

Enriching ImageNet with Human Similarity Judgments and Psychological Embeddings

Brett D. Roads & Bradley C. Love
Department of Experimental Psychology
University College London
London, United Kingdom
{b.roads,b.love}@ucl.ac.uk

Abstract

Advances in supervised learning approaches to object recognition flourished in part because of the availability of high-quality datasets and associated benchmarks. However, these benchmarks—such as ILSVRC—are relatively task-specific, focusing predominately on predicting class labels. We introduce a publicly-available dataset that embodies the task-general capabilities of human perception and reasoning. The Human Similarity Judgments extension to ImageNet (ImageNet-HSJ) is composed of a large set of human similarity judgments that supplements the existing ILSVRC validation set. The new dataset supports a range of task and performance metrics, including evaluation of unsupervised algorithms. We demonstrate two methods of assessment: using the similarity judgments directly and using a psychological embedding trained on the similarity judgments. This embedding space contains an order of magnitude more points (i.e., images) than previous efforts based on human judgments. We were able to scale to the full 50,000 image ILSVRC validation set through a selective sampling process that used variational Bayesian inference and model ensembles to sample aspects of the embedding space that were most uncertain. To demonstrate the utility of ImageNet-HSJ, we used the similarity ratings and the embedding space to evaluate how well several popular models conform to human similarity judgments. One finding is that the more complex models that perform better on task-specific benchmarks do not better conform to human semantic judgments. In addition to the human similarity judgments, pre-trained psychological embeddings and code for inferring variational embeddings are made publicly available. ImageNet-HSJ supports the appraisal of internal representations and the development of more human-like models.

1. Introduction

One interesting question is how models' internal representations compare to human-perceived similarities. While people make such judgments with little effort, human-perceived similarity flexibly adapts to different contexts, reflecting a rich understanding of the world [32, 23]. For example, people may perceive a beer bottle as similar to cigarettes because both are age-restricted, while perceiving a beer bottle as similar to a soda because both are beverages. Humans may perceive two objects as similar for many reasons, including the two objects playing related roles within encompassing systems, sharing perceptual properties, or simply interacting with one another [23]. Human-perceived similarity has been leveraged in applications such as image retrieval [11, 15] and human-in-the-loop categorization [42, 62, 65], but has been under-utilized in the general development of computer vision algorithms. The lack of research is partly due to the absence of an appropriate dataset and the technical challenges associated with collecting such a dataset. This work introduces the Human Similarity Judgments extension to ImageNet (ImageNet-HSJ), designed to include maximally informative similarity judgments for the widely-used ILSVRC dataset [10, 45].

As algorithm development shifts from learning *task-specific* representations towards *task-general* representations, the evaluation metrics used to assess models may also benefit from an equivalent shift. While task-specific metrics—such as classification accuracy—will always be relevant, complementary task-general evaluation metrics seem increasingly necessary. One strategy for creating a task-general metric is to assess how well model-perceived similarity aligns with human-perceived similarity. Comparing human and model similarity allows researchers to focus on the internal representations that precede task-specific output. Internal representation metrics create a level playing field when comparing across diverse training paradigms; such as supervised, unsupervised, and self-supervised ap-

proaches. Helping machines think in a more human-like way may also be better for human-machine interactions.

The main aims of this work are to assemble a dataset that embodies human-perceived similarity and demonstrate how the dataset can be used to assess arbitrary models. Extending previous work [56, 43], we employed *psychological embeddings* to concisely model the information contained in the similarity judgments. A psychological embedding includes an embedding of the stimuli, as well as functions that link the embedding to observed behavior. The psychological embeddings serve three different roles: as a means of modeling uncertainty and performing active learning (§ 4), assessing whether sufficient data has been collected (§ 5), and evaluating the internal representations of arbitrary models (§ 6). To handle the large number of stimuli, existing approaches are extended using variational inference and ensembles. The dataset and companion pre-trained psychological embeddings are hosted at <https://osf.io/cn2s3/>. An open-source python package for inferring variational psychological embeddings is available at github.com/roads/psiz.

2. Related work

There is a long history of using human similarity judgments to infer embeddings [57, 58]. The family of algorithms includes a number of kernel variants [13, 33, 48, 56, 60], as well as non-metric approaches [50, 51]. More recent work has seen an increase in the scale and naturalism of the stimuli being used [63, 61, 43, 44, 35]. Of particular note is recent work that collected similarity judgments for 1,854 unique images belonging to the THINGS dataset [17, 18]. In this work, we advance the size of the stimulus set by an order of magnitude, while also focusing on a dataset that is widely used in machine learning.

An enormous volume of work has conducted targeted comparisons between artificial neural networks and human behavior. Much of this work has focused on comparing human and model *classification* performance [12, 37]. Classification experiments have revealed differences in image-level confusion statistics [39], as well as the degree of shape bias [3]. Other work has examined the ability of artificial neural networks to predict human-perceived typicality ratings [28]. Most relevant is work that compared human and model-perceived *similarity*. For example, models pre-trained on ImageNet can predict both human similarity ratings [36] and similarity rankings [38, 2, 64]. Taking things further, deep neural networks can be specifically trained to predict human similarity judgments with relatively high fidelity [2, 36, 38, 46]. Similarly, substantial work has compared deep neural networks and brain activity [24, 30, 47]. Our aim is to complement existing research by providing a dataset that can serve many purposes, in addition to being an evaluation metric.

Property	Seed (v0.1)	Full (v0.2)
Unique stimuli	1,000	50,000
Judged trials	25,273	384,277
Unique participants	431	5572
Median trial duration (s)	7.85	8.30
Participant contribution		
Median	0.18%	0.01%
Max	3.48%	1.16%

Table 1. Basic properties of the ImageNet-HSJ dataset.

