

MetricOpt: Learning to Optimize Black-Box Evaluation Metrics

Chen Huang Shuangfei Zhai Pengsheng Guo Josh Susskind
Apple Inc.

{chen-huang, szhai, pengsheng-guo, jsusskind}@apple.com

Abstract

We study the problem of directly optimizing arbitrary non-differentiable task evaluation metrics such as misclassification rate and recall. Our method, named MetricOpt, operates in a black-box setting where the computational details of the target metric are unknown. We achieve this by learning a differentiable value function, which maps compact task-specific model parameters to metric observations. The learned value function is easily pluggable into existing optimizers like SGD and Adam, and is effective for rapidly finetuning a pre-trained model. This leads to consistent improvements since the value function provides effective metric supervision during finetuning, and helps to correct the potential bias of loss-only supervision. MetricOpt achieves state-of-the-art performance on a variety of metrics for (image) classification, image retrieval and object detection. Solid benefits are found over competing methods, which often involve complex loss design or adaptation. MetricOpt also generalizes well to new tasks and model architectures.

1. Introduction

In real-world vision applications, machine learning models are usually evaluated on a variety of complex evaluation metrics. For example, one may evaluate a classification model using Mis-Classification Rate (MCR), and evaluate a ranking model using recall. Many of these metrics are non-continuous, non-differentiable, or non-decomposable, which poses challenges for direct metric optimization due to the difficulty of obtaining an informative gradient (e.g., it is zero almost everywhere for MCR). In other scenarios, the computational details may be unknown for a black box metric function. Hence its true gradient is simply inaccessible, which further increases the challenge of metric optimization.

As a common practice, people usually rely on a surrogate differentiable loss, which can be easily optimized with Stochastic Gradient Descent (SGD). One example is the widely used cross-entropy loss [15] for classification problems. While cross-entropy loss can be regarded as a

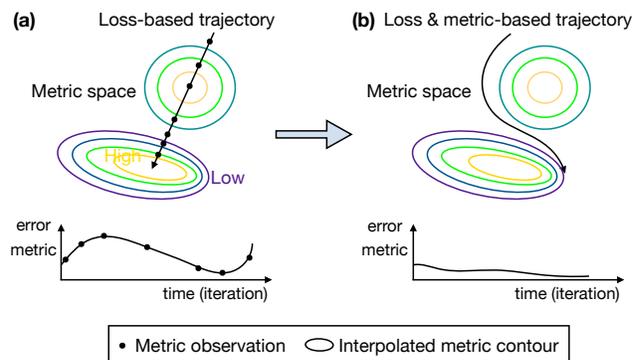


Figure 1. **Motivation** (a) Optimization with loss-only supervision may travel through several “bumps” in the metric space, and tends to converge to a suboptimal solution in terms of evaluation metric. We propose to collect a sparse set of (black-box) metrics along the optimization trajectory. Then we use the temporally interpolated metrics to meta-learn a differentiable value function, which can provide effective metric supervision to augment and improve any loss optimization process. (b) This way, we find continuous decrease in the error metric while keeping the loss from growing.

smooth relaxation of MCR, this loss is not a good proxy for other metrics like recall. When the loss does not match the target metric, inferior performance can be obtained [23]. State-of-the-art approaches follow two main paradigms to address this loss-metric mismatch issue. One is to introduce better metric-aligned surrogate losses, e.g., AUCPR loss [9]. These hand-designed losses not only require tedious manual effort and white-box metric formulation, but also tend to be specific to a given metric. Another paradigm is to learn adaptive losses in a relaxed or interpolated surrogate space [16, 23, 24, 28], which is inherently sub-optimal when compared to optimization in the original space.

In this paper, we propose to directly adapt the gradient-based optimization process to optimize *black-box* metrics, without knowing any details about the metric function. To do so, we first meta-learn a differentiable value function to model the metric observations along optimization trajectories. We focus on the model finetuning setup which can provide meaningful metrics to learn our value function. Once learned, the value function can provide useful metric super-

vision, including approximate metric gradients, to augment surrogate loss gradients. As a result, we can use the value function to finetune a new model that has been pre-trained using any given surrogate loss. Fig. 1 illustrates the high level idea. Intuitively, the value function is trained to produce meaningful adjustments to the optimization trajectory driven by loss only, leading to corrective directions on the metric landscape.

In practice, we parameterize the value function by a lightweight network for fast training and inference speeds. The input are a small set of *adapter* parameters that modulate a pre-trained model. Such a compact parameterization has been shown to suffice for task specialization [49], and removes our need to finetune the entire network. We meta-learn our value function using an enhanced ordinal regression objective, which is uncertainty-aware to avoid overconfident metric estimates. We then show that it is straightforward to apply our value function to off-the-shelf optimizers like SGD and Adam [25], and also to a learned optimizer.

The resulting method *MetricOpt* is shown to consistently improve different evaluation metrics across the tasks of (image) classification, image retrieval and object detection. *MetricOpt* not only outperforms prior methods based on strong surrogate losses (either hand-designed or adaptively learned), but also shares speed advantages as a fast finetuning method, often with no more than thousands of tuning steps. Furthermore, *MetricOpt* generalizes well to new tasks, *e.g.*, from CIFAR-10 [27] to ImageNet [7] classification. We summarize our contributions as follows:

- We introduce a differentiable value function to model a black-box evaluation metric.
- We show the value function is easily pluggable to existing optimizers, resulting in a fast finetuning approach.
- We show *MetricOpt* consistently improves over different surrogate losses without tedious loss engineering, achieves state-of-the-art performance for various tasks and metrics, and generalizes to out-of-distribution tasks and model architectures.

