

Language Aligned Visual Representations Predict Human Behavior in Naturalistic Learning Tasks

Can Demircan¹ ✉, Tankred Saanum¹, Leonardo Pettini^{2,3}, Marcel Binz¹, Blazej M Baczkowski^{3,4}, Paula Kaanders¹, Christian F Doeller^{3,5,6,7}, Mona M Garvert^{3,8,9}, and Eric Schulz¹

¹Max Planck Research Group Computational Principles of Intelligence, Max Planck Institute for Biological Cybernetics, Tübingen 72076, Germany; ²Max Planck School of Cognition, Leipzig 04103, Germany; ³Department of Psychology, Max Planck Institute for Human Cognitive & Brain Sciences, Leipzig 04103, Germany; ⁴Department of Cognitive Psychology, Universität Hamburg 20146, Germany; ⁵Kavli Institute for Systems Neuroscience, Centre for Neural Computation, The Egil & Pauline Braathen & Fred Kavli Centre for Cortical Microcircuits, Jebsen Centre for Alzheimer's Disease NTNU, Trondheim 7030, Norway; ⁶Wilhelm Wundt Institute of Psychology, Leipzig University, Leipzig 04109, Germany; ⁷Department of Psychology, Technical University Dresden, Dresden, Germany; ⁸Julius-Maximilians-Universität Würzburg, Faculty of Human Sciences, Würzburg 97070, Germany; ⁹Max Planck Research Group NeuroCode, Max Planck Institute for Human Development, Berlin 14195, Germany

Humans possess the ability to identify and generalize relevant features of natural objects, which aids them in various situations. To investigate this phenomenon and determine the most effective representations for predicting human behavior, we conducted two experiments involving category learning and reward learning. Our experiments used realistic images as stimuli, and participants were tasked with making accurate decisions based on novel stimuli for all trials, thereby necessitating generalization. In both tasks, the underlying rules were generated as simple linear functions using stimulus dimensions extracted from human similarity judgments. Notably, participants successfully identified the relevant stimulus features within a few trials, demonstrating effective generalization. We performed an extensive model comparison, evaluating the trial-by-trial predictive accuracy of diverse deep learning models' representations of human choices. Intriguingly, representations from models trained on both text and image data consistently outperformed models trained solely on images, even surpassing models using the features that generated the task itself. These findings suggest that language-aligned visual representations possess sufficient richness to describe human generalization in naturalistic settings and emphasize the role of language in shaping human cognition.

shape, or brand. People can use these features and make predictions about how tasteful that apple could be, the environmental impacts of growing it, or the significance of it in different mythological and religious settings. How can we describe this kind of generalization behavior across realistic stimuli? And which representations underlie this ability in the first place?

One way to adapt sensory representations to different situations is to learn functions over them. Cognitive psychologists have studied how people learn these functions using simple learning paradigms. In these tasks, participants are repeatedly shown abstract stimuli, such as geometric objects of different colors and shapes, and they are instructed to learn what kind of objects are more rewarding or belong to a certain category through feedback (3–5). This approach has been fruitful for understanding the kinds of learning strategies people use, such as rule-based learning (6), similarity-based learning (7), or heuristics strategies (8). However, some important aspects of real-life learning have been ignored in these tasks. First, in real-life, associations are often not learned purely through repeated trial-error, but generalizing to novel situations is needed (9). Second, in this line of work, how to represent the stimuli has largely been ignored, given that stimuli are simple and are therefore easy to represent. Therefore, how well humans can generalize in learning tasks that involve naturalistic stimuli and how humans represent these stimuli in the first place remains unknown.

To address these gaps, we designed a category learning and a reward learning task. These two tasks did not involve any repeating stimuli, requiring generalization for any successful decision. Additionally, we used naturalistic images, which include features humans acquired throughout their lifespans. Naturalistic stimuli are high dimensional (10), and it may be challenging to assign credit to the relevant stimulus dimension. We used these rich naturalistic tasks to study whether and how fast humans can discover the relevant stimulus features from limited experience and exploit this knowledge for novel decisions.

To test how people represent naturalistic stimuli, we turn to deep neural networks (DNNs). Across various domains of cognitive science, these models have

cognitive psychology | generalization | deep learning | vision models | language models

Correspondence: can.demircan@tuebingen.mpg.de

Introduction

Generalization is a notoriously difficult challenge. Especially in realistic environments, where there are infinitely many different features to describe objects, extracting the dimensions that matter becomes difficult. Nonetheless, humans and other animals can generalize efficiently from experience to novel situations, and adapt their representations to the task at hand. Consider an apple, as an example. There can be many features describing an apple such as its color, taste,




Figure 1. Task descriptions. (A) An example trial from the category learning task, where an incorrect decision is made. (B) An example trial from the reward learning task where the best option is chosen and highlighted in orange. (C) Examples of where different images fall in the three features of the embedding used to generate the task and how these features relate to category membership and associated rewards in our tasks. The feature labels are those used by (1) who constructed the embedding. The original images are replaced with copyright-free alternatives from the THINGSplus database (2).

been successful in predicting behavior and neural data in response to naturalistic stimuli. For example, when prompted with the same images, the representational hierarchy of DNNs have shown to be similar to that of the ventral visual stream in primates (11, 12). These models have also been shown to behave similarly to humans in categorization tasks (13, 14). In the auditory domain, similar hierarchical correspondence of representations have been found between DNNs and the auditory hierarchy in the human brain (15). Moving up to higher cognition, it has been shown that representations of large language models can predict brain activity in response to natural language as well as human reading times (16). Whether and what kind of DNN representations can be useful for predicting human behavior in high-dimensional learning tasks that require generalization remains untested.

The stimuli in our category and reward learning tasks were sampled from the THINGS database (10). To assign rewards and category membership to these images, we constructed sparse linear functions over an embedding built to capture human similarity judgment of the THINGS stimuli (1). Participants were exposed to a novel stimulus on each trial and could therefore not solve the task by associating specific stimuli with specific outcomes. Instead, they had to apply their knowledge about regularities between the stimuli they had seen already to solve the task efficiently by means of generalization. We found that humans learned to do this surprisingly quickly, suggesting that they could identify relevant stimulus dimensions within just a few trials and use this knowledge to guide choices. To understand the nature of the representations participants used when solving this task, we first extracted representations from DNNs of various architectures that were trained under different regimes and different modalities of data. Then, we trained linear models over these representations to predict humans' learning trajec-

ries. While all DNN representations could predict human behavior, language-aligned visual representations were consistently better at predicting human choices than their uni-modal counterparts. Our modeling results show that human learning can be modeled using simple strategies when provided with sufficiently rich representations and that DNNs trained on multi-modal data can provide these rich representations. Taken together, our paradigm and behavioral results pave the way to a deeper understanding of representational structure in the human mind.

Results

Category Learning

Human participants ($n = 91$) completed 120 trials of an online category learning task, where they were presented with a novel image in each trial. They were asked to deliver these images to one of two dinosaurs, *July* or *Foltly*, using key presses. Participants were told that the two dinosaurs had completely non-overlapping preferences for what gifts they enjoyed. After each trial, we gave participants feedback on whether their choice of delivery was correct. An example trial from the task is shown in Fig. 1A.

Participants were assigned to one of three conditions, where in each condition the category boundary was dependent on a different rule. These rules depended on a feature of an embedding that reflected human similarity judgments on the THINGS database images (1, 10). Specifically, the authors built an embedding to predict human choices in an odd-one-out task and found that an embedding with 49 human interpretable features described human choices the best. The three chosen features were those that explained the most variance in the embedding and are displayed in Fig. 1C. For each participant, 120 unique stimuli from the THINGS database were sampled. A median split over the assigned feature

