
Meta-Learning Bidirectional Update Rules

Mark Sandler¹ Max Vladymyrov¹ Andrey Zhmoginov¹ Nolan Miller¹ Andrew Jackson¹ Tom Madams¹
Blaise Agüera y Arcas¹

Abstract

In this paper, we introduce a new type of generalized neural network where neurons and synapses maintain multiple states. We show that classical gradient-based backpropagation in neural networks can be seen as a special case of a two-state network where one state is used for activations and another for gradients, with update rules derived from the chain rule. In our generalized framework, networks have neither explicit notion of nor ever receive gradients. The synapses and neurons are updated using a bidirectional Hebb-style update rule parameterized by a shared low-dimensional “genome”. We show that such genomes can be meta-learned from scratch, using either conventional optimization techniques, or evolutionary strategies, such as CMA-ES. Resulting update rules generalize to unseen tasks and train faster than gradient descent based optimizers for several standard computer vision and synthetic tasks.

1. Introduction

Neural networks revolutionized the way ML systems are built today. Advances in neural design patterns, training techniques, and hardware performance allowed ML to solve tasks that seemed hopelessly out of reach less than ten years ago. However, despite the rapid progress, their basic neuron-synapse design has remained fundamentally unchanged for nearly six decades, since the introduction of perceptron models in the 50s and 60s (Minsky & Papert, 1969; Rosenblatt, 1957) that modeled the complex biology of a synapse firing as a simple combination of a weight and a bias combined with a non-linear activation function.

With such models, the next question was “How should we best find the optimal weights and biases?” and great successes have come from the use of stochastic gradient descent,

¹Google Research. Correspondence to: Mark Sandler <sandler@google.com>.

traced back to Robbins & Monro (1951). Since its introduction, many of the more recent remarkable improvements can be attributed to improving the efficiency of the gradient signal: adjusting connectivity patterns such as in convolutional neural networks and residual layers (He et al., 2015), improved optimizer design (Kingma & Ba, 2014; Duchi et al., 2011; Schmidt et al., 2020), and normalization methods such as (Ioffe & Szegedy, 2015; Ulyanov et al., 2016). All these methods improve the learning characteristics of the networks and enable scaling to larger networks and more complex problems.

However the underlying principle behind these methods remained the same: minimize an engineered loss function using gradient descent. Instead, we propose a different approach. While we still follow the general strategy of forwards and backwards signal transmission, we learn the rules governing both forward and back-propagation of neuron activation from scratch. The key enabling factor here is a generalization where each neuron can have multiple states.

We define a *space* of possible transformations that specify the interaction between neurons’ feed-forward and feed-back signals. The matrices controlling these interactions are meta-parameters that are shared across both layers and tasks. We term these meta-parameters a “genome”. This reframing opens up a new, more generalized space of neural networks, allowing the introduction of arbitrary numbers of states and channels into neurons and synapses, which have their analogues in biological systems, such as the multiple types of neurotransmitters, or chemical vs. electrical synapse transmission.

Our framework, which we call BLUR (Bidirectional Learned Update Rules) describes a general set of multi-state update rules that are capable to train networks to learn new tasks without ever having access to explicit gradients. We demonstrate that through meta-learning BLUR can learn effective genomes with just a few training tasks. Such genomes can be learned using off-the-shelf optimizers or evolutionary strategies. We show that such genomes can train networks on unseen tasks faster than comparably sized gradient networks. The learned genomes can also generalize to architectures unseen during the meta-training.

2. Related Work

Alternatives to stochastic gradient descent Replacing the backpropagation of loss function gradients with a different signal has been explored before. One example is a family of methods building upon *difference target propagation* (Bengio, 2014; Lee et al., 2015), which uses the desired output “targets” or “errors” as the primary signals communicated during the backpropagation stage. A recent paper by Ahmad et al. (2020) builds on the target propagation approach and proposes an alternative, a more biologically plausible *local update* that is shown to be equivalent to the SGD update. Similarly, SGD can be expressed as a special case of a more general family of update rules, which we explore for the best-performing learning algorithm.

Another related line of work builds upon the *feedback alignment* method that replaces backwards weights in SGD with fixed random matrices (Lillicrap et al., 2016; Nøkland, 2016). In Akrouf et al. (2019), the authors extend this idea by separately evolving backward weights for automatic “synchronization” of forward and backward weights. In addition, Xiao et al. (2018) have demonstrated empirically that many of the biologically inspired methods listed above train *at most* as good as SGD.

The bidirectional nature of networks was explored in (Pontes-Filho & Liwicki, 2019; Adigun & Kosko, 2019), where the backward pass is treated as a generative inference rather than learning update.

The common feature of approaches above is that all these methods rely on a standard forward pass, and a hand-designed backward pass. While some variants of feedback alignment are contained within the family of update rules we explore in this paper, in contrast to them our approach learns the meta-parameters that control both the forward and backward passes.

Another line of inquiry are methods that utilize greedy layer-wise training (Bengio et al., 2006; Belilovsky et al., 2019; Löwe et al., 2019; Xiong et al., 2020), a pseudoinverse (Ranganathan & Lewandowski, 2020), or like Taylor et al. (2016) reformulate the end-to-end training as a constrained optimization problem with additional variables. While these methods do not share as many common elements with the present approach, we believe that exploring layer-wise training in conjunction with our method is a promising direction for future research.

Meta-Learning Recently, researchers have turned to meta-learning (Schmidhuber, 1987) approaches that aim on improving existing functional modules (Munkhdalai et al., 2019) or learning methods by meta-training optimal hyperparameters (often called meta-parameters) in a problem-independent way (Andrychowicz et al., 2016; Wichrowska et al., 2017; Maheswaranathan et al., 2020; Metz et al., 2019;

2020). The important difference with our method is that these work by directly using the gradient of a given loss function, whereas our method proposes a holistic learning framework that does not correspond to any known optimizer or even rely on a predefined loss function.