2. Related Work

Optimizing evaluation metrics Due to the non-differentiable nature of most real-life evaluation metrics, a large body of surrogate loss functions have been proposed as smooth metric relaxations. Examples include AUCPR loss [9], pairwise AUCROC loss [38], Lovász-Softmax loss [3] for IoU metric, and cost-sensitive classification for F-measure [37]. To remove the manual effort to design these metric-approximating losses, there has been a recent push towards learning them instead [46, 47]. However, the loss learning is still based on metric relaxation schemes, whereas our focus is on a direct metric representation to be used for optimization. Existing options for direct metric

optimization are post-shift methods [26, 33] that tune model threshold accordingly, and direct loss minimization methods [18, 43] that embed the true metric as a correction term for optimization. Their common limitation is that they require the evaluation metric to be available in closed-form. Some recent works aim to optimize black-box metrics, from learning adaptive losses [16, 23, 24, 28] to learning adaptive example weights [48]. However, these methods can be sub-optimal since they work in a relaxed surrogate space or with predefined weighting schemes. By contrast, we learn to directly model black-box metrics which can in turn adapt the optimization process.

Computer vision metrics are frequently defined as rank-based ones, such as recall for image retrieval, and Average Precision (AP) for object detection and feature matching, *etc.* Prior works approach the AP metric optimization via histogram binning approximations [4, 19, 20] and gradient approximation [22, 32, 43]. While these works mainly focus on AP relaxation, we differ in learning a fast and generic function approximation of the *original* metric. Other techniques have been proposed to learn the ranking operation with a large LSTM [10], or optimize an AP-loss using error-driven update [5], but both suffer from low efficiency. More recently, black-box differentiation [41] is used for efficient interpolation of the ranking function. More similar to our work is the deep embedding method [35] where the model and embedding space are learned alternately such that the Euclidean distance between the embedded prediction and groundtruth approximates the metric value. Our paper introduces a value function to regress metrics directly. The value function can be learned and deployed separately, runs fast and generalizes.

Meta learning The goal of meta-learning is to infer a learning strategy from a distribution of similar tasks, which facilitates efficient learning of new tasks from the same distribution. Previous works have approached this problem by learning a good initialization [12, 34, 49], preconditioning matrix [13] or mathematical update equation [2]. We use the efficient Reptile algorithm [34] to meta-learn the initialization of our value function in an online fashion.

Value function approximation Many Reinforcement Learning (RL) methods require value function approximation to estimate an action-value for fast decision making. Popular off-policy actor-critic methods [14, 17] are such examples. Their critic uses a differentiable function approximator to provide a loss, which guides the actor to update its action policy. The notion of “value function” in RL is related to the idea of learning a parametric loss as a metric approximator in the optimization field [23, 16], in order to provide some proxy gradients. Here we directly learn a differentiable mapping to evaluation metrics, bypassing the tedious step of loss function engineering.

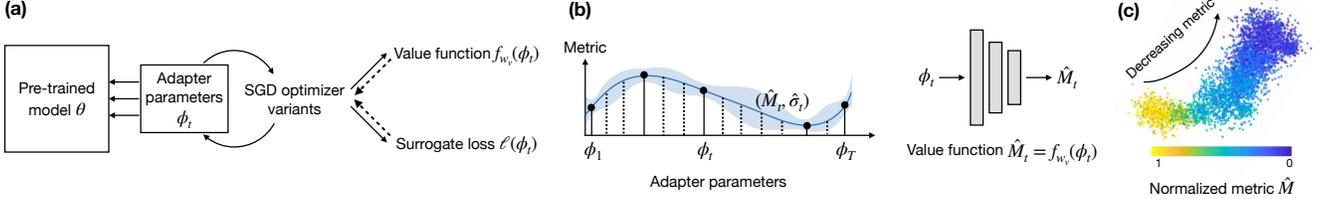


Figure 2. (a) Alongside finetuning the adapter parameters ϕ (e.g., FiLM conditioning parameters [36]) of a pre-trained model θ , we learn a value function f_{w_v} that maps ϕ to the non-differentiable metric \mathcal{M} . The learned value function is applicable to existing optimizers (e.g., SGD and Adam) in a plug-and-play manner to augment loss-based optimization. (b) We parameterize f_{w_v} by a lightweight network, which is meta-trained from different finetuning trajectories (solid stems – sparse metric observations, dotted lines – interpolated metrics $\hat{\mathcal{M}}_t$ and variances $\hat{\sigma}_t$). (c) t-SNE [44] visualization of 2D feature embeddings of f_{w_v} , where the ordinal ranks of metrics are well preserved.

3. Methodology

We consider optimizing the evaluation metric \mathcal{M} of a neural network model in a finetuning setup, where the model weights θ have been pre-trained using a user-specified surrogate loss. Instead of finetuning the high-dimensional θ which is costly, we optimize the conditioning parameters $\phi \in \Phi$ of a small adapter module that modulates θ (detailed later). Assume the surrogate loss is given as $\ell(\theta, \phi)$ for the tuple (θ, ϕ) . For brevity, we will, moving forward, use $\ell(\phi)$ in place of $\ell(\theta, \phi)$ since θ is fixed. Let $\mathcal{M}(\phi)$ be the metric subject to minimization. We also assume $\ell(\phi)$ and $\mathcal{M}(\phi)$ are normalized into the range of $[0, 1]$ for scale-independent modeling across tasks.

Generally, the metric function of \mathcal{M} is either non-differentiable (true gradient of the metric is often not meaningful), or simply unknown (black box setting with no access to the gradient). In both cases, we are faced with the great challenge of direct metric optimization. Here, we assume the only operation available to us is the ability to query metrics from a meaningful set of examples, without knowing the computational details of \mathcal{M} . Evaluation can be conducted on a held-out validation set D_{val} or on the training set D_{train} . We consider optimizing the metric in the form:

$$\phi^* = \arg \min_{\phi \in \Phi} \lambda \mathcal{M}(\phi) + \ell(\phi), \quad (1)$$

where λ is a weighting parameter of $\mathcal{M}(\phi)$, and loss $\ell(\phi)$ acts as an auxiliary signal. Without metric supervision, Eq. (1) is reduced to the standard loss optimization problem, which often leads to the loss-metric mismatch and hence suboptimal performance in evaluation metrics. On the other hand, without the surrogate loss, the supervision from metric alone is often sparse and limiting. We will show in the paper that the metric supervision always boosts the loss-based performance without any bells and whistles.