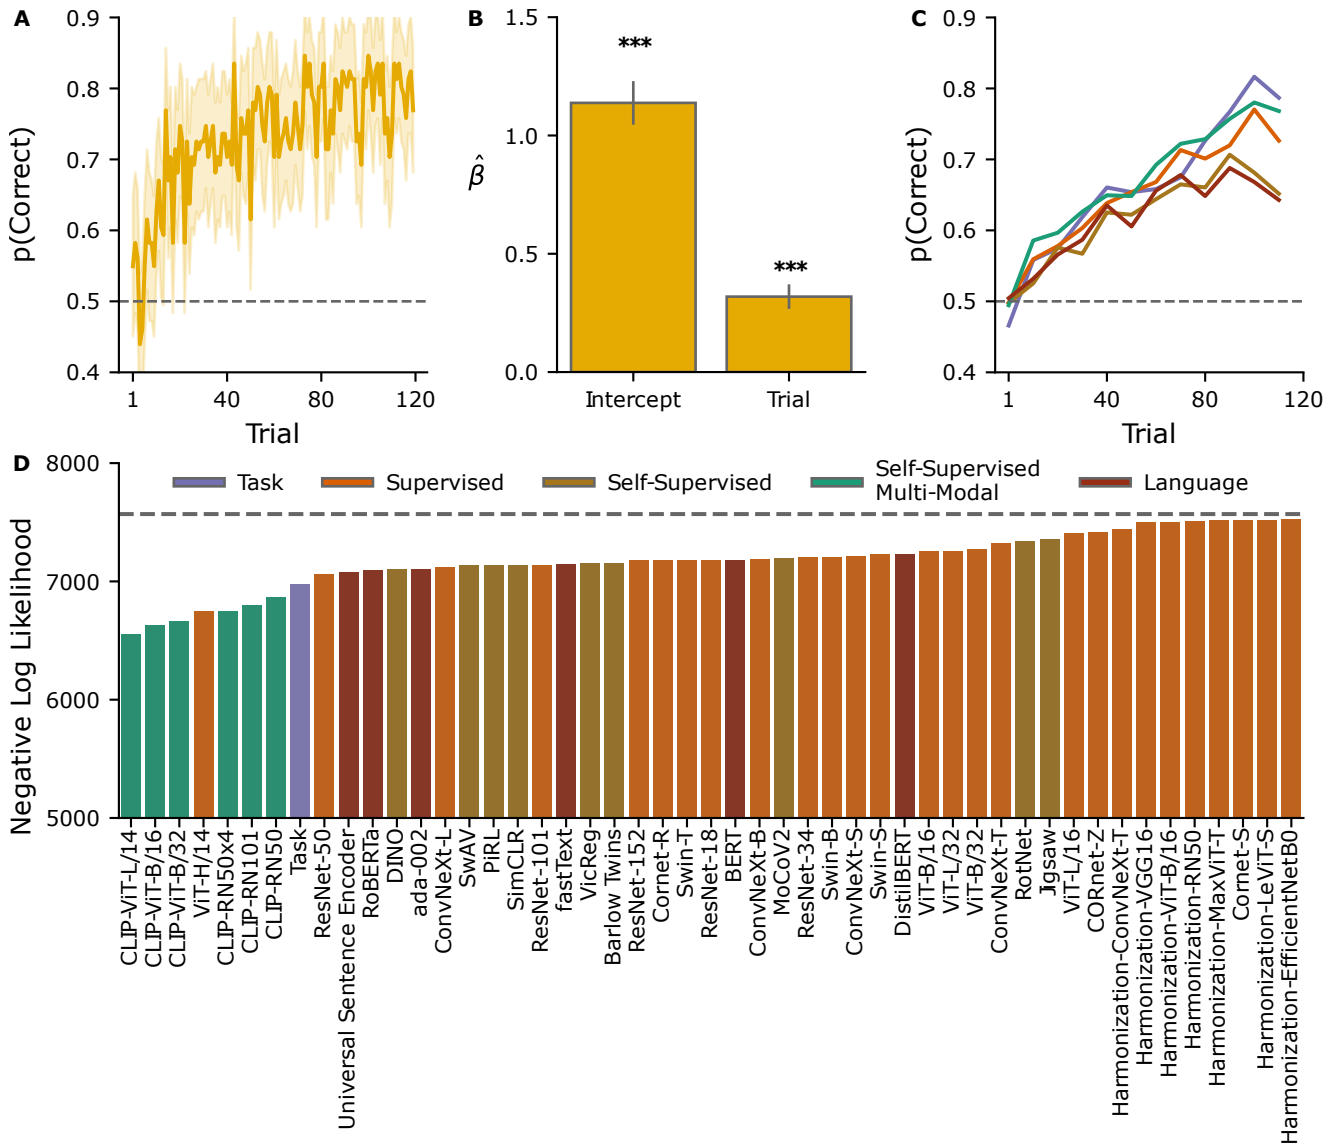


Figure 2. Category learning results. (A) Participants’ performance across trials. Shaded lines indicate 95% confidence intervals. (B) Coefficient values from the mixed-effects logistic regression analysis to predict participant choice. Error bars indicate standard errors. (C) Performance of the computational models. For each category, we show the learning trajectory of the model predicting participant behavior the best. The curves are smoothed every 10th trial. (D) Cross-validated negative log likelihoods of each representation. Lower values indicate better fits to human behavior. The dashed horizontal line indicates chance level performance.

of the sampled stimuli determined the category boundary.

We analyzed participants’ behavior (Fig. 2A), using mixed-effects logistic regression. We predicted whether a participant made the correct choice using the trial number as a fixed effect. Additionally, we fitted participant-specific random intercepts, as well as random slopes for the trial number and the assigned experimental condition of participants. We found that participants performed the task above chance level ($\hat{\beta} = 1.14 \pm 0.09$, $z = 13.18$, $P < .001$) and that their performance improved over trials ($\hat{\beta} = 0.32 \pm 0.05$, $z = 6.89$, $P < .001$), indicating a learning effect (Fig. 2B). This suggests that humans can very efficiently extract the relevant feature dimension in high-dimensional naturalistic environments despite seeing each individual stimulus only once. To characterize how quickly participants learn the task, we compared the accuracy of partici-

pants at each trial against chance level using a one-sided one-sample t-test. We found that participants performed above chance level starting from trial number 7, $t(90) = 2.25$, $p = .01$. See the *Supplementary Information* for details of testing and Fig. S1 for the results for all trials.

To gain an understanding of the representation that was guiding choices in this task, we assessed what kind of deep neural network representations can best describe humans’ choices in our task. We tested representations from 48 different models in addition to the task embedding (17–38). The models were trained on either text, images, or the two combined in order to solve common machine learning tasks such as text-generation, text-representation, or image-recognition. For the models that were trained on image data, we make the additional distinction between self-supervised and supervised learning paradigms. This is because they are

shown to learn different representations even when they share the same architecture (39). In order to extract representations from models trained on image and multi-modal data, we provided the images used in the task to these models. To extract representations from models trained on text, we provided them with the prompt *This is the image of a X*, where *X* was the category label of the task images. We trained logistic regression models in a sequential manner on the different representations and used the predictions of these models to model participants' choices (see *Methods* for details on representation extraction and modeling). The simulated behavior of these models is shown in Fig. 2C.

First, we observed that all the representations we tested can do our task and predict human behavior above chance level. This is remarkable given that these representations were obtained through training regimes that were independent of our task, and the semantic representations we used to generate the task were unknown to the other representations. Additionally, we found that 7 of the 48 candidate representations described participant behavior better than representations used to generate the task. Of these 7 representations, one was a large vision transformer, trained in a supervised manner. The other 6 were different variants of a self-supervised model, Contrastive Language-Image Pre-training (CLIP) (17), that was trained to represent pairs of images and text as similarly as possible. We saw that regardless of encoder architecture choice, this multi-modal training regime produced representations describing human choices in our task exceptionally well. The rest of the supervised and self-supervised vision models, as well as the language models, had a heterogeneous distribution in how well they predicted human behavior, as visualized in Fig. 2D. The descriptions of the models we extracted representations from are available in the Supplementary Information.

Reward Learning

To test whether our behavioral and modeling findings can generalize across learning tasks, we designed a second learning task. In our second task, human participants ($n = 82$) completed 60 trials of a reward learning paradigm, in which they were asked to maximize their accumulated reward over the course of the task. In each trial, participants were presented with two images and were asked to select one using key presses. After making a choice, the associated reward with each option was shown. An example trial from the task is shown in Fig. 1B.

Participants were assigned to one of the three conditions as in the category learning task. Stimuli were sampled in the same way as the category learning task. For each participant, the values of the task-relevant feature were re-scaled linearly between 0 and 100.

We again analyzed participants' behavior (Fig. 3A) using mixed-effects logistic regression. We predicted whether a participant chose the image on the right us-

ing the reward difference between the two options, the trial number, and the interaction of the two terms as fixed effects. We additionally fitted participant-specific random slopes for these terms and for the assigned experimental condition of participants. We found that participants used the reward difference between the options ($\hat{\beta} = 0.89 \pm 0.07$, $z = 12.56$, $P < .001$). While the trial number did not predict which option the participants chose ($\hat{\beta} = 0.002 \pm 0.03$, $z = 0.09$, $P = .93$), the interaction of the two terms revealed that participants used the reward differences between the options more effectively over trials ($\hat{\beta} = 0.34 \pm 0.04$, $z = 9.30$, $P < .001$; Fig. 3B). These results suggest that the generalization effect we found in the category-learning task transfers to other naturalistic learning tasks as well. To test how quickly participants learn the task, we compared the accuracy of participants at each trial against chance level using a one-sided one-sample t-test. We found that participants performed above chance level starting from trial number 6, $t(81) = 3.01$, $p = .002$. See the *Supplementary Information* for details of testing and Fig. S1 for the results for all trials.

The same representations extracted for the category learning task were used in the reward learning task. We trained linear regression models (again in a sequential manner) on the observations of the participants to predict the associated reward with novel images. We then regressed the reward estimates of the linear models onto participant choice in a mixed-effects logistic regression model (see *Materials and Methods* for details), whose simulated behavior is shown in Fig. 3C.

The modeling results for this task were similar to those of the category learning task. All the models described participant behavior above chance level. Again, 7 models, a large vision transformer, and all variants of CLIP predicted participant behavior better than the task features, and we observed heterogeneity in how well language models, supervised vision models, and self-supervised vision models predicted human behavior. The complete modeling results are shown in Fig. 3D.

