





Learning Symbolic Model-Agnostic Loss Functions via Meta-Learning

Christian Raymond , *Graduate Student Member, IEEE*, Qi Chen , *Member, IEEE*,
Bing Xue , *Senior Member, IEEE*, and Mengjie Zhang , *Fellow, IEEE*

Abstract—In this paper, we develop upon the emerging topic of loss function learning, which aims to learn loss functions that significantly improve the performance of the models trained under them. Specifically, we propose a new meta-learning framework for learning model-agnostic loss functions via a hybrid neuro-symbolic search approach. The framework first uses evolution-based methods to search the space of primitive mathematical operations to find a set of symbolic loss functions. Second, the set of learned loss functions are subsequently parameterized and optimized via an end-to-end gradient-based training procedure. The versatility of the proposed framework is empirically validated on a diverse set of supervised learning tasks. Results show that the meta-learned loss functions discovered by the newly proposed method outperform both the cross-entropy loss and state-of-the-art loss function learning methods on a diverse range of neural network architectures and datasets. We make our code available at **retracted**.

Index Terms—Loss function learning, meta-learning, evolutionary computation, neuro-symbolic.

I. INTRODUCTION

THE field of learning-to-learn or *meta-learning* has been an area of increasing interest to the machine learning community in recent years [1], [2]. In contrast to conventional learning approaches, which learn from scratch using a static learning algorithm, meta-learning aims to provide an alternative paradigm whereby intelligent systems leverage their past experiences on related tasks to improve future learning performances [3]. This paradigm has provided an opportunity to utilize the shared structure between problems to tackle several traditionally very challenging deep learning problems in domains where both data and computational resources are limited [4], [5].

Manuscript received 14 September 2022; revised 28 June 2023; accepted 6 July 2023. Date of publication 11 July 2023; date of current version 3 October 2023. This work was supported in part by the Marsden Fund of New Zealand Government under Grants VUW1913, VUW1914, and VUW2016, in part by the Science for Technological Innovation Challenge (SfTI) fund under Grant E3603/2903, in part by MBIE Data Science SSIF Fund under Grant RTVU1914, in part by Huayin Medical under Grant E3791/4165, and in part by MBIE Endeavor Research Programme under Grants C11X2001 and UOCX2104. Recommended for acceptance by Prof. C. G. M. Snoek. (*Corresponding author: Qi Chen.*)

The authors are with Evolutionary Computation and Machine Learning Research Group, School of Engineering and Computer Science, Victoria University of Wellington, Wellington 6140, New Zealand (e-mail: christian.raymond@ecs.vuw.ac.nz; qi.chen@ecs.vuw.ac.nz; bing.xue@ecs.vuw.ac.nz; mengjie.zhang@ecs.vuw.ac.nz).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TPAMI.2023.3294394>, provided by the authors.

Digital Object Identifier 10.1109/TPAMI.2023.3294394

Many meta-learning approaches have been proposed for optimizing various components of deep neural networks. For example, early research on the topic explored using meta learning for generating learning rules [6], [7], [8], [9]. More recent research has extended itself to learning everything from activation functions [10], shared parameter/weight initializations [11], [12], [13], [14], and neural network architectures [15], [16], [17], [18] to whole learning algorithms from scratch [19], [20] and many more.

However, one component that has been overlooked until very recently is the loss function [21]. Typically in deep learning, neural networks are trained through the backpropagation of gradients originating from a handcrafted and manually selected loss function [22]. One significant drawback of this approach is that traditionally loss functions have been designed with task-generalty in mind, i.e. large classes of tasks in mind, but the system itself is only concerned with a single instantiation or small subset of that class. However, as shown by the *No Free Lunch Theorems* [23] no algorithm is able to do better than a random strategy in expectation — this suggests that specialization to a subclass of tasks is in fact the only way that performance can be improved in general.

Given this importance, the prototypical approach of selecting a loss function heuristically from a modest set of handcrafted loss functions should be reconsidered in favor of a more principled data-informed approach. The new and emerging subfield of loss function learning offers an alternative to this, which instead aims to leverage task-specific information and past experiences to infer and discover highly performant loss functions directly from the data. Initial approaches to loss function learning have shown promise in improving various aspects of deep neural networks training. However, they have several key issues and limitations which must be addressed for meta-learned loss functions to become a more desirable alternative than handcrafted loss functions.

In particular, many loss function learning approaches use a parametric loss representation such as a neural network [24] or Taylor polynomial [25], [26], [27], which is limited as it imposes unnecessary assumptions and constraints on the structure of the learned loss function. However, the current non-parametric alternative to this is to use a two-stage discovery and optimization process, which infers both the loss function structure and parameters simultaneously using genetic programming and covariance matrix adaptation [28], and quickly become intractable for large-scale optimization problems. Subsequent work [29], [30]

has attempted to address this issue; however, they crucially omit the optimization stage, which is known to produce sub-optimal performance.

This paper aims to resolve these issues through a newly proposed framework called Evolved Model-Agnostic Loss (EvoMAL), which meta-learns non-parametric symbolic loss functions via a hybrid neuro-symbolic search approach. The newly proposed framework aims to resolve the limitations of past approaches to loss function learning by combining genetic programming [31] with an efficient gradient-based local-search procedure [32], [33]. This unifies two previously divergent lines of research on loss function learning, which prior to this method, exclusively used either a gradient-based or an evolution-based approach.

This work innovates on the prior loss function learning approaches by introducing the first computationally tractable approach to optimizing symbolic loss functions. Consequently, improving the scalability and performance of symbolic loss function learning algorithm. Furthermore, unlike prior approaches, the proposed framework is both task and model-agnostic, as it can be applied to learning algorithm trained with a gradient descent style procedure and is compatible with different model architectures. This branch of general-purpose loss function learning algorithms provides a new powerful avenue for improving a neural network's performance, which has until recently not been explored.

The performance of EvoMAL is assessed on a diverse range of datasets and neural network architectures in the direct learning and transfer learning settings, where the empirical performance is compared with the ubiquitous cross-entropy loss and other state-of-the-art loss function learning methods. Finally, an analysis of the meta-learned loss functions produced by EvoMAL is presented, where several reoccurring trends are identified in both the shape and structure. Further analysis is also given to show why meta-learned loss functions are so performant through 1) examining the loss landscapes of the meta-learned loss functions and 2) investigating the relationship between the base learning rate and the meta-learned loss functions.

A. Contributions

The key contributions of this work are as follows:

- We propose a new task and model-agnostic search space and a corresponding search algorithm for meta-learning interpretable symbolic loss functions.
- We demonstrate a simple transition procedure for converting expression tree-based symbolic loss functions into gradient trainable loss networks.
- We utilize the new loss function representation to integrate the first computationally tractable approach to optimizing symbolic loss functions into the framework.
- We evaluate the proposed framework by performing the first-ever comparison of existing loss function learning techniques in both direct learning and transfer learning settings.

- We analyze the meta-learned loss functions to highlight key trends and explore why meta-learned loss functions are so performant.

II. BACKGROUND AND RELATED WORK

The goal of loss function learning in the meta-learning context is to learn a loss function \mathcal{M}_ϕ with parameters ϕ , at meta-training time over a distribution of tasks $p(\mathcal{T})$. A *task* is defined as a set of input-output pairs $\mathcal{T} = \{(x_1, y_1), \dots, (x_N, y_N)\}$, and multiple tasks compose a *meta-dataset* $\mathcal{D} = \{\mathcal{T}_1, \dots, \mathcal{T}_M\}$. Then, at meta-testing time the learned loss function \mathcal{M}_ϕ is used in place of a traditional loss function to train a base learner, e.g. a classifier or regressor, denoted by $f_\theta(x)$ with parameters θ on a new unseen task from $p(\mathcal{T})$. In this paper, we constrain the selection of base learners to models trainable via a gradient descent style procedures such that we can optimize weights θ as follows:

$$\theta_{new} = \theta - \alpha \nabla_{\theta} \mathcal{M}_\phi(y, f_\theta(x)) \quad (1)$$

Several approaches have recently been proposed to accomplish this task, and an observable trend is that most of these methods fall into one of the following two key categories.

A. Gradient-Based Approaches

Gradient-based approaches predominantly aim to learn a loss function \mathcal{M} through the use of a meta-level neural network to improve on various aspects of the training. For example, in [34], [35], differentiable surrogates of non-differentiable performance metrics are learned to reduce the misalignment problem between the performance metric and the loss function. Alternatively, in [24], [36], [37], [38], [39], [40], [41], [42], loss functions are learned to improve sample efficiency and asymptotic performance in supervised and reinforcement learning, while in [27], [43], [44], [45], they improved on the robustness of a model.

While the aforementioned approaches have achieved some success, they have notable limitations. The most salient limitation is that they *a priori* assume a parametric form for the loss functions. For example, in [24] and [46], it is assumed that the loss functions take on the parametric form of a two hidden layer feed-forward neural network with 50 nodes in each layer and ReLU activations. However, such an assumption imposes a bias on the search, often leading to an over parameterized and sub-optimal loss function. Another limitation is that these approaches often learn black-box (sub-symbolic) loss functions, which is not ideal, especially in the meta-learning context where *post hoc* analysis of the learned component is crucial, before transferring the learned loss function to new unseen problems at meta-testing time.

B. Evolution-Based Approaches

A promising alternative paradigm is to use evolution-based methods to learn \mathcal{M}_ϕ , favoring their inherent ability to avoid local optima via maintaining a population of solutions, their ease of parallelization of computation across multiple processors,

and their ability to optimize for non-differentiable functions directly. Examples of such work include [25] and [26], which both represent \mathcal{M}_ϕ as parameterized Taylor polynomials optimized with covariance matrix adaptation evolutionary strategies (CMA-ES). These approaches successfully derive interpretable loss functions, but similar to previously, they also assume the parametric form via the degree of the polynomial.

To resolve the issue of having to assume the parametric form of \mathcal{M}_ϕ , another avenue of research first presented in [28] investigated the use of genetic programming (GP) to learn the structure of \mathcal{M}_ϕ in a symbolic form before applying CMA-ES to optimize the parameterized loss. The proposed method was effective at learning performant loss functions and clearly demonstrated the importance of local-search. However, the method had intractable computational costs as using a population-based method (GP) with another population-based method (CMA-ES) resulted in a significant expansion in the number of evaluations at meta-training time, hence it needing to be run on a supercomputer in addition to using a truncated number of training steps.

Subsequent work in [29] and [30] reduced the computational cost of GP-based loss function learning approaches by proposing time saving mechanisms such as: rejection protocols, gradient-equivalence-checking, convergence property verification and model optimization simulation. These methods successfully reduced the wall-time of GP-based approaches; however, both papers omit the use of local-search strategies, which is known to cause sub-optimal performance when using GP [47], [48], [49]. Furthermore, neither method is task and model-agnostic, limiting their utility to a narrow set of domains and applications.

