

A Generalizable Approach to Learning Optimizers

Diogo Almeida Clemens Winter Jie Tang Wojciech Zaremba
OpenAI

Abstract

A core issue with learning to optimize neural networks has been the lack of generalization to real world problems. To address this, we describe a system designed from a generalization-first perspective, learning to update optimizer hyperparameters instead of model parameters directly using novel features, actions, and a reward function. This system outperforms Adam at all neural network tasks including on modalities not seen during training. We achieve 2x speedups on ImageNet, and a 2.5x speedup on a language modeling task using over 5 orders of magnitude more compute than the training tasks.

1 Introduction

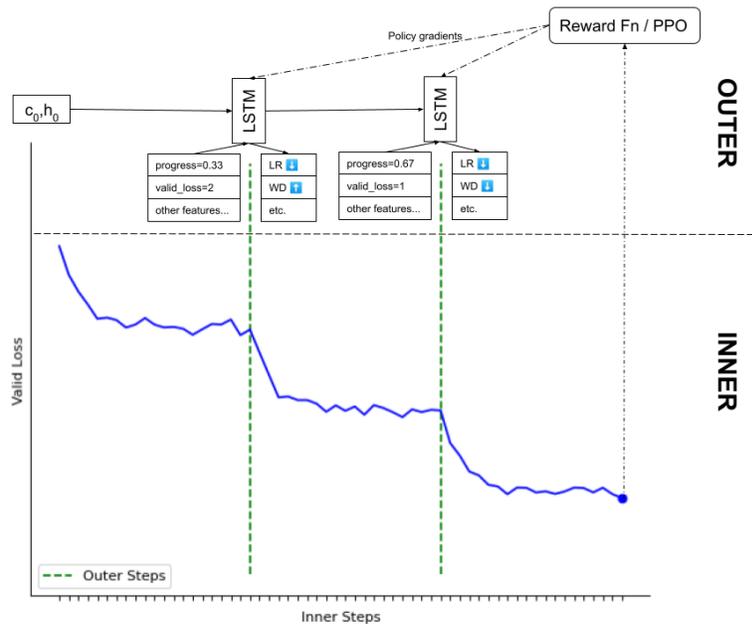


Figure 1: Simplified diagram of our LHOPT approach. Some small number of times through training, an LSTM controller is used to update optimizer hyperparameters based on input statistics from the training process. We optimize for generalization by performing hyperparameter updates instead of parameter updates, decouple update frequency from inner task size, use unitless features and relative actions, solely optimize for final performance, and use a novel reward function.

Why haven't we automated deep neural network optimization yet? We know that a problem exists: optimization settings greatly affect the quality of trained models [7]. We know that tools exist that can tackle problems like this: modern reinforcement learning methods have made great strides at solving

difficult sequential decision making problems with imperfect information, such as Dota 2 [37]. And we know that solving this problem would be impactful: the field is continuing to scale up models [21] and even small relative speedups could result in saving thousands of petaflop/s-days [8].

Previous work [3, 46, 33, 34] has proposed solutions to the problem of learning how to optimize neural networks, but despite demonstrating large speedups on problems from or close to their training distribution, all have shown limited generalization far outside of that training distribution, and not a single one is used for real use cases. While there are many problems learned optimizers have to solve, we, like previous work [34], argue that the core issue is generalization.

In this work, we provide a description of a system as well as motivation for novel design choices for learned optimizers with the focus on generalizing to and increasing computational efficiency of real world problems. We caveat that this work is preliminary with limitations outlined throughout the paper.

1.1 Contributions

Most existing methods attempt to learn a general update rule from scratch: learning parameter updates [3, 46, 33, 34]. In contrast, we leverage the priors of existing optimization algorithms as much as possible, and instead learn hyperparameter updates. We call this class of learned optimizers Learned Hyperparameter Optimizers (LHOPTs). The resulting optimizers can be interpreted as having data-driven schedules that interpolate between hand-designed optimizers.

In addition to presenting the general framework of learning to update hyperparameters, we also revisit many design decisions for learned optimizers from a generalization-first perspective and present novel approaches for actions, features, and reward functions for learned optimizers in general.

Our LHOPTs manage to generalize without tuning. They achieve approximately a 2x speedup on ImageNet [13] compared to tuning both AdamW [26] learning rate and schedule. They outperform both baselines from MLPerf [29] that are most different our training distribution: speech recognition and neural collaborative filtering. Finally, using a LHOPT’s hyperparameter schedule from a smaller language modeling task, we apply it to a well-tuned language modeling codebase resulting in 2.5x speedups. In this scenario, despite the optimizer being trained on inner optimization problems that average only 170 GPU-seconds, it generalizes to all models that take up to 300 GPU-days (See Figure 2) - generalization of over 5 orders of magnitude.

We are not claiming that using a LHOPT is always better than default optimizers for neural network tasks, nor that LHOPTs would be the best for getting state of the art results. We see the primary contribution of this work as demonstrating that learned optimizers can generalize to real neural network tasks, and hope that it will inspire further efforts which deviate from the standard approaches and focus purely on generalization.

2 Related Work

Our work is most heavily inspired by a recent body of work that uses neural networks to learn parameter updates [3, 46, 33, 34]. Each subsequent refinement achieved large improvements in generalization capabilities, but even the most recent still had issues generalizing to slightly wider networks on the same task [34].

Another similar branch of work is on using hypergradients [27, 6, 16, 35], direct gradients from loss to hyperparameters, in order to optimize those hyperparameters on-the-fly. These methods have the advantage of not requiring training a policy ahead of time, but generally suffer from short-horizon bias [48].

Most similar to our work are those that use reinforcement learning to control learning rate [12, 50, 51]. We interpret these models as learning rate-only versions of our LHOPT. Despite most of their results being on small scale datasets, the results of this paper indicate that their results could possibly generalize to large scale models.

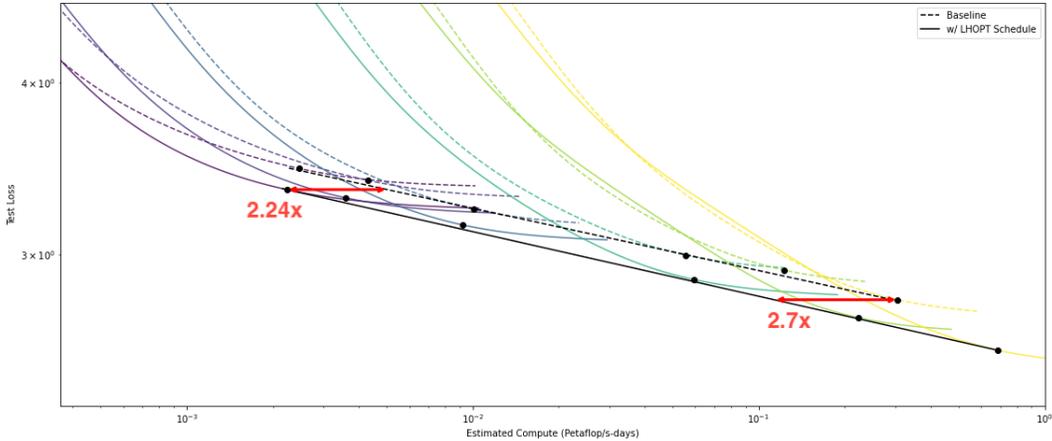


Figure 2: Test learning curves and scaling law fit of compute efficient frontier on a large well-tuned language modeling codebase. Our learned optimizers demonstrate consistent speedups $\geq 2x$, with speedup increasing as model size does with no computational overhead. Dotted lines are baselines, full lines use a LHOPT hyperparameter schedule from a similar but smaller task.

3 Learning Hyperparameter Updates

Learning to update hyperparameters as opposed to parameters affords us 2 main axes of decoupling.

The first is our policy network architecture is decoupled from the details of the inner optimization task. For example, an optimizer that learns parameter updates would at least require access to gradient values, as well as possibly parameter values and other statistics. The problem with these statistics is that they may not be robust since they can assume very different values on out-of-distribution tasks, and thus result in undefined behavior. Instead, by letting the inner optimizer (e.g. AdamW) interact with these low level statistics, we have more flexibility in the types of features our policy uses for decision making and we can trade-off between giving more information and avoiding overfitting to idiosyncrasies of our training set.

