

# Learning to Defer with Limited Expert Predictions

Patrick Hemmer\*, Lukas Thede\*, Michael Vössing, Johannes Jakubik, Niklas Kühl

Karlsruhe Institute of Technology

{patrick.hemmer, michael.voessing, johannes.jakubik, niklas.kuehl}@kit.edu, lukas.thede@alumni.kit.edu

## Abstract

Recent research suggests that combining AI models with a human expert can exceed the performance of either alone. The combination of their capabilities is often realized by *learning to defer* algorithms that enable the AI to learn to decide whether to make a prediction for a particular instance or defer it to the human expert. However, to accurately learn which instances should be deferred to the human expert, a large number of expert predictions that accurately reflect the expert’s capabilities are required—in addition to the ground truth labels needed to train the AI. This requirement shared by many learning to defer algorithms hinders their adoption in scenarios where the responsible expert regularly changes or where acquiring a sufficient number of expert predictions is costly. In this paper, we propose a three-step approach to reduce the number of expert predictions required to train learning to defer algorithms. It encompasses (1) the training of an embedding model with ground truth labels to generate feature representations that serve as a basis for (2) the training of an expertise predictor model to approximate the expert’s capabilities. (3) The expertise predictor generates artificial expert predictions for instances not yet labeled by the expert, which are required by the learning to defer algorithms. We evaluate our approach on two public datasets. One with “synthetically” generated human experts and another from the medical domain containing real-world radiologists’ predictions. Our experiments show that the approach allows the training of various learning to defer algorithms with a minimal number of human expert predictions. Furthermore, we demonstrate that even a small number of expert predictions per class is sufficient for these algorithms to exceed the performance the AI and the human expert can achieve individually.

## Introduction

Recent advances in the field of artificial intelligence (AI) have improved the accuracy of AI models to a point where they exceed the performance of human experts on an increasing number of tasks (Farrell 2021; Rajpurkar et al. 2017; Topol 2019). However, various application domains remain where AI models can not consistently outperform human experts (Cremer 2021; Raghu et al. 2019; Kühl et al. 2022). Motivated by this observation, prior work has started

to explore how the capabilities of human experts and AI models can be combined. In this context, *learning to defer* algorithms have demonstrated promising results by enabling the AI model not to make a prediction for all instances but rather to learn to defer a subset of them to the human expert—based on the AI model’s and the expert’s capabilities (Mozannar and Sontag 2020; Okati, De, and Rodriguez 2021; Raghu et al. 2019). Non-identical capabilities of the human expert and the AI can originate, e.g., from limited training data, limited model capacity, or the availability of side information only accessible to the human expert (Hemmer et al. 2022b; Charusaie et al. 2022). By compensating for the weaknesses of both the human expert and the AI through effective instance allocation, learning to defer algorithms can achieve a higher performance level than either the human or the AI model can achieve individually (Raghu et al. 2019; Wilder, Horvitz, and Kamar 2020).

However, large amounts of labeled data are required to train learning to defer algorithms. In fact, they require not only ground truth labels to train the AI model but also additional human expert predictions that accurately reflect the expert’s capabilities (Leitão et al. 2022). In this context, any new human expert aiming to collaborate with the AI model would have to provide expert predictions so that the learning to defer algorithms can understand the expert’s individual capabilities. It can become particularly costly in application domains with frequently changing human experts or where data labeling is time-consuming and requires the knowledge of highly trained specialists, e.g., in medicine. Due to this impediment, the usage of learning to defer algorithms can become infeasible in these cases—despite their potential for decision-making tasks with high costs of errors.

To reduce the number of human expert predictions required for the training of learning to defer algorithms, we propose a novel three-step approach capable of generating artificial expert predictions from only a small number of human expert predictions accurately reflecting the expert’s capabilities. (1) We train an embedding model with ground truth labels to extract feature representations which (2) serve as a basis for training an expertise predictor model that learns to approximate the human expert’s capabilities with only a small number of human expert predictions. It is trained on the available expert predictions while simultaneously leveraging instances for which no human expert pre-

\*These authors contributed equally.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

dictions are available by drawing upon an interchangeable semi-supervised learning component. (3) The expertise predictor generates artificial expert predictions for instances not labeled by the human expert. Both human and artificial expert predictions combined can then be used to train a variety of learning to defer algorithms.

We empirically demonstrate the efficiency of our approach on the CIFAR-100 dataset (Krizhevsky 2009) using “synthetically” generated human expert predictions and on the NIH chest X-ray dataset (Majkowska et al. 2020; Wang et al. 2017) that provides real-world individual radiologists’ predictions. In addition, it also contains adjudicated ground truth labels provided by a panel of radiologists that serve as a “gold standard”. For example, with our proposed approach, six expert predictions per class suffice to enable all evaluated learning to defer algorithms to outperform both the AI model and the radiologist compared to either conducting the task alone. On average, this corresponds to 98.96% of the achieved accuracy compared to the training with a complete set of human expert predictions.

To summarize, the contribution of our work is threefold: First, we propose a novel approach that learns the capabilities of individual human experts from a small set of human expert predictions. Second, we demonstrate our approach’s ability to enable state-of-the-art learning to defer algorithms to function with only a minimum of required human expert predictions while maintaining their desirable properties of achieving superior team performance. Third, we show its real-world applicability in the context of the medical domain with a dataset providing individual radiologists’ predictions together with high-quality ground truth labels.

## Related Work

Over the last years, research has investigated how human and AI capabilities can be combined to achieve superior team performance compared to both performing tasks individually (Madras, Pitassi, and Zemel 2018; Mozannar and Sontag 2020; Okati, De, and Rodriguez 2021; Raghu et al. 2019; Wilder, Horvitz, and Kamar 2020). In this context, *learning to defer* algorithms that enable the AI model to decide whether to make a prediction for a particular instance or defer it to a human expert have demonstrated promising results (Leitão et al. 2022). Raghu et al. (2019) propose to estimate the prediction confidence of the classifier and the human expert on an instance basis. Instances are deferred when the human expert is estimated to be more confident than the classifier. While this approach optimizes the classifier in isolation, Wilder, Horvitz, and Kamar (2020) propose the idea of optimizing for team performance by training the classifier to specifically complement the human expert’s capabilities. In this context, several approaches jointly train the classifier together with a deferral system (Madras, Pitassi, and Zemel 2018; Okati, De, and Rodriguez 2021; Wilder, Horvitz, and Kamar 2020). Further work utilizes objective functions with theoretical guarantees for regression (De et al. 2020) and classification tasks (De et al. 2021). Moreover, Mozannar and Sontag (2020) propose a consistent surrogate loss function inspired by cost-sensitive learning. Raman and Yee (2021) build upon their work by indi-

vidualizing the approach to a particular human expert via fine-tuning. Further research studies the assignment of instances in the context of bandit feedback (Gao et al. 2021) or suggests learning to allocate instances within teams consisting of an AI model and multiple human experts (Hemmer et al. 2022a; Keswani, Lease, and Kenthapadi 2021).

