

Forming Effective Human-AI Teams: Building Machine Learning Models that Complement the Capabilities of Multiple Experts

Patrick Hemmer^{1*}, Sebastian Schellhammer^{1,2*}, Michael Vössing¹,
Johannes Jakubik¹ and Gerhard Satzger¹

¹Karlsruhe Institute of Technology

²GESIS - Leibniz Institute for the Social Sciences

{patrick.hemmer, michael.voessing, johannes.jakubik, gerhard.satzger}@kit.edu,
sebastian.schellhammer@gesis.org

Abstract

Machine learning (ML) models are increasingly being used in application domains that often involve working together with human experts. In this context, it can be advantageous to defer certain instances to a single human expert when they are difficult to predict for the ML model. While previous work has focused on scenarios with one distinct human expert, in many real-world situations several human experts with varying capabilities may be available. In this work, we propose an approach that trains a classification model to complement the capabilities of multiple human experts. By jointly training the classifier together with an allocation system, the classifier learns to accurately predict those instances that are difficult for the human experts, while the allocation system learns to pass each instance to the most suitable team member—either the classifier or one of the human experts. We evaluate our proposed approach in multiple experiments on public datasets with “synthetic” experts and a real-world medical dataset annotated by multiple radiologists. Our approach outperforms prior work and is more accurate than the best human expert or a classifier. Furthermore, it is flexibly adaptable to teams of varying sizes and different levels of expert diversity.

1 Introduction

Over the last years, the performance of machine learning (ML) models has become comparable to the one of human experts in a rising number of application domains [Esteva *et al.*, 2017]. For example, in medicine, several studies have demonstrated that ML models can surpass the performance of radiologists for the diagnosis of pneumonia in chest X-rays [Irvin *et al.*, 2019] or achieve similar performance for the detection of diabetic retinopathy [Gulshan *et al.*, 2019]. Despite these competitive results, individual experts are often still more accurate than ML algorithms in distinct areas of the input space [Raghu *et al.*, 2019], e.g., due to limited model capacity, limited training data, or the existence

of side information that is not accessible to the ML model [Wilder *et al.*, 2020]. In this context, prior work has investigated whether these different capabilities can be leveraged to enable ML models to complement the weaknesses of a single human expert by learning to predict only a fraction of the instances and pass the remaining ones to the domain expert [Bansal *et al.*, 2021; De *et al.*, 2021; Madras *et al.*, 2018; Mozannar and Sontag, 2020; Raghu *et al.*, 2019; Wilder *et al.*, 2020]. Even though in many real-world applications multiple human experts are available for load sharing and covering different instances [Chouldechova *et al.*, 2018; Grønsund and Aanestad, 2020], the majority of these studies focus on settings in which ML models take the decision themselves for a given instance or leave it to a single human expert. Consequently, they do not leverage potential disagreement between individual human experts due to diverse capabilities. Thus, from a team perspective, the potential to improve the overall performance remains unused. Differences in the predictions of domain experts can, e.g., stem from different levels of knowledge or personal biases [Lampert *et al.*, 2016]. Additionally, the rapid developments in highly specialized fields, e.g., medicine, are making it increasingly difficult for domain experts to acquire all knowledge within a single specialization domain. This creates the need to identify the appropriate human expert to assign particular instances to when they are difficult to predict for the ML model. We, therefore, propose an approach that jointly (1) trains a classifier to specifically complement the individual weaknesses of multiple human experts and (2) trains an allocation system to assign each instance either to the classifier or one specific human expert. The setting with several available human experts and a classifier is referred to as a team. The team member to whom an instance is assigned makes then a prediction on behalf of the entire team. Overall, our work makes the following contributions:

First, we present a novel approach that optimizes the performance of human-AI teams by jointly training a classifier together with an allocation system under consideration of the individual capabilities of available human experts. While the classifier learns to accurately predict the areas of the input space in which human experts are less competent, the allocation system learns to assign each instance to the best-suited team member.

Second, we demonstrate that our approach improves the

*These authors contributed equally.

state-of-the-art by conducting experiments on two datasets with synthetically generated human experts—a dataset consisting of tweets containing hate speech and offensive language [Davidson *et al.*, 2017] and the popular CIFAR-100 dataset. We show that our approach can not only deal with different team sizes but also better leverages high human expert diversity in terms of capabilities than prior work [Keswani *et al.*, 2021]. Moreover, we visually illustrate that the classifier in our approach makes correct predictions for the instances where human experts are less accurate and that the allocation system learns to assign exactly these instances to the classifier.