Representational Similarity Analyses

Our modeling results showed that CLIP representations are robustly better than their uni-modal counterparts in describing human behavior, suggesting that they capture essential features of object representations. As the next step, we studied what makes these representations particularly successful. To get a better understanding of these results, we conducted two representational similarity analyses (RSA) (40) where we investigated the representations for all the images in the THINGS database. First, we computed Centered Kernel Alignment (CKA) (41) using a linear kernel between the task embedding and every representation we tested. While a CKA score of 0 indicates maximal dissimilarity between representations, a score of 1 indicates maximal similarity. We found that all of the 6 different CLIP representa-

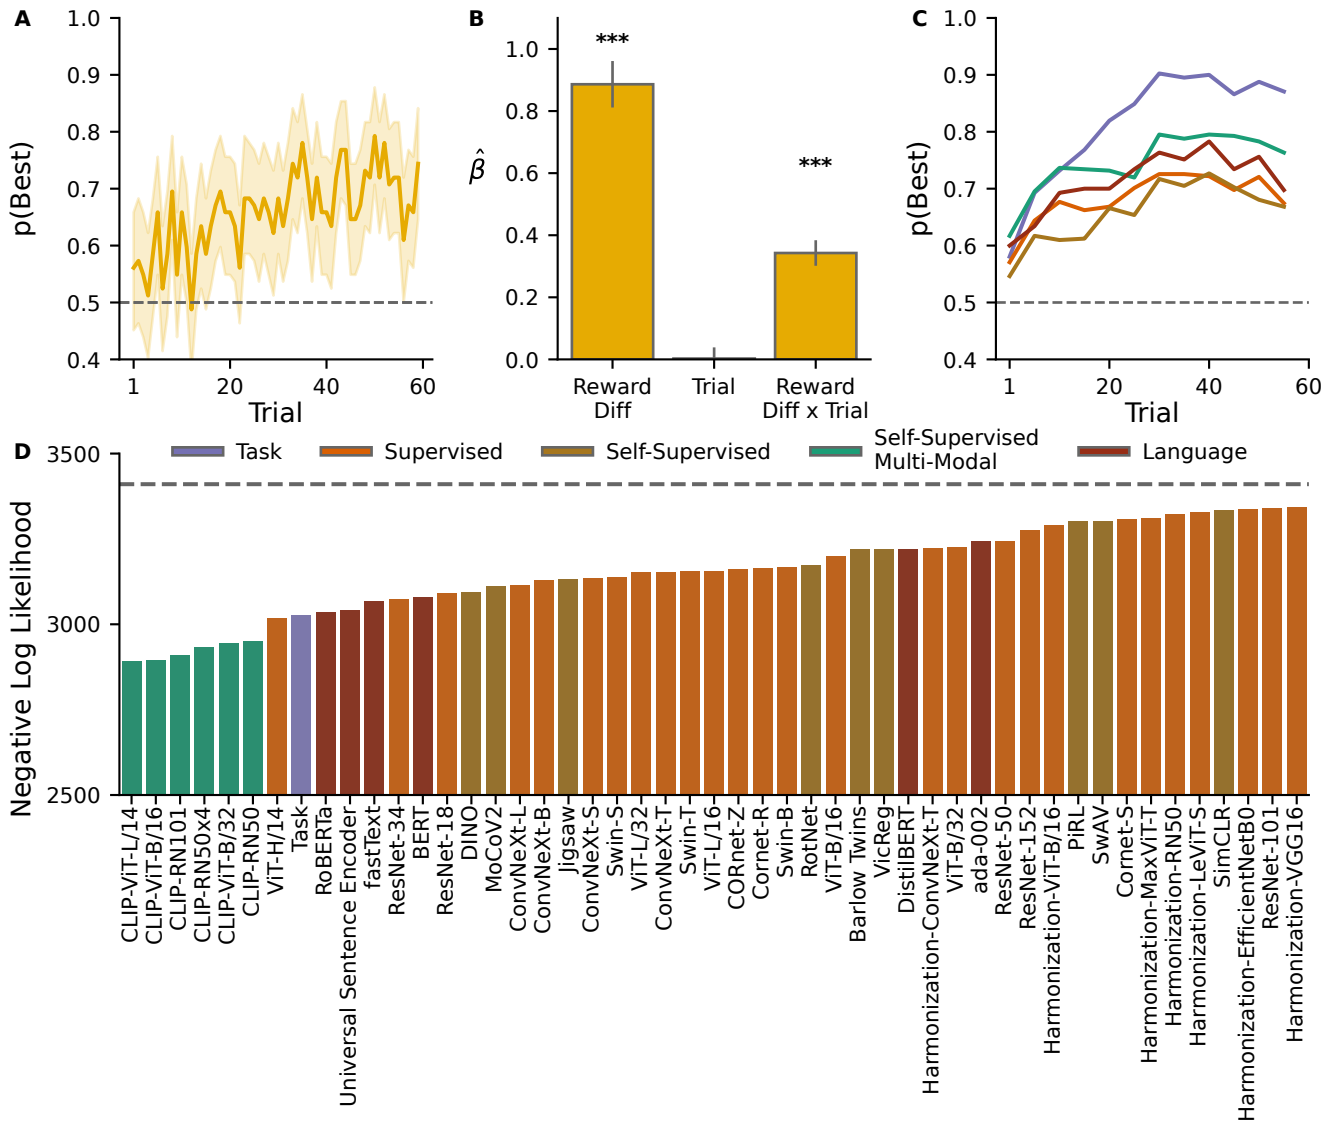


Figure 3. Reward learning results. (A) Participants’ performance across trials. Shaded lines indicate 95% confidence intervals. (B) Coefficient values from the mixed-effects logistic regression analysis to predict participant choice. Error bars indicate standard errors. (C) Performance of the computational models. For each category, we show the learning trajectory of the model predicting participant behavior the best. The curves are smoothed every 5th trial. (D) Cross-validated negative log likelihoods of each representation. Lower values indicate a better fit for human behavior. The dashed horizontal line indicates chance level performance.

tions are more similar to the embedding that was used to generate the task than any other representation we tested ($CKA = 0.61 \pm .008$) (see Fig. 4A for the full comparison). This similarity shows they provide a good representational basis to arrive at meaningful solutions in our tasks.

Does the high similarity between the CLIP representations and the task embedding mean that the CLIP is simply a good enough approximation of the task embedding, or are there also meaningful differences between the two? To answer this question, we calculated the CKA between CLIP representations and every other representation, as well as the CKA between the task embedding and every representation except CLIP. Then, we subtracted CKA values for CLIP from those for the task embedding. Here, while a positive value would indicate a representation being more similar to the task embedding than it is to CLIP representations,

a negative value would indicate an opposite relationship. We found a clear pattern, where language representations were consistently more similar to the task embedding than they were to CLIP representations. On the contrary, visual representations were more similar to CLIP representations than they were to the task embedding (Fig. 4B). These findings show that in addition to the high similarity between the task embedding and the CLIP representations, there is a meaningful difference that CLIP representations are better aligned with visual representations. This difference plays an important role in predicting participant behavior. To provide better intuition for the differences between the two models and the better predictive accuracy of CLIP models, we show example trials from the last 30 trials of the reward learning task in Fig. 5. In these trials, participants and models trained on CLIP representations made the same decision that differed from the decision of models




Figure 4. Representational similarity analyses (RSA). (A) Linear Centered Kernel Alignment (CKA) similarity between the task embedding and every representation tested. (B) Difference between the CKA of the task embedding and tested representations to the CKA of the CLIP representations and the other tested representations. Individual points indicate values for different CLIP representations. Error bars indicate 95% confidence intervals.

trained on the task embedding. These examples illustrate how CLIP representations can capture human intuition in decisions where the task embedding failed to do so. We ruled out other alternative explanations to our findings such as the type of learning algorithm used and the size of the representational space in extra analyses as shown in Fig. S2, Fig. S3, and Fig. S4.

Discussion

We developed novel category learning and reward learning tasks to test people’s abilities to generalize in high-dimensional spaces. The tasks required participants to identify relevant stimulus dimensions from feedback and use this knowledge to make correct subsequent decisions. Previous work has shown that humans can exploit relational structure for guiding choice in low dimensional physical spaces (5). Here, we observed that humans can exploit much higher dimensional abstract relational structures to make decisions. Furthermore, participants did not require any repetition of the

stimuli to generalize effectively.

We believe that the basis of this behavior lies within rich and expressive sensory representations. Over such representations, simple function learning mechanisms can identify relevant stimulus features. We trained linear models over several DNN representations to test what kind of DNN representations contain this richness and can predict human behavior. All 48 representations we tested predicted human behavior above chance level. Previous work has shown that DNN representations can predict human similarity judgments (42), performance in psychophysical tasks (43), and visual selectivity tasks (44). This line of work has been generalized to the auditory domain (45), as well as to language (16). We extended this literature by showing that DNN representations can predict human decisions in naturalistic learning tasks. We find it surprising that these representations can reveal task-relevant, semantically meaningful stimulus features after a few observations. First, because most of these representations contain thousands of features and learning which ones are relevant over



Figure 5. Example trials showing the similarity between CLIP and human decisions that show disagreement with the task embedding. Each row shows three trials from a different condition. Orange highlighted text shows the option chosen by all CLIP models and the human participant, whereas grey text shows the decision made by the task embedding. As the tasks were generated using the task embedding, all the choices shown here made by CLIP and humans are sub-optimal. Shown examples are from the second half of the task, as to eliminate the learning process as a confound. The original images are replaced with copyright-free alternatives from the THINGSplus database (2).

a few observations is a challenging learning problem. Second, it was unexpected that DNNs contained information about the tasks' generative features since they were trained on objectives independent of our experimental tasks.

Another interesting finding was the predictive success of the multi-modal representations we tested. As our tasks only used images, visual representations alone should be sufficient to solve these tasks. However, models trained on text and image data combined consistently outperformed vision-only and text-only models. Furthermore, multi-modal representations predicted participant behavior even better than the generative features of the task. We believe the multi-modal representations predicted behavior better than the generative task features because of the way in which the generative features were derived. An unsupervised similarity judgment was used to generate the task features, which does not require as fine-grained consideration as our tasks did. Additionally, our RSA results showed that multi-modal representations were better aligned with visual representations than the generative task features. These results together indicate that grounding in visual information is necessary but insufficient. We conclude that aligning visual representations with language gives rise to rich representations. These representations can then be adapted to generalize in naturalistic learning tasks as humans can. The success of the multi-modal representations offers novel evidence for the importance of language in shaping cognition, which has previously been shown through other methods (46–49).

