
Gradient-based Hyperparameter Optimization Over Long Horizons

Paul Micaelli

University of Edinburgh
{paul.micaelli}@ed.ac.uk

Amos Storkey

University of Edinburgh
{a.storkey}@ed.ac.uk

Abstract

Gradient-based hyperparameter optimization has earned a widespread popularity in the context of few-shot meta-learning, but remains broadly impractical for tasks with long horizons (many gradient steps), due to memory scaling and gradient degradation issues. A common workaround is to learn hyperparameters online, but this introduces greediness which comes with a significant performance drop. We propose forward-mode differentiation with sharing (FDS), a simple and efficient algorithm which tackles memory scaling issues with forward-mode differentiation, and gradient degradation issues by sharing hyperparameters that are contiguous in time. We provide theoretical guarantees about the noise reduction properties of our algorithm, and demonstrate its efficiency empirically by differentiating through $\sim 10^4$ gradient steps of unrolled optimization. We consider large hyperparameter search ranges on CIFAR-10 where we significantly outperform greedy gradient-based alternatives, while achieving $\times 20$ speedups compared to the state-of-the-art black-box methods. Code is available at: <https://github.com/po1o5/FDS>

1 Introduction

Deep neural networks have shown tremendous success on a wide range of applications, including classification [1], generative models [2], natural language processing [3] and speech recognition [4]. This success is in part due to effective optimizers such as SGD with momentum or Adam [5], which require carefully tuned hyperparameters for each application. In recent years, a long list of heuristics to tune such hyperparameters has been compiled by the deep learning community, including things like: how to best decay the learning rate [6], how to scale hyperparameters with the budget available [7], and how to scale learning rate with batch size [8]. Unfortunately these heuristics are often dataset specific and architecture dependent [9]. They also don't apply well to new optimizers [10], or new tools, like batch normalization which allows for larger learning rates [11].

With so many ways to choose hyperparameters, the deep learning community is at risk of adopting models based on how much effort went into tuning them, rather than their methodological insight. The field of hyperparameter optimization (HPO) aims to find hyperparameters that provide a good generalization performance automatically. Unfortunately, existing tools are rather unpopular for deep learning, largely owing to their computational cost. The method developed here is a gradient-based HPO approach; that is, it calculates hypergradients, i.e. the gradient of some validation loss with respect to each hyperparameter. Gradient-based HPO should be more efficient than black-box methods as the dimensionality of the hyperparameter space increases, since it is able to utilize gradient information rather than rely on trial and error. In practice however, learning hyperparameters with gradients has only been popular in few-shot learning tasks where the horizon is short, i.e. where the underlying task is solved with a few gradient steps. This is because long horizons cause hypergradient degradation, and incur a memory cost that makes reverse-mode differentiation prohibitive.

Greedy gradient-based methods alleviate both of these issues by calculating local hypergradients based on intermediate validation losses. Unfortunately, this introduces some bias [15] and results in a significant performance drop, which we are able to quantify in this work. We make use of forward-mode differentiation, which has been shown to offer a memory cost constant with horizon size. However, previous forward-mode methods don’t address gradient degradation explicitly and are thus limited to the greedy setting [16, 17].

We introduce FDS (Forward-mode Differentiation with hyperparameter Sharing), which to the best of our knowledge demonstrates for the first time that hyperparameters can be differentiated non-greedily over long horizons. Specifically, we make the following contributions: (1) we propose to share hyperparameters through time, both motivating it theoretically and with various experiments, (2) we combine the above in a forward-mode differentiation algorithm, and (3) we show that our method can significantly outperform various HPO algorithms, for instance when learning the hyperparameters of the SGD-momentum optimizer.

2 Related Work

There are many ways to perform hyperparameter optimization (HPO) [18], including Bayesian optimization (BO) [19], reinforcement learning [20], evolutionary algorithms [21] and gradient-based methods [22]. The state-of-the-art in HPO depends on the problem setting, but black-box methods like Hyperband (HB) [12], and its combination with BO into a method called BOHB [13] have been the most popular. Modern work in meta-learning deals with various forms of gradient-based HPO [23], but usually focuses on the few-shot regime [24] where horizons are conveniently short (~ 10 steps) while we focus on long horizons ($\sim 10^4$ steps).

Gradient-based HPO. Using the gradient of some validation loss with respect to the hyperparameters is typically the preferred choice when the underlying optimization is differentiable. This is a type of constrained optimization [25] which stems from earlier work on backpropagation through time [26] and real-time recurrent learning [27]. Unfortunately, differentiating optimization is an expensive procedure in both time and memory, and most proposed methods are limited to small models and toy datasets [28–30]. Efforts to make the problem more tractable include optimization shortcuts [31], truncation [32] and implicit gradients [33–35]. Truncation can be combined with our approach but produces biased gradients [36], while implicit differentiation is only applicable to hyperparameters that define the training loss (e.g. augmentation) but not to hyperparameters that define how the training loss is minimized (e.g. optimizer hyperparameters). Forward-mode differentiation [27] boasts a memory cost constant with horizon size, but gradient degradation has restricted its use to the greedy setting [16, 17].

Greedy Methods. One trick that prevents gradient degradation and significantly reduces compute and memory cost is to solve the bilevel optimization greedily. This has become the default trick in various long-horizon problems, including HPO over optimizers [14, 16, 17, 37], architecture search [38], dataset distillation [39] or curriculum learning [40]. Greediness refers to finding the best hyperparameters locally rather than globally. In practice, it involves splitting the inner optimization problem into smaller chunks (often just one batch), and solving for hyperparameters over these smaller horizons instead; often in an online fashion. We formalize greediness in Section 3.2. In this paper we expand upon previous observations [15] and take the view that greediness fundamentally solves

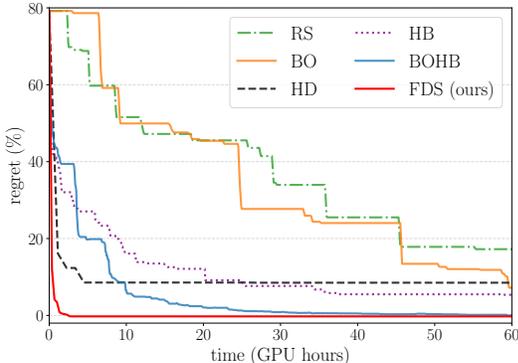


Figure 1: Performance of the most popular hyperparameter optimization methods when learning the learning rate, momentum and weight decay on 50 epochs of CIFAR-10, for a WideResNet of 16 layers. Non-greedy methods (RS, BO, HB [12], BOHB [13]) solve for global hyperparameters but rely on trial-and-error which makes them slow. Greedy gradient-based methods (e.g. HD [14]) are faster but solve for local hyperparameters which makes them suboptimal. Our method combines the strengths of these two paradigms and outperforms the next best method while converging 20 times faster. Each curve is the average of 8 seeds.

for the wrong objective. Instead, the focus of our paper is to extend forward-mode differentiation methods to the non-greedy setting.

