


Statistically efficient neural encoding of natural object variability shapes the temporal dynamics of visual processing

David M. Watson^{a,b,*} , Richard Aveyard^b, Timothy J. Andrews^{a,b}

^a Department of Psychology, University of York, York, UK

^b York Neuroimaging Centre, University of York, York, UK

ARTICLE INFO

Keywords:

Object perception
Visual perception
Neural encoding
Data-driven
EEG
MEG

ABSTRACT

Object perception unfolds dynamically over millisecond timescales, yet the organisational principles that shape the emerging neural responses are not fully understood. Traditional hypothesis-driven approaches risk constraining interpretations by focusing on pre-selected object features. To circumvent this limitation, we applied a data-driven framework to behavioural and neuroimaging data obtained from the THINGS initiative, which provides a systematic sampling of real-world objects. Behaviourally relevant stimulus dimensions were derived from prior large-scale similarity judgements, offering an unbiased, ecologically grounded representation of object space. Using Partial Least Squares Regression (PLSR), we generated neural encoding models to predict time-resolved evoked responses in EEG and MEG from these dimensions. Across both modalities, the PLSR identified a small set of latent components that reliably captured the temporal dynamics of the neural activity. These components were similar across the EEG and MEG datasets and with a prior MRI analysis. The components encoded a diverse range of object features, including visual and semantic properties, yet did not map straightforwardly onto canonical accounts of visual cortical organisation. Instead, our findings suggest that object representations in the brain are structured by principles of statistical efficiency, capturing the co-occurrence of features amongst natural variability in real-world objects to support dynamic visual processing.

1. Introduction

The perception and recognition of objects is fundamental to visual cognition, underpinning our ability to interact with the world around us (Bracci and Op de Beeck, 2023; Peelen and Downing, 2017). In the human brain, object representations emerge along a hierarchical pathway, with early stages encoding low-level features (Hubel and Wiesel, 1968; Wandell and Winawer, 2011) and later stages representing higher-level object properties (Goodale and Milner, 1992; Ungerleider and Haxby, 1994). However, the organisational principles governing these higher-level representations remain the subject of ongoing debate. While classical accounts propose categorical selectivity in higher visual areas (Kanwisher, 2010), more recent perspectives suggest that object representations may be better explained by continuous feature dimensions (Grill-Spector and Weiner, 2014; Ritchie et al., 2025), including real-world size (Konkle and Oliva, 2012), animacy (Kriegeskorte, Mur, Ruff, et al., 2008), shape (Bao et al., 2020), low-level visual features (Andrews et al., 2015; Levy et al., 2001), and semantic properties (Huth et al., 2012).

A limitation of conventional neuroimaging approaches is their reliance on limited, experimenter-defined stimulus sets, often assigned to discrete conditions. This constrains the scope of interpretations and risks obscuring the true organising principles of object representation. Data-driven approaches provide a powerful alternative by combining rich, high-dimensional stimulus spaces with computational models to reduce experimenter bias (Brunton and Beyeler, 2019; Krishnan et al., 2011). The THINGS database (Hebart et al., 2019, 2023) exemplifies this approach, offering a systematic sampling of real-world objects paired with large-scale behavioural and neural datasets. A data-driven analysis of behavioural similarity judgements derived from this database revealed a set of behaviourally relevant perceptual dimensions that characterise object space without imposing experimenter preconceptions on the dimensions selected (Hebart et al., 2020). Furthermore, recent work has demonstrated that linear encoding models can predict fMRI responses based on sparse and distributed representations of these dimensions (Contier et al., 2024).

Extending this work, we recently demonstrated that Partial Least Squares Regression (PLSR) applied to the THINGS data generates a small

* Corresponding author.

E-mail address: david.watson@york.ac.uk (D.M. Watson).

<https://doi.org/10.1016/j.neuroimage.2026.122012>

Received 25 November 2025; Received in revised form 19 May 2026; Accepted 19 May 2026

Available online 20 May 2026

1053-8119/© 2026 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

number of latent components that can predict fMRI responses across both high- and low-level visual areas using continuous and graded representations (Watson and Andrews, 2025). These latent components captured diverse object features, spanning low-level visual properties to higher-level semantic features. However, these components did not directly align with current theories on the organisation of the visual brain. Instead, the components reflected combinations of object features including (though not limited to) animacy, real-world size, shape, texture, colour, and category. This indicated that visual representations may be shaped by principles of statistical efficiency – encoding the co-occurrence of features evident in natural variability amongst real-world objects.

While fMRI reveals the spatial organisation of object representations, it lacks the temporal resolution to capture the fast dynamics of visual processing. Time-resolved methods, such as EEG and MEG, can reveal unfolding visual representations on millisecond timescales. Previous studies of object perception have demonstrated that low-level visual features dominate early EEG/MEG responses, while higher-level categorical representations emerge later (Carlson et al., 2013; Cichy et al., 2014, 2016; Coggan et al., 2016). Recently, Teichmann et al. (2026) demonstrated that stimulus dimensions derived from behavioural similarity judgements of objects from the THINGS database could be predicted by MEG responses over sustained temporal windows. Nevertheless, it remains unclear whether neural responses can be explained by a more compact, statistically efficient representation.

In this study, we address this question by applying PLSR to behavioural, EEG, and MEG data from the THINGS initiative (Grootswagers et al., 2022; Hebart et al., 2020, 2023). We identified a small number of latent components from the THINGS stimulus dimensions that accurately predicted time-resolved neural responses to naturalistic objects across both EEG and MEG. These components revealed a diverse array of visual and semantic features shaping neural activity across early and sustained temporal windows. Comparisons with prior fMRI analyses demonstrated a shared underlying structure of object representations across spatial and temporal dimensions. Together, these findings suggest that the human visual system exploits principles of statistical efficiency to encode the natural variability of objects, supporting behaviourally relevant distinctions during dynamic visual processing.

2. Methods

2.1. Datasets and preprocessing

All data were obtained from the publicly available THINGS initiative (Hebart et al., 2019; <https://things-initiative.org/>). All participants provided informed consent, and the studies were approved by the NIH Office of Human Subjects Research Protections (Hebart et al., 2020, 2023) and the University of Sydney ethics committee (Grootswagers et al., 2022).

2.1.1. EEG dataset

We obtained EEG data from the publicly available THINGS-EEG1 dataset (Grootswagers et al., 2022; <https://openneuro.org/datasets/ds003825>). This includes EEG recordings from 48 subjects viewing 22, 248 images, comprising 12 images from each of the 1854 object concepts in the THINGS stimulus set (Hebart et al., 2019). Images were presented at 10 Hz in a rapid serial visual presentation (RSVP) stream, such that images are forward and backward masked by the preceding and following images. Subjects viewed all images in a single recording session.

EEG data were recorded with a 64-channel BrainVision ActiChamp system and digitised at 1000 Hz. Full methodological details are provided in Grootswagers et al. (2022). We obtained the raw EEG recordings and applied our own preprocessing pipeline. All processing was performed in MNE-Python (Gramfort et al., 2013, 2014). Data were re-referenced to the average response over channels. We then applied

notch-filters to remove power line noise at 50 Hz and its harmonics. Bad channels were manually annotated and replaced by interpolation with neighbouring channels. We then applied independent components analysis (ICA) denoising. The ICLabel toolbox (Pion-Tonachini et al., 2019), implemented in MNE-ICLabel (Li et al., 2022), was used to automatically identify components associated with ocular, heart, and muscle artifacts and remove them from the data. The data were then bandpass filtered between 0.1 and 80 Hz. We then epoched the data from -100 to 700 ms relative to stimulus onset and downsampled to 250 Hz (201 timepoints). Because the stimuli are presented within an RSVP stream, there is no prestimulus period with which to perform a baseline correction. Instead, we removed the DC component from each epoch, equivalent to performing a baseline correction by mean subtraction over the whole epoch. Finally, evoked responses to each object concept were estimated by averaging over the 12 epochs within each concept.

2.1.2. MEG dataset

We also obtained MEG data from the publicly available THINGS-MEG dataset (Hebart et al., 2023; <https://openneuro.org/datasets/ds004212>). This includes MEG recordings from 4 subjects, viewing the same images as in the EEG dataset (22,248 images comprising 12 images from each of 1854 object concepts) obtained from the THINGS stimulus set (Hebart et al., 2019). Images were presented in an event-related design over multiple scan sessions. Sensor positions relative to the head were kept consistent over scan sessions using subject-specific head casts.

MEG data were recorded on a 275-channel CTF system and digitised at 1200 Hz. Four noisy channels were removed prior to further analysis, leaving 271 channels. We obtained data that had already undergone initial preprocessing by the THINGS initiative. In brief, data were bandpass filtered between 0.1 and 40 Hz and epoched from -100 to 1300 ms relative to stimulus onset. Data were then baseline corrected by subtracting the mean and dividing by the standard deviation of the pre-stimulus period in each epoch, and finally downsampled to 200 Hz. Full details of the data acquisition and initial preprocessing are provided in Hebart et al. (2023). We then additionally downsampled the data to 100 Hz (140 time points) to reduce computational load for the encoding analyses and estimated evoked responses to each object concept by averaging over the 12 images within each concept.

2.2. Neural encoding models

2.2.1. Overview

We employed a neural encoding approach to predict time-resolved evoked responses in the EEG and MEG datasets. Stimulus features of the object concepts were obtained from a previous data-driven analysis of odd-one-out similarity judgements between triplets of images drawn from the THINGS dataset (Hebart et al., 2020). These dimensions capture behaviourally relevant properties of the objects spanning low-, mid-, and high-level features, without imposing experimenter pre-conceptions on the selection of those properties. We used an updated version of this model comprising 66 stimulus dimensions (Hebart et al., 2023).

We applied both linear and partial least squares neural encoding models to map these stimulus features to the neural responses. Encoding models were applied within individual subjects because the EEG and MEG channels are not guaranteed to be directly comparable over subjects. The 1854 object concepts are represented as samples within a 66-dimensional space defined by the stimulus model. These same object concepts are also represented by their evoked responses within a neural space defined by the concatenation of all timepoints and channels. This neural space comprised 12,864 dimensions for the EEG dataset (64 channels \times 201 timepoints) and 37,940 dimensions for the MEG dataset (271 channels \times 140 timepoints). The 1854 object concepts were split by thirds into a training set comprising the first 1236 concepts and a test set comprising the final 618 concepts. Samples were z-scored along each

of the stimulus and neural dimensions of the training set, such that each dimension has zero mean and unit variance over the training samples and all model coefficients are in standardised units. When predicting responses for the test set, the dimensions were normalised by the means and standard deviations derived from the training set. All encoding analyses were implemented in *scikit-learn* (Abraham et al., 2014; Pedregosa et al., 2011).

2.2.2. Linear model

We first employed a linear encoding model that maps the stimulus dimensions directly to the neural responses. A linear regression was fit to

samples in the training set, including the stimulus dimensions (X_{train}) as predictor variables and the neural dimensions (Y_{train}) as outcome variables. This derives a single matrix of regression coefficients, mapping the 66 stimulus dimensions to each of the neural dimensions (Fig. 1a).

2.2.3. Partial least squares model

The linear regression provides a baseline estimate of the ability of the stimulus dimensions to predict the neural responses. However, the resulting regression coefficients are high-dimensional and contain redundancy due to correlations amongst both the stimulus and neural dimensions. To provide a more parsimonious account, we next applied a

