

# Visual Knowledge Tracing

Neehar Kondapaneni<sup>1</sup> Pietro Perona<sup>1</sup> Oisín Mac Aodha<sup>2</sup>

<sup>1</sup>Caltech <sup>2</sup>University of Edinburgh

**Abstract.** Each year, thousands of people learn new visual categorization tasks – radiologists learn to recognize tumors, birdwatchers learn to distinguish similar species, and crowd workers learn how to annotate valuable data for applications like autonomous driving. As humans learn, their brain updates the visual features it extracts and attend to, which ultimately informs their final classification decisions. In this work, we propose a novel task of tracing the evolving classification behavior of human learners as they engage in challenging visual classification tasks. We propose models that jointly extract the visual features used by learners as well as predicting the classification functions they utilize. We collect three challenging new datasets from real human learners in order to evaluate the performance of different visual knowledge tracing methods. Our results show that our recurrent models are able to predict the classification behavior of human learners on three challenging medical image and species identification tasks.

**Keywords:** visual classification, knowledge tracing, human learning

## 1 Introduction

Humans excel at learning new concepts even when they have only received limited explicit supervision [28,55]. Key to our success is our ability to extract informative and generalizable representations from the world around us and our ability to update these representations given relatively sparse feedback. This capacity, in turn, enables us to perform complex tasks such as spatial navigation and visual categorization with apparent ease.

Despite recent progress that has been made in computer vision in learning visual representations through self-supervision alone [54,9,14], large amounts of supervision are still required to make best use of the resulting features [12]. In light of this, it is important for us to better understand: (i) what are the properties that make representations learned by *humans* so effective, (ii) how are these representations learned, and (iii) can we predict the classification behavior of humans during learning? Our ultimate goal is to obtain better insight into how humans are such effective learners, which can then potentially inform new learning mechanisms for future artificial systems.

In order to attempt to address some of these questions, in this work we explore the problem of visual knowledge tracing. In the educational data mining community, knowledge tracing is the problem of monitoring and predicting the

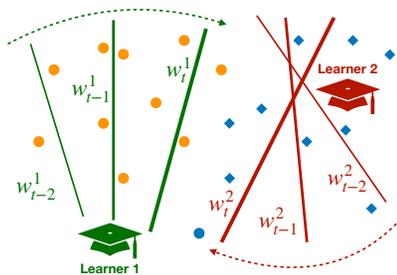


Fig. 1: We model a learner as an evolving classifier in a learned feature space. We assume the feature space is static and, as the learner is presented with images and class labels over time, their internal classification function self-updates. Here, we illustrate this for three time-steps, for two learners, learning a binary classification task (orange versus blue).

evolving knowledge state of a learner engaged in a learning task [26]. Recent work has applied advances in deep learning to knowledge tracing for question and answer-style text datasets and has investigated applications in domains such as mathematics education [34,33,35]. However, tracing the behavior of humans that are engaged in learning and performing challenging visual categorization tasks is underexplored. Most closely related is the work on deep metric learning that attempts to learn human aligned visual representations from sparse human annotations [19]. However, these works typically make simplifying assumptions, e.g. assuming the learners are not changing over time (i.e. they are ‘static’) or that the visual criteria used by the learners is the same across all learners. Instead, we explore a more challenging visual knowledge tracing setting where the learners are assumed to be non-stationary during learning, i.e. the visual features they use to perform the classification task at hand can, and likely do, change over time (see Fig. 1).

We present a recurrent neural network-based approach for visual knowledge tracing. Once trained, our models are capable of predicting the classification behavior of human learners that were not observed during training. The proposed models make use of the history of previous learner responses, images, and ground truth class labels in order to predict their future responses. Through experiments on three challenging image classification datasets we show that our models are superior to baseline approaches. Our models are capable of tracing the learning dynamics more accurately than non-recurrent baselines.

We make the following three contributions: (i) A new model for visual knowledge tracing that jointly estimates the visual features and per-time-step classification function used by non-stationary human learners. (ii) A new set of annotations for three benchmark evaluation datasets collected from humans engaged in learning challenging visual classification tasks. (iii) A detailed comparison of several visual knowledge tracing methods on these datasets.

## 2 Related Work

### 2.1 Metric Learning

The goal of metric learning is to learn perceptual embeddings such that distance in the lower dimensional embedding space encodes information related to semantic similarity. Pre-deep learning approaches to metric learning were primarily concerned with learning embeddings directly for each item in an input set. In the case of learning from human supervision, approaches that use relative similarity judgements have been shown to be effective [23,45,53,37]. These

methods have also been extended to the adaptive setting where the model can decide which items to request annotations for in order to speed up training [43].

More recently, end-to-end metric learning methods have attempted to parameterize an embedding function (e.g. a convolutional neural network) directly [38,29]. As a result, they are able to embed any new item into the embedding space, even those not seen at training time. Representative earlier applications of this line of work include image ranking [49] and face recognition [40]. The standard assumption made by the majority of these methods is that only *one* similarity criterion is being used. However, when collecting data from human annotators, different individuals may be using different visual criteria when making classification and similarity judgements, e.g. one individual could be using shape, while the other is using color. Furthermore, the same individual may change the criteria they use conditioned on the specific items they are shown, and could change to another criteria when shown another set of items at a different point in time.

There has been some work that attempts to deal with the fact that *different* similarity criteria may be being used. These range from fully supervised, where the similarity criteria is known at training time [48], through to unsupervised methods that attempt to estimate the criteria [21,30,44,52]. Attempts have also been made to probabilistically estimate item embeddings along with annotator-specific parameters representing the individual criteria they are using [52].

