

---

# Rapid Adaptation with Conditionally Shifted Neurons

---

Tsendsuren Munkhdalai<sup>1</sup> Xingdi Yuan<sup>1</sup> Soroush Mehri<sup>1</sup> Adam Trischler<sup>1</sup>

## Abstract

We describe a mechanism by which artificial neural networks can learn rapid adaptation – the ability to adapt on the fly, with little data, to new tasks – that we call *conditionally shifted neurons*. We apply this mechanism in the framework of metalearning, where the aim is to replicate some of the flexibility of human learning in machines. Conditionally shifted neurons modify their activation values with task-specific shifts retrieved from a memory module, which is populated rapidly based on limited task experience. On metalearning benchmarks from the vision and language domains, models augmented with conditionally shifted neurons achieve state-of-the-art results.

## 1. Introduction

The ability to adapt our behavior rapidly in response to external or internal feedback is a primary ingredient of human intelligence. This cognitive flexibility is commonly ascribed to prefrontal cortex (PFC) and working memory in the brain. Neuroscientific evidence suggests that these areas use incoming information to support task-specific temporal adaptation and planning (Stokes et al., 2013; Siegel et al., 2015; Miller & Buschman, 2015; Brincat & Miller, 2016). This occurs on the fly, within only a few hundred milliseconds, and supports a wide variety of task-specific behaviors (Monsell, 2003; Sakai, 2008).

On the other hand, most existing machine learning systems are designed for a single task. They are trained through one optimization phase after which learning ceases. Systems built in such a *train-and-then-test* manner do not scale to complex, realistic environments: they require gluts of single-task data and are prone to issues related to distributional shifts, such as catastrophic forgetting (Srivastava et al., 2013; Goodfellow et al., 2014; Kirkpatrick et al., 2016) and

adversarial data points (Szegedy et al., 2013).

There is growing interest and progress in building flexible, adaptive models, particularly within the framework of *metalearning* (learning to learn) (Mitchell et al., 1993; Andrychowicz et al., 2016; Vinyals et al., 2016; Bachman et al., 2017). The goal of metalearning algorithms is the ability to learn new tasks efficiently, given little training data for each individual task. Metalearning models learn this ability (to learn) by training on a distribution of related tasks.

In this work we develop a neural mechanism for metalearning via rapid adaptation that we call *conditionally shifted neurons*. Conditionally shifted neurons (CSNs), like standard artificial neurons, produce activation values based on input from connected neurons modulated by the connection weights. Additionally, they have the capacity to shift their activation values on the fly based on auxiliary conditioning information. These conditional shifts adapt model behavior to the task at hand.

A model with CSNs operates in two phases: a description phase and a prediction phase. Assume, for each task  $\tau \sim p(\tau)$ , that we have access to a description  $D_\tau$ . In the simplest case, this is a set of example datapoints and their corresponding labels:  $D_\tau = \{(x'_i, y'_i)\}_{i=1}^n$ .<sup>1,2</sup> In the description phase, the model processes  $D_\tau$  and extracts conditioning information as a function of its performance on  $D_\tau$ . Based on this information, it generates activation shifts to adapt itself to the task and stores them in a key-value memory. In the prediction phase, the model acts on unseen datapoints  $x_j \in \tau$ , from the same task, to predict their labels  $y_j$ . To improve these predictions, the model retrieves shifts from memory and applies them to the activations of individual neurons. During training, the model learns the meta procedure of *how* to extract conditioning information in the description phase and generate useful conditional shifts for the prediction phase. At test time, it uses this procedure to adapt itself to new tasks from  $p(\tau)$ .

We define and investigate two forms of conditioning information in this work (§2.3). The first uses the gradient of

---

<sup>1</sup>Microsoft Research, Montréal, Québec, Canada. Correspondence to: Tsendsuren Munkhdalai <tsendsuren.munkhdalai@microsoft.com>.

<sup>1</sup>More abstractly, a task description could be given by a set of instructions or demonstrations of expert behavior.

<sup>2</sup>In a  $C$ -way,  $k$ -shot classification task,  $n = k \times C$ .

the model’s prediction loss with respect to network pre-activations, computed on the description data; the second replaces the loss gradient with direct feedback alignment (Lillicrap et al., 2016; Nøkland, 2016). The direct feedback information is computationally cheaper and leads to competitive or superior performance in our experiments. Note that other sources of conditioning are possible.

Our proposed neuron-level adaptation has several advantages over previous methods for metalearning that adapt the connections between neurons, for instance via fast weights (Munkhdalai & Yu, 2017) or an optimizer (Finn et al., 2017; Ravi & Larochelle, 2017). First, it is more efficient computationally, since the number of neurons is generally much less than the number of weight parameters (e.g., the number of weights scales quadratically in the number of neurons per layer for a fully connected network). Second, conditionally shifted neurons can be incorporated into various neural architectures, including convolutional and recurrent networks, without special modifications to suit the structure of such models.

After describing the details of our framework, we demonstrate experimentally that ResNet (He et al., 2016) and deep LSTM (Hochreiter & Schmidhuber, 1997) models equipped with CSNs achieve 56.88% and 71.94% accuracy on the standard Mini-ImageNet 1- and 5-shot benchmarks, and 41.25%, 52.1%, and 57.8% accuracy on Penn Treebank 1-, 2-, and 3-shot language modeling tasks. These results mark a significant improvement over the previous state of the art.

