

Learning to Continually Learn Rapidly from Few and Noisy Data

Nicholas I-Hsien Kuo¹

Mehrtash Harandi²

Nicolas Fourrier³

Christian Walder^{1,4}

Gabriela Ferraro^{1,4}

Hanna Suominen^{1,4,5}

U6424547@ANU.EDU.AU

MEHRTASH.HARANDI@MONASH.EDU

NFOURRIER@GMAIL.COM

CHRISTIAN.WALDER@DATA61.CSIRO.AU

GABRIELA.FERRARO@DATA61.CSIRO.AU

HANNA.SUOMINEN@ANU.EDU.AU

¹*RSCS, The Australian National University, Canberra, ACT, Australia*

²*ECSE, Monash University, Melbourne, Victoria, Australia*

³*Research Center of Léonard de Vinci Pôle Universitaire, Paris La Défense, France*

⁴*Data61, CSIRO, Canberra, ACT, Australia*

⁵*Department of Computing, University of Turku, Turku, Finland*

Editors: Isabelle Guyon, Jan N. van Rijn, Sébastien Treguer, Joaquin Vanschoren

Abstract

Neural networks suffer from catastrophic forgetting and are unable to sequentially learn new tasks without guaranteed stationarity in data distribution. Continual learning could be achieved via replay – by concurrently training externally stored old data while learning a new task. However, replay becomes less effective when each past task is allocated with less memory. To overcome this difficulty, we supplemented replay mechanics with meta-learning for rapid knowledge acquisition. By employing a meta-learner, which *learns a learning rate per parameter per past task*, we found that base learners produced strong results when less memory was available. Additionally, our approach inherited several meta-learning advantages for continual learning: it demonstrated strong robustness to continually learn under the presence of noises and yielded base learners to higher accuracy in less updates¹.

1. We released our codes at <https://github.com/Nic5472K/MetaLearnCC2021-AAAI-MetaSGD-CL>.

1. Introduction

In standard practice, it is assumed that the examples used to train a neural network are drawn independently and identically distributed (i.i.d.) from some fixed distribution. However, in a changing environment, *continual learning* is required where an agent faces a continual stream of data and must adapt to learn new predictions. This remains a difficult challenge as neural networks are known to suffer from *catastrophic forgetting* (McCloskey and Cohen, 1989), the phenomenon in which they abruptly lose the ability to solve a previously learnt task when information for solving a new task is incorporated.

It was shown that catastrophic forgetting can be effectively alleviated by concurrently training data of past tasks. These *replay* (Robins, 1995) and *pseudo replay* (Shin et al., 2017) mechanisms have continuously been improved (Rebuffi et al., 2017; Lopez-Paz and Ranzato, 2017; Chaudhry et al., 2018); and a recent paper from Chaudhry et al. (2019) have even shown that forget prevention could be made possible by co-training as little as

one single sample of past data per parametric update. Nonetheless, we found that the replayed data was still required to be sampled from a large set of externally stored memory. Sampling from an insufficient amount of external memory would lead a network to overfit early tasks thus generalise poorly on subsequent tasks.

This paper combines the work of Chaudhry et al. (2019) with meta-learning for rapid knowledge acquisition. We extended their method with *MetaSGD* (Li et al., 2017) to maintain high performances under the constraint of learning with low memory consumption. We employed a meta-learner to assist the training procedure of a task-specific base learner of which *learns a learning rate per parameter per past task*. This hybrid approach not only produced strong results when less memory was available, but it also inherited several desirable meta-learning advantages for continual learning. Our approach demonstrated strong robustness against the realistic scenario in which we continually learn under the presence of noises, and it could also optimise base learner parameters to achieve higher accuracy in fewer iterations.

2. Preliminaries

Our narrative below mainly follows Lopez-Paz and Ranzato (2017), Kirkpatrick et al. (2017), and Chaudhry et al. (2019).

