
PAC Prediction Sets for Meta-Learning

Sangdon Park
School of Cybersecurity and Privacy
Georgia Institute of Technology
sangdon@gatech.edu

Edgar Dobriban
Dept. of Statistics & Data Science
The Wharton School
University of Pennsylvania
dobriban@wharton.upenn.edu

Insup Lee
Dept. of Computer & Info. Science
PRECISE Center
University of Pennsylvania
lee@cis.upenn.edu

Osbert Bastani
Dept. of Computer & Info. Science
PRECISE Center
University of Pennsylvania
obastani@seas.upenn.edu

Abstract

Uncertainty quantification is a key component of machine learning models targeted at safety-critical systems such as in healthcare or autonomous vehicles. We study this problem in the context of meta learning, where the goal is to quickly adapt a predictor to new tasks. In particular, we propose a novel algorithm to construct PAC prediction sets, which capture uncertainty via sets of labels, that can be adapted to new tasks with only a few training examples. These prediction sets satisfy an extension of the typical PAC guarantee to the meta learning setting; in particular, the PAC guarantee holds with high probability over future tasks. We demonstrate the efficacy of our approach on four datasets across three application domains: mini-ImageNet and CIFAR10-C in the visual domain, FewRel in the language domain, and the CDC Heart Dataset in the medical domain. In particular, our prediction sets satisfy the PAC guarantee while having smaller size compared to other baselines that also satisfy this guarantee.

1 Introduction

Uncertainty quantification is a key component for safety-critical systems such as healthcare and robotics, since it enables agents to account for risk when making decisions. Prediction sets are a promising approach, since they provide theoretical guarantees when the training and test data are i.i.d. [1, 2, 3, 4]; there have been extensions to handle covariate shift [5, 6, 7] and online learning [8].

We consider the meta learning setting, where the goal is to adapt an existing predictor to new tasks using just a few training examples [9, 10, 11, 12, 13]. Most approaches follow two steps: meta learning (i.e., learn a predictor for fast adaptation) and adaptation (i.e., adapt this predictor to the new task). While there has been work on prediction sets in this setting [14], a key shortcoming is that their guarantees do not address some important needs in meta learning; in particular, they do not hold with high probability over future tasks. For instance, in autonomous driving, the future tasks might be new regions, and correctness guarantees ought to hold for each region; in healthcare, future tasks might be new patient cohort, and correctness guarantees ought to hold for each cohort.

We propose an algorithm for constructing probably approximately correct (PAC) prediction sets for the meta learning setting. As an example, consider learning an image classifier for self-driving cars as our use case (Figure 1). Labeled images for image classifier learning are collected, with different regions representing different but related tasks (e.g., New York, Pennsylvania, and California). The

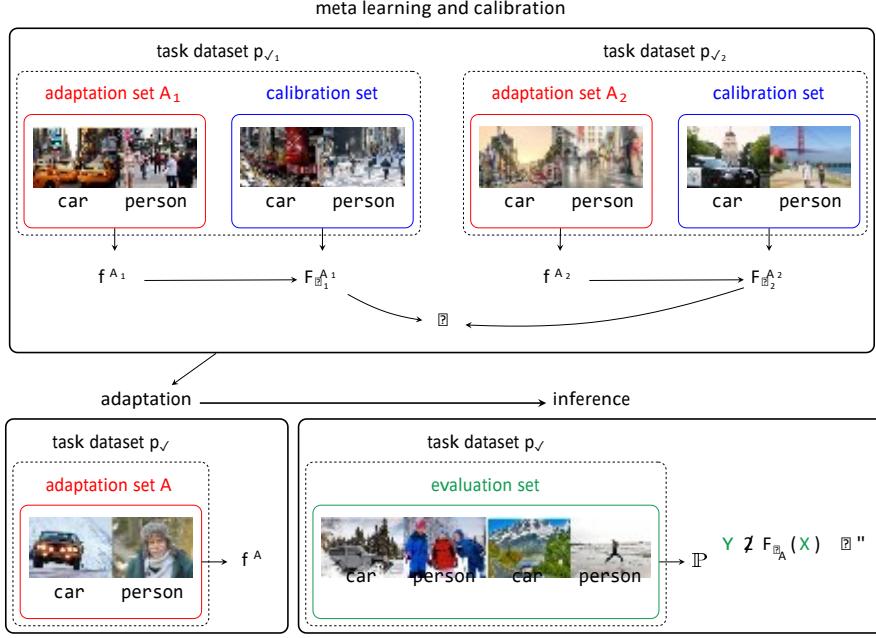


Figure 1: Learning for PAC prediction sets with fast adaptation. In meta learning and calibration, the PAC prediction set parameter θ is chosen from multiple task datasets given a meta-learned and adapted score function f^{A_i} using any meta-learning algorithm and an adapted PAC prediction set $F_{\theta_i}^{A_i}$ (e.g., an image classifier and a prediction set are learned over labeled images from various regions, including New York). In adaptation, the score function is adapted to the new task distribution f^A (e.g., the image classifier is adapted using labeled images from Alaska). In inference, by using the same parameter θ used without additional calibration, the constructed and adapted prediction set F_{θ}^A has a correctness guarantee for most of future examples (e.g., the prediction set along with the adapted image classifier has a correctness guarantee over most images from Alaska). Importantly, the correctness guarantee (error below " ϵ ") is achieved with probability at least (1) $1 - \delta$ over the calibration sets in the meta learning phase and (2) $1 - \epsilon$ over an adaptation set in the adaptation phase. This explains why we use a three-parameter PAC-style guarantee for meta learning.

goal is to deploy this learned image classifier to a new target region (e.g., Alaska) in a cost-efficient way, e.g., minimizing label gathering cost on labeled images from Alaska for the classifier adaptation.

Due to the safety-critical nature of the task, it is desirable to quantify uncertainty in a way that provides rigorous guarantees, where the guarantees are adapted to a new task (e.g., Alaska images)—i.e., they hold with high probability specifically for future observations from this new task. In contrast, existing approaches only provide guarantees on average across future tasks [14].