III. EVOLVED MODEL-AGNOSTIC LOSS

In this work, a novel hybrid neuro-symbolic search approach named *Evolved Model-Agnostic Loss (EvoMAL)* is proposed, which consolidates and extends past research on the topic of loss function learning. The proposed method learns performant and interpretable symbolic loss functions by inferring both the structure and the weights/coefficients directly from the data. The evolution-based technique GP [31] is used to solve the discrete problem of deriving the symbolic structure of the learned loss functions, while unrolled differentiation [32], [50], [51], [52], a gradient-based technique previously used in Meta-Learning via Learned Loss (ML³) [24], and sometimes referred to as Generalized Inner Loop Meta-Learning [33], is used to solve the continuous problem of optimizing their weights/coefficients.

A. Offline Loss Function Learning Setup

EvoMAL is a new approach to offline loss function learning, a meta-learning paradigm concerned with learning new and performant loss functions that can be used as a drop-in replacement for a prototypical handcrafted loss function such as the squared loss for regression or cross-entropy loss for classification [53], [54]. Offline loss function learning follows a conventional offline meta-learning setup, which partitions the learning into two sequential phases: meta-training and meta-testing.

1) *Meta-Training Phase*: The meta-training phase in EvoMAL is formulated as a bilevel optimization problem¹, where the goal of the outer optimization is to meta-learn a performant loss function \mathcal{M}_ϕ minimizing the average task loss \mathcal{L}_T (the meta-objective) across m related tasks, where \mathcal{L}_T is selected based on the desired task. The inner optimization uses \mathcal{M}_ϕ as the base loss function to train the base model parameters θ . Formally, the meta-training phase of EvoMAL is defined as follows:

$$\begin{aligned} \mathcal{M}_\phi^* &= \arg \min_{\mathcal{M}_\phi} \frac{1}{m} \sum_{i=1}^m \mathcal{L}_T(y_i, f_{\theta_i}(x_i)) \\ s.t. \quad \theta_i^*(\phi) &= \arg \min_{\theta_i} [\mathcal{M}_\phi(y, f_{\theta_i}(x))] \end{aligned} \quad (2)$$

2) *Meta-Testing Phase*: In the meta-testing phase, the best-performing loss function learned at meta-training time \mathcal{M}_ϕ^* is used directly to train and optimize the base model parameters θ .

$$\theta^* = \arg \min_{\theta} [\mathcal{M}_\phi^*(y, f_{\theta}(x))] \quad (3)$$

In contrast to online loss function learning [42], offline loss function learning does not require any alterations to the existing training pipelines to accommodate the meta-learned loss function, this makes the loss functions easily transferable and straightforward to use in existing code bases.

B. Learning Symbolic Loss Functions

To learn the symbolic structure of the loss functions in EvoMAL, we propose using GP, a powerful population-based technique that employs an evolutionary search to directly search the set of primitive mathematical operations [31]. In GP, solutions are composed of terminal and function nodes in a variable-length hierarchical expression tree-based structure. This symbolic structure is a natural and convenient way to represent loss functions, due to its high interpretability and trivial portability to new problems. Transferring a learned loss function from one problem to another requires very little effort, typically only a line or two of additional code. The task and model-agnostic loss functions produced by EvoMAL can be used directly as a drop-in replacement for handcrafted loss functions without requiring any new sophisticated meta-learning pipelines to train the loss on a per-task basis.

1) *Search Space Design*: In order to utilize GP, a search space containing promising loss functions must first be designed. When designing the desired search space, four key considerations are made — first, the search space should superset existing loss functions such as the squared error in regression and the cross entropy loss in classification. Second, the search space should be dense with promising new loss functions while also containing sufficiently simple loss functions such that cross task generalization can occur successfully at meta-testing time.

¹For simplicity, EvoMAL is presented as a bilevel optimization problem; however, since \mathcal{M} and ϕ are learned using separate processes, EvoMAL can be viewed as solving for a tri-level optimization problem with the outer optimization from Equation (2) being split into two distinct phases, 1) learning the symbolic loss functions structure, and 2) optimizing the loss functions weights/coefficients.

TABLE I
SET OF SEARCHABLE PRIMITIVE MATHEMATICAL OPERATIONS

Operator	Expression	Arity
Addition	$x_1 + x_2$	2
Subtraction	$x_1 - x_2$	2
Multiplication	$x_1 * x_2$	2
Division (AQ)	$x_1 / \sqrt{1 + x_2^2}$	2
Minimum	$\min(x_1, x_2)$	2
Maximum	$\max(x_1, x_2)$	2
Sign	$\text{sign}(x)$	1
Square	x^2	1
Absolute	$ x $	1
Logarithm	$\log(x + \epsilon)$	1
Square Root	$\sqrt{ x + \epsilon}$	1
Hyperbolic Tangent	$\tanh(x)$	1

Third, ensuring that the search space satisfies the key property of GP closure, i.e. loss functions will not cause *NaN*, *Inf*, undefined, or complex output. Finally, ensuring that the search space is both task and model-agnostic. With these considerations in mind, we present the function set in Table I. Regarding the terminal set, the loss function arguments $f_\theta(x)$ and y are used, as well as (ephemeral random) constants $+1$ and -1 . Unlike previously proposed search spaces for loss function learning, we have made several necessary amendments to ensure proper GP closure, and sufficient task and model-generality. The salient differences are as follows:

- Previous work in [28] uses unprotected operations: natural log ($\log(x)$), square root (\sqrt{x}), and division (x_1/x_2). Using these unprotected operations can result in imaginary or undefined output violating the GP closure property. To satisfy the closure property, we replace both the natural log and square root with protected alternatives, as well as replace the division operator with the analytical quotient (AQ) operator, a smooth and differentiable approximation to the division operator [55].
- The proposed search space for loss functions is both task and model-agnostic in contrast to [29] and [30], which use multiple aggregation-based and element-wise operations in the function set. These operations are suitable for object detection (the respective paper’s target domain) but are not compatible when applied to other tasks such as tabulated and natural language processing problems.

2) *Search Algorithm Design*: The symbolic search algorithm used in EvoMAL follows a prototypical implementation of GP; as shown in Fig. 1. First, initialization is performed via randomly generating a population of 25 expression trees, where the inner nodes are selected from the function set and the leaf nodes from the terminal set. Subsequently, the main loop begins by performing the loss function optimization and evaluation stages to determine each loss function’s respective fitness, discussed in Section III-C and D, respectively. Following this, a new offspring population of equivalent size is constructed via crossover, mutation, and elitism. For crossover, two loss functions are

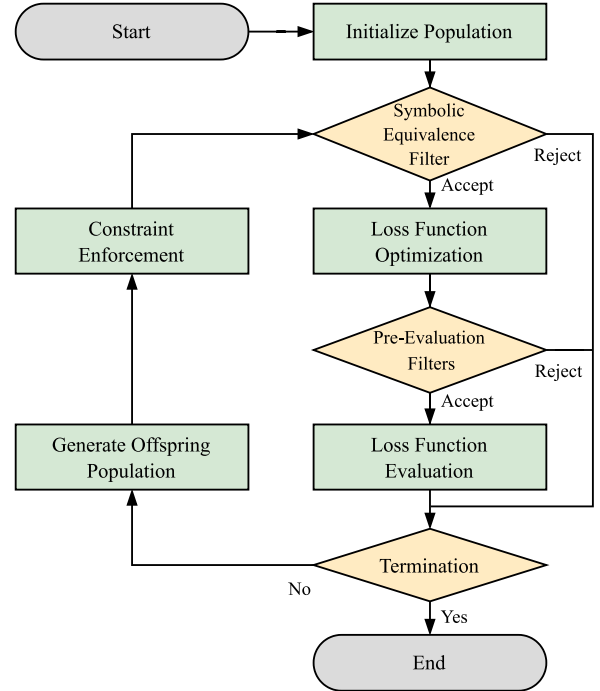


Fig. 1. Overview of the EvoMAL framework, showing the pipeline for learning loss functions at meta-training time.

selected via tournament selection and combined using a one-point crossover with a crossover rate of 70%. For mutation, a loss function is selected, and a uniform mutation is applied with a mutation rate of 25%. Finally, to ensure performance does not degrade, elitism is used to retain top-performing loss functions with an elitism rate of 5%. The main loop is iteratively repeated 50 times, and the loss function with the best fitness is selected as the final learned loss function. Note, to reduce the computational overhead of the meta-learning process, a number of time-saving measures, which we further refer to as filters, have been incorporated into the EvoMAL algorithm. This is discussed in detail in Section III-E.

3) *Constraint Enforcement*: When using GP, the evolved expressions often violate the constraint that a loss function must have as arguments $f_\theta(x)$ and y . Our preliminary investigation found that often over 50% of the loss functions in the first few generations violated this constraint. Thus far, existing methods for handling this issue have been inadequate; for example, in [28], violating loss functions were assigned the worst-case fitness, such that selection pressure would phase out those loss functions from the population. Unfortunately, this approach degrades search efficiency, as a subset of the population is persistently searching infeasible regions of the search space. To resolve this, we propose a simple but effective corrections strategy to violating loss functions, which randomly selects a terminal node and replaces it with a random binary node, with arguments $f_\theta(x)$ and y in no predetermined order.

An additional optional constraint enforceable in the EvoMAL algorithm is that the learned loss function can always return a non-negative output $\mathcal{M} : \mathbb{R}^2 \rightarrow \mathbb{R}_0^+$. This is achieved via the loss

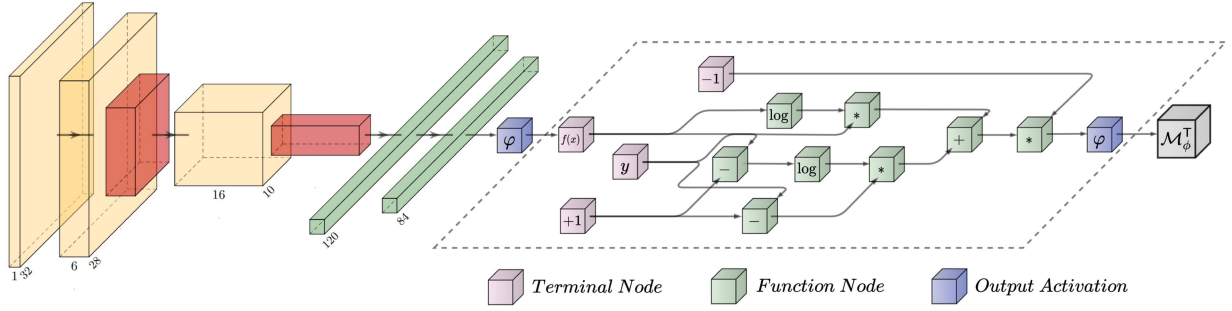


Fig. 3. Overview of the EvoMAL algorithm at *meta-testing* time, where the base network $f_{\theta}(x)$, shown (left) as the popular LeNet-5 architecture, is trained using the meta-loss network $\mathcal{M}_{\phi}^{\top}$ (right) found at *meta-training* time as the loss function.