Our primary contributions in this paper are as follows: (i) we propose a generic neural mechanism, *conditionally shifted neurons*, by which learning systems can adapt on the fly; (ii) we introduce *direct feedback* as a computationally inexpensive metalearning signal; and (iii) we implement and evaluate conditionally shifted neurons in several widely-used neural network architectures.<sup>3</sup>

## 2. Conditionally Shifted Neurons

The core idea of *conditionally shifted neurons* is to modify a network’s activation values on the fly, by shifting them as a function of auxiliary conditioning information. A layer with CSNs takes the following form:

$$h_t = \begin{cases} \sigma(a_t) + \sigma(\beta_t) & t \neq T \\ \text{softmax}(a_t + \beta_t) & t = T \end{cases} \quad (1)$$

for hidden layer  $t$  or output layer  $T$  (which represents a probability distribution). The pre-activation vector  $a_t \in \mathbb{R}^{L_t}$ , for a layer with  $L_t$  neurons, can take various forms depending on the network architecture (fully connected, convolutional, etc.). The nonlinear function  $\sigma$  computes an element-wise

activation.  $\beta_t \in \mathbb{R}^{L_t}$  is the layer-wise conditional shift vector, determined from layer-wise conditioning information  $I_t$  (defined in §2.3).

To implement a model with CSNs, we must define functions that extract and transform the conditioning information  $I_t$  into the shifts  $\beta_t$ . For this we build on the *MetaNet* architecture of Munkhdalai & Yu (2017). MetaNet consists of a base learner plus a shared meta learner with working memory. For each task  $\tau$ , MetaNet processes the task description  $D_\tau = \{(x'_i, y'_i)\}_{i=1}^n$  and stores relevant “meta information” in a key-value memory. To classify unseen examples  $x_j$  from the described task, the model queries its working memory with an attention mechanism to generate a set of fast weights; these modify the base learner, which in turn predicts labels  $y_j$ .

To begin, we describe model details for a fully connected feed-forward network (FFN) with CSNs. The architecture is depicted in Figure 1. As shown, the model factors into a *base learner*, which makes predictions on inputs, and a *meta learner*. The meta learner extracts conditioning information from the base learner and uses a key-value memory to store and retrieve activation shifts. After walking through the model details, we define error gradient and direct feedback variants of the conditioning information  $I_t$  (§2.3). In §2.4 and §2.5 we describe how CSNs can be added to ResNet and LSTM architectures, respectively.

### 2.1. Feed-Forward Networks with Conditionally Shifted Neurons

Our model operates in two phases: a description phase, wherein it processes the task description  $D_\tau = \{(x'_i, y'_i)\}_{i=1}^n$ , and a prediction phase, wherein it acts on unseen datapoints  $x_j$  to predict their labels  $y_j$ . In an episode of training or test, we sample a task from  $p(\tau)$ . The model then ingests the task description and uses what it learns therefrom, via the conditioning information, to make predictions on unseen task data.

#### 2.1.1. BASE LEARNER

The base learner maps an input datapoint to its label prediction through layers described by equation 1, where in the FFN case, the pre-activation vector  $a_t$  is given by  $a_t = W_t h_{t-1} + b_t$ . Weight matrix  $W_t$  and bias vector  $b_t$  are learned parameters.

The base learner operates similarly in both phases. During the **description phase**, the base learner’s input is a datapoint  $x'_i$  from  $D_\tau$ . Its softmax output is an estimate  $\hat{y}'_i$  for the label  $y'_i$ . The conditional shifts  $\beta_t$  in eq. 1 are set to 0 in this phase.

During the **prediction phase**, the base learner operates on inputs  $x_j$ . It receives conditional shifts  $\beta_t$  from the meta learner and applies them layer-wise according to eq. 1. Con-

<sup>3</sup>Code and data will be available at <https://aka.ms/csns>




Figure 1. Schematic illustration of our model with conditionally shifted neurons. In the description phase, the meta learner populates working memory with keys and values, based on the base learner’s performance on the task description; in the prediction phase, the meta learner retrieves task-specific shifts from memory through key-based attention and feeds them to the base learner to adapt it to the task.

ditioned on these shifts, the base learner computes an estimate  $\hat{y}_j$  for the label  $y_j$ .

### 2.1.2. META LEARNER

The meta learner’s operation is more complicated and differs more significantly from phase to phase.

During the **description phase**, as the base learner processes  $x'_i \in D_\tau$ , the meta learner extracts layer-wise conditioning information for this example,  $I_{t,i}$ , according to eq. 6 or 7. The meta learner uses the conditioning information to generate memory values. These act as “template” conditional shifts for the task, and are computed via the memory function  $g$ :

$$V_{t,i} = g(I_{t,i}), \quad (2)$$

where  $V_{t,i} \in \mathbb{R}^{L_t}$  encodes the shift template at layer  $t$  for input  $x'_i$ . There are  $n$  of these; we arrange them into matrix  $V_t \in \mathbb{R}^{n \times L_t}$  over the full task description.

For parsimony, we desire a single memory function  $g$  for all layers of the base learner, which may have different sizes  $L_t$ . Therefore, we parameterize  $g$  as a multi-layer perceptron (MLP) that operates independently on the vector of conditioning information for each neuron (defined in §2.3). More sophisticated  $L_t$ -agnostic parameterizations for  $g$  are possible, such as recurrent networks.