3.1. Meta learning differentiable value function

In Eq. (1), we aim at efficient optimization of the metric term $\mathcal{M}(\phi)$ via gradient-based methods. The core component of our approach is the value function $f_{w_v} : \Phi \rightarrow \mathbb{R}$, a differentiable mapping from input $\phi \in \Phi$ to metric

$\mathcal{M} \in \mathbb{R}$. We parameterize f by a deep neural network with weights w_v . It has the benefits of being able to output both metric estimate for any ϕ , as well as partial derivatives $\partial f / \partial \phi$ w.r.t. parameters ϕ to augment the loss gradients $\partial \ell / \partial \phi$. In other words, the value function is a generic function approximation of \mathcal{M} , which can offer useful supervision for the metric and gradient-based optimizability. As a result, finetuning can be accomplished using both the loss function and value function, without a significant computational overhead, see Fig. 2(a).

In practice, we train the value function from different optimization tasks $\mathcal{T} \sim p(\mathcal{T})$, where each task finetunes randomly initialized ϕ with random mini-batches (two sources of task variability). The finetuning process ensures the value function can be learned with meaningful metrics around converged θ , rather than with noisy metrics collected by a “from-scratch” model. This simplifies value function learning and gives rise to a fast finetuning process. Below we address two main challenges for value function learning.

Adapter parameters The high dimensionality of input parameters θ (often millions) renders value function learning parameter-inefficient. To this end, we have tried using a subset of θ in early experiments, e.g., only the layer biases or last layer. However, we found these approaches still not parameter-efficient enough, and they cannot strike a good performance-efficiency tradeoff.

In this paper, we turn to using the *adapter* modules that offer an equally simple but better performing solution. Existing adapters like conditional BatchNorm [6] and FiLM layers [36] introduce a few conditioning parameters $\phi \in \mathbb{R}^d$ to modulate the fixed θ . They have been shown to suffice for task specialization, with comparable performance to full model fine-tuning. We follow [49] to use dynamic biases at the first layer of fully connected networks, and use FiLM parameters to modulate the feature maps of convolutional networks (visualizations in supplementary materials). These two types of ϕ have small dimension d on the order of tens to hundreds, and are found effective in our experiments. We are open to other choices.

Metric interpolation Another challenge for value function training is the lack of training signals. Usually we can

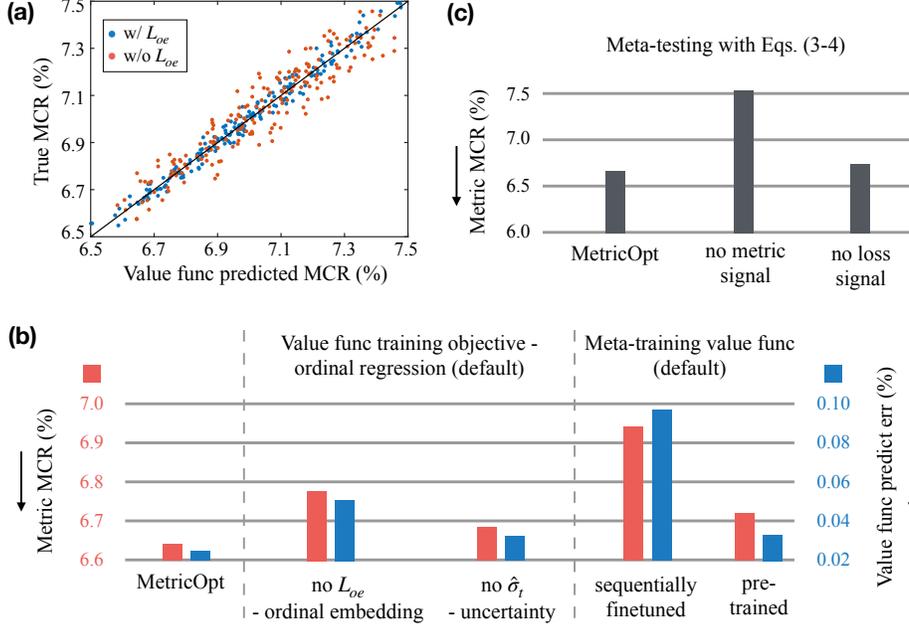


Figure 3. Optimizing Mis-Classification Rate (MCR) on CIFAR10. (a) Value function trained with and without the ordinal embedding constraint L_{oe} . The latter has more scattered predictions and higher prediction error (0.052 vs. 0.023). (b) Value function prediction error highly correlates with the final optimized metric, thus can serve as a performance indicator to diagnose *e.g.*, meta-training alternatives for the value function. (c) Comparing different meta-testing schemes.

only collect $K \ll T$ metrics $\{\mathcal{M}_k = \mathcal{M}(\phi_k)\}_{k=1}^K$ during T finetuning steps, where fast metric evaluation may not be an option. For instance, metrics like recall require evaluation on sizable data to be statistically significant, and the large evaluation cost makes frequent evaluation impractical. For this reason, we simply use the K sparse metric observations, with each \mathcal{M}_k evaluated at user-specified time $k \in [1, \dots, T]$. We further interpolate metrics over time to enrich training signals for our value function. This incurs much lower cost than using heavier metric evaluations, but has similar performance empirically. Specifically, we use a standard Gaussian process with RBF kernel [39], and interpolate metrics into a dense time series $\{(\hat{\mathcal{M}}_t, \hat{\sigma}_t)\}_{t=1}^T$ with mean $\hat{\mathcal{M}}_t$ and variance $\hat{\sigma}_t$ at time t , see Fig. 2(b). Compared to linear interpolation like line fitting, the Gaussian process can derive useful uncertainties for the interpolated metrics, which is important to avoid overconfident supervision for value function learning.

Value function parameterization Given sequence $\{\phi_t\}_{t=1}^T$ and the corresponding ‘‘labels’’ $\{(\hat{\mathcal{M}}_t, \hat{\sigma}_t)\}_{t=1}^T$, we can now train the mapping $\hat{\mathcal{M}}_t = f_{w_v}(\phi_t)$ as our value function. We parametrize f_{w_v} by an MLP (Multi-Layer Perceptron) network with architecture d–64–32–32–16–1. BatchNorm and ReLU activation are used for all layers. The network is lightweight for fast metric regression and gradient-based optimization, but can further benefit from improved design choices for particular problems or metrics.