Gradient Degradation. Gradient degradation of some scalar w.r.t a parameter is a broad issue that arises when that parameter influences the scalar in a chaotic fashion, such as through long chains of nonlinear mappings. This manifests itself in HPO as vanishing or exploding hypergradients, due to low or high curvature components of the validation loss surface. This leads to hypergradients that have a large variance, making differentiation through long horizons impractical. This is usually observed in the context of recurrent neural networks [41, 42], but also in reinforcement learning [43] and HPO [29]. Solutions like LSTMs [44] and gradient clipping [45] have been proposed, but are respectively inapplicable and insufficient to long-horizon HPO. Variational optimization [36] and preconditioning warp-layers [46] can mitigate gradient degradation, but these methods are expensive in memory and therefore are limited to small architectures and/or a few hundred inner steps. In comparison, we differentiate over $\sim 10^4$ inner steps for WideResNets [47].

3 Background

3.1 Problem statement

Consider a neural network with weights θ , trained to minimize a loss $\mathcal{L}_{\text{train}}$ over a dataset $\mathcal{D}_{\text{train}}$. This is done by taking T steps with a gradient-based optimizer Φ , which uses a collection of hyperparameters $\lambda \in \mathbb{R}^T$. For clarity of notation, consider that Φ uses one hyperparameter per step, written $\lambda_{[t]}$, where indices are shown in brackets to differentiate them from a variable evaluated at time t . We can explicitly write out this optimizer as $\Phi : \theta_{t+1} = \Phi(\mathcal{L}_{\text{train}}(\theta_t(\lambda_{[1:t]}), \mathcal{D}_{\text{train}}), \lambda_{[t+1]}) \forall t \in \{1, 2, \dots, T\}$. Note that θ_t is a function of $\lambda_{[1:t]}$, and it follows that $\theta_T = \theta_T(\lambda_{[1:T]}) = \theta_T(\lambda)$.

We would like to find the hyperparameters λ^* such that the result, at time T , of the gradient process minimizing objective $\mathcal{L}_{\text{train}}$, also minimizes some generalization loss \mathcal{L}_{val} on a validation dataset \mathcal{D}_{val} . This can be cast as the following constrained optimization:

$$\lambda^* = \arg \min_{\lambda} \mathcal{L}_{\text{val}}(\theta_T(\lambda), \mathcal{D}_{\text{val}}) \tag{1}$$

$$\text{subject to } \theta_{t+1} = \Phi(\mathcal{L}_{\text{train}}(\theta_t(\lambda_{[1:t]}), \mathcal{D}_{\text{train}}), \lambda_{[t+1]}) \tag{2}$$

The inner loop in Eq 2 expresses a constraint on the outer loop in Eq 1. In gradient-based HPO, our task is to compute the hypergradient $d\mathcal{L}_{\text{val}}/d\lambda$ and update λ accordingly. Note that some methods require $\theta_T = \arg \min_{\theta} \mathcal{L}_{\text{train}}(\theta, \mathcal{D}_{\text{train}}, \lambda)$ with $\mathcal{L}_{\text{train}}$ being a function of λ explicitly [33–35], in which case the above becomes a bilevel optimization [48]. Our algorithm waives these requirements.

3.2 Greediness

Let H be the horizon, which corresponds to the number of gradient steps taken in the inner loop (to optimize θ) before one gradient step is taken in the outer loop (to optimize λ). When solving Eq 1 non-greedily we have $H = T$. However, most modern approaches are greedy [14, 16, 37–39], in that they rephrase the above problem into a sequence of several independent problems of smaller horizons, where $\lambda_{[t:t+H]}^*$ is learned in the outer loop subject to an inner loop optimization from θ_t to θ_{t+H} with $H \ll T$. Online methods such as Hypergradient Descent (HD) [14] are an example of greedy differentiation, where $H = 1$ and $\lambda_{[t]}$ is updated at time t . Instead, non-greedy algorithms like FDS only update $\lambda_{[t]}$ at time T .

Greediness mitigates gradient degradation and lowers the memory cost of backpropagation through time (BPTT). However, greedy methods do not solve for the actual objective in Eq 1, but for some proxy objective which minimizes intermediate validation losses evaluated at $\theta_H, \theta_{2H}, \dots$. The degree to which these two objectives overlap is unknown and problem dependent. In our experiments, we found that finding competitive hyperparameters with greedy methods often revolves around tricks like online learning with a very low outer learning rate combined with hand-tuned initial hyperparameter values, to manually prevent convergence to small values. But solving the greedy objective correctly leads to poor solutions, a special case of which was previously described as the “short-horizon bias” [15]. In FDS, we learn global hyperparameters in a non-greedy fashion, which means that the hypergradient $d\mathcal{L}_{\text{val}}/d\lambda$ is only ever calculated at θ_T .

3.3 Forward-mode differentiation of modern optimizers

The vast majority of meta-learning applications use reverse-mode differentiation in the inner optimization problem (Eq 2) to optimize θ . However, the memory cost of BPTT, namely using reverse-mode differentiation for the outer optimization (Eq 1) is $\mathcal{O}(FH)$, where F is the memory used by one forward pass through the network (weights plus activations). In the non-greedy setting where $H = T$, BPTT is extremely limiting: for large networks, only $T \sim 10$ gradient steps could be solved with modern GPUs, while problems like CIFAR-10 require $T \sim 10^4$. Instead, we make use of forward-mode differentiation, which scales in memory as $\mathcal{O}(DN)$, where D is the number of weights and N is the number of learnable hyperparameters. The additional scaling with N is a limitation if we learn one hyperparameter per inner step ($N = T$). Sharing hyperparameters (section 4.1) mitigates gradient degradation, but also conveniently allows for smaller values of N .

