

Regularized Evolutionary Population-Based Training

Jason Liang
Cognizant AI Labs
San Francisco, CA, USA

Santiago Gonzalez
Cognizant AI Labs
The Univ. of Texas at Austin

Hormoz Shahrzad
Cognizant AI Labs
The Univ. of Texas at Austin

Risto Miikkulainen
Cognizant AI Labs
The Univ. of Texas at Austin

ABSTRACT

Metalearning of deep neural network (DNN) architectures and hyperparameters has become an increasingly important area of research. At the same time, network regularization has been recognized as a crucial dimension to effective training of DNNs. However, the role of metalearning in establishing effective regularization has not yet been fully explored. There is recent evidence that loss-function optimization could play this role, however it is computationally impractical as an outer loop to full training. This paper presents an algorithm called Evolutionary Population-Based Training (EPBT) that interleaves the training of a DNN's weights with the metalearning of loss functions. They are parameterized using multivariate Taylor expansions that EPBT can directly optimize. Such simultaneous adaptation of weights and loss functions can be deceptive, and therefore EPBT uses a quality-diversity heuristic called Novelty Pulsation as well as knowledge distillation to prevent overfitting during training. On the CIFAR-10 and SVHN image classification benchmarks, EPBT results in faster, more accurate learning. The discovered hyperparameters adapt to the training process and serve to regularize the learning task by discouraging overfitting to the labels. EPBT thus demonstrates a practical instantiation of regularization metalearning based on simultaneous training.

CCS CONCEPTS

• **Computing methodologies** → **Neural networks**; *Distributed artificial intelligence*; Computer vision.

KEYWORDS

Deep Learning; Evolutionary Algorithms; Metalearning; Neural Networks; Regularization

ACM Reference Format:

Jason Liang, Santiago Gonzalez, Hormoz Shahrzad, and Risto Miikkulainen. 2021. Regularized Evolutionary Population-Based Training. In *2021 Genetic and Evolutionary Computation Conference (GECCO '21)*, July 10–14, 2021, Lille, France. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3449639.3459292>

1 INTRODUCTION

Training modern deep neural networks (DNNs) often requires extensive tuning, and many seminal architectures have been developed

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
GECCO '21, July 10–14, 2021, Lille, France

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-8350-9/21/07.
<https://doi.org/10.1145/3449639.3459292>

through a hand-design process that requires extensive expertise [1, 26, 28, 37]. To make this process easier and more productive, automated methods for metalearning and optimization of DNN hyperparameters and architectures have recently been developed, using techniques such as Bayesian optimization, reinforcement learning, and evolutionary search [36, 50, 57]. At the same time, regularization during training has become an important area of research, as preventing overfitting has been identified as crucial to the generalization capabilities of DNNs [27, 51]. An interesting opportunity is therefore emerging: Using metalearning to discover improved regularization mechanisms [2, 5, 44].

One promising such approach is evolution of loss functions. Instead of optimizing network structure or weights, evolution is used to modify the gradients, making it possible to automatically regularize the learning process [9, 18]. However, as in most of the earlier metalearning methods, evolution serves as an outer loop to network training. Such an approach is computationally prohibitive since fitness evaluations in principle need full training of DNNs. Also, the approach cannot adapt loss functions to different stages of learning.

Population-Based Training (PBT) was recently proposed to overcome this limitation in metalearning [21]. PBT interleaves DNN weight training with the optimization of hyperparameters that are relevant to the training process but also have no particular fixed value (e.g., learning rate). Such online adaptation is crucial in domains where the learning dynamics are non-stationary. Therefore, PBT forms a promising starting point for making loss-function optimization practical as well.

Building on PBT, this paper develops *Evolutionary Population-Based Training* (EPBT) as such an approach through four extensions. First, powerful heuristics from evolutionary black-box optimization are employed to discover promising combinations of hyperparameters for DNN training. In particular, EPBT uses selection, mutation, and crossover operators adapted from genetic algorithms [53] to find good solutions.

Second, a recently developed loss function parameterization based on multivariate Taylor expansions called *TaylorGLO* [10] is combined with EPBT to optimize loss functions. This parameterization makes it possible to encode many different loss functions compactly and works on a variety of DNN architectures. In prior work, it was found to discover loss functions that result in faster training and better convergence than the standard cross-entropy loss.

Third, an interesting challenge emerges when both loss functions and the weights are adapted at the same time: The problem becomes inherently deceptive. Configurations that allow for fast learning in the beginning are often bad for fine tuning at the end of training. To overcome this problem, EPBT makes use of Novelty Pulsation [45, 46], a powerful heuristic for maintaining population diversity and escaping from deceptive traps during optimization.

Fourth, another challenge with the coadaptation is that training is noisy and can overfit to the validation dataset during evolution. EPBT thus introduces a new variant of knowledge distillation [16] called Population-Based Distillation. This method stabilizes training, helps reduce evaluation noise, and thus allows evolution to make more reliable progress.

Experiments on the CIFAR-10 and SVHN image classification benchmarks with several network architectures show that EPBT results in faster, more accurate learning. An analysis of the shapes of the discovered loss functions suggests that they penalize overfitting, thus regularizing the learning process automatically. Different loss shapes are most effective at different stages of the training process, suggesting that an adaptive loss functions perform better than one that remains static throughout the training, thus taking advantage of the synergy between learning and evolution.

2 BACKGROUND AND RELATED WORK

This section summarizes relevant work in metalearning and regularization of DNNs, especially through population-based methods. Previous work on loss function optimization and novelty pulsation is also reviewed.

2.1 Metalearning

Metalearning of good DNN hyperparameters and architectures is a highly active field of research. One popular approach is to use reinforcement learning to tune a controller that generates the designs [42, 57, 58]. Another approach is to make the metalearning differentiable to the performance of the DNN, and then learn by gradient descent [31, 34].

Recently, metalearning methods based on evolutionary algorithms (EA) have also gained popularity. These methods can optimize DNNs of arbitrary topology and structure [36], achieving state-of-the-art results e.g. on large-scale image classification benchmarks [44], and demonstrating good trade-offs in multiple objectives such as performance and network complexity [33]. Many of these EAs use proven and time-tested heuristics such as mutation, crossover, selection, and elitism [8, 52, 53] to perform black-box optimization on arbitrary complex objectives. Advanced EAs such as CMA-ES [13] have also optimized DNN hyperparameters successfully in high-dimensional search spaces [32] and are competitive with statistical hyperparameter tuning methods such as Bayesian optimization [24, 49, 50].