It has long been recognized that SGD is a biologically implausible mechanism (Bengio et al., 2015). Another direction that is explored in the literature is more biologically plausible mechanisms (Soltoggio et al., 2018) including those based on *Hebb’s rule* (Hebb, 1949) and its modification *Oja’s rule* (Oja, 1982). Meta-learning has been used to learn the plasticity of similar update rules (Miconi et al., 2018; 2019; Lindsey & Litwin-Kumar, 2020; Confavreux et al., 2020; Najarro & Risi, 2020) as well as different neuro-modulation mechanisms (Bengio et al., 1995; Norouzzadeh & Clune, 2016; Velez & Clune, 2017; Wilson et al., 2018). Recent work by Camp et al. (2020) followed a related direction learning the sub-structure of individual neurons by representing them as perceptron networks while keeping gradient-based backpropagation. In yet another alternative approach, Kaplanis et al. (2018) proposed to introduce more nuanced memory to synapses, while in Randazzo et al. (2020), authors meta-learn parameters of a complex message-passing learning algorithm that replaces the backpropagation while leaving the forward pass intact.

Kirsch & Schmidhuber (2020) propose a generalized learning algorithm based on a set of RNNs that, similar to our framework, does not use any gradients or explicit loss function, yet is able to approximate forward pass and backpropagation solely from forward activations of RNNs. Our system, in contrast, does not use RNNs and explicitly leaves (meta-parametrized) bidirectional update rules in place.

Different from traditional meta-learning, Real et al. (2020) and Ha et al. (2016) devise a specialized learning algorithm to directly find the optimal parameters of another target algorithm that they want to learn.

To the best of our knowledge, our paper is the first work that customizes *both* inference and learning passes by successfully finding the update rule for both forward and backward passes that does not rely on neither explicit gradients or a predefined loss function.

3. Learning a New Type of Neural Network

3.1. A generalization of gradient descent using neuronal state

To learn a new type of neural network we need to formally define the space of possible configurations. Our proposed space is a generalization of classical artificial neural networks, with inspiration drawn from biology. For the purpose of clarity, in this section we modify the notation by




Figure 1. Generalization of a three layers feed-forward neural networks as a multi-state systems. *Left:* Forward pass and chain-rule backpropagation organized as a generalized two-state network. Arrows indicate the flow of information from forward and backward passes to synapse updates. *Right:* Our proposed generalized formulation. Green nodes, defined in the genome, control the amount of mixing between the states. They are fixed during the synapse update (inner-loop) and are optimized during the meta-training (outer-loop). Grey boxes indicate multi-state variables. Orange boxes represent activation functions. Notice the symmetry between forward and backward passes.

abstracting from the standard layer structure of a neural network, and instead assume our network is essentially a *bag-of-neurons* \mathcal{N} of n neurons with a connectivity structure defined by two functions: “upstream” neurons $I(i) \subset \mathcal{N}$ that send their outputs to i , and the set of “downstream” neurons $J(i) \subset \mathcal{N}$ that receive the output of i as one of their inputs. Thus the synapse weight matrix w_{ij} can encode separate weights for forward and backward connections.

Normally we think of a neuron in an artificial neural network as having a single scalar value, it turns out that one state is not enough once we incorporate a back-propagation signal, which uses both feed-forward state and feedback state to propagate through the network. The standard forward pass over a densely connected neural network updates the state of each neuron $j \in \mathcal{N}$ according to

$$h_j = \sigma \left(\sum_{i \in I(j)} w_{ij} h_i \right), \quad (1)$$

where h_j is the activation for neuron $j \in \mathcal{N}$ resulting from applying a function $\sigma(\cdot)$ to the product of the network weights w_{ij} and the incoming stimulus from $I(j)$. Here and further we combine the bias and the weights into a single vector by adding a unit element at the end of each hidden vector. For the first layer, the activation is given by the input batch. Let us define $h'_j := \sigma' \left(\sum_{i \in I(j)} w_{ij} h_i \right)$ as a derivative of the activation with respect to its argument. Then we can write chain rule, describing the back-propagation of a loss function as:

$$\frac{\partial L}{\partial h_i} = \sum_{j \in J(i)} w_{ij} \frac{\partial L}{\partial h_j} h'_j. \quad (2)$$

For the last layer, the first step of backpropagation is given by the derivative of the loss function with respect to the last activations. After the backpropagation, using the notation above, the gradient descent updates the synapses w_{ij} using

$$w_{ij} \leftarrow w_{ij} - \tilde{\eta} \frac{\partial L}{\partial h_j} h'_j h_i, \quad (3)$$

where $\tilde{\eta}$ is a learning rate.

Notice that the update has a form of the Hebbian learning rule (Hebb, 1949) with pre-synaptic activation given by h_i and post-synaptic one given by $\frac{\partial L}{\partial h_j} h'_j$. We can make this connection even more explicit using neurons with two states, i.e., activations above are now replaced with a two-dimensional vector $\mathbf{a}_i = (a_i^{(1)}, a_i^{(2)})$. One of these states would be used for a feed-forward signal and another for a back-propagated feedback signal. During the forward pass we set $\mathbf{a}_j \leftarrow (h_j, h'_j)$ and during the backward pass the second state is updated multiplicatively using (2) as $a_i^{(2)} \leftarrow \frac{\partial L}{\partial h_i} h'_i = a_i^{(2)} \sum_{j \in J(i)} w_{ij} a_j^{(2)}$. Then the synapse update is given by $w_{ij} \leftarrow w_{ij} - \tilde{\eta} a_j^{(2)} a_i^{(1)}$. The left side of Figure 1 demonstrates the described operations as a two-state neural network.

We can further generalize the learning procedure with the following constant matrices: $\nu = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$, $\mu = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$, $\tilde{\nu} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$, $\tilde{\mu} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$ and a generalized binary activation function $\phi \left(\begin{pmatrix} x \\ y \end{pmatrix} \right) = \begin{pmatrix} \sigma(x) \\ \sigma'(y) \end{pmatrix}$. Then the operations above can be equivalently rewritten as

$$\begin{aligned} \text{Forward pass:} \quad a_j^c &\leftarrow \phi^c \left(\sum_{i \in I(j), d} w_{ij} a_i^d \nu^{cd} \right) \\ \text{Backward pass:} \quad a_i^{(2)} &\leftarrow a_i^{(2)} \sum_{j \in J(i), d} w_{ij} a_j^d \mu^d \\ \text{Weights update:} \quad w_{ij} &\leftarrow w_{ij} - \tilde{\eta} \sum_{c, d} a_j^c \tilde{\mu}^c a_i^d \tilde{\nu}^d. \end{aligned} \quad (4)$$

Here $c, d = \{1, 2\}$ represent the states of the network and we use superscript to index over the states. Thus, traditional gradient backpropagation can be expressed as a general two-state network, whose update rules are controlled by a pre-defined set of very low-dimensional matrices $\{\nu, \mu, \tilde{\nu}, \tilde{\mu}, \tilde{\eta}\}$. These matrices are fixed during the weight update phase and optimized during the meta-optimization to achieve a more general update rules.