Value function training objective consists of a regres-

sion term $L_{regress}$ and an ordinal embedding term L_{oe} :

$$L_v(w_v) = \gamma L_{regress} + L_{oe}, \quad (2)$$

$$L_{regress} = \frac{1}{\sum_{t=1}^T 1/\hat{\sigma}_t} \sum_{t=1}^T \frac{\|f_{w_v}(\phi_t) - \hat{\mathcal{M}}_t\|_2}{\hat{\sigma}_t},$$

$$L_{oe} = \frac{1}{T} \sum_{t=1}^T \log(1 + \exp(-(D_{t,t_n} - D_{t,t_p}))),$$

where γ is a weighting parameter, $D_{t,t'} = \|g_{w_v}(\phi_t) - g_{w_v}(\phi_{t'})\|_2$ is the Euclidean distance between the feature embeddings $g_{w_v}(\cdot)$ of our value function (penultimate layer), and (t, t_p, t_n) is the sampled triplet at time t for embedding learning.

Note $L_{regress}$ is aware of the uncertainty $\hat{\sigma}_t$ that helps to re-weight metric regression errors. While the L_{oe} term imposes ordinal constraints for the metric in feature space. Concretely, we sample triplets (t, t_p, t_n) , with anchor $t \in [1, \dots, T]$, its positive sample $t_p \in \mathcal{P}_t$ and negative sample $t_n \in \mathcal{N}_t$ for embedding learning. To do so, we leverage Fisher’s ratio $r_{t,t'} = (\hat{\mathcal{M}}_t - \hat{\mathcal{M}}_{t'})^2 / (\hat{\sigma}_t^2 + \hat{\sigma}_{t'}^2)$ that measures the discrimination between two metric variables. The Fisher’s ratio offers us great convenience to construct the positive set $\mathcal{P}_t = \{t' | r_{t,t'} < 2\}$ and negative set $\mathcal{N}_t = \{t' | r_{t,t'} \geq 2\}$ for every time step t . Then we perform hard mining within \mathcal{P}_t and \mathcal{N}_t to sample the required t_p and t_n , in a similar way to [45]. In general, L_{oe} provides effective regularization in feature space (see Fig. 2(c)) and eases

our value function training. Fig. 3(a-b) shows improved performance by L_{oe} . Fig. 3(b) also validates the need to encode metric uncertainties $\hat{\sigma}_t$ for value function learning.

Meta-training value function For meta-training f_{w_v} , we have access to different finetuning trajectories of tasks $\mathcal{T} \sim p(\mathcal{T})$ and their densely interpolated metrics. Task variability comes from random initialization ϕ_0 and mini-batches. Then using the given training data and objective (Eq. (2)), how do we train f_{w_v} effectively and efficiently? We follow the first-order Reptile algorithm [34] to meta-update w_v in an online fashion, see Algorithm 1 (meta-training stage). Such meta-training avoids catastrophic forgetting with sequential task finetuning by knowledge distillation into the initialization of w_v . Another competing scheme is to pre-train the value function from all the task trajectories stored offline, which however requires constructing a large dataset. Fig. 3(b) compares performance for the two schemes.

3.2. Meta-testing with learned value function

Meta-testing involves optimizing the target metric for a new task \mathcal{T} , *i.e.*, finetuning a new pre-trained model θ or with a different adapter initialization ϕ_0 . The key here is to use our value function to augment a given surrogate loss with metric information during finetuning. Recall that value function can provide direct metric supervision or explicit gradients about metric. Hence it can be readily applied to off-the-shelf optimizers in a plug-and-play manner.

Inspired by the Guided ES approach [30], we combine our value function with SGD/Adam optimizers in a random search framework. We call the resulting method Metric Optimizer (MetricOpt). The high-level idea is to perform value function-based random search around surrogate loss gradients, in order to estimate a new descent direction favoring metric. Specifically, we keep track of a subspace defined by the past k loss gradients from SGD or Adam. Then the orthonormal basis $U \in \mathbb{R}^{d \times k}$ of the subspace is derived. We define the search covariance as:

$$\Sigma = \frac{1}{2d}I + \frac{1}{2k}UU^T. \quad (3)$$

By sampling P perturbations $\{\delta_i\}_{i=1}^P \sim \mathcal{N}(0, s^2\Sigma)$ with variance s^2 within the subspace, we estimate the descent direction u_t using Evolutionary Strategies (ES) informed by our learned value function $f_{\tilde{w}_v}$:

$$u_t = \frac{1}{s^2P} \sum_{i=1}^P \delta_i [f_{\tilde{w}_v}(\phi_t + \delta_i) - f_{\tilde{w}_v}(\phi_t - \delta_i)]. \quad (4)$$

Remarks The above meta-testing method provides one simple way to combine the loss and metric signals for finetuning. As evidenced in Fig. 3(c), the metric signal plays an important role. Without it, we see a large performance drop

Algorithm 1 Pseudocode of MetricOpt (with SGD/Adam)

```

Input: Task distribution  $p(\mathcal{T})$ , hyper-parameters  $\alpha, \eta, P, s^2$ 
Input: Number of finetuning steps  $T$  per task
Input: Number of metric evaluations  $K$  per task
// Meta-training value function
1: Initialize  $w_v$  of value function  $f_{w_v}$ 
2: for  $\mathcal{T} \sim p(\mathcal{T})$  do
3:   Collect sparse evaluation metrics  $\{M(\phi_k)\}_{k=1}^K$  over
   the tuning sequence  $\{\phi_t\}_{t=1}^T$  of SGD or Adam
4:   Obtain dense interpolation  $\{(\hat{M}_t, \hat{\sigma}_t)\}_{t=1}^T$ 
5:   Update  $\tilde{w}_v = w_v - \alpha \nabla_{w_v} L_t(w_v)$  via Eq. (2)
6:   Update  $w_v \leftarrow w_v + \eta(\tilde{w}_v - w_v)$  // Reptile update
7: end for
8: return  $\tilde{w}_v$ 
// Meta-testing with Guided ES and value function
9: Initialize  $\phi_0$  for new task  $\mathcal{T}$ 
10: for  $t = 0$  to  $T - 1$  do
   // Loss signal
11:   Update search covariance  $\Sigma$  via Eq. (3)
12:   Sample  $P$  perturbations  $\{\delta_i\}_{i=1}^P \sim \mathcal{N}(0, s^2\Sigma)$ 
   // Metric signal from value function
13:   Estimate descent direction  $u_t$  using  $f_{\tilde{w}_v}$ , see Eq. (4)
14:   Update  $\phi_{t+1} \leftarrow \phi_t - \alpha_t u_t$ 
15: end for
16: return  $\phi_T$ 