Algorithm 2: Loss Function Evaluation.

In: $\mathcal{M}_{\phi}^{\top} \leftarrow$ Loss function learned by EvoMAL
 $S_{testing} \leftarrow$ Number of base testing gradient steps

```

for  $i \in \{0, \dots, |\mathcal{D}|\}$  do
   $\theta_i \leftarrow$  Initialize parameters of base learner  $f_{\theta_i}$ 
   $X_i, y_i \leftarrow$  Sample task  $\mathcal{T}_i \sim p(\mathcal{T})$ 
  for  $j \in \{0, \dots, S_{testing}\}$  do
     $\mathcal{L}_{learned} \leftarrow \mathcal{M}_{\phi}^{\top}(y_i, f_{\theta_i}(X_i))$ 
     $\theta_i \leftarrow \theta_i - \alpha \nabla_{\theta_i} \mathcal{L}_{learned}$ 
 $\mathcal{F} \leftarrow \frac{1}{|\mathcal{D}|} \sum_i \mathcal{L}_{\mathcal{P}}(y_i, f_{\theta_i}(X_i))$ 

```

η is the meta learning rate.

$$\begin{aligned} \phi_{new} &= \phi - \eta \nabla_{\phi} \mathcal{L}_{\mathcal{T}}(y, f_{\theta_{new}}(x)) \\ &= \phi - \eta \nabla_{\phi} \mathbb{E}_{X,y} [\mathcal{L}_{\mathcal{T}}(y, f_{\theta_{new}}(x))] \end{aligned} \quad (8)$$

where the gradient computation can be decomposed by applying the chain rule as shown in Equation (9) where the gradient with respect to the meta-loss network weights ϕ requires the new model parameters θ_{new} .

$$\phi_{new} = \phi - \eta \nabla_{\phi} \mathcal{L}_{\mathcal{T}} \nabla_{\theta_{new}} f_{\theta_{new}} \nabla_{\theta} \theta_{new} \quad (9)$$

This process is repeated for a predetermined number of meta gradient steps $S_{meta} = 250$, which was selected via cross-validation. Following each meta gradient step, the base learner weights θ is reset. Note that Equations (5)–(6) and (8)–(9) can alternatively be performed via automatic differentiation.

D. Loss Function Evaluation

To derive the fitness \mathcal{F} of $\mathcal{M}_{\phi}^{\top}$, a conventional training procedure is used as summarized in Algorithm 2, where $\mathcal{M}_{\phi}^{\top}$ is used in place of a traditional loss function to train $f_{\theta}(x)$ over a predetermined number of base gradient steps $S_{testing}$. This training process is identical to training at meta-testing time as shown in Fig. 3. The final inference performance of $\mathcal{M}_{\phi}^{\top}$ is assigned to \mathcal{F} , where any differentiable or non-differentiable performance metric $\mathcal{L}_{\mathcal{P}}$ can be used. For our experiments we use the error rate to compute \mathcal{F} .

E. Time Saving Measures (Filters)

Optimizing and evaluating a large number of candidate loss functions can become prohibitively expensive. Fortunately, in the case of loss function learning, a number of techniques can be employed to reduce significantly the computational overhead of the otherwise very costly meta-learning process.

1) *Symbolic Equivalence Filter*: For the GP-based symbolic search, a loss archival strategy based on a key-value pair structure with $\Theta(1)$ lookup is used to ensure that symbolically equivalent loss functions are not reevaluated. Two expression trees are said to be symbolically equivalent when they contain identical operations (nodes) in an identical configuration [61], [62], [63]. Loss functions that are identified by the symbolic equivalence filter skip both the loss function optimization and evaluation stage and are placed directly in the offspring population with their fitness cached.

2) *Pre-Evaluation Filter - Poor Training Dynamics*: Following loss function optimization, the candidate solution's fitnesses are evaluated. This costly fitness evaluation can be obviated in many cases since many of the loss functions found, especially in the early generations, are non-convergent and produce poor training dynamics. We use the loss rejection protocol employed in [30] as a filter to identify candidate loss functions that should skip evaluation and be assigned the worst-case fitness automatically.

The loss rejection protocol takes a batch of B randomly sampled instances from \mathcal{D}_{Train} and using an untrained network $f_{\theta_0}(x)$ produces a set of predictions and their corresponding true target values $\{(\hat{y}_b, y_b)\}_{b=1}^B$. As minimizing the proper loss function $\mathcal{M}_{\phi}^{\top}$ should correspond closely with optimizing the performance metric $\mathcal{L}_{\mathcal{P}}$, a correlation g between $\mathcal{M}_{\phi}^{\top}$ and $\mathcal{L}_{\mathcal{P}}$ can be calculated.

$$g(\mathcal{M}_{\phi}^{\top}) = \sum_{b=1}^B \left[\mathcal{L}_{\mathcal{P}}(\hat{y}_b, y_b) - \mathcal{L}_{\mathcal{P}}(\hat{y}_b^*(\mathcal{M}_{\phi}^{\top}), y_b) \right], \quad (10)$$

where \hat{y}_b^* is the network predictions optimized with the candidate loss function $\mathcal{M}_{\phi}^{\top}$. Importantly, optimization is performed directly to \hat{y}_b^* , as opposed to the network parameters θ , thus omitting any base network computation (i.e. no training of the

base network).

$$\hat{y}_b^*(\mathcal{M}_\phi^\top) = \arg \min_{\hat{y}_b} \mathcal{M}_\phi^\top(\hat{y}_b^*, y_b) \quad (11)$$

A large positive correlation indicates that minimizing \mathcal{M}_ϕ^\top corresponds to minimizing the given performance metric $\mathcal{L}_\mathcal{P}$ (assuming both \mathcal{M}_ϕ^\top and $\mathcal{L}_\mathcal{P}$ are for minimization). In contrast to this, if $g \leq 0$, then \mathcal{M}_ϕ^\top is regarded as being unpromising and should be assigned the worst-case fitness and not evaluated. The underlying assumption here is that if a loss function cannot directly optimize the labels, it is unlikely to be able to optimize the labels through the model weights θ successfully.

3) *Pre-Evaluation Filter - Gradient Equivalence*: Many of the candidate loss functions found in the later generations, as convergence is approached, have near-identical gradient behavior (i.e. functionally equivalence). To address this, the gradient equivalence checking strategy from [30] is employed as another filter to identify loss functions that have near-identical behavior to those seen previously. Using the prediction from previously, the gradient norms are computed.

$$\{\|\nabla_{\hat{y}_b} \mathcal{M}_\phi^\top\|_2\}_{b=1}^B \quad (12)$$

If, for all of the B samples, two-loss functions have the same gradient norms within two significant digits, they are considered functionally equivalent, and their fitness is cached.

4) *Partial Training Sessions*: For the remaining loss functions whose fitness evaluation cannot be fully obviated, we compute the fitness \mathcal{F} using a truncated number of gradient steps $\mathcal{S}_{testing} = 500$. As noted in [64], [65], performance at the beginning of training is correlated with the performance at the end of training; consequently, we can obtain an estimate of what \mathcal{F} would be by performing a partial training session of the base model. Preliminary experiments with EvoMAL showed minimal short-horizon bias [66], and the ablation study found in [25] indicated that 500 gradient steps during loss function evaluation is a good trade-off between final base-inference performance and meta-training time. In addition to significantly reducing the run-time of EvoMAL, reducing the value of $\mathcal{S}_{testing}$ has the effect of implicitly optimizing for the base-training convergence and sample-efficiency, as mentioned in [3].

IV. EXPERIMENTAL SETUP

In this section, the performance of EvoMAL is evaluated. A wide range of experiments are conducted across four datasets and numerous popular network architectures, with the performance contrasted against a representative set of benchmark methods implemented in DEAP [67], PyTorch [68] and Higher [33].

A. Benchmark Methods

The selection of benchmark methods is intended to showcase the performance of the newly proposed algorithm against the current state-of-the-art. Additionally, the selected methods enable direct comparison between EvoMAL and its derivative methods, which aids in validating the effectiveness of hybridizing the approaches into one unified framework.

- *Baseline* – Directly using $\mathcal{L}_\mathcal{T}$ as the loss function, i.e. using the squared error loss (regression) or cross-entropy loss (classification) and a prototypical training loop (i.e. no meta-learning).
- *ML³ Supervised* – Gradient-based method proposed in [24], which uses a parametric loss function defined by a two hidden layer feed-forward network trained with generalized inner loop meta-learning, i.e. the method shown in Section III-C.
- *TaylorGLO* — Evolution-based method proposed by [25], which uses a third-order Taylor-polynomial representation for the meta-learned loss functions, optimized via covariance matrix adaptation evolution strategy.
- *GP-LFL* – A proxy method used to aggregate previous GP-based approaches for loss function learning *without* any local-search mechanisms, using an identical setup to EvoMAL excluding Section III-C.

Where possible, hyper-parameter selection has been standardized across the benchmark methods to allow for a fair comparison. For example, in TaylorGLO, GP-LFL, and EvoMAL we use an identical population size = 25 and number of generations = 50. For unique hyper-parameters, the suggested values from the respective publications are utilized.

B. Benchmark Problems

Regarding the problem domains, seven datasets have been selected. Three tabulated regression tasks are initially used: Diabetes, Boston Housing, and California Housing, all taken from the UCI’s dataset repository [69]. Following this, analogous to the prior literature [24], [25], [28], both MNIST [70] and CIFAR-10 [71] are employed to evaluate the benchmark methods. Finally, experiments are conducted on the more challenging but related domains of SVHN [72] and CIFAR-100 [71], respectively.

For the three tabulated regression tasks, the datasets are partitioned 60:20:20 for training, validation, and testing. Furthermore, to improve the training dynamics, both the features and labels are normalized. For the remaining datasets, the original training-testing partitioning is used, with 10% of the training instances allocated for validation. In addition, data augmentation techniques consisting of normalization, random horizontal flips, and cropping are applied to the training data during meta and base training.

