

On the Utility of Prediction Sets in Human-AI Teams

Varun Babbar¹, Umang Bhatt^{1,2}, Adrian Weller^{1,2}

¹University of Cambridge

²The Alan Turing Institute

{vb395, usb20, aw665}@cam.ac.uk

Abstract

Research on human-AI teams usually provides experts with a single label, which ignores the uncertainty in a model’s recommendation. Conformal prediction (CP) is a well established line of research that focuses on building a theoretically grounded, calibrated prediction set, which may contain multiple labels. We explore how such prediction sets impact expert decision-making in human-AI teams. Our evaluation on human subjects finds that set valued predictions positively impact experts. However, we notice that the predictive sets provided by CP can be very large, which leads to unhelpful AI assistants. To mitigate this, we introduce D-CP, a method to perform CP on some examples and defer to experts. We prove that D-CP can reduce the prediction set size of non-deferred examples. We show how D-CP performs in quantitative and in human subject experiments (n=120). Our results suggest that CP prediction sets improve human-AI team performance over showing the top-1 prediction alone, and that experts find D-CP prediction sets are more useful than CP prediction sets.

1 Introduction

Human-AI collaboration is of increasing importance. Several works have shown the benefits of human-AI collaboration in boosting accuracy, fairness, and compatibility [Madras *et al.*, 2018; Bansal *et al.*, 2019; Mozannar and Sontag, 2020]. One form of collaboration in the medical domain is the effect of AI explanations on team performance [Lundberg *et al.*, 2018], where team performance improves if model explanations are provided. Another form of human-AI collaboration develops techniques for models to defer to an expert. Prior literature exploring both these forms of collaboration has mainly considered models which output singular predictions. However, this does not allow experts to gauge and interpret the predictive uncertainty of the model, which can prevent deployment in high risk settings. A solution to this is for the model to display *set valued predictions*. We define a set valued model prediction Γ as a mapping from the input space \mathcal{X} to the power set of the label space \mathcal{Y} , i.e. $\Gamma : \mathcal{X} \rightarrow 2^{\mathcal{Y}}$. One way to construct a set valued predictor is through a technique called Conformal




Figure 1: An AI assistant working alongside an expert can output one of three things: the most likely label, a set valued prediction with a predetermined error probability, or a deferral token indicating that the example should be labelled by the expert. The precise nature of the AI’s prediction should be dependent on acquired knowledge of the expert’s capabilities. Generally speaking, because the size of the predictive set is a reflection of the model’s confidence, deferring examples on which an expert is more confident than the model would prevent an expert from using large, incoherent prediction sets.

Prediction (CP) [Vovk *et al.*, 2005]. CP generates a prediction set that may contain multiple labels, but contains the true label with a user defined error probability. The goal of CP is to construct predictive sets that are sufficiently small but have high probability of containing the true label.

One problem often associated with CP sets is that they can be quite large, which can limit their usefulness in time and cost sensitive domains such as medical diagnostics, where it is crucial to narrow down the list of possible diagnoses. Previous work such as [Bellotti, 2021] and [Stutz *et al.*, 2022] have devised surrogate loss functions for minimizing set sizes whilst maintaining coverage guarantees. [Angelopoulos *et al.*, 2020] regularize the low scores of unlikely classes to provide small, stable sets. However, CP literature in general has given little consideration given to a) how useful such predictive sets are in human-AI teams and b) how human expertise could be leveraged to get smaller predictive sets [Sadinle *et al.*, 2016; Romano *et al.*, 2020; Angelopoulos and Bates, 2021]. Recently, [Straitouri *et al.*, 2022] explored improving expert predictions using CP, with a focus on finding the optimal error tolerance parameter α that benefits the expert. In this concurrent work, we fix α and use deferral as a mechanism to provide sets that are smaller and hence more useful to an expert. Our contributions are the following:

- Through human subject experiments on CIFAR-100, we

first show that CP sets result in higher levels of reported trust and utility as compared to Top-1 predictions.

- We empirically demonstrate one limiting aspect of CP sets: they can be very large for some examples. In order to mitigate this, we propose combining set valued classification and learning to defer. Through a toy example, we show the utility of learning to defer for the predictive set size of non-deferred examples.
- We then introduce D-CP, a general practical scheme that learns to defer on some examples and perform conformal prediction on others. We prove that, under mild conditions, the set size associated with non-deferred examples will be smaller than that of the original dataset.
- Through further human subject experiments on the CIFAR-100 dataset, we discover two benefits of D-CP: smaller predictive sets and improved team performance. We find that D-CP leads to higher reported utility and expert accuracy compared to showing only CP sets.