A number of active learning paradigms exist for efficiently collecting data [22, 41, 53]. This work builds on an approach specifically catered to collecting human similarity judgments [56, 43]. While Roads *et al.* [43] used MCMC to obtain posterior samples, we directly sample from an approximate posterior learned using variational inference. This approach reduces computational costs, helping us scale up to larger stimulus sets. In addition to computational savings, variational inference also eliminates a problem associated with drawing posterior samples from an embedding. Embeddings typically have some degree of invariance, such as being rotation invariant in Euclidean space. Consequently, sampling the posterior using MCMC can artificially inflate the uncertainty of peripheral points as the sampling chain rotates in space. While various techniques can be used to limit this issue [14], variational inference naturally addresses the problem.

3. ImageNet-HSJ dataset

The Human Similarity Judgments extension to ImageNet (ImageNet-HSJ) is a versioned dataset composed of ordinal human similarity judgments intended to supplement the ILSVRC 2012 validation dataset. This work focuses exclusively on the validation set because evaluation is a prerequisite for model selection paradigms. The dataset includes a large number of quality-controlled similarity judgments, collected from a wide range of participants (Table 1).

ImageNet is an ideal dataset for enrichment because it is widely used and captures a spectrum of everyday visual experience. In addition to exhibiting both conceptual breadth and depth, each class is made up of highly variable images. ImageNet also strikes a balance in size. While small relative to more recent datasets, it is still challenging enough to remain relevant. The remainder of this section describes the human task, stimuli, and the versioning strategy.

3.1. Human similarity judgment task

Human similarity judgments were collected using a web-based application. Participants were recruited via Amazon Mechanical Turk (AMT) and completed a *session* in exchange for monetary compensation (approximately 8.00

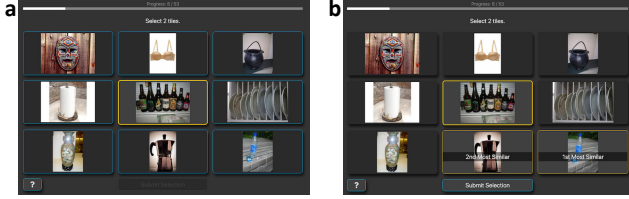


Figure 1. An example similarity judgment trial shown to human participants. (a) On initial presentation, nice images are shown with the query image placed in the center. (b) When participants click a tile, it is highlighted and marked as their first or second choice. After participants select two reference images, they are allowed to advance to the next trial. In this case, the participant thought the beer bottles were most similar to a plastic water bottle, and second most similar to a moka pot.

USD per hour). Each session was composed of 50 trials and designed to take 10 minutes.

On each trial, participants were presented with a display composed of nine images arranged in a grid (Figure 1). The center image is the *query* and the surrounding images are *references*. Participants were instructed to select the two reference images most similar to the query. Based on the order of their selection, participants also indicated which reference they thought was most similar and second most similar. Participants could unselect a reference by clicking it a second time. Once satisfied, participants clicked a button to submit the trial and the next trial was displayed.

In general, a trial can be composed R references and participants tasked with making $R - 1$ choices [61, 63, 43]. The 8-rank-2 trial format was chosen because most participants find it an easy task and it strikes a reasonable balance between throughput and noise [9, 29].

3.2. Stimulus set

Human similarity judgments were collected for all 50,000 images of the ILSVRC validation set. From the *full* dataset, we selected a subset of stimuli to serve as the *seed* subset. Data collection began by focusing exclusively on the seed subset. Once convergence was achieved (see § 5), collection began for the full dataset. The seed subset contained 1,000 images with a single representative image per class. The representative images were selected using a pre-trained VGG19 network [54]. For each class, the image with the highest probability of correct classification was chosen, with ties broken randomly. Images listed on the ILSVRC 2014 exclusion list were not eligible for the seed subset.

3.3. Versioning

ImageNet-HSJ follows a versioning scheme. Version 0.1 includes similarity judgments for the seed subset. Version 0.2 expands coverage to the full stimulus set and includes

all judgments collected at the time of submission. Once convergence is achieved for the full dataset (see § 5), a final version 1.0 will be released.

4. Dataset collection

The primary goal of the data collection process is to assemble a set of observations that informatively probes human-perceived similarity. Given a stimulus set of 50,000 images and an 8-rank-2 trial format, there are more than 10^{36} potential trials to show participants. Given the conceptual breadth of ImageNet, most trials are likely to be composed of images that are all highly dissimilar (e.g., firetruck, banana, squid, baseball, parrot, elephant, tiara, ladybug, and pretzel). On these trials participants are likely to make highly idiosyncratic or random choices. Instead, we would like to prioritize trials that reveal consistent perceptual beliefs.

In lieu of random sampling, we use an active learning paradigm to identify promising trials across multiple iterations. We extend existing active learning paradigms [53, 56, 43] using variational Bayesian inference. Each iteration of the active learning procedure is composed of three steps: trial selection, judgment collection and embedding inference (Figure 2). In the first step, an existing psychological embedding is used to select a new set of trials that maximizes expected information gain (§ 4.3). In the absence of data, the psychological embedding is governed by a prior distribution and trials are selected randomly. In the second step, human similarity judgments are collected (§ 4.1). In the third step, all available observations are used to infer a new ensemble of psychological embeddings (§ 4.2).



Figure 2. Outline of the active learning procedure. A set of trials is selected from a set of stimuli. Following human judgment, variational inference is used to obtain posterior distributions for an ensemble of psychological embeddings. The posterior distributions are used to select the next set of trials that maximize expected information gain. After each iteration, the increase in information will result in increasingly precise estimates of where stimuli reside in the embedding.

4.1. Judgment collection

Trials selected by the active learning procedure were randomly partitioned into 72 sessions composed of 50 trials each. During a 50-trial session, there were four *catch trials* designed to assess quality and earnest effort. Catch trials

were identical to randomly sampled trials, except that one of the reference images was a mirror-image of the query image. Catch trials were strategically placed such that two catch trials occurred in the first 20 trials and two in the last 20 trials.