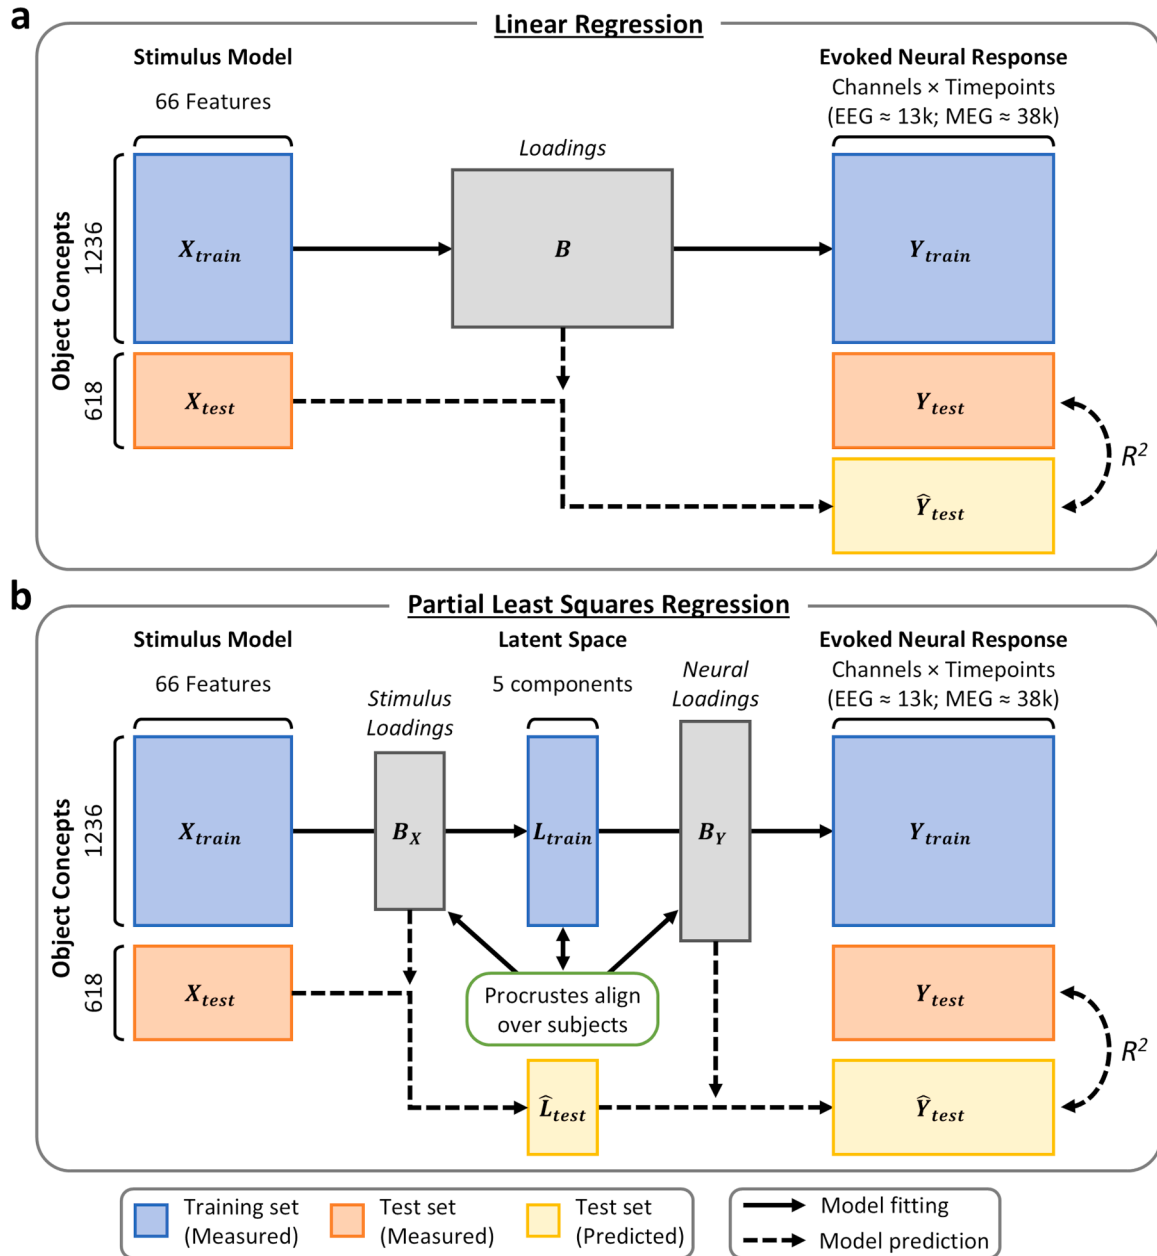


Fig. 1. Schematic of neural encoding models. Samples comprise 1854 object concepts split into a training set of 1236 concepts and a test set of 618 concepts. Each object concept is measured by its representation across the 66 features of the stimulus model and by its evoked neural response across the concatenation of EEG/MEG channels and timepoints. Neural encoding models are fit to object concepts in the training set to map the stimulus model (X_{train}) to the neural responses (Y_{train}). The fitted encoding models then project object concepts in the test set from the stimulus model (X_{test}) to the neural space (Y_{test}). Prediction accuracy is measured by the coefficient of determination (R^2) between predicted and measured neural responses in the test set. (a) Linear regression derives a direct mapping from the stimulus model to the neural responses. (b) Partial least squares regression derives an intermediate low-dimensional ($k = 5$) latent space (L_{train}). Latent components comprise linear combinations of the stimulus and neural features which maximise prediction of the neural responses. The components can be aligned over subjects via Procrustes analysis of the latent scores to facilitate examination of the latent scores and stimulus and neural loadings.

neural encoding approach using Partial Least Squares Regression (PLSR), which incorporates dimensionality reduction into the regression process (Fig. 1b). PLSR inserts an additional low-dimensional latent space (L_{train}) between the stimulus and neural spaces. The components of this space are linear combinations of both the stimulus and neural dimensions and are defined to capture dimensions of cross-covariance between both feature spaces which optimally predict the neural responses. The components are ordered by their predictive power of the neural response, such that the first component is maximally predictive, the second component is second most predictive, and so forth. Here, we chose to retain the first 5 latent components following testing of the prediction accuracies and cross-subject alignment of components over varying numbers of components (see below). The PLSR derives a matrix of stimulus loadings (B_X) describing the relationship between the 66 stimulus dimensions and the 5 latent components, and a matrix of neural loadings (B_Y) describing the relationship between the 5 latent components and the neural dimensions.

Although we expect representations within the latent space to show some commonality between subjects, the components are not guaranteed to be directly comparable across subjects. For instance, components may emerge in a different order, with a different sign, or a given component in one subject might correspond to a mixture of components in another subject. To align the components over subjects, we applied a Generalised Procrustes analysis or “hyperlignment” procedure (Gower, 1975; Haxby et al., 2011) to the latent scores in the training set. Initially, the first subject is defined as a reference subject. Procrustes analysis (allowing translations, rotations, reflections, and scaling) aligns the latent scores in each subject to the first subject. A new reference is then defined by averaging the aligned scores over subjects, and the latent scores in each subject are realigned to this new reference. This process is repeated until the change in disparity over iterations falls below a threshold. Note that no information is lost in this transformation – the components are simply rotated within the latent space. The resulting rotations can then be further applied to the stimulus and neural loadings to bring the components into alignment over subjects. The aligned stimulus loadings and latent scores were then averaged over subjects. We did not average the neural loadings because the EEG and MEG channels are not directly comparable over subjects.

We performed two analyses to validate the number of latent components to retain. First, we measured the model prediction accuracies across a 3-fold cross-validation nested within the main training set, varying the number of components from 1 to 10 (Supplementary Figure 1a-b). On each iteration, we measured the maximum R^2 score (over channels and timepoints) between the predicted and measured neural responses, representing the peak performance over all neural dimensions. For the EEG dataset the results indicated generally increasing prediction accuracies with more components, but did not reveal any clear knee-points. For the MEG dataset the values indicated knee-points between 3 and 6 components across subjects. Second, we measured the correlations of latent scores within and between subjects and components following Procrustes alignment (Supplementary Figure 1c-d). In both datasets, retaining 4 or 5 components yielded good cross-subject alignment, indicated by high within-component correlations and low between-component correlations. We therefore chose to retain 5 latent components. This is also consistent with our previous analysis of the THINGS-fMRI dataset (Watson and Andrews, 2025).

2.2.4. Out-of-sample prediction accuracies

For both the linear and partial least squares encoding models, the regression coefficients derived from the training set were applied to samples in the test set to predict the evoked neural responses (\hat{Y}_{test}) from the stimulus dimensions (X_{test}). Prediction accuracy was assessed by computing the coefficient of determination (R^2) between the predicted (\hat{Y}_{test}) and measured (Y_{test}) evoked responses to object concepts in the test set at each channel and timepoint. This provides an independent

estimate of the model prediction accuracy in samples held out from the model fitting.

Statistical significance was assessed using a maximum-statistic permutation procedure which automatically controls the familywise error rate. On each permutation, we randomly shuffled the order of object concepts within the training set for the stimulus model space and refit the encoding models (linear and PLSR). These models were then used to predict the evoked responses to the unpermuted object concepts in the test set, such that the models were fit to permuted data but tested on unpermuted data. We then calculated the R^2 values between the predicted and measured evoked responses for the linear and PLSR models, as well as the difference between the models. We retained the largest R^2 value over channels and timepoints for the linear and PLSR maps, and the largest absolute R^2 value for the difference map. This procedure was repeated 1000 times to generate an empirical null distribution for each map. Building null distributions from the signed values of the linear and PLSR maps provides one-tailed tests, while using the absolute values of the difference map provides a two-tailed test. The 95th percentile of these distributions provides the threshold for statistical significance while controlling the familywise error rate (one-tailed $p < .05$ for the linear and PLSR models, and two-tailed $p < .05$ for the difference).

2.3. Source localisation of MEG components

For the MEG dataset, we derived source localisations of the PLSR latent components. We used Freesurfer (Dale et al., 1999) to reconstruct cortical surfaces from T1-weighted anatomical MRI scans of each subject. Freesurfer’s watershed algorithm (Ségonne et al., 2004) segmented the outer brain surface, inner and outer skull surfaces, and outer scalp surface. MEG sensors were manually co-registered to the head using the marker coils attached to the head casts over the nasion and left and right preauricular pits. A forward solution was computed sampling vertices along the white-matter cortical surface and using a one-layer boundary element model constructed from the inner skull surface. The noise covariance between sensors was computed from the prestimulus periods of the epoched data. We then computed an inverse solution loosely constrained by the cortical surface orientation (weighted 80% towards the surface orientation, and 20% towards free orientations). Source localisations were then computed using a minimum norm estimate including dSPM noise normalisation (Dale et al., 2000; Hämäläinen and Ilmoniemi, 1994).

The neural loadings represent the expected standard deviations of change in the evoked response given a one-unit change in the latent score. Consequently, positive and negative loadings do not have the same interpretation as in the evoked response itself and thus are not suitable for source localisation directly. Instead, we computed source localisations of the evoked responses to the 1236 object concepts in the training set. For each latent component, the source-space evoked responses to each object concept were weighted according to their corresponding latent scores and averaged together. This provides a parametric contrast of the evoked responses along each component: responses to positively scoring concepts are positively upweighted, responses to negatively scoring concepts are negatively upweighted, and responses to low scoring concepts are downweighted towards zero. These weighted activations were then transformed from each subject’s cortical surface to the *fsaverage* surface using a surface-based registration (Fischl et al., 1999) and averaged over subjects.

Source localisations were visualised on cortical flat patches using *pycortex* (Gao et al., 2015). To aid visualisation, we annotated regions of interest for V1 and V2 taken from a retinotopic atlas (Benson et al., 2014), and scene- and face-selective regions (PPA, OPA, RSC, FFA, OFA, pSTS) taken from an independent functional localiser (Noad et al., 2024).

2.4. Object property ratings

To provide additional descriptives of the PLSR latent components, we obtained ratings for 12 object properties from the THINGSplus metadata (Stoinski et al., 2023). These provide ratings for the object size, and how natural, living, moving, manmade, precious, heavy, graspable, holdable, moveable, pleasant, and arousing the object is. Ratings were averaged over images at the object-concept level. We then computed Spearman's correlations between these ratings and the latent scores along each component for the 1236 object concepts in the training set. A Holm-Bonferroni correction was applied across the resulting 60 comparisons (Holm, 1979).

2.5. DCNN representations

To provide an external validation of the PLSR models, we conducted representational similarity analyses (RSAs; Kriegeskorte, Mur, and Bandettini, 2008) comparing the representations of object concepts in the latent space to activations within layers of an Alexnet deep convolutional network (DCNN) pre-trained for object classification on the ImageNet database (Krizhevsky et al., 2017). Images from the THINGS stimulus set (Hebart et al., 2019) for each of the 1236 object concepts in the training set were passed through the DCNN. The corresponding activations were extracted from each convolutional and fully connected layer and averaged over images within each object concept. Finally, pairwise correlation distances were measured between each object concept, producing a 1236-by-1236 representational dissimilarity matrix (RDM) for each DCNN layer.

We next calculated RDMs for the object concept representations within the PLSR latent space. First, we computed an RDM for the overall space by measuring pairwise correlation distances between object concepts from the 5-dimensional latent scores. Second, we computed per-component RDMs by measuring the pairwise absolute differences (equivalent to Euclidean distances) between object concepts along each component in turn. The PLSR RDMs (for the full space and each component) were then correlated with the RDMs for each DCNN layer. A maximum statistic permutation procedure was used to assess statistical significance. On each permutation, the order of the object concepts was randomly shuffled (corresponding to reordering the rows and columns of the RDMs). The RSA correlations were then recomputed for all DCNN layers (and PLSR components, if applicable), and the maximum correlation over all comparisons was recorded. This procedure was repeated 5000 times to derive an empirical null distribution. The 95th percentile of this distribution provides the threshold for statistical significance (one-tailed $p < .05$), controlling the familywise error rate over DCNN layers (and PLSR components, if applicable).

2.6. Comparing components between datasets

Finally, to examine the robustness of the PLSR models, we compared the latent components between the THINGS-EEG1, THINGS-MEG, and THINGS-fMRI datasets. Latent scores for the THINGS-fMRI dataset were obtained from our previous analysis of whole-brain fMRI responses (Watson and Andrews, 2025). In each case, latent scores for object concepts within the training set were correlated within and between components across the datasets. Comparisons between the EEG and MEG dataset included all 1236 training concepts in the current study, while comparisons with the fMRI dataset were restricted to the 480 concepts used in the training set of that dataset. Because components are not guaranteed to emerge in the same order or direction across datasets, we adopted an iterative procedure to reorder and sign-flip components to match them between datasets. For a given pairing of datasets, we identified the largest absolute value within the correlation matrix. If necessary, the matrix columns were swapped to place this value on the matrix diagonal. In addition, if the value was negative then the values within the matrix column were sign-flipped. The corresponding matrix

row and column were then removed from the search, and the process was repeated for the remaining components. This places the largest correlations along the matrix diagonal where possible.

3. Results

3.1. Reliability of encoding models

Time-resolved evoked neural responses to 1854 object concepts from the THINGS initiative were measured using EEG and MEG (Grootswagers et al., 2022; Hebart et al., 2019, 2023). Each object concept was additionally described by 66 stimulus dimensions obtained from a data-driven analysis of behavioural perceptual similarity judgements (Hebart et al., 2020). Partial least squares regression (PLSR) was then used to generate 5 latent components that predicted the evoked responses from the stimulus features. To evaluate the efficacy of the PLSR, we also compared its reliability against a linear regression model using all 66 stimulus dimensions.