The second is that we decouple the frequency at which our policy has to run. By definition, learned parameter optimizers have to run on every update, but LHOPTs can be run at any frequency, independent of the number of updates of the problem. This benefits both generalization and solves the problem of short-horizon bias [48]: it allows us to both use the same number of steps for extremely large problems as that of smaller ones, and also optimize solely for final reward, as opposed to optimizing for reward after some fixed number of unrolled steps. This has the additional benefit of greatly decreasing the computational overhead of learned optimizers.

Given these benefits, we formulate the problem of learning hyperparameter optimizers as that of reinforcement learning, optimizing solely for final reward (Section 3.4). We choose PPO [41] out of convenience. The LHOPT formulation allows us much flexibility, and we describe our more details for our choices of input features (Section 3.2) and actions (Section 3.3), though we do not claim any design decision to be optimal and haven't tested alternatives unless otherwise specified.

3.1 Problem Setup

Our LHOPTs are designed to be applied to an inner task - a standard optimization problem, which uses an inner optimizer to perform updates every inner step (this would be the standard training step/update). Some predefined number of times throughout inner training, we also take an outer step where features from the inner task are summarized and passed into policy network π , which then outputs hyperparameter updates for the inner optimizer. These updates are applied and inner training resumes until the next outer step.

When training our learned optimizers, at the end of training each inner task, we map the final validation loss into a specified reward, which the PPO algorithm uses to update the weights of our policy network. Figure 1 shows a simplified view of the training setup.

At evaluation time, the LHOPT’s behavior is defined by the learned weights of its policy network, as well as its predefined feature space and action space.

3.2 Feature Space

Our guiding principle for designing the feature space was to prevent the policy from inferring the inner optimization task. To ensure good generalization, we wanted every feature to look similar from toy problems all the way to massive scale models beyond which we could afford to train on. As such, we focused on "unitless features".

Existing learned optimizers often use statistics such as the gradient, square of the gradient, loss values, and exponential moving averages (EMAs) of these statistics, which can be highly specific to the inner optimization task. In contrast, the majority of our features were transformations: log-ratios, cosine similarities, averages of booleans, or what we call CDF features.

Some examples of unitless features that we use are log-ratio between training and validation loss, log of the noise scale, cosine similarity between gradient and momentum tensors, how often the gradient is clipped, whether or not the loss is NaN, progress (how far along through inner training we are), an encoding of the previously sampled action, and CDF features of the validation loss. For a complete list of features of our final optimizers, see Appendix J.

The motivations for CDF features are that (1) there are use-cases for which we would like to know relative values of a feature within an inner task and (2) ranking features should be invariant to details of the inner task. One example is to know if validation loss is plateauing without seeing its exact value. To achieve this, we calculate an estimate of that value’s mean and variance (see Appendix E for how we do this in a robust way), and map that feature with a gaussian cumulative distribution function onto the interval [0,1]. Note that these features lose a lot of information: the first feature value is always 0.5, the second feature is 1 if larger than the first or 0 if smaller, etc.

One detail related to using PPO is that the value function V can be given non-robust features [39] in addition to the robust features given to the policy π .

As a final step of preprocessing, all input features are normalized by subtracting an exponential moving average and dividing by an exponential moving standard deviation (both with $\beta = 0.999$). We additionally clip every input feature element-wise to $[-2, 2]$ with the purpose of not being able to differentiate outliers.

Limitations. (1) One requirement that prevents plug-and-play usage of our LHOPTs is they require progress, training loss, and validation loss as an input from the user at evaluation time. Note that it is entirely possible to train a LHOPT with any or all of these features missing at the expense of performance. (2) As a result of making the validation loss observable by the policy, it is theoretical possible to overfit the validation dataset.. This can be solved with a separate validation set for this loss, but in practice, this doesn’t appear to be an issue (Appendix D). (3) Some features can be expensive to calculate on every inner step - for those features (specified in Appendix J), we instead calculate them intermittently every 4 inner steps to reduce the overhead.

3.3 Action Space

For our actions, we limited our models to just two ways of updating hyperparameters: scaling (multiplying by a constant) for most hyperparameters, and logit shifting (applying the logit function, adding a constant, then applying a sigmoid) where we want to be careful when changing values near 0 and 1. By limiting our policy to relative actions and including randomness in the initial hyperparameters, models can never know the current values of hyperparameters and we make our model more robust by forcing them to react to the underlying inner task and not simply set hyperparameters to a specific value. Our hypothesis is that one set of hyperparameters on one task might behave very similarly to another set of hyperparameters on a different task.

We also had one additional class of actions that were not hyperparameter updates but fit in nicely within the existing framework: learning to restart from checkpoints. There are many motivations for such an action: (1) ideally learned optimizers would be able to handle all the task-specific tuning that a practitioner would have to do and restarting on divergence is one such tasks, (2) previous work has noted that SGD often works best with the highest possible stable learning rate [43] and it may

not be possible to determine that value without venturing into unstable territory, (3) sophisticated hyperparameter optimizations algorithms such as Population Based Training [20] could be learned from such a simple action, and finally (4) even if restarting was never used by a trained model, it could greatly help with exploration while training - to both decrease the length of credit assignment paths and also make it less punishing for models to sample suboptimal settings.

All actions were represented as discrete choices rather than real values, as this helps the stability of the RL training. In order to aid with exploration, we removed the identity action for many hyperparameters to encourage setting their values dynamically - in other words, by forcing the controller to always change the hyperparameter, we force it to be reactive to the inner problem instead of sticking with a safe initial value.

Limitations. (1) Despite the motivation, the policy could theoretically find the initial hyperparameters because most hyperparameters are clipped to minimum and maximum values. The optimizer could find when that clipping occurs, then restart to the beginning with that knowledge. We did not observe this behavior in our models, but it is possible that, as the techniques become more sophisticated, this design decision would need to be revisited. (2) In practice, our models did nothing close to the sophistication of PBT with restart actions, though they did use restarts effectively when divergence occurred, especially when given highly unstable initial hyperparameters. We hypothesize that the policy may be undertrained (see Section 3.6) and these sophisticated behaviors could emerge with larger-scale training.

3.4 Reward Function

There are many desirable properties of a reward function: handling problems with very different losses, returning reward values that correspond to how difficult improvement was on a problem, being able to extrapolate beyond any baseline performance, and being computable on-the-fly to enable a potentially infinite inner problem distribution. Previous work mostly sacrifices the last two properties by computing many baselines and normalizing by the best one [34].

Instead, we use the following reward function: we first evaluate a baseline on the exact same problem, including all task specific values and initial weights. Then, we take all the points on the learning curve of that baseline and fit a power law curve to it. The result is a function where 0 corresponds to the worst loss, 1 to the best loss, and values greater than 1 take into account how difficult that problem is (i.e., we can extrapolate past the best loss of a baseline). While any baseline could be used here, we use an approach inspired by self-play [44, 42] and use the same learned optimizer policy that we were training but using an EMA of the weights, thus allowing our reward function to improve throughout outer training.

The main downside of naive application of this reward function is that it doubles the overhead of inner training. Instead, we choose to repeat each task four times to amortize the overhead of computing a baseline.

3.5 Inner Optimizer

While the LHOPT formulation would apply to any optimizer¹, we choose to focus on our own Customizable Inner Adaptive Optimizer (CIAO). The reason for this is that existing optimizers have been designed to work with mostly static hyperparameters and minimal hyperparameter tuning. By employing an inner optimizer designed to work with a LHOPT policy, hyperparameters turn from a weakness into a strength, and we can more easily create optimization techniques without worrying about how to tune them.

As such, our inner optimizer not only generalizes many existing standard and non-standard optimizers, we also introduce several novel techniques to improve optimization quality, including one to automate clipping gradients based on an exponentially weighted moving maximum of the gradient norm. Source code for our inner optimizer will be released. Additionally, the default settings for our inner optimizer involve using the denominator of Adamax [22] and the trust ratio of LAMB [53]. See Appendix I for more inner optimizer details and Appendix C for rationale on why our inner optimizers are use the LAMB trust ratio.