For a given session, catch trial performance determined the sample weight of the trials. For each catch trial, a participant selected the mirrored query as the first choice, second choice, or not at all; resulting in a grade of 1, .5, and 0 respectively. The sample weight of each trial was defined as the mean grade of all the catch trials.

Depending on the session’s average grade, a session was marked *premium* ($> .875$), *satisfactory* ($[\cdot5, .875]$), or *unsatisfactory* ($< .5$). Unsatisfactory sessions were dropped from the final dataset. If a participant received a premium grade, they remained eligible to complete more sessions. If a participant received a non-premium grade, they were added to an ineligible list. A scheduler created HITs on AMT in order to obtain one premium rating for each session.

As an additional precaution, we also dropped any trials with a duration less than 1 s, which was extremely rare. We did not require trials be under a specific duration since there are many valid reasons for a long trial that should not impact quality.

4.2. Embedding inference

Given a training dataset composed of trials, observed outcomes, and sample weights $\mathcal{D} = \{\mathcal{T}, \mathbf{y}, \mathbf{w}\}$, our goal is to infer an ensemble of psychological embeddings. The remainder of this section details inference for a single embedding $\mathbf{Z} \in \mathbb{R}^{n \times d}$, where n indicates the number of stimuli and d is a hyperparameter that determines the dimensionality of the embedding. For improved readability, \mathbf{Z}_i is used to denote the i th row vector (*i.e.*, the i th stimulus embedding).

The set of m trials is defined as $\mathcal{T} = \{\mathcal{T}_1, \dots, \mathcal{T}_m\}$. For each trial, $\mathcal{T}_i = \{q_i, \mathcal{R}_i\}$, where q_i indicates the query stimulus index and \mathcal{R}_i indicates an ordered n -tuple of reference stimuli indices. The references are ordered so that a deterministic rule can be applied to enumerate all possible outcomes associated with a particular trial. In general, a trial with r references requiring c ranked choices has $k = \frac{r!}{(r-c)!}$ distinct outcomes. Thus y_i indicates the categorical outcome of the i th trial and w_i the corresponding sample weight. Sample weights are determined by a trial’s corresponding session grade.

Assuming independence between trials, the likelihood function is

$$p(\mathcal{D}|\mathbf{Z}) = \prod_i^m p(y_i|\mathcal{T}_i, \mathbf{Z})^{w_i}. \quad (1)$$

Sample weights are integrated into the likelihood so that

the weights linearly scale each trial’s contribution in logarithmic space. Following [43], the embedding is linked to observed behavior using a Luce’s ratio of strengths formulation [31] that is widely used in psychology [48, 49, 33, 34]. For a single 8-rank-2 trial, the likelihood is

$$p(y_i|\mathcal{T}_i, \mathbf{Z}) = \frac{s(\mathbf{Z}_{q_i}, \mathbf{Z}_a)}{\sum_{r \in \mathcal{R}_i} s(\mathbf{Z}_{q_i}, \mathbf{Z}_r)} \frac{s(\mathbf{Z}_{q_i}, \mathbf{Z}_b)}{\sum_{r \in \mathcal{R}_i - a} s(\mathbf{Z}_{q_i}, \mathbf{Z}_r)}. \quad (2)$$

The subscripts a and b are semantic pointers to reference indices in \mathcal{R}_i . The subscript a indicates the first choice reference and b indicates the second choice reference. Motivated by psychological theory [48, 49, 33, 34], this work assumes the similarity function

$$s(\mathbf{Z}_i, \mathbf{Z}_j) = \exp(-\beta \|\mathbf{Z}_i - \mathbf{Z}_j\|), \quad (3)$$

although alternative functions could also be used [1, 56, 60]. In this work, β serves as a convenience parameter and is set to 10. The prior distribution of the embedding is

$$p(\mathbf{Z}_i) = \mathcal{N}(0, \sigma \mathbf{I}), \quad (4)$$

where σ is a scaling parameter fit with the other free parameters.

To perform variational inference, the true posterior $p(\mathbf{Z}|\mathcal{D})$ is approximated with the variational distribution $q(\mathbf{Z}|\mathcal{D})$ that assumes

$$\mathbf{Z}_i \sim \mathcal{N}(\boldsymbol{\mu}^{(i)}, \boldsymbol{\Sigma}^{(i)}). \quad (5)$$

To limit model complexity, $\boldsymbol{\Sigma}^{(i)}$ is constrained to be a diagonal covariance matrix. The set of all free parameters is denoted by $\Theta = \{\boldsymbol{\mu}^{(i)}, \boldsymbol{\Sigma}^{(i)}, \sigma | 1 \leq i \leq n\}$. The optimization objective is to minimize the variational free energy

$$L(\mathcal{D}, \Theta) = D_{\text{KL}}(q(\mathbf{Z}|\Theta) \| p(\mathbf{Z})) - \mathbb{E}_{q(\mathbf{Z}|\Theta)} \log p(\mathcal{D}|\mathbf{Z}), \quad (6)$$

where the first term is the complexity cost and the second term is the likelihood cost [19, 40]. After minimizing the variational free energy, the approximate posterior distribution provides a concise description of the uncertainty associated with each embedding point. It is worth noting that the posterior distribution reflects uncertainty due to a lack of data as well as uncertainty due to inter- and intra-participant inconsistency.

Embedding algorithms are sensitive to initial conditions and prone to discovering local optima. In a low-data regime, different weight initializations can result in solutions that provide equally good fits to the training data, but exhibit qualitatively different structure. Since the embeddings are used downstream to select future trials, we limit the influence of any single model by using an ensemble composed of three equally weighted embeddings. Similar to standard

bootstrap aggregating, variability is encouraged by withholding a different 5% subset of the training data for each embedding.