These approaches have in common that they require human expert predictions to recognize an expert’s strengths and weaknesses when deciding whether an instance shall be deferred. Since the capabilities of human experts can vary in different areas of the feature space, e.g., due to different levels of knowledge or biases (Lampert, Stumpf, and Gançarski 2016), this requires the availability of individual predictions for each expert working with the AI model, which impedes the applicability of learning to defer algorithms, especially in application domains with frequently changing human experts or high acquisition costs for expert predictions.

Our approach addresses this limitation through the generation of artificial expert predictions only from a small number of predictions provided by an individual human expert. Concurrent work designs an active learning scheme to reduce the number of required predictions (Charusaie et al. 2022). On the contrary, our approach does not require iteratively identifying instances for which human expert predictions are queried. Instead, given a small number of human expert predictions, it learns to infer artificial ones for unlabeled instances in the training dataset that can then be used for the training of learning to defer algorithms.

Besides their benefits with regard to the overall team performance, research has also investigated the effect of learning to defer algorithms in the context of fairness, as AI models are known to be capable of generating biased predictions against protected groups (Buolamwini and Gebru 2018). Particularly, when AI models can abstain from uncertain predictions, e.g., in selective classification, this can result in amplifying existing biases (Jones et al. 2021). A similar phenomenon can also be observed in the behavior of humans since their decisions are neither free from prejudices (Tversky and Kahneman 1974). To prevent possible amplifications of AI model and human biases, Madras, Pitassi, and Zemel (2018) propose a learning to defer algorithm that considers both overall team performance and fairness by introducing a regularized loss function combining the error rate with a fairness metric. They find that their approach can result in a system that is more accurate and less biased. Moreover, Keswani, Lease, and Kenthapadi (2021) propose to ensure that their learning to defer algorithm produces unbiased predictions by balancing error rates for all protected groups and drawing upon the minimax Pareto fairness concept.

Lastly, another related stream of work focuses on efficiently combining predictions of multiple human annotators to derive high-quality labels, e.g., from crowdsourced predictions (Branson, Van Horn, and Perona 2017; Guan et al. 2018; Liao, Kar, and Fidler 2021; Welinder and Perona 2010). Whereas these approaches focus on ground truth label quality using a larger worker pool, we are interested in learning the capabilities of an individual human expert only from a minimal number of expert predictions.

## Problem Formulation

In this section, we consider the learning to defer setting for classification tasks. Given an instance  $x_i \in \mathcal{X}$ , the goal is to learn to predict its ground truth label  $y_i \in \mathcal{Y} = \{1, \dots, k\}$  with  $k$  denoting the number of classes. The indices of the respective instances are denoted as  $i \in N = \{1, \dots, n\}$ . In addition to  $y_i$ , we assume to have access to the prediction of a human expert  $h_i \in \mathcal{Y}$ . The expert’s predictions can deviate from the ground truth as individual human experts’ capabilities vary in different areas of the feature space. Possible reasons can be different levels of background knowledge, experience, or biases (Lampert, Stumpf, and Gançarski 2016). Moreover, ground truth labels are often curated in various application domains through multiple human experts to ensure high label quality (Guan et al. 2018). For example, in the medical domain, a panel of radiologists is often involved in determining the prevalence of a certain clinical finding (Majkowska et al. 2020; Wang et al. 2017). In this context, these experts most likely differ from the ones that actually cooperate with the deployed AI model in practice.

As it is often accompanied by high costs to acquire a large number of individual human expert predictions, particularly when the expert that collaborates with the AI model changes frequently, we assume individual human expert predictions to be available only for a small subset of instances  $L \subseteq N$  with  $|L| \ll |N|$ . For the remaining subset  $U = N \setminus L$ , no individual human expert predictions are available. Hence, we define the dataset with ground truth labels and human expert predictions as  $D^l = \{(x_i, y_i, h_i)\}_{i=1}^l$  with  $l = |L|$ . Similarly, we define the dataset without such predictions but with ground truth labels as  $D^u = \{(x_j, y_j)\}_{j=1}^u$  with  $u = |U|$  such that  $D^l \cap D^u = \emptyset$ .

The general idea of learning to defer algorithms is to train a deferral system  $A : \mathcal{X} \rightarrow \{0, 1\}$  that decides whether to defer an instance to the human expert, with  $A(x_i) = 1$  denoting the deferral decision and  $A(x_i) = 0$  referring to make a prediction using a classifier  $F : \mathcal{X} \rightarrow \mathcal{Y}$ . However, jointly optimizing the deferral system  $A$  and the classifier  $F$  requires the availability of a human expert’s prediction  $h_i$  and ground truth label  $y_i$  for each instance  $x_i$  which is only met for the small subset  $D^l$ . With our proposed approach, we learn to generate an artificial expert prediction  $\hat{h}_j \in \mathcal{Y}$  for each instance in  $D^u$  that accurately reflects the expert’s capabilities and enables the training of  $A$  and  $F$  using all instances from  $D = D^l \cup D^u$ . Keeping the approach independent in the form of a prior step makes it applicable to any learning to defer algorithm requiring both human expert predictions and ground truth labels for the training.

## Approach

In this section, we describe our proposed approach for learning the capabilities of an individual human expert to generate artificial expert predictions. It consists of three steps: (1) Training of an *embedding model* with ground truth labels, which is used to extract feature representations. (2) They serve as input for the training of an *expertise predictor model* to approximate the human expert’s capabilities. (3) The expertise predictor model generates *artificial expert*

*predictions* for the instances not labeled by the human expert. Both human and artificial expert predictions can then be used for the training of learning to defer algorithms.

## Training of Embedding Model

The first component of the approach embeds each instance into a feature representation using a feature extractor. We denote the model as  $\Phi_{emb} : \mathcal{X} \rightarrow \mathbb{R}^d$  with  $d$  being the length of the feature vectors. The embedding model is trained together with a classifier  $\Omega : \mathbb{R}^d \rightarrow \mathcal{Y}$ . The classifier uses the feature representations of the embedding model to learn the classification task based on all data instances of  $D$  and their respective ground truth labels. The embedding model is trained using the loss:

$$\mathcal{L} = \frac{1}{n} \sum_{i=1}^n H(y_i, \Omega(\Phi_{emb}(x_i))) \quad (1)$$

with  $H$  being the cross-entropy between a ground truth label and the classifier’s prediction. Using the trained embedding model  $\Phi_{emb}$ , we can transpose each instance  $x_i$  into a feature vector  $f_i \in \mathbb{R}^d \forall i \in N$  thereby constricting the hypothesis space for the subsequent task of approximating the human expert’s capabilities. This is an approach commonly used in few-shot learning (Koch et al. 2015; Scott, Ridgeway, and Mozer 2018; Snell, Swersky, and Zemel 2017; Tian et al. 2020; Vinyals et al. 2016) to transfer prior knowledge from the task of predicting ground truth labels—to the novel task of generating artificial expert predictions.