In parallel during the description phase, the meta learner constructs an embedded representation of the input that it uses to key the memory. This is the objective of the key function,  $f$ , which we parameterize here as an MLP with a linear output layer. The key function generates, for each description input, a  $d$ -dimensional key vector  $k'_i = f(x'_i)$ .

At **prediction** time, the meta learner generates a memory query  $k_j$  from input  $x_j$  using the key function. It uses  $k_j$  to recall layer-wise shifts  $\beta_t$  from memory via soft attention:

$$\alpha = \underset{i}{\text{softmax}}(\cos(k_j, k'_i)), \quad (3)$$

$$\beta_t = \alpha^\top V_t. \quad (4)$$

Note that keys correspond to inputs, not base-learner layers. The meta learner finally feeds the layer-wise shifts  $\beta_t$  to the base learner to condition the computation of  $\hat{y}_j$ .

## 2.2. Training and Test

We train and test the model in episodes. For each episode: we sample a training or test task from  $p(\tau)$ , process its description  $D_\tau$ , and then feed its unseen data forward to obtain their label predictions. Training and test tasks are both drawn from the same distribution, but crucially, we partition the data such that the classes seen at training time do not overlap with those seen at test time.

Over a collection of training episodes, we optimize the model parameters end-to-end via stochastic gradient descent (SGD). Gradients are taken with respect to the (cross-entropy) task losses,  $\mathcal{L}_\tau = \sum_j \mathcal{L}_{\text{CE}}(\hat{y}_j, y_j)$ . In this scheme, the model’s computational graph contains the operations for processing the description, like the transformation of conditioning information and the generation of memory keys and values. Parameters of these operations are also optimized.

### 2.3. Conditioning Information

**Error gradient information** Inspired by the success of MetaNets, we first consider gradients of the base learner’s loss on the task description as the conditioning information. To compute these error gradients we apply the chain rule and the standard backpropagation algorithm to the base learner. Given a true label  $y'_i$  from the task description and the model’s corresponding label prediction  $\hat{y}'_i$ , we obtain loss gradients for base-learner neurons at layer  $t$  as

$$\nabla_{t,i} = \frac{\partial \mathcal{L}(\hat{y}'_i, y'_i)}{\partial a_t}, \quad (5)$$

where  $a_t$  is the  $L_t$ -dimensional vector of pre-activations at layer  $t$ ,  $\nabla_{t,i}$  has the same size, and we denote with  $\mathcal{L}(\cdot)$  a loss function (such as the cross entropy loss on the labels). Note that  $\mathcal{L}$  here is *not* the target of optimization via SGD.

We obtain the conditioning information  $I_{t,i,\ell}$  for each neuron (indexed by  $\ell$ ) using the gradient preprocessing formula of Andrychowicz et al. (2016):

$$I_{t,i,\ell} = \begin{cases} \left( \frac{\log(|\nabla_{t,i,\ell}|)}{p}, \text{sgn}(\nabla_{t,i,\ell}) \right) & \text{if } |\nabla_{t,i,\ell}| \geq e^{-p} \\ (-1, e^p \nabla_{t,i,\ell}) & \text{otherwise} \end{cases} \quad (6)$$

where  $\text{sgn}$  is the signum function and we set  $p = 7$ . We use this preprocessing to smooth variation in  $I_{t,i,\ell}$ , since gradients with respect to different base-learner activations can have very different magnitudes. By eq. 6, each neuron obtains a 2-dimensional vector of conditioning information. In this case, we can interpret eq. 1 as a one-step, transformed gradient update on the neuron activations via  $\beta_t$ . “Raw” gradients are transformed through preprocessing, the memory read and write operations, and the nonlinearity  $\sigma$ .

Because backpropagation is inherently sequential, this information is expensive to compute. It becomes increasingly costly for deeper networks, such as RNNs processing long sequences.

**Direct feedback information** Direct feedback (DF) information is inspired by feedback alignment methods (Lillicrap et al., 2016; Nøkland, 2016) and biologically plausible deep learning (Bengio et al., 2015). We obtain the DF information for base-learner neurons at layer  $t$  as

$$I_{t,i,\ell} = \sigma'(a_{t,\ell}) \cdot (\hat{y}'_i - y'_i), \quad (7)$$

where  $\sigma'(\cdot)$  represents the derivative of the nonlinear activation function  $\sigma$  and  $(\hat{y}'_i - y'_i)$  is the derivative of the cross entropy loss with respect to the softmax input. Thus, the DF conditioning information for each neuron is the derivative of the loss function scaled by the derivative of the activation function. In the DF case, each neuron obtains a  $C$ -dimensional vector of information, with  $C$  the number of output classes. We can compute this conditioning information for all neurons in a network simultaneously, with a single multiplication. This is more efficient than sequentially locked backpropagation-based error gradients. Furthermore, to obtain DF information, it is sufficient that only the loss and neuron activation functions are differentiable. This is more relaxed than for backpropagation methods. We demonstrate the effectiveness of both conditioning variants in §4.

