
A Modern Self-Referential Weight Matrix That Learns to Modify Itself

Kazuki Irie¹ Imanol Schlag¹ Róbert Csordás¹ Jürgen Schmidhuber^{1,2}

Abstract

The weight matrix (WM) of a neural network (NN) is its program. The programs of many traditional NNs are learned through gradient descent in some error function, then remain fixed. The WM of a self-referential NN, however, can keep rapidly modifying all of itself during runtime. In principle, such NNs can meta-learn to learn, and meta-meta-learn to meta-learn to learn, and so on, in the sense of recursive self-improvement. While NN architectures potentially capable of implementing such behaviour have been proposed since the '90s, there have been few if any practical studies. Here we revisit such NNs, building upon recent successes of fast weight programmers and closely related linear Transformers. We propose a scalable self-referential WM (SRWM) that learns to use outer products and the delta update rule to modify itself. We evaluate our SRWM in supervised few-shot learning and in multi-task reinforcement learning with procedurally generated game environments. Our experiments demonstrate both practical applicability and competitive performance of the proposed SRWM. Our code is public[†].

1. Introduction

The program of a neural network (NN) is its weight matrix (WM) (Schmidhuber, 1990). With prediction tasks, for example, starting from random values, an NN training procedure based on gradient descent might update the WM to minimize an error function that favors compression of given input-output observations (Solomonoff, 1964). The WM becomes permanent once training ends, and its usefulness

¹The Swiss AI Lab, IDSIA, USI & SUPSI, Lugano, Switzerland ²AI Initiative, King Abdullah University of Science and Technology (KAUST), Thuwal, Saudi Arabia. Correspondence to: <{kazuki, imanol, robert, juergen}@idsia.ch>. An earlier version was presented at NeurIPS 2021 Deep RL Workshop.

is evaluated with respect to its generalisation capability on yet unseen data.

Many environments, however, continue to evolve after training has halted (e.g., Lazaridou et al. (2021); Lin et al. (2021)), and the test setting may deviate from training in ways that exceed the NN's generalisation capability. Then human intervention might be required to re-train or fine-tune the model. Instead, more general and autonomous systems should learn to update their own programs in the light of new experience without such intervention. Especially in multi-task learning and meta-learning (learning to learn; Schmidhuber (1987)), it may be useful to learn how to keep changing and fine-tuning the model in a way that quickly adapts to various situations and new challenges (Hochreiter et al., 2001; Finn et al., 2017).

In principle, a WM could learn by itself a way of executing rapid WM adaptations in task-dependent and context-dependent fashion through a generic mechanism for recursive self-modification. Various self-modifying NNs have been proposed previously (see Sec. 6). Here we revisit the self-referential WM (Schmidhuber, 1992a; 1993a;b;c) from the '90s in the light of modern techniques for updating and generating weights. In particular, we leverage mechanisms which are now well established in the context of Fast Weight Programmers (FWPs, Schmidhuber (1991; 1992b; 1993d); reviewed in Sec. 2). FWPs have recently seen advancements in terms of performance and scalability, inspired by their formal equivalence (Schlag et al., 2021) to linear variants (Katharopoulos et al., 2020; Choromanski et al., 2021; Peng et al., 2021) of the popular Transformer (Vaswani et al., 2017).

Here we derive a new type of self-referential WM (SRWM) which naturally emerges as an extension to recent works on FWPs. We evaluate the proposed SRWM in three settings. We start by demonstrating that the proposed model is effectively capable of generating useful self-modifications by showing that the model achieves competitive performance on standard few-shot learning benchmarks. Second, by extending the few-shot learning setting to a sequential multi-task learning setting, we test the SRWM's ability to sequentially adapt itself to changes of the task at runtime. Finally, by using ProcGen (Cobbe et al., 2020), we evaluate it in a multi-task reinforcement learning (RL) setting

with procedurally generated game environments. Overall, we demonstrate both practical applicability and competitive performance of the proposed method.

2. Background on Fast Weight Programmers

Here we briefly review the essential components of fast weight programmers (FWPs; Schmidhuber (1991; 1992b; 1993d)) which our model is built upon (Sec. 3). In what follows, let t , d_{in} , d_{out} , and d_{key} denote positive integers. FWPs have a *slow* NN which can rapidly modify weights of another *fast* NN. The concept has seen a recent revival, in particular in light of its direct formal connection (Schlag et al., 2021) to linear variants (Katharopoulos et al., 2020; Choromanski et al., 2021; Peng et al., 2021; Irie et al., 2021) of the popular Transformer (Vaswani et al., 2017) when the weight generation is based on outer products between keys and values generated by the slow NN (Schmidhuber, 1991). Recent work augmented the basic FWPs (Schmidhuber, 1991; 1992b) with an improved elementary programming instruction or update rule invoked by the slow NN to reprogram the fast NN, called *delta update rule* (akin to the delta rule by Widrow & Hoff (1960)). The resulting ‘‘DeltaNet’’ (Schlag et al., 2021) is a general purpose auto-regressive NN with linear complexity w.r.t. input sequence length, which transforms the input $\mathbf{x}_t \in \mathbb{R}^{d_{\text{in}}}$ to the output $\mathbf{y}_t \in \mathbb{R}^{d_{\text{out}}}$ as follows:

$$\mathbf{k}_t, \mathbf{v}_t, \mathbf{q}_t, \beta_t = \mathbf{W}_{\text{slow}} \mathbf{x}_t \quad (1)$$

$$\bar{\mathbf{v}}_t = \mathbf{W}_{t-1} \phi(\mathbf{k}_t) \quad (2)$$

$$\mathbf{W}_t = \mathbf{W}_{t-1} + \sigma(\beta_t)(\mathbf{v}_t - \bar{\mathbf{v}}_t) \otimes \phi(\mathbf{k}_t) \quad (3)$$

$$\mathbf{y}_t = \mathbf{W}_t \phi(\mathbf{q}_t) \quad (4)$$

where \otimes denotes the outer product, σ is a sigmoid function, and ϕ is an element-wise activation function whose output elements are all positive and sum up to one (e.g. softmax). In Eq. 1, the input \mathbf{x}_t is first projected to key $\mathbf{k}_t \in \mathbb{R}^{d_{\text{key}}}$, value $\mathbf{v}_t \in \mathbb{R}^{d_{\text{out}}}$, query $\mathbf{q}_t \in \mathbb{R}^{d_{\text{key}}}$ vectors and a scalar $\beta_t \in \mathbb{R}$ using a trainable weight matrix $\mathbf{W}_{\text{slow}} \in \mathbb{R}^{(d_{\text{out}}+2*d_{\text{key}}+1) \times d_{\text{in}}}$. The generated key vector \mathbf{k}_t and a learning rate β_t (also generated by the slow NN) are used to update the *fast weight matrix* \mathbf{W}_{t-1} using the delta rule expressed in Eqs. 2-3. The fast weight matrix is typically initialized to zero i.e. $\mathbf{W}_0 = 0$. The final output \mathbf{y}_t is obtained by querying the updated fast weight matrix \mathbf{W}_t using a generated query vector \mathbf{q}_t (Eq. 4). We note that the use of ϕ function for both writing to (Eq. 3) and reading from (Eqs. 2 and 4) the fast weights is crucial for stability when the delta rule is used (we refer the readers to Schlag et al. (2021) for detailed explanations).

In practice, we use the multi-head version (Vaswani et al., 2017) of the computation above, that is, after the projection (Eq. 1), the vectors \mathbf{k}_t , \mathbf{v}_t , \mathbf{q}_t are split into equally sized H sub-vectors, and the operations in Eqs. 2-4 are conducted by H computation heads independently. Also, the layer above

is typically used as a replacement for self-attention layers in a regular Transformer architecture (Vaswani et al., 2017) while preserving other components such as feedforward blocks, layer-norm and residual connections across layers.

The slow NN or the *programmer* (here a one-layer feedforward NN; Eq. 1) with *slow weights* \mathbf{W}_{slow} learns by gradient descent to continuously modify or *program* the fast NN (here also a one-layer feedforward NN; Eq. 4) with fast weights \mathbf{W}_t as it continually receives a stream of inputs. For further extensions of this concept to more complex slow and fast NN architectures such as recurrent NNs, we refer the readers to another recent study (Irie et al., 2021).

FWPs are of great interest from the perspective of context-sensitive information processing, since the fast WM is completely context-dependent: while processing some sequence, a continually changing custom fast NN is built on the fly.

Here we leverage this mechanism to design a new kind of FWP which programs itself. It can be naturally derived from the operations described above, resulting in a ‘‘modern’’ version of the self-referential weight matrix (Schmidhuber, 1992a; 1993a;b;c) of the ‘90s.

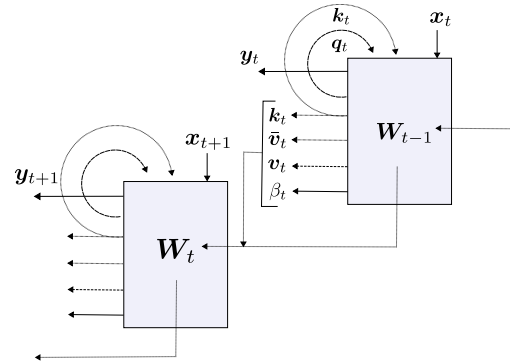


Figure 1. A ‘‘modern’’ self-referential weight matrix (SRWM).

3. A Modern Self-Referential Weight Matrix

Our ‘‘modern’’ self-referential weight matrix (SRWM) learns to train itself through self-invented key/value ‘‘training’’ patterns and learning rates, invoking sequences of elementary programming instructions based on outer products and the delta update rule, as in the recently proposed variants (Schlag et al., 2021) of FWPs (Sec. 2).

