The footprint of colour in EEG signal

Anonymous authors
Paper under double-blind review

Abstract

Our perception of the world is inherently colourful, and colour provides well-documented benefits for vision: it helps us see things quicker and remember them better. We hypothesised that colour is not only central to perception but also a rich, decodable source of information in electroencephalography (EEG) signals recorded non-invasively from the scalp. While previous work has shown that brain activity carries colour information for simple, uniform stimuli, it remains unclear whether this extends to natural, complex images with no explicit colour cueing. To investigate this, we analysed the THINGS EEG dataset, which contains 64-channel recordings from participants viewing 1,800 distinct objects (16,740 images) presented for 100 ms each, yielding over 82,000 trials. We established a perceptual colour ground truth through a psychophysical experiment in which participants viewed each image for 100 ms and selected the perceived colours from a 13-option palette. An artificial neural network trained to predict these scene-level colour distributions directly from EEG signals showed that colour information was robustly decodable (average F-score of 0.5). We further examined the effect of colour features on object decoding. Using a contrastive learning framework, we modelled colour–object perception with the Segment Anything Model (SAM), in which all pixels within a segment were replaced with their average colour, followed by standard feature extraction using CLIP vision encoders. We trained an EEG encoder, CUBE (ColoUr and oBjEct decoding), to align features in both object and colour spaces. Across EEG and MEG datasets in a 200-class recognition task, incorporating colour improved decoding accuracy by approximately 5%. Together, these findings demonstrate that EEG signals recorded during natural vision carry substantial colour information that interacts with object perception. Modelling this interaction enhances the power of neural decoding.

1 How strong is the colour signal in neuroimaging?

Our visual system makes sense of a scene with remarkable speed. In just a fleeting glance, as brief as 13 ms, we can attach a simple description such as “green tree” to what we have seen (Potter et al., 2014). This raises a critical question: what neural representations emerge within such a brief window, and to what extent can they be captured in neuroimaging signals? Here we turn our attention to colour, an effortless and ever-present aspect of vision. Colour not only shapes how we perceive objects (Tanaka et al., 2001; Bramão et al., 2011), but also enhances memorability (Gegenfurtner & Rieger, 2000; Wichmann et al., 2002) and speeds up recognition (Møller & Hurlbert, 1996; Rosenthal et al., 2018).

Colour decoding from neuroimaging has a long history (Regan, 1970; Paulus et al., 1984). Brain activity carries information about chromaticity, luminance, and saturation (Sutterer et al., 2021; Hermann et al., 2022; Pennock et al., 2023; Rozman et al., 2024), the hue circle (Hajonides et al.,

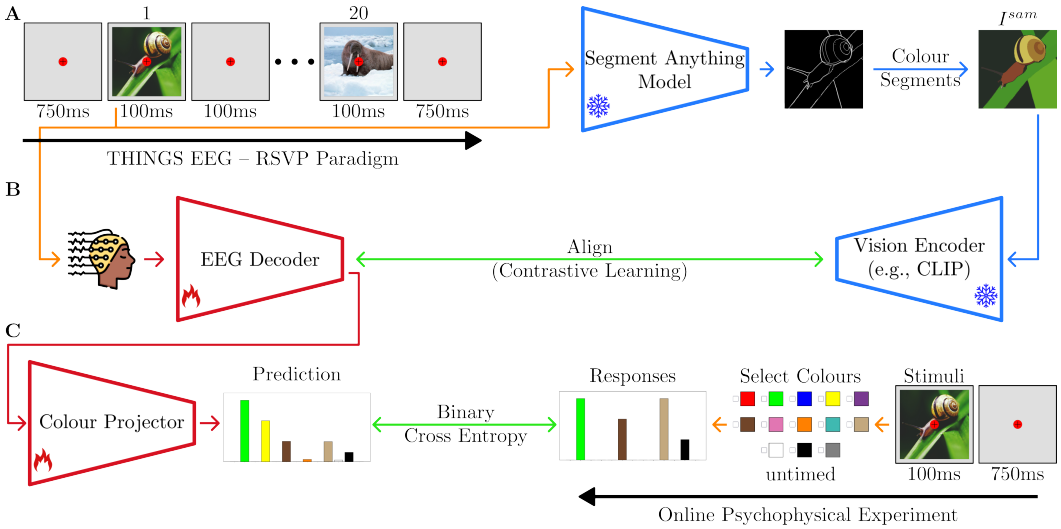

Figure 1: The overview of ColoUr and oBjEct decoding—CUBE. A: The RSVP paradigm used to collect the THINGS EEG dataset (Gifford et al., 2022). B: An EEG decoder aligns brain activity with features from a pretrained vision encoder applied to colour-segmented images. C: A linear projection layer maps the aligned representation onto the behavioural colour responses.

2021), the geometry of colour space (Rosenthal et al., 2021), and even unique hues (Chauhan et al., 2023). Fewer studies, however, have examined how colour interacts with object processing. One study suggests that while both shape and colour can be decoded as early as 60–70 ms after stimulus onset, shape–colour congruency emerges later, around 200 ms (Teichmann et al., 2020). These findings are valuable but mostly derive from simplified displays of uniform coloured patches on plain backgrounds. It remains unclear whether colour can be reliably decoded from brain signals when viewing rich, natural scenes—this is the challenge addressed in the present study.

Artificial intelligence and large datasets have recently accelerated progress in decoding. The NSD dataset (Allen et al., 2022), for example, provides large-scale fMRI data covering about 10,000 natural images. Similarly, the THINGS EEG (Gifford et al., 2022) and MEG (Hebart et al., 2023) datasets offer recordings of comparable scale, enabling new opportunities to investigate how natural images are represented in the brain. Alongside these resources, contrastive learning (Radford et al., 2021) has emerged as a powerful tool for decoding. It has already shown strong performance across modalities, from speech recognition (Défossez et al., 2023) to visual object recognition in fMRI (Scotti et al., 2024), EEG (Song et al., 2024), and MEG (Wu et al., 2025).

1.1 CUBE (ColoUr and oBjEct decoding)

We adopted the same general framework of large datasets and contrastive learning to investigate how colour is represented in brain signals during natural image viewing. Our focus here is on EEG, which, despite its low spatial resolution, offers high temporal resolution, affordability, portability, and the possibility of real-time decoding (Benchetrit et al., 2023; Robinson et al., 2023). Instead of collecting a new dataset, we created colour annotations for the THINGS EEG dataset (Gifford et al., 2022) through a large-scale psychophysical experiment designed to mimic the conditions of the original recordings. Participants viewed an image for 100 ms and then selected all perceived colours from a palette of 13 options (see Figure 1, panel C).

We trained an EEG decoder, implemented as a simple artificial neural network (ANN) with two linear layers and a residual connection, to align EEG representations with visual features extracted from pretrained CLIP networks. To better capture perceptual colour structure, which operates at the object rather than pixel level (Gegenfurtner, 2025), we processed each image I^{org} using the Segment Anything Model (SAM) (Kirillov et al., 2023). For each segmented region, we averaged pixel colours to create uniformly colour-segmented images I^{sam} (see Figure 1, panel B), providing a closer approximation to colour perception at a glance.