### 2.4. Deep Residual Networks with CSNs

For ResNets (He et al., 2016) we incorporate conditionally shifted neurons into the output of a residual block. Let us denote the residual block as ResBlock, which is defined as follows:

$$\begin{aligned} h^1 &= \text{ReLU}(\text{conv}(x)) \\ h^2 &= \text{ReLU}(\text{conv}(h^1)) \\ h^3 &= \text{conv}(h^2) \\ h^4 &= \text{conv}(x) \\ a_t &= h^3 + h^4 \end{aligned}$$

where  $x$  and  $a_t$  are the inputs to the block and the output pre-activations, respectively. Function  $\text{conv}$  denotes a convolutional layer, which may optionally be followed by a batch normalization (Ioffe & Szegedy, 2015) layer. The activations  $h_t$  of the CSNs for the ResBlock are computed as:

$$h_t = \sigma(a_t) + \sigma(\beta_t)$$

where  $\beta_t$  is the task-specific shift retrieved from the memory, constructed based on the activation values  $a_t$  analogously to the FFN case; i.e., the conditioning information is computed for neurons at the output of each residual block. We stack several residual blocks with CSNs to construct a deep adaptive ResNet model. We use the ReLU function as the nonlinearity  $\sigma$  in this model.

### 2.5. Long Short-Term Memory Networks with CSNs

Given the current input  $x_t$ , the previous hidden state  $h_{t-1}$ , and the previous memory cell state  $c_{t-1}$ , an LSTM model with CSNs computes its gates, new memory cell states, and

hidden states at time step  $t$  with the following update rules:

$$\begin{aligned} i_t &= \text{Sigmoid}(W_i[x_t; h_{t-1}] + b_i) \\ f_t &= \text{Sigmoid}(W_f[x_t; h_{t-1}] + b_f) \\ o_t &= \text{Sigmoid}(W_o[x_t; h_{t-1}] + b_o) \\ c_t &= \sigma(W_v[x_t; h_{t-1}] + b_v) \odot i_t + c_{t-1} \odot f_t \\ h_t &= (\sigma(c_t) + \sigma(\beta_t)) \odot o_t \end{aligned}$$

where  $\odot$  represents element-wise multiplication,  $[\cdot; \cdot]$  is concatenation, and  $\beta_t$  is the task-specific shift from the memory. In the LSTM case, the memory is constructed by processing conditioning information extracted from the memory cell  $c_t$ . By stacking such layers together we build a deep LSTM model that adapts across both depth and time. We use the tanh function as the nonlinearity  $\sigma$  in this model.

### 3. Related Work

Among the many problems in supervised, reinforcement, and unsupervised learning that can be framed as metalearning, few-shot learning has emerged as a natural and popular test bed. Few-shot supervised learning refers to a scenario where a learner is introduced to a sequence of tasks, where each task entails multi-class classification given a single or very few labeled examples per class. A key challenge in this setting is that the classes or concepts vary across the tasks; thus, models require a capacity for rapid adaptation in order to recognize new concepts on the fly.

Few-shot learning problems were previously addressed using metric learning methods (Koch, 2015). Recently, there has been a shift towards building flexible models for these problems within the learning-to-learn paradigm (Mishra et al., 2017; Santoro et al., 2016). Vinyals et al. (2016) unified the training and testing of a one-shot learner under the same procedure and developed an end-to-end, differentiable nearest-neighbor method for one-shot learning. More recently, one-shot optimizers were proposed by Ravi & Larochelle (2017); Finn et al. (2017). The MAML framework (Finn et al., 2017) learns a parameter initialization from which a model can be adapted rapidly to a given task using only a few steps of gradient updates. To learn this initialization it makes use of more sophisticated second-order gradient information. Here we harness only first-order gradient information, or the simpler direct feedback information.

As highlighted, the architecture of our model with conditionally shifted neurons is closely related to Meta Networks (Munkhdalai & Yu, 2017). The MetaNet modifies synaptic connections (weights) between neurons using fast weights (Schmidhuber, 1987; Hinton & Plaut, 1987) to implement rapid adaptation. While MetaNet’s fast weights enable flexibility, it is very expensive to modify these weights when the connections are dense. Neuron-level adaptation as proposed in this work is significantly more efficient while lending

itself to a range of network architectures, including ResNet and LSTM. Other previous work on metalearning has also formulated the problem as two-level learning: specifically, “slow” learning of a meta model across several tasks, and “fast” learning of a base model that acts within each task (Schmidhuber, 1987; Bengio et al., 1990; Hochreiter et al., 2001; Mitchell et al., 1993; Vilalta & Drissi, 2002; Mishra et al., 2017). Schmidhuber (1993) discussed the use of network weight matrices themselves for continuous adaptation in dynamic environments.

Viewed as a form of feature-wise transformation, CSNs are closely related to conditional normalization techniques (Lei Ba et al., 2015; Dumoulin et al., 2016; Ghiasi et al., 2017; De Vries et al., 2017; Perez et al., 2017). FiLM (Perez et al., 2017), the most similar such approach which was inspired by Dumoulin et al. (2016) and Ghiasi et al. (2017), modulates CNN *feature maps* using global scale and shift operations conditioned on an auxiliary input modality. In contrast, CSNs apply shifts to individual neurons’ activations, locally, and this modification is based on the model’s behavior on the task description rather than the input itself.

In the case of gradient-based conditioning information, our approach can be viewed as a synthesis of a conditional normalization model (in the style of FiLM) with a learned optimizer (in the style of Andrychowicz et al. (2016)). Specifically, the learned memory and key functions,  $g$  and  $f$ , transform error gradients into the conditioning shifts  $\beta_t$ , which are then applied like a one-step update to the activation values. A CSN model uses this learned optimizer on the fly.

