and end of each run. Subjects maintained fixation on a central spot, as in the video presentation experiment.

Before the imagery experiment, subjects practiced associating each target video with its verbal description, viewing the pairs during inter-run rest periods to aid memory. The 72 videos were randomly distributed among six runs, with each set of six runs containing all videos in a pseudo-random order.

Retinotopy and functional localizer experiments. In addition to the main experiments, we conducted a retinotopy experiment and three functional localizer experiments (visual category, MT+, and language area localizers) to delineate visual areas and localize regions of interest (ROIs).

Retinotopy. We followed the Human Connectome Project 7T Retinotopy Dataset protocol (*72*), using dynamic colorful textures through moving apertures (wedge, ring, bar) in eight 300-s runs. This identified retinotopic maps (V1, V2, V3, V3A, V3B, hV4, and V7) on cortical surfaces using *fsfast* retinotopy analysis [*Freesurfer* (*73*)] and population receptive field analysis (code is available at https://kendrickkay.net/analyzePRF/) (*74*).

Visual category localizer. Following the fLoc protocol (*75*) (stimuli and code are available at http://vpnl.stanford.edu/fLoc/), we presented images from word, body, face, place, and object categories in eight 300-s runs (48 blocks each, 6 s per stimulus or blank, with 6-s initial and final rests). Additional intact and scrambled object conditions were included to localize object-selective areas (original images are available at tarrlab, https://sites.google.com/andrew.cmu.edu/tarrlab/stimuli). The contrasts between word/body/face/place and others (from these four categories) were used to define visual category-selective areas (word: visual word form area [VWFA], occipital word form area [OWFA]; body: extrastriate body area [EBA], fusiform body area [FBA]; face: fusiform face area [FFA]; occipital face area [OFA]; place: parahippocampal place area [PPA], occipital place area [OPA], medial place area [MPA] consisted of the retrosplenial cortex [RSC] and parieto-occipital sulcus [POS]). The contrast between intact and scrambled objects was used to define an object-selective area (lateral occipital complex [LOC]).

MT+ localizer. Following Tootell et al. (1995) (*76*), we presented random dot stimuli in three conditions (moving, dynamic, static) in four 232-s runs (13 blocks each, 12 s per stimulus, with 12-s initial and final rests). The contrast between moving and dynamic/static conditions defined the visual motion area MT+.

Language area localizer. We followed protocols by Fedorenko et al. (2010) (*77*) and Scott et al. (2017) (*78*), modifying to include both visual and auditory stimuli and to use Japanese stimuli in each of eight 358-s runs (19 blocks each, 18 s per stimulus or 14-s blank, with 14-s initial and final rests). Subjects read sentences or nonword sequences and listened to intact or degraded auditory passages. The contrasts between sentence/intact and nonword/degraded conditions defined language-sensitive areas in the temporal and frontal cortices.

Voxels from V1, V2, and V3 were combined as the lower visual cortex (LVC); voxels from VWFA, OWFA, EBA, FBA, FFA, OFA, PPA, OPA, MPA, LOC, and MT+ were combined as the higher visual cortex (HVC); voxels from temporal and frontal language areas were combined as the language network. Overlapping voxels with LVC were excluded from HVC.

MRI acquisition

MRI data were collected using a 3.0-Tesla Siemens MAGNETOM Prisma scanner located at the WPI-IRCN Human fMRI Core, the University of Tokyo Institutes for Advanced Studies. An interleaved T2*-weighted gradient-echo echo planar imaging (EPI) scan was performed to acquire functional images covering the entire brain (TR, 1000 ms; TE, 30 ms; flip angle, 65 deg; FOV, 192 × 192 mm; voxel size, 2 × 2 × 2 mm; slice gap, 0 mm; number of slices, 72; multiband factor, 6). T1-weighted (T1w) magnetization-prepared rapid acquisition gradient-echo (MP-RAGE) fine-structural images of the entire head were also acquired (TR, 2000 ms; TE, 1.97 ms; TI, 900 ms; flip angle, 10 deg; FOV, 256 × 256 mm; voxel size, 1.0 × 1.0 × 1.0 mm).

MRI data preprocessing

For anatomical data, we first used *SPM12* to preprocess each of the T1w anatomical

images of individual subjects for bias-field correction and for re-defining its origin and orientation to be set on the anterior commissure and the anterior commissure - posterior commissure (AC-PC) line, respectively. Cortical surface meshes were generated from the processed T1w images using *Freesurfer* (version 7.3.2) (*73*) with manual corrections for anatomical segmentations. Analytical results were visualized on flattened cortical surfaces, created by making relaxation cuts in each hemisphere, with functional data aligned and projected using *Pycortex* (*79*).

For the functional data from each run, we performed the MRI data preprocessing through the pipeline provided by *fMRIPrep* (version 20.2.7) (*80*). First, a BOLD reference image was generated using a custom methodology of *fMRIPrep*. A field map (B0-nonuniformity map) estimated based on a phase-difference map was used to estimate susceptibility distortion and to correct the BOLD reference for a more accurate co-registration with the anatomical reference. The BOLD reference was then co-registered to the T1w reference using *bbregister* (FreeSurfer; version 7.3.2), which implements boundary-based registration (*81*). BOLD runs were slice-time corrected using 3dTshift from AFNI 20160207 (*82*), and the BOLD time series were resampled onto their original, native space (2 × 2 × 2 mm voxels) by applying a single, composite transform to correct for head-motion and susceptibility distortions using antsApplyTransforms from ANTs (version 2.3.3) with Lanczos interpolation.

To create data samples, we first discarded the first 8-s scans of the preprocessed BOLD signals from each run to avoid MRI scanner instability. We then regressed out nuisance parameters from each voxel amplitude for each run, including a constant baseline, a linear trend, and 24 head-motion parameters (three rotations, three translations, their temporal derivatives, and quadratic terms) and 12 global signals (mean amplitudes within cerebrospinal fluid, white matter, gray matter, their temporal derivatives, and quadratic terms). The data samples were temporally shifted by 4 s to account for hemodynamic delays, despiked to reduce extreme values (beyond ± 3 SD for each run), and averaged within each stimulus and imagery block. Finally, each voxel's amplitude was z-scored within each run to eliminate potential non-stationarities and scanner-specific biases.

For data from the training session of the video presentation experiment (training perception data), we discarded samples from the last two runs, in which videos used in the test session and the imagery experiment were presented, to ensure generalization to new stimuli. For test data from the video presentation experiment (test perception data) and the imagery experiment (test imagery data), we averaged samples of identical video clips (5 repetitions) to increase the signal-to-noise ratio of the fMRI signals unless otherwise stated.

fMRI brain activity data from Natural Scenes Dataset

We additionally employed the Natural Scenes Dataset (NSD) (*39*), which contains high-resolution fMRI signals from eight subjects measured across 30 to 40 recording sessions. Our analysis focused on data from four subjects (Subjects 1, 2, 5, and 7) who completed all 40 sessions. We analyzed image-induced fMRI activity while these subjects viewed images from the MS-COCO (*83*). We utilized pre-processed fMRI signals in 1.8 mm native volume space corresponding to "nsdgeneral" brain areas, which included approximately 15,000 voxels per subject in the posterior cortex responsive to visual stimuli. Using the code from Ferrante et al. (2023) (*18*) (https://github.com/enomodnara/BrainCaptioning), we constructed fMRI data samples of image-induced activity corresponding to 8,859 images for training and 982 images for testing. Since these stimulus images were presented up to three times, we averaged the fMRI signals from multiple trials to enhance the signal-to-noise ratio. For the

analysis with NSD, we used the API of ChatGPT (gpt-4o-mini; https://chat.openai.com/) to proofread and refine the captions, as the MS-COCO captions often contain typos and grammatical errors.

Caption annotation for visual stimuli

We used Amazon Mechanical Turk (AMT) to collect written captions for stimulus video clips, following the procedure used for Microsoft COCO Captions (*84*). Multiple workers viewed each video to provide a detailed sentence (over eight words) describing the visual content. The captions were manually checked for quality and proofread with the assistance of ChatGPT (GPT-3.5; https://chat.openai.com/; prompt: "Proofread the following:") to correct typos and remove incorrect or unintelligible sentences. We collected twenty unique captions per video, matching the number in the MSR-VTT dataset (*85*). The collected video captions are available from our repository (https://github.com/horikawa-t/MindCaptioning).

Feature computation by deep neural network models

We used deep neural network models pre-trained for language (e.g., DeBERTa-large) or vision (TimeSformer) (*55*) tasks to compute semantic and visual features, respectively. Additionally, we used the image encoder of a multi-modal model (CLIP) (*57*) to compute visuo-semantic features. To mitigate biases arising from baseline differences across model units, we applied *z*-score normalization to the feature values using means and standard deviations estimated from respective training data for each analysis.

Semantic features. To extract semantic features from video captions, we used 42 pre-trained language models (LMs; available at Hugging Face's Transformers library, version 4.30.2) (*86*). These models cover a range of architectures (e.g., BERT and GPT2) and sizes (e.g., base, large, and xlarge). Each input sequence was tokenized and processed by an LM to produce vector embeddings for each token across multiple layers. Following Reimers & Gurevych (2019) (*87*), we averaged the embeddings over tokens, excluding special tokens (e.g., <cls> token), in each layer. The averaged embeddings from multiple layers were used as semantic features for the input sequence. For each video, we computed semantic features for 20 annotated captions and averaged them to construct the final semantic features for the video.

The 42 LMs used in the present study were based on the following model families: BERT (*35*), RoBERTa (*36*), DeBERTa (*37*), ALBERT (*88*), OpenAI-GPT (*89*), GPT2 (*90*), Sentence GPT (*91*), XLnet (*92*), DistilBERT (*93*), DistilGPT2 (*86*), T5 (*94*), BART (*95*), CTRL (*96*), XLM (*97*), XLM-RoBERTa (*98*), ELECTRA (*99*), CLIP (*57*) text encoder. See fig. S8 for the full list of the LMs.

We hypothesized that an LM closely aligned with the human brain will provide more effective intermediate representations for translating visual semantic information in the brain into text. We thus performed a cross-validation encoding analysis using semantic features from each of the 42 LMs within the training perception data. Based on the results of the validation analysis (fig. S2A), we selected the DeBERTa-large model, which demonstrated the highest performance.

To construct semantic features without structured semantic information (fig. S3C), we used captions with randomly shuffled word orders. For each caption, we created up to 1,000 word-order shuffled variants, computed their semantic features, and averaged these features across all the variants. These features were then averaged across twenty captions for each video to obtain semantic features for each video.

Vision model features. To extract visual features from video stimuli, we used a TimeSformer model (*55*) (model and code are available at https://github.com/facebookresearch/TimeSformer) pre-trained for object and action recognition using *ImageNet* (*100*) and Kinetics-400 (*101*). This model has demonstrated high performance in predicting video-induced brain activity (*102*). For each video, we resized its spatial size to 224 pixels while preserving the aspect ratio, and selected frames at intervals of 32. If the video had fewer than eight temporal positions at 32-frame intervals, we uniformly selected eight frames to cover the entire video length. We computed feature vectors for each layer from these frames and averaged them over the spatial dimension within each layer. This procedure was repeated for three spatial crops (left-center-right or top-center-bottom). Finally, we averaged the visual features over the temporal dimension and the three spatial positions, resulting in 768-dimensional features for each of the 12 layers.

Visuo-semantic model features. To extract visuo-semantic features from video stimuli, we used the image encoder of CLIP (ViT-B/16) developed by OpenAI (available at Hugging Face's Transformers library, version 4.30.2) (*86*). We primarily employed this model to replicate the non-linear image captioning-based decoding method, known as *Brain Captioning* (*18*). Since this approach was originally designed for image-induced brain activity, we adapted it for video-induced brain activity by averaging features across time (or multiple frames). Specifically, we extracted features from the final layer of the vision encoder for each frame using the default preprocessing and concatenated the resulting 768-dimensional embeddings from 197 tokens, including the [CLS] token and 196 patch tokens. We then averaged these frame-wise features, resulting in a 151,296-dimensional visuo-semantic feature vector for each video.

Voxel-wise encoding modeling analysis

We performed voxel-wise encoding modeling analysis by constructing encoding models that predict signal amplitudes of individual voxels from a feature vector in each model layer using the L2 regularized linear regression algorithm (ridge regression). The analysis was performed using both cross-validation and generalization approaches. In the cross-validation analysis, we used 6-fold cross-validation on the training perception data (58 runs divided into five sets of 10 runs and one set of 8 runs). In the generalization analysis, we trained models on all of the training perception data and tested them on the test perception data (fig. S8). We evaluated the model performance of each voxel by calculating Pearson correlation coefficients between measured and predicted brain activities of that voxel.

The regularization parameters of ridge regression were determined separately for each layer, model, and subject by considering the performance of all voxels on the respective training data. Models for individual voxels were trained using ten possible regularization coefficients (log spaced between 10 and 10000). The regularization parameters that produced the maximal model performance (mean correlation coefficients averaged across all voxels) on the training data were used for predictions on the test data. In the cross-validation analysis, we used a 5-fold cross-validation (inner loop) nested within a 6-fold cross-validation (outer loop). In the generalization analysis, we used 6-fold cross-validation on the training perception data.

For each model, predictions from multiple layers were integrated by selecting the best layer for each voxel based on model performance in the training data. In each fold of the cross-validation analysis, we determined the best layer per voxel from the 5-fold nested cross-validation loops (cf., Fig. 3F). We aggregated predictions from these best layers for each left-out set to construct final predictions for all data samples. In the generalization analysis, the best layer per voxel was determined from the 6-fold cross-

validation on the entire training perception data, and predictions from the best layers were used for the test perception data.

Encoding accuracies of semantic and visual models were compared using slope angles of the best linear fit estimated by Deming regression (*103*), which accounts for observation errors on both axes. The slopes were converted to angles, then subtracted from 45 degrees to obtain deviations from parity (Fig. 3E).

Feature decoding analysis

We performed feature decoding analysis by constructing a set of L2 regularized linear regression models (decoders) that predict feature values for each layer of each model from fMRI activity patterns (one decoder for each model unit). Both cross-validation and generalization analyses were performed to produce predictions for all training perception data samples as well as the test perception and imagery data samples. Decoders were trained using whole-brain fMRI voxel patterns (unless otherwise stated), selecting up to 50,000 voxels best predicted by the target feature set in cross-validation (or nested cross-validation) encoding analysis on the respective training data. Performance was evaluated by calculating Pearson correlation coefficients between feature values computed by an LM and predicted from the brain for each model unit.

In the cross-validation analysis, we used 6-fold cross-validation on the training perception data, with model training and the regularization parameter estimation conducted using a 5-fold cross-validation procedure nested within the 6-fold cross-validation, similar to the encoding analysis. In the generalization analysis, models were trained on all the training perception data and tested on the test perception and test imagery data, with regularization parameters determined using a 6-fold cross-validation within the entire training perception data. Regularization parameters of the ridge regression models were optimized based on model performance on the respective training data, with coefficients estimated separately for each layer, model, and subject considering the performance of all units.

Text generation analysis

To generate descriptive text based on a set of semantic features (target features), we conducted an iterative optimization of descriptions, in which semantic features of candidate descriptions were progressively aligned with target features through the iterative replacement and interpolation of tokens (referred to as "words" for simplicity) within the candidate descriptions. Each step of the optimization process consisted of three stages: masking, unmasking, and candidate selection.