The 5% withheld data determined the dimensionality hyperparameter d . For a given iteration, independent ensembles were inferred with different dimensionality values. The ensemble with the best categorical cross entropy determined the selected dimensionality. Since inferring models is somewhat costly, this process was not performed at every iteration. Instead, the dimensionality was set based on the chosen dimensionality of the previous iteration.

At the beginning of each iteration, the ensemble was partially initialized using weights from the previous iteration. Using categorical cross-entropy on the validation set, the two top performing models resumed training from the weights of the previous iteration. The worst-performing model was re-initialized with a fresh set of weights. Reusing weights from the previous iteration allowed the active learning procedure to build upon previous choices, while reducing the risk of becoming stuck in local optima.

4.3. Trial selection

Given an ensemble of psychological embeddings, the goal is to select trials that will maximize expected information gain. Given the enormous number of possible trials, it is not feasible to compute the information gain for all trials. Instead, the search space is reduced using a set of heuristics. First, we describe how information gain is computed for an arbitrary trial. Second, we describe the heuristics used to subsample a set of candidate trials.

The expected information gain of a candidate trial \mathcal{T}_{m+1} is the difference between the Shannon entropy of the current embedding (using all observations collected so far) and the expected Shannon entropy if the embedding was inferred with the candidate trial

$$I(\mathbf{Z}, y|\mathcal{D}, \mathcal{T}_{m+1}) = H(\mathbf{Z}|\mathcal{D}) - H(\mathbf{Z}|\mathcal{D}, \mathcal{T}_{m+1}, y). \quad (7)$$

The discrete random variable y represents all possible outcomes associated with the candidate trial. Using the identity for mutual information and sample-based integral approximations, the expected information gain for a candidate trial can be computed using samples from the posterior [43]. For a given candidate trial, the ensemble-level expected information gain is the mean expected information gain of all models.

The set of candidate trials is constrained using two heuristics that operate in tandem. The first heuristic influences which stimuli serve as a query stimulus. Query stimuli are stochastically sampled in proportion to frequency-adjusted stimulus entropy,

$$P(q) \propto \frac{1}{c_q + 1} \tilde{H}(\mathbf{Z}_q|\mathcal{D}), \quad (8)$$

where c_q indicates the number of times stimulus q has already served as a query. The tilde notation is used to indicate measures derived from an equal-weight ensemble average. For each iteration, 828 query stimuli are stochastically sampled without replacement.

The second heuristic influences which stimuli serve as a reference stimuli for the chosen queries. Given a query, references are stochastically sampled without replacement, in proportion to their ensemble-average expected similarity,

$$P(r|q) \propto \mathbb{E}\tilde{s}(\mathbf{Z}_q, \mathbf{Z}_r). \quad (9)$$

The set of eligible reference stimuli was restricted to the 500 nearest neighbor stimuli (1% of the dataset), where nearest neighbors was also determined by ensemble-average, expected similarity. For each query, 10,000 candidate trials were assembled by choosing from the eligible set of reference stimuli. The eligible reference stimuli were restricted in order to ensure that there was sufficient sampling of high similarity trials. Without this threshold, it is very unlikely that the sampled trials will adequately probe *fine-grained* similarity structure. The ensemble-average information gain was computed for all candidate trials, and the top three trials for each query were retained (2484 trials). Three trials were retained in order to lessen the influence of any one participant, while also providing different information.

While the selection heuristics curb computational cost, they risk driving the active learning procedure into a local optimum. Given the large number of stimuli, any given stimulus is at risk of being embedded among highly dissimilar neighbors. Purely by chance, some stimuli may be repeatedly embedded among high dissimilar neighbors, preventing the detection of a signal that could move the stimulus to more similar neighbors. To offset this risk, 828 additional *confirmation* trials were included in each iteration. A confirmation trial was assembled by first stochastically sampling a query in proportion to the ensemble-average entropy. Next, two references were randomly selected from among the 500 nearest neighbors and six references were randomly selected from stimuli outside of the 500-nearest neighborhood. If the stimulus was embedded in an appropriate neighborhood, participants will select the two references drawn from the 500 nearest neighbors. If not, participant choices will help move the stimulus to a more appropriate neighborhood.

The goal was not to minimize uncertainty associated with a particular trial, but to minimize uncertainty associated with the embedding as a whole. On average, expected information gain is higher for trials that have never been asked. Rather than require that n observations be obtained for each trial, we used an aggressive threshold for catch trial performance and let the active learning procedure decide which trials were most useful. Consequently, the majority

of trials were only shown once.

5. Evaluating dataset convergence

Since the human similarity judgments and companion embeddings are used to evaluate the internal representations of other models, it is critical that a sufficient amount of data has been collected. Convergence is primarily assessed by examining how the psychological embeddings change as more data is collected. We focus on three measures of the psychological embeddings: *coarse-grained loss*, *within-ensemble agreement*, and *across-iteration agreement*. These measures and corresponding results are described in the remainder of this section.

5.1. Coarse-grained loss

A simple measure of convergence is to evaluate the ability of an ensemble to correctly predict a *coarse-grained set* of similarity judgments. It is misleading to call this set of observations a test set. A basic requirement is that a test set be drawn from the same distribution as the training set. Since the coarse-grained trials are sampled randomly, they are unlikely to contain trials composed of highly similar images. As a consequence, this set will contain few observations that probe fine-grained similarity. We therefore refer to the corresponding loss (*i.e.*, categorical cross-entropy) as coarse-grained loss.

By computing coarse-grained loss for ensembles trained at different iterations, we can see how generalization performance changes with additional data. Coarse-grained loss that has asymptoted at a minimum value provides some evidence that the psychological embeddings have converged.