Third, we demonstrate the real-world applicability of our approach with an example from the medical domain. The Chest X-ray dataset [Majkowska *et al.*, 2020; Wang *et al.*, 2017] of the NIH Clinical Center provides a suitable setting as the annotations do not only include radiologist-adjudicated labels, which serve as a high-quality “gold standard”, but also the annotations on an individual radiologist level. In this context, prior work has conducted experimental evaluations on datasets that either provide only a distribution of human expert labels, e.g., CIFAR-10H [Peterson *et al.*, 2019] or labels that can not be traced back to individual human experts, e.g., Galaxy Zoo dataset [Lintott *et al.*, 2008].

We provide our code and the Appendix of this paper at <https://github.com/ptrckhmmr/human-ai-teams>.

2 Related Work

In recent years, there has been growing interest in optimizing the performance of human-AI teams—teams consisting of both a human expert and a classifier [Hemmer *et al.*, 2021]—by allocating a subset of the instances to a single human expert [De *et al.*, 2021; Keswani *et al.*, 2021; Madras *et al.*, 2018; Mozannar and Sontag, 2020; Okati *et al.*, 2021; Raghu *et al.*, 2019; Wilder *et al.*, 2020]. In this context, Raghu *et al.* [2019] separately train a classifier and a model for predicting whether a human expert errs. By deferring based on which of the two models has higher uncertainty, superior team performance can be achieved that surpasses both full automation (i.e., the classifier predicts all instances) and sole human effort (i.e., the human expert predicts all instances). While the authors train the classifier in isolation to optimize its performance, different studies emphasize that the classifier could also be trained to specifically complement the human expert’s strengths and weaknesses. Several approaches propose to jointly train a classifier and a deferral system [Madras *et al.*, 2018; Wilder *et al.*, 2020]. Other approaches utilize objective functions with theoretical guarantees [De *et al.*, 2021; Mozannar and Sontag, 2020]. De *et al.* [2021] focus on support vector machines and maximize sub-modular functions to optimize team performance. Mozannar and Sontag [2020] optimize a consistent surrogate loss derived from cost-sensitive learning. Further related work has investigated optimal triage policies [Okati *et al.*, 2021], user-initiated defer options [Bansal *et al.*, 2021], or learning complementarity from bandit feedback [Gao *et al.*, 2021]. In general, these approaches focus on combining a classifier with a single human expert. However, in reality, more than one

human expert is often available [Chouldechova *et al.*, 2018; Grønsund and Aanestad, 2020]. Hereby, it can be of particular interest in resource-scarce domains to consult exactly the right expert to avoid potentially costly mistakes. Our approach extends the idea that the classifier should complement the strengths and weaknesses of a single human expert to a setting with multiple human experts. The system allocates instances to one expert of a team consisting of multiple human experts and a complementary classifier. Concurrent work proposes a multi-label classification approach in which a classifier either makes a prediction or defers the instance to one of multiple human experts [Keswani *et al.*, 2021]. However, the approach does not leverage the human experts’ disagreement to enhance the overall team performance but solely relies on identifying the best overall team member instead.

3 Problem Formulation

In this section, we introduce the problem of optimizing the performance of human-AI teams on classification tasks. Given a training example $\mathbf{x} \in \mathcal{X}$, we learn to predict its ground truth label $y \in \mathcal{Y} = \{1, \dots, k\}$ where k denotes the number of classes. In addition to y , we assume to have access to the predictions $\mathbf{h} \in \mathcal{H}$ of m human experts for all training instances to implicitly learn their strengths and weaknesses. Each \mathbf{h} is a m -dimensional vector with $\mathbf{h} = [h_1, \dots, h_m]$ with $h_j \in \mathcal{Y} \forall j \in \{1, \dots, m\}$. We combine this data to the training dataset $\mathcal{D} = \{(\mathbf{x}, y, \mathbf{h})\}_1^N \sim P$, with N denoting the number of data instances, to jointly train an allocation system $A : \mathcal{X} \rightarrow \mathbb{R}^{m+1}$ and a classifier $F : \mathcal{X} \rightarrow \mathbb{R}^k$. Here, P is an (unknown) joint distribution, $m + 1$ denotes the total size of the human-AI team, and the outputs of A and F refer to the penultimate layer of the respective models. The classifier predicts $\hat{y}_{clf} = \arg \max F(\mathbf{x})$ and the allocation system indicates the competence of each team member including the classifier with $\mathbf{a} = A(\mathbf{x})$ for an instance \mathbf{x} . Combining the individual predictions of all team members $\hat{\mathbf{y}} = [h_1, \dots, h_m, \hat{y}_{clf}]$ and the index of the most competent team member $j_{max} = \arg \max \mathbf{a}$ results in the team prediction $\hat{y}_{team} = \hat{y}_{j_{max}}$. The goal is then to minimize the team loss

$$\mathcal{L}_{team}(F, A, \mathbf{x}, y, \mathbf{h}) = \mathbf{E}_{(\mathbf{x}, y, \mathbf{h}) \sim P} [l(y, \hat{y}_{team})]. \quad (1)$$

To minimize the loss, the classifier needs to be accurate on those instances where the human experts are less accurate. The allocation system’s function is to assign each instance to the team member with the highest probability of a correct prediction to answer on behalf of the team.