In the masking stage, for each of the current candidate descriptions (e.g., "metal shapes"), we first generated an exhaustive list of masked candidates by replacing each word or a sequence of words (three words at the maximum) with a mask token (e.g., "<MASK> shapes," "metal <MASK>," and "<MASK>") or interpolating a mask token between words or at the top or bottom of the description (e.g., "metal <MASK> shapes," "<MASK> metal shapes," and "metal shapes <MASK>"). This masking procedure was repeatedly applied on the newly generated masked candidates (two times at the maximum) to generate masked candidates with multiple masks (e.g., "<MASK> <MASK> shapes" and "metal <MASK> <MASK>"). From all the generated masked candidates for each original candidate, we randomly selected five masked candidates to be processed in the next unmasking stage.

In the unmasking stage, we used an LM pre-trained for masked language modeling (MLM model; not necessarily the same as the model used for the feature computation) to generate alternative words to fill in the masks in the masked candidates within the

context of surrounding words. For each mask within a masked candidate description, we generated five alternative words for the mask to create five new candidates by random sampling from a categorical distribution likelihood estimated by the MLM model. In cases where a masked candidate had multiple masks, we processed the masks sequentially from the top until all masks were updated. These procedures yielded five new candidates from one masked candidate. In the main analysis, we used the pre-trained RoBERTa-large model (vocabulary size, 50265; subword segmentation based on byte pair encoding) for the guide of the text generation because this model consistently demonstrated stable performance in optimizing descriptions in a validation analysis (fig. S2B).

In the candidate selection stage, we computed semantic features of all new and original candidate descriptions using an LM from which the target features originated. We then calculated Pearson correlation coefficients between those candidate features and target features for all layers and averaged those correlations over layers to score the candidate descriptions. To enhance the conciseness of generated descriptions, reduce the computational cost of handling long descriptions, and avoid overfitting to noises on brain-decoded features, we added an exponential penalty to the length of candidate descriptions as described by $s = \frac{r}{l^{\alpha}}$, where s is the similarity score used to rank candidates, r is the mean correlation coefficient between candidate features and target features averaged across layers, l is the length (or the number of tokens) of the candidate, and α is the parameter for the length penalty. We chose the parameter $\alpha = 0.1$ based on the validation analyses using six possible penalty parameters (0, 0.05, 0.1, 0.15, 0.2, and 0.25) with a randomly selected subset of 50 samples from the training perception data (fig. S2C–E). After computing similarity scores for all candidates, we ranked them and selected the top five candidates to proceed with further optimization.

During each step of the optimization process, the maximum search width was 130 (five new candidates for each of the five masked candidates derived from five original candidates, in addition to the five original candidates themselves).

We repeated these optimization stages 100 times, and the obtained description was taken to be the text describing the semantic information represented in the target features, or the brain. Since the masking and unmasking stages involve randomness in the optimization process, to avoid local optima we repeated the same process five times for each data sample to select the description showing the highest similarity scores with the target features as the final prediction. For all the analysis, we began the optimization process from a non-informative initial state (i.e., unknown token, <unk> for the tokenizer of the RoBERTa-large) to avoid incorporating any prior assumptions for description generation.

Unlike autoregressive LMs with causal attention often used for linguistic information decoding (*3, 4*), MLM models with bidirectional attention have the substantial advantage of incorporating contextual information from all surrounding words (*35*). This characteristic makes our optimization process more suitable for decoding visual information, which lacks specific directionality, in contrast to linguistic information.

Database-search-based description prediction analysis

Among the notable strengths of our method are its flexibility in optimizing descriptions at the word (or token) level and its ability to generate word sequences that do not currently exist in the databases. To assess its effectiveness, we used the database-search (DB-search) method (*17*) to set a baseline performance for description prediction from the brain. This analysis involved searching for captions with the highest feature correlations with target brain-decoded features from large databases of image and video

captions [MS-COCO (*83*); GCC (*104*); MSR-VTT (*85*)]. We computed semantic features for all captions (~4.1M) using the DeBERTa-large model. For each fMRI data sample, we predicted semantic features at multiple layers using trained decoders and computed correlation coefficients between the decoded features and features computed from all database captions. The caption with the highest mean correlations across all layers was selected as the prediction for the sample. Note that the brain-decoded features used in this analysis were the same as those used in the main analysis (i.e., mind captioning; e.g., Fig. 2).

We also performed the same analysis on fMRI data from the NSD (*39*) (fig. S7), employing candidate captions from the GCC image database (*104*) to match the experimental conditions used in the previous study (*17*).

Non-linear image captioning-based description prediction analysis

Another strength of our method is that it does not require training an additional module, such as non-linear captioning models, by directly optimizing word sequences to match brain-decoded features. To examine the effectiveness of this aspect, we compared its performance with another brain-to-text decoding approach based on a non-linear image captioning model (*Brain Captioning*) (*18*). We implemented this method using code from the authors' repository (https://github.com/enomodnara/BrainCaptioning). This approach utilizes an independently trained image captioning model, the *Generative Image-to-text Transformer* (GIT) (*58*), which generates descriptions by processing visuo-semantic features extracted from CLIP (*57*). We computed visuo-semantic features for all stimuli (videos or images) using CLIP. For each fMRI data sample, we predicted visuo-semantic features from the final layer using trained decoders and inputted these features into GIT to produce a description.

Evaluation of the similarity between generated descriptions and references

We used multiple metrics to evaluate the similarity of generated descriptions to reference captions annotated to the viewed or recalled videos. These metrics include Pearson correlation (feature correlation), BLEU, METEOR, ROUGE-L, CIDEr, and three variants of BERTScore. For each predicted description, we computed scores for each metric against twenty reference captions of the corresponding video, selecting the highest score as the final score. BLEU and METEOR were computed using the NLTK toolbox (*105*), while ROUGE-L and CIDEr were computed using code from https://github.com/salaniz/pycocoevalcap. BERTScore was computed using code from https://github.com/Tiiiger/bert_score.

Feature correlation. We defined the feature correlation as the mean of Pearson correlation coefficients between semantic features of a reference caption or target features and those of a generated description averaged across multiple layers. We computed semantic features using the same LM as in the decoding analysis, unless otherwise stated. Feature correlation was used in all analyses where the similarity metric was not explicitly specified.

BLEU. BLEU computes the precision scores by comparing predicted n-grams with reference captions while considering a brevity penalty. We used the 4-gram variant (BLEU-4) with a smoothing method (*106*).

METEOR. METEOR computes scores based on unigram matching in predicted and reference sequences using precision and recall while considering word variations like stemming and synonymy.

ROUGE-L. ROUGE-L computes scores by emphasizing recall, measuring the overlap in words between the prediction and reference, and primarily focusing on their longest shared sequence.

CIDEr. CIDEr measures the similarity between predicted descriptions and references through n-gram-based comparison while considering consensus across multiple references.

BERTScore. BERTScore computes similarity using contextualized embeddings from individual tokens by a bidirectional transformer LM. We used the 17th layer of the RoBERTa-large model with baseline rescaling to compute scores (default of the official implementation) and evaluated three variants: P (precision), R (recall), and F1. To evaluate the token-wise precision of generated descriptions, we used the precision (P) weighted by the inverse of document frequency (IDF) estimated from captions in multiple databases (MSCOCO, GCC, and MSR-VTT) and captions collected in this study. The IDF-weighted P was only used to highlight tokens with high precision in generated descriptions, but not used in quantitative evaluations.

Based on these metrics, we evaluated the discriminability of the generated descriptions. For each description, we calculated similarity scores between the generated description and reference captions of the corresponding viewed or recalled videos (correct) and irrelevant videos (incorrect; $N = 2{,}179$). Discriminability was defined as the difference between the score for the correct video and the mean score averaged across all incorrect videos.

Text-based video identification analysis

To evaluate description generation performance, we performed a video identification analysis based on the similarity between generated descriptions and reference captions using multiple metrics. For each data sample, we computed similarity scores between the generated description and all reference captions of all videos using a specific metric. We compared the similarity scores to the correct reference captions (20 captions for the target video) with those to incorrect reference captions (43,580 captions for 2,179 irrelevant videos). For the analysis with feature correlation and BERTScore, we used mean similarity scores averaged across multiple captions per video. For BLEU, METEOR, ROUGE-L, and CIDEr, we used the highest matching scores among multiple captions per video. The analysis was conducted with varying numbers of candidates, ranging from 2 (chance level, 50%) to 100 (chance level, 1%), selecting the video with the highest similarity score as the prediction. Accuracy was defined as the proportions of correct video identification. For two candidates (one correct and the other incorrect), we performed identification for all combinations of correct and incorrect candidates. For more than two candidates, we randomly selected incorrect candidates, repeating the analysis 100 times to estimate the mean accuracy averaged across repetitions.

Feature-based video identification analysis

To evaluate the generalizability of feature decoders trained on stimulus-induced brain activity for decoding imagery-induced activity, we performed a feature-based video identification analysis and compared accuracies between perception and imagery conditions (cf. Fig. 4E). The analysis was conducted separately for each model and layer by comparing brain-decoded feature vectors with feature vectors derived from video captions or visual stimuli using three types of models (visual, visuo-semantic, and semantic). For each data sample, we calculated Pearson correlation coefficients between the decoded feature vector and the model-derived feature vectors of all candidate videos ($N = 2{,}180$). Identification was performed in a pairwise manner by determining, for

each pair consisting of the correct video and one of the 2,179 incorrect videos, whether the correlation with the correct video was higher (chance level = 50%). Accuracy was defined as the proportion of pairs in which the correct video had a higher correlation than the incorrect one. For each model layer, we assessed the generalizability from perception to imagery by measuring the angle between the parity line (i.e., the diagonal line indicating equal accuracy for the two conditions) and a line connecting the chance level to the observed accuracies for the perception and imagery conditions. Smaller angles indicate stronger generalizability.

Evaluation of the relational information in generated descriptions

To assess whether the generated descriptions accurately represented visual relationships among individual components in viewed or recalled videos, we evaluated the effect of shuffling word orders of generated descriptions on discriminability and video identification performance. For example, distinctions such as "a bird eats a snake" versus "a snake eats a bird" or "some grass in a mug" versus "a mug in some grass" (*43*) are sensitive to word order and thus indicative of relational structure. Disrupting this word order through shuffling can therefore help reveal whether such structured information is accurately described in the generated text. We randomly shuffled the word order of each generated description to create word-shuffled variants. The shuffling was performed at the word level, not the token level, to maintain minimal coherence within individual words. For each original description, we created up to 1,000 shuffled variants by shuffling all words or only nouns, excluding descriptions with only one noun from the noun-shuffling analysis. Nouns were identified using part-of-speech tagging with *spaCy* (*107*) (version 2.2.4). We then computed semantic features for these shuffled descriptions.

If generated descriptions accurately capture visual relations among individual components in videos, we should observe higher similarity to correct captions with the original descriptions compared to the shuffled ones while maintaining differences with irrelevant captions. To examine this, we computed feature correlation scores between reference captions and both original and shuffled descriptions to examine if the original descriptions exhibited higher discriminability (e.g., Fig. 2F right). Given the varying levels of sentence structure disruption among the shuffled descriptions, we also conducted the same analysis using the least disrupted shuffled sentences to rigorously assess the superiority of the original descriptions. Specifically, we selected the shuffled sentence with the highest pseudo-log-likelihood score—a fluency (or linguistic acceptability) metric computed by MLM scoring (*44*)—from 1,000 shuffled variants for each generated description (e.g., fig. S4F). We employed an adapted metric proposed by Kauf & Ivanova (2023) (*45*). Additionally, to quantify the impact of shuffling on identification accuracy, we performed video identification analysis using both original and shuffled descriptions (e.g., Fig. 2F left).

We also examined if the word order of generated descriptions was unduly influenced by the MLM model used to support the text generation, potentially diverging from the information represented in the brain. We reasoned that if a generated description closely matches the brain representation and contains semantic information uniquely depicted by the generated word order—rather than alternative arrangements of the same words—then brain-decoded features should exhibit higher similarity to features of the generated description than features of shuffled variants. To test this, we computed feature correlation scores between target brain-decoded features and both original and shuffled descriptions to see if the original exhibited greater scores. To consider the degree of meaning changes introduced by shuffling, we also computed feature correlation scores between original and shuffled descriptions, defining correlation distance as one minus the feature correlation (Fig. 2G and fig. S4G).

Evaluation of the diversity in generated descriptions

To evaluate the diversity of generated descriptions (fig. S5C, D), we used *Self-BLEU* (*108*)—a metric assessing the diversity of the generated text data—to compute the sentence (dis)similarity either across subjects for each video or across videos for each subject. For each video/subject, we can compute a BLEU score by regarding one description from a subject/video as a hypothesis and descriptions from the other subjects/videos as references. To evaluate the diversity across subjects/videos, we computed BLEU scores for every generated description from different subjects/videos and defined the average BLEU scores across subjects/videos as the Self-BLEU of the video/subject, respectively. A higher Self-BLEU score indicates less diversity in the descriptions generated for the video/subject. In this analysis, we used the same analytical settings as we did when using the BLEU for the similarity evaluation (BLEU-4 with smoothing).

Evaluation of the consistency between generated descriptions and subjective perception

The reference captions in this study were collected from subjects in an independent online experiment, not from our fMRI subjects. Therefore, not all captions may precisely match the subjective perception of our fMRI subjects, even though they were used as "correct" references in the evaluation. To examine if the generated descriptions from video-induced brain activity were consistent with the subjective perceptions of individual subjects, we investigated the relationship between subjective ratings from the video presentation experiment and the similarity of generated descriptions to evaluated captions. We focused on 212 videos evaluated during the training perception data collection, analyzing descriptions generated from the cross-validation decoding analysis using feature correlations as the similarity metric. For each generated description, we computed feature correlations against five evaluated captions, yielding 1060 scores per subject. These scores were classified according to individual ratings to explore if higher-rated captions had higher scores. Additionally, we calculated Pearson correlation coefficients between the feature correlation scores and subjective ratings to determine the presence of positive correlations (fig. S5E).

Illustrations

Due to copyright restrictions, the images shown in this paper are not actual frames from the video stimuli used in our experiments. Instead, they are schematic illustrations manually created by a professional scientific illustration company (Medical Education, Tokyo, Japan), based on the captions of the original videos. All illustration copyrights have been transferred to the author.

Statistical Analysis

Statistical analysis was performed individually unless otherwise stated, with results from six subjects considered replications (*109*). We reported quantitative results for each subject and averages across subjects, except for the validation analysis (fig. S2C–E). To account for multiple comparisons, we used the Benjamini-Hochberg method (*110*) to control the false discovery rate (FDR) and provided this information where applicable.

Statistical significance between results of pre-trained and untrained MLM models was tested using a one-tailed Wilcoxon signed-rank test on feature correlations of generated descriptions after 100 optimization iterations ($N = 72$; fig. S4A).

Discriminability based on generated descriptions was also tested using a one-tailed Wilcoxon signed-rank test ($N = 72$; e.g., fig. S4B). The effect size of discriminability (e.g., Fig. 2D) was computed by calculating similarity scores between a generated

description and both sets of correct ($N = 20$) and incorrect captions ($N = 43{,}580$). These scores were averaged separately for each set. These averaged scores for correct and incorrect sets from all test samples ($N = 72$) were used to estimate means and standard deviations for computing Cohen's d for discriminability.