To evaluate colour decoding from EEG signals, we added a linear projection layer atop the CLIP-aligned features to output a 13-dimensional vector matching the behavioural colour palette. Performance was measured with the F-score for this multi-class task. Results show reliable decoding, with an average F-score of 0.50 across participants—well above chance (0.17). This constitutes our first contribution: demonstrating colour decoding from EEG during natural image viewing with 100 ms exposure. Notably, the noise ceiling in this rapid serial visual presentation (RSVP) paradigm is 0.64, indicating the decoder approaches average human agreement.

Encouraged by these results, we hypothesised that incorporating colour features into the contrastive alignment framework could improve object decoding. Colour and object perception are closely intertwined, both behaviourally (Bramão et al., 2011; Gegenfurtner, 2025) and neurally (Rosenthal et al., 2018; Tanaka et al., 2001). To model this, we aligned the EEG decoder simultaneously to CLIP features from the original images I^{org} and colour-segmented images I^{sam} (Figure 2), which capture object–colour associations more directly, particularly under the brief 100 ms exposure.

We term this framework CUBE (ColoUr and oBjEct decoding), as it explicitly leverages the interaction between colour and object representations in the brain. Experimentally, CUBE improves state-of-the-art object recognition decoding by a consistent 5% across all participants and in both EEG and MEG. This underscores the importance of colour–object interactions in neuroimaging decoding and points to a strongly shared representational space for colour and object in the brain.

2 Method

We primarily focused on the THINGS EEG dataset and, secondarily, on the MEG dataset, both derived from a subset of the THINGS collection (Hebart et al., 2019), a high-quality set comprising 1,854 diverse object concepts. We generated colour annotations for the images through an online psychophysical experiment and employed a contrastive learning framework to train our networks.

2.1 Neuroimaging datasets

THINGS EEG (Gifford et al., 2022): Recordings were collected from 10 participants using an RSVP paradigm (Intraub, 1981), where each image was shown for 100 ms, followed by a 100 ms blank (Figure 1, Panel A). EEG was recorded with a 64-channel cap. The training set included 1,654 concepts (10 images per concept, 4 repetitions per image), and the test set 200 unseen concepts (1 image per concept, 80 repetitions per image). Preprocessing followed the original paper: signals were epoched 0–1000 ms post-stimulus, downsampled to 250 Hz, and reduced to 17 occipito-parietal channels most relevant to vision[1]. To improve signal-to-noise ratio, repetitions of the same image were averaged, yielding 16,540 training samples and 200 test samples per participant.

THINGS MEG (Hebart et al., 2023): Recordings from 4 participants with 271 channels, each image presented for 500 ms followed by a 1000 ± 200 ms interval. The training set included 1,854 concepts (12 images per concept, 1 repetition each), and the test set comprised 200 concepts (1

[1]P7, P5, P3, P1, Pz, P2, P4, P6, P8, PO7, PO3, POz, PO4, PO8, O1, Oz, O2

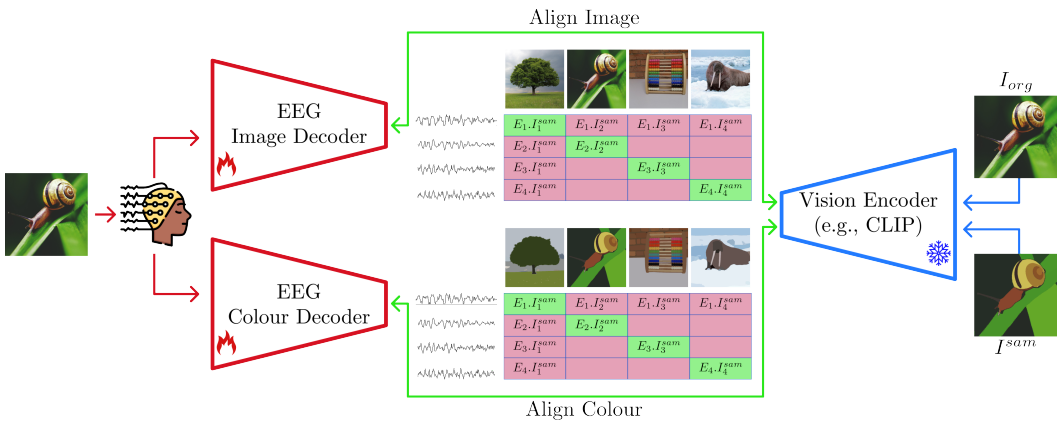

Figure 2: CUBE: Incorporating colour features into object decoding. Two EEG decoders are trained to align with (1) the original RGB images viewed by participants and (2) colour-segmented versions of the same images. Colour alignment follows contrastive learning (Radford et al., 2021), while image alignment employs the Uncertainty-aware Blur Prior algorithm (Wu et al., 2025).

image per concept, 12 repetitions each). Test concepts were excluded from training for zero-shot EEG evaluation. Preprocessing followed the same pipeline as EEG, with MEG signals downsampled to 200 Hz.

2.2 Psychophysical experiment on colour perception under brief exposure

When an image is shown for only 100 ms, a key question is how much of the scene can be understood and accessed consciously (Keysers et al., 2001). To approximate participants' colour perception in the EEG experiment, we conducted a psychophysical study with a similar presentation rate. A blank screen with a central fixation cross was shown for 750 ms, followed by the image for 100 ms. Participants, instructed via written guidelines, selected all perceived colours from a thirteen-colour palette (Figure 1, Panel C).

To annotate the full THINGS EEG dataset (16,740 images), the experiment was run online via Prolific (Palan & Schitter, 2018). Each participant saw 450 images. For training images, participants had 5 seconds per trial due to the large volume; for test images, no time limit was imposed to capture all perceived colours. In total, 133 participants from diverse cultural and linguistic backgrounds participated and were compensated monetarily.

Subjectively, this fast-paced paradigm shows that participants tend to recall only a few foreground objects and their colours. Strong colour–object associations emerge, while background colours are poorly remembered unless covering large uniform regions. This is reflected in the responses: participants selected on average 2.1 colours per image. Beyond the THINGS EEG colour annotations, this large-scale experiment provides a valuable dataset for studying colour perception under brief exposures in ecologically relevant settings, which is released for the community.

2.3 Visual neural decoding

To decode colour and object from neuronal data, we adopted a contrastive learning paradigm (Radford et al., 2021), which has been widely used in neuroimaging research (Song et al., 2024; Li et al., 2024; Scotti et al., 2024). In this framework, an EEG decoder is trained to align its output

representations with those of a pretrained vision encoder, often a variant of CLIP (Figure 2). Similar strategies have also been applied with other modalities, such as language encoders (Akbarinia, 2024) and depth encoders (Zhang et al., 2025).