Given an input $\mathbf{x}_t \in \mathbb{R}^{d_{\text{in}}}$ at time t , our SRWM $\mathbf{W}_{t-1} \in \mathbb{R}^{(d_{\text{out}}+2*d_{\text{in}}+1) \times d_{\text{in}}}$ produces four variables $[\mathbf{y}_t, \mathbf{q}_t, \mathbf{k}_t, \beta_t]$ where $\mathbf{y}_t \in \mathbb{R}^{d_{\text{out}}}$ is the output of this layer at the current time step, $\mathbf{q}_t \in \mathbb{R}^{d_{\text{in}}}$ and $\mathbf{k}_t \in \mathbb{R}^{d_{\text{in}}}$ are query and key vectors, and $\beta_t \in \mathbb{R}$ is the self-invented learning rate to be used by the delta rule. In analogy to the terminology introduced by the original SRWM papers (Schmidhuber,

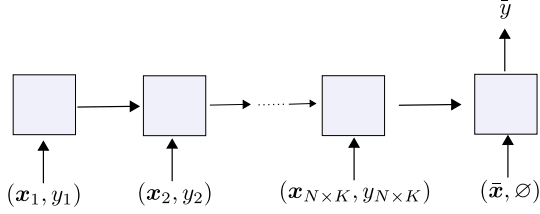


Figure 2. Synchronous-label setting for N -way K -shot learning (like in Mishra et al. (2018)). The correct label is fed together with the corresponding input for the first $N \times K$ tokens. Label prediction is only done for the $(NK+1)$ -th input which is presented without the label.

1992a; 1993a;b;c), $\mathbf{k}_t \in \mathbb{R}^{d_{in}}$ is the *modifier*-key vector, representing the key whose current value in the SRWM has to be modified, and $\mathbf{q}_t \in \mathbb{R}^{d_{in}}$ is the *analyser*-query which is again fed to the SRWM to retrieve a new “value” vector to be associated with the modifier-key.

The overall dynamics can be expressed as simply as follows:

$$\mathbf{y}_t, \mathbf{k}_t, \mathbf{q}_t, \beta_t = \mathbf{W}_{t-1} \phi(\mathbf{x}_t) \quad (5)$$

$$\bar{\mathbf{v}}_t = \mathbf{W}_{t-1} \phi(\mathbf{k}_t) \quad (6)$$

$$\mathbf{v}_t = \mathbf{W}_{t-1} \phi(\mathbf{q}_t) \quad (7)$$

$$\mathbf{W}_t = \mathbf{W}_{t-1} + \sigma(\beta_t)(\mathbf{v}_t - \bar{\mathbf{v}}_t) \otimes \phi(\mathbf{k}_t) \quad (8)$$

where the value vectors have dimensions: $\mathbf{v}_t, \bar{\mathbf{v}}_t \in \mathbb{R}^{(d_{out}+2*d_{in}+1)}$. Figure 1 illustrates the model.

Importantly, the initial values of the SRWM \mathbf{W}_0 are the only parameters in this layer which are trained by gradient descent. In practice, we extend the output dimension of the matrix from “3D+1” ($d_{out}+2*d_{in}+1$) to “3D+4” ($d_{out}+2*d_{in}+4$) to generate four different, self-invented, time-varying learning rates $\beta_t \in \mathbb{R}^4$ to be used in Eq. 8 for the four sub-matrices of $\mathbf{W}_{t-1} = [\mathbf{W}_{t-1}^y, \mathbf{W}_{t-1}^q, \mathbf{W}_{t-1}^k, \mathbf{W}_{t-1}^\beta]$ used to produce $\mathbf{y}_t, \mathbf{q}_t, \mathbf{k}_t$, and β_t in Eq. 5. For efficient computation, we also make use of multi-head computation as is done in regular Transformers (Vaswani et al., 2017; Katharopoulos et al., 2020). Please refer to Appendix A for the full description.

The SRWM described above can potentially be used to replace any regular WM. Here we mainly focus on a model which can be obtained by replacing Eqs. 1-4 in the baseline DeltaNet by the corresponding SRWM equations Eqs. 5-8. In Appendix C.1, we further show preliminary results for another type of model incorporating an SRWM, which is also based on the DeltaNet but where we replace its slow weight matrix (Eq. 1) by an SRWM (Eqs. 5-8).

4. Experiments

Can the self-referential dynamics described in Eqs. 5-8 generate useful self-modifications? The overall goal of our

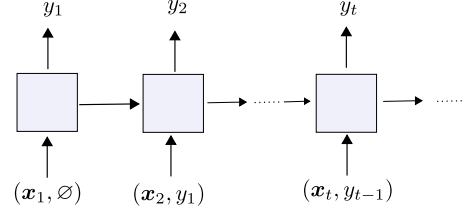


Figure 3. Delayed-label setting (like in Hochreiter et al. (2001); Santoro et al. (2016)). The correct label is fed one step after the corresponding input. Predictions take place at each step.

experiments is to evaluate the proposed SRWM on various types of tasks which require “good” self-modifications. We conduct experiments on standard supervised few-shot learning tasks (Sec. 4.1) and multi-task reinforcement learning (RL) in game environments (Sec. 4.3). In addition, we show how our SRWM can be trained to efficiently adapt itself in a sequential multi-task few-shot learning setting (Sec. 4.2). Multi-task settings are particularly relevant for evaluating SRWMs, since unlike the baseline DeltaNet (which uses the same fixed slow weights to generate fast weights for all tasks), SRWMs can potentially learn to even adapt the way it adapts itself to each task as it receives task-specific inputs.

4.1. Standard Few-Shot Learning

We start with evaluating the proposed SRWM’s capability to generate useful self-modifications on standard supervised few-shot image classification tasks. We conduct experiments on the classic Omniglot (Lake et al., 2015) and Mini-ImageNet (Vinyals et al., 2016; Ravi & Larochelle, 2017) datasets. For further details on the datasets, we refer to the respective references and Appendix B where we also provide extra experimental results on the Fewshot-CIFAR100 (Oreshkin et al., 2018) dataset.

Let N , K and C denote positive integers. The task of few-shot image classification, or N -way K -shot image classification based on a dataset containing C classes, is structured through so-called *episodes*. In each episode, N distinct classes are randomly drawn from C . The resulting N classes are re-labelled such that each class is assigned to one out of N distinct random label indices. For each of these N classes, K examples are randomly sampled. The set of resulting $N \times K$ labelled images is called the *support set*. The goal of the task is to predict the label of another image (a *query image* that is not in the support set) which is sampled from one of the N classes, based on the information available in the support set.

While there are several ways of approaching this problem, we are evaluating our SRWM in the *sequential learning* approach (Santoro et al., 2016; Hochreiter et al., 2001). That is, the image/label pairs of the support set are randomly ordered

Table 1. Single task, 5-way, few-shot classification test accuracies (%) on Omniglot and Mini-ImageNet. The bottom section shows the generic sequence models, while the top part shows other related approaches. The numbers marked by * are taken from the corresponding papers. Following the standard convention (Ravi & Larochelle, 2017), we report 95% confidence interval computed over five sets of test episodes. For further details, we refer to Appendix B.

	Omniglot	Mini-ImageNet		
	1-shot	Backend	1-shot	5-shot
MAML* (Finn et al., 2017)	98.7 ± 0.4	Conv-4-32	48.7 ± 1.8	63.1 ± 0.9
fwCNN-Hebb* (Munkhdalai & Trischler, 2018)	99.4 ± 0.1	Conv-5-64	50.2 ± 0.4	64.8 ± 0.5
HyperTransformer* (Zhmoginov et al., 2022)	96.2	Conv-4-32	53.8	67.1
SNAIL* (Mishra et al., 2018)	99.1 ± 0.2	Conv-4-32	45.1	55.2
Delta Net	97.2 ± 0.0	Conv-4-32	47.0 ± 0.1	62.7 ± 0.1
SRWM	97.4 ± 0.0	Conv-4-32	47.0 ± 0.2	61.4 ± 0.1

to form a sequence which is read by a sequence-processing NN (e.g., a recurrent NN). By encoding the support set information into its internal state, the corresponding NN predicts the label of the query image. In the case of our SRWM, the model generates updates of its own weights as it reads the sequence of support set items. The generated weights are used to compute the final prediction for the query image. To fully specify this approach, we also need to explain how input image/label pairs are fed to the model. Here we follow the approach used by Mishra et al. (2018) which we refer to as the *synchronous-label setting* illustrated in Figure 2. This strategy is specifically designed for N -way K -shot learning, which consists in feeding the input and its label to the model at the same time for $N \times K$ items in the support set. The model only predicts the label of the $(N \times K + 1)$ -th input which is the query image presented without label. An alternative approach (Santoro et al., 2016) which we call the *delayed label setting* (Figure 3) is used later in Sec. 4.2.

Mishra et al. (2018)’s SNAIL model serves as a baseline in our experiments as it is a generic Transformer-like sequence processing model where the regular feedforward block is replaced by 1D convolution (Waibel et al., 1989). Other related approaches include the fwCNN-Hebb model by Munkhdalai & Trischler (2018) and the recently proposed HyperTransformer (Zhmoginov et al., 2022), as well as the standard MAML (Finn et al., 2017; Finn & Levine, 2018). Both fwCNN-Hebb and HyperTransformer generate fast weights of an NN based on the support set information. The fwCNN-Hebb is an outer product based fast weight programmer (similar to ours but using the purely additive update rule like in the standard Linear Transformers) in which the fast net is a linear layer before the final softmax layer. A crucial difference to our generic FWPs is that in the fwCNN-Hebb, key vectors are generated from the images while the values are generated from the labels, which is a specific design choice reflecting the task of few-shot learning. In the HyperTransformers, the fast net is a

convolutional NN whose weights are parameterised by the Transformer encoder architecture taking images and their labels as inputs. We note that these two methods are specifically designed for few-shot learning, unlike our models and SNAIL which are generic sequence processing NNs directly applicable beyond few-shot learning e.g., to RL as in Sec. 4.3. MAML optimises the initial weights for future fine-tuning via gradient descent. In the SRWM, the initial weights also learn and encode its own self-modification algorithm. All models presented here use the same vision backend: the 4-layer convolutional net (Vinyals et al., 2016) with 32 channels for Mini-ImageNet (Conv-4-32), and 64 channels for Omniglot, except the fwCNN-Hebb which uses 5 layers and 64 channels (Conv-5-64) in both cases. Additional details of the model architecture are explained in Appendix B.