For clarity of notation, we consider forward-mode hypergradients for the general case of using one hyperparameter per step, i.e. $\lambda \in \mathbb{R}^T$. First, we use the chain rule to write $d\mathcal{L}_{\text{val}}/d\lambda = (\partial\mathcal{L}_{\text{val}}/\partial\theta_T)(d\theta_T/d\lambda)$ where the direct gradient has been dropped since $\partial\mathcal{L}_{\text{val}}/\partial\lambda = 0$ for optimizer hyperparameters. The first term on the RHS is trivial and can be obtained with reverse-mode differentiation as usual. The second term is more problematic because $\theta_T = \theta_T(\theta_{T-1}(\theta_{T-2}(\dots), \lambda_{[T-1]}), \lambda_{[T]})$. We use the chain rule again to calculate this term recursively:

$$\frac{d\theta_t}{d\lambda} = \frac{\partial\theta_t}{\partial\theta_{t-1}} \Big|_{\lambda} \frac{d\theta_{t-1}}{d\lambda} + \frac{\partial\theta_t}{\partial\lambda} \Big|_{\theta_{t-1}} \quad \text{which we write as } \mathbf{Z}_t = \mathbf{A}_t \mathbf{Z}_{t-1} + \mathbf{B}_t \quad (3)$$

where $\theta_t \in \mathbb{R}^D$, $\lambda \in \mathbb{R}^T$, $\mathbf{Z}_t \in \mathbb{R}^{D \times T}$, $\mathbf{A}_t \in \mathbb{R}^{D \times D}$ and $\mathbf{B}_t \in \mathbb{R}^{D \times T}$. Note that the columns of \mathbf{Z} at step t are zeros for indices $t+1, t+2, \dots, T$, since the hyperparameters at those steps haven't been used yet.

The expressions for \mathbf{A}_t and \mathbf{B}_t depend on the specific hyperparameters we are differentiating. In this work, we consider the most popular optimizer, namely SGD with momentum and weight decay. To the best of our knowledge, previous work focuses on simpler versions of this optimizer, usually by removing momentum and weight decay, and only learns the learning rate, greedily. We use Pytorch's update rule for SGD [49], namely $\theta_t = \Phi(\theta_{t-1}) = \theta_{t-1} - \alpha_t \mathbf{v}_t$ with learning rate α_t , momentum β_t , weight decay ξ_t and velocity $\mathbf{v}_t = \beta_t \mathbf{v}_{t-1} + (\partial\mathcal{L}_{\text{train}}/\partial\theta_{t-1}) + \xi_t \theta_{t-1}$. Consider the case when we learn the learning rate schedule, namely $\lambda = \alpha$. If we use the update rule without momentum [17], \mathbf{B}_t is conveniently sparse: it is a $D \times T$ matrix that only has one non-zero column at index t corresponding to $\partial\theta_t/\partial\alpha_t$. However, we include terms like momentum and therefore the velocity depends on the hyperparameters of previous steps. In that case, a further recursive term $\mathbf{C}_t = (\partial\mathbf{v}_t/\partial\lambda)$ must be considered to get exact hypergradients. Putting it together (see Appendix A) we obtain:

$$\begin{cases} \mathbf{A}_t^\alpha = \mathbf{1} - \alpha_t \left(\frac{\partial^2 \mathcal{L}_{\text{train}}}{\partial\theta_{t-1}^2} + \xi_t \mathbf{1} \right) \\ \mathbf{B}_t^\alpha = -\beta_t \alpha_t \mathbf{C}_{t-1}^\alpha - \delta_t^\otimes \left(\beta_t \mathbf{v}_{t-1} + \frac{\partial \mathcal{L}_{\text{train}}}{\partial\theta_{t-1}} + \xi_t \theta_{t-1} \right) \\ \mathbf{C}_t^\alpha = \beta_t \mathbf{C}_{t-1}^\alpha + \left(\xi_t \mathbf{1} + \frac{\partial^2 \mathcal{L}_{\text{train}}}{\partial\theta_{t-1}^2} \right) \mathbf{Z}_{t-1}^\alpha \end{cases} \quad (4)$$

where $\mathbf{1}$ is a $D \times D$ identity matrix, and $\delta_t^\otimes(\mathbf{q})$ turns a vector \mathbf{q} of size D into a matrix of size $D \times T$, whose t -th column is set to \mathbf{q} and other columns to $\mathbf{0}$ s. While the matrices in Eq 4 are updated online, the hyperparameters aren't. This is because in the non-greedy setting we don't have access to the hypergradients until we have computed \mathbf{Z}_T^α . A similar technique can be applied to momentum and weight decay to get \mathbf{Z}_T^β and \mathbf{Z}_T^ξ (see Appendix A). All hypergradient derivations in this paper were checked with finite differences. Note that we focus on learning the hyperparameters of SGD with momentum because it is the most common optimizer in the deep learning literature. However, just like reverse-mode differentiation, forward-mode differentiation can be applied to any differentiable hyperparameter, albeit with appropriate \mathbf{A}_t and \mathbf{B}_t matrices.

4 Non-greedy Differentiation Over Long Horizons

4.1 Hyperparameter sharing: trading off noise reduction with bias increase

The main challenge of doing non-greedy meta-learning over long horizons is gradient degradation. In HPO this arises because a small change in $\mathcal{L}_{\text{train}}(\theta_t)$ can cascade forward into a completely different θ_T , resulting in large fluctuations of the hypergradients. This noise makes the generalization loss hard to minimize (Eq 1). We find that the two main causes for this noise are the ordering of the training minibatches, and the weight initialization θ_0 . Ideally, the hyperparameters we learn should be agnostic to both of these factors, and so we would like to average out their effect on hypergradients.

Ensemble averaging The most obvious way to address the above is to compute all hypergradients across several random seeds, where each seed corresponds to a different dataset ordering and weight initialization. We could then obtain an average hypergradient μ_t for each inner step t . Here, $\mu \in \mathbb{R}^T$ is often called an ensemble average in statistical mechanics. The issue with ensemble averaging in our setting is its computational and memory cost, since each random seed requires differentiating through T unrolled inner steps for T hyperparameters. This makes both reverse-mode and forward-mode differentiation intractable. We consider the ensemble average as optimal in our analysis, which allows us to derive an expression for the mean square error between our hypergradients estimate and μ .

Time averaging In FDS, we use the long horizon to our advantage and propose to do time averaging instead of ensemble averaging, i.e. *we average out hypergradients across the inner loop (Eq 2) rather than the outer loop (Eq 1)*. More specifically, we sum hypergradients from W neighbouring time steps in the inner loop, which is exactly equivalent to sharing one hyperparameter over all these steps. The average is then obtained trivially by dividing by W . For instance, when learning the learning rate schedule α for $H = 10^4$ inner steps, we can learn $\alpha \in \mathbb{R}^{10}$ where each learning rate is used for a window of $W = 10^3$ contiguous gradient steps. A stochastic system where the time average is equivalent to the ensemble average is called *ergodic* [50] and has been the subject of much research in thermodynamics [51] and finance [52]. In our case, hypergradients aren't generally ergodic, and so using a single average hypergradient for W contiguous steps can introduce a bias. Informally, time averaging contiguous hypergradients leads to both noise reduction and bias increase, and we flesh out the nature of this trade-off in Theorem 4.1.

