
Learned Optimizers that Scale and Generalize

Olga Wichrowska¹ Niru Maheswaranathan^{2,3} Matthew W. Hoffman⁴ Sergio Gómez Colmenarejo⁴
Misha Denil⁴ Nando de Freitas⁴ Jascha Sohl-Dickstein¹

Abstract

Learning to learn has emerged as an important direction for achieving artificial intelligence. Two of the primary barriers to its adoption are an inability to scale to larger problems and a limited ability to generalize to new tasks. We introduce a learned gradient descent optimizer that generalizes well to new tasks, and which has significantly reduced memory and computation overhead. We achieve this by introducing a novel hierarchical RNN architecture, with minimal per-parameter overhead, augmented with additional architectural features that mirror the known structure of optimization tasks. We also develop a meta-training ensemble of small, diverse, optimization tasks capturing common properties of loss landscapes. The optimizer learns to outperform RMSProp/ADAM on problems in this corpus. More importantly, it performs comparably or better when applied to small convolutional neural networks, despite seeing no neural networks in its meta-training set. Finally, it generalizes to train Inception V3 and ResNet V2 architectures on the ImageNet dataset for thousands of steps, optimization problems that are of a vastly different scale than those it was trained on.

1. Introduction

Optimization is a bottleneck for almost all tasks in machine learning, as well as in many other fields, including engineering, design, operations research, and statistics. Advances in optimization therefore have broad impact. Historically, optimization has been performed using hand-designed algorithms. Recent results in machine

learning show that, given sufficient data, well-trained neural networks often outperform hand-tuned approaches on supervised tasks. This raises the tantalizing possibility that neural networks may be able to outperform hand-designed optimizers.

Despite the promise in this approach, previous work on learned RNN optimizers for gradient descent has failed to produce neural network optimizers that generalize to new problems, or that continue to make progress on the problems for which they were meta-trained when run for large numbers of steps (see Figure 2). Current neural network optimizers are additionally too costly in both memory and computation to scale to larger problems.

We address both of these issues. Specifically, we improve upon existing learned optimizers by:

1. Developing a meta-training set that consists of an ensemble of small tasks with diverse loss landscapes
2. Introducing a hierarchical RNN architecture with lower memory and compute overhead, and which is capable of capturing inter-parameter dependencies.
3. Incorporating features motivated by successful hand-designed optimizers into the RNN, so that it can build on existing techniques. These include dynamically adapted input and output scaling, momentum at multiple time scales, and a cross between Nesterov momentum and RNN attention mechanisms.
4. Improving the meta-optimization pipeline, for instance by introducing a meta-objective that better encourages exact convergence of the optimizer, and by drawing the number of optimization steps during training from a heavy tailed distribution.

2. Related work

Learning to learn has a long history in psychology (Ward, 1937; Harlow, 1949; Kehoe, 1988; Lake et al., 2016). Inspired by it, machine learning researchers have proposed meta-learning techniques for optimizing the process of learning itself. Schmidhuber (1987), for example, considers networks that are able to modify their own weights.

¹Google Brain ²Work done during an internship at Google Brain. ³Stanford University ⁴Deepmind. Correspondence to: Olga Wichrowska <olganw@google.com>.

This leads to end-to-end differentiable systems which allow, in principle, for extremely general update strategies to be learned. There are many works related to this idea, including (Sutton, 1992; Naik & Mammone, 1992; Thrun & Pratt, 1998; Hochreiter et al., 2001; Santoro et al., 2016).

A series of papers from Bengio et al. (1990; 1992; 1995) presents methods for learning parameterized local neural network update rules that avoid back-propagation. Runarsson & Jonsson (2000) extend this to more complex update models. The result of meta learning in these cases is an algorithm, i.e. a local update rule.

Andrychowicz et al. (2016) learn to learn by gradient descent by gradient descent. Rather than trying to distill a global objective into a local rule, their work focuses on learning how to integrate gradient observations over time in order to achieve fast learning of the model. The component-wise structure of the algorithm allows a single learned algorithm to be applied to new problems of different dimensionality. While Andrychowicz et al. (2016) consider the issue of transfer to different datasets and model structures, they focus on transferring to problems of the same class. In fact, they report negative results when transferring optimizers, meta-trained to optimize neural networks with logistic functions, to networks with ReLU functions.

Li & Malik (2017) proposed an approach similar to Andrychowicz et al. (2016), around the same time, but they rely on policy search to compute the meta-parameters of the optimizer. That is, they learn to learn by gradient descent by reinforcement learning.

Zoph & Le (2017) also meta-train a controller RNN, but this time to produce a string in a custom domain specific language (DSL) for describing neural network architectures. An architecture matching the produced configuration (the “child” network) is instantiated and trained in the ordinary way. In this case the meta-learning happens only at the network architecture level.

Ravi & Larochelle (2017) modify the optimizer of Andrychowicz et al. (2016) for 1 and 5-shot learning tasks. They use test error to optimize the meta learner. These tasks have the nice property that the recurrent neural networks only need to be unrolled for a small number of steps.