The PLSR and linear encoding models were trained on 1236 object concepts, then we quantified the reliability by measuring each model's accuracy at predicting evoked responses to 618 object concepts held out from the model fitting. Prediction accuracy was measured by computing the coefficient of determination (R^2) between predicted and measured evoked responses for each channel and timepoint. Statistical significance, controlling for the familywise error rate, was determined by maximum-statistic permutation tests. Fig. 2 illustrates the prediction accuracies of the linear and PLSR encoding models, and the difference between them. The scalp topographies illustrate R^2 values at discrete time points for two subjects from the EEG (Fig. 2a) and MEG datasets (Fig. 2b). Animations of the R^2 scalp topographies are shown for 4 EEG subjects in Supplementary Video 1 and for all MEG subjects in Supplementary Video 2. Both the linear and PLSR encoding models accurately predicted EEG and MEG evoked responses. In the EEG dataset, both models achieved significant prediction accuracies within a relatively brief period from approximately 100 to 300 ms after stimulus onset. For the MEG dataset, both models achieved significant prediction accuracies across a longer period from approximately 100 to 800 ms after stimulus onset. The more transient nature of the encoding performance in the EEG dataset likely reflects forward and backward masking effects from the rapid serial visual presentation stream used to present the stimuli. In contrast, the event-related design used in the MEG dataset elicited more sustained responses.

Furthermore, the PLSR achieved comparable performance to the linear regression despite the substantial dimensionality reduction from 66 stimulus features down to just 5 latent components. Kernel density estimates show the distributions of the difference in R^2 values between the linear and PLSR models in each subject in the EEG (Fig. 2c) and MEG datasets (Fig. 2d), for all channels and timepoints within windows when significant prediction accuracies were observed. The linear and PLSR models achieved comparable performance for all subjects in both the EEG and MEG datasets. There was a slight bias towards better performance for the PLSR model, however none of difference values reached significance in the EEG dataset, and only one value reached significance in the MEG dataset. This demonstrates the PLSR successfully captured a low-dimensional representation of the stimulus features which was predictive of the neural response.

3.2. Object representations on PLSR components

3.2.1. EEG dataset

We next inspected the object features represented on each PLSR latent component, beginning with the EEG dataset. Components were aligned across subjects by Procrustes analysis of the latent scores, and the aligned stimulus loadings and latent scores were then averaged over subjects. The stimulus loadings (matrix B_X in Fig. 1) indicate how the 66 dimensions in the stimulus model relate to the latent components.

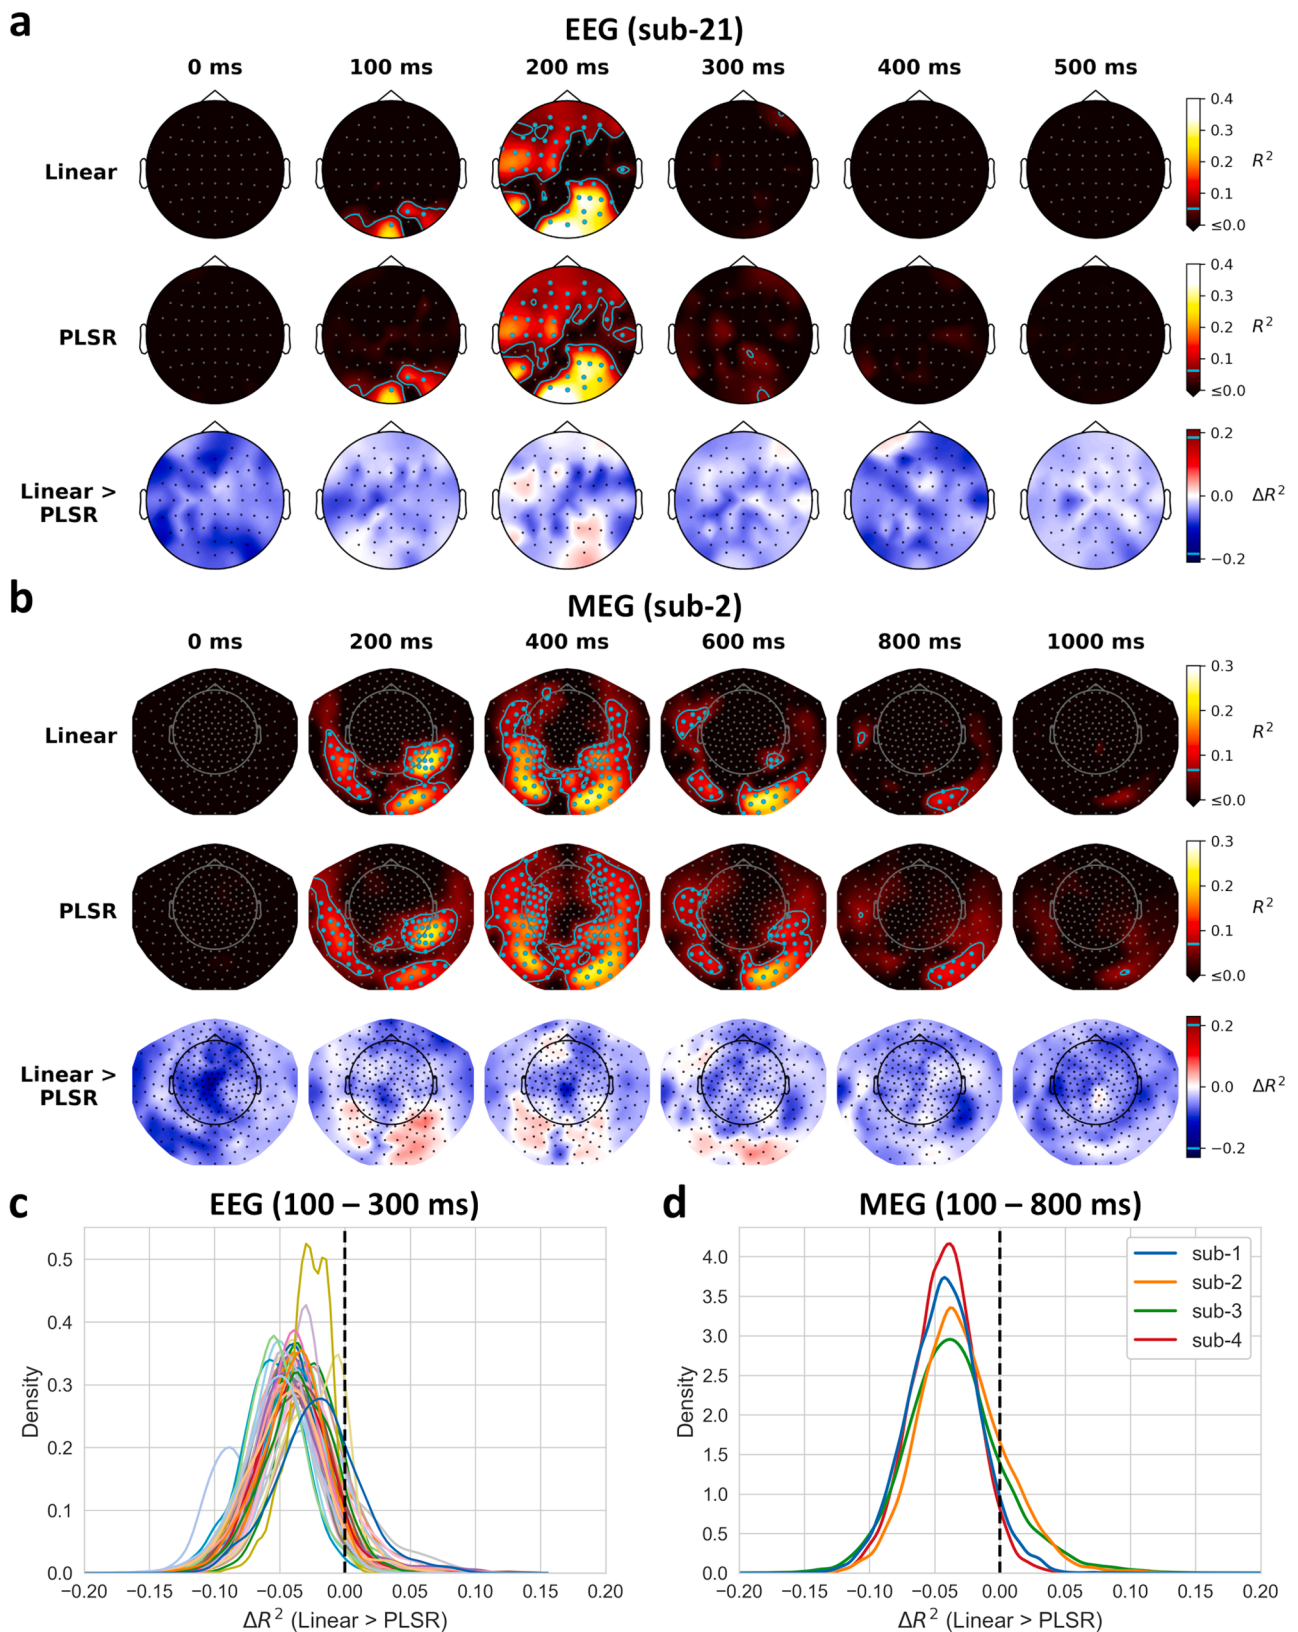


Fig. 2. Encoding model prediction accuracies (coefficients of determination; R^2) between predicted and measured evoked responses to object concepts held out from the model fitting. (a, b) Scalp topographies for two example subjects from the (a) EEG and (b) MEG datasets, illustrating R^2 values for the linear and PLSR encoding models and the difference between them. Positive difference values indicate better performance for the linear model. Cyan outlines and markers indicate channels and timepoints differing significantly from zero, as determined by permutation testing (one-tailed $p < .05$ for linear and PLSR R^2 ; two-tailed $p < .05$ for R^2 differences; all FWER corrected). Note that the R^2 difference values did not reach significance. (c, d) Distributions of R^2 difference values for each subject in the (c) EEG and (d) MEG datasets, across all channels and timepoints within time windows when prediction accuracies are high.

Table 1

Top and bottom five stimulus model dimensions loading on each PLSR component in the EEG dataset. Loading values are indicated in parentheses.

		Component				
		1	2	3	4	5
Top 5	1	Plant-related (0.35)	Body- / people-related (0.34)	Household (0.21)	Colourful / playful (0.26)	Outdoors (0.27)
	2	Food-related (0.32)	Animal-related (0.32)	Fluid- / drink-related (0.20)	Coarse pattern / many things (0.25)	Long / thin (0.24)
	3	Repetitive / spiky (0.20)	Body part-related (0.31)	White (0.20)	Feminine (stereotypical) (0.22)	Beams- / mesh-related (0.19)
	4	Coarse pattern / many things (0.18)	Head-related (0.22)	Bathroom- / wetness-related (0.18)	Oriented / many things (0.18)	Tools / handheld / elongated (0.19)
	5	Animal-related (0.18)	Foot- / walking-related (0.18)	Thin / flat / wrapping (0.17)	Stick-shaped / container (0.14)	Pointed / spiky (0.18)
Bottom 5	5	Head-related (-0.18)	Oriented / many things (-0.19)	Weapon- / danger-related (-0.14)	Water-related (-0.15)	Household (-0.17)
	4	Body part-related (-0.19)	Box-related / container (-0.23)	Farm-related / historical (-0.16)	Weapon- / danger-related (-0.15)	Sweet / dessert-related (-0.17)
	3	Body- / people-related (-0.26)	Electronics / technology (-0.23)	Sports- / playing-related (-0.16)	Outdoors (-0.15)	Circular / round (-0.17)
	2	Textile (-0.27)	House- / furnishing-related (-0.25)	Animal-related (-0.16)	Animal-related (-0.15)	Fluid- / drink-related (-0.21)
	1	Metallic / artificial (-0.28)	Metallic / artificial (-0.29)	Transportation- / movement-related (-0.19)	Transportation- / movement-related (-0.19)	Food-related (-0.22)

Table 1 shows the top and bottom five stimulus dimensions loading on each component (see Supplementary Tables 1–5 for expanded lists of the top, middle, and bottom 10 stimulus dimensions). Meanwhile, the latent scores (matrix L_{train} in Fig. 1) represent the scoring of each object concept in the training set on the 5 latent components. Fig. 3 shows the top and bottom five scoring object concepts on each component (see Supplementary Figures 2–6 for expanded lists of the top, middle, and bottom 10 object concepts). To further quantify the behavioural relevance of each component, we also obtained human-generated ratings of 12 key object properties for each of the object concepts in the training set from the THINGSplus dataset (Stoinski et al., 2023). We then correlated these ratings with the latent scores to help provide additional descriptive details of each latent component (Fig. 4a).

The first component was positively associated with images containing highly textured and repetitive patterns, particularly those relating to plant- and food-related dimensions. Consistent with this, correlations with the object property ratings indicated a positive association with natural, living objects. The component was negatively associated with more tangible singular objects. The most negative scoring concepts included some people- and body-related objects; nevertheless, the stimulus loadings and property ratings indicated a broader association with larger, manmade, and artificial objects.

The second component was positively associated with objects and stimulus features related to bodies and faces, and with objects rated as more natural and living. Conversely, it was negatively associated with images of inanimate objects relating to artificial stimulus features, and objects rated as larger, heavier, and man-made.

The third component identified positive associations with small and artificial objects, particularly relating to household contexts, embodying manipulable object properties. The component was negatively associated with larger objects, particularly including objects relating to transportation and movement-related features, and objects rated as heavier and moving.

The fourth component was also positively correlated with smaller objects embodying manipulable properties, but now more so with colourful and textured images, as well as objects rated as more pleasant. The component was negatively correlated with objects in outdoor contexts, including large, heavy, and moving property ratings.

Finally, the fifth component was positively associated with thin, spindly, or spiky objects, particularly those in outdoor contexts. This included an association with ratings of how large, moving, heavy, and arousing the objects were. Conversely, the component was negatively associated with rounder objects, and food and drink-related features, as

well as ratings of how manipulable and pleasant the objects were.

