

# Multidimensional Belief Quantification for Label-Efficient Meta-Learning

Deep Shankar Pandey    Qi Yu  
Rochester Institute of Technology  
{dp7972, qi.yu}@rit.edu

## Abstract

*Optimization-based meta-learning offers a promising direction for few-shot learning that is essential for many real-world computer vision applications. However, learning from few samples introduces uncertainty, and quantifying model confidence for few-shot predictions is essential for many critical domains. Furthermore, few-shot tasks used in meta training are usually sampled randomly from a task distribution for an iterative model update, leading to high labeling costs and computational overhead in meta-training. We propose a novel uncertainty-aware task selection model for label efficient meta-learning. The proposed model formulates a multidimensional belief measure, which can quantify the known uncertainty and lower bound the unknown uncertainty of any given task. Our theoretical result establishes an important relationship between the conflicting belief and the incorrect belief. The theoretical result allows us to estimate the total uncertainty of a task, which provides a principled criterion for task selection. A novel multi-query task formulation is further developed to improve both the computational and labeling efficiency of meta-learning. Experiments conducted over multiple real-world few-shot image classification tasks demonstrate the effectiveness of the proposed model.*

## 1. Introduction

Deep learning (DL) models have achieved state-of-the-art performance for many computer vision applications. However, the effectiveness of DL models is challenged by some specialized domains (*e.g.* medicine, biology, and security intelligence), in which labeled data for model training may be scarce. Unlike the DL models, human beings can learn efficiently from limited training samples by using the prior knowledge stored in their brains and applying it to new tasks. For example, once a child learns how to distinguish between lions and tigers, s/he can quickly generalize the concept to distinguish lions and cats with little or no additional training. Inspired by such human learning, various few-shot learning techniques [6, 24, 37] have been devel-

oped, which provide a promising approach to address the label scarcity problem for DL models.

In recent successful few-shot learning approaches, the model is trained from multiple few-shot tasks comprised of few labeled examples instead of one large dataset as in the traditional setting. By learning from many similar tasks, the model can accumulate the shared knowledge among tasks. After training, it uses the knowledge gained from similar tasks as the prior knowledge to perform well on new unseen few-shot tasks. Meta-learning is one popular approach for few-shot learning where the model learns at two stages: *rapid learning* within a new task, which is guided by prior knowledge gained from *gradual learning* across tasks [30]. In meta-learning, the model is trained on a large number of few-shot tasks to learn the shared inter-task knowledge. The learned model is evaluated based on its generalization capabilities on unseen few-shot tasks.

Few-shot tasks have limited data to learn from (in some cases just 1 example/class). So, some model predictions may not be reliable. For critical applications (*e.g.* autonomous driving), it is essential to quantify the prediction uncertainty. Some existing approaches indirectly provide uncertainty information of few-shot tasks by learning a posterior predictive distribution for testing data samples [7, 10, 12, 17, 28]. However, they usually suffer from a high computational cost and rely on assumptions/approximations that may be invalid in practice.

Additionally, few-shot tasks used in meta training are usually sampled randomly from a task distribution formed using a large pool of labeled data samples. Thus, meta-training for many optimization-based meta-learning approaches is computationally expensive, requiring evaluation of the second-order derivative (*i.e.*, Hessian) of the (global) model parameters over each of the sampled tasks. Furthermore, the large number of tasks leads to high labeling costs in many real-world problems. However, not all the tasks contribute equally to the learning of the (global) model parameters, and evaluating the Hessian over these tasks can significantly slow down the meta training process.

In this paper, we present a novel *Uncertainty-aware task selection model for efficient meta-learning* (referred to as

Units-ML) that provides uncertainty estimation to quantify the model confidence in few-shot predictions. Building upon the theory of subjective logic [13], we formulate a multidimensional belief measure, including vacuous, conflicting, and incorrect beliefs, which can quantify both the *known uncertainty (KUN)* and *unknown uncertainty (UUN)* of a given task. However, evaluating incorrect belief relies on the labels of a query set in a task, making the UUN not accessible during task selection. We address this issue by proving a novel relationship between conflicting belief and incorrect belief, which allows us to bridge the gap to UUN. A novel task selection function is designed accordingly that integrates both KUN and UUN for *belief-oriented* label-efficient meta-learning.

We summarize our key contributions below: (1) a novel computationally and label-efficient meta-learning model that can estimate uncertainty in few-shot tasks, (2) a multidimensional belief measure to quantify the KUN and lower bound the UUN of a given task, (3) theoretical justification that conflicting belief lower bounds incorrect belief, which allows UUN estimation without label information, and (4) an uncertainty-aware task selection criterion and a novel multi-query task formulation to improve both computation and label efficiency of meta-learning. We conduct intensive experiments over multiple real-world image datasets to demonstrate the effectiveness of the proposed Units-ML model in terms of accurate uncertainty estimation, computationally effective task selection, and label-efficient learning under a limited annotation budget.

## 2. Related Work

In meta-learning, the meta-model aims to learn (prior) knowledge shared by relevant tasks over multiple training episodes so that the model can perform well in new few-shot tasks. The prior knowledge can be learned through embedding functions and similarity metrics as in metric-based models [5, 8, 33, 35, 37]. The prior knowledge can also be captured by a deep neural network that maps a training dataset to parameters of the task-specific meta-model as in model-based meta-learning approaches [9, 11, 16, 23].

In optimization-based meta-learning [29], task-specific model parameters are learned from the meta-dataset using an optimization procedure such that the model can adapt quickly with only a few examples from a new task. In Model-Agnostic Meta-Learning (MAML) [6], a good global initialization is learned from which the model can adapt quickly to new tasks using only a few data samples using a few gradient descent steps. Some improvements for MAML include MetaSGD [20] and MAML++ [2] that help further improve the generalization and stability of MAML. First-order approaches such as Reptile [26] have also been developed to address high computational costs in meta-training of MAML. Meta-learning has recently been

extended to the Bayesian setting [7, 10, 12, 41] to develop uncertainty awareness. We discuss these relevant uncertainty-aware meta-learning works in the Appendix.