Wang et al. (2016) show that it is possible to learn to solve reinforcement learning tasks by reinforcement learning. They demonstrate their approach on several examples from the bandits and cognitive science literature. A related approach was proposed by Duan et al. (2016).

Finally, Chen et al. (2016) also learn reinforcement learning, but by supervised meta-training of the meta-learner. They apply their methods to black-box function optimiza-

tion tasks, such as Gaussian process bandits, simple low-dimensional controllers, and hyper-parameter tuning.

3. Architecture

At a high level, a hierarchical RNN is constructed to act as a learned optimizer, with its architecture matched to the parameters in the target problem. The hierarchical RNN’s parameters (called meta-parameters) are shared across all target problems, so despite having an architecture that adapts to the target problem, it can be applied to new problems. At each optimization step, the learned optimizer receives the gradients for every parameter along with some additional quantities derived from the gradients, and outputs an update to the parameters. Figure 1 gives an overview.

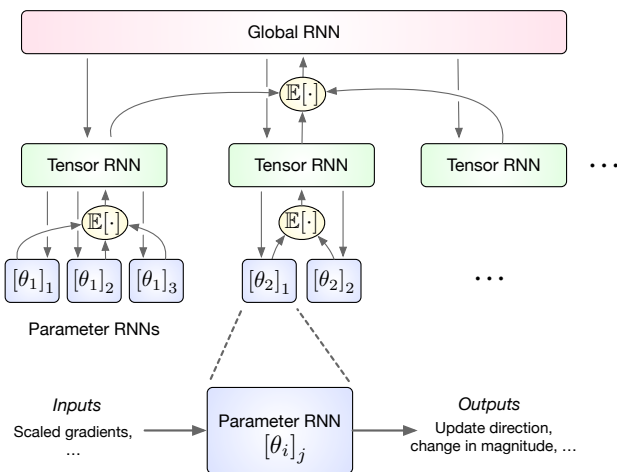


Figure 1. Hierarchical RNN architecture. At the lowest level, a small Parameter RNN processes the inputs and outputs (Section 3.3) for every parameter ($\theta_{i,j}$) in the target problem. At the intermediate level, a medium-sized Tensor RNN exists for every parameter tensor (denoted by θ_i) in the target problem. It takes as input the average latent state across all Parameter RNNs belonging to the same tensor. Its output enters those same Parameter RNNs as a bias term. At the top level, a single Global RNN receives as input the average hidden state of all Parameter RNNs, and its output enters the Tensor RNNs as a bias term and is added to the Parameter RNN bias term. This architecture has low per-parameter overhead, while the Tensor RNNs are able to capture inter-parameter dependencies, and the Global RNN is able to capture inter-tensor dependencies.

3.1. Hierarchical architecture

In order to effectively scale to large problems, the optimizer RNN must stay quite small while maintaining enough flexibility to capture inter-parameter dependencies that shape the geometry of the loss surface. Optimizers that account for this second order information are often particularly effective (e.g. quasi-Newton approaches). We propose a novel hierarchical architecture to enable both low per-

parameter computational cost, and aggregation of gradient information and coordination of update steps across parameters (Figure 1). At the lowest level of the hierarchy, we have a small *Parameter RNN* that receives direct per-parameter (scalar) gradient inputs. One level up, we have an intermediate *Tensor RNN* that incorporates information from a subset of the Parameter RNNs (where the subsets are problem specific). For example, consider a feedforward fully-connected neural network. There would be a Tensor RNN for each layer of the network, where each layer contains an $(n \times m)$ weight matrix and therefore nm Parameter RNNs.

At the highest level of the hierarchy is a *Global RNN* which receives output from every Tensor RNN. This allows the Parameter RNN to have very few hidden units with larger Tensor and Global RNNs keeping track of problem-level information. The Tensor and Global RNNs can also serve as communication channels between Parameter and Tensor RNNs respectively. The Tensor RNN outputs are fed as biases to the Parameter RNN, and the new parameter state is averaged and fed as input to the Tensor RNN. Similarly, the Global RNN state is fed as a bias to each Tensor RNN, and the output of the Tensor RNNs is averaged and fed as input to the Global RNN (Figure 1).

The architecture used in the experimental results has a Parameter RNN hidden state size of 10, and a Tensor and Global RNN state size of 20 (the architecture used by Andrychowicz et al. (2016) had a two layer RNN for each parameter, with 20 units per layer). These sizes showed the best generalization to ConvNets and other complex test problems. Experimentally, we found that we could make the Parameter RNN as small as 5, and the Tensor RNN as small as 10 and still see good performance on most problems. We also found that the performance decreased slightly even on simple test problems if we removed the Global RNN entirely. We used a GRU architecture (Cho et al., 2014) for all three of the RNN levels.