Theorem 4.1. Let each time step $t \in \{1, 2, \dots, T\}$ have a corresponding hyperparameter λ_t and non-greedy hypergradient $g_t = \partial L_{\text{val}}(\theta_T)/\partial \lambda_t$. Each g_t is viewed as a random variable due to the sampling process of the weight initialization θ_0 and the inner loop minibatch selection. Let $\mathbf{g} = [g_1, g_2, \dots, g_T]$ be sufficiently well approximated by a Gaussian distribution $\mathbf{g} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, with mean $\boldsymbol{\mu} = [\mu_1, \mu_2, \dots, \mu_T]$ and covariance matrix $\boldsymbol{\Sigma}$, where $\boldsymbol{\mu}$ corresponds to the optimal hypergradients. Assume that the changes in the values of $\boldsymbol{\mu}$ over time are bounded, i.e. $\boldsymbol{\mu}$ can be written as the ϵ -Lipschitz function $\mu_{t+1} = \mu_t + \epsilon_t$, where $\epsilon_t \in [-\epsilon, \epsilon]$. Finally, let $c \in [0, 1]$ denote the maximum absolute correlation between the values of \mathbf{g} , i.e. $c \geq |\Sigma_{tt'}|/\sqrt{\Sigma_{tt}\Sigma_{t't'}} \forall t \neq t'$. Then, we show that the mean squared error of the hypergradients with respect to $\boldsymbol{\mu}$ when sharing W contiguous hyperparameters has an upper bound:

$$\text{MSE}_W \leq \frac{(1 + c(W - 1))}{W} \text{MSE}_1 + \epsilon^2 \frac{(W^2 - 1)}{12} \quad (5)$$

$$\text{where } \text{MSE}_1 = \frac{1}{T} \sum_t \Sigma_{tt} \quad (6)$$

and so for sufficiently small ϵ and c we have with certainty, $\text{MSE}_W < \text{MSE}_1$ for some $W > 1$, where MSE_1 is the hypergradient error without any sharing (See Appendix B for a proof).

Discussion Theorem 4.1 demonstrates how sharing W contiguous hyperparameters has two effects: 1) it reduces the hypergradient noise by a factor $\mathcal{O}(W/(1 + cW))$ due to the averaging of noisy hypergradients, and 2) it increases the error up to an amount $\mathcal{O}(\epsilon^2 W^2)$ due to an induced bias. Intuitively, averaging contiguous hypergradients maximally reduces noise when they aren't correlated ($c = 0$), and minimally increases bias when they are drawn from distributions of similar means ($\epsilon = 0$). In the simpler case where hypergradients are iid and have the same variance at each step, namely $\mathbf{g} \sim \mathcal{N}(\boldsymbol{\mu}, \sigma^2 \mathbf{1})$, the expressions above become $\text{MSE}_1 = \sigma^2$ and $\text{MSE}_W \leq \text{MSE}_1/W + \epsilon^2(W^2 - 1)/12$. Note that in all cases, the upper bound on MSE_W has a single minimum W^* corresponding to the optimal trade-off between noise reduction and bias increase.

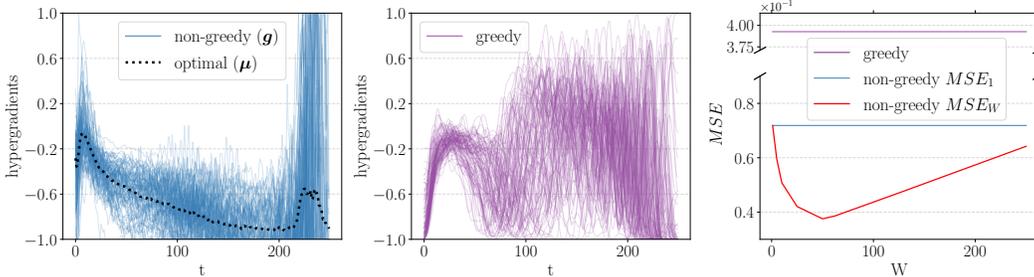


Figure 2: Hypergradients of α on SVHN for 100 seeds in the non-greedy (left) and greedy (middle) setting. The mean squared error is also shown (right), and is calculated with respect to the optimal hypergradient (μ), i.e. the ensemble average of non-greedy hypergradients (dotted line in left figure). We can see that sharing hyperparameters over W steps lowers the MSE even for the small value of $T = 250$ used here. The best trade-off between noise reduction and bias increase is $W = 50$.

4.2 The FDS algorithm

As it is presented in Section 3.3, forward-mode differentiation would still have a memory cost that scales as $\mathcal{O}(DT)$ since we are learning one hyperparameter per step. However, addressing gradient degradation by sharing hyperparameters also conveniently reduces that memory cost by a factor W down to $\mathcal{O}(DN)$, where $N = T/W$ is the number of unique hyperparameter values we learn. This is because we can safely average hypergradients without calculating them individually, by reusing the same column of Z for W contiguous steps. This is shown in Algorithm 1, when learning a schedule of N^α learning rates with FDS. Here, $Z_{[i]}^\alpha$ refers to the i -th column of matrix Z^α . We don't need to store \mathcal{H} or A^α in memory since we calculate the Hessian matrix product $\mathcal{H}Z^\alpha$ directly. Most importantly, note that hyperparameters aren't updated greedily or online, but are updated once per outer step, which corresponds to differentiating through the entire unrolled inner loop optimization and getting the exact hypergradients $\partial\mathcal{L}_{\text{val}}(\theta_T)/\partial\alpha$.

The main cost of Algorithm 1 comes from calculating $\mathcal{H}Z^\alpha$. There exists several methods to approximate this product, but we found them too crude for long horizons. This includes first-order approximations or truncation [32], which can be adapted to FDS trivially. One could also use functional forms for more complex schedules to be learned in terms of fewer hyperparameters, but this typically makes stronger assumptions about the shape of each hyperparameter schedule, which can easily cloud the true performance of HPO algorithms. In practice, we calculate $\mathcal{H}Z^\alpha$ exactly, and use a similar form to Algorithm 1 to learn α , β and ξ .