4 Approach

In this section, we present our approach that jointly trains a classifier to complement the capabilities of multiple human experts together with an allocation system that assigns instances to either one of the human experts or the classifier. We first introduce the surrogate loss used to train both components followed by formalizing the approach as an algorithm.

Our approach draws upon the mixture of experts (MoE) framework that combines multiple classifiers with a gating network and trains them simultaneously [Jacobs *et al.*, 1991].

The gating network combines the outputs of the classifiers and partitions the input space such that each classifier is trained to predict the outputs for one partition. However, contrary to the MoE framework, we combine the outputs of one classifier F with the fixed predictions of the human experts. Besides using the allocation system A as a gating network, we further employ it to allocate instances among team members. Given the output of the allocation system $\mathbf{a} = A(\mathbf{x})$, we obtain the probability w_j that team member j produces the correct output for an instance \mathbf{x} using

$$w_j = \frac{e^{a_j}}{\sum_{l=1}^{m+1} e^{a_l}} \quad j = 1, \dots, m+1. \quad (2)$$

The classifiers' conditional probability distribution over the ground truth labels \mathcal{Y} is given by

$$c_i = \frac{e^{z_i}}{\sum_{l=1}^k e^{z_l}} \quad i = 1, \dots, k, \quad (3)$$

where $\mathbf{z} = F(\mathbf{x})$ is the output of the classifier. We then combine the conditional probabilities \mathbf{c} with the one-hot encoded predictions of the human experts to

$$\mathbf{T} = [\mathbf{h}_1, \dots, \mathbf{h}_m, \mathbf{c}], \quad (4)$$

with $\mathbf{h}_j \in \{0, 1\}^k$. Using the predictions of all team members \mathbf{T} and the probabilities \mathbf{w} , we derive the conditional probabilities on a team level

$$P_{team}(Y = i|\mathbf{x}) = \sum_{j=1}^{m+1} w_j T_{i,j} \quad i = 1, \dots, k. \quad (5)$$

Finally, using the team prediction and the one-hot encoded ground truth label \mathbf{y} , we calculate the cross-entropy loss

$$\mathcal{L}_{team}(F, A, \mathbf{x}, \mathbf{y}, \mathbf{h}) = - \sum_{i=1}^k y_i \log(P_{team}(Y = i|\mathbf{x})) \quad (6)$$

for each instance to jointly train the allocation system and the classifier. Using the predictions of all team members weighted with the softmax probabilities of the allocation system in Equation (5) is a soft relaxation of the team prediction used as input in Equation (1) to ensure differentiability. We formalize our approach in Algorithm 1 and provide an intuition about the loss function's inner workings based on an analysis of the partial derivatives with respect to the classifier and allocation system in Appendix B.

5 Experiments

In this section, we evaluate the performance of our approach for different settings with synthetic and real-world human experts. First, we show that our approach is applicable to teams of varying sizes on a hate speech detection dataset from Twitter [Davidson *et al.*, 2017] and on the CIFAR-100 dataset [Krizhevsky, 2009]. Second, we investigate the influence of human experts' diversity with regard to their capabilities on the team performance using CIFAR-100 and visually illustrate the inner workings of our approach. The presence of 100 sub- and 20 superclasses in CIFAR-100 allows us to model

Algorithm 1 Training of our approach

Require: $\mathcal{D} = \{(\mathbf{x}, \mathbf{y}, \mathbf{h}_1, \dots, \mathbf{h}_m)\}_1^N$, the training data with \mathbf{y} and \mathbf{h}_j being one-hot encoded
Require: F , the classifier with $F : \mathcal{X} \rightarrow \mathbb{R}^k$
Require: A , the allocation system with $A : \mathcal{X} \rightarrow \mathbb{R}^{m+1}$

for number of iterations **do**
 sample a minibatch $B \subseteq \{1, \dots, N\}$
for $i \in B$ **do**
 $\mathbf{a} \leftarrow A(\mathbf{x}^{(i)})$
for $j = 1, \dots, m+1$ **do**
 $w_j \leftarrow \frac{e^{a_j}}{\sum_{l=1}^{m+1} e^{a_l}}$
end for
 $\mathbf{z} \leftarrow F(\mathbf{x}^{(i)})$
for $l = 1, \dots, k$ **do**
 $c_l \leftarrow \frac{e^{z_l}}{\sum_{o=1}^k e^{z_o}}$
end for
 $\mathbf{T} \leftarrow [\mathbf{h}_1^{(i)}, \dots, \mathbf{h}_m^{(i)}, \mathbf{c}]$
for $l = 1, \dots, k$ **do**
 $P_{team}(Y = l|\mathbf{x}^{(i)}) \leftarrow \sum_{j=1}^{m+1} w_j T_{l,j}$
end for
 $\mathcal{L}^{(i)} = - \sum_{l=1}^k y_l^{(i)} \log(P_{team}(Y = l|\mathbf{x}^{(i)}))$
end for
 Backpropagate $\mathcal{L}_{team} = \frac{1}{|B|} \sum_{i \in B} \mathcal{L}^{(i)}$
end for

human experts' capabilities with more granularity. Third, we assess its real-world applicability by presenting results on a chest X-ray dataset from the NIH Clinical Center [Majkowska *et al.*, 2020; Wang *et al.*, 2017] containing real radiologists' annotations. Additional experimental results and implementation details are provided in Appendix A.