2 Related Work

2.1 Conformal Prediction

There is growing interest in conformal prediction [Vovk *et al.*, 2005] as a method of rigorous uncertainty quantification. Given a test example X_{test} and its (hidden) true label Y_{test} , this method allows the user to construct sets $\Gamma(X_{test})$ that control for the binary risk, i.e. the error probability $\alpha = P(Y_{test} \notin \Gamma(X_{test}))$. This is done by performing a statistical test for each label in order to decide whether the label should be present in the set. In particular, we define a *conformity score* $\tau(X_{test}, y) : \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{R}$ that determines how different example (X_{test}, y) is from already observed data $\{(X_i, Y_i)\}_{i=1}^n$. This is a design choice and several papers explore different choices of conformity functions [Sadinle *et al.*, 2016; Angelopoulos *et al.*, 2020; Romano *et al.*, 2020]. To include label y in a predictive set, we require that the conformity score $\tau(X_{test}, y)$ is at least α -common with respect to conformity scores on previously observed data, i.e. $\text{Quantile}(\tau(X_{test}, y), \{\tau(X_i, Y_i)\}) \geq \alpha$.

This is equivalent to learning a threshold conformity score for including labels in a set. Defining the threshold as $\tau_{cal} = \text{Quantile}(\alpha, \{\tau(X_i, Y_i)\})$, we require $\tau(X_{test}, y) \geq \tau_{cal}$ to include the label y in the predictive set for X_{test} . The conformal set is therefore defined as: $\Gamma(X_{test}, \tau_{cal}) = \{y : \tau(X_{test}, y) \geq \tau_{cal}\}$. In this paper, we employ a computationally efficient scheme called Inductive Conformal Prediction (ICP) [Papadopoulos, 2008]. This requires an additional **calibration dataset** $\mathcal{D}_{cal} = \{(X_i, Y_i)\}_{i=1}^n$ drawn from the same distribution as training and validation sets. After training a classifier on a training dataset, we can use this calibration dataset to choose the α Quantile threshold τ_{cal} .

None of the works that have previously studied CP [Sadinle *et al.*, 2016; Angelopoulos *et al.*, 2020; Romano *et al.*, 2020; Stutz *et al.*, 2022] involved experts in the loop or even considered the utility of prediction sets generated in the context of human-AI teams.

2.2 Learning to Defer

Many works have studied the idea of learning a model that adapts to an underlying expert. One approach is to learn a rejector and a classifier, wherein, given a cost of deferring c , one learns a binary classifier that rejects whenever it is less than $1 - c$ confident [Bartlett and Wegkamp, 2008; Cortes *et al.*, 2016]. For multi-class problems, [Mozannar and Sontag, 2020] learn a model that predicts the true label whenever the expert is wrong and defers otherwise. Similarly, [Okati *et al.*, 2021] develop a method for exact triage under multiple expert annotations and prove its optimality under conditions where there is expert disagreement. [Wilder *et al.*, 2021], on the other hand, develop a decision theoretic approach, training 3 probabilistic models representing the AI, expert, and joint human-AI to maximize utility. However, all these approaches only examine settings where the AI makes point predictions whereas we aim to defer some examples and provide principled, calibrated prediction sets on others.

3 Are Prediction Sets better for Human-AI teams than Top-1 Predictions?

Our first study focuses on establishing the value of set valued predictions. For our experiments, we focus on one particular CP scheme called Regularised Adaptive Prediction Sets (RAPS) [Angelopoulos *et al.*, 2020]. We recruit 30 participants on Prolific, paying them at a rate of £10 per hour prorated, and divide them into 2 equal groups. The first group is shown 18 images from the CIFAR-100 dataset alongside the model’s most probable prediction (Top-1). The second group is shown the same images but alongside a RAPS prediction set with error rate $\alpha = 0.1$. To understand the effect of set valued predictions on examples of varying difficulty, we divide the CIFAR-100 test dataset into 3 difficulty quantiles, where difficulty is defined as the entropy of the model predictive distribution. We select 6 images from each difficulty quantile. For each quantile, we show 2 images whose Top-1 prediction is incorrect but whose RAPS set contains the true label. This is consistent with the accuracy of the model ($\approx 65\%$) and lets us determine the effect of set valued predictions on examples on which the model is *almost* correct. Given these model predictions for each image, we ask participants in both groups to predict the correct class, rate their confidence in their predictions, and rate how useful they found the model predictions on that example. At the end of the survey, we ask participants to rate their overall trust in the model’s predictions. All ratings are on a scale from 1 – 10. We employ preliminary attention checks by first asking them to classify 3 easy examples, rejecting any participants who classify these examples incorrectly. We evaluate the statistical significance of our results using a two sample t-test.