2.2 Population-based Training

One challenge shared by every DNN metalearning algorithm is deciding the right amount of training required to evaluate a network architecture and hyperparameter configuration on a benchmark task. Many algorithms simply stop training prematurely, assuming that the partially trained performance is correlated with the true performance [30, 36]. Other methods rely on weight sharing, where many candidate architectures share model layers [42], thus ensuring that the training time is amortized among all solutions being evaluated.

PBT [21] uses the weight sharing approach, which is more computationally efficient. PBT works by alternating between training models in parallel and tuning the model’s hyperparameters through

an exploit-and-explore strategy. During exploitation, the hyperparameters and weights of well-performing models are duplicated to replace the worst performing ones. During exploration, hyperparameters are randomly perturbed within a constrained search space. Because PBT never retrains models from scratch, its computational complexity scales only with the population size and not with the total number of hyperparameter configurations searched. Besides tuning training hyperparameters such as the learning rate, PBT has successfully discovered data augmentation schedules [17]. Therefore, PBT serves as a promising basis for the design of EPBT.

2.3 Regularization of DNNs

While metalearning seeks to find good DNN architectures, regularization is concerned about preventing DNNs from overfitting during training or optimization. Besides classic penalty-based approaches such as weight decay [38], there are methods that leverage the structure of DNN layers. One simple but popular approach is dropout [51], which randomly sets the outputs of a layer to zero. This approach helps prevent overfitting by forcing subsequent layers to adapt to the noise generated by the previous layers. A related technique is batch normalization [20], which normalizes the outputs of layers and prevents exploding gradients. These approaches work universally on almost all DNN architectures and problem domains and can even be combined.

Recently, much attention has been focused on manipulating training data to help regularize network training. Advanced data augmentation techniques such as cutout [6], and cutmix [55] purposely create more diverse distributions of the input data to improve generalization and avoid overfitting. Similarly, adversarial examples [12] are another way to regularize by training the network with inputs that are particularly difficult for it to get right. Techniques such as label smoothing [40] and knowledge distillation/self-distillation [16, 23] soften the training targets to ensure more properly behaved gradients, resulting in better generalization. Along the same lines, loss-function optimization is an alternative method to achieve better regularization through modification of gradients, as will be reviewed next.

2.4 Loss Function Optimization

DNNs are trained through the backpropagation of gradients that originate from a loss function [28]. Loss functions represent the primary training objective for a neural network. The choice of the loss function can have a significant impact on a network’s performance [4, 9, 22]. Recently, Genetic Loss Optimization (GLO) [9] was proposed as a new kind of metalearning, making it possible to automatically discover novel loss functions that can be used to train higher-accuracy neural networks in less time.

In GLO, loss functions are represented as trees and optimized through genetic programming [3]. This approach has the advantage of allowing arbitrarily complex loss functions. However, there are pathological functions in this search space with undesirable behaviors, such as discontinuities. To avoid these issues, this paper uses a loss-function representation based on multivariate Taylor expansions [10]. TaylorGLO parameterization is smoother, has guaranteed continuity and adjustable complexity, and is easier to

implement. How these loss functions are optimized with EPBT is described in the next section.

2.5 Novelty Selection and Pulsation

Novelty Selection is a form of novelty search that augments the original fitness-based selection with novelty and thus improves the quality-diversity of the population [29, 45]. Novelty is measured through a behavioral description of the individuals, i.e. a phenotypical feature vector that is not related to fitness. An initial set of m elite candidates is first selected based on fitness and sorted according to their novelty score S_i , measured as the sum of pairwise distances d of the individual’s behavior vector b_i to those of all other individuals j in the set:

$$S_i = \sum_{j=1}^m d(b_i, b_j). \quad (1)$$

The top k candidates from this set are selected as elites, skipping candidates that represent the same cluster.

In principle, fitness-based selection is more greedy than novelty selection, and could result in faster convergence. On the other hand, novelty selection explores more diverse candidates, which could help discover better regularization. In Novelty Pulsation [46], such exploitation and exploration are both leveraged by switching Novelty Selection on and off for every p generations, resulting in faster convergence and more reliable solutions. Novelty Pulsation plays an important role in keeping the population diverse enough to avoid deceptive interactions between weight and loss-function adaptation in EPBT. Furthermore, unlike other novelty-based methods such as Map Elites [39], Novelty Pulsation can be easily integrated with EPBT due to the simplicity of the algorithm and lack of computational overhead from maintaining an archive of individuals.

3 THE EBPT METHOD

This section describes in detail how EPBT implements regularization metalearning. EPBT utilizes genetic operators from black-box optimization to enhance hyperparameter metalearning in PBT, and combines it with loss-function metalearning. Their deceptive and overfitting interactions are mitigated through a selection heuristic based on quality-diversity, and through knowledge distillation.

3.1 Overview

The core concept of EPBT is intuitive and builds on extensive work already done with evolutionary optimization of DNNs [52]. Figure 1 shows how EPBT maximizes the fitness of a population of candidate solutions (individuals) over multiple iterations (generations). As a black-box method, EPBT requires no gradient information but only the fitness value of each individual. With EPBT, it is thus possible to apply metalearning to tasks where meta-gradients are not available.

At the beginning of generation g , the population \mathbb{M}_g consists of individuals M_{gi} . Each $M_{gi} = \{\mathbf{D}_{gi}, \mathbf{h}_{gi}, f_{gi}\}$ where \mathbf{D}_{gi} is a DNN model (defined as the weights of a user-specified architecture), \mathbf{h}_{gi} is a set of hyperparameters, and f_{gi} is a real-valued scalar fitness. In the Step 1 of the generation, f_{gi} is used to select promising individuals \hat{M}_{gi} to form a parent set $\hat{\mathbb{M}}_g$, where $\hat{\mathbb{M}}_g \subset \mathbb{M}_g$. In the Step 2, $\hat{\mathbb{M}}_g$ is used to create a set \mathbb{N}_g , which contains new individuals N_{gi} . Each of these new individuals inherits \mathbf{D}_{gi} unchanged from its