¹See Appendix F for an older model which was used to set LR for both SGD and Adam.

Limitations. (1) Despite the inner optimizer is closer to LAMB or Adamax than Adam, we still compare to Adam for most of the results. We choose to do so because of the prevalence of use of Adam as well as having a better understanding of how to tune learning rates and schedules for Adam. (2) Having CIAO automate more than Adam is desirable property from a usage perspective, but an undesirable one from the perspective of scientific comparison: it would be interesting future work to disentangle how much benefit comes from learning to optimize versus having more functionality. Section 4.5 shows some evidence that majority of the benefit comes from learning to optimize.

3.6 Training Details

Inner tasks were a combination of MLPs, CNNs, RNNs, and Transformers (Appendix H goes into more detail of the inner task distribution). Tasks had between 2 and 128 outer steps and had a median, mean, and max runtime of 149s, 170s, and 220s, respectively. Default settings of PPO were used, other than an maximum environment reuse of 4 (each rollout included in up to 4 batches for PPO training). Both the policy π and value function V used the same architecture, and our default architecture was a single hidden layer LSTM [14] with 256 hidden units. We trained all models for 16k PPO iterations resulting in less than 244k inner problems over the course of slightly under 512 GPU-days for our default model. Models were only optimized for final reward with discount factor $\gamma = 1$, but we also performed reward shaping [36] using intermediate validation losses to aid with optimization.

4 Results

Limitations. One difficulty when evaluating schedules (dynamic or otherwise) is that they optimize for a chosen completion time: a single run only provides an approximate upper or lower bound of the speedup and cannot give you an exact measurement of a speedup, because the latter would involve changing that number of updates. Measuring precise speedups would require an iterative search over many different end times for a schedule. As such, for most of our evaluations, we only compare to with the same and half as many iterations as baselines to verify $\geq 1x$ and $\geq 2x$ speedups, respectively.

4.1 On a Distribution of Small Problems

Our main development evaluation metric compares our trained models on almost 4000 tasks similar to those in the training distribution to a set of 35 baselines using a very simple metric: the fraction of tasks where the model gets a $\geq 2x$ speedup. We compute this metric by running each baseline with twice as many updates, and compare the model’s performance on every task to that of every baseline on that task. The baselines are all AdamW-based and combinations of 5 learning rates ($1e^{-4}$, $3e^{-4}$, $1e^{-3}$, $3e^{-3}$, $1e^{-2}$) and 7 commonly used schedules (constant, multi-step, linear decay, quadratic decay, exponential decay, cosine [25] to 0, cosine to 0.1 of original LR). Early on in development, we looked at the fraction of tasks where the model gets a $\geq 1x$ speedup (i.e., fraction of tasks where the model beats baselines), but the models quickly dominated all the baselines at that threshold.

Table 1 shows overall scores for several models, as well as scores for different subsets of tasks. While the metrics are meaningless for real world performance, they do show the relative strength of learned optimizers, as well as some amount of generalization: the models generalization to Transformers [45] on algorithmic tasks despite never being trained on them.

Limitations. (1) These tasks have some small similarity with those in the training distribution. Due to the large amount of randomness of our training distribution, it would be possible but very unlikely that any of the exact settings of our test set were sampled during training. (2) Some of the 35 baselines might be much weaker than one would use on a real problem - as such this metric almost certainly overestimates the rate at which our model would get $\geq 2x$ speedups even with perfect generalization. (3) This metric ignores how much better or worse the models are than baselines. (4) It is quite possible that a smaller subset of tasks provides the same information and this metric may be unnecessarily computationally expensive, costing approximately 64 V100-days.

Table 1: Fraction of test tasks where the model gets a $\geq 2x$ speedup

	Default	OnlyMNISTs	OnlyNQM	Large
Overall	0.691	0.631	0.294	0.725
MNISTs + MLP	0.443	0.479	0.073	0.457
MNISTs + CNN	0.777	0.829	0.095	0.811
Algo + RNN	0.817	0.637	0.469	0.893
Text + RNN	0.827	0.719	0.478	0.853
Algo + Transformer	0.572	0.580	0.340	0.599
Text + Transformer	0.713	0.543	0.309	0.738

4.2 GPT-2 on 1 Epoch of WikiText-103

We test our model’s ability to generalize to large language models by using it to train a 760M parameter GPT2-Large [40] model on a single epoch of WikiText-103 [32] using a standard HuggingFace [47] setup. Because even the largest language modeling tasks are trained for less than an epoch [8], we choose to train for only a single epoch to evaluate performance in an underfitting regime. For this task, AdamW was used as a baseline and grid search was used over both learning rate and schedule (over constant, linear, and cosine). The LHOPT was not tuned at all. Table 2 shows the relative performance with our LHOPT greatly outperforming at 1x time and almost showing a 2x speedup.

Figure 3 shows the learning curves for the LHOPTs and best baseline. An interesting observation that we will see repeated throughout the paper is that despite being capable of achieving a lower loss earlier, the chosen hyperparameters tend to underperform the best possible loss for that compute, presumably to achieve a better loss later. It’s unclear how necessary it is trade-off early performance for later, but many successful hand-made schedules tend to do this: multi-step schedules tend to stay at the same learning rate long after they’ve hit a plateau and cosine schedules tend to decay their learning rates much less aggressively than other commonly used schedules.

Limitations. Despite the LHOPT not being tuned at all, this evaluation was used a decent amount during development, and it’s possible that some implicit overfitting has occurred.

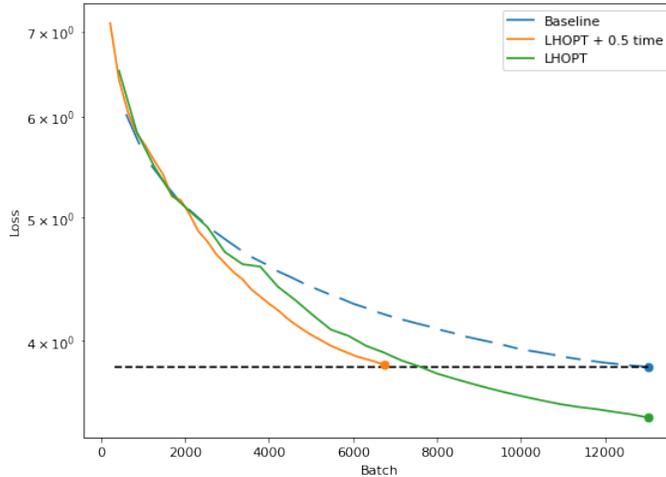


Figure 3: Performance of learned optimizers on optimizing 1 epoch of GPT2-Large on WikiText-103. Our learned optimizers get almost 2x speedups on this task despite being over 2 magnitudes larger than training tasks.

4.3 ResNets on ImageNet

We test our model’s ability to generalize to novel architectures by training ResNets [18] on ImageNet. Table 3 shows performance relative to an AdamW baseline tuned with grid search over both learning rate and schedule (over constant, cosine, and multi-step) and Table 5 also includes a tuned SGD

Table 2: Language modeling performance on WikiText-103

Model	Test Perplexity
GPT2-Large	45.6
+ LHOPT + 0.5 time	46.1
+ LHOPT	32.5
+ LHOPT + OM + 0.5 time	63.4
+ LHOPT + OM	40.9

Table 3: ImageNet performance of adaptive optimizers + SGD baselines for reference

Model	Epochs	Acc@1	Acc@5	Test Loss
Resnet18 + AdamW	90	67.32	87.52	1.39
Resnet18 + LHOPT	45	67.24	87.54	1.39
Resnet18 + LHOPT	90	68.89	88.43	1.31
Resnet18 + LHOPT + OM	45	66.54	87.09	1.41
Resnet18 + LHOPT + OM	90	68.34	88.54	1.31
Resnet50 + AdamW	90	71.42	89.66	1.35
Resnet50 + LHOPT	45	72.88	91.17	1.14
Resnet50 + LHOPT	90	73.52	91.38	1.07
Resnet18 + SGD	90	69.10	88.62	1.26
Resnet50 + SGD	90	74.66	92.09	1.01

baseline. Performance is reported on half of the validation set, where the other half is used for either tuning for the baseline or as an input validation loss for the LHOPT. Despite performing worse than SGD, our LHOPTs greatly outperform baselines resulting in an approximately 2x speedup despite having no ResNets on our training set.