Metric	Top-1	RAPS	<i>p</i> value	Effect Size
Accuracy	0.76 ± 0.05	0.76 ± 0.05	0.999	0.000
Reported Utility	5.43 ± 0.69	6.94 ± 0.69	0.003	1.160
Reported Confidence	7.21 ± 0.55	7.88 ± 0.29	0.082	0.674
Reported Trust	5.87 ± 0.81	8.00 ± 0.69	< 0.001	1.487

Table 1: Top-1 vs RAPS: All Examples

Metric	Top-1	RAPS	<i>p</i> value	Effect Size
Accuracy	0.90 ± 0.05	0.87 ± 0.07	0.486	0.273
Reported Utility	6.07 ± 0.94	6.35 ± 1.00	0.438	0.195
Reported Confidence	7.88 ± 0.65	8.82 ± 0.31	0.013	1.019

Table 2: Top-1 vs RAPS: Lowest Difficulty Quantile

Metric	Top-1	RAPS	<i>p</i> value	Effect Size
Accuracy	0.64 ± 0.07	0.66 ± 0.10	0.828	0.068
Reported Utility	5.30 ± 0.75	7.28 ± 0.69	0.001	1.432
Reported Confidence	6.64 ± 0.64	6.96 ± 0.78	0.4888	0.280

Table 3: Top-1 vs RAPS: Highest Difficulty Quantile

From Table 1, we see that Top-1 predictions result in statistically significant lower levels of trust ($p < 0.001$) and perceived utility ($p = 0.003$) compared to RAPS. However, both schemes result in similar accuracy and confidence in predictions. We also see that users find Top-1 and RAPS predictions equally useful for easy examples (Table 2). This makes sense because in such cases, the predictive set will be small and therefore comparable to a Top-1 prediction. However, users are more confident about their answers when they observe RAPS predictions. On the other hand, RAPS sets are perceived to be much more useful on hard examples, where Top-1 predictions will often be wrong.

Takeaway: *While there is seen to be no significant difference in team accuracy when shown either Top-1 or set-valued predictions, displaying set-valued predictions in human-AI teams results in higher reported utility of predictions as well as higher reported overall trust in the model.*

4 Proposed Approach

4.1 The Problem with CP Sets

In our experiments above, we showed users examples where the set sizes on CIFAR-100 are small enough to be considered useful. However, this may not always be the case, especially on tasks with large label spaces. For instance, a standard WideResNet model trained on CIFAR-100 ($\approx 65\%$ accuracy) with APS conformal prediction yields prediction sets with greater than 15 labels for 20% of examples. One option to mitigate this issue is to defer examples with the largest CP set sizes to an expert. However, this provides no guarantee that the expert will be able to classify them with sufficient accuracy. Furthermore, we also lose the finite sample coverage guarantees provided by contemporary CP methods, i.e. we cannot ascertain that $P(Y_{test} \notin \Gamma(X_{test})) \leq \alpha$.

4.2 Our Scheme

Our scheme, described in Algorithm 1, is centered around two components: a deferral policy $\pi(x) : \mathcal{X} \rightarrow \{0, 1\}$ and a CP method. The deferral policy is based on our knowledge of the expert’s strengths either acquired during training or a-priori. For example, if an expert is better at identifying brain tumors than our model, our policy should learn to defer those examples with high probability. Using this black box policy, we first prune our calibration dataset, removing all examples where our deferral policy recommends deferral. One could use any scheme in [Mozannar and Sontag, 2020; Okati *et al.*, 2021;

Wilder *et al.*, 2021] to learn a deferral policy. While Algorithm 1 specifies a deferral policy as an input, for some deferral methods (such as [Mozannar and Sontag, 2020]), the policy is trained alongside the model. In others, such as [Okati *et al.*, 2021], the policy is applied post-hoc. In this paper, we consider the former deferral policy: the D-CP algorithm for this is outlined in Algorithm 2 in the Supplementary Material. After training a model and a suitable deferral policy, we perform conformal calibration on this pruned dataset of non-deferred examples. In this procedure, for any predictive set $\Gamma(X_{test}, \tau_{cal})$ for an example X_{test} we can guarantee that:

$$1 - \alpha \leq P(Y \in \Gamma(X_{test}, \tau_{cal}) | \pi(X_{test}) = 0) \quad (1)$$

where 1 represents the action of deferral. From [Angelopoulos *et al.*, 2020], when the conformity scores are known to be almost surely distinct and continuous, we can also guarantee:

$$P(Y \in \Gamma(X_{test}, \tau_{cal}) | \pi(X_{test}) = 0) \leq 1 - \alpha + \frac{1}{n + 1} \quad (2)$$

where n is the size of the non-deferred calibration dataset. Because the deferral policy π probabilistically decides which unseen examples to defer, all non-deferred examples can be thought of as being generated from a data generating distribution $X \sim p(X | \pi(X) = 0)$. Any new test example X_{test} that is not deferred is therefore independently drawn from this distribution. Thus, $\{X_i\}_{i=1}^n \cup \{X_{test}\} \sim p(X_1, \dots, X_{test} | \pi(X_1), \dots, \pi(X_{test}) = 0)$ are exchangeable, thereby satisfying the coverage guarantee in Equation 1. To