Name	Dimension	Description
Constants		
n	-	total number of neurons.
k	-	total number of states.
Network params		
a_i^c	$i \in [n], c \in [k]$	state c of neuron i .
w_{ij}^c	$i, j \in [n], c \in [k]$	channel c of synapse between i and j .
Meta-learning params (genome)		
f, η	1	neuron forget and update gates.
$\tilde{f}, \tilde{\eta}$	1	synapses forget and update gate.
ν^{cd}, μ^{cd}	$c \in [k], d \in [k]$	forward/backward neuron transform matrix.
$\tilde{\nu}^{cd}, \tilde{\mu}^{cd}$	$c \in [k], d \in [k]$	pre- and post-synaptic transform matrix.

Table 1. Description and dimensions of variables.

3.2. Multi-state bidirectional update rules

The two-state interpretation of the backpropagation algorithm outlined above is asymmetrical and contains several potentially biologically implausible design details like the use of the same weight matrix on the forward and backward passes and a multiplicative update during the backpropagation phase. Being inspired by this update mechanism, we propose a general family of bidirectional learned update rules (BLUR) that: (a) use multi-channel asymmetrical synapses, (b) use the same update mechanisms on the forward and backward paths, and finally (c) allow for information mixing between different channels of each neuron. In its final state, this family can be described by the following equations:

$$\text{Forward pass: } a_j^c \leftarrow \sigma(f a_j^c + \eta \sum_{i \in I(j), d} w_{ij}^c \nu^{cd} a_i^d) \quad (5)$$

$$\text{Backward pass: } a_i^c \leftarrow \sigma(\tilde{f} a_i^c + \eta \sum_{j \in J(i), d} w_{ji}^c \mu^{cd} a_j^d) \quad (6)$$

$$\text{Weights update: } w_{ij}^c \leftarrow \tilde{f} w_{ij}^c + \tilde{\eta} \sum_{e, d} a_i^e \tilde{\nu}^{ec} \cdot \tilde{\mu}^{cd} a_j^d. \quad (7)$$

The right side of Fig. 1 demonstrates the proposed framework and Table 1 tracks the description and dimensionality of variables used in the formulae above. The matrices $\{f, \tilde{f}, \nu, \tilde{\nu}, \mu, \tilde{\mu}, \eta, \tilde{\eta}\}$ used in (5–7) form our complete genome \mathcal{G} .

Here we make the following generalizations with respect to the backpropagation update rules:

- The neuron transform matrices ν, μ and synapse transform matrices $\tilde{\nu}, \tilde{\mu}$ all have dimension $k \times k$ and allow

for mixing of every input state to every output state as well as possibility using more than two states in the genome.

- We expand the genome to include $f, \eta, \tilde{f}, \tilde{\eta}$ that control how much of the information is *forgotten* and how much is being *updated* after each step. A similar approach has been studied in Ravi & Larochelle (2017), however we learn these scalars directly and do not model them as a function of a previous iteration.
- We propose an additive update for both neurons and synapses. Note that in order to generalize to backpropagation, an additive update for the backward pass has to be replaced with a multiplicative one and applied only to the second state. Experimentally, we discovered that both additive and multiplicative updates perform similarly.
- We extend the activation function to be applied on both forward and backward pass and, to make things simple, make it unitary (same function applied to every state).
- We generalize the synapse matrices to be asymmetric for a forward and backward pass ($w_{ij} \neq w_{ji}$) as well as contain more than one state. Symmetric weight matrices are ordinarily used for deep learning, but distinct weight matrices are more biologically plausible.
- The synapse update has a general form of a Hebbian-update rule mixing pre- and post-synaptic activity according to the synapse transform matrices $\tilde{\nu}, \tilde{\mu}$.

In addition to generalizing existing gradient learning, not relying on gradients in backpropagation has additional benefits. For example, the network doesn’t need to have an explicit notion of a final loss function. The feedback (ground truth) can be fed directly into the last layer (e.g. by an update to the second state or simply by replacing the second state altogether) and the backward pass would take care of backpropagating it through the layers.

Notice that the genome is defined at the level of individual neurons and synapses and is independent from the network architecture. Thus, the same genome can be trained for different architectures and, more generally, genome trained on one architecture can be applied to a genome with different architectures. We show some examples of this in the experimental section.

Since the proposed framework can use more than two states, we hypothesize that just as the number of layers relates to the complexity of learning required for an individual task (inner loop of the meta-learning), the number of states might be related to complexity of learning behaviour across the task (outer loop). More informally: synapses regulate “how

hard is a given task” vs genome’s “how hard it is to learn the task given a variety of other tasks available”.

Our resulting genome now completely describes the communication between individual neurons. The neurons themselves can be arranged in any of the familiar ways – in convolutional layers, residual blocks, etc. For the rest of the paper we focus on the simplest types of networks consisting of one or more fully connected layers.

3.3. Meta-learning the genome

Once we have defined the space of possible update rules, the next step is to design an algorithm to find a useful genome capable of successful training and generalization. In this work we concentrate on meta-learning genomes that can solve classification problems with multiple hidden layers.

For a d -class classification problem with l -dimensional input, we use the first layer as input and the last layer with d neurons as predictors. We denote those neurons as $x_1 \dots x_l$ and $y_1 \dots y_d$ respectively.

During the learning process, in the forward pass we apply equation (5) to compute a logit prediction for a given class i to the *first* state of the last layer $a_{y_i}^{(1)}$. During the backward pass, we set the *second* state $a_{y_i}^{(2)}$ as 1 or -1 from the ground truth based on the class attribute. In experiments with more than two states per neuron, we fill the other states of the last layer with zeros. We then use equations (6) and (7) to compute updated synapse state.

To evaluate the quality of a genome we apply equations (5–7) for multiple unroll steps and then test learned synapses on a previously unseen set of inputs. To meta-learn using SGD we use standard softmax-cross entropy loss: $L_{\text{meta}}(\mathcal{G}) = \mathbf{E}_s \left[p_i(s) \log a_{y_i}^{(0)}(s) \right]$ where $p_i(s)$ is one-hot vector representing the true category for a sample s , and $a_{y_i}^{(0)}(s)$ is the prediction of that sample from the forward pass, after applying our forward and backward updates for a given number of unroll steps. We then can minimize this function with standard off-the-shelf optimizers to find an optimal genome.