Figure 4 shows learning curves for ResNet18 models, similarly showing learned have the capacity to achieve lower loss earlier in training.

Limitations. (1) Note that we choose to compare to AdamW instead of SGD because we constrain our optimizer to be adaptive. While past work has noted that a more general optimizer should be able to outperform one they can approximate [10], the bounds on our inner optimizer’s hyperparameter values prevent it from doing so (Appendix I). Appendix F also contains a comparison to LHOPTs using SGD as an inner optimizer. (2) AdamW is a weak baseline for ResNets on ImageNet.

4.4 Additional Out-of-distribution Evaluations

To additionally test the limits of generalization of our learned optimizers, we also evaluated them on 2 tasks in the MLPerf suite: the NeuMF model [19] for recommendation and Deep Speech 2 [2] for speech recognition. These 2 tasks were chosen as the most different from tasks in our training set which had an available PyTorch [38] implementation from the MLPerf training reference implementations.

4.4.1 Neural Collaborative Filtering

We trained both the Generalized Matrix Factorization (GMF) and Neural Matrix Factorization (NeuMF) models from scratch on the MovieLens 1M Dataset. Table 4 shows the results for both models. Interestingly, the LHOPT underperformed the baseline for the GMF models, but beat the baseline for the NeuMF model. This is the only instance we have found where a LHOPT underperforms a baseline. We hypothesize that this is due to our model’s behavior specializing in optimizing deep networks and the GMF model being shallow (no hidden nonlinearities). As an additional surprise, the LHOPT worked without any tuning on the NeuMF model, despite their default settings failing to train and their README recommending tuning weight decay in order to get it working.

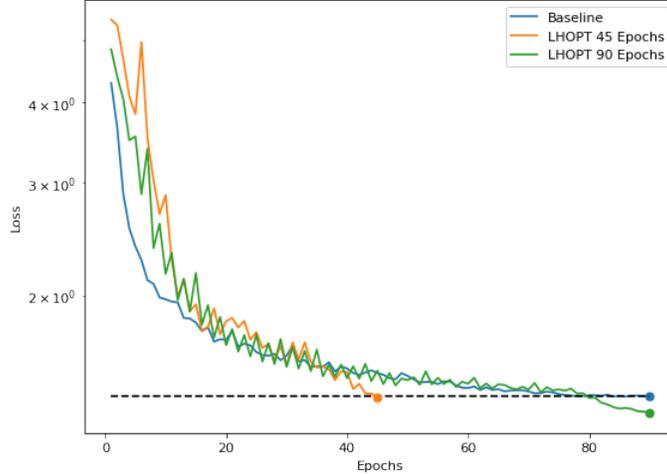


Figure 4: Test learning curve of LHOPTs and a tuned AdamW optimizing ResNet18 on ImageNet. Learned optimizers outperform the baseline adaptive optimizer despite never encountering residual architectures in its training set.

Table 4: NCF Performance. Larger is better for both metrics.

Model	NCDG	Hit Ratio
GMF	0.3677	0.6397
GMF + LHOPT	0.3134	0.5553
NeuMF	0.3859	0.6584
NeuMF + LHOPT	0.3932	0.6705

Figure 5 shows validation metrics over time where the performance plots show the learned optimizer training to look comparatively less stable. One potential benefit of using LHOPTs might be in being able to train with less stable dynamics without divergence.

Because their codebase doesn’t expose a evaluation loss similar to their training loss, we simply pass in negative NCDG as our validation loss to the LHOPT. The fact that this still works demonstrates the flexibility of our unitless features.

Limitations. (1) Despite having to do minimal tuning of weight decay to get the NeuMF model, we only compare to the baseline learning rates and do no tuning of learning rate or schedule of our own. We assume that being part of MLPerf’s training suite, the model has already been tuned, but we are unsure of how thoroughly. (2) The presented metrics are on the validation set, which the LHOPT could theoretically overfit to.

4.4.2 Speech Recognition

Figure 6 shows our LHOPT outperforming the baseline MLPerf optimization. This demonstrates generalization to a modality that our models have never been exposed to as well as working with half-precision floating point.

Limitations. (1) Similarly to NCF above, we only compare to a single baseline and do no tuning of our own. (2) The presented metrics are on the validation set, which the LHOPT could theoretically overfit to.

4.5 Fixed Schedules for Large Language Models

To explore another axis of generalization, we tested our models on a well-tuned language modeling codebase [8]. To avoid the software complexity of making our optimizer work in a distributed setting, we instead investigated the effectiveness of transferring a fixed hyperparameter schedule created by

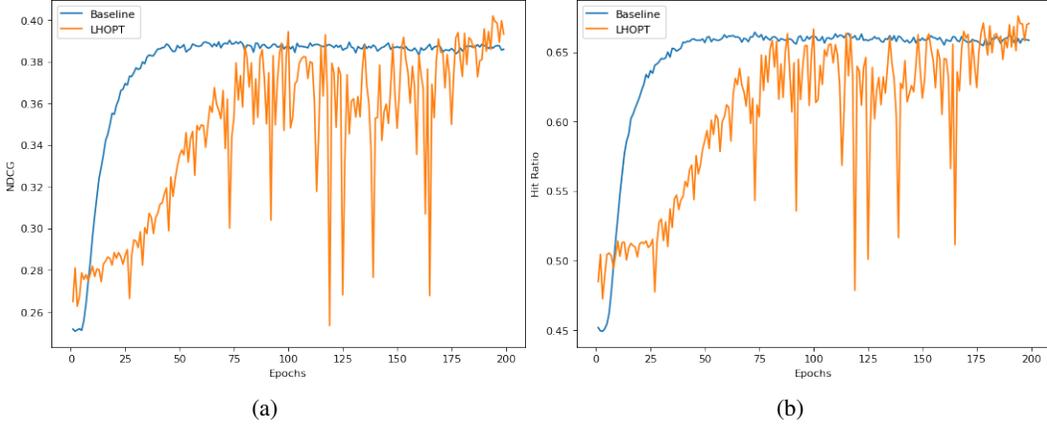


Figure 5: Comparison on NeuMF model of (5a) NDCG and (5b) hit ratio. These demonstrate generalization to out of distribution architecture, input modality, and loss function without tuning. Hyperparameters are updated 200 times (every epoch).

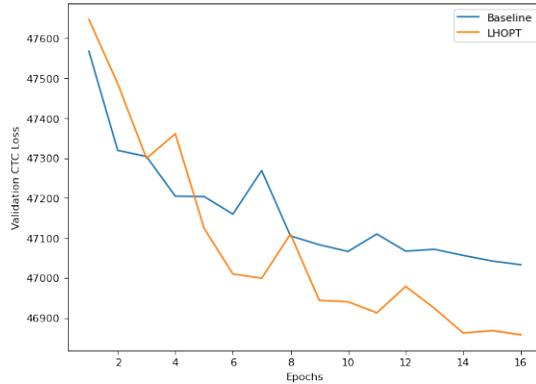


Figure 6: Loss curves for Deep Speech 2 on LibriSpeech. Similar to Figure 5, this plot demonstrates generalization of our learned optimizers to inner tasks with out of distribution size, input modality, loss function, and fp16 training without tuning. Hyperparameters are updated 16 times (every epoch) demonstrating robustness to outer step frequency.

our optimizer on a similar task at a smaller scale. Experiments on WikiText-103 seemed to indicate that majority of the benefit of LHOPTs could be had by running the optimizer once on a similar problem and using the resulting schedule retains the majority of performance.

We thus retrained a new LHOPT that uses an inner optimizer with the exact same hyperparameters as those available in the distributed codebase instead of CIAO with only 4 hyperparameters to control: learning rate, ϵ , β_2 , and weight decay. This model had slightly degraded performance compared to our defaults on the metric from Section 4.1: 0.673 for this LHOPT versus 0.691 for the default. We ran this model on the same setting as in Section 4.2, achieving a slightly worse perplexity of 37.7, and used the resulting schedule for the remainder of this section. See Appendix B for this schedule.