Recent works have also attempted to show the effectiveness of task selection for meta-learning in reinforcement learning problems [15, 22]. In terms of task selection for few-shot classification problems, MTL [34] and GCP [21] share similar motivations to our approach. In MTL, a two-stage hard-task scheme is introduced where the model is first trained on a batch of tasks and a list of failure classes (based on query set loss) is maintained. In the second stage, the model trains from hard tasks created using the failure classes that lead to better generalization. In GCP, a class-pair-based task sampling scheme is developed as an effective alternative to existing uniform sampling for meta-learning. In GCP, the class-pair potential matrix is used to sample the training tasks. GCP requires keeping track of pairwise potential among all training classes, and might not scale well when the number of training classes is large, or when new training classes are introduced as training progresses. These approaches can be applied complementary to our approach as they focus on determining the most informative classes (instead of tasks as in our model) from which the task is to be generated. With these approaches, once the candidate classes are determined, our model could be applied to formulate a multi-query task and select the most informative task for effective meta-learning.

Our method is an instance of an optimization-based meta-learning with uncertainty awareness, *i.e.*, our model outputs the uncertainty estimates along with the predictions for few-shot tasks. In contrast to the probabilistic meta-learning approaches discussed above, our method does not add any significant computation overhead. Also, by further leveraging the predicted multidimensional belief (*i.e.*, vacuous, conflicting, and incorrect), we perform belief-oriented task selection for uncertainty-aware meta-learning with faster and better convergence, augmented with multidimensional belief based uncertainty quantification.

## 3. Methodology

In this section, we present the proposed belief-oriented task selection for efficient meta-learning (Units-ML). We first describe the standard problem setup for few-shot learning. We then provide an analysis on the computational cost of MAML that motivates the need to perform belief-oriented task selection that sets the stage for us to describe the proposed Units-ML model in detail.

**Problem setup.** We focus on few-shot classification problems and follow the episodic training procedure introduced in [37]. In particular, multiple tasks sampled from a task distribution are considered where each task consists of a support set  $S$  and a query set  $Q$ . Specifically, a task in a

$N$ -way  $K$ -shot classification problem is defined as

$$\begin{aligned} \mathcal{T} &= \{S, Q\} \\ S &= \{X_S, Y_S\} = \{(\mathbf{x}_1, l_1), \dots, (\mathbf{x}_{N_s}, l_{N_s})\} \\ Q &= \{X_Q, Y_Q\} = \{(\mathbf{x}_1, l_1), \dots, (\mathbf{x}_{N_q}, l_{N_q})\} \end{aligned} \quad (1)$$

Support set  $S$  has a total of  $N_s = N \times K$  instances with  $K$  examples/class, and query set  $Q$  has  $N_q$  new examples belonging to one of the  $N$  classes. During meta-training, both support and query sets are used to train the model; during meta-testing, the model performs adaptation using the support set and is evaluated on the query set.

Besides the standard problem setup as above, we also consider the setting where only limited samples can be annotated due to a limited labeling budget and the goal is to train the meta-model in a label-efficient way. We assume each task consists of small support set with limited labeled samples along with an unlabeled query set with varied sizes. We want the meta-model to be trained such that it can perform well on any new samples of the task (*i.e.*, any query set) after learning from the knowledge of the support set.

**Analysis of MAML.** MAML aims to learn a good initialization over multiple meta-iterations using the support-query setup discussed above. In each meta-iteration, a batch of tasks updates the model’s global parameters. The updates in MAML can be summarized in two iterative steps: a local update using the support set and a global update using the query set. For each task, the local update proceeds as:

$$\begin{aligned} \theta_0 &= \theta \quad [\text{make copy of global parameters}] \\ \theta_1 &= \theta_0 - \alpha \nabla_{\theta_0} \mathcal{L}[f(\theta_0, X_S), Y_S] \quad \dots \\ \theta_M &= \theta_{M-1} - \alpha \nabla_{\theta_{M-1}} \mathcal{L}[f(\theta_{M-1}, X_S), Y_S] \end{aligned} \quad (2)$$

Here, the model  $f$  outputs the predictions for support set input  $X_S$  based on the parameters  $\theta_m, m \in [1, M]$ . The support set prediction  $f(\theta_m, X_S)$  and support set ground truth  $Y_S$  is used to compute the loss  $\mathcal{L}$  and the  $m^{\text{th}}$  local update is done based on this loss. After  $M$  local updates, the global parameters are updated using the query set input  $X_Q$  and the query set ground truth  $Y_Q$  as:

$$\theta_{\text{new}} = \theta - \beta \nabla_{\theta} \mathcal{L}[f(\theta_M, X_Q), Y_Q] \quad (3)$$

Denote the query set loss  $\mathcal{L}[f(\theta_M, X_Q), Y_Q]$  by  $\mathcal{L}_M^Q$  and support set loss  $\mathcal{L}[f(\theta_M, X_S), Y_S]$  by  $\mathcal{L}_M^S$ . After  $M$  local updates using the support set  $S$ , we update global parameters using the query set  $Q$  as:

$$\begin{aligned} \theta_{\text{new}} &= \theta - \beta \nabla_{\theta} \mathcal{L}_M^Q = \theta - \beta \nabla_{\theta_M} \mathcal{L}_M^Q \times \nabla_{\theta} [\theta_M] \\ &= \theta - \beta \nabla_{\theta_M} \mathcal{L}_M^Q \times \left( \prod_{m=M}^1 \nabla_{\theta_{m-1}} [\theta_m] \right) \times \nabla_{\theta} [\theta_0] \\ &= \theta - \beta \nabla_{\theta_M} \mathcal{L}_M^Q (I - H_{M-1}) \dots \times (I - H_0) \times I \end{aligned} \quad (4)$$

where  $\nabla_{\theta_M} \mathcal{L}_M^Q$  is a vector of length same as the number of parameters in  $\theta$ ,  $I$  is the identity matrix, and  $H_m = \nabla_{\theta_m} [\nabla_{\theta_m} \mathcal{L}_m^S]$  is the Hessian matrix. As shown above, the global parameter update is through the loss over the query set samples  $\mathcal{L}_M^Q$ , with  $\theta_m$  implicitly capturing the support set information. To achieve label-efficient meta-learning, we need to quantify the informativeness of a task through its query set. Furthermore, global parameter update involves multiple Hessian-gradient products, which are computationally expensive. In standard meta-learning, a large number of tasks need to be labeled and then used for episodic training to find good global parameters. This not only incurs a high annotation cost but also takes a long time to converge. The proposed Units-ML model aims to *select the most informative tasks* for efficient meta-learning to reduce both the label and computational cost.