An additional important takeaway from our findings is that simple learning strategies can be very effective when modeling human learning in naturalistic cognitive tasks. It has been shown in previous work that linear learning strategies can successfully predict participant behavior in learning tasks that use simple stimuli (4, 50). We showed that the same strategies can generalize to higher dimensional settings as well, indicating that simple learning strategies can be effective in naturalistic settings too.

Our findings have implications both for cognitive psy-

chology and for computer science. By showing that people can do learning tasks with naturalistic stimuli and that we can model these processes, our findings create the opportunity to study exploration-exploitation (51), contextual learning (52), and learning functions of different structures (53) in more naturalistic settings. In computer science, the attempt to build models aligned with humans has been increasing (34, 44, 54, 55). Our tasks and modeling approach offer a new way to measure the human alignment of DNN representations and to use this as a metric while building human-aligned DNNs. Previous work in this domain has focused either strictly on psychophysical tasks (43, 56) or on similarity judgments between stimuli (42, 57). We suggest that representations should not only be aligned at these levels but also translate to higher-level cognition, as measured by our tasks. Alignment at this level can pave the way for artificial systems that can generalize across semantically rich tasks, making them more robust and powerful.

Our work on understanding human representations in cognitive psychology tasks can be extended in multiple ways. First, we have only tested two paradigms and focused on learning problems. Within the learning paradigms, the same approach can be used to test whether these representations can predict human behavior when task rules are determined through non-linear functions over the embedding. Moving beyond learning paradigms, our approach can be used to test and model other cognitive functions such as memory, attention, and inhibition control among others. Second, we have focused on learning from visual observations. Future work can provide participants with text descriptions instead of images and test whether multi-modal representations are still needed to predict behavior or whether text-based representations are sufficient if there are no images. Lastly, in our modeling work, we have considered models of different architectures, different modalities of training data, and different regimes. However, we have not considered the effect of different learning rules DNNs use on the representations. This is a factor that can be investigated in future work.

Previous work has investigated decision-making in naturalistic settings, such as purchasing decisions and ratings of goods. These studies have improved our understanding of various questions such as how to explore in the real world (58) and how to represent tabular data (59). We extend our understanding of naturalistic decision-making, by showing that people can quickly adapt the relational knowledge they have about natural objects. This allows them to learn very quickly in naturalistic learning tasks and generalize effectively. We also show that DNNs are powerful tools for modeling the structure of these representations. This could open up the door for a whole new cognitive psychology that uses naturalistic tasks and environments and thereby increase the validity of the cognitive sciences more generally.

Methods

Participants

For the category learning task, we recruited 98 participants (48 females, 50 males, mean age = 28.92y, SD = 7.32) on the Prolific platform. Participants with less than 50% accuracy were excluded from the analyses, leaving us with 91 participants. A base payment of £ 1.50 was made, and participants could earn an additional bonus of £ 6.00. The median completion time was 12 minutes and 38 seconds. The inclusion criteria included having a minimum approval rate of 97%, and a minimum number of 15 previous submissions on Prolific. Participation in the reward learning study was an exclusion criterion. For the reward learning task, 99 participants were recruited (49 females, 49 males, 1 other, mean age = 27.9 y, SD = 9.13). After applying the 50% accuracy criteria, we were left with 82 participants. A base payment of £ 2.00 was made, and an additional performance-dependent bonus of £4.00 was offered. The median completion time was 9 minutes and 26 seconds. The inclusion criteria included having a minimum approval rate of 95%, and a minimum number of 10 previous submissions on Prolific. All participants agreed to their anonymized data being used for research. The study was approved by the ethics committee of the medical faculty of the University of Tübingen (number 701/2020BO). Participants gave consent for their data to be anonymously analyzed by agreeing to a data protection sheet approved by the data protection officer of the MPG (Datenschutzbeauftragte der MPG, Max-Planck-Gesellschaft zur Förderung der Wissenschaften).

Tasks and Stimuli

Both tasks were run online in forced full-screen mode. Participants were shown written instructions and were asked to complete comprehension check questions before they could start the tasks. In both tasks, participants were given unlimited time to make decisions. In the category learning task, binary (correct versus

wrong) feedback was given for 2s. In the reward learning task, the associated reward with the stimuli was shown for 1.5s, and there was an inter-trial interval of 1s where participants were shown a blank screen. Throughout both tasks, the estimated total payment of participants was shown on the upper part of the screen. At the end of the tasks, participants were asked whether they think their data should be used for analysis. Across both tasks, all but one participant responded saying their data should be analyzed, whose data was anyway excluded due to poor performance. The category learning task was programmed using jsPsych (60), whereas the reward learning task was programmed in plain JavaScript.

For each participant, 120 stimuli were sampled independently from the THINGS database. Because the loadings of the features were not uniformly distributed, we made 5 equally sized bins of the loadings for the assigned feature and sampled object categories uniformly from these bins. From these object categories, the specific images were assigned randomly. For details on the used features and the embedding, see Hebart et al. (1).

Extracting Representations

All the visual and multi-modal models were given the task images as inputs. The representations were extracted from the penultimate layer if the models had a classification layer, and from the final layer otherwise. For the vision transformer (ViT) models, the [class] token representations were extracted. We used the THINGSVision (39) toolbox for the steps above.

To extract representations from language models, we provided them with the prompt A photo of a X where X was the category label of the task image. fastText was only provided with the category label instead. ada-002 and fastText provided representations as outputs, whereas from the other models, we extracted the [class] token representations. All the resulting representational matrices had different observations as rows and different features as columns.

Modeling

To model participants' learning trajectories, we trained linear models on the different representations separately. The models were conditioned on the observations \mathbf{X} made until trial $t - 1$, and outputted estimates for the novel observation \mathbf{x} on trial t . For both tasks, we report results from L2 regularised models because overall they provided the best fit for human data. See Fig. S2 for results from sparse linear models.

For the category learning task, we used the following linear model to estimate the probability that a given observation \mathbf{x} belonged to the category $C = 1$.

$$p(C = 1) = h(\mathbf{x}) = \frac{1}{1 + e^{-\beta^T \mathbf{x}}} \quad (1)$$

where β was estimated by minimising the loss:

$$\mathcal{L} = -\sum_i^{t-1} [y_i \log h(\mathbf{x}_i) + (1 - y_i) \log(1 - h(\mathbf{x}_i))] + \alpha \|\beta\|^2 \quad (2)$$

The penalty term α was determined with grid search to maximize the task performance of each model on a participant basis. Up until observations from both categories were made, the models were conditioned on a single uninformative pseudo-observation for each category. We regressed the probability estimates of the linear learning model onto participant choice using a mixed-effects logistic regression model in order to estimate participants' policies. The probability estimates were the only fixed and random predictors used. The negative log likelihoods were obtained through leave-one-trial-out cross-validation, where each choice in the task served once as the test set.

For the reward learning task, we modeled data using a probabilistic linear model with centered spherical Gaussian priors over model weights, scaled by λ . The reward estimate \hat{r} was computed as follows:

$$\hat{r}(\mathbf{x}) = \left(\sigma^{-2} \left(\sigma^{-2} \mathbf{X}^T \mathbf{X} + \lambda \mathbf{I} \right)^{-1} \mathbf{X}^T \mathbf{r} \right)^T \mathbf{x} \quad (3)$$

where λ and observation noise σ were fitted to maximize the log marginal likelihood of the task performance. We then regressed the reward estimate differences between the left and the right options onto participant choice. This was the only fixed and random predictor in the model. For all learning models and mixed-effects models, we centered the training data and divided it by its standard deviation, and we applied the same scaling parameters to the test data. The learning models were constructed using `scikit-learn` (61), and the mixed-effects models were fitted using `lme4` (62).

RSA

We calculated pairwise similarities between different representations. First, the representations were mean-centered. Then, the linear CKA between two representations \mathbf{A} and \mathbf{B} was calculated as follows:

$$CKA(\mathbf{A}, \mathbf{B}) = \frac{\|\mathbf{B}^T \mathbf{A}\|_F^2}{\|\mathbf{A}^T \mathbf{A}\|_F \|\mathbf{B}^T \mathbf{B}\|_F} \quad (4)$$

where $\|\cdot\|_F$ denotes the Frobenius norm.

Data, Materials, and Software Availability

The code for the current study is available through the GitHub repository <https://github.com/candemircan/NaturalCogSci>. The data are available on the OSF from the following link: <https://osf.io/h3t52/>.

Acknowledgements

This work was funded by the Max Planck Society, the Volkswagen Foundation, as well as the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy-EXC2064/1-390727645.