We then trained a range of model sizes to compute scaling laws [21] for both baselines and models trained with the LHOPT schedule and present the results on Figure 2. The LHOPT schedule demonstrates consistent speedup over baselines with a slightly steeper slope. We can estimate what a constant speedup would be for this range of points by assume their scaling law slopes are equal and from this calculate a 2.5x speedup. To take the change in slope into account as well, we extrapolate the curves to 175 billion parameters (same size as GPT-3) and at that size, the estimated speedup would be 3.6x.

Note that this result is despite the codebase doing multiple optimization techniques that our LHOPT would have no way of taking into account: gradient clipping to a fixed value and gradually increasing batch size.

Limitations. In order to run the schedule on this task, one important hyperparameter is how far we set the end of the schedule. Since we are comparing curves of the compute efficient frontier, we cannot simply set it to a single point that we care about. Instead, we choose to set that end point to twice the number of updates as that of the efficient point of the baseline. We found performance to be robust to how this end point was set, but given how our models tend to not optimize for performance before their end point, these frontiers almost certainly underestimate their performance.

4.6 How Important is Task Diversity?

Previous work has emphasized the importance of task diversity [34] in the context of learned parameter optimizers, and we would like to explore its importance for our generalization-first approach. To do so, we train additional LHOPTs on only MNIST-like datasets (OM) and on only NQMs [54]. Table 1 shows that training on NQMs only causes a large drop in performance but training on only MNISTs² recovers majority of that performance. Notably, we see the latter model generalize well to all other tasks. We test the OM model on the same setups as Sections 4.2 and 4.3 and include their results in Tables 2 and 3 respectively. On the language modeling task, the OM model does worse than the LHOPT trained on the full training set, but still much better than the baseline despite never being trained on anything other than 28x28 images and never being exposed to an attention mechanism, and on ImageNet, the OM does somewhat comparably to having the full dataset.

We believe that these results imply that for LHOPTs constructed in the way described, task diversity is not as important as it is for previous work, though these results are definitely not conclusive. Another possibility is that the flexibility in our reward function allowed us to include more randomness in task specification, and thus provided enough diversity from just those 3 MNIST-like datasets: there were over 50 billion different task settings we were sampling from, not including sampling initial hyperparameters.

4.7 Computational Overhead

Computational overhead for running the learned optimizer depends on the size for the inner problem. Running the policy network involves a constant amount of calculation to run the LSTM network for a fixed number of outer steps which would be insignificant for most real world problems, but the majority of computational overhead comes from computing features for the policy, which using our default setup would be proportional to the number of parameters of the inner problem. This overhead ranges from less than 2% on ResNets to 20% on Transformers, which could be reduced by recalculating inner-features less frequently - by default, we only update inner-feature statistics every 4 optimizer steps or simply using fixed schedules.

The memory overhead of our default setup is similar to Adam and other adaptive methods: 2 accumulators per parameter. There is additional constant memory cost to run the LSTM controller and keep track of the features we compute, but this is negligible for practical problems. This is significantly lower than existing learned optimizers which require many times the memory of Adam [34] due to having a per-parameter neural network running.

5 Conclusion

We present a novel approach for learning optimizers that prioritizes generalization above all. We demonstrate successful generalization over a variety of tasks very distinct from the training distribution including 2x speedups on ImageNet, 2.5x speedups on large scale language modeling, and outperforming tuned Adam baselines on a large set of practical neural network task.

The failures of the OnlyNQM LHOPT to generalize in Section 4.6 as well as the GMF task in Section 4.4.1 being the only task we could find where our LHOPT underperformed baselines lead us to believe that there is some commonality in optimizing specifically deep networks that our optimizers exploit. We believe that these results show a hopeful picture for learning optimizers that generalize.

²Note that MNISTs refer to the original MNIST [24], KMNIST [11], and Fashion-MNIST [49].

6 Future Work

We view this work as preliminary work on getting learned optimizers to generalize rather than a thorough analysis of all the components of our system. Many ablations could be done to better understand how impactful each of the individual design decisions were, such as applying our reward model and inner task distribution to existing learned parameter optimizers, viewing LHOPT performance with and without the constraints of unitless features and/or relative hyperparameter updates, analyzing the importance of initial hyperparameter randomness, and controlling for the effect of inner optimizer.

While this work focuses on a more standard training setup, we believe it would be incredibly useful to practitioners to extend it to other setups including but not limited to GANs, reinforcement learning, fine-tuning, and multi-loss training.

LHOPTs could also be used in combination with learned parameter optimizers: an outer-most LHOPT could control hyperparameters related to generalization and short horizon bias, which would allow learned parameter optimizers to greatly increase the per-step learning capacity, getting the best of both worlds.

We also believe that LHOPTs could greatly assist other subfields of research: there have been many promising optimization primitives emerging that so far have been unsuccessful but perhaps just need the right hyperparameters to get working such as SVRG [5], hypergradients [27], and second-order methods [17, 4, 28]. In fact, there could be many opportunities to improve performance using dynamic hyperparameter-controlled architectures [1]: hyperparameters need not be a bad thing if the user doesn't have to set them.

7 Acknowledgements

We would like to thank Luke Metz, Sander Dieleman, Alec Radford, Arthur Caillau, and Ryan Lowe for discussion related to this work and everyone at OpenAI, especially those with unused compute capacity that is made available for low priority jobs. We would like to especially thank Luke Metz, whose discussions and previous work in the field of learned optimizers inspired our generalization-first direction.

References

- [1] D. Almeida and N. Sauder. Gradnets: Dynamic interpolation between neural architectures. *arXiv preprint arXiv:1511.06827*, 2015.
- [2] D. Amodei, S. Ananthanarayanan, R. Anubhai, J. Bai, E. Battenberg, C. Case, J. Casper, B. Catanzaro, Q. Cheng, G. Chen, et al. Deep speech 2: End-to-end speech recognition in english and mandarin. In *International conference on machine learning*, pages 173–182. PMLR, 2016.
- [3] M. Andrychowicz, M. Denil, S. Gomez, M. W. Hoffman, D. Pfau, T. Schaul, B. Shillingford, and N. De Freitas. Learning to learn by gradient descent by gradient descent. *arXiv preprint arXiv:1606.04474*, 2016.
- [4] R. Anil, V. Gupta, T. Koren, K. Regan, and Y. Singer. Second order optimization made practical. *arXiv preprint arXiv:2002.09018*, 2020.
- [5] R. Babanezhad, M. O. Ahmed, A. Virani, M. Schmidt, J. Konečný, and S. Sallinen. Stop wasting my gradients: Practical svrg, 2015.
- [6] A. G. Baydin, R. Cornish, D. M. Rubio, M. Schmidt, and F. Wood. Online learning rate adaptation with hypergradient descent, 2018.
- [7] Y. Bengio. Practical recommendations for gradient-based training of deep architectures, 2012.
- [8] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray,