Formally, let $\mathbf{z}_i^{EEG}$ denote the feature representation predicted by the EEG decoder for sample i, and $\mathbf{z}_i^{IMG}$ the corresponding representation from the vision encoder. The goal of contrastive learning is to maximise similarity between matching pairs $(\mathbf{z}_i^{EEG}, \mathbf{z}_i^{IMG})$ while minimising similarity with all non-matching pairs in the batch. This is achieved with a symmetric cross-entropy objective:

$$\mathcal{L}_{\text{CLIP}} = -\frac{1}{N}\sum_{i=1}^{N}\left[\log\frac{\exp\left(\text{sim}(\mathbf{z}_i^{EEG}, \mathbf{z}_i^{IMG})/\tau\right)}{\sum_{j=1}^{N}\exp\left(\text{sim}(\mathbf{z}_i^{EEG}, \mathbf{z}_j^{IMG})/\tau\right)} + \log\frac{\exp\left(\text{sim}(\mathbf{z}_i^{IMG}, \mathbf{z}_i^{EEG})/\tau\right)}{\sum_{j=1}^{N}\exp\left(\text{sim}(\mathbf{z}_i^{IMG}, \mathbf{z}_j^{EEG})/\tau\right)}\right], \tag{1}$$

where N is the batch size, τ is a learnable temperature parameter, and $\text{sim}(\cdot,\cdot)$ denotes the cosine similarity. This loss encourages the EEG and image representations of the same stimulus to be close in the embedding space, while separating them from mismatched pairs.

One persistent challenge in neuroimaging applications is dataset size: current datasets are relatively small for deep learning, leading to overfitting during training and poor generalisation at test time. A recently proposed technique, the Uncertainty-aware Blur Prior (Wu et al., 2025), mitigates this by introducing a foveated blur to the original images I^{org}, simulating how participants perceive stimuli. By suppressing high-frequency details, this strategy reduces one of the main drivers of overfitting. We adopt this approach in our training framework, processing I^{org} with the foveation blur described in Wu et al. (2025).

Building on this idea, we propose a colour-aware contrastive learning framework, in which the decoder is additionally aligned with CLIP features extracted from colour-segmented images, denoted as I^{sam}. These images are derived from the original input I^{org} using SAM-1 (Kirillov et al., 2023) with default global segmentation parameters, except for an increased resolution of 64 points per side and a stability score threshold of 0.92. We hypothesise that colour-segmented images more closely resemble participants' perceived colours and their object associations during brief exposures (100 ms). Consequently, introducing a contrastive loss term $\mathcal{L}_{\text{CLIP}}$ between $\mathbf{z}_i^{EEG}$ and $\mathbf{z}_i^{SAM}$ is expected to boost object decoding.

To evaluate this, we conducted pilot experiments in which participants viewed I^{org} for 100 ms, followed by a 750 ms grey screen. After this interval, I^{sam} was presented either alone or alongside a greyscale version of I^{org}. Participants were instructed to click on pixels whose colour values were inconsistent with the original scene viewed for 100 ms. Only a small number of mismatches were reported, suggesting that colour-segmented images provide a close approximation of perceived colours, objects, and their associations under such brief viewing conditions.

The EEG Decoder in CUBE follows Wu et al. (2025): two linear layers with GELU activation and a residual connection. The Colour Projector comprises two linear layers with ReLU, mapping CLIP-aligned features to a thirteen-colour palette. In all experiments, OpenCLIP (Cherti et al., 2023) was used to extract features from different vision encoders. Networks were trained for 50 epochs with batch size 1024 using the AdamW optimiser (Loshchilov & Hutter, 2017), with learning rate 1×10^{-4} and weight decay 1×10^{-4}. Other configurations followed Wu et al. (2025).

For all experiments, two types of networks were trained. Intra-participant networks were trained and evaluated on the EEG data of the same participant, whereas inter-participant networks were trained on all participants except one, which was held out for testing.

3 Colour decoding

Colour decoding is inherently a multi-class task, as multiple colours may co-occur within a single image. Unlike object recognition, colour perception shows striking individual differences, particularly under brief viewing (Mollon et al., 2017; Bosten, 2022). For example, one participant may label wood as brownish, whereas another may choose a more beige shade (Lafer-Sousa et al., 2015). To accommodate this multi-class structure and the variability across observers, we quantified agreement between two human responses using the F-score:

$$F = \frac{TP}{2TP + FP + FN}, \tag{2}$$

where TP, FP, and FN denote true positives, false positives, and false negatives, respectively. We chose the F-score over the closely related Jaccard index (Jaccard, 1901), which is also used for set comparison, because the F-score assigns greater weight to TPs. This emphasis is better suited to colour fidelity, as it highlights dominant colours selected by participants. We directly compared behavioural and neural data by applying the same metric to quantify agreement between EEG-decoded colours and the average human responses. Because neither the human averages nor the model predictions are binary, both were thresholded. All reported F-scores use a threshold of $\frac{1}{3}$, based on the rationale that at least one-third of participants agreed on a colour for a given image.

The results of colour decoding on the THINGS EEG dataset are shown in Figure 3. Overall, the CUBE model achieves an F-score above 0.50, approaching the noise ceiling (0.64)—the average agreement among participants—and well above chance (0.17), estimated over 10,000 iterations using two randomly selected colours per trial to match typical participant responses. Performance also exceeds a baseline of 0.23, computed similarly but sampling colours from the training-set distribution. Inter-participant models reach lower F-scores (0.33) yet still far exceed chance, indicating a shared representation of colour across individuals (Gegenfurtner, 2003).

Excluding object alignment from training results in colour decoding with an F-score of 0.46, significantly exceeding chance levels. This shows that the network can extract meaningful colour information directly from EEG signals without relying on any additional source of information. Nevertheless, for most participants, colour decoding improves significantly (Student's t-test) when object alignment is included, consistent with evidence that object and contextual features influence colour perception (Witzel & Gegenfurtner, 2018; Tanaka & Presnell, 1999; Gegenfurtner, 2025). Although the reported results use the CoCa-ViT-L-14 architecture (Yu et al., 2022), no significant differences were found when using alternative architectures.

The F-scores vary by 7% between the best and worst participants (53% for participant 06 and 46% for participant 04), yet the distributions across the test set appear qualitatively similar. This consistency likely reflects the fact that the colour ground truth is based on average human responses rather than each participant's individual colour perception during the EEG experiment (Bosten, 2022). Consequently, higher decoding accuracy might be achievable if the behavioural ground truth corresponded to the neural data of the same individual.

The examples in Figure S1 show that neurally decoded colours often remain plausible even in trials with low F-scores. For instance, the Flax Seed image is behaviourally labelled brown–beige, while the decoded hues are neighbouring orange–red, indicating only a small mismatch. Similar swaps between red and orange appear for the Omelette and Fruit images—a known challenge even for computer vision models (Parraga & Akbarinia, 2020). In other cases, such as the Elephant image, mechanisms like colour constancy (Akbarinia & Parraga, 2017; Gil Rodríguez et al., 2024) or memory (Hansen et al., 2006) may drive participants to report grey despite physically yellowish–beige pixels, reflecting processes requiring longer integration than the 100 ms of neural data. Overall, the neurally decoded colours qualitatively appear both meaningful and perceptually coherent.