### 4. Experimental Evaluation

We evaluate the proposed CSNs on tasks from the vision and language domains. Below we describe the datasets we evaluate on and the according preprocessing steps, followed by test results and an ablation study.

#### 4.1. Few-shot Image Classification

In the vision domain, we used two widely adopted few-shot classification benchmarks: the Omniglot and Mini-ImageNet datasets.

**Omniglot** consists of images from 1623 classes from 50 different alphabets, with only 20 images per class (Lake et al., 2015). As in previous studies, we randomly selected 1200 classes for training and 423 for testing and augmented the training set with 90, 180 and 270 degree rotations. We resized the images to  $28 \times 28$  pixels for computational efficiency.

For the Omniglot benchmark we performed 5- and 20-way classification tests, each with one or five labeled examples from each class as the description  $D_\tau$ . We use a convolu-

Table 1. Omniglot few-shot classification test accuracy for error gradient ( $\nabla$ ) and direct feedback (DF) conditioning information.

Model	5-way		20-way	
	1-shot	5-shot	1-shot	5-shot
Siamese Net (Koch, 2015)	97.3	98.4	88.2	97.0
MANN (Santoro et al., 2016)	82.8	94.9	-	-
Matching Nets (Vinyals et al., 2016)	98.1	98.9	93.8	98.5
MAML (Finn et al., 2017)	<b>98.7 <math>\pm</math> 0.4</b>	<b>99.9 <math>\pm</math> 0.3</b>	95.8 $\pm$ 0.3	98.9 $\pm$ 0.2
MetaNet (Munkhdalai & Yu, 2017)	<b>98.95</b>	-	97.0	-
TCML (Mishra et al., 2017)	<b>98.96 <math>\pm</math> 0.2</b>	99.75 $\pm$ 0.11	<b>97.64 <math>\pm</math> 0.3</b>	<b>99.36 <math>\pm</math> 0.18</b>
adaCNN ( $\nabla$ )	98.41 $\pm$ 0.16	99.27 $\pm$ 0.12	95.95 $\pm$ 0.43	98.48 $\pm$ 0.06
adaCNN (DF)	98.42 $\pm$ 0.21	99.37 $\pm$ 0.28	96.12 $\pm$ 0.31	98.43 $\pm$ 0.05

 Table 2. Mini-ImageNet few-shot classification test accuracy for error gradient ( $\nabla$ ) and direct feedback (DF) conditioning information.

Model	5-way	
	1-shot	5-shot
Matching Nets (Vinyals et al., 2016)	43.6	55.3
MetaLearner LSTM (Ravi & Larochelle, 2017)	43.4 $\pm$ 0.77	60.2 $\pm$ 0.71
MAML (Finn et al., 2017)	48.7 $\pm$ 1.84	63.1 $\pm$ 0.92
MetaNet (Munkhdalai & Yu, 2017)	49.21 $\pm$ 0.96	-
adaCNN ( $\nabla$ )	48.26 $\pm$ 0.63	62.80 $\pm$ 0.41
adaCNN (DF)	48.34 $\pm$ 0.68	62.00 $\pm$ 0.55
TCML (Mishra et al., 2017)	55.71 $\pm$ 0.99	68.88 $\pm$ 0.92
adaResNet ( $\nabla$ )	56.62 $\pm$ 0.69	71.69 $\pm$ 0.67
adaResNet (DF)	<b>56.88 <math>\pm</math> 0.62</b>	<b>71.94 <math>\pm</math> 0.57</b>

tional network (CNN) with 64 filters as the base learner. This network has 5 convolutional layers, each of which uses  $3 \times 3$  convolutions followed by the ReLU nonlinearity and a  $2 \times 2$  max-pooling layer. Convolutional layers are followed by a fully connected (FC) layer with softmax output. Another CNN with the same architecture is used for the key function  $f$ . We use CSNs in the last four layers of the CNN components, referring to this model as “adaCNN.” Full implementation details can be found in Appendix A.

Table 1 shows that our adaCNN model achieves competitive, though not state-of-the-art, results on the Omniglot tasks. There is an obvious ceiling effect among the best performing models as accuracy saturates near 100%.

**Mini-ImageNet** features  $84 \times 84$ -pixel color images from 100 classes (64/16/20 for training/validation/test splits) and each class has 600 exemplar images. We ran our experiments on the class subset released by Ravi & Larochelle (2017). Compared to Omniglot, Mini-ImageNet has fewer classes (100 vs 1623) with more labeled examples provided of each class (600 vs 20). Given this larger number of examples, we evaluated a similar adaCNN model with 32 filters as well as a model with more sophisticated ResNet components (“adaResNet”) on the Mini-ImageNet 5-way classification tasks. The ResNet architecture follows that of TCML (Mishra et al., 2017) with two exceptions due to

memory constraints. Instead of two  $1 \times 1$  convolutional layers with 2048 and 512 filters we use only a single such layer with 1024 filters, and the ReLU nonlinearity instead of its leaky variant. We incorporate CSNs into the last two residual blocks as well as the two fully connected output layers. Full implementation details can be found in Appendix A.