- B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei. Language models are few-shot learners, 2020.
- [9] K. Cho, B. van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation, 2014.
- [10] D. Choi, C. J. Shallue, Z. Nado, J. Lee, C. J. Maddison, and G. E. Dahl. On empirical comparisons of optimizers for deep learning, 2020.
- [11] T. Clauwat, M. Bober-Irizar, A. Kitamoto, A. Lamb, K. Yamamoto, and D. Ha. Deep learning for classical japanese literature. *arXiv preprint arXiv:1812.01718*, 2018.
- [12] C. Daniel, J. Taylor, and S. Nowozin. Learning step size controllers for robust neural network training. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- [13] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [14] F. A. Gers, J. Schmidhuber, and F. Cummins. Learning to forget: Continual prediction with lstm. 1999.
- [15] X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256. JMLR Workshop and Conference Proceedings, 2010.
- [16] R. Grazi, L. Franceschi, M. Pontil, and S. Salzo. On the iteration complexity of hypergradient computation, 2020.
- [17] V. Gupta, T. Koren, and Y. Singer. Shampoo: Preconditioned stochastic tensor optimization. In *International Conference on Machine Learning*, pages 1842–1850. PMLR, 2018.
- [18] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition, 2015.
- [19] X. He, L. Liao, H. Zhang, L. Nie, X. Hu, and T.-S. Chua. Neural collaborative filtering. In *Proceedings of the 26th international conference on world wide web*, pages 173–182, 2017.
- [20] M. Jaderberg, V. Dalibard, S. Osindero, W. M. Czarnecki, J. Donahue, A. Razavi, O. Vinyals, T. Green, I. Dunning, K. Simonyan, C. Fernando, and K. Kavukcuoglu. Population based training of neural networks, 2017.
- [21] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, and D. Amodei. Scaling laws for neural language models, 2020.
- [22] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [23] G. Leclerc and A. Madry. The two regimes of deep network training, 2020.
- [24] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [25] I. Loshchilov and F. Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016.
- [26] I. Loshchilov and F. Hutter. Fixing weight decay regularization in adam. 2018.
- [27] D. Maclaurin, D. Duvenaud, and R. P. Adams. Gradient-based hyperparameter optimization through reversible learning, 2015.
- [28] J. Martens and R. Grosse. Optimizing neural networks with kronecker-factored approximate curvature, 2020.

- [29] P. Mattson, C. Cheng, C. Coleman, G. Diamos, P. Micikevicius, D. Patterson, H. Tang, G.-Y. Wei, P. Bailis, V. Bittorf, et al. Mlperf training benchmark. *arXiv preprint arXiv:1910.01500*, 2019.
- [30] S. McCandlish, J. Kaplan, D. Amodei, and O. D. Team. An empirical model of large-batch training. *arXiv preprint arXiv:1812.06162*, 2018.
- [31] S. Merity. Single headed attention rnn: Stop thinking with your head, 2019.
- [32] S. Merity, C. Xiong, J. Bradbury, and R. Socher. Pointer sentinel mixture models, 2016.
- [33] L. Metz, N. Maheswaranathan, J. Nixon, D. Freeman, and J. Sohl-Dickstein. Understanding and correcting pathologies in the training of learned optimizers. In *International Conference on Machine Learning*, pages 4556–4565. PMLR, 2019.
- [34] L. Metz, N. Maheswaranathan, C. D. Freeman, B. Poole, and J. Sohl-Dickstein. Tasks, stability, architecture, and compute: Training more effective learned optimizers, and using them to train themselves, 2020.
- [35] T. Moskovitz, R. Wang, J. Lan, S. Kapoor, T. Miconi, J. Yosinski, and A. Rawal. First-order preconditioning via hypergradient descent, 2020.
- [36] A. Y. Ng, D. Harada, and S. Russell. Policy invariance under reward transformations: Theory and application to reward shaping. In *Icml*, volume 99, pages 278–287, 1999.
- [37] OpenAI, :, C. Berner, G. Brockman, B. Chan, V. Cheung, P. Dębiak, C. Dennison, D. Farhi, Q. Fischer, S. Hashme, C. Hesse, R. Józefowicz, S. Gray, C. Olsson, J. Pachocki, M. Petrov, H. P. d. O. Pinto, J. Raiman, T. Salimans, J. Schlatter, J. Schneider, S. Sidor, I. Sutskever, J. Tang, F. Wolski, and S. Zhang. Dota 2 with large scale deep reinforcement learning, 2019.
- [38] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *arXiv preprint arXiv:1912.01703*, 2019.
- [39] L. Pinto, M. Andrychowicz, P. Welinder, W. Zaremba, and P. Abbeel. Asymmetric actor critic for image-based robot learning, 2017.
- [40] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- [41] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms, 2017.
- [42] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- [43] L. N. Smith. Cyclical learning rates for training neural networks, 2017.
- [44] G. Tesauro. Temporal difference learning and td-gammon. *Communications of the ACM*, 38(3): 58–68, 1995.
- [45] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need, 2017.
- [46] O. Wichrowska, N. Maheswaranathan, M. W. Hoffman, S. G. Colmenarejo, M. Denil, N. Freitas, and J. Sohl-Dickstein. Learned optimizers that scale and generalize. In *International Conference on Machine Learning*, pages 3751–3760. PMLR, 2017.
- [47] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, et al. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*, 2019.
- [48] Y. Wu, M. Ren, R. Liao, and R. Grosse. Understanding short-horizon bias in stochastic meta-optimization, 2018.

- [49] H. Xiao, K. Rasul, and R. Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms, 2017.
- [50] C. Xu, T. Qin, G. Wang, and T.-Y. Liu. Reinforcement learning for learning rate control, 2017.
- [51] Z. Xu, A. M. Dai, J. Kemp, and L. Metz. Learning an adaptive learning rate schedule. *arXiv preprint arXiv:1909.09712*, 2019.
- [52] Y. You, I. Gitman, and B. Ginsburg. Large batch training of convolutional networks, 2017.
- [53] Y. You, J. Li, S. Reddi, J. Hseu, S. Kumar, S. Bhojanapalli, X. Song, J. Demmel, K. Keutzer, and C.-J. Hsieh. Large batch optimization for deep learning: Training bert in 76 minutes, 2020.
- [54] G. Zhang, L. Li, Z. Nado, J. Martens, S. Sachdeva, G. E. Dahl, C. J. Shallue, and R. Grosse. Which algorithmic choices matter at which batch sizes? insights from a noisy quadratic model, 2019.

Appendix A Scaling Laws for Large Language Models

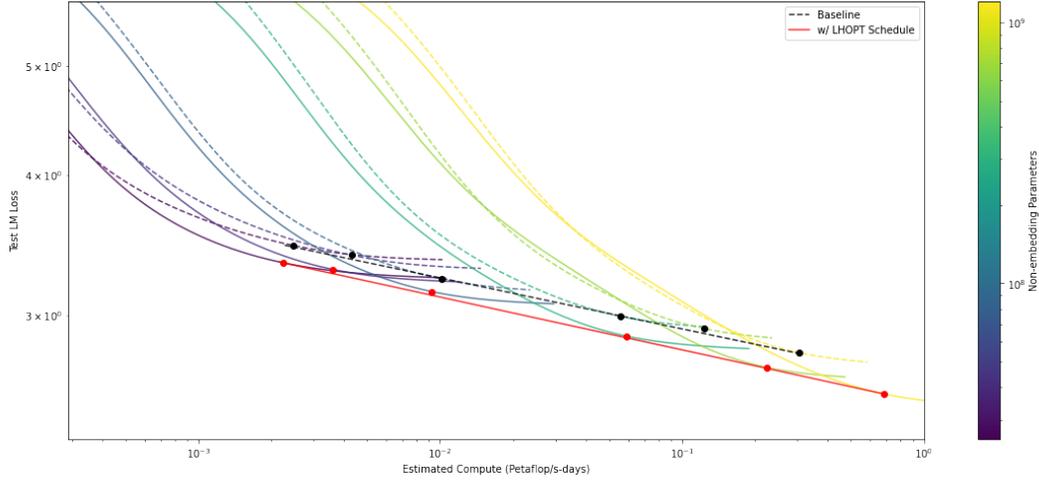


Figure 7: Same data as Figure 2 but also including parameter count colorbar.

Appendix B Fixed Schedule for Language Modeling

Figure 8 shows the values of the 4 hyperparameters over time used for all model sizes. β_1 was also available as a hyperparameter for the inner optimizer, but our experiments showed β_1 to be difficult to set robustly from small models - in particular, we had found that the ideal β_1 was a function of transformer context length and thus just kept it set at the default value (0.9).

