BOLD Moments: modeling short visual events through a video fMRI dataset and metadata

Benjamin Lahner[1,*], Kshitij Dwivedi[2,3,^], Polina Iamshchinina[2], Monika Graumann[2], Alex Lascelles[1], Gemma Roig[3,5], Alessandro Thomas Gifford[2], Bowen Pan[1], SouYoung Jin[1], N. Apurva Ratan Murty[6], Kendrick Kay[4], Aude Oliva[1,+], Radoslaw Cichy[2,+]

[1] Computer Science and Artificial Intelligence Laboratory, MIT, Cambridge, MA, USA
[2] Department of Education and Psychology, Freie Universität Berlin, Germany
[3] Department of Computer Science, Goethe University Frankfurt, Germany
[4] Center for Magnetic Resonance Research (CMRR), Department of Radiology, University of Minnesota, Minneapolis, MN, USA
[5] The Hessian Center for AI (hessian.AI), Darmstadt, Germany
[6] Department of Brain and Cognitive Science, MIT, Cambridge, MA, USA

[+] denotes equal senior author contribution
[*] correspondence: blahner@mit.edu

Abstract

Studying the neural basis of human dynamic visual perception requires extensive experimental data to evaluate the large swathes of functionally diverse brain neural networks driven by perceiving visual events. Here, we introduce the BOLD Moments Dataset (BMD), a repository of whole-brain fMRI responses to over 1,000 short (3s) naturalistic video clips of visual events across ten human subjects. We use the videos' extensive metadata to show how the brain represents word- and sentence-level descriptions of visual events and identify correlates of video memorability scores extending into the parietal cortex. Furthermore, we reveal a match in hierarchical processing between cortical regions of interest and video-computable deep neural networks, and we showcase that BMD successfully captures temporal dynamics of visual events at second resolution. With its rich metadata, BMD offers new perspectives and accelerates research on the human brain basis of visual event perception.

^ work was done prior to joining Amazon

Introduction

Understanding visual events is a hallmark of human intelligence that engages a distributed and functionally diverse network of cortical regions. For example, extracting the meaning of even a simple event, such as a person opening a door, requires at a minimum parsing the scene into relevant components (a person, door, indoor room) (Carandini, 2005; DeYoe & Van Essen, 1988; DiCarlo et al., 2012; Felleman & Van Essen, 1991; Logothetis & Sheinberg, 1996; Ress & Heeger, 2003) and integrating these components over time for accurate action recognition (the door is being opened, not closed) (Fairhall et al., 2014; Hasson, Yang, et al., 2008; Lingnau & Downing, 2015; Orlov & Zohary, 2018; Wurm & Caramazza, 2022), including socially relevant cues (is the person opening the door angry or friendly?) (McMahon et al., 2023; Pitcher et al., 2011; Pitcher & Ungerleider, 2021). Part of the event is also encoded into brain memory regions for later recall (Bainbridge, 2019; Bylinskii et al., 2022; J. Han et al., 2015; Hasson, Furman, et al., 2008; Schneider, 2013). Due to the richness of this cortical network, ecologically-valid visual event understanding is challenging to study with experimental rigor.

To advance empirical evaluation of this network, we introduce the BOLD Moments Dataset (BMD), a dataset of fMRI responses to 1,102 3-second naturalistic videos from ten human adult participants. The videos span a range of events that humans may witness in daily life in order to evoke a representative range of neural responses. Each video clip in BMD is human-annotated with object, scene, action, sentence, and memorability labels to relate brain activity to meaningful aspects of visual events.

Naturalistic videos of a few seconds length strike a good balance between ecological validity and experimental control, making them well suited for investigating human perception of dynamic visual content. Short video stimuli supplement work using still images by driving a greater extent (Bartels & Zeki, 2004; Konen & Kastner, 2008; Press et al., 2001; Schultz & Pilz, 2009; Yildirim et al., 2019) and different pattern (Buccino, Binkofski, Fink, et al., 2004; Kret et al., 2011; Lingnau & Downing, 2015; Wurm & Caramazza, 2022) of neural responses. Short visual events also complement work using longform movies by isolating meaningful actions from long term contextual effects (Hasson, Furman, et al., 2008; Hasson, Landesman, et al., 2008; Roberts et al., 2013).

Below, we detail BMD's experimental design and investigate human dynamic perception using a diverse set of modeling approaches. Aided by BMD's large number of stimuli and widespread reliable activity, we use a video-computable deep neural network (DNN) trained on an action recognition task to predict brain responses through all cortex, including the dorsal visual and parietal cortices. We

additionally input frame-shuffled videos and the first/last second of video frames into DNNs to highlight the temporal content captured by the brain responses: the shuffled video inputs significantly reduce the DNN's prediction accuracy in most regions of interest, and the first (last) second of video frames best explains the early (late) BOLD signal estimates in early and ventral visual cortex. A representational similarity analysis (Kriegeskorte, 2008) comparing language model embeddings of BMD's rich stimuli metadata with multivariate brain activity reveals that the sentence descriptions correlate with ventral and dorsal visual cortex at nearly twice the amount of the object, scene, and action labels. Finally, we show the videos' memorability scores correlate with brain responses not only in the visual cortex, as seen with images (Bainbridge et al., 2017; Bainbridge & Rissman, 2018; Mohsenzadeh et al., 2019), but also in the parietal cortex to further reveal how dynamic visual perception and memory interface.

Together, BMD's 1,102 thoroughly sampled short video stimuli isolate brain responses to ecologically-valid but experimentally controlled dynamic content. This dataset is well-suited to address open questions as diverse as developing methodologies to model rapid event BOLD signals (Misaki et al., 2013; Prince et al., 2022; Wittkuhn & Schuck, 2021), characterizing interactions between visual processing pathways (Lingnau & Downing, 2015; Mineault et al., 2021; Pitcher & Ungerleider, 2021; Wurm & Caramazza, 2022), and bridging the gap between still image and longform movie perception (Aliko et al., 2020; Allen et al., 2022; Hanke et al., 2016; Hebart et al., 2019) (see Supplementary 1 section, "The added value of brain responses to a short video dataset versus a static image dataset" for further discussion). Crucially, BMD's shared experimental design across subjects enables robust and generalizable conclusions, and its preparation in BIDS format facilitates easy adoption among researchers. Its scale of stimuli and subjects, reliable brain activity, and stimuli metadata enable a wide range of interdisciplinary analyses to reveal the neural mechanisms underlying visual event understanding.

Results

Sampling brain activity for 1,102 distinct visual events

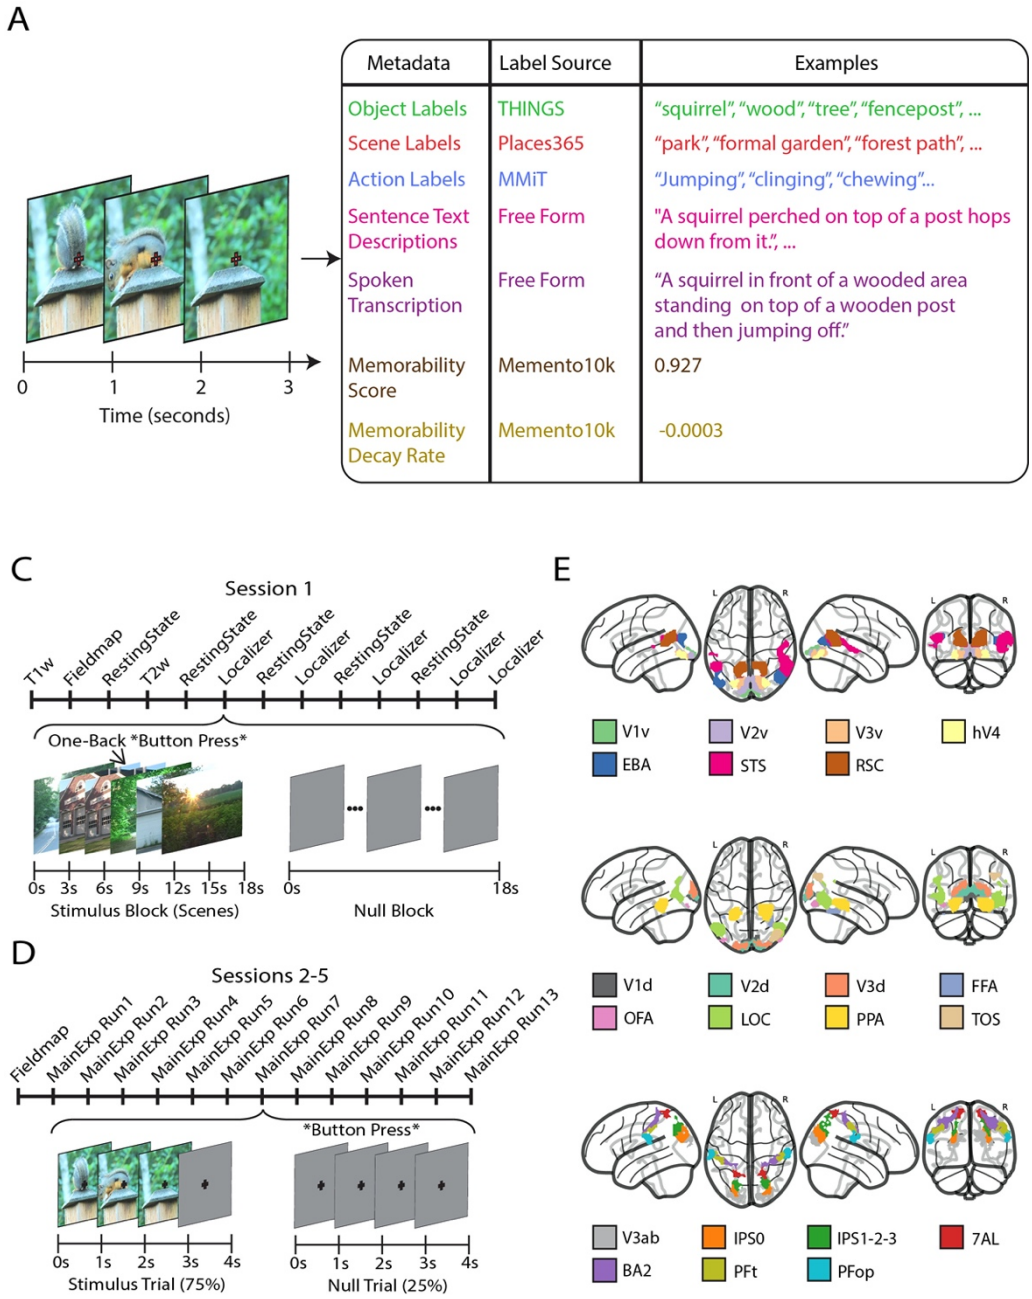

Figure 1: Experimental design and data acquisition.

(A) Stimuli and metadata: The 1,102 3-second video stimuli are sampled from the Moments in Time dataset (Monfort et al., 2020, 2022) and annotated with: 15 object labels, 5 scene labels, 5 action labels, 5 sentence

descriptions, 1 spoken transcription, 1 memorability score, and 1 memorability decay rate. **(B) fMRI session 1:** Subjects underwent T1- and T2-weighted structural runs interspersed with resting state and functional localizer functional runs to define category selective regions. **(C) fMRI sessions 2-5:** Sessions were identical in design and consisted of a fieldmap run followed by thirteen runs (ten training runs and three testing runs) in random order. The 3-second videos were presented on a gray background at a five degree visual angle overlaid with a red fixation cross. Stimuli presentation was followed by a 1-second intertrial interval composed of a red fixation cross on a gray background. Participants reported luminance changes of the fixation cross occurring irregularly between videos with a button press and did so reliably (hit rate 0.964 +/- 0.014 (mean ± SD)). **(D) Region of Interest Definitions.** Regions of interest for a representative subject (subject 1), functionally (from session 1) and anatomically defined.

We sampled 1,102 naturalistic 3-second videos depicting diverse dynamic visual events from the Memento10k video memorability dataset (Newman et al., 2020), which is a subset of the 1 million video large Moments in Time (MiT) and Multi-Moments in Time (MMiT) datasets (Monfort et al., 2020, 2022). As used here, the term visual event is a 3-second video depicting someone or something performing an action in the real world. Understanding a visual event reflects a viewer's ability to parse and integrate the video's spatiotemporal information to report its relevant happenings, such as in a sentence description. The 3-second events from Moments in Time datasets were typically cropped from a longer in-the-wild video (e.g., a 3-second segment of a dance recital) and contain no emotionally jarring or inappropriate content. Their naturalistic content engages rich, real-life neural processes (Berkes et al., 2011; Olshausen & Field, 1996a, 1996b; Smyth et al., 2003). The 3-second video duration, approximately the duration of human working memory (Baddeley, 1992; Barrouillet et al., 2004; Schneider, 2013), is an ideal length to capture brain responses to meaningful events in context (e.g., opening a door, cooking BBQ food, playing tennis). Videos shorter than 3 seconds risk capturing incomplete actions (e.g. extending the arm) while videos longer than 3 seconds risk undesired complexity by capturing a sequence of actions (e.g., opening a door, then walking into the kitchen, and finally removing a jacket) that introduce highly correlated events that are difficult to disentangle from brain responses.

The experimental design weighted benefits of both a large stimulus set and several stimulus repetitions to enable downstream analyses that depend more strongly on one or the other. We thus divided the 1,102 stimuli into two non-overlapping sets, one consisting of 1,000 stimuli with three repetitions per subject and the other consisting of 102 stimuli with ten repetitions per subject. We use terminology common to encoding studies, naming the former set the training set and the latter set the testing set. Each participant saw the same set of 1,102 videos. This amounts to 4,020 unique fMRI trials per

participant, and 40,200 unique fMRI trials across the entire dataset. The large number of trials and participants enable direct comparisons of results across participants and invite potential integration of data across participants (Hasson et al., 2004; Hasson, Furman, et al., 2008; Haxby et al., 2011, 2020).

Functional brain responses were recorded with whole-brain 3T fMRI at 2.5 x 2.5 x 2.5 mm voxel resolution to densely sample the widely-distributed cortical responses to video across the whole cortex (Buccino, Binkofski, Fink, et al., 2004; Buccino, Binkofski, & Riggio, 2004; Buccino, Vogt, et al., 2004; Gazzola & Keysers, 2009; Kret et al., 2011; Rizzolatti & Sinigaglia, 2010; Yildirim et al., 2019). The fMRI experiment was split into 5 sessions. Session 1 (Figure 1B) contained crucial auxiliary brain measurements, interleaving high-resolution T1- and T2-weighted structural, video-based functional localizer (Lafer-Sousa et al., 2016), and resting state scans (Hutchison et al., 2013; S. M. Smith et al., 2013). Sessions 2 to 5 (Figure 1C) contained the main experiment and had identical structure. Each video trial was 4s long, consisting of a 3s silent video presentation followed by a 1s intertrial interval.

Semantic and behavioral metadata on visual events

Revealing how the brain mediates visual event understanding benefits from detailed, human-labeled descriptions of a visual event. Thus, we used human crowd-sourced experiments to annotate each clip with five word-level scene, object, and action labels, five sentence-level text descriptions, one spoken transcription, one behavioral memorability score, and one memorability decay rate (Figure 1A).

Each category of metadata labels was collected in separate experiments. The possible scene, object, and action labels were sampled from the Places365 (B. Zhou et al., 2018), THINGS (Hebart et al., 2019, 2023), and Multi-Moments in Time (Monfort et al., 2022) datasets, respectively, for their carefully crafted label coverage and extensive overlap with computer science resources. Due to most videos clearly containing more than one object, the annotators in the object label experiment were instructed to label three different objects for a total of 15 object labels per video. The 1,102 BMD stimuli cover 305 of 365 possible scene labels, 1,002 of 1,854 possible object labels, and 261 of 292 possible action labels. The sentence text descriptions (13.06 mean +/- 2.800 std words per sentence) were typed free-form by the annotator to detail salient interactions in the video and summarize the pertinent content. The spoken transcriptions, collected through free-form audio recordings and transcribed (audio not available due to privacy considerations), tend to be more verbose with additional emotional and linguistic subtleties present in speech but often not in text (26.72 mean +/- 17.55 std words per transcription) (Monfort et al., 2021). Finally, to determine how visual event understanding interfaces with memory, we

use a crowd-sourced memory game to behaviorally measure if participants recognize a visual event at a later time period (memorability: 0.8422 mean +/- 0.0888 std) and how this recognition performance fades over time (memorability decay rate: -0.0014 mean +/- 0.0011 std) (Newman et al., 2020).

This metadata allows the use of grouping, subdivision, or other transforms of these measures to investigate additional perceptual and cognitive processes related to visual event understanding (see Supplementary Figure S8 for metadata distributions between train and test splits). We provide first example analyses on this basis in the following sections.

(f)MRI data processing, response modeling, and ROI definition

We provide raw as well as preprocessed versions of the MRI data in the community-backed and standardized BIDS format (K. J. Gorgolewski et al., 2016) to ensure transparent quality assessment, reproducible preprocessing, and easy-to-share results. The raw data gives researchers control over preprocessing to pursue research questions at any stage of the analysis pipeline. The preprocessed data allows for analyses into the spatial and temporal neural dynamics underlying visual event understanding, at both the group or single-participant level.

We first processed the data using the standardized fMRIPrep tool (Esteban et al., 2019) to achieve reproducible and transparent results. During acquisition, we sampled fMRI data at a TR of 1.75s to achieve a densely sampled time course with respect to the onset of the stimulus (K. Kay et al., 2020); since 1.75 does not evenly divide into the 4 second trial length, the BOLD response to a single video across three repetitions might be sampled at onset delays of [0.5, 2.25, 4, 5.75,…], [0, 1.75, 3.5, 5.25, …], and [1, 2.75, 4.5, 6.25,...] seconds for the three (or ten) repetitions. In the preprocessing we then temporally resampled the data from a TR of 1.75s to a TR of 1s in order to acquire BOLD responses exactly time-locked to the onset of each trial (stimulus presentation happens every 4s, i.e., a multiple of 1s, but not 1.75s). To accommodate variability in the hemodynamic responses at different locations in cortex, to resolve temporal structure of the stimuli, and to address response overlap from the rapid event-related design, we used Finite Impulse Response (FIR) basis functions. We modeled the hemodynamic response to visual events using data from 1-9s after stimulus onset (to account for the hemodynamic lag) in 1s steps (i.e. 9 bins of 1s length each) for each trial separately. The FIR model's flexible BOLD estimates, randomization of stimulus presentations across runs and sessions, and ability to average over multiple repetitions reduce potential unwanted memory effects.

To guide analysis in a region-specific manner, we used the auxiliary functional localizer data from session 1 to define a set of 22 regions of interest (ROIs) previously reported to be involved in visual perception of natural images, natural video, or motion (Figure 1D) (Gazzola & Keysers, 2009; Le et al., 2017; Logothetis & Sheinberg, 1996; Rizzolatti & Sinigaglia, 2010; Silver & Kastner, 2009; VanRullen & Thorpe, 2001). The set includes early visual, category-selective ventral visual, dorsal visual, and parietal regions.

All together, this preprocessing suite ensures a low threshold for researchers to interact with the brain data at their desired processing level. All results presented here in the main text use version A of the dataset. We provide another preprocessed version of BMD (version B) to make it available in more output spaces (version B is detailed in Supplementary 2).

(f)MRI image scans of high quality across subjects and task

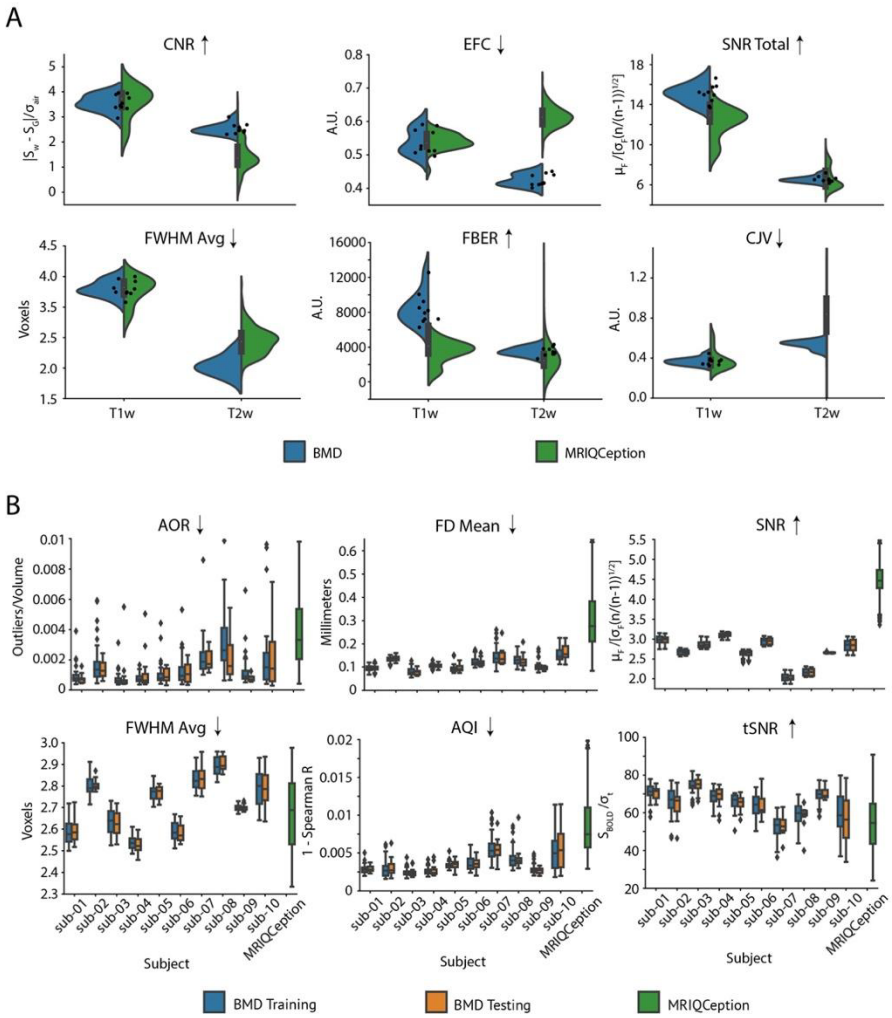

Figure 2: Preprocessed scan quality measures

A. T1- and T2-weighted structural scan data quality: Violin plots compare Image Quality Metrics (IQMs) of the T1- and T2-weighted structural scans from our BMD dataset (T1-weighted, n=10; T2-weighted, n=10) with a collection of anonymous data pulled from the MRIQCeption API (T1-weighted, n=219; T2-weighted, n=695). We report the Signal to Noise Ratio (SNR Total), Contrast to Noise Ratio (CNR), Coefficient of Joint Variation (CJV), Entropy Focus Criterion (EFC), Average Full-Width Half Maximum Smoothness (FWHM Avg), and Foreground-Background Energy Ration (FBER) as a summary of the structural scan quality. Points correspond to the data value from an individual subject in the BMD dataset. An up arrow after the IQM title means a higher value corresponds to better data quality, while a down arrow after the IQM title means a lower value corresponds to better data quality. A.U. stands for "arbitrary units." **B. Training and testing functional scans data quality:** Panels compare boxplots of IQM values for each subject for the training and testing functional tasks in the BMD

dataset (per subject: training, n=40; testing, n=12) with anonymous BOLD data pulled from the MRIQCeption API (BOLD, n=687). We report the Signal to Noise Ratio (SNR), Temporal Signal to Noise Ratio (tSNR), Mean Framewise Displacement (FD Mean), Average Full-Width Half Maximum Smoothness (FWHM Avg), AFNI Outlier Ratio (AOR), and AFNI Quality Index (AQI) to summarize the functional scan quality. The boxplot extends 1.5 times the high and low quartiles, with outliers defined as a scan with a value outside that range and denoted by diamonds. An up arrow after the IQM title means a higher value corresponds to better data quality, while a down arrow after the IQM title means a lower value corresponds to better data quality.

