BAYESIAN LEARNING TO OPTIMIZE: QUANTIFYING THE OPTIMIZER UNCERTAINTY

Anonymous authors
Paper under double-blind review

ABSTRACT

Optimizing an objective function with uncertainty awareness is well-known to improve the accuracy and confidence of optimization solutions. Meanwhile, another relevant but very different question remains yet open: how to model and quantify the uncertainty of an optimization algorithm itself? To close such a gap, the prerequisite is to consider the optimizers as sampled from a distribution, rather than a few pre-defined and fixed update rules. We first take the novel angle to consider the algorithmic space of optimizers, each being parameterized by a neural network. We then propose a Boltzmann-shaped posterior over this optimizer space, and approximate the posterior locally as Gaussian distributions through variational inference. Our novel model, *Bayesian learning to optimize* (**BL2O**) is the first study to recognize and quantify the uncertainty of the optimization algorithm. Our experiments on optimizing test functions, energy functions in protein-protein interactions and loss functions in image classification and data privacy attack demonstrate that, compared to state-of-the-art methods, BL2O improves optimization and uncertainty quantification (UQ) in aforementioned problems as well as calibration and out-of-domain detection in image classification.

1 INTRODUCTION

Computational models of many real-world applications involve optimizing non-convex objective functions. As the non-convex optimization problem is NP-hard, no optimization algorithm (or optimizer) could guarantee the global optima in general and instead, their solutions' usefulness (sometimes based on their proximity to the optima), when the optima are unknown, can be very uncertain. Being able to quantify such uncertainty is important to not only assessing the solution uncertainty after optimization but also enhancing the search efficiency during optimization. For instance, reliable and trustworthy machine learning models demand uncertainty awareness and quantification during training (optimizing) such models, whereas in reality deep neural networks without proper modeling of uncertainty suffer from overconfidence and miscalibration (Guo et al., 2017). In another application example of protein docking, although there exists epistemic uncertainty of the objective function and the aleatoric uncertainty of the protein structure data (Cao & Shen, 2020), state-of-the-art methods only predict several single solutions (Porter et al., 2019) without any associated uncertainty, which makes those predictions hard to interpret by the end users of protein docking methods (such as biologists).

Various optimization methods have been proposed in response to the need of uncertainty awareness. Stochastic optimization methods like random search (Zhigljavsky, 2012), simulated annealing (Kirkpatrick et al., 1983), genetic algorithms (Goldenberg, 1989) and particle swarm optimization (Kennedy & Eberhart, 1995) injected the randomness into the algorithms in order to reduce uncertainties. However, these methods do not provide the uncertainty quantification (UQ) of solutions. Recently, there has been growing interests in applying inference-based methods to optimization problems (Brochu et al., 2010; Shapiro, 2000; Pelikan et al., 1999). Generally, they transfer the uncertainties within the data and model into the final solution by modelling the posterior distribution over the global optima. For instance, Bijl et al. (2016) uses sequential Monte Carlo to approximate the distribution over the optima with Thompson sampling as the search strategy. Hernández-Lobato et al. (2014) uses kernel approximation for modelling the posterior over the optimum under Gaussian process. Ortega et al. (2012); Cao & Shen (2020) directly model the posterior over the optimum

as a Boltzmann distribution. They not only surpass the previous methods in accuracy and efficiency, but also provide easy-to-interpret uncertainty quantification.

Despite progress in optimization with uncertainty-awareness, significant open questions remain. Existing methods consider uncertainty either within the data or the model (including objective functions) (Kendall & Gal, 2017; Ortega et al., 2012; Cao & Shen, 2020). *However, no attention was ever paid to the uncertainty arising from the optimizer that is directly responsible for deriving the end solutions with given data and model.* The optimizer is usually pre-defined and fixed. For instance, there are several popular update rules in Bayesian optimization, such as expected improvement Vazquez & Bect (2010) or upper confidence bound Srinivas et al. (2009), that are chosen and fixed for the entire process. For Bayesian neural networks training, the update rule is usually chosen off-the-shelf, such as Adam, SGD, or RMSDrop. *The uncertainty in the optimizer is intrinsically defined over the optimizer space and important to uncertainty-aware optimization and solution UQ, and unwittingly ignored when the optimizer is treated as a fixed sample in the space.*

To fill the aforementioned gap, the **core intellectual value** of this work is to recognize and quantify a new form of uncertainty, that lies in the optimization algorithm (optimizer), besides the classical data- or model- based uncertainties (also known as epistemic and aleatoric uncertainties). The **underlying innovation** is to treat an optimizer as sample from an *algorithmic space*, i.e., a distribution of optimizers, rather than one of a few hand-crafted update rules. The **key enabling technique** is to consider an algorithmic space of optimizers, each of them being parameterized by a neural network. We then leverage a Boltzmann-shaped posterior over the optimizers, and approximate the posterior locally as Gaussian distributions through variational inference. Our approach, *Bayesian learning to optimize* (**BL2O**), **for the first time** addresses the modeling of the optimizer-based uncertainty. Extensive experiments on optimizing test functions, energy functions in a bioinformatics application, and loss functions in the image classification and data privacy attack demonstrate that compared to the start-of-art methods, BL2O substantially improves the performance of optimization and uncertainty quantification, as well as calibration and out-of-domain detection in classification.

2 RELATED WORK

Many works (Wang & Jegelka, 2017; Hennig & Schuler, 2012) studied optimization with uncertainty quantification under the framework of Bayesian Optimization (Shahriari et al., 2016; Brochu et al., 2010). In Bayesian optimization, multiple objectives are sampled from the posterior over the objectives ($p(f|\mathcal{D})$), where $\mathcal{D}$ is the observed data. Each sampled objective is optimized for obtaining samples of the global optima: $\boldsymbol{w}^*$ so that the empirical distribution over $\boldsymbol{w}^*$ can be built. However, since this approach needs optimization for each sample, which is extreme time-consuming, significant approximation is much needed. For instance, Henrández-Lobato et al. (2014) uses kernel approximation to approximate the posterior distirbution.

Another line of work uses various sampling schemes for estimating the density of posterior distributions. For instance, Bijl et al. (2016) uses sequential Monte Carlo sampling. De Bonet et al. (1997) designs a randomized optimization algorithm that directly samples global optima. These methods are much more efficient, but their performance heavily depends on the objective landscapes.

Finally, there are approaches (Ortega et al., 2012; Cao & Shen, 2020) that directly model the shape of posterior as the Boltzmann distributions: $p(\boldsymbol{w}^*|D) \propto exp(-\alpha f(\boldsymbol{w}))$, where α is the scheduled temperature constant. The algorithm automatically adjusts α during the search in order to balance the exploration-exploitation tradeoff. Those works beat previous work in terms of both efficiency and accuracy.