Regarding the base models, a diverse set of neural network architectures are utilized to evaluate the selected benchmark methods. For Diabetes, Boston Housing, and California Housing, a simple Multi-Layer Perceptron (MLP) taken from [73] with 1000 hidden nodes and ReLU activations are employed. For MNIST, Logistic Regression [74], MLP, and the well-known LeNet-5 architecture [70] are used. While on CIFAR-10 AlexNet [75], VGG-16 [76], AllCNN-C [77], ResNet-18 [78], Preactivation ResNet-101 [79], WideResNet 28-10 [80] and SqueezeNet [81] are used. For the remaining datasets, WideResNet 28-10 is again used, as well as PyramidNet [82] on CIFAR-100. All models are trained using stochastic gradient descent

TABLE II
CLASSIFICATION RESULTS REPORTING THE MEAN \pm STANDARD DEVIATION FINAL INFERENCE ERROR RATE ACROSS 5 INDEPENDENT EXECUTIONS OF EACH ALGORITHM ON EACH TASK + MODEL PAIR

Task and Model	Baseline	ML ³	TaylorGLO	GP-LFL	EvoMAL (Ours)
MNIST					
Logistic ¹	0.0787 \pm 0.0009	0.0768 \pm 0.0061	0.0725\pm0.0013	0.0781 \pm 0.0052	0.0721\pm0.0017
MLP ²	0.0247 \pm 0.0005	0.0201 \pm 0.0081	0.0151\pm0.0013	0.0156\pm0.0012	0.0152\pm0.0012
LeNet-5 ³	0.0203 \pm 0.0025	0.0135 \pm 0.0039	0.0137 \pm 0.0038	0.0115 \pm 0.0015	0.0100\pm0.0010
CIFAR-10					
AlexNet ⁴	0.1544 \pm 0.0012	0.1450\pm0.0028	0.1499 \pm 0.0075	0.1506 \pm 0.0047	0.1437\pm0.0033
VGG-16 ⁵	0.0771 \pm 0.0023	0.0700\pm0.0006	0.0700\pm0.0022	0.0686\pm0.0014	0.0687\pm0.0016
AIICNN-C ⁶	0.0761 \pm 0.0015	0.0712 \pm 0.0043	0.0735 \pm 0.0030	0.0701\pm0.0022	0.0697\pm0.0010
ResNet-18 ⁷	0.0658 \pm 0.0019	0.0584 \pm 0.0022	0.0546 \pm 0.0033	0.0818 \pm 0.0391	0.0528\pm0.0015
PreResNet ⁸	0.0661\pm0.0015	0.0660\pm0.0016	0.0660\pm0.0027	0.0658\pm0.0023	0.0655\pm0.0018
WideResNet ⁹	0.0548 \pm 0.0016	0.0549 \pm 0.0040	0.0493\pm0.0023	0.0489\pm0.0014	0.0484\pm0.0018
SqueezeNet ¹⁰	0.0838 \pm 0.0013	0.0800\pm0.0012	0.0800\pm0.0025	0.0810\pm0.0016	0.0796\pm0.0017
CIFAR-100					
WideResNet ⁹	0.2293\pm0.0017	0.2299\pm0.0027	0.2347 \pm 0.0077	0.3382 \pm 0.1406	0.2276\pm0.0033
PyramidNet ¹¹	0.2527\pm0.0028	0.2792 \pm 0.0226	0.3064 \pm 0.0549	0.2747 \pm 0.0087	0.2664 \pm 0.0063
SVHN					
WideResNet ⁹	0.0340\pm0.0005	0.0335\pm0.0003	0.0343 \pm 0.0016	0.0340\pm0.0015	0.0329\pm0.0013

Network architecture references: ¹ McCullagh et al. (2019) ² Baydin et al. (2018) ³ LeCun et al. (1998) ⁴ Krizhevsky et al. (2012) ⁵ Simonyan and Zisserman (2015) ⁶ Springenberg et al. (2015) ⁷ He et al. (2015) ⁸ He et al. (2016) ⁹ Zagoruyko and Komodakis (2016) ¹⁰ N. Iandola et al. (2016) ¹¹ Han et al. (2017)

The best performance and those within its standard deviation are bolded. Loss functions are directly meta-learned and applied to the same respective task.

(SGD) with momentum. The model hyper-parameters are selected using their respective values from the literature in an identical setup to [24], [25].

Finally, due to the stochastic nature of the benchmark methods, we perform five independent executions of each method on each dataset + model pair. Furthermore, we control for the base initializations such that each method gets identical initial conditions across the same random seed; thus, any difference in variance between the methods can be attributed to the respective algorithms.

V. RESULTS AND ANALYSIS

The results and analysis are approached from three distinct angles. First, the experimental results reporting the final inference meta-testing performance when using meta-learned loss functions for base-training are presented in Section V-A. Following this, the performance of the loss function learning algorithms themselves at meta-training time is compared in Section V-B. The focus is turned towards an analysis of the meta-learned loss functions themselves to highlight some of the loss functions developed by EvoMAL on both classification and regression tasks in Section V-C. Finally, an analysis is given in Section V-D and E, which explore two of the central hypotheses for why meta-learned loss functions are so performant compared to their handcrafted counterparts.

A. Meta-Testing Performance

A summary of the final inference testing results reporting the average error rate (classification) or mean squared error (regression) across the seven tested datasets is shown in

Tables II and III respectively, where the same dataset and model pair are used for both meta-training and meta-testing. The results show that meta-learned loss functions consistently produce superior performance compared to the baseline handcrafted squared error loss and cross-entropy loss. Large performance gains are made on the California Housing, MNIST, and CIFAR-10 datasets, while more modest gains are observed on Diabetes, WideResNet CIFAR-100 and SVHN, and worse performance on Boston Housing MLP and CIFAR-100 PyramidNet, a similar finding to that found in [25].

Contrasting the performance of EvoMAL to the benchmark loss function learning methods, it is shown that EvoMAL consistently meta-learns more performant loss functions, with better performance on all task + model pairs except for on CIFAR-10 VGG-16, where performance is comparable to the next best method. Furthermore, compared to its derivative methods ML³ and GP-LFL, EvoMAL successfully meta-learns loss functions on more complex tasks, i.e. CIFAR-100 and SVHN, whereas the other techniques often struggle to improve upon the baseline. These results empirically confirm the benefits of unifying existing approaches to loss function learning into one unified framework. Furthermore, the results clearly show the necessity for integrating local-search techniques in GP-based loss function learning.

Compared to prior research on loss function learning, our method exhibits a relatively smaller improvement when comparing the use of meta-learned loss functions to using the cross-entropy loss. For example, prior research has reported increasing the accuracy of a classification model by up to 5% in some cases when using meta-learned loss functions compared to the cross-entropy loss. However, with heavily tuned baselines,

TABLE III
REGRESSION RESULTS REPORTING THE MEAN \pm STANDARD DEVIATION FINAL INFERENCE TESTING MEAN SQUARED ERROR ACROSS 5 EXECUTIONS OF EACH ALGORITHM ON EACH TASK + MODEL PAIR

Task and Model	Baseline	ML ³	TaylorGLO	GP-LFL	EvoMAL (Ours)
Diabetes					
MLP ¹	0.3829 \pm 0.0065	0.4278 \pm 0.1080	0.3620\pm0.0096	0.3746\pm0.0577	0.3573\pm0.0198
Boston					
MLP ¹	0.1304\pm0.0057	0.2474 \pm 0.0526	0.1349\pm0.0074	0.1341\pm0.0082	0.1314\pm0.0131
California					
MLP ¹	0.2326 \pm 0.0019	0.2438 \pm 0.0189	0.1794 \pm 0.0156	0.2071 \pm 0.0609	0.1723\pm0.0048

Network architecture references: ¹ Baydin et al. (2018)

The best performance and those within its standard deviation are bolded. Loss functions are directly meta-learned and applied to the same respective task.

TABLE IV
LOSS FUNCTION TRANSFER RESULTS REPORTING THE MEAN \pm STANDARD DEVIATION FINAL INFERENCE TESTING ERROR RATE ACROSS 5 INDEPENDENT EXECUTIONS OF EACH ALGORITHM ON EACH TASK + MODEL PAIR

Task and Model	Baseline	ML ³	TaylorGLO	GP-LFL	EvoMAL (Ours)
CIFAR-100					
AlexNet ¹	0.5262\pm0.0094	0.7735 \pm 0.0295	0.5543 \pm 0.0138	0.5329\pm0.0037	0.5324\pm0.0031
VGG-16 ²	0.3025\pm0.0022	0.3171 \pm 0.0019	0.3155 \pm 0.0021	0.3171 \pm 0.0041	0.3115 \pm 0.0038
AllCNN-C ³	0.2830\pm0.0021	0.2817\pm0.0032	0.4191 \pm 0.0058	0.2849 \pm 0.0012	0.2807\pm0.0028
ResNet-18 ⁴	0.2474 \pm 0.0018	0.6000 \pm 0.0173	0.2436 \pm 0.0032	0.2373 \pm 0.0013	0.2326\pm0.0014
PreResNet ⁵	0.2908 \pm 0.0065	0.2838\pm0.0019	0.2993 \pm 0.0030	0.2839\pm0.0025	0.2899 \pm 0.0024
WideResNet ⁶	0.2293 \pm 0.0017	0.2448 \pm 0.0063	0.2285 \pm 0.0031	0.2276 \pm 0.0028	0.2238\pm0.0017
SqueezeNet ⁷	0.3178\pm0.0015	0.3402 \pm 0.0057	0.3178\pm0.0012	0.3343 \pm 0.0054	0.3166\pm0.0026

Network architecture references: ¹ Krizhevsky et al. (2012) ² Simonyan and Zisserman (2015) ³ Springenberg et al. (2015) ⁴ He et al. (2015) ⁵ He et al. (2016) ⁶ Zagoruyko and Komodakis (2016) ⁷ N. Iandola et al. (2016)

The best performance and those within its standard deviation are bolded. Loss functions are meta-learned on CIFAR-10 with the respective model, and then transferred to CIFAR-100 using that same model.

optimizing for both α and $\mathcal{S}_{testing}$, such performance gains were very hard to obtain. This suggests that a proportion of the performance gains reported previously by loss function learning methods likely come from an implicit tuning effect on the training dynamics as opposed to a direct effect from using a different loss function. Implicit tuning is not a drawback of loss function learning as a paradigm; however, it is essential to disentangle the effects. In Section V-C, further experiments are given to isolate the effects.

Finally, although the performance gains are at times modest, it is impressive that performance gains can be made by a straightforward change to the selection of the loss function. Many previous developments in neural networks have added many millions of parameters or incorporated complex training routines to achieve similar improvements in performance [83], [84]. Furthermore, our results show that performance gains can be made upon well-tuned models which already employ a multitude of regularization techniques. This suggests that meta-learned loss functions can learn a distinct form of regularization that existing techniques such as data augmentation, dropout, weight decay, batch norm, skip layers, hyper-parameter optimization, etc., fail to capture.