To assess the quality of the raw and preprocessed MRI data, we used the open-source and community-based MRIQC analysis package (Esteban et al., 2017). This yielded a comprehensive set of 44 functional and 68 structural image quality metrics (IQMs) (full report available), of which we present a representative set of six structural (Figure 2A) and six functional MRI metrics (Figure 2B) (see Supplementary for more details on the IQMs and Supplementary Figure S1 for the resting state and functional localizer scans). In brief, since no set of metrics can completely describe data quality, the selection of the IQMs to display here considered metrics especially relevant to functional or structural scans (e.g., Temporal SNR for functional scans and Contrast to Noise ratio for structural scans), metrics shared between functional and structural scans for a more cohesive set (e.g., SNR and Full-Width Half Maximum Smoothness), and metrics common across other literature and analysis packages to increase familiarity with readers (e.g., SNR, Framewise Displacement, AFNI Outlier Ratio, AFNI Quality Index). Since most quality metrics do not have a ground-truth reference value to compare against, we contextualize our values against the values of hundreds of anonymized studies with similar scanner parameters (1 < Tesla < 3, 1 <= TR < 3; aggregated with the MRIQCeption API).

Taken together, the IQMs indicate BMD's structural and functional scans are of excellent quality. The distribution of all but one IQM for each subject falls within or is noticeably better than the distribution seen in similar fMRI studies (Figure 2, green). The strong results of the IQMs tSNR (a measure of SNR over time), aor (an indicator of the number of outliers per fMRI volume), and aqi (a correlational measure of quality per volume) assure satisfactory functional SNR quality in light of the below-typical per volume SNR IQM. We highlight that within each participant, the range of quality metric values was especially consistent between the training and testing sets (Figure 2, blue and orange boxplots). This shows that the training and testing sets are of comparable data quality, facilitating analyses that depend on this split. Further, none of the 10 participants were outliers as indicated by consistently lower within- than between-participant variability, encouraging group result inference.

Reliable univariate and multivariate fMRI response profiles

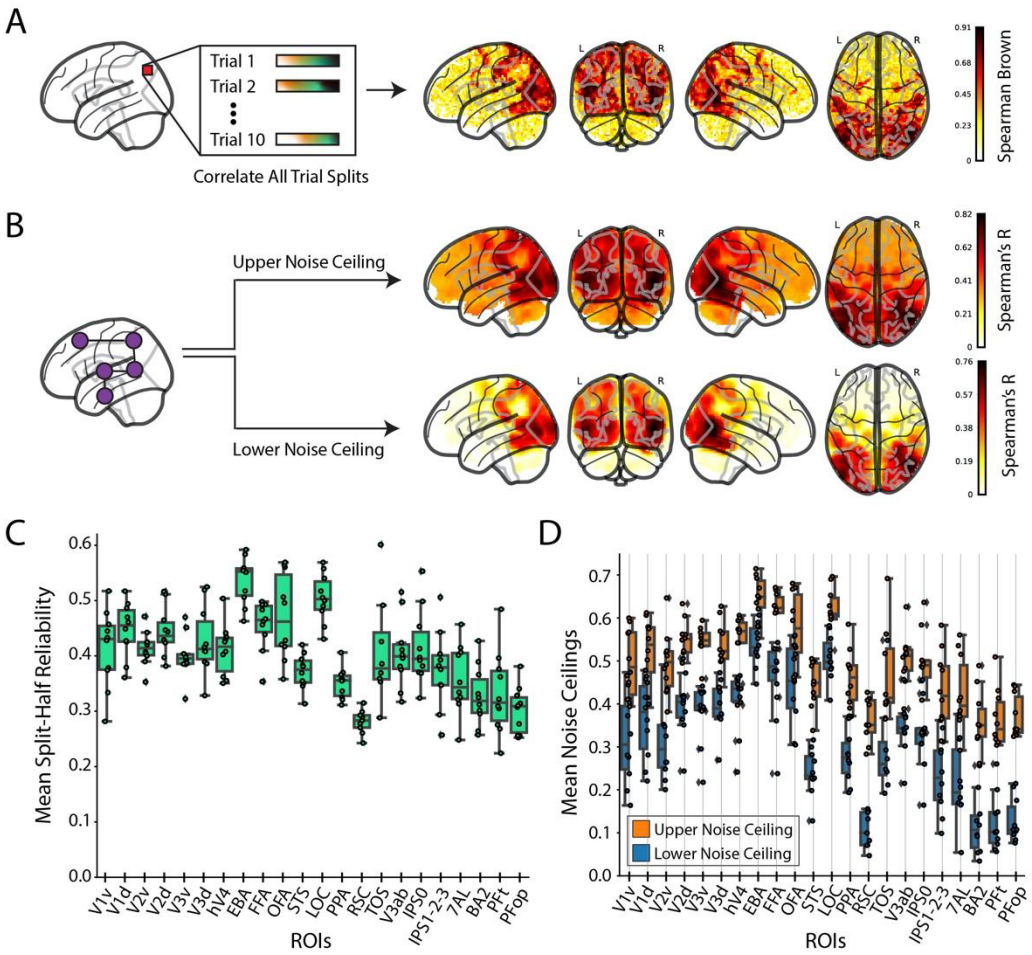

Figure 3: Whole-brain and ROI reliability after response modeling

(A) Whole-brain single-subject split-half reliability analysis: We perform a voxelwise split-half reliability analysis and present the voxels that pass the reliability criteria (p < 0.05, Spearman-Brown ($2\rho/(1+\rho)$)) for a representative subject (subject 1). **(B) Whole brain searchlight noise-ceiling analysis:** We estimate the upper and lower noise-ceilings across the whole brain for a representative subject (subject 1) from a searchlight representational dissimilarity analysis (RSA). **C) ROI-based group split-half reliability:** For each of the 22 ROIs, we present the mean split-half reliability across participants from the voxels passing the reliability criteria. **(D) ROI-based group searchlight noise ceilings:** For each of the 22 ROIs, we calculate the upper (orange) and lower (blue) noise ceilings across participants. The box plot encompasses the first and third data quartiles and the median (horizontal line). The whiskers extend to the minimum and maximum values within 1.5 times the interquartile range, and values falling outside that range are considered outliers (denoted by a diamond). The overlaid points show the value at each observation (n=10 for all ROIs except TOS (n=8) RSC (n=9)). The brain

responses used for the reliability analyses are the beta values averaged over TRs 5-9 (the peak of the BOLD signal) from the testing set.

We provide both univariate and multivariate reliability measures to evaluate the reliability in the context of the two main fMRI analysis traditions today: the univariate framework that focuses on local information at a single voxel scale (Friston et al., 1994; M. Khosla et al., 2022; Naselaris et al., 2011; Ratan Murty et al., 2021; Schrimpf, Kubilius, Hong, et al., 2020), and the multivariate analysis framework that emphasizes the distributed nature of information in population codes (Haxby, 2012; Haynes, 2015; Kriegeskorte, 2008). Neural encoding and decoding analyses often use single-voxel responses to measure information content in an ROI (Naselaris et al., 2011), and hypothesis models of neural representations derived from computational or behavioral models can be easily compared to the brain in multivariate analysis frameworks (Kriegeskorte & Kievit, 2013). These reliability measures provide intuition on data quality and can used at the discretion of the researcher to normalize results with respect to the noise in BMD's data.

To assess univariate reliability, we identify voxels whose beta value estimates satisfy a Spearman-Brown split-half reliability criterion (mean reliability of all combinations of 5-trial averaged splits from the 10-trial testing set) of $p<0.05$ (assessed by stimulus label permutation).

To assess multivariate reliability, we perform representational similarity analysis (RSA) in a searchlight approach (Kriegeskorte et al., 2006) to determine the upper (subject-to-group RDM correlation per voxel) and lower (leave-one-out RDM correlation per voxel) estimate of the noise ceilings. We report the univariate and multivariate reliability results across the whole brain (Figure 3AB for a representative subject, see Supplementary Figures S2 and S3 for all subjects) and in ROIs (Figure 3CD).

In both the whole-brain univariate and multivariate reliability analyses, we observe statistically significant reliability values across the occipital, temporal, and parietal cortices, even extending into the frontal lobe. The ROI analyses yield equivalent results, showing high explainable variance in a functionally diverse set of ROIs responsible for visual event understanding. This demonstrates that BMD is well suited for comprehensive and advanced analysis at both the single- and multi-voxel spatial scale.

Modeling visual event understanding with a video-computable deep neural network for action recognition

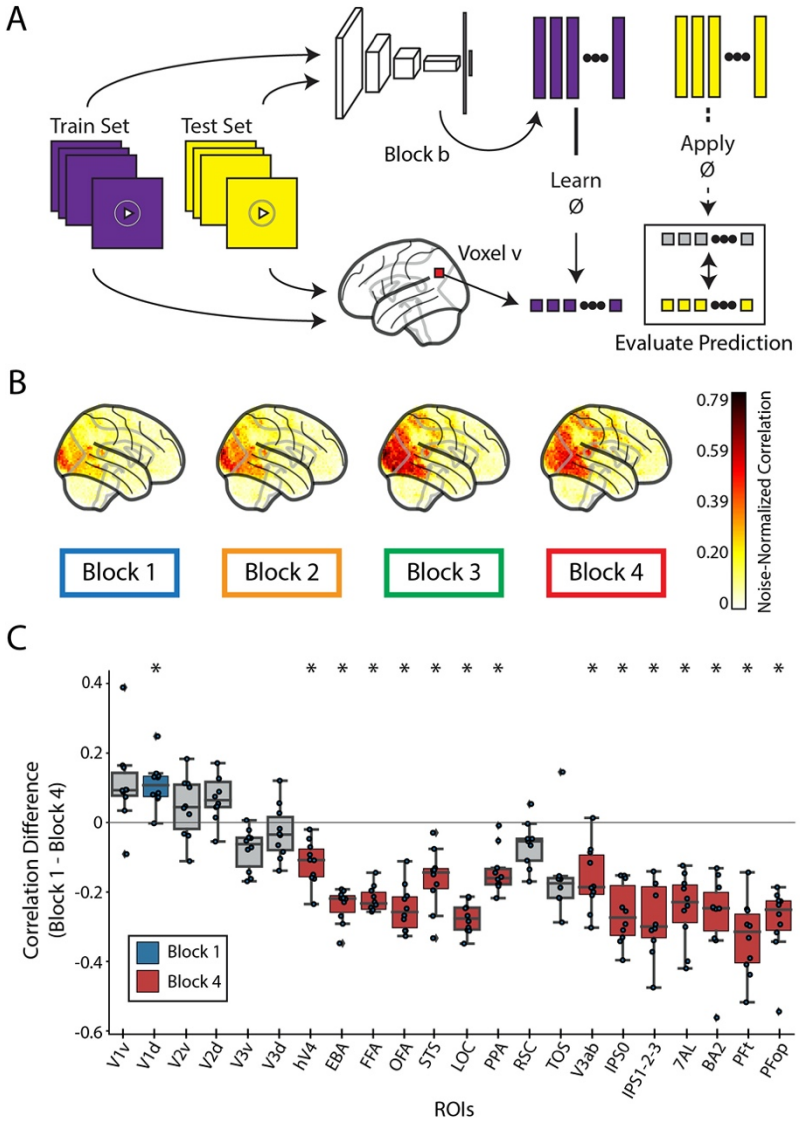

Figure 4: Evaluation of biologically-similar video-based encoding model

(A) Voxelwise encoding model procedure: All videos are shown to both a TSM ResNet50 DNN and a human. Training set video embeddings are extracted from a block *b* of the DNN and used to learn a voxelwise mapping function to the human responses. This mapping is then applied to the testing set video embeddings to predict the brain response at each voxel. **(B) Whole-brain encoding accuracy across blocks:** We use the encoding model procedure with each of the four blocks of a TSM ResNet50 model trained to recognize actions in videos to predict the neural response at each voxel in the whole brain. The brain figures show the subject-average noise-normalized predictive correlation (divided by the voxel's upper noise ceiling) at each voxel. **(C) ROI-based**

encoding accuracy difference: Difference in predictive performance between block 1 and block 4 at each of the 22 ROIs. Predictive performance at each voxel is measured as the noise-normalized correlation between the brain responses and the predicted responses, averaged over all reliable voxels in each ROI. Significant ROIs are denoted with an asterisk and a color (blue for Block 1, red for Block 4, gray is not significant) corresponding to the significant layer (one sample two-sided t-test against a population mean of 0, Bonferroni corrected across 22 ROIs, $p < 0.05$). The box plot encompasses the first and third data quartiles and the median (horizontal line). The whiskers extend to the minimum and maximum values within 1.5 times the interquartile range, and values falling outside that range are considered outliers (denoted by a diamond). The overlaid points show the value at each observation (n=10 for all ROIs except TOS (n=8) RSC (n=9)).

A major goal in explaining brain responses is providing explicit quantitative models that predict the underlying computations and their cortical organizations (Dayan & Abbott, 2001). Deep Neural Network (DNN)-based modeling has emerged as a dominant form of scientific modeling in visual neuroscience, due to their image-computable design, biologically-inspired architecture, and high neural prediction performance (Allen et al., 2022; Chang et al., 2019; Rajalingham et al., 2018; Schrimpf, Kubilius, Hong, et al., 2020; Schrimpf, Kubilius, Lee, et al., 2020; D. L. Yamins et al., 2013).

However, modeling visual events has been limited by the lack of a suitable dataset that accounts for complex distributed processes across the whole brain (Yildirim et al., 2019), drastic differences to image understanding (Buccino, Binkofski, Fink, et al., 2004; Krekelberg et al., 2003; Kret et al., 2011; Lingnau & Downing, 2015; Schultz & Pilz, 2009; Senior et al., 2000; Shirai & Imura, 2014; Wurm & Caramazza, 2022), and the temporal boundaries of a visual event (Aliko et al., 2020; Hasson, Yang, et al., 2008; Nishimoto et al., 2011; Seeliger et al., 2019). BMD helps with its large number of short video stimuli and whole-brain responses.

Towards the goal of modeling visual events, we used a video-computable DNN following biological constraints. The DNN uses a recurrent ResNet50 backbone (He et al., 2016) that mimics the biological recurrent computations essential for human motion perception and categorization (Kietzmann et al., 2019; Koivisto et al., 2011; Pascual-Leone & Walsh, 2001; Silvanto, Cowey, et al., 2005; Silvanto, Lavie, et al., 2005) and builds on the ResNet family's strong neural predictivity performance seen for still images (Kubilius et al., 2019; Schrimpf, Kubilius, Lee, et al., 2020). The ResNet50's four recurrent blocks are connected using a Temporal Shift Network (Lin et al., 2019) designed to process video input in the natural, uni-directional temporal order. We train the model on an action recognition task using the same dataset from which the BMD stimuli were sampled, the Moments in Time dataset (Monfort et al., 2020, 2022) (BMD stimuli were excluded from model training; see Methods section "Action Recognition

TSM ResNet50 Model Training" for training details). We release this model to aid investigations of visual event understanding (model available at: https://github.com/pbw-Berwin/M4-pretrained).

We queried the relationship between each of the model's blocks and BMD. Using a voxelwise encoding model approach (Naselaris et al., 2011) (Figure 6A), we observe a correspondence between DNN block depth and predictivity performance along the visual processing hierarchy and beyond (see Methods section "DNN Block to Cortex Correspondence Procedure" for encoding model details). The predictivity of each of the DNN's four blocks progressively spreads across the brain from posterior to anterior (Figure 6B). As a proxy for low-level and high-level features, we extract the video features at early (Block 1) and late (Block 4) blocks of the DNN, respectively (Cichy, Khosla, et al., 2016; Kubilius et al., 2019; D. L. Yamins et al., 2013; D. L. K. Yamins et al., 2014). We then estimate the dominance of low-level and high-level feature processing across ROIs by computing the difference in predictivity between DNN Block 1 and DNN Block 4. We observe that predictivity of DNN Block 4 becomes significantly greater than DNN Block 1 beginning in the ventral visual cortex and extending into dorsal visual cortex and parietal cortex (Figure 6C), reflecting the increase of feature complexity in the representations across the visual processing hierarchy. We repeat this encoding analysis with different architectures and training diets to see how this result generalizes across models. Specifically, we extract features from a TSM MobileNetV2 (Lin et al., 2019; Sandler et al., 2018) and a TimeSformer (Bertasius et al., 2021) architectures trained on Kinetics-400 (W. Kay et al., 2017) and HowTo100M (Miech et al., 2019) datasets, respectively. We find that both are good predictors of dorsal visual and parietal regions, but the TSM MobileNetV2 does not show increasing dominance of high-level model features along the cortical hierarchy (see supplementary Figure S7 for results on all architectures and blocks).

This result extends previous research demonstrating a hierarchical correspondence between DNNs and brains from still image stimuli (Cichy, Pantazis, et al., 2016; Eickenberg et al., 2017; Kriegeskorte, 2015; Kubilius et al., 2019; D. L. K. Yamins et al., 2014) to dynamic video stimuli, a non-trivial outcome given that many cortical regions in the ventral visual and temporal cortex and beyond respond to stimulus features uniquely present in videos and not images (e.g., movement kinematics, temporal interactions) (Lingnau & Downing, 2015; Pitcher et al., 2011; Wurm et al., 2017; Wurm & Caramazza, 2022). The finding that a transformer-based TimeSformer model follows a similar pattern to the TSM ResNet50 while the TSM MobileNetV2 does not invites further inquiry into the effect of training diet, parameter count, and architecture on visual event understanding. Finally, these results also help clarify previously conflicting results about whether or not DNNs trained on action recognition tasks accurately

predict dorsal stream regions (Bakhtiari et al., 2021; Güçlü & van Gerven, 2017; Mineault et al., 2021), showing that all three architectures accurately predict responses not only in the dorsal visual stream but also in the parietal cortex.

fMRI responses capture temporal event structure

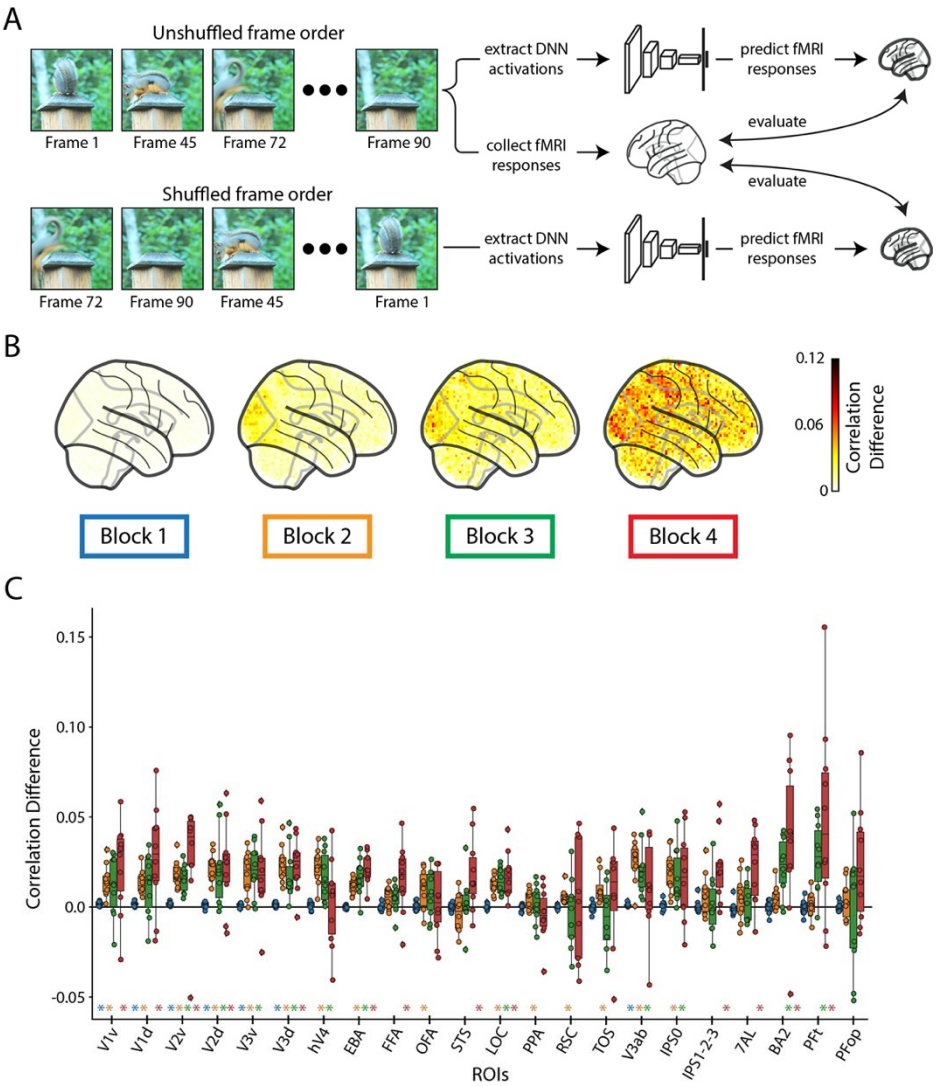

Figure 5: Importance of temporal order in predicting fMRI responses.