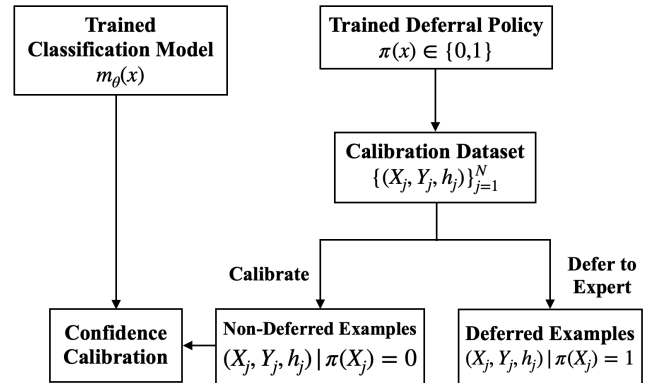

 Figure 2: D-CP: Test Phase given a deferral policy $\pi(X)$


Figure 3: D-CP: Training and Calibration Phase

show the utility of our scheme, a good deferral policy would guarantee that resulting predictive sets on non-deferred examples will contain fewer incorrect labels than before. We prove

Algorithm 1 General D-CP

Input: Classifier $m_\theta(x) \in \mathcal{R}^{|\mathcal{Y}|}$, Deferral Policy $\pi(x) \in \{0, 1\}$, Training Set \mathcal{D} , Expert $h(x) \in \mathcal{Y}$, Calibration Set \mathcal{D}_{cal} , Validation Set \mathcal{D}_{val} , Test Example x_{test} , Conformity Score Function $\tau(X, y)$, Loss function $l(m_\theta(x), y, h(x))$

Parameter: Number of Epochs N , Learning Rate γ , Error Tolerance α

```

1: for  $i \in \{1, \dots, N\}$  do
2:   for Batch  $\mathcal{B} \in \mathcal{D}$  do
3:      $\theta = \theta - \gamma \mathbb{E}_{(x,y) \in \mathcal{B}} [\nabla_{\theta} l(m_\theta(x), y, h(x))]$ 
4:   end for
5: end for
6:  $\mathcal{D}'_{\text{cal}} = \{(X, Y) | \pi(X) = 0, (X, Y) \in \mathcal{D}_{\text{cal}}\}$ 
7:  $\tau_{\text{cal}} = \text{Quantile}(\alpha, \{\tau(X_i, Y_i) | (X_i, Y_i) \in \mathcal{D}'_{\text{cal}}\})$ 
8: if  $\pi(X_{\text{test}}) = 0$  then
9:   return  $\Gamma(X_{\text{test}}, \tau_{\text{cal}}) = \{y | \tau(X_{\text{test}}, y) \geq \tau_{\text{cal}}\}$ 
10: else if  $\pi(X_{\text{test}}) = 1$  then
11:   Defer to Expert  $h(X_{\text{test}})$ 
12: end if
    
```

this formally in Theorem 1 in the Supplementary Material. While this technically applies to conformity scores $\tau(x, y)$ that are monotonic with respect to softmax probabilities (such as LAC), our subsequent experiments with other CP methods such as RAPS and APS suggest that our scheme generalises well across other classes of conformity score functions.

5 Toy Example

One way to combine conformal prediction and deferral is to only perform CP on “easy” examples and defer the “hard” examples. An “easy” example would be one which the model is confident on and a “hard” example is the converse. This can lead to smaller sets. To demonstrate the intuition, we generate equiprobable synthetic data using a Mixture of Gaussians (MoG) model. Each datapoint is generated from one of 4 Gaussians - $\mathcal{N}(1, 1)$, $\mathcal{N}(1, -1)$, $\mathcal{N}(-1, 1)$, and $\mathcal{N}(-1, -1)$ - and we wish to infer class memberships. We first train a multi-layer perceptron (MLP) on 1000 training samples (not shown) to infer the decision boundaries. Then, using a held out calibration set, we perform CP with error tolerance $\alpha = 5\%$ using the Least Ambiguous Classifiers (LAC) method [Sadinle *et al.*, 2016]. In this method, we use the model softmax probabilities $p(y|x) = \tau(x, y)$ as conformity scores. Figure 5 (top) shows a 1-D scatter plot of conformity scores assigned to ground truth labels in the toy dataset.

Figure 4 shows the resulting test datapoints colored according to their true classes with model decision boundaries overlaid. We see that points closer to the decision boundary have larger predictive set sizes, reflecting their inherent uncertainty. If we defer points with conformity scores in the bottom 15th percentile (naive decision policy) as in Figure 5, the α threshold conformity score will increase. From Figure 5 (Right), for non-deferred examples, this increases the threshold for including labels in the set, resulting in more confident sets for the same error control. However, this naive deferral method, whilst ensuring small set sizes on the remaining examples, does not take into account the expertise of the expert involved.