Our proposed PAC prediction set starts with a meta-learned and potentially adaptable score function as our base predictor (using a standard meta learning or few shot classification approach). Then, we learn the parameter of a prediction set for meta “calibration” by using calibration task distributions (e.g., labeled images from Texas and Massachusetts). Given a test task distribution (e.g., labeled images from Alaska), we only have a few labeled examples for adapting the score function to this target task without further updating the prediction set parameter, which potentially saves additional labeled examples. We prove that the proposed algorithm satisfies the PAC guarantee tailored to the meta learning setup. Additionally, we demonstrate the validity of the proposed algorithm over four datasets across three application domains: mini-ImageNet [10] and CIFAR10-C [15] in the visual domain, FewRel [16] in the language domain, and CDC Heart Dataset [17] in the medical domain.

2 Related Work

Meta learning for few-shot learning. Meta-learning follows the paradigm of learning-to-learn [9], usually formulated as few-shot learning: learning or adapting a model to a new task or distribution

by leveraging a few examples from the distribution. Existing approaches propose neural network predictors or adaptation algorithms [10, 11, 12, 13] for few-shot learning. However, there is no guarantee that these approaches can quantify uncertainty properly, and thus they do not provide guarantees for probabilistic decision-making.

Conformal prediction. The classical goal in conformal prediction (and inductive conformal prediction) [18, 2, 19] is to find a prediction set that contains the label with a given probability, with respect to randomness in both calibration examples and a test example, where a score function (or non-conformity score) is given for the inductive conformal prediction. This property is known as marginal coverage. If the labeled examples are independent and identically distributed (or more precisely exchangeable), prediction sets constructed via inductive conformal prediction have marginal coverage [19, 2]. However, the identical distribution assumption can fail if the covariates of a test datapoint are differently distributed from the training covariates, i.e., covariate shift occurs. Assuming that the likelihood ratio is known, the weighted split conformal prediction set [5] satisfies marginal coverage under covariate shift. So far, we have considered only two distributions, a calibration distribution and test distribution. In meta learning, a few examples from multiple distributions are given, providing a different setup. Assuming exchangeability of these distributions, conformal prediction sets for meta learning have been constructed, covering over a random sample from a new randomly drawn distribution, exchangeable with the training distributions [14]. In statistics, two-layer hierarchical models consider examples drawn from multiple distributions, while assuming examples and distributions are independent and identically distributed; these are identical to the model of meta-learning we study. Several approaches to construct prediction sets (e.g., double conformal prediction) satisfy the marginal coverage guarantee [20].

PAC prediction sets. Marginal coverage is unconditional over calibration data, and does not hold conditionally over calibration data (i.e., with specified probability for a fixed calibration set). Training conditional conformal prediction was introduced as a way to construct prediction sets with PAC guarantees [21]; Specifically, for independent and identically distributed calibration examples, they construct a prediction set that contains a true label with high probability with respect to calibration examples; thus, these prediction sets cover the labels for most future test examples. This approach applies classical techniques for constructing tolerance regions [1, 22] to the non-conformity score to obtain these guarantees. Recent work casts this technique in the language of learning theory [3], informing our extension to the meta-learning setting; thus, we adopt this formalism in our exposition. PAC prediction sets can be constructed in other learning settings—e.g., under covariate shift (assuming the distributions are sufficiently smooth [6] or well estimated [7]). In the meta learning setting, the meta conformal prediction set approach also has a PAC property [14], but does not control error over the randomness in the adaptation test examples. In contrast, we propose a prediction set satisfying a meta-PAC property—i.e., separately conditional with respect to the calibration data and the test adaptation examples; see Section 4.1 for details.

3 Background

3.1 PAC Prediction Sets

We describe a training-conditional inductive conformal prediction strategy for constructing probably approximately correct (PAC) prediction sets [21], using the reformulation as a learning theory problem in [3]. Let X be a set of examples or features and Y be a set of labels. We consider a given score function $f : X \rightarrow Y \setminus R_0$, which is higher if $y \in Y$ is the true label for $x \in X$. We consider a parametrized prediction set; letting $\theta \in R_0$, we define a prediction set $F_\theta : X \rightarrow 2^Y$ as follows:

$$F_\theta(x) := \{y \in Y \mid f(x, y) \geq \theta\}.$$

Consider a family of distributions $(p_\sqrt{})$ over $Z := X \rightarrow Y$, parametrized by $\sqrt{} \in \mathcal{V}$. For an unknown $\sqrt{}$, we observe independent and identically distributed (i.i.d.) calibration datapoints $S \sim p_\sqrt{}^n$. We define a prediction set, or equivalently the corresponding parameter θ , to be " $\sqrt{}$ -correct", as follows:

Definition 1. A parameter $\theta \in R_0$ is " $\sqrt{}$ -correct" for $p_\sqrt{}$ if

$$P\{Y \in F_\theta(X)\} \geq 1 - \epsilon, \quad (1)$$

where the probability is taken over $(X, Y) \sim p_\sqrt{}$.

Letting $T^*(\sqrt{\epsilon})$ be the set of θ satisfying (1) for $p_\sqrt{\epsilon}$, we can write $\theta \in T^*(\sqrt{\epsilon}) \cap R_0$ if θ is " $\sqrt{\epsilon}$ -correct" for $p_\sqrt{\epsilon}$. If an estimator of θ is " $\sqrt{\epsilon}$ -correct" on most random calibration sets $S \leftarrow p_\sqrt{\epsilon}^1$, we say it is PAC.

Definition 2. An estimator $\hat{\theta} : Z^n \rightarrow R_0$ is " $\sqrt{\epsilon}$ -probably approximately correct (PAC)" for $p_\sqrt{\epsilon}$ if

$$P\{\hat{\theta}(S) \in T^*(\sqrt{\epsilon})\} \geq 1 - \epsilon,$$

where the probability is taken over S .