A) **Frame shuffling procedure:** We predict the real fMRI responses using both the DNN activations of the original (unshuffled) frame order and the DNN activations of the randomly shuffled frame order. A difference in prediction accuracy between the DNN activations of the unshuffled and shuffled frames indicates the preservation of correct temporal order in the fMRI response. **B**) **Whole-brain prediction difference:** Difference in the correlation averaged over participants between the shuffled framed prediction accuracy and unshuffled frame

prediction accuracy across the whole brain at different DNN layers (TSM model). **C**) **ROI-based prediction difference:** Difference in the correlation between the shuffled frame prediction accuracy and unshuffled frame prediction accuracy at different ROIs and DNN layers (TSM model). A colored asterisk along the x-axis indicates significant difference between the unshuffled and shuffled prediction accuracy at that DNN block (one sample two-sided t-test against a population mean of 0, FDR correction across 22 ROIs x 4 blocks = 88 comparisons, $p < 0.05$). The box plot encompasses the first and third data quartiles and the median (horizontal line). The whiskers extend to the minimum and maximum values within 1.5 times the interquartile range, and values falling outside that range are considered outliers (denoted by a diamond). The overlaid points show the value at each observation (n=10 for all ROIs except TOS (n=8) RSC (n=9)).

A real world event unfolds in a systematic, spatiotemporal sequence. Is the temporal structure of such an event important for its cortical representation? Shuffling the frames of a video effectively destroys its meaningful temporal structure. We reasoned that if a voxel captured an event's spatiotemporal relationships, unshuffled (meaningfully ordered) video input would correspond better to the brain responses than shuffled video input. The shuffled and unshuffled video inputs are both temporally dynamic and contain identical frame-averaged spatial content, thereby isolating the effects of the encoding of ordered temporal content.

To investigate, we measured the neural prediction performance of the DNN model introduced above when given unshuffled video input and shuffled video input (Figure 5A). Since this DNN is engineered to process video input in a unidirectional temporal order, we know that shuffling the video frames will affect the DNN activation in at least one of the four blocks. By using the shuffled and unshuffled activations from these blocks to predict the brain responses, we can assess the effect of correct temporal ordering on our BMD brain responses (see Methods section "Shuffling Analysis to Determine Importance of Temporal Order" for details).

Our results show the frame-shuffled input decreased DNN prediction accuracy across most of the visual system. The whole-brain analysis (Figure 5B) shows that this decrease in DNN prediction accuracy is most pronounced in both magnitude and coverage with increasing DNN blocks. The ROI analysis (Figure 5C) reveals that early visual ROIs were sensitive to differences in all DNN blocks, while ventral and dorsal ROIs were most affected by differences in DNN blocks 3 and 4. These results suggest functionally-specific sensitivities to shuffled input along the visual system in line with known feature preferences in the visual system and DNN computations (Honey et al., 2012; Kubilius et al., 2019; Logothetis & Sheinberg, 1996).

In order to further elucidate the type of temporal dynamics in the brain that this frame shuffling analysis is affecting, we perform an additional small-scale experiment comparing the difference in prediction accuracy between a Temporal Shift Module (TSM) ResNet50 (Lin et al., 2019) and a Temporal Segment Network (TSN) (L. Wang et al., 2016) ResNet50. Since the only difference between the two models is how spatial information is shared between frames, this analysis effectively isolates temporal integration. We find that implementation of TSM significantly improves encoding accuracy in early visual ROIs primarily from DNN blocks 3 and 4 (see Supplementary 1 Figure S9 for results). This pattern is different from the one observed from the frame shuffling analysis (Figure 5C), suggesting that the brain activity is capturing various forms of temporal dynamics.

Lastly, we repeat this frame shuffling analysis on a TSM MobileNetV2 (Lin et al., 2019; Sandler et al., 2018) and TimeSformer (Bertasius et al., 2021) architecture trained on Kinetics-400 (W. Kay et al., 2017) and HowTo100M (Miech et al., 2019) datasets, respectively, to see if the frame-shuffling results generalize to models of varying architectures, training diets, parameter counts, and task performance. We find that the TSM ResNet50 is the only model of the three that sees robust effects in most ROIs and DNN blocks (see Supplementary 1 Figure S10 for all results), implying that biological similarity of the model may be closely tied to the effects of frame shuffling.

Together this battery of analyses demonstrates that most visual regions captured meaningful temporal structure of the video stimuli, providing the necessary background for investigations into the processing of high-level dynamic concepts, such as event categories and actions. The use of our video-computable DNN model allows the extraction of feature spaces at specific stages of video processing, inviting precise inquiry into an ROI's function in visual event understanding.

BMD tracks the temporal dynamics unfolding within events

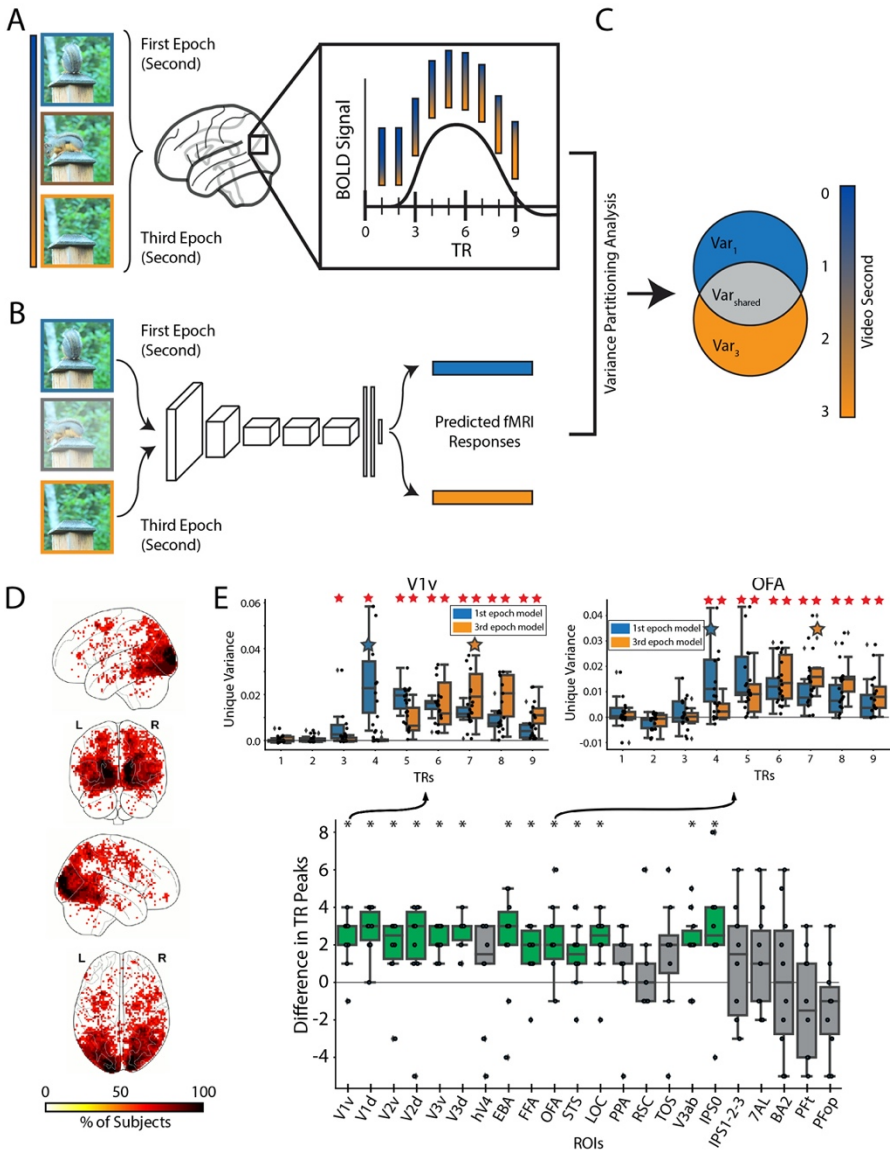

Figure 6: Encoding the temporal dynamics of the BOLD signal.

(A) TR estimated fMRI responses: We estimate the video-evoked brain response (beta values) of the first 9 TRs of the BOLD signal. **(B) DNN predicted fMRI responses:** We use an encoding model procedure to extract two sets of DNN activations, one from the first video second (first epoch) and the other from the third video second (third epoch). We then predict fMRI responses using the two sets. **(C) Variance partitioning analysis:** We calculate the unique variance explained by the first and third video epochs' predicted fMRI responses at each TR. **(D) Whole brain analysis:** Each voxel shows the percentage of subjects with a TR peak difference of 1 to 3 TRs at that specific voxel. Only significant voxels are plotted ($p < 0.05$, binomial test, FDR corrected). **(E) ROI analysis:** Upper ROI panels: We show the unique variance explained by the predicted fMRI responses of the first

(blue) and third (orange) video epoch at TRs 1-9 in representative ROIs V1v and OFA. Red asterisks along the top of the plot indicate significant unique variance greater than 0 ($p<0.05$, one-sample one-side t-test, FDR corrected across 9 TRs x 2 video epochs = 18 comparisons, assuming positive correlation). Blue/orange stars indicate the TR with the maximum subject averaged unique variance for the first/third video epochs, respectively. Main ROI Panel: We compute each subject's TR peak difference by subtracting the first video epoch's maximum TR from the third video epoch's maximum TR at each ROI. Green bars with asterisks indicate TR peak differences significantly greater than 0 ($p<0.05$, one-sample two-sided t-test). The box plot encompasses the first and third data quartiles and the median (horizontal line). The whiskers extend to the minimum and maximum values within 1.5 times the interquartile range, and values falling outside that range are considered outliers (denoted by a diamond). The overlaid points show the value at each observation (n=10 for all ROIs except TOS (n=8) RSC (n=9)).

Does the temporally sluggish BOLD signal capture any temporal information within events? One possibility is that the BOLD signal only captures a global representation of the event that has no temporal structure itself, akin to a time-less semantic label of "a person opens a door". Another possibility is that instead the BOLD signal captures delayed but temporally-resolved information, where different time points of the BOLD signal capture local snapshots of the changing event (Figure 6A).

To test the latter hypothesis, we extracted two sets of activations from a DNN, one set using only the first second (first epoch) of the videos and the other set using only the last second (third epoch) of the videos (Figure 6B). We utilized an encoding model procedure to measure the two sets' predictive performance at each of the first 9 seconds (corresponding to TRs in acquisition) (Figure 6A). This was done with a feed-forward DNN trained on object categorization from still images to avoid any confounds from temporal integration. We expected that if a voxel captures snapshots of the changing event, the best prediction accuracy of the first video epoch encoding model is at least one TR (one TR = 1s) earlier than the best prediction accuracy of the third video epoch encoding model. We used variance partitioning to identify the unique contribution of the first and third video epochs' predictions to the real fMRI responses (Figure 6C) (see Methods section "Encoding and Variance Partitioning Analysis Procedure" for details).

A whole-brain voxel-wise analysis (Figure 6D) revealed the percentage of subjects at each voxel that show a significant 1-3s delay between the best predicted time point (TR) using a video's first epoch and using a video's third epoch (only significant voxels are plotted). Results highlight significant temporal delays throughout the ventral and dorsal cortex, but most pronounced in the early visual cortex. An equivalent ROI-based analysis (Figure 6E) yielded a similar result pattern. ROIs in the visual brain (14

of the 22 total), mostly in the early visual and ventral stream, showed a significant timing difference (black asterisks) between the time points at which fMRI responses are most related to the contents of the first and the third epoch of video (Figure 6E, main bottom panel).

Together, these results support the hypothesis that early and late TRs of the BOLD signal better code early and late video snapshots, respectively. Previous work showed that the content and sequence order of distinct images presented in under a second can be reliably decoded in the BOLD signal (Wittkuhn & Schuck, 2021). Here we add to this line of work to show that BMD can differentiate between two visually and conceptually similar “images” of one second duration (i.e., first and third video second) separated by another highly similar “image” of one second duration (the second video second) with the most pronounced effects focused in the early visual cortex (Fairhall et al., 2014; Kiebel et al., 2008). This invites future research to use BMD’s temporally well-defined stimuli to explore how visual event information is integrated over shorter time periods, bridging an important gap to temporal integration studies of longform movies and BOLD encoding of rapidly presented stimuli.

Semantic metadata reveal strong similarity between sentence-level descriptions and visual brain activity

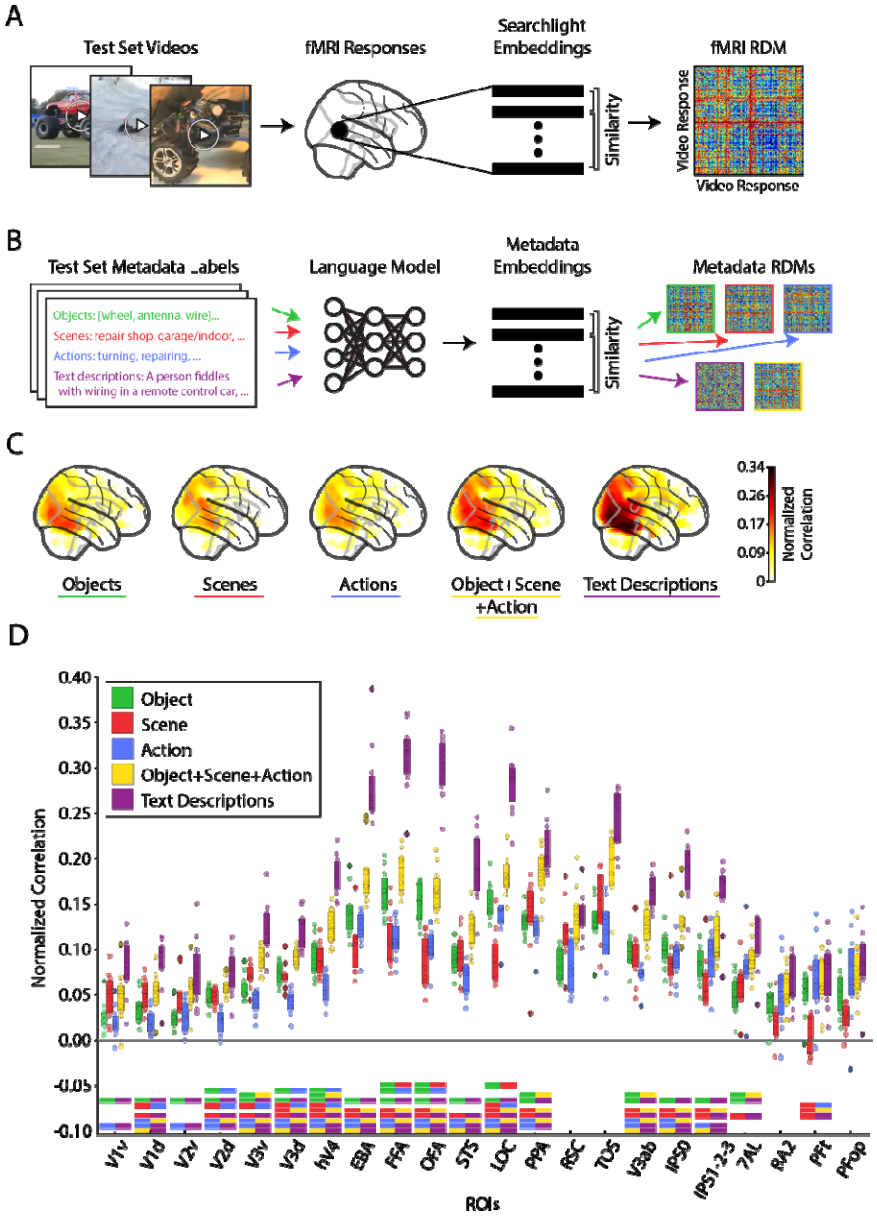

Figure 7: Metadata-driven analysis of the semantics of visual event encoding.

(A) Searchlight RDM methodology: The video-evoked brain responses to the 102 test set videos were extracted within a spherical searchlight to produce a vector embedding at each voxel. The distance (1-Pearson's R) between the searchlight brain embeddings were computed to create a representational dissimilarity matrix (RDM) of size n_stimuli x n_stimuli at each voxel in the brain. **(B) RSA with metadata methodology:** Video

metadata were fed into a language model to produce a vector embedding. The scene and actions metadata were fed into a FastText language model, and the “text caption” metadata was fed into a Sentence-BERT model to produce vector embeddings. Similar to the searchlight RDM computation, the pairwise distance (cosine distance) between the metadata embeddings were calculated to create an RDM for the object, scene, action, object+scene+action, and text description metadata. The RDMs are visualized here with rank-normalized values (0-100% of maximum value rank). **(C) Whole-brain correlation of metadata RDMs with searchlight-based RDMs:** The metadata RDM was correlated (Spearman’s R) with each searchlight-based RDM at each voxel in each subject. After statistical analysis (one-sample two-sided t-test against a null correlation of zero, FDR correction with q=0.05), dividing each voxel by the subject’s upper noise ceiling, and averaging across subjects, the results were plotted in a glass brain to show each metadata’s different pattern and magnitude of whole-brain responses. Only significant voxels are shown. **(D) ROI-based correlation:** The mean noise-normalized correlations are shown within each subject’s ROI. At each ROI, a one-way ANOVA test compared the mean noise normalized correlation between the 5 conditions (objects, scenes, actions, objects+scenes+actions, text descriptions). The ANOVA test was Bonferroni corrected with 22 ROIs (alpha = 0.05/22). If the ANOVA test was significant at an ROI, a Tukey’s Honestly Significant Difference test determined pairwise significance (family-wise error rate = 0.05; significance of pairs denoted by the dual-colored bars under each ROI). The box plot encompasses the first and third data quartiles and the median (horizontal line). The whiskers extend to the minimum and maximum values within 1.5 times the interquartile range, and values falling outside that range are considered outliers (denoted by a diamond). The overlaid points show the value at each observation (n=10 for all ROIs except TOS (n=8) RSC (n=9)).

Varying levels of semantic information content, from static objects and scenes (e.g., “duck”, “water”) to temporal actions (e.g., “swimming”) to complex relations between parts (e.g., “the duck is swimming on the water"), can describe a visual event. It is unclear how these varying levels of complexity and content are reflected in ROIs while viewing visual events, especially given that regions throughout the ventral visual, dorsal visual, and parietal cortices have all been implicated in processing temporal aspects of videos (Buccino, Binkofski, Fink, et al., 2004; Konen & Kastner, 2008; Lingnau & Downing, 2015; Silver & Kastner, 2009; Wurm & Caramazza, 2022) but also diverse feature preferences (Buccino, Binkofski, Fink, et al., 2004; Kanwisher, 2010; Kanwisher et al., 1997; Konen & Kastner, 2008; Lafer-Sousa et al., 2016; Ratan Murty et al., 2021). We leverage our metadata labels to construct five semantic video descriptions – objects, scenes, actions, objects+scenes+actions, and sentence text descriptions - to explore how semantics of varying information content and granularity are represented across brain regions during visual event perception. The word-level object, scene, and action labels provide low information descriptions of the video’s static (object and scene) and temporal (action) content. The sentence text descriptions not only describe the video’s object, scene, and temporal content but also how they relate. The combined objects+scenes+actions description concatenates the

individual object, scene, and action label information to serve as an intermediate between the word-level and sentence-level descriptions (see Methods section “Metadata” for details on the metadata collection).

We compute a neural representational dissimilarity matrix (RDM) at each voxel by computing the pairwise distances between vector embeddings obtained from a searchlight procedure (four voxel radius) (Kriegeskorte et al., 2006) (Figure 7A). We compute one RDM for each metadata category by calculating the pairwise distances between vector embeddings obtained by feeding the text-based labels through natural language processing models (FastText, Bojanowski et al., 2017, for the single-word object, scene, and action labels, and Sentence-BERT, Reimers & Gurevych, 2019, for sentence-level text descriptions) (Figure 7B). We use representational similarity analysis (RSA) (Kriegeskorte, 2008) to correlate the metadata representations (Figure 7C) with neural representations to measure how similarly the different metadata descriptions are reflected in brain activity of dynamic videos (see methods section “Metadata RSA Analysis Procedure” for more information).

Whole-brain results show the metadata labels differ in strength and pattern of similarity with brain activity (Figure 7C). ROI results (Figure 7D) compare the correlational strength between the five metadata descriptions at each ROI and depict a clear dominance of sentence text descriptions throughout cortex (purple bar). Taken individually, the object (green) and scene (red) description correlations are highest in their respective category-selective regions (LOC, EBA, FFA, OFA for object labels and PPA, RSC, and TOS for scene labels) as expected (Downing et al., 2001; Epstein & Kanwisher, 1998; Grill-Spector et al., 2001; Kanwisher, 2010; Kanwisher et al., 1997) (note that the object metadata can and often does describe people and animals present in the videos). Action labels (blue) correlate with ventral and dorsal visual regions (Lingnau & Downing, 2015; Wurm et al., 2017) more strongly than parietal regions (Hardwick et al., 2018). In all ROIs, the combined object+scene+action description (yellow) correlated with each ROI at a level between or equal to individual object, scene, and action descriptions and the text description.

Overall, the sentence text description results in stronger (or equally strong) correlation values than the other four semantic descriptions in all ROIs. Additionally, the three concatenated single-word labels (object+scene+action) results in stronger (or equally strong) correlation values than the individual single-word labels across all regions, even in category-selective regions. Both of these results are consistent with the idea that complex scene analysis, rather than simpler tasks such as object recognition, is the objective of the visual brain (Doerig et al., 2022). One might suspect that the

category-selective ventral regions would best correlate with their respective metadata (e.g., PPA, RSC, and TOS for scene metadata), reasoning that the text description and object+scene+action labels, while including the pertinent category information, contain mostly irrelevant and distracting extra-category content (Kosakowski et al., 2022; Ratan Murty et al., 2021, but see Bonner & Epstein, 2021).

BMD's reveal of visuo-semantic representations throughout cortex promises to better inform theories of the visual system's objective and action encoding. This methodology can easily be extended to externally collected metadata to interrogate other facets of visual event understanding.