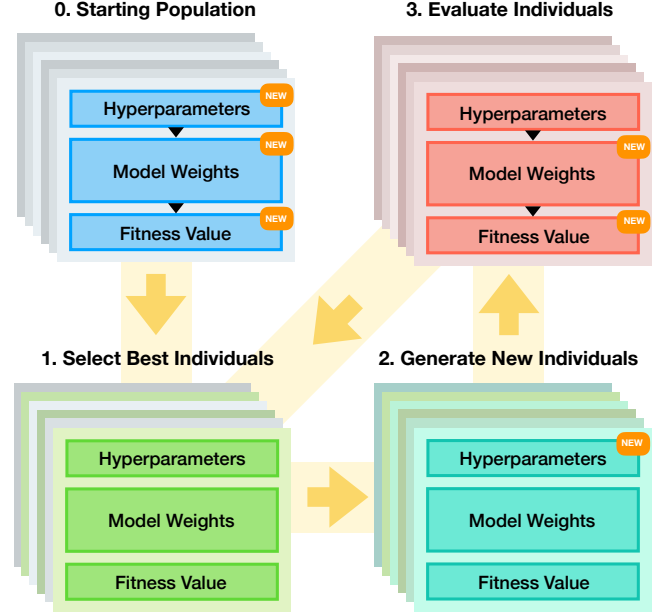


Figure 1: An overview of the EPBT evolutionary loop. EPBT begins by randomly initializing individuals, which are composed of hyperparameters, model weights, and fitness values. Next, EPBT runs for multiple generations in a three-step loop: (1) selection of the best individuals, (2) generation of new individuals, and (3) evaluation of these individuals. In Step 1, promising individuals are selected using a heuristic. In Step 2, new individuals with updated hyperparameters are created, but the weights and fitness are inherited. In Step 3, these individuals are evaluated on a task and have their model weights and fitness (i.e., performance in the task) updated. Thus, EPBT makes it possible to simultaneously train the network and evolve loss function parameterizations.

parent \hat{M}_{gi} , but with modified hyperparameters $\hat{\mathbf{h}}_{gi}$. The genetic operators used for generating \mathbb{N}_g will be described in more detail in the next subsection. Finally, in Step 3, each N_{gi} is evaluated by training \mathbf{D}_{gi} on a task or dataset, thereby creating a model with updated weights $\hat{\mathbf{D}}_{gi}$. The validation performance of $\hat{\mathbf{D}}_{gi}$ is used to determine a new fitness value \hat{f}_{gi} . Thus, by the end of generation g , the population pool contains the evaluated individuals $\hat{N}_{gi} \in \hat{\mathbb{N}}_g$, where $\hat{N}_{gi} = \{\hat{\mathbf{D}}_{gi}, \hat{\mathbf{h}}_{gi}, \hat{f}_{gi}\}$. This process is repeated for multiple generations until the fitness of the best individual in the population converges.

Within the core metalearning evolutionary loop, EPBT contains several other components. They include: (1) A collection of genetic operators specifically chosen for the task of hyperparameter optimization. (2) The Novelty Selection and Pulsation search heuristic that improves population diversity by preserving the most novel elites. (3) A TaylorGLO representation of loss functions with parameters that EPBT can optimize. (4) Population-Based Distillation method (PBD) that uses the best model in the population to help train other networks. Each of these components will be described in more detail below.

Algorithm 1 EPBT

Input: max generations n , initial population \mathbb{M}_0 , genetic operators τ, γ, ξ

for $g = 0$ **to** $n - 1$ **do**

1. Select $\hat{M}_{gi} = \{\mathbf{D}_{gi}, \mathbf{h}_{gi}, f_{gi}\}$ using τ
- 2a. Set $\hat{\mathbf{h}}_{gi} = \xi(\gamma(\mathbf{h}_{gi}))$
- 2b. Set $N_{gi} = \{\mathbf{D}_{gi}, \hat{\mathbf{h}}_{gi}\}$
- 3a. Evaluate \hat{N}_{gi} , set $\hat{N}_{gi} = \{\hat{\mathbf{D}}_{gi}, \hat{\mathbf{h}}_{gi}, \hat{f}_{gi}\}$
- 3b. Set $\hat{\mathbb{M}}_g$ to top k M_{gi} from \mathbb{M}_g
- 3c. Set $\mathbb{M}_{g+1} = \hat{N}_g \cup \hat{\mathbb{M}}_g$

end for

Note that since the evaluation of an individual does not depend on other individuals, the entire EPBT process can be parallelized. In the current implementation of EPBT, fitness evaluations are mapped onto a multi-process pool of workers on a single machine. Each worker has access to a particular GPU of the machine, and if there are multiple GPUs available, every GPU will be assigned to at least one worker. For the experiments in this paper, a single worker does not fully utilize the GPU and multiple workers can be trained in parallel without any slowdown.

3.2 Genetic Operators

EPBT uses standard evolutionary black-box optimization operators [53] to tune individuals. This subsection details how EPBT is initialized and how these operators are utilized through the three septs of each generation, with a summary in Algorithm 1.

Initialization: A population with P individuals is created as \mathbb{M}_0 . For each $M_{0i} \in \mathbb{M}_0$, \mathbf{D}_{0i} is set to a fixed DNN architecture and its weights are randomly initialized. Also, each variable in \mathbf{h}_{0i} is uniformly sampled from within a fixed range and f_{0i} is set to zero.

Step 1 – Tournament Selection: Using the tournament selection operator τ , t individuals are repeatedly chosen at random from \mathbb{M}_g . Each time, the individuals are compared and the one with the highest fitness is added to $\hat{\mathbb{M}}_g$. This process is repeated until $|\hat{\mathbb{M}}_g| = |\mathbb{M}_g| - k$, where k is the number of elites. The value $t = 2$ is commonly used in EA literature and in the experiments in this paper also.

Step 2 – Mutation and Crossover: For each \hat{M}_{gi} , a uniform mutation operator γ is applied by introducing multiplicative Gaussian noise independently to each variable in \mathbf{h}_{gi} . The mutation operator can randomly and independently reinitialize every variable as well. This approach allows for the exploration of novel combinations of hyperparameters. After mutation, a uniform crossover operator ξ is applied, where each variable in \mathbf{h}_{gi} is randomly swapped (50% probability) with the same variable from another individual in $\hat{\mathbb{M}}_g$, resulting in the creation of $\hat{\mathbf{h}}_{gi}$. \mathbf{D}_{gi} is copied from \hat{M}_{gi} and combined with $\hat{\mathbf{h}}_{gi}$ to form the unevaluated individual N_{gi} .