Video identification analysis results were presented with 95% C.I. across samples ($N = 72$) to determine if the mean accuracy exceeded the chance level (e.g., Fig. 2E).

We used one-tailed Wilcoxon signed-rank tests to evaluate the impact of word-order shuffling on discriminability ($N = 72$; e.g., Fig. 2F) and the diversity differences in generated descriptions ($N = 72$ for self-BLEU across subjects; $N = 6$ for self-BLEU across videos; fig. S5C, D).

Correlations between discriminability and encoding accuracy across multiple LMs were evaluated using one-tailed t-tests after Fisher's z transform ($N = 42$; fig. S8B). The correlation between layer depth and both discriminability and its drop caused by shuffling were tested using a one-tailed t-test (fig. S9).

The variability of the generalization angle was estimated using a jackknife resampling procedure. Specifically, one of the 72 samples—defined as the identification accuracy for each pair of viewed and imagined videos—was sequentially left out, and the angle was computed from the remaining 71 samples. The standard error was then estimated from the distribution of these leave-one-out angle values (e.g., Fig. 4E).

To evaluate the similarity between the generated descriptions and rated captions, we pooled results from all six subjects to ensure sufficient data samples to detect differences, while also accounting for variations in the number of samples across rating levels and subjects (fig. S5E). Differences in feature correlations across subjective ratings were tested using one-tailed t-tests after Fisher's z transform. Correlation between ratings and feature correlations between generated descriptions and rated captions were tested using a one-tailed t-test after applying Fisher's z transform ($N = 6{,}360$). Interactions between text generation methods and ratings were tested using ANOVA.

For encoding analysis, the correlation between measured and predicted fMRI signals for each voxel was tested using a one-tailed t-test after Fisher's z transform ($N = 2{,}108$; Fig. 3A, B). Mean encoding accuracy within each brain area was presented with a 95% C.I. across voxels (Fig. 3D). Comparisons between semantic and visual encoding models were based on the slope angles of linear fits, converted to deviations from parity (Fig. 3E). Statistical significance of differences in the best layers exhibiting highest encoding accuracy was tested using a one-tailed Wilcoxon rank-sum test (Fig. 3F).

Data availability

The experimental data that support the findings of this study are available from open data repositories (raw data in OpenNeuro: https://doi.org/10.18112/openneuro.ds005191.v1.0.2; preprocessed data in figshare: https://doi.org/10.6084/m9.figshare.25808179).

Code availability

The code that supports the findings of this study is publicly available at our repository (https://github.com/horikawa-t/MindCaptioning) and archived on Zenodo (doi: 10.5281/zenodo.15686864).

References

1. A. Défossez, C. Caucheteux, J. Rapin, O. Kabeli, J. King, Decoding speech perception from non-invasive brain recordings. *Nat. Mach. Intell.* 5, 1097–1107 (2023).
2. S.L. Metzger, K.T. Littlejohn, A.B. Silva, D.A. Moses, M.P. Seaton, R. Wang, M.E. Dougherty, J.R. Liu, P. Wu, M. Berger, I. Zhuravleva, A. Tu-Chan, K. Ganguly, G.K. Anumanchipalli, E.F. Chang, A high-performance neuroprosthesis for speech decoding and avatar control. *Nature* 620, 1037–1046 (2023).
3. J. Tang, A. LeBel, S. Jain, A.G. Huth, Semantic reconstruction of continuous language from non-invasive brain recordings. *Nat. Neurosci*. 26, 858–866 (2023).
4. F.R. Willett, E.M. Kunz, C. Fan, D.T. Avansino, G.H. Wilson, E.Y. Choi, F. Kamdar, L.R. Hochberg, S. Druckmann, K. V. Shenoy, J. M. Henderson, A high-performance speech neuroprosthesis. *Nature* 620, 1031–1036 (2023).
5. T.B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D.M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, D. Amodei, Language Models are Few-Shot Learners. https://doi.org/10.48550/arXiv.2005.14165 (2020).
6. Gemini Team Google, Gemini: A Family of Highly Capable Multimodal Models. https://doi.org/10.48550/arXiv.2312.11805 (2023).
7. T. Horikawa, Y. Kamitani, Generic decoding of seen and imagined objects using hierarchical visual features. *Nat. Commun.* 8, 15037 (2017a).
8. A.G. Huth, T. Lee, S. Nishimoto, N.Y. Bilenko, A.T. Vu, J.L. Gallant, Decoding the Semantic Content of Natural Movies from Human Brain Activity. *Front. Syst. Neurosci.* 10, 81 (2016).
9. S. Nishida, S. Nishimoto, Decoding naturalistic experiences from human brain activity via distributed representations of words. *NeuroImage*, 180 232–242 (2018).
10. T. Horikawa, Y. Kamitani, Hierarchical Neural Representation of Dreamed Objects Revealed by Brain Decoding with Deep Neural Network Features. *Front. Comput. Neurosci.* 11, 4 (2017b).
11. C. Baldassano, D.M. Beck, L. Fei-Fei, Human-Object Interactions Are More than the Sum of Their Parts. *Cereb. Cortex* 27, 2276–2288 (2017).
12. E. McMahon, L. Isik, Seeing social interactions. *Trends Cogn. Sci.* 27, 1165–1179 (2023).
13. L. Tarhan, T. Konkle, Sociality and interaction envelope organize visual action representations. *Nat. Commun.* 11, 3002 (2020).
14. S. Karakose-Akbiyik, A. Caramazza, M.F. Wurm, A shared neural code for the physics of actions and object events. *Nat. Commun.* 14, 3316 (2023).
15. M.F. Wurm, A. Caramazza, A. Lingnau, Action categories in lateral occipitotemporal cortex are organized along sociality and transitivity. *J. Neurosci*. 37, 562–575 (2017).
16. M.F. Wurm, A. Caramazza, Two 'what' pathways for action and object recognition. *Trends Cogn. Sci.* 26, 103–116 (2022).
17. A. Doerig, T.C. Kietzmann, E.J. Allen, Y. Wu, T. Naselaris, K.N. Kay, I. Charest, Semantic scene descriptions as an objective of human vision. https://doi.org/10.48550/arXiv.2209.11737 (2022).
18. M. Ferrante, F. Ozcelik, T. Boccato, R.V. Rullen, N. Toschi, Brain Captioning: Decoding human brain activity into images and text. https://doi.org/10.48550/arXiv.2305.11560 (2023).
19. W. Huang, H. Yan, K. Cheng, Y. Wang, C. Wang, J. Li, C. Li, C. Li, Z. Zuo, H. Chen, A dual-channel language decoding from brain activity with progressive transfer training. *Hum. Brain map.* 42, 5089–5100 (2021a).

20. W. Huang, H. Yan, K. Cheng, C. Wang, J. Li, Y. Wang, C. Li, C. Li, Y. Li, Z. Zuo, H. Chen, A neural decoding algorithm that generates language from visual activity evoked by natural images. *Neural Netw.* 144, 90–100 (2021b).
21. Y. Kamitani, F. Tong, Decoding the visual and subjective contents of the human brain. *Nat. Neurosci.* 8, 679–685 (2005).
22. T. Naselaris, K.N. Kay, S. Nishimoto, J.L. Gallant, Encoding and decoding in fMRI. *NeuroImage*, 56, 400–410 (2011).
23. N. Kriegeskorte, R.A. Kievit, Representational geometry: integrating cognition, computation, and the brain. *Trends in Cogn. Sci.* 17, 401–412 (2013).
24. J.B. Ritchie, D.M. Kaplan, C. Klein, Decoding the Brain: Neural Representation and the Limits of Multivariate Pattern Analysis in Cognitive Neuroscience. *Br. J. Philos. Sci.* 70, 581–607 (2019).
25. O. Vinyals, A. Toshev, S. Bengio, D. Erhan, Show and tell: A neural image caption generator. *CVPR*, 3156–3164. (2014).
26. N. Dijkstra, S.E. Bosch, M.A. J. van Gerven, Shared Neural Mechanisms of Visual Perception and Imagery. *Trends Cogn. Sci.*, 23, 423–434 (2019).
27. T. Horikawa, M. Tamaki, Y. Miyawaki, Y. Kamitani, Neural decoding of visual imagery during sleep. *Science* 340, 639–642 (2013).
28. J. Pearson, T. Naselaris, E.A. Holmes, S.M. Kosslyn, Mental Imagery: Functional Mechanisms and Clinical Applications. *Trends Cogn. Sci.* 19, 590–602 (2015).
29. C. Caucheteux, J.-R. King, Brains and algorithms partially converge in natural language processing. *Commun. Biol.* 5, 134 (2022).
30. S. Jain, A.G. Huth, Incorporating context into language encoding models for fMRI. *NeurIPS* 31, 6629–6638 (2018).
31. A. LeBel, S. Jain, A.G. Huth, Voxelwise encoding models show that cerebellar language representations are highly conceptual. *J. Neurosci.* 41, 10341–10355 (2021).
32. M. Schrimpf, I. A. Blank, G. Tuckute, C. Kauf, E.A. Hosseini, N.G. Kanwisher, J.B. Tenenbaum, E. Fedorenko, The neural architecture of language: integrative modeling converges on predictive processing. *Proc. Natl. Acad. Sci. USA* 118, e2105646118 (2021).
33. M. Toneva, L. Wehbe, Interpreting and improving natural-language processing (in machines) with natural language-processing (in the brain). *NeurIPS*, 32 14928–14938 (2019).
34. M. Toneva, T.M. Mitchell, L. Wehbe, Combining computational controls with natural text reveals aspects of meaning composition. *Nat. Comput. Sci.* 2, 745–757 (2022).
35. J. Devlin, M.W. Chang, K. Lee, K. Toutanova, BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *In Proc. Conference of the North American Chapter of the Association for Computational Linguistics* 1, 4171–4186 (2019).
36. A.S. Cowen, D. Keltner, Self-report captures 27 distinct categories of emotion bridged by continuous gradients. *Proc. Natl. Acad. Sci. USA* 114, E7900–E7909 (2017).
37. P. He, X. Liu, J. Gao, W. Chen, DeBERTa: Decoding-enhanced BERT with Disentangled Attention. https://doi.org/10.48550/arXiv.2006.03654 (2020).
38. Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, V. Stoyanov, RoBERTa: A Robustly Optimized BERT Pretraining Approach. https://doi.org/10.48550/arXiv.1907.11692 (2019).
39. E.J. Allen, G. St-Yves, Y. Wu, J. L. Breedlove, J. S. Prince, L.T. Dowdle, M. Nau, B. Caron, F. Pestilli, I. Charest, J. B. Hutchinson, T. Naselaris, K. Kay, A massive 7T fMRI dataset to bridge cognitive neuroscience and artificial intelligence. *Nat. Neurosci.*, 25, 116–126 (2022).

40. M. Schrimpf, J. Kubilius, H. Hong, N.J. Majaj, R. Rajalingham, E.B. Issa, K. Kar, P. Bashivan, J. Prescott-Roy, K. Schmidt, D. Yamins, J.J. DiCarlo, Brain-Score:Which Artificial Neural Network for Object Recognition is most Brain-Like? https://www.biorxiv.org/content/10.1101/407007v1 (2018).
41. S. Zhang, S. Roller, N. Goyal, M. Artetxe, M. Chen, S. Chen, C. Dewan, M.T. Diab, X. Li, X.V. Lin, T. Mihaylov, M. Ott, S. Shleifer, K. Shuster, D. Simig, P.S. Koura, A. Sridhar, T. Wang, L. Zettlemoyer, OPT: Open Pre-trained Transformer Language Models. https://doi.org/10.48550/arXiv.2205.01068 (2022).
42. H. Touvron, T. Lavril, G. Izacard, X. Martinet, M. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, A. Rodriguez, A. Joulin, E. Grave, G. Lample, LLaMA: Open and Efficient Foundation Language Models. https://doi.org/10.48550/arXiv.2302.13971 (2023).
43. T. Thrush, R. Jiang, M. Bartolo, A. Singh, A. Williams, D. Kiela, C. Ross, Winoground: Probing Vision and Language Models for Visio-Linguistic Compositionality. *CVPR*, 5228–5238 (2022).
44. J. Salazar, D. Liang, T.W. Nguyen, K. Kirchhoff, Masked Language Model Scoring. *Annual Meeting of the Association for Computational Linguistics* (2019).
45. C. Kauf, A.A. Ivanova, A Better Way to Do Masked Language Model Scoring. https://doi.org/10.48550/arXiv.2305.10588 (2023).
46. J.R. Binder, R.H. Desai, W.W. Graves, L.L. Conant, Where is the semantic system? A critical review and meta-analysis of 120 functional neuroimaging studies. *Cerebral cortex* 19, 2767–2796 (2009).
47. B.J. Devereux, A. Clarke, A. Marouchos, L.K. Tyler, Representational similarity analysis reveals commonalities and differences in the semantic processing of words and objects. *J. Neurosci.* 33, 18906–18916 (2013).
48. A.G. Huth, S. Nishimoto, A.T. Vu, J.L. Gallant, A continuous semantic space describes the representation of thousands of object and action categories across the human brain. *Neuron* 76, 1210–1224 (2012).
49. A.G. Huth, W.A. de Heer, T.L. Griffiths, F.E. Theunissen, J.L. Gallant, Natural speech reveals the semantic maps that tile human cerebral cortex. *Nature* 532, 453–458 (2016).
50. C. Pallier, A.D. Devauchelle, S. Dehaene, Cortical representation of the constituent structure of sentences. *Proc. Natl. Acad. Sci. USA* 108, 2522–2527 (2011).
51. F. Mollica, M. Siegelman, E. Diachek, S.T. Piantadosi, Z. Mineroff, R. Futrell, H. Kean, P. Qian, E. Fedorenko, Composition is the Core Driver of the Language-selective Network. *Neurobiol. Lang. (Camb)* 1, 104–134 (2020).
52. M. Toneva, T.M. Mitchell, L. Wehbe, Combining computational controls with natural text reveals aspects of meaning composition. *Nat. Comput. Sci.* 2, 745–757 (2022).
53. A.A. Ivanova, Z. Mineroff, V.C. Zimmerer, N.G. Kanwisher, R.A. Varley, E. Fedorenko, The Language Network Is Recruited but Not Required for Nonverbal Event Semantics. *Neurobiol. Lang. (Camb)* 2, 176–201 (2021).
54. Y. Sueoka, A.M. Paunov, A.A. Ivanova, I.A. Blank, E. Fedorenko, The language network reliably 'tracks' naturalistic meaningful non-verbal stimuli. *Neurobiol. Lang.* (2024).
55. G. Bertasius, H. Wang, L. Torresani, Is Space-Time Attention All You Need for Video Understanding? https://doi.org/10.48550/arXiv.2102.05095 (2021).
56. E. Fedorenko, R. Varley, Language and thought are not the same thing: evidence from neuroimaging and neurological patients. *Ann. N. Y. Acad. Sci.* 1369, 132–153 (2016).
57. A. Radford, J.W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, I. Sutskever, Learning Transferable Visual Models From Natural Language Supervision. *ICML* (2021).
58. J. Wang, Z. Yang, X. Hu, L. Li, K. Lin, Z. Gan, Z. Liu, C. Liu, L. Wang, GIT: A