The results are summarised in Table 1. Overall, the proposed SRWM performs well among the generic approaches. Comparing the SRWM to the SNAIL baseline, the SRWM achieves very competitive performance on Mini-ImageNet¹. DeltaNet and SRWM tend to have similar performance. This is a satisfactory result, as it shows that a single self-modifying WM (instead of separate slow and fast nets) remains competitive in this single task scenario.

While we find the HyperTransformer to outperform all other models considered here, our performance is respectable compared to that of MAML without requiring bi-level optimisation, and fwCNN-Hebb without inductive bias on the key/value generation from image/label, respectively. Although the SRWM is a very generic approach, its overall performance is thus competitive, demonstrating the effectiveness of the proposed self-referential dynamics (the main goal of this experiment).

¹On Omniglot, we did not manage to reach > 99.0% performance of SNAIL. We might be missing some technical details/tricks but the results are nonetheless respectable.

Table 2. Total and instance-level accuracies (%) for **sequential multi-task few-shot learning** experiments (Sec. 4.2). Regarding the instance-level accuracy, column $k \in \{1, 2, 3, 5, 10\}$ shows the percentage of correctly predicted k -th instances from each class. This is for the test time scenario where the model is first tasked to learn to predict Omniglot and then Mini-ImageNet.

Task	Model	1	2	3	5	10	Total
Omniglot	DeltaNet	39.1	91.2	93.9	95.8	96.8	92.2
	SRWM	40.6	92.1	94.8	96.3	96.7	92.3
Mini-ImageNet	DeltaNet	20.8	43.5	48.2	51.3	54.8	50.4
	SRWM	20.9	46.0	50.3	54.2	58.0	53.3

4.2. Sequential Multi-Task Adaptation

The basic few-shot learning experiments presented in the previous section demonstrate that the very generic SRWM can effectively generate useful weight updates, achieving competitive performance on standard benchmarks. Now we are interested in testing its self-referential dynamics on a task which requires adaptation to environmental changes at runtime. We introduce two modifications to the few-shot learning setting above. First, instead of specifically training the model for N -way K -shot classification using the synchronous-label setting (Figure 2), we train our model in the delayed-label setting (Hochreiter et al., 2001) as illustrated in Figure 3. Here the model makes prediction at each time step by receiving an input image to be classified and the correct label of the previous input (the label feeds are thus shifted/delayed by one time step). This setting is convenient for evaluating the model on a continuous stream of predictions/solutions². Second, the sequence of images to be predicted is constructed by concatenating two image sequences sampled from two different datasets: Omniglot and Mini-ImageNet. The model first receives a stream of images from one of the datasets; at some point, the dataset suddenly changes, to simulate a change of environment. The model has to learn to adapt itself to this shift without human intervention, in a continual execution of its program.

Note that our goal is to construct a task which requires adaptation to sudden changes during the model’s runtime, which is different from *continual few-shot learning* (e.g. Yap et al. (2021)) whose goal is to successively meta-train on multiple few-shot learning tasks.

We conduct experiments in a 5-way classification setting with concatenation of Omniglot and Mini-ImageNet segments containing up to 15 examples per class in each segment. The concatenation order is alternated for each batch,

²The synchronous-label setting could also be modified to support such usage, but that would require to feed the same input image twice, once without the correct label (for prediction) and once with it (for learning). While such a setting may be of interest as a way of clearly separating the learning mode from the prediction mode, it is out of the scope of this work.

and training segment lengths are randomised by trimming. Regardless of model type, we find that training models in the delayed-label setting is more difficult than in the synchronous-label setting. We observe that in many configurations, the model gets stuck at a sub-optimal behaviour where it learns to improve its class-averaged zero-shot accuracy (apparently by learning to output one of the unused labels for a new class appearing in the sequence for the first time), but fails to learn to properly learn from the feedback at each step. The most crucial hyper-parameter we identified was the large enough batch size. Additional details on practicalities are shared in Appendix B.5.

In the end, we successfully trained both the DeltaNet baseline and the SRWM on this sequential adaptation task. Figure 4 shows the evolution of the SRWM’s test time accuracy as it gets more inputs. In this test setting, the model starts with receiving a stream of samples from Omniglot. At step 74, the task changes; the model now has to classify images sampled from Mini-ImageNet. The accuracy obviously drops due to this change, since the model can not know which class the new datapoint belongs to, but it effectively adapts itself and starts learning the second task. Table 2 compares the DeltaNet to the SRWM. While their performance is similar in the single task scenario (Table 1), the SRWM achieves a higher accuracy in this multi-task setting, demonstrating its rapid adaptation capability. In Appendix B.5, we also provide results for the test scenario with the reversed testing order, i.e., Mini-ImageNet then Omniglot.

4.3. Multi-Task Reinforcement Learning (RL)

We finally evaluate the proposed model in a multi-task RL setting using procedurally generated game environments of ProcGen (Cobbe et al., 2020). The corresponding setting is illustrated in Figure 5. Unlike other popular game environments such as Atari (Bellemare et al., 2016), ProcGen provides various procedurally different levels. This allows for creating clean train/test splits. In addition, working with diverse levels is especially relevant to our setting, as we wish to build models which are adaptive to changes across game types, as well as diversity within the same game.

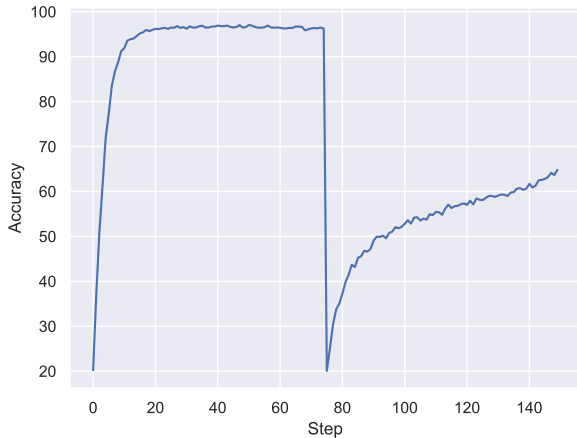


Figure 4. Average test accuracy (%) of the SRWM-based model as a function of the number of steps in the sequential multi-task adaptation setting (Sec. 4.2). A stream of datapoints is fed to the model in delayed-label fashion (Figure 3). The datapoints are sampled from **Omniglot** until step 74 (where the accuracy drops), then from **Mini-ImageNet**.

General Settings. In our main experiment, we jointly train on 6 environments, namely, *Bigfish*, *Fruitbot*, *Maze*, *Leaper*, *Plunder*, and *Starpilot* in the **easy** distribution. We conduct distributed training using the standard IMPALA (Espeholt et al., 2018) architecture implemented in Torchbeast (Küttler et al., 2019). We use 48 actors (i.e. 8 actors per environment). All our models use the common *large* architecture of Espeholt et al. (2018) which consists of a 15-layer residual convolutional vision model. They differ from each other by the “memory” module inserted between the vision stem and the output layer. In addition to our SRWM model, we train the baseline IMPALA feed-forward and LSTM (Gers et al., 2000; Hochreiter & Schmidhuber, 1997) models, as well as two additional baselines: the DeltaNet (Schlag et al., 2021) and a “Fake SR” model which is the SRWM model without the self-modification mechanism. (i.e. we only keep the “y”-part in Eq. 5). We set a backpropagation span of 50 steps to train self-modification³ as well as LSTM and DeltaNet baselines. In all cases, the memory states (including the SRWM weight changes) are only reset at episode boundaries for all stateful models (LSTM, DeltaNet, and SRWM). The LSTM model has one layer with 256 nodes as in the IMPALA baseline. Both DeltaNet and SRWM have two layers with a hidden size of 128, following the setting used by previous work (Irie et al., 2021) training the DeltaNet on Atari.

These 6 environments are known for not explicitly requir-

³Truncated backpropagation through time (Williams & Peng, 1990) can be used to train the SRWM thanks to the additive nature of Eq. 8.

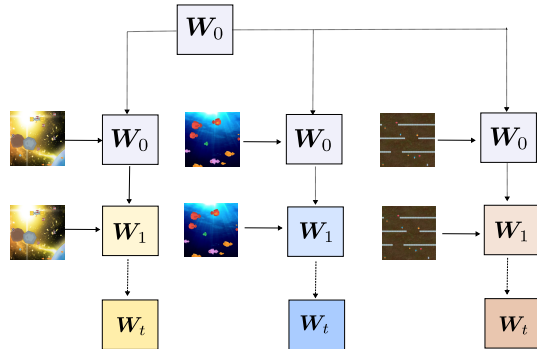


Figure 5. Illustration for multi-task RL. The initial weight matrix W_0 is common to all tasks and episodes. The effective weight matrix is a function of task/episode specific input streams.

ing “memory” to perform the task (Cobbe et al. (2020)); we also confirm this trend by comparison to our baseline feed-forward and LSTM RNN models). In principle, this allows for evaluating the effect of self-modifications in isolation (although it is difficult to completely dissociate self-modification from the concept of “memory”). We jointly train on 6 environments in the easy distribution for a total of 300 M steps (ca. 50 M per environment). In Appendix C.1, we also present an extra experiment using 4 environments from the *memory* distribution (*Dodgeball*, *Heist*, *Maze*, *Miner*) to evaluate our models also in partially observable settings, confirming the effectiveness of SRWM.

Train/Test split. Following Cobbe et al. (2020), we use 200 levels (level ID 0 to 199) to train in the easy distribution. For evaluation, instead of randomly sampling the test levels as is commonly done for ProcGen, we consistently use the same set of 3 distinct test splits for all models. Each of our test splits contains 200 levels, respectively including levels 1000 to 1199, 1200 to 1399, and 1400 to 1599. We use 3 test splits and report an average score to take into account the performance variability across the choice of test levels. The performance on the training set is computed using all 200 training levels. We train each model three times, reporting training performance averaged over three runs, and test performance over 9 data points (3 test splits for 3 training runs).