One common assumption made by the above methods is that the learner is stationary, i.e. they are used a predefined and fixed similarity criteria, or small set of criteria, which do not evolve or change over time. This is a reasonable assumption to make when dealing with common everyday object categories where the annotators will likely be familiar with the objects depicted and have an a priori understanding of how the visual features of the objects may vary. However, this assumption is violated in cases where the annotator is in the process of *learning* the visual concepts of interest. In this work, we address this non-stationary setting and show that by doing this, we can more accurately predict the visual classification behavior of real human learners.

## 2.2 Human Category Representation

Existing models of human category representation and learning can be clustered into five major groups: rule-based, prototype-based, exemplar-based, knowledge / theory-based, and decision boundary-based approaches [56,4]. The current consensus is that humans likely use multiple different category-learning systems depending on the specific nature of the task at hand [5,3]. For example, in rule-based tasks, the optimal policy may be easy to verbalize and thus efficiently encoded via a set of rules. In practice however, perceptual tasks such as fine-grained visual categorization can be much harder to represent in this way [8]. Several works have attempted to extract perceptual embeddings that align with human similarity judgements from coarse [15,36] and more fine-grained [32] image collections. [6] showed that with simple linear transformations, pre-trained deep image classifiers can be predictive of human similarity judgements. Relevant

to our work, [7] investigated whether human learning dynamics mimic gradient descent in Artificial Neural Networks when learning visual categories. However, they assumed they had access to the feature space used by the human, whereas we instead attempt to learn this. In this work, we also aim to extract human-aligned representations, but in the more challenging setting whereby our learners are not static, but instead are in the process of learning the categorization task.

### 2.3 Knowledge Tracing

The problem of modelling the hidden state of dynamic learners as they interact with a learning task has also been tackled in the knowledge tracing literature. Bayesian Knowledge Tracing-based methods model a learner’s knowledge state by assuming that the learner can be represented as a Markov process which updates online during learning [13]. Building on this line of work, Deep Knowledge Tracing (DKT) instead uses a recurrent neural network as the underlying tracing model [34], and fully self-attention-based methods have also been proposed [33,35]. It is important to note that conventional knowledge tracing attempts to model knowledge acquisition as a binary variable at the ‘skill’ level (i.e. the visual class) as opposed to the ‘instance’ level (i.e. a specific image). In contrast, our approach jointly learns an image embedding function in addition to being able to capture and predict the individual learning trajectories of multiple different learners. A detailed comparison of the original DKT model and our model is provided in the supplementary material (see Sec. A.2).

### 2.4 Machine Teaching

Estimating the representations used by humans is an important component for developing automated teaching algorithms and systems. Machine teaching algorithms address the teaching problem by generating sequences of instructional examples to show to novice learners in order to improve their ability on a given task [57]. Machine teaching has applications in crowdsourcing where the aim is to efficiently train crowdworkers, in addition to education where the goal is to train new experts, e.g. in medical image analysis [10,2].

There is a growing body of work in computer vision that attempts to teach visual concepts to human learners, e.g. [42,18,27,50,51]. However, many of these works assume a fixed feature space that is generated before teaching begins [42,18,27]. In one experiment, [27] showed that representations that are better aligned with human perception result in improved learner performance on the downstream teaching task. While we do not explicitly investigate teaching algorithms in this work, we instead explore a setting where data is collected from humans engaged in learning a visual categorization task with instructional images selected by a ‘random’ teacher. Importantly, the representations and learner parameterizations extracted by our model can be used directly with computer assisted teaching methods.

### 3 Method

Our goal is to estimate the image classification function used by a human learner that has been provided with a sequence of images and corresponding ground truth class labels as training data. We begin by outlining the problem, and then present our approach to human visual knowledge tracing.

#### 3.1 Problem Setup

Given an image  $\mathbf{x}$  as input, we model a human learner as a classification function that returns a response,  $r^k = \operatorname{argmax}_c P(c|\mathbf{x}, \theta^k)$ . Here,  $r$  is a discrete class label representing the class response for learner  $k$ , i.e.  $r \in \{1, \dots, C\}$ , where  $C$  is the number of possible classes, and  $\theta^k$  are unobserved parameters representing the state of the learner. Our learners are not stationary as their internal ‘knowledge state’ changes depending on the information they have previously been exposed to that is relevant to the task. As a result, for a given learner  $k$  we model their classification function at time  $t$  as  $r_{t+1}^k = \operatorname{argmax}_c P(c|\mathbf{x}, \theta_t^k, \mathbf{x}_{1:t}^k, y_{1:t}^k, r_{1:t}^k)$ . Here,  $(\mathbf{x}_{1:t}^k, y_{1:t}^k, r_{1:t}^k)$  is the history of images, ground truth class labels, and responses that a learner  $k$  has seen, and provided, up to and including time  $t$ .

Specifically, at each training time-step, a learner  $k$  is presented with an image  $\mathbf{x}$ , they provide their response  $r$ , and are given feedback in the form of the correct class label  $y$  (Fig. 2B). However, fitting a model for an individual learner with a single response per time-step is difficult. Alternatively, requesting more responses per time-step would reduce the number of teaching examples presented to the learner in the same amount of time. Instead, to overcome this limited information setting, we train a model  $\phi$  across many learners, allowing the model to discover knowledge states and learning rules shared across all learners. Once trained, our model can make predictions for how a learner, who was not observed during training, will classify an image based on their prior classification behavior.

#### 3.2 Tracing Human Learners

Our tracing model  $\phi$  can be decomposed into a feature extractor  $f$ , a classification function  $\psi$ , and a non-learned and non-linear transformation  $\sigma$  (softmax). We explore how to represent the feature extractor and classification function.