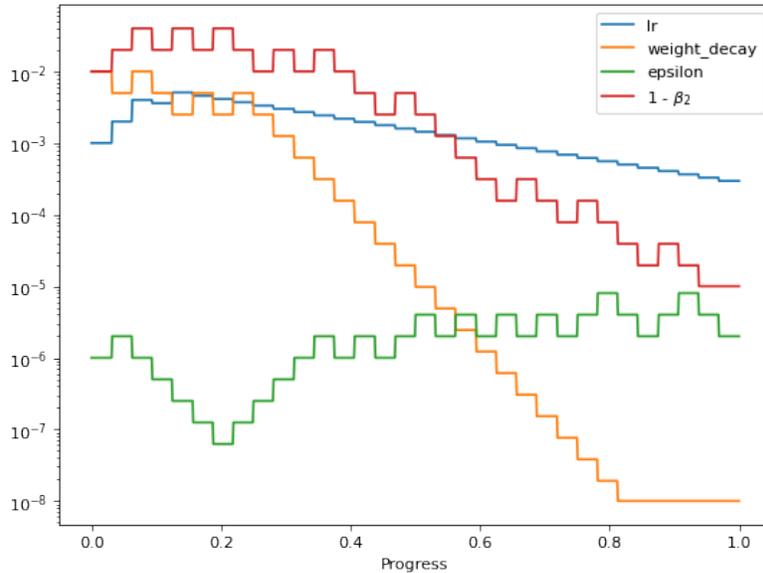


Figure 8: Hyperparameter settings for the fixed language modeling schedule.

Appendix C Why Use the LAMB Trust Ratio

Basing our inner optimizer on LAMB may be seen as an unusual choice, especially because the LAMB optimizer has mostly been advertised as for training with very large batches [53] while instead our inner training is on incredibly small models and batch sizes.

The reason we chose LAMB is, like many other of our design decisions, for generalization. For optimizers in the family as Adam [22], the numerator and denominator cancel out gradient scale and we are left with updates of approximately the same size as the learning rate. The issue is that past research has found that parameter scale is quite important to optimization [15] and that this scale is a function of network width. Reasonably, past work has also chosen to set learning rate as a function of model size [21].

These issues combined make it difficult for learned optimizers to robustly set learning rates for Adam-like optimizers because it would involve them having to learn to extrapolate beyond the training set. LAMB-like optimizers do not have this issue.

While the above is the main reason for using LAMB, there are several other benefits as well: (1) Even if one set the learning rate optimally based on the width of the network, there are some parameters, such as biases and normalization scales, which don't have this property and setting learning rate in this way effectively slows down the learning of these parameters as model scale increases. (2) Having hyperparameters behave independently is a useful property for being able to control optimizer behavior, and while past research has noted the importance of ϵ for Adam [10], extremely large learning rates need to be used with it. For Adam, increasing both learning rate and ϵ together allow the latter to dominate the second moment term in the denominator and simply behave like SGD-M. For LAMB, ϵ can be set independently, and as it approaches infinity, the optimizer updates simply approach LARS [52] - no learning rate tuning needed.

Appendix D Do LHOPTs Overfit The Validation Set?

Since they take in validation loss, one potential concern is overfitting that validation set and requiring an additional set for evaluating models trained with LHOPTs. In order to investigate this, we look at the distributions of the relative difference between test and validation scores for ResNet18. This dataset-model combination was chosen because we had the most data for baseline and LHOPT runs. We find that despite theoretically being able to overfit the validation set, these models generally do not for loss ($p = 0.14$) (see Figure 9), top-1 accuracy ($p = 0.68$), or top-5 accuracy ($p = 0.36$) and the relative differences generally appear to be from the same distribution. For reference, comparing AdamW to SGD shows that SGD overfits significantly more at all 3 metrics: loss ($p = 1.49 \times 10^{-22}$), top-1 accuracy ($p = 7.32 \times 10^{-7}$), and top-5 accuracy ($p = 0.04$).

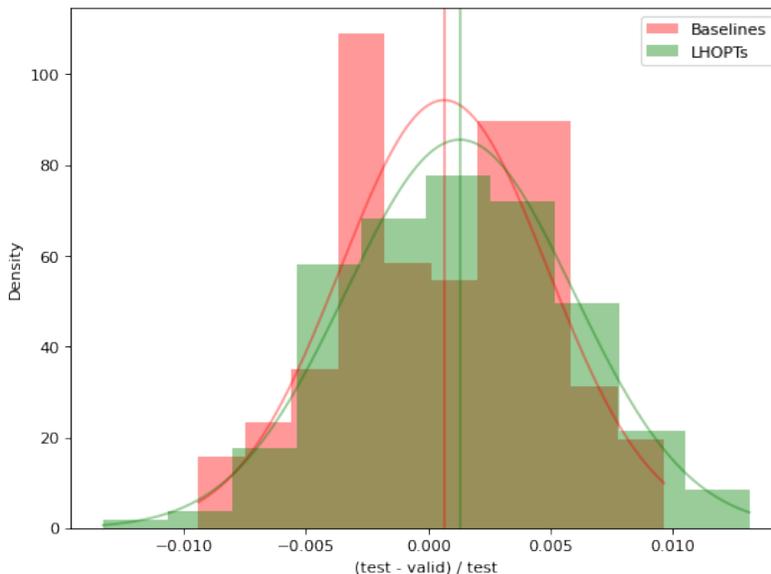


Figure 9: Showing the distribution of relative difference between test and validation loss on ResNet18. We find that LHOPTs do not significantly overfit to the validation set.

One limitation of this analysis is that we look at the entire curve, and as we have noted, our models tend to solely optimize for final performance, and thus the overfitting might be crowded out by not

doing so before the final point(s). This may still be possible, but the p -values are even higher when looking at the final data points ($p > 0.8$), but this may be due to having much fewer data points for comparing distributions.

Note that the p -values are similar if we look at absolute difference between test and valid metrics.

Appendix E Integral CDF Features

Given a mean μ and standard deviation σ , it's a single line of code to convert a value into a z-score and apply a cdf to that z-score (As mentioned in Section 3.2). One issue is in calculating robust μ and σ . For example, if one used a global μ and σ , these statistics would change faster at the beginning of inner training, and slower at the end. Another possibility is using an exponential moving μ and σ , which solves the previous problem but introduces a new problem: that how quickly the statistics react is a function of how frequently the statistics are updated. For example, a point half-way through training would have much more weight on the final μ and σ if we update the statistics twice as frequently (i.e., we double the number of outer steps). Our final solution to this is what we call Integral CDF Features: we replace the discrete sampling of points at each outer step by linearly interpolating between them, assign the weight of each point with the continuous version of an exponential moving average. Specifically, given observations y_1, y_2, \dots, y_N at progress values of t_1, t_2, \dots, t_N , we would calculate weighted average μ as:

$$\mu = \frac{\sum_{n=1}^{N-1} \int_{t=t_n}^{t_{n+1}} e^{t \cdot \log(b)} \left[y_n \frac{t-t_n}{t_{n+1}-t_n} + y_{n+1} \frac{t_{n+1}-t}{t_{n+1}-t_n} \right] dt}{\int_{t=t_1}^{t_N} e^{t \cdot \log(b)} dt}$$

And using the same weights for σ . b here is specified as a hyperparameter corresponding to how much more weight a feature at the end of inner training ($t = 1$) has than one at the beginning of training ($t = 0$). Rather than tune b , we instead make a copy for each feature for several different values of b . For this paper, we use $b \in [1.25, 2.5, 5, 10, 20]$.

Appendix F ImageNet w/ SGD

When we first started this project, our earlier models were inner optimizer-agnostic and as part of the task, we also chose a random inner optimizer between SGD/Adam/Adamax. These models actually had very strong performance, outperforming our newer ones on ImageNet (Table 5) and beating learning rate and schedule tuned SGD on the task.

The reason we moved away from these models and instead chose to specialize on an inner optimizer was because this allow us to control more hyperparameters, and the problems that require the largest use of compute generally use adaptive optimizers.