3.2. Features inspired by optimization literature

The best performing neural networks often have knowledge about task structure baked into their design. Examples of this include convolutional models for image processing (Krizhevsky et al., 2012; He et al., 2016), causal models (RNNs) for modeling causal time series data, and the merging of neural value functions with Monte Carlo tree search in AlphaGo (Silver et al., 2016).

We similarly incorporate knowledge of effective strategies for optimization into our network architecture. We emphasize that these are not arbitrary design choices. The features below are motivated by results in optimization and recurrent network literature. They are also individually important to the ability of the learned optimizer to generalize

to new problems, as is illustrated by the ablation study in Section 5.5 and Figure 6.

Let $L(\theta)$ be the loss of the target problem, where $\theta = \{\theta_1, \dots, \theta_{N_T}\}$ is the set of all parameter tensors θ_t (e.g. all weight matrices and bias vectors in a neural network). At each training iteration n , each parameter tensor t is updated as $\theta_t^{n+1} = \theta_t^n + \Delta\theta_t^n$, where the update step $\Delta\theta_t^n$ is set by the learned optimizer (Equation 5 below).

3.2.1. ATTENTION AND NESTEROV MOMENTUM

Nesterov momentum (Nesterov, 1983a) is a powerful optimization approach, where parameter updates are based not on the gradient evaluated at the current iterate θ^n , but rather at a location ϕ^n which is extrapolated ahead of the current iterate. Similarly, attention mechanisms have proven extremely powerful in recurrent translation models (Bahdanau et al., 2015), decoupling the iteration n of RNN dynamics from the observed portion of the input sequence. Motivated by these successes, we incorporate an attention mechanism that allows the optimizer to explore new regions of the loss surface by computing gradients away (or ahead) from the current parameter position. At each training step n the attended location is set as $\phi_t^{n+1} = \theta_t^n + \Delta\phi_t^n$, where the offset $\Delta\phi_t^n$ is further described by Equation 6 below. Note that the attended location is an offset from the previous parameter location θ^n rather than the previous attended location ϕ^n .

The gradient \mathbf{g}^n of the loss $L(\theta)$ with respect to the attended parameter values ϕ^n will provide the only input to the learned optimizer, though it will be further transformed before being passed to the hierarchical RNN. For every parameter tensor t , $\mathbf{g}_t^n = \frac{\partial L}{\partial \phi_t^n}$.

3.2.2. MOMENTUM ON MULTIPLE TIMESCALES

Momentum with an exponential moving average is typically motivated in terms of averaging away minibatch noise or high frequency oscillations, and is often a very effective feature (Nesterov, 1983b; Tseng, 1998). We provide the learned optimizer with exponential moving averages $\bar{\mathbf{g}}_{ts}$ of the gradients on several timescales, where s indexes the timescale of the average. The update equation for the moving average is

$$\bar{\mathbf{g}}_{ts}^{n+1} = \bar{\mathbf{g}}_{ts}^n \sigma(\beta_{gt}^n)^{2^{-s}} + g_t^n \left(1 - \sigma(\beta_{gt}^n)^{2^{-s}}\right), \quad (1)$$

where the σ indicates the sigmoid function, and where the momentum logit β_{gt}^n for the shortest $s = 0$ timescale is output by the RNN, and the remaining timescales each increase by a factor of two from that baseline.

By comparing the moving averages at multiple timescales, the learned optimizer has access to information about how rapidly the gradient is changing with training time (a mea-

sure of loss surface curvature), and about the degree of noise in the gradient.

3.2.3. DYNAMIC INPUT SCALING

We would like our optimizer to be invariant to parameter scale. Additionally, RNNs are most easily trained when their inputs are well conditioned, and have a similar scale as their latent state. In order to aid each of these goals, we rescale the average gradients in a fashion similar to what is done in RMSProp (Tieleman & Hinton, 2012), ADAM (Kingma & Ba, 2015), and SMORMS3 (Funk, 2015),

$$\lambda_{ts}^{n+1} = \lambda_{ts}^n \sigma(\beta_{\lambda t}^n)^{2^{-s}} + (\bar{\mathbf{g}}_{ts}^n)^2 \left(1 - \sigma(\beta_{\lambda t}^n)^{2^{-s}}\right) \quad (2)$$

$$\mathbf{m}_{ts}^n = \frac{\bar{\mathbf{g}}_{ts}^n}{\sqrt{\lambda_{ts}^n}}, \quad (3)$$

where λ_{ts}^n is a running average of the square *average* gradient, \mathbf{m}_{ts}^n is the scaled averaged gradient, and the momentum logit $\beta_{\lambda t}^n$ for the shortest $s = 0$ timescale will be output by the RNN, similar to how the timescales for momentum are computed in the previous section.

It may be useful for the learned optimizer to have access to how gradient magnitudes are changing with training time. We therefore provide as further input a measure of *relative* gradient magnitudes at each averaging scale s . Specifically, we provide the relative log gradient magnitudes,

$$\gamma_{ts}^n = \log \lambda_{ts}^n - \mathbb{E}_s[\log \lambda_{ts}^n]. \quad (4)$$

3.2.4. DECOMPOSITION OF OUTPUT INTO DIRECTION AND STEP LENGTH

Another aspect of RMSProp and ADAM is that the learning rate corresponds directly to the characteristic step length. This is true because the gradient is scaled by a running estimate of its standard deviation, and after scaling has a characteristic magnitude of 1. The length of update steps therefore scales linearly with the learning rate, but is invariant to any scaling of the gradients.