Baselines. We compare the performance of our approach (*Classifier & Expert Team*) to multiple baselines. The first is the Joint Sparse Framework (*JSF*) by Keswani *et al.* [2021]. In contrast to our approach, the authors pursue a *multi-label* classification approach by training $m+1$ binary sigmoid classifiers and assigning each instance to the team member with the highest sigmoid weight. The second is a single classifier using a cross-entropy loss to predict all instances (*One Classifier*) with the same architecture as the classifier in our approach. For the third baseline, we allocate instances randomly to one human expert from the team (*Random Expert*). For the fourth baseline, we allocate all instances to the most accurate human expert in the team (*Best Expert*).

Moreover, we report two additional baselines to investigate the contribution of the classifier and human experts to the team performance. Both baselines are simplified versions of the approach outlined in this work. For the first, we replace all m human experts with the same number of classifiers resulting in a *Classifier Team* of size $m+1$. These have the same architecture as the classifier in our approach and are jointly trained with the allocation system according to Algorithm 1. For the second, we omit the classifier resulting in an *Expert Team* of size m . The human experts are then used to train the allocation system using Algorithm 1.

5.1 Hate Speech and Offensive Language Dataset

We first conduct experiments on a Twitter dataset containing hate speech and offensive language [Davidson *et al.*, 2017].

Data and Expert Generation. The dataset consists of 24,783 tweets labeled as hate speech, offensive language, or neither. We follow the procedure for dataset preprocessing as described in Keswani *et al.* [2021]. We transform it into a binary classification task with 1 indicating hate speech or offensive language and 0 indicating neither of them. The tweets are additionally annotated with a second label indicating their dialect (i.e., African-American English (AAE) or non-African-American English (non-AAE)). We create m synthetic experts using this dialect label. For each expert, we sample the quantities $p, q \sim \mathcal{U}(\cdot)$, with $\mathcal{U}(\cdot)$ being the uniform distribution. The probability of correctly classifying non-AAE tweets is denoted by p , whereas q represents the probability of correctly classifying AAE tweets. For $\lfloor \frac{3m}{4} \rfloor$ experts, we define $p \sim \mathcal{U}(0.6, 1)$ and $q \sim \mathcal{U}(0.6, p)$. For the remaining $\lceil \frac{m}{4} \rceil$ experts, we set $q \sim \mathcal{U}(0.6, 1)$ and $p \sim \mathcal{U}(0.6, q)$. As a result, $\lfloor \frac{3m}{4} \rfloor$ experts are more accurate on non-AAE tweets and $\lceil \frac{m}{4} \rceil$ experts are more accurate on AAE tweets.

Model. We use pretrained GloVe word embeddings [Pennington *et al.*, 2014] to generate 100-dimensional feature vectors. These features are used to train our approach and the algorithmic baselines. The classifier and allocation system are modeled by a neural network consisting of a single hidden layer with 50 units followed by a ReLU activation.

Experimental Setup. To investigate the influence of the team size, we generate 2 to 20 synthetic human experts. The dataset is divided into a train, validation, and test split with 80%, 10%, and 10% of the data. We train the models for 100 epochs using Adam optimizer with a learning rate of $5 \cdot 10^{-3}$, a cosine annealing learning rate scheduler, and a batch size of 512. We apply early stopping on the validation split and report the results on the test split. We repeat the experiments 5 times with different seeds.

Results. We display the team accuracy of our approach and the respective baselines over the number of human experts in Figure 1. We see that all methods except the *Random Expert* baseline drawing on human experts exhibit an increase in

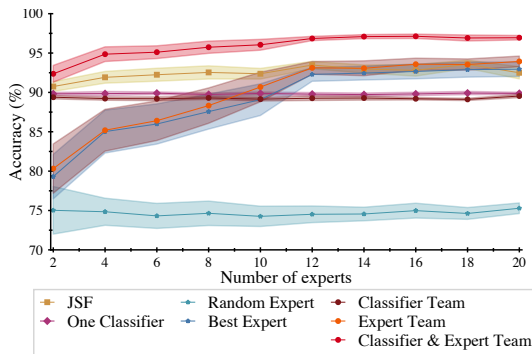


Figure 1: Team accuracy of our approach and the baselines with increasing team size on the Twitter dataset. Shaded regions display standard errors.