5.2. Within-ensemble agreement

As an alternative to generalization performance, we can examine the learned similarity structure of the psychological embeddings [43]. If a sufficient amount of data has been collected, then all models *within* an ensemble should have comparable embeddings. The degree of agreement is quantified by comparing the implied pairwise similarity matrix of one model with the similarity matrix of another model. Since we have access to the posterior distribution, we can assemble the expected similarity matrix for each model. By computing the Pearson correlation between the upper diagonal elements of the similarity matrix, we can quantify the degree that the two models agree [52, 25]. Since there are three models, within-ensemble correlation is defined as the average over all pairwise model correlations.

5.3. Consecutive-ensemble agreement

Analogous to within-ensemble correlation, we can examine how the learned similarity structure differs *across* iterations. If a sufficient amount of data has been collected,

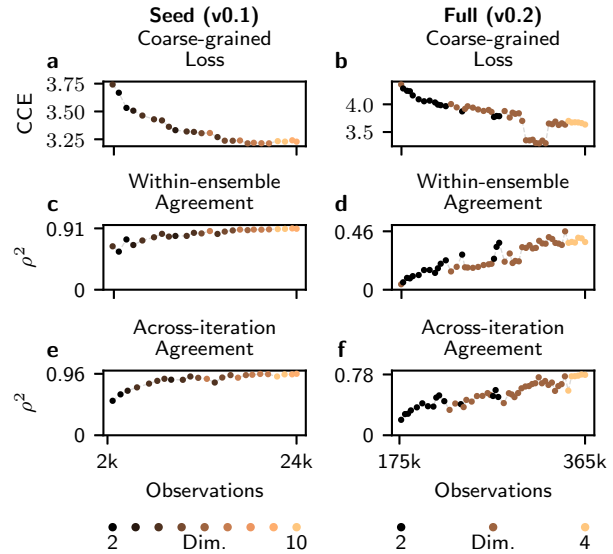


Figure 3. Measures for assessing the convergence of ImageNet-HSJ. The left and right columns show measures for seed (v0.1) and full (v0.2) dataset respectively. Each panel plots the performance of a measure (ordinate) as the number of observations increases (abscissa). The first row (a,b), demonstrates how categorical cross-entropy for a novel set of observations decreases as the amount of training data increases. It should be noted that the results in the two columns are not comparable since they use different coarse-grained sets. The second row (c,d), shows that models within an ensemble become increasingly consistent as measured by the squared Pearson correlation between similarity matrices. Likewise, the third row shows that ensembles of consecutive iterations become increasingly stable as measured by the squared Pearson correlation between ensemble similarity matrices. Each point represents a metric evaluation of a single iteration (not all iterations are shown). Points are color-coded based on the selected dimensionality of the ensemble.

then including more data should not alter the inferred ensemble. To compute consecutive-ensemble correlations, the Pearson correlation is computed between the expected similarity matrix of ensembles belonging to consecutive iterations.

5.4. Convergence results

All measures indicate that the psychological embeddings for the seed subset have converged, while the full dataset has not converged (Figure 3). It should be noted that it is not appropriate to directly compare Figure 3a and b. A different coarse-grained set is used for the seed and full dataset since since these two versions contain a different number of unique stimuli (1,000 vs 50,000 respectively). The dip in Figure 3b clearly demonstrates how susceptible embeddings are to local optima.

6. Using the dataset to evaluate target models

Given a set of high-quality human similarity judgments, it is now possible to assess the human-likeness of representations from arbitrary models. While there are a variety of approaches for performing this assessment, we describe two straightforward metrics for evaluating a *target model*. The first metric converts the 8-rank-2 trials into a set of triplet inequality relations and counts the proportion of triplets satisfied by the target model. The second metric examines the second order isomorphism between the internal representations of the target model and a psychological embedding [52, 25].

6.1. Triplet accuracy

The first evaluation metric provides a minimalist approach for assessing the representations of a target model. This metric does not take advantage of the inferential machinery of a psychological embedding, but uses the similarity judgments directly. There are two primary advantages associated with this evaluation metric. First, this metric is less sensitive to issues of dataset convergence. While the psychological embeddings may change as more data is collected, the veracity of the observations does not change. Second, the researcher is only required to select a function for computing distances between representations of the target model.

Since each observed similarity judgment trial implies a set of triplet similarity relations [61], we can convert an arbitrary observation into an equivalent set of triplet observations. For example, consider an 8-rank-2 trial observation denoted as $q : a > b > [c, d, e, f, g, h]$, where q indicates the query, a indicates the first choice, b indicates the second choice, and $c-h$ the remaining unselected references. This trial implies 13 triplet observations, e.g., $q : a > b$, $q : a > c$.

The set of triplet observations is then used to compute the triplet accuracy for a target model. Given a triplet observation of the form $q : a > b$, and a distance function $d(\cdot, \cdot)$, we can determine if the target model correctly ranks a triplet,

$$d(\mathbf{Z}_q, \mathbf{Z}_a) < d(\mathbf{Z}_q, \mathbf{Z}_b), \quad (10)$$

where \mathbf{Z} indicates the stimulus representation of the target model. The mean number of correct predictions yields the triplet accuracy.

Given the nature of human similarity judgments, we do not expect a model to correctly rank all triplets. Different participants may not agree on perceived similarity and individuals may not be metric consistent [59].

6.2. Psychological embedding correlation

While triplet accuracy provides a simple method for assessing a target model, it is a relatively coarse metric that

does not fully account for the similarity structure implied by the observations. Given sufficient observations, a stronger approach is to compare the internal representations of the target model to the internal representations of a psychological embedding. In order to execute this evaluation scheme, the researcher must make two choices. First, the researcher must select a psychological embedding. Second, the researcher must decide which similarity function to use when assembling a similarity matrix for the target model.