A natural choice for the feature extractor  $f$  is a Convolutional Neural Network (CNN). Given an image  $\mathbf{x}$  as input, the feature extractor outputs a  $D$  dimensional vector  $\mathbf{z} = f(\mathbf{x})$ . We will assume that all learners use the same underlying feature extractor which remains constant over time i.e.  $f = f_t^k$ , and that they simply differ in the relative importance they place on different visual features. While these are both big assumptions to make, they are not overly restrictive. For example, a novice and an expert might be engaged in the same visual classification task but differ in the set of visual features they select in order to make their decision. Furthermore, while we assume that the feature extractor remains constant over the time interval of our experiments, we do not assume that the classification function  $\psi$  used by a learner remains static. In the next sections

we will explore different choices for this classification function, comparing simple static classifiers with more expressive recurrent models.

**Static Tracing Model.** The first model we explore is the simplest. Here we assume that all learners use the same classifier which does not vary over time. In this setting,  $\psi$  is a multi-class linear classifier with a weight matrix  $\mathbf{w}$  and per-class biases  $\mathbf{b}$ ,

$$\phi_{static}(\mathbf{x}) = \sigma(\psi(f(\mathbf{x}))) = \sigma(\mathbf{w}^\top f(\mathbf{x}) + \mathbf{b}). \quad (1)$$

This model is similar to conventional metric learning approaches which do not attempt to capture any annotator specific differences related to individual biases or temporal changes. At training time we simply estimate one set of parameters for all learners. This model does not take the response history into account.

**Time-Sensitive Tracing Model.** One obvious limitation of the static tracing model is that it does not take into account the fact that a learner will likely change over time, i.e. they may be much worse at a new classification task early on, but may improve over time as they are shown sequences of example images along with their associated ground truth class labels. A more advanced model, that captures this temporal evolution, is one that has a different classifier for each time-step,

$$\phi_{static\_time}(\mathbf{x}) = \sigma(\mathbf{w}_t^\top f(\mathbf{x}) + \mathbf{b}_t). \quad (2)$$

Again, the same classifiers are shared across all learners, but in this case the weights and biases are different at each time-step, i.e.  $\mathbf{w}_t \neq \mathbf{w}_{t-1}$ .

### 3.3 Recurrent Tracing Models

The previous tracing models do not account for the fact that individual learners may start with different levels of ability and update their internal knowledge state in different ways depending on the information that they are provided with. [34] showed that recurrent networks could be used to track the skill acquisition of human learners engaged in learning math quiz questions. Direct application of their model to our visual categorization setting is not possible as they assume one hot encodings of the query and learner responses as input. Their approach also uses large training sets – on the order of thousands of learners and tens of thousands of interactions. Furthermore, they model knowledge acquisition at the ‘concept’ (i.e. visual category) and not ‘instance’ (i.e. a specific image) level, and thus their approach is not capable of making predictions for items not seen during training. We build on [34] and adapt it to our visual category learning setting by presenting two different recurrent-based models for human visual knowledge tracing.

**Direct Response Model.** Our first model uses a recurrent network to directly predict the responses of a learner given their previous response history,

$$\phi_{direct}(\mathbf{x}) = \sigma(\psi_{rnn}(\mathbf{z}_{1:t}^k, y_{1:t}^k, r_{1:t}^k, \mathbf{z}, y)). \quad (3)$$

Here,  $\psi_{rnn}$  is a recurrent network (in practice we represent this using an LSTM [16]) and  $\mathbf{z}_t = f(\mathbf{x}_t)$  are visual features extracted from our CNN.

This model assumes that a learner’s knowledge state at time  $t$  is defined by the images they have previously seen and their past classification responses. Recurrent models can produce unique transformations for individual learners by conditioning on their hidden states. In this case, after the shared feature extractor transforms an image into a feature vector, the model modifies the feature vector with a series of non-linear transformations conditioned on the learner’s hidden state. The final linear layer transforms the feature vector into a predicted response. Note that this model is also conditioned on the current query image  $\mathbf{z} = f(\mathbf{x})$  and the corresponding ground truth class label  $y$ .

**Classifier Prediction Model.** Our second recurrent model attempts to provide a more interpretable approximation of human classification. In this case, instead of letting the recurrent model directly predict the probability for each response, it instead attempts to approximate the weights of a linear classifier used internally by the learner. Importantly, the values this classifier takes will depend on the response history of a given learner and will differ at each time-step,

$$\mathbf{w}_t^k, \mathbf{b}_t^k = \psi_{rnn}(\mathbf{z}_{1:t}^k, y_{1:t}^k, r_{1:t}^k, y), \quad (4)$$

$$\phi_{cls\_pred}(\mathbf{x}) = \sigma(\mathbf{w}_t^T \mathbf{z} + \mathbf{b}_t). \quad (5)$$

Unlike the previous direct prediction recurrent model, we now explicitly represent the classification function used by an individual learner. Also, here the features  $\mathbf{z}$  for the query image are not processed by the recurrent network in Eqn. 4. Instead they are evaluated using the much simpler predicted classifier weights in Eqn. 5. This decoupling is advantageous in applications like machine teaching where we have to query the tracing model multiple times at each time-step in order to determine the next image to show learners. Reducing the computation required to perform these queries will result in faster teacher algorithms.