Generative Image-to-text Transformer for Vision and Language. https://doi.org/10.48550/arXiv.2205.14100 (2022).
59. Z. He, T. Sun, K. Wang, X. Huang, X. Qiu, DiffusionBERT: Improving Generative Masked Language Models with Diffusion Models. *Annual Meeting of the Association for Computational Linguistics*. (2022).
60. S.S. Sahoo, M. Arriola, Y. Schiff, A. Gokaslan, E. Marroquin, J.T. Chiu, A. Rush, V. Kuleshov, Simple and Effective Masked Diffusion Language Models. https://doi.org/10.48550/arXiv.2406.07524 (2024).
61. S. Nie, F. Zhu, Z. You, X. Zhang, J. Ou, J. Hu, J. Zhou, Y. Lin, J. Wen, C. Li, Large Language Diffusion Models. https://doi.org/10.48550/arXiv.2502.09992 (2025).
62. R. Fukuma, T. Yanagisawa, S. Nishimoto, H. Sugano, K. Tamura, S. Yamamoto, Y. Iimura, Y. Fujita, S. Oshino, N. Tani, N. Koide-Majima, Y. Kamitani, H. Kishima, Voluntary control of semantic neural representations by imagery with conflicting visual stimulation. *Commun. Biol.* 5, 214 (2022).
63. S.F. Popham, A.G. Huth, N.Y. Bilenko, F. Deniz, J.S. Gao, A.O. Nunez-Elizalde, J.L. Gallant, Visual and linguistic semantic representations are aligned at the border of human visual cortex. *Nat. Neurosci.* 24, 1628–1636 (2021).
64. A. Steel, M.M. Billings, E.H. Silson, C.E. Robertson, A network linking scene perception and spatial memory systems in posterior cerebral cortex. *Nat. Commun.* 12 (2021).
65. P. Achlioptas, M. Ovsjanikov, K. Haydarov, M. Elhoseiny, L.J. Guibas, ArtEmis: Affective Language for Visual Art. *CVPR* 11564–11574 (2021).
66. T. Horikawa, A.S. Cowen, D. Keltner, Y. Kamitani, The Neural Representation of Visually Evoked Emotion Is High-Dimensional, Categorical, and Distributed across Transmodal Brain Regions. *iScience*, 23 (2019).
67. J. Wang, V. Cherkassky, Y. Yang, K.K. Chang, R. Vargas, N. Diana, M.A. Just, Identifying thematic roles from neural representations measured by functional magnetic resonance imaging. *Cogn. Neuropsychol.* 33, 257–264 (2016).
68. R. Navigli, S. Conia, B. Ross, Biases in Large Language Models: Origins, Inventory, and Discussion. *ACM Journal of Data and Information Quality* 15, 1–21 (2023).
69. P.S. Scotti, M. Tripathy, C.K. Villanueva, R. Kneeland, T. Chen, A. Narang, C. Santhirasegaran, J. Xu, T. Naselaris, K.A. Norman, T. M. Abraham, MindEye2: Shared-Subject Models Enable fMRI-To-Image With 1 Hour of Data. https://doi.org/10.48550/arXiv.2403.11207 (2024).
70. H. Wang, J.K. Ho, F. Cheng, S.C. Aoki, Y. Muraki, M. Tanaka, Y. Kamitani, Inter-individual and inter-site neural code conversion and image reconstruction without shared stimuli. https://doi.org/10.48550/arXiv.2403.11517 (2024).
71. S. Goering, E. Klein, L. Specker Sullivan, A. Wexler, B. Agüera y Arcas, G. Bi, J.M. Carmena, J.J. Fins, P. Friesen, J.L. Gallant, J.E. Huggins, P. Kellmeyer, A.H. Marblestone, C. Mitchell, E. Parens, M. Pham, A. Rubel, N. Sadato, M. Teicher, D.T. Wasserman, M. Whittaker, J.R. Wolpaw, R. Yuste, Recommendations for Responsible Development and Application of Neurotechnologies. *Neuroethics* 14, 365–386 (2021).
72. N.C. Benson, K.W. Jamison, M. Arcaro, A.T. Vu, M.F. Glasser, T.S. Coalson, D.C. Van Essen, E. Yacoub, K. Uğurbil, J.A. Winawer, K.N. Kay, The Human Connectome Project 7 Tesla retinotopy dataset: Description and population receptive field analysis. *J. Vis*. 18, (2018).
73. B. Fischl, FreeSurfer. *Neuroimage* 62, 774–781 (2012).
74. K.N. Kay, J. Winawer, A. Mezer, B. A. Wandell, Compressive spatial summation in human visual cortex. *J. neurophysiol.* 110, 481–494 (2013).
75. A. Stigliani, K.S. Weiner, K. Grill-Spector, Temporal processing capacity in high-level visual cortex is domain specific. *J. Neurosci.* 35, 12412–12424 (2015).

76. R.B. Tootell, J.B. Reppas, K.K. Kwong, R. Malach, R.T. Born, T.J. Brady, B.R. Rosen, J.W. Belliveaul, R.R. Benson, J.R. Baker, A. Jiang, C.E. Stern, H.C. Breiter, M. Cohen, R. O 'sullivan, T.C. Kennedy, L.A. Kobierski, A. Guimares, C.K. We, R.A. Andersen, B. Geeselman, R. Siegal, Functional analysis of human MT and related visual cortical areas using magnetic resonance imaging. *J. Neurosci.* 15, 3215–3230 (1995).
77. E. Fedorenko, P.J. Hsieh, A. Nieto-Castañón, S. Whitfield-Gabrieli, N. Kanwisher, New method for fMRI investigations of language: defining ROIs functionally in individual subjects. *J. Neurophysiol*. 104, 1177–1194 (2010).
78. T.L. Scott, J. Gallée, E. Fedorenko, A new fun and robust version of an fMRI localizer for the frontotemporal language system. *Cogn. Neurosci.* 8, 167–176 (2017).
79. J.S. Gao, A.G. Huth, M.D. Lescroart, J.L. Gallant, Pycortex: an interactive surface visualizer for fMRI. *Front. Neuroinform.* 9, 23. (2015).
80. O. Esteban, C.J. Markiewicz, R.W. Blair, C.A. Moodie, A.I. Isik, A. Erramuzpe, J.D. Kent, M. Goncalves, E. Dupre, M. Snyder, H. Oya, S.S. Ghosh, J. Wright, J. Durnez, R.A. Poldrack, K.J. Gorgolewski, fMRIPrep: a robust preprocessing pipeline for functional MRI. *Nat. Methods* 16, 111–116 (2019).
81. D.N. Greve, B. Fischl, Accurate and robust brain image alignment using boundary-based registration. *NeuroImage* 48, 63–72 (2009).
82. R. Cox, J. Hyde, Software tools for analysis and visualization of fMRI data. *NMR in Biomedicine* 10 (1997).
83. T. Lin, M. Maire, S.J. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, C. L. Zitnick, Microsoft COCO: Common Objects in Context. *European Conference on Computer Vision* (2014).
84. X. Chen, H. Fang, T. Lin, R. Vedantam, S. Gupta, P. Dollár, C.L. Zitnick, Microsoft COCO Captions: Data Collection and Evaluation Server. https://doi.org/10.48550/arXiv.1504.00325 (2015).
85. J. Xu, T. Mei, T. Yao, Y. Rui, MSR-VTT: A Large Video Description Dataset for Bridging Video and Language. *CVPR*, 5288–5296 (2016).
86. T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. V. Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T.L. Scao, S. Gugger, M. Drame, Q. Lhoest, A. M. Rush, Transformers: State-of-the-Art Natural Language Processing. *Conference on Empirical Methods in Natural Language Processing.* (2019).
87. N. Reimers, I. Gurevych, Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. *Conference on Empirical Methods in Natural Language Processing.* (2019).
88. Z. Lan, M. Chen, S. Goodman, K. Gimpel, P. Sharma, R. Soricut, ALBERT: A Lite BERT for Self-supervised Learning of Language Representations. https://doi.org/10.48550/arXiv.1909.11942 (2019).
89. A. Radford, K. Narasimhan, T. Salimans, I. Sutskever, "Improving language understanding with unsupervised learning." (OpenAI, 2018); https://openai.com/index/language-unsupervised/.
90. A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever, "Language Models are Unsupervised Multitask Learners." (OpenAI, 2019); https://d4mucfpksywv.cloudfront.net/better-language-models/language-models.pdf.
91. N. Muennighoff, SGPT: GPT Sentence Embeddings for Semantic Search. https://doi.org/10.48550/arXiv.2202.08904 (2022).
92. Z. Yang, Z. Dai, Y. Yang, J.G. Carbonell, R. Salakhutdinov, Q.V. Le, XLNet: Generalized Autoregressive Pretraining for Language Understanding. https://doi.org/10.48550/arXiv.1906.08237 (2019).
93. V. Sanh, L. Debut, J. Chaumond, T. Wolf, DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter.

https://doi.org/10.48550/arXiv.1910.01108 (2019).
94. C. Raffel, N.M. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, P. J. Liu, Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *J. Mach. Learn. Res.,* 21, 1–140 (2019).
95. M. Lewis, Y. Liu, N. Goyal, M. Ghazvininejad, A. Mohamed, O. Levy, V. Stoyanov, L. Zettlemoyer, BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. *Annual Meeting of the Association for Computational Linguistics*. (2019).
96. N.S. Keskar, B. McCann, L.R. Varshney, C. Xiong, R. Socher, CTRL: A Conditional Transformer Language Model for Controllable Generation. https://doi.org/10.48550/arXiv.1909.05858 (2019).
97. G. Lample, A. Conneau, Cross-lingual Language Model Pretraining. https://doi.org/10.48550/arXiv.1901.07291 (2019).
98. A. Conneau, K. Khandelwal, N. Goyal, V. Chaudhary, G. Wenzek, F. Guzmán, E. Grave, M. Ott, L. Zettlemoyer, V. Stoyanov, Unsupervised Cross-lingual Representation Learning at Scale. *Annual Meeting of the Association for Computational Linguistics.* (2019).
99. K. Clark, M. Luong, Q.V. Le, C.D. Manning, ELECTRA: Pre-training Text Encoders as Discriminators Rather Than Generators. *ICLR* (2020).
100. J. Deng, W. Dong, R. Socher, L. Li, K. Li, L. Fei-Fei, ImageNet: A large-scale hierarchical image database. *CVPR.* 248–255 (2009).
101. J. Carreira, A. Zisserman, Quo Vadis, action recognition? A new model and the kinetics dataset. *CVPR*, 4724–4733 (2017).
102. S. Nishimoto, Modeling movie-evoked human brain activity using motion-energy and space-time vision transformer features. https://www.biorxiv.org/content/10.1101/2021.08.22.457251v1 (2021).
103. P.J. Cornbleet, N. Gochman, Incorrect Least-Squares Regression Coefficients in Method-Comparison Analysis. *Clin. Chem.* 25, 432–438 (1979).
104. P. Sharma, N. Ding, S. Goodman, R. Soricut, Conceptual Captions: A Cleaned, Hypernymed, Image Alt-text Dataset For Automatic Image Captioning. *Annual Meeting of the Association for Computational Linguistics.* (2018).
105. S. Bird, E. Loper, NLTK: The Natural Language Toolkit. *In Proc. ACL Interactive Poster and Demonstration Sessions (ed. Scott, D.)* 214–217 (2004).
106. C. Lin, F. J. Och, Automatic Evaluation of Machine Translation Quality Using Longest Common Subsequence and Skip-Bigram Statistics. *Annual Meeting of the Association for Computational Linguistics* (2004).
107. M. Honnibal, I. Montani, spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. (Sentometrics Research, 2017); https://sentometrics-research.com/publication/72/.
108. Y. Zhu, S. Lu, L. Zheng, J. Guo, W. Zhang, J. Wang, Y. Yu, Texygen: A Benchmarking Platform for Text Generation Models. *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval.* (2018).
109. R.A.A. Ince, J.W. Kay, P.G. Schyns, Within-participant statistics for cognitive science. *Trends Cogn. Sci.* 26, 626–630 (2022).
110. Y. Benjamini, Y. Hochberg, Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. R. Statist. Soc. ser. B* 57, 298–300 (1995).

Acknowledgments: The author thanks Yuki Honda for assistance with MRI data collection and cleaning caption data, Scinob Kuroki, Shimpei Yamagishi, Hiromi Narimatsu, and Yuta Suzuki for assistance with scanner operation, Mitsuaki Tsukamoto, Misato Tanaka, and Yukiyasu Kamitani for assistance with preliminary investigations, and Shogo Kajimura for discussions. We

appreciate the support of the WPI-IRCN Human fMRI Core, the University of Tokyo Institutes for Advanced Studies.

Funding: This research was supported by grants from JST PRESTO Grant Number JPMJPR185B Japan and JSPS KAKENHI Grant Number JP21H03536.

Author contributions:

Conceptualization: TH

Methodology: TH

Software: TH

Validation: TH

Formal Analysis: TH

Investigation: TH

Resources: TH

Data curation: TH

Writing – original draft: TH

Writing – review & editing: TH

Visualization: TH

Project administration: TH

Funding acquisition: TH

Supervision: TH

Competing interests: The author declares that he has no competing interests.

Data and materials availability: All data needed to evaluate the conclusions in the paper are present in the paper and/or the Supplementary Materials.

Supplementary Materials

Figs. S1 to S14

References (*72–110*)

Supplementary Materials for

Mind Captioning: Evolving descriptive text of mental content from human brain activity

Tomoyasu Horikawa

*Corresponding author. Email: horikawa.t@gmail.com

This PDF file includes:

Figs. S1 to S14
References (*72–110*)

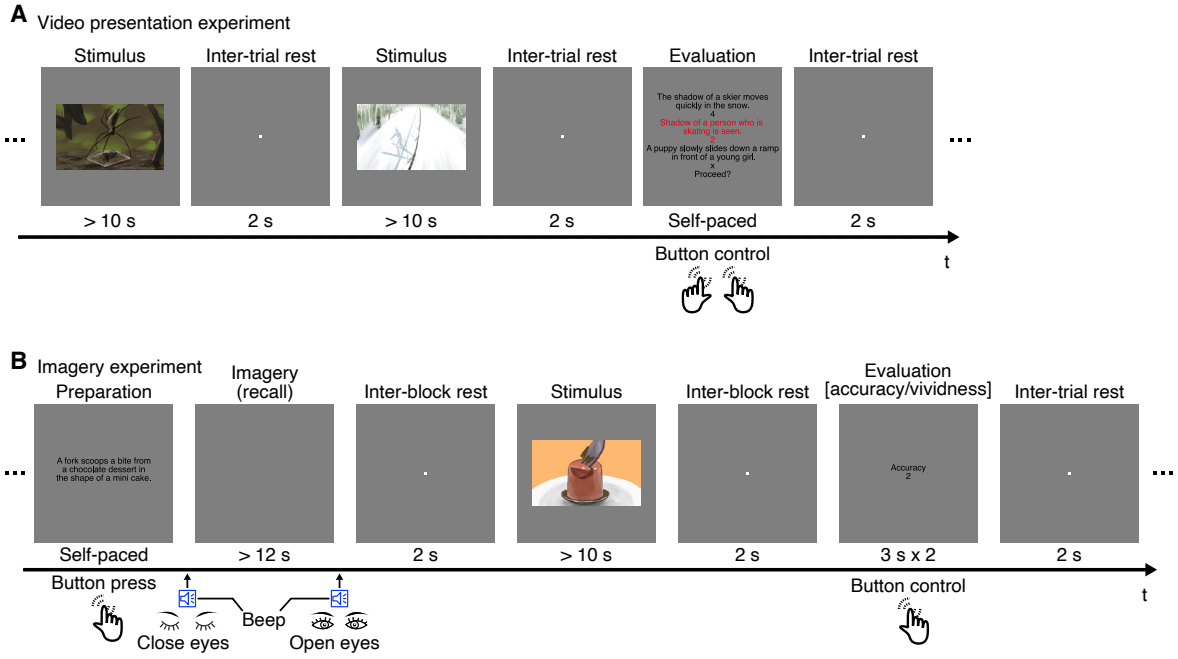

Fig. S1. Overview of experiments.
We conducted two types of experiments. (**A**) Video presentation experiment (training and test sessions). Subjects were allowed to view videos without fixation. Occasionally, we presented five descriptions depicting the visual content of the preceding video to the subjects, asking them to rate the consistency between each description and their subjective perception. We used these ratings to evaluate the consistency between descriptions generated from the brain and the subjective perception of each individual (fig. S5E). (**B**) Imagery experiment. Subjects were well-trained to associate descriptions and videos before the experiment. They were required to visually imagine (recall) a video based on a description of the target video with their eyes closed according to beeps. Each imagery block was followed by a stimulus block, during which the target video was presented to allow the subjects to confirm the validity of their imagery. An evaluation block followed, asking the subjects to evaluate the accuracy and vividness of their imagery. Data samples were constructed by averaging fMRI volumes during each stimulus/imagery block. Training data contained samples for 2,108 videos, each presented once. Test and imagery data contained samples for the same 72 videos (not used during training), averaged across five repetitions.