5 Experiments

Our experiments show how FDS mitigates gradient degradation and outperforms competing HPO methods for tasks with a long horizon. In Sections 5.1 and 5.2 we consider small datasets (MNIST and SVHN) and a small network (LeNet) to make reverse-mode differentiation tractable, so that the effect of hyperparameter sharing can be directly measured. In Sections 5.3 and 5.4, we then showcase FDS on CIFAR-10 where only forward-mode differentiation is

Algorithm 1 Simplified FDS algorithm when learning N^α learning rates for the SGD optimizer with momentum. Each learning rate is shared over W contiguous time steps

Initialize: $N^\alpha, W = T/N^\alpha, \alpha = \mathbf{0}^{N^\alpha}$

#outer loop

for o in $1, 2, \dots$ **do**

Initialize: $\mathcal{D}_{\text{train}}, \mathcal{D}_{\text{val}}, \theta_0 \in \mathbb{R}^D$,
 $Z^\alpha = \mathbf{0}^{D \times N^\alpha}, C^\alpha = \mathbf{0}^{D \times N^\alpha}$,

#inner loop

for t in $1, 2, \dots, T$ **do**

$\mathbf{x}_{\text{train}}, \mathbf{y}_{\text{train}} \sim \mathcal{D}_{\text{train}}$
 $\mathbf{g}_{\text{train}} = \partial\mathcal{L}_{\text{train}}(\mathbf{x}_{\text{train}}, \mathbf{y}_{\text{train}})/\partial\theta$

#hyperparameter sharing

$i = \lceil t/W \rceil$
 $\mathcal{H}Z_{[1:i]}^\alpha = \partial(\mathbf{g}_{\text{train}} Z_{[1:i]}^\alpha)/\partial\theta$
 $Z_{[1:i]}^\alpha = A^\alpha Z_{[1:i]}^\alpha + B_{[1:i]}^\alpha$

update C^α (Eq 4)
 $\theta_{t+1} = \Phi(\theta_t, \mathbf{g}_{\text{train}})$

end for

$\mathbf{g}_{\text{val}} = \partial\mathcal{L}_{\text{val}}(\mathcal{D}_{\text{val}})/\partial\theta$
 $\alpha \leftarrow \alpha - 0.1 \times \mathbf{g}_{\text{val}} Z^\alpha / W$

end for

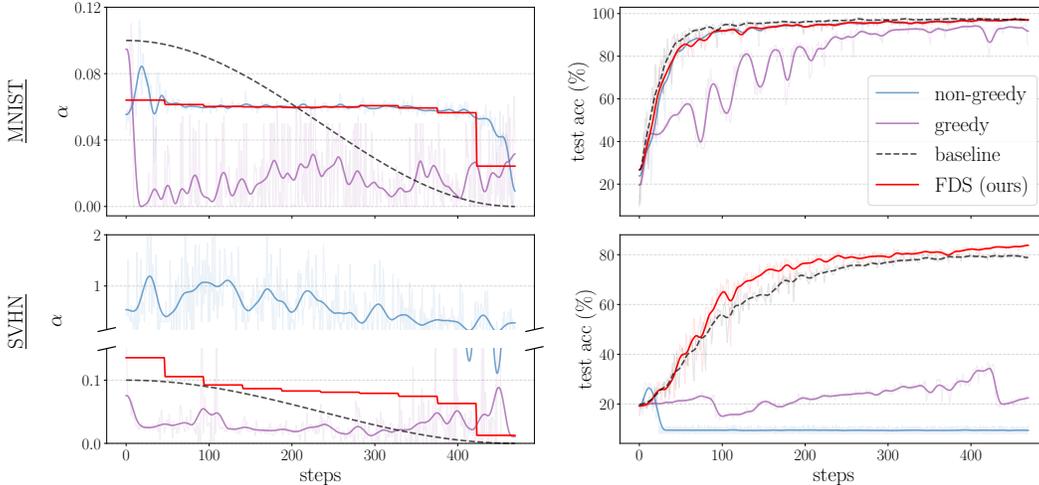


Figure 3: The learning rate schedule α learned for the MNIST and SVHN datasets using a LeNet architecture over one epoch. We observe that on real-world datasets like SVHN, both greedy and non-greedy hyperparameter optimizations fail to learn decent learning rate schedules when using one hypergradient per inner step. However, sharing learning rates over contiguous steps stabilizes non-greedy hypergradients and allows us to find learning rates that can even outperform reasonable off-the-shelf schedules in this setting.

tractable. All experiments are carried out on a single GTX 2080 GPU. More implementation details can be found in Appendix C.

5.1 The effect of hyperparameter sharing on hypergradient noise

We consider a LeNet network trained on an inner loop of $T = 250$ gradient steps, and calculate the hypergradients of the learning rate α across 100 seeds. Each seed is evaluated at the same value of α , but corresponds to a different training dataset ordering and weight initialization. We can calculate the hypergradients greedily ($H = 1$) and non-greedily ($H = T$). The optimal hypergradients are considered to be the ensemble average of the non-greedy seeds as per Section 4.1, which allows an MSE to be calculated for each method. These results are shown in Figure 2 for a learning rate schedule initialized to small values. We observe that greediness is a poor approximation to the optimal hypergradients, and that time averaging contiguous hypergradients in the non-greedy case can significantly reduce the MSE even when averaging over only $W = 50$ steps.

The simplicity of this problem allows us to calculate Σ , c and ϵ as defined in Theorem 4.1, to verify that our upper bound holds and that its shape as a function of W is realistic. In the setting shown in Figure 2 we find $\epsilon = \max |\mu_{t+1} - \mu_t| \sim 0.08$ and $(1/T) \sum_t \Sigma_{tt} \sim 0.25$. The value of the maximum correlation between any two steps, c , can be quite high which makes the upper bound loose. In practice however, the shape of the upper bound in Theorem 4.1 as a function of W (as plotted in Appendix B) closely matches that of the measured MSE shown in Figure 2. Note that the values of ϵ , Σ , and c can vary depending on the value of α . As illustrated in Theorem 4.1, we find that settings that have a smaller values of ϵ_t benefit from a larger W and reduce the MSE more (see Appendix D for more examples).

5.2 The effect of hyperparameter sharing on HPO

In the section above, we considered hypergradient noise around a fixed hyperparameter setting. In this section, we consider how that noise and its mitigation translate to meta-learning hyperparameters over several outer steps.