Along with the PAC guarantee, the prediction set should be small in size. Here, by increasing the scalar parameterization, we see that the expected prediction set size is decreasing with an increasing prediction set error (i.e., $\mathbb{P}_{X,Y}\{Y \notin F_\theta(X)\}$). Thus, a practical " $\sqrt{\epsilon}$ -PAC" algorithm that implements the estimator $\hat{\theta}$ maximizes θ while satisfying a PAC constraint to minimize the expected prediction set size [3].

3.2 Meta Learning

Meta learning aims to learn a score function that can be conveniently adapted to a new distribution, also called a task. Consider a task distribution p_τ over \mathcal{T} , and M training task distributions $p_{\sqrt{1}}, \dots, p_{\sqrt{M}}$, where $\sqrt{1}, \dots, \sqrt{M}$ is an i.i.d. sample from p_τ . Here, we mainly consider meta learning with adaptation (e.g., few-shot learning) to learn the score function, while having learning without adaptation (e.g., zero-shot learning) as a special case of our main setup.

For meta learning without adaptation, labeled examples are drawn from each training task distribution to learn the score function $f : X \rightarrow Y \setminus R_0$, and the same score function is used for inference (i.e., at test time). For meta learning with adaptation, t adaptation examples are also drawn from each training task distribution to learn the score function $f : Z^t \rightarrow X \rightarrow Y \setminus R_0$, and the same number of adaptation examples are drawn from a new test task distribution for adaptation. Given an adaptation set $A \in Z^t$, the adapted score function is denoted by $f^A : X \rightarrow Y \setminus R_0$, where $f^A(x, y) = f(A, x, y)$. We consider Y large enough to include all label classes, some of which possibly seen during testing but not during training. This formulation includes learning setups (e.g., few-shot learning) where unseen label classes are observed during inference, by considering different distributions over labeled examples during training and inference.

4 Meta PAC Prediction Sets

We propose a PAC prediction set for meta learning. In particular, we consider a meta-learned score function, which is typically trained to be a good predictor after adaptation to the test distribution. Our prediction set has a correctness guarantee for future test distributions.

There are three sources of randomness: (1) an adaptation set and a calibration set for each calibration task, (2) an adaptation set for a new task, and (3) an evaluation set for the same new task. We control the three sources of randomness and error by ϵ , δ , and η , respectively, aiming to construct a prediction set that contains labels from a new task. Given a meta-learned model, the proposed algorithm consists of two steps: constructing a prediction set—captured by a specific parameter to be specified below—for each calibration task (related to ϵ and δ), and then constructing a prediction interval over parameters (related to δ and η).

Suppose we have prediction sets for each task, parametrized by some threshold parameters. Suppose that, for each task, the parameters only depend on an adaptation set for the respective task. As tasks are i.i.d., the parameters also turn out to be i.i.d.. The distribution of parameters represents a possible range of the parameter for a new task. In particular, if for each task, the prediction set contains the true label a fraction $1 - \epsilon$ of the time, the distribution of the parameters suggests how one may achieve " ϵ -correct" prediction sets for a new task. Based on this observation, we construct a prediction interval that contains a fraction $1 - \delta$ of " ϵ -correct" prediction set parameters. However, the parameters of prediction sets need to be estimated from calibration tasks, and the estimated parameters depend on the randomness of calibration tasks. Thus, we construct a prediction interval over parameters which accounts for this randomness. The interval is chosen to be correct (i.e., to contain " ϵ -correct" prediction set parameters for a fraction $1 - \eta$ of new tasks) with probability at least $1 - \delta$ over the randomness of calibration tasks.

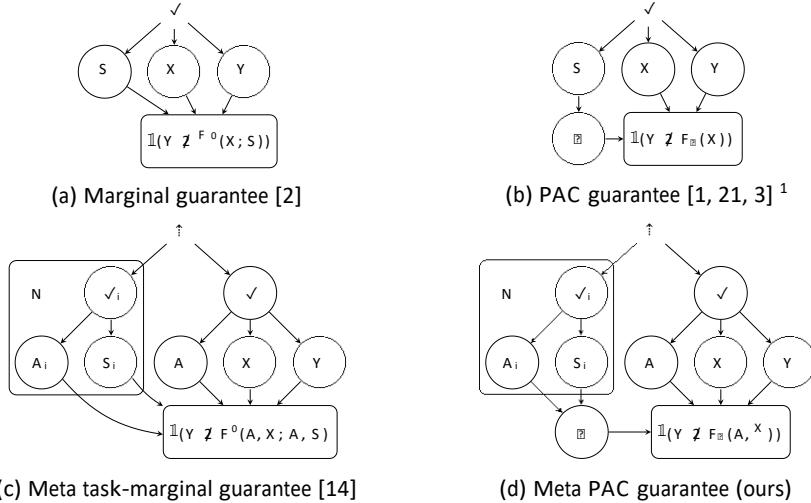


Figure 2: Comparison among different guarantees on prediction sets. An edge $A \rightarrow B$ between random variables A, B represents that B is drawn conditionally on A , and a plate with the symbol N represents that there are N copies of the variables in it, indexed by $i = 1, \dots, N$. Each stripe pattern with a different color indicates that the corresponding random variables are simultaneously conditioned on in the correctness guarantee; F^0 denotes a conformal prediction set. Thus, for instance, in (d), there is a $1 - \epsilon$ probability of correctness with respect to the test datapoint (x, y) , a $1 - \epsilon$ probability of correctness with respect to the test task and its adaptation data (\checkmark, A) , and a $1 - \epsilon$ probability of correctness with respect to the calibration tasks, their calibration data, and their adaptation data (\checkmark_i, S_i, A_i) for $i = 1, \dots, N$. The proposed meta PAC prediction set (2d) has the most fine-grained control over the correctness guarantee conditioning over calibration tasks, the test task and adaptation set, and the test data.