2.1. Framework of Continual Learning

In continual learning, we employ a base learner F to sequentially learn the n tasks of $\omega_1, \dots, \omega_n \in \Omega$ in a *continuum* of data

$$\Xi = \{(x_1^1, \omega_1, y_1^1), \dots, (x_\gamma^i, \omega_i, y_\gamma^i), \dots, (x_\Gamma^n, \omega_n, y_\Gamma^n)\} \quad (1)$$

to map inputs x_γ^i to labels y_γ^i from datasets

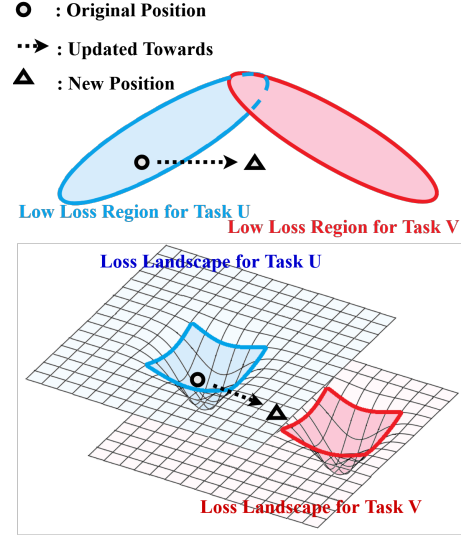


Figure 1: Exiting a Local Minimum While Continual Learning

$(x_\gamma^i, y_\gamma^i) \in \mathcal{D}_{\omega_i}$ for all Γ data where $\gamma = 1, \dots, \Gamma$. The aim is to yield a learner that generalises well on subsequent tasks while it maintains high accuracy for learnt tasks.

Continual learning is challenging (Kemker et al., 2018). When a learner *only* observes data of a new task, its optimised parameters for an old task will be modified to suit the learning objective of the new task. We elaborate this with the scenario presented in Figure 1 where a learner optimised for Task U is now sequentially updated for an independent Task V. Naïvely training the learner with gradient descent will cause the learner to leave the minimum (low loss region) on the loss landscape of Task U in pursuit for that of the new task (Kirkpatrick et al., 2017)².

2.2. Replay with Episodic Memory

Catastrophic forgetting can be alleviated by replaying all data of past tasks but this inevitably creates a large computational and memory burden. Several studies have shown

² Inspiration taken from Figure 1 of Kirkpatrick et al. (2017).

```

Define memory storage
 $\mathcal{M}_{\omega_i} \leftarrow \{\} \quad \forall i \text{ such that } \omega_i \in \Omega$ 
for  $i = 1, \dots, |\Omega|$  do
    for  $(x_\gamma^i, \omega_i, y_\gamma^i) \in \Xi$  do
        Update memory
         $\mathcal{M}_{\omega_i} \leftarrow \mathcal{M}_{\omega_i} \cup (x_\gamma^i, y_\gamma^i)$ 
        if  $i > 2$  then
            Sample from all memory
             $B_{\mathcal{M}} \sim \bigcup_{u=1}^{i-1} \mathcal{M}_{\omega_u}$ 
            Update with sampled memory
             $\theta \leftarrow \text{SGD}((x_\gamma^i, y_\gamma^i) \cup B_{\mathcal{M}}, \theta)$ 
        else
            Conventional update
             $\theta \leftarrow \text{SGD}((x_\gamma^1, y_\gamma^1), \theta)$ 
        end
    end
end
end
    
```

Algorithm 1: Experience Replay Recipe

that memory-based strategy does not require all past training data. Rebuffi et al. (2017) prevented forgetting by constructing an exemplar set of past data via learning representation; while Lopez-Paz and Ranzato (2017) preconditioned the gradients of the current learning task with small *episodic memories* of past tasks. This paper will focus specifically on episodic memory-based methods as Lopez-Paz and Ranzato’s original work was continually streamlined by the important works of Chaudhry et al. (2018) and Chaudhry et al. (2019).

There are two steps to the episodic memory mechanism. Before the sequential training phase, Lopez-Paz and Ranzato first define the size of the memory storage \mathcal{M}_{ω_i} allocated for each task. Then upon training a new task, all stored data from the external memory was used to precondition the gradients of the current task. Though this method yielded very strong results and only stored a section of all past data, it still had difficulties in scaling well as the amount of subsequent

tasks increased.