Video memorability is reflected in high-level visual and parietal cortices

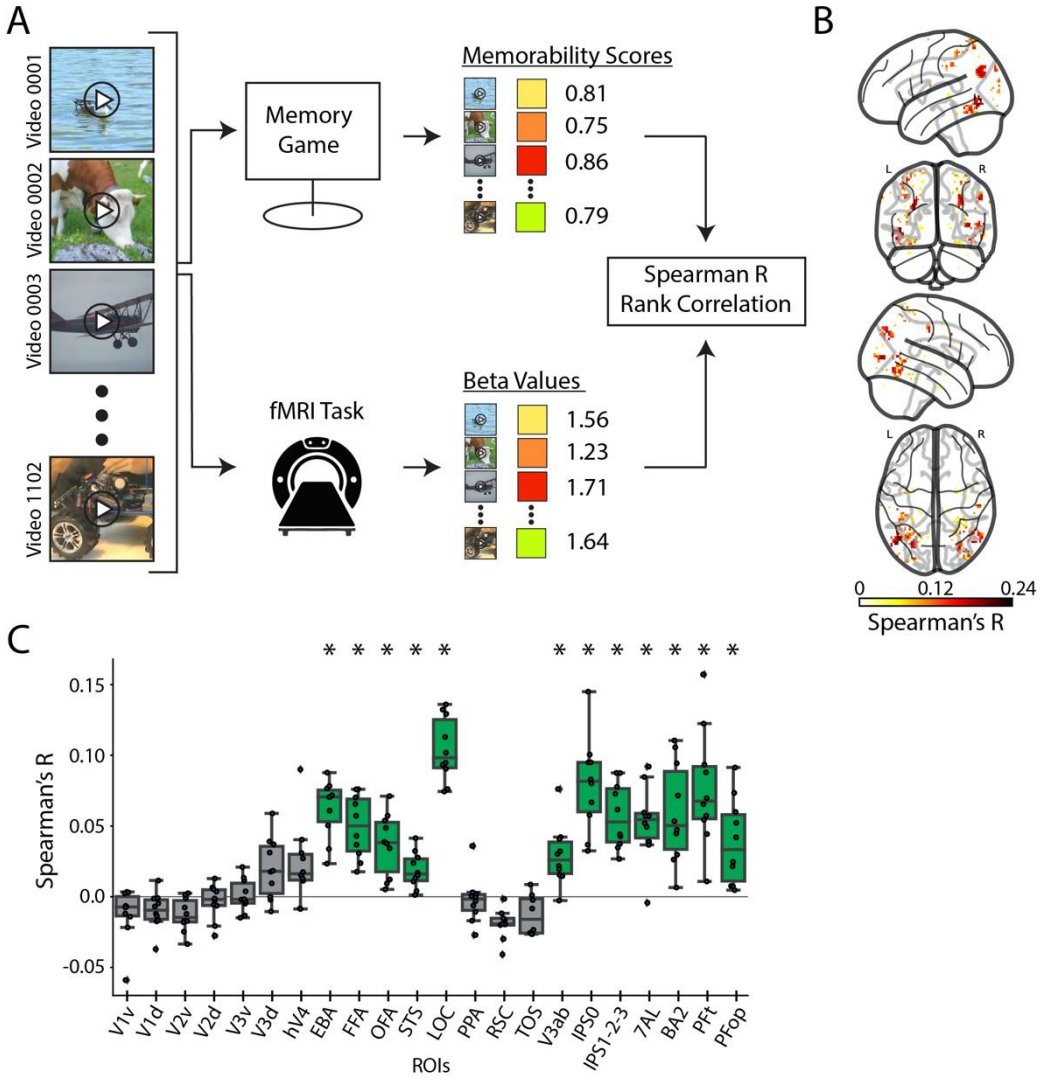

Figure 8: Correlates of video memorability in cortex

(A) Memorability analysis procedure: For each of the 1,102 video stimuli, we collect a memorability score (values between 0-1, collected via a behavioral memory game in an independent study by Newman et al., 2020) and brain response (beta value averaged across repetitions). We rank and correlate (Spearman's R) a vector of the stimulus memorability scores with a vector of the stimulus-evoked beta values at each voxel to obtain a correlation coefficient and corresponding p-value. **(B) Whole-brain memorability correlation results:** We show significant whole brain voxel-wise correlations (Spearman's R) between brain activity and memorability scores. Significant voxels are determined by a t-test across subjects (one-sample, one-sided) and FDR correction (q=0.05, positive correlation assumption). **(C) ROI memorability correlation results:** We show the Spearman correlation value between brain activity and memorability scores for each ROI. Significance is determined by a t-

test against a null population mean correlation of 0 (one sample, one-sided, Bonferroni corrected across 22 ROIs, $p < 0.05$). Significant ROIs are denoted with a green box and black asterisk. The box plot encompasses the first and third data quartiles and the median (horizontal line). The whiskers extend to the minimum and maximum values within 1.5 times the interquartile range, and values falling outside that range are considered outliers (denoted by a diamond). The overlaid points show the value at each observation (n=10 for all ROIs except TOS (n=8) RSC (n=9)).

Some images and videos are predisposed to be recognized better than others (Goetschalckx et al., 2018; J. Han et al., 2015; Isola et al., 2011; A. Khosla et al., 2015; Newman et al., 2020). However, fMRI studies have mostly focused on the memorability of still images (but see Han et al., 2015). Here, we evaluate where video memorability is reflected in the brain to give insights into the relation between stimuli perception and memory (Bainbridge, 2019; Bylinskii et al., 2022).

Under the hypothesis that stimuli with higher memorability scores elicit a greater magnitude of brain response (Bainbridge et al., 2017; Bainbridge & Rissman, 2018; Jaegle et al., 2019), we correlate a vector of video memorability scores with a vector of each voxel's corresponding brain responses (beta values) (Figure 8A) (see Methods section "Memorability Analysis Procedure" for details). Whole-brain (Figure 8B) and ROI-based (Figure 8C) analyses converged in revealing significant correlations in ventral visual, dorsal visual, and parietal cortex, all regions involved in video perception (green colored bars in Figure 8C). This result contrasts with previous work using images, where the effects were largely relegated to the ventral visual cortex associated with image perception (Bainbridge et al., 2017; Bainbridge & Rissman, 2018; Jaegle et al., 2019; Mohsenzadeh et al., 2019).

BMD allows the study of the neural correlates of memorability into the video domain, inviting further work into how visual perception and memory formation share computational resources (Cohen et al., 1997; Martin & Chao, 2001; Riou et al., 2011; Slotnick et al., 2012; Vermeulen et al., 2008; Weinberger, 2004) and bridging the study of static image memory with longform movie memory (Bainbridge & Baker, 2022; Furman et al., 2007; Hasson, Furman, et al., 2008).

Discussion

Disentangling the intricacies of human visual perception requires datasets varying in stimuli type (Boyle et al., 2020; Chang et al., 2019; M. Zhou et al., 2023), experimental design (Aliko et al., 2020; Horikawa & Kamitani, 2017; Seeliger et al., 2019), and neuroimaging acquisition efforts (Allen et al., 2022; Hanke

et al., 2016; Hebart et al., 2023) (see Supplementary Table S1 for a detailed comparison between currently available large-scale visual fMRI datasets). Here, we contribute the BOLD Moments Dataset (BMD), a resource of human fMRI brain responses to 1,102 annotated short videos to model human dynamic visual perception. The 3 second, naturalistic video stimuli balance the ecological validity of longform movies with the experimental control of still images to better model discrete events. BMD's ten subjects, high-repetition subset, and rich stimuli metadata accommodate uni- or multi-variate and single-subject or group-level modeling approaches, and its auxiliary resting state and dynamic functional localizer scans supplement insights into cortical activity. Together, BMD invites a uniquely comprehensive set of analyses into human dynamic visual perception. BMD can help bridge the research communities of computer science and computational neuroscience and reveal how perceptual and cognitive brain systems extract information from complex visual inputs.

For example, we use DNN models to interrogate perception of dynamic visual events. Our results demonstrating a correspondence in processing stages between video-computable DNN model blocks and cortical regions both extend previous studies that analyzed responses to still images (Kubilius et al., 2019; Schrimpf, Kubilius, Lee, et al., 2020; D. L. Yamins et al., 2013) and show we can accurately predict brain activity in the dorsal visual and parietal regions that are largely driven a stimulus's dynamic properties (Gazzola & Keysers, 2009; Heim et al., 2012; Konen & Kastner, 2008; R. Peeters et al., 2009; R. R. Peeters et al., 2013). Furthermore, we demonstrate that despite the sluggishness of the BOLD response, BMD allows tracking of visual information processing at the level of seconds (Hasson et al., 2004; Hasson, Yang, et al., 2008; Kiebel et al., 2008; Murray et al., 2014; Orlov & Zohary, 2018; Piasini et al., 2021). Thus, BMD presents a unique opportunity to probe brain function by leveraging existing modeling tools that capture facets of visual events, such as social expressions (Hu et al., 2022; Kahou et al., 2016; Tzirakis et al., 2018), action recognition (Carreira & Zisserman, 2017; Feichtenhofer et al., 2016; Monfort et al., 2022), integration of temporal features (Bertasius et al., 2021; Ji et al., 2013; Y. Wang et al., 2023), and object detection (Fan et al., 2021; Shafiee et al., 2017).

We have shown how our provided metadata can link brain activity to meaningful features of a visual event to yield theoretically relevant insights. For example, the higher correlations between brain activity and complex sentence descriptions over word-level labels (see Figure 7) suggest that the function of these brain regions extends beyond object recognition (Doerig et al., 2022). Additionally, the significant correlations between the crowd-sourced behavioral data on memorability and brain activity distributed in ventral visual, dorsal visual, and parietal regions (see Figure 8) imply the memorability effect occurs in regions implicated in perception.

BMD's unique combination of stimuli type, experimental design, and auxiliary scans makes it well-suited for a diverse set of analyses extending beyond those demonstrated here, such as video reconstruction (Chen et al., 2023; Kupershmidt et al., 2022; Nishimoto et al., 2011), video synthesis for cortical discovery (Gu et al., 2022; Luo et al., 2023), encoding and decoding models (K. Han et al., 2019; Naselaris et al., 2011), spatio-temporal integration (Fairhall et al., 2014; Hu et al., 2022), and action observation networks (Buccino, Vogt, et al., 2004; Shmuelof & Zohary, 2005; Spunt et al., 2011; Urgen et al., 2019). BMD is thus well positioned to probe the functionally diverse and widespread cortical network recruited during dynamic visual perception.

Methods

Participants

Ten healthy volunteers (6 female, mean age ± SD = 27.01 ± 3.96 years) with normal or corrected-to-normal vision participated in the experiment. All participants gave informed consent, were screened for MRI safety, and were compensated for their time. The experiment was conducted in accordance with the Declaration of Helsinki and approved by the local ethics committee (Institutional Review Board of Massachusetts Institute of Technology, approval code: 1510287948).

Stimuli

The stimulus set consisted of 1,102 videos in total. The videos were sampled from the Memento10k dataset (Newman et al., 2020), which is a subset of the Moments in Time dataset (Monfort et al., 2020) and Multi-Moments in Time dataset (Monfort et al., 2022). Each video was square-cropped and resized to 268x268 pixels. Videos had a duration of 3 seconds and frame rates ranging from 15-30 frames per second (mean = 28.3). The 1,102 videos were manually selected from the Memento10k dataset by two human observers to encompass videos that contained movement (i.e. no static content), were filmed in a natural context, and represented a wide selection of possible events a human might witness. Additional criteria were to be free of post-processing effects, textual overlays, excessive camera movements, blur, and objectionable or inappropriate content.

The 1,102 videos selected for the main experiment were split into "training" and "testing" sets; 102 videos were chosen for the testing set, and the remaining 1,000 videos formed the training set.

Specifically, the testing set videos were chosen randomly from the 1,102 videos, and then checked manually to ensure no semantic overlap, in terms of objects plus actions, occurred between any pair of testing set videos. If semantic overlap was found between a pair of videos as determined by an author, one of these videos was swapped with a video randomly selected from the pool of remaining videos and incorporated into the testing set. This was repeated until a semantically-diverse testing set was formed. The training and testing sets are only intended to differ in the number of repetitions shown to the participant. In this way, the BMD dataset contains a low repetition (3 repetitions) training set of 1,000 videos and a high repetition (10 repetitions) testing set of 102 videos to facilitate analyses dependent on either large number of stimuli or large number of repetitions. The training and testing sets are additionally intended to mirror training and testing sets common in machine learning applications, reflecting its potential use for model building and evaluation.

(f)MRI experimental design

The fMRI data collection procedure was as follows: subjects completed a total of 5 separate fMRI sessions on separate days. Session 1 consisted of structural scans, functional localizer runs, and functional resting state scans all interspersed (Figure 1B). Sessions 2-5 consisted of the main functional experimental runs where the subjects viewed the training and testing set videos (Figure 1C). Throughout the whole experiment for a given subject, each training set video was shown a total of 3 times and each testing set video was shown a total of 10 times, resulting in 3,000 and 1,020 trials in training and testing sets, respectively. We organized trials into experimental runs such that they either only contained testing set videos (test runs) or only contained training set videos (training runs).

Session 1

Functional localizers:

Subjects completed five functional localizer runs (Figure 1B). Subjects freely viewed colored, naturalistic videos (18s length, composed of 6 3-second videos) corresponding to one of five categories (faces, bodies, scenes, objects, and scrambled objects) in order to functionally localize each subject's category selective regions (Julian et al., 2012; Lafer-Sousa et al., 2016; Pitcher et al., 2011; Ratan Murty et al., 2021). Subjects viewed the videos freely and performed a one-back vigilance task to ensure attention to the task (videos are available with the dataset release). Each stimulus category included 48 individual stimuli. Each run consisted of 5 blocks of fixation baseline (null) and 20 blocks of stimulus presentation for a total of 25 blocks. Baseline blocks occurred every sixth block, with 5

stimulus blocks of each category presented in between baseline blocks in a randomized order. The duration of each video was 3 seconds, with each block lasting 18 seconds. Each category stimulus block included 5 unique videos chosen randomly from the 48 stimuli plus one one-back stimulus repetition. Subject accuracy on the one-back task was 0.941 +/- 0.011 (mean ± SD). Each run began and ended with a 18 second fixation (null) block. The end of the run contained an additional 19 seconds of gray screen (not modeled in the GLM) for a total duration of 268 volumes.

Resting state:

Resting state data was obtained across 5 runs (Figure 1B). Within each run, participants were instructed to keep their eyes closed, to not think of anything specific, but to remain awake. The duration of resting state runs is 212 volumes.

Session 2-5

Each of the 4 sessions after session 1 had identical structure. Since fixation against a dynamic video background is difficult, each session began with a 2.5 minute fixation training outside the scanner to provide the subjects with real-time feedback of any eye movements, voluntary or involuntary (Guzman-Martinez et al., 2009). In this fixation training, subjects viewed a black-and-white random dot display flickering in counter phase. With fixation, this display evoked the illusion of a uniform gray display. Breaking fixation with eye movements disrupts this illusion, and a vivid black-and-white random dot display is perceived. Subjects were instructed to keep their eyes fixated as to minimize the disruption of the illusion.

Main experiment:

Inside the scanner, videos were presented at the center of the screen subtending 5 degrees visual angle and overlaid with a central red fixation cross (0.52 degree visual angle). Subjects were instructed to focus on the fixation cross for the duration of the main experiment. The experiment consisted of both video presentation trials (occurring 75% of the time) and null trials (occurring 25% of the time). Both trial types were 4 seconds long. The video presentation trial consisted of a 3 sec video presentation followed by a 1 sec intertrial interval. The videos were presented in random order with the constraint of no consecutive repetition. The null trial consisted of the presentation of a gray screen. During the null trial the fixation cross turned darker for 1,000 ms and participants reported the change with a button press. Participant accuracy in this task was 0.964 +/- 0.014 (mean ± SD).

The testing and training set videos were presented within testing and training runs, respectively, where each test run consisted of 113 trials and each train run consisted of 100 trials. The order of these runs was randomized except for the constraint that two testing runs cannot be shown back-to-back in the same session in order to reduce potential memory effects caused by the same stimulus being presented within a short period of time. Each session contained 3 test runs and 10 training runs and lasted approximately 100 minutes. Each run began with 4 seconds of fixation and ended with 13 seconds (testing runs) or 12.5 seconds (training runs) of fixation. The 0.5 second difference of the duration of the ending fixation period between the testing and training runs is due to their different number of trials (113 and 100) and the acquisition TR of 1.75 seconds not being an even factor the 4 second trial length. In total, testing and training runs consisted of 268 and 238 volumes, respectively.

Metadata

Visual events consist of complex combinations of objects, locations, actions, and more. To capture the many dimensions of visual events, we characterized each video with a set of seven metadata categories: object labels, scene labels, action labels, text descriptions, a spoken transcription, a memorability score, and an index of memorability decay rate (Figure 1A). Five object, scene, action, and text description labels were collected for each stimulus to ensure comprehensive coverage (the five annotators' labels reflect their unique interpretations) and form a group consensus (the five annotators' labels can be used to converge on a single label). Labels for each metadata category were collected in different human crowd-sourced experiments. An annotator was allowed to label up to 20 different videos but no more than one label per video to encourage a diverse sampling of human annotators within videos and throughout the stimulus set. While crowd-source workers were not restricted from participating in multiple metadata experiments, this was unlikely due to the experiments being collected at different times and the large population pool from which the crowd-source workers were drawn.

Object labels

For each video, we obtained at least 5 sets of up to three different object labels in a human crowd-sourced experiment on Prolific. Each annotator was instructed to select up to three different object labels visible in the video. They selected at least one object label, and if they believed no more objects were present in the video, they were allowed to select a "No more objects in the video" option up to two times. This option thus encouraged accurate labels and carried information on the density of objects in the video. Each object label was one of 1,854 possible labels from the THINGS dataset (Hebart et al., 2023) to encourage overlap with computational neuroscience work and leverage the additional THINGS

metadata on each label (e.g., animacy, size, indoor/outdoor). The object label selections can be different or the same across annotators. One author manually reviewed the labels to ensure the labels assigned to the video were sensical (i.e. participants were not choosing labels at random).

Scene labels

For each video, we obtained at least 5 scene labels using a human crowd-sourced experiment on Prolific. Each of the five different annotators were instructed to select a scene label that best describes the scene of the video. All scene labels came from the Places365 dataset (B. Zhou et al., 2018) for its broad scene coverage and overlap with computer vision resources. The scene label selections can be different or the same across videos. One author manually reviewed the labels to ensure the labels assigned to the video were sensical (i.e. participants were not choosing labels at random).

Action labels

The 5 action labels were selected by workers on the crowd-sourcing platform Prolific. We restricted the possible action labels to be from one of 292 possible action labels that broadly encompass meaningful human actions (Monfort et al., 2022). The participant viewed these 292 possible action labels, watched the video, and selected one action label that best described the video. Each of the 5 action labels per video were produced by different participants. Given that an action, by definition, unfolds across time and the BMD stimuli have a short 3 second duration, the majority of the videos contain one primary action. Additionally, the limited number of stimulus repetitions and the central fixation further limits the fMRI scanner participant's ability to perceive a video's small, inconsequential actions if they are present. In the edge case that a complex video captures multiple salient actions, the five annotations can capture these multiple actions. Note that annotators labeled up to 3 objects in a video (described above) because objects, unlike actions, have clear visual boundaries, often occur in multiple instances in a video, and have no temporal dimension. Two authors manually reviewed the labels to ensure the labels assigned to the video were sensical (i.e. participants were not choosing labels at random).

Text descriptions

We obtained five high-level descriptions of the content of each video in the form of written descriptions. The five text descriptions were human-generated from participants on the crowd-sourcing platform Amazon Mechanical Turk (AMT). Their task was to watch the video and type a 10-15 word caption in complete English sentences. This instruction was not used as acceptance criteria. Each of the 5 text descriptions were produced by a different human participant, and each participant was allowed to

annotate multiple videos. The authors manually checked the text descriptions to ensure the text descriptions pertained to the video (i.e. participants followed instructions) and to correct obvious typos.

Spoken transcriptions

We obtained one spoken description per video to capture emotional and descriptive nuances typically conveyed in speech but not present in typing. The spoken description was collected via human participants on AMT, as in Monfort et al., 2021. Participants were instructed to watch the video and verbally describe it. They were given no instructions pertaining to the length of the description. We record the audio file and use Google's speech-to-text transcription to generate a text transcription. The transcription was manually checked to ensure it pertained to the video (i.e. participants followed instructions) and to correct obvious typos. We release the text transcription but not the original audio file for privacy purposes.

Memorability score and decay rate

The memorability score and decay rate were measured by Newman et al., 2020, where AMT human participants played a video memory game. The game consisted of a continuous video stream where the participant pressed the spacebar upon seeing a repeated video. Repeated videos were presented at various delays, from 30 seconds to ten minutes. The participant's responses were then used to calculate a video's memorability score from 0 (no recall) to 1 (perfect recall) and memorability decay rate from 0 (no decay) to -inf (instantaneous decay). Specifically, a video's memorability score is the fraction of correct identifications in the memory game normalized to a lag of 80 videos in between consecutive presentations. A video's memorability decay rate describes how a video's memorability score changes over different lags. The memorability decay rate was regressed from the output of SemanticMemNet (Newman et al., 2020), an Inflated 3D (I3D) Convolutional Neural Network trained to predict a video's memorability score and text caption from the video frames and optical flow. A video's memorability score (m) and decay rate (α) can be used to predict its memorability at any time lag t (the number of videos in between the first and second presentation) according to the following equation (Newman et al., 2020):

$$m_t = m_T + \alpha(\mathrm{t} - \mathrm{T})$$

where T is a set lag of 80. The memorability scores given in the stimuli metadata were computed at a lag of t=80.

Note that Newman and colleagues (2020) experimentally observed memorability decay rates above 0. Since a positive memorability decay rate may reflect a currently unknown stimulus or cognitive feature, we preserve them in BMD's stimuli metadata. One may wish to clip them to maximum value of 0 depending on the analysis.

Motion energy features

We provide video-computable motion features of each video using a motion energy model (Adelson & Bergen, 1985; Nishimoto et al., 2011; Watson & Ahumada, 1985). The motion energy model uses a set of spatial and temporal Gabor filters to extract a video's motion and direction. We use these features to predict brain activity in motion-selective regions of MT (Born & Bradley, 2005; Nishimoto & Gallant, 2011), hV4 (Kamitani & Tong, 2006; Roe et al., 2012), V3AB (Konen & Kastner, 2008; A. T. Smith et al., 1998), and IPS0 (Konen & Kastner, 2008). This analysis uses version B of the dataset and is detailed in Supplementary 2 section 5, "Motion energy features computation and encoding model."

fMRI data preprocessing and analysis

fMRI data acquisition

The MRI data were acquired with a 3T Trio Siemens scanner using a 32-channel head coil. During the experimental runs, T2*-weighted gradient-echo echo-planar images (EPI) were collected (TR = 1750 ms, TE = 30 ms, flip angle = 71°, FOV read = 190 mm, FOV phase = 100%, bandwidth = 2268 Hz/Px, resolution = 2.5 × 2.5 × 2.5 mm, slice gap = 10%, slices = 54, multi-band acceleration factor = 2, ascending interleaved acquisition). Additionally, a T1-weighted image (TR = 1900 ms, TE = 2.52 ms, flip angle = 9°, FOV read = 256mm, FOV phase = 100%, bandwidth = 170 Hz/px, resolution = 1.0 x 1.0 x 1.0 mm, slices = 176 sagittal slices, multi-slice mode= single shot, ascending) and T2-weighted image (TR = 7970ms, TE = 120 ms, flip angle = 90°, FOV read = 256 mm, FOV phase = 100%, bandwidth = 362 Hz/Px, resolution = 1.0 × 1.0 × 1.1 mm, slice gap = 10%, slices = 128, multi-slice mode = interleaved, ascending) were obtained as high-resolution anatomical references. We acquired resting state and functional localizer data using acquisition parameters identical to the main experimental runs. Dual echo fieldmaps (TR = 636 ms, TE1 = 5.72 ms, TE2 = 8.18 ms, flip angle = 60°, FOV read = 190 mm, FOV phase = 100%, bandwidth = 260 Hz/Px, resolution = 2.5 × 2.5 × 2.5 mm, slice gap = 10%, slices = 54, ascending interleaved acquisition) were acquired at the beginning of every session to post-hoc correct for spatial distortion of functional scans induced by magnetic field inhomogeneities.