Step 3 – Fitness Evaluation with Elitism: The evaluation process proceeds as described above and results in evaluated individuals \hat{N}_g . After evaluation, EPBT uses an elitism heuristic to preserve progress. In elitism without any Novelty Selection, \mathbb{M}_g is sorted by f_{gi} and the best k performing individuals $\hat{\mathbb{M}}_g \subset \mathbb{M}_g$ are preserved and combined with \hat{N}_g to form \mathbb{M}_{g+1} , the population for the next

generation. With Novelty Selection, the k individuals are selected based on a combination of fitness and novelty, as was described in Section 2.5. By default, k is set to half of the population size, which is a popular value in literature. In the same way that mutation and crossover encourage exploration of a search space, elitism allows for the exploitation of promising regions in the search space.

3.3 Loss Function Parameterization

Loss functions are represented by leveraging the TaylorGLO parameterization. This parameterization is defined as a fixed set of continuous values, in contrast to the original GLO parameterization based on trees [9]. TaylorGLO loss functions have several functional advantages over GLO: they are inherently more stable, smooth, and lack discontinuities [10]. Furthermore, because of their simple and compact representation as a continuous vector, TaylorGLO functions can be easily tuned using black-box methods. Specifically in this paper, a third-order TaylorGLO loss function with parameters $\theta_0 \dots \theta_7$, is used:

$$\mathcal{L}(y, \hat{y}) = -\frac{1}{n} \sum_{i=1}^n \left[\theta_2(\hat{y}_i - \theta_1) + \frac{1}{2}\theta_3(\hat{y}_i - \theta_1)^2 + \frac{1}{6}\theta_4(\hat{y}_i - \theta_1)^3 + \theta_5(y_i - \theta_0)(\hat{y}_i - \theta_1) + \frac{1}{2}\theta_6(y_i - \theta_0)(\hat{y}_i - \theta_1)^2 + \frac{1}{2}\theta_7(y_i - \theta_0)^2(\hat{y}_i - \theta_1) \right], \quad (2)$$

where y is the sample’s true label in one-hot form, and \hat{y} is the network’s prediction (i.e., scaled logits). The eight parameters ($\theta_0 \dots \theta_7$) are stored in \mathbf{h}_{gi} and optimized using EPBT.

3.4 Novelty Selection and Pulsation

Novelty Pulsation works by turning Novelty Selection on and off at each pulsation cycle interval p , changing how elite individuals are chosen during Step 3. When Novelty Selection is on, $\hat{\mathbb{M}}_g$ is first increased to include the most fit $m > k$ candidates, and then filtered down to k most novel elites as described in Section 2.5. For the experiments in this paper, the behavior metric used to compute novelty is a binary vector indicating whether the candidate correctly predicts the classes of a randomly chosen N -sized subset of the validation data. This behavior metric encourages evolution to discover models that can perform well generally and not just overfit to a few classes in particular. Preliminary experiments showed that setting $m = 3/2k$, $N = 400$, and $p = 5$ works well and helps protect against premature convergence.

3.5 Population-Based Distillation

To avoid overfitting during evolution, model evaluation includes a variant of knowledge distillation [16]. The main idea in knowledge distillation is to construct training targets as a linear combination of sample labels and the predictions of a better performing teacher model. A model from the previous epoch can be used as a teacher, resulting in a strong regularizing effect [23]. Thus, in PBD, the loss for an individual’s model \mathbf{D}_{gi} is computed as

$$\mathcal{L}(y, \hat{y}) = \mathcal{L}(\hat{\alpha} * \hat{y}_p + (1 - \hat{\alpha}) * y, \hat{y}), \quad (3)$$

where \hat{y}_p are the predictions of the best individual in the previous generation. $\hat{\alpha} = \alpha * (t/T)$ where t is the number of training epochs elapsed, T the total number of training epochs, and α is a scalar between 0 and 1. As in prior work [23], to minimize the effect of inaccurate teacher models in the beginning, $\hat{\alpha}$ is initialized at zero




Figure 2: Experiments on CIFAR-10 with ResNet-32. Each line represents the test classification accuracy (y -axis) of the method over the number of epochs of training (x -axis). All results are averaged over five runs with error bars shown. The top plot is a zoomed-in version of the bottom plot. EPBT outperforms all baselines by a significant margin.

and then increased linearly with the number of training epochs. Thus, PBD regularizes against overfitting to the validation dataset by smoothing the target labels and making for better-behaved gradients, allowing evolution to proceed more reliably.

4 EXPERIMENTAL RESULTS

To show the effectiveness of EPBT, the algorithm was applied to optimize loss functions for two popular image classification datasets: CIFAR-10 and SVHN. Experimental results and comparisons to multiple baselines are presented below.

4.1 CIFAR-10

CIFAR-10 [25] is a widely used image classification dataset consisting of 60,000 natural images in ten classes. The dataset is composed of a training set of 50,000 images and a test set of 10,000 images. For evaluating individuals in EPBT, the training set was split into a separate validation set of 1,250 images and smaller training set with 48,750 images. To control noise during evaluations, the validation dataset was artificially enlarged to 25,000 images through data augmentation. The fitness was calculated by finding the classification accuracy of the trained model on this set. The test accuracies of each individual’s model after evaluation were also recorded for comparison purposes only.

To understand the improvements brought about by EPBT, two baselines were created. The first baseline is a model trained without EPBT: a 32-layer residual network (ResNet-32) with 0.47 million weights that was initialized with the He method [14, 15]. The model was trained using stochastic gradient descent (SGD) for 200 epochs on all 50,000 training images with a batch size of 128, momentum of

0.9, and cross-entropy loss. A fixed learning rate schedule that starts at 0.1 and decays by a factor of 5 at 60, 120, and 160 epochs was used. Input images were normalized to have unit pixel variance and a mean pixel value of zero before training while data augmentation techniques such as random flips, translations, and cutout [6] were applied during training.

The second baseline is a reimplemention of the original PBT algorithm that was used to optimize DNNs on the CIFAR-10 dataset [21]. The training setup was similar to the first baseline but with learning rate as an evolvable hyperparameter. Unlike EPBT, PBT only makes uses of truncation selection, where the weights and loss parameters from the top 25% of the population were copied over to the bottom 25% every generation, and simple mutation, where the hyperparameters were tuned using a mixture of both random resets and multiplicative perturbations of magnitude 1.2. PBT was run for 25 generations, each with eight epochs of training (for a total of 200 epochs), and with a population size of 40. The learning rate was randomly initialized between 0.1 and 0.0001.