### 3.1. Multidimensional Task Belief Quantification

We formulate a novel multidimensional belief-based measure to quantify different types of task uncertainty in meta-learning by leveraging the formalism from Subjective Logic (SL) [13]. SL considers  $N$  opinions (corresponding to  $N$  classes) and assigns belief masses to each opinion ( $b_1, b_2, \dots, b_N$ ) along with an overall uncertainty mass  $u$ . The belief masses represent the total evidence from the model whereas the uncertainty mass represents the vacuity (*i.e.*, *lack of evidence*), and the two masses sum to 1:

$$\sum_{n=1}^N b_n + u = 1, \forall n : 0 \leq b_n \leq 1, 0 \leq u \leq 1 \quad (5)$$

By explicitly considering the uncertainty mass and using evidence-based measure (vacuity) to quantify it, we can obtain the model’s *vacuous belief* on given tasks. By learning tasks with a high vacuity, the model can gain the lacking knowledge. Furthermore, we can also capture the uncertainty due to the *conflicting belief* using the dissonance (*dis*) that is complementary to the vacuity ( $u$ ):

$$dis = \sum_{n=1}^N \left( b_n \frac{\sum_{j \neq n} b_j \text{Bal}(b_j, b_n)}{\sum_{j \neq n} b_j} \right), \quad (6)$$

$$\text{Bal}(b_j, b_n) = \begin{cases} 1 - \frac{|b_j - b_n|}{b_j + b_n}, & \text{if } b_j b_n > 0 \\ 0, & \text{otherwise} \end{cases} \quad (7)$$

where  $\text{Bal}(\cdot, \cdot)$  is the relative mass balance function between two belief masses. By learning tasks with a high dissonance, the model can correct its acquired conflicting knowledge to ensure more accurate predictions. Additional discussions about SL is presented in the Appendix.

The theory of SL can be conveniently embedded in a standard (non-Bayesian) neural network, making it computationally attractive. In particular, a neural network can

form multinomial opinions in classification by replacing the final softmax layer with a non-negative activation layer [31]. As a result, the network is trained to predict an evidence vector  $e = (e_1, e_2, \dots, e_N)$  for a given input  $x$ . The belief and vacuity are then computed as

$$b_n = \frac{e_n}{S}, \quad u = \frac{N}{S}, \quad \text{where } S = \sum_{n=1}^N (e_n + 1), \quad (8)$$

By setting  $\alpha_n = e_n + 1$ , the probability of assigning  $x$  to the  $n$ -th class is  $\frac{\alpha_n}{S}$ . If we use a set of categorical random variables  $(p_1, \dots, p_N)^\top$  to represent the class assignment probabilities, then  $\alpha$ 's are essentially the concentration parameters of a Dirichlet prior  $\text{Dir}[(p_1, \dots, p_N)^\top | (\alpha_1, \dots, \alpha_N)^\top]$ .

**Multidimensional Belief for Task Uncertainty Quantification.** To quantify the vacuous belief (*i.e.*, vacuity) and conflicting belief (*i.e.*, dissonance) for a given task  $t$  that consists of a support set with limited labeled instances along with an unlabeled query set  $q$ , we propose to perform meta-testing on each of the sampled tasks. In particular, the model first adapts to the task by using the support set  $S$ . Then, task vacuity and dissonance are evaluated using the unlabeled query set  $Q$ . We compute the vacuity and dissonance using (8) and (6) for each data sample in the query set of a task. The *vacuous belief*  $vb^t$  and *conflicting belief*  $cb^t$  of task  $t$  are computed as the average of vacuity and dissonance of query set samples:

$$vb^t = \frac{1}{N_q} \sum_{q=1}^{N_q} u_q^t \quad \text{(Vacuous Belief)} \quad (9)$$

$$cb^t = \frac{1}{N_q} \sum_{q=1}^{N_q} dis_q^t \quad \text{(Conflicting Belief)} \quad (10)$$

where  $u_q^t$  and  $dis_q^t$  are the vacuity and the dissonance for the  $q^{th}$  query sample in task  $t$ . Since vacuity reflects the model's lack of evidence on a data sample, vacuous belief indicates the model's overall lack of knowledge on a task. Therefore, selecting a task with a high  $vb^t$  and performing meta-training using (4) can adjust the global parameters to effectively learn the missing knowledge. As a result, the model is expected to perform well on similar unseen few-shot tasks in the future. While vacuous belief captures one source of uncertainty that indicates the model's lack of knowledge for the task, conflicting belief helps identify the difficult tasks, *i.e.*, the tasks in which the model gets confused between different classes. Learning from these tasks can help adjust the global parameters such that the model can correct the acquired confusing knowledge. As a result, the model is expected to be able to better differentiate different classes within a task.