1) *Loss Function Transfer*: To further validate the performance of EvoMAL, the much more challenging task of loss function transfer is assessed, where the meta-learned loss

functions learned in one source domain are transferred to a new but related target domain. In our experiments, the meta-learned loss functions are taken directly from CIFAR-10 (in the previous section) and then transferred to CIFAR-100 with no further computational overhead, using the same model as the source. To ensure a fair comparison is made to the baseline, only the best-performing loss function found across the 5 random seeds from each method on each task + model pair are used. A summary of the final inference testing error rates when performing loss function transfer is given in Table IV.

The results show that even when using a single-task meta-learning setup where cross-task generalization is not explicitly optimized for, meta-learned loss functions can still be transferred with some success. In regards to meta-generalization, it is observed that EvoMAL and GP-LFL transfer their relative performance the most consistently to new tasks. In contrast, ML³, which uses a neural network-based representation, often fails to generalize to the new tasks, a finding similar to [3] which found that symbolic representations often generalize better than sub-symbolic representations. Another notable result is that on both AlexNet and VGG-16 in the direct meta-learning setup, large gains in performance compared to the baseline are observed, as shown in Table II. Conversely, when the learned loss functions are transferred to CIFAR-100, all the loss function learning methods perform worse than the baseline as shown in

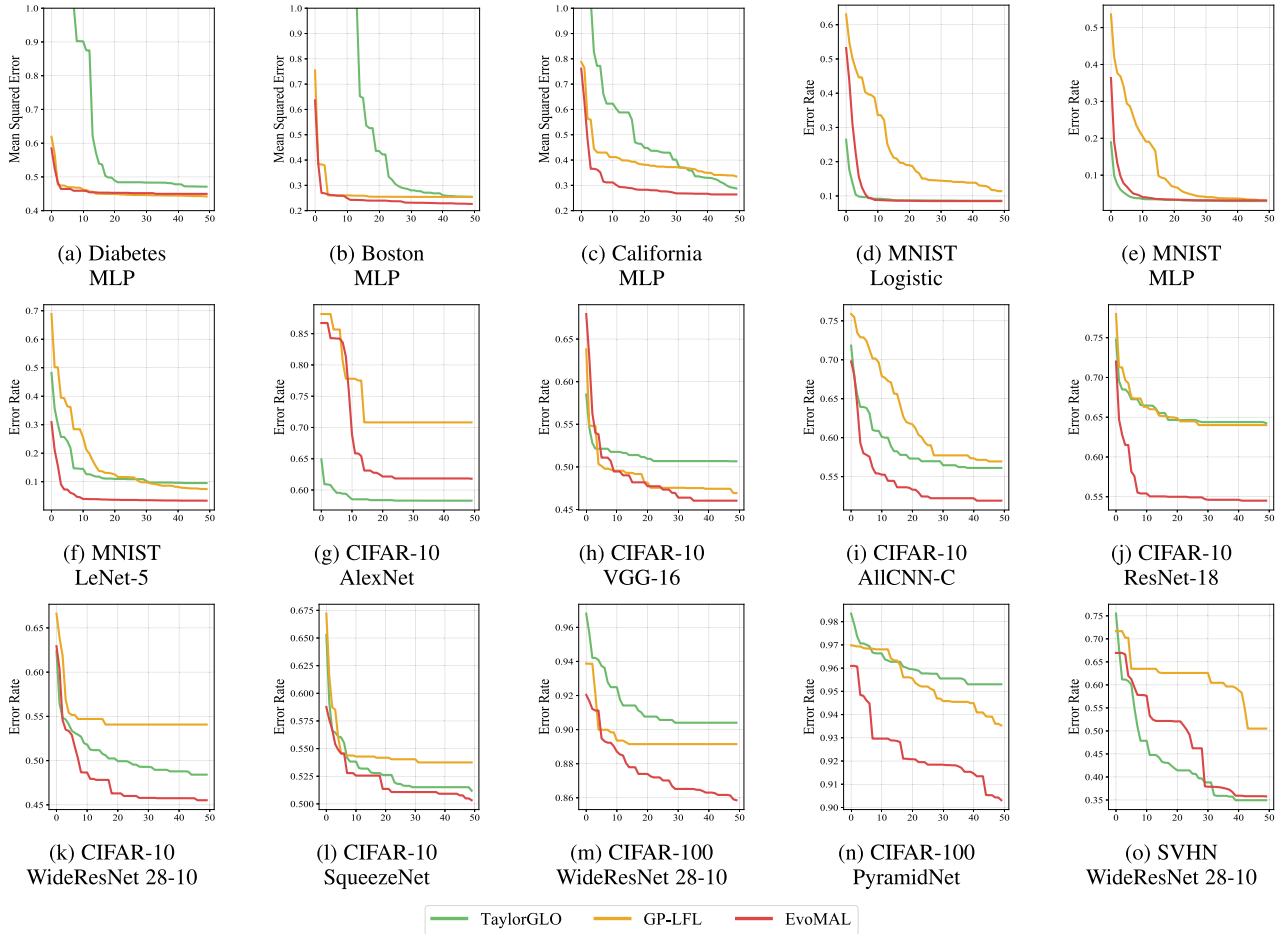


Fig. 4. Mean *meta-training* learning curves across 5 independent executions of each algorithm, showing the fitness score (y -axis) against generations (x -axis), where a generation is equivalent to 25 evaluations. Best viewed in color.

Table IV. These results suggest that the learned loss functions have likely been meta-overfitted to the source task and that meta-regularization is an important aspect to consider for loss function transfer.

B. Meta-Training Performance

The meta-training learning curves are given in Fig. 4, where the search performance of EvoMAL is compared to TaylorGLO and GP-LFL at each iteration and the performance is quantified by the fitness function using partial training sessions. Based on the results, it is very evident that adding local-search mechanisms into the EvoMAL framework dramatically increases the search effectiveness of GP-based loss function learning. EvoMAL is observed to consistently attain better performance than GP-LFL in a significantly shorter number of iterations. In almost all tasks, EvoMAL is shown to find better performing loss functions in the first 10 generations compared to those found by GP-LFL after 50.

Contrasting EvoMAL to TaylorGLO, it is generally shown again that for most tasks, EvoMAL produces better solutions in a smaller number of iterations. Furthermore, performance does

not appear to prematurely converge on the more challenging tasks of CIFAR-100 and SVHN compared to TaylorGLO and GP-LFL. Interestingly, on both CIFAR-10 AlexNet and SVHN WideResNet, TaylorGLO is able to achieve slightly better final solutions compared to EvoMAL on average; however, as shown by the final inference testing error rates in Tables II and III, these don't necessarily correspond to better final inference performance. This discrepancy between meta-training curves and final inference is also observed in the inverse case, where EvoMAL is shown to have much better learning curves than both TaylorGLO and GP-LFL, e.g. in CIFAR-10 AliCNN-C and ResNet-18. However, the final inference error rates of EvoMAL in Table II are only marginally better than those of TaylorGLO and GP-LFL.

This phenomenon is likely due to some of the meta-learned loss functions implicitly tuning the learning rate (discussed further in V-C). Implicit learning rate tuning can result in increased convergence capabilities, i.e., faster learning which results in better fitness when using partial training sessions, but does not necessarily imply a strongly generalizing and robustly trained model at meta-testing time when using full training sessions.

TABLE V
AVERAGE RUN-TIME OF THE *META-TRAINING* PROCESS FOR EACH OF THE BENCHMARK METHODS

Task and Model	ML ³	TaylorGLO	GP-LFL	EvoMAL
Diabetes				
MLP	0.01	0.83	0.53	1.93
Boston				
MLP	0.01	0.85	0.46	1.59
California				
MLP	0.01	0.94	0.84	2.06
MNIST				
Logistic	0.02	1.44	0.61	2.45
MLP	0.03	2.06	0.77	5.31
LeNet-5	0.03	2.30	0.82	3.29
CIFAR-10				
AlexNet	0.03	3.75	1.04	5.90
VGG-16	0.12	4.67	0.89	9.12
AlIcNN-C	0.12	4.55	0.90	8.77
ResNet-18	0.60	9.14	1.02	54.22
PreResNet	0.41	8.40	0.96	41.73
WideResNet	0.63	12.85	0.98	66.72
SqueezeNet	0.12	4.95	0.41	11.18
CIFAR-100				
WideResNet	0.20	20.34	1.34	57.61
PyramidNet	0.20	24.83	1.32	49.89
SVHN				
WideResNet	0.20	41.57	1.09	67.28

Each algorithm is run on a single nvidia RTX A5000, and the results are reported in hours.

1) *Meta-Training Run-Time*: The use of a two-stage discovery process by EvoMAL enables the development of highly effective loss functions, as shown by the meta-training and meta-testing results. Producing on average models with superior inference performance compared to TaylorGLO and ML³ which only optimize the coefficients/weights of fixed parametric structures, and GP-LFL which uses no local-search techniques. However, this bi-level optimization procedure where both the model structure and parameters are inferred adversely affects the computational efficiency of the meta-learning process, as shown in Table V, which reports the average run-time (in hours) of meta-training for each loss function learning method.

The results show that EvoMAL is more computationally expensive than ML³ and GP-LFL and approximately twice as expensive as TaylorGLO on average. Although EvoMAL is computationally expensive, it should be emphasized that this is still dramatically more efficient than GLO [28], the bi-level predecessor to TaylorGLO, whose costly meta-learning procedure required a supercomputer for even very simple datasets such as MNIST.

The bi-level optimization process of EvoMAL is made computationally tractable by replacing the costly CMA-ES loss optimization stage from GLO with a significantly more efficient gradient-based procedure. On CIFAR-10, using the relatively




Fig. 5. Frequencies of occurrence of the time saving filters, and the corresponding frequencies of loss optimization and evaluation throughout the symbolic search process. Reporting the average frequencies across all task + model pairs and independent executions of EvoMAL.

small network, PreResNet-20 GLO required 11,120 partial training sessions and approximately 171 GPU days of computation [25], [28] compared to EvoMAL, which only needed on average 1.7 GPU days. In addition, the reduced run times of EvoMAL can also be partially attributed to the application of time-saving filters, which enables a subset of the loss optimizations and fitness evaluations to be either cached or obviated entirely.