3.2.2. MEG dataset

We next turned to the latent components for the MEG dataset. As before, the components were aligned over subjects by Procrustes analysis of the latent scores, then the stimulus loadings and latent scores were averaged over subjects. Table 2 shows the top and bottom five stimulus dimensions loading on each component (see Supplementary Tables 6–10 for expanded lists of the top, middle, and bottom 10 stimulus dimensions). Fig. 5 shows the top and bottom five object concepts scoring on each component (see Supplementary Figures 7–11 for expanded lists of the top, middle, and bottom 10 object concepts). Fig. 4b illustrates correlations between behavioural ratings of 12 key object properties and the latent scores along each component.

Similar to the EEG dataset, the first component was positively associated with highly textured objects and repetitive patterns, particularly including food- and plant-related objects, and was positively correlated with natural, living, and pleasant object property ratings. The component was again negatively associated with more tangible singular objects. The top scoring concepts included some objects related to faces and bodies, but the stimulus loading and property ratings again indicated a broader association with man-made and artificial objects.

Also similar to the EEG dataset, the second component was positively associated with animate objects, and with natural, living, and moving property ratings. However, unlike the EEG, this association was now more closely related to animals than humans. The component was also negatively associated with inanimate objects and stimulus features, and with objects rated as more manmade and manipulable.

The third component was positively correlated with long, thin, and stick-shaped or spiky objects and stimulus features, as well as properties relating to how manipulable the objects were. It was negatively correlated with colourful or textile-related objects and stimulus features, and more broadly with objects rated as having more manipulable and unpleasant properties.

The fourth component was positively associated with small objects, particularly those related to health, medicine, or body parts, and with objects rated as more manipulable. It was also negatively associated with artificial objects, especially those in outdoor contexts, and rated as being larger and heavier.

Finally, the fifth component was positively correlated with objects related to households or interior scene contexts, and negatively correlated with objects containing colourful and textured patterns and with thin and spiky objects. The component was less strongly correlated with

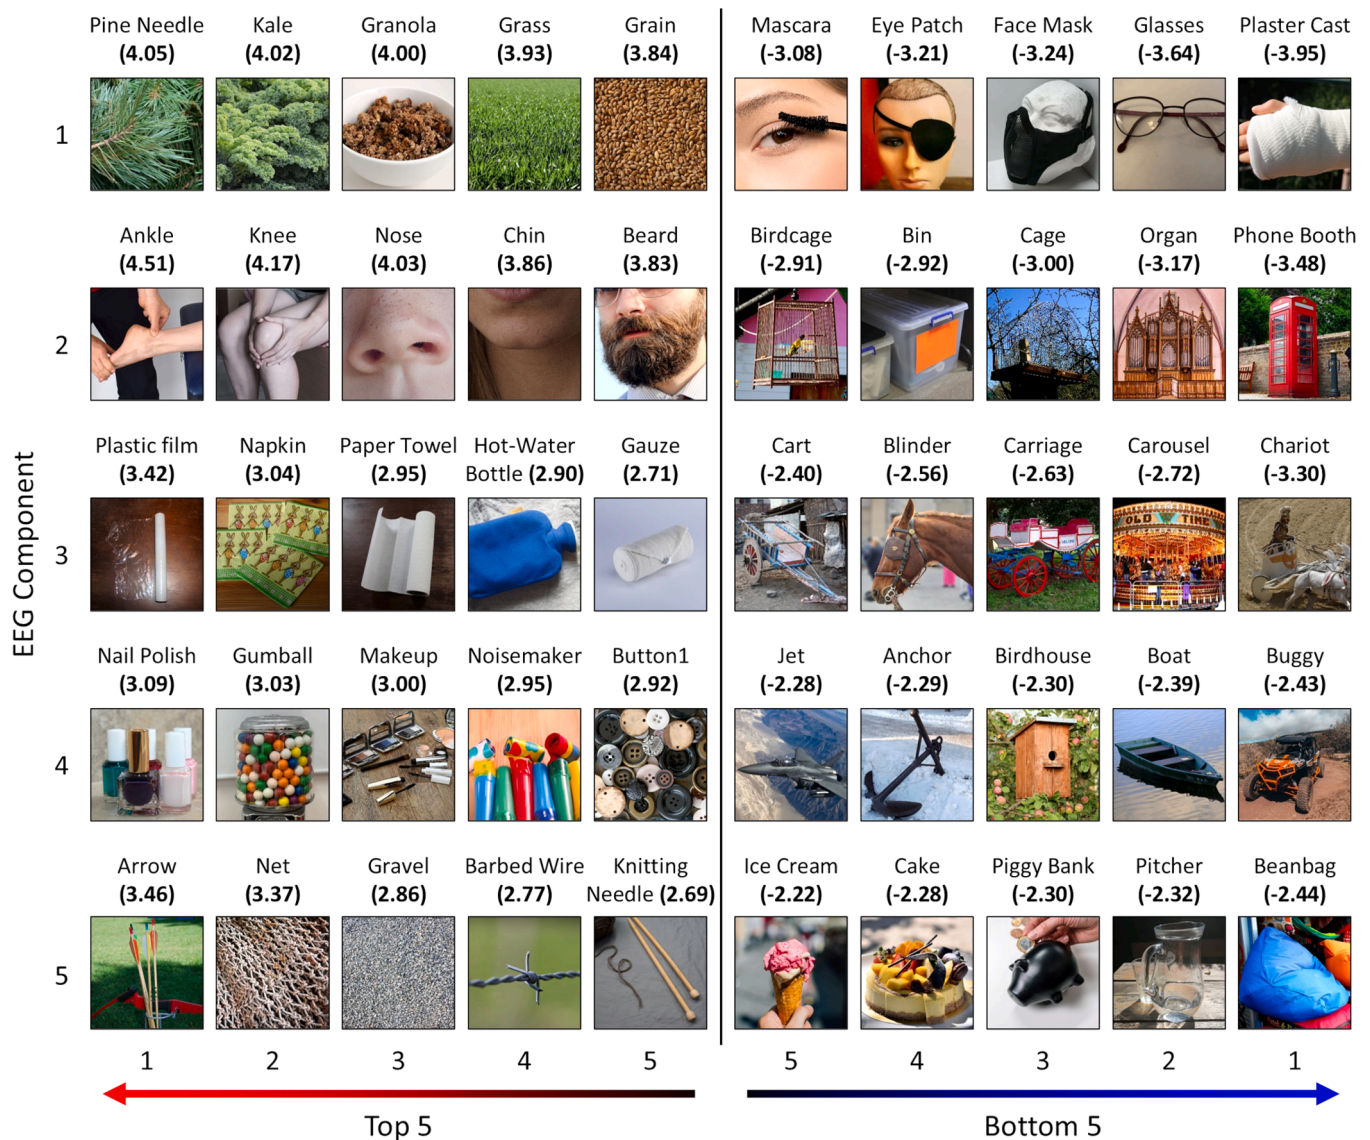


Fig. 3. Top and bottom five object concepts (within the training set) scoring on each PLSR component in the EEG dataset. Representative images are illustrated for each concept. Latent scores are indicated in parentheses. The images shown are distributed under a CC0 licence and are representative of those in the THINGS database.

the object property ratings, but nevertheless indicated moderate positive associations with small, natural, and living properties, and negative associations with manmade and manipulable properties.

Taken together, these results indicate that the PLSR components in both the EEG and MEG datasets captured variation in meaningful and tangible features of the objects. Such features included visual properties such as texture, size, shape, colour, and animacy, as well as more semantic and abstract properties. However, each component often embodied multiple object features, and equally different features were often evident across multiple components.

3.3. Neural representations of PLSR components

We next considered the Procrustes-aligned neural loadings (matrix B_T in Fig. 1), describing the relationship between the latent components and the evoked neural responses. Fig. 6 illustrates the timecourses of component loadings, along with scalp topographies at example time points, for representative subjects from the EEG (Fig. 6a) and MEG (Fig. 6b) dataset. Supplementary Video 3 shows animated scalp topographies for four representative subjects from the EEG dataset, and

Supplementary Video 4 shows animations for all MEG subjects.

Clear modulations in the loadings are evident over time. Loadings appear close to zero in the prestimulus period then show deflections from approximately 100 ms post stimulus onset. It is important to note that positive and negative loadings do not have the same interpretation as positive and negative potentials in the evoked responses themselves. Rather, they reflect the sensitivity to each latent component at specific time points. Modulations in the EEG dataset are often relatively transient due to the forward and backward masking effects of the RSVP stream. However, loadings for the MEG dataset often display a more sustained response including modulations into later time periods.

For the MEG dataset, we additionally source localised representations of the PLSR components. Evoked responses for each of the 1236 object concepts in the training set were projected to the cortical surface via a minimum norm estimate. The source space responses were then weighted by the corresponding latent scores on a given component and averaged over subjects. This represents a parametric contrast of responses along the component – positive and negative weights indicate stronger responses to concepts scoring more positively or negatively along the component respectively. These weighted activation maps were

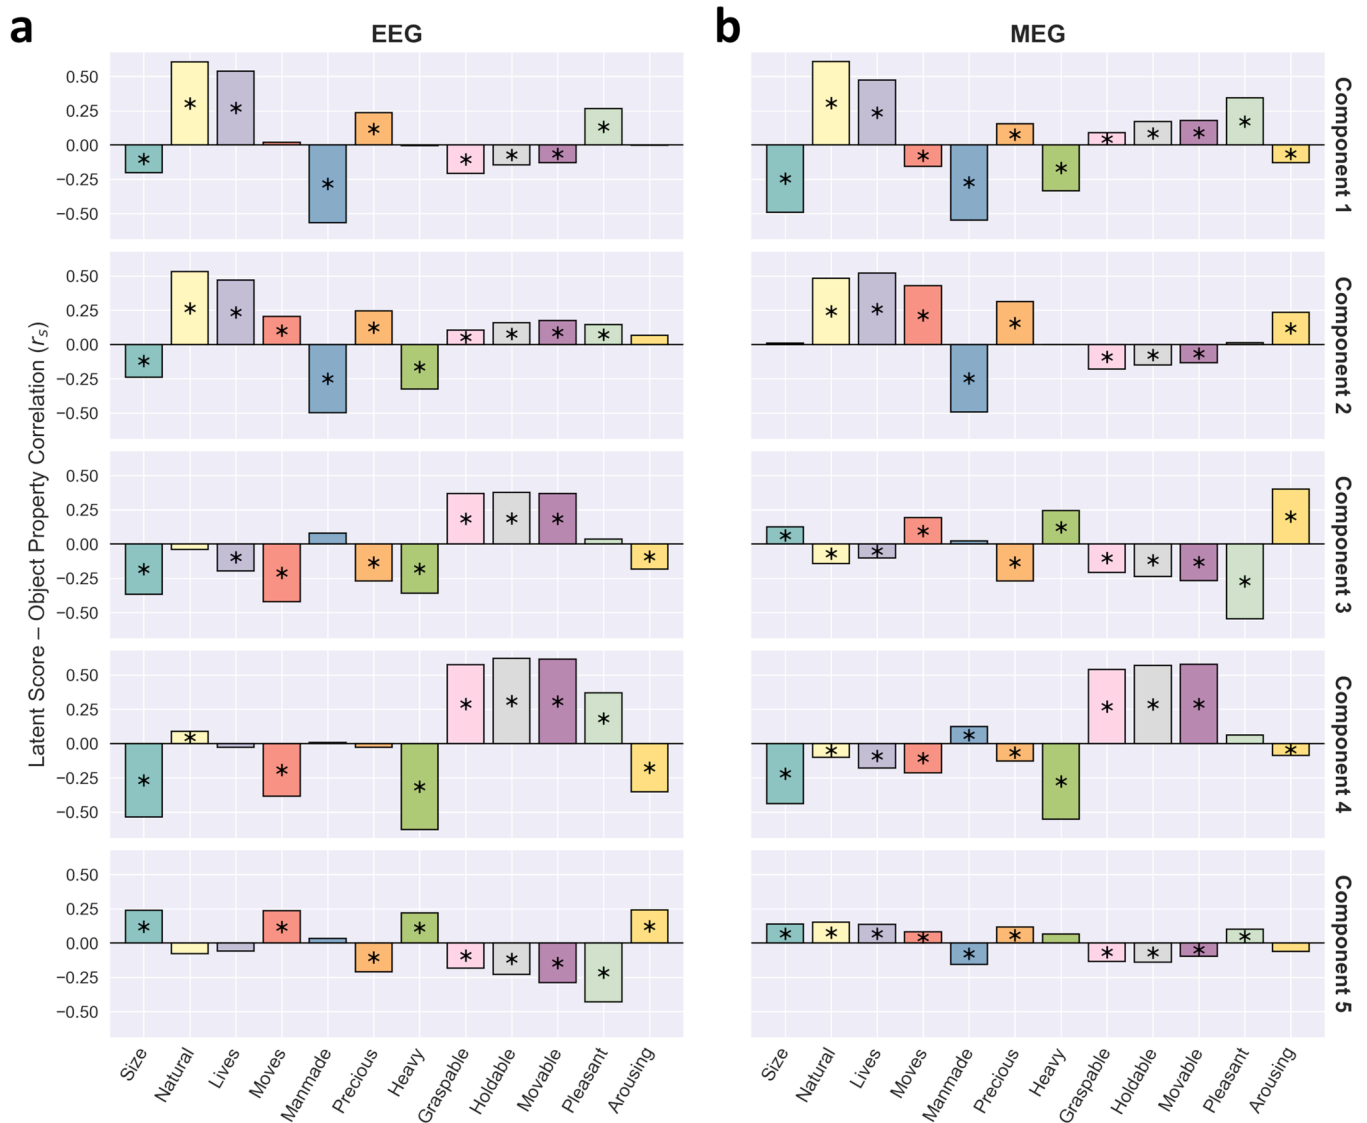


Fig. 4. Correspondence between PLSR components and key object properties for (a) EEG and (b) MEG datasets. Plots illustrate Spearman correlations between group average latent scores on each component and 12 object property ratings obtained from the THINGSplus metadata for object concepts in the training set. Note that components are not necessarily directly comparable between datasets. Bars marked with asterisks indicate significant correlations ($p < .05$, FWER corrected).

then averaged over subjects.