However, as revealed in the aforementioned gap, none of these methods consider the uncertainty within the optimizer. In the following sections, we first formally define the problem of optimization with uncertainty quantification. And then we briefly describe how to address the optimizer uncertainty in classical optimizations. Third, we propose our novel model, **BL2O**. And lastly, we compare our **BL2O** with both Bayesian and non-Bayesian competing methods on extensive test functions and real-world applications.

3 METHODS

3.1 PROBLEM STATEMENT

We use a bold-faced uppercase letter to denote a matrix (e.g. $\boldsymbol{W}$), a bold-faced lowercase letter to denote a vector (e.g. $\boldsymbol{w}$), and a normal lowercase letter to denote a scalar (e.g. w). The goal of optimization is to find the global optima for an objective function $f(\boldsymbol{w})$ w.r.t. $\boldsymbol{w}$:

$$\boldsymbol{w}^* = \arg\min_{\boldsymbol{w}} f(\boldsymbol{w}). \tag{1}$$

Such problems are usually solved by iterative optimization algorithms, whose update rules could be expressed as: $\boldsymbol{w}^{t+1} = \boldsymbol{w}^t + \delta\boldsymbol{w}^t$, where $\boldsymbol{w}^t$ and $\delta\boldsymbol{w}^t$ are the parameter vector and update vector at iteration t, respectively. $\delta\boldsymbol{w}^t$ is usually a function of past gradients:

$$\delta\boldsymbol{w}^t = g(\{\nabla\ell(\boldsymbol{w}^\tau)\}_{\tau=1}^t) \tag{2}$$

where $g(\cdot)$ is a pre-defined update rule (optimizer). For instance, in gradient descent, $g(\{\nabla\ell(\boldsymbol{w}^\tau)\}_{\tau=1}^t) = -\alpha\nabla\ell(\boldsymbol{w}^t)$, where α is the step size.

Once we obtain the final solution $\hat{\boldsymbol{w}}$, it is important to assess the quality of the solution as $||\hat{\boldsymbol{w}} - \boldsymbol{w}^*||$. However, since $\boldsymbol{w}^*$ is unknown, instead, we can provide the following probably approximately correct assessment of $\hat{\boldsymbol{w}}$:

$$P(||\hat{\boldsymbol{w}} - \boldsymbol{w}^*|| \leqslant r_\sigma|\mathcal{D}) = \sigma \tag{3}$$

where r_σ is the upper bound of $||\hat{\boldsymbol{w}} - \boldsymbol{w}^*||$ at σ confidence level, and $\mathcal{D} = \{\boldsymbol{w}_i, f(\boldsymbol{w}_i)\}_{i=1}^T$ are the observed data. It is straightforward to first model the posterior over the global optima ($p(\boldsymbol{w}^*|\mathcal{D})$) and then sample from the posterior to obtain Eq 3. If the optimizer g is fixed, then we can put g into the condition as $p(\boldsymbol{w}^*|\mathcal{D}, g)$.

3.2 OPTIMIZER UNCERTAINTY IN CLASSICAL METHODS

A simple way for considering the optimizer uncertainty for fixed optimizers (e.g. Adam, Gradient Descent) is running algorithms for multiple times with different hyperparameters and(or) different starting points for obtaining $\hat{\boldsymbol{w}}$ and modelling $p(\boldsymbol{w}^*|\mathcal{D})$. Such strategy can reduce the uncertainty within the optimizer to some extent. However, it heavily relies on the pre-defined distributions (discrete grids) over the hyperparameters and(or) the start points. More importantly, an optimizer cannot be just represented by its hyperparameters or starting points. The intrinsic uncertainty that lies in its update rule is not explored at all in this strategy. Such statements can be clearly demonstrated by extensive experiments in Sec. 4.

3.3 TREATING AN OPTIMIZER AS A SAMPLE FROM AN ALGORITHMIC SPACE

In order to consider the intrinsic uncertainty within the update rule, we assume that there exists an algorithm space $\mathcal{G}$ where each point $g \in \mathcal{G}$ is an optimizer. In order to quantify the uncertainty over the optimizer, a good way is to first model the posterior distribution over $\mathcal{G}$ as $p(g(\cdot)|\mathcal{D})$. And then integrate this posterior into the final posterior through:

$$p(\boldsymbol{w}^*|\mathcal{D}) = \int p(g(\cdot)|\mathcal{D})p(\boldsymbol{w}^*|\mathcal{D}, g)dg \tag{4}$$

However, modelling $p(g(\cdot)|\mathcal{D})$ seems intractable as there is no parameterization of the algorithmic space. In the next subsection, we will parameterize this algorithmic space by the parameters of a neural network so that modelling $p(g(\cdot)|\mathcal{D})$ becomes feasible.

3.4 PARAMETERIZING THE SPACE THROUGH A NEURAL NETWORK

Boltzmann-shaped posterior In order to reasonably and feasibly model the intrinsic uncertainty within the optimizer $g(\cdot)$, we parameterize $g(\cdot)$ as a neural network: $g_{\boldsymbol{\theta}}(\cdot)$, where $\boldsymbol{\theta}$ are the parameters in the neural network. Then $p(g(\cdot)|\mathcal{D})$ becomes $p(\boldsymbol{\theta}|\mathcal{D})$. We use the LSTM architecture from Andrychowicz et al. (2016) as our optimizer's architecture. To make $g_{\boldsymbol{\theta}}(\cdot)$ has good optimization performance, it is important to optimize the following loss:

$$F(\boldsymbol{\theta}) = \sum_{j=1}^{N} f_j(\hat{\boldsymbol{w}}_j) + \lambda||\boldsymbol{\theta}||_2, \tag{5}$$

where $f_1(\cdot), f_2(\cdot), ..., f_N(\cdot)$ are N objective functions used for training the neural network; $\hat{\boldsymbol{w}}_j$ is the final solution of applying $g_{\boldsymbol{\theta}(\cdot)}$ on jth objective and λ is the L2 regularization constant.