4.1 Setup and Problem

We describe the problem on PAC prediction set construction for meta learning with adaptation; see Appendix A for a simpler problem setup without adaptation, a special case of our main problem.

Setup. In the setting of Section 3.2, the training samples from some training distributions $p_i, i = 1, \dots, M$ are used to learn a score function via an arbitrary meta learning method, including metric-based meta-learning [11] and optimization-based meta-learning [12]. We denote the given meta-learned score function by f . Additionally, we have N calibration distributions $p_{\checkmark_1}, \dots, p_{\checkmark_N}$, where $\checkmark_1, \dots, \checkmark_N$ is an i.i.d. sample from p_{\checkmark} , i.e., $\checkmark := (\checkmark_1, \dots, \checkmark_N) \sim p_{\checkmark}^N$. From each calibration distribution p_{\checkmark_i} , we draw n i.i.d. labeled examples for a calibration set, i.e., $S_i = (S_{i1}, \dots, S_{iN}) \sim p_{\checkmark_i}^n \cdots p_{\checkmark_N}^n$ and $S \supseteq S_i$. Additionally, we draw t i.i.d. labeled examples for an adaptation set, i.e., $A = (A_1, \dots, A_N) \sim p_{\checkmark_1}^t \cdots p_{\checkmark_N}^t$ and $A \supseteq A_i$; we use them to construct the meta PAC prediction set. Finally, we consider a test distribution p_{\checkmark} where $\checkmark \sim p_{\checkmark}^N$. From it, we draw i.i.d. labeled adaptation examples A , i.e., $A \sim p_{\checkmark}^t$, and i.i.d. evaluation examples.

Problem. Our goal is to find a prediction set $F_{\checkmark} : Z^t \rightarrow X \times \{0, 1\}$ parametrized by $\checkmark \in \mathcal{R}_0$ which satisfies a meta-learning variant of the PAC definition. In particular, consider an estimator $\checkmark \hat{\rightarrow} S \rightarrow A \times \mathcal{R}_0$ that estimates \checkmark for a prediction set, given an augmented calibration set (S, A) . Let $T^*(\checkmark, A)$ be a set of all "-correct parameters \checkmark with respect to f adapted using a test adaptation set A . We also consider $F_{\checkmark}^A(x, y) \triangleq F_{\checkmark}(A, x, y)$.

We say that a prediction set F_{\checkmark} parametrized by \checkmark is "-correct for p_{\checkmark} for a fixed adaptation set A , if $\checkmark \in T^*(\checkmark, A)$. Moreover, we want the prediction set F_{\checkmark} where \checkmark is estimated via \checkmark to be correct on most test sets A for adaptation and for most test tasks \checkmark , i.e., given $S \supseteq S_i$ and $A \supseteq A_i$,

$$\mathbb{P}^n_{\checkmark}(S, A) \supseteq T^*(\checkmark, A) \quad 1 - \epsilon,$$

where the probability is taken over \mathcal{V} and A . Intuitively, we want a prediction set, adapted via A , to be correct on future labeled examples from $p_{\mathcal{V}}$. Finally, we want the above to hold for most augmented calibration sets (S, A) from most calibration tasks $\mathcal{V}_1, \dots, \mathcal{V}_N$, leading to the following definition.

Definition 3. An estimator $\hat{\mathbb{P}}: S \rightarrow A \rightarrow R$ is (\cdot, ϵ, \cdot) -meta probably approximately correct (mPAC) for $p_{\mathcal{V}}$ if

$$\mathbb{P} \mathbb{P} \hat{\mathbb{P}}(S, A) \geq T^*(\mathcal{V}, A) \quad 1 - \epsilon \quad 1 \quad ,$$

where the outer probability is taken over \mathcal{V} , S and A , and the inner probability is taken over \mathcal{V} and A .

Notably, Definition 3 is tightly related to the PAC definition of meta conformal prediction [14]. We illustrate the difference in Figure 2; the figure represents the different guarantees of prediction sets via graphical models. Each stripe pattern indicates the random variables that are conditioned on in the respective guarantee. Figure 2d represents the proposed guarantee, having Figure 2b and 7b as special cases. Figure 2c shows the guarantee of meta conformal prediction proposed in [14], also having Figure 2a and 7a as special cases. Figure 2d shows that the proposed correctness guarantee is conditioned over calibration tasks, an adaptation set A from a test task, and test data (x, y) . This implies that the correctness guarantee of an adapted prediction set (thus conditioned on the adaptation set) holds on most future test datasets from the same test task. In contrast, the guarantee in Figure 2c is conditioned on the entire test task. Thus, the correctness guarantee does not hold conditionally over adaptation sets; this means that the adapted prediction set (thus conditioned on an adaptation set) does not need to be correct over most future datasets in the same task.

4.2 Algorithm

Algorithm 1 Meta-PS: PAC prediction set for meta-learning. Internally, any PAC prediction set algorithms are used to implement an estimator $\hat{\mathbb{P}}(\cdot, \cdot)$; we use PS-BINOM (Algorithm 2) in Appendix B.

```

1: procedure META-PS( $S, A, f, g, \epsilon, \cdot$ )
2:   for  $i \in \{1, \dots, N\}$  do
3:      $\hat{\mathbb{P}}_i \leftarrow \text{PS-BINOM}(S_i, f^{A_i}, \epsilon/2)$  (.) Use Algorithm 2 in Appendix B
4:      $S^0 \leftarrow ((\hat{\mathbb{P}}_1, 1), \dots, (\hat{\mathbb{P}}_N, 1))$ 
5:    $\hat{\mathbb{P}} \leftarrow \text{PS-BINOM}(S^0, g, \epsilon/2, \cdot)$ 
6:   return  $\hat{\mathbb{P}}$ 

```

Next, we propose our meta PAC prediction set algorithm for an estimator $\hat{\mathbb{P}}$ that satisfies Definition 3 while aiming to minimize the prediction set size, by leveraging the scalar parameterization of a prediction set. The algorithm consists of two steps: find a prediction set parameter over labeled examples S_i and A_i from the i -th calibration task distribution, and find a prediction set over the parameter of prediction sets in the previous step. At inference, the constructed prediction set is adapted to a test task distribution via an adaptation set A . Importantly, in finding prediction sets (i.e., implementing an (\cdot, \cdot) -PAC estimator $\hat{\mathbb{P}}(\cdot, \cdot)$), we can use any PAC prediction set algorithms (e.g., Algorithm 2 in Appendix B).