the overall performance with growing team size. This aligns with our expectation as with increasing team size more accurate human experts might enter the team and the covered feature space in which at least one team member is likely to perform an accurate prediction potentially increases. When taking a closer look, we find that our approach (*Classifier & Expert Team*) outperforms all baselines. The average improvement across all team sizes is 3.66% over the *JSF* baseline, 6.76% over the *One Classifier* baseline, 7.92% over the *Best Expert* baseline, 28.38% over the *Random Expert* baseline, and 7.44% over the *Classifier Team*. From the 6.96% performance improvement of our approach over the *Expert Team*, it becomes evident that the classifier in our approach learns to accurately predict the areas of the feature space that are not covered by human experts’ competencies.

5.2 CIFAR-100 Dataset

The analysis in this section is threefold and uses the CIFAR-100 dataset [Krizhevsky, 2009]. First, we provide further experimental evidence on the performance of our approach with varying team sizes. Second, we investigate the influence of the diversity of human experts’ capabilities on our approach. Third, we visualize how our approach distributes instances among classifier and human experts relative to each team member’s respective competencies.

Dataset and Expert Generation. The dataset consists of 60,000 images drawn from 100 subclasses which are grouped into 20 superclasses of equal size. Using the 20 superclasses for the classification task and the 100 subclasses for synthetic human expert generation allows us to model their competencies with more granularity compared to the Twitter dataset. For the first analysis, we simulate human experts in the following way. By design, for each of m human experts in the team, we draw $k_j \sim \mathcal{N}(70, 5) \forall j \in \{1, \dots, m\}$ from a normal distribution with $\lfloor k_j \rfloor$ denoting the number of subclasses that each human expert predicts perfectly. For each of m human experts, the set of perfectly predicted subclasses is sampled uniformly at random without replacement from all 100 subclasses and mapped to the respective superclasses. For the remaining subclasses, each expert predicts uniformly at random across all superclasses. For the second analysis, we

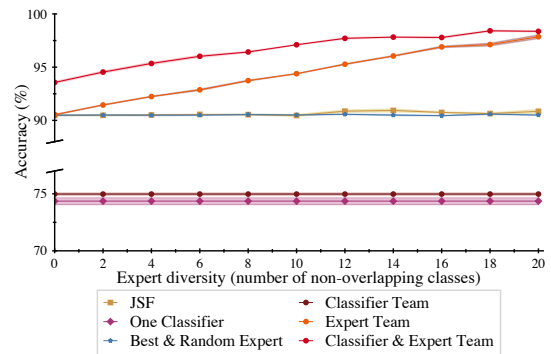


Figure 2: Team accuracy of our approach and the baselines with increasing human expert diversity (number of non-overlapping classes) on CIFAR-100. Shaded regions display standard errors.




Figure 3: Visualization of the team members’ performance—human experts’ accuracy in (a) and classifier accuracy in (b)—and instance assignment by the allocation system in (c) on a subset of the CIFAR-100 input space.

select $m = 2$ human experts over $i \in \mathbb{N}$ runs. To isolate the influence of diversity, we vary the capabilities of the second human expert per run while keeping the ones of the first fixed. For $i = 1$, we select the first 90 subclasses $\{1, \dots, 90\}$ in alphabetical order for which both human experts predict the corresponding superclass perfectly. For the remaining subclasses, both experts predict uniformly at random across all superclasses. We then vary the competencies of the second human expert such that for run i the perfectly predicted subclasses are $\{i, \dots, 89 + i\}$. By increasing i per run by 1, the number of subclasses in which only one of both human experts predicts perfectly increases by 2 subclasses resulting in a diversity increase. For the third analysis, we resort to one of these diversity scenarios.

Model. We use a ResNet-18 model [He *et al.*, 2016] pre-trained on ImageNet as a fixed feature extractor with 512 output units. These features are used to train our approach and the algorithmic baselines. The classifier and allocation system are modeled by a neural network consisting of a single hidden layer with 100 units followed by a ReLU activation.

Experimental Setup. For the first analysis, we assess the performance of teams with 2 to 10 human experts with a normally distributed number of perfectly predicted subclasses as previously described. For the second analysis, we vary i from 1 to 11 for the second human expert to gradually increase the number of non-overlapping subclasses. As a result, diversity is 0 for run $i = 1$ and 20 for run $i = 11$. For the third analysis, we select one of the diversity scenarios—namely where the first human expert is defined as $\{1, \dots, 90\}$ and the second one as $\{6, \dots, 95\}$ —to visualize how instances are distributed among classifier and human experts relative to each team member’s respective competencies.