The experiments with EPBT were run using a similar training setup as described above. Like the PBT baselines, EPBT was run for 25 generations of eight epochs each and with a population size of 40. Besides the TaylorGLO loss function parameterization, EPBT optimized the hyperparameters for SGD learning rate schedule and momentum as well. The learning rate schedule was based on the one used by the first baseline but with a tunable scaling and decay factor. This search space allows for the exploration of novel schedules but can rediscover the original schedule if necessary. EPBT was configured similarly as the PBT baseline, but with an elitism size of $k = 20$ and with the initial TaylorGLO parameters sampled uniformly between -10 and 10 .

The test accuracies of each baseline and best model in EPBT’s population, averaged over five independent runs with standard error, are shown in Figure 2. EPBT converges rapidly to the highest test accuracy and outperforms all other baselines. Baseline 2 (PBT) results in the worst performance, followed by the Baseline 1, which does not use a population at all. The relatively poor performance of PBT might be because it does not explicitly optimize for learning rate schedules. Previous work has shown that decaying the learning rate at crucial moments during training is important [54]. On the other hand, EPBT’s more advanced parameterization and Novelty Pulsation heuristic allows for more principled, yet diverse learning rate schedules to be discovered.

EPBT can also be scaled up to larger DNN architectures with more weights. In Figure 3, Baseline 1 and EPBT were used to train four different wide residual networks [56] with different number of layers, but with similar number of parameters (11 million). The four architectures, in order of increasing depth, are: WRN-10-12, WRN-16-8, WRN-22-6, and WRN-28-5. EPBT is able to achieve noticeable improvements over Baseline 1 for all of the networks, reaching a better test accuracy at a faster pace. However, the final performance difference between the baseline and EPBT does decrease as the architecture becomes more complicated. Recent work has shown that DNNs are implicitly self-regularizing and thus might benefit less from external regularization as they increase in depth and size [35].

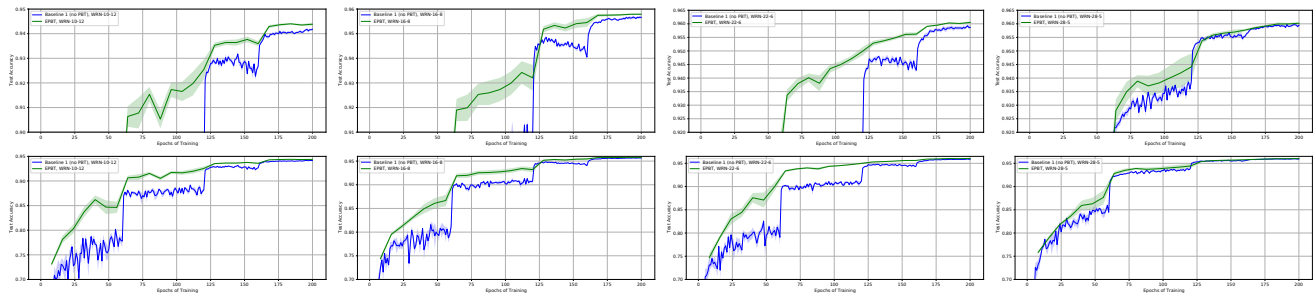


Figure 3: Experiments on CIFAR-10 with WRN-10-12, WRN-16-8, WRN-22-6, and WRN-28-5 (left to right). All results are averaged over five runs with error bars shown. The top plot is a zoomed-in version of the bottom plot. EPBT outperforms the baselines for all each of the architectures.

ALGORITHM	CIFAR-10, ResNet-32	CIFAR-10, WRN-10-12	CIFAR-10, WRN-16-8	CIFAR-10, WRN-22-6	CIFAR-10, WRN-28-5	SVHN, ResNet-32
BASELINE 1 (no PBT)	92.42 (0.21)	94.18 (0.09)	95.66 (0.13)	95.87 (0.18)	95.95 (0.07)	97.81 (0.04)
BASELINE 2 (PBT)	91.53 (0.56)	–	–	–	–	–
EPBT	92.79 (0.15)	94.38 (0.14)	95.79 (0.08)	96.05 (0.09)	96.02 (0.12)	98.08 (0.08)

Table 1: Mean and standard deviation (over five runs) of final test accuracies on the CIFAR-10 and SVHN datasets. EPBT achieves better results (bold) compared to the baselines. Results are reported in percentage.




Figure 4: Experiments on SVHN with ResNet-32. All results are averaged over five runs with error bars shown. The top plot is a zoomed-in version of the bottom plot. EPBT outperforms the baseline, which uses cross-entropy loss to train.

4.2 SVHN

To demonstrate that loss function optimization scales with dataset size, EPBT was applied to SVHN [41], a larger image classification task. This dataset is composed of around 600,000 training images and 26,000 testing images. Following existing practices [19], the dataset was normalized but no data augmentation was used during training. The baseline model was optimized with SGD on the full training set for a total of 40 epochs, with the learning rate decaying from 0.1 by a factor of 10 at 20 and 30 epochs. EPBT was run for 40

generations, each with one epoch of training, and a validation set of 30,000 images was separated for evaluating individuals. Otherwise, the experiment setup was identical to the CIFAR-10 domain.

Figure 4 gives a comparison of EPBT against Baseline 1 in the SVHN domain. As expected, EPBT achieves higher test accuracy than the baseline. Like in the earlier experiments with CIFAR-10 and ResNet-32, both EPBT variants learn faster and converge to a high test accuracy at the end. Interestingly, while the baseline begins to overfit and drop in accuracy at the end of training, EPBT’s regularization mechanisms allow it to avoid this effect and maintain performance.

5 ANALYSIS OF RESULTS

This section analyzes the performance and computational complexity of EPBT, as well as the loss functions and learning rate schedules that EPBT discovered.

5.1 Performance

A summary of the final test accuracies at the end of training for EPBT and the baselines is shown in Table 1. The results show that EPBT achieves the best results for multiple datasets and model architectures. Another noticeable benefit provided by EPBT is the ability to train models to convergence significantly faster than non-population based methods, especially with a limited number of training epochs. This is because multiple models are simultaneously trained with EPBT, each with different loss functions. If progress is made in one of the models, its higher fitness leads to that model’s loss function or weights being shared among the rest of the models, thus lifting their performance as well.