In their paper, Chaudhry et al. (2019) streamlined Lopez-Paz and Ranzato’s method with *Experience Replay* (ER), and showed that catastrophic forgetting can be alleviated without utilising all data stored within episodic memories. Algorithm 1 shows their ER approach; and the core mechanism of ER is to *sample* mini-batches $B_{\mathcal{M}}$ from all stored data of past tasks (see line 8) to effectively mitigate the computational burden from employing episodic memories. The sampled content was then concurrently trained with the existing batch of the current learning task and they used SGD to update parameters θ of base learner F (see line 10). Impressively, they found that strong results for forget prevention can be achieved even when $B_{\mathcal{M}}$ only sampled 1 data from all memory $\bigcup \mathcal{M}_{\omega_i}$.

2.3. Learning with Episodic Memory Sampled from a Tiny Memory Storage

To employ episodic memory, the practitioner would need to explicitly define the size of the external memory storage;

but what are the disadvantages of selecting small sizes?

A small memory storage meant that the replay mini-batch $B_{\mathcal{M}}$ would have a higher chance of re-sampling the same past data for concurrent training (see line 11 of Algorithm 1). A base learner would hence overfit on earlier tasks through the consecutive re-exposure to a small set of past data and thus underfit on future tasks. This issue was lightly investigated in Chaudhry et al. (2019)³, but they only compared medium-size to large-size memories⁴. In this current work, we limit the memory size to only a handful of

3. See Section 5.5 and Appendix C of Chaudhry et al. (2019).

4. See Table 7 in Chaudhry et al. (2019), where they compared 1000 to 20000 memory units. What happens when we use 100?

data per past task to study its consequences for continual learning and propose meta-learning as a remedy.

3. Meta-learning

Meta-learning studies rapid knowledge acquisition and has contributed to few shot learning (Ravi and Larochelle, 2017) where a base learner is required to learn from a low resource environment. In this paper, we formulate continual learning with tiny episodic memory storage as a low resource problem. Though sequential learning does not lack data on subsequent tasks, our base learner is still required to learn from a small memory unit for past tasks.

3.1. Learning to Optimise with MetaSGD

One archetype of meta-learning focuses on optimisation. This type of approach is also known as *learning to learn* (Andrychowicz et al., 2016). Optimisation is cast as the learning problem of a meta-learner of which is used to configure the weights of the base learner network. It replaces the classic gradient descent of

$$\theta_{t+1} = \theta_t - \alpha \nabla_{\theta} \mathcal{L}_{t+1}(\theta_t) \quad (2)$$

for a base learner F with parameters θ_t at time t ; where α is the learning rate and that $\nabla_{\theta} \mathcal{L}_{t+1}(\theta_t)$ is the gradient of the learner loss with respect to the parameters.

One particularly simple method called *MetaSGD* (Li et al., 2017) proposes to update learner parameters via

$$\theta_{t+1} = \theta_t - \beta \odot \nabla_{\theta} \mathcal{L}_{t+1}(\theta_t), \quad (3)$$

where $\beta \in \mathbb{R}^{|\theta_t|}$ and that symbol \odot represents the element-wise product. Learning rates β are the meta-learning target of MetaSGD, and hence a unique learning rate is learnt for every parameter to enable important fea-

tures to be updated more quickly. In their paper, Li et al. showed that MetaSGD achieved strong results for few shot learning; and Kuo et al. (2020) showed that MetaSGD could converged base learners much more rapidly than SGD⁵. We will now simplify our notation by denoting $\nabla_{\theta} \mathcal{L}_{t+1}(\theta_t)$ as $\nabla \mathcal{L}$.

3.2. An Analysis on the Update of Experience Replay for Continual Learning

Under the conventional setup, loss \mathcal{L} aggregates every individual cost C therein a mini-batch B_n with

$$\mathcal{L} = \frac{1}{|B_n|} \sum_{i=1}^{|B_n|} C(\hat{y}_i, y_i) \quad (4)$$

where \hat{y}_i is the prediction of learner F and that y_i is the corresponding true label. Hence under the ER regime, every batch used in continual learning will be substituted with $B_n \leftarrow B^*$ where $B^* = B_n \cup B_{\mathcal{M}}$ (see line 11 of Algorithm 1). Then because the linearity of summation, the new loss \mathcal{L}^* can be written as

$$\mathcal{L}^* = \frac{1}{|B^*|} \left[\sum_{i=1}^{|B_n|} C(\hat{y}_i, y_i) + \sum_{j=1}^{|B_{\mathcal{M}}|} C(\hat{y}_j, y_j) \right] \quad (5)$$