## Training of Expertise Predictor Model

The second component is the expertise predictor model. It learns to approximate the human expert’s capabilities reflected in the expert’s labeling behavior. For a multi-class classification task with  $k > 2$ , there are  $k - 1$  possibilities for an incorrect prediction. In case a human expert does not systematically confound certain classes when making incorrect predictions, identifying the exact class of an incorrect prediction can make the task of learning the expert’s labeling behavior unnecessarily complex. Therefore, we propose to reduce the complexity to a binary task—learning whether the expert makes a correct or incorrect prediction. We use the following function to translate a multi-class expert prediction  $h_i$  into a binary expert prediction  $h_i^{bin}$ :

$$\Psi(h_i, y_i) = \begin{cases} 1, & \text{if } h_i = y_i \\ 0, & \text{otherwise} \end{cases}, \quad \forall i \in L. \quad (2)$$

The resulting binary expert predictions  $h_i^{bin}$  are then used together with the feature representations  $f_i \forall i \in L$  as input to train the expertise predictor model  $\Phi_{ex} : \mathbb{R}^d \rightarrow \{0, 1\}$ .

In order to leverage both labeled and unlabeled instances for this task, we propose to incorporate an interchangeable semi-supervised learning component for training the expertise predictor model. We instantiate this component with two state-of-the-art semi-supervised learning algorithms to demonstrate its flexibility. In detail, we use FixMatch and CoMatch as both have demonstrated that they can outperform other semi-supervised approaches (Sohn et al. 2020;

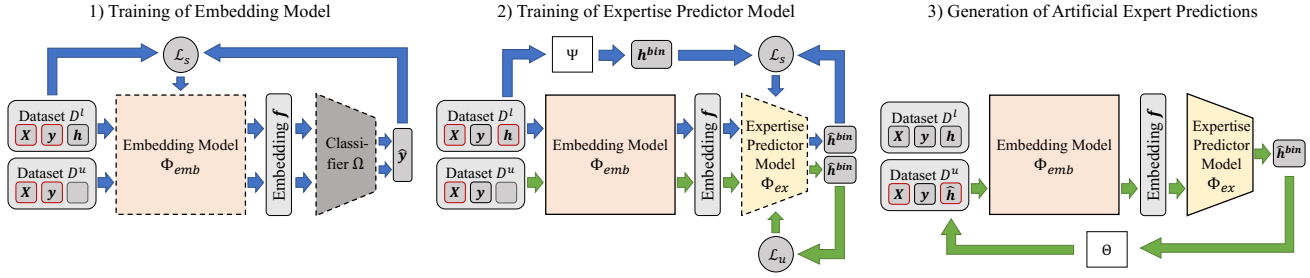


Figure 1: The proposed approach for learning the capabilities of a human expert. We use blue arrows to visualize the path of supervised instances, while green arrows represent unsupervised instances. Red borders denote the relevant components of the dataset  $D^l$  and  $D^u$  that are used at the respective step.

Li, Xiong, and Hoi 2021). We denote these two instantiations of the proposed approach as *Embedding-FixMatch* and *Embedding-CoMatch* respectively. Both algorithms combine methods from consistency regularization and entropy minimization. The loss function of both algorithms is displayed in Equation (3). Whereas FixMatch’s loss function consists of a supervised and an unsupervised loss term, CoMatch includes an additional contrastive loss term:

$$\begin{aligned}
 \mathcal{L} = & \underbrace{\frac{1}{l} \sum_{i=1}^l H(h_i^{bin}, \Phi_{ex}(\Phi_{emb}(\text{Aug}_w(x_i))))}_{\mathcal{L}_s} + \\
 & \lambda_u * \underbrace{\frac{1}{u} \sum_{j=1}^u \mathbb{1}(\max(q_j) \geq \tau) H(\hat{q}_j, \Phi_{ex}(\Phi_{emb}(\text{Aug}_s(x_j))))}_{\mathcal{L}_u} + \\
 & \underbrace{\lambda_c * \frac{1}{u} \sum_{j=1}^u H(\hat{W}_j^q, \hat{W}_j^z)}_{\mathcal{L}_c \text{ (CoMatch only)}}.
 \end{aligned} \tag{3}$$

In both cases, the first part of the loss  $\mathcal{L}_s$  represents the cross-entropy  $H$  between a binary expert prediction  $h_i^{bin}$  and the model’s prediction for a weakly augmented input  $\text{Aug}_w(x_i)$  with  $i \in L$ . For the second part  $\mathcal{L}_u$ , the predicted class distribution of the model  $q_j = \Phi_{ex}(\Phi_{emb}(\text{Aug}_w(x_j)))$  is computed given an unlabeled weakly augmented image  $x_j$  with  $j \in U$ . FixMatch calculates  $\hat{q}_j = \arg \max(q_j)$  to obtain a hard pseudo-label. The cross-entropy compares the pseudo-label against the model prediction for a strong augmentation of the same instance  $\text{Aug}_s(x_j)$ , with  $\tau$  referring to a threshold above which the pseudo-label is retained. In contrast, CoMatch does not convert  $q_j$  to a hard pseudo-label  $\hat{q}_j$  but performs memory-smoothed pseudo-labeling to obtain a pseudo-label. For the third part  $\mathcal{L}_c$ , CoMatch constructs a normalized pseudo-label graph  $\hat{W}_j^q$  and a normalized embedding graph  $\hat{W}_j^z$ . The first graph measures the similarity of instances in the label space. It is used as the target for the training of an embedding graph that quantifies the similarity of strongly-augmented instances in the embedding space. For the embedding graph, CoMatch requires the expertise

predictor model to have an additional projection head transforming the feature representation  $f_j$  into a low-dimensional embedding via two fully connected layers (Li, Xiong, and Hoi 2021). Both graphs are used to calculate a contrastive loss minimizing the cross-entropy between them for all unlabeled instances  $j \in U$ . Lastly, the second and third components of the loss are controlled by the fixed hyperparameters  $\lambda_u$  and  $\lambda_c$  respectively. For further technical details, we refer to Sohn et al. (2020) for FixMatch and Li, Xiong, and Hoi (2021) for CoMatch.