Table 2 details how many epochs of training are required for EPBT to surpass the fully trained performance of the baselines. As expected, EPBT outperforms the baselines on most architectures after training for 80% the total number of epochs. The results are remarkable considering that Baseline 1 is trained on the full training

ALGORITHM	CIFAR-10, ResNET-32	CIFAR-10, WRN-10-12	CIFAR-10, WRN-16-8	CIFAR-10, WRN-22-6	CIFAR-10, WRN-28-5	SVHN, ResNET-32
BASILINE 1 (NO PBT)	168	168	168	176	184	21
BASILINE 2 (PBT)	128	–	–	–	–	–

Table 2: Number of training epochs required for EPBT to exceed the final test accuracy of the baselines. The baselines were trained for 200 epochs in CIFAR-10 and 40 epochs in SVHN. EPBT surpasses most of the baselines at around 80% into the run.

set, while EPBT is not. These experiments thus demonstrate the power of EPBT in not just training better models but doing it faster too.

The experiments also suggest that EPBT’s main components, i.e. the metalearning evolutionary loop, genetic operators, Novelty Selection and Pulsation, TaylorGLO parameterization, and PBD, serve as powerful metalearning and regularization tools when combined. While each component has its own weaknesses, other components can help mitigate them. For example, TaylorGLO is useful for its regularizing effects, but the search space is large and potentially deceptive. However, Novelty Selection and Pulsation can help overcome this deception by maintaining population diversity. Similarly, PBD is a powerful general regularization tool that requires a good teacher model to work. Conveniently, such a model is provided by the elite individuals in EPBT’s population.

5.2 Computational Complexity

Compared to simpler hyperparameter tuning methods that do not interleave training and optimization, EPBT is much more efficient. On the CIFAR-10 dataset, EPBT discovered 40 new loss functions during the first generation and an additional 20 loss functions every subsequent generation. EPBT was run for 25 generations and thus was able to explore up to 520 unique TaylorGLO parameterizations. This process is efficient despite how large the search space is; if grid search is performed at intervals of 1.0, a total of 21^8 (38 billion) unique loss function parameterizations will have to be evaluated.

Furthermore, the computational complexity of EPBT scales linearly with the population size and not with the number of loss functions explored. Loss function evaluation is efficient in EPBT because it is not necessary to retrain the model from scratch whenever a new loss function is discovered; the model’s weights are copied over from an existing model with good performance. If each of the 520 discovered loss functions was used to fully train a model from random initialization, over 100,000 epochs of training would be required, much higher than the 8,000 epochs EPBT needed.

Because EPBT evaluates all the individuals in the population in parallel, the real-time complexity of each generation is not significantly higher than training a single model for the same number of epochs. Furthermore, the amount of time spent in Steps 1 and 2 to generate new individuals is negligible compared to Step 3, where model training occurs. The EPBT experiments in this paper were run on a machine with eight NVIDIA V100 GPUs and utilized several GPU-days worth of compute. When comparing computational cost to EPBT, Baseline 1 used 40 times less compute while Baseline 2 (PBT) required the same amount.

5.3 Loss Functions

Do the loss functions discovered by EPBT stay static or adapt to the current stage of the training process? An analysis of functions

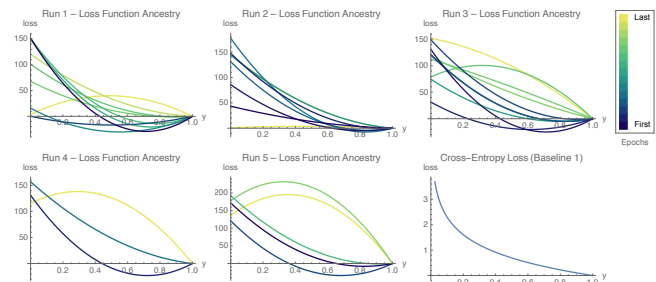


Figure 5: EPBT loss function ancestries for the best candidates across five different runs on CIFAR-10 with ResNet-32. Their shapes are simplified into a 2D binary classification loss [9] for visualization purposes. Cross-entropy loss is shown in the bottom right plot for comparison. Loss functions in the starting generations with fewer training epochs are darker, while functions from later generations with more training are lighter. Some runs have fewer ancestors because elitism allows the same loss function to be reused for multiple generations. Across all runs, there is a temporal pattern in the loss function ancestry. Early loss functions tend to regularize more (indicated by a positive slope at $y = 1$), while later functions encourage more accurate fitting to the ground-truth labels.

discovered by EPBT during an experiment indicates that they do change significantly over the generations.

To characterize how the loss functions adapt with increased training, the ancestries for the final top-performing functions across five separate runs of EPBT (CIFAR-10, ResNet-32) are shown in Figure 5. The cross-entropy loss (as used in Baseline 1) is plotted for comparison as well. Ancestry is determined by tracing the sequence of individuals $M_{0i} \dots M_{ni}$, where $M_{(g-1)i}$ is the parent whose $D_{(g-1)i}$ and $h_{(g-1)i}$ were used to create M_{gi} . The sequence is simplified by removing any duplicate individuals that do not change between generations due to elitism, thus causing some runs to have shorter ancestries.

Because the loss functions are multidimensional, graphing them is not straightforward. However, for visualization purposes, the losses can be simplified into a 2D binary classification modality where $y = 1$ represents a perfect prediction, and $y = 0$ represents a completely incorrect prediction [9]. This approach makes it clear that the loss generally decreases as the predicted labels become more accurate and closer to the ground-truth labels.

There is an interesting trend across all five runs: the loss functions optimized by EPBT are not all monotonically-decreasing. Instead, there are parabolic losses that have a minimum of around 0.7 and rises slightly as y approaches 1. Such concavity is likely a form of regularization that prevents the network from overfitting




Figure 6: Visualization of how the learning rate and momentum of the best individual in the population changes during an EPBT run. The right plot is a log-scale version of the left plot. A cyclic and decaying learning rate pattern is present, while no such pattern exists for momentum.

to the training data by penalizing low-entropy prediction distributions centered around $y = 1$. Similar behavior was observed when training using GLO [9].

The plots also show that the loss functions change shape as training progresses. As the number of epochs increases, the slope near $y = 1$ becomes increasingly positive, suggesting less regularization would occur. This result is consistent with recent research that suggests regularization is most important during a critical period early in the training process [7]. If regularization is reduced or removed after this critical period, generalization sometimes may even improve. In EPBT, this principle was discovered and optimized without any prior knowledge as part of the metalearning process. EPBT thus provides an automatic way for exploring metaknowledge that could be difficult to come upon manually.