We enforce a similar decomposition of the parameter updates into update directions \mathbf{d}_{θ}^n and \mathbf{d}_{ϕ}^n for parameters and attended parameters, with corresponding step lengths $\exp(\eta_{\theta}^n)$ and $\exp(\eta_{\phi}^n)$,

$$\Delta \theta_t^n = \exp(\eta_{\theta t}^n) \frac{\mathbf{d}_{\theta t}^n}{\|\mathbf{d}_{\theta t}^n\|/N_t}, \quad (5)$$

$$\Delta \phi_t^n = \exp(\eta_{\phi t}^n) \frac{\mathbf{d}_{\phi t}^n}{\|\mathbf{d}_{\phi t}^n\|/N_t}, \quad (6)$$

where N_t is the number of elements in the parameter tensor θ_t . The directions $\mathbf{d}_{\theta t}^n$ and $\mathbf{d}_{\phi t}^n$ are read directly out of the RNN (though see A.1 for subtleties).

Relative learning rate We want the performance of the optimizer to be invariant to parameter scale. This requires that the optimizer judge the correct step length from the history of gradients, rather than memorizing the range of step lengths that were useful in its meta-training ensemble. The RNN therefore controls step length by outputting a multiplicative (additive after taking a logarithm) change, rather than by outputting the step length directly,

$$\eta_{\theta}^{n+1} = \Delta \eta_{\theta}^n + \bar{\eta}_{\theta}^{n+1}, \quad (7)$$

$$\bar{\eta}_{\theta}^{n+1} = \gamma \bar{\eta}_{\theta}^n + (1 - \gamma) \eta_{\theta}^{n+1}, \quad (8)$$

where for stability reasons, the log step length η_{θ}^n is specified relative to an exponential running average $\bar{\eta}_{\theta}^n$ with meta-learned momentum γ . The attended parameter log step length η_{θ}^n is related to η_{θ}^n by a meta-learned constant offset c ,

$$\eta_{\phi}^n = \eta_{\theta}^n + c. \quad (9)$$

To further force the optimizer to dynamically adapt the learning rate rather than memorizing a learning rate trajectory, the learning rate is initialized from a log uniform distribution from 10^{-6} to 10^{-2} . We emphasize that the RNN has no direct access to the learning rate, so it must adjust it based purely on its observations of the statistics of the gradients.

In order to aid in coordination across parameters, we do provide the RNN as an input the *relative* log learning rate of each parameter, compared to the remaining parameters, $\eta_{\text{rel}}^n = \eta_{\theta}^n - \mathbb{E}_{ti}[\eta_{\theta ti}^n]$.

3.3. Optimizer inputs and outputs

As described in the preceding sections, the full set of Parameter RNN inputs for each tensor t are $\mathbf{x}_t^n = \{\mathbf{m}_t^n, \gamma_t^n, \eta_{\text{rel}}^n\}$, corresponding to the scaled averaged gradients, the relative log gradient magnitudes, and the relative log learning rate.

The full set of Parameter RNN outputs for each tensor t are $\mathbf{y}_t^n = \{\mathbf{d}_{\theta t}^n, \mathbf{d}_{\phi t}^n, \Delta \eta_{\theta t}^n, \beta_{gt}^n, \beta_{\lambda t}^n\}$, corresponding to the parameter and attention update directions, the change in step length, and the momentum logits. Each of the outputs in \mathbf{y}_t^n is read out via a learned affine transformation of the Parameter RNN hidden state. The readout biases are clamped to 0 for \mathbf{d}_{θ}^n and \mathbf{d}_{ϕ}^n . The RNN update equations are then:

$$\mathbf{h}_{\text{Param}}^{n+1} = \text{ParamRNN}(\mathbf{x}^n, \mathbf{h}_{\text{Param}}^n, \mathbf{h}_{\text{Tensor}}^n, \mathbf{h}_{\text{Global}}^n) \quad (10)$$

$$\mathbf{h}_{\text{Tensor}}^{n+1} = \text{TensorRNN}(\mathbf{x}^n, \mathbf{h}_{\text{Param}}^{n+1}, \mathbf{h}_{\text{Tensor}}^n, \mathbf{h}_{\text{Global}}^n) \quad (11)$$

$$\mathbf{h}_{\text{Global}}^{n+1} = \text{GlobalRNN}(\mathbf{x}^n, \mathbf{h}_{\text{Param}}^{n+1}, \mathbf{h}_{\text{Tensor}}^{n+1}, \mathbf{h}_{\text{Global}}^n) \quad (12)$$

$$\mathbf{y}^n = \mathbf{W} \mathbf{h}_{\text{Param}}^n + \mathbf{b}, \quad (13)$$

where \mathbf{h}^n is the hidden state for each level of the RNN, as described in Section 3.1, and \mathbf{W} and \mathbf{b} are learned weights

of the affine transformation from the lowest level hidden state to output.