We divide the 50,000 training images into a training and validation split of 80% and 20% and use the 10,000 test images as the test split. We train the models for 100 epochs using Adam optimizer with a learning rate of $5 \cdot 10^{-3}$, a cosine annealing learning rate scheduler, weight decay of $5 \cdot 10^{-4}$, and a batch size of 512. We apply early stopping on the validation split and report the results on the test split. We repeat the experiments 5 times with different seeds.

Results. First, consistent with the results from Section 5.1, we find that our approach benefits from an increasing number of human experts. This is in line with our expectations as the probability that at least one expert has strengths in distinct subclasses or on all subclasses of a superclass is higher with

an increasing number of human experts. In both cases, the classifier can focus on an even smaller area of the input space. On average, our approach (*Classifier & Expert Team*) outperforms the *JSF* baseline by 19.77%, the *Best Expert* baseline by 19.73%, the *One Classifier* baseline by 25.86%, the *Random Expert* baseline by 30.75%, and the *Classifier Team* by 24.88%. Even though the contribution from the classifier in our approach diminishes the more human experts become part of the team, it leverages each individual team member’s capabilities best. This can be observed from a performance gain of 0.98% compared to the *Expert Team*. We refer to Figure A1 in Appendix A for an additional visualization.

Second, we find that for the *Classifier & Expert Team*, a higher diversity of human team members’ capabilities contributes to higher team performance as displayed in Figure 2. Even though the contribution of the classifier diminishes in the presence of increasingly diverse human experts, it adds an average performance gain of 2.40% across all runs compared to the *Expert Team*. Contrary, the *JSF* baseline does not benefit from a diversity increase but instead assigns the instances to the best human expert. Note that we aggregate the *Best Expert* and *Random Expert* baseline in Figure 2 as both human experts have the same accuracy by design.

Finally, we illustrate the capabilities of our approach in Figure 3: (1) the classifier’s compensation for human experts’ weaknesses and (2) the allocation system’s assignment of instances relative to each team member’s competencies. For the visualization, we reduce the hidden layer’s features to 2 dimensions using the UMAP algorithm [McInnes *et al.*, 2018] and aggregate a set of images in each hexagon. Figure 3a and Figure 3b show the accuracy of 2 human experts and the classifier, respectively. In combination, the illustrations reveal that the classifier primarily learns to optimize its performance on those images (green hexagons in Figure 3b) where human experts tend to be less accurate (yellow and red hexagons in Figure 3a). Figure 3c displays that the allocation system assigns those instances to the classifier where the human experts are less accurate. Here, dark areas indicate the assignment of the images to the classifier and brighter areas denote the allocation to one of the human experts.

5.3 NIH Dataset

Lastly, we evaluate our approach on a real-world chest X-ray dataset collected by the NIH Clinical Center [Majkowska *et al.*, 2020; Wang *et al.*, 2017]. In the previous experiments, we induced a synthetic connection between certain features

Method	Team Accuracy (%)			
	[<i>F</i> , 357, 117]	[<i>F</i> , 357, 121]	[<i>F</i> , 249, 124]	[<i>F</i> , 249, 296]
Baselines:				
- JSF	94.69 (± 0.25)	90.14 (± 0.17)	88.69 (± 0.32)	90.92 (± 0.13)
- One Classifier	83.13 (± 0.21)	84.70 (± 0.13)	84.59 (± 0.12)	83.63 (± 0.08)
- Random Expert	89.52 (± 0.29)	89.05 (± 0.19)	88.45 (± 0.19)	87.20 (± 0.36)
- Best Expert	95.34 (± n/a)	91.01 (± n/a)	88.84 (± n/a)	91.31 (± n/a)
- Classifier Team	83.44 (± 0.16)	85.09 (± 0.17)	84.83 (± 0.23)	83.69 (± 0.27)
- Expert Team	95.34 (± 0.04)	91.08 (± 0.04)	88.33 (± 0.11)	91.29 (± 0.11)
Our Approach:				
- C. & E. Team	95.45 (± 0.06)	91.72 (± 0.12)	90.36 (± 0.15)	91.51 (± 0.13)

Table 1: Team accuracies of our approach and the baselines including standard errors for 4 different human-AI teams consisting of the classifier F and two radiologists on the NIH dataset.

(i.e., the tweet dialect or the image subclass) and the human experts’ capabilities. Contrary, for the NIH dataset, it remains ex-ante unclear to which degree features are associated with the human experts’ competencies.

Dataset. The dataset comprises radiologists’ annotations for 4 radiographic findings on 4,374 chest X-ray images from the ChestX-ray8 dataset [Majkowska *et al.*, 2020; Wang *et al.*, 2017]. We focus on the occurrence of the clinically important finding *airspace opacity* which has a prevalence (i.e., the percentage of patients in the dataset that is affected with this disease) of 49.50%. For each X-ray image, the ground-truth label was adjudicated by a panel of 3 radiologists serving as the “gold standard” label. Additionally, the annotations of each of the three radiologists from a cohort of 22 human experts are reported per image. Each radiologist can be identified by a unique ID of which we list the last three digits due to space constraints.