### 3.4 Training Tracing Models

We jointly estimate the parameters of the feature extractor  $f$  and classification functions  $\phi$  for each of the above models using a standard cross entropy loss,

$$\mathcal{L} = -\frac{1}{KT} \sum_{k=1}^K \sum_{t=1}^T \log(\phi(\mathbf{x}_t^k)_r). \quad (6)$$

Here,  $\phi(\mathbf{x}_t^k)_r$  indicates the predicted probability from a model  $\phi$  choosing class  $r$ , for learner  $k$ , at time-step  $t$ . The training objective aims to minimize the

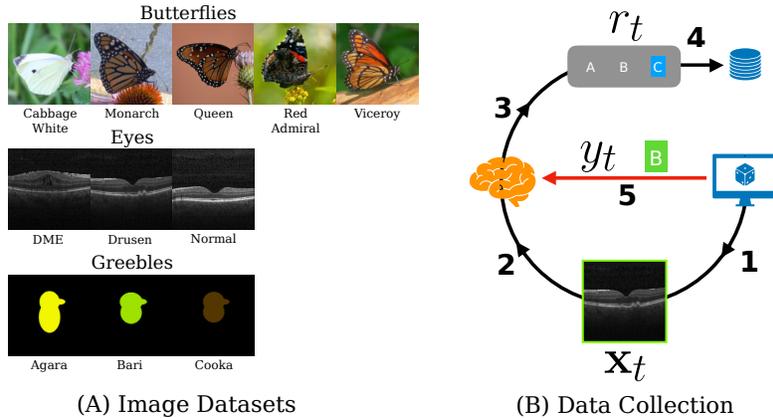


Fig. 2: **(A)** Example images from the three different datasets from our experiments. ‘Butterflies’ contains images of five different species and was originally presented in [27]. ‘Eyes’ contains optical coherence tomography images of the human retina from [20], and features two diseased classes and one normal one. ‘Greebles’ is a synthetic dataset we created where the three object classes vary in terms of shape and color. **(B)** Data collection pipeline. A random image is selected (1), shown to the learner (2), and the learner provides a response (3). Their response is stored (4) and the correct class label is provided to them (5).

difference between the learner responses from our training set and the tracing model outputs.

## 4 Experiments

In this section we evaluate the different proposed models for visual knowledge tracing on data we collected from real human participants<sup>1</sup>.

### 4.1 Datasets

Traditional image classification datasets mostly contain labels produced by annotators familiar with the subject material, e.g. [25,46,24], or they have at least received detailed instructions and examples on how to annotate them, e.g. [39]. As a result, these datasets do not contain annotations from learners engaged in learning a task and are thus not suitable for evaluating visual knowledge tracing. While some work has focused on teaching crowd learners (e.g. [42,27,50]), they often use teaching image sequences that are determined offline and fixed. For our tracing experiments, we require unbiased sequences of images that are randomly selected for each learner. Some of these existing works compare their approaches to a random image selection baseline, but the size of these random

<sup>1</sup> Code and dataset - <https://github.com/nkondapa/VisualKnowledgeTracing>

teaching subsets is insufficient for thorough evaluation of our different tracing approaches, e.g. [27] have random selection data from only  $\sim 40$  participants. Due to these limitations, we collected annotations from human learners for three challenging fine-grained visual classification datasets.

**Image Data.** We selected three different image datasets that cover three distinct domains: artificial data where we have full knowledge of the underlying distribution, medical image data, and images of different wildlife species. The first two datasets in particular are representative of the types of visual identification tasks that many humans are interested in learning.

Our first dataset, ‘Butterflies’, contains images from five different common species of North American butterflies. The ‘Cabbage White’ class is immediately recognizable, ‘Red Admiral’ can be learned relatively easily, and the remaining three are difficult to discriminate. This dataset was originally used in [27] and contains between 386 and 481 images per class, for a total of 2,224 images. Our second dataset, ‘Eyes’, is a three-class subset of a large collection of publicly available images of the human retina from [20]. It contains two diseased classes, ‘Diabetic Macular Edema’ (DME) and ‘Drusen’, and one ‘Normal’ class. We manually selected 200 images from each class. The third dataset is a challenging synthetic one we created called ‘Greebles’. It contains three classes, where the underlying feature space used to generate the images is known by design. The relevant features are the body length and color, but the images also includes some irrelevant variation in the form of the head size and body width. The distinctions between the classes can be subtle, making the task challenging. It contains 1,200 images in total, with an even number per class. Visual examples for each of the datasets are presented in Fig. 2A. Note that the single examples in Fig. 2A do not convey the visual diversity of the datasets.

**Human Data Collection.** For each of the previously described datasets, we collected data from human participants that were engaged in learning the classification task. Each learner was presented with 30 training images and 15 test images. During the training phase they were provided with ground truth feedback indicating the correct class labels. This feedback was not provided in the test phase. For each image, we asked learners to rank the top three classes in order of most likely to least likely to be correct (see Sec. F.1 for further information). The training and testing examples were randomly selected for each learner. An overview of the data collection process for one iteration, for one learner, is shown in Fig. 2B.

The data was collected using a custom built a web application and the participants were recruited through the crowd sourcing app Prolific [1]. In total, we collected data from 150 learners for each dataset, where each individual could only do the task once. The median time spent on the Butterflies task, including training and testing, was  $\sim 11.7$  minutes with a median of 11 correct on the test phase. The corresponding statistics for the Eyes dataset was  $\sim 12.1$  minutes

with a median of 13 correct, and for the Grebbles dataset is was  $\sim 8.8$  minutes with a median of 9 correct.