3.4. Compute and memory cost

The computational cost of the learned optimizer is $\mathcal{O}(N_P B + N_P K_P^2 + N_T K_T^2 + K_G^2)$, where B is the minibatch size, N_P is the total number of parameters, N_T is the number of parameter tensors, and K_P , K_T , and K_G are the latent sizes for Parameter, Tensor, and Global RNNs respectively. Typically, we are in the regime where $N_P K_P^2 \gg N_T K_T^2 > K_G^2$, in which case the computational cost simplifies to $\mathcal{O}(N_P B + N_P K_P^2)$. Note that as the minibatch size B is increased, the computational cost of the learned optimizer approaches that of vanilla SGD, as the cost of computing the gradient dominates the cost of computing the parameter update.

The memory cost of the learned optimizer is $\mathcal{O}(N_P + N_P K_P + N_T K_T + K_G)$, which similarly to computational cost typically reduces to $\mathcal{O}(N_P + N_P K_P)$. So long as the latent size K_P of the Parameter RNN can be kept small, the memory overhead will also remain small.

We show experimental results for computation time in Section 5.6.

4. Meta-training

The RNN optimizer is meta-trained by a standard optimizer on an ensemble of target optimization tasks. We call this process meta-training, and the parameters of the RNN optimizer the meta-parameters.

4.1. Meta-training set

Previous learned optimizers have failed to generalize beyond the problem on which they were meta-trained. In order to address this, we meta-train the optimizer on an ensemble of small problems, which have been chosen to capture many commonly encountered properties of loss landscapes and stochastic gradients. By meta-training on small toy problems, we also avoid memory issues we would encounter by meta-training on very large, real-world problems.

Except where otherwise indicated, all target problems were designed to have a global minimum of zero (in some cases a constant offset was added to make the minimum zero). The code defining each of these problems will be open sourced shortly.

4.1.1. EXEMPLAR PROBLEMS FROM LITERATURE

We included a set of 2-dimensional problems which have appeared in optimization literature (Surjanovic &

Bingham, 2013) as toy examples of various loss landscape pathologies. These consisted of Rosenbrock, Ackley, Beale, Booth, Styblinski-Tang, Matyas, Branin, Michalewicz, and log-sum-exp functions.

4.1.2. WELL BEHAVED PROBLEMS

We included a number of well-behaved convex loss functions, consisting of quadratic bowls of varying dimension with randomly generated coupling matrices, and logistic regression on randomly generated, generally linearly separable data. For the logistic regression problem, when the data is not fully linearly separable, the global minimum is greater than 0.

4.1.3. NOISY GRADIENTS AND MINIBATCH PROBLEMS

For problems with randomly generated data, such as logistic regression, we fed in minibatches of various sizes, from 10 to 200. We also used a minibatch quadratic task, where the minibatch loss consisted of the square inner product of the parameters with random input vectors.

For full-batch problems, we sometimes added normally distributed noise with standard deviations from 0.1 to 2.0 in order to simulate noisy minibatch loss.

4.1.4. SLOW CONVERGENCE PROBLEMS

We included several tasks where optimization could proceed only very slowly, despite the small problem size. This included a many-dimensional oscillating valley whose global minimum lies at infinity, and a problem with a loss consisting of a very strong coupling terms between parameters in a sequence. We additionally included a task where the loss only depends on the minimum and maximum valued parameter, so that gradients are extremely sparse and the loss has discontinuous gradients.

4.1.5. TRANSFORMED PROBLEMS

We also included a set of problems which transform the previously defined target problems in ways which map to common situations in optimization.

To simulate problems with sparse gradients, one transformation sets a large fraction of the gradient entries to 0 at each training step. To simulate problems with different scaling across parameters, we added a transformation which performs a linear change of variables so as to change the relative scale of parameters. To simulate problems with different steepness-profiles over the course of learning, we added a transformation which applied monotonic transformations (such as raising to a power) to the final loss. Finally, to simulate complex tasks with diverse parts, we added a multi-task transformation, which summed the loss and concatenated the parameters from a diverse set of prob-

lems.

4.2. Meta-objective

For the meta-training loss, used to train the meta-parameters of the optimizer, we used the average log loss across all training problems,

$$L(\psi) = \frac{1}{N} \sum_{n=1}^N (\log(\ell(\theta^n(\psi)) + \epsilon) - \log(\ell(\theta^0) + \epsilon)), \quad (14)$$

where the second term is a constant, and where ψ is the full set of meta-parameters for the learned optimizer, consisting of $\psi = \{\psi_{\text{P-RNN}}, \psi_{\text{T-RNN}}, \psi_{\text{G-RNN}}, \gamma, c\}$, where $\psi_{\bullet\text{-RNN}}$ indicates the GRU weights and biases for the Parameter, Tensor, or Global RNN, γ is the learning rate momentum and c is the attended step offset (Section 3.2.4).