```

as expected. Whereas using metric alone (no auxiliary loss) barely hurts performance. The downside is the reduced efficiency – tuning with metric only will incur a larger cost for random search (over $5\times$ slower in the CIFAR10 case). Therefore we always treat our value function as a metric improver for surrogate losses, including those well-designed and learned losses. Also, the value function can be easily applied to other gradient-based optimization methods. Supplementary materials give one example of combining value function with a learned optimizer that achieves state-of-the-art performance.

3.3. Computational complexity

Algorithm 1 shows meta-training of our value function over three main steps: obtain the finetuning trajectory for task \mathcal{T} with task-dependent tuning steps T , metric interpolation for task \mathcal{T} , and meta-update of the value function. The task finetuning is very efficient since in practice we find it only involves hundreds of tuning steps over the compact adapter parameters ϕ . The overhead of metric interpolation is negligible. Overall, our value function can be trained for 0.5k to 2k meta-iterations (tasks), which requires less than half a GPU day for all experiments reported. For meta-testing, finetuning with Guided ES and value function is

Table 1. Mis-Classification Rate (MCR) on CIFAR-10. The average and standard deviation are from 10 runs. Note MetricOpt (SGD) is our default approach, while MetricOpt (learned) involves joint learning of value function and optimizer (see text for details).

Method	MCR (%) ↓
Cross-entropy loss	7.51
Large-margin softmax [29]	7.01
L2T-DLF [46]	6.95
Example weighting (MOEW) [48]	6.71±0.08
Adaptive Loss Alignment (ALA) [23]	6.79±0.07
Black-box differentiation [41]	6.84±0.05
Cross-entropy finetune + dense metric eval	7.19±0.16
MetricOpt (SGD) + cross-entropy	6.63±0.05
MetricOpt (learned) + cross-entropy	6.58±0.06
MetricOpt (SGD) + large-margin softmax [29]	6.56±0.06
MetricOpt (learned) + large-margin softmax [29]	6.47±0.07

only 22% to 34% slower than standard loss optimization. The overhead is reasonable given the few finetuning steps.

4. Experiments

We aim to answer the following questions in our experiments: 1) Does the value function consistently improve over surrogate losses on problems with only black-box access to the metric? 2) How does MetricOpt compare to state-of-the-art methods for metric optimization, and why the difference? 3) Can MetricOpt generalize to new tasks and model architectures? To answer these questions, we consider optimizing various evaluation metrics on a diverse set of problems, including (image) classification, image retrieval and object detection.

Hyper-parameters In all experiments we set $\gamma = 10$ to make the two terms in Eq. (2) roughly equally weighted. Results are insensitive to γ in a wide range. The number of finetuning steps T is task-dependent. For meta-testing with Guided ES [30], we use the same parameters—search within a subspace of past $k = 3$ loss gradients (Eq. (3)), and sample $P = 3$ perturbations with variance $s^2 = 0.01$ (Eq. (4)). In Algorithm 1, the inner learning rate $\alpha = 0.005$, and meta learning rate $\eta = 1$ with a linear decay to 0.

4.1. Image classification

We optimize the Mis-Classification Rate (MCR) and Area Under the Precision Recall Curve (AUCPR) on CIFAR-10 dataset [27]. ResNet-32 [21] and the network in [9] are used for MCR and AUCPR respectively, for fair comparisons with methods that use the same networks. We use the train/validation/test split of sizes 45k/5k/10k for both MetricOpt and those methods [1, 23, 31, 48] whose training also requires querying the validation statistics (loss or metric). MetricOpt is trained to optimize $d = 128$ conditioning parameters ϕ of the FiLM layers for considered ConvNets. During meta-training, we finetune $T = 200$

Table 2. Area Under the Precision Recall Curve (AUCPR) on CIFAR-10. The average and standard deviation are from 10 runs. MetricOpt (SGD) is our default approach, while MetricOpt (learned) involves joint learning of value function and optimizer (see text for details).

Method	AUCPR (%) ↑
Cross-entropy loss	84.6
Pairwise AUCROC loss [38]	94.2
AUCPR loss [9]	94.2
Example weighting (MOEW) [48]	94.6±0.11
Adaptive Loss Alignment (ALA) [23]	94.9±0.14
Black-box differentiation [41]	94.4±0.07
Cross-entropy finetune + dense metric eval	90.8±0.18
MetricOpt (SGD) + cross-entropy	95.7±0.12
MetricOpt (learned) + cross-entropy	96.2±0.09
MetricOpt (SGD) + AUCPR loss [9]	96.6±0.10
MetricOpt (learned) + AUCPR loss [9]	97.2±0.08

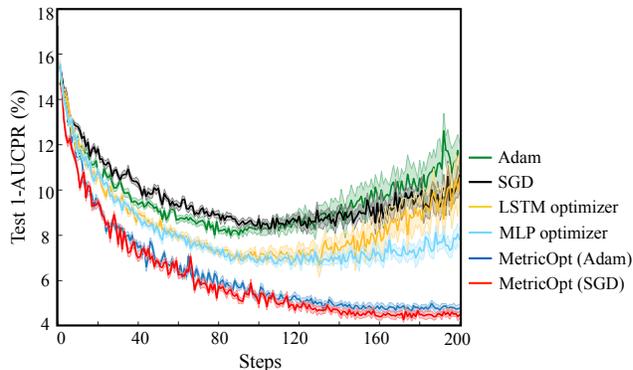


Figure 4. 10-run results of optimizing AUCPR on CIFAR-10. The metric undergoes a $1 - x$ conversion, so the lower the better. Our MetricOpt, when combined with either SGD or Adam, outperforms both the SGD variants and learned optimizers on surrogate cross-entropy loss.

steps and randomly collect $K = 5\%T$ evaluation metrics on validation data D_{val} along the finetuning trajectory. We report 10-run results of meta-testing on D_{test} .