Furthermore, we assumed access to ground truth labels for test examples, which is not practical. We can engage the expert in a better manner and approximate the idea of the toy example by learning a *deferral policy* which incorporates estimates of expert ability as well as machine difficulty. This scheme makes an implicit assumption that the expert is a) either better than the model on average or b) not necessarily better than the model on average, but is proficient in classifying certain subgroups of examples. In these situations, our deferral policy is more likely to defer examples that a model is less confident on. Given these assumptions about an expert, we are assured of lower predictive set sizes on non deferred examples.




Figure 4: (Left): Toy dataset comprising of datapoints belonging to one of 4 classes along with overlaid model decision boundaries. **The size of the datapoints indicates their predictive set sizes.** (Right) We defer the $\beta = 0.15$ proportion of examples with the lowest ground truth conformity scores. Doing so increases the value of the 5th percentile conformity score of the remaining examples, causing CP set sizes of examples to be smaller. Note that we have not altered the model in this process.




Figure 5: (Left): 1-D scatter plot of all ground truth conformity scores $\tau = p(Y_i | X_i)$ on the toy calibration dataset in Figure 4. We defer $\beta = 15\%$ of examples with the lowest τ . Both values of τ_{cal} maintain 95% coverage on their respective datasets. (Right): Class probabilities for the test **green starred example**. For the predictive set, we include all scores which are greater than the threshold τ_{cal} . Thus, the predictive set $\{1, 2, 3\}$ gives 95% coverage for the original dataset. On the non-deferred dataset, the set $\{1\}$ gives 95% coverage.

6 Experiments with D-CP

To validate our approach, we perform experiments with synthetic expert labels on the CIFAR-100 dataset and real expert labels on the CIFAR-10H [Peterson *et al.*, 2019] dataset¹. Because the CIFAR-10H dataset contains expert labels only on

¹Our code is hosted at <https://github.com/cambridge-mlg/d-cp>.

the CIFAR-10H validation set, we employ the approach in [Mozannar and Sontag, 2020] and train a binary classifier to predict examples where the expert is correct. We then provide synthetic expert labels $\mathcal{I}_{h(x)=y}$ or $\mathcal{I}_{h(x)\neq y}$ for examples the training set according to whether the expert errs on them. Note that, in line with the assumption made in this paper, the experts chosen in this setting are, on average, better than the model trained. We consider 2 scenarios:

- We have access to a single expert’s annotations. For CIFAR-100, we generate a synthetic expert with 70% accuracy. To motivate this choice, we ran a control study where we asked 20 participants to classify 15 randomly chosen CIFAR-100 examples. We found participants had average accuracy of 69% with a standard error of $\approx 2.5\%$. For the CIFAR-10H dataset, we randomly sample a label from the predictive distribution provided.
- We have access to multiple expert annotations. This is an ensemble of the above experts, and the predicted class is chosen through majority voting for both datasets. For the CIFAR-100, we generate predictions from 5 experts.

Our deferral policy is based on the loss function in [Mozannar and Sontag, 2020]. We train a WideResNet [Zagoruyko and Komodakis, 2016] classifier $m_\theta(x) : \mathcal{X} \rightarrow \mathcal{Y} \cup \perp$ on CIFAR-10H and CIFAR-100 for 5 and 10 epochs respectively using the learning rate schedule in [Mozannar and Sontag, 2020]. \perp represents the action of deferral to an expert $h(x)$. We modify the loss function used in this work as below:

$$\mathcal{L}_{CE}(h, x, y, m_\theta) = -(\mathcal{I}_{h(x)\neq y} + \alpha \mathcal{I}_{h(x)=y}) \log m_\theta(x)_y - \beta \mathcal{I}_{h(x)=y} \log m_\theta(x)_\perp$$

where we set $\alpha = 1$ and vary the $\beta \in [0, 1]$ penalty term to control the deferral rate. The policy $\pi(x)$ is therefore:

$$\pi(x) = \begin{cases} 1 & \text{argmax}_{y \in \mathcal{Y} \cup \perp} m_\theta(x)_y = |\mathcal{Y} \cup \perp| \\ 0 & \text{otherwise} \end{cases}$$

To compute conformity scores, we renormalize the softmax probabilities for examples where $\pi(x) = 0$ using Bayes’ rule:

$$p(y|x, \pi(x) = 0, \theta) = \frac{p(y \neq |\mathcal{Y} \cup \perp| | x, y, \theta) p(y|x, \theta)}{p(y \neq |\mathcal{Y} \cup \perp| | x, \theta)} = \frac{p(y|x, \theta)}{p(y \neq |\mathcal{Y} \cup \perp| | x, \theta)}$$




Figure 6: Cumulative CP and D-CP Set Size Distribution of Non-Deferred Examples in the CIFAR-100 dataset ($\alpha = 0.05$, deferral rate $b = 0.2$, Single Expert).