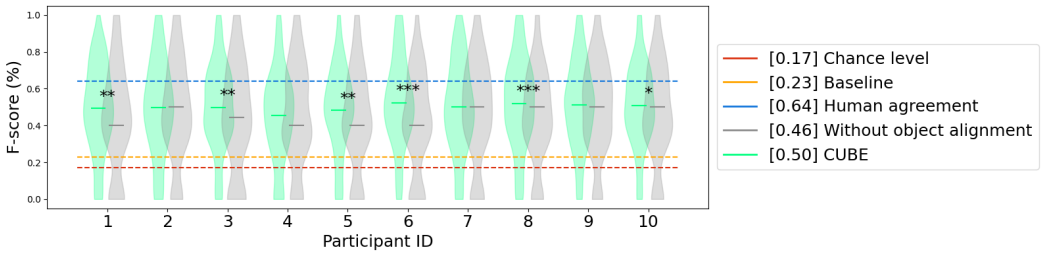

Method	Intra-participant											Inter-participant										
	S01	S02	S03	S04	S05	S06	S07	S08	S09	S10	AVG	S01	S02	S03	S04	S05	S06	S07	S08	S09	S10	AVG
CUBE	.484	.495	.502	.460	.486	.530	.512	.523	.513	.520	**.502**	.318	.400	.282	.341	.322	.276	.332	.335	.324	.363	**.329**

Figure 3: Colour decoding performance of CUBE. Top: F-score distributions over 200 test images. Chance level is shown in red and the noise ceiling from behavioural agreement in blue. Asterisks mark significant differences between models with and without object alignment. Bottom: Average F-scores for within- (intra) and across-participant (inter) training. Table cells are colour-coded from green to yellow as F-scores decline.

4 Object decoding

Object decoding in the THINGS EEG and MEG datasets is formulated as a retrieval task. Each network is evaluated in a zero-shot setting, where, among 200 candidate images, the one with the highest cosine similarity to the decoded EEG features is taken as the predicted object category.

Table 1 reports the object decoding accuracy of CUBE and several comparison models on the THINGS EEG dataset (Gifford et al., 2022). CUBE achieves 57% top-1 and 86% top-5 accuracy, representing a 6% improvement over UBP (Wu et al., 2025). This gain is consistent across all participants, indicating robust decoding boost. Accuracy peaks at 66% top-1 and 93% top-5 for the best participant—remarkable given the inherently noisy nature of EEG signals. In inter-participant evaluation, the improvement is more modest—just over 1% in both top-1 and top-5 accuracy—reflecting the substantial challenges of cross-subject decoding, including variability in neural responses (Wei et al., 2021) and individual differences in visual processing (De Haas et al., 2019).

We next examined whether CUBE's improvements generalise across different vision encoders. Figure 4 shows that CUBE yields a statistically significant 5% increase in object decoding accuracy across all seven OpenCLIP encoders (Cherti et al., 2023). To test whether this boost arises solely from the semantic structure of I^{sam} rather than from colour features, we trained CUBE variants that aligned EEG with visual features from greyscale I^{sam}, where colour was removed but semantics preserved. These models achieved only a modest 1% average gain, which was inconsistent across encoders and even reduced performance for the CoCa-B32 encoder. Together, these results indicate that colour features provide a substantive and reliable contribution to object decoding.

Table 2 reports object decoding accuracy for CUBE and two comparison models on the THINGS MEG dataset (Hebart et al., 2023). The results closely parallel the EEG findings: CUBE improves intra-participant accuracy by roughly 5% for both top-1 and top-5, and inter-participant accuracy by about 1%. These findings demonstrate that the decoding boost provided by colour features generalises beyond EEG to other neuroimaging modalities.

Table 1: Object decoding performance of CUBE on the THINGS EEG dataset (Gifford et al., 2022) across 200 object categories. Comparison methods from the literature include BraVL (Du et al., 2023), NICE (Song et al., 2024), ATM (Li et al., 2024), IDES (Akbarinia, 2024), VE-SDN (Chen et al., 2024), and UBP (Wu et al., 2025). Table cells are colour-coded from green to yellow as accuracies decrease.

	Top-1											Top-5										
Method	S01	S02	S03	S04	S05	S06	S07	S08	S09	S10	**AVG**	S01	S02	S03	S04	S05	S06	S07	S08	S09	S10	**AVG**
Intra-participant																						
BraVL	.061	.049	.056	.050	.040	.060	.065	.088	.043	.070	.058	.179	.149	.174	.151	.134	.182	.204	.237	.140	.197	.175
NICE	.132	.135	.145	.206	.101	.165	.170	.229	.154	.174	.161	.395	.403	.427	.527	.315	.440	.421	.561	.416	.458	.436
ATM	.256	.220	.250	.314	.129	.213	.305	.388	.344	.291	.285	.604	.545	.624	.609	.430	.511	.615	.720	.515	.635	.604
IDES	.355	.330	.305	.365	.225	.345	.300	.440	.355	.370	.339	.645	.650	.695	.715	.570	.735	.700	.755	.705	.705	.688
VE-SDN	.326	.344	.387	.398	.294	.345	.345	.493	.390	.398	.372	.637	.699	.735	.720	.586	.688	.683	.798	.696	.753	.699
UBP	.412	.512	.512	.511	.422	.575	.490	.586	.451	.615	.509	.705	.809	.820	.769	.728	.835	.799	.858	.762	.882	.797
CUBE	.460	.565	.615	.605	.450	.595	.530	.630	.555	.655	**.566**	.770	.855	.895	.835	.790	.880	.845	.920	.855	.930	**.858**
Inter-participant																						
BraVL	.023	.015	.014	.017	.015	.018	.021	.022	.016	.023	.018	.080	.063	.059	.067	.056	.072	.081	.076	.064	.085	.070
NICE	.076	.059	.060	.063	.044	.056	.056	.063	.057	.084	.062	.228	.205	.223	.207	.183	.222	.197	.220	.176	.283	.214
ATM	.105	.071	.119	.147	.070	.111	.161	.150	.049	.205	.118	.268	.248	.338	.394	.239	.358	.435	.403	.227	.465	.337
IDES	.090	.165	.090	.135	.085	.100	.075	.135	.100	.190	.117	.280	.340	.215	.355	.260	.350	.235	.315	.295	.365	.320
UBP	.115	.155	.098	.130	.088	.117	.102	.122	.155	.160	.124	.297	.400	.270	.323	.338	.310	.238	.322	.405	.435	.334
CUBE	.140	.175	.080	.140	.135	.130	.095	.125	.195	.155	**.137**	.335	.410	.235	.350	.310	.355	.335	.385	.415	.420	**.355**

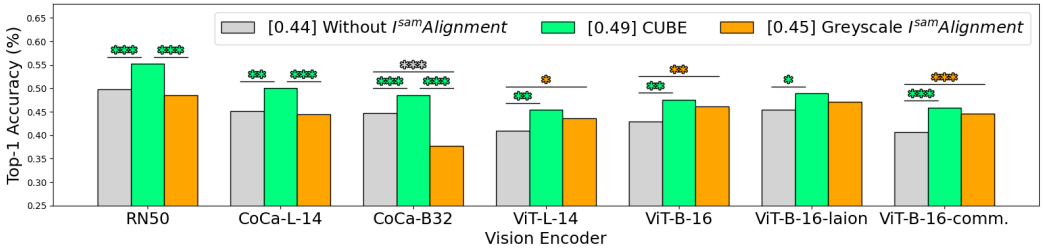

Figure 4: The impact of colour features on object–recognition decoding. Decoding performance for seven vision encoders on the THINGS EEG dataset (Gifford et al., 2022), evaluated across 200 object categories. All encoders were pretrained using OpenCLIP (Cherti et al., 2023). Asterisks denote significant differences between the compared conditions.