If the training objective $f_j(\cdot)$ is the loss function of a neural network classification task, then for the jth mini-batch samples consisting of $\{\boldsymbol{x}_{ij}, \boldsymbol{y}_{ij}\}_{i=1}^{M}$, we have the following training objective assuming independence among samples:

$$f_j(\hat{\boldsymbol{w}}_j) = -\sum_{i=1}^{M} \log p(\boldsymbol{y}_{ij}|\boldsymbol{x}_{ij}, \hat{\boldsymbol{w}}_j, \boldsymbol{\theta}), \tag{6}$$

Then we can re-write Eq 5 as:

$$F(\boldsymbol{\theta}) = -\sum_{j=1}^{N}\sum_{i=1}^{M} \log p(\boldsymbol{y}_{ij}|\boldsymbol{x}_{ij}, \hat{\boldsymbol{w}}_j, \boldsymbol{\theta}) + \lambda||\boldsymbol{\theta}||_2 \tag{7}$$

We consider the first term of Eq 7 as the negative log likelihood term, while the second term as the negative log-prior term. Then we could regard $F(\boldsymbol{\theta})$ as the negative logarithm of un-normalized posterior distribution over $\boldsymbol{\theta}$: $F(\boldsymbol{\theta}) \propto -\log p(\boldsymbol{\theta}|\mathcal{D})$. We then obtain the posterior as the following Boltzmann distribution:

$$p(\boldsymbol{\theta}|\mathcal{D}) \propto \exp(-F(\boldsymbol{\theta})), \tag{8}$$

In the general case of optimizing a function $f(\cdot)$ besides classification loss, we use a Boltzmann-shaped posterior where $F(\boldsymbol{\theta})$ is as defined in Eq. (5). Similar ideas were proposed in Ortega et al. (2012); Cao & Shen (2020).

Local Approximation and Bayesian Loss However, the above posterior distribution involves an integral in the normalization constant which is computationally intractable. Moreover, the architecture of $F(\boldsymbol{\theta})$ is so complicated that it is impossible to directly sample from the posterior distribution. In order to overcome the aforementioned challenges, we would like to learn a distribution function $q(\boldsymbol{\theta}|\boldsymbol{\phi})$ that has the analytic form and is easy to be sampled, where $\boldsymbol{\phi}$ is the parameter vector in $q(\boldsymbol{\theta}|\boldsymbol{\phi})$, to approximate the real posterior $p(\boldsymbol{\theta}|\mathcal{D})$.

Furthermore, due to the high dimensions of $\boldsymbol{\theta}$ and the complicated landscape of the posterior, it is impossible to approximate $p(\boldsymbol{\theta}|\mathcal{D})$ at every position in the $\boldsymbol{\theta}$ space. We then consider to approximate it locally around $\boldsymbol{\theta}^*$, an optimum of interest for $F(\boldsymbol{\theta})$.

We denote the local region as Θ^*, a neighborhood around $\boldsymbol{\theta}^*$, and re-normalization constant $C = \int_{\boldsymbol{\theta}\in\Theta^*} p(\boldsymbol{\theta}|\mathcal{D})d\boldsymbol{\theta}$. Then the local posterior will be a conditioned (re-scaled) version of $p(\boldsymbol{\theta}|\mathcal{D})$:

$$p'(\boldsymbol{\theta}|\mathcal{D}) = p(\boldsymbol{\theta}|\mathcal{D})/C, \ \boldsymbol{\theta} \in \Theta^* \tag{9}$$

In order to make $q(\boldsymbol{\theta}|\boldsymbol{\phi}) \approx p'(\boldsymbol{\theta}|\mathcal{D})$, we calculate the KL-divergence between these two:

$$\begin{aligned} \mathrm{KL}(q(\boldsymbol{\theta}|\boldsymbol{\phi})||p'(\boldsymbol{\theta}|\mathcal{D})) &= \int_{\boldsymbol{\theta}\in\Theta^*} q(\boldsymbol{\theta}|\boldsymbol{\phi}) \log \frac{q(\boldsymbol{\theta}|\boldsymbol{\phi})}{p'(\boldsymbol{\theta}|\mathcal{D})} d\boldsymbol{\theta} = \int_{\boldsymbol{\theta}\in\Theta^*} q(\boldsymbol{\theta}|\boldsymbol{\phi}) \log \frac{q(\boldsymbol{\theta}|\boldsymbol{\phi})}{p(\boldsymbol{\theta}|\mathcal{D})/C} d\boldsymbol{\theta} \\ &= \int_{\boldsymbol{\theta}\in\Theta^*} q(\boldsymbol{\theta}|\boldsymbol{\phi}) \log \frac{q(\boldsymbol{\theta}|\boldsymbol{\phi})}{\exp(-F(\boldsymbol{\theta}))} d\boldsymbol{\theta} + \int_{\boldsymbol{\theta}\in\Theta^*} q(\boldsymbol{\theta}|\boldsymbol{\phi}) \log(ZC) d\boldsymbol{\theta}, \end{aligned} \tag{10}$$

where $Z = \int \exp(-F(\boldsymbol{\theta}))d\boldsymbol{\theta}$ is the normalization constant. The second term in the above equation equals to $\log(ZC)$, a constant w.r.t. $\boldsymbol{\phi}$, thus could be ignored during optimization.

We then propose our Bayesian loss as:

$$\begin{aligned} F_{\mathrm{B}}(\boldsymbol{\phi}) &= \int_{\boldsymbol{\theta}\in\Theta^*} q(\boldsymbol{\theta}|\boldsymbol{\phi}) \log q(\boldsymbol{\theta}|\boldsymbol{\phi}) d\boldsymbol{\theta} + \int_{\boldsymbol{\theta}\in\Theta^*} q(\boldsymbol{\theta}|\boldsymbol{\phi}) F(\boldsymbol{\theta}) d\boldsymbol{\theta} \\ &= -H(q(\boldsymbol{\theta}|\boldsymbol{\phi})) + E_{q(\boldsymbol{\theta}|\boldsymbol{\phi})}[F(\boldsymbol{\theta})], \end{aligned} \tag{11}$$

where the first term of F_{B} measures the negative entropy of our approximated posterior, and the second term is the expectation of the loss function over of posterior.

Gaussian Posterior We consider $\boldsymbol{\phi} = (\boldsymbol{\mu}, \boldsymbol{\Sigma})$ and $q(\boldsymbol{\theta}|\boldsymbol{\phi}) = \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, where $\boldsymbol{\mu}$ is the mean vector and $\boldsymbol{\Sigma}$ is the covariance matrix of a normal distribution. For simplicity, we consider $\boldsymbol{\Sigma}$ to be a diagonal matrix: $\boldsymbol{\Sigma} = \text{diag}(\sigma_1^2, \sigma_2^2, \sigma_3^2, ...)$. The second term in Eq (11) involves the integral over $F(\boldsymbol{\theta})$, which is intractable. Therefore, we use Monte Carlo sampling through $q(\boldsymbol{\theta}|\boldsymbol{\phi})$ to replace the integral there. However, the direct sampling of the posterior parameters makes it difficult for the optimization as it is inaccessible to get the gradient w.r.t. $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$. Moreover, the standard deviation σ_1, σ_2, ... must be non-negative, making the optimization constrained.

To overcome those two challenges, we use the trick introduced in (Blundell et al., 2015) to shift sampling from $q(\boldsymbol{\theta}|\boldsymbol{\phi})$ to sampling from a standard normal distribution $\mathcal{N}(\mathbf{0}, \boldsymbol{I})$. And we reparameterize standard deviation σ_i to ρ_i as $\sigma_i = \log(1 + \exp(\rho_i))$. Then for any $\boldsymbol{\epsilon}$ sampled from $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{I})$, we could calculate $\boldsymbol{\theta}$ as $\boldsymbol{\theta} = \boldsymbol{u} + \log(1 + \exp(\boldsymbol{\rho}))$, where $\boldsymbol{\rho} = (\rho_1, \rho_2, ...)$.