Meta calibration. For the first step, for each $i \in \{1, \dots, N\}$, we estimate a prediction set parameter $\hat{\mathbb{P}}(\cdot, \epsilon/2)(S_i)$, satisfying a $(\cdot, \epsilon/2)$ -PAC property using a calibration set S_i and a score function f adapted via an adaptation set A_i , leading to f^{A_i} , by using any meta learning and adaptation algorithm. We consider the N learned prediction set parameters—with a dummy label 1 to fit the convention of the PAC prediction set algorithm—to form a new labeled sample, i.e., $S^0 = ((\hat{\mathbb{P}}(\cdot, \epsilon/2)(S_1), 1), \dots, (\hat{\mathbb{P}}(\cdot, \epsilon/2)(S_N), 1))$.

In the second part, we construct an $(\epsilon/2, \cdot)$ -PAC prediction set for the threshold $\hat{\mathbb{P}}$ using S^0 and a score function g defined as $g(\hat{\mathbb{P}}, y) = \mathbb{P}[\mathbb{I}(y = 1)]$, where $\mathbb{I}(\cdot)$ is the indicator function². Thus, the constructed prediction set can be equivalently viewed as an interval $[\hat{\mathbb{P}}_{\epsilon/2}, (S^0), 1]$ where the thresholds $\hat{\mathbb{P}}$ of most calibration distributions are included; see Section 4.3 for a formal statement along with intuition on our design choice, and see Algorithm 1 for an implementation.

¹It is also called a training-set conditional guarantee in [21].

²The constructed prediction set parameter is equivalent to $\hat{\mathbb{P}}$ satisfying $\mathbb{P}[\mathbb{P}[\hat{\mathbb{P}}(\cdot, \epsilon/2)(S_i), 1] \geq \epsilon/2]$ with probability at least 1 over the randomness of S_1, \dots, S_N used to estimate $\hat{\mathbb{P}}$.

Adaptation and inference. We then use labeled examples A from a test distribution to adapt the meta-learned score function f as in the meta calibration part. Then, we use the same parameter $\hat{\alpha}(S, A)$ of the meta prediction set along with the adapted score function as our final meta prediction set $F_{\hat{\alpha}(S, A)}$ for inference in the future.

Reusing the same meta prediction set parameter after adaptation is valid, as the meta prediction set parameter is also chosen after adaptation via a hold-out adaptation set from the same distribution as each respective calibration set. Moreover, by reusing the same parameter, we only need a few labeled examples from the test distribution for an adaptation set, without requiring additional labeled examples for calibration on the test distribution. Finally, we rely on a PAC prediction set algorithm that aims to minimize the expected prediction set size [21, 3]; Algorithm 1 relies on this property, thus making the meta prediction set size small.

4.3 Theory

The proposed Algorithm 1 outputs a threshold for a prediction set that satisfies the mPAC guarantee, as stated in the following result (see Appendix C for a proof):

Theorem 1. The estimator $\hat{\alpha}$ implemented in Algorithm 1 is (ϵ, δ, η) -mPAC for p_{α} .

Intuitively, the algorithm implicitly constructs an interval $[\hat{\alpha}_{1/2}, \hat{\alpha}_{1/2}^0, 1]$ over scalar parameters of prediction sets. In particular, for the i -th distribution, the Algorithm in Line 3 finds a scalar parameter of a prediction set for the i -th distribution, which forms an empirical distribution over the scalar parameters. Due to the i.i.d. nature of \sqrt{i} , $i = 1, \dots, N$, the parameter of a prediction set for a test distribution follows the same distribution over scalar parameters. Thus, choosing a conservative scalar parameter as in Line 5 of the Algorithm suffices to satisfy the mPAC guarantee.

5 Experiments

We demonstrate the efficacy of the proposed meta PAC prediction set based on four few-shot learning datasets across three application domains: mini-ImageNet [10] and CIFAR10-C [15] for the visual domain, FewRel [16] for the language domain, and CDC Heart Dataset [17] for the medical domain. See Appendix E for additional experiments, including the results for CIFAR10-C.

5.1 Experimental Setup

Few-shot learning setup. We consider k -shot c -way learning except for the CDC Heart Dataset; In particular, there are c classes for each task dataset, and k adaptation examples for each class. Thus, we have $t := kc$ labeled examples to adapt a model to a new task. As the CDC Heart dataset is a label unbalanced dataset, we do not assume equal shots per class.

Score functions. We use a prototypical network [11] as a score function. It uses t adaptation examples to adapt the network to a given task dataset, and predicts labels on query examples.

Baselines. We compare four approaches, including our proposed approach.

- PS: We consider a naive application of PAC prediction sets [3]. We pool all nN labeled examples from calibration datasets into one calibration set to run the PAC prediction set algorithm.
- PS-Test: We construct PAC prediction sets by using 20 shots for each class directly drawn from the test distribution.
- Meta-CP: We use the PAC variant of meta conformal prediction [14], which trains a quantile predictor along with a score function, while our approach only needs a score function. We run the authors' code with the same evaluation parameters for comparison.
- Meta-PS: This is the proposed meta PAC prediction set from Algorithm 1.