In section 4.4 we also perform experiments using CMA-ES (Hansen & Ostermeier, 1996) where we use training accuracy as a fitness metric.

3.4. Activation normalization and synapse saturation

Synapse updates that rely on Hebb’s rule alone ($\Delta w_{ij} \sim a_i a_j$) are generally unstable, as network weights w grow monotonically with the training steps. One way to alleviate this issue while also reducing the sensitivity of network outputs to small synapse perturbations is to use *activation normalization*. Normalization techniques are known to be used by certain biological systems (Carandini & Heeger,

2012) to calibrate neuron activation to the optimal firing regime and are widely used in conventional deep neural architectures (Ioffe & Szegedy, 2015; Ulyanov et al., 2016). In most of our experiments, we used per-channel normalization similar to batch-normalization to normalize a pre-nonlinearity activation distribution into one with a learnable mean and deviation. To maintain the symmetry between the forward and backward pass we apply normalization for both forward and backward activations. This not only helps in training deeper models, but also allows learned update rules to generalize to different input sizes and number of classes.

However, activation normalization alone does not always prevent an unbounded growth of synapse weights, and so another mechanism for weight saturation is necessary. One such approach is based on using Oja’s update rule (Oja, 1982) that modifies Hebb’s rule with an additional component that by itself leads to the decay of the singular components of the weight matrix. One of the most commonly used forms of Oja’s update rule is $\Delta w_{ij} = \gamma a_i a_j - \gamma a_j^2 w_{ij}$ however there also exist other forms with similar properties. In linear systems, the interplay of the excitatory and inhibitory driving forces leads to synapse saturation (Oja, 1982). But in our case, the linear component of the update rule $\tilde{f}w$ along with the usage of nonlinearities (like tanh) or activation normalization may prevent Oja’s original rule from saturating model weights. In Appendix B, we show that the same principles that were used to derive Oja’s rule can also be applied to our system and result in the following inhibitory Oja-like update:

$$(\Delta w_{ij}^c)^{\text{Oja}} = -(\tilde{f} - 1)w_{ij}^c \sum_r (w_{rj}^c)^2 - \tilde{\eta} w_{ij}^c \sum_r w_{rj}^c \sum_{e,d} a_r^e \tilde{\nu}^{ec} \cdot \tilde{\mu}^{cd} a_j^d. \quad (8)$$

The first component of this update usually dominates the second component, so in our experiments we only used this term with an additional learnable multiplier.

Oja’s term is generally derived as an inhibitory additive component that acts as a “counterweight” to the Hebbian synapse update and keeps the weight norm fixed. Instead of using such inhibitory terms, we could instead apply normalization and saturating nonlinearities (like $\alpha \tanh(x/\alpha)$ with a learnable α coefficient) directly to the synapses. In our experiments, we empirically validated that both approaches lead to very similar results and could thus potentially be used interchangeably.

3.5. Is this update rule still a gradient descent?

Once we have used meta-learning to identify a promising genome, one might ask if the resulting training algorithm is in fact identical to a conventional gradient descent with some unknown loss function L_{equiv} . In this section we




Figure 2. Natural logarithm of $|\partial\Delta w_k/\partial w_m - \partial\Delta w_m/\partial w_k|$ computed numerically for flattened and concatenated forward weight matrices in a two-state model learning a 2-input Boolean task (e.g. `xor`). The model has a single hidden layer of size 20 and thus 204 total forward weights and biases. The magnitude of the numerical error is estimated to be below e^{-4} . Bands in the left-upper corner correspond to asymmetries between first-layer connections to the same hidden neurons; other structures appear to arise because of the asymmetries between the first- and the second-layer weights.

empirically demonstrate that the answer is generally no. Consider a full-batch training scenario and let $\Delta w_{ij}^c(\hat{w}; \mathbf{a})$ be the weight update rule defined by our learned genome. Equivalence to the gradient descent would then mean that

$$\Delta w_{ij}^c = -\gamma \frac{\partial L_{\text{equiv}}}{\partial w_{ij}^c} \quad (9)$$

and therefore

$$\frac{\partial \Delta w_{ij}^c}{\partial w_{mn}^d} = -\gamma \frac{\partial^2 L_{\text{equiv}}}{\partial w_{ij}^c \partial w_{mn}^d}.$$

Since the partial derivatives are symmetric, we see that

$$\frac{\partial \Delta w_{ij}^c}{\partial w_{mn}^d} = \frac{\partial \Delta w_{mn}^d}{\partial w_{ij}^c} \quad (10)$$

is a necessary condition for the existence of the loss L_{equiv} satisfying Eq. (9). This condition can be tedious to verify analytically, but we can instead check it numerically. Computing $|\partial \Delta w_{ij}^c / \partial w_{mn}^d - \partial \Delta w_{mn}^d / \partial w_{ij}^c|$ in a simple experiment with Boolean functions and a single hidden layer of size 20, we verified that the discovered update rules do not satisfy condition (10) (see Fig. 2) and therefore L_{equiv} does not generally exist for our update rule family.

The observation that learning trajectories obtained with our update rule cannot be recovered using conventional gradient descent does not rule out a potential equivalence to other learning algorithms such as, e.g., a gradient descent with a Riemannian metric $\hat{g}(\mathbf{w})$ defined via $\Delta w_i = -\gamma \sum_j g_{ij} \cdot (\partial L_{\text{equiv}} / \partial w_j)$ (for details see Appendix A).

It is also worth noticing that the question of existence of a loss function that is monotonically non-increasing along the training trajectories is directly related to the well-established theory of Lyapunov functions that currently includes many general existence theorems (Conley, 1978; 1988; Farber et al., 2003; Franks, 2017) and is subject of future work.

4. Experiments

In this section we describe experimental evaluation of update rules using BLUR. Our code uses tensorflow (Abadi et al., 2015) and JaX (Bradbury et al., 2018) libraries. All our experiments run on GPU. Typically a single genome can be trained on a single GPU in between 30 minutes and 20 hours depending on configuration. For some experiments we run multiple identical runs to estimate variance. For all our experiments we use the same set of basic parameters as described in Appendix C.1. In section 4.5 we study several alternative formulations to show their impact on convergence and stability. Code for the paper is available at <https://github.com/google-research/google-research/tree/master/blur>

4.1. Meta-learning simple functions

We first consider a very simple setup where we try to learn toy-examples with two variable inputs. The datasets are shown in Fig. 3. The tasks we use for training are and, `xor`, `two-moon` (Pedregosa et al., 2011) and several others. We then validate it on held out datasets shown in Fig. 3. These datasets include both in-domain (e.g. other 2-class functions), five-class function of two variables `blobs` (Pedregosa et al., 2011) and out-of-domain MNIST (Lecun et al., 1998). We train a two-layer network using a two-state genome, and use the same two-layer architecture for meta-training and meta-validation. The results are shown in Fig. 4.