Overall Performance. Table 3 shows the aggregated normalised scores. First of all, comparing the LSTM and feed-forward baselines, we confirm Cobbe et al. (2020)’s finding that the LSTM layer does not provide any improvements regarding the test performance (while some improvements are obtained on the train set). The two fast weight models (which can adapt to each task based on the task specific inputs), DeltaNet and SRWM, clearly outperform the feed-forward and LSTM baselines, as well as the Fake SR model,

Table 3. ProcGen normalised aggregated scores (multiplied by 100) over **6 environments** (Bigfish, Fruitbot, Maze, Leaper, Plunder, and Starpilot) in the **easy distribution**. The models are trained in a multi-task setting. Normalisation constants are taken from the original ProcGen paper (Cobbe et al., 2020). Results are computed from 3 independent training runs for 300 M steps in the easy distribution. The test scores are averaged over 3 distinct sets of 200 fixed test levels (i.e., the mean/std computed from 9 data points). For further details, see tables in Appendix C.2 where we provide scores obtained for each game. Results on memory distribution can be found in Appendix C.1. The number of trainable parameters are 626 K for the feedforward baseline (FF), 959 K for Fake SR, 1.2 M for LSTM, 1.05 M for DeltaNet and 968 K for SRWM.

	FF	LSTM	Fake SR	DeltaNet	SRWM
Train	22.5 (2.6)	28.3 (1.4)	27.0 (1.8)	35.0 (1.6)	34.6 (1.8)
Test	16.4 (1.6)	15.7 (1.6)	15.3 (1.9)	18.6 (1.7)	20.0 (1.8)

which is the SRWM without self-modification. The SRWM achieves a slightly better test score than the DeltaNet. Overall, the proposed very generic SRWM based model achieves very competitive performance.

Comparison to Expert Models. We observe that the performance gains achieved by the SRWM over the baselines are particularly large for two of the environments, *Bigfish* and *Starpilot*. Here we study these two cases in isolation. In Table 4, we compare multi-task agents presented above to expert agents trained specifically on one environment for 50 M steps. On *Starpilot*, we observe that the self-modification mechanism yields improvements even in the single task case. The case of *Bigfish* is more interesting: the performances of the agents with and without self-modification capability are close in the expert training case. However, the self-modifying agent achieves a much better score in the multi-task setting, where the baseline agent’s performance drops by a large margin. This indicates the usefulness of the SRWM’s ability to adapt itself to each environment in the multi-task scenario.

Table 4. Comparison between multi-task vs expert agent performance. Raw scores obtained in the easy distribution of ProcGen.

Env	Split	Weight Update			
		No (Fake SR)		Yes (SRWM)	
		Multi-6	Expert	Multi-6	Expert
Bigfish	Train	11.6 (5.7)	28.9 (0.9)	20.1 (2.4)	28.5 (1.2)
	Test	4.7 (2.4)	15.8 (1.7)	9.0 (2.0)	14.2 (2.0)
Starpilot	Train	55.0 (1.3)	59.8 (0.7)	61.3 (2.0)	64.0 (1.9)
	Test	49.6 (2.1)	52.9 (1.2)	54.6 (2.4)	57.3 (1.6)

Ablation on State Reset. The SRWM models presented above are trained by carrying over the weight modifications across entire episodes whose lengths are variable—often episodes are getting longer during training as the agent becomes better at the task. We were initially uncertain about the empirical stability of the dynamics described by Eqs. 5-8 in such a scenario. As an ablation study, we trained and evaluated an SRWM agent by resetting the weight update after every fixed time span (whose length was the backpropagation span). Such models failed to leverage the SRWM mechanism, obtaining scores of 28.5 (1.2) and 16.1 (2.2) on the train and test splits respectively, similar to those of the baseline without self-modification (Table 3).

5. Discussion

Here we discuss interesting aspects (ignored in the previous sections) of the proposed SRWM and the experimental settings.

Reducing the ratio between learning complexity and number of time-varying variables. Our SRWM inherits a remarkable property of the 1993 system (Schmidhuber, 1993d): it has many more temporal variables under massively parallel end-to-end differentiable control than what’s possible in standard RNNs of the same size: $O(d^2)$ instead of $O(d)$, where d is the number of hidden units.

Interpretability. In general, interpreting NNs is not straightforward. The values of β_t of Eq. 8 in the RL setting intuitively define the strength of the weight modifications. We observe that the values of all four components of β_t vary between 0.50 and 0.65 depending on the input, instead of covering the full range of sigmoid values between 0 and 1. We find it difficult to derive any further interpretation beyond these statistics. Note, however, that β_t by itself does not fully describe the self-modification effects of Eq. 8, which also depend on the actual values of key and query.

Implementation/Limitation. Similar to recent works on fast weight programmers (Schlag et al., 2021; Katharopoulos et al., 2020; Irie et al., 2021), our SRWM is implemented as a custom CUDA kernel. While this approach yields competitive computation time and memory-efficient custom backpropagation, its flexibility is limited. For instance, to replace all weight matrices in an RL agent’s vision module or in the feature extractor for few-shot learning, a custom implementation for convolution would be required, although in principle the SRWM above could replace any regular weight matrix. In contrast, vision models that are fully based on linear layers such as the MLP-Mixer (Tolstikhin et al., 2021) can be straightforwardly combined with the SRWM. Regarding speed, the feedforward and LSTM baselines process about 3,500 steps per second, while DeltaNet

and SRWM do 2,300 and 1,700 steps per second respectively on a single P100 GPU in the RL experiments which require slow state copying due to separate interaction and training modes. In supervised few-shot learning settings, the speeds of LSTM, DeltaNet and SRWM are comparable. With a batch size of 128, they process about 8,000 images per second, using the same Conv-4 backend on 1-shot Omniglot on a single P100 GPU.

Limitation of this paper’s scope. The main tasks in our experiments are limited to those solvable by feed-forward NNs: image classification and RL in fully observable environments. Appendix C.1 presents promising preliminary experiments on multi-task RL in partially observable environments using the memory distribution of ProcGen. We found that augmenting the DeltaNet (which already has a short-term memory) with an SRWM by replacing the slow weight matrix in the DeltaNet by an SRWM yields performance improvements. Further experiments are needed to test the SRWM on tasks which themselves require sequence processing, while involving adaptation to different types or domains of sequences, e.g., automated domain adaptation in language modelling (Irie et al., 2018; Lazaridou et al., 2021).

Other perspectives. While we motivated our work purely from the perspective of context-adaptive self-modifying NNs, never-ending self-reconfiguration could also be motivated by analogy to “livewired” synaptic connections in biological neural networks (Eagleman, 2020).

6. Related Work

Original Self-Referential Weight Matrix. The *original* SRWM was proposed in the ’90s as a framework for self-improving recurrent NNs (Schmidhuber, 1992a; 1993a;b;c). Such an RNN has special input and output units to directly address, read, and modify any of its own current weights through an index for each weight of its weight matrix (i.e., for a weight matrix with an input/output dimension N , the weight index ranges from 0 to $N^2 - 1$ which is encoded as a binary vector). In contrast, our self-modification is based on *key/value associations*, i.e., to encode a WM modification, our NN generates a key vector, value vectors, and a temporary learning rate which allows for the rapid modification of an entire rank at a time (Schmidhuber, 1991; 1993d). This design is reinforced by the recent success of linear Transformers and fast weight programmers (Katharopoulos et al., 2020; Schlag et al., 2021; Irie et al., 2021). In this sense, our SRWM is a *modern* approach to self-modification, even if the use of outer products to parameterise fast weight generation itself is not (e.g. Schmidhuber (1991; 1993d)).

Other Self-Modifying Neural Networks. There are also more recent works on self-modifying NNs. Neuromodulated

plasticity is a Hebbian-style self-modification (Miconi et al., 2018; 2019; Schmidgall, 2020; Najarro & Risi, 2020) which also makes use of outer products to generate a modulation term which is added to the base weights. The corresponding computations can also be interpreted as key/value/query association operations. However, the key, query, and value patterns are hard-coded to be one of the input/output pairs of the corresponding layer at each time step. While this circumvents the necessity to allocate parameters for generating those vectors, it is known that the resulting program can be expressed in terms of an unnormalised attention (Schmidhuber, 1993d) over the past outputs (Ba et al., 2016). In contrast, in our model, all these patterns are arbitrary as they are generated from learned transformations whose parameters are themselves self-modifying.

Hierarchical Fast Weight Programmers. As reviewed in Sec. 2, an FWP is an NN which learns to generate, update, and maintain weights of another NN. However, a typical FWP has a slow NN with a weight matrix that remains fixed after training. Previous work (Irie et al., 2021) has proposed to go one step further by parameterising the slow weights in the Delta Net with another FWP to obtain the DeltaDelta Net. However, such a hierarchy has no end, as the highest level programmer would still have a fixed weight matrix. In this work, we follow the spirit of early work (Schmidhuber, 1992a; 1993a;b;c;d) and collapse these potentially hierarchical meta-levels into one single self-referential weight matrix.

Fixed Weight Meta-RNNs. Learning learning dynamics using a fixed-weight NN (typically an RNN) has become a common approach (Hochreiter et al., 2001; Cotter & Conwell, 1990; 1991; Santoro et al., 2016; Wang et al., 2017; Duan et al., 2016; Munkhdalai & Yu, 2017; Munkhdalai et al., 2019; Kirsch & Schmidhuber, 2021; Sandler et al., 2021). A truly self-referential weight matrix, however, would allow for modifying *all* of its own components. The only thing that’s trained by gradient descent are the SRWM’s *initial* weights at the beginning—all of them, however, may rapidly change during sequence processing, in a way that’s driven by the SRWM itself.

Dynamic Evaluation. Auto-regressive language modelling is an interesting task in the context of sequence processing with error feedback. In the standard testing scenario, language models receive a label feedback after each prediction at each time step. Such feedback may not only act as error feedback and influence NN activations like in fixed weight meta-RNNs: in the common method of dynamic evaluation (Mikolov et al., 2010; Krause et al., 2018), it is also used to adapt the weights at test time using the fixed gradient descent algorithm. A self-referential approach could also adapt the corresponding adaptation algorithm.