Since different stages of EPBT utilize different types of loss functions, it is possible that a single static loss might not be optimal for the entire training process. Furthermore, loss functions that change makes sense considering that the learning dynamics for some DNNs are non-stationary or unstable [21]. For example, adaptive losses might improve the training of generative adversarial networks [11, 43].

5.4 Learning Rate Schedules

In addition to the TaylorGLO parameters, EPBT also optimized SGD hyperparameters such as learning rate schedule and momentum. Figure 6 shows how these parameters of the best individual in the population changes over the course of an EPBT run (CIFAR-10, ResNet-32). While momentum remains mostly the same with some occasional dips and bumps, there is a clear downward, decaying trend for learning rate. The discovered learning rate schedule shares several similarities with hand designed schedules: both use high learning rates early in training for rapid learning but lower learning rates later to fine tune the model weights. More interestingly, there appears to be several cycles in the EPBT optimized schedule where the learning rate repeatedly goes up and down. This might be due the beneficial effects of cyclic learning rates in helping SGD escape from saddle points in the loss landscape during training [48].

6 DISCUSSION AND FUTURE WORK

When viewed from the EA perspective, EPBT is a more complex variant of PBT [21]. Mutation corresponds to the *explore* step and elitism corresponds to the *exploit* step in PBT. Besides Novelty Pulsation, EPBT improves upon PBT in two major ways. First, EPBT makes use of uniform Gaussian mutation (compared to the deterministic mutation in PBT) and uniform crossover. These biologically

inspired heuristics allow PBT to scale better to higher dimensions. In particular, the crossover operator plays an important role in discovering good global solutions in large search spaces [8, 53]. Second, EPBT utilizes tournament selection, a heuristic that helps prevent premature convergence to a local optimum [47].

One possible direction of future work is to allow the loss function parameterization to take the current state of training as input, potentially resulting in more refined training and better performance. Alternatively, domain information could be taken into account, allowing the learned loss function schedules to be more easily transferred to different tasks. EPBT could also be extended to include neural architecture search to jointly optimize network structure and loss functions. A more detailed study on how network architecture impacts EPBT performance and what components are necessary for training a particular DNN architectures could also lead to better models in the future. Another compelling direction would be to understand the synergies of the various components of EPBT. Ablation studies can be used to characterize the contribution that each component provides. Lastly, more experimental runs can be performed to improve statistical analysis of EPBT and the baselines.

7 CONCLUSION

This paper presents EPBT, an evolutionary algorithm for regularization metalearning. EPBT first improves upon PBT by introducing more advanced genetic operators. It then focuses it on regularization by evolving TaylorGLO loss functions. The deceptive interactions of weight and loss adaptation require more diversity, which is achieved through Novelty Pulsation, and more careful avoidance of overfitting, which is achieved through Population-Based Distillation. On the CIFAR-10 and SVHN image classification benchmarks with several ResNet and Wide Resnet architectures, EPBT achieved faster and better model training. An analysis of the optimized loss functions suggests that these advantages stem from discovering strong regularization automatically. Furthermore, an adaptive loss function schedule naturally emerges as a likely key to achieving such performance. EPBT thus forms a practical method for regularization metalearning in deep networks.

REFERENCES

- [1] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473* (2014).
- [2] Yogesh Balaji, Swami Sankaranarayanan, and Rama Chellappa. 2018. Metareg: Towards domain generalization using meta-regularization. *Advances in Neural Information Processing Systems* 31 (2018), 998–1008.
- [3] Wolfgang Banzhaf, Peter Nordin, Robert E Keller, and Frank D Francone. 1998. *Genetic programming: an introduction*. Vol. 1. Morgan Kaufmann San Francisco.
- [4] Anna Sergeevna Bosman, Andries Engelbrecht, and Mardé Helbig. 2019. Visualising Basins of Attraction for the Cross-Entropy and the Squared Error Neural Network Loss Functions. *arXiv preprint arXiv:1901.02302* (2019).
- [5] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. 2018. Autoaugment: Learning augmentation policies from data. *arXiv preprint arXiv:1805.09501* (2018).
- [6] Terrance DeVries and Graham W Taylor. 2017. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552* (2017).
- [7] Aditya Sharad Golatkar, Alessandro Achille, and Stefano Soatto. 2019. Time Matters in Regularizing Deep Networks: Weight Decay and Data Augmentation Affect Early Learning Dynamics, Matter Little Near Convergence. In *Advances in Neural Information Processing Systems* 32. 10677–10687.