First off, Fig. 4 uses the AUCPR case to compare our MetricOpt against existing popular optimizers. AUCPR is a non-differentiable metric and hard to optimize directly. Hence existing optimizers often choose the widely used cross-entropy loss as a surrogate, and we will discuss about advanced surrogates later. For best performance, we tune the learning rates of hand-crafted optimizers SGD and Adam [25], and tune the meta learning rates of the learned LSTM optimizer [1] and MLP optimizer [31] by a grid search over the range of $[10^{-4}, 1]$. We can see from the figure that existing loss optimizers suffer from the absence of metric supervision and do not optimize the metric well. By contrast, MetricOpt helps SGD/Adam to converge to better metrics by the guidance of value function.

In the following, **we choose MetricOpt (SGD) as our default approach**. Tables 1 and 2 compare MetricOpt with

Table 3. Binary classification on A9A and CoverType (Cov) datasets. We compare MetricOpt (with cross-entropy loss) against learned Surrogate Loss (SL) and various metric-specific losses. Specifically, we compare against CE (Cross-Entropy) and SL for MCR (Mis-Classification Rate), LO (Lovász-Softmax loss) and SL for JAC (Jaccard Index), and CS (Cost-Sensitive) and SL for F-measure.

Data	MCR ↓			JAC ↑			F-measure ↑		
	CE	SL [16]	MetricOpt	LO [3]	SL [16]	MetricOpt	CS [37]	SL [16]	MetricOpt
A9A	0.1520	0.1502	0.1369	0.8461	0.8488	0.8691	0.6823	0.6866	0.6934
Cov	0.2149	0.2198	0.2057	0.7406	0.7808	0.7758	0.7695	0.7898	0.8013

Table 4. Top-1 and Top-5 classification accuracies (%) on ImageNet. All methods use the same NASNet-A network.

Method	Top-1	Top-5
RMSProp + cross-entropy	73.5	91.5
PowerSign-cd [2] (CIFAR-10 transfer)	73.9	91.9
ALA [23] (CIFAR-10 transfer)	74.3	92.1
MetricOpt (transfer from ImageNet training with ResNet-18)	74.2	92.3
MetricOpt (transfer from CIFAR-10 training with NASNet-A)	74.4	92.2
MetricOpt (ImageNet training with NASNet-A)	74.9	92.7
MetricOpt + 2nd run (finetune θ and ϕ)	75.1	93.0

advanced loss functions and state-of-the-art black-box metric optimization methods. We have several observations:

- Surrogate loss functions (top cell) are suboptimal, although they take lots of time for manual design or online learning (*e.g.*, L2T-DLF).
- MOEW and ALA learn dynamic weighting schemes for data and loss respectively, both adapted to black-box metric observations. But MOEW is limited by using a predefined weighting scheme, and ALA still works in a relaxed surrogate space. In [41], gradient interpolation is performed via black-box differentiation. Our MetricOpt is more competitive by directly adapting the optimization process based on explicit metric modeling.
- With MetricOpt, our value function consistently improves over surrogate losses, including cross-entropy loss and advanced losses (large-margin softmax or AUCPR loss). The gains over advanced losses are smaller because they are often better aligned with metrics already. However, designing advanced losses requires tremendous human efforts. Hence in the following experiments, we mainly combine our value function with the default loss function for each problem, which waives the need for tedious loss engineering.
- A learned optimizer from the loss and value functions finds extra boosts. Supplementary materials introduce the learning mechanism and more results.
- We further compare with a simple finetuning baseline: it finetunes adapter parameters just like MetricOpt, but uses cross-entropy loss only and evaluates metrics at every step with the best picked. We see the resulting performance is far from that of MetricOpt. This suggests MetricOpt can find a truly different solution off loss-based trajectories. The solution has better test metric without requiring excessive metric evaluations along finetuning.

4.2. Binary classification on non-image data

Experiments are conducted on the A9A and CoverType datasets from the UCI Machine Learning Repository [8]. We randomly split data into 70%/10%/20% as the train/validation/test sets for both datasets. All methods use class-balanced mini-batch sampling, and train a Leaky ReLU-activation MLP with 100-30-10-1 neurons. Batch-Norm and dropout are applied at each layer. Our MetricOpt finetunes $d = 16$ dynamic biases added at the input layer for $T = 50$ steps. We still use cross-entropy loss as the baseline loss for model pre-training and MetricOpt-based finetuning, arriving at MetricOpt (SGD) + cross-entropy again.

Table 3 compares our MetricOpt with competing methods for 3 evaluation metrics—MCR, JAC and F-measure. Specifically, we compare with hand-designed metric-specific losses CE, LO [3] and CS [37]. We also compare with the learned Surrogate Loss (SL) [16], a representative of recent adaptively learned surrogate losses [16, 23, 24, 28]. The table generally confirms our advantages over these surrogate loss methods. This demonstrates the versatility of MetricOpt on black-box metrics.

4.3. Generalization test on ImageNet

This section compares different methods on the large-scale ImageNet dataset [7]. All methods use the same NASNet-A network architecture [50] unless otherwise noted. Evaluation metrics are the top 1 and top 5 classification accuracies. Table 4, top cell, outlines the full model training baselines, while the middle cell lists our MetricOpt-based finetuning variants. All MetricOpt variants finetune $d = 256$ conditioning parameters of FiLM layers for $T = 500$ steps. They aim at augmenting the cross-entropy loss to improve the MCR metric. Note for fair comparison, the number of model pre-training plus finetun-

ing steps of MetricOpt is the same as the total number of training steps for full training methods.