Minimizing the average log function value, rather than the average function value, better encourages exact convergence to minima and precise dynamic adjustment of learning rate based on gradient history (Figure 6). The average logarithm also more closely resembles minimizing the *final* function value, while still providing a meta-learning signal at every training step, since very small values of $\ell(\theta^n)$ make an outsized contribution to the average after taking the logarithm.

4.3. Partial unrolling

Meta-learning gradients were computed via backpropagation through partial unrolling of optimization of the target problem, similarly to Andrychowicz et al. (2016). Note that Andrychowicz et al. (2016) dropped second derivative terms from their backpropagation, due to limitations of Torch. We compute the full gradient in TensorFlow, including second derivatives.

4.4. Heavy-tailed distribution over training steps

In order to encourage the learned optimizer to generalize to long training runs, both the number of partial unrollings, and the number of optimization steps within each partial unroll, was drawn from a heavy tailed exponential distribution. The resulting distribution is shown in Appendix B.1

4.5. Meta-optimization

The optimizers were meta-trained for at least 40M meta-iterations (each meta-iteration consists of loading a random problem from the meta-training set, running the learned optimizer on that target problem, computing the meta-gradient, and then updating the meta-parameters). The meta-objective was minimized with asynchronous RMSProp across 1000 workers, with a learning rate of 10^{-6} .




Figure 2. Training loss versus number of optimization steps on MNIST for the Learned optimizer in this paper compared to the L2L optimizer from Andrychowicz et al. (2016), ADAM (learning rate $2e-3$), and RMSProp (learning rate $1e-2$). The L2L optimizer from previous work was meta-trained on a 2-layer, fully-connected network with sigmoidal nonlinearities. The test problems were a 2-layer fully-connected network and a 2-layer convolutional network. In both cases, ReLU activations and mini-batches of size 64 was used.


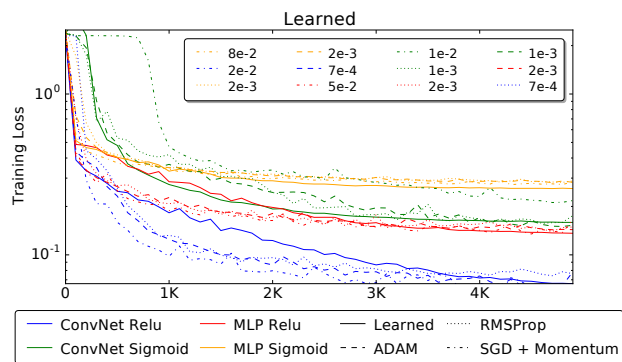


Figure 3. Three sample problems from the meta-training corpus on which the learned optimizer outperforms RMSProp and ADAM. The learning rates for RMSProp ($1e-2$) and ADAM ($2e-3$) were chosen for good average performance across all problem types in the training and test set. The learned optimizer generally beats the other optimizers on problems in the training set.


5. Experiments

5.1. Failures of existing learned optimizers

Previous learned optimizer architectures like Andrychowicz et al. (2016) perform well on the problems on which they are meta-trained. However, they do not generalize well to new architectures or scale well to longer timescales. Figure 2 shows the performance of an optimizer meta-trained on a 2-layer perceptron with sigmoid activations on



(a) Learned optimizer matches performance of ADAM, RMSProp, and SGD with momentum on four problems never seen in the meta-training set. For the non-learned optimizer, the optimal learning rate for each problem was chosen from a sweep over learning rates from 10^{-9} to 0.1. Actual learning rates used are shown in the inset legend.



(b) Training loss on ImageNet data in early training as a function of number of training examples seen (accounting for varying minibatch sizes). While other optimizer performance is highly dependent on hyperparameters, learned optimizer performance is similar to the best tuned optimizers (though in late training, the learned optimizer loss increases again). In both cases the learned optimizer was used for distributed, synchronized learning with an effective minibatch size of 800. The Inception V3 plot was generated from a newer version of the codebase, with small improvements described in Appendix C. On Inception V3, other optimizers used a learning rate of 0.045 and an effective minibatch size of 1600 (the optimal hyperparameters for the RMSProp optimizer from the original paper). On Resnet, other optimizers used a learning rate of 0.1 and an effective minibatch size of 256 (the optimal hyperparameters for the SGD + momentum optimizer from the original paper).

Figure 4. The learned optimizer generalizes to new problem types unlike any in the meta-training set, and with many more parameters.

the same problem type with ReLU activations and a new problem type (a 2-layer convolutional network). In both cases, the same dataset (MNIST) and minibatch size (64) was used. In contrast, our optimizer, which has not been meta-trained on this dataset or *any* neural network problems, shows performance comparable with ADAM and RMSProp, even for numbers of iterations not seen during meta-training (Section 4.4).