Table 2: Object decoding performance of CUBE on the THINGS MEG dataset (Hebart et al., 2023) across 200 object categories. Comparison methods from the literature include BraVL (Du et al., 2023), NICE (Song et al., 2024), and UBP (Wu et al., 2025). Table cells are colour-coded from green to yellow as accuracies decrease.

	Top-1					Top-5				
Method	S01	S02	S03	S04	**AVG**	S01	S02	S03	S04	**AVG**
Intra-participant										
NICE	.096	.185	.142	.090	.128	.278	.478	.416	.266	.360
UBP	.150	.460	.273	.185	.267	.380	.805	.590	.435	.552
CUBE	.165	.520	.335	.235	**.314**	.435	.850	.625	.480	**.598**
Inter-participant										
UBP	.020	.015	.027	.025	.022	.057	.172	.100	.080	.104
CUBE	.020	.050	.035	.025	**.033**	.075	.175	.100	.105	**.114**

5 Discussion

EEG, whose core technology dates back nearly a century (Berger, 1929), measures tiny fluctuations in ionic potentials to non-invasively record brain activity. The resulting signal is notoriously noisy and has low spatial resolution, reflecting the aggregate activity of billions of neurons (Azevedo et al., 2009; Goriely, 2025). Despite these limitations, EEG has long been an invaluable tool—both clinically and for advancing our understanding of the brain. Recent work suggests that the decoding capabilities of EEG, and neuroimaging more broadly, are undergoing a major leap forward, enabled by AI and large-scale datasets. In particular, EEG benefits from its exceptionally high temporal resolution. We can now decode speech from three seconds of EEG with remarkable accuracy (Défossez et al., 2023), and in the visual domain, emerging work is progressing toward 3D object reconstruction (Guo et al., 2025) and even video decoding (Liu et al., 2024).

Here, we showed that object decoding reaches a remarkable 57% accuracy—far above the 1/200 chance level—from just one second of EEG. Likewise, we demonstrate for the first time that perceived colours in complex natural images can be decoded with high reliability (F-score = 0.5). It is striking that EEG recorded during natural viewing—without any colour cues—can recover colours with reliability approaching that of average behavioural responses. One might expect decoding performance to improve even further (Robinson et al., 2023) if neural and behavioural data were obtained from the same individuals, allowing models to more precisely capture individual differences in perception and colour (Bosten, 2022; De Haas et al., 2019).

5.1 The interaction between colour and object

Colour information plays an important role in object recognition (Bramão et al., 2011; Rosenthal et al., 2018), and, conversely, object and scene semantics influence perceived colours (Bloj et al., 1999; Hansen et al., 2006; Akbarinia, 2025). Nevertheless, the interaction between colour and object remains surprisingly little understood (Teichmann et al., 2020; Taylor & Xu, 2021; Gegenfurtner, 2025). Our results from colour and object decoding further support this bidirectional relationship: object alignment boosts colour decoding by 4%, and colour features boost object decoding by 5%.

A central question in visual neuroscience is whether colour is encoded first—with object boundaries emerging later, as in bottom-up region growing—or whether objects are parsed first and colours filled in afterwards, reflecting a more top-down process. To explore this, we compared colour and object decoding within CUBE. A direct comparison, however, is not straightforward. First, evaluation metrics differ: F-score for multi-class colour decoding versus accuracy for single-class object decoding. Second, the ground truths differ: object labels are objective ("this is a snail"), whereas colour annotations represent subjective averages. Third, colour and object are tightly intertwined, creating potential confounds. Colour-diagnostic objects can enhance colour decoding (e.g., bananas are yellow), while natural colour statistics can bias object recognition (Tanaka & Presnell, 1999; Therriault et al., 2009). For instance, a uniformly orange sphere may be decoded as an orange, whereas a striped orange sphere may instead be classified as a volleyball. Fourth, material properties add further complexity (Schmidt et al., 2025), since some materials (e.g., wood) have characteristic colour–texture associations.

Despite methodological challenges, within our framework colour decoding performs relatively better when normalised to chance and noise ceilings: the colour-decoding F-score reaches 0.70 on a 0–1 scale, exceeding the object-decoding accuracy of 0.57. To examine temporal dynamics, we trained models on 100 ms EEG segments from different intervals (e.g., $[0, 100)$, $[100, 200)$, $[900, 1000)$) and also tested shorter windows from stimulus onset up to time t. These analyses, shown in Figure 5, reveal that colour decoding becomes statistically significant substantially earlier than

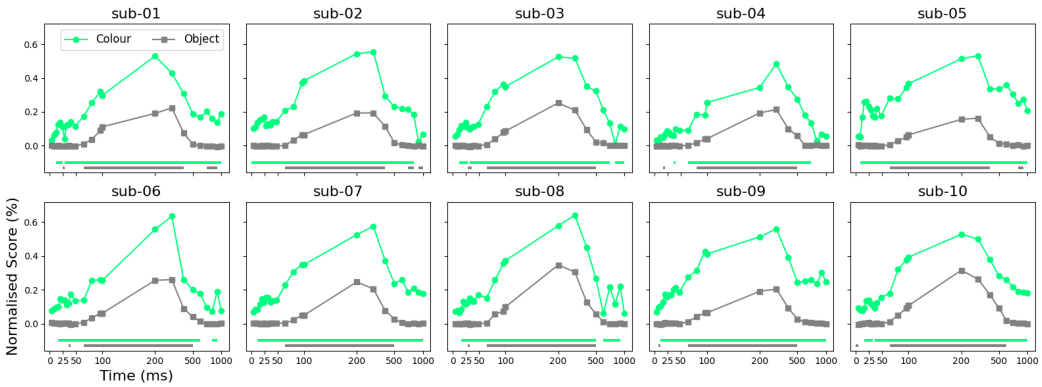

Figure 5: Colour and object decoding across temporal epochs. The x-axis (0 ms) marks stimulus onset; the y-axis shows normalised F-scores (colour) and top-1 accuracy (object), scaled by chance level and noise ceiling. Data points after 100 ms use EEG intervals $[t-100, t)$, and points before 100 ms from $[0, t)$. Significance lines indicate intervals with decoding above chance ($p < 0.01$).

object decoding—on average at 18 ms for colour versus 41 ms for object. This temporal advantage suggests that colour may play a more prominent role in the earliest stages of visual perception.