### Generation of Artificial Expert Predictions

After the training of both components, we obtain an artificial binary expert prediction  $\hat{h}_j^{bin}$  by transposing  $x_j$  into a feature representation  $f_j$  using the embedding model together with the expertise predictor model  $\hat{h}_j^{bin} = \arg \max \Phi_{ex}(\Phi_{emb}(x_j)) \forall j \in U$ . Finally, we translate all artificial binary expert predictions  $\hat{h}_j^{bin}$  into multi-class artificial expert predictions  $\hat{h}_j \forall j \in U$  in the following way:

$$\Theta(\hat{h}_j^{bin}, y_j) = \begin{cases} y_j, & \text{if } \hat{h}_j^{bin} = 1 \\ p, & \text{otherwise} \end{cases}, p \sim \mathcal{U}(\{1, \dots, k\} \setminus y_j). \tag{4}$$

Figure 1 provides an overview of the proposed approach with the FixMatch loss terms as an exemplary instantiation of the semi-supervised learning component. We formalize the approach in Algorithm A1 in the Appendix.

## Experiments

We present the experimental evaluation of our proposed approach on CIFAR-100 (Krizhevsky 2009) with synthetic human expert predictions and on the NIH dataset (Majkowska et al. 2020; Wang et al. 2017) with real-world radiologists’ predictions. Further implementation details and results are presented in the Appendix, which we provide together with the code at <https://github.com/ptrckhmmr/learning-to-defer-with-limited-expert-predictions>.

### Experimental Setup

The performance of the approach is evaluated in two steps. First, we evaluate how well the artificial binary expert pre-

	$l$	40	80	120	200	400	1000	5000
$H_{60}$	FixMatch	68.81 ( $\pm 0.51$ )	71.72 ( $\pm 1.73$ )	74.93 ( $\pm 2.42$ )	75.42 ( $\pm 2.00$ )	78.57 ( $\pm 1.90$ )	79.62 ( $\pm 1.72$ )	82.67 ( $\pm 0.44$ )
	CoMatch	75.00 ( $\pm 2.49$ )	75.97 ( $\pm 1.20$ )	79.17 ( $\pm 2.32$ )	81.46 ( $\pm 1.36$ )	84.61 ( $\pm 0.86$ )	85.33 ( $\pm 0.67$ )	87.62 ( $\pm 0.51$ )
	Embedding-NN	75.74 ( $\pm 1.57$ )	77.57 ( $\pm 1.73$ )	79.32 ( $\pm 1.65$ )	79.46 ( $\pm 2.08$ )	80.78 ( $\pm 0.49$ )	81.47 ( $\pm 0.24$ )	83.77 ( $\pm 0.13$ )
	Embedding-SVM	74.49 ( $\pm 1.48$ )	79.49 ( $\pm 1.49$ )	81.61 ( $\pm 1.31$ )	81.81 ( $\pm 1.61$ )	84.69 ( $\pm 1.15$ )	<b>86.41</b> ( $\pm 0.37$ )	<b>89.16</b> ( $\pm 0.33$ )
	<b>Embedding-FixMatch</b>	<b>79.69</b> ( $\pm 0.02$ )	80.52 ( $\pm 0.04$ )	<b>84.06</b> ( $\pm 0.36$ )	<b>83.82</b> ( $\pm 1.15$ )	<b>85.50</b> ( $\pm 0.07$ )	86.04 ( $\pm 0.34$ )	86.95 ( $\pm 0.15$ )
	<b>Embedding-CoMatch</b>	79.65 ( $\pm 0.02$ )	<b>80.88</b> ( $\pm 0.02$ )	84.01 ( $\pm 0.27$ )	83.74 ( $\pm 1.17$ )	85.38 ( $\pm 0.14$ )	85.95 ( $\pm 0.36$ )	86.87 ( $\pm 0.05$ )
$H_{90}$	FixMatch	89.51 ( $\pm 1.16$ )	90.23 ( $\pm 0.72$ )	91.44 ( $\pm 0.92$ )	93.21 ( $\pm 0.63$ )	93.88 ( $\pm 0.26$ )	94.65 ( $\pm 0.46$ )	95.84 ( $\pm 0.06$ )
	CoMatch	87.24 ( $\pm 2.02$ )	90.09 ( $\pm 0.75$ )	91.72 ( $\pm 0.85$ )	92.82 ( $\pm 0.45$ )	93.69 ( $\pm 0.64$ )	95.15 ( $\pm 0.32$ )	<b>96.62</b> ( $\pm 0.12$ )
	Embedding-NN	92.10 ( $\pm 0.82$ )	92.35 ( $\pm 0.67$ )	92.70 ( $\pm 0.65$ )	92.86 ( $\pm 0.79$ )	93.48 ( $\pm 0.22$ )	93.63 ( $\pm 0.22$ )	93.68 ( $\pm 0.20$ )
	Embedding-SVM	92.08 ( $\pm 0.14$ )	93.00 ( $\pm 0.62$ )	93.46 ( $\pm 0.51$ )	93.78 ( $\pm 0.52$ )	94.21 ( $\pm 0.35$ )	95.53 ( $\pm 0.29$ )	96.38 ( $\pm 0.24$ )
	<b>Embedding-FixMatch</b>	92.44 ( $\pm 0.01$ )	93.82 ( $\pm 0.17$ )	95.34 ( $\pm 0.02$ )	95.02 ( $\pm 0.02$ )	<b>95.70</b> ( $\pm 0.02$ )	<b>95.81</b> ( $\pm 0.02$ )	96.25 ( $\pm 0.02$ )
	<b>Embedding-CoMatch</b>	<b>94.75</b> ( $\pm 0.02$ )	<b>95.10</b> ( $\pm 0.01$ )	<b>95.46</b> ( $\pm 0.01$ )	<b>95.11</b> ( $\pm 0.01$ )	95.58 ( $\pm 0.01$ )	95.63 ( $\pm 0.01$ )	96.14 ( $\pm 0.02$ )

Table 1:  $F_{0.5}$ -score (mean and standard deviation) of the artificial binary expert predictions generated by the expertise predictor model using different numbers of  $l$  available human expert predictions for the synthetic experts  $H_{60}$  and  $H_{90}$  on CIFAR-100.

dictions generated by the expertise predictor model represent the true binary predictions of the human experts. Second, we assess the performance difference of multiple learning to defer algorithms trained with both human and artificial expert predictions compared to the training with the entire set of human expert predictions. Both steps are evaluated with different numbers of instances for which  $l = |D^l|$  human expert predictions are provided to learn the expert’s capabilities with  $l = m * k$  and  $m \in \{2, 4, 6, 10, 20, 50, 250\}$ .

**Evaluation of Artificial Expert Predictions.** For each  $l$ , we draw instances randomly from the training data while ensuring balanced class proportions. We use the  $F_{0.5}$ -score as the metric to evaluate the performance of the artificial binary expert predictions, as the transformation according to Equation (2) can result in an imbalanced class distribution with regard to whether the human expert provides a correct or incorrect prediction.