Metric. We evaluate prediction sets mainly via their empirical prediction set error over a test sample, i.e., having learned a parameter $\hat{\alpha}$, we find $1/|E| \sum_{(x, y) \in E} \mathbb{1}(y \notin F_{\hat{\alpha}}^A(x))$, where A is an adaptation set, and E is an evaluation set drawn from each test task distribution. We desire this to be at most "given fixed calibration and test datasets. This should hold with probability at least $1 - \delta$ over the randomness due to the calibration and adaptation data; and with probability at least $1 - \epsilon$ over the randomness due to the test adaptation set. To evaluate this, we conduct 100 random trials, drawing independent calibration datasets, and for each fixed calibration dataset, we conduct 50 random trials

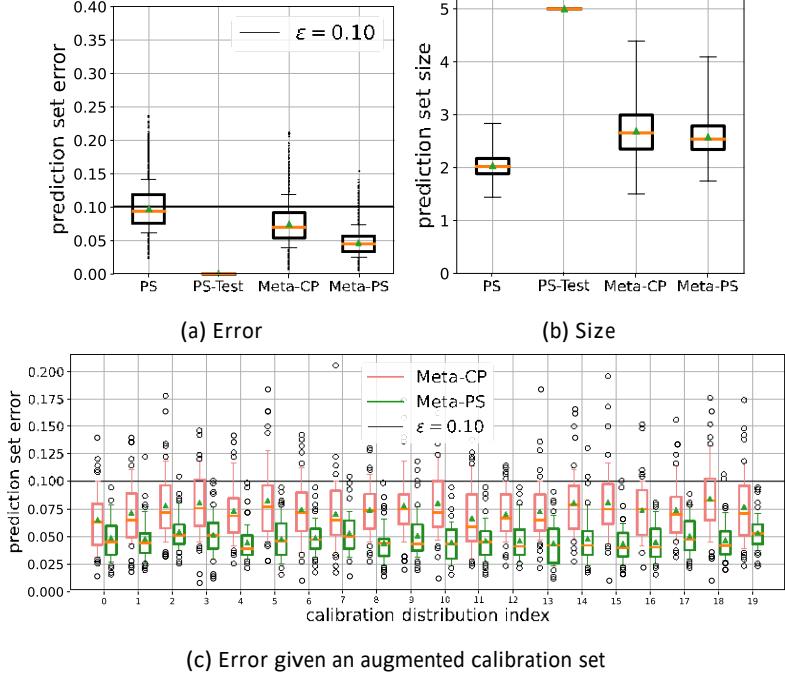


Figure 3: Prediction set error (3a), size (3a), and error given a calibration set (3c) on mini-ImageNet. Parameters are $N = 500$, $n = 2500$, $t = 25$, $\epsilon = 0.1$, $\epsilon = 0.1$, $\epsilon = 10^{-5}$ for Meta-PS and $\epsilon = 0.1$, $\epsilon = 10^{-5}$ for the other methods.

drawing independent test datasets. The prediction set size evaluation is done similarly to the error; we compute the empirical prediction set size $1/|E| \sum_{(x,y) \in E} S(F_{\pi}^A(x))$, where S is a size measure, e.g., set cardinality for classification.

5.2 Mini-ImageNet

Mini-ImageNet [10] is a smaller version of the ImageNet dataset for few-shot learning.

Setup. Mini-ImageNet consists of 100 classes with 64 classes for training, 16 classes for calibration, and 20 classes for testing; each class has 600 images. Following a standard setup, we consider five randomly chosen classes as one task and 5-way classification. In training, we follow 5-shot 5-way learning, considering $M = 800$ training task distribution randomly drawn from the $\binom{64}{5}$ possible tasks. In calibration, we have $N = 500$ calibration task datasets randomly drawn from the $\binom{16}{5}$ possible tasks, and use 5 shots for adaptation and 500 shots for calibration (i.e., $t = 25$ and $n = 2500$). In evaluation, we have 50 test task datasets randomly drawn from the $\binom{20}{5}$ possible tasks and use 5 shots from each of 5 classes for adaptation and 100 shots from each 5 classes for evaluation.

Results. Figure 3 shows the prediction set error and size of the various approaches in box plots. Whiskers in each box plot for the error cover the range between the 10-th and 90-th percentiles of the empirical error distribution, while whiskers in each box plot for the size represents the minimum and maximum of the empirical size distribution. The randomness of error and size is due to the randomness of the augmented calibration (S, A) and a test adaptation set A ; but since $\epsilon = 10^{-5}$, the randomness is mostly due to the adaptation set.

The prediction set error of the proposed approach is below the desired level $\epsilon = 0.1$ at least a fraction $1 - \epsilon = 0.9$ of the time. This supports that the proposed prediction set satisfies the meta PAC criterion. However, PS and Meta-CP do not empirically satisfy the meta PAC criterion, and PS-Test is too conservative. In particular, Meta-CP does not empirically control the randomness due to the test distribution. To make this point precise, Figure 3c shows a prediction set error distribution over different test distributions given a fixed calibration dataset. As can be seen, the proposed Meta-PS approach empirically satisfies the ϵ error criterion at least a fraction $1 - \epsilon = 0.9$ of the time (as

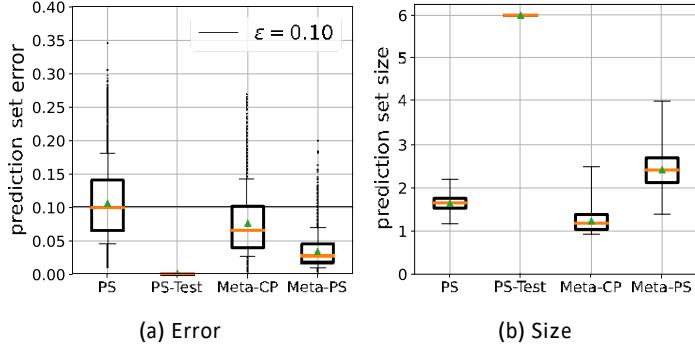


Figure 4: Prediction set error and size on FewRel. Parameters are $N = 500$, $n = 2500$, $t = 25$, $\epsilon = 0.1$, $\alpha = 0.1$, $\beta = 10^{-5}$ for Meta-PS and $\epsilon = 0.1$, $\beta = 10^{-5}$ for other methods.

shown in the whiskers), but Meta-CP does not empirically control it. Figure 3b demonstrates that the proposed Meta-PS produces smaller prediction sets compared to the PS-Test approach that also satisfies the " ϵ " error constraint.