Since both vacuous belief and conflicting belief can be quantified without knowing the labels of the query set, they are instances of *known uncertainty (KUN)*. There is another source of uncertainty, referred to as *unknown uncertainty (UUN)*, that the model is unaware of. UUN usually leads to a highly confident wrong prediction that can cause more severe consequences in critical domains (*e.g.*, autonomous driving). This type of uncertainty is essentially caused by model overfitting, which can be quite common when applying deep learning models to few-shot problems. It is critical to train the meta-learning model to minimize the unknown uncertainty so that the model can avoid making over-fitted predictions in the future. UUN can be captured by a third type of belief, referred to as *incorrect belief*:

$$ib^t = \frac{1}{N_q} \sum_{q=1}^{N_q} \|b_q^t \odot (\mathbf{1} - y_q^t)\|_1 \quad \text{(Incorrect Belief)} \quad (11)$$

where  $y_q^t = (y_{q,1}^t, \dots, y_{q,N}^t)^T$  is the one-hot vector representing the ground-truth label of the  $q^{th}$  query set sample,  $b_q^t = (b_{q,1}^t, \dots, b_{q,N}^t)^T$  is the  $N$ -dimensional belief vector,  $\odot$  represents element-wise multiplication, and  $\|\cdot\|_1$  is the  $l_1$  norm. Intuitively, when the model is wrongly confident, *i.e.*, the model places a strong belief in a class that is different from the true class label, it will contribute a large incorrect belief component. The task-level incorrect belief aggregates these components to reflect the overall UUN of the task.

However, a key limitation of incorrect belief is that computing incorrect belief requires the query set labels, making it infeasible to be used in task selection for label-efficient meta-learning. We address this issue through an important theoretical result as presented in the following theorem. The theorem establishes an important relationship between incorrect belief and conflicting belief, which essentially bridges the gap between unknown uncertainty and known uncertainty.

**Theorem 1 (Lower bound of incorrect belief).** *Consider an unlabeled task  $t$  with (unknown) incorrect belief of  $ib^t$  and conflicting belief  $cb^t$ . Then, incorrect belief is lower bounded by a half of the conflicting belief on the same task*

$$ib^t \geq \frac{1}{2} cb^t \quad \text{where } 0 \leq cb^t \leq 1, 0 \leq ib^t \leq 1 \quad (12)$$

*Proof sketch.* We first consider a sample within the task for which the model outputs an  $N$ -dimensional belief vector. We consider the analytical expression of the conflicting belief, simplify the relative mass balance between the beliefs, and expand the different belief terms of the conflicting belief. After expanding and rearranging, we find an upper bound for each of the terms in the conflicting belief expression that proves that for any sample, the incorrect belief is lower bounded by half the conflicting belief. Finally, we generalize the relationship to be true for any task.  $\square$

Due to space limitations, the complete proof of the theorem is provided in the Appendix. Ideally, we want to select tasks with high incorrect belief that encourages the model to correct the model’s incorrect knowledge. Since the conflicting belief provides a lower bound for incorrect belief, it provides a way to estimate the incorrect belief (and reduce UUN) without the label information, which is instrumental for active task selection.

### 3.2. Belief-Oriented Task Selection and Training

The multidimensional belief provides a principled way to quantify both the KUN and UUN of different tasks without the query set labels. With this, the most informative tasks are the ones with the greatest overall uncertainty, including both KUN and UUN. For the former, it is captured by two different types of belief: vacuous and conflicting. As for the latter, we can obtain its lower bound through conflicting belief. Since conflicting belief is used to quantify both KUN and UUN, we propose a task selection function that integrates vacuous beliefs (vacuity) and conflicting beliefs (dissonance) to estimate the total task uncertainty:

$$\text{unc}^t = \lambda(vb^t) + (1 - \lambda)(cb^t) \quad (\text{Task Uncertainty}) \quad (13)$$

where  $\lambda$  is a balancing term that determines the relative importance between these two types of belief. Intuitively, the tasks with high vacuous belief represent new/unseen tasks on which the model is not able to make confident predictions, whereas the tasks with high conflicting belief represent challenging tasks on which the model struggles to confidently discriminate among the classes. We start with a relatively large  $\lambda$  in the early phase of meta-learning so that the model can better explore the task space to fill out the knowledge gap. Then, the focus will shift to the conflicting belief to fine-tune the model on the more difficult tasks or tasks that the model has an incorrect knowledge.

**Multi-Query Tasks.** With the above selection score ( $\text{unc}^t$ ), we propose to conduct uncertainty-aware task selection with a novel task formulation strategy for label-efficient meta-learning. A straightforward method for task selection is to sample a large number of tasks (say  $J$  tasks) from a task distribution  $p(\mathcal{T})$  and use task selection criteria to select  $I$  tasks to be labeled and perform meta-learning. We refer to this strategy as Units-ST (see Figure 1). In Units-ST, for each discarded task, the model needs to adapt to the support set to determine the informativeness which is a waste of computation and support set labels. To further improve efficiency, we propose to formulate *multi-query* tasks, where each task consists of a shared support set and multiple query sets. In this new formulation, referred to as Units-ML (see Figure 1), the model will adapt to the support set in the task and choose the most informative query set to label. Other unlabeled query sets will be discarded and not used

for meta-learning. Such Multi-Query Tasks can be an ideal choice for limited budget real-world few-shot problems.

**Belief Regularized Model Training.** We aim to train the meta-model to learn a good initialization such that for a new task, after learning from limited data of the support set, the meta-model can make a prediction as well as output the confidence in the prediction (the uncertainty information). To this end, we assume that the label for each sample is obtained from a generative process with a Dirichlet prior and a multinomial likelihood as specified through the SL framework. The parameters for the Dirichlet prior express the vacuity and belief masses for uncertainty estimation. Further, we leverage the conjugacy between the Dirichlet prior and the multinomial likelihood. With this, we can learn these parameters by minimizing a loss between the multinomial output and the ground truth labels.