5.2 Limitations of the EEG decoding

What are the limitations of current EEG decoding frameworks? A major challenge is the substantially lower cross-participant performance compared with within-participant decoding, driven by large individual differences in visual processing across both sensory and perceptual levels (De Haas et al., 2019; Bosten, 2022). Signal quality is further influenced by technical factors such as electrode placement and impedance, which can vary across sessions and participants. One promising direction is to pretrain models on large, diverse EEG datasets (Huang et al., 2025), analogous to large language models, and then fine-tune them for individual participants. This strategy may help bridge the gap between generalisation and personalisation, and could be particularly valuable for practical neuroimaging applications, such as brain–machine interfaces for individuals with severe motor impairments (Chaudhary et al., 2015).

6 Conclusion

In this article, we introduced CUBE (ColoUr and oBjEct decoding) and highlighted the importance of jointly representing colour and object features in neuroimaging decoding. Our results show that EEG signals contain reliable, decodable colour information—even during object recognition tasks with no explicit colour cues and under very brief viewing conditions (100 ms). We further demonstrated that incorporating colour features into a standard contrastive-learning alignment framework boosts object decoding by about 5% across participants in both EEG and MEG. Overall, our findings open a novel avenue for future work: theoretically, enabling the investigation of individual colour perception in more ecological settings by decoding colour from neuroimaging signals under naturalistic viewing; and practically, offering potential benefits for applications in brain–computer interfaces and clinical psychology.

References

Arash Akbarinia. Optimising eeg decoding with refined sampling and multimodal feature integration. arXiv preprint arXiv:2409.20086, 2024.

Arash Akbarinia. Exploring the categorical nature of colour perception: Insights from artificial networks. Neural Networks, 181:106758, 2025.

Arash Akbarinia and C Alejandro Parraga. Colour constancy beyond the classical receptive field. IEEE transactions on pattern analysis and machine intelligence, 40(9):2081–2094, 2017.

Emily J Allen, Ghislain St-Yves, Yihan Wu, Jesse L Breedlove, Jacob S Prince, Logan T Dowdle, Matthias Nau, Brad Caron, Franco Pestilli, Ian Charest, et al. A massive 7t fmri dataset to bridge cognitive neuroscience and artificial intelligence. Nature neuroscience, 25(1):116–126, 2022.

Frederico AC Azevedo, Ludmila RB Carvalho, Lea T Grinberg, José Marcelo Farfel, Renata EL Ferretti, Renata EP Leite, Wilson Jacob Filho, Roberto Lent, and Suzana Herculano-Houzel. Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain. Journal of Comparative Neurology, 513(5):532–541, 2009.

Yohann Benchetrit, Hubert Banville, and Jean-Rémi King. Brain decoding: toward real-time reconstruction of visual perception. International Conference on Representation Learning, 2023.

Hans Berger. Über das elektroenkephalogramm des menschen. Archiv für psychiatrie und nervenkrankheiten, 87(1):527–570, 1929.

Marina G Bloj, Daniel Kersten, and Anya C Hurlbert. Perception of three-dimensional shape influences colour perception through mutual illumination. Nature, 402(6764):877–879, 1999.

Jenny M Bosten. Do you see what i see? diversity in human color perception. Annual review of vision science, 8(1):101–133, 2022.

Inês Bramão, Alexandra Reis, Karl Magnus Petersson, and Luís Faísca. The role of color information on object recognition: A review and meta-analysis. Acta psychologica, 138(1):244–253, 2011.

Ujwal Chaudhary, Niels Birbaumer, and Marco Rocha Curado. Brain-machine interface (bmi) in paralysis. Annals of physical and rehabilitation medicine, 58(1):9–13, 2015.

Tushar Chauhan, Ivana Jakovljev, Lindsay N Thompson, Sophie M Wuerger, and Jasna Martinovic. Decoding of eeg signals reveals non-uniformities in the neural geometry of colour. NeuroImage, 268:119884, 2023.

Hongzhou Chen, Lianghua He, Yihang Liu, and Longzhen Yang. Visual neural decoding via improved visual-eeg semantic consistency. arXiv preprint arXiv:2408.06788, 2024.

Mehdi Cherti, Romain Beaumont, Ross Wightman, Mitchell Wortsman, Gabriel Ilharco, Cade Gordon, Christoph Schuhmann, Ludwig Schmidt, and Jenia Jitsev. Reproducible scaling laws for contrastive language-image learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 2818–2829, 2023.

Benjamin De Haas, Alexios L Iakovidis, D Samuel Schwarzkopf, and Karl R Gegenfurtner. Individual differences in visual salience vary along semantic dimensions. Proceedings of the National Academy of Sciences, 116(24):11687–11692, 2019.

Alexandre Défossez, Charlotte Caucheteux, Jérémy Rapin, Ori Kabeli, and Jean-Rémi King. Decoding speech perception from non-invasive brain recordings. Nature Machine Intelligence, 5(10), 2023.

Changde Du, Kaicheng Fu, Jinpeng Li, and Huiguang He. Decoding visual neural representations by multimodal learning of brain-visual-linguistic features. IEEE Transactions on Pattern Analysis and Machine Intelligence, 45(9):10760–10777, 2023.

Karl R Gegenfurtner. Cortical mechanisms of colour vision. Nature Reviews Neuroscience, 4(7): 563–572, 2003.

Karl R. Gegenfurtner. The verriest lecture: Color vision from pixels to objects. J. Opt. Soc. Am. A, 42(5):B313–B328, 2025.

Karl R Gegenfurtner and Jochem Rieger. Sensory and cognitive contributions of color to the recognition of natural scenes. Current Biology, 10(13):805–808, 2000.

Alessandro T Gifford, Kshitij Dwivedi, Gemma Roig, and Radoslaw M Cichy. A large and rich eeg dataset for modeling human visual object recognition. NeuroImage, 264:119754, 2022.

Raquel Gil Rodríguez, Laysa Hedjar, Matteo Toscani, Dar'ya Guarnera, Giuseppe Claudio Guarnera, and Karl R Gegenfurtner. Color constancy mechanisms in virtual reality environments. Journal of Vision, 24(5):6–6, 2024.

Alain Goriely. Eighty-six billion and counting: do we know the number of neurons in the human brain? Brain, 148(3):689–691, 2025.

Zhanqiang Guo, Jiamin Wu, Yonghao Song, Jiahui Bu, Weijian Mai, Qihao Zheng, Wanli Ouyang, and Chunfeng Song. Neuro-3d: Towards 3d visual decoding from eeg signals. In Proceedings of the Computer Vision and Pattern Recognition Conference, pp. 23870–23880, 2025.

Jasper E Hajonides, Anna C Nobre, Freek van Ede, and Mark G Stokes. Decoding visual colour from scalp electroencephalography measurements. NeuroImage, 237:118030, 2021.