The psychological embedding correlation is computed in a similar manner as described above and in RSA [25]. After assembling a pair-wise similarity matrix for the psychological embedding and the target model, we compute the Spearman correlation coefficient between the upper triangular portions of the two similarity matrices. Spearman correlation is used since we cannot assume the target model will yield similarities that are linearly related to similarities generated from the psychological embedding.

6.3. Target model results

The two evaluation metrics are demonstrated by assessing the representations of some popular and successful target models. The selected supervised target models (and corresponding layer) are VGG16 (fc2), VGG19 (fc2), ResNet50 (avg_pool) [16], ResNet101 (avg_pool), ResNet152 (avg_pool), ResNet50V2 (avg_pool), Xception (avg_pool) [6], DenseNet121 (avg_pool) [21], InceptionV3 (avg_pool) [55], Inception-ResNetV2 (avg_pool) [8], MobileNet (avg_pool) [20]. In addition to the supervised models, we include DeepCluster (VGG16, fc2) [4], since unsupervised models may be more capable of learning task-general representations. With the exception of DeepCluster, the pre-trained weights for the target models are obtained through Keras Applications [7]. In addition to model-specific preprocessing, images are appropriately resized from the largest centered crop that preserves the aspect ratio. The weights and preprocessing procedure for DeepCluster are obtained from the DeepCluster GitHub repository.

Computing the triplet accuracy for the target models reveals that all models are better than chance (Table 6.3). Since computing triplet accuracy requires a choice of distance function, triplet accuracy is computed using L1, L2, and cosine distance. In all cases, cosine distance outperforms L1 and L2 distance. ResNet50 has the highest triplet accuracy for both the seed and full dataset.

For comparison, we also compute triplet accuracy for a psychological embedding. The psychological embeddings perform the best at predicting the triplet ratings, despite being blind to pixel information. In addition, the other target models can leverage approximately one million images worth of training experience. While not a completely fair comparison since the psychological embedding is ef-

Target Model	Seed (v0.1)			Full (v0.2)		
	L1	L2	cos	L1	L2	cos
Psych. Emb.	–	81.7	–	–	80.7	–
VGG16	61.1	57.8	69.6	67.4	63.1	74.4
VGG19	60.9	57.7	69.8	67.3	63.1	74.6
ResNet50	62.3	62.5	70.9	68.6	69.0	75.2
ResNet101	60.8	61.6	70.2	68.0	68.8	75.0
ResNet152	60.7	61.3	70.4	68.1	68.8	75.2
ResNet50V2	57.6	57.9	67.0	65.4	66.7	72.9
Xception	55.3	53.9	56.3	66.2	65.1	67.2
DenseNet121	64.2	62.7	67.3	70.9	70.0	73.3
InceptionV3	54.9	54.2	55.6	65.1	64.7	65.5
Inception						
ResNetV2	53.2	53.0	54.0	64.8	64.8	65.2
MobileNet	64.7	64.3	67.5	69.9	69.5	72.8
DeepCluster						
VGG16	56.8	58.0	61.9	59.0	60.6	65.8

Table 2. Triplet accuracy for various target models. Triplet accuracy measures the ability of a target model to correctly predict implicit triplet inequalities derived from the 8-rank-2 similarity judgments. The best performing result for each model is emphasized in bold. The high triplet accuracy of the psychological embedding demonstrates that there is substantial room for improvement.

fectively trained on the triplet data, one could imagine performing a hold-one-trial-out test procedure, which would result in a very similar triplet accuracy.

The psychological embedding correlations are computed by reusing the ensembles fit during the active learning procedure. Since computing embedding correlation requires a choice of similarity function, results are presented using a dot product and cosine similarity. Mirroring the triplet accuracy results, ResNet50 is the top performer. In the majority of cases, cosine similarity yields the best correlation for each target model. For comparison, if we compute the rank correlation between two psychological embeddings trained on the same data, the Spearman correlation between the corresponding expected similarity matrices is 0.81 and 0.62 for the seed and full dataset respectively.

7. Discussion

Assessing how consistent a model’s internal representations are with human judgments can offer a useful evaluation metric. Here, we presented a method to efficiently infer embedding spaces from human judgments for a large number of items. We were able to consider a problem an order of magnitude larger than previous work by leveraging variational inference, ensemble models, and active learning.

We demonstrated that human similarity judgments collected by our method are useful in evaluating model representations. Critically, these metrics can be used to assess models trained using supervised or unsupervised algo-

Target Model	Seed (v0.1)		Full (v0.2)	
	dot	cos	dot	cos
VGG16	0.25	0.28	0.25	0.27
VGG19	0.27	0.29	0.26	0.27
ResNet50	0.23	0.36	0.25	0.30
ResNet101	0.18	0.33	0.23	0.29
ResNet152	0.18	0.34	0.24	0.29
ResNet50V2	0.26	0.30	0.23	0.24
Xception	0.01	0.06	0.06	0.04
DenseNet121	0.17	0.26	0.19	0.22
InceptionV3	0.01	0.06	0.04	0.03
Inception				
ResNetV2	0.00	-0.00	-0.00	0.02
MobileNet	0.23	0.29	0.19	0.23
DeepCluster				
VGG16	0.01	0.11	-0.02	0.07

Table 3. The psychological embedding correlation for various target models. Embedding correlation computes the Spearman correlation between the similarity matrices of an ensemble of psychological embeddings and a target model.

rithms. We expected that representations from unsupervised models would be more general and conform better to human judgment, but, interestingly, DeepCluster was one of the worst performing models. Although we only evaluated one unsupervised model, this result highlights the challenging *semantic gap* that any learning algorithm must bridge, which we hope Imagenet-HSJ will help close. Future work will need to assess other promising models such as CORnet-S [27] and SwAV [5].

Within the set of supervised models considered, the results reflected the history of the field. Initially, classification models that performed better on ImageNet also better corresponded with human behavioral and brain data [26]. Although newer models continue to advance classification performance, we found that correspondence with human similarity judgments has not increased. An analogous trend can be observed in the model rankings on BrainScore [47], which evaluates correspondences between brain activity and model representations.