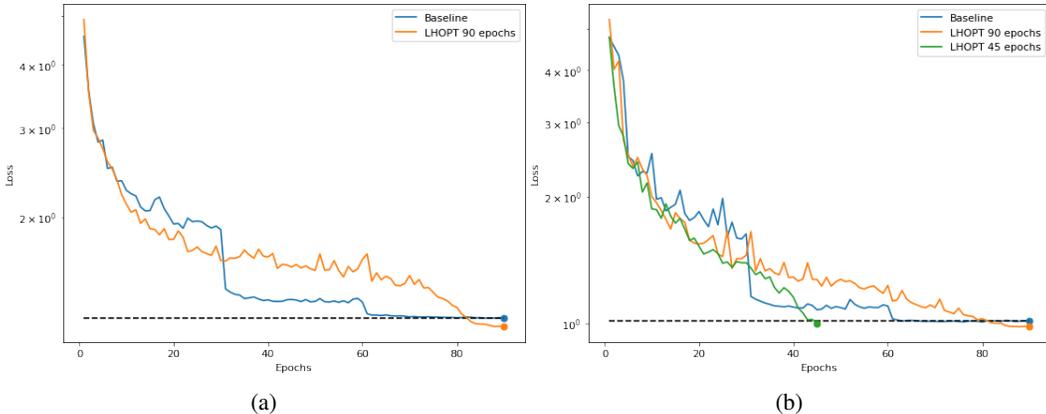


Figure 10: Test learning curve of a tuned SGD baseline and older versions of our LHOPT that were trained to be robust to inner optimizer on (10a) Resnet18 and (10b) ResNet50.

Table 5: Imagenet Performance using SGD.

Model	Epochs	Acc@1	Acc@5	Loss
Resnet18	90	69.10	88.62	1.26
Resnet18 + LHOPT	90	69.97	89.31	1.21
Resnet50	90	74.66	92.09	1.01
Resnet50 + LHOPT	45	74.79	92.18	1.00
Resnet50 + LHOPT	90	75.50	92.49	0.98

Table 6: Inner Task Distribution.

	Ratio
Algo + RNN	0.27
MNISTs + CNN	0.068
MNISTs + MLP	0.2
NQM	0.014
Text + RNN	0.14
Text + Transformer	0.31

Appendix G How Important are Larger Models?

One question we had was if scaling up our controller would be beneficial, since our default one was so small. We show results on Table 1 for a large version of the controller that uses 2 hidden LSTM layers and increasing with to 512. We can see from that table that the larger model performs better on the distribution of tasks than the default, but from the limited testing we performed, the larger model seemed to perform quite similarly on real-world tasks and thus don't include the results. One possible reason for this could be that the larger models could be doing better at tasks with unusual settings in the task distribution, but doing better at those may not apply to real world tasks.

Appendix H Inner Task Distribution

Table 6 shows the breakdown of our training distribution. None of the task ratios were tuned - the unusual values came from incrementally adding more datasets to the mix over time. Performance rarely improved when adding new datasets, but we chose to do so anyway because our prior is that they would be helpful. Section 4.6 was our attempt at answering whether or not it was important to do so retrospectively, and it seems like the additional inner tasks were not very helpful.

The full DSL for our training distribution will be made available, but we provide some additional high-level details: RNN in the table refers to randomly using a simple RNN, LSTM, or GRU [9]. For CNNs and MLPs, we randomize over 5 different activation functions (ReLU, leaky ReLU, very leaky ReLU, ELU, PReLU) and normalizations (none, BN, LN). For MLPs, we additionally randomize over loss functions (CCE, MAE, MSE, Huber). Algo + RNN is on XOR and binary addition tasks. Text + RNN is classification on the SST dataset. Text + Transformer is language modeling over 31 different simple datasets.

Appendix I Initial Inner Hyperparameters, Noise, and Actions

I.1 Initial Inner Hyperparameters

Table 7 shows the initial hyperparameter. Note that we normally represent moving average rates as $1 - \beta$ to make multiplicative scaling actions easier to apply. We also include several novel hyperparameters:

1. Grad clip fraction: instead of gradient clipping to an absolute value (which cannot be robust to task), we instead clip gradients as a fraction of an exponentially weighted moving maximum of gradient norms.

2. $\beta_{gradclip}$: this is the rate at which the maximum decays: $G_t = \max(G_{t-1} * \beta_{gradclip}, ||grad_t||)$
3. β_{LAMB} : Instead of normalizing by the update norm in LAMB, we instead normalize by an exponential moving average of the update norm - which allows us to recover relative update norm information.

Table 7: Initial Inner Hyperparameters.

Hyperparameter	Value
Learning rate	$1e^{-3}$
$1 - \beta_1$	0.1
$1 - \beta_2$	$1e^{-2}$
ϵ	$1e^{-6}$
Weight decay	$1e^{-2}$
Grad clip fraction	0.8
$1 - \beta_{gradclip}$	$1e^{-2}$
Denominator	Adamax
Use LAMB	True
LAMB minimum trust [31]	$1e^{-3}$
$1 - \beta_{LAMB}$	0.05

I.2 Actions and Noise

In addition to the actions in Table 8, we also had restart actions corresponding to 3 checkpoints with 3 actions per checkpoint: save current state to the checkpoint, load state at the checkpoint, and swap states between current and checkpoint. This results in a total of 10 restart actions ($3 * 3 + 1$ for doing nothing).

In the table, noise is represented by an upper and lower bound that is applied by the corresponding action type. For example, initial learning rate is scaled by some value sampled in $[1e^{-2}, 1e^2]$ when the task is sampled.

Table 8: Actions and Noise.

Hyperparameter	Action Type	Action Values	Noise
Learning rate	Scale	[0.5, 0.707, 0.9, 1, 1.1, 1.414, 2.0]	$[1e^{-2}, 1e^2]$
Weight decay	Scale	[0.5, 2.0]	[0.1, 10]
ϵ	Scale	[0.5, 2.0]	[0.1, 10]
$1 - \beta_1$	Scale	[0.5, 2.0]	[0.1, 10]
$1 - \beta_2$	Scale	[0.5, 2.0]	[0.1, 10]
Grad clip fraction	Logit Shift	[-1, -0.3, 0.3, 1]	[-1, 1]
$1 - \beta_{gradclip}$	Scale	[0.5, 1.0, 2.0]	[0.5, 2]
$1 - \beta_{LAMB}$	Scale	$[\frac{1}{1.5}, 1.0, 1.5]$	[0.5, 2]

Appendix J π and V Features

J.1 π Features

For some features, we will mention transforms applied to them. `is_nan` refers to a feature on whether or not that value is NaN. NaN's are all converted to a default value to avoid passing them in to the controller. CDF features refers to the 5 features in Appendix E.

Features updated at each outer step:

1. Progress $\in [0, 1]$
2. Previous action representation

3. Log loss ratio (train and valid): is_nan, tanh, CDF features
4. Log train loss: is_nan, is_inf, CDF features
5. Log valid loss: is_nan, is_inf, CDF features
6. Percentile of loss at each checkpoint
7. Progress at each checkpoint
8. For all hyperparameter actions: current value / initial value
9. Log ratio of previous and current param norm: tanh, CDF features
10. Log ratio of update and previous param norm: tanh, CDF features

Features updated every 4 inner steps (for all of these features, we include the raw value as well as the CDF features):

1. Fraction of gradients clipped
2. Fraction of $\sqrt{\hat{v}} \geq \epsilon$
3. Average per-parameter update magnitude (pre-LR)
4. Log scale from [23]
5. Log ratio of update norm and param norm
6. Log of noise scale [30]
7. Log LAMB trust ratio
8. Cosine similarity of gradient and momentum
9. Logit of CDF cosine similarity of gradient and momentum
10. Cosine similarity of gradient and update
11. Logit of CDF cosine similarity of gradient and update
12. Cosine similarity of gradient and parameter
13. CDF cosine similarity of gradient and parameter

One issue with cosine similarity is that it is not robust to vector size: the expected cosine similarity between random vectors is a function of their size. In order to make this more robust, we additionally introduce a "CDF cosine similarity", where we apply a gaussian CDF to a cosine similarity times the square root of the dimensionality of the vector. We additionally apply the logit function on top of this when we expect the value to be very close to 1.

J.2 V Features

In addition to all the features for π , we pass in the following additional features to our value function:

1. A vector representation of the task: all high-level task choices are converted into a set of values. For example, if we allow between 1 and 7 layers for a MLP and the task chooses 4, this would be encoded as a one-hot vector.
2. A vector representation of the random initial hyperparameters
3. Reward
4. Log train and eval loss
5. Statistics about the reward baseline optimizer's losses (From Section 3.4)