In Figure 3, we initialize $\alpha = \mathbf{0}$ and train a LeNet network over 1 epoch ($T \sim 500$ gradient steps) on MNIST and SVHN. We compare the maximally greedy setting ($H = 1$), the non-greedy setting ($H = T$), and FDS (also $H = T$ but with sharing of $W \sim 50$ contiguous hyperparameters). In

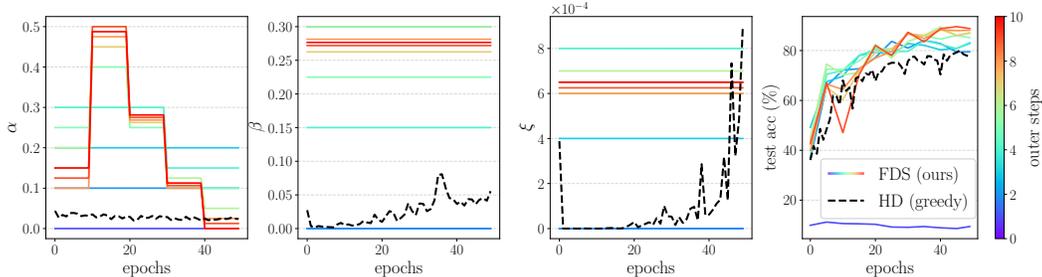


Figure 4: FDS applied to the SGD optimizer to learn (from left to right) the learning rate schedule α , the momentum β , and weight decay ξ for a WRN-16-1 on CIFAR-10. For each outer step (color) we solve CIFAR-10 from scratch for 50 epochs, and update all hyperparameters such that the final weights minimize some validation loss. We use hyperparameter sharing over $W = 10, 50$ and 50 epochs for α, β and ξ respectively. All hyperparameters are initialized to zero and converge within just 10 outer steps to values that significantly outperform Hypergradient Descent (HD) [14], the greedy alternative. We match the performance of the best known hyperparameters in that setting and do so much faster than state-of-the-art black-box methods (see Section 5.4).

all cases we take 50 outer steps per hyperparameter. We make greediness more transparent by not using tricks as in Hypergradient Descent [14], where α_t is set to α_{t-1} before its hypergradients are calculated. As previously observed by [15], greedy optimization leads to poor solutions with learning rates that are always too small. While the non-greedy setting without sharing works well for simple datasets like MNIST, it fails for real-world datasets like SVHN, converging to much higher learning rates than reasonable. This is due to gradient degradation, whose negative effect can compound during outer optimization, as a single large learning rate can increase the hypergradient noise for neighbouring steps. As we share hyperparameters in FDS, we reduce and stabilize the outer optimization. This allows us to learn a much more sensible learning rate schedule, which can even outperform a reasonable cosine annealing off-the-shelf schedule on SVHN.

5.3 FDS on CIFAR-10

We demonstrate that our algorithm can be used to learn the learning rate, momentum and weight decay over 50 epochs of CIFAR-10 ($H \sim 10^4$) for a WideResNet of 16 layers (WRN-16-1). We choose not to use larger architectures or more epochs to enable compute time for more extensive comparisons and because hyperparameters matter most for fewer epochs. Note also that we are not interested in finding the architecture with the best performance, but rather in finding the best hyperparameters given an architecture. For the learning rate, we choose W such that the ratio of T/W is similar to the optimal one found in 5.1, and we set $W = T$ for the momentum and weight decay since only a single value is commonly used for these hyperparameters. The schedules learned are shown in Figure 4, which demonstrates that FDS converges in just 10 outer steps to hyperparameters that are very different to online greedy differentiation [14], and correspond to significantly better test accuracy performances. Note that having the maximum learning rate be large and occur half way during training is reminiscent of the one-cycle schedule [53].

5.4 FDS compared to other HPO methods

A common theme in meta-learning research has been the lack of appropriate baselines, with researchers often finding that random search (RS) can outperform complex search algorithms, for instance in NAS [54] or automatic augmentation [55]. In this section we compare FDS to several competing HPO methods on CIFAR-10, in both the performance of the hyperparameters found and the time it takes to find them. We consider Random Search (RS), Bayesian Optimization (BO), Hypergradient Descent (HD) [14], HyperBand (HB) [12] and the combination of BO and HB, namely BOHB [13]. The latter is typically regarded as state-of-the-art in HPO. We use the official HpBandster [56] implementation of these algorithms, except for HD which we re-implemented ourselves.

The performance of HPO methods is often clouded by very small search ranges. For instance, in [13] the learning rate is searched over the range $[10^{-6}, 10^{-1}]$. In the case of DARTS [38], expanding the search space to include many poor architectures has helped diagnose issues with its search algorithm [57]. For these reasons, we assume only weak priors and consider large search ranges: $\alpha \in [-1, 1]$, $\beta \in [-1.5, 1.5]$, and $\xi \in [-4 \times 10^{-3}, 4 \times 10^{-3}]$, which includes many poor hyperparameter values. In FDS we can specify search ranges by using a fixed update size γ to the hyperparameters, which decays by 2 every time the hypergradient flips sign, which is common in sign based optimizers [58–61]. We found that black-box HPO methods did not scale well to more than ~ 10 hyperparameters for such large search ranges, and so considered learning 7 learning rates, 1 momentum and 1 weight decay over 50 epochs. Note that increasing these numbers doesn't affect the accuracy of FDS but significantly reduces that of RS, BO, HB and BOHB. The performance over time of the hyperparameters found by each method is shown in Figure 1. As is common in HPO, we plot regret (in test accuracy) with respect to the best hyperparameters known in this setting, which we obtained from an expensive grid search around the most common hyperparameters used in the literature (more details in Appendix F). To squeeze the optimal performance out of FDS in this experiment, we match the process used in HB and BOHB, namely using smaller budgets for some runs, in particular early ones. We can see that our method reaches zero regret in just 3 hours, while the next best method (BOHB) reaches zero regret in 60 hours. Other methods did not reach zero regret in the maximum of 100 hours they were run for. Note also that we retain the convergence speed of online greedy differentiation [14] while outperforming its regret by $\sim 10\%$ test accuracy.

6 Discussion

FDS is significantly better than modern HPO alternatives in the case of learning differentiable hyperparameters over long horizons. Its main advantage over black-box methods is its convergence speed, and its main advantage over other gradient-based methods is that it's non-greedy. Furthermore, forward-mode differentiation can provide exact hypergradients for any differentiable hyperparameter, contrary to some other approaches like implicit differentiation which approximates hypergradients, and can do so for certain types of differentiable hyperparameters only.

Nonetheless, it is worth pointing out the main limitations of FDS. First, black-box methods are slower but can readily tackle non-differentiable hyperparameters, while FDS would need to use relaxation techniques to differentiate through, say, discrete variables. Then, the computational and memory cost of FDS scales linearly with the number of hyperparameters. For instance, for a WideResNet of 16 layers we are limited in memory to $\sim 10^3$ hyperparameters on a single 12 GB GPU, which falls short of reverse-mode-based greedy methods which scale to millions of hyperparameters. Finally, while automatic forward-mode differentiation is becoming available for some toolboxes like JAX [62], it remains unavailable for the most popular deep learning libraries, which means that some pen-and-paper work is still required to derive hypergradients for given hyperparameters (Eq 4).