Fig. 7 illustrates group average weighted activations on cortical flat patches at 3 example time points. Supplementary Video 5 shows animations of the activations on the inflated cortical surface. All components were associated with strong weights in occipital regions at early time points around 100 ms post stimulus onset, followed by topographic patterns of weights across visual and non-visual cortices at later time points. For instance, within 300 ms of stimulus onset, component 1 identified a gradient between positive weights in early visual cortices and negative weights across ventrotemporal regions. Component 2 was associated with positive parietal and occipital weights alongside alternating patterns of positive and negative weights in ventral-temporal regions from approximately 200 ms post onset.

Meanwhile, component 3 identified patterns of positive and negative weights across ventrotemporal regions from approximately 200 ms post onset. Component 4 indicated positive weights in parietal regions, and negative weights in occipital, ventrotemporal, and ventral-frontal regions by 500 ms post onset. Finally, component 5 was associated with negative occipital weights and positive weights in temporal and parietal regions from approximately 300 ms post onset.

Taken together, these results demonstrate that the latent components

track modulations in the evoked neural responses over time, and these are associated with gradients of function across visual and non-visual brain regions.

3.4. Similarity between PLSR components and artificial neural networks

To provide a further external validation of the PLSR models, we next compared the features represented by the components to representations in a deep convolutional neural network (DCNN) (Alexnet; Krizhevsky et al., 2017) pretrained for object identification. We performed a series of representational similarity analyses (RSAs; Kriegeskorte, Mur, and Bandettini, 2008) comparing the PLSR latent scores to activations in each of the DCNN layers for the 1236 object concepts in the training set (Fig. 8).

We first performed RSAs for the full latent space. Representational dissimilarity matrices (RDMs) were constructed by measuring the correlation distance between object concepts from their 5-dimensional latent scores. RDMs for each DCNN layer were also constructed by averaging activations over images within each object concept, then measuring the correlation distance between concepts from the averaged activation vectors. The PLSR RDMs for the EEG and MEG datasets were

Table 2

Top and bottom five stimulus model dimensions loading on each PLSR component in the MEG dataset. loading values are indicated in parentheses.

		Component				
		1	2	3	4	5
Top 5	1	Food-related (0.39)	Animal-related (0.45)	Tools / handheld / elongated (0.30)	Body part-related (0.32)	White (0.19)
	2	Plant-related (0.29)	Body- / people-related (0.24)	Tubular (0.23)	Feminine (stereotypical) (0.29)	Household (0.19)
	3	Coarse pattern / many things (0.24)	Bug-related / non-mammalian / disgusting (0.24)	Metallic / artificial (0.23)	Stick-shaped / container (0.25)	House- / furnishing-related (0.16)
	4	Sweet / dessert-related (0.19)	Flying- / sky-related (0.16)	Construction- / craftsmanship-related (0.22)	Head-related (0.23)	Fluid- / drink-related (0.16)
	5	Repetitive / spiky (0.17)	Farm-related / historical (0.15)	Long / thin (0.20)	Medicine- / health-related (0.22)	Bathroom- / wetness-related (0.14)
Bottom 5	5	Sports- / playing-related (-0.20)	Oriented / many things (-0.21)	Valuable / precious (-0.17)	Plant-related (-0.14)	Weapon- / danger-related (-0.18)
	4	Electronics / technology (-0.20)	Household (-0.23)	Colourful / playful (-0.19)	Upright / elongated / voluminous (-0.15)	Beams- / mesh-related (-0.20)
	3	Transportation- / movement-related (-0.23)	Metallic / artificial (-0.24)	Feminine (stereotypical) (-0.19)	Beams- / mesh-related (-0.18)	Colourful / playful (-0.23)
	2	Body- / people-related (-0.24)	Paper-related / flat (-0.26)	Textile (-0.22)	Transportation- / movement-related (-0.24)	Oriented / many things (-0.24)
	1	Metallic / artificial (-0.28)	House- / furnishing-related (-0.28)	Body- / people-related (-0.23)	Outdoors (-0.29)	Coarse pattern / many things (-0.26)

then correlated with each of the DCNN layer RDMs. The RSA correlations for the EEG and MEG datasets are shown in Fig. 8a and Fig. 8b respectively. For both datasets, significant RSA correlations were observed between the PLSR latent space and all DCNN layers. Correlations increased over convolutional layers and peaked in the first fully connected layer, suggesting the latent space most prominently captured mid- to high-level visual features.

We also performed RSAs for each of the latent components individually. Per-component RDMs were calculated by measuring the Euclidean distance of latent scores between pairwise combinations of object concepts and then correlated with the DCNN RDMs. RSA correlations are shown for the EEG and MEG datasets in Fig. 8c and Fig. 8d respectively. The highest correlations were observed for the first two latent components, although all components displayed significant correlations with multiple DCNN layers. Correlations were generally highest for later convolutional layers and early fully connected layers, though the pattern varied over components and datasets. Taken together, these results demonstrate the latent components embody features which are diagnostic for object identification.

3.5. Similarity of PLSR components between datasets

Finally, to examine the robustness of the PLSR components, we correlated the latent scores within and between components across the EEG and MEG datasets for object concepts in the training set. We also computed the similarity for both datasets with latent scores for a subset of 480 object concepts obtained from our previous analysis of whole-brain fMRI responses (Watson and Andrews, 2025). Because components are not guaranteed to emerge in the same order across datasets, the components were reordered and sign-flipped as necessary to achieve the best correspondence between datasets. The resulting correlations are shown Fig. 9.

We observed strong correlations along the matrix diagonals, indicating clear commonalities between the components identified in each dataset, particularly for the first two components. A series of independent-samples *t*-tests confirmed that on-diagonal correlations were significantly higher than off-diagonal correlations (MEG-EEG: $t(6.00) = 4.74, p = .006$, Cohen's $d_s = 2.52$; fMRI-EEG: $t(5.34) = 5.96, p = .005$, Cohen's $d_s = 0.94$; fMRI-MEG: $t(5.23) = 3.50, p = .016$, Cohen's $d_s = 2.16$). High correlations were also frequently observed in the off-diagonal elements, indicating that features represented on a given component in one dataset were often shared with multiple components in the other datasets. Thus, there is a clear correspondence in the object

features captured by the latent spaces in each dataset, although there is not an exact one-to-one match between the components.

4. Discussion

We applied a data-driven approach to identify the stimulus dimensions underlying the perception and neural representation of real-world objects. Using publicly available datasets from the THINGS initiative (Grootswagers et al., 2022; Hebart et al., 2020, 2023), we used Partial Least Squares Regression (PLSR) to derive a low-dimensional representation of objects that reliably predicted time-resolved neural responses in both EEG and MEG. Despite differences in participants, imaging modalities, and task designs, PLSR identified shared representational structures across EEG, MEG, and fMRI datasets. This convergence suggests that the extracted latent dimensions capture general aspects of object representation in the human brain.

The PLSR model reduced the stimulus dimensions and neural responses down to 5 latent components, providing an efficient and powerful predictive account of the evoked neural responses. Prediction accuracy was robust across data splits and reached comparable explanatory power to linear models (Teichmann et al., 2026), but with greater parsimony. Thus, it was possible to explain a large proportion of variance in the neural response (up to almost 40%) from only 5 latent components. Importantly, unlike other dimensionality reducing methods such as Principal Components Analysis (Huth et al., 2012), PLSR retains the mapping between the stimulus samples and latent components. This allowed us to describe the selectivity of each of the latent components for different perceptual and neural properties. Furthermore, the latent components in PLSR are explicitly defined to predict the neural response and thus embody direct representations of the neural response. By comparison, a principal component decomposition – either of the stimulus features themselves or of the weights in a linear encoding model (Huth et al., 2012) – would not produce components directly linking the stimulus features to the neural response, and so would not provide this same direct interpretation. Hidden Markov models are a further data-driven approach often applied to MEG and EEG data to identify dynamic brain states (Baker et al., 2014). However, this approach operates entirely on the neuronal signal, whereas encoding models – including our PLSR method – allow explicitly linking stimulus features to neural representations.

In PLSR, the components are ordered by predictive power, so earlier components would retain the same information even if additional components were included. We used 5 latent components as these

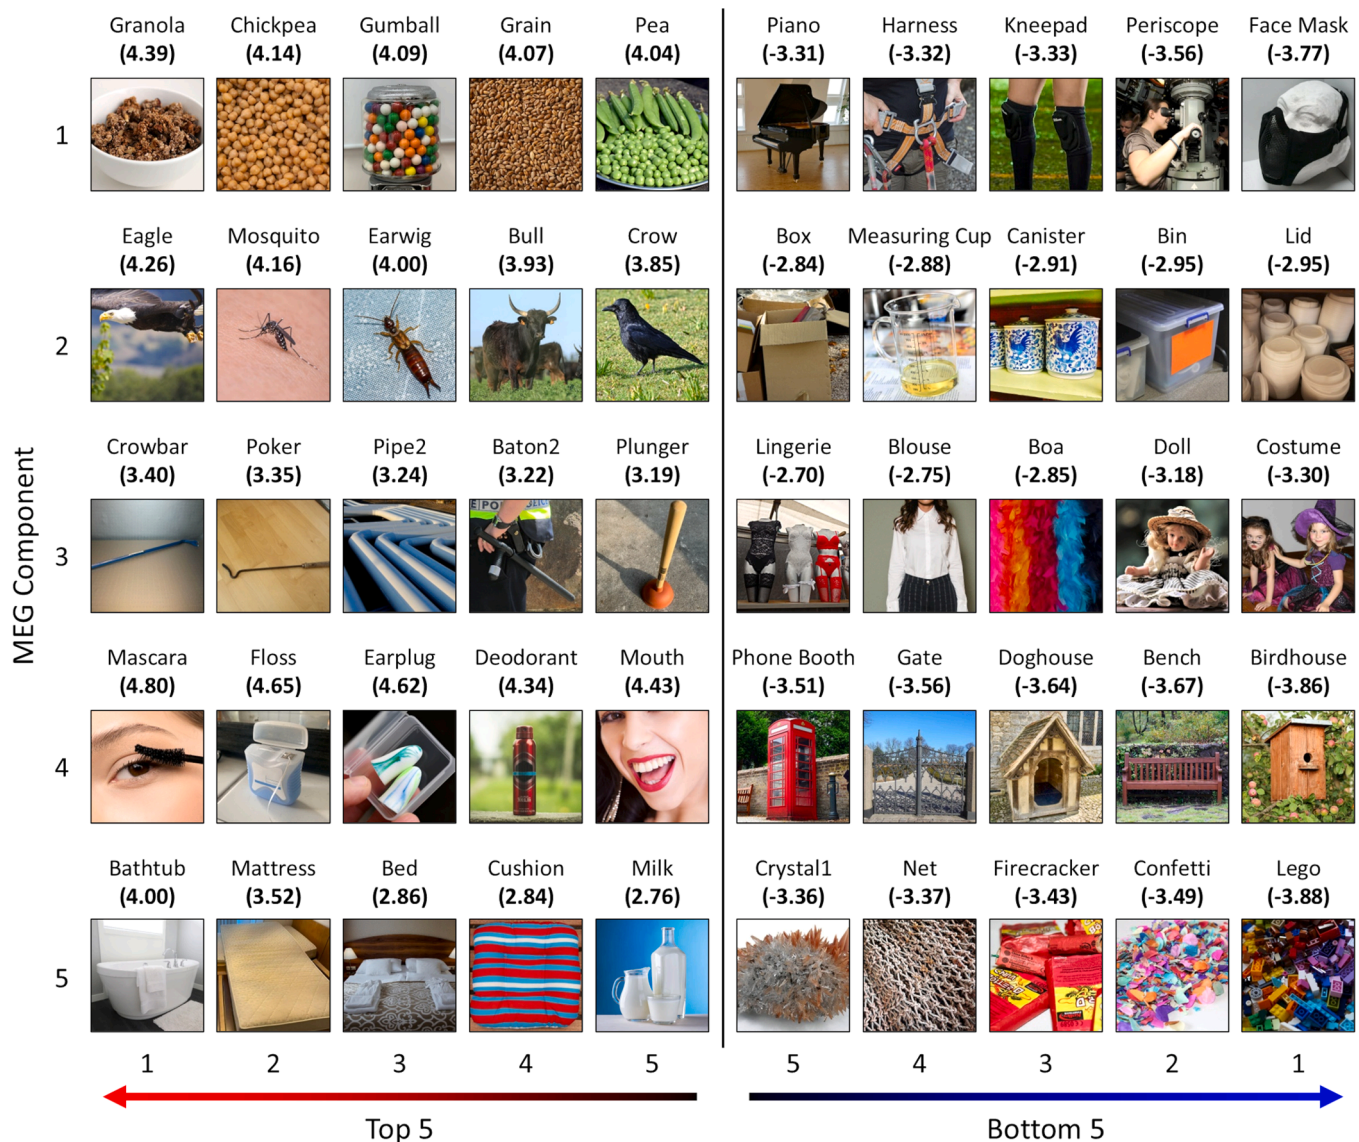


Fig. 5. Top and bottom five object concepts (within the training set) scoring on each PLSR component in the MEG dataset. Representative images are illustrated for each concept. Latent scores are indicated in parentheses. The images shown are distributed under a CC0 licence and are representative of those in the THINGS database.

explained a large proportion of variance in the neural response and showed good alignment over subjects following the Procrustes analysis. This selection of components was also consistent with our previous analysis of the THINGS-fMRI dataset (Watson and Andrews, 2025). That we were able to explain neural responses with a small set of components is consistent with previous accounts emphasising low-dimensional representations of object features in the visual brain (Chen and Bonner, 2025; Haxby et al., 2011; Huth et al., 2012). The components we retained identified key object features represented in the dynamics of the neural response, showed good correspondence with a deep convolutional neural network and with behavioural ratings of object properties, and were consistent between datasets. Nevertheless, our results do not preclude the possibility that the brain may encode additional stimulus features represented in later latent components (Gauthaman et al., 2025; Han and Bonner, 2026; Stringer et al., 2019).