Additionally, while the incorrect belief can only be estimated through its lower bound during the task selection phase, once the task is selected, the labels of its query set will be collected. Consequently, the incorrect belief can be accurately quantified, which can be used to guide the model training (to minimize the incorrect belief). To this end, we propose a belief regularized loss function

$$\mathcal{L}_i = -\ln \int \text{Mult}(\mathbf{y}_i|\mathbf{p}_i)\text{Dir}(\mathbf{p}_i|\boldsymbol{\alpha}_i)d\mathbf{p}_i + \eta R_{ib} \quad (14)$$

$$R_{ib} = \mathbf{b}_i \odot (\mathbf{1} - \mathbf{y}_i) \quad (15)$$

where  $\mathcal{L}_i$  is the loss on the  $i$ -th data sample with one hot label  $\mathbf{y}_i$ ,  $R_{ib}$  is the incorrect belief regularization for the sample, and  $\eta$  is a regularization coefficient that balances between minimizing the incorrect belief and maximizing the log likelihood. Moreover, the model outputs  $N$ -dimensional evidence  $\mathbf{e}_i$  from which the belief  $\mathbf{b}_i$  and dirichlet parameters  $\boldsymbol{\alpha}_i$  are obtained. Limited by space, we present descriptions of SL and the additional details about the incorrect belief regularization, loss function, design choices, and hyperparameter settings in the Appendix.

## 4. Experiments

We first carry out experiments to show the accurate multidimensional belief quantification, which empirically validates our theoretical results. We then conduct intensive experiments on real-world few-shot image classification tasks to demonstrate the effectiveness of the Units-ML model on (1) accurate uncertainty estimation for few-shot learning, (2) fast convergence with limited label budget, and (3) competitive meta-learning performance in terms of generalization of the learned model and flexibility of adjusting prediction results based on model confidence. To demonstrate the general applicability of the proposed model, we extend the MetaSGD models to make them uncertainty-aware and also

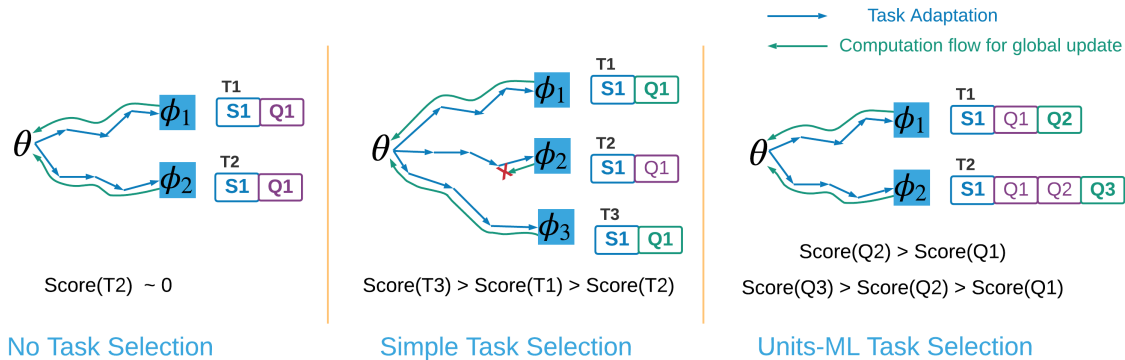


Figure 1. The left figure shows the idea of optimization-based meta-learning, the middle figure shows the idea of task selection (Units-ST), and the right figure illustrates the idea of our proposed task selection method with multi-query tasks. The forward arrows indicate the task adaptation (2), and the backward arrows represent the computations for global parameter update (4).

conduct additional experiments on any-shot classification using *mini*-ImageNet/CifarFS and in multi-dataset setting on Meta-Dataset [36] as proposed by Bayesian TAML [19]. Limited by space, we report these results along with a detailed experiment setup in the Appendix.

**Datasets.** We evaluate our proposed method on three real-world benchmark image datasets: Omniglot [18], *mini*-ImageNet [37], and CifarFS [4]. Details of the datasets are summarized in Table 2 of the Appendix.

#### 4.1. Details of Comparison Baselines

**Comparison models.** Our comparison includes optimization-based meta-learning models (MAML [6], MUMOMAML [38], CAVIA [42], MetaSGD [20], Reptile [26], HSML [40]) and models with uncertainty quantification capabilities (PLATIPUS [7], VERSA [10], BMAML [41], LLAMA [12], and ABML [28]). A description of each comparison baseline is provided in the Appendix. As some of these models do not release their source codes, we refer to the existing literature to report their performance. Thus, the results may not be available for all three datasets for some models.

**Experiment setup.** We experiment with few-shot classification problems, where we consider  $N$ -way  $K$ -shot tasks with  $q$  instances/class in the query set. In such a setting, tasks are created by randomly sampling  $N$  classes and then sampling  $K + q$  instances from each class. The  $N \times K$  instances make the support set of the task, and there are  $N \times q$  instances in the query set. The models are trained using Adam optimizer with the outer loop learning rate of 0.001 and evaluated on 600 validation set tasks. We train the model for 100 epochs where each epoch consists of 500 meta-iterations, and average the final test set performance across 3 independent runs. In each meta-iteration, the model is trained with 8 tasks for Omniglot, 4 tasks for CifarFS [4], 4 tasks for 5-way 1-shot *mini*-ImageNet,

and 2 tasks for 5-way 5-shot *mini*-ImageNet. We use the same standard 4-module convolutional architecture similar to ALFA [3], Antoniou et al. [2]. For limited labeling budget experiments, we train the models for 50 epochs with 100 meta-iterations where each meta-iteration consists of 2 tasks/batch (total of 10,000 tasks) for all the models. In the multi-query tasks, each task has 8 unlabeled query sets sharing the same support set. Additional implementation details are provided in the Appendix.