We can see from Table 4 that the RMSProp optimizer achieves suboptimal results since it merely optimizes the cross-entropy loss. PowerSign-cd [2] and ALA [23] transfer their update policy and loss controlling policy learned from CIFAR-10 to ImageNet, both showing strong results of out-of-distribution training. Our MetricOpt considers direct metric optimization during finetuning. In doing so, we obtain even better results than the full training baselines, although we only finetune a small number of parameters. We further study the transferability of MetricOpt by reusing our value function in two settings: learned with a smaller base network (ResNet-18) on the same ImageNet data, and learned with the same NASNet-A network on a smaller dataset (CIFAR-10). MetricOpt achieves better generalization results for a larger problem or model architecture, demonstrating the flexibility of the approach.

Given the good finetuning performance of MetricOpt, the bottom cell of Table 4 studies the following question: what if we also finetune the pre-trained network θ instead of keeping it fixed and only finetuning the modulating parameters ϕ ? To answer the question, we build on the finetuned ϕ in the first run, and continue to fine-tune θ (using cross-entropy loss) followed by a quick adjustment of ϕ via MetricOpt for a second run. This variant leads to marginal gains, but has much lower time efficiency than just tuning ϕ since finetuning the high-dimensional θ is very expensive. Moreover, we found diminishing returns when we alternate between finetuning θ and ϕ for a couple of more rounds.

4.4. Image retrieval

The image retrieval task is a perfect testbed for optimizing rank-based metrics like recall. We conduct experiments on the Stanford Online Products dataset [42] which contains 120,053 images of 22,634 classes. We follow the standard train/test splits and image data preparation procedure in [42]. A pretrained ResNet-50 [21] is used with a fully connected embedding layer. The embedding has a fixed size at 512 with L_2 normalization. For a fair comparison with related methods, we use the same mini-batch sampling strategy with batch size 128. The only difference in experimental setup is that we focus on finetuning from different pretrained models. For finetuning, instead of using SGD/Adam optimizers for a continued loss optimization, we use MetricOpt to finetune $d = 128$ conditioning parameters of FiLM layers for only $T = 200$ steps. Such finetuning relies on a value function optimized for the evaluation metric of average Recall@k.

For good performance, we use our value function to augment state-of-the-art surrogate losses – FastAP [4] and Recall loss [41]. They are derived by histogram binning approximation and ranking function interpolation respec-

Table 5. Recall(%)@k on Stanford Online Products dataset.

k	1	10	100	1000
FastAP [4]	76.4	89.1	95.4	98.5
Recall loss [41]	78.6	90.5	96.0	98.7
MetricOpt + FastAP [4]	79.6	91.2	96.3	98.9
MetricOpt + Recall loss [41]	80.4	91.8	96.5	99.0

Table 6. Object detection performance in terms of AP^{50} on Pascal VOC 07 test set. CE stands for Cross-Entropy loss.

Training	CE	AP loss [41]	MetricOpt + CE	MetricOpt + AP loss [41]
07	74.2	75.7	77.8	78.3
07+12	80.4	81.4	82.3	82.6

tively. Table 5 confirms our consistent improvements over both surrogate losses due to direct metric optimization.

4.5. Object detection

We adopt the commonly-used Faster R-CNN [40] framework for object detection. Experiments are conducted on the Pascal VOC dataset [11] where the VOC 07 or VOC 07+12 trainval sets are used for training, and VOC 07 test set is used for testing (4952 images with 20 object categories). We use the ResNet-50 [21] backbone and standard hyperparameters (*e.g.*, batch size, anchor setting) for Faster R-CNN. Our MetricOpt finetunes $d = 128$ conditioning parameters of FiLM layers for $T = 200$ steps. The finetuning goal is to improve the AP^{50} metric (average precision for boxes with at least 50% IoU against groundtruth).

Table 6 shows MetricOpt-based finetuning is able to obtain notable gains in AP^{50} over two baselines, using either the standard cross entropy loss or state-of-the-art AP loss [41]. This validates again the effectiveness of MetricOpt as a consistent metric improver but with low cost (hundreds of finetuning steps). Supplementary materials include an interesting observation that our metric-driven finetuning has non-trivial impact on the two-staged Faster R-CNN detector (*e.g.*, improves recall of the region proposal module). One future work is a detailed investigation of different impacts from varying metrics.

5. Conclusion

We present a new method to optimize black-box evaluation metrics using a differentiable value function. The value function is meta-learned to provide useful supervision or gradients about metric to augment any user-specified loss. We show it is easy to apply the value function to existing optimizers for gradient-based optimization, which consistently improves over the given loss. The resulting MetricOpt approach is efficient and achieves state-of-the-art performance in various metrics. MetricOpt also generalizes well across tasks and model architectures.