In our experiments, we did not notice any statistically significant difference in accuracy of non-deferred examples or predictive set sizes when employing multiple experts as opposed to a singular expert, at least in the deferral rate regimes tested. Because we are performing experiments in the low deferral rate regime, it is likely that the deferral scheme defers similar examples to both expert types - examples the model is sure the expert(s) will get right. Thus, in Table 6, the classifier accuracy and predictive set sizes are representative for both singular and multiple experts. However, we benefit from increased team accuracy by using ensemble voting across multiple experts. In addition, per Table 6 and Figure 6, our scheme ensures smaller set sizes across all conformal methods and deferral rates tested. Increasing the deferral rate reduces the predictive set size. In Figure 7, the model and expert have a mutually beneficial relationship: the model provides smaller predictive sets on examples the expert is more uncertain on and defers examples it is less certain of than an expert.

Takeaway: *D-CP provides smaller predictive set sizes on non-deferred examples for the same level of coverage. For the policy in [Mozannar and Sontag, 2020], while the number of experts does not make a difference in the resulting predictive set size of non-deferred examples, having more experts predict through majority voting improves the team accuracy.*

7 Evaluation on Experts

Our second human subject experiment focuses on establishing the value of smaller set predictions and learning to defer - the 2 promises of D-CP. We choose another set of 15 examples from the CIFAR-100 test set for which we generate RAPS prediction sets with error rate $\alpha = 0.1$ and D-RAPS prediction sets with deferral rate 0.2 and error rate $\alpha = 0.1$. We select 12 non-deferred examples at random wherein the D-RAPS predictive set is smaller than the RAPS predictive set, but the ground truth labels are contained in both sets. Lastly, we choose the remaining 3 deferred examples where the model is underconfident, i.e. the ground truth label is not in the RAPS set. This aims to establish the value of deferral in situations where the model may provide misleading predictions. We ask participants the same questions as in Section 3 and follow a similar recruitment procedure as in Section 3 (60 participants total, 2 groups, reward of £10 per hour prorated).

Metric	D-RAPS	RAPS	p value	Effect Size
Accuracy	0.76 ± 0.08	0.67 ± 0.05	0.002	0.832
Reported Utility	7.93 ± 0.39	6.32 ± 0.60	< 0.001	1.138
Reported Confidence	7.31 ± 0.29	7.28 ± 0.29	0.862	0.046
Reported Trust	8.00 ± 0.45	6.87 ± 0.61	0.006	0.754

Table 4: D-RAPS vs RAPS: All Examples

Metric	D-RAPS	RAPS	p value	Effect Size
Accuracy	0.88 ± 0.05	0.81 ± 0.04	0.058	0.508
Reported Utility	7.93 ± 0.39	6.19 ± 0.62	< 0.001	1.211
Reported Confidence	7.78 ± 0.33	7.31 ± 0.34	0.059	0.507

Table 5: D-RAPS vs RAPS: Non-Deferred Examples

Tables 4 and 5 suggest that there is a statistically significant increase in expert accuracy when the D-CP scheme is

Deferral Rate b	Classifier Accuracy	Team Accuracy (Single Expert)	Team Accuracy (Multiple Experts)	Predictive Set Size of Non-Deferred Examples		
				RAPS	APS	LAC
0	65.18 \pm 0.30	65.18 \pm 0.30	65.18 \pm 0.30	3.75 \pm 0.06	4.61 \pm 0.08	3.26 \pm 0.03
0.05	68.39 \pm 0.31	68.04 \pm 0.32	68.91 \pm 0.33	3.22 \pm 0.05	4.16 \pm 0.06	2.48 \pm 0.03
0.10	69.92 \pm 0.24	69.95 \pm 0.31	71.53 \pm 0.35	2.81 \pm 0.05	4.05 \pm 0.06	2.13 \pm 0.04
0.20	72.98 \pm 0.30	72.25 \pm 0.30	78.99 \pm 0.40	2.36 \pm 0.07	2.93 \pm 0.10	2.07 \pm 0.03

Deferral Rate b	Classifier Accuracy	Team Accuracy (Single Expert)	Team Accuracy (Multiple Experts)	Predictive Set Size of Non-Deferred Examples		
				RAPS	APS	LAC
0	82.02 \pm 0.55	82.02 \pm 0.55	82.02 \pm 0.55	1.91 \pm 0.03	2.83 \pm 0.05	2.47 \pm 0.12
0.05	84.41 \pm 0.69	84.31 \pm 0.65	84.64 \pm 0.30	1.87 \pm 0.08	2.76 \pm 0.10	2.25 \pm 0.08
0.10	86.12 \pm 0.67	86.53 \pm 0.68	88.12 \pm 0.61	1.73 \pm 0.08	2.56 \pm 0.07	1.90 \pm 0.15
0.20	88.97 \pm 0.50	89.43 \pm 0.64	91.46 \pm 0.32	1.49 \pm 0.06	2.13 \pm 0.11	1.51 \pm 0.09

Table 6: Set Size, Overall Team Accuracy, and Classifier Accuracy on non-deferred examples on the CIFAR-100 (top) and CIFAR-10H (bottom) datasets for the deferral scheme in [Mozannar and Sontag, 2020] with $\alpha = 0.1$ (5 Trials, 95% CI). Even with low deferral rates, we not only obtain smaller set sizes, but also benefit from increased human-AI team accuracy compared to the baseline ($b = 0$). While having multiple experts does not further improve the predictive set size for this deferral policy, we benefit from further improved team accuracy.