Recursive Self-Improvements. Beyond the scope of NNs, recursive self-modification (Good, 1965; Eden et al., 2013) is of general interest when considering autonomous, self-improving machines (Schmidhuber, 1987; 2006; Wang, 2007; Nivel & Thórisson, 2009; Steunebrink et al., 2016; Wang et al., 2018). While we proposed a WM that can recursively modify itself, our first set of experiments is limited to studying its practical ability on well known supervised and reinforcement learning tasks. In future work, we intend to define additional tasks for specifically measuring the ability to self-improve. This work may also have to consider standard limitations of NNs such as their lack of systematic generalisation (Fodor et al., 1988), including length generalisation (e.g., see recent work on Transformers (Csordás et al., 2022)), in the context of self-improving NNs. For example, in the long run, there is no guarantee that the SRWM’s self-modifications preserve the original objective function it is trained for (Hubinger et al., 2019).

7. Conclusion

We proposed a new type of self-referential weight matrix (SRWM) with a modern mechanism for self-modification. Our self-modifying neural networks (NNs) learn to generate patterns of keys and values and learning rates, translating these patterns into rapid changes of their own weight matrix through sequential outer products and invocations of the delta update rule. In a set of three experiments, we demonstrated that our generic SRWM is practical and performs well in both supervised few-shot learning and multi-task reinforcement learning settings, using procedurally generated game environments. Our promising results encourage further investigations of self-improving NNs.

Acknowledgements

We would like to thank Karl Cobbe for answering some practical questions about ProcGen. Kazuki Irie wishes to thank Anand Gopalakrishnan for letting him know about ProcGen. This research was partially funded by ERC Advanced grant no: 742870, project AlgoRNN, and by Swiss National Science Foundation grant no: 200021_192356, project NEUSYM. We are thankful for hardware donations from NVIDIA & IBM. The resources used for the project were partially provided by Swiss National Supercomputing Centre (CSCS) project d115.

References

Ba, J., Hinton, G. E., Mnih, V., Leibo, J. Z., and Ionescu, C. Using fast weights to attend to the recent past. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, pp. 4331–4339, Barcelona, Spain, December 2016.

Bellemare, M. G., Ostrovski, G., Guez, A., Thomas, P. S., and Munos, R. Increasing the action gap: New operators for reinforcement learning. In *Proc. AAAI Conf. on Artificial Intelligence*, pp. 1476–1483, Phoenix, AZ, USA, February 2016. AAAI Press.

Chen, Y., Liu, Z., Xu, H., Darrell, T., and Wang, X. Meta-baseline: exploring simple meta-learning for few-shot learning. In *Proc. IEEE Int. Conf. on Computer Vision (ICCV)*, pp. 9062–9071, Virtual only, March 2021.

Choromanski, K., Likhoshesterov, V., Dohan, D., Song, X., Gane, A., Sarlos, T., Hawkins, P., Davis, J., Mohiuddin, A., Kaiser, L., et al. Rethinking attention with performers. In *Int. Conf. on Learning Representations (ICLR)*, Virtual only, 2021.

Cobbe, K., Hesse, C., Hilton, J., and Schulman, J. Leveraging procedural generation to benchmark reinforcement learning. In *Proc. Int. Conf. on Machine Learning (ICML)*, pp. 2048–2056, Virtual only, July 2020.

Cotter, N. E. and Conwell, P. R. Fixed-weight networks can learn. In *Proc. Int. Joint Conf. on Neural Networks (IJCNN)*, pp. 553–559, San Diego, CA, USA, June 1990.

Cotter, N. E. and Conwell, P. R. Learning algorithms and fixed dynamics. In *Proc. Int. Joint Conf. on Neural Networks (IJCNN)*, pp. 799–801, Seattle, WA, USA, July 1991.

Csordás, R., Irie, K., and Schmidhuber, J. The neural data router: Adaptive control flow in transformers improves systematic generalization. In *Int. Conf. on Learning Representations (ICLR)*, Virtual only, April 2022.

Deleu, T., Würfl, T., Samiei, M., Cohen, J. P., and Bengio, Y. Torchmeta: A meta-learning library for PyTorch. *Preprint arXiv:1909.06576*, 2019.

Duan, Y., Schulman, J., Chen, X., Bartlett, P. L., Sutskever, I., and Abbeel, P. RL²: Fast reinforcement learning via slow reinforcement learning. *Preprint arXiv:1611.02779*, 2016.

Eagleman, D. *Livewired: The inside story of the ever-changing brain*. 2020.

Eden, A. H., Moor, J. H., Soraker, J. H., and Steinhart, E. *Singularity Hypotheses: A Scientific and Philosophical Assessment*. Springer, 2013.

Espeholt, L., Soyer, H., Munos, R., Simonyan, K., Mnih, V., Ward, T., Doron, Y., Firoiu, V., Harley, T., Dunning, I., Legg, S., and Kavukcuoglu, K. IMPALA: scalable distributed deep-RL with importance weighted actor-learner architectures. In *Proc. Int. Conf. on Machine Learning (ICML)*, pp. 1406–1415, Stockholm, Sweden, July 2018.

- Finn, C. and Levine, S. Meta-learning and universality: Deep representations and gradient descent can approximate any learning algorithm. In *Int. Conf. on Learning Representations (ICLR)*, Vancouver, Canada, April 2018.
- Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proc. Int. Conf. on Machine Learning (ICML)*, pp. 1126–1135, Sydney, Australia, August 2017.
- Fodor, J. A., Pylyshyn, Z. W., et al. Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28 (1-2):3–71, 1988.
- Gers, F. A., Schmidhuber, J., and Cummins, F. Learning to forget: Continual prediction with LSTM. *Neural computation*, 12(10):2451–2471, 2000.
- Good, I. Speculations concerning the first ultraintelligent machine. *Advances in Computers*, pp. 31–88, 1965.
- Hanson, S. J. A stochastic version of the delta rule. *Physica D: Nonlinear Phenomena*, 42(1-3):265–272, 1990.
- 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 *Proc. Int. Conf. on Artificial Neural Networks (ICANN)*, volume 2130, pp. 87–94, Vienna, Austria, August 2001.
- Hubinger, E., van Merwijk, C., Mikulik, V., Skalse, J., and Garrabrant, S. Risks from learned optimization in advanced machine learning systems. *Preprint arXiv:1906.01820*, 2019.
- Irie, K., Kumar, S., Nirschl, M., and Liao, H. RADMM: Recurrent adaptive mixture model with applications to domain robust language modeling. In *Proc. IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6079–6083, Calgary, Canada, April 2018.
- Irie, K., Schlag, I., Csordás, R., and Schmidhuber, J. Going beyond linear transformers with recurrent fast weight programmers. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, Virtual only, 2021.
- Katharopoulos, A., Vyas, A., Pappas, N., and Fleuret, F. Transformers are RNNs: Fast autoregressive transformers with linear attention. In *Proc. Int. Conf. on Machine Learning (ICML)*, Virtual only, July 2020.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *Preprint arXiv:1412.6980*, 2014.
- Kirsch, L. and Schmidhuber, J. Meta-learning backpropagation and improving it. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, Virtual only, 2021.
- Krause, B., Kahembwe, E., Murray, I., and Renals, S. Dynamic evaluation of neural sequence models. In *Proc. Int. Conf. on Machine Learning (ICML)*, pp. 2771–2780, Stockholm, Sweden, July 2018.
- Krizhevsky, A. Learning multiple layers of features from tiny images. Master’s thesis, Computer Science Department, University of Toronto, 2009.
- Küttler, H., Nardelli, N., Lavril, T., Selvatici, M., Sivakumar, V., Rocktäschel, T., and Grefenstette, E. Torchbeast: A PyTorch platform for distributed RL. *Preprint arXiv:1910.03552*, 2019.
- Lake, B. M., Salakhutdinov, R., and Tenenbaum, J. B. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.
- Lazaridou, A., Kuncoro, A., Gribovskaya, E., Agrawal, D., Liska, A., Terzi, T., Gimenez, M., d’Autume, C. d. M., Ruder, S., Yogatama, D., et al. Pitfalls of static language modelling. *Preprint arXiv:2102.01951*, 2021.
- Lin, Z., Shi, J., Pathak, D., and Ramanan, D. The CLEAR Benchmark: Continual LEARNING on Real-World Imagery. In *Conference on Neural Information Processing Systems (NeurIPS), Track on Datasets and Benchmarks*, Virtual only, December 2021.
- Miconi, T., Stanley, K., and Clune, J. Differentiable plasticity: training plastic neural networks with backpropagation. In *Proc. Int. Conf. on Machine Learning (ICML)*, pp. 3559–3568, Stockholm, Sweden, July 2018.
- Miconi, T., Rawal, A., Clune, J., and Stanley, K. O. Backpropamine: training self-modifying neural networks with differentiable neuromodulated plasticity. In *Int. Conf. on Learning Representations (ICLR)*, New Orleans, LA, USA, May 2019.
- Mikolov, T., Karafiát, M., Burget, L., Cernocký, J., and Khudanpur, S. Recurrent neural network based language model. In *Proc. Interspeech*, pp. 1045–1048, Makuhari, Japan, September 2010.
- Mishra, N., Rohaninejad, M., Chen, X., and Abbeel, P. A simple neural attentive meta-learner. In *Int. Conf. on Learning Representations (ICLR)*, Vancouver, Canada, 2018.
- Munkhdalai, T. and Trischler, A. Metalearning with hebbian fast weights. *Preprint arXiv:1807.05076*, 2018.
- Munkhdalai, T. and Yu, H. Meta networks. In *Proc. Int. Conf. on Machine Learning (ICML)*, pp. 2554–2563, Sydney, Australia, August 2017.