Our human learners demonstrated learning across all three datasets. In Fig. 3 we plot two histograms for each dataset, the first histogram displays the number of correct responses during the training (i.e. teaching) phase and the second reports the same for the test phase. For all of the datasets, we see a right skew in the histogram of correct responses in the testing phase, indicating that most learners are providing correct answers after training. The skewness for each dataset is -0.45 for Butterflies, -0.81 for Eyes, and -0.37 for Grebbles. The skewness is correlated to the amount of improvement learners showed on a given dataset. The numbers imply that the Grebbles task is the most difficult and the learners demonstrate the least improvement. The Eyes dataset, on the other hand, is clearly the easiest. We provide additional analysis in the supplementary material (Sec. C).

## 4.2 Implementation Details

All models we are considering primarily consist of a feature extractor and a classification function. The feature extractor is implemented as a CNN with eight layers (two convolutional, two max-pool, four linear), and is the same across all models (Sec. B). The feature extractor produces a 16 dimensional embedding for an input image which is then processed by the respective classification function. We evaluated one of our models on higher dimensional feature spaces, but found no impact on performance (Sec. A.7). For the two recurrent models,  $\phi_{direct}$  and  $\phi_{cls\_pred}$ , we use a three layer LSTM-based [16] fully connected network with a hidden dimension of size 128. The output of the LSTM is passed to a small two-layer network that transforms the output into the desired representation (either a response or a classifier (see Sec. 3.3)). We provide a detailed description of the architectures and inputs in the supplementary material (Sec. F.2).

All models are trained using mini-batch stochastic gradient descent with a batch size of 16 using the Adam optimizer [22]. For the feature extractor, we use a learning rate of 1e-5, and we use a learning rate of 1e-3 for each of the different classification functions. We train using a cross entropy loss. Models are trained with early stopping. Training ends when the best validation loss does not improve for 35 epochs. The upper limit on the number of training epochs is 400. Data is split into train (70%), validation (13.3%), and test (16.7%) splits at the learner level, i.e. sequences from the same individual learner cannot be in more than one split. To ensure more robust results, we validate models by re-shuffling the data five times and report the averages across the splits using micro and macro average precision.

## 4.3 Results

**Tracing Human Learners.** We compute the precision and recall curves for each of the visual knowledge tracing models on the held-out learners and report

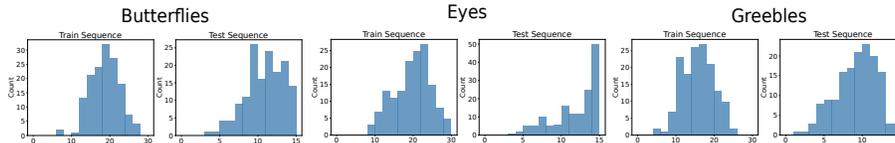


Fig. 3: Human learner performance on our three datasets. For each dataset we provide histograms of learner performance on the respective training and testing sequences. The training results are always worse as it include responses from all time-steps, including when the learner has just started the task and is unfamiliar with the classes. For ‘Greebles’, the worse performance on the test set compared to the other datasets indicates that learners find this task more challenging.

the average precision for each. We also include one additional model for completeness, the ground truth baseline model (GT Label). This baseline does not fit any parameters and simply predicts the corresponding ground truth label of the image for all learners at all time-steps (instead of predicting the learner’s response). Results are summarized in Table 1. Standard deviations are between 0.01 and 0.03 for all models, and do not change the interpretation of the results.

We observe that the recurrent models,  $\phi_{direct}$  and  $\phi_{cls\_pred}$ , out-perform the baselines in tracing the learner on the Butterflies ( $\sim 8\%$ ) and Eyes datasets ( $\sim 6\%$ ). However, on the Greebles dataset, the static tracing model,  $\phi_{static}$ , out-performs the recurrent models by ( $\sim 3\%$ ). Both recurrent models are comparable in terms of performance, indicating the reduction in computation described in Sec. 3.3 does not come with a reduction in performance. The time-sensitive tracing model,  $\phi_{static\_time}$ , is clearly the worst at tracing learners, likely owing to the unrealistic assumptions it makes about them.

Finally we explore the differences between how different models trace human learners. Fig. 4 shows the average probability of predicting an image correctly for each time-step conditioned on each class for the Butterflies dataset. In the top row, we present the training and test-split average accuracy on each of the five classes over time for the human learners. As previously noted, the training sequences contains 30 randomly selected images and the test sequence contains 15 images. We can see that on average, some classes are much easier than others. In the bottom row, we average model-predicted probabilities for  $\sim 50$  images in each class. To produce the probabilities for the recurrent  $\phi_{cls\_pred}$  model, we process the sequential data of the same set of learners in the top panel. The static tracing model,  $\phi_{static}$ , estimates a hyperplane for each class that roughly tracks the average probability of being correct per class. As expected, this model is not capable of capturing any learning behavior. In contrast,  $\phi_{cls\_pred}$  more faithfully traces how the average probability evolves over time. Note, that we process each of the test images independently for  $\phi_{cls\_pred}$ .

Table 1: Performance of different visual knowledge tracing approaches on data from human learners. We observe that our two recurrent based models, the direct response  $\phi_{direct}$  and the classifier prediction  $\phi_{cls\_pred}$ , perform best on the Butterflies and Eyes dataset but are worse on the synthetic Greebles task. Learners found the Greebles task the most challenging, and as a result, there was much less learning occurring compared to the first two datasets. ‘GT Label’ is an additional baseline that uses the corresponding ground truth class label  $y$  as the prediction of the learner’s response  $r$ .