Preprocessing

All MRI data was first converted to BIDS format (K. J. Gorgolewski et al., 2016), and the T1-weighted structural scans were anonymized using PyDeface (https://github.com/poldracklab/pydeface) before public release. All data from all sessions was then preprocessed using the standardized fMRIPrep preprocessing pipeline. For all results unless stated otherwise, we use the standard MNI152NLin2009cAsym volumetric space for its frequent use in other work and preservation of natural left-right asymmetries in the brain. As recommended by fMRIPrep to increase transparency and reproducibility in MRI preprocessing, we copy their generated preprocessing text in its entirety below:

“”

Results included in this manuscript come from preprocessing performed using *fMRIPrep* 20.2.1 (Esteban et al., 2019, 2022; RRID:SCR_016216), which is based on *Nipype* 1.5.1 (Esteban, Oscar et al., 2022; K. Gorgolewski et al., 2011; RRID:SCR_002502).

Anatomical data preprocessing

A total of 1 T1-weighted (T1w) images were found within the input BIDS dataset. The T1-weighted (T1w) image was corrected for intensity non-uniformity (INU) with N4BiasFieldCorrection (Tustison et al., 2010), distributed with ANTs 2.3.3 (Avants et al., 2008; RRID:SCR_004757), and used as T1w-reference throughout the workflow. The T1w-reference was then skull-stripped with a *Nipype* implementation of the antsBrainExtraction.sh workflow (from ANTs), using OASIS30ANTs as target template. Brain tissue segmentation of cerebrospinal fluid (CSF), white-matter (WM) and gray-matter (GM) was performed on the brain-extracted T1w using fast (FSL 5.0.9, RRID:SCR_002823, Zhang et al., 2001). Brain surfaces were reconstructed using recon-all (FreeSurfer 6.0.1, RRID:SCR_001847, Dale et al., 1999), and the brain mask estimated previously was refined with a custom variation of the method to reconcile ANTs-derived and FreeSurfer-derived segmentations of the cortical gray-matter of Mindboggle (RRID:SCR_002438, Klein et al., 2017). Volume-based spatial normalization to one standard space (MNI152NLin2009cAsym) was performed through nonlinear registration with antsRegistration (ANTs 2.3.3), using brain-extracted versions of both T1w reference and the T1w template. The following template was selected for spatial normalization: *ICBM 152 Nonlinear Asymmetrical template version 2009c* [Fonov et al., 2009, RRID:SCR_008796; TemplateFlow ID: MNI152NLin2009cAsym].

Functional data preprocessing

For each of the 62 BOLD runs found per subject (across all tasks and sessions), the following preprocessing was performed. First, a reference volume and its skull-stripped version were generated using a custom methodology of *fMRIPrep*. A B0-nonuniformity map (or *fieldmap*) was estimated based on a phase-difference map calculated with a dual-echo GRE (gradient-recall echo) sequence, processed with a custom workflow of *SDCFlows* inspired by the epidewarp.fsl script and further improvements in HCP Pipelines (Glasser et al., 2013). The *fieldmap* was then co-registered to the target EPI (echo-planar imaging) reference run and converted to a displacements field map (amenable to registration tools such as ANTs) with FSL's fugue and other *SDCflows* tools. Based on the estimated susceptibility distortion, a corrected EPI (echo-planar imaging) reference was calculated for a more accurate co-registration with the anatomical reference. The BOLD reference was then co-registered to the T1w reference using bbregister (FreeSurfer) which implements boundary-based registration (Greve & Fischl, 2009). Co-registration was configured with six degrees of freedom. Head-motion parameters with respect to the BOLD reference (transformation matrices, and six corresponding rotation and translation parameters) are estimated before any spatiotemporal filtering using mcflirt (FSL 5.0.9, Jenkinson et al., 2002). BOLD runs were slice-time corrected using 3dTshift from AFNI 20160207 (Cox & Hyde, 1997, RRID:SCR_005927). The BOLD time-series (including slice-timing correction when applied) were resampled onto their original, native space by applying a single, composite transform to correct for head-motion and susceptibility distortions. These resampled BOLD time-series will be referred to as *preprocessed BOLD in original space*, or just *preprocessed BOLD*. The BOLD time-series were resampled into standard space, generating a *preprocessed BOLD run in MNI152NLin2009cAsym space*. First, a reference volume and its skull-stripped version were generated using a custom methodology of *fMRIPrep*. Several confounding time-series were calculated based on the *preprocessed BOLD*: framewise displacement (FD), DVARS and three region-wise global signals. FD was computed using two formulations following Power (absolute sum of relative motions, Power et al., 2014) and Jenkinson (relative root mean square displacement between affines, Jenkinson et al., 2002). FD and DVARS are calculated for each functional run, both using their implementations in *Nipype* (following the definitions by Power et al., 2014). The three global signals are extracted within the CSF,

the WM, and the whole-brain masks. Additionally, a set of physiological regressors were extracted to allow for component-based noise correction (*CompCor*, Behzadi et al., 2007). Principal components are estimated after high-pass filtering the *preprocessed BOLD* time-series (using a discrete cosine filter with 128s cut-off) for the two *CompCor* variants: temporal (tCompCor) and anatomical (aCompCor). tCompCor components are then calculated from the top 2% variable voxels within the brain mask. For aCompCor, three probabilistic masks (CSF, WM and combined CSF+WM) are generated in anatomical space. The implementation differs from that of Behzadi et al. in that instead of eroding the masks by 2 pixels on BOLD space, the aCompCor masks are subtracted a mask of pixels that likely contain a volume fraction of GM. This mask is obtained by dilating a GM mask extracted from the FreeSurfer's *aseg* segmentation, and it ensures components are not extracted from voxels containing a minimal fraction of GM. Finally, these masks are resampled into BOLD space and binarized by thresholding at 0.99 (as in the original implementation). Components are also calculated separately within the WM and CSF masks. For each CompCor decomposition, the *k* components with the largest singular values are retained, such that the retained components' time series are sufficient to explain 50 percent of variance across the nuisance mask (CSF, WM, combined, or temporal). The remaining components are dropped from consideration. The head-motion estimates calculated in the correction step were also placed within the corresponding confounds file. The confound time series derived from head motion estimates and global signals were expanded with the inclusion of temporal derivatives and quadratic terms for each (Satterthwaite et al., 2013). Frames that exceeded a threshold of 0.5 mm FD or 1.5 standardised DVARS were annotated as motion outliers. All resamplings can be performed with *a single interpolation step* by composing all the pertinent transformations (i.e. head-motion transform matrices, susceptibility distortion correction when available, and co-registrations to anatomical and output spaces). Gridded (volumetric) resamplings were performed using antsApplyTransforms (ANTs), configured with Lanczos interpolation to minimize the smoothing effects of other kernels (Lanczos, 1964). Non-gridded (surface) resamplings were performed using mri_vol2surf (FreeSurfer).

Many internal operations of *fMRIPrep* use *Nilearn* 0.6.2 (Abraham et al., 2014, RRID:SCR_001362), mostly within the functional processing workflow. For more details

of the pipeline, see the section corresponding to workflows in *fMRIPrep*'s documentation.

“”

Next, the fMRI data from the main experimental runs (sessions 2-5) underwent temporal resampling. In detail, we resampled each voxel's time series using cubic interpolation to change the acquisition TR of 1.75 second to a new TR of 1 second. The 1.75 second acquisition TR combined with a 4 second trial length allowed for a dense sampling of the BOLD response relative to stimulus onset and thereby a good estimate of the BOLD response shape. However, as on every trial the timing of MR acquisition with respect to the trial was different, analysis of the BOLD response time-locked to the onset of each trial was cumbersome. The interpolation to a 1 second TR achieved a time series with regular sampling of the BOLD response relative to stimulus onset, enabling analysis time-locked to the onset of the videos.

Functional localizer scans were subsequently smoothed with a 9 mm full width half maximum of the Gaussian kernel. The main experimental functional runs use unsmoothed data.

General linear model

Functional localizer

To model the hemodynamic response to the localizer videos, the preprocessed fMRI data, video, and fixation baseline onsets and durations were included in a general linear model (GLM). The fixation and five category (faces, objects, scenes, bodies, and scrambled objects) blocks were included as regressors of interest. Motion and run regressors were included as regressors of no interest. All regressors were convolved with a hemodynamic response function (canonical HRF) to calculate beta estimates.

Main experiment

We modeled the BOLD-signal of each voxel in the preprocessed fMRI data of each participant as a weighted combination of simple Finite Impulse Response (FIR) basis functions. We modeled the BOLD response with respect to each video onset from 1 to 9 seconds in 1 second steps (corresponding to the resolution of the resampled time series). The stimulus presentation trials, but not the null response trials, were included as regressors of interest. Within this time interval the voxel-wise time course of activation was high-pass filtered (removing signal with f<1/128 Hz) and serial correlations due to

aliased biorhythms or unmodelled neuronal activity were accounted for using an autoregressive AR(1) model.

Using FIR in the way described above, we modeled every trial in the experimental run of each session, simultaneously capturing the spatial variability and the temporal evolution of brain responses underlying visual event understanding. For every session we generated separate FIR models for training and testing sets. Overall, for each video condition in the testing set we extracted 10 (repetitions) x 9 (seconds) estimated beta values, and for each video condition in the training set we extracted 3 (repetitions) x 9 (seconds) beta estimates.

ROI definitions of early visual cortex and ventral visual stream

We computed five t-contrasts (FWE corrected at p=0.05) per subject based on the beta values from the localizer experiment to quantify category-specific voxel activations. The t-contrasts were used to localize voxels in early visual regions (objects & scrambled objects > baseline; V1v, V1d, V2v, V2d, V3v, V3d, hV4), a body-selective region (bodies > objects; EBA), an object-selective region (objects > scrambled objects; LOC), face-selective regions (faces > objects; FFA, OFA, STS), and scene-selective regions (scenes > objects; PPA, RSC, TOS). Together, these 15 ROIs cover brain regions along the ventral visual pathway thought to transform low-level visual features into complex, semantic representations useful for object recognition.

The ROI masks for each subject were defined using the subject's t-contrast maps. Voxels outside the anatomical region of interest were manually set to 0, and ROIs were limited to 1,000 voxels in size. In the event that two or more ROIs overlapped, we used a probability map computed from all ten subjects' t-contrasts to assign each overlapping voxel to the ROI it most likely belonged to. In detail:

We first created a probabilistic map for each of the 8 category-selective and 7 early visual ROIs, using data from all 10 subjects. Each subject's five t-contrast maps from the localizer experiment were family-wise error (FWE) corrected at p = 0.05. We then added each subject's binarized version of the FWE corrected t-contrast maps (1 if the voxel passed FWE correction, 0 if not) and divided by the total number of subjects. This resulted in 5 probabilistic maps corresponding to each of the 5 t-contrasts, where each voxel had a value from 0 to 1 in steps of 0.1 (0, 0.1, 0.2, …, 1.0) representing the decimal percentage of the number of subjects where that voxel was significant. We used the "objects & scrambled objects > baseline" t-contrast probability map to define a probabilistic map for each of the 7

early visual ROIs by masking the t-contrast probability map with the early visual masks defined by L. Wang et al., 2015. We used the remaining 4 category-selective t-contrast probabilistic maps to define a probabilistic map for each of the 8 category-selective ROIs by visually inspecting the appropriate t-contrast probability map (e.g. inspecting the scenes > objects t-contrast probability map to define the scene-selective ROIs) and manually setting the voxels clearly outside the region of interest to 0. This way we defined a total of 15 separate probabilistic maps.

Next, we defined 15 ROI masks for each subject. For each subject, we masked their own appropriate FWE corrected t-contrast map with the corresponding binarized t-contrast probability map. The ROIs were limited to the top 1,000 voxels. This process resulted in fifteen subject-specific ROI masks.

Lastly, we assigned any overlapping voxels to a single ROI. In the event two or more voxels overlapped, we assigned the overlapping voxels to a single ROI based on the ROI probability map in step 1. Exemplified for two ROIs A and B, if ROI A and ROI B overlapped on voxel x, we indexed into ROI A's probability map at voxel x and ROI B's probability map at voxel x. We assigned voxel x to the ROI with the higher probability. If ROI A and B have equal probabilities at voxel x, we grew a patch by one voxel along each dimension and compared the ROIs' mean probabilities within that patch. This process was repeated until there was no longer a tie between ROIs and each ROI was non-overlapping.

Subject 6 did not show any responses in TOS and RSC, and subject 7 did not show any responses in RSC. The ROI masks were used in subsequent analyses to extract active voxels within a subject's ROI during the main experiment.

ROI definitions of dorsal visual stream

We additionally defined dorsal regions of interest using anatomical landmarks, since our functional localizer was designed to define early visual and ventral category-selective regions only. All dorsal regions were defined the same way for all subjects. Specifically, we use the maximum probability map in L. Wang and colleagues (2015) to anatomically define area V3ab (grouping areas V3a and V3b), IPS0, and IPS1-2-3 (grouping areas IPS1, IPS2, and IPS3). We additionally define 7AL, BA2, PFt, and PFop in more superior regions of the dorsal stream using the atlas described in Glasser et al., 2016. Atlas versions registered in MNI volume space were downloaded and resampled to BMD's functional

voxel resolution with nearest neighbor interpolation using SPM 12's "reslice" function. The resampled atlases were then used to index ROIs in the BMD volume data.

Univariate split-half reliability analysis

To select voxels with high signal-to-noise ratio, we defined a selection criteria based on split-half trial reliability (Lage-Castellanos et al., 2019; Schrimpf, Kubilius, Hong, et al., 2020). Our assumption behind the criteria was that the voxel responses on different trials corresponding to the same video should be more correlated with each other than voxel responses to different videos. As the voxel response to a video, we used the beta estimates from the FIR model averaged over time points TR 5-9 (representing the peak of a typical BOLD signal). We then z-scored the voxel responses across the videos.

We divided the voxel responses of the testing set stimuli trials (n=10 trials) into two equal splits and calculated the Pearson correlation (ρ) between the splits. The split-half reliability was calculated using the Spearman-Brown formula ($2\rho/(1+\rho)$), where the maximum reliability is 1. We calculated the split-half reliability for all possible combinations of splits and used the mean reliability as the reliability for that voxel.

To assess if reliability is better than chance, we first estimated chance-level reliability. For each voxel, we calculated the split-half reliability for all possible combinations of splits while randomly permuting the video indices for one of the two splits. This process was repeated 100 times with a different video index permutation each time. This procedure resulted in 100 random reliability values for each voxel, which was used to calculate a p-value. The voxels that satisfy our reliability criteria ($p < 0.05$) are referred to as "reliable voxels". Due to the testing set's high number of repetitions, the reliable voxels were defined using data from the testing set runs.

Multivariate searchlight-based reliability analysis

We computed the upper and lower noise ceilings at each voxel in the whole brain using the testing set and present subject-specific and subject-averaged results (Supplementary Figure S3). For each subject separately, the raw FIR beta estimates to the video stimuli at each voxel were z-scored across stimuli, averaged across TRs 5-9, and averaged across stimuli repetitions to result in a (n_stimuli x 1) vector of beta values. A spherical searchlight with radius 4 voxels was defined and centered on a voxel *v*. The (n_stimuli x 1) vector of beta values over all voxels contained within the searchlight sphere compose a (n_stimuli x n_voxel) matrix. The 1-Pearson correlation of the matrix resulted in a single

Representational Dissimilarity Matrix (RDM) of size (n_stimuli x n_stimuli) for the voxel *v*. The searchlight then centered on the next voxel, and the procedure was repeated until an RDM was calculated for each voxel. This procedure was repeated for each of the 10 subjects.

The upper and lower noise ceilings of the multivariate searchlight results were computed (Carlin & Kriegeskorte, 2017; Nili et al., 2014). To compute the upper noise ceiling at voxel *v*, the voxel's searchlight-computed RDM from one subject was correlated (Spearman's R) with the 10-subject group averaged RDM. The average of each subject's correlation to the 10-subject group averaged RDM is the upper noise ceiling for voxel *v*. This process was repeated over all voxels to result in the upper noise ceiling values throughout the whole brain displayed in Figure 3B. The upper noise ceiling estimates the highest correlation that a model can be expected to obtain given the noise in the data.

To compute the lower noise ceiling at voxel *v*, the voxel's searchlight-computed RDM from one left-out subject was correlated (Spearman's R) with the remaining 9-subject group averaged RDM. The average of each left-out subject's correlation to the corresponding 9-subject group averaged RDM is the lower noise ceiling for voxel *v*. This process was repeated over all voxels to estimate the lower noise ceiling values throughout the whole brain.

Action recognition TSM ResNet50 model training

The model adopts the architecture of a Temporal Shift Module (TSM) (Lin et al., 2019), with ResNet50 as the backbone network. We trained our model on the M4 (Multi-Moments minus Memento) training dataset for 120 epochs by using LSEP (log-sum-exp pairwise) loss (Monfort et al., 2022). LSEP loss was first proposed in (Li et al., 2017) and modified in (Monfort et al., 2022) as an appropriate loss function to train on multi-label and class imbalanced datasets, such as actions. The M4 training dataset consists of 1,012,169 videos which are in the Multi-Moments in Time dataset but not in the Memento dataset to ensure no overlap with the 1,102 BMD stimuli. Our model was initialized with the weights of the ResNet50 trained on ImageNet-1k dataset. We chose the model hyperparameters to closely follow those used in Lin and colleagues (2019). Specifically, during the training phase, our model split the input video into 8 segments and sampled 1 frame from each segment. We used SGD optimizer to optimize our model. The learning rate followed the cosine learning rate schedule and was initialized as 0.02. The weight decay was set to be 0.0001 and the batch size 128. The model achieved a precision-at-one score 0.593, a precision-at-five score of 0.829, and a mAP score of 0.636 (loss of 2.75054).

Model training took 3 months on 16 V100 GPUs and is available here: https://github.com/pbw-Berwin/M4-pretrained.

DNN block to cortex correspondence procedure

We used an encoding model procedure to quantify the correspondence between DNN Blocks and regions of cortex (Kriegeskorte & Douglas, 2019; Naselaris et al., 2011; Schrimpf, Kubilius, Hong, et al., 2020). For the DNN, we train a Temporal Shift Module (TSM) network (Lin et al., 2019) with a ResNet50 backbone on the M4 dataset (Multi-Moments in Time Minus Memento10k). In this way, we achieved a model that consecutively processed video frames (as opposed to frame averages), incorporated biologically necessary recurrent computations, and learned to perform a video-based task (i.e. action recognition) from the same set of short, natural videos of which we sample the BMD stimuli.

We ran inference on the TSM ResNet50 model using the 1,102 videos used in the fMRI experiment and extracted the activations for each video. The activations for a given block were extracted after the nonlinearity. We then used an encoding model procedure. In detail, we standardized the DNN activations (using the mean and standard deviation of the training videos) and performed principal component analysis (PCA) using the top 100 components of the DNN activations to ensure fair comparison of activations of different embedding sizes. The PCA procedure was fit on the 1,000 training set video activations and applied to the 102 testing set video activations. For each voxel *v*, we fit a linear model from the training set DNN activations (size (n_training_videos x n_PCA_components)) to the training set fMRI responses (beta values z-scored across stimuli and averaged over TRs 5-9) averaged over the 3 trial repetitions (size (n_training_videos x n_voxels)). We then predicted testing set voxel responses (size (n_testing_videos x n_voxels)) by applying the linear fit on the testing set DNN activations (size (n_testing_videos x n_PCA_components)). We evaluated the performance of the prediction by correlating (Pearson) the predicted testing set voxel responses with the true testing set fMRI responses of each of the 10 testing set repetitions (size (n_testing x n_voxels)). The final performance of the prediction is the average correlation of the 10 repetitions. The noise-normalized correlation is this 10-repetition average Pearson correlation divided by the voxel's split-half correlation value (the pearson correlation value before Spearman-Brown). We only predicted the values of the voxels that met the split-half reliability criteria, as described in the "Univariate Split-Half Reliability Analysis" methods section above, in order to model meaningful signal.

In this way, we obtained an encoding model accuracy (correlation) at each voxel in the whole brain for each of the four ResNet50 Blocks and for each subject. We averaged the noise normalized correlation

at each voxel across subjects for each of the four Blocks and displayed the results in a whole-brain volume (Figure 4B).

We then computed the difference in encoding accuracy between Block 1 and Block 4 at each of the 22 ROIs to determine if a region's brain responses were predicted significantly better by activations of early or late DNN Blocks. For each subject, we computed both Block 1 and Block 4's average noise-normalized correlation (encoding accuracy) within each ROI and took the difference (Block 1 - Block 4). We then performed a t-test (one-sample, two-sided) against a null hypothesis of zero correlation and corrected for multiple comparisons across ROIs (Bonferroni, $p < 0.05$ with $n=22$). We plotted the subject-averaged Block 1 - Block 4 noise-normalized correlation differences and denoted the significant ROIs with an asterisk and a color corresponding to significance with Block 1 (blue) and BLock 4 (red) (Figure 4C). See supplementary Figure S7 to see each Block's encoding accuracy at each ROI.

We note that while the temporal dimension in videos invites exciting modeling opportunities, it also adds complexities in the fMRI data that may make modeling difficult. For example, regions may reconfigure their roles over the duration of the video or integrate features over time in a manner that cannot be resolved with fMRI. Thus, the extent that models can predict fMRI brain responses to videos may be inherently limited by the temporal resolution of fMRI and best be modeled alongside millisecond-level temporal resolution neuroimaging data (M/EEG).