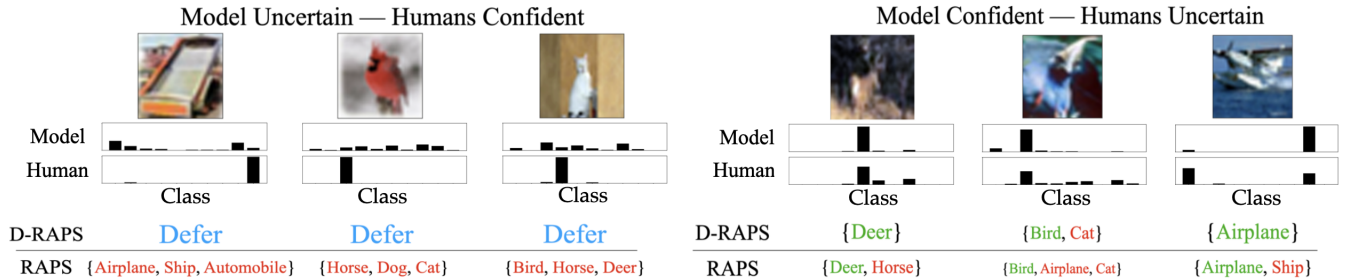


Figure 7: D-RAPS vs RAPS on CIFAR-10H examples ($\alpha = 0.05$, $b = 0.2$). Deferring whenever experts are more confident than the model yields smaller sets on examples where the model is more confident than the expert. Thus, we leverage both the model and the expert's strengths.

employed, with borderline significance on non-deferred examples. Even though participants did not perform as well on deferred examples in general, we noticed that their accuracy was still higher than when they were shown CP sets, which contained misleading labels. Equally interestingly, on examples where both RAPS and D-RAPS sets contain the ground truth label (i.e. the non-deferred examples), the perceived utility of D-CP sets is higher ($p < 0.001$). As D-RAPS sets are smaller, this shows that, for the same confidence level, smaller set sizes are more useful to experts and therefore a preferred choice for human-AI teams. Table 4 also shows a statistically significant difference in reported trust in the model between D-RAPS and RAPS. These are important considerations for real world human-AI teams. We provide further results in the Supplementary Material which show that participants display a bias towards incorrect predictions shown in larger CP sets, warranting caution when deploying models with large CP sets in human-AI teams.

Takeaway: *There are statistically significant improvements increases in reported utility, trust, and accuracy in the model when the D-CP scheme is employed.*

8 Conclusion

In this paper, we explored the importance of set valued predictions for human-AI teams. We first showed experts find CP predictive sets more useful than Top-1 predictions. However, CP set sizes can be very large for some examples, especially in large label spaces. Thus, we motivate the need for combining the ideas of learning to defer and set valued predictions. We introduce D-CP, a general practical scheme that defers some examples and performs CP on others. Empirical and theoretical evidence shows that the scheme provides smaller set sizes on non-deferred examples for any CP method compared to performing CP on all examples. The scheme allows the model and expert to have a mutually beneficial relationship by leveraging the expert and the model's respective strengths. Our human subject experiments show that, compared to CP, experts find the smaller D-CP predictive sets more useful, the model more trustworthy, and are more accurate. We hope that this informs a) future research on improved deferral policies that consider the predictive uncertainty of the model and b) larger scale human evaluations that uncover specific, desirable properties of a predictive set.

Acknowledgements

The authors thank the anonymous reviewers for their thorough feedback. UB acknowledges support from DeepMind and the Leverhulme Trust via the Leverhulme Centre for the Future of Intelligence (CFI), and from the Mozilla Foundation. AW acknowledges support from a Turing AI Fellowship under grant EP/V025279/1, The Alan Turing Institute, and the Leverhulme Trust via CFI.