The focus of our future work will be on adapting FDS to the many meta-learning applications that rely on greedy optimization, such as differentiable architecture search, to improve their performance by making them non-greedy.

7 Conclusion

This work makes an important step towards gradient-based HPO for long horizons by introducing FDS, which enables non-greediness through the mitigation of gradient degradation. More specifically, we theorize and demonstrate that sharing hyperparameters over contiguous time steps is a simple yet efficient way to reduce the error in their hypergradients; a setting which naturally lands itself to forward-mode differentiation. Finally, we show that FDS outperform greedy gradient-based alternatives in the quality of hyperparameters found, while being significantly faster than all state-of-the-art black-box methods. We hope that our work encourages the community to reconsider gradient-based HPO in terms of non-greediness, and pave the way towards a universal hyperparameter solver.

Acknowledgements The authors would like to thank Joseph Mellor, Antreas Antoniou, Miguel Jaques, Luke Darlow and Benjamin Rhodes for their useful feedback throughout this project. Our work was supported in part by the EPSRC Centre for Doctoral Training in Data Science, funded by the UK Engineering and Physical Sciences Research Council (grant EP/L016427/1) and the University of Edinburgh, as well as a Huawei DDMPLab Innovation Research Grant.

References

- [1] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016.
- [2] A. Brock, J. Donahue, and K. Simonyan, “Large scale GAN training for high fidelity natural image synthesis,” in *International Conference on Learning Representations*, 2019.
- [3] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” *arXiv preprint arXiv:1810.04805*, 2018.
- [4] A. v. d. Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, “Wavenet: A generative model for raw audio,” *arXiv preprint arXiv:1609.03499*, 2016.
- [5] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *International Conference on Learning Representations* (Y. Bengio and Y. LeCun, eds.), 2015.
- [6] I. Loshchilov and F. Hutter, “Sgdr: Stochastic gradient descent with warm restarts,” in *International Conference on Learning Representations*, 2017.
- [7] M. Li, E. Yumer, and D. Ramanan, “Budgeted training: Rethinking deep neural network training under resource constraints,” in *International Conference on Learning Representations*, 2020.
- [8] P. Goyal, P. Dollár, R. Girshick, P. Noordhuis, L. Wesolowski, A. Kyrola, A. Tulloch, Y. Jia, and K. He, “Accurate, large minibatch sgd: Training imagenet in 1 hour,” *arXiv preprint arXiv:1706.02677*, 2017.
- [9] X. Dong, M. Tan, A. W. Yu, D. Peng, B. Gabrys, and Q. V. Le, “Autohas: Differentiable hyper-parameter and architecture search,” 2020.
- [10] I. Loshchilov and F. Hutter, “Decoupled weight decay regularization,” in *International Conference on Learning Representations*, 2019.
- [11] S. Ioffe and C. Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” in *International conference on Machine Learning* (F. Bach and D. Blei, eds.), Proceedings of Machine Learning Research, (Lille, France), PMLR, 07–09 Jul 2015.
- [12] L. Li, K. G. Jamieson, G. DeSalvo, A. Rostamizadeh, and A. Talwalkar, “Hyperband: A novel bandit-based approach to hyperparameter optimization,” *J. Mach. Learn. Res.*, vol. 18, pp. 185:1–185:52, 2017.
- [13] S. Falkner, A. Klein, and F. Hutter, “BOHB: Robust and efficient hyperparameter optimization at scale,” in *Proceedings of the 35th International Conference on Machine Learning*, pp. 1436–1445, 2018.
- [14] A. G. Baydin, R. Cornish, D. M. Rubio, M. Schmidt, and F. Wood, “Online learning rate adaptation with hypergradient descent,” in *International Conference on Learning Representations*, 2018.
- [15] Y. Wu, M. Ren, R. Liao, and R. Grosse., “Understanding short-horizon bias in stochastic meta-optimization,” in *International Conference on Learning Representations*, 2018.
- [16] L. Franceschi, M. Donini, P. Frasconi, and M. Pontil, “Forward and reverse gradient-based hyperparameter optimization,” in *International conference on Machine Learning* (D. Precup and Y. W. Teh, eds.), Proceedings of Machine Learning Research, (International Convention Centre, Sydney, Australia), PMLR, 06–11 Aug 2017.
- [17] M. Donini, L. Franceschi, M. Pontil, O. Majumder, and P. Frasconi, “Scheduling the Learning Rate via Hypergradients: New Insights and a New Algorithm,” *arXiv e-prints*, Oct. 2019.
- [18] M. Feurer and F. Hutter, *Chapter 1: Hyperparameter Optimization*. Cham: Springer International Publishing, 2019.