5.3 FewRel

FewRel [16] is an entity relation dataset, where each example is the tuple of the first and second entities in a sentence along with their relation as a label.

Setup. FewRel consists of 100 classes, where 64 classes are used for training and 16 classes are originally used validation; each class has 700 examples. The detailed setup for training, calibration, and test are the same as those of mini-ImageNet, except that the test task distributions are randomly drawn from the original validation split.

Results. Figure 4 presents the prediction set error and size of the approaches. The trend of the results is similar to the mini-ImageNet result. Ours empirically satisfies the desired prediction set error specified by " ϵ " with empirical probabilities of at least $1 - \alpha$ over the test adaptation set and $1 - \beta$ over the (augmented) calibration data. In contrast, the other approaches either do not empirically satisfy the meta PAC criterion or are highly conservative. Figure 6a also confirms that the proposed approach empirically controls the error with probability at least $1 - \alpha$ over the randomness in the test dataset.

5.4 CDC Heart Dataset

The CDC Behavioral Risk Factor Surveillance System [17] consists of health-related telephone surveys started in 1984. The survey data is collected from US residents. We use the data collected from 2011 to 2020 and consider predicting the risk of heart attack given other features (e.g., level of exercise or weight). We call this curated dataset the CDC Heart Dataset. See Appendix D.4 for the details of the experiment setup.

Setup. The 10 years of the CDC Heart Dataset consists of 530 task distributions from various US states. We consider data from 2011-2014 as the training task distributions, consisting of 212 different tasks and 1,805,073 labeled examples, and data from 2015-2019 as calibration task distributions, consisting of 265 different tasks and 2,045,981 labeled examples. We use the data from 2020 as a test task distribution, consisting of 53 different tasks and 367,511 labeled examples. We use $M = 200$, $N = 250$, $n = 3000$, and $t = 10$; see Appendix D.4 for the details.

Results. Figure 5 includes the prediction set error and size of each approach. The trend is as before. The proposed Meta-PS empirically satisfies the " $\epsilon = 0.1$ " constraint (i.e., the top of the whisker is below of the " $\epsilon = 0.1$ " line), while PS and Meta-CP violate this constraint. Meanwhile, PS-Test conservatively satisfies it. As before, Figure 6b empirically justifies that the proposed approach controls the prediction set error due to the randomness over the test adaptation data.

6 Conclusion

We propose a PAC prediction set algorithm for meta learning, which satisfies a PAC property tailored to a meta learning setup. The efficacy of the proposed algorithm is demonstrated in four datasets across three application domains. Specifically, we observe that the proposed algorithm finds a

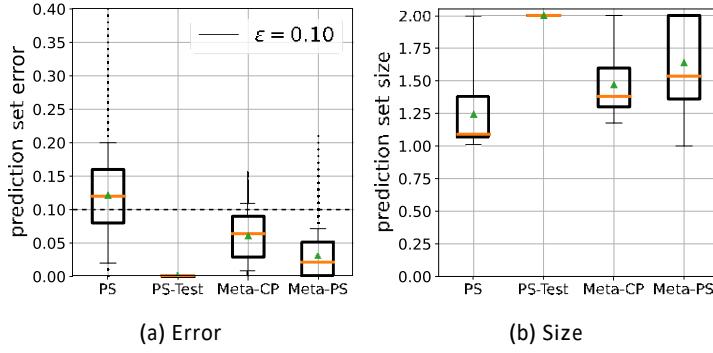


Figure 5: Prediction set error and size on CDC Heart Dataset. Parameters are $N = 250$, $n = 3000$, $t = 10$, $\gamma = 0.1$, $\epsilon = 0.1$, $\delta = 10^{-5}$ for Meta-PS and $\gamma = 0.1$, $\delta = 10^{-5}$ for other methods.

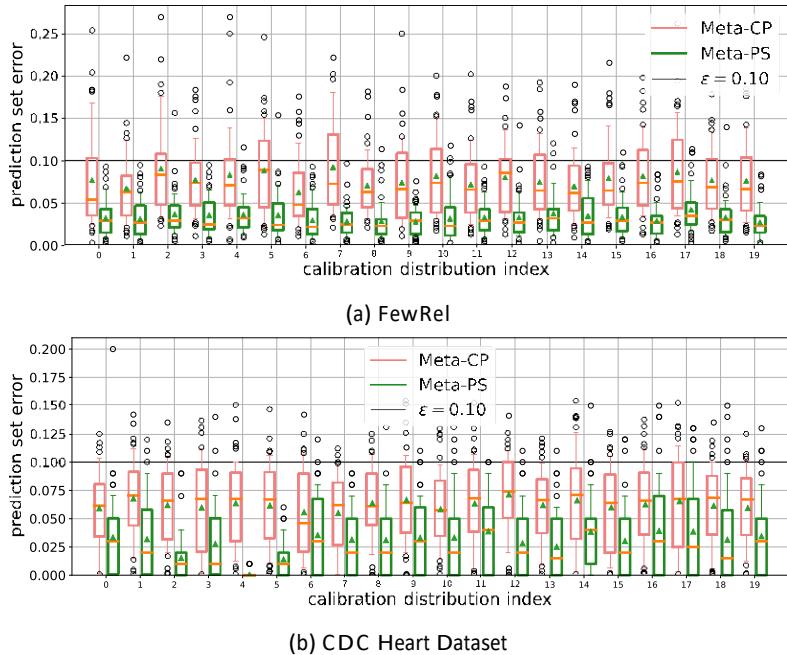


Figure 6: Meta-PS and Meta-CP prediction set error for various calibration sets. Given one calibration set (i.e., a fixed calibration distribution index), the box plot with whiskers represents the (ϵ, ϵ) -guarantee. The proposed Meta-PS satisfies the $\epsilon = 0.1$ constraints at least 1 $\epsilon = 0.9$ of the time, as the whisker is below the dotted line indicating ϵ . However, Meta-CP does not satisfy this constraint, showing that the proposed approach can control the correctness on future tasks.

prediction set that satisfies a specified PAC guarantee, while producing a small set size. A limitation of the current approach is that it requires enough calibration datapoints to satisfy the PAC guarantee, see [3] for an analysis. A potential societal impact is that a user of the proposed algorithm might misuse the guarantee without a proper understanding of assumptions; see Appendix F for details.