#### 4.2. Multidimensional Belief Quantification

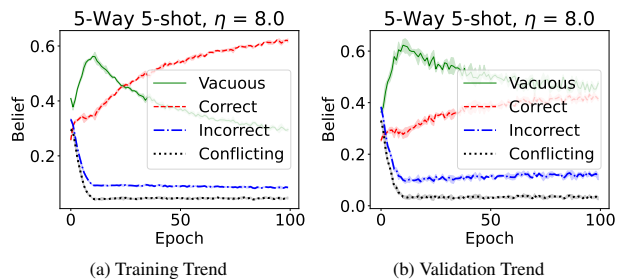


Figure 2. Multidimensional belief trend for CifarFS dataset

We present the training and validation trends of the multidimensional belief for 5-way 5-shot CifarFS dataset in Figure 2. There is low average vacuous belief and high incorrect belief at the initial training phase. This is most likely due to the model’s overfitting on the limited training data, which could also indicate the importance of applying the proposed multidimensional belief to quantify both KUN and UUN under a limited labeling budget when overfitting is more likely to occur. In the early phase, since the model knows less, it also under-estimates the vacuous belief. In the immediate next few epochs, the model starts to make an accurate adjustment to its multidimensional beliefs as it is exposed to more samples. This leads to an increase in the correct belief and a decrease in all other beliefs as desired.

It should be noted that the conflicting belief closely trails the incorrect belief in both the training and test tasks in all training phases. Thus, the figures empirically validate our theorem that the conflicting belief lower bounds twice the incorrect belief.

### 4.3. Uncertainty Estimation Performance

We then carry out experiments to assess the effectiveness of the proposed Units-ML model for uncertainty estimation in few-shot learning. For few-shot samples that remain new to the meta-model after being adapted to the samples in the support set, the model is expected to report a high vacuity; otherwise, the predicted vacuity should be low, reflecting high model confidence. Figure 3 shows the predicted vacuity for a query set in a 5-way 1-shot Omniglot test task. As the query set images are rotated (by angles from  $0^\circ$  to  $90^\circ$  as indicated by R), the model starts to make mistakes (indicated by red) and becomes more uncertain in its predictions (as indicated by vacuity). Furthermore, when tested using MNIST characters as out-of-distribution (OOD) samples, the model accurately outputs a large vacuity, which shows its potential for OOD detection in few-shot learning. Figure 4 shows our model’s performance on a 5-way 4-shot ImageNet test task with 3 in-distribution images, 2 open-set/OOD images (a cake image from a different *mini*-ImageNet class, and a bird image from the CUB dataset [39]) in the query set. For in-distribution samples, the model prediction is correct, and the prediction confidence (indicated by the vacuity and dissonance) is reasonable. For the OOD/open-set images, the model outputs high vacuity showing our model’s potential in OOD/open-set detection. Furthermore, all the belief mass for OOD samples contributes to the incorrect belief (*i.e.*, there is no correct belief for OOD samples). Our model outputs low dissonance for confidently correct predictions and comparably higher conflicting belief (the dissonance) for confusing samples, a highly desirable characteristic of a model with accurate uncertainty awareness. Additional illustrative examples and comparisons demonstrating our model’s potential in open-set/OOD detection, and the model training process is provided in the Appendix.

We further study the relationship between vacuity and the prediction accuracy for query set predictions to assess the reliability of the uncertainty. Figure 5 visualizes how the prediction accuracy varies with vacuity using 5-way 1-shot and 5-way 5-shot tasks from CifarFS. For instance, in 5-way 5-shot CifarFS, by setting the vacuity threshold to 0.2 (considering prediction accuracy for samples whose vacuity is less than 0.2), the model’s prediction accuracy reaches around 85%, which is around 10% better than making predictions without consulting vacuity. The results for other datasets and settings are presented in the appendix that show a similar trend. Such flexibility can effectively avoid mak-




Figure 3. Uncertainty prediction in a 5-w 1-s task

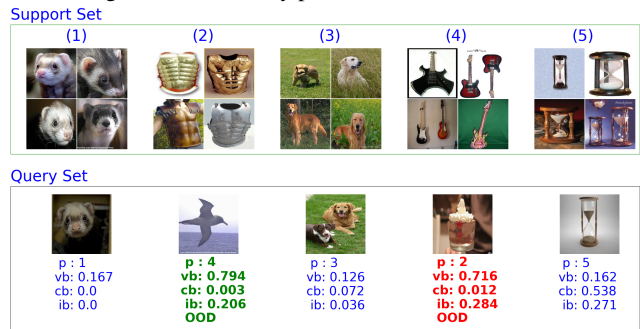


Figure 4. Uncertainty prediction in a 5-w 4-s *mini*-ImageNet test task with ood/open-set instances in query set. vb, cb, and ib represent the vacuous belief, conflicting belief, and incorrect belief, respectively for the prediction p.

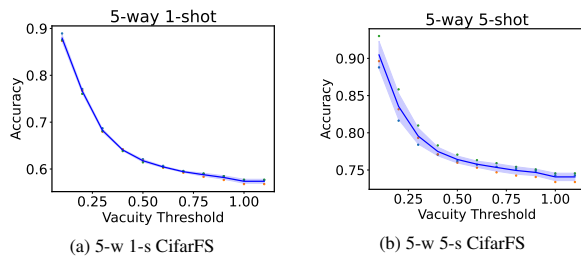


Figure 5. Prediction accuracy at different vacuity thresholds

ing less reliable predictions, a highly desirable property to facilitate decision-making in critical domains.

### 4.4. Active Task Selection

Next, to demonstrate the label efficiency of our proposed model, we experiment under a limited labeling budget scenario. We consider that the model has access to the small labeled support set and the model needs to decide the tasks to be labeled (limited labeled budget). We consider a labeling

budget of 10,000 tasks and compare our proposed belief-oriented task selection with the uncertainty aware Versa model that uses the model’s predicted uncertainty on the query set for task selection and MAML model (which is not uncertainty aware) that randomly selects the query set of the task to be labeled. Our model uses a novel task uncertainty score in (13) to select the tasks with the greatest task-level uncertainty to be labeled. Figure 6 shows the results where we outperform the baselines in the early phase of the meta-learning process under limited task budget scenario. For instance, in 20-way 5-shot omniglot experiments, our model converges to more than 90% accuracy only after 1000 iterations whereas the baseline models take significantly longer to converge to similar levels. Additional results are presented in the Appendix.