3.5 Bayesian Averaging

We recall our goal to build the posterior over the global optimum: $p(\boldsymbol{w}^*|\mathcal{D})$ through Eq 4. We consider using Monte Carlo sampling to approximate the integral as:

$$p(\boldsymbol{w}^*|\mathcal{D}) = \int p(g(\cdot)|\mathcal{D}) p(\boldsymbol{w}^*|\mathcal{D}, g) dg \approx \int_{\boldsymbol{\theta} \in \Theta} q(\boldsymbol{\theta}|\boldsymbol{\phi}) p(\boldsymbol{w}^*|g_{\boldsymbol{\theta}}(\cdot), \mathcal{D}) d\boldsymbol{\theta} \approx \sum_{i=1}^{N} p(\boldsymbol{w}^*|g_{\boldsymbol{\theta}_i}(\cdot), \mathcal{D}) \tag{12}$$

where $\boldsymbol{\theta}_i$ is sampled from $q(\boldsymbol{\theta}|\boldsymbol{\phi})$ and $N = 10,000$.

3.6 Training Stages and Model Implementation

As mentioned before, our posterior is local around $\boldsymbol{\theta}^*$, an optimum of interest. To obtain $\boldsymbol{\theta}^*$, we first pre-train our model in a non-Bayesian way through optimizing the loss in Eq 5. We then use $\boldsymbol{\theta}^*$ as the warm start for $\boldsymbol{\mu}$, and start the second Bayesian training stage. The model is implemented in Tensorflow 1.13 (Abadi et al., 2016) and optimized by Adam (Kingma & Ba, 2014). Due to the coordinate-wise LSTM (Andrychowicz et al., 2016), our BL2O model only contains 10,282 free parameters. For all experiments, the length of LSTM is set to be 20. Both training stages include 5,000 training epochs.

3.7 Computational Complexity

The time complexity for BL2O is $O(KBN_e + KN_eH^2)$, where K is the number of sampling trajectories, B is the minibatch size, N_e is the number of objective parameters, and H is the hidden size of LSTM ($H = 20$ in the study). As the batch size increases, the computational cost is close to the traditional Bayesian neural networks trained through SGD. Due to the coordinate-wise LSTM, the space coomplexity (memory cost) of BL2O is only $O(H^2)$, which remains the same as the number of objective parameters varies. Both the time and the space complexity of BL2O are the same as DM_LSTM (Abadi et al., 2016), while those of Adam are $O(KBN_e)$ and $O(N_e)$, respectively.

4 Experiments

We test our BL2O model extensively on optimizing: non-convex test functions, energy functions in protein-protein interactions, loss functions in image classification and loss functions in data privacy attack. We compare BL2O to three non-Bayesian methods: Adam, Particle Swarm Optimization (PSO) (Kennedy & Eberhart, 1995), DM_LSTM (Andrychowicz et al., 2016) and a recently published Bayesian method, BAL (Cao & Shen, 2020). All algorithms are running for 10,000 times with random initializing points to obtain the empirical posterior distributions. During each run, the hyperparameters in Adam and PSO are sampled from Table 4 in Appendix A. Out of 10,000 solutions we choose the one with the lowest function value to be the final solution ($\hat{\boldsymbol{w}}$).

Generally, for optimization performance, we assess the distance between the final solution and the global optima: $||\hat{\boldsymbol{w}} - \boldsymbol{w}^*||$. The lower the distance is, the better the solution quality is. For uncertainty quantification, we assess the upper bound r_σ and the real confidence ϵ_σ given a fixed confidence level σ. The real confidence ϵ_σ is defined by the fraction of 10,000 solutions that actually fall in the bounded region. The lower the r_σ is, the tighter the confidence interval is. And the closer of ϵ_σ to σ is, the more accurate the confidence estimate is.

Comparison in optimizing test functions. We first test the performance on test functions in the global optimization benchmark set (Jamil & Yang, 2013). We choose three extremely rugged, non-convex functions: Rastrigin, Ackley and Griewank in 5 dimensions: 6D, 12D, 18D, 24D, 30D. For each function, we create a diverse, broad family of similar functions $f_j(\boldsymbol{w})$ as the meta-training set used for training DM_LSTM and BL2O. The analytical forms and the meta-training sets of those functions are shown in Table 5 in Appendix B.

We compare BL2O with all 4 competing methods. The optimization and UQ performances are shown in Fig. 1. In all three cases and 5 dimensions, BL2O has led to the best solution quality. In terms of UQ, BL2O has shown the most accurate confidence estimation ($\epsilon_\sigma \approx \sigma$) when $\sigma = 0.9$ and $\sigma = 0.8$, while BAL was the second best. And BL2O has shown much tighter confidence intervals r_σ against BAL. In some cases, although DM_LSTM has lower r_σ than BL2O, it has much lower confidence level, indicating that this tight upper bound in DM_LSTM is miscalibrated. As a result, BL2O has shown the best performance in both optimization and UQ.

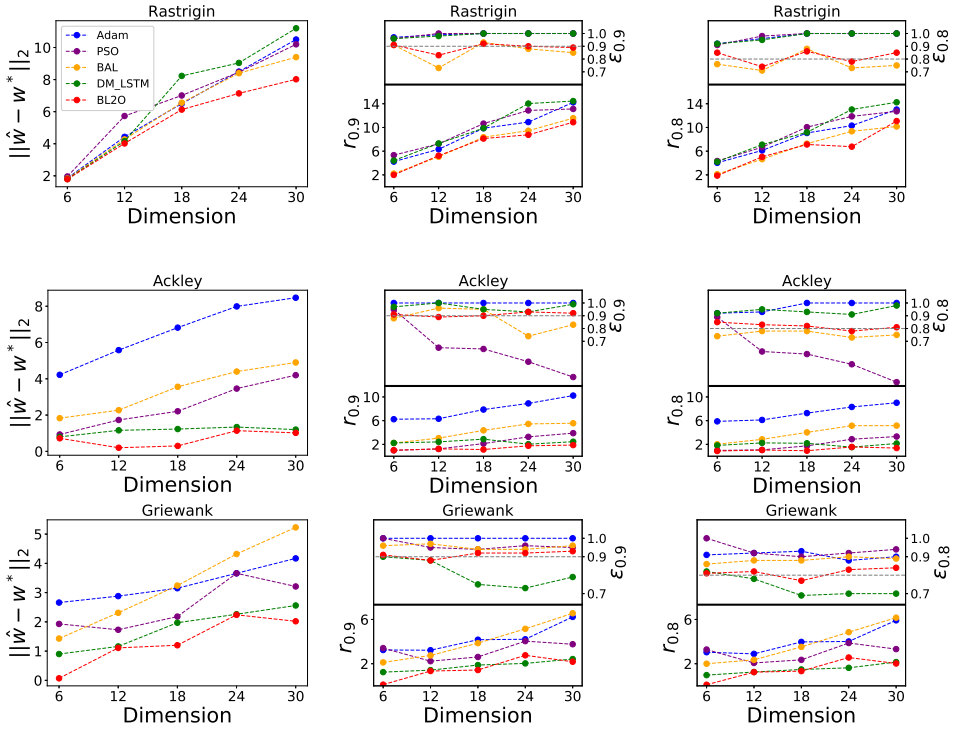

Figure 1: The optimization performance (left) and the UQ performance (r_σ and ϵ_σ) of different methods on three test functions.