**Evaluation of Learning to Defer Algorithms.** As our approach can be applied independently as a prior step to learning to defer algorithms, we demonstrate its applicability to the approaches of Mozannar and Sontag (2020), Raghu et al. (2019), and Okati, De, and Rodriguez (2021). Mozannar and Sontag (2020) jointly train the classification and deferral decision with a consistent surrogate loss function by adding an additional deferral class to the classification problem. Raghu et al. (2019) train a classifier and a human error model separately. Each instance is allocated by comparing both models’ confidences. Okati, De, and Rodriguez (2021) propose to assign each instance via a negative log-likelihood loss calculated on an instance basis for both the classifier and expert. The classifier is then trained to complement the human expert by minimizing the negative log-likelihood loss only on the instances assigned to the classifier.

As both the CIFAR-100 and the NIH datasets are balanced across classes, we use the system accuracy to evaluate the performance of the learning to defer algorithms trained on both the available human expert predictions  $l$  and the artificial expert predictions. For each  $l$ , we report their performances relative to the performance of the respective approach trained on a complete set of expert predictions  $|D^l| = |D|$  and  $D^u = \emptyset$  (*Complete Expert Pre-*

*dictions*), which serves as upper boundary. Furthermore, we consider the performance of the human expert and the AI alone as two lower boundaries (*Human Expert Alone* and *Classifier Alone*) that have to be exceeded to demonstrate our approach’s applicability to learning to defer algorithms.

**Baselines.** We benchmark our proposed embedding semi-supervised learning (*Embedding-SSL*) approaches *Embedding-FixMatch* and *Embedding-CoMatch* against the following baselines: First, we use the algorithms *FixMatch* (Sohn et al. 2020) and *CoMatch* (Li, Xiong, and Hoi 2021) without an embedding model as semi-supervised learning baselines (*SSL*). Second, we consider two embedding supervised learning (*Embedding-SL*) baselines that use an embedding model combined with a supervised classifier as an expertise predictor model. Specifically, we implement an *Embedding-NN* baseline, which uses a single-layer neural network as an expertise predictor model, and an *Embedding-SVM* baseline, where the expertise predictor model consists of a support vector classifier.

## CIFAR-100

**Dataset and Expert Generation.** The CIFAR-100 dataset (Krizhevsky 2009) consists of 60,000 images from 100 subclasses which are grouped into 20 superclasses of equal size. We use the 20 superclasses for the classification task and the 100 subclasses for the generation of synthetic human experts. This allows us to model the expert’s capabilities on a fine-granular level. To model a synthetic expert, by design, we draw a subset of the 100 subclasses as the expert’s strengths, i.e., subclasses for which the expert’s predictions are correct. The remaining subclasses are defined as the expert’s weaknesses, i.e., subclasses for which the expert’s predictions are incorrect. The subset of the expert’s strengths subclasses is defined by selecting one of the subclasses uniformly at random as a strength base. The remaining strength subclasses are drawn randomly, weighted by their similarity to the strength base. We calculate the similarities between the subclasses as the average cosine distance between the feature representations of the instances, which are obtained from an EfficientNet-B1 model (Tan and Le 2019) pretrained on ImageNet. For the remaining weakness sub-

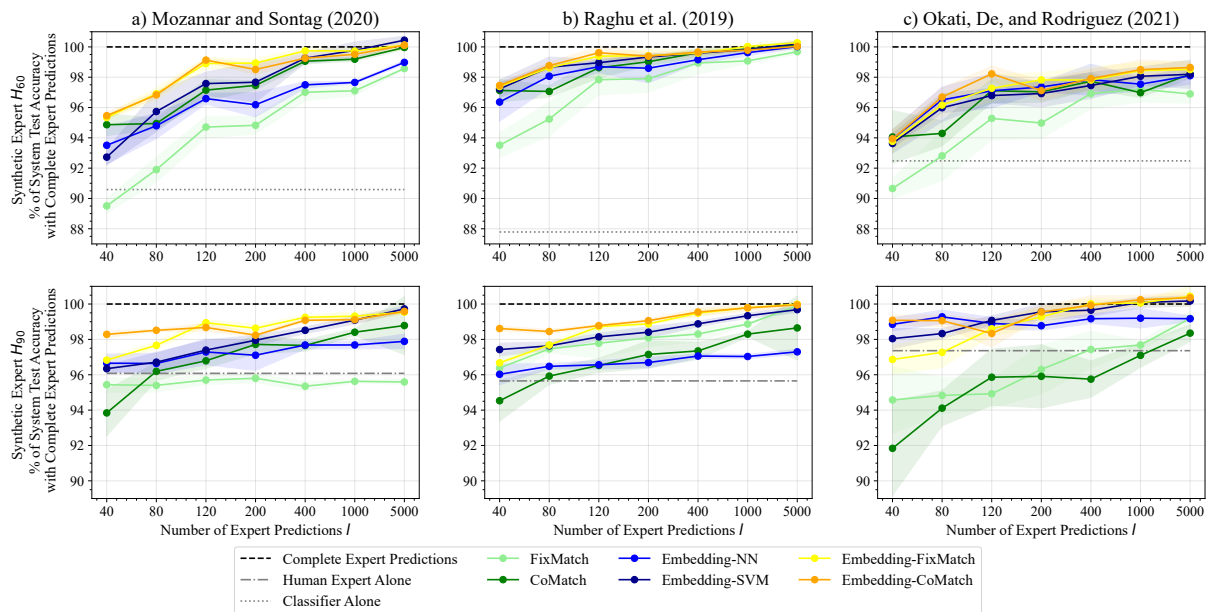


Figure 2: Experimental results for the learning to defer algorithms of a) Mozannar and Sontag (2020), b) Raghu et al. (2019), and c) Okati, De, and Rodriguez (2021) trained on different numbers of human expert predictions  $l$  and artificial expert predictions for  $H_{60}$  and  $H_{90}$  on CIFAR-100. Artificial expert predictions are generated using different numbers of human expert predictions  $l$ . We report mean and standard deviation over five seeds and omit the lower boundary for better visualization.

classes, the expert predicts randomly across all subclasses weighted by the similarity to the true subclass. Finally, the respective subclasses are mapped to the corresponding superclasses for the actual classification task.

**Implementation.** We allocate 40,000 images to the training and 10,000 images to the validation split while reserving 10,000 images for the test split. As an embedding model, we use an EfficientNet-B1 (Tan and Le 2019). The expertise predictor model consists of a single fully connected layer with a softmax activation. For the learning to defer algorithms, we use a WideResNet (Zagoruyko and Komodakis 2016) as backbone. We consider a strong synthetic expert  $H_{90}$  (90 of 100 subclasses as strengths) as well as a weaker one  $H_{60}$  (60 of 100 subclasses as strengths) and evaluate all approaches with the following number of available expert predictions:  $l \in \{40, 80, 120, 200, 400, 1000, 5000\}$ . We train the embedding model for 200 epochs using SGD as an optimizer with Nesterov momentum and a learning rate of 0.1. Each expertise predictor model of our *Embedding-SSL* approaches is trained for 50 epochs using SGD with a learning rate of 0.03. We repeat each experiment five times with different random seeds.