Shuffling analysis to determine importance of temporal order

To determine whether the temporal order of visual information is preserved in fMRI responses of the human visual system, we compared the encoding performance of the TSM model trained on the M4 dataset (Multi-Moments in Time Minus Memento10k) with preserved order of visual information with the model with a randomly shuffled order of visual information. TSM requires eight frames as the input to the model. These frames were sampled uniformly. In the original (unshuffled) order, the order of frames was preserved for all the videos. In the shuffled case, the indices of the frames were shuffled randomly, and then for all the videos, the same order of shuffled indices was used to create the input to the TSM model. We used ten such random shuffles of indices to introduce more randomness.

We first extracted the activations from the four blocks of the TSM model for the unshuffled case and each shuffled case. Then we performed PCA to extract the top 100 components of each block's

activations to reduce the number of features and equate their dimensionality while preserving the variance in the activations. Then, we performed an encoding model procedure to predict the fMRI responses of the testing videos. We repeated the shuffling ten times and then took the mean encoding correlation across ten shuffles to compare with the encoding results using unshuffled order of frames. As brain responses, we used beta values z-scored across stimuli and averaged over TRs 5-9 (TRs 5-9 reflect the peak of a typical BOLD response). We only predicted the “reliable” voxels, as defined in the “Univariate Split-Half Reliability Analysis” methods section above.

At each of the four blocks, we computed the difference in encoding accuracy (correlation) between activations that used the unshuffled and shuffled video input (unshuffled minus shuffled). We visualized the difference in correlation at each block in a whole-brain volume (Figure 5B). Within each Block, we then computed the difference in encoding accuracy (unshuffled minus shuffled) within each ROI. We performed a t-test (one-sample, two-sided) between the subject-averaged difference in correlation and a null hypothesis correlation of zero. We corrected for multiple comparisons across ROIs (FDR, assuming positive correlation, $p < 0.05$ with $n=88$ comparisons). The bar plot in Figure 5C displays the subject-averaged unshuffled minus shuffled correlation difference at each TSM ResNet50 Block and ROI. Significant blocks were marked with an appropriately colored asterisk.

We additionally correlate (Pearson) the unshuffled activations with each of the ten shuffled activations at each of the four blocks to examine the effect of frame shuffling on the model itself without the brain data. We average the correlations across activation units and the ten random seeds and find the following block correlations:

Block 1: 0.998
Block 2: 0.984
Block 3: 0.928
Block 4: 0.808

These results show decreasing activation similarity between the unshuffled and shuffled activations through the four blocks, with a notably steep drop in activation similarity by block 4. This result suggests that block 4 is most impacted by frame shuffling. Subsequently, block 4 also has the largest impact in brain prediction performance across cortex (figure 5B) and ROIs (figure 5C).

Encoding and variance partitioning analysis procedure

We use an encoding model and variance partitioning analysis to identify any unique variance explained by the first and third video seconds in the brain activity. In this way, we measure if the brain activity captures temporal content. The encoding algorithm involved two steps. In the first step, we fed the first and third video second frames of the 1000 training and 102 testing videos to an AlexNet architecture (Krizhevsky et al., 2017) pre-trained on the ILSVRC-2012 image classification challenge (Russakovsky et al., 2015), and we extracted the corresponding activations at each layer. We then applied the following operations to the activations of both video seconds (first and third), independently: we appended the feature maps of all layers, averaged them across frames, standardized them (using the mean and standard deviation of the training videos feature maps) and downsampled them to 100 components through principal components analysis (PCA) (computed on the training videos feature maps). This resulted in the training activations of size 1000 x 100 x 2 and testing activations of size 102 x 100 x 2 (number of videos x features x video seconds). In the second step, we linearly mapped the stimuli videos feature space onto voxel space, thus predicting the fMRI responses to videos. For each combination of subjects (10), fMRI voxels (N), and fMRI TRs (9), we trained the weights of a linear regression model to predict the fMRI training data (averaged over the three repetitions) using the training activations of both video seconds independently as predictors, and then multiplied the learned weights with the test activations. This resulted in two synthetic fMRI test data instances (size 10 subjects × 102 test videos × 9 fMRI TRs × N fMRI voxels), one for each video second.

To test our hypothesis that the brain activity captures information unique to the first and third video seconds, we ran a variance partitioning analysis between the biological fMRI test data and the two instances of synthetic fMRI test data. At each subject, TR, and voxel we ran a searchlight (Kriegeskorte et al., 2006) to calculate the portion of the biological fMRI test data (averaged over the ten test set repeats) uniquely explained by, respectively, the synthetic fMRI test data of the first or third video seconds. The unique biological fMRI variance explained by the first/third video second fMRI synthetic data consisted in the adjusted R^2 score of a linear regression trained to predict the biological fMRI data using both synthetic fMRI data instances as predictors, minus the adjusted R^2 score of a linear regression trained to predict the biological fMRI data using only the third/first synthetic fMRI data instance as predictor. We then observed at which TRs the unique variance explained by the two versions of synthetic test data peaked, and subtracted the peak TR of the first video second synthetic data from the peak TR of the third video second synthetic data. Next, we created subject wise binary whole brain masks with ones in voxels that show TR peak differences in the range 1 to 3 and zeros

elsewhere, summed the binary masks across subjects, and performed a binomial test with FDR correction to remove the non-significant voxels.

The variance partitioning analysis for the ROIs was similar but performed on the reliable (split-half reliability $p<0.05$) voxels within each ROI. Again, this results in time courses that reveal how well the synthetic fMRI test data from either the first or third video second explains the real fMRI data at each of the nine TRs. To quantify this difference, we again subtracted the peak TRs of the first and third video second synthetic data.

Metadata RSA analysis procedure

We performed Representational Similarity Analysis (RSA) (Kriegeskorte, 2008) between metadata RDMs and neural RDMs to examine the extent that representations defined by the semantic metadata of varying information content are reflected in neural activity. The analysis broadly consisted of correlating a Representational Dissimilarity Matrix (RDM) defined by the metadata (Figure 7B) with a RDM at each voxel in the brain defined by the brain responses (Figure 7A).

We defined five metadata RDMs – object, scene, action, object+scene+action, and text description RDMs - from each of the 102 testing set videos. The object, scene, action, and text description RDM was defined by first indexing the annotations from the first five annotators (if the video contained annotations from more than five annotators). We then feed the 5 object, 5 scene, 5 action, and 5 text caption metadata from each of the 102 testing set videos into a language model to generate vector embeddings for each label. In the case of the object labels, since each annotator labeled up to three different objects, we computed the word embedding of each object label individually then averaged them (object labels corresponding to “no more objects in video” were skipped) to obtain one object embedding per annotator. The object, scene, and action labels were fed into the FastText model (Bojanowski et al., 2017) to compute single-word embeddings (length 300) and the text descriptions were fed into the Sentence-BERT (Reimers & Gurevych, 2019) model to compute sentence-level embeddings (length 512). To minimize the effect of noise in the metadata labels, we averaged the 3 most similar vector embeddings together to result in a single vector embedding that represents the object, scene, action, or text caption for that video. At this step, the object+scene+action embedding was created by concatenating the individual object, scene, and action vector embeddings (length 900). We then computed the pairwise cosine distance between each video’s vector embedding to produce a single 102 x 102 Representational Dissimilarity Matrix (RDM) for the object, scene, action,

object+scene+action, and text caption metadata. Figure 7B shows the rank-normalized (rank each distance value and divide by the maximum rank) RDM for the object, scene, action, object+scene+action, and text description RDMs.

We compute the correlations (Spearman) between the metadata RDMs to measure the similarity of their information content. As expected, the correlations between the individual single-word object, scene, and action labels were generally lowest, since these labels highlight explicitly different components of the video. The correlations between the individual single-word object, scene, and action labels with the text descriptions were generally next highest, also expected because the text description likely contains information about each of the single-word labels plus extra information. The combined object+scene+action RDM with the text description RDM was higher still given even more overlapping information. The highest similarities were between the single-word object, scene, and action RDMs and the combined object+scene+action RDM because of explicitly overlapping information content (computed using the same embedding) but less extra information than the text description.

[scene, object]: 0.1808
[scene, action] 0.1133
[scene, scene+object+action]: 0.6641
[scene, text description]: 0.1969
[object, action]: 0.0884
[object, scene+object+action]: 0.6137
[object, text description]: 0.2618
[action, scene+object+action]: 0.5238
[action, text description]: 0.1560
[scene+object+action, text description]: 0.3145

To define the RDMs at each voxel in the brain for each subject, we perform a searchlight analysis in the way described in the “multivariate searchlight-based reliability analysis” methods section. To summarize, we center a sphere (radius of 4 voxels) around voxel *v* and extract the beta values (TRs 5-9 averaged over repetitions and z-scored across conditions) for all testing set conditions at all voxels encompassed in the sphere. Each stimulus thus has a corresponding vector of beta values, one from each voxel within the searchlight sphere. We compute the 1 - Pearson R correlation between all pairs of stimulus vectors to obtain an RDM at the centered voxel *v*. We repeat this process for all voxels in the whole brain for each subject (Figure 7A).

We then correlate (Spearman's R) the metadata RDM (cosine-distance, not rank-normalized) with the searchlight-based RDMs at each voxel for each of the 10 subjects separately. For the whole-brain analysis (Figure 7C), we compute a t-test (1-sample, 2-sided) against a null hypothesis of a correlation of 0 at each voxel then perform FDR correction (q=0.05, assuming positive correlation) on all p-values in the whole brain to obtain a set of significant voxels. We compute the noise-normalized correlation by dividing the correlation with the voxel's upper noise ceiling and plot the 10-subject average noise-normalized correlation at each significant voxel (Figure 7C). For the ROI-based analysis (Figure 7D), after we correlate (Spearman's R) the metadata RDM with the searchlight-based RDM at each voxel, we compute the average noise-normalized correlation within each ROI. For each ROI, we compute a one-way ANOVA test between the average noise-normalized correlations corresponding to the five semantic metadata models. If the p-value of the ANOVA test is significant ($p<0.05$, Bonferroni corrected with n=22 ROIs), we perform a pairwise Tukeys Honestly Significant Difference test (alpha=0.05). Significant differences between a pair of metadata models are reported with the dual-colored bars under the ROI name in Figure 7D.

Memorability analysis procedure

For each subject, we averaged the beta values (z-scored across conditions) over TRs 5-9 and over repetitions (3 repetitions per training set stimuli and 10 repetitions per testing set stimuli) to obtain one beta value per video. From the memory game implemented by Newman et al., 2020, we had one memorability score per video. Under the hypothesis that the magnitude of brain response positively correlates with stimuli memorability (Bainbridge et al., 2017; Bainbridge & Rissman, 2018; Jaegle et al., 2019), we performed a ranked correlation (Spearman's R) between the vector of memorability scores (size (1102 x 1)) and the vector of beta values (size (1102 x 1)) at each voxel for each subject. In this way, we obtained a correlation value at each voxel in the brain for each subject.

For the whole-brain analysis, we first performed a t-test (one-sample, one-sided) at each voxel against a null hypothesis of zero correlation. We then performed FDR correction on the p-values (q=0.05, assuming positive correlation). We visualized the subject-averaged correlations of the significant voxels that passed FDR correction in the whole-brain volume (Figure 8B).

For the ROI analysis, we computed the average correlation within each ROI for each subject. We then computed a t-test (one-sample, one-sided) for each ROI against a null hypothesis of zero average

correlation and corrected for multiple comparisons (Bonferroni, $p<0.05$ with $n=22$). We plotted the subject-average correlation at each ROI (Figure 8C) and denoted significance with an asterisk and a green colored bar. A one-sided, as opposed to two-sided, test against a correlation of 0 was computed because there exists clear a-priori hypotheses that relate memorability effects to larger (not smaller) magnitudes of response (Bainbridge et al., 2017; Bainbridge & Rissman, 2018; Bylinskii et al., 2022; Jaegle et al., 2019).

Data and Code Availability

The data and core analysis code will be available in a public repository upon peer-reviewed publication.

Author Contributions

R.C. and A.O. conceptualized the study and supervised the work. R.C., A.O., and P.I. designed the MRI protocols. A.L. and A.O. selected the stimuli. B.L., A.L., P.I., M.G., and R.M. collected the data. B.L., P.I., M.G., and K.D. preprocessed the data. B.L., K.D., and A.G. analyzed and visualized the data. B.P. trained the TSM model. S.J. ran crowd-sourced stimulus labeling experiments. B.L., S.J., and A.L. quality checked the stimuli metadata. B.L. and R.C. wrote the original draft of the manuscript. All authors reviewed and edited the manuscript.

Acknowledgements

This research was funded by DFG (CI-241/1-1, CI241/1-3, CI-241/1-7) and ERC grant (ERC-2018-StG) to R.M.C.; the Vannevar Bush Faculty Fellowship program funded by the ONR (N00014-16-1-3116) to A.O.; the Alfons and Gertrud Kassel foundation to G.R.; the EECS MathWorks Fellowship to B.L.. We also thank the MIT-IBM Watson AI Lab for support. The experiments were conducted at the Athinoula A. Martinos Imaging Center at the McGovern Institute for Brain Research, Massachusetts Institute of Technology, on a Siemens PrismaFit 3T scanner (Erlangen, Germany) supported with funding from a NIH Shared Instrumentation Grant (1S10OD021569). We would like to thank Santani Teng and Emilie Josephs for their valuable writing feedback.

References

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

Adelson, E. H., & Bergen, J. R. (1985). Spatiotemporal energy models for the perception of motion. *JOSA A*, *2*(2), 284–299. https://doi.org/10.1364/JOSAA.2.000284

Aliko, S., Huang, J., Gheorghiu, F., Meliss, S., & Skipper, J. I. (2020). A naturalistic neuroimaging database for understanding the brain using ecological stimuli. *Scientific Data*, *7*(1), Article 1. https://doi.org/10.1038/s41597-020-00680-2

Allen, E. J., St-Yves, G., Wu, Y., Breedlove, J. L., Prince, J. S., Dowdle, L. T., Nau, M., Caron, B., Pestilli, F., Charest, I., Hutchinson, J. B., Naselaris, T., & Kay, K. (2022). A massive 7T fMRI dataset to bridge cognitive neuroscience and artificial intelligence. *Nature Neuroscience*, *25*(1), Article 1. https://doi.org/10.1038/s41593-021-00962-x

Avants, B. B., Epstein, C. L., Grossman, M., & Gee, J. C. (2008). Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain. *Medical Image Analysis*, *12*(1), 26–41. https://doi.org/10.1016/j.media.2007.06.004

Baddeley, A. (1992). Working Memory. *Science*, *255*(5044), 556–559. https://doi.org/10.1126/science.1736359

Bainbridge, W. A. (2019). Chapter One - Memorability: How what we see influences what we remember. In K. D. Federmeier & D. M. Beck (Eds.), *Psychology of Learning and Motivation* (Vol. 70, pp. 1–27). Academic Press. https://doi.org/10.1016/bs.plm.2019.02.001

Bainbridge, W. A., & Baker, C. I. (2022). Multidimensional memory topography in the medial parietal cortex identified from neuroimaging of thousands of daily memory videos. *Nature Communications*, *13*(1), Article 1. https://doi.org/10.1038/s41467-022-34075-1

Bainbridge, W. A., Dilks, D. D., & Oliva, A. (2017). Memorability: A stimulus-driven perceptual neural signature distinctive from memory. *NeuroImage*, *149*, 141–152. https://doi.org/10.1016/j.neuroimage.2017.01.063

Bainbridge, W. A., & Rissman, J. (2018). Dissociating neural markers of stimulus memorability and subjective recognition during episodic retrieval. *Scientific Reports*, *8*(1), Article 1. https://doi.org/10.1038/s41598-018-26467-5

Bakhtiari, S., Mineault, P., Lillicrap, T., Pack, C., & Richards, B. (2021). The functional specialization of visual cortex emerges from training parallel pathways with self-supervised predictive learning. *Advances in Neural Information Processing Systems*, *34*, 25164–25178. https://proceedings.neurips.cc/paper/2021/hash/d384dec9f5f7a64a36b5c8f03b8a6d92-Abstract.html

Barrouillet, P., Bernardin, S., & Camos, V. (2004). Time Constraints and Resource Sharing in Adults' Working Memory Spans. *Journal of Experimental Psychology: General*, *133*(1), 83–100. https://doi.org/10.1037/0096-3445.133.1.83

Bartels, A., & Zeki, S. (2004). Functional brain mapping during free viewing of natural scenes. *Human Brain Mapping*, *21*(2), 75–85. https://doi.org/10.1002/hbm.10153

Behzadi, Y., Restom, K., Liau, J., & Liu, T. T. (2007). A component based noise correction method (CompCor) for BOLD and perfusion based fMRI. *NeuroImage*, *37*(1), 90–101. https://doi.org/10.1016/j.neuroimage.2007.04.042

Berkes, P., Orbán, G., Lengyel, M., & Fiser, J. (2011). Spontaneous Cortical Activity Reveals Hallmarks of an Optimal Internal Model of the Environment. *Science*, *331*(6013), 83–87. https://doi.org/10.1126/science.1195870

Bertasius, G., Wang, H., & Torresani, L. (2021). Is Space-Time Attention All You Need for Video Understanding? *Proceedings of the 38th International Conference on Machine Learning*, 813–824. https://proceedings.mlr.press/v139/bertasius21a.html

Bojanowski, P., Grave, E., Joulin, A., & Mikolov, T. (2017). Enriching Word Vectors with Subword Information. *Transactions of the Association for Computational Linguistics*, *5*, 135–146. https://doi.org/10.1162/tacl_a_00051

Bonner, M. F., & Epstein, R. A. (2021). Object representations in the human brain reflect the co-occurrence statistics of vision and language. *Nature Communications*, *12*(1), Article 1. https://doi.org/10.1038/s41467-021-24368-2

Born, R. T., & Bradley, D. C. (2005). Structure and Function of Visual Area Mt. *Annual Review of Neuroscience*, *28*(1), 157–189. https://doi.org/10.1146/annurev.neuro.26.041002.131052

Boyle, J. A., Pinsard, B., Boukhdhir, A., Belleville, S., Brambatti, S., Chen, J., & Bellec, P. (2020). The Courtois project on neuronal modelling-first data release. *26th Annual Meeting of the Organization for Human Brain Mapping*.

Buccino, G., Binkofski, F., Fink, G. R., Fadiga, L., Fogassi, L., Gallese, V., Seitz, R. J., Zilles, K., Rizzolatti, G., & Freund, H.-J. (2004). Action Observation Activates Premotor and Parietal Areas in a Somatotopic Manner: An fMRI Study. In *Social Neuroscience*. Psychology Press.

Buccino, G., Binkofski, F., & Riggio, L. (2004). The mirror neuron system and action recognition. *Brain and Language*, *89*(2), 370–376. https://doi.org/10.1016/S0093-934X(03)00356-0

Buccino, G., Vogt, S., Ritzl, A., Fink, G. R., Zilles, K., Freund, H.-J., & Rizzolatti, G. (2004). Neural Circuits Underlying Imitation Learning of Hand Actions: An Event-Related fMRI Study. *Neuron*, *42*(2), 323–334. https://doi.org/10.1016/S0896-6273(04)00181-3

Bylinskii, Z., Goetschalckx, L., Newman, A., & Oliva, A. (2022). Memorability: An Image-Computable Measure of Information Utility. In B. Ionescu, W. A. Bainbridge, & N. Murray (Eds.), *Human Perception of Visual Information* (pp. 207–239). Springer International Publishing. https://doi.org/10.1007/978-3-030-81465-6_8

Carandini, M. (2005). Do We Know What the Early Visual System Does? *Journal of Neuroscience*, *25*(46), 10577–10597. https://doi.org/10.1523/JNEUROSCI.3726-05.2005

Carlin, J. D., & Kriegeskorte, N. (2017). Adjudicating between face-coding models with individual-face fMRI responses. *PLOS Computational Biology*, *13*(7), e1005604. https://doi.org/10.1371/journal.pcbi.1005604

Carreira, J., & Zisserman, A. (2017). *Quo Vadis, Action Recognition? A New Model and the Kinetics Dataset*. 6299–6308. https://openaccess.thecvf.com/content_cvpr_2017/html/Carreira_Quo_Vadis_Action_CVPR_2017_paper.html

Chang, N., Pyles, J. A., Marcus, A., Gupta, A., Tarr, M. J., & Aminoff, E. M. (2019). BOLD5000, a public fMRI dataset while viewing 5000 visual images. *Scientific Data*, *6*(1), Article 1. https://doi.org/10.1038/s41597-019-0052-3

Chen, Z., Qing, J., & Zhou, J. H. (2023). *Cinematic Mindscapes: High-quality Video Reconstruction from Brain Activity* (arXiv:2305.11675). arXiv. https://doi.org/10.48550/arXiv.2305.11675

Cichy, R. M., Khosla, A., Pantazis, D., Torralba, A., & Oliva, A. (2016). Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence. *Scientific Reports*, *6*(1), 27755. https://doi.org/10.1038/srep27755

Cichy, R. M., Pantazis, D., & Oliva, A. (2016). Similarity-Based Fusion of MEG and fMRI Reveals Spatio-Temporal Dynamics in Human Cortex During Visual Object Recognition. *Cerebral Cortex*, *26*(8), 3563–3579. https://doi.org/10.1093/cercor/bhw135

Cohen, J. D., Perlstein, W. M., Braver, T. S., Nystrom, L. E., Noll, D. C., Jonides, J., & Smith, E. E. (1997). Temporal dynamics of brain activation during a working memory task. *Nature*, *386*(6625), Article 6625. https://doi.org/10.1038/386604a0

Cox, R. W., & Hyde, J. S. (1997). Software tools for analysis and visualization of fMRI data. *NMR in Biomedicine*, *10*(4–5), 171–178. https://doi.org/10.1002/(SICI)1099-1492(199706/08)10:4/5<171::AID-NBM453>3.0.CO;2-L

Dale, A. M., Fischl, B., & Sereno, M. I. (1999). Cortical Surface-Based Analysis. *NeuroImage*, *9*(2), 179–194. https://doi.org/10.1006/nimg.1998.0395

Dayan, P., & Abbott, L. F. (2001). *Theoretical neuroscience: Computational and mathematical modeling of neural systems*. Massachusetts Institute of Technology Press.