For every 400 training tasks, we tested the model for another 400 tasks sampled from the validation set. If the model performance exceeded the previous best validation result, we applied it to the test set. Following previous approaches that we compare with in Table 2, we sampled another 400 tasks randomly from the test set to report model accuracy.

Unlike Omniglot, there remains significant room for improvement on Mini-ImageNet. As shown in Table 2, on this more challenging task, CNN-based models with conditionally shifted neurons achieve performance just below that of the best CNN-based approaches like MAML and MetaNet (recall that these modify weight parameters rather than activation values). The more sophisticated adaResNet model, on the other hand, achieves state-of-the-art results. The best-performing adaResNet (DF) yields almost 10% improvement over the corresponding adaCNN model and improves over the previous best result of TCML by 1.16% and 3.06% on the one and five shot 5-way classification tasks, respectively. Note that TCML likewise uses a ResNet

Table 3. Penn Treebank few-shot classification test accuracy for error gradient ( $\nabla$ ) and direct feedback (DF) conditioning information.

Model	5-way (400 random/all-inclusive)		
	1-shot	2-shot	3-shot
LSTM-LM oracle (Vinyals et al., 2016)	72.8	72.8	72.8
Matching Nets (Vinyals et al., 2016)	32.4	36.1	38.2
2-layer LSTM + adaFFN ( $\nabla$ )	32.55/33.2	44.15/46.0	50.4/51.7
1-layer adaLSTM ( $\nabla$ )	36.55/37.7	43.25/44.6	50.7/52.1
2-layer adaLSTM ( $\nabla$ )	<b>43.1/43.0</b>	<b>52.05/54.2</b>	<b>57.35/58.4</b>
2-layer LSTM + adaFFN (DF)	33.65/35.3	46.6/47.8	51.4/52.6
1-layer adaLSTM (DF)	36.35/36.3	41.6/43.4	49.1/50.1
2-layer adaLSTM (DF)	<b>41.25/43.2</b>	<b>52.1/52.9</b>	<b>57.8/58.8</b>

architecture. The best accuracy among five different seed runs of the adaResNet with DF conditioning was 72.91% on the five-shot task.

#### 4.2. Few-shot Language Modeling

To evaluate the effectiveness of recurrent models with conditionally shifted neurons, we ran experiments on the few-shot **Penn Treebank** (PTB) language modeling task introduced by Vinyals et al. (2016).

In this task, a model is given a query sentence with one missing word and a support set (i.e., description) of one-hot-labeled sentences that also have one missing word each. One of the missing words in the description set is identical to that missing from the query sentence. The model must select the label of this corresponding sentence.

Following Vinyals et al. (2016), we split the PTB sentences into training and test such that, for the test set, target words for prediction and the sentences in which they appear are unseen during training. Concretely, we removed the test target words as well as sentences containing those words from the training data. This process necessarily reduces the training data and increases out-of-vocabulary (OOV) test words. We used the same 1000 target words for testing as provided by Vinyals et al. (2016).

We evaluated two models with conditionally shifted neurons on 1-, 2-, and 3-shot language modelling (LM) tasks. In both cases, we represent words with randomly initialized dense embeddings. For the first model we stacked a 3-layer feed-forward net with CSNs (adaFFN) on top of an LSTM network (LSTM+adaFFN) at each prediction timestep. In this model, only the adaFFN can adapt to the task while it processes the hidden state of the underlying LSTM. The LSTM encoder builds up the context for each word and provides a generic (non-task-specific) representation to the adaFFN. Both components are trained jointly.

The second model we propose for this task is more flexible, an LSTM with conditionally shifted neurons in the recurrence (adaLSTM). This entire model is adapted with

task-specific shifts at every time step. For few-shot classification output, a softmax layer with CSNs is stacked on top of the adaLSTM. Comparing LSTM+adaFFN and adaLSTM, the former is much faster since we only adapt the activations of the three feedforward layers, but it lacks full flexibility since the LSTM is unaware of the current task information. We also evaluated deep (2-layer) versions of both LSTM+adaFFN and adaLSTM models. Full implementation details can be found in Appendix A.

We used two different methods to form test tasks for evaluation. First, we randomly sampled 400 tasks from the test data and report the average accuracy. Second, we make sure to include all test words in the task formulation. We randomly partition the 1000 target words into 200 groups and solve each group as a task. In the random approach there is a chance that a word could be missed or included multiple times in different tasks. However, the random approach also enables formulation of an exponential number of test tasks.

Table 3 summarizes our results. The approximate upper bound achieved by the oracle LSTM-LM of Vinyals et al. (2016) is 72.8%. Our best accuracy – around 58% on the 3-shot task – comes using a 2-layer adaLSTM, and improves over the Matching Nets results by 11.1%, 16.0% and 19.6% for 1-, 2-, and 3-shot tasks, respectively. Comparing model variants, adaLSTM consistently outperforms the standard LSTM augmented with a conditionally shifted output FFN, and deeper models yield higher accuracy. Providing more sentences for the target word increases performance, as expected. These results indicate that our model’s few-shot language modelling capabilities far exceed those of Matching Networks (Vinyals et al., 2016). Some of this improvement surely arises from adaLSTM’s recurrent structure, which is known to apply well to sequence-based tasks and in the language domain. However, it is one of the strengths of conditionally shifted neurons that they can be ported easily to various neural architectures.