5.2. Performance on training set problems

The learned optimizer matches or outperforms ADAM and RMSProp on problem types from the meta-training set (Figure 3). The exact setup for each problem type can be seen in the python code in the supplementary materials.

5.3. Generalization to new problem types

The meta-training problem set did not include any convolutional or fully-connected layers. Despite this, we see comparable performance to ADAM, RMSProp, and SGD with momentum on simple convolutional multi-layer networks and multi-layer fully connected networks both in terms of final loss and number of iterations to convergence (Figure 4a and Figure 2).

We also tested the learned optimizer on Inception V3 (Szegedy et al., 2016) and on ResNet V2 (He et al., 2016). Figure 4b shows the learned optimizer is able to stably train these networks for the first 10K to 20K steps, with performance similar to traditional optimizers tuned for the spe-

cific problem. Unfortunately, we find that later in training the learned optimizer stops making effective progress, and the loss approaches a constant (approximately 6.5 for Inception V3). Addressing this issue would be a goal of future work.

5.4. Performance is robust to choice of learning rate




Figure 5. Learned optimizer performance is robust to learning rate hyperparameter. Training curves on a randomly generated quadratic loss problem with different learning rate initializations.

One time-consuming aspect of training neural networks with current optimizers is choosing the right learning rate for the problem. While the learned optimizer is also sensitive to initial learning rate, it is much more robust. Figure 5 shows the learned optimizer’s training loss curve on a quadratic problem with different initial learning rates com-

pared to those same learning rates on other optimizers.

5.5. Ablation experiments




Figure 6. Ablation study demonstrating importance of design choices on a small ConvNet on MNIST data. DEFAULT is the optimizer with all features included.

The design choices described in Section 3 matter for the performance of the optimizer. We ran experiments in which we removed different features and re-meta-trained the optimizer from scratch. We kept the features which, on average, made performance better on a variety of test problems. Specifically, we kept all of the features described in 3.2 such as attention (3.2.1), momentum on multiple timescales (gradient scl) (3.2.2), dynamic input scaling (variable scl decay) (3.2.3), and a relative learning rate (relative lr) (3.2.4). We found it was important to take the logarithm of the meta-objective (log obj) as described in 4.2. In addition, we found it helpful to let the RNN learn its own initial weights (trainable weight init) and an accumulation decay for multiple gradient timescales (inp decay). Though all features had an effect, some features were more crucial than others in terms of consistently improved performance. Figure 6 shows one test problem (a 2-layer convolutional network) on which all final features of the learned optimizer matter.

5.6. Wall clock comparison

In experiments, for small minibatches, we significantly underperform ADAM and RMSProp in terms of wall clock time. However, consistent with the prediction in 3.4, since our overhead is constant in terms of minibatch we see that the overhead can be made small by increasing the minibatch size.

6. Conclusion

We have shown that RNN-based optimizers meta-trained on small problems can scale and generalize to early train-




Figure 7. Wall clock time in seconds to run a single gradient and update step for a 6-layer ConvNet architecture on an HPz440 workstation with an NVIDIA Titan X GPU. As batch size increases, the total computation time for the Learned optimizer approaches ADAM.

ing on large problems like ResNet and Inception on the ImageNet dataset. To achieve these results, we introduced a novel hierarchical architecture that reduces memory overhead and allows communication across parameters, and augmented it with additional features shown to be useful in previous optimization and recurrent neural network literature. We also developed an ensemble of small optimization problems that capture common and diverse properties of loss landscapes. Although the wall clock time for optimizing new problems lags behind simpler optimizers, we see the difference decrease with increasing batch size. Having shown the ability of RNN-based optimizers to generalize to new problems, we look forward to future work on optimizing the optimizers.