Inspection of the stimulus loadings and latent scores revealed that the components captured a broad spectrum of visual and semantic features. The components also correlated with behavioural ratings of key object properties (Stoinski et al., 2023). While some features aligned with prior accounts of visual cortical organisation (Grill-Spector and

Weiner, 2014), the correspondence with existing models was not straightforward. Rather, multiple features were represented across multiple components, and each component embodied multiple features. This is consistent with a hypothesis that neural representations in visual cortex are tuned to efficiently represent statistical variation of features in real-world objects. Future research might also consider how components are modulated by task demands. All the datasets used passive viewing tasks: the THINGS-fMRI and THINGS-MEG datasets employed an oddball-detection task, while the THINGS-EEG dataset used an orthogonal fixation task. Although tasks requiring participants to actively attend different object features can lead to changes in neural responses throughout early and high-level visual brain regions (Harel et al., 2014), it remains unclear whether this would change the underlying stimulus representations.

The PLSR method provides a data-driven approach for identifying latent components underlying the stimulus dimensions and neural responses. However, a necessary consequence of this method is that the components can only identify features represented amongst the original stimulus dimensions and neural responses. The unconstrained nature of the similarity judgement task employed by Hebart et al. (2020) allowed

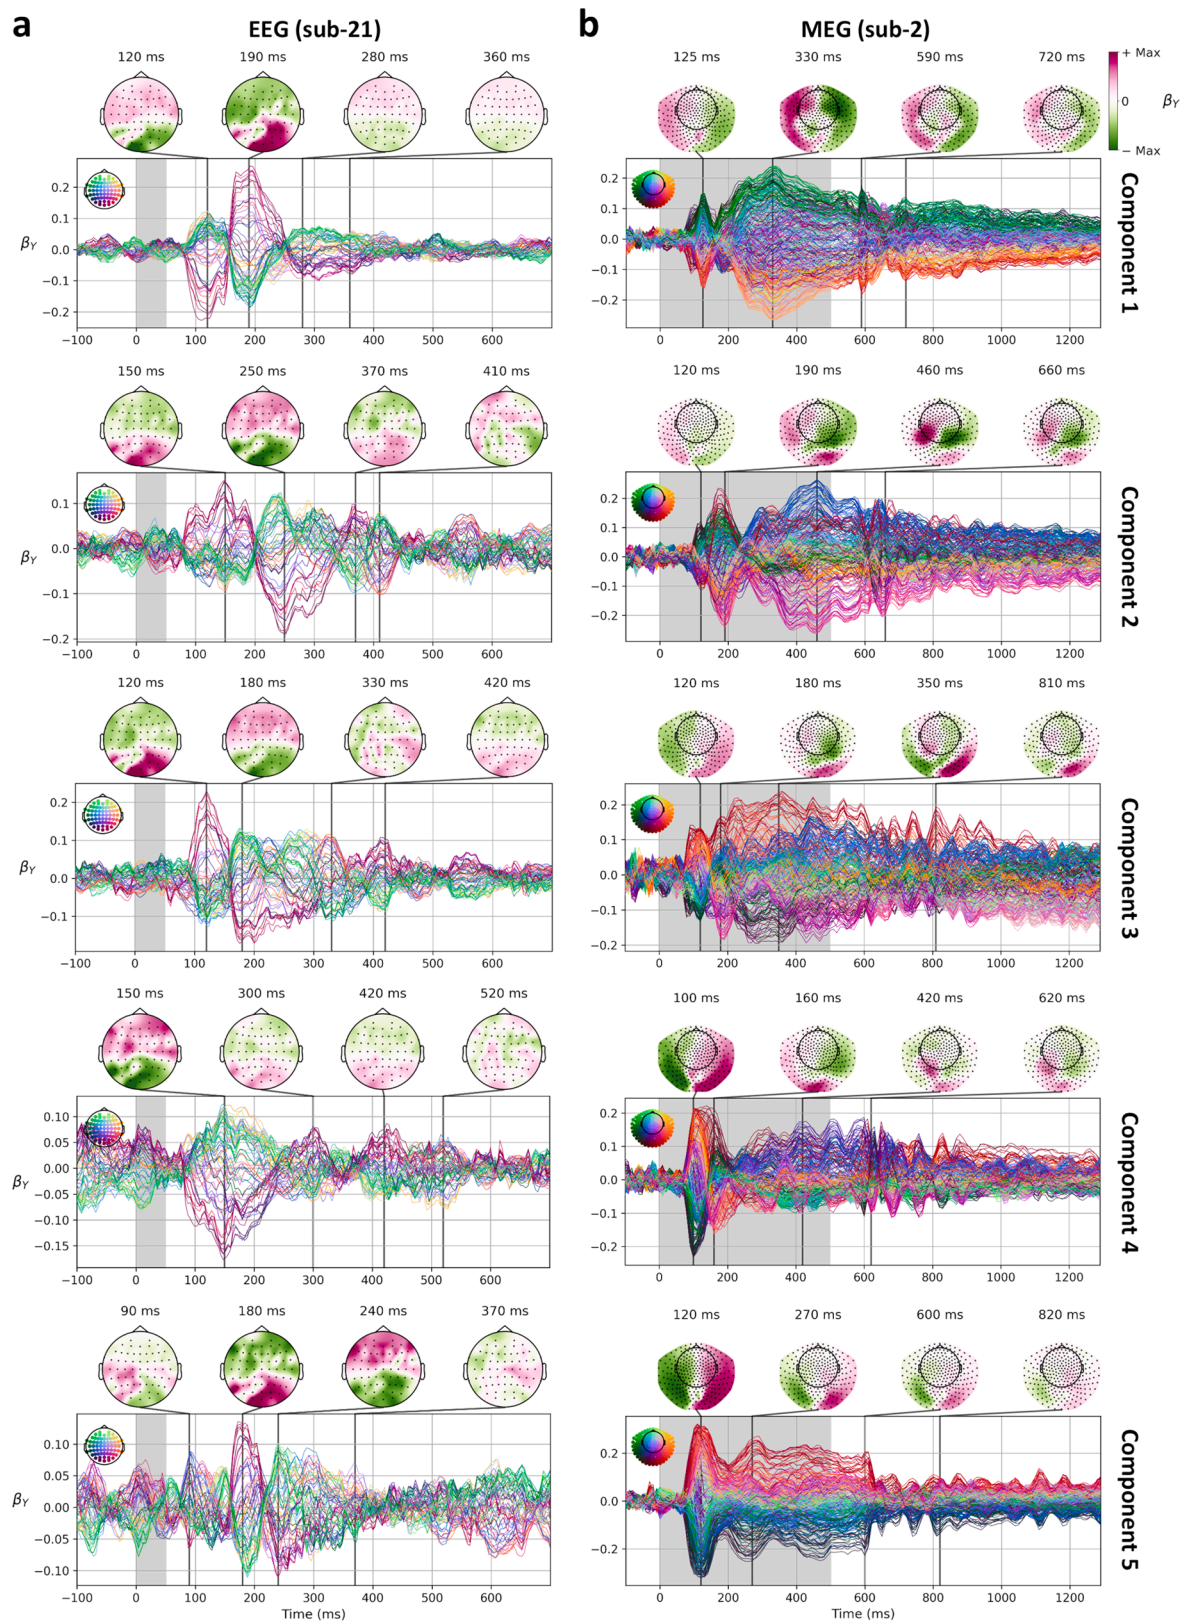


Fig. 6. Neural loadings on each PLSR component for two example subjects from the (a) EEG and (b) MEG datasets. Note that components are not necessarily directly comparable between datasets. Butterfly plots illustrate the time-course of loadings across the epoch for each channel, and scalp topographies illustrate the pattern of loadings over channels at example timepoints. Grey background highlights indicate the stimulus duration. Loadings represent the expected standard deviations of change in the evoked response for one unit of change in the latent score. Note that positive and negative values do not have the same interpretation as in evoked responses themselves.

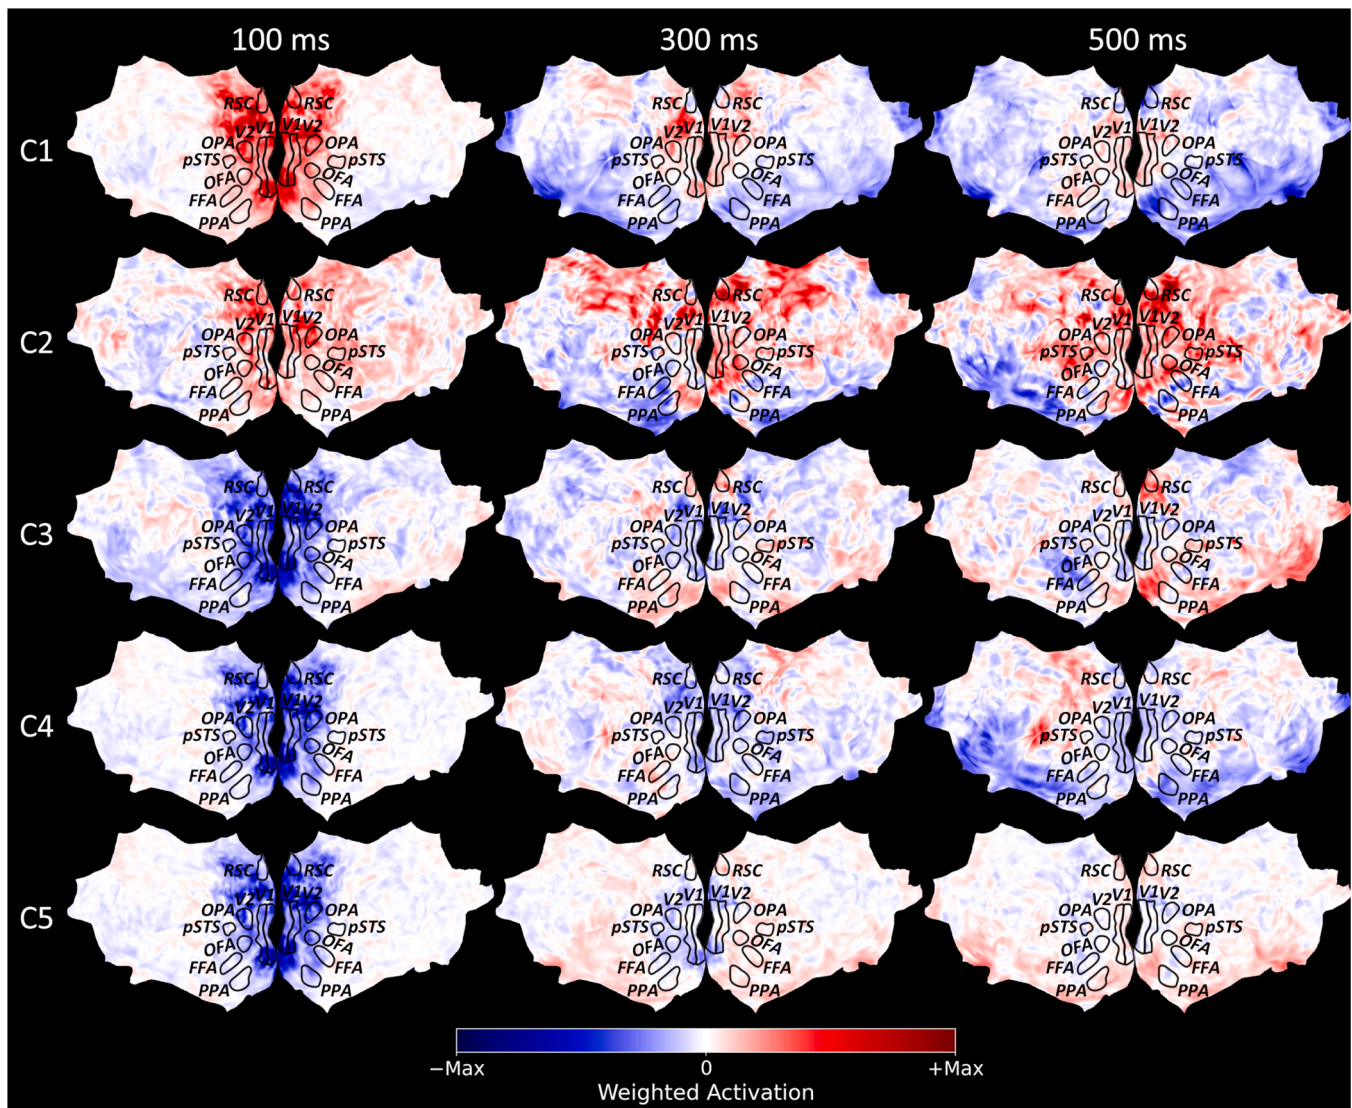


Fig. 7. Group average source localisations of evoked responses in the MEG dataset weighted by latent scores on each PLSR component. Flat patches illustrate whole-brain responses at three example timepoints. Positive and negative activations indicate stronger evoked responses to object concepts scoring more positively or negatively along each component, respectively. Annotations indicate locations of early visual (V1, V2) and category-selective (PPA, RSC, OPA, FFA, OFA, pSTS) regions obtained from independent datasets (Benson et al., 2014; Noad et al., 2024).

them to identify a varied spectrum of behaviourally relevant stimulus features, without imposing experimenter preconceptions on the selection of those features. These include lower-level visual properties such as shape, colour, and texture; mid-level features such as animacy; and higher-level abstract and semantic properties such as how valuable or flammable an object is. Correspondingly, our latent components also captured a wide range of object features and were able to predict a substantial portion of variance in the neural response. Nevertheless, it is possible that features not accounted for amongst the original stimulus dimensions (or our latent components) also contribute to the neural response. Furthermore, our approach required averaging responses over images within each object concept, thus limiting the ability to capture image-level variation. Indeed, both the linear and PLSR encoding models only significantly predicted responses from approximately 100 ms after stimulus onset, which is later than the earliest visually evoked potentials (Inui and Kakigi, 2006; Jeffreys and Axford, 1972), suggesting neither model fully accounted for more basic visual features. However, the PLSR modelling approach is highly flexible, and future research could readily apply it to examine neural representations of other dimensions captured by alternative stimulus models.