Thorsten Hansen, Maria Olkkonen, Sebastian Walter, and Karl R Gegenfurtner. Memory modulates color appearance. Nature neuroscience, 9(11):1367–1368, 2006.

Martin N Hebart, Adam H Dickter, Alexis Kidder, Wan Y Kwok, Anna Corriveau, Caitlin Van Wicklin, and Chris I Baker. Things: A database of 1,854 object concepts and more than 26,000 naturalistic object images. PloS one, 14(10):e0223792, 2019.

Martin N Hebart, Oliver Contier, Lina Teichmann, Adam H Rockter, Charles Y Zheng, Alexis Kidder, Anna Corriveau, Maryam Vaziri-Pashkam, and Chris I Baker. Things-data, a multimodal collection of large-scale datasets for investigating object representations in human brain and behavior. Elife, 12, 2023.

Katherine L Hermann, Shridhar R Singh, Isabelle A Rosenthal, Dimitrios Pantazis, and Bevil R Conway. Temporal dynamics of the neural representation of hue and luminance polarity. Nature communications, 13(1):661, 2022.

Shuai Huang, Huan Luo, Haodong Jing, Qixian Zhang, Litao Chang, Yating Feng, Xiao Lin, Chendong Qin, Han Chen, Shuwen Jia, Siyi Sun, and Yongxiong Wang. Need: Cross-subject and cross-task generalization for video and image reconstruction from eeg signals. Advances in Neural Information Processing Systems, 36, 2025.

Helene Intraub. Rapid conceptual identification of sequentially presented pictures. Journal of Experimental Psychology: Human Perception and Performance, 7(3):604, 1981.

Paul Jaccard. Étude comparative de la distribution florale dans une portion des alpes et des jura. Bull Soc Vaudoise Sci Nat, 37:547–579, 1901.

Christian Keysers, D-K Xiao, Peter Földiák, and David I Perrett. The speed of sight. Journal of cognitive neuroscience, 13(1):90–101, 2001.

Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In Proceedings of the IEEE/CVF international conference on computer vision, pp. 4015–4026, 2023.

Rosa Lafer-Sousa, Katherine L Hermann, and Bevil R Conway. Striking individual differences in color perception uncovered by 'the dress' photograph. Current Biology, 25(13):R545–R546, 2015.

Dongyang Li, Chen Wei, Shiying Li, Jiachen Zou, and Quanying Liu. Visual decoding and reconstruction via eeg embeddings with guided diffusion. Advances in Neural Information Processing Systems, 2024.

Xuan-Hao Liu, Yan-Kai Liu, Yansen Wang, Kan Ren, Hanwen Shi, Zilong Wang, Dongsheng Li, Bao-Liang Lu, and Wei-Long Zheng. Eeg2video: Towards decoding dynamic visual perception from eeg signals. Advances in Neural Information Processing Systems, 37:72245–72273, 2024.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. International Conference on Representation Learning, 2017.

Per Møller and Anya C Hurlbert. Psychophysical evidence for fast region-based segmentation processes in motion and color. Proceedings of the National Academy of Sciences, 93(14):7421–7426, 1996.

John D Mollon, Jenny M Bosten, David H Peterzell, and Michael A Webster. Individual differences in visual science: What can be learned and what is good experimental practice? Vision research, 141:4–15, 2017.

Stefan Palan and Christian Schitter. Prolific. ac—a subject pool for online experiments. Journal of behavioral and experimental finance, 17:22–27, 2018.

C. Alejandro Parraga and Arash Akbarinia. Color name applications in computer vision. Encyclopedia of Color Science and Technology, pp. 1–7, 2020.

WM Paulus, V Hömberg, K Cunningham, AM Halliday, and N Rohde. Colour and brightness components of foveal visual evoked potentials in man. Electroencephalography and clinical neurophysiology, 58(2):107–119, 1984.

Ian ML Pennock, Chris Racey, Emily J Allen, Yihan Wu, Thomas Naselaris, Kendrick N Kay, Anna Franklin, and Jenny M Bosten. Color-biased regions in the ventral visual pathway are food selective. Current Biology, 33(1):134–146, 2023.

Mary C Potter, Brad Wyble, Carl Erick Hagmann, and Emily S McCourt. Detecting meaning in rsvp at 13 ms per picture. Attention, Perception, & Psychophysics, 76(2):270–279, 2014.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. International conference on machine learning, pp. 8748–8763, 2021.

D Regan. Evoked potential and psychophysical correlates of changes in stimulus colour and intensity. Vision Research, 10(2):163–178, 1970.

Amanda K Robinson, Genevieve L Quek, and Thomas A Carlson. Visual representations: insights from neural decoding. Annual Review of Vision Science, 9(1):313–335, 2023.

Isabelle Rosenthal, Sivalogeswaran Ratnasingam, Theodros Haile, Serena Eastman, Josh Fuller-Deets, and Bevil R Conway. Color statistics of objects, and color tuning of object cortex in macaque monkey. Journal of vision, 18(11):1–1, 2018.

Isabelle A Rosenthal, Shridhar R Singh, Katherine L Hermann, Dimitrios Pantazis, and Bevil R Conway. Color space geometry uncovered with magnetoencephalography. Current Biology, 31 (3):515–526, 2021.

Ana Rozman, Tushar Chauhan, and Jasna Martinovic. Characterising representations of hue and saturation in the cortex using information decoding. bioRxiv, pp. 2024–12, 2024.

Filipp Schmidt, Martin N Hebart, Alexandra C Schmid, and Roland W Fleming. Core dimensions of human material perception. Proceedings of the National Academy of Sciences, 122(10): e2417202122, 2025.

Paul Scotti, Atmadeep Banerjee, Jimmie Goode, Stepan Shabalin, Alex Nguyen, Aidan Dempster, Nathalie Verlinde, Elad Yundler, David Weisberg, Kenneth Norman, et al. Reconstructing the mind's eye: fmri-to-image with contrastive learning and diffusion priors. Advances in Neural Information Processing Systems, 36, 2024.

Yonghao Song, Bingchuan Liu, Xiang Li, Nanlin Shi, Yijun Wang, and Xiaorong Gao. Decoding natural images from eeg for object recognition. International Conference on Representation Learning, 2024.

David W Sutterer, Andrew J Coia, Vincent Sun, Steven K Shevell, and Edward Awh. Decoding chromaticity and luminance from patterns of eeg activity. Psychophysiology, 58(4):e13779, 2021.

James Tanaka, Daniel Weiskopf, and Pepper Williams. The role of color in high-level vision. Trends in cognitive sciences, 5(5):211–215, 2001.

James W Tanaka and Lynn M Presnell. Color diagnosticity in object recognition. Perception & Psychophysics, 61(6):1140–1153, 1999.

JohnMark Taylor and Yaoda Xu. Joint representation of color and form in convolutional neural networks: A stimulus-rich network perspective. PLoS One, 16(6):e0253442, 2021.