Comparison in optimizing energy functions for protein docking. We then apply BL2O to a bioinformatics application: predicting the 3D structures of protein-complexes (Smith & Sternberg, 2002), called protein docking. *Ab initio* protein docking can be recast as optimizing a noisy and expensive energy function in a high-dimensional conformational space (Cao & Shen, 2020): $\boldsymbol{x}^* = \arg\min_{\boldsymbol{x}} f(\boldsymbol{x})$. While solving such optimization problems still remains difficult, quantifying the uncertainty of resulting optima (docking solutions) is even more challenging. In this section, we apply our BL2O to optimization and uncertainty quantification in protein docking and compare with a state-of-the-art method BAL (Cao & Shen, 2020).

We describe the detailed settings of BL2O on protein docking in Appendix C. From BL2O, we obtain a posterior distribution $p(\boldsymbol{w}^*|D)$ over the native structure $\boldsymbol{w}^*$ and the lowest energy structure, $\hat{\boldsymbol{w}}$. In protein docking, the quality of a predicted structure is based on the distance to the native structure (the global optimum): $||\hat{\boldsymbol{w}} - \boldsymbol{w}^*||$. For UQ, we assess the two-sided confidence interval at $\sigma = 0.9$ as $P(l_{0.9} \leqslant ||\hat{\boldsymbol{w}} - \boldsymbol{w}^*|| \leqslant r_{0.9}) = 0.9$.

In Table 1, we assess $||\hat{\boldsymbol{w}} - \boldsymbol{w}^*||$, $r_{0.9} - l_{0.9}$ and whether $||\hat{\boldsymbol{w}} - \boldsymbol{w}^*||$ is within the confidence interval. For optimization, BL2O clearly outperforms BAL in two medium cases while performing slightly

worse in the other cases. Yet for UQ, BL2O shows clearly superior performance over BAL in all cases, with accurate or/and tight confidence intervals. We also visulize the posterior distributions over $||\hat{\boldsymbol{w}} - \boldsymbol{w}^*||$ for protein 1JMO_4. As shown in Fig 2, we can see compared to that of BAL, BL2O's distribution has real $||\hat{\boldsymbol{w}} - \boldsymbol{w}^*||$ within the 90% C.I. and smaller variance. More posterior distributions are shown in Appendix D.

Table 1: Performances in optimization and uncertainty quantification on 5 docking cases.

	$\|\|\hat{\boldsymbol{w}} - \boldsymbol{w}^*\|\|$		$r_{0.9} - l_{0.9}$ (Å)		$\|\|\hat{\boldsymbol{w}} - \boldsymbol{w}^*\|\| \in [l_{0.9}, r_{0.9}]$?	
Target (docking diffficulty)	BAL	BL2O	BAL	BL2O	BAL	BL2O
1AHW_3 (easy)	1.89	2.07	2.20	1.98	No	No
1AK4_7 (easy)	2.45	2.70	1.93	1.66	Yes	Yes
3CPH_7 (medium)	3.89	3.21	1.70	2.20	No	Yes
1HE8_3 (medium)	3.05	2.32	2.24	1.61	Yes	Yes
1JMO_4 (difficult)	1.45	1.55	2.90	1.26	No	Yes

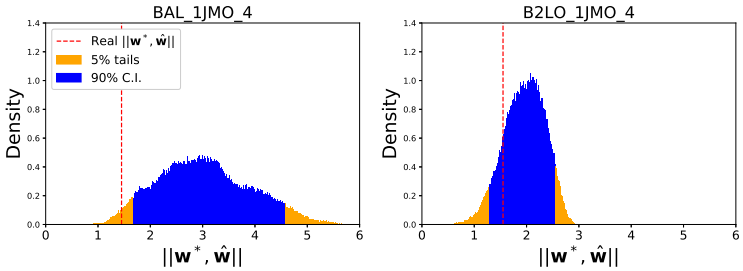

Figure 2: Visualizations of estimated posterior distributions and confidence intervals.

Comparison in optimizing loss functions in image classification. We then test the performance of optimizing the loss function in image classification on the MNIST dataset. We apply a 2-layers MLP network as the classifier. The competing methods include Adam, DM_LSTM and a Bayesian neural network method: variational inference (VI) (Blundell et al., 2015). Moreover, for DM_LSTM and BL2O, we apply a trick during the optimizer training called curriculum learning (CL) and introduce it in detail in Appendix E for training over long-term iterations. We call DM_LSTM with CL as DM_LSTM_C and BL2O with CL as BL2O_C.

The assessment of the optimization and UQ for this machine learning task is different from that for optimization before. In terms of optimization, we assess the classification accuracy on the test set. In terms of UQ, we measure two metrics that assess the robustness and trustworthiness of the classifier: the in-domain calibration error and the out-of-domain detection rate.

We first compare the accuracy on the testing set among different methods. As shown in Table 2, Adam, DM_LSTM_C and BL2O_C have almost the same best performance. The significant improvement from DM_LSTM to DM_LSTM_C, and from BL2O to BL2O_C shows the big advantage of curriculum learning in learning to optimize. In conclusion, BL2O_C had on par accuracy with Adam and DM_LSTM_C on the MNIST dataset.

However, classification models must not only be accurate, but also indicate when they are likely to be incorrect. Confidence calibration, the probability that estimates the true likelihood of each prediction is also important for classification models. In the ideal case, the maximum output probability (MaxConfidence) for each test sample should be equal to the prediction accuracy for that sample. To assess the calibration of each methods, we split the test set into 20 equal-sized bins and assess the **calibration error** as the average discrepancy between accuracy and MaxConfidence in each bin. As seen in Table 2, among all methods compared, BL2O_C and BL2O had the least calibration error. The figure of Acc. vs MaxConf. is also shown in Fig. 4 in Appendix E.