**Artificial Expert Prediction Results.** We report the performances of the artificial binary expert predictions of the expertise predictor model against the binary expert predictions for both synthetic experts in Table 1. Overall, the proposed *Embedding-SSL* approaches outperform all baselines for  $l < 1,000$ . More specifically, the results show that for  $l < 1,000$  our proposed *Embedding-SSL* approaches outperform the *SSL* baselines, on average, by 8.13% (for  $H_{60}$ )

and 3.80% (for  $H_{90}$ ). Moreover, the proposed *Embedding-SSL* approaches are able to outperform the *Embedding-SL* baselines for  $l < 1,000$ , on average, by 4.09% (for  $H_{60}$ ) and 1.97% (for  $H_{90}$ ). This demonstrates the benefit of using semi-supervised learning within the embedding space.

**Learning to Defer Results.** Figure 2 visualizes the results of the experiments with the approaches of Mozannar and Sontag (2020), Raghu et al. (2019), and Okati, De, and Rodriguez (2021). They show the general feasibility of training learning to defer algorithms using artificial expert predictions generated by our proposed approach, as all three outperform the individual team members in all experiments. Furthermore, the results indicate that even with only a few human expert predictions, the performances of the learning to defer algorithms converge quickly towards the upper boundary, matching the accuracy achieved with a fully labeled dataset. Regarding the approach of Mozannar and Sontag (2020), the *Embedding-SSL* approaches reach, on average, 99.02% (for  $H_{60}$ ) and 98.81% (for  $H_{90}$ ) of the upper boundary with only six expert predictions per class. Furthermore, the results support the previous statements concerning the superiority of the *Embedding-SSL* approaches. For  $l < 400$ , they outperform the *SSL* baselines, on average, by 3.27% (for  $H_{60}$ ) and 2.47% (for  $H_{90}$ ), and the *Embedding-SL* baselines, on average, by 1.99% (for  $H_{60}$ ) and 1.25% (for  $H_{90}$ ). The results of the approach of Raghu et al. (2019) confirm the advantage of the *Embedding-SSL* approaches. For  $l < 400$ , they improve the *SSL* baselines, on average, by 1.78% (for  $H_{60}$ ) and 1.68% (for  $H_{90}$ ). Regarding the *Embedding-SL*, an average improvement of 0.53%

$l$	4	8	12	20	40	100	500
FixMatch	56.97 ( $\pm 7.26$ )	59.43 ( $\pm 4.68$ )	72.63 ( $\pm 4.43$ )	62.09 ( $\pm 3.62$ )	56.39 ( $\pm 3.43$ )	83.40 ( $\pm 1.34$ )	87.25 ( $\pm 1.05$ )
CoMatch	70.88 ( $\pm 5.50$ )	69.17 ( $\pm 3.45$ )	76.76 ( $\pm 2.01$ )	67.86 ( $\pm 1.07$ )	69.79 ( $\pm 1.92$ )	84.45 ( $\pm 1.42$ )	89.77 ( $\pm 0.33$ )
Embedding-NN	85.87 ( $\pm 4.72$ )	88.61 ( $\pm 1.75$ )	89.03 ( $\pm 1.44$ )	91.14 ( $\pm 0.91$ )	91.35 ( $\pm 1.37$ )	91.27 ( $\pm 0.52$ )	92.29 ( $\pm 0.53$ )
Embedding-SVM	59.86 ( $\pm 33.44$ )	86.34 ( $\pm 1.63$ )	86.89 ( $\pm 2.32$ )	90.27 ( $\pm 0.95$ )	91.49 ( $\pm 0.98$ )	92.07 ( $\pm 0.55$ )	<b>92.74</b> ( $\pm 0.96$ )
<b>Embedding-FixMatch</b>	<b>90.74</b> ( $\pm 0.12$ )	<b>92.29</b> ( $\pm 0.00$ )	92.56 ( $\pm 0.00$ )	<b>91.65</b> ( $\pm 0.03$ )	<b>92.69</b> ( $\pm 0.13$ )	<b>92.22</b> ( $\pm 0.08$ )	91.96 ( $\pm 0.17$ )
<b>Embedding-CoMatch</b>	88.56 ( $\pm 1.34$ )	92.28 ( $\pm 0.12$ )	<b>92.63</b> ( $\pm 0.11$ )	91.57 ( $\pm 0.20$ )	90.84 ( $\pm 0.44$ )	91.93 ( $\pm 0.00$ )	91.07 ( $\pm 0.20$ )

Table 2:  $F_{0.5}$ -score (mean and standard deviation) of the artificial expert predictions generated by the expertise predictor model using different numbers of  $l$  available human expert predictions of radiologist 4295342357 on the NIH dataset.




Figure 3: Experimental results for the learning to defer algorithms of a) Mozannar and Sontag (2020), b) Raghu et al. (2019), and c) Okati, De, and Rodriguez (2021) trained on different numbers of human expert predictions  $l$  and artificial expert predictions for radiologist 4295342357. Artificial expert predictions are generated using different numbers of human expert predictions  $l$ . We report mean and standard deviation over five seeds.

(for  $H_{60}$ ) and 1.22% (for  $H_{90}$ ) can be achieved for  $l < 400$ . Lastly, for the approach of Okati, De, and Rodriguez (2021), the *Embedding-SSL* approaches outperform the *SSL* baselines, on average, by 1.96% (for  $H_{60}$ ) and 3.92% (for  $H_{90}$ ) for  $l < 400$ . Regarding the *Embedding-SL* baselines, our *Embedding-SSL* approaches achieve an average improvement of 0.40% (for  $H_{60}$ ) and 0.06% (for  $H_{90}$ ) over all  $l$ .

## NIH Chest X-rays

**Dataset.** The NIH dataset consists of radiologists’ predictions for four radiographic findings on 4,374 chest X-ray images from the ChestX-ray8 dataset (Majkowska et al. 2020; Wang et al. 2017). Contrary to the previous experiment, where we induced a synthetic connection between certain features, the degree to which the features are associated with experts’ capabilities remains ex-ante unclear for the NIH dataset. Each image is labeled by a panel of three radiologists denoting the “gold standard” ground truth labels. Moreover, the predictions of each of the three individual radiologists are reported for each image. We select the radiologist that provides the largest number of expert predictions (labeler-id: 4295342357) and focus on the diagnosis of the clinical finding *airspace opacity* for our experiments due to its balanced prevalence in the dataset. The experimental results using a further radiologist are reported in the Appendix.