DeYoe, E. A., & Van Essen, D. C. (1988). Concurrent processing streams in monkey visual cortex. *Trends in Neurosciences*, *11*(5), 219–226. https://doi.org/10.1016/0166-2236(88)90130-0

DiCarlo, J. J., Zoccolan, D., & Rust, N. C. (2012). How Does the Brain Solve Visual Object Recognition? *Neuron*, *73*(3), 415–434. https://doi.org/10.1016/j.neuron.2012.01.010

Doerig, A., Kietzmann, T. C., Allen, E., Wu, Y., Naselaris, T., Kay, K., & Charest, I. (2022). *Semantic scene descriptions as an objective of human vision*. https://doi.org/10.48550/ARXIV.2209.11737

Downing, P. E., Jiang, Y., Shuman, M., & Kanwisher, N. (2001). A Cortical Area Selective for Visual Processing of the Human Body. *Science*, *293*(5539), 2470–2473. https://doi.org/10.1126/science.1063414

Eickenberg, M., Gramfort, A., Varoquaux, G., & Thirion, B. (2017). Seeing it all: Convolutional network layers map the function of the human visual system. *NeuroImage*, *152*, 184–194. https://doi.org/10.1016/j.neuroimage.2016.10.001

Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature*, *392*(6676), 598–601. https://doi.org/10.1038/33402

Esteban, O., Birman, D., Schaer, M., Koyejo, O. O., Poldrack, R. A., & Gorgolewski, K. J. (2017). MRIQC: Advancing the automatic prediction of image quality in MRI from unseen sites. *PLOS ONE*, *12*(9), e0184661. https://doi.org/10.1371/journal.pone.0184661

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

Esteban, O., Markiewicz, C. J., Goncalves, M., Provins, C., Kent, J. D., DuPre, E., Salo, T., Ciric, R., Pinsard, B., Blair, R. W., Poldrack, R. A., & Gorgolewski, K. J. (2022). *fMRIPrep: A robust preprocessing pipeline for functional MRI* [Computer software]. Zenodo. https://doi.org/10.5281/zenodo.7430291

Esteban, Oscar, Markiewicz, Christopher J., Burns, Christopher, Goncalves, Mathias, Jarecka, Dorota, Ziegler, Erik, Berleant, Shoshana, Ellis, David Gage, Pinsard, Basile, Madison, Cindee, Waskom, Michael, Notter, Michael Philipp, Clark, Daniel, Manhães-Savio, Alexandre, Clark, Dav, Jordan, Kesshi, Dayan, Michael, Halchenko, Yaroslav O., Loney, Fred, … Ghosh, Satrajit. (2022). *nipy/nipype: 1.8.3* (1.8.3) [Computer software]. Zenodo. https://doi.org/10.5281/ZENODO.596855

Fairhall, S. L., Albi, A., & Melcher, D. (2014). Temporal Integration Windows for Naturalistic Visual Sequences. *PLoS ONE*, *9*(7), e102248. https://doi.org/10.1371/journal.pone.0102248

Fan, L., Zhang, T., & Du, W. (2021). Optical-flow-based framework to boost video object detection performance with object enhancement. *Expert Systems with Applications*, *170*, 114544. https://doi.org/10.1016/j.eswa.2020.114544

Feichtenhofer, C., Pinz, A., & Zisserman, A. (2016). *Convolutional Two-Stream Network Fusion for Video Action Recognition*. 1933–1941. https://openaccess.thecvf.com/content_cvpr_2016/html/Feichtenhofer_Convolutional_Two-Stream_Network_CVPR_2016_paper.html

Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex (New York, N.Y.*, *1*(1), 1–47. https://doi.org/10.1093/cercor/1.1.1-a

Fonov, V., Evans, A., McKinstry, R., Almli, C., & Collins, D. (2009). Unbiased nonlinear average age-appropriate brain templates from birth to adulthood. *NeuroImage*, *47*, S102. https://doi.org/10.1016/S1053-8119(09)70884-5

Friston, K. J., Holmes, A. P., Worsley, K. J., Poline, J.-P., Frith, C. D., & Frackowiak, R. S. J. (1994). Statistical parametric maps in functional imaging: A general linear approach. *Human Brain Mapping*, *2*(4), 189–210. https://doi.org/10.1002/hbm.460020402

Furman, O., Dorfman, N., Hasson, U., Davachi, L., & Dudai, Y. (2007). They saw a movie: Long-term memory for an extended audiovisual narrative. *Learning & Memory*, *14*(6), 457–467. https://doi.org/10.1101/lm.550407

Gazzola, V., & Keysers, C. (2009). The Observation and Execution of Actions Share Motor and Somatosensory Voxels in all Tested Subjects: Single-Subject Analyses of Unsmoothed fMRI Data. *Cerebral Cortex*, *19*(6), 1239–1255. https://doi.org/10.1093/cercor/bhn181

Glasser, M. F., Coalson, T. S., Robinson, E. C., Hacker, C. D., Harwell, J., Yacoub, E., Ugurbil, K., Andersson, J., Beckmann, C. F., Jenkinson, M., Smith, S. M., & Van Essen, D. C. (2016). A multi-modal parcellation of human cerebral cortex. *Nature*, *536*(7615), 171–178. https://doi.org/10.1038/nature18933

Glasser, M. F., Sotiropoulos, S. N., Wilson, J. A., Coalson, T. S., Fischl, B., Andersson, J. L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J. R., Van Essen, D. C., & Jenkinson, M. (2013). The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage*, *80*, 105–124. https://doi.org/10.1016/j.neuroimage.2013.04.127

Goetschalckx, L., Moors, P., & Wagemans, J. (2018). Image memorability across longer time intervals. *Memory*, *26*(5), 581–588. https://doi.org/10.1080/09658211.2017.1383435

Gorgolewski, K., Burns, C. D., Madison, C., Clark, D., Halchenko, Y. O., Waskom, M. L., & Ghosh, S. S. (2011). Nipype: A Flexible, Lightweight and Extensible Neuroimaging Data Processing Framework in Python. *Frontiers in Neuroinformatics*, *5*. https://doi.org/10.3389/fninf.2011.00013

Gorgolewski, K. J., Auer, T., Calhoun, V. D., Craddock, R. C., Das, S., Duff, E. P., Flandin, G., Ghosh, S. S., Glatard, T., Halchenko, Y. O., Handwerker, D. A., Hanke, M., Keator, D., Li, X., Michael,

Z., Maumet, C., Nichols, B. N., Nichols, T. E., Pellman, J., … Poldrack, R. A. (2016). The brain imaging data structure, a format for organizing and describing outputs of neuroimaging experiments. *Scientific Data*, *3*(1), 160044. https://doi.org/10.1038/sdata.2016.44

Greve, D. N., & Fischl, B. (2009). Accurate and robust brain image alignment using boundary-based registration. *NeuroImage*, *48*(1), 63–72. https://doi.org/10.1016/j.neuroimage.2009.06.060

Grill-Spector, K., Kourtzi, Z., & Kanwisher, N. (2001). The lateral occipital complex and its role in object recognition. *Vision Research*, *41*(10), 1409–1422. https://doi.org/10.1016/S0042-6989(01)00073-6

Gu, Z., Jamison, K. W., Khosla, M., Allen, E. J., Wu, Y., St-Yves, G., Naselaris, T., Kay, K., Sabuncu, M. R., & Kuceyeski, A. (2022). NeuroGen: Activation optimized image synthesis for discovery neuroscience. *NeuroImage*, *247*, 118812. https://doi.org/10.1016/j.neuroimage.2021.118812

Güçlü, U., & van Gerven, M. A. J. (2017). Increasingly complex representations of natural movies across the dorsal stream are shared between subjects. *NeuroImage*, *145*, 329–336. https://doi.org/10.1016/j.neuroimage.2015.12.036

Guzman-Martinez, E., Leung, P., Franconeri, S., Grabowecky, M., & Suzuki, S. (2009). Rapid eye-fixation training without eyetracking. *Psychonomic Bulletin & Review*, *16*(3), 491–496. https://doi.org/10.3758/PBR.16.3.491

Han, J., Chen, C., Shao, L., Hu, X., Han, J., & Liu, T. (2015). Learning Computational Models of Video Memorability from fMRI Brain Imaging. *IEEE Transactions on Cybernetics*, *45*(8), 1692–1703. https://doi.org/10.1109/TCYB.2014.2358647

Han, K., Wen, H., Shi, J., Lu, K.-H., Zhang, Y., Fu, D., & Liu, Z. (2019). Variational autoencoder: An unsupervised model for encoding and decoding fMRI activity in visual cortex. *NeuroImage*, *198*, 125–136. https://doi.org/10.1016/j.neuroimage.2019.05.039

Hanke, M., Adelhöfer, N., Kottke, D., Iacovella, V., Sengupta, A., Kaule, F. R., Nigbur, R., Waite, A. Q., Baumgartner, F., & Stadler, J. (2016). A studyforrest extension, simultaneous fMRI and eye gaze recordings during prolonged natural stimulation. *Scientific Data*, *3*(1), Article 1. https://doi.org/10.1038/sdata.2016.92

Hardwick, R. M., Caspers, S., Eickhoff, S. B., & Swinnen, S. P. (2018). Neural correlates of action: Comparing meta-analyses of imagery, observation, and execution. *Neuroscience & Biobehavioral Reviews*, *94*, 31–44. https://doi.org/10.1016/j.neubiorev.2018.08.003

Hasson, U., Furman, O., Clark, D., Dudai, Y., & Davachi, L. (2008). Enhanced Intersubject Correlations during Movie Viewing Correlate with Successful Episodic Encoding. *Neuron*, *57*(3), 452–462. https://doi.org/10.1016/j.neuron.2007.12.009

Hasson, U., Landesman, O., Knappmeyer, B., Vallines, I., Rubin, N., & Heeger, D. J. (2008). Neurocinematics: The Neuroscience of Film. *Projections*, *2*(1), 1–26. https://doi.org/10.3167/proj.2008.020102

Hasson, U., Nir, Y., Levy, I., Fuhrmann, G., & Malach, R. (2004). Intersubject Synchronization of Cortical Activity During Natural Vision. *Science*, *303*(5664), 1634–1640. https://doi.org/10.1126/science.1089506

Hasson, U., Yang, E., Vallines, I., Heeger, D. J., & Rubin, N. (2008). A Hierarchy of Temporal Receptive Windows in Human Cortex. *Journal of Neuroscience*, *28*(10), 2539–2550. https://doi.org/10.1523/JNEUROSCI.5487-07.2008

Haxby, J. V. (2012). Multivariate pattern analysis of fMRI: The early beginnings. *NeuroImage*, *62*(2), 852–855. https://doi.org/10.1016/j.neuroimage.2012.03.016

Haxby, J. V., Guntupalli, J. S., Connolly, A. C., Halchenko, Y. O., Conroy, B. R., Gobbini, M. I., Hanke, M., & Ramadge, P. J. (2011). A Common, High-Dimensional Model of the Representational Space in Human Ventral Temporal Cortex. *Neuron*, *72*(2), 404–416. https://doi.org/10.1016/j.neuron.2011.08.026

Haxby, J. V., Guntupalli, J. S., Nastase, S. A., & Feilong, M. (2020). Hyperalignment: Modeling shared information encoded in idiosyncratic cortical topographies. *eLife*, *9*, e56601. https://doi.org/10.7554/eLife.56601

Haynes, J.-D. (2015). A Primer on Pattern-Based Approaches to fMRI: Principles, Pitfalls, and Perspectives. *Neuron*, *87*(2), 257–270. https://doi.org/10.1016/j.neuron.2015.05.025

He, K., Zhang, X., Ren, S., & Sun, J. (2016). *Deep Residual Learning for Image Recognition*. 770–778. https://openaccess.thecvf.com/content_cvpr_2016/html/He_Deep_Residual_Learning_CVPR_2016_paper.html

Hebart, M. N., Contier, O., Teichmann, L., Rockter, A. H., Zheng, C. Y., Kidder, A., Corriveau, A., Vaziri-Pashkam, M., & Baker, C. I. (2023). THINGS-data, a multimodal collection of large-scale datasets for investigating object representations in human brain and behavior. *eLife*, *12*, e82580. https://doi.org/10.7554/eLife.82580

Hebart, M. N., Dickter, A. H., Kidder, A., Kwok, W. Y., Corriveau, A., Van Wicklin, C., & Baker, C. I. (2019). THINGS: A database of 1,854 object concepts and more than 26,000 naturalistic object images. *PLOS ONE*, *14*(10), e0223792. https://doi.org/10.1371/journal.pone.0223792

Heim, S., Amunts, K., Hensel, T., Grande, M., Huber, W., Binkofski, F., & Eickhoff, S. (2012). The Role of Human Parietal Area 7A as a Link between Sequencing in Hand Actions and in Overt Speech Production. *Frontiers in Psychology*, *3*. https://www.frontiersin.org/articles/10.3389/fpsyg.2012.00534

Honey, C. J., Thesen, T., Donner, T. H., Silbert, L. J., Carlson, C. E., Devinsky, O., Doyle, W. K., Rubin, N., Heeger, D. J., & Hasson, U. (2012). Slow Cortical Dynamics and the Accumulation of Information over Long Timescales. *Neuron*, *76*(2), 423–434. https://doi.org/10.1016/j.neuron.2012.08.011

Horikawa, T., & Kamitani, Y. (2017). Generic decoding of seen and imagined objects using hierarchical visual features. *Nature Communications*, *8*(1), Article 1. https://doi.org/10.1038/ncomms15037

Hu, M., Ge, P., Wang, X., Lin, H., & Ren, F. (2022). A spatio-temporal integrated model based on local and global features for video expression recognition. *The Visual Computer*, *38*(8), 2617–2634. https://doi.org/10.1007/s00371-021-02136-z

Hutchison, R. M., Womelsdorf, T., Allen, E. A., Bandettini, P. A., Calhoun, V. D., Corbetta, M., Della Penna, S., Duyn, J. H., Glover, G. H., Gonzalez-Castillo, J., Handwerker, D. A., Keilholz, S., Kiviniemi, V., Leopold, D. A., de Pasquale, F., Sporns, O., Walter, M., & Chang, C. (2013). Dynamic functional connectivity: Promise, issues, and interpretations. *NeuroImage*, *80*, 360–378. https://doi.org/10.1016/j.neuroimage.2013.05.079

Isola, P., Parikh, D., Torralba, A., & Oliva, A. (2011). Understanding the Intrinsic Memorability of Images. *Advances in Neural Information Processing Systems*, *24*. https://proceedings.neurips.cc/paper/2011/hash/286674e3082feb7e5afb92777e48821f-Abstract.html

Jaegle, A., Mehrpour, V., Mohsenzadeh, Y., Meyer, T., Oliva, A., & Rust, N. (2019). Population response magnitude variation in inferotemporal cortex predicts image memorability. *eLife*, *8*, e47596. https://doi.org/10.7554/eLife.47596

Jenkinson, M., Bannister, P., Brady, M., & Smith, S. (2002). Improved Optimization for the Robust and Accurate Linear Registration and Motion Correction of Brain Images. *NeuroImage*, *17*(2), 825–841. https://doi.org/10.1006/nimg.2002.1132

Ji, S., Xu, W., Yang, M., & Yu, K. (2013). 3D Convolutional Neural Networks for Human Action Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *35*(1), 221–231. https://doi.org/10.1109/TPAMI.2012.59

Julian, J. B., Fedorenko, E., Webster, J., & Kanwisher, N. (2012). An algorithmic method for functionally defining regions of interest in the ventral visual pathway. *NeuroImage*, *60*(4), 2357–2364. https://doi.org/10.1016/j.neuroimage.2012.02.055

Kahou, S. E., Bouthillier, X., Lamblin, P., Gulcehre, C., Michalski, V., Konda, K., Jean, S., Froumenty, P., Dauphin, Y., Boulanger-Lewandowski, N., Chandias Ferrari, R., Mirza, M., Warde-Farley, D., Courville, A., Vincent, P., Memisevic, R., Pal, C., & Bengio, Y. (2016). EmoNets: Multimodal deep learning approaches for emotion recognition in video. *Journal on Multimodal User Interfaces*, *10*(2), 99–111. https://doi.org/10.1007/s12193-015-0195-2

Kamitani, Y., & Tong, F. (2006). Decoding seen and attended motion directions from activity in the human visual cortex. *Current Biology*, *16*(11), 1096–1102. https://doi.org/10.1016/j.cub.2006.04.003

Kanwisher, N. (2010). Functional specificity in the human brain: A window into the functional architecture of the mind. *Proceedings of the National Academy of Sciences*, *107*(25), 11163–11170. https://doi.org/10.1073/pnas.1005062107

Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The Fusiform Face Area: A Module in Human Extrastriate Cortex Specialized for Face Perception. *The Journal of Neuroscience*, *17*(11), 4302–4311. https://doi.org/10.1523/JNEUROSCI.17-11-04302.1997

Kay, K., Jamison, K. W., Zhang, R.-Y., & Uğurbil, K. (2020). A temporal decomposition method for identifying venous effects in task-based fMRI. *Nature Methods*, *17*(10), Article 10. https://doi.org/10.1038/s41592-020-0941-6

Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., Suleyman, M., & Zisserman, A. (2017). *The Kinetics Human Action Video Dataset* (arXiv:1705.06950). arXiv. https://doi.org/10.48550/arXiv.1705.06950

Khosla, A., Raju, A. S., Torralba, A., & Oliva, A. (2015). *Understanding and Predicting Image Memorability at a Large Scale*. 2390–2398. https://openaccess.thecvf.com/content_iccv_2015/html/Khosla_Understanding_and_Predicting_ICCV_2015_paper.html

Khosla, M., Ratan Murty, N. A., & Kanwisher, N. (2022). A highly selective response to food in human visual cortex revealed by hypothesis-free voxel decomposition. *Current Biology*, *32*(19), 4159-4171.e9. https://doi.org/10.1016/j.cub.2022.08.009

Kiebel, S. J., Daunizeau, J., & Friston, K. J. (2008). A Hierarchy of Time-Scales and the Brain. *PLoS Computational Biology*, *4*(11), e1000209. https://doi.org/10.1371/journal.pcbi.1000209

Kietzmann, T. C., Spoerer, C. J., Sörensen, L. K. A., Cichy, R. M., Hauk, O., & Kriegeskorte, N. (2019). Recurrence is required to capture the representational dynamics of the human visual system. *Proceedings of the National Academy of Sciences*, *116*(43), 21854–21863. https://doi.org/10.1073/pnas.1905544116

Klein, A., Ghosh, S. S., Bao, F. S., Giard, J., Häme, Y., Stavsky, E., Lee, N., Rossa, B., Reuter, M., Chaibub Neto, E., & Keshavan, A. (2017). Mindboggling morphometry of human brains. *PLOS Computational Biology*, *13*(2), e1005350. https://doi.org/10.1371/journal.pcbi.1005350

Koivisto, M., Railo, H., Revonsuo, A., Vanni, S., & Salminen-Vaparanta, N. (2011). Recurrent Processing in V1/V2 Contributes to Categorization of Natural Scenes. *Journal of Neuroscience*, *31*(7), 2488–2492. https://doi.org/10.1523/JNEUROSCI.3074-10.2011

Konen, C. S., & Kastner, S. (2008). Representation of Eye Movements and Stimulus Motion in Topographically Organized Areas of Human Posterior Parietal Cortex. *Journal of Neuroscience*, *28*(33), 8361–8375. https://doi.org/10.1523/JNEUROSCI.1930-08.2008

Kosakowski, H. L., Cohen, M. A., Takahashi, A., Keil, B., Kanwisher, N., & Saxe, R. (2022). Selective responses to faces, scenes, and bodies in the ventral visual pathway of infants. *Current Biology*, *32*(2), 265-274.e5. https://doi.org/10.1016/j.cub.2021.10.064

Krekelberg, B., Dannenberg, S., Hoffmann, K.-P., Bremmer, F., & Ross, J. (2003). Neural correlates of implied motion. *Nature*, *424*(6949), 674–677. https://doi.org/10.1038/nature01852

Kret, M. E., Pichon, S., Grèzes, J., & de Gelder, B. (2011). Similarities and differences in perceiving threat from dynamic faces and bodies. An fMRI study. *NeuroImage*, *54*(2), 1755–1762. https://doi.org/10.1016/j.neuroimage.2010.08.012

Kriegeskorte, N. (2008). Representational similarity analysis – connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*. https://doi.org/10.3389/neuro.06.004.2008

Kriegeskorte, N. (2015). Deep Neural Networks: A New Framework for Modeling Biological Vision and Brain Information Processing. *Annual Review of Vision Science*, *1*(1), 417–446. https://doi.org/10.1146/annurev-vision-082114-035447

Kriegeskorte, N., & Douglas, P. K. (2019). Interpreting encoding and decoding models. *Current Opinion in Neurobiology*, *55*, 167–179. https://doi.org/10.1016/j.conb.2019.04.002

Kriegeskorte, N., Goebel, R., & Bandettini, P. (2006). Information-based functional brain mapping. *Proceedings of the National Academy of Sciences*, *103*(10), 3863–3868. https://doi.org/10.1073/pnas.0600244103

Kriegeskorte, N., & Kievit, R. A. (2013). Representational geometry: Integrating cognition, computation, and the brain. *Trends in Cognitive Sciences*, *17*(8), 401–412. https://doi.org/10.1016/j.tics.2013.06.007

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2017). ImageNet classification with deep convolutional neural networks. *Communications of the ACM*, *60*(6), 84–90. https://doi.org/10.1145/3065386

Kubilius, J., Schrimpf, M., Kar, K., Rajalingham, R., Hong, H., Majaj, N., Issa, E., Bashivan, P., Prescott-Roy, J., Schmidt, K., Nayebi, A., Bear, D., Yamins, D. L., & DiCarlo, J. J. (2019). Brain-Like Object Recognition with High-Performing Shallow Recurrent ANNs. *Advances in Neural Information Processing Systems*, *32*. https://proceedings.neurips.cc/paper/2019/hash/7813d1590d28a7dd372ad54b5d29d033-Abstract.html

Kupershmidt, G., Beliy, R., Gaziv, G., & Irani, M. (2022). *A Penny for Your (visual) Thoughts: Self-Supervised Reconstruction of Natural Movies from Brain Activity*. https://doi.org/10.48550/ARXIV.2206.03544

Lafer-Sousa, R., Conway, B. R., & Kanwisher, N. G. (2016). Color-Biased Regions of the Ventral Visual Pathway Lie between Face- and Place-Selective Regions in Humans, as in Macaques. *Journal of Neuroscience*, *36*(5), 1682–1697. https://doi.org/10.1523/JNEUROSCI.3164-15.2016

Lage-Castellanos, A., Valente, G., Formisano, E., & Martino, F. D. (2019). Methods for computing the maximum performance of computational models of fMRI responses. *PLOS Computational Biology*, *15*(3), e1006397. https://doi.org/10.1371/journal.pcbi.1006397

Lanczos, C. (1964). Evaluation of noisy data. *Journal of the Society for Industrial and Applied Mathematics, Series B: Numerical Analysis*, *1*(1), 76–85.