To highlight task specificity on the past $V-1$ tasks, we can rewrite the Equation as

$$\mathcal{L}^* = \frac{1}{|B^*|} \left[\sum_{i=1}^{|B_n|} C(\hat{y}_i, y_i) + \sum_{u=1}^{V-1} \sum_{j=1}^{|B_{\mathcal{M}_u}|} C(\hat{y}_j, y_j) \right] \quad (6)$$

where $B_{\mathcal{M}} = \cup_{u=1}^{V-1} B_{\mathcal{M}_u}$. Then with \mathbb{L} as the

5. See Figure 5 of Kuo et al. (2020) where MetaSGD-like methods decreased updates from 37500 to 10000 steps.




Figure 2: Task-wise Directional Updates

task-wise loss we have

$$\mathcal{L}^* = \frac{1}{|B^*|} \left[|B_n| \mathbb{L}_V + \sum_{u=1}^{V-1} |B_{\mathcal{M}_u}| \mathbb{L}_u \right]; \quad (7)$$

and this will create the standard parametric update of

$$\theta_{t+1} = \theta_t - \left(\frac{\alpha}{|B^*|} |B_n| \right) \nabla \mathbb{L}_V - \sum_{u=1}^{V-1} \left(\frac{\alpha}{|B^*|} |B_{\mathcal{M}_u}| \right) \nabla \mathbb{L}_u \quad (8)$$

of which we could employ MetaSGD to introduce β as

$$\theta_{t+1} = \theta_t - \beta_V \odot \nabla \mathbb{L}_V - \sum_{u=1}^{V-1} \beta_u \odot \nabla \mathbb{L}_u \quad (9)$$

and learn a learning rate per parameter per past task.

To elaborate, consider the scenario in Figure 2 where a learner observes its third task. The learner faces different directions of update $\nabla \mathbb{L}$ from the competing learning objectives of different tasks. *Interference* (Riemer et al., 2018) hence occurs and transfer in knowledge is difficult; this is further complicated as learner F could easily have more than thousands of parameters in practice (see Szegedy et al. (2015) for instance). Thus

instead of applying handcrafted rules such as SGD to update a base learner, we employ **MetaSGD for Continual Learning** (MetaSGD-CL) to explore a stream of changes in a series of task structures.

4. Implementation Details on MetaSGD-CL

Upon learning a new task V (including task 1), MetaSGD-CL initialises a trainable vector $b_V = 0.01 \cdot \mathbf{\hat{1}} \in \mathbb{R}^{|\theta|}$. In the context of Equation (9), we set

$$\beta_{V_k} = \text{ReLU}(\min(b_{V_k}, \kappa)); \quad (10)$$

where κ is a postive constant hyper-parameter. Thus every individual learning rate β_{V_k} lies within the range of $[0, \kappa]$. In this paper, we default κ as 0.02 unless specified otherwise.

When sequential learning a new task V (excluding task 1), MetaSGD-CL freezes past learning rates β_u for all $u < V$ and updates base learner F with the formulation below

$$\theta_{t+1} = \theta_t - \beta_V \odot \nabla \mathbb{L}_V - \frac{1}{V-1} \sum_{u=1}^{V-1} \beta_u \odot \nabla \mathbb{L}_u. \quad (11)$$

Equation (11) accounts for alignments in $\nabla \mathbb{L}$ due to similarity among arbitrary task structures; $\frac{1}{V-1}$ hence prevents the over-parametrisation in any given feature dimension.

The loss function of MetaSGD-CL is adapted from Andrychowicz et al. (2016)

$$\mathcal{L}(\beta_V) = \mathcal{L}(F_{\theta_{t+1}}(x_t), y_t | \beta_V, \{\beta_u \text{ for } u = 1, \dots, (V-1)\}), \quad (12)$$

based on an updated learner $F_{\theta_{t+1}}$ with original data x_t . Thus, the objective for MetaSGD-CL is to further fine-tune learner F . The meta-learning targets β s are updated with Adam (Kingma and Ba, 2015) with learning rate 0.01.