To summarize the effects of the time-saving filters, a set of histograms are presented in Fig. 5 showing the frequency of occurrence throughout the evolutionary process. Examining the symbolic equivalence filter results, it is observed that, on average, $\sim 10\%$ of the loss functions are identified as being symbolically equivalent at each generation; consequently, these loss functions have their fitness cached, and $\sim 90\%$ of the loss functions progress to the next stage and are optimized. Regarding the pre-evaluation filters, the rejection protocol initially rejects the majority of the optimized loss functions early in the search for being unpromising, automatically assigning them the worst-case fitness. In contrast, in the late stages of the symbolic search, this filter occurs incrementally less frequently, suggesting that convergence is being approached, further supported by the frequency of occurrence of the gradient equivalence filter, which caches few loss functions at the start of the search, but many near the end. Due to this aggressive filtering, only $\sim 25\%$ of the population at each generation have their fitness evaluated, which helps to reduce the run-time further. Note that




Fig. 6. Example loss functions meta-learned by EvoMAL. The left plot shows classification loss functions, and right shows regression loss functions. Best viewed in color.

TABLE VI

EXAMPLE LOSS FUNCTIONS META-LEARNED BY EVOMAL, WHERE SOLUTIONS HAVE BEEN NUMERICALLY AND ALGEBRAICALLY SIMPLIFIED

a.	$ \log(\sqrt{y \cdot f_\theta(x)} + \epsilon) $
b.	$\log(y \cdot f_\theta(x)) + \sqrt{\log(y \cdot f_\theta(x))}^2$
c.	$((y - f_\theta(x))/f_\theta(x))^4$
d.	$y^3 + f_\theta(x)^3 + f_\theta(x) \cdot y^2 + f_\theta(x)^2 \cdot y$
e.	$ \log((y \cdot f_\theta(x))^2 + \epsilon) $
f.	$(f_\theta(x) - \min(\max(y, -1), 1))^2$
g.	$ y \cdot (\sqrt{ y/\sqrt{1 + f_\theta(x)^2} }) $
h.	$ (y/(1 + \log(y - 1)^2)) - f_\theta(x) $
i.	$\sqrt{ y \cdot (y - f_\theta(x)) }$
j.	$\sqrt{ y - f_\theta(x) }$

Furthermore, the parameters ϕ have also been omitted for improved clarity and parsimony.

the run-time of EvoMAL can be further reduced for large-scale optimization problems through parallelization to distribute loss function optimization and evaluation across multiple GPUs or clusters.

C. Meta-Learned Loss Functions

To better understand why the meta-learned loss functions produced by EvoMAL are so performant, an analysis is conducted on a subset of the interesting loss functions found throughout the experiments. In Fig. 6, examples of the meta-learned loss functions produced by EvoMAL are presented. The corresponding loss functions are also given symbolically in Table VI.

1) *Learned Loss Functions for Classification Tasks:* The classification loss functions meta-learned by EvoMAL appear to converge upon three classes of loss functions. First are cross-entropy loss variants such as loss functions a) and b), which closely resemble the cross-entropy loss functionally and symbolically. Second, are loss functions that have similar characteristics to the parametric focal loss [85], such as loss functions c) and d). These loss functions recalibrate how easy and hard samples are prioritized; in most cases, very little or no loss is attributed to high-confidence correct predictions, while significant loss is attributed to high-confidence wrong predictions. Finally, loss

functions such as e) which demonstrate unintuitive behavior, such as assigning more loss to confident and correct solutions relative to unconfident and correct solutions, a characteristic which induces implicit label smoothing regularization [86], [87], see analysis in Appendix A, available online.

2) *Learned Loss Functions for Regression Tasks:* Analyzing the regression loss functions it is observed that there are several unique behaviors particularly around incorporating strategies for improving robustness to outliers. For example in f) the loss function takes on the form of the squared loss; however, it incorporates a thresholding operation via the max operator for directly limiting the size of the ground truth label. Alternatively in loss functions g), i), and j), there is frequent use of the square root operator, which decreases the loss attributed to increasingly large errors. Finally, we observe that many of the learned loss functions, take on shapes similar to well-known handcrafted loss functions; for example, h) appears to be a cross between the absolute loss and the Cauchy (Lorentzian) loss [88].

3) *General Observations:* In addition to the trends identified for classification and regression, respectively, there are several more noteworthy trends:

- Structurally complex loss functions typically perform worse relative to simpler loss functions, potentially due to the increased meta or base optimization difficulties. This is a similar finding to what was found when meta-learning activation functions [10].
- Many of the learned loss functions in classification are asymmetric, producing different loss values for false positive and false negative predictions, often caused by exploiting $f_\theta(x)$ softmax output activation, where the sum of the class-wise outputs is required to equal to 1.

D. Loss Landscapes Analysis

Loss function learning as a paradigm has consistently shown to be an effective way of improving performance; however, it is not yet fully understood what exactly meta-learned loss functions are learning and why they are so performant compared to their handcrafted counterparts. In [25] and [27], it was found that the loss landscapes of models trained with learned loss functions produce flatter landscapes relative to those trained with




Fig. 7. Average 1D loss landscapes generated on the CIFAR-10 dataset, the top figure shows the loss landscapes generated on AllCNN-C, and the bottom figure shows the loss landscapes generated on AlexNet. The landscapes show the average (mean) loss taken across the 5 independent executions of each algorithm on each task+model pair. Best viewed in color.

the cross-entropy loss. The flatness of a loss landscape has been hypothesized to correspond closely to a model’s generalization capabilities [89], [90], [91], [92]; thus, they conclude that meta-learned loss functions improve generalization. These findings are independently reproduced and are shown in Fig. 7. The loss landscapes are generated using the filter-wise normalization method [91], which plots a normalized random direction of the weight space θ .

The loss landscapes visualizations show that the loss functions developed by EvoMAL can produce flatter loss landscapes on average in contrast to those produced by the cross-entropy, ML^3 , TaylorGLO, and GP-LFL, as shown in the top figure which shows the landscapes generated using AllCNN-C. In contrast to prior findings, we also find that in some cases the meta-learned loss functions produced in our experiments show relatively sharper loss landscapes compared to those produced using the cross-entropy loss, as shown in the bottom figure which shows loss landscapes generated using the base model AlexNet. These findings suggest that the relative flatness of the loss landscape does not fully explain why meta-learned loss functions can produce improved performance, especially since there is evidence that sharp minima can generalize well i.e. “*flat vs sharp*” debate [93].

E. Implicit Learning Rate Tuning

Another explanation for why meta-learned loss functions improve performance over handcrafted loss functions is that they can implicitly tune the (base) learning rate since for some suitably expressive representation of \mathcal{M} , since

$$\exists \alpha \exists \phi : \theta - \alpha \nabla_{\theta} \mathcal{L}_{\mathcal{T}} \approx \theta - \nabla_{\theta} \mathcal{M}_{\phi}^{\top}. \quad (13)$$

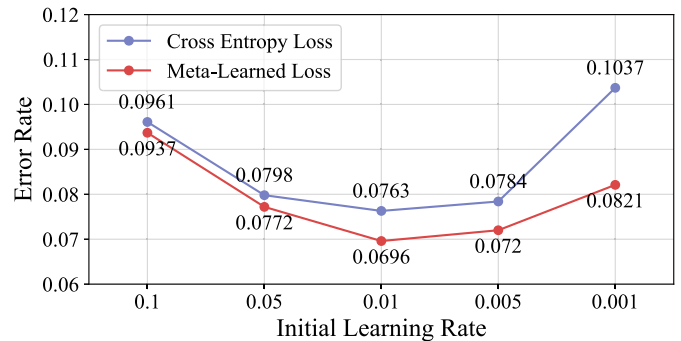


Fig. 8. Grid search comparing the average error rate of the baseline cross-entropy loss and EvoMAL on CIFAR-10 AllCNN-C across a set of base learning rate values, where 5 executions of each algorithm are performed on each learning rate value.

Thus, performance improvement when using meta-learned loss functions may be the indirect result of a change in the learning rate, scaling the resulting gradient of the loss function. To validate the implicit learning rate hypothesis, a grid search is performed over the base-learning rate using the cross-entropy loss and EvoMAL on CIFAR-10 AllCNN-C. The results are shown in Fig. 8.

The results show that the base learning rate α is a crucial hyper-parameter that influences the performance of both the baseline and EvoMAL. However, in the case of EvoMAL, it is found that when using a relatively small α value, the base learning rate is implicitly tuned, and the loss function learning algorithm achieves an artificially large performance margin compared to the baseline. Implicit learning rate tuning of a similar magnitude is also observed when using relatively large α values; however, the algorithm’s stability is inconsistent, with some runs failing to converge. Finally, when a near-optimal α value is used, performance improvement is consistently better than the baseline. These results indicate two key findings: (1) meta-learned loss functions improve upon handcrafted loss functions and that the performance improvement when using meta-learned loss functions is not primarily a result of implicit learning rate tuning when α is tuned. (2) The base learning rate α can be considered as part of the initialization of the meta-learned loss function, as it determines the initial scale of the loss function.

VI. CONCLUSIONS AND FUTURE WORK

This work presents a new framework for meta-learning symbolic loss function via a hybrid neuro-symbolic search approach called Evolved Model-Agnostic Loss (EvoMAL). The proposed technique uses genetic programming to learn a set of expression tree-based loss functions, which are subsequently transformed into a new network-style representation using a newly proposed transitional procedure. This new representation enables the integration of a computationally tractable gradient-based local-search approach to enhance the search capabilities significantly. Unlike previous approaches, which stack evolution-based

techniques, EvoMAL's efficient local-search enables loss function learning on commodity hardware.

The experimental results confirm that EvoMAL consistently meta-learns loss functions which can produce more performant models compared to those trained with conventional handcrafted loss functions, as well as other state-of-the-art loss function learning techniques. Furthermore, analysis of some of the meta-learned loss functions reveals several key findings regarding common loss function structures and how they interact with the models trained by them. Finally, automating the task of loss function selection has shown to enable a diverse and creative set of loss functions to be generated, which would not be replicable through a simple grid search over handcrafted loss functions.

There are many promising future research directions as a consequence of this work. Regarding algorithmic extensions, the EvoMAL framework is general in its design. It would be interesting to extend it to different meta-learning applications such as gradient-based optimizers or activation functions. In terms of loss function learning, a natural extension of the work would be to meta-learn the loss function and other deep neural network components simultaneously, similar to the methods presented in [17], [94], [95], [96]. For example, the meta-learned loss functions could consider additional arguments such as the timestep or model weights, which would implicitly induce learning rate scheduling or weight regularization, respectively. Another example would be to combine neural architecture search with loss function learning, as the experiments in this work use handcrafted neural network architectures which are biased towards the squared error loss or cross-entropy loss since they were designed to optimize for that specific loss function. Larger performance gains may be achieved using custom neural network architectures explicitly designed for the meta-learned loss functions.