	Greebles				Eyes				Butterflies			
	Train		Test		Train		Test		Train		Test	
	Micro	Macro										
GT Label	0.48	0.51	0.58	0.61	0.56	0.56	0.69	0.69	0.45	0.44	0.50	0.49
$\phi_{static}$	0.63	0.52	0.67	0.58	0.60	0.59	0.67	0.68	<b>0.55</b>	<b>0.53</b>	<b>0.64</b>	<b>0.61</b>
$\phi_{static.time}$	0.52	0.44	0.49	0.40	0.34	0.34	0.33	0.34	0.54	0.52	0.61	0.59
$\phi_{direct}$	0.70	0.59	<b>0.77</b>	0.64	<b>0.66</b>	<b>0.65</b>	<b>0.75</b>	<b>0.74</b>	<b>0.55</b>	<b>0.53</b>	0.60	0.57
$\phi_{cls\_pred}$	<b>0.71</b>	<b>0.62</b>	<b>0.77</b>	<b>0.65</b>	0.65	<b>0.65</b>	0.74	<b>0.74</b>	0.54	0.52	0.60	0.57

#### 4.4 Discussion

**Comparing Models.** We observe that our recurrent models are quite effective at tracing human learner knowledge on visual categorization tasks for datasets where there is a clear learning signal. On the Greebles dataset, which is the most challenging and displays the least amount of learning, we see that the simple static tracing model  $\phi_{static}$  is less prone to over-fitting and is thus marginally better. The time-sensitive tracing model  $\phi_{static.time}$  performs the worst overall. Unlike the recurrent methods, it is unable to share information between time-steps, forcing it to fit a classifier at each time-step with only  $\sim 90$  training points (the number of learners in the training set after we split out the validation and test sets). This makes it extremely prone to over-fitting on the limited training data that is available.

The direct response model  $\phi_{direct}$  and the classifier prediction model  $\phi_{cls\_pred}$  differ in their outputs and their inputs.  $\phi_{direct}$  includes  $\mathbf{z}$ , the representation of the query image  $\mathbf{x}$  to be classified by the learner, as input to the LSTM.  $\phi_{cls\_pred}$  outputs weights of a multi-class linear classifier that is used to classify the representation  $\mathbf{z}$ . Unlike  $\phi_{direct}$ ,  $\mathbf{z}$  is not an input to the LSTM for  $\phi_{cls\_pred}$ . However, both models incorporate the ground truth label  $y$  for  $\mathbf{x}$  as input. These structural differences between the models make little difference to the tracing performance, but the  $\phi_{cls\_pred}$  model is more efficient if it needs perform evaluation multiple times for new query items. In supplementary experiments (A.1), we explore the impact of removing the ground truth class label  $y$  for the query image from the input. We observe that this results in a large decrease in performance for both models, suggesting that this class information is valuable to the model when making future predictions. Also in the supplement, we compare two cognitive models (prototype and exemplar [31]) (A.3), a Transformer model [11] (A.4), a pre-trained ResNet feature extractor (A.5), and input meta-information to the tracing models (A.6). We find that the cognitive, Transformer, and ResNet

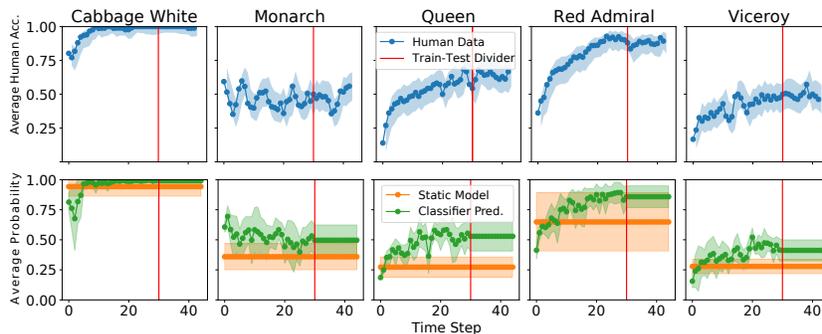


Fig. 4: **(Top)** The smoothed average human learner accuracy for class over time from the Butterflies dataset. The shadowed regions indicate confidence intervals as the number of samples in each time and class bin are not guaranteed to be the same. **(Bottom)** The average probability of having a class correctly predicted by the static  $\phi_{static}$  model (orange) and the recurrent  $\phi_{cls\_pred}$  model (green). At each time-step, for each learner in the test set, the models predict class probabilities for  $\sim 50$  images per class. The probabilities are averaged (solid line) and the shadows indicate one standard deviation. In both rows, the red line indicates the point at which the learners switch from training to testing. After that, the models will continue to produce the same probabilities on the test images for the remaining time-steps as the sampled images do not change.

models do not out-perform the recurrent architectures, but are worth further exploration. Additionally, including meta-information in the input vectors results in a performance increase across all models we tested, suggesting this a promising direction for future work. Finally, we explore the representations learned by the recurrent models (Secs. D, E).

**Limitations and Future Work.** Currently, we train our feature encoder  $f$  from scratch for each task on relatively small amounts of image data. One source of improvement would be to pre-train the feature space so that it better reflects human visual similarity judgements. Such an improved feature space would provide a better starting point for task specific finetuning. Replacing the LSTM with an appropriately designed Transformer network [47] is another change that could result in greater flexibility for the model. Transformers are better able to capture long-term dependencies and have been shown to be useful in knowledge tracing on non-visual educational datasets [33,11]. However, it remains to be seen if this would be valuable for visual knowledge tracing.

Our approaches do not explicitly model memory decay – the phenomenon of memory ‘fading’ due to the passage of time [41,17]. This is likely to be more of a problem when tracing over longer time horizons, e.g. days or weeks, as opposed to the multiple minute long sessions that our learners engage in. Similarly, given the short time durations of our teaching sessions, we assume that no significant

‘feature learning’ is happening for individual learners. Instead we model learners as attending to a subset of the different visual features that are captured in our joint embedding space. In future work, it would be interesting to further explore if these two assumptions are valid.