5. Experiments

This section includes some benchmarks, demonstrates the benefits of MetaSGD-CL over ER, and an ablation study.

5.1. Dataset

We conducted experiments on *permuted MNIST* (Kirkpatrick et al., 2017) and *incremental Cifar100* (Lopez-Paz and Ranzato, 2017). Previously, Lopez-Paz and Ranzato (2017) and Chaudhry et al. (2019) showed that memory-based approaches outperformed continual learning techniques of other archetypes and achieved state of the art status. However, we will show that when the amount of memory decreases, ER performances can drop sharply on the considerably simple permuted MNIST tasks.

MNIST (LeCun et al., 1998) contains 60,000 training and 10,000 test images in 10 classes of 28×28 grey-scale handwritten digits. Permuted MNIST first reshaped each image as a vector of 784 ($=28 \times 28$) pixels, then applied a unique permutation on all vectors of that task. In our experiments, we sequentially trained over 10 permuted MNIST tasks.

Cifar100 (Krizhevsky, 2009) is a coloured dataset with 100 classes of objects. There are 50,000 training and 10,000 test images all presented on 32×32 pixel boxes. In our incremental class experiments, we sequentially introduced 5 classes per task drawn without replacement from the dataset. We reserved the first 50 classes for pre-training a classifier, and tested on the remaining 50 classes, hence there were 10 tasks in total. See the next subsection for more details.

5.2. Base Learners

For permuted MNIST, we followed Lopez-Paz and Ranzato (2017) and trained MLPs as base learners. As for incremental Cifar100,

we followed Mai et al. (2020)⁶ and used their best practice which employed a fixed feature extractor and only sequentially trained the classifier networks. We pre-trained a ResNet18 (He et al., 2016) as feature extractor on the first 50 classes of Cifar100; then updated MLP base learners as classifiers. All MLP base learners of our study had 2 hidden layers with 100 dimensions followed by a softmax layer.

Following Lopez-Paz and Ranzato, we used batch size 10. Each task of permuted MNIST observed 100 batches; and it was 1 epoch each for incremental Cifar100. Our MetaSGD-CL was described in the last section; and all baselines were updated with SGD with learning rate 0.01.

5.3. Metrics

The 2 metrics of which we use in this paper are based on a matrix $\mathcal{R} \in \mathbb{R}^{n \times n}$ of which we employ to store all accuracy during the entire course of continual learning. Row i recorded the accuracy after training on the i th dataset; and column m recorded the accuracy of the m th task.

The severity of catastrophic forgetting can be reflected via

Final Accuracy of Task 1

(FA1) $\mathcal{R}_{n,1}$, ; while

Average accuracy

(ACC) $\frac{1}{n} \sum_{i=1}^n \mathcal{R}_{n,i}$

can show the overall performance and the severity of overfitting of the base learner. For both metrics, it is the larger the better.

5.4. Continual Learning Techniques

Our benchmark experiments included 6 setups. *Singular* was the scenario where we naïvely sequentially trained a base learner. Then, we tested our *MetaSGD-CL* against the baseline *ER* (Chaudhry et al., 2019). In addition, we considered the memory-based

6. Mai et al. (2020) won the CVPR CLVision 2020 challenge.




Figure 3: Performances for Permuted MNIST. Top row reflects catastrophic forgetting; and bottom row reflects overfitting.

Methods	FA1 (%)	ACC (%)
MetaSGD-CL (Ours)	81.02	82.19
ER	80.60	69.42
GEM	80.71	79.43
EWC	64.80	71.84
HAT	74.03	76.17
Singular	62.18	71.34

Table 1: Metrics for Permuted MNIST in Figure 3

Gradient Episodic Memory (GEM) of Lopez-Paz and Ranzato (2017), the regularisation-based *Elastic Weight Consolidation* (EWC) of Kirkpatrick et al. (2017), and the dynamic architectural-based *Hard Attention to the Task* (HAT) of Serra et al. (2018).