4.2. Generalization capabilities

In this section we explore the ability of genomes that we found to generalize to new datasets. We use MNIST as a meta-training dataset. Since MNIST generally requires more than 10 steps to converge we use a variant of curriculum training to gradually increase the number of unrolls. We start with 8-identical randomly initialized genomes and train them for 10,000 steps with 10 unrolls. Then we increase the unroll number by 5 for each consecutive 10,000 steps and synchronize genomes across all runs. The example of meta-training training accuracy is shown in Appendix. In Fig. 5 we show the performance of a curriculum-trained genome. The genome was trained on 3 tasks: full MNIST, and two half-size MNIST datasets, one with digits from 0 to 4 and another from 5 to 9. Interestingly, even this naive setup that uses just three slightly different tasks shows improved gen-

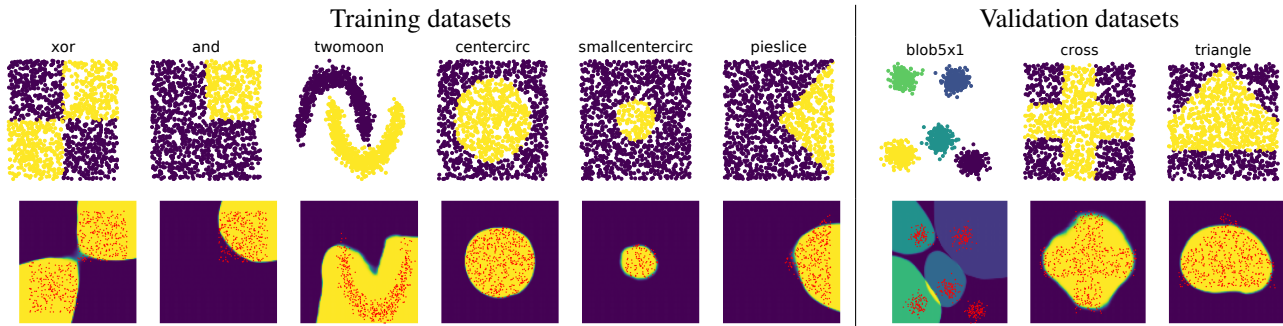


Figure 3. Toy datasets. *Left side*: meta-training datasets, *right side*: meta-validation datasets. *Top row*: training data, *bottom row*: prediction of the trained genome produces on a dense and enlarged grid.

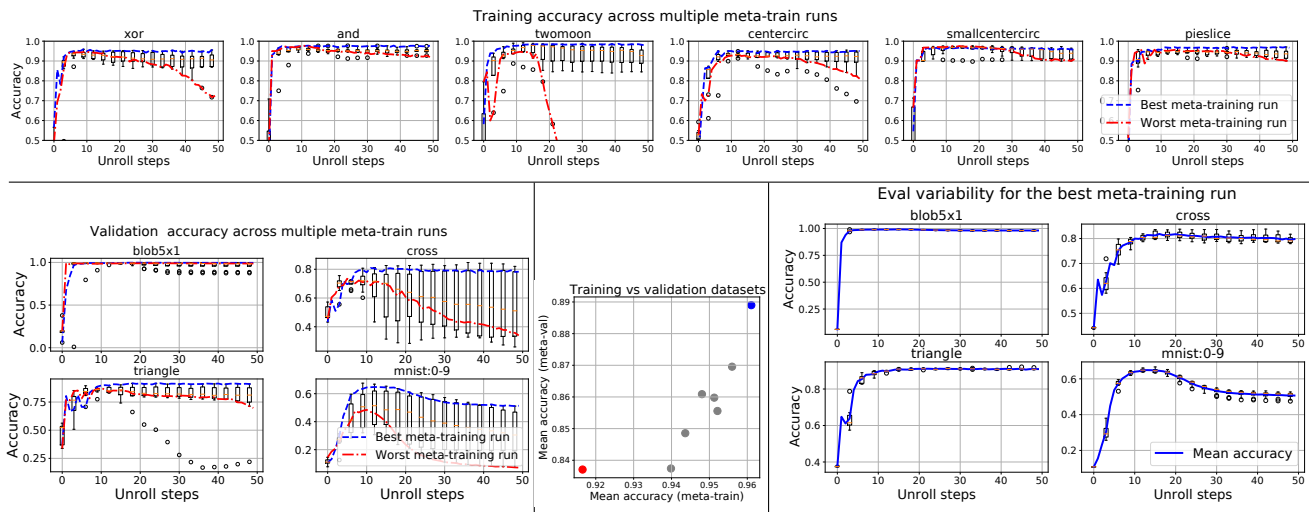


Figure 4. Meta-learning on toy datasets. *Top row*: the validation performance on training datasets. The spread indicates variation across multiple meta-training runs. *Bottom left*: shows the variation on validation datasets. *Bottom right*: the variation within a single run, but for different synapse initializations. *Bottom middle*: the correlation between mean meta-validation and meta-training accuracy across different runs at unroll 10. For all the graphs the blue line shows the run with the highest average meta-training accuracy and the red one the lowest. Since genomes are selected based on their *average* meta-training, it is expected that for some tasks they are not the worst.

eralization abilities. For meta-training we used the MNIST dataset cropped to 20x20 and resized to 10x10. For meta-validation we always used 28x28 datasets. Specifically we used MNIST, a 10-class letter subset of E-MNIST (Cohen et al., 2017), Fashion MNIST (Xiao et al., 2017), and the full 62-category E-MNIST. This shows that the meta-learned update rules can successfully learn unrelated tasks without ever having access to gradient functions. We didn’t include the graph of evaluation accuracy for *meta-training* variant of MNIST that used cropped and down-sampled digits, but we note that it generally produced results that were about 1% below 28x28 MNIST, thus suggesting excellent meta-generalization.

Finally, on the same Figure 5 we show an example of an MNIST-trained genome applied to an out-of-domain Boolean task.