We also inspect the out-of-domain detection of BL2O, BL2O_C and competing methods. We train all models on the data belonging to the first 5 classes in the MNIST training dataset (the last layer of the optimizee is modified to have 5 rather than 10 neurons) and test them on the remaining samples from the other 5 classes. An ideal model would predict a uniform distribution over the 5 wrong classes.

Table 2: Performance of classification on the MNIST test set.

Models	In-Domain Accuracy (%)	In-Domain Calibration Error	Out-of-Domain Detection	
			$q_{.4}(\%)$	$q_{.5}(\%)$
Adam	93.2	5.0E-4	0.7	2.8
DM_LSTM	81.0	4.2E-3	0.4	2.0
DM_LSTM_C	93.4	9.9E-4	4.6	10.6
VI	87.8	4.4E-3	4.8	10.6
BL2O	90.1	**4.9E-4**	**29.9**	**31.8**
BL2O_C	**93.5**	**4.3E-4**	**12.4**	**20.9**

Therefore, we define the **out-of-domain detection rate** at threshold t, q_t, as the percentage of test samples with max class confidence below t. The larger the q_t, the better out-of-domain detection is. As shown in Table 2, BL2O and BL2O_C shows superior performance with all competing methods.

Comparison in optimizing loss functions for data privacy attack. We finally apply our model to an application that critically needs UQ. As many machine learning models are deployed publicly, it is important to avoid leaking private sensitive information, such as financial data, health data and so on. Data privacy attack (Nasr et al., 2018) studies this problem by playing the role of hackers and attacking the machine-learning models to quantify the risk of privacy leakage. Better attacks would help models to be better prepared for privacy defense.

We use the model and dataset in (Cao et al., 2019), where each input has 9 features involving patient genetic information and the output p is the probability of the clinical significance (having cancer or not) for a patient. We study the following model inversion attack (Fredrikson et al., 2015): by giving 5 features $\boldsymbol{w}' \in [0,1]^5$ out of 9 and the label p of each patient, we want to recover the rest 4 features $\boldsymbol{w}^* \in [0,1]^4$ (potentially sensitive patient information). Therefore, for each patient, the objective is

$$\boldsymbol{w}^* = \arg\min_{\boldsymbol{w}\in[0,1]^4} (m(\boldsymbol{w}', \boldsymbol{w}) - p)^2 \tag{13}$$

where $\boldsymbol{w}^*$ is the ground-truth of $\boldsymbol{w}$ and m is the trained predictive model. The closeness between the predicted and the real input features can quantify the risk of information leakage and the quality of the attack. We compare BL2O with Adam, PSO, BAL and DM_LSTM on optimization and UQ on all test cases in (Cao et al., 2019). The meta-training objectives for BL2O and DM_LSTM are the training set in (Cao et al., 2019).

As shown in Table 3, BL2O has shown the best performance in both optimization and UQ compared to all competing methods. It is noteworthy that learned optimizers (DM_LSTM and BL2O) had much better optimization performance than pre-defined optimizers. And the Bayesian methods (BAL and BL2O) had significantly better UQ performance than non-Bayesian methods. BL2O possessed the advantages of both learned and Bayesian optimizers to achieve the best performance.

Table 3: The optimization and UQ performance of different methods on data privacy attack.

	$\|\|\hat{\boldsymbol{w}} - \boldsymbol{w}^*\|\|$	$r_{0.9}$	$\|\epsilon_{0.9} - 0.9\|$	$r_{0.8}$	$\|\epsilon_{0.8} - 0.8\|$
Adam	0.45	0.74	0.10	0.63	0.20
PSO	0.32	0.82	0.10	0.72	0.20
BAL	0.34	0.53	0.06	0.41	0.08
DM_LSTM	0.20	0.52	0.10	0.47	0.19
BL2O	**0.17**	**0.43**	**0.05**	**0.32**	**0.04**

5 Conclusion

Current optimization algorithms, even with uncertainty-awareness, do not address the uncertainty arising within the optimizer itself. To close this gap, we parameterize the update rule as a neural network and build a Boltzmann-shaped posterior over the algorithmic space. We apply our Bayesian Learning-to-Optimize (BL2O) framework to optimize test functions, energy functions in protein docking, loss functions in image classification and loss functions in data privacy attack. The empirical results demonstrate that BL2O outperforms the state-of-the-art methods in both optimization and uncertainty quantification, as well as the calibration and out-of-domain detection in classification.

REFERENCES

Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, pp. 265–283, 2016.

Marcin Andrychowicz, Misha Denil, Sergio Gomez, Matthew W Hoffman, David Pfau, Tom Schaul, Brendan Shillingford, and Nando De Freitas. Learning to learn by gradient descent by gradient descent. In *Advances in Neural Information Processing Systems*, 2016.

Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pp. 41–48, 2009.

Hildo Bijl, Thomas B Schön, Jan-Willem van Wingerden, and Michel Verhaegen. A sequential monte carlo approach to thompson sampling for bayesian optimization. *arXiv preprint arXiv:1604.00169*, 2016.

Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural networks. *arXiv preprint arXiv:1505.05424*, 2015.

Eric Brochu, Vlad M Cora, and Nando De Freitas. A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*, 2010.

Yue Cao and Yang Shen. Bayesian active learning for optimization and uncertainty quantification in protein docking. *Journal of chemical theory and computation*, 16(8):5334–5347, 2020.

Yue Cao, Yuanfei Sun, Mostafa Karimi, Haoran Chen, Oluwaseyi Moronfoye, and Yang Shen. Predicting pathogenicity of missense variants with weakly supervised regression. *Human mutation*, 40(9):1579–1592, 2019.

Jeremy S De Bonet, Charles Lee Isbell Jr, and Paul A Viola. Mimic: Finding optima by estimating probability densities. In *Advances in neural information processing systems*, pp. 424–430, 1997.

Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. Model inversion attacks that exploit confidence information and basic countermeasures. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pp. 1322–1333, 2015.

David E Goldenberg. Genetic algorithms in search, optimization and machine learning, 1989.

Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 1321–1330. JMLR. org, 2017.

Philipp Hennig and Christian J Schuler. Entropy search for information-efficient global optimization. *The Journal of Machine Learning Research*, 13(1):1809–1837, 2012.

José Miguel Henrández-Lobato, Matthew W. Hoffman, and Zoubin Ghahramani. Predictive Entropy Search for Efficient Global Optimization of Black-box Functions. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 1*, NIPS'14, pp. 918–926, Cambridge, MA, USA, 2014. MIT Press. URL `http://dl.acm.org/citation.cfm?id=2968826.2968929`.

José Miguel Hernández-Lobato, Matthew W Hoffman, and Zoubin Ghahramani. Predictive entropy search for efficient global optimization of black-box functions. In *Advances in neural information processing systems*, pp. 918–926, 2014.