5.5. Remark 1: Preventing Catastrophic Forgetting

Our results on permuted MNIST are in Figure 3 and Table 1; and those for incremental Cifar100 are in Figure 4⁷ and Table 2 overleaf. Based on the FA1 metric in the 2 tables, we validated Lopez-Paz and Ranzato (2017) and Chaudhry et al. (2019)’s finding such that memory-based approaches were more effective in preventing catastrophic forgetting.

7. Refer to our codes (link in Abstract) to reformat Figure 4 as Figure 3 for further inspection.

Our memory-based approach used *Hard Storage* (Lopez-Paz and Ranzato, 2017) of which assigned 250 memory units to each task. For each update, the less computationally efficient GEM used all stored data to prevent forgetting; while both MetaSGD-CL and ER randomly sampled 10 data from each task-wise storage to append them as $B_{\mathcal{M}}$ to the original batch B_n (see Equation (5)). The samples of old tasks do not change throughout training.

As an additional note, the base learners of HAT had hidden dimensions of 400 instead. With only 100 hidden dimensions, HAT could not prevent forgetting and behaved like that in Ahn et al. (2019).

5.6. Remark 2: 3 Advantages of Meta-SGD over ER

5.6.1. REMARK 2.1: DOES NOT OVERFIT ON OLD TASKS

From Figure 3 and Table 1, ER was comparable to MetaSGD-CL in FA1 but much lower in ACC. This served as a clear indication of ER’s overfitting problem. For incremental Cifar100 and Table 2, Mai et al. (2020)’s experimental setup limited some extent of overfitting by only continually training the classifier network. However, the bottom subplot

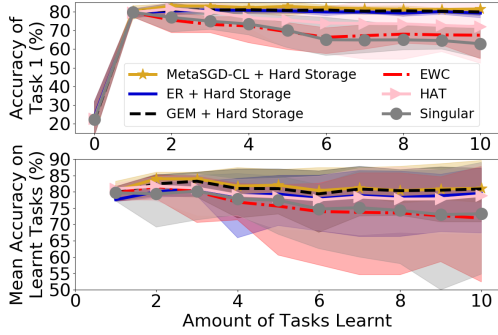


Figure 4: Performances for Incremental Cifar100

Methods	FA1 (%)	ACC (%)
MetaSGD-CL (Ours)	81.52	81.06
ER	79.28	79.50
GEM	79.92	80.88
EWC	67.46	71.98
HAT	71.98	78.86
Singular	62.86	73.27

Table 2: Metrics for Incremental Cifar100 in Figure 4.

of Figure 4 showed that ER’s variability was much larger than MetaSGD-CL’s; and hence overfitting persisted.

We further tested overfitting with the alternative memory storage of *Ring Buffer* (Chaudhry et al., 2019). Ring buffer shares Φ memory units **for all tasks**⁸. By design, the buffer is under-utilised in the early phases of sequential learning; and that stored samples do not change. Compared to hard storage, the ring buffer is much more computationally efficient. However, it is also more likely to incur overfitting because data of all past tasks are now sharing the same buffer.

Figure 5 shows the results for MetaSGD-CL and ER on permuted MNIST with ring buffers of 1000, 250, and 100 units. Note, these settings were much smaller than the scenarios tested in Chaudhry et al. (2019)⁴. Similar to Remark 1, each iteration randomly

8. If $\Phi = 250$ and that $n = 10$ tasks, then it is 25 units per task.




Figure 5: Ring Buffer of Different Sizes for Permuted MNIST

sampling 10 past data as B_M but from one unified storage (see line 11 of Algorithm 1).

Unsurprisingly, the performances of either method decreased as the sizes of memory dropped. This was a natural consequence of overfitting by repeatedly re-sampling from a small set of past data. ER is one of the state of the art techniques; but still, there was a significant amount of lesion on the considerably simple permuted MNIST benchmark when the size of ring buffer changed from 1000 to 100. MetaSGD-CL was more capable of maintaining higher accuracy.

5.6.2. REMARK 2.2: RAPID KNOWLEDGE ACQUISITION




Figure 6: Permuted MNIST but with Iteration Reduction

We also tested 2 realistic continual learning scenarios. First, we considered when there was a limited amount of source data; and second, when data was collected from a noisy environment.