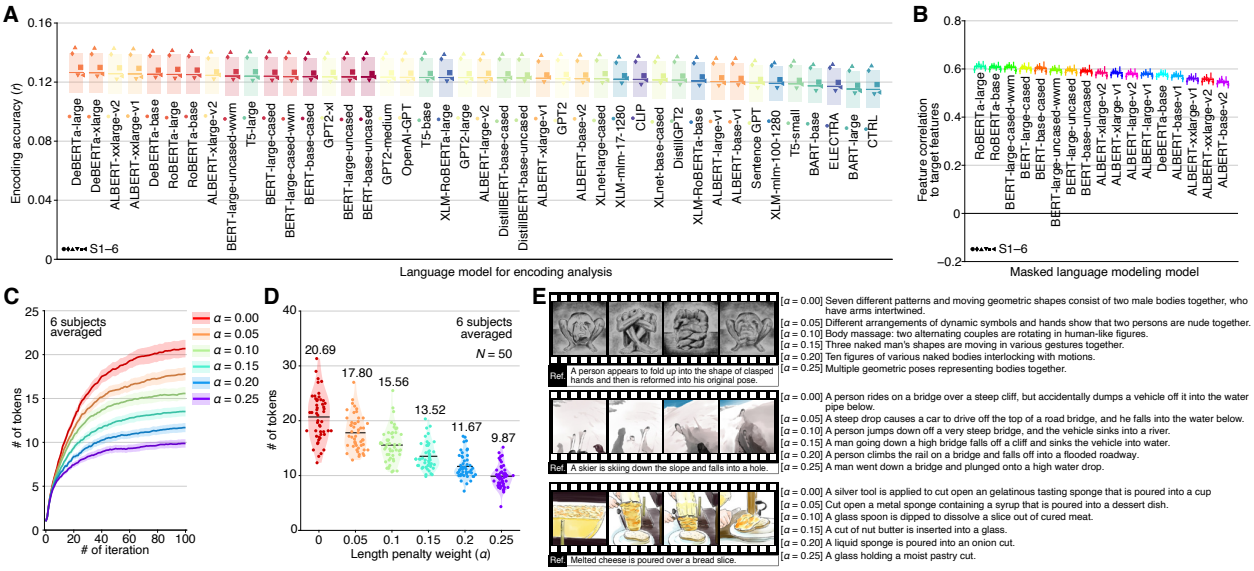

Fig. S2. Validation of model and parameter determinations.

Validation analyses were performed using cross-validation within the training perception data to determine LMs used in feature decoding and text generation analyses and a length penalty parameter. (**A**) Encoding accuracy from different LMs (averaged across voxels). We hypothesized that an LM with better performance in predicting brain responses to visual stimuli would provide more effective intermediate representations for translating visual semantic information in the brain into text. Accordingly, we performed an encoding analysis using features from each of the 42 LMs and found that the DeBERTa-large model consistently exhibited the best performance. Therefore, we used it in the main analysis. (**B**) Feature correlations between target decoded features and features of descriptions generated with different MLM models. Different MLM models were used in the text generation analysis (the DeBERTa-large model was used for feature decoding and evaluation). All models exhibited reliable performance in generating descriptions aligned with target decoded features. We decided to use the RoBERTa-large model, as it exhibited the best performance. (**C**) The number of tokens in generated descriptions through optimization. (**D** and **E**) The number of tokens in generated descriptions (D) and the example descriptions (E) obtained with varying strength of length penalty. We evaluated the similarity between target brain-decoded features and features of candidate descriptions based on Pearson correlation coefficients with an exponential penalty to the length of candidate descriptions (see Materials and Methods for details). To determine the length penalty parameter, we performed a validation analysis with a randomly selected subset of 50 samples from the training perception data using six possible penalty parameters (0, 0.05, 0.1, 0.15, 0.2, and 0.25). While the length of generated descriptions varied depending on the length penalty, accurate descriptions of viewed content were consistently generated. We decided to use $\boldsymbol{\alpha} = 0.1$, as the mean length of generated descriptions with this parameter (mean = 15.56) was comparable with that of reference captions in the training perception data (mean = 15.94). In (B−E), the text generation analysis was performed for 50 videos randomly selected from the training perception data (DeBERTa-large; decoded from whole brain activity). Shades in (A, B) indicate 95% C.I. across subjects ($N = 6$). Error bars in (A, B) and shades in (C) indicate 95% C.I. across samples ($N = 50$).

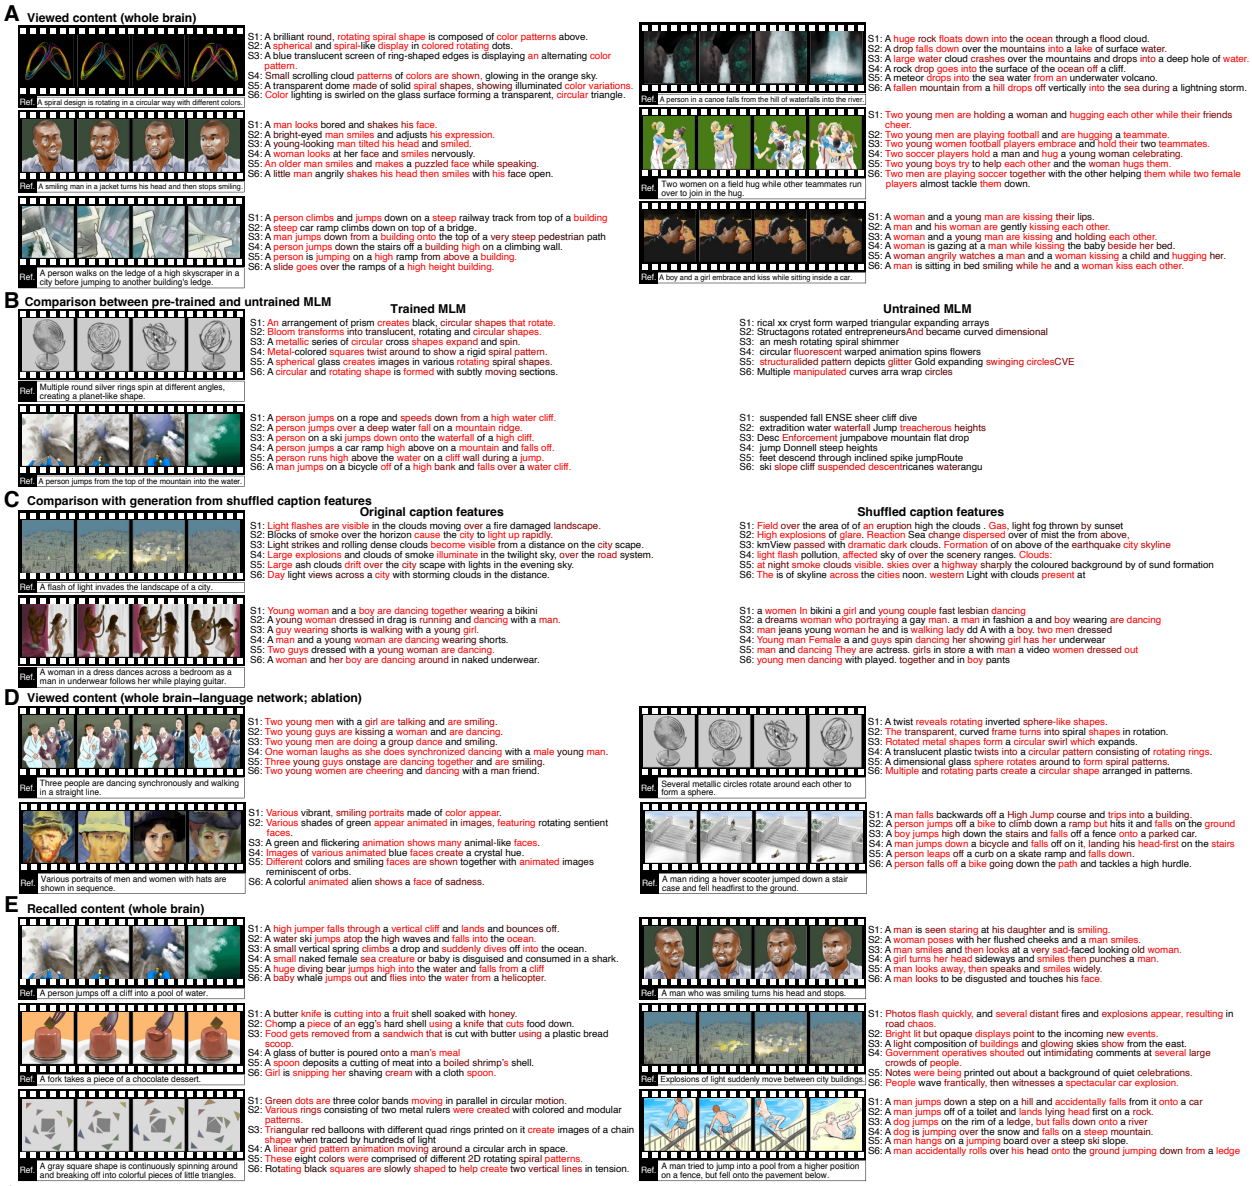

Fig. S3. Examples of generated descriptions. (A) Descriptions of viewed content generated from whole-brain activity. (B) Descriptions of viewed content generated with the support of either the pre-trained or untrained MLM model. (C) Descriptions of viewed content generated with features predicted by decoders trained based on original captions or those based on captions whose word order was randomly shuffled within each caption. Notably, the descriptions generated for the shuffled condition contained words representing individual components in viewed videos, indicating that semantic features of shuffled captions still convey the word-level semantic information. However, these descriptions lacked the coherence to accurately describe the relations among individual components. (D) Descriptions of viewed content generated from the activity of the whole brain except for the language network. (E) Descriptions of recalled content generated from whole-brain activity. Conventions are the same as for Fig. 2A.

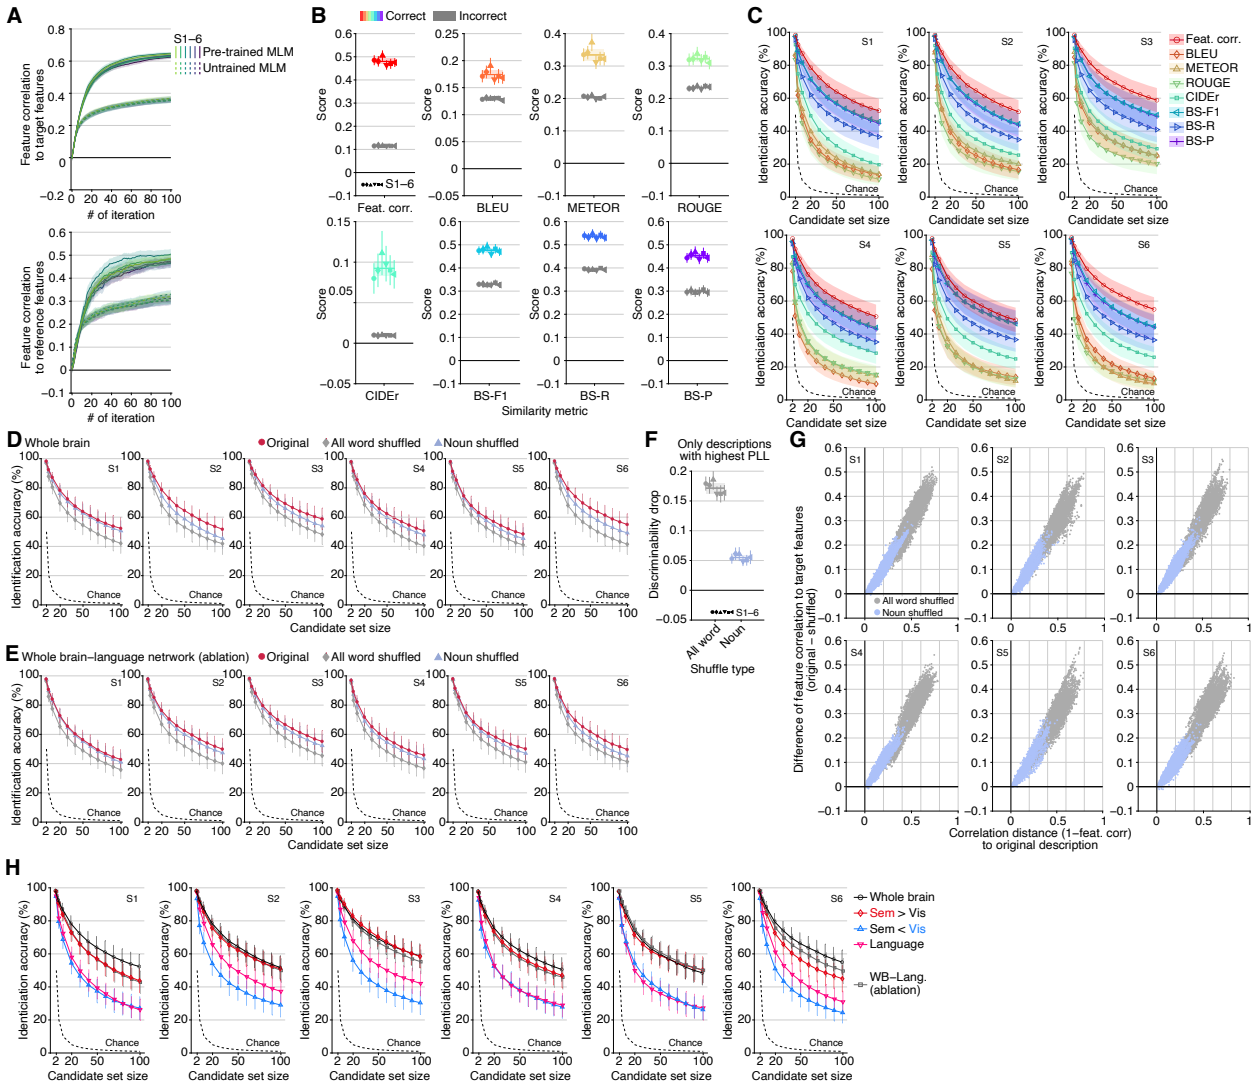

Fig. S4. Text generation performance of viewed content for individual subjects. (**A**) Feature correlations between semantic features of generated descriptions and those decoded from the brain, as well as those computed from correct references. The generated descriptions supported by the pre-trained MLM model outperformed those supported by the untrained MLM model for all subjects in both similarity measures (Wilcoxon signed-rank test, one-tailed, $P < 0.01$, FDR corrected across subjects; 100 iterations). (**B**) Raw scores of the similarity between generated descriptions and captions for correct and incorrect references. (**C**) Video identification accuracy. (**D** and **E**) Effects of word-order shuffling on video identification accuracy applied to descriptions generated from whole-brain activity (D) and the activity of the whole brain except the language network (E). (**F**) Effects of word-order shuffling on discriminability using minimally disrupted shuffled descriptions. For each generated description, reductions in discriminability were evaluated by comparing it to the shuffled variants that had the highest pseudo-log-likelihood (PLL) scores, as determined by MLM scoring (*44, 45*). These descriptions were selected from a pool of up to 1,000 shuffled variants created under the all-word or noun-only shuffling conditions. (**G**) Scatterplot of the correlation distances (one minus feature correlation) between the original and shuffled descriptions against the difference in feature correlations to target features between the original and shuffled descriptions. (**H**) Video identification accuracy obtained from different brain areas. Shades in (A, C), and error bars in (B, D, F, H) indicate 95% C.I. across samples ($N = 72$). Shades in (B, F) indicate 95% C.I. across subjects ($N = 6$).