References

- [Angelopoulos and Bates, 2021] Anastasios N Angelopoulos and Stephen Bates. A Gentle Introduction to Conformal Prediction and Distribution-Free Uncertainty Quantification. *arXiv preprint arXiv:2107.07511*, 2021.
- [Angelopoulos *et al.*, 2020] Anastasios Nikolas Angelopoulos, Stephen Bates, Michael Jordan, and Jitendra Malik. Uncertainty Sets for Image Classifiers Using Conformal Prediction. In *International Conference on Learning Representations*, 2020.
- [Bansal *et al.*, 2019] Gagan Bansal, Besmira Nushi, Ece Kamar, Daniel S Weld, Walter S Lasecki, and Eric Horvitz. Updates in Human-AI Teams: Understanding and Addressing the Performance/Compatibility Tradeoff. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 2429–2437, 2019.
- [Bartlett and Wegkamp, 2008] Peter L Bartlett and Marten H Wegkamp. Classification with a reject option using a hinge loss. *Journal of Machine Learning Research*, 9(8), 2008.
- [Bellotti, 2021] Anthony Bellotti. Optimized conformal classification using gradient descent approximation. *CoRR*, abs/2105.11255, 2021.
- [Bondi *et al.*, 2022] Elizabeth Bondi, Raphael Koster, Hannah Sheahan, Martin Chadwick, Yoram Bachrach, Taylan Cemgil, Ulrich Paquet, and Krishnamurthy Dvijotham. Role of Human-AI Interaction in Selective Prediction. *Proceedings of the AAAI Conference on Artificial Intelligence*, 2022.
- [Cortes *et al.*, 2016] Corinna Cortes, Giulia DeSalvo, and Mehryar Mohri. Learning with Rejection. In Ronald Ortner, Hans Ulrich Simon, and Sandra Zilles, editors, *Algorithmic Learning Theory*, pages 67–82, Cham, 2016. Springer International Publishing.
- [Lei *et al.*, 2016] Jing Lei, Max G’Sell, Alessandro Rinaldo, Ryan J. Tibshirani, and Larry Wasserman. Distribution-Free Predictive Inference For Regression. *Journal of the American Statistical Association*, 113:1094–1111, 4 2016.
- [Lundberg *et al.*, 2018] Scott M. Lundberg, Bala Nair, Monica S. Vavilala, Mayumi Horibe, Michael J. Eisses, Trevor Adams, David E. Liston, Daniel King-Wai Low, Shu-Fang Newman, Jerry Kim, and Su-In Lee. Explainable Machine-Learning Predictions for the Prevention of Hypoxaemia During Surgery. *Nature Biomedical Engineering*, 2(10):749–760, October 2018.
- [Madras *et al.*, 2018] David Madras, Toniann Pitassi, and Richard Zemel. Predict Responsibly: Improving Fairness and Accuracy by Learning to Defer. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 6150–6160, 2018.
- [Mozannar and Sontag, 2020] Hussein Mozannar and David Sontag. Consistent Estimators for Learning to Defer to an Expert. In *International Conference on Machine Learning*, pages 7076–7087. PMLR, 2020.
- [Okati *et al.*, 2021] Nastaran Okati, Abir De, and Manuel Gomez-Rodriguez. Differentiable Learning Under Triage. In *Advances in Neural Information Processing Systems*, 2021.
- [Papadopoulos, 2008] Harris Papadopoulos. Inductive Conformal Prediction: Theory and Application to Neural Networks. In Paula Fritzsche, editor, *Tools in Artificial Intelligence*, chapter 18. IntechOpen, Rijeka, 2008.
- [Peterson *et al.*, 2019] Joshua Peterson, Ruairidh Battleday, Thomas Griffiths, and Olga Russakovsky. Human Uncertainty Makes Classification More Robust. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9616–9625, 2019.
- [Romano *et al.*, 2020] Yaniv Romano, Matteo Sesia, and Emmanuel Candes. Classification with Valid and Adaptive Coverage. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 3581–3591. Curran Associates, Inc., 2020.
- [Sadinle *et al.*, 2016] Mauricio Sadinle, Jing Lei, and Larry Wasserman. Least Ambiguous Set-Valued Classifiers with Bounded Error Levels. *Journal of the American Statistical Association*, 114:223–234, 9 2016.
- [Straitouri *et al.*, 2022] Eleni Straitouri, Lequn Wang, Nastaran Okati, and Manuel Gomez Rodriguez. Provably improving expert predictions with conformal prediction. *CoRR*, abs/2201.12006, 2022.
- [Stutz *et al.*, 2022] David Stutz, Krishnamurthy Dj Dvijotham, Ali Taylan Cemgil, and Arnaud Doucet. Learning optimal conformal classifiers. In *ICLR*, 2022.
- [Tibshirani *et al.*, 2019] Ryan J. Tibshirani, Rina Foygel Barber, Emmanuel J. Candes, and Aaditya Ramdas. Conformal Prediction Under Covariate Shift. *Advances in Neural Information Processing Systems*, 32, 4 2019.
- [Vovk *et al.*, 2005] Vladimir Vovk, Alex Gammerman, and Glenn Shafer. *Algorithmic Learning in a Random World*. Springer, 01 2005.
- [Wilder *et al.*, 2021] Bryan Wilder, Eric Horvitz, and Ece Kamar. Learning to Complement Humans. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI’20*, 2021.
- [Zagoruyko and Komodakis, 2016] Sergey Zagoruyko and Nikos Komodakis. Wide Residual Networks. In Edwin R. Hancock Richard C. Wilson and William A. P. Smith, editors, *Proceedings of the British Machine Vision Conference (BMVC)*, pages 87.1–87.12. BMVA Press, September 2016.