References

- [1] Marcin Andrychowicz, Misha Denil, Sergio Gómez, Matthew W Hoffman, David Pfau, Tom Schaul, Brendan Shillingford, and Nando de Freitas. Learning to learn by gradient descent by gradient descent. In *NeurIPS*, 2016.
- [2] Irwan Bello, Barret Zoph, Vijay Vasudevan, and Quoc V. Le. Neural optimizer search with reinforcement learning. In *ICML*, 2017.
- [3] Maxim Berman, Amal Rannen Triki, and Matthew B Blaschko. The lovász-softmax loss: A tractable surrogate for the optimization of the intersection-over-union measure in neural networks. In *CVPR*, 2018.
- [4] Fatih Cakir, Kun He, Xide Xia, Brian Kulis, and Stan Sclaroff. Deep metric learning to rank. In *CVPR*, 2019.
- [5] K. Chen, J. Li, W. Lin, J. See, J. Wang, L. Duan, Z. Chen, C. He, and J. Zou. Towards accurate one-stage object detection with AP-loss. In *CVPR*, 2019.
- [6] Harm de Vries, Florian Strub, Jeremie Mary, Hugo Larochelle, Olivier Pietquin, and Aaron C Courville. Modulating early visual processing by language. In *NeurIPS*, 2017.
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009.
- [8] Dheeru Dua and Casey Graff. UCI machine learning repository, 2017.
- [9] Elad Eban, Mariano Schain, Alan Mackey, Ariel Gordon, Ryan Rifkin, and Gal Elidan. Scalable learning of non-decomposable objectives. In *AISTATS*, 2017.
- [10] Martin Engilberge, Louis Chevallier, Patrick Perez, and Matthieu Cord. Sodeep: A sorting deep net to learn ranking loss surrogates. In *CVPR*, 2019.
- [11] Mark Everingham, Luc Van Gool, Christopher K. I. Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 88:303–308, 2010.
- [12] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017.
- [13] Sebastian Flennerhag, Andrei A. Rusu, Razvan Pascanu, Hujun Yin, and Raia Hadsell. Meta-learning with warped gradient descent. In *ICLR*, 2020.
- [14] Scott Fujimoto, Herke Hoof, and David Meger. Addressing function approximation error in actor-critic methods. In *ICML*, 2018.
- [15] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- [16] Josif Grabocka, Randolph Scholz, and Lars Schmidt-Thieme. Learning surrogate losses. *CoRR*, abs/1905.10108, 2019.
- [17] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *ICML*, 2018.
- [18] Tamir Hazan, Joseph Keshet, and David A. McAllester. Direct loss minimization for structured prediction. In *NeurIPS*, 2010.
- [19] K. He, F. Cakir, S. A. Bargal, and S. Sclaroff. Hashing as tie-aware learning to rank. In *CVPR*, 2018.
- [20] Kun He, Yan Lu, and Stan Sclaroff. Local descriptors optimized for average precision. In *CVPR*, 2018.
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [22] Paul Henderson and Vittorio Ferrari. End-to-end training of object class detectors for mean average precision. In *ACCV*, 2016.
- [23] Chen Huang, Shuangfei Zhai, Walter Talbott, Miguel Ángel Bautista, Shih-Yu Sun, Carlos Guestrin, and Joshua M. Susskind. Addressing the loss-metric mismatch with adaptive loss alignment. In *ICML*, 2019.
- [24] Qijia Jiang, Olaoluwa Adigun, Harikrishna Narasimhan, Mahdi Milani Fard, and Maya Gupta. Optimizing black-box metrics with adaptive surrogates. In *ICML*, 2020.
- [25] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- [26] Oluwasanmi O Koyejo, Nagarajan Natarajan, Pradeep K Ravikumar, and Inderjit S Dhillon. Consistent binary classification with generalized performance metrics. In *NeurIPS*, 2014.
- [27] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- [28] Lanlan Liu, Mingzhe Wang, and Jia Deng. A unified framework of surrogate loss by refactoring and interpolation. In *ECCV*, 2020.
- [29] Weiyang Liu, Yandong Wen, Zhiding Yu, and Meng Yang. Large-margin softmax loss for convolutional neural networks. In *ICML*, 2016.
- [30] Niru Maheswaranathan, Luke Metz, George Tucker, Dami Choi, and Jascha Sohl-Dickstein. Guided evolutionary strategies: augmenting random search with surrogate gradients. In *ICML*, 2019.
- [31] Luke Metz, Niru Maheswaranathan, Jeremy Nixon, Daniel Freeman, and Jascha Sohl-Dickstein. Understanding and correcting pathologies in the training of learned optimizers. In *ICML*, 2019.
- [32] P. Mohapatra, M. Rolinek, C. V. Jawahar, V. Kolmogorov, and M. P. Kumar. Efficient optimization for rank-based loss functions. In *CVPR*, 2018.
- [33] Harikrishna Narasimhan, Rohit Vaish, and Shivani Agarwal. On the statistical consistency of plug-in classifiers for non-decomposable performance measures. In *NeurIPS*, 2014.
- [34] Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *CoRR*, abs/1803.02999, 2018.
- [35] Yash Patel, Tomas Hodan, and Jiri Matas. Learning surrogates via deep embedding. In *ECCV*, 2020.
- [36] Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron C. Courville. Film: Visual reasoning with a general conditioning layer. In *AAAI*, 2018.
- [37] Shameem Puthiya Parambath, Nicolas Usunier, and Yves Grandvalet. Optimizing F-measures by cost-sensitive classification. In *NeurIPS*, 2014.

- [38] Alain Rakotomamonjy. Optimizing Area Under Roc Curve with SVMs. In *ROCAI*, 2004.
- [39] CE. Rasmussen and CKI. Williams. *Gaussian Processes for Machine Learning*. Adaptive Computation and Machine Learning. MIT Press, 2006.
- [40] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *NeurIPS*, 2015.
- [41] Michal Rolinek, Vit Musil, Anselm Paulus, Marin Vlastelica, Claudio Michaelis, and Georg Martius. Optimizing rank-based metrics with blackbox differentiation. In *CVPR*, 2020.
- [42] Hyun Oh Song, Yu Xiang, Stefanie Jegelka, and Silvio Savarese. Deep metric learning via lifted structured feature embedding. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4004–4012, 2016.
- [43] Yang Song, G. Alexander Schwing, S. Richard Zemel, and Raquel Urtasun. Training deep neural networks via direct loss minimization. In *ICML*, 2016.
- [44] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9:2579–2605, 2008.
- [45] Xun Wang, Xintong Han, Weilin Huang, Dengke Dong, and Matthew R Scott. Multi-similarity loss with general pair weighting for deep metric learning. In *CVPR*, 2019.
- [46] Lijun Wu, Fei Tian, Yingce Xia, Yang Fan, Tao Qin, Lai Jian-Huang, and Tie-Yan Liu. Learning to teach with dynamic loss functions. In *NeurIPS*, 2018.
- [47] Haowen Xu, Hao Zhang, Zhiting Hu, Xiaodan Liang, Ruslan Salakhutdinov, and Eric Xing. Autoloss: Learning discrete schedule for alternate optimization. In *ICLR*, 2019.
- [48] Sen Zhao, Mahdi Milani Fard, Harikrishna Narasimhan, and Maya Gupta. Metric-optimized example weights. In *ICML*, 2019.
- [49] Luisa M Zintgraf, Kyriacos Shiarlis, Vitaly Kurin, Katja Hofmann, and Shimon Whiteson. Fast context adaptation via meta-learning. In *ICML*, 2019.
- [50] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V. Le. Learning transferable architectures for scalable image recognition. In *CVPR*, June 2018.