Bibliography

- Martin N. Hebart, Charles Y. Zheng, Francisco Pereira, and Chris I. Baker. Revealing the multidimensional mental representations of natural objects underlying human similarity judgements. *Nature Human Behaviour*, 4(11):1173–1185, November 2020. ISSN 2397-3374. doi: 10.1038/s41562-020-00951-3. URL <https://www.nature.com/articles/s41562-020-00951-3>. Number: 11 Publisher: Nature Publishing Group.
- Laura Mai Stoinski, Jonas Perkuhn, and Martin N. Hebart. THINGSplus: New Norms and Metadata for the THINGS Database of 1,854 Object Concepts and 26,107 Natural Object Images, July 2022. URL <https://psyarxiv.com/exu9f/>.
- Roger N. Shepard, Carl I. Hovland, and Herbert M. Jenkins. Learning and memorization of classifications. *Psychological Monographs: General and Applied*, 75: 1–42, 1961. ISSN 0096-9753. doi: 10.1037/h0093825. Place: US Publisher: American Psychological Association.
- Yael Niv, Reka Daniel, Andra Geana, Samuel J. Gershman, Yuan Chang Leong, Angela Radulescu, and Robert C. Wilson. Reinforcement Learning in Multidimensional Environments Relies on Attention Mechanisms. *Journal of Neuroscience*, 35(21): 8145–8157, May 2015. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.2978-14.2015. URL <https://www.jneurosci.org/content/35/21/8145>. Publisher: Society for Neuroscience Section: Articles.
- Mona M. Garvert, Tankred Saanum, Eric Schulz, Nicolas W. Schuck, and Christian F. Doeller. Hippocampal spatio-predictive cognitive maps adaptively guide reward generalization. *Nature Neuroscience*, 26(4):615–626, April 2023. ISSN 1546-1726. doi: 10.1038/s41593-023-01283-x. URL <https://www.nature.com/articles/s41593-023-01283-x>. Number: 4 Publisher: Nature Publishing Group.
- J. Douglas Carroll. Functional Learning: The Learning of Continuous Functional Mappings Relating Stimulus and Response Continua. *ETS Research Bulletin Series*, 1963(2):i-144, 1963. ISSN 2333-8504. doi: 10.1002/j.2333-8504.1963.tb00958.x. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/j.2333-8504.1963.tb00958.x>. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/j.2333-8504.1963.tb00958.x>.
- Edward L. DeLosh, Jerome R. Busemeyer, and Mark A. McDaniel. Extrapolation: The sine qua non for abstraction in function learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23:968–986, 1997. ISSN 1939-1285. doi: 10.1037/0278-7393.23.4.968. Place: US Publisher: American Psychological Association.
- Marcel Binz, Samuel J. Gershman, Eric Schulz, and Dominik Endres. Heuristics from bounded meta-learned inference. *Psychological Review*, 129(5):1042–1077, October 2022. ISSN 1939-1471. doi: 10.1037/rev0000330.
- Charley M Wu, Eric Schulz, Maarten Speekenbrink, Jonathan D Nelson, and Björn Meder. Generalization guides human exploration in vast decision spaces. *Nature human behaviour*, 2(12):915–924, 2018.
- Martin N. Hebart, Adam H. Dickter, Alexis Kidder, Wan Y. Kwok, Anna Coriveau, 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, October 2019. ISSN 1932-6203. doi: 10.1371/journal.pone.0223792. URL <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0223792>. Publisher: Public Library of Science.
- Seyed-Mahdi Khaligh-Razavi and Nikolaus Kriegeskorte. Deep Supervised, but Not Unsupervised, Models May Explain IT Cortical Representation. *PLOS Computational Biology*, 10(11):e1003915, November 2014. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1003915. URL <https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1003915>. Publisher: Public Library of Science.
- Chengxu Zhuang, Siming Yan, Aran Nayebi, Martin Schrimpf, Michael C. Frank, James J. DiCarlo, and Daniel L. K. Yamins. Unsupervised neural network models of the ventral visual stream. *Proceedings of the National Academy of Sciences*, 118(3): e2014196118, January 2021. doi: 10.1073/pnas.2014196118. URL <https://www.pnas.org/doi/10.1073/pnas.2014196118>. Publisher: Proceedings of the National Academy of Sciences.
- Ruairidh M. Battleday, Joshua C. Peterson, and Thomas L. Griffiths. Capturing human categorization of natural images by combining deep networks and cognitive models. *Nature Communications*, 11(1):5418, October 2020. ISSN 2041-1723. doi: 10.1038/s41467-020-18946-z. URL <https://www.nature.com/articles/s41467-020-18946-z>. Number: 1 Publisher: Nature Publishing Group.
- Joshua C. Peterson, Joshua T. Abbott, and Thomas L. Griffiths. Evaluating (and Improving) the Correspondence Between Deep Neural Networks and Human Representations. *Cognitive Science*, 42(8):2648–2669, November 2018. ISSN 1551-6709. doi: 10.1111/cogs.12670.

15. Greta Tuckute, Jenelle Feather, Dana Boebinger, and Josh H. McDermott. Many but not all deep neural network audio models capture brain responses and exhibit hierarchical region correspondence, November 2022. URL <https://www.biorxiv.org/content/10.1101/2022.09.06.506680v3>. Pages: 2022.09.06.506680 Section: New Results.
16. Martin Schrimpf, Idan Asher Blank, Greta Tuckute, Carina Kauf, Eghbal A. Hosseini, Nancy Kanwisher, Joshua B. Tenenbaum, and Evelina Fedorenko. The neural architecture of language: Integrative modeling converges on predictive processing. *Proceedings of the National Academy of Sciences*, 118(45):e2105646118, November 2021. doi: 10.1073/pnas.2105646118. URL <https://www.pnas.org/doi/10.1073/pnas.2105646118>. Publisher: Proceedings of the National Academy of Sciences.
17. Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning Transferable Visual Models From Natural Language Supervision, February 2021. URL <http://arxiv.org/abs/2103.00020>. arXiv:2103.00020 [cs].
18. Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale, June 2021. URL <http://arxiv.org/abs/2010.11929>. arXiv:2010.11929 [cs].
19. Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. RoBERTa: A Robustly Optimized BERT Pretraining Approach, July 2019. URL <http://arxiv.org/abs/1907.11692>. arXiv:1907.11692 [cs].
20. Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, Yun-Hsuan Sung, Brian Strope, and Ray Kurzweil. Universal Sentence Encoder, April 2018. URL <http://arxiv.org/abs/1803.11175>. arXiv:1803.11175 [cs].
21. Tomas Mikolov, Edouard Grave, Piotr Bojanowski, Christian Puhrsch, and Armand Joulin. Advances in Pre-Training Distributed Word Representations, December 2017. URL <http://arxiv.org/abs/1712.09405>. arXiv:1712.09405 [cs].
22. Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition, December 2015. URL <http://arxiv.org/abs/1512.03385>. arXiv:1512.03385 [cs].
23. Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, May 2019. URL <http://arxiv.org/abs/1810.04805>. arXiv:1810.04805 [cs].
24. Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging Properties in Self-Supervised Vision Transformers, May 2021. URL <http://arxiv.org/abs/2104.14294>. arXiv:2104.14294 [cs].
25. Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved Baselines with Momentum Contrastive Learning, March 2020. URL <http://arxiv.org/abs/2003.04297>. arXiv:2003.04297 [cs].
26. Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A ConvNet for the 2020s, March 2022. URL <http://arxiv.org/abs/2201.03545>. arXiv:2201.03545 [cs].
27. Mehdi Norouzi and Paolo Favaro. Unsupervised Learning of Visual Representations by Solving Jigsaw Puzzles, August 2017. URL <http://arxiv.org/abs/1603.09246>. arXiv:1603.09246 [cs].
28. Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin Transformer: Hierarchical Vision Transformer using Shifted Windows, August 2021. URL <http://arxiv.org/abs/2103.14030>. arXiv:2103.14030 [cs].
29. Jonas Kubilius, Martin Schrimpf, Aran Nayebi, Daniel Bear, Daniel L. K. Yamins, and James J. DiCarlo. CORnet: Modeling the Neural Mechanisms of Core Object Recognition, September 2018. URL <https://www.biorxiv.org/content/10.1101/408385v1>. Pages: 408385 Section: New Results.
30. Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised Representation Learning by Predicting Image Rotations, March 2018. URL <http://arxiv.org/abs/1803.07728>. arXiv:1803.07728 [cs].
31. Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow Twins: Self-Supervised Learning via Redundancy Reduction, June 2021. URL <http://arxiv.org/abs/2103.03230>. arXiv:2103.03230 [cs, q-bio].
32. Adrien Bardes, Jean Ponce, and Yann LeCun. VICReg: Variance-Invariance-Covariance Regularization for Self-Supervised Learning, January 2022. URL <http://arxiv.org/abs/2105.04906>. arXiv:2105.04906 [cs].
33. Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter, February 2020. URL <http://arxiv.org/abs/1910.01108>. arXiv:1910.01108 [cs].
34. Thomas Fel, Ivan Felipe, Drew Linsley, and Thomas Serre. Harmonizing the object recognition strategies of deep neural networks with humans, November 2022. URL <http://arxiv.org/abs/2211.04533>. arXiv:2211.04533 [cs].
35. Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language Models are Few-Shot Learners, July 2020. URL <http://arxiv.org/abs/2005.14165>. arXiv:2005.14165 [cs].
36. Ishan Misra and Laurens van der Maaten. Self-Supervised Learning of Pretext-Invariant Representations, December 2019. URL <http://arxiv.org/abs/1912.01991>. arXiv:1912.01991 [cs].
37. Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised Learning of Visual Features by Contrasting Cluster Assignments, January 2021. URL <http://arxiv.org/abs/2006.09882>. arXiv:2006.09882 [cs].
38. Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A Simple Framework for Contrastive Learning of Visual Representations, June 2020. URL <http://arxiv.org/abs/2002.05709>. arXiv:2002.05709 [cs, stat].
39. Lukas Muttenthaler and Martin N. Hebart. THINGSvision: A Python Toolbox for Streamlining the Extraction of Activations From Deep Neural Networks. *Frontiers in Neuroinformatics*, 15, 2021. ISSN 1662-5196. URL <https://www.frontiersin.org/articles/10.3389/fninf.2021.679838>.
40. Nikolaus Kriegeskorte, Marieke Mur, and Peter Bandettini. Representational similarity analysis - connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, 2, 2008. ISSN 1662-5137. URL <https://www.frontiersin.org/articles/10.3389/neuro.06.004.2008>.
41. Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of Neural Network Representations Revisited, July 2019. URL <http://arxiv.org/abs/1905.00414>. arXiv:1905.00414 [cs, q-bio, stat].
42. Lukas Muttenthaler, Jonas Dippel, Lorenz Linhardt, Robert A. Vandermeulen, and Simon Kornblith. Human alignment of neural network representations, April 2023. URL <http://arxiv.org/abs/2211.01201>. arXiv:2211.01201 [cs, q-bio].
43. Robert Geirhos, Kantharaju Narayanappa, Benjamin Mitzkus, Tizian Thieringer, Matthias Bethge, Felix A. Wichmann, and Wieland Brendel. Partial success in closing the gap between human and machine vision, October 2021. URL <http://arxiv.org/abs/2106.07411>. arXiv:2106.07411 [cs, q-bio].
44. Thomas Langlois, Haicheng Zhao, Erin Grant, Ishita Dasgupta, Tom Griffiths, and Nori Jacoby. Passive attention in artificial neural networks predicts human visual selectivity. In *Advances in Neural Information Processing Systems*, volume 34, pages 27094–27106. Curran Associates, Inc., 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/e360367584297ee8d2d5afa709cd440e-Abstract.html>.
45. Alexander J. E. Kell, Daniel L. K. Yamins, Erica N. Shook, Sam V. Norman-Haignere, and Josh H. McDermott. A Task-Optimized Neural Network Replicates Human Auditory Behavior, Predicts Brain Responses, and Reveals a Cortical Processing Hierarchy. *Neuron*, 98(3):630–644.e16, May 2018. ISSN 0896-6273. doi: 10.1016/j.neuron.2018.03.044. URL [https://www.cell.com/neuron/abstract/S0896-6273\(18\)30250-2](https://www.cell.com/neuron/abstract/S0896-6273(18)30250-2). Publisher: Elsevier.
46. Lera Boroditsky. How Language Shapes Thought. *Scientific American*, 304(2):62–65, 2011. ISSN 0036-8733. URL <https://www.jstor.org/stable/26002395>. Publisher: Scientific American, a division of Nature America, Inc.
47. Asifa Majid, Melissa Bowerman, Sotaro Kita, Daniel B. M. Haun, and Stephen C. Levinson. Can language restructure cognition? The case for space. *Trends in Cognitive Sciences*, 8(3):108–114, March 2004. ISSN 1364-6613. doi: 10.1016/j.tics.2004.01.003. URL <https://www.sciencedirect.com/science/article/pii/S1364661304000208>.
48. Felice van 't Wout and Christopher Jarrold. The role of language in novel task learning. *Cognition*, 194:104036, January 2020. ISSN 0010-0277. doi: 10.1016/j.cognition.2019.104036. URL <https://www.sciencedirect.com/science/article/pii/S0010027719302094>.
49. Benedetto De Martino, Dharshan Kumaran, Ben Seymour, and Raymond J. Dolan. Frames, Biases, and Rational Decision-Making in the Human Brain. *Science (New York, N.Y.)*, 313(5787):684–687, August 2006. ISSN 0036-8075. doi: 10.1126/science.1128356. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2631940/>.
50. Maarten Speekenbrink and David R. Shanks. Learning in a changing environment. *Journal of Experimental Psychology: General*, 139:266–298, 2010. ISSN 1939-2222. doi: 10.1037/a0018620. Place: US Publisher: American Psychological Association.
51. Franziska Brändle, Marcel Binz, and Eric Schulz. Exploration Beyond Bandits. In Irene Cogliati Dezza, Eric Schulz, and Charley M. Wu, editors, *The Drive for Knowledge: The Science of Human Information Seeking*, pages 147–168. Cambridge University Press, Cambridge, 2022. ISBN 978-1-316-51590-7. doi: 10.1017/9781009026949.008. URL <https://www.cambridge.org/core/books/drive-for-knowledge/exploration-beyond-bandits/F5A30142D51738132E582FDA06C3CD7D>.
52. Eric Schulz, Emmanouil Konstantinidis, and Maarten Speekenbrink. Putting bandits into context: How function learning supports decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 44(6):927–943, June 2018. ISSN 1939-1285. doi: 10.1037/xlm0000463.
53. Eric Schulz, Josh Tenenbaum, David K Duvenaud, Maarten Speekenbrink, and Samuel J Gershman. Probing the Compositionality of Intuitive Functions. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. URL https://papers.nips.cc/paper_files/paper/2016/hash/49ad23d1ec9fa4bd8d77d02681df5cfa-Abstract.html.
54. Philippe Esling, Axel Chemla-Romeu-Santos, and Adrien Bitton. Generative timbre spaces: regularizing variational auto-encoders with perceptual metrics, October 2018. URL <http://arxiv.org/abs/1805.08501>. arXiv:1805.08501 [cs, eess].
55. Pranay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised Contrastive Learning. In *Advances in Neural Information Processing Systems*, volume 33, pages 18661–18673. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/d89a66c7c80a299b1bdbab0f2a1a94af8-Abstract.html>.
56. Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric, April 2018. URL <http://arxiv.org/abs/1801.03924>. arXiv:1801.03924 [cs].
57. Brett D. Roads and Bradley C. Love. Enriching ImageNet with Human Similarity Judgments and Psychological Embeddings, November 2020. URL <http://arxiv.org/abs/2011.11015>. arXiv:2011.11015 [cs].
58. Eric Schulz, Rahul Bhui, Bradley C. Love, Bastien Brier, Michael T. Todd, and