REFERENCES

- [1] J. Vanschoren, "Meta-learning: A survey," *CoRR*, vol. abs/1810.03548, 2018. [Online]. Available: <https://arxiv.org/abs/1810.03548>
- [2] H. Peng, "A comprehensive overview and survey of recent advances in meta-learning," *CoRR*, vol. abs/2004.11149, 2020. [Online]. Available: <https://arxiv.org/abs/2004.11149>
- [3] T. Hospedales, A. Antoniou, P. Micaelli, and A. Storkey, "Meta-learning in neural networks: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 9, pp. 5149–5169, Sep. 2022.
- [4] H. Altae-Tran, B. Ramsundar, A. S. Pappu, and V. Pande, "Low data drug discovery with one-shot learning," *ACS Central Sci.*, vol. 3, no. 4, pp. 283–293, 2017.
- [5] A. Ignatov et al., "AI benchmark: All about deep learning on smartphones in 2019," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. Workshop*, 2019, pp. 3617–3635.
- [6] J. Schmidhuber, "Evolutionary principles in self-referential learning," Ph.D. dissertation, Technische Universität München, 1987.
- [7] J. Schmidhuber, "Learning to control fast-weight memories: An alternative to dynamic recurrent networks," *Neural Comput.*, vol. 4, no. 1, pp. 131–139, 1992.
- [8] S. Bengio, Y. Bengio, and J. Cloutier, "Use of genetic programming for the search of a new learning rule for neural networks," in *Proc. IEEE 1st Conf. Evol. Comput. World Congr. Comput. Intell.*, 1994, pp. 324–327.
- [9] M. Andrychowicz et al., "Learning to learn by gradient descent by gradient descent," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 3981–3989.
- [10] P. Ramachandran, B. Zoph, and Q. V. Le, "Searching for activation functions," in *Proc. 6th Int. Conf. Learn. Representations*, 2018.
- [11] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *Proc. Int. Conf. Mach. Learn.*, PMLR, 2017, pp. 1126–1135.
- [12] A. Nichol, J. Achiam, and J. Schulman, "On first-order meta-learning algorithms," *CoRR*, vol. abs/1803.02999, 2018. [Online]. Available: <http://arxiv.org/abs/1803.02999>
- [13] A. Rajeswaran, C. Finn, S. M. Kakade, and S. Levine, "Meta-learning with implicit gradients," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 113–124.
- [14] X. Song, W. Gao, Y. Yang, K. Choromanski, A. Pacchiano, and Y. Tang, "ES-MAML: Simple hessian-free meta learning," in *Proc. 8th Int. Conf. Learn. Representations*, 2020.
- [15] J. Kim et al., "Auto-meta: Automated gradient based meta learner search," *CoRR*, vol. abs/1806.06927, 2018. [Online]. Available: <http://arxiv.org/abs/1806.06927>
- [16] K. O. Stanley, J. Clune, J. Lehman, and R. Miikkulainen, "Designing neural networks through neuroevolution," *Nat. Mach. Intell.*, vol. 1, no. 1, pp. 24–35, 2019.
- [17] T. Elsken, B. Staffler, J. H. Metzen, and F. Hutter, "Meta-learning of neural architectures for few-shot learning," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 12 365–12 375.
- [18] Y. Ding et al., "Learning to learn by jointly optimizing neural architecture and weights," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 129–138.
- [19] E. Real, C. Liang, D. So, and Q. Le, "AutoML-Zero: Evolving machine learning algorithms from scratch," in *Proc. Int. Conf. Mach. Learn.*, PMLR, 2020, pp. 8007–8019.
- [20] J. D. Co-Reyes et al., "Evolving reinforcement learning algorithms," in *Proc. 9th Int. Conf. Learn. Representations*, 2021.
- [21] Q. Wang, Y. Ma, K. Zhao, and Y. Tian, "A comprehensive survey of loss functions in machine learning," *Ann. Data Sci.*, vol. 9, no. 2, pp. 187–212, 2022.
- [22] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning representations by back-propagating errors," *Nature*, vol. 323, no. 6088, pp. 533–536, 1986.
- [23] D. H. Wolpert and W. G. Macready, "No free lunch theorems for optimization," *IEEE Trans. Evol. Comput.*, vol. 1, no. 1, pp. 67–82, Apr. 1997.
- [24] S. Bechtle et al., "Meta-learning via learned loss," in *Proc. 25th Int. Conf. Pattern Recognit.*, 2021, pp. 4161–4168.
- [25] S. Gonzalez and R. Miikkulainen, "Optimizing loss functions through multi-variate Taylor polynomial parameterization," in *Proc. Genet. Evol. Comput. Conf.*, 2021, pp. 305–313.
- [26] B. Gao, H. Gouk, and T. M. Hospedales, "Searching for robustness: Loss learning for noisy classification tasks," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 6670–6679.
- [27] B. Gao, H. Gouk, Y. Yang, and T. Hospedales, "Loss function learning for domain generalization by implicit gradient," in *Proc. Int. Conf. Mach. Learn.*, PMLR, 2022, pp. 7002–7016.
- [28] S. Gonzalez and R. Miikkulainen, "Improved training speed, accuracy, and data utilization through loss function optimization," in *Proc. IEEE Congr. Evol. Comput.*, 2020, pp. 1–8.
- [29] P. Liu et al., "Loss function discovery for object detection via convergence-simulation driven search," in *Proc. Int. Conf. Learn. Representations*, 2020.
- [30] H. Li, T. Fu, J. Dai, H. Li, G. Huang, and X. Zhu, "Autoloss-zero: Searching loss functions from scratch for generic tasks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 1009–1018.
- [31] J. R. Koza, *Genetic Programming: On the Programming of Computers by Means of Natural Selection*, vol. 1. Cambridge, MA, USA: MIT Press, 1992.
- [32] D. Maclaurin, D. Duvenaud, and R. Adams, "Gradient-based hyperparameter optimization through reversible learning," in *Proc. Int. Conf. Mach. Learn.*, PMLR, 2015, pp. 2113–2122.
- [33] E. Grefenstette et al., "Generalized inner loop meta-learning," *CoRR*, vol. abs/1910.01727, 2019. [Online]. Available: <http://arxiv.org/abs/1910.01727>
- [34] J. Grabocka, R. Scholz, and L. Schmidt-Thieme, "Learning surrogate losses," vol. abs/1910.01727, 2019. [Online]. Available: <http://arxiv.org/abs/1910.01727>
- [35] C. Huang et al., "Addressing the loss-metric mismatch with adaptive loss alignment," in *Proc. Int. Conf. Mach. Learn.*, PMLR, 2019, pp. 2891–2900.