- Munkhdalai, T., Sordoni, A., Wang, T., and Trischler, A. Metalearned neural memory. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, pp. 13310–13321, Vancouver, Canada, December 2019.
- Najarro, E. and Risi, S. Meta-learning through hebbian plasticity in random networks. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, Virtual only, December 2020.
- Nivel, E. and Thórisson, K. R. Self-programming: Operationalizing autonomy. In *Proc. Conf. on Artificial General Intelligence (AGI)*, pp. 150–155, Arlington, VA, USA, March 2009.
- Oreshkin, B. N., López, P. R., and Lacoste, A. TADAM: task dependent adaptive metric for improved few-shot learning. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, pp. 719–729, Montréal, Canada, December 2018.
- Peng, H., Pappas, N., Yogatama, D., Schwartz, R., Smith, N. A., and Kong, L. Random feature attention. In *Int. Conf. on Learning Representations (ICLR)*, Virtual only, 2021.
- Ravi, S. and Larochelle, H. Optimization as a model for few-shot learning. In *Int. Conf. on Learning Representations (ICLR)*, Toulon, France, April 2017.
- Sandler, M., Vladymyrov, M., Zhmoginov, A., Miller, N., Madams, T., Jackson, A., and y Arcas, B. A. Meta-learning bidirectional update rules. In *Proc. Int. Conf. on Machine Learning (ICML)*, pp. 9288–9300, Virtual only, July 2021.
- Santoro, A., Bartunov, S., Botvinick, M., Wierstra, D., and Lillicrap, T. P. Meta-learning with memory-augmented neural networks. In *Proc. Int. Conf. on Machine Learning (ICML)*, pp. 1842–1850, New York City, NY, USA, June 2016.
- Schlag, I., Irie, K., and Schmidhuber, J. Linear Transformers are secretly fast weight programmers. In *Proc. Int. Conf. on Machine Learning (ICML)*, Virtual only, July 2021.
- Schmidgall, S. Adaptive reinforcement learning through evolving self-modifying neural networks. In *Proc. Genetic and Evolutionary Computation Conference (GECCO), Companion Volume*, pp. 89–90, Cancún, Mexico, July 2020.
- Schmidhuber, J. Evolutionary principles in self-referential learning, or on learning how to learn: the meta-meta-... hook. Institut für Informatik, Technische Universität München, 1987. <http://www.idsia.ch/~juergen/diploma.html>.
- Schmidhuber, J. Making the world differentiable: On using fully recurrent self-supervised neural networks for dynamic reinforcement learning and planning in non-stationary environments. Technical Report FKI-126-90, [http://people.idsia.ch/~juergen/FKI-126-90_\(revised\)bw_ocr.pdf](http://people.idsia.ch/~juergen/FKI-126-90_(revised)bw_ocr.pdf), Tech. Univ. Munich, 1990.
- Schmidhuber, J. Learning to control fast-weight memories: An alternative to recurrent nets. Technical Report FKI-147-91, Institut für Informatik, Technische Universität München, March 1991.
- Schmidhuber, J. Steps towards “self-referential” learning. Technical Report CU-CS-627-92, Dept. of Comp. Sci., University of Colorado at Boulder, November 1992a.
- Schmidhuber, J. Learning to control fast-weight memories: An alternative to dynamic recurrent networks. *Neural Computation*, 4(1):131–139, 1992b.
- Schmidhuber, J. An introspective network that can learn to run its own weight change algorithm. In *Proc. IEE Int. Conf. on Artificial Neural Networks*, pp. 191–195, Brighton, UK, May 1993a.
- Schmidhuber, J. A self-referential weight matrix. In *Proc. Int. Conf. on Artificial Neural Networks (ICANN)*, pp. 446–451, Amsterdam, Netherlands, September 1993b.
- Schmidhuber, J. A neural network that embeds its own meta-levels. In *Proc. IEEE Int. Conf. on Neural Networks (ICNN)*, San Francisco, CA, USA, March 1993c.
- Schmidhuber, J. Reducing the ratio between learning complexity and number of time varying variables in fully recurrent nets. In *International Conference on Artificial Neural Networks (ICANN)*, pp. 460–463, Amsterdam, Netherlands, September 1993d.
- Schmidhuber, J. Gödel machines: Fully self-referential optimal universal self-improvers. In *Artificial General Intelligence*. Springer, 2006.
- Solomonoff, R. J. A formal theory of inductive inference. Part I. *Information and control*, 7(1):1–22, 1964.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- Steunebrink, B. R., Thórisson, K. R., and Schmidhuber, J. Growing recursive self-improvers. In *Proc. Conf. on Artificial General Intelligence (AGI)*, pp. 129–139, New York, NY, USA, July 2016.

- Tolstikhin, I. O., Houlsby, N., Kolesnikov, A., Beyer, L., Zhai, X., Unterthiner, T., Yung, J., Steiner, A., Keysers, D., Uszkoreit, J., et al. MLP-mixer: An all-MLP architecture for vision. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, Virtual only, December 2021.
- Tompson, J., Goroshin, R., Jain, A., LeCun, Y., and Bregler, C. Efficient object localization using convolutional networks. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pp. 648–656, Boston, MA, USA, June 2015.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, pp. 5998–6008, Long Beach, CA, USA, December 2017.
- Vinyals, O., Blundell, C., Lillicrap, T., Kavukcuoglu, K., and Wierstra, D. Matching networks for one shot learning. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, pp. 3630–3638, Barcelona, Spain, December 2016.
- Waibel, A., Hanazawa, T., Hinton, G., Shikano, K., and Lang, K. J. Phoneme recognition using time-delay neural networks. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 37(3):328–339, 1989.
- Wang, J., Kurth-Nelson, Z., Soyer, H., Leibo, J. Z., Tirumala, D., Munos, R., Blundell, C., Kumaran, D., and Botvinick, M. M. Learning to reinforcement learn. In *Proc. Annual Meeting of the Cognitive Science Society (CogSci)*, London, UK, July 2017.
- Wang, P. The logic of intelligence. In *Artificial general intelligence*, pp. 31–62. Springer, 2007.
- Wang, P., Li, X., and Hammer, P. Self in NARS, an AGI system. *Frontiers in Robotics and AI*, 5:20, 2018.
- Widrow, B. and Hoff, M. E. Adaptive switching circuits. In *Proc. IRE WESCON Convention Record*, pp. 96–104, Los Angeles, CA, USA, August 1960.
- Williams, R. J. and Peng, J. An efficient gradient-based algorithm for on-line training of recurrent network trajectories. *Neural computation*, 2(4):490–501, 1990.
- Yap, P. C., Ritter, H., and Barber, D. Addressing catastrophic forgetting in few-shot problems. In *Proc. Int. Conf. on Machine Learning (ICML)*, pp. 11909–11919, Virtual only, July 2021.
- Zhmoginov, A., Sandler, M., and Vladymyrov, M. Hypertransformer: Model generation for supervised and semi-supervised few-shot learning. *Preprint arXiv:2201.04182*, 2022.

A. SRWM Model Details

Equations for the four-learning rate case. In Sec. 3, for the purpose of clarity, we presented the equations for our SRWM model in the case where we only have a single learning rate β_t . Here we provide a complete description of an SRWM with a separate self-invented learning rate for each component. As we noted in Sec. 3, the SRWM can be split into sub-matrices: $\mathbf{W}_{t-1} = [\mathbf{W}_{t-1}^y, \mathbf{W}_{t-1}^q, \mathbf{W}_{t-1}^k, \mathbf{W}_{t-1}^\beta]$ according to the sub-components used to produce \mathbf{y}_t , \mathbf{q}_t , \mathbf{k}_t , and β_t in Eq. 5. In case where we use separate learning rates, we need separate equations to describe the update of each sub-matrix. For example, for the “y”-part \mathbf{W}_{t-1}^y , while keeping the same equation for the first projection (Eq. 5), the rest becomes:

$$\mathbf{y}_t^k = \mathbf{W}_{t-1}^y \phi(\mathbf{k}_t) \quad (9)$$

$$\mathbf{y}_t^q = \mathbf{W}_{t-1}^y \phi(\mathbf{q}_t) \quad (10)$$

$$\mathbf{W}_t^y = \mathbf{W}_{t-1}^y + \sigma(\beta_{y,t})(\mathbf{y}_t^q - \mathbf{y}_t^k) \otimes \phi(\mathbf{k}_t) \quad (11)$$

where \mathbf{y}_t^k and \mathbf{y}_t^q are the “y”-part of $\bar{\mathbf{v}}_t$ and \mathbf{v}_t in Eq. 6 and 7 respectively, and $\beta_{y,t} \in \mathbb{R}$ is one of four learning rates dedicated to the “y”-part. The equations for other sub-matrices \mathbf{W}_{t-1}^q , \mathbf{W}_{t-1}^k , \mathbf{W}_{t-1}^β are analogous.

Use of multiple heads. We inserted the SRWM between other layers with learned parameters and configurable dimensionalities. This allows for efficient computation using multiple heads as follows. Given a number of heads H used in the SRWM layer, the model dimensions are configured such that the input dimension to an SRWM layer d_{in} is divisible by H . The input is then split into H equally sized components, and each head executes separate SRWM operations (Eqs. 5-8) on one of the input components. In consequence, an SRWM has fewer parameters than a DeltaNet with the same model hyper-parameters. For example, if $d_{\text{in}} = d_{\text{key}}$, the common head dimension is $d = d_{\text{in}}/H$, the parameter shape of key projection in the SRWM is (H, d, d) while it is $(d_{\text{in}}, d_{\text{in}}) = (H * d, H * d)$ for the DeltaNet. This is an important remark as for some tasks (such as language modelling), the number of parameters can be the dominant factor for good performance. If the input size of the SRWM layer is not configurable, this option has to be disabled and a single head version should be used.

B. Experimental Details for Few-Shot Learning

B.1. Datasets

We conduct few-shot image classification experiments using the Omniglot (Lake et al., 2015), and Mini-ImageNet (Vinyals et al., 2016; Ravi & Larochelle, 2017) datasets. Extra experiments using the Fewshot-CIFAR100 dataset (FC100 for short; Oreshkin et al. (2018)) are presented in Appendix B.4 below. We use `torchmeta` by Deleu et al. (2019) which implements all common settings used with these datasets. For each dataset, classes are split into train, validation and test for few-shot learning settings.

Omniglot images are grayscale hand-written characters from 50 different alphabets, and the dataset contains 1632 different classes with 20 examples per class. The original setting (Lake et al., 2015) splits these 1632 classes into 1200 for training and 432 for testing without validation set. Instead, we use Vinyals et al. (2016)’s 1028/172/432-split for the train/validation/test set and their data augmentation methods based on rotation (90, 180, and 270 degrees). The images are typically resized to 28×28 . In the sequential multi-task experiments (Sec. 4.2), to jointly train on Omniglot and Mini-ImageNet, we resized Omniglot images to 84×84 and duplicated the channels to match the number of RGB color channels for Mini-ImageNet, such that the same feature extractor (here, Conv-4) can process both of them.