**Implementation.** We split the 2,350 available images into a training and test set by randomly selecting 20% of the

patients as test patients. For the embedding model and the learning to defer algorithms, we use a ResNet18 (He et al. 2016) pretrained on the CheXpert dataset—a different X-ray dataset (Irvin et al. 2019). The expertise predictor model consists of a single fully connected layer with a softmax activation. We consider the following numbers of expert predictions:  $l \in \{4, 8, 12, 20, 40, 100, 500\}$ . We train the embedding model for 200 epochs using SGD as an optimizer with Nesterov momentum and a learning rate of 0.001. The expertise predictor models of our *Embedding-SSL* approaches are trained for 25 epochs using SGD with a learning rate of 0.03. We repeat the experiment five times with different random seeds.

**Artificial Expert Prediction Results.** The results of the artificial expert predictions are reported in Table 2. They demonstrate that, on average, the proposed *Embedding-SSL* approaches outperform both *SSL* baselines by 38.83% for  $l < 100$ . This confirms the result of the experiments conducted on CIFAR-100 regarding the impact of the embedding model. Furthermore, regarding the *Embedding-SL* baselines, the proposed *Embedding-SSL* approaches achieve an average performance improvement of 7.04% for  $l < 100$ , while for  $l \geq 100$ , the *Embedding-SL* baselines are on par with the performances of the proposed approaches.

**Learning to Defer Results.** Figure 3 displays the results of the learning to defer algorithms trained on human and artificial expert predictions for different numbers of available

human expert predictions on the NIH dataset. These results underline the feasibility of using artificial expert predictions generated by the *Embedding-SSL* approaches as the respective system accuracies exceed the lower baselines for all  $l$ . For the algorithm of Mozannar and Sontag (2020), the proposed *Embedding-SSL* approaches outperform the *SSL* baselines for  $l < 100$ , on average, by 1.97%. Compared to the *Embedding-SL* baselines, our *Embedding-SSL* approaches achieve an average improvement of 2.10% for  $l < 100$ . A closer look at the algorithm of Raghu et al. (2019) reveals that the proposed *Embedding-SSL* approaches outperform the *SSL* and *Embedding-SL* baselines, on average, by 1.67% and 1.06% for  $l < 100$  respectively. For the algorithm of Okati, De, and Rodriguez (2021), the *Embedding-SSL* approaches improve the *SSL* baselines, on average, by 2.42% and the *Embedding-SL* baselines by 0.69% for  $l < 100$ . For example, with only six expert predictions per class, the *Embedding-SSL* approaches enable the algorithm of Mozannar and Sontag (2020) to achieve, on average, 99.12% and the ones of Raghu et al. (2019) and Okati, De, and Rodriguez (2021) to achieve 98.83% of the performance compared to when they are trained with all available expert predictions.

**Bias.** Approximating the capabilities of a human expert through artificial expert predictions can entail the risk of inducing or amplifying prediction performance disparities between groups within a population. Therefore, for all learning to defer algorithms, we evaluate whether our proposed approach amplifies accuracy disparities with respect to patients' gender and age. For each number of human expert predictions  $l$  used to generate artificial ones, we calculate bias as the absolute difference in the accuracy between male and female patients. We can then compare it with the bias of the upper boundary (*Complete Expert Predictions*). We also apply this procedure to patients' age split into five bins. Here, we calculate bias as the mean absolute deviation of the accuracy of each age bin from the overall accuracy. For the algorithm of Mozannar and Sontag (2020), the average bias difference over all human expert predictions  $l$  between the proposed *Embedding-SSL* approaches and the upper boundary decreases by 0.72 percentage points (pp) for gender and increases by 2.63 pp for age. Similarly, for the algorithm of Raghu et al. (2019), we find a decrease of 1.53 pp for gender and an increase of 1.74 pp for age. Lastly, the algorithm of Okati, De, and Rodriguez (2021) has a 0.29 pp increase in bias for gender and a 0.8 pp increase in bias for age. We refer to the Appendix for additional visualizations for different numbers of human expert predictions  $l$  including an analysis for an additional radiologist. To summarize, we do not find evidence that the artificial expert predictions systematically amplify the bias of the evaluated learning to defer algorithms.

## Conclusion

In this work, we present a novel approach to reduce the number of human expert predictions required for the training of learning to defer algorithms. Our approach accurately learns the capabilities of an individual expert. On that basis, it generates artificial expert predictions that serve together with

the available expert predictions as training data. The result of our empirical evaluation shows that even a small number of expert predictions per class are sufficient for generating artificial expert predictions that enable learning to defer algorithms to exceed the individual performance of both human and AI. Currently, the data instances used to learn the expert's capabilities are selected randomly. Future work could explore advanced strategies that determine representative instances for which an expert should provide predictions. This could further improve the effectiveness of our approach and potentially make it easier to onboard new experts to collaborate with learning to defer algorithms.

## Ethics Statement

We hope that the proposed approach will contribute to the application of learning to defer algorithms in practice by reducing the required number of human expert predictions in addition to ground truth labels as a considerable implementation barrier. However, generating artificial expert predictions that are used together with a small number of human expert predictions to train learning to defer algorithms can entail the risk of potentially resulting in an overall system that is discriminatory against certain protected groups. For this reason, we want to raise awareness of the importance of monitoring fairness in the application settings in which the learning to defer algorithms are employed.

## Acknowledgments

This work was supported by the Federal Ministry for Economic Affairs and Climate Action of Germany in the project Smart Design and Construction (project number 01MK20016F).

## References

- Branson, S.; Van Horn, G.; and Perona, P. 2017. Lean crowdsourcing: combining humans and machines in an online system. In *Conference on Computer Vision and Pattern Recognition*, 7474–7483.
- Buolamwini, J.; and Gebru, T. 2018. Gender shades: intersectional accuracy disparities in commercial gender classification. In *Conference on Fairness, Accountability and Transparency*, 77–91.
- Charusaie, M.-A.; Mozannar, H.; Sontag, D.; and Samadi, S. 2022. Sample efficient learning of predictors that complement humans. In *International Conference on Machine Learning*, 2972–3005.
- Cremer, C. Z. 2021. Deep limitations? Examining expert disagreement over deep learning. *Progress in Artificial Intelligence*, 10(4): 449–464.
- De, A.; Koley, P.; Ganguly, N.; and Gomez-Rodriguez, M. 2020. Regression under human assistance. In *AAAI Conference on Artificial Intelligence*, 2611–2620.
- De, A.; Okati, N.; Zarezade, A.; and Gomez-Rodriguez, M. 2021. Classification under human assistance. In *AAAI Conference on Artificial Intelligence*, 5905–5913.