Comparing direct feedback information to the gradient-based variant across the full suite of experiments, we observe overall that DF information performs competitively

well. Even more positively, DF information speeds up the the model runtime considerably. For example, the 2-layer adaLSTM processed 400 test episodes in 200 seconds using gradient information vs. 160 seconds for DF information, representing a speedup of about 25%. In Appendix B, we compare runtimes for the adaCNN variants and MetaNet model on the Mini-ImageNet task. Even for the shallow adaCNN model, we observe 2-3 ms/task speed-up with the DF conditioning information.

### 4.3. Ablation Study

To better understand our model, we performed an ablation study on adaCNN trained on Mini-ImageNet and the 1-layer adaLSTM trained on PTB. Results are shown in Figure 2 and Figure 3, respectively.




Figure 2. Model ablation for adaCNN tested on the Mini-ImageNet one-shot task. Blue:  $g$  as a scalar multiplier of  $\nabla_{t,i}$  (gradient case) or perceptron (DF case); Green:  $\beta$  without normalization; Red: baseline model.




Figure 3. Model ablation on the one-shot language modeling task, for a single layer adaLSTM. We report average accuracy on 400 random test tasks. Yellow:  $g$  as a scalar multiplier of  $\nabla_{t,i}$  (gradient case) or perceptron (DF case); Violet:  $\beta$  without normalization; Grey: baseline model.

Our first ablation was to determine the effect of normalizing the task shifts  $\beta_t$  through the nonlinear activation function

$\sigma$ . Ablating the activation function from eq. 1 and simply adding  $\beta_t$  resulted in a slight performance drop on the Mini-ImageNet task and a significant decrease (around 7%) on the one-shot LM task, for both variants of conditioning information. We similarly tried adding  $\beta_t$  directly to the pre-activations  $a_t$  inside the nonlinearity  $\sigma$ . On Mini-ImageNet, this variant of adaCNN achieved 47.8% and 48.09% accuracy with gradient and DF conditioning information, respectively, which is competitive with the baseline. However, adaLSTM performance decreased more significantly to 30.7/31.8% and 33.25/34.05% on the LM task. We conclude that squashing the conditional shift  $\beta_t$  to the same range as the neuron’s standard activation value is beneficial.

Our second ablation evaluates variations on the function  $g$  for transforming the conditioning information  $I_{t,i}$  into the memory values  $V_{t,i}$ . In the case of gradient-based conditioning information, we remove the preprocessing of eqn. 6 and replace the learned MLP with a simple learned scalar,  $\lambda \in \mathbb{R}$ , that multiplies the gradient vector  $\nabla_{t,i}$ . In this case we can more clearly interpret the conditional shift as a one-step gradient update on the activation values (although this update is still modulated by the memory read procedure and the function  $\sigma$ ). As shown in Figure 2, the adaCNN model with learned scaling loses about 3 percentage points of accuracy on the image classification task. However, as per Figure 3, adaLSTM performance plummets with learned scaling, dropping to 20% test accuracy (random chance).

For direct feedback conditioning information, we cannot use a scalar parameter because we have a  $C$ -dimensional information vector for each neuron (recall §2.3). We therefore parameterize  $g$  as a one-layer perceptron in this case rather than a deep MLP. Using a simple perceptron to process the direct feedback information decreased test accuracy significantly on the Mini-ImageNet and LM tasks (drops of over 10%). This highlights that a deeper mapping function is crucial for processing DF conditioning information.

## 5. Conclusion

We introduced *conditionally shifted neurons*, a mechanism for rapid adaptation in neural networks. Conditionally shifted neurons are generic and easily incorporated into various neural architectures. They are also computationally efficient compared to alternative metalearning methods that adapt synaptic connections between neurons. We proposed two variants of conditioning information for use with CSNs, one based on error gradients and another based on feedback alignment methods. The latter is more efficient because it does not require a sequential backpropagation procedure, and achieves competitive performance with the former. We demonstrated empirically that models with conditionally shifted neurons improve the state of the art on metalearning benchmarks from the vision and language domains.