References

- Andrychowicz, Marcin, Denil, Misha, Gomez, Sergio, Hoffman, Matthew W, Pfau, David, Schaul, Tom, Shillingford, Brendan, and de Freitas, Nando. Learning to learn by gradient descent by gradient descent. In *Advances in Neural Information Processing Systems*, 2016.
- Bahdanau, Dzmitry, Cho, Kyunghyun, and Bengio, Yoshua. Neural machine translation by jointly learning to align and translate. *iclr*, 2015.
- Bengio, S., Bengio, Y., and Cloutier, J. On the search for new learning rules for ANNs. *Neural Processing Letters*, 2(4):26–30, 1995.
- Bengio, Yoshua, Bengio, Samy, and Cloutier, Jocelyn. *Learning a synaptic learning rule*. Université de Montréal, Département d’informatique et de recherche opérationnelle, 1990.
- Bengio, Yoshua, Bengio, Samy, Cloutier, Jocelyn, and Gecsei, Jan. On the optimization of a synaptic learning rule. In *Conference on Optimality in Biological and Artificial Networks*, 1992.
- Chen, Yutian, Hoffman, Matthew W., Colmenarejo, Sergio Gomez, Denil, Misha, Lillicrap, Timothy P., and de Freitas, Nando. Learning to learn for global optimization of black box functions. arXiv Report 1611.03824, 2016.
- Cho, Kyunghyun, Van Merriënboer, Bart, Bahdanau, Dzmitry, and Bengio, Yoshua. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*, 2014.
- Duan, Yan, Schulman, John, Chen, Xi, Bartlett, Peter, Sutskever, Ilya, and Abbeel, Pieter. RI^2 : Fast reinforcement learning via slow reinforcement learning. Technical report, UC Berkeley and OpenAI, 2016.
- Funk, Simon. RMSprop loses to SMORMS3 - beware the epsilon!, 2015. URL sifter.org/~simon/journal/20150420.html.
- Harlow, Harry F. The formation of learning sets. *Psychological review*, 56(1):51, 1949.
- He, Kaiming, Zhang, Xiangyu, Ren, Shaoqing, and Sun, Jian. Identity mappings in deep residual networks. In *European Conference on Computer Vision*, pp. 630–645. Springer, 2016.
- Hochreiter, Sepp, Younger, A Steven, and Conwell, Peter R. Learning to learn using gradient descent. In *International Conference on Artificial Neural Networks*, pp. 87–94. Springer, 2001.
- Kehoe, E James. A layered network model of associative learning: learning to learn and configuration. *Psychological review*, 95(4):411, 1988.
- Kingma, Diederik and Ba, Jimmy. Adam: A method for stochastic optimization. *iclr*, 2015.
- Krizhevsky, Alex, Sutskever, Ilya, and Hinton, Geoffrey E. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- Lake, Brenden M, Ullman, Tomer D, Tenenbaum, Joshua B, and Gershman, Samuel J. Building machines that learn and think like people. arXiv Report 1604.00289, 2016.
- Li, SKe and Malik, Jitendra. Learning to optimize. In *International Conference on Learning Representations*, 2017.
- Naik, Devang K and Mammone, RJ. Meta-neural networks that learn by learning. In *International Joint Conference on Neural Networks*, volume 1, pp. 437–442. IEEE, 1992.
- Nesterov, Yurii. A method of solving a convex programming problem with convergence rate $o(1/k^2)$. In *Soviet Mathematics Doklady*, volume 27, pp. 372–376, 1983a.
- Nesterov, Yurii. A method of solving a convex programming problem with convergence rate $o(1/k^2)$. In *Soviet Mathematics Doklady*, volume 27, pp. 372–376, 1983b.
- Ravi, Sachin and Larochelle, Hugo. Optimization as a model for few-shot learning. In *International Conference on Learning Representations*, 2017.
- Runarsson, Thomas Philip and Jonsson, Magnus Thor. Evolution and design of distributed learning rules. In *IEEE Symposium on Combinations of Evolutionary Computation and Neural Networks*, pp. 59–63. IEEE, 2000.
- Santoro, ADAM, Bartunov, Sergey, Botvinick, Matthew, Wierstra, Daan, and Lillicrap, Timothy. Meta-learning with memory-augmented neural networks. In *International Conference on Machine Learning*, 2016.
- Schmidhuber, Jurgen. *Evolutionary Principles in Self-Referential Learning. On Learning how to Learn: The Meta-Meta-Meta...-Hook*. PhD thesis, Institut f. Informatik, Tech. Univ. Munich, 1987.
- Silver, David, Huang, Aja, Maddison, Chris J, Guez, Arthur, Sifre, Laurent, Van Den Driessche, George, Schrittwieser, Julian, Antonoglou, Ioannis, Panneershelvam, Veda, Lanctot, Marc, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.

- Surjanovic, Sonja and Bingham, Derek. Optimization test functions and datasets, 2013. URL <http://www.sfu.ca/ssurjano/optimization.html>.
- Sutton, Richard S. Adapting bias by gradient descent: An incremental version of delta-bar-delta. In *Association for the Advancement of Artificial Intelligence*, pp. 171–176, 1992.
- Szegedy, Christian, Vanhoucke, Vincent, Ioffe, Sergey, Shlens, Jon, and Wojna, Zbigniew. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2818–2826, 2016.
- Thrun, Sebastian and Pratt, Lorien. *Learning to learn*. Springer Science and Business Media, 1998.
- Tieleman, Tijmen and Hinton, Geoffrey. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural Networks for Machine Learning*, 4:2, 2012.
- Tseng, Paul. An incremental gradient (-projection) method with momentum term and adaptive stepsize rule. *Journal on Optimization*, 8(2):506–531, 1998.
- Wang, Jane X., Kurth-Nelson, Zeb, Tirumala, Dhruva, Soyer, Hubert, Leibo, Joel Z., Munos, Rémi, Blundell, Charles, Kumaran, Dharshan, and Botvinick, Matt. Learning to reinforcement learn. arXiv Report 1611.05763, 2016.
- Ward, Lewis B. Reminiscence and rote learning. *Psychological Monographs*, 49(4):i, 1937.
- Zoph, Barret and Le, Quoc V. Neural architecture search with reinforcement learning. In *International Conference on Learning Representations*, 2017.