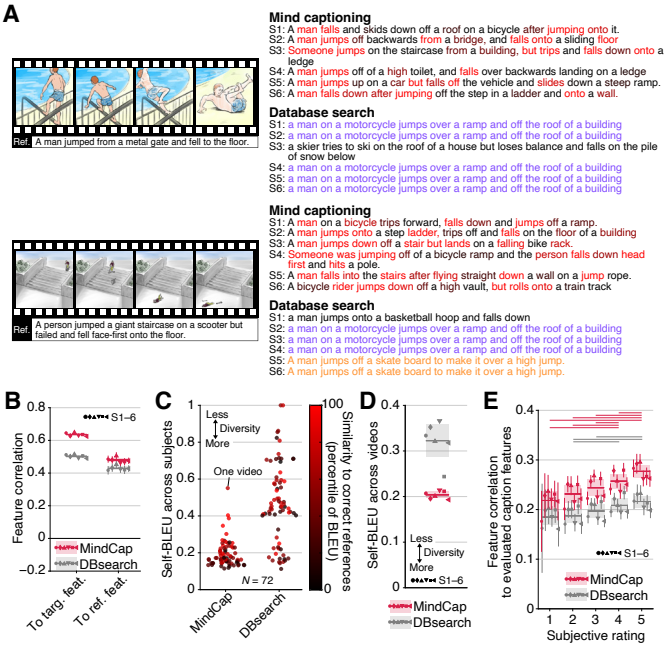

Fig. S5. Comparison with database-search method.
Text generation performance for viewed content was compared between our method and the database (DB) search method (*17*). (**A**) The DB search produced moderately accurate descriptions, but it often produced identical ones for different individuals and videos. (**B**) Our method outperformed DB-search, showing higher feature correlations with both target and reference features (Wilcoxon signed-rank test, one-tailed, $P < 0.01$, FDR corrected across subjects). (**C** and **D**) Diversity (self-BLEU) across subjects (C) and videos (D). Our method exhibited greater diversity than the DB search both across subjects and videos (Wilcoxon signed-rank test, one-tailed, $P < 0.05$). (**E**) Feature correlations at each subjective rating. Both methods showed higher scores for captions with higher ratings (lines above, $P < 0.01$, FDR corrected across pairs, six subjects pooled). Our method showed a higher correlation between ratings and feature correlations than the DB search (mean $r = 0.184$ and 0.134 for mind captioning and DB search, respectively; six subjects averaged; $P < 0.01$, six subjects pooled), demonstrating higher flexibility in generating descriptions better aligned with subjective perceptions (ANOVA, interaction between methods and ratings, $P < 0.01$; six subjects pooled). Shades and error bars indicate 95% C.I. across subjects and samples ($N = 6$ and 72, respectively).

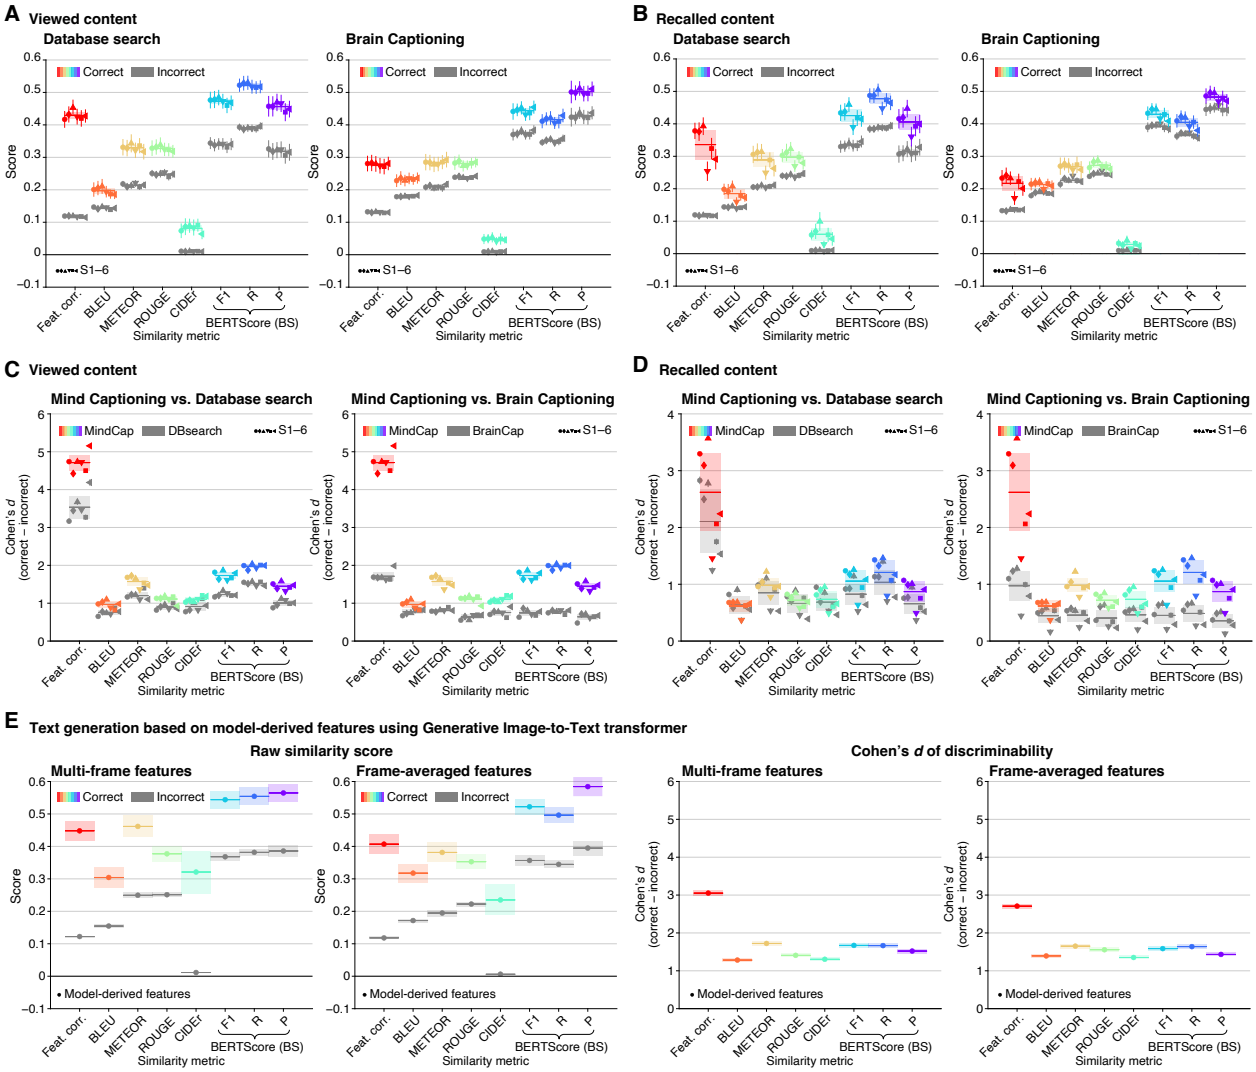

Fig. S6. Performance comparison of brain-to-text decoding methods on video data. We compared the performance of three brain-to-text decoding methods—mind captioning, the database (DB) search (*17*), and brain captioning (*18*)—using our video fMRI data. Brain captioning is a decoding approach that employs a non-linear image captioning model, the *Generative Image-to-text Transformer* (GIT) (*58*) (see Materials and Methods). (**A** and **B**) Raw similarity scores between generated descriptions and captions of correct and incorrect references for viewed (A) and recalled (B) content. (**C** and **D**) Comparisons of Cohen's *d* of discriminability between methods for viewed (C) and recalled (D) content. Overall, mind captioning achieved the highest discriminability in both viewed and recalled conditions, suggesting that it generates descriptions that are more distinguishable from those of irrelevant videos. While brain captioning and DB-search achieved higher raw scores on certain precision-oriented metrics (e.g., BLEU; see also Fig. S4B), this may be due to their tendency to produce descriptions with less lexical and syntactic variability, which surface-level similarity metrics may favor. In contrast, mind captioning generated more semantically discriminative descriptions, even with lower surface similarity. (**E**) Raw similarity scores and Cohen's *d* of discriminability of descriptions generated by GIT using model-derived (not brain-decoded) features. While GIT supports both image and video captioning—processing either single images or multiple frames—brain captioning is designed for single-image inputs. Therefore, we compared two input configurations for GIT: (1) multi-frame features from uniformly sampled frames and (2) frame-averaged features, obtained by averaging features across all video frames. Although the latter is suboptimal for GIT, its performance was comparable, supporting the validity of using frame-averaged features in brain captioning. For (A) and (B), and (C) and (D), conventions follow those in Fig. 2D and fig. S4B, respectively. Shades in (E) indicate 95% C.I. across samples ($N = 72$).

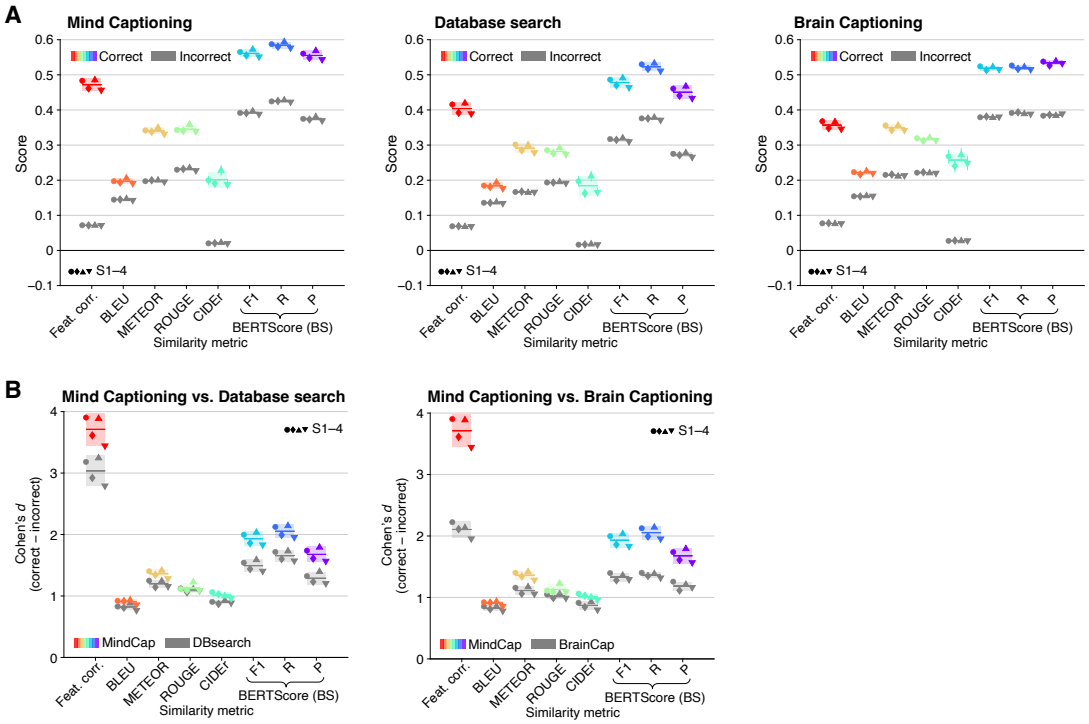

Fig. S7. Performance comparison of brain-to-text decoding methods on NSD image data.
Both the database (DB) search (*17*) and brain captioning (*18*) were originally introduced for decoding image-based brain activity; however, in fig. S6, we evaluated them on video-induced data. To complement that comparison, we further evaluated the performance of all three methods, including mind captioning, using image-induced fMRI data from four subjects in the *Natural Scenes Dataset* (NSD) (*39*) (see Materials and Methods). (**A**) Raw similarity scores between generated descriptions and captions of correct and incorrect references for viewed content. (**B**) Cohen's *d* of discriminability between methods for viewed content. Mind captioning achieved the highest discriminability among these methods, demonstrating strong performance even for captioning image-induced brain activity. These results emphasize the superiority of our method and support its robustness across stimulus modalities. For (A) and (B), conventions follow those in Fig. 2D and fig. S4B, respectively.