- Farrell, C.-J. 2021. Identifying mislabelled samples: machine learning models exceed human performance. *Annals of Clinical Biochemistry*, 58(6): 650–652.
- Gao, R.; Saar-Tsechansky, M.; De-Arteaga, M.; Han, L.; Lee, M. K.; and Lease, M. 2021. Human-AI collaboration with bandit feedback. In *International Joint Conference on Artificial Intelligence*, 1722–1728.
- Guan, M.; Gulshan, V.; Dai, A.; and Hinton, G. 2018. Who said what: modeling individual labelers improves classification. In *AAAI Conference on Artificial Intelligence*, 3109–3118.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Conference on Computer Vision and Pattern Recognition*, 770–778.
- Hemmer, P.; Schellhammer, S.; Vössing, M.; Jakubik, J.; and Satzger, G. 2022a. Forming effective human-AI teams: building machine learning models that complement the capabilities of multiple experts. In *International Joint Conference on Artificial Intelligence*, 2478–2484.
- Hemmer, P.; Schemmer, M.; Kühn, N.; Vössing, M.; and Satzger, G. 2022b. On the effect of information asymmetry in human-AI teams. In *CHI Conference on Human Factors in Computing Systems, Workshop on Human-Centered Explainable AI*.
- Irvin, J.; Rajpurkar, P.; Ko, M.; Yu, Y.; Ciurea-Ilicus, S.; Chute, C.; Marklund, H.; Haghighi, B.; Ball, R.; Shpan-skaya, K.; et al. 2019. Chexpert: a large chest radiograph dataset with uncertainty labels and expert comparison. In *AAAI Conference on Artificial Intelligence*, 590–597.
- Jones, E.; Sagawa, S.; Koh, P. W.; Kumar, A.; and Liang, P. 2021. Selective classification can magnify disparities across groups. In *International Conference on Learning Representations*.
- Keswani, V.; Lease, M.; and Kenthapadi, K. 2021. Towards unbiased and accurate deferral to multiple experts. In *AAAI/ACM Conference on AI, Ethics, and Society*, 154–165.
- Koch, G.; Zemel, R.; Salakhutdinov, R.; et al. 2015. Siamese neural networks for one-shot image recognition. In *International Conference on Machine Learning, Deep Learning Workshop*.
- Krizhevsky, A. 2009. Learning multiple layers of features from tiny images. Technical report, University of Toronto.
- Kühl, N.; Goutier, M.; Baier, L.; Wolff, C.; and Martin, D. 2022. Human vs. supervised machine learning: who learns patterns faster? *Cognitive Systems Research*, 76: 78–92.
- Lampert, T. A.; Stumpf, A.; and Gançarski, P. 2016. An empirical study into annotator agreement, ground truth estimation, and algorithm evaluation. *IEEE Transactions on Image Processing*, 25(6): 2557–2572.
- Leitão, D.; Saleiro, P.; Figueiredo, M. A.; and Bizarro, P. 2022. Human-AI collaboration in decision-making: beyond learning to defer. In *International Conference on Machine Learning, Workshop on Human-Machine Collaboration and Teaming*.
- Li, J.; Xiong, C.; and Hoi, S. C. 2021. CoMatch: semi-supervised learning with contrastive graph regularization. In *International Conference on Computer Vision*, 9475–9484.
- Liao, Y.-H.; Kar, A.; and Fidler, S. 2021. Towards good practices for efficiently annotating large-scale image classification datasets. In *Conference on Computer Vision and Pattern Recognition*, 4350–4359.
- Madras, D.; Pitassi, T.; and Zemel, R. 2018. Predict responsibly: improving fairness and accuracy by learning to defer. In *Advances in Neural Information Processing Systems*.
- Majkowska, A.; Mittal, S.; Steiner, D. F.; Reicher, J. J.; McKinney, S. M.; Duggan, G. E.; Eswaran, K.; Cameron Chen, P.-H.; Liu, Y.; Kalidindi, S. R.; et al. 2020. Chest radiograph interpretation with deep learning models: assessment with radiologist-adjudicated reference standards and population-adjusted evaluation. *Radiology*, 294(2): 421–431.
- Mozannar, H.; and Sontag, D. 2020. Consistent estimators for learning to defer to an expert. In *International Conference on Machine Learning*, 7076–7087.
- Okati, N.; De, A.; and Rodriguez, M. 2021. Differentiable learning under triage. *Advances in Neural Information Processing Systems*, 9140–9151.
- Raghu, M.; Blumer, K.; Corrado, G.; Kleinberg, J.; Obermeyer, Z.; and Mullainathan, S. 2019. The algorithmic automation problem: prediction, triage, and human effort. *arXiv preprint arXiv:1903.12220*.
- Rajpurkar, P.; Irvin, J.; Zhu, K.; Yang, B.; Mehta, H.; Duan, T.; Ding, D.; Bagul, A.; Langlotz, C.; Shpan-skaya, K.; et al. 2017. Chexnet: radiologist-level pneumonia detection on chest x-rays with deep learning. *arXiv preprint arXiv:1711.05225*.
- Raman, N.; and Yee, M. 2021. Improving learning-to-defer algorithms through fine-tuning. In *Advances in Neural Information Processing Systems, Workshop on Human and Machine Decisions*.
- Scott, T.; Ridgeway, K.; and Mozer, M. C. 2018. Adapted deep embeddings: a synthesis of methods for k-shot inductive transfer learning. In *Advances in Neural Information Processing Systems*.
- Snell, J.; Swersky, K.; and Zemel, R. 2017. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*.
- Sohn, K.; Berthelot, D.; Carlini, N.; Zhang, Z.; Zhang, H.; Raffel, C. A.; Cubuk, E. D.; Kurakin, A.; and Li, C.-L. 2020. FixMatch: simplifying semi-supervised learning with consistency and confidence. In *Advances in Neural Information Processing Systems*, 596–608.
- Tan, M.; and Le, Q. 2019. EfficientNet: rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, 6105–6114.
- Tian, Y.; Wang, Y.; Krishnan, D.; Tenenbaum, J. B.; and Isola, P. 2020. Rethinking few-shot image classification: a good embedding is all you need? In *European Conference on Computer Vision*, 266–282.
- Topol, E. J. 2019. High-performance medicine: the convergence of human and artificial intelligence. *Nature Medicine*, 25(1): 44–56.

- Tversky, A.; and Kahneman, D. 1974. Judgment under uncertainty: heuristics and biases. *Science*, 185(4157): 1124–1131.
- Vinyals, O.; Blundell, C.; Lillicrap, T.; Kavukcuoglu, K.; and Wierstra, D. 2016. Matching networks for one shot learning. In *Advances in Neural Information Processing Systems*.
- Wang, X.; Peng, Y.; Lu, L.; Lu, Z.; Bagheri, M.; and Summers, R. M. 2017. Chestx-ray8: hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In *Conference on Computer Vision and Pattern Recognition*, 2097–2106.
- Welinder, P.; and Perona, P. 2010. Online crowdsourcing: rating annotators and obtaining cost-effective labels. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition - Workshops*, 25–32.
- Wilder, B.; Horvitz, E.; and Kamar, E. 2020. Learning to complement humans. In *International Joint Conference on Artificial Intelligence*, 1526–1533.
- Zagoruyko, S.; and Komodakis, N. 2016. Wide residual networks. In *British Machine Vision Conference*, 87.1–87.12.