Le, A., Vesia, M., Yan, X., Crawford, J. D., & Niemeier, M. (2017). Parietal area BA7 integrates motor programs for reaching, grasping, and bimanual coordination. *Journal of Neurophysiology*, *117*(2), 624–636. https://doi.org/10.1152/jn.00299.2016

Li, Y., Song, Y., & Luo, J. (2017). *Improving Pairwise Ranking for Multi-Label Image Classification*. 3617–3625. https://openaccess.thecvf.com/content_cvpr_2017/html/Li_Improving_Pairwise_Ranking_CVPR_2017_paper.html

Lin, J., Gan, C., & Han, S. (2019). *TSM: Temporal Shift Module for Efficient Video Understanding*. 7083–7093. https://openaccess.thecvf.com/content_ICCV_2019/html/Lin_TSM_Temporal_Shift_Module_for_Efficient_Video_Understanding_ICCV_2019_paper.html

Lingnau, A., & Downing, P. E. (2015). The lateral occipitotemporal cortex in action. *Trends in Cognitive Sciences*, *19*(5), 268–277. https://doi.org/10.1016/j.tics.2015.03.006

Logothetis, N. K., & Sheinberg, D. L. (1996). Visual object recognition. *Annual Review of Neuroscience*, *19*, 577–621. https://doi.org/10.1146/annurev.ne.19.030196.003045

Luo, A. F., Henderson, M. M., Wehbe, L., & Tarr, M. J. (2023). *Brain Diffusion for Visual Exploration: Cortical Discovery using Large Scale Generative Models* (arXiv:2306.03089). arXiv. https://doi.org/10.48550/arXiv.2306.03089

Martin, A., & Chao, L. L. (2001). Semantic memory and the brain: Structure and processes. *Current Opinion in Neurobiology*, *11*(2), 194–201. https://doi.org/10.1016/S0959-4388(00)00196-3

McMahon, E., Bonner, M. F., & Isik, L. (2023). Hierarchical organization of social action features along the lateral visual pathway. *Current Biology*, *33*(23), 5035-5047.e8. https://doi.org/10.1016/j.cub.2023.10.015

Miech, A., Zhukov, D., Alayrac, J.-B., Tapaswi, M., Laptev, I., & Sivic, J. (2019). *HowTo100M: Learning a Text-Video Embedding by Watching Hundred Million Narrated Video Clips*. 2630–2640. https://openaccess.thecvf.com/content_ICCV_2019/html/Miech_HowTo100M_Learning_a_Text-Video_Embedding_by_Watching_Hundred_Million_Narrated_ICCV_2019_paper.html

Mineault, P., Bakhtiari, S., Richards, B., & Pack, C. (2021). Your head is there to move you around: Goal-driven models of the primate dorsal pathway. *Advances in Neural Information Processing Systems*, *34*, 28757–28771. https://proceedings.neurips.cc/paper/2021/hash/f1676935f9304b97d59b0738289d2e22-Abstract.html

Misaki, M., Luh, W.-M., & Bandettini, P. A. (2013). Accurate decoding of sub-TR timing differences in stimulations of sub-voxel regions from multi-voxel response patterns. *NeuroImage*, *66*, 623–633. https://doi.org/10.1016/j.neuroimage.2012.10.069

Mohsenzadeh, Y., Mullin, C., Oliva, A., & Pantazis, D. (2019). The perceptual neural trace of memorable unseen scenes. *Scientific Reports*, *9*(1), Article 1. https://doi.org/10.1038/s41598-019-42429-x

Monfort, M., Jin, S., Liu, A., Harwath, D., Feris, R., Glass, J., & Oliva, A. (2021). *Spoken Moments: Learning Joint Audio-Visual Representations From Video Descriptions*. 14871–14881. https://openaccess.thecvf.com/content/CVPR2021/html/Monfort_Spoken_Moments_Learning_Joint_Audio-Visual_Representations_From_Video_Descriptions_CVPR_2021_paper.html

Monfort, M., Pan, B., Ramakrishnan, K., Andonian, A., McNamara, B. A., Lascelles, A., Fan, Q., Gutfreund, D., Feris, R. S., & Oliva, A. (2022). Multi-Moments in Time: Learning and Interpreting Models for Multi-Action Video Understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *44*(12), 9434–9445. https://doi.org/10.1109/TPAMI.2021.3126682

Monfort, M., Vondrick, C., Oliva, A., Andonian, A., Zhou, B., Ramakrishnan, K., Bargal, S. A., Yan, T., Brown, L., Fan, Q., & Gutfreund, D. (2020). Moments in Time Dataset: One Million Videos for Event Understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *42*(2), 502–508. https://doi.org/10.1109/TPAMI.2019.2901464

Murray, J. D., Bernacchia, A., Freedman, D. J., Romo, R., Wallis, J. D., Cai, X., Padoa-Schioppa, C., Pasternak, T., Seo, H., Lee, D., & Wang, X.-J. (2014). A hierarchy of intrinsic timescales across primate cortex. *Nature Neuroscience*, *17*(12), 1661–1663. https://doi.org/10.1038/nn.3862

Naselaris, T., Kay, K. N., Nishimoto, S., & Gallant, J. L. (2011). Encoding and decoding in fMRI. *NeuroImage*, *56*(2), 400–410. https://doi.org/10.1016/j.neuroimage.2010.07.073

Newman, A., Fosco, C., Casser, V., Lee, A., McNamara, B., & Oliva, A. (2020). Multimodal Memorability: Modeling Effects of Semantics and Decay on Video Memorability. In A. Vedaldi, H. Bischof, T. Brox, & J.-M. Frahm (Eds.), *Computer Vision – ECCV 2020* (pp. 223–240). Springer International Publishing. https://doi.org/10.1007/978-3-030-58517-4_14

Nili, H., Wingfield, C., Walther, A., Su, L., Marslen-Wilson, W., & Kriegeskorte, N. (2014). A Toolbox for Representational Similarity Analysis. *PLOS Computational Biology*, *10*(4), e1003553. https://doi.org/10.1371/journal.pcbi.1003553

Nishimoto, S., & Gallant, J. L. (2011). A Three-Dimensional Spatiotemporal Receptive Field Model Explains Responses of Area MT Neurons to Naturalistic Movies. *Journal of Neuroscience*, *31*(41), 14551–14564. https://doi.org/10.1523/JNEUROSCI.6801-10.2011

Nishimoto, S., Vu, A. T., Naselaris, T., Benjamini, Y., Yu, B., & Gallant, J. L. (2011). Reconstructing Visual Experiences from Brain Activity Evoked by Natural Movies. *Current Biology*, *21*(19), 1641–1646. https://doi.org/10.1016/j.cub.2011.08.031

Olshausen, B. A., & Field, D. J. (1996a). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, *381*(6583), Article 6583. https://doi.org/10.1038/381607a0

Olshausen, B. A., & Field, D. J. (1996b). Natural image statistics and efficient coding. *Network: Computation in Neural Systems*, *7*(2), 333–339. https://doi.org/10.1088/0954-898X_7_2_014

Orlov, T., & Zohary, E. (2018). Object Representations in Human Visual Cortex Formed Through Temporal Integration of Dynamic Partial Shape Views. *Journal of Neuroscience*, *38*(3), 659–678. https://doi.org/10.1523/JNEUROSCI.1318-17.2017

Pascual-Leone, A., & Walsh, V. (2001). Fast backprojections from the motion to the primary visual area necessary for visual awareness. *Science*, *292*(5516), 510–512.

Peeters, R. R., Rizzolatti, G., & Orban, G. A. (2013). Functional properties of the left parietal tool use region. *NeuroImage*, *78*, 83–93. https://doi.org/10.1016/j.neuroimage.2013.04.023

Peeters, R., Simone, L., Nelissen, K., Fabbri-Destro, M., Vanduffel, W., Rizzolatti, G., & Orban, G. A. (2009). The Representation of Tool Use in Humans and Monkeys: Common and Uniquely

Human Features. *Journal of Neuroscience*, *29*(37), 11523–11539. https://doi.org/10.1523/JNEUROSCI.2040-09.2009

Piasini, E., Soltuzu, L., Muratore, P., Caramellino, R., Vinken, K., Op de Beeck, H., Balasubramanian, V., & Zoccolan, D. (2021). Temporal stability of stimulus representation increases along rodent visual cortical hierarchies. *Nature Communications*, *12*(1), Article 1. https://doi.org/10.1038/s41467-021-24456-3

Pitcher, D., Dilks, D. D., Saxe, R. R., Triantafyllou, C., & Kanwisher, N. (2011). Differential selectivity for dynamic versus static information in face-selective cortical regions. *NeuroImage*, *56*(4), 2356–2363. https://doi.org/10.1016/j.neuroimage.2011.03.067

Pitcher, D., & Ungerleider, L. G. (2021). Evidence for a Third Visual Pathway Specialized for Social Perception. *Trends in Cognitive Sciences*, *25*(2), 100–110. https://doi.org/10.1016/j.tics.2020.11.006

Power, J. D., Mitra, A., Laumann, T. O., Snyder, A. Z., Schlaggar, B. L., & Petersen, S. E. (2014). Methods to detect, characterize, and remove motion artifact in resting state fMRI. *NeuroImage*, *84*, 320–341. https://doi.org/10.1016/j.neuroimage.2013.08.048

Press, W. A., Brewer, A. A., Dougherty, R. F., Wade, A. R., & Wandell, B. A. (2001). Visual areas and spatial summation in human visual cortex. *Vision Research*, *41*(10), 1321–1332. https://doi.org/10.1016/S0042-6989(01)00074-8

Prince, J. S., Charest, I., Kurzawski, J. W., Pyles, J. A., Tarr, M. J., & Kay, K. N. (2022). Improving the accuracy of single-trial fMRI response estimates using GLMsingle. *eLife*, *11*, e77599. https://doi.org/10.7554/eLife.77599

Rajalingham, R., Issa, E. B., Bashivan, P., Kar, K., Schmidt, K., & DiCarlo, J. J. (2018). Large-Scale, High-Resolution Comparison of the Core Visual Object Recognition Behavior of Humans, Monkeys, and State-of-the-Art Deep Artificial Neural Networks. *Journal of Neuroscience*, *38*(33), 7255–7269. https://doi.org/10.1523/JNEUROSCI.0388-18.2018

Ratan Murty, N. A., Bashivan, P., Abate, A., DiCarlo, J. J., & Kanwisher, N. (2021). Computational models of category-selective brain regions enable high-throughput tests of selectivity. *Nature Communications*, *12*(1), 5540. https://doi.org/10.1038/s41467-021-25409-6

Reimers, N., & Gurevych, I. (2019). *Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks*. https://doi.org/10.48550/ARXIV.1908.10084

Ress, D., & Heeger, D. J. (2003). Neuronal correlates of perception in early visual cortex. *Nature Neuroscience*, *6*(4), Article 4. https://doi.org/10.1038/nn1024

Riou, B., Lesourd, M., Brunel, L., & Versace, R. (2011). Visual memory and visual perception: When memory improves visual search. *Memory & Cognition*, *39*(6), 1094–1102. https://doi.org/10.3758/s13421-011-0075-2

Rizzolatti, G., & Sinigaglia, C. (2010). The functional role of the parieto-frontal mirror circuit: Interpretations and misinterpretations. *Nature Reviews Neuroscience*, *11*(4), Article 4. https://doi.org/10.1038/nrn2805

Roberts, J., Wallis, G., & Breakspear, M. (2013). Fixational eye movements during viewing of dynamic natural scenes. *Frontiers in Psychology*, *4*. https://www.frontiersin.org/articles/10.3389/fpsyg.2013.00797

Roe, A. W., Chelazzi, L., Connor, C. E., Conway, B. R., Fujita, I., Gallant, J. L., Lu, H., & Vanduffel, W. (2012). Toward a Unified Theory of Visual Area V4. *Neuron*, *74*(1), 12–29. https://doi.org/10.1016/j.neuron.2012.03.011

Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., & Fei-Fei, L. (2015). ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*, *115*(3), 211–252. https://doi.org/10.1007/s11263-015-0816-y

Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., & Chen, L.-C. (2018). *MobileNetV2: Inverted Residuals and Linear Bottlenecks*. 4510–4520. https://openaccess.thecvf.com/content_cvpr_2018/html/Sandler_MobileNetV2_Inverted_Residuals_CVPR_2018_paper.html

Satterthwaite, T. D., Elliott, M. A., Gerraty, R. T., Ruparel, K., Loughead, J., Calkins, M. E., Eickhoff, S. B., Hakonarson, H., Gur, R. C., Gur, R. E., & Wolf, D. H. (2013). An improved framework for confound regression and filtering for control of motion artifact in the preprocessing of resting-state functional connectivity data. *NeuroImage*, *64*, 240–256. https://doi.org/10.1016/j.neuroimage.2012.08.052

Schneider, W. X. (2013). Selective visual processing across competition episodes: A theory of task-driven visual attention and working memory. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *368*(1628), 20130060. https://doi.org/10.1098/rstb.2013.0060

Schrimpf, M., Kubilius, J., Hong, H., Majaj, N. J., Rajalingham, R., Issa, E. B., Kar, K., Bashivan, P., Prescott-Roy, J., Geiger, F., Schmidt, K., Yamins, D. L. K., & DiCarlo, J. J. (2020). *Brain-Score: Which Artificial Neural Network for Object Recognition is most Brain-Like?* (p. 407007). bioRxiv. https://doi.org/10.1101/407007

Schrimpf, M., Kubilius, J., Lee, M. J., Ratan Murty, N. A., Ajemian, R., & DiCarlo, J. J. (2020). Integrative Benchmarking to Advance Neurally Mechanistic Models of Human Intelligence. *Neuron*, *108*(3), 413–423. https://doi.org/10.1016/j.neuron.2020.07.040

Schultz, J., & Pilz, K. S. (2009). Natural facial motion enhances cortical responses to faces. *Experimental Brain Research*, *194*(3), 465–475. https://doi.org/10.1007/s00221-009-1721-9

Seeliger, K., Sommers, R. P., Güçlü, U., Bosch, S. E., & Gerven, M. A. J. van. (2019). *A large single-participant fMRI dataset for probing brain responses to naturalistic stimuli in space and time* (p. 687681). bioRxiv. https://doi.org/10.1101/687681

Senior, C., Barnes, J., Giampietroc, V., Simmons, A., Bullmore, E. T., Brammer, M., & David, A. S. (2000). The functional neuroanatomy of implicit-motion perception or 'representational momentum.' *Current Biology*, *10*(1), 16–22. https://doi.org/10.1016/S0960-9822(99)00259-6

Shafiee, M. J., Chywl, B., Li, F., & Wong, A. (2017). *Fast YOLO: A Fast You Only Look Once System for Real-time Embedded Object Detection in Video*. https://doi.org/10.48550/ARXIV.1709.05943

Shirai, N., & Imura, T. (2014). Implied motion perception from a still image in infancy. *Experimental Brain Research*, *232*(10), 3079–3087. https://doi.org/10.1007/s00221-014-3996-8

Shmuelof, L., & Zohary, E. (2005). Dissociation between Ventral and Dorsal fMRI Activation during Object and Action Recognition. *Neuron*, *47*(3), 457–470. https://doi.org/10.1016/j.neuron.2005.06.034

Silvanto, J., Cowey, A., Lavie, N., & Walsh, V. (2005). Striate cortex (V1) activity gates awareness of motion. *Nature Neuroscience*, *8*(2), Article 2. https://doi.org/10.1038/nn1379

Silvanto, J., Lavie, N., & Walsh, V. (2005). Double Dissociation of V1 and V5/MT activity in Visual Awareness. *Cerebral Cortex*, *15*(11), 1736–1741. https://doi.org/10.1093/cercor/bhi050

Silver, M. A., & Kastner, S. (2009). Topographic maps in human frontal and parietal cortex. *Trends in Cognitive Sciences*, *13*(11), 488–495. https://doi.org/10.1016/j.tics.2009.08.005

Slotnick, S. D., Thompson, W. L., & Kosslyn, S. M. (2012). Visual memory and visual mental imagery recruit common control and sensory regions of the brain. *Cognitive Neuroscience*, *3*(1), 14–20. https://doi.org/10.1080/17588928.2011.578210

Smith, A. T., Greenlee, M. W., Singh, K. D., Kraemer, F. M., & Hennig, J. (1998). The Processing of First- and Second-Order Motion in Human Visual Cortex Assessed by Functional Magnetic Resonance Imaging (fMRI). *Journal of Neuroscience*, *18*(10), 3816–3830. https://doi.org/10.1523/JNEUROSCI.18-10-03816.1998

Smith, S. M., Vidaurre, D., Beckmann, C. F., Glasser, M. F., Jenkinson, M., Miller, K. L., Nichols, T. E., Robinson, E. C., Salimi-Khorshidi, G., Woolrich, M. W., Barch, D. M., Uğurbil, K., & Van Essen, D. C. (2013). Functional connectomics from resting-state fMRI. *Trends in Cognitive Sciences*, *17*(12), 666–682. https://doi.org/10.1016/j.tics.2013.09.016

Smyth, D., Willmore, B., Baker, G. E., Thompson, I. D., & Tolhurst, D. J. (2003). The Receptive-Field Organization of Simple Cells in Primary Visual Cortex of Ferrets under Natural Scene Stimulation. *Journal of Neuroscience*, *23*(11), 4746–4759. https://doi.org/10.1523/JNEUROSCI.23-11-04746.2003

Spunt, R. P., Satpute, A. B., & Lieberman, M. D. (2011). Identifying the What, Why, and How of an Observed Action: An fMRI Study of Mentalizing and Mechanizing during Action Observation. *Journal of Cognitive Neuroscience*, *23*(1), 63–74. https://doi.org/10.1162/jocn.2010.21446

Tustison, N. J., Avants, B. B., Cook, P. A., Zheng, Y., Egan, A., Yushkevich, P. A., & Gee, J. C. (2010). N4ITK: Improved N3 Bias Correction. *IEEE Transactions on Medical Imaging*, *29*(6), 1310–1320. https://doi.org/10.1109/TMI.2010.2046908

Tzirakis, P., Zhang, J., & Schuller, B. W. (2018). End-to-End Speech Emotion Recognition Using Deep Neural Networks. *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 5089–5093. https://doi.org/10.1109/ICASSP.2018.8462677

Urgen, B. A., Pehlivan, S., & Saygin, A. P. (2019). Distinct representations in occipito-temporal, parietal, and premotor cortex during action perception revealed by fMRI and computational modeling. *Neuropsychologia*, *127*, 35–47. https://doi.org/10.1016/j.neuropsychologia.2019.02.006

VanRullen, R., & Thorpe, S. J. (2001). The Time Course of Visual Processing: From Early Perception to Decision-Making. *Journal of Cognitive Neuroscience*, *13*(4), 454–461. https://doi.org/10.1162/08989290152001880

Vermeulen, N., Corneille, O., & Niedenthal, P. M. (2008). Sensory load incurs conceptual processing costs. *Cognition*, *109*(2), 287–294. https://doi.org/10.1016/j.cognition.2008.09.004

Wang, L., Mruczek, R. E. B., Arcaro, M. J., & Kastner, S. (2015). Probabilistic Maps of Visual Topography in Human Cortex. *Cerebral Cortex*, *25*(10), 3911–3931. https://doi.org/10.1093/cercor/bhu277

Wang, L., Xiong, Y., Wang, Z., Qiao, Y., Lin, D., Tang, X., & Van Gool, L. (2016). Temporal Segment Networks: Towards Good Practices for Deep Action Recognition. In B. Leibe, J. Matas, N. Sebe, & M. Welling (Eds.), *Computer Vision – ECCV 2016* (pp. 20–36). Springer International Publishing. https://doi.org/10.1007/978-3-319-46484-8_2

Wang, Y., Jiang, L., Yang, M.-H., Li, L.-J., Long, M., & Fei-Fei, L. (2023, January 23). *Eidetic 3D LSTM: A Model for Video Prediction and Beyond*. International Conference on Learning Representations. https://openreview.net/forum?id=B1lKS2AqtX

Watson, A. B., & Ahumada, A. J. (1985). Model of human visual-motion sensing. *JOSA A*, *2*(2), 322–342. https://doi.org/10.1364/JOSAA.2.000322

Weinberger, N. M. (2004). Specific long-term memory traces in primary auditory cortex. *Nature Reviews Neuroscience*, *5*(4), Article 4. https://doi.org/10.1038/nrn1366

Wittkuhn, L., & Schuck, N. W. (2021). Dynamics of fMRI patterns reflect sub-second activation sequences and reveal replay in human visual cortex. *Nature Communications*, *12*(1), Article 1. https://doi.org/10.1038/s41467-021-21970-2

Wurm, M. F., & Caramazza, A. (2022). Two 'what' pathways for action and object recognition. *Trends in Cognitive Sciences*, *26*(2), 103–116. https://doi.org/10.1016/j.tics.2021.10.003

Wurm, M. F., Caramazza, A., & Lingnau, A. (2017). Action Categories in Lateral Occipitotemporal Cortex Are Organized Along Sociality and Transitivity. *Journal of Neuroscience*, *37*(3), 562–575. https://doi.org/10.1523/JNEUROSCI.1717-16.2016

Yamins, D. L., Hong, H., Cadieu, C., & DiCarlo, J. J. (2013). Hierarchical Modular Optimization of Convolutional Networks Achieves Representations Similar to Macaque IT and Human Ventral Stream. *Advances in Neural Information Processing Systems*, *26*. https://proceedings.neurips.cc/paper/2013/hash/9a1756fd0c741126d7bbd4b692ccbd91-Abstract.html

Yamins, D. L. K., Hong, H., Cadieu, C. F., Solomon, E. A., Seibert, D., & DiCarlo, J. J. (2014). Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences*, *111*(23), 8619–8624. https://doi.org/10.1073/pnas.1403112111

Yildirim, I., Wu, J., Kanwisher, N., & Tenenbaum, J. (2019). An integrative computational architecture for object-driven cortex. *Current Opinion in Neurobiology*, *55*, 73–81. https://doi.org/10.1016/j.conb.2019.01.010

Zhang, Y., Brady, M., & Smith, S. (2001). Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Transactions on Medical Imaging*, *20*(1), 45–57. https://doi.org/10.1109/42.906424

Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., & Torralba, A. (2018). Places: A 10 Million Image Database for Scene Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *40*(6), 1452–1464. https://doi.org/10.1109/TPAMI.2017.2723009

Zhou, M., Gong, Z., Dai, Y., Wen, Y., Liu, Y., & Zhen, Z. (2023). A large-scale fMRI dataset for human action recognition. *Scientific Data*, *10*(1), Article 1. https://doi.org/10.1038/s41597-023-02325-6