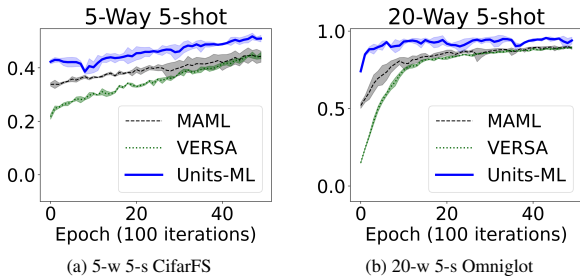


Figure 6. Meta Learning under limited labeling budget

#### 4.5. Meta-Learning Performance Comparison

Meta-learning performance on the three datasets and comparison with state-of-the-art competitive models are presented in Table 1. For the comparison, we present the Units-ML (the proposed task selection) and Units-NTS (stands for No Task Selection). As can be seen, the model with task selection achieves improvement over no task selection in almost all experiments except 5-way 1-shot mini-ImageNet experiments where the performance is close. We also present Units-ML with different uncertainty thresholds (e.g., Units-ML 0.2) to show the flexibility and effectiveness of the predicted uncertainty. For example, Units-ML 0.2 considers the predictive performance using the samples for which the model’s predicted uncertainty is less than 0.2. When considering only the confident predictions by adjusting the uncertainty threshold, our model achieves considerably higher accuracy demonstrating the effectiveness of uncertainty threshold.

### 5. Conclusion

In this paper, we propose an uncertainty-aware optimization-based meta-learning model for few-shot learning. Building upon the theory of subjective logic, the proposed Units-ML model successfully identifies the known uncertainty using vacuity and dissonance and identifies unknown uncertainty using (incorrect) belief mass, respectively. We design a novel task uncertainty score to choose

Table 1. Meta-learning Performance Comparison

<b>Omniglot</b>	20-way 1-shot(%)	20-way 5-shot(%)
MAML	95.8±0.3	98.9±0.2
Reptile	89.43±0.14	97.12±0.32
VERSA	97.66±0.29	98.77±0.18
Units-NTS	91.96±0.48	97.38±0.07
Units-ML	93.17±0.25	97.72±0.18
Units-ML 0.2	96.83±0.48	99.04±0.13
Units-ML 0.1	98.85±0.83	99.42±0.25
<b>mini-ImageNet</b>	5-way 1-shot(%)	5-way 5-shot(%)
MAML	48.70±1.84	63.15±0.91
MetaSGD	50.47±1.87	64.03±0.94
MUMOMAML	49.86±1.85	-
HSML	50.38±1.85	-
CAVIA	51.82±0.65	65.85±0.55
LLAMA	49.40±1.83	-
PLATIPUS	50.13±1.86	-
BMAML	53.17±0.87	-
VERSA	53.40±1.82	67.37±0.86
ABML	45.0±0.6	-
Units-NTS	51.38±0.33	66.75±0.42
Units-ML	50.86±0.67	68.16±0.72
Units-ML 0.2	61.25±2.89	80.70±0.93
Units-ML 0.1	82.26±4.77	91.07±0.26
<b>CifarFS</b>	5-way 1-shot (%)	5-way 5-shot (%)
MAML	58.9±1.9	71.5±1.0
MetaSGD*	57.77±0.17	71.16±0.21
Reptile	55.86±1.00	71.08±0.74
VERSA*	60.6±0.68	74.69±0.29
Units-NTS	59.80±0.31	76.15±0.35
Units-ML	59.84±0.11	76.69±0.44
Units-ML 0.2	76.62±0.42	83.54±1.73
Units-ML 0.1	87.92±0.71	90.46±1.82

\* Indicates results from local reproduction

the most informative tasks for meta-training. Our approach achieves comparable performance to many state-of-the-art optimization-based meta-learning methods. We further show the potential of our model for out-of-distribution distribution detection and label efficient task selection. Also, by adjusting the uncertainty threshold, Units-ML can provide a much more reliable prediction performance, which is essential to support decision-making in critical domains. As future work, we plan to extend our framework to metric-based and other meta-learning approaches that train in an episodic fashion.

### Acknowledgements

This research was supported in part by an NSF IIS award IIS-1814450 and an ONR award N00014-18-1-2875. The views and conclusions contained in this paper are those of the authors and should not be interpreted as representing any funding agency. We would also like to thank the anonymous reviewers for their constructive comments.