4.3. Comparison with SGD

In figure 6 we compare the convergence performance on MNIST dataset of BLUR vs. SGD with and without momentum with different values of learning rate spanning 4 orders of magnitude.

4.4. Training using evolution strategies

Using the evolution strategies to train genomes are appealing for two reasons. First, it enables us to train networks with a much larger number of unrolls. Second, it allows us to find genomes with non-differentiable objectives.

In this section we use a simplified setup without curriculum learning. We train a network with a single hidden layer of 512 channels and 2 states on MNIST downsampled to 14x14. Each meta-learning step represents training the network on




Figure 5. Generalization of a genome meta-trained with 2-layer/4-state architecture on MNIST to other datasets. Note that our search only explored up to 50 unroll steps. Our method converges much faster than SGD in the explored training trajectory, however after about 100 steps SGD reaches the same accuracy and continues to grow. The rightmost visualization applies the same genome to learn a XOR function.




Figure 6. Performance of 4-state BLUR network vs. SGD with different learning rates.




Figure 7. Using CMA to learn best non-differentiable accuracy objective, against gradient based meta training. *Left*: meta-training trajectory *right*: fully trained genome learning MNIST. Both methods show best run out of 8.

15 batches of 128 inputs, then evaluating its accuracy on 20 batches of 128 inputs. We use the CMA-ES/pycma (Hansen et al., 2019) library to optimize the network genome with respect to the (non-differentiable) 20 batch accuracy. Using a population of 80 parallel experiments per time step, CMA-ES is able to learn a genome that achieves 0.84 accuracy in 400 steps. Accuracy as a function of meta-learning steps can be seen in Fig. 7. Evaluating the generalization of this genome over different numbers of unrolls reveals that it reaches full accuracy in 15 unroll steps as expected, and plateaus without significant decay and results in comparable performance as gradient-based search.




Figure 8. Importance of neuron normalization. The genomes that do not use normalization produced did not generalize as well to different resolutions and deeper networks.

4.5. Ablation study

In this section we explore the importance of several critical parameter choices. We will be training with the same setup as in section 4.1, but instead of verifying it on the meta-validation dataset, we will measure the impact on MNIST.

Normalization Normalization plays a crucial role in the stability of our meta-training process and improves the final training accuracy. Fig. 8 shows a meta-training accuracy comparison of identical runs with forward/backward normalization turned off.

Impact of non-linearity In contrast with chain-rule back-propagation, our learning algorithm uses symmetric non-linearity for simplicity. Curiously, it appears that the independent choice of non-linearity on the forward and backward pass has relatively little impact on our ability to find genomes that can learn. Generally the space seems to be insensitive to the choice of non-linearity used on either forward or backward pass. In Appendix we include a table showing the variation in validation accuracy across different non-linearities.

Symmetry of the synapses Here we compare our framework across combinations of three different dimensions: (a) using symmetric or asymmetric synapses for backward and forward passes, (b) using single or multi-state synapses and

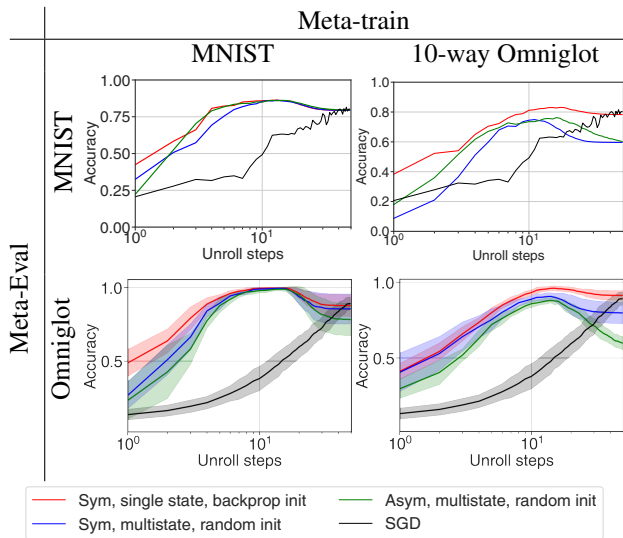


Figure 9. Accuracy of different variants of BLUR trained on MNIST (left) or 10 episodes from Omniglot (right) to 10 unrolls and evaluated on MNIST (top) and Omniglot (bottom). Black line corresponds to the SGD. Errorbars show standard deviation of the accuracy over 10 different subsets of Omniglot. Only the best result of 8 runs is plotted.

(c) initializing genome close to backpropagation (i.e. using $\nu = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$, $\mu = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$, $\tilde{\nu} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$, $\tilde{\mu} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ as defined in Sec. 3.1). We trained eight different combinations above on the MNIST dataset and on a 10-way episode learning from Omniglot (at every meta-iteration sampling 10 different random classes from the 1200 class subset of Omniglot). We then evaluate the resulting genomes on test set of MNIST or 10 different subsets of Omniglot tasks that were not part of the training set. In Fig. 9 we show three best variants of these parameters (the other five had much worse results and available in the Appendix) as well as comparison with SGD trained using cross entropy loss and a learning rate 0.001. We noticed that while having asymmetric forward and backward synapses allows our system to be strictly more general, it does not always lead to better generalization. The simplest variant with symmetric, single-state synapses tends to be the winner, however it has to be initialized from the backprop genome. Other two variants that performed well: symmetric, multi-state genomes and the most general asymmetric, multi-state genome.

Learning genomes for deeper and wider networks In this experiment we meta-train our genome on networks that contain 2, 3, and 4 hidden layers and explore their ability to generalize across changes in layer sizes and architecture, as measured by their performance on MNIST after 10 steps. We use the same setup as in section 4.1. The results are shown in Fig. 10. Interestingly we discover that genomes generalize from more complex architectures to less complex

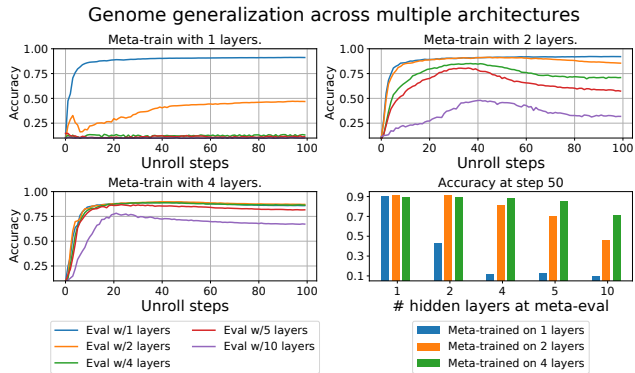


Figure 10. Genome generalization across architectures. Here we show that the genomes that were trained using deeper architectures work well on shallower architectures but not vice versa.

architectures, but not vice versa! For instance genomes trained using a 2-layer network performed well when tasked to use a single-hidden layer. However they diverged when training a 4-layer network. A 4-layer genome was able to train networks both with shallower (e.g. 1 to 3 layers) and deeper (10) architectures.