Howook Hwang, Thom Vreven, Joël Janin, and Zhiping Weng. Protein-Protein Docking Benchmark Version 4.0. *Proteins*, 78(15):3111–3114, November 2010. ISSN 0887-3585. doi: 10.1002/prot.22830. URL `https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2958056/`.

Momin Jamil and Xin-She Yang. A literature survey of benchmark functions for global optimisation problems. *International Journal of Mathematical Modelling and Numerical Optimisation*, 4(2): 150–194, 2013.

Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision?, 2017.

J Kennedy and R Eberhart. Particle swarm optimization, proceedings of ieee international conference on neural networks (icnn'95) in, 1995.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Scott Kirkpatrick, C Daniel Gelatt, and Mario P Vecchi. Optimization by simulated annealing. *science*, 220(4598):671–680, 1983.

Iain H. Moal and Paul A. Bates. SwarmDock and the Use of Normal Modes in Protein-Protein Docking. *International Journal of Molecular Sciences*, 11(10):3623–3648, September 2010. ISSN 1422-0067. doi: 10.3390/ijms11103623. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2996808/.

Milad Nasr, Reza Shokri, and Amir Houmansadr. Comprehensive privacy analysis of deep learning: Stand-alone and federated learning under passive and active white-box inference attacks. *arXiv preprint arXiv:1812.00910*, 2018.

Pedro Ortega, Jordi Grau-Moya, Tim Genewein, David Balduzzi, and Daniel Braun. A nonparametric conjugate prior distribution for the maximizing argument of a noisy function. In *Advances in Neural Information Processing Systems*, pp. 3005–3013, 2012.

Martin Pelikan, David E Goldberg, Erick Cantú-Paz, et al. Boa: The bayesian optimization algorithm. In *Proceedings of the genetic and evolutionary computation conference GECCO-99*, volume 1, pp. 525–532. Citeseer, 1999.

Brian G. Pierce, Kevin Wiehe, Howook Hwang, Bong-Hyun Kim, Thom Vreven, and Zhiping Weng. ZDOCK server: interactive docking prediction of protein–protein complexes and symmetric multimers. *Bioinformatics*, 30(12):1771–1773, 02 2014. ISSN 1367-4803. doi: 10.1093/bioinformatics/btu097. URL https://doi.org/10.1093/bioinformatics/btu097.

K. A. Porter, I. Desta, D. Kozakov, and S. Vajda. What method to use for protein-protein docking? *Curr. Opin. Struct. Biol.*, 55:1–7, 04 2019.

B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas. Taking the Human Out of the Loop: A Review of Bayesian Optimization. *Proceedings of the IEEE*, 104(1):148–175, January 2016. ISSN 0018-9219. doi: 10.1109/JPROC.2015.2494218.

Alexander Shapiro. Probabilistic constrained optimization: Methodology and applications. *Statistical inference of stochastic optimization problems*, pp. 282–304, 2000.

Graham R Smith and Michael JE Sternberg. Prediction of protein–protein interactions by docking methods. *Current opinion in structural biology*, 12(1):28–35, 2002.

Niranjan Srinivas, Andreas Krause, Sham M Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv preprint arXiv:0912.3995*, 2009.

Emmanuel Vazquez and Julien Bect. Convergence properties of the expected improvement algorithm with fixed mean and covariance functions. *Journal of Statistical Planning and inference*, 140(11):3088–3095, 2010.

Zi Wang and Stefanie Jegelka. Max-value entropy search for efficient bayesian optimization. *arXiv preprint arXiv:1703.01968*, 2017.

Anatoly A Zhigljavsky. *Theory of global random search*, volume 65. Springer Science & Business Media, 2012.

A Optimizer distribution settings for Adam and PSO

Methods	Optimizer Distribution Settings
Adam	lr $\sim$ U[0.01, 0.1]
PSO	w $\sim$ U[0.5, 1.5], C1 $\sim$ U[1.5, 2.5], C2 $\sim$ U[1.5, 2.5]

Table 4: The optimizer distributions over hyperparameters in Adam and PSO.

B ANALYTIC FORMS AND META-TRAINING SETS OF TEST FUNCTIONS

Function Name	Analytic form
Rastrigin	$f_j(\boldsymbol{w}) = ||\boldsymbol{w}||_2^2 - \sum_{i=1}^{n} 10\cos(2\pi w_i) + 10n$
Ackley	$f_j(\boldsymbol{w}) = -20\exp(-0.2\sqrt{0.5||\boldsymbol{w}||_2^2}) - \sum_{i=1}^{n} \exp(\cos(2\pi w_i)/n) + e + 20$
Griewank	$f_j(\boldsymbol{w}) = 1 + \frac{1}{4000}\sum_{i=1}^{n} ||\boldsymbol{w}||_2^2 - \prod_{i=1}^{n} \cos(w_i)$
	Meta-training set
Rastrigin	$f_j(\boldsymbol{w}) = ||\boldsymbol{A}_j\boldsymbol{w} - \boldsymbol{b}_j||_2^2 - 10\boldsymbol{c}_j\cos(2\pi\boldsymbol{w}) + 10n$
Ackley	$f_j(\boldsymbol{w}) = -20\exp(-0.2\sqrt{0.5||\boldsymbol{A}_j\boldsymbol{w} - \boldsymbol{b}_j||_2^2}) - \exp(\boldsymbol{c}_j\cos(2\pi\boldsymbol{w})/n)$
Griewank	$f_j(\boldsymbol{w}) = 1 + \frac{1}{4000}\sum_{i=1}^{n} ||\boldsymbol{A}_j\boldsymbol{w} - \boldsymbol{b}_j||_2^2 - \prod_{i=1}^{n}(\cos(w_i) + c_{ji} - 1)$

Table 5: The analytic forms of test functions and the meta-training sets used for training DM_LSTM and BL2O, where n is the dimension and $\boldsymbol{A}_j \in \mathbb{R}^{n\times n}$, $\boldsymbol{b}_j \in \mathbb{R}^{n\times 1}$ and $\boldsymbol{c}_j \in \mathbb{R}^{n\times 1}$ are parameters whose elements are sampled from i.i.d. normal distributions. It is obvious that the test functions are the special cases in the training sets, respectively.

C SETTINGS FOR PROTEIN DOCKING EXPERIMENTS

We calculate the energy function (objective function $f(\boldsymbol{x})$) in a CHARMM 19 force field as in (Moal & Bates, 2010). 25 protein-protein complexes are chosen from the protein docking benchmark set 4.0 (Hwang et al., 2010) as the training set, which is shown in Table 6. For each target, we choose 5 starting points (top-5 models from ZDOCK (Pierce et al., 2014)). In total, our training set includes 125 samples. Moreover, we parameterize the search space as $\mathbb{R}^{12}$ as in BAL (Cao & Shen, 2020). The resulting $f(\boldsymbol{x})$ is fully differentiable in the search space. We only consider 100 interface atoms due to the computational concern. The number of iterations for one training epoch is 600 and in total we have 5000 training epochs. Both BL2O and BAL have 600 iterations during the testing stage. For fair comparison, after optimization, we rescore the BL2O samples for UQ using the scoring function (random forest) in BAL.