- Samuel J. Gershman. Structured, uncertainty-driven exploration in real-world consumer choice. *Proceedings of the National Academy of Sciences*, 116(28): 13903–13908, July 2019. doi: 10.1073/pnas.1821028116. URL <https://www.pnas.org/doi/10.1073/pnas.1821028116>. Publisher: Proceedings of the National Academy of Sciences.
59. Sudeep Bhatia and Neil Stewart. Naturalistic multiattribute choice. *Cognition*, 179: 71–88, October 2018. ISSN 0010-0277. doi: 10.1016/j.cognition.2018.05.025. URL <https://www.sciencedirect.com/science/article/pii/S0010027718301513>.
60. Joshua R. de Leeuw. jsPsych: A JavaScript library for creating behavioral experiments in a Web browser. *Behavior Research Methods*, 47(1):1–12, March 2015. ISSN 1554-3528. doi: 10.3758/s13428-014-0458-y. URL <https://doi.org/10.3758/s13428-014-0458-y>.
61. Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12(85):2825–2830, 2011. ISSN 1533-7928. URL <http://jmlr.org/papers/v12/pedregosa11a.html>.
62. Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, 67:1–48, October 2015. ISSN 1548-7660. doi: 10.18637/jss.v067.i01. URL <https://doi.org/10.18637/jss.v067.i01>.
63. David Wipf and Srikantan Nagarajan. A New View of Automatic Relevance Determination. In *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc., 2007. URL https://papers.nips.cc/paper_files/paper/2007/hash/9c01802ddb981e6bcfbec0f0516b8e35-Abstract.html.
64. Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention Is All You Need, December 2017. URL <http://arxiv.org/abs/1706.03762>. arXiv:1706.03762 [cs].
65. Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, June 2009. doi: 10.1109/CVPR.2009.5206848. ISSN: 1063-6919.
66. Drew Linsley, Dan Shiebler, Sven Eberhardt, and Thomas Serre. Learning what and where to attend, June 2019. URL <http://arxiv.org/abs/1805.08819>. arXiv:1805.08819 [cs].
67. Leo Gao. On the Sizes of OpenAI API Models, May 2021. URL <https://blog.eleuther.ai/gpt3-model-sizes/>.

Supplementary Information

Behavioral Analyses

Testing How Fast Participants Learn

For each trial in each task, we conducted a one-sided 1 sample t-test of choice accuracy against chance level performance. The results from these analyses are displayed in Fig. S1.

Alternative Learning Models

For each task, we tested two extra sets of learning models. Overall, we found that the models reported in the main text describe participant behavior the best. See Table 1 for a comparison of all learning models for each representation.

Sparse Models

In the main text, we reported results from L2 regularised learning models. In addition to that, we tested sparse linear models. For the category learning task, this was simply done by changing the penalty term to consider the L1 norm in the loss function:

$$\mathcal{L} = - \sum_i^{t-1} [y_i \log h(\mathbf{x}_i) + (1 - y_i) \log(1 - h(\mathbf{x}_i))] + \alpha \|\beta\|_1 \quad (\text{S1})$$

For the reward learning task, the prior over the weights was set differently from the model described in the main text. Here, instead of using a spherical Gaussian, each weight's prior had a different standard deviation. The standard deviations over the weights' priors and observation noise were estimated using evidence maximization (63). The results from sparse models are shown in Fig S2.

PCA Models

Different representations we tested had different numbers of features. To mitigate any negative effect of the number of features on model performance, we trained our linear models on Principal Component Analysis (PCA) transformed representations. For each representation and each participant, we took the first 49 principal components over all the images shown to the participant. We selected a 49 dimensional space because it matches the size of the embedding used to generate the task. Then, we trained our linear models on these representations as described in the main text. The results from these analyses are shown in Fig S3. Overall, these models do not perform better than those trained on the original representations.

The Effect of Number of Features

In addition to the PCA analysis described above, we tested whether the number of features had an effect on how well they described participant behavior through a correlation analysis. We found no significant correlation between the number of features and the sum of the negative log likelihoods across the two tasks ($\tau_b = .19, p = .07$).