Acknowledgements

This work was supported in part by DARPA/AFRL FA8750-18-C-0090, ARO W911NF-20-1-0080, NSF TRIPODS 1934960, NSF DMS 2046874 (CAREER), NSF-Simons 2031895 (MoDL). Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the Air Force Research Laboratory (AFRL), the Army Research Office (ARO), the Defense Advanced Research Projects Agency (DARPA), or the Department of Defense, or the United States Government.

References

- [1] Samuel S Wilks. Determination of sample sizes for setting tolerance limits. *The Annals of Mathematical Statistics*, 12(1):91–96, 1941.
- [2] Vladimir Vovk, Alex Gammerman, and Glenn Shafer. *Algorithmic learning in a random world*. Springer Science & Business Media, 2005.
- [3] Sangdon Park, Osbert Bastani, Nikolai Matni, and Insup Lee. Pac confidence sets for deep neural networks via calibrated prediction. In *International Conference on Learning Representations*, 2020.
- [4] Stephen Bates, Anastasios Angelopoulos, Lihua Lei, Jitendra Malik, and Michael I Jordan. Distribution-free, risk-controlling prediction sets. *arXiv preprint arXiv:2101.02703*, 2021.
- [5] Ryan J Tibshirani, Rina Foygel Barber, Emmanuel Candes, and Aaditya Ramdas. Conformal prediction under covariate shift. *Advances in Neural Information Processing Systems*, 32:2530–2540, 2019.
- [6] Sangdon Park, Edgar Dobriban, Insup Lee, and Osbert Bastani. PAC prediction sets under covariate shift. In *International Conference on Learning Representations*, 2022.
- [7] Hongxiang Qiu, Edgar Dobriban, and Eric Tchetgen Tchetgen. Distribution-free prediction sets adaptive to unknown covariate shift. *arXiv preprint arXiv:2203.06126*, 2022.
- [8] Isaac Gibbs and Emmanuel Candès. Adaptive conformal inference under distribution shift, 2021.
- [9] Jurgen Schmidhuber. *Evolutionary Principles in Self-Referential Learning*. On Learning now to Learn: The Meta-Meta-Meta...-Hook. Diploma thesis, Technische Universität München, Germany, 14 May 1987.
- [10] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *Advances in neural information processing systems*, volume 29, 2016.
- [11] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. *Advances in neural information processing systems*, 30, 2017.
- [12] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pages 1126–1135. PMLR, 2017.
- [13] Kwonjoon Lee, Subhransu Maji, Avinash Ravichandran, and Stefano Soatto. Meta-learning with differentiable convex optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10657–10665, 2019.
- [14] Adam Fisch, Tal Schuster, Tommi Jaakkola, and Regina Barzilay. Few-shot conformal prediction with auxiliary tasks, 2021.
- [15] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *Proceedings of the International Conference on Learning Representations*, 2019.
- [16] Xu Han, Hao Zhu, Pengfei Yu, Ziyun Wang, Yuan Yao, Zhiyuan Liu, and Maosong Sun. FewRel: A large-scale supervised few-shot relation classification dataset with state-of-the-art evaluation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4803–4809, Brussels, Belgium, October–November 2018. Association for Computational Linguistics.
- [17] Centers for Disease Control and Prevention (CDC). Behavioral Risk Factor Surveillance System. https://www.cdc.gov/brfss/annual_data/annual_data.htm, 1984.
- [18] A Gammerman, V Vovk, and V Vapnik. Learning by transduction. In *Proceedings of the Fourteenth conference on Uncertainty in artificial intelligence*, pages 148–155, 1998.
- [19] Harris Papadopoulos, Kostas Proedrou, Volodya Vovk, and Alex Gammerman. Inductive confidence machines for regression. In *European Conference on Machine Learning*, pages 345–356. Springer, 2002.
- [20] Robin Dunn, Larry Wasserman, and Aaditya Ramdas. Distribution-free prediction sets for two-layer hierarchical models. *Journal of the American Statistical Association*, pages 1–12, 2022.

- [21] Vladimir Vovk. Conditional validity of inductive conformal predictors. *Machine learning*, 92(2-3):349–376, 2013.
- [22] I. Guttman. Statistical Tolerance Regions: Classical and Bayesian. Griffin’s statistical monographs & courses. Hafner Publishing Company, 1970.
- [23] Charles J Clopper and Egon S Pearson. The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika*, 26(4):404–413, 1934.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? **[Yes]**
 - (b) Did you describe the limitations of your work? **[Yes]** See Section 6.
 - (c) Did you discuss any potential negative societal impacts of your work? **[Yes]** See Section 6
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? **[Yes]**
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? **[Yes]** See Section 4.1.
 - (b) Did you include complete proofs of all theoretical results? **[Yes]** See Section C.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **[No]** Code will be released once accepted along with data and precise instructions to run it.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **[Yes]** See Section 5 and D.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **[Yes]** We use box plots to represent the randomness on error.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **[Yes]** See Section D.1.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? **[Yes]**
 - (b) Did you mention the license of the assets? **[Yes]** See Section D.2.
 - (c) Did you include any new assets either in the supplemental material or as a URL? **[No]**
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? **[Yes]** See Section D.2.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **[Yes]** See Section D.2.
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? **[N/A]**
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? **[N/A]**
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **[N/A]**