5. Conclusions and Future Work

In this work, we define a general protocol for updating nodes in a neural network, yielding a domain of “genomes” describing many possible update rules, of which gradient descent is one example. Useful genomes are identified by training networks on training tasks, and then their generalization is evaluated on unseen tasks. We have shown that it is possible to learn an entirely new type of neural network that can be trained to solve complex tasks faster than traditional neural networks of equivalent size, but without any notion of gradients. Our approach can be combined with many existing model representations with differentiable or non-differentiable components.

There are many interesting directions for future exploration. Perhaps, the most important one is the question of scale. Here one intriguing direction is the connection between the number of states and the learning capabilities. Another possible approach is extending the space of update rules, such as allowing injection of randomness for robustness, or providing an ability for neurons to self-regulate based on current state. Finally the ability to *extend* existing genomes to produce ever better learners, might help us scale even further. Another intriguing direction is incorporating the weight updates on both forward and backward passes. The former can be seen as a generalization of unsupervised learning, thus merging both supervised and unsupervised learning in one gradient-free framework.

References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL <http://tensorflow.org/>. Software available from tensorflow.org.
- Adigun, O. and Kosko, B. Bidirectional backpropagation. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 50(5):1982–1994, 2019.
- Ahmad, N., van Gerven, M. A. J., and Ambrogioni, L. GAIT-prop: A biologically plausible learning rule derived from backpropagation of error. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- Akrouf, M., Wilson, C., Humphreys, P. C., Lillicrap, T. P., and Tweed, D. B. Deep learning without weight transport. In Wallach, H. M., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E. B., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 974–982, 2019.
- Andrychowicz, M., Denil, M., Colmenarejo, S. G., Hoffman, M. W., Pfau, D., Schaul, T., and de Freitas, N. Learning to learn by gradient descent by gradient descent. In Lee, D. D., Sugiyama, M., von Luxburg, U., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pp. 3981–3989, 2016.
- Belilovsky, E., Eickenberg, M., and Oyallon, E. Greedy layerwise learning can scale to imagenet. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pp. 583–593. PMLR, 2019.
- Bengio, S., Bengio, Y., and Cloutier, J. On the search for new learning rules for ANNs. *Neural Process. Lett.*, 2(4): 26–30, 1995. doi: 10.1007/BF02279935.
- Bengio, Y. How auto-encoders could provide credit assignment in deep networks via target propagation. *CoRR*, abs/1407.7906, 2014.
- Bengio, Y., Lamblin, P., Popovici, D., and Larochelle, H. Greedy layer-wise training of deep networks. In Schölkopf, B., Platt, J. C., and Hofmann, T. (eds.), *Advances in Neural Information Processing Systems 19, Proceedings of the Twentieth Annual Conference on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 4-7, 2006*, pp. 153–160. MIT Press, 2006.
- Bengio, Y., Lee, D., Bornschein, J., and Lin, Z. Towards biologically plausible deep learning. *CoRR*, abs/1502.04156, 2015.
- Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., and Zhang, Q. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
- Camp, B., Mandivarapu, J. K., and Estrada, R. Continual learning with deep artificial neurons. *CoRR*, abs/2011.07035, 2020.
- Carandini, M. and Heeger, D. J. Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, 13(1):51–62, 2012.
- Cohen, G., Afshar, S., Tapson, J., and van Schaik, A. EMNIST: an extension of MNIST to handwritten letters. *CoRR*, abs/1702.05373, 2017. URL <http://arxiv.org/abs/1702.05373>.
- Confavreux, B., Zenke, F., Agnes, E. J., Lillicrap, T. P., and Vogels, T. P. A meta-learning approach to (re)discover plasticity rules that carve a desired function into a neural network. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *NeurIPS*, 2020.
- Conley, C. The gradient structure of a flow: I. *Ergodic Theory and Dynamical Systems*, 8(8*):11–26, 1988.
- Conley, C. C. *Isolated invariant sets and the Morse index*. American Mathematical Soc., 1978.
- Duchi, J., Hazan, E., and Singer, Y. Adaptive subgradient methods for online learning and stochastic optimization. *J. Mach. Learn. Res.*, 12(null):21212159, July 2011. ISSN 1532-4435.
- Farber, M., Kappeler, T., Latschev, J., and Zehnder, E. Smooth lyapunov 1-forms. *arXiv preprint math/0304137*, 2003.
- Franks, J. Notes on chain recurrence and lyapunov functions. *arXiv preprint arXiv:1704.07264*, 2017.