Model Descriptions

Below, we briefly describe all the neural networks from which we extracted representations. For more detailed descriptions, see the original associated work.

Task Embedding (1)

This is a representational embedding that describes the images in the THINGS database (10). The authors of the model first presented human participants with an unsupervised task where they were asked to choose the odd image in a triplet. Then, they built a model to predict human choice in this task. The objects are provided as one-hot vectors, and the model learns to project them into a latent space to predict participant choice. The authors found that the objects could be described in a 49 dimensional latent space, and they verified that these were semantically meaningful dimensions through further experiments.

ViT (18)

Vision transformer (ViT) is a model that applies the Transformer architecture used in language models such as BERT (23) to image recognition (64). Images are initially divided into multiple 2D patches. These patches and their positions are linearly embedded in a latent space. Additionally, a [class] embedding is learned to represent the entire image. The sequence of these embeddings is passed through a Transformer encoder. The model learns a linear mapping from the output of the Transformer encoder for the [class] embedding to the object categories. The models we used were trained on ImageNet 1k (65). We tested three 3 variants of ViT that differed in size. This architecture is also used in some CLIP and Harmonization models.

ResNet (22)

ResNet is a type of convolutional neural network. In addition to a standard convolutional architecture, it introduces residual connections. These extra connections propagate the forward signal by skipping convolutional layers in between, allowing to train networks deeper than what was possible before. The models we used were trained on ImageNet 1k. We tested 5 different variants of ResNet that differed in size. This architecture is also used in some CLIP, and Harmonization models, as well as all the self-supervised models we tested.

Swin (28)

Swin is a type of Transformer used for vision tasks. It introduces hierarchical feature maps that increase in resolution through the network's layers. Additionally, the attention windows are shifted from one transformer block to the next, in order to provide connections between the separate windows of the previous layer. The models we used were trained on ImageNet 1k. We tested 3 different variants of Swin that differed in size.

ConvNeXt (26)

ConvNeXt is a recent convolutional neural network variant. It introduces a better FLOPs/accuracy trade-off compared to a ResNet by using depthwise convolution. The model uses an inverted bottleneck, where block sizes initially get larger and then smaller, and applies convolution over larger distances. In addition to these major changes, ReLU activation functions are replaced with GERU, and batch normalization is replaced by layer normalization. The models we used were trained on ImageNet 1k. We tested 4 different variants of ConvNeXt that differed in size. This architecture is also used by one of the Harmonization models.

CORnet (29)

CORnet is a deep neural network model of the primate ventral visual stream. It explicitly models the areas V1, V2, V4, and IT by assuming identical circuitry of varying sizes. CORnet-Z uses single convolution, followed by a ReLU activation and max pooling. CORnet-R additionally introduces recurrent connections in a biologically plausible manner. CORnet-S builds on top of CORnet-R by adding residual feedforward connections as seen in ResNet. These models were trained on ImageNet-1k.

Harmonization (34)

This is a family of vision models that aligns deep neural networks with human vision. In addition to the standard supervised training, the models are trained to use the same visual features of images that humans do. Roughly, the second part is achieved by aligning a function of the networks' saliency maps with feature importance maps obtained from human judgment. This results in networks that perform better in ImageNet and that are aligned with humans. These models were trained on ImageNet-1k, as well as the *ClickMe* dataset (66). We tested 6 different architectures trained under this regime, as provided by the authors of the original work. These networks included various convolutional neural networks and transformers.

DINO (24)

DINO is a self-supervised vision model. The training setup includes a Student network and a Teacher network. Both networks are provided by augmented images from ImageNet-1k. While the Teacher is provided with Global views, the Student is only provided with Local views. The Student learns to use the local views to predict the global features, which are extracted from the Teacher network. The Teacher is a network whose weights are an exponentially weighted average of the Student network. We used the ResNet-50 architecture for this model.

SwAV (37)

SwAV is a type of self-supervised learning that combines clustering and contrastive learning in an efficient way via online learning. Two augmented versions of an image are passed through an encoder network. Using prototype vectors, these latent representations are mapped onto "codes". The network is trained to predict the code of one

augmented image from the other augmented image's latent representation. The model we used had a ResNet-50 architecture and was trained on ImageNet-1k.

PIRL (36)

PIRL is a self-supervised learning recipe that promotes transformation invariant representation learning of images. In this approach, a neural network is trained to maximize the similarity of the latent representation of an image and the latent representations of the augmented versions of that image. Additionally, the model keeps a memory bank and learns to make the latent representation of a given image less similar to those of other images in the memory bank, while keeping it similar to the previous latent representation found in the memory bank. The model we used had a ResNet-50 architecture and was trained on ImageNet-1k.

SimCLR (38)

In the SimCLR framework, two different sets of transformations are applied to a batch of images. These images are passed through an encoder network to form latent representations. The network learns to maximize the similarity of the augmented images' latent representations that come from the same original image while minimizing the similarity to those that originate from different images. The model we used had a ResNet-50 architecture and was trained on ImageNet-1k.

VicReg (32)

VicReg is another self-supervised learning method used for vision. Here, the encoder network is given two different views of an image. The goal is to minimize the distance between the embedding vectors while keeping the standard deviation of each variable above a threshold and pushing the covariance between the embedding variables toward zero. The model we used had a ResNet-50 architecture and was trained on ImageNet-1k.

Barlow Twins (31)

Barlow Twins passes two distorted versions of the same images through an encoder network to obtain latent representations. In the latent space, it learns to make the cross-correlation matrix between the two groups of images as close to an identity matrix as possible. This approach makes the representations of noisy variants of the same image similar, while at the same time minimizing the redundancy between the two representations. The model we used had a ResNet-50 architecture and was trained on ImageNet-1k.

MoCo v2 (25)

MoCo v2 is another setup for contrastive learning, which treats the encoder training as a dictionary look-up task. Augmented images are passed through an encoder network to obtain queries. In parallel, a momentum encoder network is used to represent a dictionary of images, where the momentum encoder is a momentum-based updated version of the encoder network. Here, the dictionary is maintained as a queue, where the representations of the current mini-batch are enqueued, and the oldest ones are dequeued. The goal is to maximize the similarity of the query representation to the key that has the lowest distance from the query while maximizing the dissimilarity to the other keys. In this framework, the query and the current key that is most similar to the query are treated as different augmentations of the same image. The model we used had a ResNet-50 architecture and was trained on ImageNet-1k.

RotNet (30)

In the RotNet framework, rotated versions of images are passed through an encoder network. The rotation is selected from a discrete set of rotations. Given an image, the goal of RotNet is to classify the degree of rotation that has been applied to the image correctly. The model we used had a ResNet-50 architecture and was trained on ImageNet-1k.

Jigsaw (27)

An early self-supervised learning approach is to teach neural networks to solve jigsaw puzzles. Here, images are divided into a number of 2D tiles and are shuffled. Then, an encoder network learns tile-specific feature maps. Using these local features, it rearranges the tiles in the correct order. The model we used had a ResNet-50 architecture and was trained on ImageNet-1k.

CLIP (17)

CLIP is a multi-modal self-supervised learning framework. In this approach, two encoders are trained simultaneously, one for images and the other one for text. The model is trained on text image pairs. The goal of the model is to

predict the correct pairs of images and text, by learning a similar latent representation for the pairs using the two encoders. At the same time, it learns to make the latent representations of incorrect pairs dissimilar. For the text encoder, a Transformer architecture is used. For the image encoder, we tested both ResNet and ViT variants. CLIP models we used were those provided by the original authors. The training data of the models are not fully specified, however, it is reported in the original paper that the model is trained on 400 million image-text pairs.

Universal Sentence Encoder (20)

Universal Sentence Encoder is a language model that outputs vector representations for any text. The model uses a deep averaging network, where the input word embeddings and bi-grams are initially averaged and later passed through a feed-forward deep neural network. The model is trained on a range of natural language tasks. The training data is not fully disclosed. However, it is reported to include text from Wikipedia, web news, web question-answer pages and discussion forums.

BERT (23)

BERT is a pre-trained language model that uses a Transformer architecture. The model is given text as input, where some part of the text is masked, and the model learns to predict the masked text, given the rest of the input, which acts as context. An additional `class` token is added to the beginning of each input. The representations learned for this token are used for downstream tasks such as classification. The model we used was trained on BookCorpus, which is a dataset that consists of more than 10000 unpublished books, as well as English articles in Wikipedia.

DistilBERT (33)

DistilBERT is a distilled version of the BERT model, whose size is 40% and performance is 97% of the original model. The model is trained through distillation, which uses a teacher-student framework. Here, the teacher is the original BERT model. DistilBERT is trained with the objective to behave as similarly as possible to BERT, to predict masked parts of the input given the rest of the context, and to generate hidden states that were as similar as possible to that of the BERT model. The model we used was trained on the same data as BERT, described above.

RoBERTa (19)

RoBERTa is a replication of BERT, which documents the impacts of hyperparameter choice and training data size. It outperforms BERT in multiple language tasks. The model we used was trained on a superset of the BERT's training data described above. Additional sources included CC-News, OpenWebText, and Stories.

ada-002 (35)

ada-002 is a text embedding model deployed by OpenAI. It is given text as input and generates a vector representation of the text. Unfortunately, details for this model are not available. However, it is believed to be one of the encoding models that is available for GPT-3 (35).

fastText (21)

fastText is another model that uses word and sub-word level information to turn text into vector representations. We used a skipgram version of this model, which was trained on 600 billion tokens extracted from Common Crawl. In the skipgram framework, a word in a given sentence is masked. The model is trained to learn representations such that when provided with the surrounding contextual words separately as input, it can predict the masked word.