**Applications.** Successful tracing of human learners has implications for crowd-sourcing annotations, metric learning, and machine teaching. Early detection of poor annotators in crowd sourcing would reduce monetary and time costs in labelling datasets. Additionally, identifying annotators with ‘specialist’ knowledge could allow for targeted crowd sourcing, tailoring to the abilities of each individual annotator. A successful tracing algorithm should be able to predict future performance with increasing confidence as the learners are being trained on the annotation task.

The most impactful domain for a successful tracing algorithm is automated teaching, e.g. teaching medical image interpretation skills [10,2]. Teaching humans is challenging because their knowledge state is unobserved, it changes over time, and the information they provide using current interfaces can be limited. In this work, we show that with a reasonable amount of training data (i.e. data from  $\sim 150$  learners), and only a single response at each time-step, we are able to capture information about the learner’s current knowledge for visual classification tasks. The types of visual knowledge tracing approaches presented could be used in conjunction with machine teaching methods. The more a teaching algorithm knows about the learner, the more effective it can be when selecting examples to present to them.

## 5 Conclusion

More accurate models of human visual classification will lead to improved methods for crowd annotation collection, better techniques for automatically teaching visual knowledge to human learners, and perhaps provide us with insight into how we can build future artificial systems that are more data efficient. To this end, in this work we explored the problem of visual knowledge tracing – the task of predicting the internal, potentially time varying, image classification function used by human learners. To do this, we presented a series of models that range in complexity from basic static linear classifiers all the way to recurrent models that take a learner’s prior response history into account when making predictions about their future behavior. We collected new annotations for three challenging visual classification tasks from humans engaged in a visual learning task in order to benchmark the performance of these different models. Our results show that our recurrent neural network-based models resulted in the most faithful reproductions of unobserved learner predictions on real image datasets. Finally, we outlined limitations of our work and pointed to open questions that require further investigation.

**Acknowledgements:** Thanks to the anonymous reviews for their valuable feedback. This work was in part supported by the Turing 2.0 ‘Enabling Advanced Autonomy’ project funded by the EPSRC and the Alan Turing Institute and also by the Simons Collaboration on the Global Brain.

## References

1. Prolific, [www.prolific.co](http://www.prolific.co), accessed Mar 7 2022
2. Amiri, E., Sha, P., Palmer, E.M.: Training novices to discriminate retinal diseases using perceptual learning. In: Proceedings of the Human Factors and Ergonomics Society Annual Meeting. pp. 1456–1460 (2020)
3. Ashby, F.G., Maddox, W.T.: Human category learning. *Annu. Rev. Psychol.* **56**, 149–178 (2005)
4. Ashby, F.G., Maddox, W.T.: Human category learning 2.0. *Annals of the New York Academy of Sciences* (2011)
5. Ashby, F.G., O’Brien, J.B.: Category learning and multiple memory systems. *Trends in cognitive sciences* **9**(2), 83–89 (2005)
6. Attarian, M., Roads, B.D., Mozer, M.C.: Transforming neural network visual representations to predict human judgments of similarity. arXiv:2010.06512 (2020)
7. Barry, D.N., Love, B.C.: Human learning follows the dynamics of gradient descent. *PsyArXiv* (2021)
8. Biederman, I., Shiffrar, M.M.: Sexing day-old chicks: a case study and expert systems analysis of a difficult perceptual-learning task. *Journal of Experimental Psychology: Learning, memory, and cognition* **13**(4), 640 (1987)
9. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: *ICML* (2020)
10. Cheng, C.T., Chen, C.C., Fu, C.Y., Chaou, C.H., Wu, Y.T., Hsu, C.P., Chang, C.C., Chung, I., Hsieh, C.H., Hsieh, M.J., et al.: Artificial intelligence-based education assists medical students’ interpretation of hip fracture. *Insights into Imaging* **11**(1), 1–8 (2020)
11. Choi, Y., Lee, Y., Cho, J., Baek, J., Kim, B., Cha, Y., Shin, D., Bae, C., Heo, J.: Towards an appropriate query, key, and value computation for knowledge tracing. In: *Proceedings of the Seventh ACM Conference on Learning@ Scale* (2020)
12. Cole, E., Yang, X., Wilber, K., Mac Aodha, O., Belongie, S.: When does contrastive visual representation learning work? In: *CVPR* (2022)
13. Corbett, A.T., Anderson, J.R.: Knowledge tracing: Modeling the acquisition of procedural knowledge. *User modeling and user-adapted interaction* **4**(4), 253–278 (1994)
14. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: *CVPR* (2020)
15. Hebart, M.N., Zheng, C.Y., Pereira, F., Baker, C.I.: Revealing the multidimensional mental representations of natural objects underlying human similarity judgements. *Nature human behaviour* **4**(11), 1173–1185 (2020)
16. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural computation* **9**(8), 1735–1780 (1997)
17. Hunziker, A., Chen, Y., Mac Aodha, O., Gomez Rodriguez, M., Krause, A., Perona, P., Yue, Y., Singla, A.: Teaching multiple concepts to a forgetful learner. *NeurIPS* (2019)
18. Johns, E., Mac Aodha, O., Brostow, G.J.: Becoming the expert-interactive multi-class machine teaching. In: *CVPR*. pp. 2616–2624 (2015)
19. Kaya, M., Bilge, H.Ş.: Deep metric learning: A survey. *Symmetry* (2019)
20. Kermany, D.S., Goldbaum, M., Cai, W., Valentim, C.C., Liang, H., Baxter, S.L., McKeown, A., Yang, G., Wu, X., Yan, F., et al.: Identifying medical diagnoses and treatable diseases by image-based deep learning. *Cell* **172**(5), 1122–1131 (2018)