- Ha, D., Dai, A., and Le, Q. V. Hypernetworks. *arXiv preprint arXiv:1609.09106*, 2016.
- Hansen, N. and Ostermeier, A. Adapting arbitrary normal mutation distributions in evolution strategies: the covariance matrix adaptation. In *Proceedings of IEEE International Conference on Evolutionary Computation*, pp. 312–317, 1996. doi: 10.1109/ICEC.1996.542381.
- Hansen, N., Akimoto, Y., and Baudis, P. CMA-ES/pycma on Github. Zenodo, DOI:10.5281/zenodo.2559634, February 2019. URL <https://doi.org/10.5281/zenodo.2559634>.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015.
- Hebb, D. O. *The organization of behavior: a neuropsychological theory*. Science editions, 1949.
- Ioffe, S. and Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *CoRR*, abs/1502.03167, 2015. URL <http://arxiv.org/abs/1502.03167>.
- Kaplanis, C., Shanahan, M., and Clopath, C. Continual reinforcement learning with complex synapses. In Dy, J. G. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pp. 2502–2511. PMLR, 2018.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014.
- Kirsch, L. and Schmidhuber, J. Meta learning backpropagation and improving it. *arXiv preprint arXiv:2012.14905*, 2020.
- Lecun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. doi: 10.1109/5.726791.
- Lee, D., Zhang, S., Fischer, A., and Bengio, Y. Difference target propagation. In Appice, A., Rodrigues, P. P., Costa, V. S., Soares, C., Gama, J., and Jorge, A. (eds.), *Machine Learning and Knowledge Discovery in Databases - European Conference, ECML PKDD 2015, Porto, Portugal, September 7-11, 2015, Proceedings, Part I*, volume 9284 of *Lecture Notes in Computer Science*, pp. 498–515. Springer, 2015. doi: 10.1007/978-3-319-23528-8_31.
- Lee, J. M. *Introduction to Smooth Manifolds*. Springer, 2013.
- Lillicrap, T. P., Cownden, D., Tweed, D. B., and Akerman, C. J. Random synaptic feedback weights support error backpropagation for deep learning. *Nature communications*, 7(1):1–10, 2016.
- Lindsey, J. and Litwin-Kumar, A. Learning to learn with feedback and local plasticity. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- Löwe, S., O’Connor, P., and Veeling, B. S. Putting an end to end-to-end: Gradient-isolated learning of representations. In Wallach, H. M., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E. B., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 3033–3045, 2019.
- Maheswaranathan, N., Sussillo, D., Metz, L., Sun, R., and Sohl-Dickstein, J. Reverse engineering learned optimizers reveals known and novel mechanisms. *CoRR*, abs/2011.02159, 2020.
- Metz, L., Maheswaranathan, N., Nixon, J., Freeman, D., and Sohl-Dickstein, J. Understanding and correcting pathologies in the training of learned optimizers. In *International Conference on Machine Learning*, pp. 4556–4565. PMLR, 2019.
- Metz, L., Maheswaranathan, N., Freeman, C. D., Poole, B., and Sohl-Dickstein, J. Tasks, stability, architecture, and compute: Training more effective learned optimizers, and using them to train themselves. *arXiv preprint arXiv:2009.11243*, 2020.
- Miconi, T., Stanley, K. O., and Clune, J. Differentiable plasticity: training plastic neural networks with backpropagation. In Dy, J. G. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pp. 3556–3565. PMLR, 2018.
- Miconi, T., Rawal, A., Clune, J., and Stanley, K. O. Backpropamine: training self-modifying neural networks with differentiable neuromodulated plasticity. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.
- Minsky, M. and Papert, S. *Perceptrons: An Introduction to Computational Geometry*. MIT Press, Cambridge, MA, USA, 1969.

- Munkhdalai, T., Sordoni, A., Wang, T., and Trischler, A. Metalearned neural memory. *arXiv preprint arXiv:1907.09720*, 2019.
- Najarro, E. and Risi, S. Meta-learning through hebbian plasticity in random networks. *arXiv preprint arXiv:2007.02686*, 2020.
- Nøkland, A. Direct feedback alignment provides learning in deep neural networks. In Lee, D. D., Sugiyama, M., von Luxburg, U., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pp. 1037–1045, 2016.
- Norouzzadeh, M. S. and Clune, J. Neuromodulation improves the evolution of forward models. In *Proceedings of the Genetic and Evolutionary Computation Conference 2016*, pp. 157–164, 2016.
- Oja, E. Simplified neuron model as a principal component analyzer. *Journal of mathematical biology*, 15(3):267–273, 1982.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Pontes-Filho, S. and Liwicki, M. Bidirectional learning for robust neural networks. In *2019 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8. IEEE, 2019.
- Randazzo, E., Niklasson, E., and Mordvintsev, A. Mplp: Learning a message passing learning protocol. *arXiv preprint arXiv:2007.00970*, 2020.
- Ranganathan, V. and Lewandowski, A. ZORB: A derivative-free backpropagation algorithm for neural networks. *CoRR*, abs/2011.08895, 2020.
- Ravi, S. and Larochelle, H. Optimization as a model for few-shot learning. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017.
- Real, E., Liang, C., So, D., and Le, Q. Automl-zero: evolving machine learning algorithms from scratch. In *International Conference on Machine Learning*, pp. 8007–8019. PMLR, 2020.
- Robbins, H. and Monro, S. A stochastic approximation method. *Ann. Math. Statist.*, 22(3):400–407, 09 1951. doi: 10.1214/aoms/1177729586. URL <https://doi.org/10.1214/aoms/1177729586>.
- Rosenblatt, F. The perceptron: A perceiving and recognizing automaton. Report 85-460-1, Project PARA, Cornell Aeronautical Laboratory, Ithaca, New York, January 1957.
- Schmidhuber, J. *Evolutionary principles in self-referential learning, or on learning how to learn: the meta-meta... hook*. PhD thesis, Technische Universität München, 1987.
- Schmidt, R. M., Schneider, F., and Hennig, P. Descending through a crowded valley - benchmarking deep learning optimizers. *CoRR*, abs/2007.01547, 2020.
- Soltoggio, A., Stanley, K. O., and Risi, S. Born to learn: the inspiration, progress, and future of evolved plastic artificial neural networks. *Neural Networks*, 108:48–67, 2018.
- Taylor, G., Burmeister, R., Xu, Z., Singh, B., Patel, A. B., and Goldstein, T. Training neural networks without gradients: A scalable ADMM approach. In Balcan, M. and Weinberger, K. Q. (eds.), *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pp. 2722–2731. JMLR.org, 2016.
- Ulyanov, D., Vedaldi, A., and Lempitsky, V. S. Instance normalization: The missing ingredient for fast stylization. *CoRR*, abs/1607.08022, 2016. URL <http://arxiv.org/abs/1607.08022>.
- Velez, R. and Clune, J. Diffusion-based neuromodulation can eliminate catastrophic forgetting in simple neural networks. *PloS one*, 12(11):e0187736, 2017.
- Wichrowska, O., Maheswaranathan, N., Hoffman, M. W., Colmenarejo, S. G., Denil, M., de Freitas, N., and Sohl-Dickstein, J. Learned optimizers that scale and generalize. In *Proceedings of the 34th International Conference on Machine Learning*, 2017.
- Wilson, D. G., Cussat-Blanc, S., Luga, H., and Harrington, K. I. Neuromodulated learning in deep neural networks. *CoRR*, abs/1812.03365, 2018.
- Xiao, H., Rasul, K., and Vollgraf, R. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *CoRR*, abs/1708.07747, 2017. URL <http://arxiv.org/abs/1708.07747>.
- Xiao, W., Chen, H., Liao, Q., and Poggio, T. Biologically-plausible learning algorithms can scale to large datasets. *arXiv preprint arXiv:1811.03567*, 2018.

Xiong, Y., Ren, M., and Urtasun, R. LoCo: Local contrastive representation learning. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.