Model. We use a ResNet-18 model pretrained on the CheXpert dataset—a different X-ray dataset for chest radiograph interpretation [Irvin *et al.*, 2019]—which serves as a fixed feature extractor with 512 output units. For our approach and the algorithmic baselines, we use these features for model training. The classifier and allocation system are modeled by a single hidden layer neural network with 30 units and a ReLU activation function.

Experimental Setup. We conduct experiments with 4 pairs of radiologists that share more than 800 images labeled by both radiologists. Due to insufficient intersection, a third radiologist could not be included in the team. We perform a 10-fold cross-validation while ensuring that images from one patient are not spread across different folds. To keep the per fold performance of the radiologists close to their overall performance, we use stratification. For each cross-validation run, we use 7 folds for training and 2 folds for validation to perform early stopping as the team performance on one validation fold is not always a good proxy for the team performance on the test fold due to the small dataset size. The remaining fold serves as the test data. We train the models for 20 epochs using Adam optimizer with a learning rate of $1 \cdot 10^{-3}$, weight decay of $5 \cdot 10^{-4}$, and a batch size of 64. We repeat the experiments 5 times with different seeds.

Results. Table 1 compares the performance of our approach with the baselines. In detail, the performance of our approach (*Classifier & Expert Team*) outperforms the other baselines.

Instances allocated to	Individual Accuracy (%)		
	Classifier F	ID = 357	ID = 117
Classifier F	97.73 (± 1.76)	78.93 (± 2.47)	96.43 (± 1.86)
ID = 357	31.88 (± 4.73)	98.78 (± 0.32)	98.55 (± 0.34)
ID = 117	61.69 (± 1.15)	79.90 (± 0.29)	94.67 (± 0.02)

Table 2: Individual accuracy of the classifier, radiologist ID = 357, and radiologist ID = 117 on the subset of instances assigned to each of them. Additionally, the team members’ performances on the subsets not assigned to them are displayed.

Particularly, in high-stakes decision-making (e.g., medicine), the performance improvements can contribute to further reducing potentially severe consequences of misdiagnoses. A closer look at the subset of images assigned to each human expert and the classifier shows that the performance of each team member is highest on the subset assigned to them. Table 2 illustrates this allocation by the diagonal values using the example of the human-AI team [F , 357, 117]. Thus, the classifier learns and is assigned to an area of the input space where it is more accurate than the radiologists. Furthermore, the allocation system assigns the remaining instances between the two radiologists based on their competencies.

6 Conclusion

In this work, we leverage the complementary capabilities of human experts and AI models by jointly training a classifier together with an allocation system. While the classifier augments human experts’ weaknesses, the allocation system assigns instances to either one of the human experts or the classifier based on their individual capabilities. We provide experimental evidence with synthetic and real human expert annotations for the ability of such teams to not only outperform prior work but also to achieve superior performance results none of the team members could have accomplished alone. Furthermore, we show that our approach can cope with different degrees of human expert diversity and generally benefits from a higher number of human team members. As we do not constrain the availability and capacity of experts at run time, opportunities for future research could explore load balancing strategies or alternative assignment strategies if the most accurate team member is not available. Lastly, we assume access to the human experts’ predictions for all training data, implying a significant manual labeling effort. To address this limitation, further work could extend our approach to leverage data that is not annotated by all human experts in a team.

Ethical Statement

As this work pursues the training of an algorithm that augments the weaknesses of human experts, it is subject to important ethical considerations. Our goal is to realize improved performance compared to full automation and sole human effort which can be beneficial in high-stakes decision-making. However, the approach could theoretically be used to estimate the contribution of individual team members to the team’s overall performance. This could be used to exclude insufficiently performing experts from the team. However, identifying the weaknesses of individual experts also provides an opportunity to educate them through targeted training.