Mini-ImageNet contains 100 color image classes with 600 examples for each class. They are typically resized to 84×84 . The standard class train/valid/test splits of 64/16/20 are used (Ravi & Larochelle, 2017).

FC100 is based on CIFAR100 (Krizhevsky, 2009). 100 color image classes (600 images per class, each of size 32×32) are split into train/valid/test classes of 60/20/20 (Oreshkin et al., 2018).

B.2. Model and Training Details

Vision feature extractors. We mainly evaluated our models on few-shot learning using the standard convolution-based vision feature extractor, Conv-4, proposed by Vinyals et al. (2016). It has four blocks, each consisting of one 3×3 2D-convolutional layer, batch normalisation, max-pooling of size 2, and a ReLU activation layer. We use 32 channels in each layer. The feature dimension of this encoder’s output is 64, 800, and 128 for Omniglot, Mini-ImageNet, and FC100, respectively. Dropout (Hanson, 1990; Srivastava et al., 2014) is applied after each max-pooling layer.

Sequence processing components. Both the SRWM and DeltaNet-based models used in this work follow the basic Transformer architecture (Vaswani et al., 2017) where the self-attention layers are replaced by the corresponding DeltaNet (Eqs. 1-4) and SRWM (Eqs. 5-8) operations. For supervised tasks, no activation function is applied to x_t in Eq. 5 of the SRWM, while softmax is applied for the RL experiments.

Training hyper-parameters. For Omniglot, we use two layers of size 256 using 16 computational heads and 1024 ($4 * 256$) dimensional feed-forward inner dimensions. We train with a learning rate of $1e-3$ with a batch size of 128 for 300 K steps and validate every 1000 steps. For Mini-ImageNet, we conduct hyper-parameter search for the SRWM and the DeltaNet as follows: a number of layers $l \in \{2, 3, 4\}$, a hidden size $d_{\text{model}} \in \{128, 256\}$, two dropout rates $p_{\text{vision}}, p \in \{0.0, 0.1, 0.2, 0.3\}$ (separately for the vision and sequence processing components) and a learning rate $\eta \in \{1e-3, 3e-4, 1e-4\}$ or the standard Transformer “warmup” learning rate scheduling (Vaswani et al., 2017). The number of heads is fixed to 16. We set a feed-forward inner dimension to $d_{\text{ff}} = m * d_{\text{model}}$ where $m \in \{4, 8\}$. For the 1-shot Mini-ImageNet, the best SRWM was obtained for ($l = 3, d_{\text{model}} = 256, d_{\text{ff}} = 1024, p_{\text{vision}} = 0.1, p = 0.1, \eta = \text{“warmup”}$) and the best DeltaNet was obtained for ($l = 4, d_{\text{model}} = 256, d_{\text{ff}} = 2048, p_{\text{vision}} = 0.3, p = 0.1, \eta = 1e-4$). For the 5-shot Mini-ImageNet, the best SRWM was obtained for ($l = 3, d_{\text{model}} = 256, d_{\text{ff}} = 2048, p_{\text{vision}} = 0.1, p = 0.1, \eta = \text{“warmup”}$) and the best DeltaNet was obtained for ($l = 4, d_{\text{model}} = 256, d_{\text{ff}} = 2048, p_{\text{vision}} = 0.2, p = 0.0, \eta = 3e-4$). In both cases, the number of trainable parameters is 3.6 M for the SRWM and 5.5 M for the DeltaNet (see comments in Appendix A above). For Mini-ImageNet, we use a batch size of 16 and train for 600 K steps (470 K steps roughly correspond to one epoch, i.e., covering all class combinations for the choice of 5 classes for few-shot learning). All models are trained using the Adam optimiser.

B.3. Evaluation Procedure

Following the standard evaluation setting of Ravi & Larochelle (2017), we report the mean and 95% confidence interval computed over multiple sets of test episodes. We use 5 different sets consisting of 16 K random test episodes each. However, we also note that this metric does not show the performance variance across seeds (which is typically high on these tasks—this is not specific to our approach). In fact, for many configurations, with some seeds, the model performance can remain around the random guessing accuracy of 20% for the entire duration of training.

B.4. Extra Results

Few-Shot CIFAR100 (FC100). Here we provide extra experimental results on the FC100 dataset. We conduct a hyper-parameter search similar to the one described above for Mini-ImageNet. The best SRWM is obtained for ($l = 3, d_{\text{model}} = 256, d_{\text{ff}} = 1024, p_{\text{vision}} = 0.1, p = 0.1, \eta = \text{“warmup”}$) and the best DeltaNet is obtained for ($l = 3, d_{\text{model}} = 256, d_{\text{ff}} = 2048, p_{\text{vision}} = 0.2, p = 0.0, \eta = \text{“warmup”}$). Table 5 shows the results. First of all, we confirm again that the SRWM and the DeltaNet achieve similar performance. However, our generic models underperform the TADAM (Oreshkin et al., 2018) method by a large margin.⁴ To match the vision model architecture, we also test the SRWM using the Res-12 architecture of Chen et al. (2021) which has four residual blocks consisting of three 3×3 convolutional layers. The Res-12-256 architecture uses the following numbers of channels: 64, 96, 128, 256 in the respective blocks (Mishra et al., 2018), and spatial dropout (Tompson et al., 2015) is used. While we achieve slight improvements through this Res-12-256 vision model, no further improvement is obtained on this task by increasing the number of channels as in the Res-12-512 architecture of TADAM.

We observe that the SRWM outperforms the ProtoNet variant (ProtoNet Cosine; Oreshkin et al. (2018)) using the cosine similarity without a trainable scaling factor inside the softmax (like in the output layer of our models), but underperforms the one with scaling (ProtoNet Cosine Scaled). This result seems to indicate that additional tuning specific to few-shot learning is necessary to achieve competitive performance on this more challenging task.

Tuning the LSTM baseline. It is typically found difficult to achieve competitive performance on few-shot image classification tasks using an LSTM (Mishra et al., 2018). Indeed, we find that for many configurations, the LSTM’s performance remains at chance level. For some seeds and with longer training, however, the LSTM can achieve competitive performance. Table 6 shows the best obtained results. Lack of training stability prevents it from being a reliable baseline (e.g., we did not manage to successfully train any LSTM model on FC100), but we show that we can obtain good LSTM results in case of fully successful training.

⁴In a previous preprint of this paper, we reported 57.1% (instead of 44.5 %) for the Conv-4-32 backend in this setting, due to a bug in the meta-test data shuffling process. Details of this correction can be found in our public code on GitHub. We re-ran all important experiments after this fix to produce the numbers presented in the present version.

Table 5. 5-way 5-shot image classification experiments on **FC100**. All numbers marked by * are taken from Oreshkin et al. (2018).

Model	Backend	Accuracy
ProtoNet Cosine*	Res12-512	40.9 ± 0.6
ProtoNet Cosine Scaled*		51.0 ± 0.6
TADAM*		56.1 ± 0.4
DeltaNet	Conv-4-32	43.2 ± 0.1
SRWM	Conv-4-32	44.5 ± 0.1
	Res-12-256	46.3 ± 0.1

The best LSTM configuration for Mini-ImageNet has one layer with a hidden size of 1024 for both 1-shot and 5-shot cases (7.5 M parameters), with a vision dropout rate of 0.1. The warmup learning rate scheduling also works well for the LSTM. The model is trained for 1.2 M steps in the 5-shot case (while 600 K steps are enough for the DeltaNet and SRWM). We use the Adam optimiser (Kingma & Ba, 2014) and a batch size of 16. For Omniglot, the model has two layers with a hidden size of 512, which is trained with a learning rate of 1e-3 and a batch size of 128 for 300 K steps similarly to other models.

Table 6. Test accuracies (%) of LSTM on single task, 5-way, few-shot classification on Omniglot and Mini-ImageNet. We report 95% confidence intervals computed over five sets of test episodes. “Steps” refer to training steps; each is a batch of 16 episodes.

	Omniglot		Mini-ImageNet			
	1-shot	Backend	1-shot		5-shot	
			600 K steps	1.2 M steps	600 K steps	1.2 M steps
LSTM	96.3 ± 0.1	Conv-4-32	< 38	46.6 ± 0.1	< 55	59.0 ± 0.1

Extra results on the sequential multi-task setting. In Table 2, we present results for sequential multi-task few-shot learning experiments where the test sequences start with Omniglot and then change to Mini-ImageNet. Here we reverse the order of testing: Mini-ImageNet, then Omniglot. Table 7 shows the results. We observe that improvements by the SRWM on the (second) Omniglot part are slightly larger in this ordering. We also show the corresponding test accuracy curves in Figures 6 (Omniglot first) and 7 (Mini-ImageNet first). Note that the SRWM part of Figure 6 is already present in Figure 4 (where the DeltaNet part is omitted for clarity).

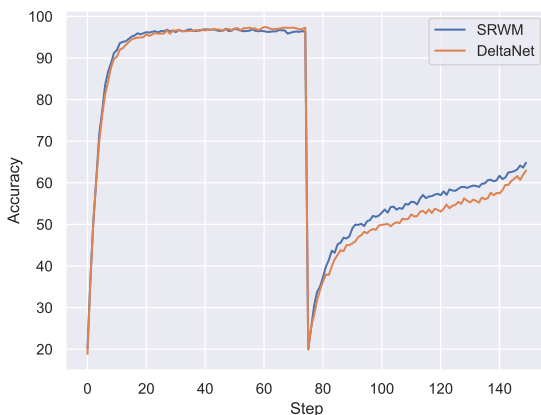


Figure 6. Average test accuracy (%) as a function of the number of steps in the sequential multi-task adaptation setting. The datapoints are sampled from **Omniglot** until step 74 (where the accuracy drops), then from **Mini-ImageNet**.