Table 6: 4-letter ID of proteins used in the training set.

Difficulty level	Protein Data Bank (PDB) code
Rigid	1N8O , 7CEI , 1DFJ , 1AVX , 1BVN , 1IQD , 1CGI , 1MAH , 1EZU , 1JPS , 1PPE , 1R0R , 2I25 , 2B42 , 1EAW , 2JEL , 1BJ1 , 1KXQ , 1EWY
Medium	1XQS , 1M10, 1IJK ,1GRN
Flexible	1IBR , 1ATN

D POSTERIOR DISTRIBUTIONS IN PROTEIN DOCKING

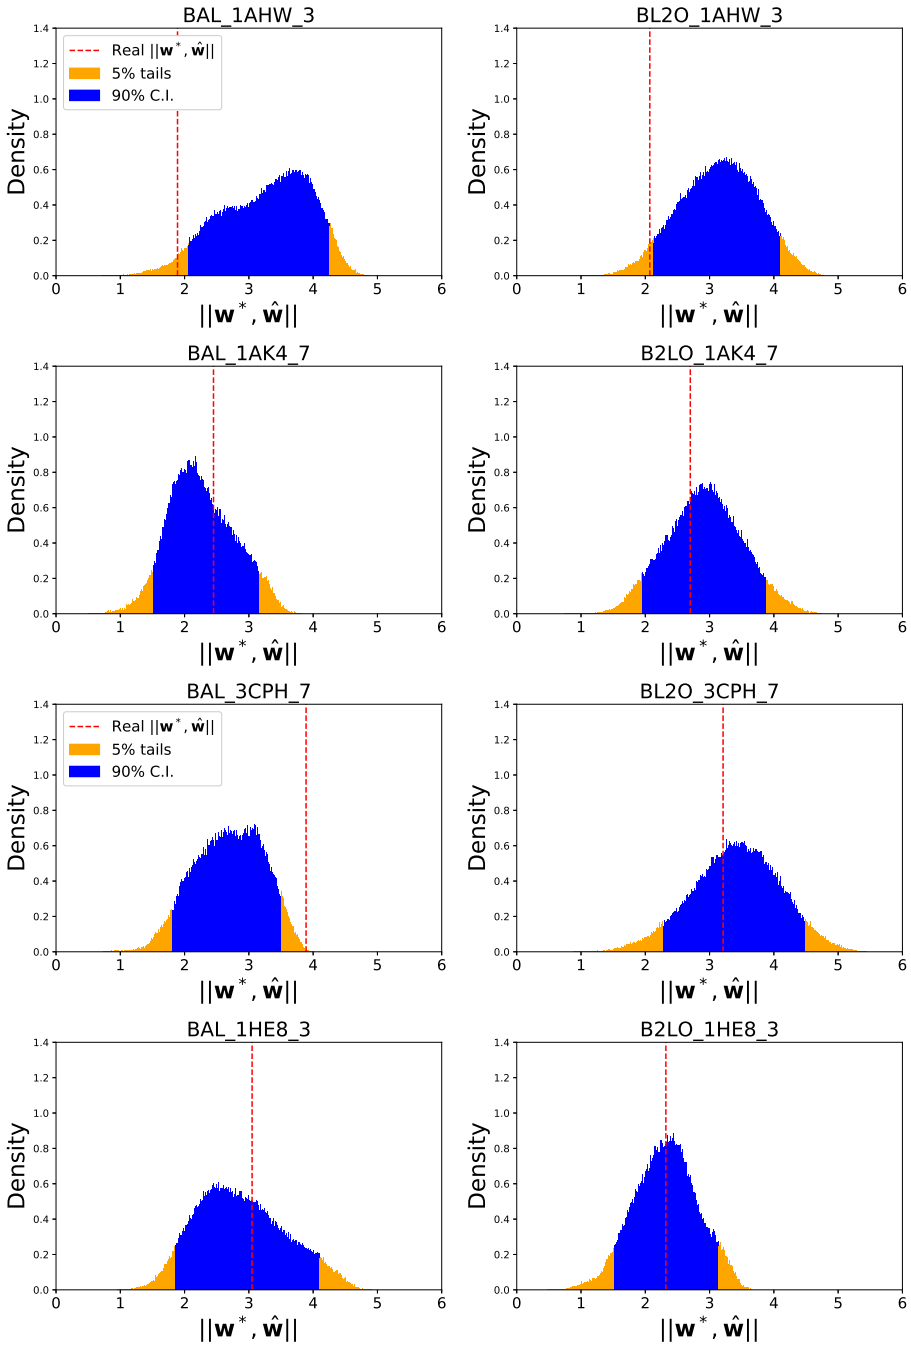

Figure 3: Visualizations of estimated posterior distributions and confidence intervals for more docking cases.

E CURRICULUM LEARNING AND CALIBRATION FIGURE IN IMAGE CLASSIFICATION

A common issue in learning to optimize for neural network training is that: the optimizer training usually takes hundreds of iterations, while training a neural network usually costs thousands or tens of thousand iterations. For instance, for a MNIST training dataset consisting of 50,000 images, training a neural network for 100 epochs corresponds to almost 20,000 iterations with a batch size of 128.

Hundreds of iterations are good for the first few epochs during optimizer training, since we would like to only focus on decreasing the loss in the first hundreds of iterations. However, as training goes on, we would like to see that the later iterations could also decrease the training loss. During this stage, only hundreds of iterations are clearly not enough.

In order to overcome this issue, we bring the idea from (Bengio et al., 2009). Specifically, We set a list for the number of iterations: [100, 200, 500, 1000, 1500, 2000, 2500, 3000] and we gradually increase the number of iterations following the list for every 100 epochs in optimizer training if the optimizee loss is decreasing. Once the number reaches 3000, it will not change any more until the training ends.

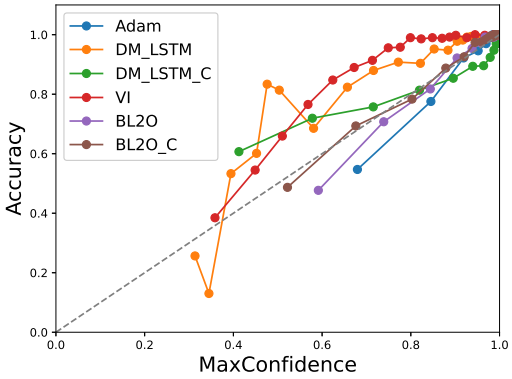

Figure 4: The calibration figure in image classification.