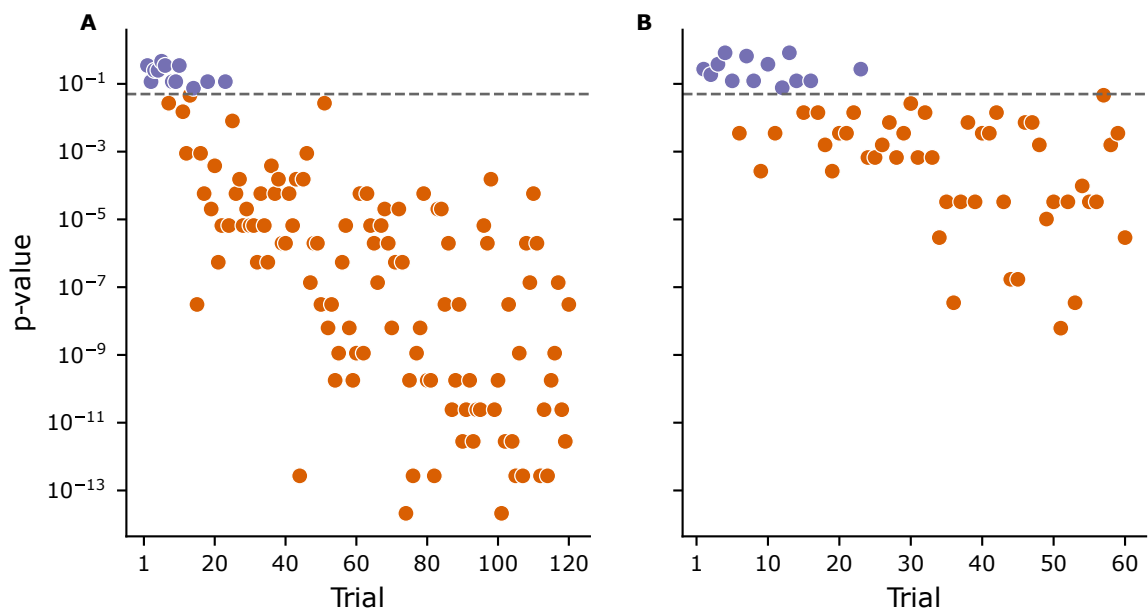


Figure S1. Participant Performance Against Chance Level at Each Trial. Trial-by-trial p-values from 1 sample t-tests testing accuracy against chance level for (A) category learning task and the (B) reward learning task.




Figure S2. Performance of the Sparse Learning Models. Cross-validated negative log-likelihoods for the (A) category learning task and the (B) reward learning task are shown. Lower values indicate better fits to human behavior. The dashed horizontal line indicates chance level performance.




Figure S3. Performance of the Learning Models trained on PCA representations. Cross-validated negative log-likelihoods for the (A) category learning task and the (B) reward learning task are shown. Lower values indicate better fits to human behavior. The dashed horizontal line indicates chance level performance.




Figure S4. Relationship Between Model Fit and the Size of the Representational Space. We observe that the number of features does not have an impact on how well different representations can explain participant behavior.

Model Name	Type	Parameters (Million)	Category Learning NLL			Reward Learning NLL		
			Main	Sparse	PCA	Main	Sparse	PCA
CLIP-ViT-L/14	Self-Supervised Multi-Modal	427.62	6556.57	7048.70	7350.88	2893.09	3293.56	3235.11
CLIP-ViT-B/16	Self-Supervised Multi-Modal	149.62	6629.96	6996.63	7341.26	2896.44	3215.89	3229.61
CLIP-ViT-B/32	Self-Supervised Multi-Modal	151.28	6660.77	7047.67	7363.74	2946.11	3197.59	3221.68
CLIP-RN50x4	Self-Supervised Multi-Modal	178.3	6746.92	6930.88	7332.79	2932.43	3074.20	3053.85
CLIP-RN101	Self-Supervised Multi-Modal	119.69	6802.21	7000.14	7307.13	2909.52	3047.88	3048.06
ViT-H/14	Supervised	633.47	6746.31	7200.56	7256.95	3019.41	3211.69	3177.17
Task	Task	0.17*	6976.13	6987.36	6876.38	3025.50	2970.23	2968.43
Universal Sentence Encoder	Language	8.54	7080.84	7161.20	7404.60	3040.20	3126.09	3143.28
RoBERTa	Language	124.65	7098.94	7197.73	7344.64	3035.65	3171.26	3178.74
DINO	Self-Supervised	23.51	7106.35	7262.30	7447.87	3094.62	3233.45	3165.34
fastText	Language	NA	7146.80	7299.53	7323.10	3066.87	3236.48	3180.89
ConvNext-L	Supervised	197.77	7120.20	7268.95	7387.76	3114.50	3203.61	3205.56
BERT	Language	109.48	7184.10	7261.03	7422.32	3080.85	3169.49	3178.61
ResNet-18	Supervised	11.69	7185.59	7304.01	7370.80	3091.07	3194.43	3156.70
ResNet-34	Supervised	21.8	7205.51	7355.84	7378.90	3073.46	3225.14	3161.01
ResNet-50	Supervised	25.56	7058.73	7249.82	7410.53	3244.52	3312.13	3180.09
MoCoV2	Self-Supervised	23.51	7198.66	7262.55	7470.35	3112.54	3198.73	3160.91
ConvNext-B	Supervised	88.59	7185.77	7282.16	7434.23	3128.76	3174.86	3167.18
Swin-T	Supervised	28.29	7180.41	7417.50	7432.17	3156.21	3271.34	3254.49
Cornet-R	Supervised	5.21	7179.79	7335.57	7370.35	3165.64	3219.80	3210.31
ConvNext-S	Supervised	50.22	7214.27	7332.72	7464.86	3134.08	3215.19	3197.22
ada-002	Language	350**	7106.84	7264.09	7433.11	3243.41	3407.04	3253.14
Swin-S	Supervised	49.61	7228.70	7423.84	7374.90	3136.99	3261.96	3270.30
VicReg	Self-Supervised	23.51	7154.49	7270.59	7453.93	3219.45	3229.88	3188.65
Swin-B	Supervised	87.77	7205.77	7343.31	7374.76	3168.50	3189.88	3261.50
Barlow Twins	Self-Supervised	23.51	7156.34	7267.88	7455.53	3219.42	3239.21	3196.55
ViT-L/32	Supervised	306.54	7259.64	7404.67	7432.01	3151.54	3211.50	3271.45
PIRL	Self-Supervised	23.51	7133.95	7216.29	7482.89	3303.51	3277.18	3179.97
SwAV	Self-Supervised	23.51	7133.95	7221.40	7479.26	3303.51	3277.18	3177.67
DistilBERT	Language	66.36	7229.13	7288.39	7476.50	3220.91	3325.47	3308.40
ResNet-152	Supervised	60.19	7178.87	7316.28	7442.82	3276.84	3282.15	3222.57
ViT-B/16	Supervised	86.57	7258.40	7414.03	7455.65	3201.05	3274.24	3296.59
SimCLR	Self-Supervised	23.51	7135.90	7221.14	7483.27	3333.31	3268.50	3191.66
ResNet-101	Supervised	44.55	7136.56	7245.44	7405.02	3339.84	3292.09	3203.80
ConvNext-T	Supervised	28.59	7324.95	7362.84	7454.36	3151.89	3239.63	3261.15
Jigsaw	Self-Supervised	23.51	7360.46	7351.58	7511.62	3133.06	3209.14	3244.48
ViT-B/32	Supervised	88.22	7271.29	7466.39	7402.88	3225.13	3280.83	3285.52
RoFormer	Self-Supervised	23.51	7343.46	7430.61	7496.19	3173.14	3210.07	3249.09
ViT-L/16	Supervised	304.33	7405.88	7470.41	7439.83	3156.49	3213.36	3283.52
CORnet-Z	Supervised	2.07	7414.85	7432.33	7458.41	3160.30	3243.43	3231.07
Harmonization-ConvNext-T	Supervised	28.59	7440.67	7460.92	7482.18	3221.63	3291.88	3284.43
Harmonization-ViT-B/16	Supervised	86.57	7501.41	7498.24	7511.05	3291.06	3328.95	3342.78
Harmonization-MaxViT-T	Supervised	30.96	7515.66	7524.84	7515.16	3309.47	3339.27	3333.26
Cornet-S	Supervised	53.42	7520.58	7522.31	7512.14	3307.13	3332.08	3336.97
Harmonization-RN50	Supervised	25.61	7508.76	7501.03	7496.49	3323.09	3314.34	3323.73
Harmonization-VGG16	Supervised	138.36	7501.40	7512.61	7522.21	3342.20	3440.21	3292.75
Harmonization-LeViT-S	Supervised	8.88	7521.85	7520.23	7483.63	3327.08	3344.02	3337.67
Harmonization-EfficientNetB0	Supervised	5.33	7524.88	7523.05	7498.16	3338.08	3337.96	3321.88

Table 1. Details on Models and Model Performance. Here, we report the model size, the type, and the negative log-likelihoods (NLL) for different tasks and different learning models. For a given representation and task, the best-performing learning model's NLL is highlighted in bold. The models are ordered such that the ones reported higher have explained participant performance better across the two tasks, as measured through the performance of the linear models reported in the main text. * The reported number is the initialized number of weights. It is a sparse model, and the resultant non-zero weights are around 90846. ** OpenAI does not report the model size. The reported number is an estimated number for GPT-3 that uses the ada-002 encoder (67).