- [36] R. Houthoof et al., “Evolved policy gradients,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2018, pp. 5405–5414.
- [37] L. Wu et al., “Learning to teach with dynamic loss functions,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2018, pp. 6467–6478.
- [38] A. Antoniou and A. J. Storkey, “Learning to learn by self-critique,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 9940–9950.
- [39] L. Kirsch, S. van Steenkiste, and J. Schmidhuber, “Improving generalization in meta reinforcement learning using learned objectives,” in *Proc. 8th Int. Conf. Learn. Representations*, 2020.
- [40] A. Collet, A. Bazco-Nogueras, A. Banchs, and M. Fiore, “Loss meta-learning for forecasting,” 2022. [Online]. Available: <https://openreview.net/forum?id=rczz7TUKIIB>
- [41] Z. Leng et al., “PolyLoss: A polynomial expansion perspective of classification loss functions,” in *Proc. 10th Int. Conf. Learn. Representations*, 2022.
- [42] C. Raymond, Q. Chen, B. Xue, and M. Zhang, “Online loss function learning,” *CoRR*, vol. abs/2301.13247, 2023. [Online]. Available: <https://doi.org/10.48550/arXiv.2301.13247>
- [43] Y. Balaji, S. Sankaranarayanan, and R. Chellappa, “MetaReg: Towards domain generalization using meta-regularization,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2018, pp. 998–1008.
- [44] J. T. Barron, “A general and adaptive robust loss function,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 4331–4339.
- [45] Y. Li, Y. Yang, W. Zhou, and T. Hospedales, “Feature-critic networks for heterogeneous domain generalisation,” in *Proc. Int. Conf. Mach. Learn.*, PMLR, 2019, pp. 3915–3924.
- [46] A. F. Psaros, K. Kawaguchi, and G. E. Karniadakis, “Meta-learning PINN loss functions,” *J. Comput. Phys.*, vol. 458, 2022, Art. no. 111121.
- [47] A. Topchy and W. F. Punch, “Faster genetic programming based on local gradient search of numeric leaf values,” in *Proc. Genet. Evol. Comput. Conf.*, Morgan Kaufmann, 2001, pp. 155–162.
- [48] W. Smart and M. Zhang, “Applying online gradient descent search to genetic programming for object recognition,” in *Proc. 2nd Workshop Australas. Inf. Secur. Data Mining Web Intell. Softw. Internationalisation*, 2004, pp. 133–138.
- [49] M. Zhang and W. Smart, “Learning weights in genetic programs using gradient descent for object recognition,” in *Proc. Workshops Appl. Evol. Comput.*, Springer, 2005, pp. 417–427.
- [50] R. E. Wengert, “A simple automatic derivative evaluation program,” *Commun. ACM*, vol. 7, no. 8, pp. 463–464, 1964.
- [51] J. Domke, “Generic methods for optimization-based modeling,” in *Proc. Artif. Intell. Statist.*, PMLR, 2012, pp. 318–326.
- [52] C.-A. Deledalle, S. Vaïter, J. Fadili, and G. Peyré, “Stein unbiased gradient estimator of the risk (sugar) for multiple parameter selection,” *SIAM J. Imag. Sci.*, vol. 7, no. 4, pp. 2448–2487, 2014.
- [53] S. Gonzalez et al., “Improving deep learning through loss-function evolution,” Ph.D. dissertation, Dept. Comput. Sci., Univ. Texas Austin, 2020.
- [54] B. Gao, “Meta-learning to optimise: Loss functions and update rules,” Ph.D. dissertation, School Inf., Univ. Edinburgh, 2023.
- [55] J. Ni, R. H. Driberg, and P. I. Rockett, “The use of an analytic quotient operator in genetic programming,” *IEEE Trans. Evol. Comput.*, vol. 17, no. 1, pp. 146–152, Feb. 2013.
- [56] Q. Chen, B. Xue, and M. Zhang, “Generalisation and domain adaptation in GP with gradient descent for symbolic regression,” in *Proc. IEEE Congr. Evol. Comput.*, 2015, pp. 1137–1144.
- [57] L. Franceschi, M. Donini, P. Frasconi, and M. Pontil, “Forward and reverse gradient-based hyperparameter optimization,” in *Proc. Int. Conf. Mach. Learn.*, PMLR, 2017, pp. 1165–1173.
- [58] L. Franceschi, P. Frasconi, S. Salzo, R. Grazzi, and M. Pontil, “Bilevel programming for hyperparameter optimization and meta-learning,” in *Proc. Int. Conf. Mach. Learn.*, PMLR, 2018, pp. 1568–1577.
- [59] A. Shaban, C.-A. Cheng, N. Hatch, and B. Boots, “Truncated back-propagation for bilevel optimization,” in *Proc. 22nd Int. Conf. Artif. Intell. Statist.*, PMLR, 2019, pp. 1723–1732.
- [60] D. Scieur, Q. Bertrand, G. Gidel, and F. Pedregosa, “The curse of unrolling: Rate of differentiating through optimization,” *NIPS*, 2022.
- [61] P. Wong and M. Zhang, “Algebraic simplification of GP programs during evolution,” in *Proc. 8th Annu. Conf. Genet. Evol. Comput.*, 2006, pp. 927–934.
- [62] D. Kinzett, M. Johnston, and M. Zhang, “Numerical simplification for bloat control and analysis of building blocks in genetic programming,” *Evol. Intell.*, vol. 2, pp. 151–168, 2009.
- [63] N. Javed, F. Gobet, and P. Lane, “Simplification of genetic programs: A literature survey,” *Data Mining Knowl. Discov.*, vol. 36, no. 4, pp. 1279–1300, 2022.
- [64] J. J. Grefenstette and J. M. Fitzpatrick, “Genetic search with approximate function evaluations,” in *Proc. Int. Conf. Genet. Algorithms Appl.*, 1985, pp. 112–120.
- [65] Y. Jin, “Surrogate-assisted evolutionary computation: Recent advances and future challenges,” *Swarm Evol. Comput.*, vol. 1, no. 2, pp. 61–70, 2011.
- [66] Y. Wu, M. Ren, R. Liao, and R. Grosse, “Understanding short-horizon bias in stochastic meta-optimization,” in *Proc. 6th Int. Conf. Learn. Representations*, 2018.
- [67] F.-A. Fortin, F.-M. De Rainville, M.-A. G. Gardner, M. Parizeau, and C. Gagné, “DEAP: Evolutionary algorithms made easy,” *J. Mach. Learn. Res.*, vol. 13, no. 1, pp. 2171–2175, 2012.
- [68] A. Paszke et al., “PyTorch: An imperative style, high-performance deep learning library,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 8026–8037.
- [69] M. Lichman, “UCI machine learning repository.” [Online]. Available: <http://archive.ics.uci.edu/ml>
- [70] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, “Gradient-based learning applied to document recognition,” *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998.
- [71] A. Krizhevsky and G. Hinton, “Learning multiple layers of features from tiny images,” Univ. Toronto, Tech. Rep., 2009. [Online]. Available: <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf>
- [72] Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, and A. Y. Ng, “Reading digits in natural images with unsupervised feature learning,” in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2011.
- [73] A. G. Baydin, R. Cornish, D. M. Rubio, M. Schmidt, and F. Wood, “Online learning rate adaptation with hypergradient descent,” in *Proc. 6th Int. Conf. Learn. Representations*, 2018.
- [74] P. McCullagh et al., *Generalized Linear Models*. Evanston, IL, USA: Routledge, 1989.
- [75] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “ImageNet classification with deep convolutional neural networks,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2012, pp. 1097–1105.
- [76] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” in *Proc. 3th Int. Conf. Learn. Representations*, 2015.
- [77] J. T. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller, “Striving for simplicity: The all convolutional net,” in *Proc. 3th Int. Conf. Learn. Representations*, 2015.
- [78] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.
- [79] K. He, X. Zhang, S. Ren, and J. Sun, “Identity mappings in deep residual networks,” in *Proc. Eur. Conf. Comput. Vis.*, Springer, 2016, pp. 630–645.
- [80] S. Zagoruyko and N. Komodakis, “Wide residual networks,” in *Proc. British Mach. Vis. Conf.*, 2016.
- [81] F. N. Iandola, S. Han, M. W. Moskewicz, K. Ashraf, W. J. Dally, and K. Keutzer, “SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and < 0.5 mb model size,” *CoRR*, vol. abs/1602.07360, 2016. [Online]. Available: <http://arxiv.org/abs/1602.07360>
- [82] D. Han, J. Kim, and J. Kim, “Deep pyramidal residual networks,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 5927–5935.
- [83] W. Liu, Z. Wang, X. Liu, N. Zeng, Y. Liu, and F. E. Alsaadi, “A survey of deep neural network architectures and their applications,” *Neurocomputing*, vol. 234, pp. 11–26, 2017.
- [84] J. Gu et al., “Recent advances in convolutional neural networks,” *Pattern Recognit.*, vol. 77, pp. 354–377, 2018.
- [85] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, “Focal loss for dense object detection,” in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 2980–2988.
- [86] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, “Rethinking the inception architecture for computer vision,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 2818–2826.
- [87] R. Müller, S. Kornblith, and G. E. Hinton, “When does label smoothing help?,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 4694–4703.
- [88] M. J. Black and A. Rangarajan, “On the unification of line processes, outlier rejection, and robust statistics with applications in early vision,” *Int. J. Comput. Vis.*, vol. 19, no. 1, pp. 57–91, 1996.
- [89] S. Hochreiter and J. Schmidhuber, “Flat minima,” *Neural Comput.*, vol. 9, no. 1, pp. 1–42, 1997.
- [90] N. S. Keskar, D. Mudigere, J. Nocedal, M. Smelyanskiy, and P. T. P. Tang, “On large-batch training for deep learning: Generalization gap and sharp minima,” in *Proc. 5th Int. Conf. Learn. Representations*, 2017.

- [91] H. Li, Z. Xu, G. Taylor, C. Studer, and T. Goldstein, "Visualizing the loss landscape of neural nets," in *Proc. Adv. Neural Inf. Process. Syst.*, 2018, pp. 6391–6401.
- [92] P. Chaudhari et al., "Entropy-SGD: Biasing gradient descent into wide valleys," *J. Stat. Mechanics Theory Experiment*, vol. 2019, no. 12, 2019, Art. no. 124018.
- [93] L. Dinh, R. Pascanu, S. Bengio, and Y. Bengio, "Sharp minima can generalize for deep nets," in *Proc. Int. Conf. Mach. Learn.*, PMLR, 2017, pp. 1019–1028.
- [94] S. Ravi and H. Larochelle, "Optimization as a model for few-shot learning," in *Proc. Int. Conf. Learn. Representations*, 2017.
- [95] Z. Li, F. Zhou, F. Chen, and H. Li, "Meta-SGD: Learning to learn quickly for few-shot learning," *CoRR*, vol. abs/1707.09835, 2017. [Online]. Available: <http://arxiv.org/abs/1707.09835>
- [96] S. Baik, J. Choi, H. Kim, D. Cho, J. Min, and K. M. Lee, "Meta-learning with task-adaptive loss function for few-shot learning," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 9465–9474.



Mengjie Zhang (Fellow, IEEE) received the BE and ME degrees from the Artificial Intelligence Research Center, Agricultural University of Hebei, Hebei, China, and the PhD degree in computer science from RMIT University, Melbourne, VIC, Australia, in 1989, 1992, and 2000, respectively. He is currently a professor of Computer Science, head of the Evolutionary Computation Research Group, and the associate dean (Research and Innovation) in the Faculty of Engineering. His current research interests include evolutionary computation, particularly genetic programming, particle swarm optimization, and learning classifier systems with application areas of image analysis, multiobjective optimization, feature selection and reduction, job shop scheduling, and transfer learning. He has published more than 600 research papers in refereed international journals and conferences. He is a fellow of the Royal Society of New Zealand and has been a panel member of the Marsden Fund (New Zealand Government Funding), and a member of ACM. He was the chair of the IEEE CIS Intelligent Systems and Applications Technical Committee, and chair for the IEEE CIS Emergent Technologies Technical Committee, the Evolutionary Computation Technical Committee, and a member of the IEEE CIS Award Committee. He is a vice-chair of the IEEE CIS Task Force on Evolutionary Feature Selection and Construction, a vice-chair of the Task Force on Evolutionary Computer Vision and Image Processing, and the founding chair of the IEEE Computational Intelligence Chapter in New Zealand. He is also a committee member of the IEEE NZ Central Section.



Christian Raymond (Graduate Student Member, IEEE) received the BSc and BSc (Hons) degrees from the School of Engineering and Computer Science, Victoria University of Wellington, New Zealand, in 2020 and 2021 respectively. He is currently working toward the PhD degree in artificial intelligence with the Evolutionary Computation Research Group (ECRG), Victoria University of Wellington. His research interests include meta-learning, hyperparameter optimization, and few-shot learning.



Qi Chen (Member, IEEE) received the BE degree in automation from the University of South China, Hunan, China in 2005, the ME degree in software engineering from the Beijing Institute of Technology, Beijing, China in 2007, and the PhD degree in computer science from the Victoria University of Wellington, New Zealand, in 2018. Since 2014, she has joined the Evolutionary Computation Research Group, Victoria University of Wellington (VUW). Currently, she is a lecturer with the School of Engineering and Computer Science, VUW. Qi's current

research mainly focuses on genetic programming for symbolic regression. Her research interests include machine learning, evolutionary computation, feature selection, feature construction, transfer learning, domain adaptation, and statistical learning theory. She serves as a reviewer of international conferences, including IEEE Congress on Evolutionary Computation, and international journals, including *IEEE Transactions on Cybernetics* and *IEEE Transactions on Evolutionary Computation*.



Bing Xue (Senior Member, IEEE) received the BSc degree from the Henan University of Economics and Law, Zhengzhou, China, in 2007, the MSc degree in management from Shenzhen University, Shenzhen, China, in 2010, and the PhD degree in computer science from the Victoria University of Wellington (VUW), New Zealand, in 2014. She is currently a professor of Computer Science, and program director of Science with the School of Engineering and Computer Science, VUW. She has more than 300 papers published in fully refereed international journals and

conferences, and her research focuses mainly on evolutionary computation, machine learning, classification, symbolic regression, feature selection, evolving deep neural networks, image analysis, transfer learning, and multi-objective machine learning. He is currently the chair of IEEE Computational Intelligence Society (CIS) Task Force on Transfer Learning and Transfer Optimization, vice-chair of IEEE CIS Evolutionary Computation Technical Committee, editor of IEEE CIS Newsletter, vice-chair of IEEE Task Force on Evolutionary Feature Selection and Construction, and vice-chair IEEE CIS Task Force on Evolutionary Deep Learning and Applications. She has also served as associate editor of several international journals, such as *IEEE Computational Intelligence Magazine* and *IEEE Transactions on Evolutionary Computation*.