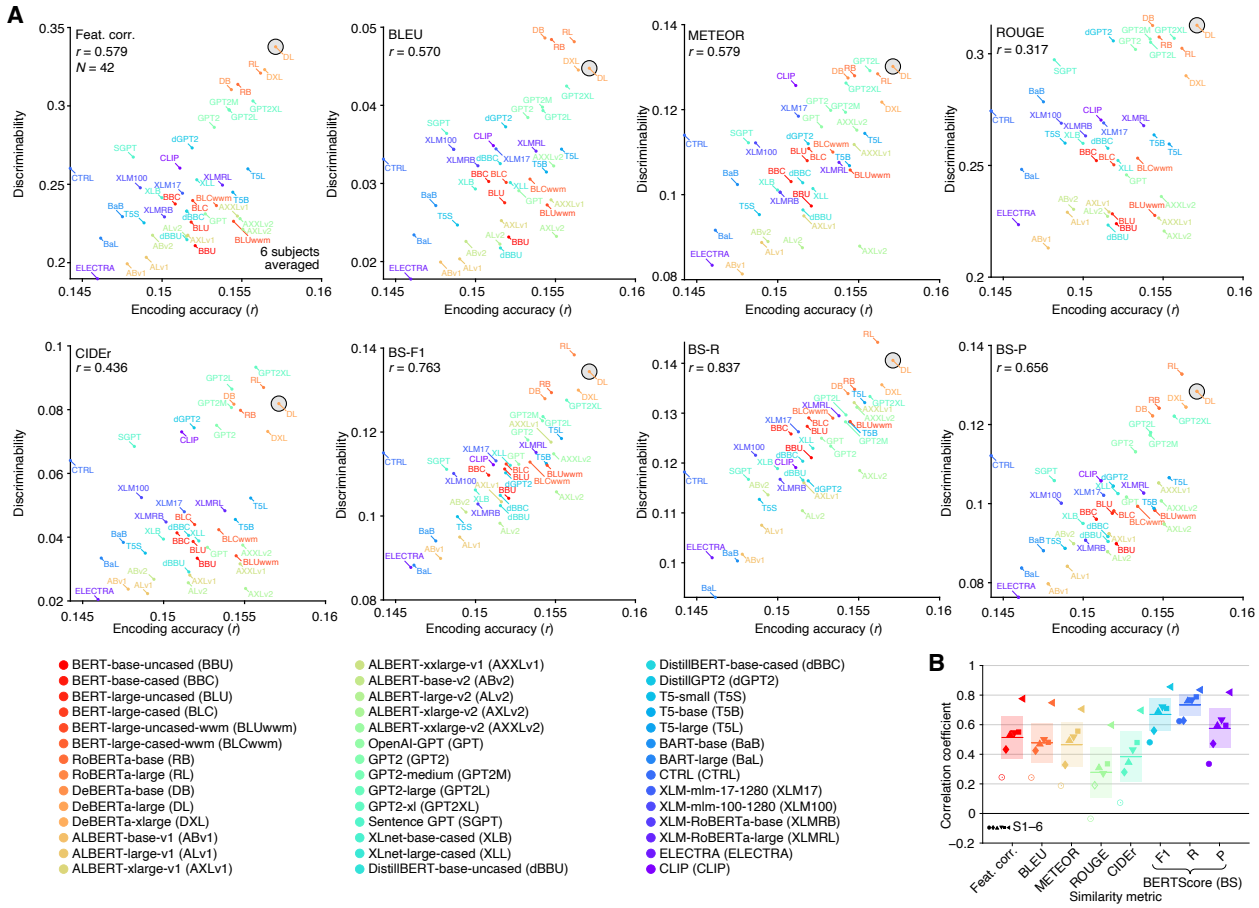

Fig. S8. Relations between encoding accuracy and text generation performance of multiple LMs.
We performed decoding and text generation analyses using features from 42 LMs. We consistently used the DeBERTa-large model for feature correlation evaluation and the RoBERTa-large model to support text generation, respectively. Encoding analysis was performed by training models on the training perception data and testing them on the test perception data. (**A**) Scatterplot of encoding accuracy against discriminability (N = 42; dark circle, the model used in the main analysis). For each model, discriminability scores were averaged across samples, and encoding accuracy was averaged across voxels in the whole brain. The discriminability scores were positive for all models, indicating the robustness of our method to the choice of LMs. (**B**) Correlation coefficients between discriminability and encoding accuracy for all models. We found significantly positive correlations for most metrics and subjects, except S1 (filled markers, t-test, one-tailed, $P < 0.01$, FDR corrected across metrics and subjects), suggesting that LMs with features more aligned with brain representations could help generate descriptions that accurately capture the viewed video content. Shades and error bars indicate 95% C.I. across samples (N = 72) and subjects (N = 6), respectively.

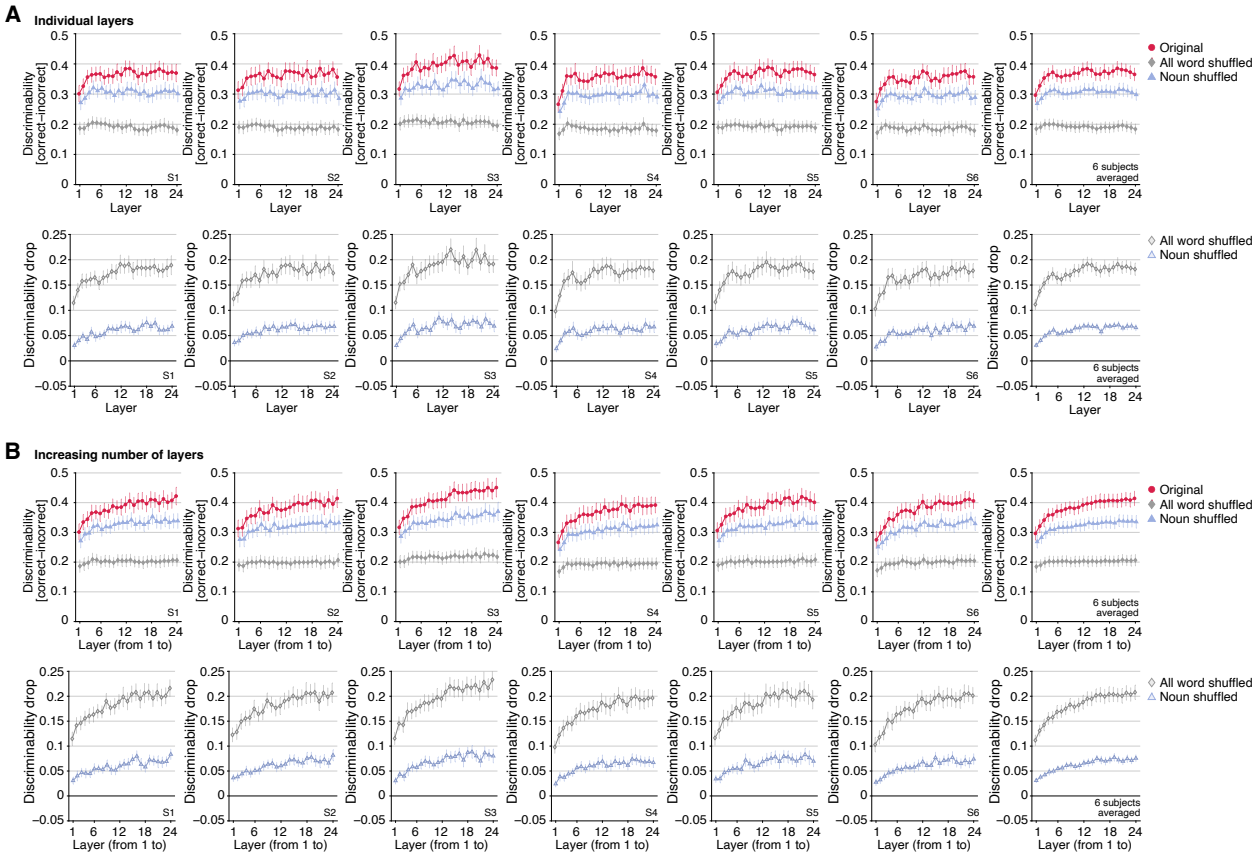

Fig. S9. Effects of shuffling on discriminability with varying numbers of layers. The text generation analysis was conducted using decoded features with varying numbers of LM layers (24 layers from DeBERTa-large; individual layers or increasing number of layers). We evaluated the impact of shuffling on the discriminability of the descriptions generated with the varying number of layers. To compare the results from different layer conditions, we consistently used the feature correlations computed using all layers of the DeBERTa-large model as a fixed metric. (**A** and **B**) Discriminability and the drop in discriminability caused by shuffling for descriptions generated with individual layers (A) and with increasing number of layers (B). Overall, text generation with deeper layers tended to result in higher discriminability (mean $r = 0.087$ for (A); mean $r = 0.189$ for (B); *t*-test, one-tailed, $P < 0.05$, FDR corrected across subjects). Furthermore, the drop in discriminability caused by shuffling became significantly larger when deeper layers were included in the analysis. Error bars indicate 95% C.I. across samples ($N = 72$; mean $r = 0.180$ [all] and 0.187 [noun] for (A); mean $r = 0.259$ [all] and 0.214 [noun] for (B); *t*-test, one-tailed, $P < 0.05$, FDR corrected across subjects).

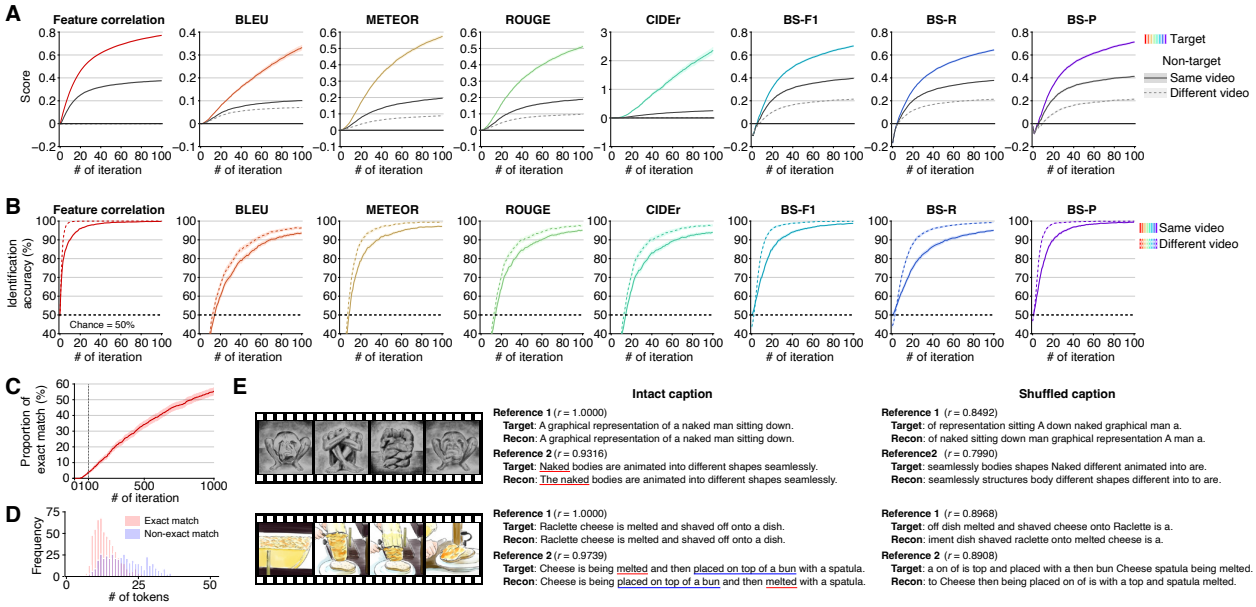

Fig. S10. Performance of caption reconstructions based on features computed from reference captions. To examine how accurately our text optimization method can reconstruct original descriptions from semantic features, we performed a text generation analysis using model-derived (not brain-decoded) features from reference captions of 50 randomly selected videos in the training perception data (20 captions per video; N = 1,000). (**A**) Raw similarity scores between reconstructed descriptions and target captions, along with non-target captions from the same and different videos. (**B**) Pairwise caption identification accuracy based on reconstructed descriptions using candidate captions from the same or different videos. Similarity scores to target captions increased during optimization, with identification accuracy—particularly with feature correlation—approaching 100% after 100 iterations, even for captions from the same video (99.83%), justifying 100 iterations as a computationally efficient and effective setting. (**C**) Proportion of exact matches between reconstructed and original captions across 1,000 iterations. The proportion of exact matches was 4.0% after 100 iterations but increased to 55.4% after 1,000 iterations, indicating that extended optimization can improve reconstruction fidelity. (**D**) Distributions of token counts in target captions. Longer captions were generally harder to reconstruct exactly. (**E**) Examples of reconstructed captions from intact and shuffled references. Exact reconstructions yielded a feature correlation (r) of 1.0, whereas minor word omissions or reordering reduced correlations, demonstrating that the semantic feature space retains sufficient information for accurate reconstruction, even with token-averaged features (see Materials and Methods). However, reconstructions of shuffled captions yielded only 0.2% exact matches even after 1,000 iterations, while recovering most target words and reproducing incoherent structure. These results suggest that while reconstruction difficulty increases with longer or less coherent captions—reflecting optimization limitations and residual MLM priors rather than information loss—our method is highly effective in generating descriptions aligned with target features through iterative feature-guided optimization.

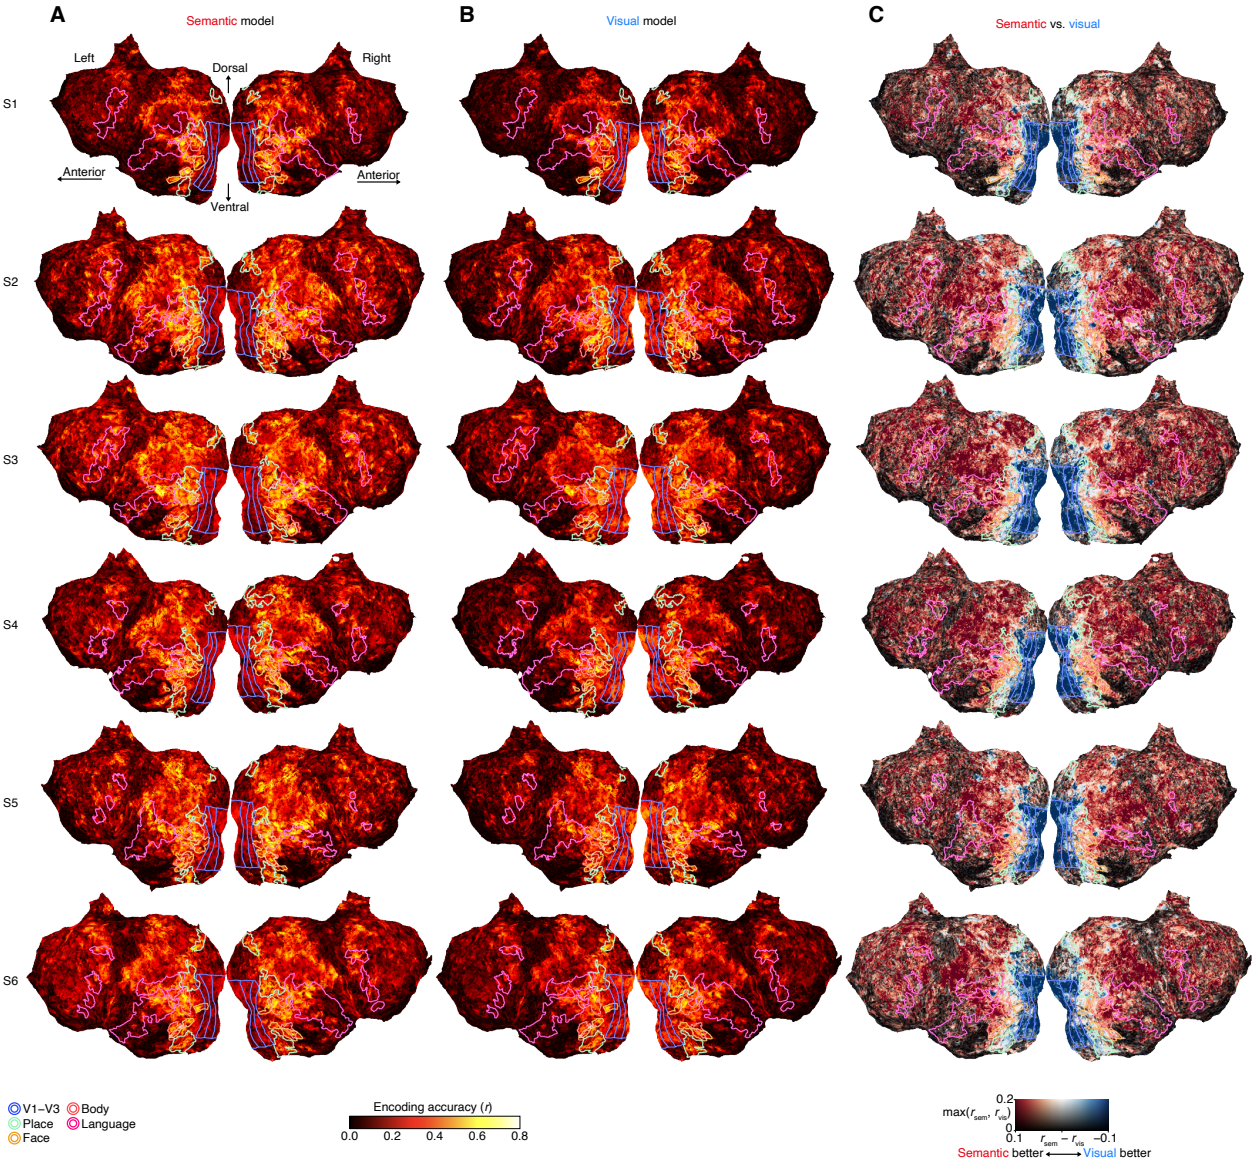

Fig. S11. Voxel-wise encoding model performance.
(**A**) Encoding accuracy of the semantic encoding model. (**B**) Encoding accuracy of the visual encoding model. (**C**) Comparison of encoding accuracy between the semantic encoding model and visual encoding model. Across all subjects, a consistent pattern emerged in which voxels in the occipital visual areas were more accurately predicted by the visual model, whereas voxels in more anterior regions were better predicted by the semantic model. The shift in model superiority occurs at the midpoint of the category-selective regions, between their posterior and anterior halves.