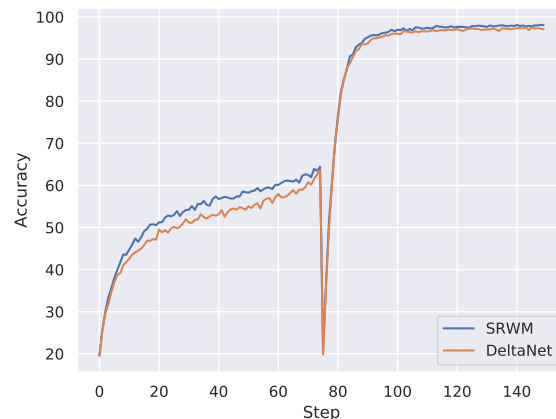


Figure 7. Average test accuracy (%) as a function of the number of steps in the sequential multi-task adaptation setting. The datapoints are sampled from **Mini-ImageNet** until step 74 (where the accuracy drops), then from **Omniglot**.

Table 7. Total and instance-level accuracies (%) for **sequential multi-task few-shot learning** experiments (Sec. 4.2). Regarding the instance-level accuracy, column $k \in \{1, 2, 3, 5, 10\}$ shows the percentage of correctly predicted k -th instances from each class. In this test time scenario the model is first tasked to learn to predict **Mini-ImageNet**, then **Omniglot**.

Task	Model	1	2	3	5	10	Total
Mini-ImageNet	DeltaNet	20.3	44.7	48.9	51.8	54.8	50.7
	SRWM	22.4	46.4	50.3	54.3	57.9	53.5
Omniglot	DeltaNet	39.2	90.9	93.8	95.7	96.9	92.2
	SRWM	39.6	91.2	94.6	96.6	97.6	92.9

B.5. Training in Delayed Label Setting

In the sequential multi-task setting of Sec. 4.2, for both the SRWM and the DeltaNet, we set the number of layers to three, with a hidden layer size of 256 using 16 computational heads, and a feed-forward dimension of 256. We train with a batch size of 32 (crucial for successful training) and a learning rate of $3e-4$. During training, each Omniglot and Mini-ImageNet segment is constructed using up to 15 examples per class (which yields the maximum segment length of 75 images in this 5-way classification setting). For each batch, the number of positions to be trimmed is randomly sampled between 1 and 60 for Omniglot and Mini-ImageNet segments during training, to prevent training sequences from having always the same number of examples per class. As described in the main text (Sec. 4.2), bad configurations or unlucky seeds get stuck with sub-optimal behaviour. Here we show some training curves of such behaviours in Figure 8, as well as successful ones in Figure 9.

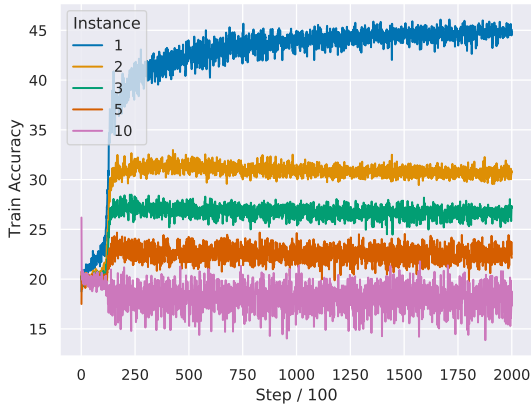


Figure 8. Example **unsuccessful** training on the sequential multi-task few-shot learning task. Average training accuracies for different instances are shown for the Omniglot part.

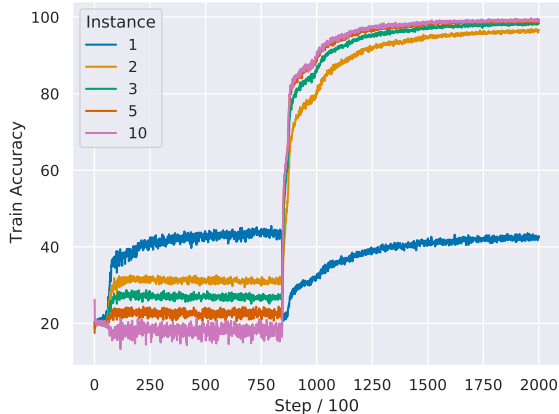


Figure 9. Example **successful** training on the sequential multi-task few-shot learning task. Average training accuracies for different instances are shown for the Omniglot part.

C. Additional Results for Reinforcement Learning Experiments

C.1. Experiments on ProcGen Memory Distribution

Sec. 4.3 presents our experimental results on 6 environments in the easy distribution. Here we present an extra experiment using 4 environments in the *memory* distribution (*Dodgeball*, *Heist*, *Maze*, *Miner*) to evaluate our models also in partially observable settings. In ProcGen memory distributions, the world size is increased and the observations are restricted to a small patch of space around the agent (Cobbe et al., 2020). In such partially observable environments, the baseline models have to be sequence processing NNs with memory, such as the DeltaNet. Here we are interested in augmenting such a model with an extra self-referential mechanism. We thus replace the slow weight matrix of DeltaNet (Eq. 1) by an SRWM (Eqs. 5-8). The resulting model thus learns to generate and update a fast weight matrix as a short-term memory, while it also

Table 8. ProcGen normalised aggregated scores (multiplied by 100) over **4 environments** (Dodgeball, Heist, Maze, Miner) in the **memory** distribution. The models are trained in a multi-task setting. We use the normalisation constants for the hard distribution provided in the original ProcGen paper. Results are derived from 3 independent training runs with 800 M steps each. The test scores are averaged over 3 distinct sets of 200 fixed test levels (i.e. the mean/std computed from 9 data points). For further details, see tables in Appendix C.2 where we provide scores obtained for each game.

	DeltaNet	SR-DeltaNet
Train	51.8 (2.6)	59.0 (2.1)
Test	38.0 (4.1)	38.5 (3.2)

learns to modify itself. We denote this model “SR-DeltaNet.”

We train with a backpropagation span of 100 steps for this memory distribution setting. While we train for a total of 300 M steps (ca. 50 M per environment) for the joint training on 6 environments in the easy distribution, here we train for 800 M steps (200 M steps per environment) on 4 environments.

Since there is no standard convention (Cobbe et al., 2020) for the memory distribution setting, we opt for training on 500 training levels (level ID 0 to 499) as recommended for the “hard” distribution. For testing, we use the exact same setting as in our easy distribution setting described in Sec. 4.3 (i.e. 3 training runs and 3 different test sets).

Table 8 shows the results. While having a similar parameter count, the SRWM variant achieves a better training score than the DeltaNet baseline, while the test scores are rather close.

We note, however, that this is only a preliminary result in the memory setting, as the model size is rather small: we use the same 2-layer architecture used in the experiments with the easy distribution. Further scaling up the model size (e.g., more layers) should lead to further performance improvements as it would allow for handling longer contexts.

C.2. Extra result tables

Table 9. Performance on ProcGen game environments. Multi-task training in 6 environments in the **easy distribution**. The three SRWM variants are as follows: *True*: the SRWM model, *Fake*: the SRWM without self-modification mechanism, and *Reset*: the SRWM trained and evaluated with weight update reset.

Env	Split	FF	LSTM	Delta	SRWM		
					True	Fake	Reset
Bigfish	Train	8.3 (3.9)	6.5 (2.0)	19.6 (4.0)	20.1 (2.4)	11.6 (5.7)	15.7 (2.8)
	Test	4.3 (2.3)	3.2 (1.1)	7.8 (1.5)	9.0 (2.0)	4.7 (2.4)	5.8 (1.3)
Fruitbot	Train	29.2 (0.2)	27.8 (0.5)	28.8 (0.9)	28.7 (0.2)	27.8 (1.3)	29.2 (0.2)
	Test	25.6 (1.1)	24.8 (0.7)	24.5 (1.5)	25.5 (1.0)	24.6 (1.2)	25.2 (1.4)
Leaper	Train	3.3 (0.2)	3.3 (0.2)	3.5 (0.4)	3.5 (0.2)	3.3 (0.3)	3.4 (0.2)
	Test	3.4 (0.4)	3.6 (0.4)	3.3 (0.4)	3.4 (0.4)	3.6 (0.4)	3.5 (0.3)
Maze	Train	1.9 (0.3)	3.1 (0.7)	3.8 (0.2)	3.6 (0.5)	3.2 (0.2)	2.9 (0.2)
	Test	1.4 (0.3)	1.6 (0.4)	1.7 (0.2)	1.8 (0.5)	1.3 (0.3)	1.5 (0.3)
Plunder	Train	3.2 (0.2)	3.2 (0.4)	3.3 (0.2)	3.1 (0.0)	3.1 (0.4)	3.1 (0.1)
	Test	3.2 (0.3)	2.9 (0.4)	3.3 (0.2)	3.0 (0.2)	3.1 (0.5)	3.0 (0.2)
Starpilot	Train	57.6 (0.9)	56.0 (1.5)	60.3 (0.4)	61.3 (2.0)	55.0 (1.3)	55.0 (1.9)
	Test	53.0 (1.7)	48.3 (2.0)	53.9 (2.4)	54.6 (2.4)	49.6 (2.1)	48.6 (1.9)
Aggregated	Train	22.5 (2.6)	28.3 (1.4)	35.0 (1.6)	27.0 (1.8)	34.6 (1.8)	28.5 (1.2)
	Test	16.4 (1.6)	15.7 (1.6)	18.6 (1.7)	20.0 (1.8)	15.3 (1.9)	16.1 (2.2)

Table 10. Performance on ProcGen game environments. Multi-task training in 4 environments in the **memory distribution**.

Env	Split	DeltaNet	SRM-Delta
Dodgeball	Train	7.1 (0.2)	7.1 (0.6)
	Test	6.4 (0.3)	6.2 (0.6)
Heist	Train	1.0 (0.3)	1.5 (0.1)
	Test	0.8 (0.2)	1.1 (0.3)
Maze	Train	5.3 (0.4)	5.9 (0.2)
	Test	3.3 (0.6)	3.3 (0.4)
Miner	Train	32.3 (0.4)	34.5 (0.8)
	Test	29.2 (1.1)	29.4 (0.7)
Aggregated	Train	51.8 (2.6)	59.0(2.1)
	Test	38.0 (4.1)	38.5(3.2)