- [19] J. Snoek, O. Rippel, K. Swersky, R. Kiros, N. Satish, N. Sundaram, M. Patwary, M. Prabhat, and R. Adams, “Scalable bayesian optimization using deep neural networks,” in *International conference on Machine Learning*, 2015.
- [20] B. Zoph and Q. V. Le, “Neural architecture search with reinforcement learning,” in *International Conference on Learning Representations*, 2017.
- [21] M. Jaderberg, V. Dalibard, S. Osindero, W. M. Czarnecki, J. Donahue, A. Razavi, O. Vinyals, T. Green, I. Dunning, K. Simonyan, *et al.*, “Population based training of neural networks,” *arXiv preprint arXiv:1711.09846*, 2017.
- [22] Y. Bengio, “Gradient-based optimization of hyperparameters,” *Neural Comput.*, Aug. 2000.
- [23] T. Hospedales, A. Antoniou, P. Micaelli, and A. Storkey, “Meta-learning in neural networks: A survey,” *arXiv preprint arXiv:2004.05439*, 2020.
- [24] C. Finn, P. Abbeel, and S. Levine, “Model-agnostic meta-learning for fast adaptation of deep networks,” in *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017* (D. Precup and Y. W. Teh, eds.), vol. 70 of *Proceedings of Machine Learning Research*, pp. 1126–1135, PMLR, 2017.
- [25] L. Franceschi, P. Frasconi, S. Salzo, R. Grazzi, and M. Pontil, “Bilevel programming for hyperparameter optimization and meta-learning,” in *International conference on Machine Learning*, 2018.
- [26] P. J. Werbos, “Backpropagation through time: what it does and how to do it,” *Proceedings of the IEEE*, 1990.
- [27] R. J. Williams and D. Zipser, “A learning algorithm for continually running fully recurrent neural networks,” *Neural Computation*, 1989.
- [28] J. Domke, “Generic methods for optimization-based modeling,” in *International Conference on Artificial Intelligence and Statistics* (N. D. Lawrence and M. Girolami, eds.), *Proceedings of Machine Learning Research*, (La Palma, Canary Islands), PMLR, 21–23 Apr 2012.
- [29] D. Maclaurin, D. Duvenaud, and R. P. Adams, “Gradient-based Hyperparameter Optimization through Reversible Learning,” *arXiv e-prints*, Feb. 2015.
- [30] F. Pedregosa, “Hyperparameter optimization with approximate gradient,” in *International conference on Machine Learning*, 2016.
- [31] J. Fu, H. Luo, J. Feng, K. H. Low, and T.-S. Chua, “Drmad: Distilling reverse-mode automatic differentiation for optimizing hyperparameters of deep neural networks,” in *IJCAI*, 2016.
- [32] A. Shaban, C.-A. Cheng, N. Hatch, and B. Boots, “Truncated back-propagation for bilevel optimization,” in *AISTATS*, 2019.
- [33] J. Larsen, L. K. Hansen, C. Svarer, and M. Ohlsson, “Design and regularization of neural networks: the optimal use of a validation set,” in *Neural Networks for Signal Processing VI. Proceedings of the 1996 IEEE Signal Processing Society Workshop*, pp. 62–71, 1996.
- [34] A. Rajeswaran, C. Finn, S. M. Kakade, and S. Levine, “Meta-learning with implicit gradients,” in *Advances in Neural Information Processing Systems*, Curran Associates, Inc., 2019.
- [35] J. Lorraine, P. Vicol, and D. Duvenaud, “Optimizing Millions of Hyperparameters by Implicit Differentiation,” *arXiv e-prints*, Nov. 2019.
- [36] L. Metz, N. Maheswaranathan, J. Nixon, D. Freeman, and J. Sohl-Dickstein, “Understanding and correcting pathologies in the training of learned optimizers,” vol. 97, pp. 4556–4565, 09–15 Jun 2019.
- [37] J. Luketina, M. Berglund, K. Greff, and T. Raiko, “Scalable gradient-based tuning of continuous regularization hyperparameters,” in *International conference on Machine Learning*, 2016.
- [38] H. Liu, K. Simonyan, and Y. Yang, “DARTS: Differentiable architecture search,” in *International Conference on Learning Representations*, 2019.
- [39] T. Wang, J. Zhu, A. Torralba, and A. A. Efros, “Dataset distillation,” *CoRR*, vol. abs/1811.10959, 2018.
- [40] M. Ren, W. Zeng, B. Yang, and R. Urtasun, “Learning to reweight examples for robust deep learning,” in *International conference on Machine Learning*, 2018.

- [41] Y. Bengio, P. Frasconi, and P. Simard, “The problem of learning long-term dependencies in recurrent networks,” in *IEEE International Conference on Neural Networks*, 1993.
- [42] Y. Bengio, P. Simard, and P. Frasconi, “Learning long-term dependencies with gradient descent is difficult,” *IEEE Transactions on Neural Networks*, 1994.
- [43] P. Parmas, C. E. Rasmussen, J. Peters, and K. Doya, “PIPPS: Flexible model-based policy search robust to the curse of chaos,” vol. 80 of *Proceedings of Machine Learning Research*, (Stockholmsmässan, Stockholm Sweden), pp. 4065–4074, PMLR, 10–15 Jul 2018.
- [44] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural computation*, 1997.
- [45] R. Pascanu, T. Mikolov, and Y. Bengio, “On the difficulty of training recurrent neural networks,” in *International conference on Machine Learning*, 2013.
- [46] S. Flennerhag, A. A. Rusu, R. Pascanu, F. Visin, H. Yin, and R. Hadsell, “Meta-learning with warped gradient descent,” in *International Conference on Learning Representations*, 2020.
- [47] S. Zagoruyko and N. Komodakis, “Wide residual networks,” in *BMVC*, 2016.
- [48] H. Stackelberg, *The Theory Of Market Economy*. Oxford University Press, 1952.
- [49] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimeshain, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, “Pytorch: An imperative style, high-performance deep learning library,” in *Advances in Neural Information Processing Systems*, Curran Associates, Inc., 2019.
- [50] P. Walters, *An Introduction to Ergodic Theory*. Graduate texts in mathematics, Springer-Verlag, 1982.
- [51] L. Boltzmann, *Vorlesungen uber Gastheorie*. J.A. Barth, 1896.
- [52] O. Peters, “The ergodicity problem in economics,” *Nature Physics*, vol. 15, pp. 1216–1221, Dec 2019.
- [53] L. N. Smith and N. Topin, “Super-convergence: Very fast training of residual networks using large learning rates,” *CoRR*, vol. abs/1708.07120, 2017.
- [54] L. Li and A. Talwalkar, “Random search and reproducibility for neural architecture search,” in *UAI*, 2019.
- [55] E. D. Cubuk, B. Zoph, J. Shlens, and Q. V. Le, “Randaugment: Practical automated data augmentation with a reduced search space,” *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 3008–3017, 2020.
- [56] A. Klein. <https://automl.github.io/HpBandSter/build/html/quickstart.html>, 2019.
- [57] A. Zela, T. Elsken, T. Saikia, Y. Marrakchi, T. Brox, and F. Hutter, “Understanding and robustifying differentiable architecture search,” in *International Conference on Learning Representations*, 2020.
- [58] “Increased rates of convergence through learning rate adaptation,” *Neural Networks*, vol. 1, no. 4, pp. 295–307, 1988.
- [59] M. Riedmiller and H. Braun, “A direct adaptive method for faster backpropagation learning: the rprop algorithm,” in *IEEE International Conference on Neural Networks*, pp. 586–591 vol.1, 1993.
- [60] J. Bernstein, Y. Wang, K. Aizzadenesheli, and A. Anandkumar, “SIGNSGD: compressed optimisation for non-convex problems,” in *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018* (J. G. Dy and A. Krause, eds.), vol. 80 of *Proceedings of Machine Learning Research*, pp. 559–568, PMLR, 2018.
- [61] M. Safaryan and P. Richtárik, “On Stochastic Sign Descent Methods,” *arXiv e-prints*, p. arXiv:1905.12938, May 2019.
- [62] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, “JAX: composable transformations of Python+NumPy programs,” 2018.