Human similarity judgments are a rich source of information. They illuminate how humans structure the world and flexibly shift between organization schemes depending on context. We offered a method to infer such embedding spaces. Although we focused on using the embedding spaces to evaluate model representations, we expect there will be other potential applications for our method and dataset.

Acknowledgments: NIH grant no. 1P01HD080679, Wellcome Trust Investigator Award no. WT106931MA and a Royal Society Wolfson Fellowship 183029 to B.C.L.

References

- [1] Sameer Agarwal, Josh Wills, Lawrence Cayton, Gert Lanckriet, David Kriegman, and Serge Belongie. Generalized non-metric multidimensional scaling. In Marina Meila and Xiaotong Shen, editors, *Proceedings of the Eleventh International Conference on Artificial Intelligence and Statistics*, volume 2 of *Proceedings of Machine Learning Research*, pages 11–18, San Juan, Puerto Rico, 21–24 Mar 2007. PMLR.
- [2] M. Attarian, B. D. Roads, and M. C. Mozer. Transforming neural network representations to predict human judgments of similarity. In *Workshop on Shared Visual Representations in Human and Machine Intelligence (SVRHM 2020)*, 2020.
- [3] Nicholas Baker, Hongjing Lu, Gennady Erlikhman, and Philip J Kellman. Deep convolutional networks do not classify based on global object shape. *PLoS Computational Biology*, 14(12):e1006613, 2018.
- [4] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Computer Vision – ECCV 2018*. Springer International Publishing, September 2018.
- [5] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [6] F. Chollet. Xception: Deep learning with depthwise separable convolutions. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1800–1807, 2017.
- [7] François Chollet et al. Keras. <https://keras.io>, 2015.
- [8] Szegedy Christian, Ioffe Sergey, Vanhoucke Vincent, and AA Alexander. Inception-v4 inception-resnet and the impact of residual connections on learning. In *AAAI*, volume 4, 2017.
- [9] Çağatay Demiralp, Michael S Bernstein, and Jeffrey Heer. Learning perceptual kernels for visualization design. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1933–1942, 2014.
- [10] 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. IEEE, 2009.
- [11] I. El-Naqa, Yongyi Yang, N. P. Galatsanos, R. M. Nishikawa, and M. N. Wernick. A similarity learning approach to content-based image retrieval: application to digital mammography. *IEEE Transactions on Medical Imaging*, 23(10):1233–1244, 2004.
- [12] Tal Golan, Prashant C. Raju, and Nikolaus Kriegeskorte. Controversial stimuli: pitting neural networks against each other as models of human recognition, 2020.
- [13] J. C. Gower. Some distance properties of latent root and vector methods used in multivariate analysis. *Biometrika*, 53:325–338, 1966.
- [14] Quentin F Gronau and Michael D Lee. Bayesian inference for multidimensional scaling representations with psychologically-interpretable metrics, Jan 2020.
- [15] Guo-Dong Guo, A. K. Jain, Wei-Ying Ma, and Hong-Jiang Zhang. Learning similarity measure for natural image retrieval with relevance feedback. *IEEE Transactions on Neural Networks*, 13(4):811–820, 2002.
- [16] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778. IEEE, June 2016.
- [17] 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, 10 2019.
- [18] 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 Behavior*, pages 1173–1185, 2020.
- [19] Matthew D. Hoffman, David M. Blei, Chong Wang, and John Paisley. Stochastic variational inference. *J. Mach. Learn. Res.*, 14(1):1303–1347, May 2013.
- [20] Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications, 2017.
- [21] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [22] Kevin G Jamieson, Lalit Jain, Chris Fernandez, Nicholas J Glattard, and Rob Nowak. Next: A system for real-world development, evaluation, and application of active learning. In *Advances in Neural Information Processing Systems*, pages 2656–2664, 2015.
- [23] Matt Jones and Bradley C. Love. Beyond common features: The role of roles in determining similarity. *Cognitive Psychology*, 55(3):196 – 231, 2007.
- [24] Seyed-Mahdi Khaligh-Razavi and Nikolaus Kriegeskorte. Deep supervised, but not unsupervised, models may explain it cortical representation. *PLOS Computational Biology*, 10(11):1–29, 11 2014.
- [25] Nikolaus Kriegeskorte, Marieke Mur, and Peter Bandettini. Representational similarity analysis - connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, 2:4, 2008.
- [26] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. Burges, L. Bottou, , and K. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, page 1097–1105. Curran Associates, Inc., 2012.
- [27] Jonas Kubilius, Martin Schrimpf, Kohitij Kar, Rishi Rajalingham, Ha Hong, Najib Majaj, Elias Issa, Pouya Bashivan, Jonathan Prescott-Roy, Kailyn Schmidt, Aran Nayebi, Daniel Bear, Daniel L Yamins, and James J DiCarlo. Brain-like object recognition with high-performing