21. Kim, K.H., Mac Aodha, O., Perona, P.: Context embedding networks. In: CVPR. pp. 8679–8687 (2018)
22. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv:1412.6980 (2014)
23. Kruskal, J.B., Wish, M.: Multidimensional scaling. No. 11, Sage (1978)
24. Kuznetsova, A., Rom, H., Alldrin, N., Uijlings, J., Krasin, I., Pont-Tuset, J., Kamali, S., Popov, S., Mallocci, M., Kolesnikov, A., et al.: The open images dataset v4. IJCV (2020)
25. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV (2014)
26. Liu, Q., Shen, S., Huang, Z., Chen, E., Zheng, Y.: A survey of knowledge tracing. arXiv:2105.15106 (2021)
27. Mac Aodha, O., Su, S., Chen, Y., Perona, P., Yue, Y.: Teaching categories to human learners with visual explanations. In: CVPR. pp. 3820–3828 (2018)
28. Markman, E.M.: Categorization and naming in children: Problems of induction. mit Press (1989)
29. Musgrave, K., Belongie, S., Lim, S.N.: A metric learning reality check. In: ECCV. pp. 681–699 (2020)
30. Nigam, I., Tokmakov, P., Ramanan, D.: Towards latent attribute discovery from triplet similarities. In: ICCV. pp. 402–410 (2019)
31. Nosofsky, R.M., Meagher, B., Kumar, P.: Contrasting exemplar and prototype models in a natural-science category domain. In: CogSci (2020)
32. Nosofsky, R.M., Sanders, C.A., Meagher, B.J., Douglas, B.J.: Toward the development of a feature-space representation for a complex natural category domain. Behavior Research Methods **50**(2), 530–556 (2018)
33. Pandey, S., Karypis, G.: A self-attentive model for knowledge tracing. arXiv:1907.06837 (2019)
34. Piech, C., Spencer, J., Huang, J., Ganguli, S., Sahami, M., Guibas, L., Sohl-Dickstein, J.: Deep knowledge tracing. NeurIPS (2015)
35. Pu, S., Yudelson, M., Ou, L., Huang, Y.: Deep knowledge tracing with transformers. In: International Conference on Artificial Intelligence in Education. pp. 252–256 (2020)
36. Roads, B.D., Love, B.C.: Enriching imagenet with human similarity judgments and psychological embeddings. In: CVPR (2021)
37. Roads, B.D., Mozer, M.C.: Predicting the ease of human category learning using radial basis function networks. Neural Computation (2021)
38. Roth, K., Milbich, T., Sinha, S., Gupta, P., Ommer, B., Cohen, J.P.: Revisiting training strategies and generalization performance in deep metric learning. In: ICML. pp. 8242–8252 (2020)
39. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. IJCV (2015)
40. Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: A unified embedding for face recognition and clustering. In: CVPR. pp. 815–823 (2015)
41. Settles, B., Meeder, B.: A trainable spaced repetition model for language learning. In: ACL (2016)
42. Singla, A., Bogunovic, I., Bartók, G., Karbasi, A., Krause, A.: Near-optimally teaching the crowd to classify. In: ICML. pp. 154–162 (2014)
43. Tamuz, O., Liu, C., Belongie, S.J., Shamir, O., Kalai, A.: Adaptively learning the crowd kernel. In: ICML (2011)

44. Tan, R., Vasileva, M.I., Saenko, K., Plummer, B.A.: Learning similarity conditions without explicit supervision. In: ICCV. pp. 10373–10382 (2019)
45. Van Der Maaten, L., Weinberger, K.: Stochastic triplet embedding. In: International Workshop on Machine Learning for Signal Processing. pp. 1–6 (2012)
46. Van Horn, G., Branson, S., Farrell, R., Haber, S., Barry, J., Ipeirotis, P., Perona, P., Belongie, S.: Building a bird recognition app and large scale dataset with citizen scientists: The fine print in fine-grained dataset collection. In: CVPR (2015)
47. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. NeurIPS (2017)
48. Veit, A., Belongie, S., Karaletsos, T.: Conditional similarity networks. In: CVPR. pp. 830–838 (2017)
49. Wang, J., Song, Y., Leung, T., Rosenberg, C., Wang, J., Philbin, J., Chen, B., Wu, Y.: Learning fine-grained image similarity with deep ranking. In: CVPR. pp. 1386–1393 (2014)
50. Wang, P., Nagrecha, K., Vasconcelos, N.: Gradient-based algorithms for machine teaching. In: CVPR. pp. 1387–1396 (2021)
51. Wang, P., Vasconcelos, N.: A machine teaching framework for scalable recognition. In: ICCV. pp. 4945–4954 (2021)
52. Welinder, P., Branson, S., Perona, P., Belongie, S.: The multidimensional wisdom of crowds. NeurIPS (2010)
53. Wilber, M., Kwak, I., Belongie, S.: Cost-effective hits for relative similarity comparisons. In: Proceedings of the AAAI Conference on Human Computation and Crowdsourcing. vol. 2, pp. 227–233 (2014)
54. Wu, Z., Xiong, Y., Yu, S.X., Lin, D.: Unsupervised feature learning via non-parametric instance discrimination. In: CVPR (2018)
55. Xu, F., Tenenbaum, J.B.: Word learning as bayesian inference. Psychological review (2007)
56. Zeithamova, D.: Category learning systems. The University of Texas at Austin (2008)
57. Zhu, X., Singla, A., Zilles, S., Rafferty, A.N.: An overview of machine teaching. arXiv:1801.05927 (2018)