## References

- Andrychowicz, M., Denil, M., Gomez, S., Hoffman, M. W., Pfau, D., Schaul, T., and de Freitas, N. Learning to learn by gradient descent by gradient descent. In *Advances in Neural Information Processing Systems*, pp. 3981–3989, 2016.
- Bachman, P., Sordoni, A., and Trischler, A. Learning algorithms for active learning. In Precup, D. and Teh, Y. W. (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 301–310, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR. URL <http://proceedings.mlr.press/v70/bachman17a.html>.
- Bengio, Y., Bengio, S., and Cloutier, J. *Learning a synaptic learning rule*. Université de Montréal, Département d’informatique et de recherche opérationnelle, 1990.
- Bengio, Y., Lee, D.-H., Bornschein, J., Mesnard, T., and Lin, Z. Towards biologically plausible deep learning. *arXiv preprint arXiv:1502.04156*, 2015.
- Brincat, S. L. and Miller, E. K. Prefrontal cortex networks shift from external to internal modes during learning. *Journal of Neuroscience*, 36(37):9739–9754, 2016.
- De Vries, H., Strub, F., Mary, J., Larochelle, H., Pietquin, O., and Courville, A. C. Modulating early visual processing by language. In *Advances in Neural Information Processing Systems*, pp. 6597–6607, 2017.
- Dumoulin, V., Shlens, J., and Kudlur, M. A learned representation for artistic style. *CoRR*, *abs/1610.07629*, 2(4): 5, 2016.
- Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. In Precup, D. and Teh, Y. W. (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1126–1135, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR. URL <http://proceedings.mlr.press/v70/finn17a.html>.
- Ghiasi, G., Lee, H., Kudlur, M., Dumoulin, V., and Shlens, J. Exploring the structure of a real-time, arbitrary neural artistic stylization network. *arXiv preprint arXiv:1705.06830*, 2017.
- Goodfellow, I. J., Mirza, M., Xiao, D., Courville, A., and Bengio, Y. An empirical investigation of catastrophic forgetting in gradient-based neural networks. In *ICLR 2014*, 2014.
- He, K., Zhang, X., Ren, S., and Sun, J. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pp. 1026–1034, 2015.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.
- Hinton, G. E. and Plaut, D. C. Using fast weights to deblur old memories. In *Proceedings of the ninth annual conference of the Cognitive Science Society*, pp. 177–186, 1987.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Hochreiter, S., Younger, A. S., and Conwell, P. R. Learning to learn using gradient descent. In *International Conference on Artificial Neural Networks*, pp. 87–94. Springer, 2001.
- Ioffe, S. and Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32Nd International Conference on International Conference on Machine Learning - Volume 37, ICML’15*, pp. 448–456. JMLR.org, 2015. URL <http://dl.acm.org/citation.cfm?id=3045118.3045167>.
- Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., et al. Overcoming catastrophic forgetting in neural networks. *arXiv preprint arXiv:1612.00796*, 2016.
- Koch, G. *Siamese neural networks for one-shot image recognition*. PhD thesis, University of Toronto, 2015.
- Lake, B. M., Salakhutdinov, R., and Tenenbaum, J. B. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.
- Lei Ba, J., Swersky, K., Fidler, S., et al. Predicting deep zero-shot convolutional neural networks using textual descriptions. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 4247–4255, 2015.
- 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, 2016.
- Miller, E. K. and Buschman, T. J. Working memory capacity: Limits on the bandwidth of cognition. *Daedalus*, 144(1):112–122, 2015.

- Mishra, N., Rohaninejad, M., Chen, X., and Abbeel, P. Meta-learning with temporal convolutions. *arXiv preprint arXiv:1707.03141*, 2017.
- Mitchell, T. M., Thrun, S. B., et al. Explanation-based neural network learning for robot control. *Advances in neural information processing systems*, pp. 287–287, 1993.
- Monsell, S. Task switching. *Trends in cognitive sciences*, 7(3):134–140, 2003.
- Munkhdalai, T. and Yu, H. Meta networks. In Precup, D. and Teh, Y. W. (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 2554–2563, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR. URL <http://proceedings.mlr.press/v70/munkhdalai17a.html>.
- Nøklund, A. Direct feedback alignment provides learning in deep neural networks. In *Advances in Neural Information Processing Systems*, pp. 1037–1045, 2016.
- Perez, E., Strub, F., De Vries, H., Dumoulin, V., and Courville, A. Film: Visual reasoning with a general conditioning layer. *arXiv preprint arXiv:1709.07871*, 2017.
- Ravi, S. and Larochelle, H. Optimization as a model for few-shot learning. In *ICLR 2017*, 2017.
- Sakai, K. Task set and prefrontal cortex. *Annu. Rev. Neurosci.*, 31:219–245, 2008.
- Santoro, A., Bartunov, S., Botvinick, M., Wierstra, D., and Lillicrap, T. Meta-learning with memory-augmented neural networks. In *Proceedings of The 33rd International Conference on Machine Learning*, pp. 1842–1850, 2016.
- Schmidhuber, J. *Evolutionary principles in self-referential learning*. PhD thesis, Technical University of Munich, 1987.
- Schmidhuber, J. A ‘self-referential’ weight matrix. In *IN PROCEEDINGS OF THE INTERNATIONAL CONFERENCE ON ARTIFICIAL NEURAL NETWORKS*, pp. 446–451. Springer, 1993.
- Siegel, M., Buschman, T. J., and Miller, E. K. Cortical information flow during flexible sensorimotor decisions. *Science*, 348(6241):1352–1355, 2015.
- Srivastava, R. K., Masci, J., Kazerounian, S., Gomez, F., and Schmidhuber, J. Compete to compute. In *Advances in neural information processing systems*, pp. 2310–2318, 2013.
- Stokes, M. G., Kusunoki, M., Sigala, N., Nili, H., Gaffan, D., and Duncan, J. Dynamic coding for cognitive control in prefrontal cortex. *Neuron*, 78(2):364–375, 2013.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., and Fergus, R. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- Tokui, S., Oono, K., Hido, S., and Clayton, J. Chainer: a next-generation open source framework for deep learning. In *Proceedings of Workshop on Machine Learning Systems (LearningSys) in The Twenty-ninth Annual Conference on Neural Information Processing Systems (NIPS)*, 2015. URL [http://learningsys.org/papers/LearningSys\\_2015\\_paper\\_33.pdf](http://learningsys.org/papers/LearningSys_2015_paper_33.pdf).
- Vilalta, R. and Drissi, Y. A perspective view and survey of meta-learning. *Artificial Intelligence Review*, 18(2): 77–95, 2002.
- Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al. Matching networks for one shot learning. In *Advances in Neural Information Processing Systems*, pp. 3630–3638, 2016.