The temporal dynamics of the neural loadings revealed smooth transitions over time. As the analysis does not impose any spatiotemporal structure on the loadings, this is consistent with a biologically plausible, graded representation in the brain. Previous studies have shown that the visual processing of objects evolves dynamically, with early time periods associated with representations of simple image-level features, while higher-level visual representations emerge in later time windows (Carlson et al., 2013; Cichy et al., 2014, 2016; Coggan et al., 2016). The sustained prediction accuracies we observed across earlier and later time points are therefore consistent with the PLSR capturing a range of lower and higher-level object features.

Source-localisations of the MEG data further mapped the components onto cortical topographies, revealing gradients from early visual to more anterior regions over time. These localisations indicated strong weightings associated with occipital cortices early in the epoch, before expanding to include more anterior visual and non-visual cortices in later time points. These loadings often revealed topographic gradients across the cortex. For instance, the first component was associated with a gradient distinguishing between early and high-level visual regions. Meanwhile, the second component revealed alternating patterns of

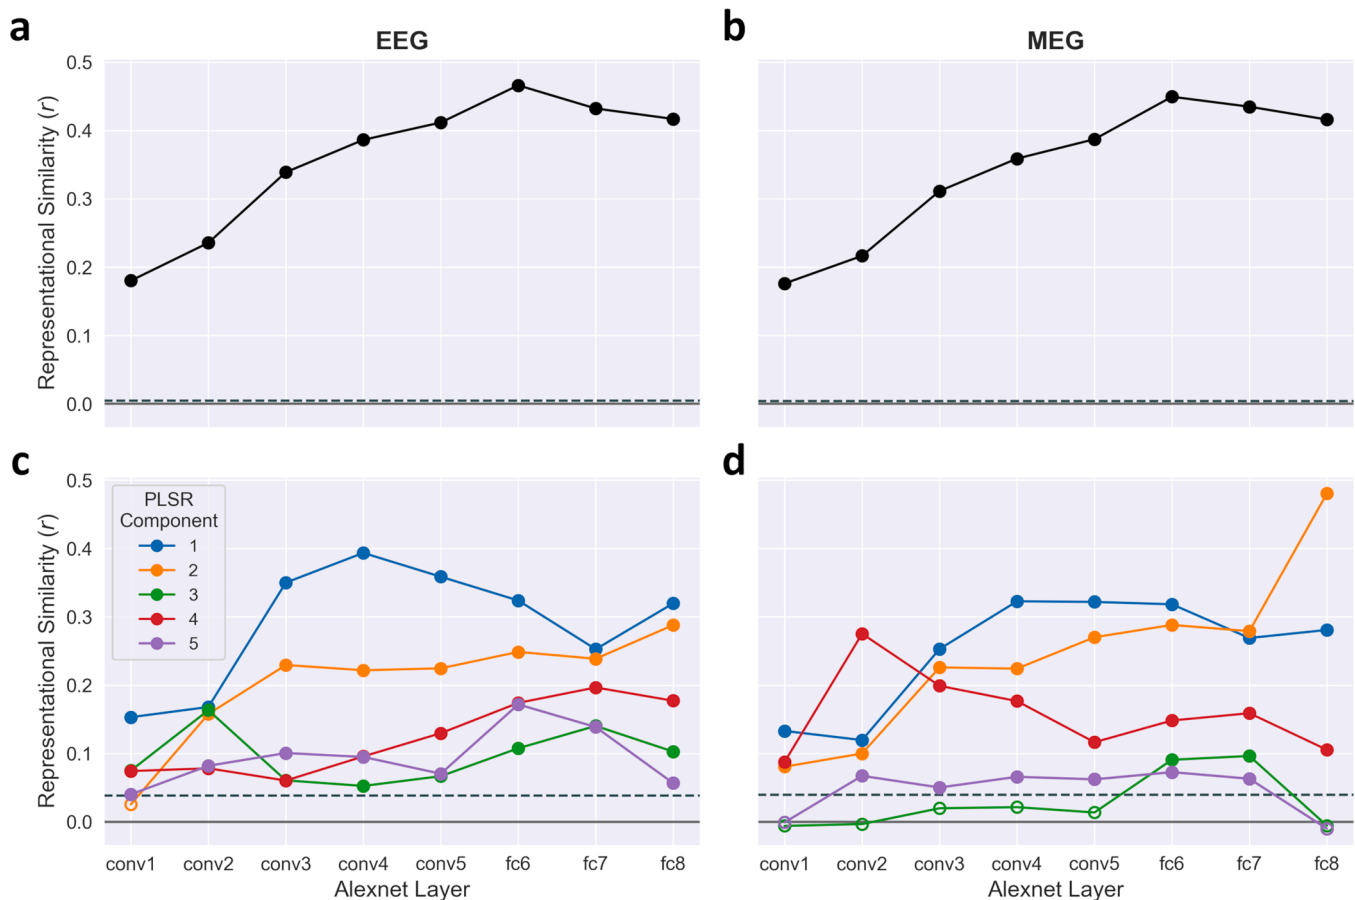


Fig. 8. Representational similarity between PLSR and DCNN representations of objects concepts in the training set. Distances between pairs of object concepts were measured from activations in each layer of an Alexnet DCNN trained for object identification. These were then correlated with corresponding distances measured from the group average PLSR latent scores in the (a, c) EEG and (b, d) MEG datasets, for both (a, b) the full 5-dimensional latent spaces and (c, d) each component separately. Note that components are not necessarily directly comparable between datasets. Filled symbols indicate significant correlations and dashed lines indicate significance thresholds ($p < .05$; FWER corrected), as determined by permutation testing.

positive and negative weights across ventrottemporal cortex. Several components highlighted the role of the parietal lobe at later timepoints, in line with growing evidence for dorsal stream contributions to object processing (Ayzenberg and Behrmann, 2022; Jeong and Xu, 2016).

An interesting question is whether features encoded along the range of each component emerge in the neural response simultaneously or hierarchically across time. For instance, the neural loadings of many components displayed oscillations over time – rather than a simple increase then decrease in magnitude – suggesting a dynamic progression of those representations. The MEG source localisations often indicated a hierarchical progression of the representations, with one end of the component more strongly associated with earlier activations in early visual cortex, and the other end more associated with later activations in higher-level regions. Future research could apply more hypothesis-driven investigations to disentangle the progression of these neural representations, for instance by employing decoding approaches to contrast between stimulus features (Kim et al., 2026; Yeh et al., 2025), or examining representations across frequency bands to disambiguate the influence of feedforward versus feedback signals (Stecher et al., 2025).

Deep convolutional neural networks (DCNNs) are increasingly employed as computational models to investigate object perception, given their capacity to generate neural representations that closely resemble those observed in the human visual system (Khaligh-Razavi and Kriegeskorte, 2014; Cichy et al., 2016; Simony et al., 2024). In the current study, we evaluated the correspondence between representational spaces derived from PLSR and those extracted from a widely used

DCNN architecture, trained for object recognition (Alexnet; Krizhevsky et al., 2017). Representational similarity analysis revealed significant alignment across all layers of the network, with the strongest correspondence observed in higher-order convolutional layers and the initial fully connected layers, suggesting a convergence between computational and neural representations at intermediate-to-late stages of visual object processing.

We also compared the latent scores between the EEG and MEG datasets, and with PLSR components derived from a previous analysis of the THINGS-fMRI dataset (Hebart et al., 2023; Watson and Andrews, 2025). In all datasets, each component correlated strongly with at least one component in the other datasets. Components often displayed high correlations with multiple components in other datasets, indicating there was not always a strict one-for-one correspondence in components between datasets. Nevertheless, there was clear overlap in the latent representations captured in each dataset. This is noteworthy given that each dataset used different imaging modalities – EEG and MEG offer different sensitivities to different neural sources (Lopes da Silva, 2013), while fMRI and M/EEG provide substantially different spatial and temporal sensitivities. Moreover, the datasets have different experimental paradigms (rapid serial visual presentation for EEG, event-related for MEG and fMRI), participant groups, and sample sizes (1854 object concepts in EEG and MEG, 720 object concepts in fMRI). This indicates the PLSR identified robust neural representations shared over subjects, imaging modalities, and experimental paradigms.




Fig. 9. Similarity of PLSR components between THINGS-MEG, THINGS-EEG1, and THINGS-fMRI datasets. Group average latent scores were correlated within and between components across the datasets. Components in the matrix columns were reordered and sign-flipped as necessary (indicated by x-axis tick labels) to best match the order of components in the matrix rows. Latent scores for the THINGS-fMRI dataset were obtained from a previous analysis (Watson and Andrews, 2025).

5. Conclusion

In conclusion, we applied a data-driven approach to predict time-resolved neural responses in EEG and MEG from behaviourally relevant stimulus features of naturalistic objects. A neural encoding model based on partial least squares regression identified a small set of latent components that accurately predicted evoked neural responses across sustained periods including earlier and later time points in the epoch. The components themselves identified many visual and non-visual features of these objects, but did not straightforwardly correspond to previous perspectives on object perception. Instead, our results indicate that the visual brain encodes information in a statistically efficient way that reflects the co-occurrence of features amongst natural variation in real-world objects.

Data availability

The behavioural, MEG, and EEG data are already publicly available from the THINGS initiative (<https://things-initiative.org/>). Our analysis code is available on the Open Science Framework (<https://osf.io/6w5ef>

).

CRediT authorship contribution statement

David M. Watson: Writing – review & editing, Writing – original draft, Visualization, Methodology, Investigation, Formal analysis, Conceptualization. **Richard Aveyard:** Writing – review & editing, Writing – original draft, Methodology, Conceptualization. **Timothy J. Andrews:** Writing – review & editing, Writing – original draft, Supervision, Methodology, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Given his role as an Editorial Intern at NeuroImage, David Watson had no involvement in the peer review of this article and had no access to information regarding its peer review. Full responsibility for the editorial process for this article was delegated to another journal editor. If there are other authors, they declare that they have no known

competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at [doi:10.1016/j.neuroimage.2026.122012](https://doi.org/10.1016/j.neuroimage.2026.122012).