## References

- [1] Alexander Amini, Wilko Schwarting, Ava Soleimany, and Daniela Rus. Deep evidential regression. *arXiv preprint arXiv:1910.02600*, 2019.
- [2] Antreas Antoniou, Harrison Edwards, and Amos Storkey. How to train your MAML. In *International Conference on Learning Representations*, 2019.
- [3] Sungyong Baik, Myungsub Choi, Janghoon Choi, Heewon Kim, and Kyoung Mu Lee. Meta-learning with adaptive hyperparameters. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 20755–20765. Curran Associates, Inc., 2020.
- [4] Luca Bertinetto, Joao F. Henriques, Philip Torr, and Andrea Vedaldi. Meta-learning with differentiable closed-form solvers. In *International Conference on Learning Representations*, 2019.
- [5] Yinbo Chen, Xiaolong Wang, Zhuang Liu, Huijuan Xu, and Trevor Darrell. A new meta-baseline for few-shot learning. *arXiv preprint arXiv:2003.04390*, 2020.
- [6] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1126–1135. JMLR. org, 2017.
- [7] Chelsea Finn, Kelvin Xu, and Sergey Levine. Probabilistic model-agnostic meta-learning. In *Advances in Neural Information Processing Systems*, pages 9516–9527, 2018.
- [8] Victor Garcia and Joan Bruna Estrach. Few-shot learning with graph neural networks. In *6th International Conference on Learning Representations, ICLR 2018*, 2018.
- [9] Marta Garnelo, Dan Rosenbaum, Christopher Maddison, Tiago Ramalho, David Saxton, Murray Shanahan, Yee Whye Teh, Danilo Rezende, and SM Ali Eslami. Conditional neural processes. In *International Conference on Machine Learning*, pages 1704–1713. PMLR, 2018.
- [10] Jonathan Gordon, John Bronskill, Matthias Bauer, Sebastian Nowozin, and Richard Turner. Meta-learning probabilistic inference for prediction. In *International Conference on Learning Representations*, 2019.
- [11] Jonathan Gordon, Wessel P. Bruinsma, Andrew Y. K. Foong, James Requeima, Yann Dubois, and Richard E. Turner. Convolutional conditional neural processes. In *International Conference on Learning Representations*, 2020.
- [12] Erin Grant, Chelsea Finn, Sergey Levine, Trevor Darrell, and Thomas Griffiths. Recasting gradient-based meta-learning as hierarchical bayes. In *International Conference on Learning Representations*, 2018.
- [13] Audun Jøsang. *Subjective logic*. Springer, 2016.
- [14] Audun Josang, Jin-Hee Cho, and Feng Chen. Uncertainty characteristics of subjective opinions. In *2018 21st International Conference on Information Fusion (FUSION)*, pages 1998–2005. IEEE, 2018.
- [15] J Kaddour, S Sæmundsson, and MP Deisenroth. Probabilistic active meta-learning. In *NeurIPS*, volume 33. NeurIPS Proceedings, 2020.
- [16] Hyunjik Kim, Andriy Mnih, Jonathan Schwarz, Marta Garnelo, Ali Eslami, Dan Rosenbaum, Oriol Vinyals, and Yee Whye Teh. Attentive neural processes. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.
- [17] Jongmin Kim, Taesup Kim, Sungwoong Kim, and Chang D Yoo. Edge-labeling graph neural network for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11–20, 2019.
- [18] Brenden M Lake, Ruslan Salakhutdinov, and Joshua B Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.
- [19] Hae Beom Lee, Hayeon Lee, Donghyun Na, Saehoon Kim, Minseop Park, Eunho Yang, and Sung Ju Hwang. Learning to balance: Bayesian meta-learning for imbalanced and out-of-distribution tasks. In *Eighth International Conference on Learning Representations, ICLR 2020*. ICLR, 2020.
- [20] Zhenguo Li, Fengwei Zhou, Fei Chen, and Hang Li. Meta-sgd: Learning to learn quickly for few-shot learning. *arXiv preprint arXiv:1707.09835*, 2017.
- [21] Chenghao Liu, Zhihao Wang, Doyen Sahoo, Yuan Fang, Kun Zhang, and Steven CH Hoi. Adaptive task sampling for meta-learning. *arXiv preprint arXiv:2007.08735*, 2020.
- [22] Ricardo Luna Gutierrez and Matteo Leonetti. Information-theoretic task selection for meta-reinforcement learning. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 20532–20542. Curran Associates, Inc., 2020.
- [23] Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. A simple neural attentive meta-learner. In *International Conference on Learning Representations*, 2018.
- [24] Tsendsuren Munkhdalai and Hong Yu. Meta networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2554–2563. JMLR. org, 2017.
- [25] Cuong Nguyen, Thanh-Toan Do, and Gustavo Carneiro. Uncertainty in model-agnostic meta-learning using variational inference. In *The IEEE Winter Conference on Applications of Computer Vision*, pages 3090–3100, 2020.
- [26] Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018.
- [27] Barak A Pearlmutter. Fast exact multiplication by the hessian. *Neural computation*, 6(1):147–160, 1994.
- [28] Sachin Ravi and Alex Beatson. Amortized bayesian meta-learning. In *International Conference on Learning Representations*, 2019.
- [29] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *International Conference on Learning Representations (ICLR)*, 2017.
- [30] Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. Meta-learning with memory-augmented neural networks. In *International conference on machine learning*, pages 1842–1850, 2016.
- [31] Murat Sensoy, Lance Kaplan, and Melih Kandemir. Evidential deep learning to quantify classification uncertainty. In

- Advances in Neural Information Processing Systems*, pages 3179–3189, 2018.
- [32] Weishi Shi, Xujiang Zhao, Feng Chen, and Qi Yu. Multi-faceted uncertainty estimation for label-efficient deep learning. *Advances in neural information processing systems*, 33, 2020.
- [33] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*, pages 4077–4087, 2017.
- [34] Qianru Sun, Yaoyao Liu, Tat-Seng Chua, and Bernt Schiele. Meta-transfer learning for few-shot learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 403–412, 2019.
- [35] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1199–1208, 2018.
- [36] Eleni Triantafillou, Tyler Zhu, Vincent Dumoulin, Pascal Lamblin, Utku Evci, Kelvin Xu, Ross Goroshin, Carles Gelada, Kevin Swersky, Pierre-Antoine Manzagol, and Hugo Larochelle. Meta-dataset: A dataset of datasets for learning to learn from few examples. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020.
- [37] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *Advances in neural information processing systems*, pages 3630–3638, 2016.
- [38] Risto Vuorio, Shao-Hua Sun, Hexiang Hu, and Joseph J. Lim. Multimodal model-agnostic meta-learning via task-aware modulation. In *Neural Information Processing Systems*, 2019.
- [39] P. Welinder, S. Branson, T. Mita, C. Wah, F. Schroff, S. Belongie, and P. Perona. Caltech-UCSD Birds 200. Technical Report CNS-TR-2010-001, California Institute of Technology, 2010.
- [40] Huaxiu Yao, Ying Wei, Junzhou Huang, and Zhenhui Li. Hierarchically structured meta-learning. In *Proceedings of the 36th International Conference on Machine Learning*, 2019.
- [41] Jaesik Yoon, Taesup Kim, Ousmane Dia, Sungwoong Kim, Yoshua Bengio, and Sungjin Ahn. Bayesian model-agnostic meta-learning. In *Advances in Neural Information Processing Systems*, pages 7332–7342, 2018.
- [42] Luisa M Zintgraf, Kyriacos Shiarlis, Vitaly Kurin, Katja Hofmann, and Shimon Whiteson. Fast context adaptation via meta-learning. 2019.