Lina Teichmann, Genevieve L Quek, Amanda K Robinson, Tijl Grootswagers, Thomas A Carlson, and Anina N Rich. The influence of object-color knowledge on emerging object representations in the brain. Journal of Neuroscience, 40(35):6779–6789, 2020.

David J Therriault, Richard H Yaxley, and Rolf A Zwaan. The role of color diagnosticity in object recognition and representation. Cognitive processing, 10(4):335–342, 2009.

Chun-Shu Wei, Corey J Keller, Junhua Li, Yuan-Pin Lin, Masaki Nakanishi, Johanna Wagner, Wei Wu, Yu Zhang, and Tzyy-Ping Jung. Inter-and intra-subject variability in brain imaging and decoding, 2021.

Felix A Wichmann, Lindsay T Sharpe, and Karl R Gegenfurtner. The contributions of color to recognition memory for natural scenes. Journal of Experimental Psychology: Learning, Memory, and Cognition, 28(3):509, 2002.

Christoph Witzel and Karl R Gegenfurtner. Color perception: Objects, constancy, and categories. Annual review of vision science, 4(1):475–499, 2018.

Haitao Wu, Qing Li, Changqing Zhang, Zhen He, and Xiaomin Ying. Bridging the vision-brain gap with an uncertainty-aware blur prior. In Proceedings of the Computer Vision and Pattern Recognition Conference, pp. 2246–2257, 2025.

Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. arXiv preprint arXiv:2205.01917, 2022.

Kaifan Zhang, Lihuo He, Xin Jiang, Wen Lu, Di Wang, and Xinbo Gao. Cognitioncapturer: Decoding visual stimuli from human eeg signal with multimodal information. Proceedings of the AAAI Conference on Artificial Intelligence, 39(13):14486–14493, 2025.

A Qualitative examples

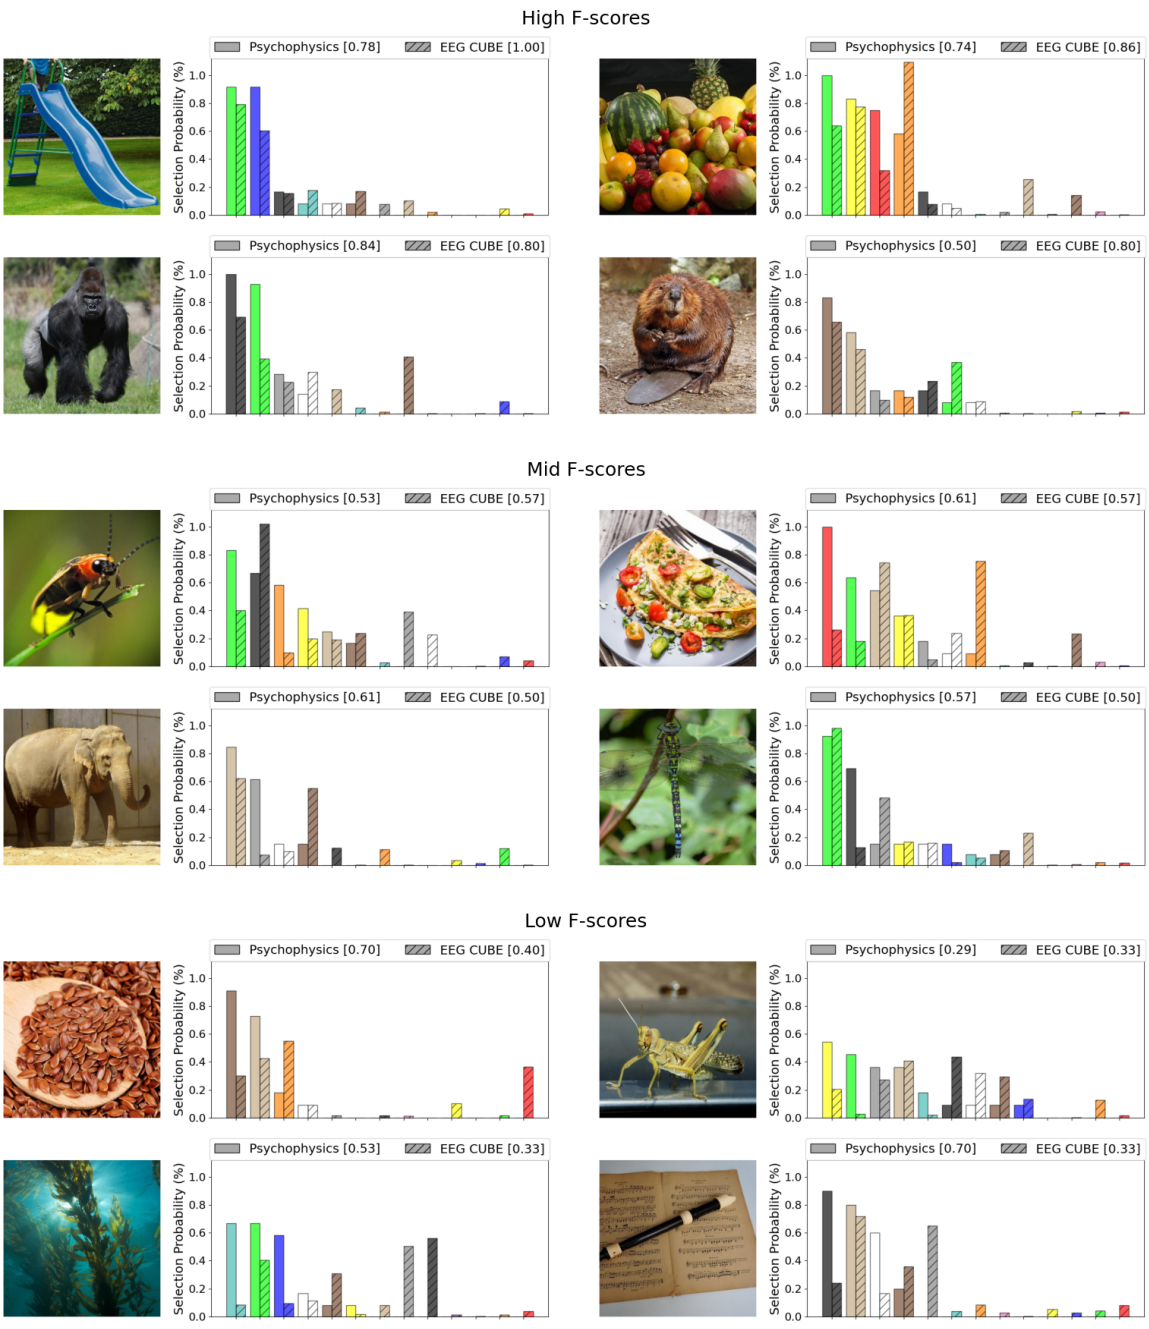

Figure S1: Examples of test images alongside the corresponding colour selections made by human participants in a psychophysical experiment and the colours decoded by CUBE from EEG signals. The reported F-score for the psychophysical data reflects the average inter-participant agreement computed using a leave-one-out strategy, whereas the EEG CUBE F-score represents the agreement between the model's predictions and the average human selections.