References

- [Bansal *et al.*, 2021] Gagan Bansal, Besmira Nushi, Ece Kamar, Eric Horvitz, and Daniel S. Weld. Is the most accurate ai the best teammate? optimizing ai for teamwork. In *AAAI*, 2021.
- [Chouldechova *et al.*, 2018] Alexandra Chouldechova, Diana Benavides Prado, Oleksandr Fialko, and Rhema Vaithianathan. A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions. In *FAT*, 2018.
- [Davidson *et al.*, 2017] Thomas Davidson, Dana Warmusley, Michael W. Macy, and Ingmar Weber. Automated hate speech detection and the problem of offensive language. In *ICWSM*, 2017.
- [De *et al.*, 2021] Abir De, Nastaran Okati, Ali Zarezade, and Manuel Gomez-Rodriguez. Classification under human assistance. In *AAAI*, 2021.
- [Esteva *et al.*, 2017] Andre Esteva, Brett Kuprel, Roberto A. Novoa, Justin M. Ko, Susan M. Swetter, Helen M. Blau, and Sebastian Thrun. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 2017.
- [Gao *et al.*, 2021] Ruijiang Gao, Maytal Saar-Tsechansky, Maria De-Arteaga, Ligong Han, Min Kyung Lee, and Matthew Lease. Human-ai collaboration with bandit feedback. In *IJCAI*, 2021.
- [Grønsund and Aanestad, 2020] Tor Grønsund and Margunn Aanestad. Augmenting the algorithm: emerging human-in-the-loop work configurations. *The Journal of Strategic Information Systems*, 2020.
- [Gulshan *et al.*, 2019] Varun Gulshan, Renu P. Rajan, Kasumi Widner, Derek J. Wu, Peter Wubbels, Tyler Rhodes, Kira Whitehouse, Marc Coram, Greg S Corrado, Kim Ramasamy, Rajiv Raman, Lily H. Peng, and Dale R. Webster. Performance of a deep-learning algorithm vs manual grading for detecting diabetic retinopathy in india. *JAMA Ophthalmology*, 2019.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [Hemmer *et al.*, 2021] Patrick Hemmer, Max Schemmer, Michael Vössing, and Niklas Kühl. Human-ai complementarity in hybrid intelligence systems: a structured literature review. In *PACIS*, 2021.
- [Irvin *et al.*, 2019] Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silviana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghighi, Robyn Ball, et al. Chexpert: a large chest radiograph dataset with uncertainty labels and expert comparison. In *AAAI*, 2019.
- [Jacobs *et al.*, 1991] Robert A. Jacobs, Michael I. Jordan, Steven J. Nowlan, and Geoffrey E. Hinton. Adaptive mixtures of local experts. *Neural Computation*, 1991.
- [Keswani *et al.*, 2021] Vijay Keswani, Matthew Lease, and Krishnaram Kenthapadi. Towards unbiased and accurate deferral to multiple experts. In *AIES*, 2021.
- [Krizhevsky, 2009] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- [Lampert *et al.*, 2016] Thomas A Lampert, André Stumpf, and Pierre Gançarski. An empirical study into annotator agreement, ground truth estimation, and algorithm evaluation. *IEEE Transactions on Image Processing*, 2016.
- [Lintott *et al.*, 2008] Chris J. Lintott, Kevin Schawinski, Anže Slosar, Kate R. Land, Steven Bamford, Daniel I. Thomas, M. Jordan Raddick, Robert Nichol, Alexander S. Szalay, Daniel Andreescu, P. G. Murray, and Jan van den Berg. Galaxy zoo: morphologies derived from visual inspection of galaxies from the sloan digital sky survey. *Monthly Notices of the Royal Astronomical Society*, 2008.
- [Madras *et al.*, 2018] David Madras, Toniann Pitassi, and Richard S. Zemel. Predict responsibly: improving fairness and accuracy by learning to defer. In *NeurIPS*, 2018.
- [Majkowska *et al.*, 2020] Anna Majkowska, Sid Mittal, David F Steiner, Joshua J Reicher, Scott Mayer McKinney, Gavin E Duggan, Krish Eswaran, Po-Hsuan Cameron Chen, Yun Liu, Kalidindi, et al. Chest radiograph interpretation with deep learning models: assessment with radiologist-adjudicated reference standards and population-adjusted evaluation. *Radiology*, 2020.
- [McInnes *et al.*, 2018] Leland McInnes, John Healy, and James Melville. Umap: uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.
- [Mozannar and Sontag, 2020] Hussein Mozannar and David A. Sontag. Consistent estimators for learning to defer to an expert. In *ICML*, 2020.
- [Okati *et al.*, 2021] Nastaran Okati, Abir De, and Manuel Rodriguez. Differentiable learning under triage. In *NeurIPS*, 2021.
- [Pennington *et al.*, 2014] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. Glove: global vectors for word representation. In *EMNLP*, 2014.
- [Peterson *et al.*, 2019] Joshua C. Peterson, Ruairidh M. Battleday, Thomas L. Griffiths, and Olga Russakovsky. Human uncertainty makes classification more robust. In *ICCV*, 2019.
- [Raghu *et al.*, 2019] Maithra Raghu, Katy Blumer, Greg Corrado, Jon Kleinberg, Ziad Obermeyer, and Sendhil Mullainathan. The algorithmic automation problem: prediction, triage, and human effort. *arXiv preprint arXiv:1903.12220*, 2019.
- [Wang *et al.*, 2017] Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, Mohammadhadi Bagheri, and Ronald M. Summers. Chestx-ray8: hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In *CVPR*, 2017.
- [Wilder *et al.*, 2020] Bryan Wilder, Eric Horvitz, and Ece Kamar. Learning to complement humans. In *IJCAI*, 2020.