- [8] David E Goldberg and John H Holland. 1988. Genetic algorithms and machine learning. *Machine learning* 3, 2 (1988), 95–99.
- [9] Santiago Gonzalez and Risto Miikkulainen. 2019. Improved Training Speed, Accuracy, and Data Utilization Through Loss Function Optimization. *arXiv preprint arXiv:1905.11528* (2019).
- [10] Santiago Gonzalez and Risto Miikkulainen. 2020. Evolving Loss Functions With Multivariate Taylor Polynomial Parameterizations. *arXiv preprint arXiv:2002.00059* (2020).
- [11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in Neural Information Processing Systems*. 2672–2680.
- [12] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. 2014. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572* (2014).
- [13] Nikolaus Hansen and Andreas Ostermeier. 1996. Adapting arbitrary normal mutation distributions in evolution strategies: The covariance matrix adaptation. In *Proceedings of IEEE international conference on evolutionary computation*. IEEE, 312–317.
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*. 1026–1034.
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016), 770–778.
- [16] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531* (2015).
- [17] Daniel Ho, Eric Liang, Ion Stoica, Pieter Abbeel, and Xi Chen. 2019. Population based augmentation: Efficient learning of augmentation policy schedules. *arXiv preprint arXiv:1905.05393* (2019).
- [18] Rein Houthooft, Yuhua Chen, Phillip Isola, Bradley Stadie, Filip Wolski, OpenAI Jonathan Ho, and Pieter Abbeel. 2018. Evolved policy gradients. In *Advances in Neural Information Processing Systems*. 5400–5409.
- [19] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. 2017. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 4700–4708.
- [20] Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*. PMLR, 448–456.
- [21] Max Jaderberg, Valentin Dalibard, Simon Osindero, Wojciech M Czarnecki, Jeff Donahue, Ali Razavi, Oriol Vinyals, Tim Green, Iain Dunning, Karen Simonyan, et al. 2017. Population based training of neural networks. *arXiv preprint arXiv:1711.09846* (2017).
- [22] Katarzyna Janocha and Wojciech Marian Czarnecki. 2017. On loss functions for deep neural networks in classification. *arXiv preprint arXiv:1702.05659* (2017).
- [23] Kyungyul Kim, ByeongMoon Ji, Doyoung Yoon, and Sangheum Hwang. 2020. Self-knowledge distillation: A simple way for better generalization. *arXiv preprint arXiv:2006.12000* (2020).
- [24] Aaron Klein, Stefan Falkner, Simon Bartels, Philipp Hennig, and Frank Hutter. 2017. Fast bayesian optimization of machine learning hyperparameters on large datasets. In *Artificial Intelligence and Statistics*. PMLR, 528–536.
- [25] Alex Krizhevsky and Geoffrey Hinton. 2009. Learning multiple layers of features from tiny images.
- [26] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems* 25, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger (Eds.). Curran Associates, Inc., 1097–1105.
- [27] Jan Kukačka, Vladimir Golkov, and Daniel Cremers. 2017. Regularization for deep learning: A taxonomy. *arXiv preprint arXiv:1710.10686* (2017).
- [28] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *Nature* 521, 7553 (2015), 436.
- [29] Joel Lehman and Kenneth O Stanley. 2008. Exploiting open-endedness to solve problems through the search for novelty. In *ALIFE*. Citeseer, 329–336.
- [30] Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar. 2017. Hyperband: A novel bandit-based approach to hyperparameter optimization. *The Journal of Machine Learning Research* 18, 1 (2017), 6765–6816.
- [31] Hanxiao Liu, Karen Simonyan, and Yiming Yang. 2018. Darts: Differentiable architecture search. *arXiv preprint arXiv:1806.09055* (2018).
- [32] Ilya Loshchilov and Frank Hutter. 2016. CMA-ES for hyperparameter optimization of deep neural networks. *arXiv preprint arXiv:1604.07269* (2016).
- [33] Zhichao Lu, Ian Whalen, Vishnu Boddeti, Yashesh Dhebar, Kalyanmoy Deb, Erik Goodman, and Wolfgang Banzhaf. 2018. NSGA-NET: a multi-objective genetic algorithm for neural architecture search. *arXiv preprint arXiv:1810.03522* (2018).
- [34] Dougal Maclaurin, David Duvenaud, and Ryan Adams. 2015. Gradient-based hyperparameter optimization through reversible learning. In *International Conference on Machine Learning*. 2113–2122.
- [35] Charles H Martin and Michael W Mahoney. 2018. Implicit self-regularization in deep neural networks: Evidence from random matrix theory and implications for learning. *arXiv preprint arXiv:1810.01075* (2018).
- [36] Risto Miikkulainen, Jason Liang, Elliot Meyerson, Aditya Rawal, Daniel Fink, Olivier Francon, Bala Raju, Hormoz Shahrzad, Arshak Navruzyan, Nigel Duffy, et al. 2019. Evolving deep neural networks. In *Artificial Intelligence in the Age of Neural Networks and Brain Computing*. Elsevier, 293–312.
- [37] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518, 7540 (2015), 529.
- [38] John Moody, Stephen Hanson, Anders Krogh, and John A Hertz. 1995. A simple weight decay can improve generalization. *Advances in neural information processing systems* 4, 1995 (1995), 950–957.
- [39] Jean-Baptiste Mouret and Jeff Clune. 2015. Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909* (2015).
- [40] Rafael Müller, Simon Kornblith, and Geoffrey Hinton. 2019. When does label smoothing help? *arXiv preprint arXiv:1906.02629* (2019).
- [41] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. 2011. Reading digits in natural images with unsupervised feature learning.
- [42] Hieu Pham, Melody Y Guan, Barret Zoph, Quoc V Le, and Jeff Dean. 2018. Efficient neural architecture search via parameter sharing. *arXiv preprint arXiv:1802.03268* (2018).
- [43] Alec Radford, Luke Metz, and Soumith Chintala. 2015. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434* (2015).
- [44] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. 2019. Regularized evolution for image classifier architecture search. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 4780–4789.
- [45] Hormoz Shahrzad, Daniel Fink, and Risto Miikkulainen. 2018. Enhanced Optimization with Composite Objectives and Novelty Selection. 616–622. https://doi.org/10.1162/isaal_a_00113
- [46] Hormoz Shahrzad, Babak Hodjat, Camille Dollé, Andrei Denissov, Simon Lau, Donn Goodhew, Justin Dyer, and Risto Miikkulainen. 2020. Enhanced Optimization with Composite Objectives and Novelty Pulsation. 275–293. https://doi.org/10.1007/978-3-030-39958-0_14
- [47] Anupriya Shukla, Hari Mohan Pandey, and Deepti Mehrotra. 2015. Comparative review of selection techniques in genetic algorithm. In *2015 International Conference on Futuristic Trends on Computational Analysis and Knowledge Management (ABLAZE)*. IEEE, 515–519.
- [48] Leslie N Smith. 2017. Cyclical learning rates for training neural networks. In *2017 IEEE winter conference on applications of computer vision (WACV)*. IEEE, 464–472.
- [49] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. 2012. Practical bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems*. 2951–2959.
- [50] Jasper Snoek, Oren Rippel, Kevin Swersky, Ryan Kiros, Nadathur Satish, Narayanan Sundaram, Mostofa Patwary, Mr Prabhat, and Ryan Adams. 2015. Scalable bayesian optimization using deep neural networks. In *International conference on machine learning*. 2171–2180.
- [51] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* 15, 1 (2014), 1929–1958.
- [52] Kenneth O. Stanley, Jeff Clune, Joel Lehman, and Risto Miikkulainen. 2019. Designing neural networks through neuroevolution. *Nature Machine Intelligence* 1, 1 (2019), 24–35. <https://doi.org/10.1038/s42256-018-0006-z>
- [53] Darrell Whitley. 1994. A genetic algorithm tutorial. *Statistics and computing* 4, 2 (1994), 65–85.
- [54] Kaichao You, Mingsheng Long, Jianmin Wang, and Michael I Jordan. 2019. How does learning rate decay help modern neural networks? *arXiv preprint arXiv:1908.01878* (2019).
- [55] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. 2019. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 6023–6032.
- [56] Sergey Zagoruyko and Nikos Komodakis. 2016. Wide residual networks. *arXiv preprint arXiv:1605.07146* (2016).
- [57] Barret Zoph and Quoc V Le. 2016. Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578* (2016).
- [58] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. 2018. Learning transferable architectures for scalable image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 8697–8710.