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. *Front. Neuroinform.* 8, 14. <https://doi.org/10.3389/fninf.2014.00014>. February.
- Andrews, T.J., Watson, D.M., Rice, G.E., Hartley, T., 2015. Low-level properties of natural images predict topographic patterns of neural response in the ventral visual pathway. *J. Vis.* 15 (7), 3. <https://doi.org/10.1167/15.7.3>.
- Ayzenberg, V., Behrmann, M., 2022. The dorsal visual pathway represents object-centered spatial relations for object recognition. *J. Neurosci.* 42 (23), 4693–4710. <https://doi.org/10.1523/JNEUROSCI.2257-21.2022>.
- Baker, A.P., Brookes, M.J., Rezek, I.A., Smith, S.M., Behrens, T., Probert Smith, P.J., Woolrich, M., 2014. Fast transient networks in spontaneous human brain activity. *Elife* 3, e01867. <https://doi.org/10.7554/eLife.01867>.
- Bao, P., She, L., McGill, M., Tsao, D.Y., 2020. A map of object space in primate inferotemporal cortex. *Nature* 583 (7814), 103–108. <https://doi.org/10.1038/s41586-020-2350-5>.
- Benson, N.C., Butt, O.H., Brainard, D.H., Aguirre, G.K., 2014. Correction of distortion in flattened representations of the cortical surface allows prediction of V1-V3 functional organization from anatomy. *PLoS Comput. Biol.* 10 (3), e1003538. <https://doi.org/10.1371/journal.pcbi.1003538>.
- Bracci, S., Op de Beeck, H.P., 2023. Understanding Human object vision: a picture is worth a thousand representations. *Annu. Rev. Psychol.* 74 (1), 113–135. <https://doi.org/10.1146/annurev-psych-032720-041031>.
- Brunton, B.W., Beyeler, M., 2019. Data-driven models in human neuroscience and neuroengineering. *Curr. Opin. Neurobiol.* 58, 21–29. <https://doi.org/10.1016/j.conb.2019.06.008>.
- Carlson, T., Tovar, D.A., Kriegeskorte, N., 2013. Representational dynamics of object vision: the first 1000 ms. *J. Vis.* 13 (10), 1–19. <https://doi.org/10.1167/13.10.1>.
- Chen, Z., Bonner, M.F., 2025. Universal dimensions of visual representation. *Sci. Adv.* 11 (27), eadw7697. <https://doi.org/10.1126/sciadv.adw7697>.
- 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. *Sci. Rep.* 6, 27755. <https://doi.org/10.1038/srep27755>. June 2016.
- Cichy, R.M., Pantazis, D., Oliva, A., 2014. Resolving human object recognition in space and time. *Nat. Neurosci.* 17 (3), 455–462. <https://doi.org/10.1038/nn.3635>.
- Coggan, D.D., Baker, D.H., Andrews, T.J., 2016. The role of visual and semantic properties in the emergence of category-specific patterns of neural response in the Human brain. *eNeuro* 3 (4). <https://doi.org/10.1523/ENEURO.0158-16.2016>. ENEURO.0158-16.2016 ENEURO.0158-16.2016.
- Contier, O., Baker, C.I., Hebart, M.N., 2024. Distributed representations of behaviour-derived object dimensions in the human visual system. *Nat. Hum. Behav.* 8 (11), 2179–2193. <https://doi.org/10.1038/s41562-024-01980-y>.
- Dale, A.M., Fischl, B., Sereno, M.I., 1999. Cortical surface-based analysis I. segmentation and surface reconstruction. *Neuroimage* 9 (2), 179–194. <https://doi.org/10.1006/nimg.1998.0395>.
- Dale, A.M., Liu, A.K., Fischl, B.R., Buckner, R.L., Belliveau, J.W., Lewine, J.D., Halgren, E., 2000. Dynamic statistical parametric mapping. *Neuron* 26 (1), 55–67. [https://doi.org/10.1016/S0896-6273\(00\)81138-1](https://doi.org/10.1016/S0896-6273(00)81138-1).
- Fischl, B., Sereno, M.I., Tootell, R.B.H., Dale, A.M., 1999. High-resolution intersubject averaging and a coordinate system for the cortical surface. *Hum. Brain Mapp.* 8 (4), 272–284. [https://doi.org/10.1002/\(SICI\)1097-0193\(1999\)8:4%3C272::AID-HBM10%3E3.0.CO;2-4](https://doi.org/10.1002/(SICI)1097-0193(1999)8:4%3C272::AID-HBM10%3E3.0.CO;2-4).
- Gao, J.S., Huth, A.G., Lescroart, M.D., Gallant, J.L., 2015. Pycortex: an interactive surface visualizer for fMRI. *Front. Neuroinform.* 9. <https://doi.org/10.3389/fninf.2015.00023>. September.
- Gauthaman, R.M., Ménard, B., Bonner, M.F., 2025. Universal scale-free representations in human visual cortex. *PLoS Comput. Biol.* 21 (11), e1013714. <https://doi.org/10.1371/journal.pcbi.1013714>.
- Goodale, M.A., Milner, A.D., 1992. Separate visual pathways for perception and action. *Trends Neurosci.* 15 (1), 20–25. [https://doi.org/10.1016/0166-2236\(92\)90344-8](https://doi.org/10.1016/0166-2236(92)90344-8).
- Gower, J.C., 1975. Generalized procrustes analysis. *Psychometrika* 40 (1), 33–51. <https://doi.org/10.1007/BF02291478>.
- Gramfort, A., Luessi, M., Larson, E., Engemann, D.A., Strohmeier, D., Brodbeck, C., Goj, R., Jas, M., Brooks, T., Parkkonen, L., Hämäläinen, M., 2013. MEG and EEG data analysis with MNE-Python. *Front. Neurosci.* 7. <https://doi.org/10.3389/fnins.2013.00267>, 7 DEC.
- Gramfort, A., Luessi, M., Larson, E., Engemann, D.A., Strohmeier, D., Brodbeck, C., Parkkonen, L., Hämäläinen, M.S., 2014. MNE software for processing MEG and EEG data. *Neuroimage* 86, 446–460. <https://doi.org/10.1016/j.neuroimage.2013.10.027>.
- Grill-Spector, K., Weiner, K.S., 2014. The functional architecture of the ventral temporal cortex and its role in categorization. *Nat. Rev. Neurosci.* 15 (8), 536–548. <https://doi.org/10.1038/nrn3747>.
- Grootswagers, T., Zhou, I., Robinson, A.K., Hebart, M.N., Carlson, T.A., 2022. Human EEG recordings for 1854 concepts presented in rapid serial visual presentation streams. *Sci. Data* 9 (3). <https://doi.org/10.1038/s41597-021-01102-7>.
- Hämäläinen, M.S., Ilmoniemi, R.J., 1994. Interpreting magnetic fields of the brain: minimum norm estimates. *Med. Biol. Eng. Comput.* 32 (1), 35–42. <https://doi.org/10.1007/BF02512476>.
- Han, C., Bonner, M.F., 2026. High-dimensional structure underlying individual differences in naturalistic visual experience. *Curr. Biol.* 36 (3), 723–733.e6. <https://doi.org/10.1016/j.cub.2025.12.039> e6.
- Harel, A., Kravitz, D.J., Baker, C.I., 2014. Task context impacts visual object processing differentially across the cortex. In: *Proceedings of the National Academy of Sciences*, pp. E962–E971. <https://doi.org/10.1073/pnas.1312567111>.
- 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>.
- Hebart, M.N., Contier, O., Teichmann, L., Rockter, A.H., Zheng, C.Y., Kidder, A., Corrivéau, 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., Corrivéau, A., Van Wicklin, C., & Baker, C.I. (2019). THINGS: a database of 1854 object concepts and >26,000 naturalistic object images. *PLoS One*, 14(10), e0223792. <https://doi.org/10.1371/journal.pone.0223792>.
- Hebart, M.N., Zheng, C.Y., Pereira, F., Baker, C.I., 2020. Revealing the multidimensional mental representations of natural objects underlying human similarity judgements. *Nat. Hum. Behav.* 4 (11), 1173–1185. <https://doi.org/10.1038/s41562-020-00951-3>.
- Holm, S., 1979. A simple sequentially rejective multiple test procedure. *Scandinavian J. Stat.* 6 (2), 65–70.
- Hubel, D.H., Wiesel, T.N., 1968. Receptive fields and functional architecture of monkey striate cortex. *J. Physiol. (Lond.)* 195, 215–243.
- Huth, A.G., Nishimoto, S., Vu, A.T., Gallant, J.L., 2012. A continuous semantic space describes the representation of thousands of object and action categories across the Human brain. *Neuron* 76 (6), 1210–1224. <https://doi.org/10.1016/j.neuron.2012.10.014>.
- Inui, K., Kakigi, R., 2006. Temporal analysis of the flow from V1 to the extrastriate cortex in humans. *J. Neurophysiol.* 96 (2), 775–784. <https://doi.org/10.1152/jn.00103.2006>.
- Jeffreys, D.A., Axford, J.G., 1972. Source locations of pattern-specific components of human visual evoked potentials. I. component of striate cortical origin. *Exp. Brain Res.* 16 (1), 1–21. <https://doi.org/10.1007/BF00233371>.
- Jeong, S.K., Xu, Y., 2016. Behaviorally relevant abstract object identity representation in the Human parietal cortex. *J. Neurosci.* 36 (5), 1607–1619. <https://doi.org/10.1523/JNEUROSCI.1016-15.2016>.
- Kanwisher, N., 2010. Functional specificity in the human brain: a window into the functional architecture of the mind. *Proc. Natl. Acad. Sci. USA* 107 (25), 11163–11170. <https://doi.org/10.1073/pnas.1005062107>.
- Kim, A.H., Quek, G.L., Moerel, D., Gorton, O.K., & Carlson, T.A. (2026). *Disentangling objects' contextual associations from perceptual and conceptual attributes using time-resolved neural decoding* (p. 2025.05.29.656895). *bioRxiv*. <https://doi.org/10.1101/2025.05.29.656895>.
- Khaligh-Razavi, S.-M., Kriegeskorte, N., 2014. Deep supervised, but not unsupervised, models may explain IT cortical representation. *PLoS Comput. Biol.* 10 (11), e1003915. <https://doi.org/10.1371/journal.pcbi.1003915>.
- Konkle, T., Oliva, A., 2012. A real-world size organization of object responses in occipitotemporal cortex. *Neuron* 74 (6), 1114–1124. <https://doi.org/10.1016/j.neuron.2012.04.036>.
- Kriegeskorte, N., Mur, M., Bandettini, P.A., 2008a. Representational similarity analysis—connecting the branches of systems neuroscience. *Front. Syst. Neurosci.* 2 (4), 1–28. <https://doi.org/10.3389/fnins.2008.004.2008>.
- Kriegeskorte, N., Mur, M., Ruff, D., Kiani, R., Bodurka, J., Esteky, H., Tanaka, K., Bandettini, P., 2008b. Matching categorical object representations in inferior temporal cortex of man and monkey. *Neuron* 60 (6), 1126–1141. <https://doi.org/10.1016/j.neuron.2008.10.043>.
- Krishnan, A., Williams, L.J., McIntosh, A.R., Abdi, H., 2011. Partial Least Squares (PLS) methods for neuroimaging: a tutorial and review. *Neuroimage* 56 (2), 455–475. <https://doi.org/10.1016/j.neuroimage.2010.07.034>.
- Krizhevsky, A., Sutskever, I., Hinton, G.E., 2017. ImageNet classification with deep convolutional neural networks. *Commun. ACM* 60 (6), 84–90. <https://doi.org/10.1145/3065386>.
- Levy, I., Hasson, U., Avidan, G., Hendler, T., Malach, R., 2001. Center – periphery organization of human object areas. *Nat. Neurosci.* 4 (5), 533–539.
- Li, A., Feitelberg, J., Saini, A.P., Höchenberger, R., Scheltienne, M., 2022. MNE-ICLabel: automatically annotating ICA components with ICLabel in python. *J. Open. Source Softw.* 7 (76), 4484. <https://doi.org/10.21105/joss.04484>.
- Lopes da Silva, F., 2013. EEG and MEG: relevance to neuroscience. *Neuron* 80 (5), 1112–1128. <https://doi.org/10.1016/j.neuron.2013.10.017>.
- Noad, K.N., Watson, D.M., Andrews, T.J., 2024. Familiarity enhances functional connectivity between visual and nonvisual regions of the brain during natural viewing. *Cerebral Cortex* 34 (7), bhae285. <https://doi.org/10.1093/cercor/bhae285>.

- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E., 2011. Scikit-learn: machine learning in Python. *J. Mach. Learn. Res.* 12 (85), 2825–2830.
- Peelen, M.V., Downing, P.E., 2017. Category selectivity in human visual cortex: beyond visual object recognition. *Neuropsychologia* 105, 177–183. <https://doi.org/10.1016/j.neuropsychologia.2017.03.033>. March.
- Pion-Tonachini, L., Kreutz-Delgado, K., Makeig, S., 2019. ICLabel: an automated electroencephalographic independent component classifier, dataset, and website. *Neuroimage* 198, 181–197. <https://doi.org/10.1016/j.neuroimage.2019.05.026>.
- Ritchie, J.B., Wardle, S.G., Vaziri-Pashkam, M., Kravitz, D.J., Baker, C.I., 2026. Rethinking category-selectivity in human visual cortex. *Cogn. Neurosci.* 17, 49–76. <https://doi.org/10.1080/17588928.2025.2543890>.
- Ségonne, F., Dale, A.M., Busa, E., Glessner, M., Salat, D., Hahn, H.K., Fischl, B., 2004. A hybrid approach to the skull stripping problem in MRI. *Neuroimage* 22 (3), 1060–1075. <https://doi.org/10.1016/j.neuroimage.2004.03.032>.
- Simony, E., Grossman, S., Malach, R., 2024. Brain-machine convergent evolution: why finding parallels between brain and artificial systems is informative. *Proc. of the Nat. Acad. Sci.* 121 (41), e2319709121. <https://doi.org/10.1073/pnas.2319709121>.
- Stecher, R., Cichy, R.M., Kaiser, D., 2025. Decoding the rhythmic representation and communication of visual contents. *Trends. Neurosci.* 48 (3), 178–188. <https://doi.org/10.1016/j.tins.2024.12.005>.
- Stoinski, L.M., Perkuhn, J., Hebart, M.N., 2023. THINGSplus: new norms and metadata for the THINGS database of 1854 object concepts and 26,107 natural object images. *Behav. Res. Methods* 56 (3), 1583–1603. <https://doi.org/10.3758/s13428-023-02110-8>.
- Stringer, C., Pachitariu, M., Steinmetz, N., Carandini, M., Harris, K.D., 2019. High-dimensional geometry of population responses in visual cortex. *Nature* 571 (7765), 361–365. <https://doi.org/10.1038/s41586-019-1346-5>.
- Teichmann, L., Hebart, M.N., Baker, C.I., 2026. Dynamic representation of multidimensional object properties in the human brain. *J. Neurosci.* 46, e1057252026. <https://doi.org/10.1523/JNEUROSCI.1057-25.2026>.
- Ungerleider, L., Haxby, J.V., 1994. What' and 'where' in the human brain. *Curr. Opin. Neurobiol.* 4 (2), 157–165. [https://doi.org/10.1016/0959-4388\(94\)90066-3](https://doi.org/10.1016/0959-4388(94)90066-3).
- Wandell, B.A., Winawer, J., 2011. Imaging retinotopic maps in the human brain. *Vision. Res.* 51 (7), 718–737. <https://doi.org/10.1016/j.visres.2010.08.004>.
- Watson, D.M., Andrews, T.J., 2025. A data-driven analysis of the perceptual and neural responses to natural objects reveals organizing principles of Human visual cognition. *J. Neurosci.* 45 (2), e1318242024. <https://doi.org/10.1523/JNEUROSCI.1318-24.2024>.
- Yeh, L.C., Peelen, M.V., & Kaiser, D. (2025). *Spatiotemporal representations of contextual associations for real-world objects* (p. 2025.10.20.683392). bioRxiv. <https://doi.org/10.1101/2025.10.20.683392>.