- shallow recurrent anns. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32, pages 12805–12816. Curran Associates, Inc., 2019.
- [28] Brenden M Lake, Wojciech Zaremba, Rob Fergus, and Todd M Gureckis. Deep neural networks predict category typicality ratings for images. In *CogSci*, 2015.
- [29] L. Li, V. Malave, A. Song, and A. Yu. Extracting human face similarity judgments: Pairs or triplets? *Journal of Vision*, 16:719, 2016.
- [30] Zhe Li, Wieland Brendel, Edgar Walker, Erick Cobos, Taliah Muhammad, Jacob Reimer, Matthias Bethge, Fabian Sinz, Zachary Pitkow, and Andreas Tolias. Learning from brains how to regularize machines. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [31] R Duncan Luce. *Individual choice behavior: A theoretical analysis*. Wiley, New York, NY, 1959.
- [32] Gregory L Murphy and Douglas L Medin. The role of theories in conceptual coherence. *Psychological Review*, 92(3):289–316, 1985.
- [33] Robert M Nosofsky. Overall similarity and the identification of separable-dimension stimuli: A choice model analysis. *Perception & Psychophysics*, 38(5):415–432, 1985.
- [34] Robert M Nosofsky. Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, 115:39–57, 1986.
- [35] Robert M. Nosofsky, Craig A. Sanders, Brian J. Meagher, and Bruce J. Douglas. Toward the development of a feature-space representation for a complex natural category domain. *Behavior Research Methods*, 50:530–556, 2018.
- [36] 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, 2018.
- [37] Joshua C. Peterson, Ruairidh M. Battleday, Thomas L. Griffiths, and Olga Russakovsky. Human uncertainty makes classification more robust. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [38] Ekta Prashnani, Hong Cai, Yasamin Mostofi, and Pradeep Sen. Pieapp: Perceptual image-error assessment through pairwise preference. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [39] Rishi Rajalingham, Elias B. Issa, Pouya Bashivan, Kohitij Kar, Kailyn Schmidt, and James J. DiCarlo. Large-scale, high-resolution comparison of the core visual object recognition behavior of humans, monkeys, and state-of-the-art deep artificial neural networks. *Journal of Neuroscience*, 38(33):7255–7269, 2018.
- [40] Rajesh Ranganath, Sean Gerrish, and David Blei. Black box variational inference. In *Artificial Intelligence and Statistics*, pages 814–822. PMLR, 2014.
- [41] Martina A Rau, Blake Mason, and Robert Nowak. How to model implicit knowledge? similarity learning methods to assess perceptions of visual representations. *International Educational Data Mining Society*, 2016.
- [42] Brett D. Roads and Michael C. Mozer. Improving human-machine cooperative classification via cognitive theories of similarity. *Cognitive Science: An Multidisciplinary Journal*, 41:1394–1411, 2017.
- [43] Brett D. Roads and Michael C. Mozer. Obtaining psychological embeddings through joint kernel and metric learning. *Behavior Research Methods*, 2019.
- [44] Brett D. Roads, Buyun Xu, June K. Robinson, and James W. Tanaka. The easy-to-hard training advantage with real-world medical images. *Cognitive Research: Principles and Implications*, 3(38), 2018.
- [45] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.
- [46] Craig A Sanders and Robert M Nosofsky. Training deep networks to construct a psychological feature space for a natural-object category domain. *Computational Brain & Behavior*, 3:229–251, 2020.
- [47] Martin Schrimpf, Jonas Kubilius, Michael J. Lee, N. Apurva Ratan Murty, Robert Ajemian, and James J. DiCarlo. Integrative benchmarking to advance neurally mechanistic models of human intelligence. *Neuron*, 108(3):413 – 423, 2020.
- [48] Roger N Shepard. Stimulus and response generalization: A stochastic model relating generalization to distance in psychological space. *Psychometrika*, 22(4):325–345, 1957.
- [49] Roger N Shepard. Stimulus and response generalization: tests of a model relating generalization to distance in psychological space. *Journal of Experimental Psychology*, 55(6):509–523, 1958.
- [50] Roger N Shepard. The analysis of proximities: Multidimensional scaling with an unknown distance function. i. *Psychometrika*, 27(2):125–140, 1962.
- [51] Roger N Shepard. The analysis of proximities: Multidimensional scaling with an unknown distance function. ii. *Psychometrika*, 27(3):219–246, 1962.
- [52] Roger N Shepard and Susan Chipman. Second-order isomorphism of internal representations: Shapes of states. *Cognitive Psychology*, 1(1):1 – 17, 1970.
- [53] Scott Sievert, Daniel Ross, Lalit Jain, Kevin Jamieson, Rob Nowak, and Robert Mankoff. Next: A system to easily connect crowdsourcing and adaptive data collection. In *Proceedings of the 16th Python in Science Conference*, pages 113–119, 2017.
- [54] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, 2015.
- [55] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.

- [56] Omer Tamuz, Ce Liu, Serge Belongie, Ohad Shamir, and Adam Tauman Kalai. Adaptively learning the crowd kernel. *arXiv preprint arXiv:1105.1033*, 2011.
- [57] Warren S Torgerson. Multidimensional scaling: I. theory and method. *Psychometrika*, 17(4):401–419, 1952.
- [58] W. S. Torgerson. *Theory and methods of scaling*. New York: Wiley, 1958.
- [59] Amos Tversky. Features of similarity. *Psychological Review*, 84(4):327–352, 1977.
- [60] L. van der Maaten and K. Weinberger. Stochastic triplet embedding. In *Machine Learning for Signal Processing (MLSP), 2012 IEEE International Workshop on*, pages 1–6, Sept 2012.
- [61] Catherine Wah, Grant Van Horn, Steve Branson, Subhansu Maji, Pietro Perona, and Serge Belongie. Similarity comparisons for interactive fine-grained categorization. In *Computer Vision and Pattern Recognition (CVPR)*, Columbus, OH, June 2014.
- [62] Catherine Wah, Subhansu Maji, and Serge Belongie. Learning localized perceptual similarity metrics for interactive categorization. In *Applications of Computer Vision (WACV), 2015 IEEE Winter Conference on*, pages 502–509. IEEE, 2015.
- [63] Michael J Wilber, Iljung S Kwak, and Serge J Belongie. Cost-effective hits for relative similarity comparisons. In *Second AAAI Conference on Human Computation and Crowdsourcing*, 2014.
- [64] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [65] Bin Zheng, Amy Lu, Lara A. Hardesty, Jules H. Sumkin, Christiane M. Hakim, Marie A. Ganott, and David Gur. A method to improve visual similarity of breast masses for an interactive computer-aided diagnosis environment. *Medical Physics*, 33(1):111–117, 2006.