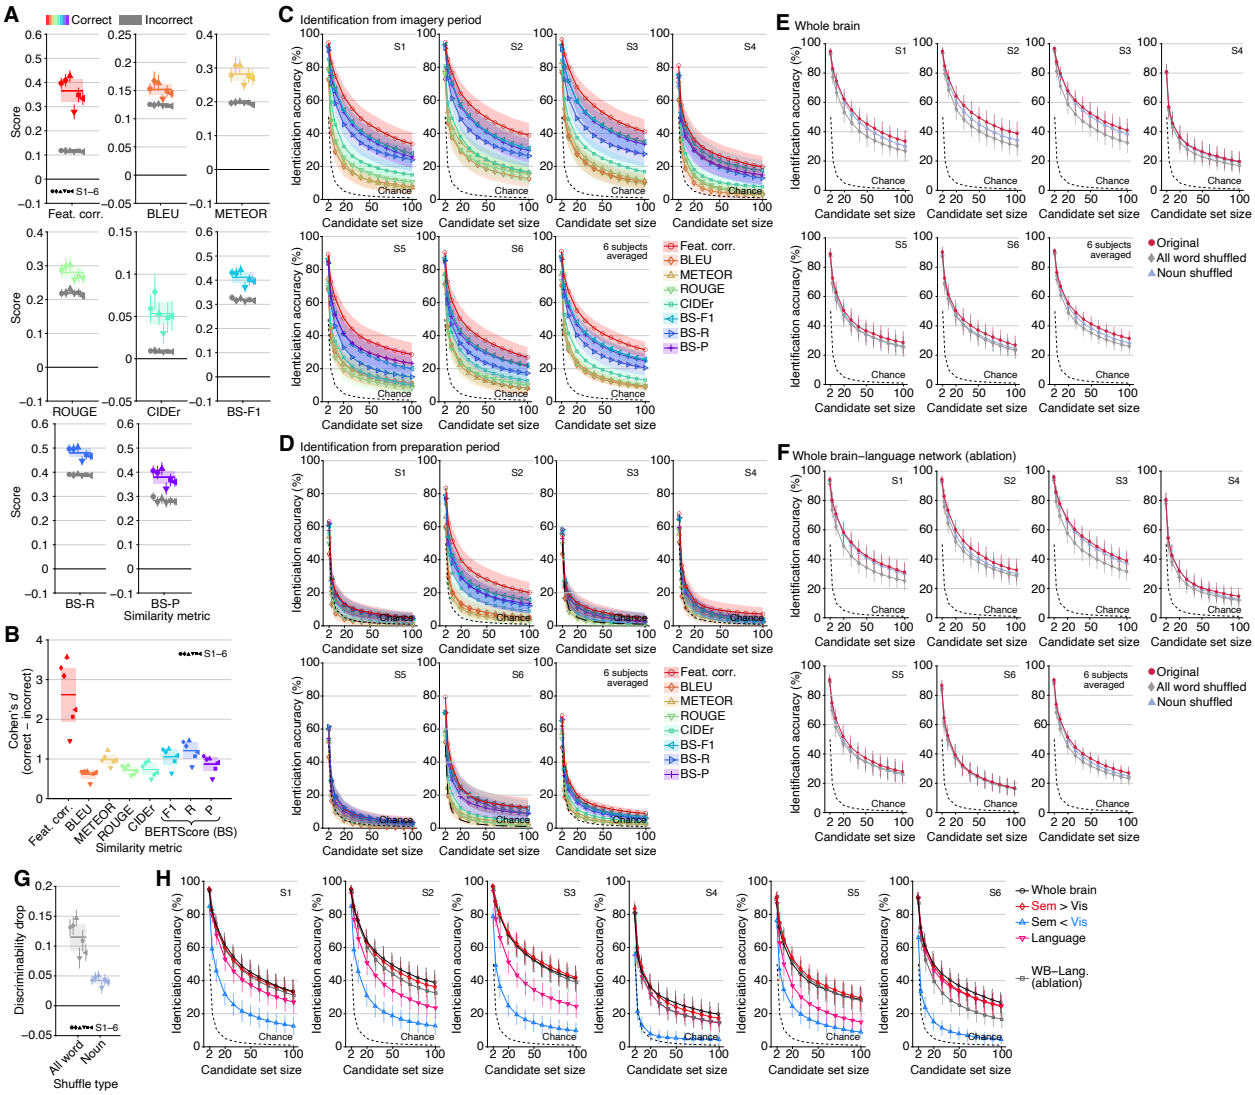

Fig. S12. Text generation performance of recalled content for individual subjects. (**A**) Raw scores for the similarity between generated descriptions and captions for correct and incorrect references based on multiple similarity metrics. (**B**) Cohen's *d* of the discriminability. The generated descriptions showed significantly high discriminability across all metrics and subjects (Wilcoxon signed-rank test, one-tailed, $P < 0.01$, FDR corrected across metrics and subjects). (**C** and **D**) Identification accuracy of recalled videos from brain activity during the imagery period (C) and the preparation period (D). (**E** and **F**), Effects of word-order shuffling on identification accuracy of recalled videos applied to descriptions generated from whole-brain activity (E) and the activity of the whole brain except the language network (F). (**G**) Effects of word-order shuffling on discriminability. The word-order shuffling resulted in significant drops in discriminability even without using the language network (Wilcoxon signed-rank test, one-tailed, $P < 0.01$, FDR corrected across subjects). (**H**) Identification accuracy of recalled videos from different brain areas. Shades in (C, D) and error bars in (A, E–H) indicate 95% C.I. across samples ($N = 72$). Shades in (A, B, G) indicate 95% C.I. across subjects ($N = 6$).

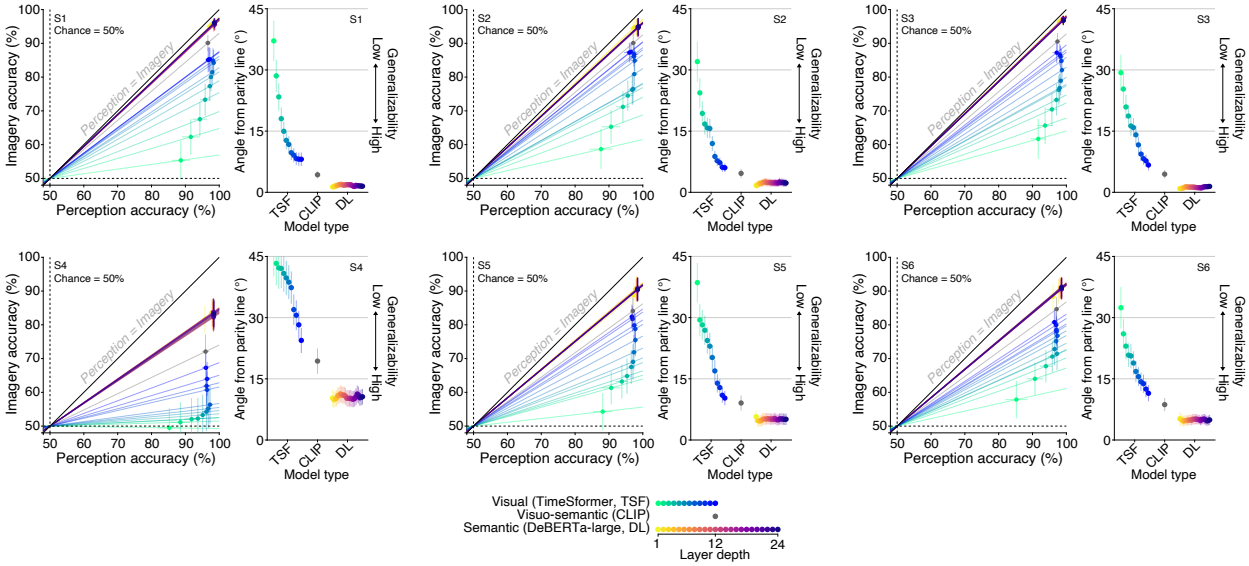

Fig. S13. Comparison of feature-based pairwise video identification accuracy between perception and imagery for individual subjects.
Across all subjects, generalizability—defined as the angle between the parity line and the line connecting chance level (50%) with the observed accuracies—was consistently highest for semantic features across all layers. These results provide robust evidence that semantic features enable perception-trained decoders to effectively generalize for decoding mental content, even at the single-subject level. Conventions are the same as for Fig. 4E.

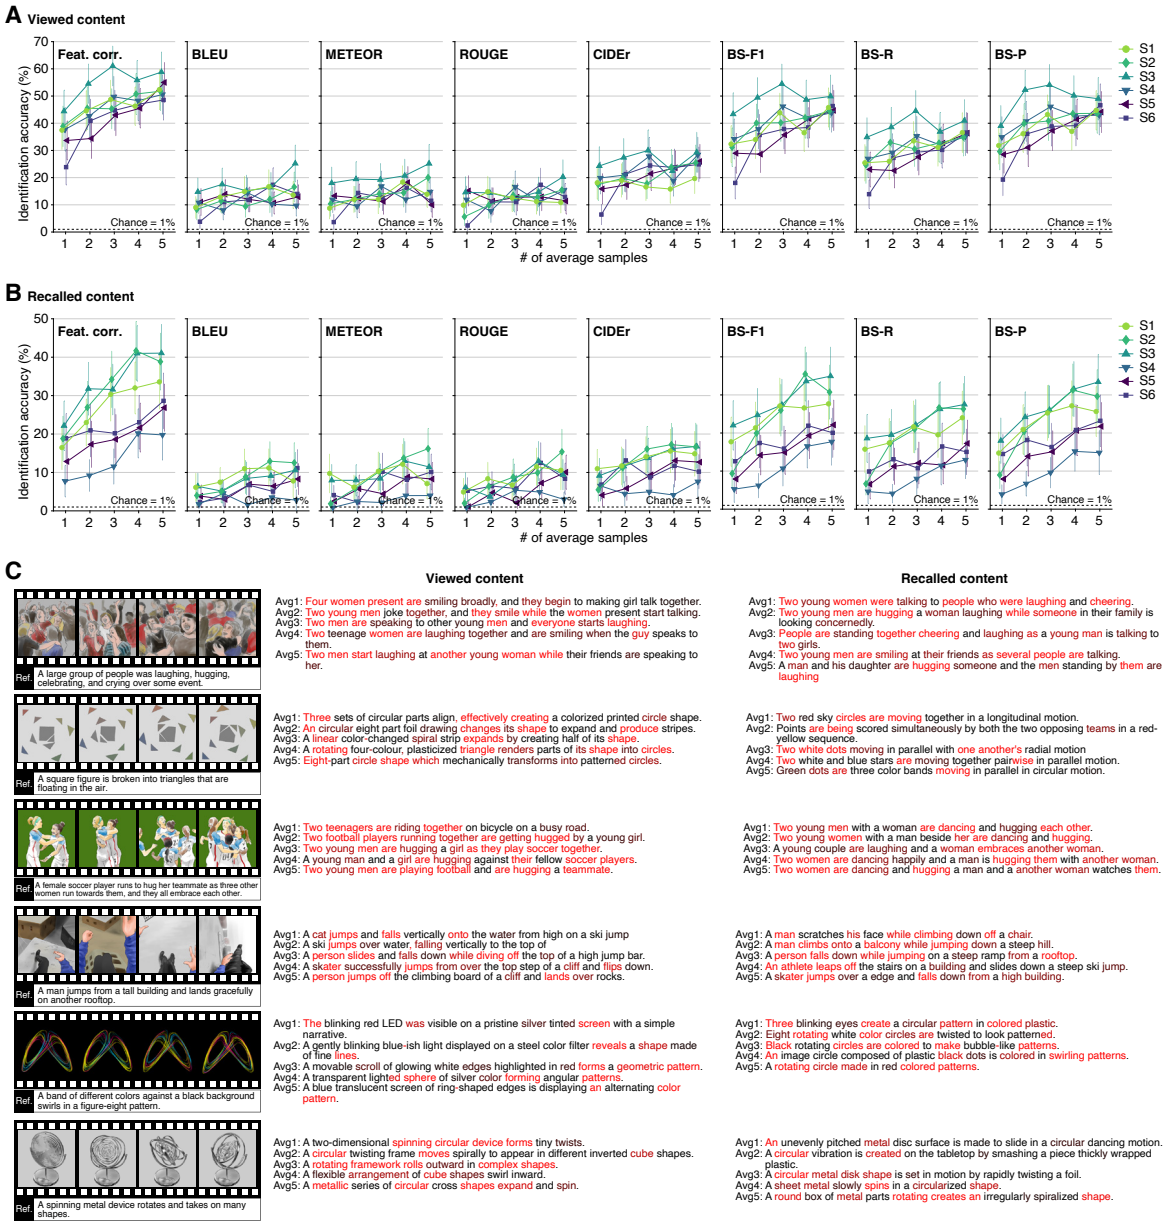

Fig. S14. Performance of text generation with different numbers of averaged samples.
The feature decoding and text generation analyses were performed with brain activity averaged across different numbers of samples. (**A** and **B**) Video identification accuracy of viewed (A) and recalled (B) content with varying numbers of averaged samples. The video identification analysis was performed with 100 candidates (chance = 1%; error bars, 95% CI across samples). (**C**) Descriptions generated using different numbers of averaged samples. For each video, generated descriptions from the same subjects are shown for both viewed and recalled content. The analysis revealed that the quality of the generated descriptions improved as the number of averaged samples increased, indicating that the performance was limited by noise in the fMRI signal of the test data. Meanwhile, the descriptions generated from fMRI activity in single trials reasonably captured the content of both viewed and recalled videos. These results suggest that our method is capable of generating moderately accurate descriptions of both viewed and recalled content, even when using single-trial fMRI activity.