In Figure 6, we tested MetaSGD-CL and ER on permuted MNIST with ring buffer but reduced the available training data from 100 iterations to 25 iterations per task. Unsur-

prisingly, the test accuracy had dropped; in addition, overfitting continued to occur in ER. Our MetaSGD-CL approach was able to leverage meta-learning and it yielded higher accuracy in few iterations; also, its mean accuracy concurrently increased with its accuracy for task 1.

5.6.3. REMARK 2.3:

STRONG ROBUSTNESS TO NOISE




Figure 7: Permuted MNIST but with Noise Injection.

In Figure 7, we tested MetaSGD-CL and ER on permuted MNIST with a ring buffer, and we demonstrated the consequences of *noise injection* (NI) in both the training data and stored memory. NI randomly shuffled a section of pixels in the MNIST images. Our lowest setting injected 10% noise, and the highest setting injected up to 50% noise.

We found that MLPs trained with MetaSGD-CL were less susceptible to NI than their counterparts trained with ER; there were higher accuracy and smaller variability. This was likely because the learning rates β of Equation (9) assigned higher learning rates to features that were less prone to NI.

5.7. Remark 3:

Behaviours on Learning Rate β

5.7.1. REMARK 3.1:

CONVERGENCE IN β VALUES

As we previously mentioned, the maximal value of β was defaulted as $\kappa = 0.02$ (see Equation (10)). The value for κ was purposely set low because we hypothesised that

Proportion of $\beta > 0.05$ (in %)					
Tasks	2	4	6	8	10
Layers					
1	1.39	0.63	0.47	0.46	0.36
2	2.26	1.41	1.33	1.32	1.30
Softmax	3.64	2.28	1.88	1.76	1.74
Proportion of $\beta < 0.02$ (in %)					
Tasks	2	4	6	8	10
Layers					
1	1.43	3.90	6.56	8.31	8.79
2	4.29	6.59	7.71	10.51	11.27
Softmax	6.02	7.50	8.38	8.92	11.58

Table 3: Ratios of Extreme Task-wise β with $\kappa = 0.1$




Figure 8: Transition in Task-wise β Magnitudes

there existed conflicted objectives between meta-learning and continual learning.

As shown in Figure 6, meta-learning is naturally capable of rapid knowledge acquisition. That is, we hypothesised that β would grow unboundedly with large κ values; and this would hence indirectly cause overfitting on old tasks.

To our surprise, $\kappa = 0.1$ yielded 76.95 ACC and this was a minor lesion to $\kappa = 0.02$'s 82.19 ACC shown in Table 1. We found that it was because that extreme β values converged as the base learner observed more tasks. As shown in Table 3 and Figure 8, only a minor proportion of β values were lower than 0.02 and larger than 0.05. We found that as the base learner observed more tasks, the proportion of large magnitude task-wise β decreased, while that of small magnitude task-wise β increased. That is, less parametric updates were made to the

Methods	FA1	ACC
Un-ablated MetaSGD-CL	81.02	82.19
Singular	62.18	71.34
Replace all old β with 0s	69.81	74.41
Replace all old β with 0.01s	73.48	77.16
Replace all old β with 0.1s	72.57	75.80

Both FA1 and ACC are in %.

Table 4: Metrics for permuted MNIST in Figure 3.

base learner from data of later tasks.

Furthermore, a larger proportion of large values were found in the task-wise β of deeper base learner layers. Thus MetaSGD-CL prioritised the update of weights of the deeper layers rather than their counterparts in the shallower layers.

5.7.2. REMARK 3.2: ABLATION STUDY

To further understand MetaSGD-CL, we reported 3 ablation studies in Table 4. When learning a new task, we removed the learnt learning rates β , and replaced them with 0, 0.01, and 0.1.

When β of the old tasks were replaced with 0, catastrophic forgetting occurred. However, these ablated performances were still much better than the naïvely implemented sequential learning. Interestingly, when β of the old tasks were replaced with non-zero values, the ablated performances were similar. We attributed this to formulation Equation (9); as long as information of the past was given, MetaSGD-CL was able to learn appropriate β for the new task to indirectly balance base learner parametric configurations for all tasks.

6. Related Studies and Future Work

Like EWC and *Learning without Forgetting* (LwF) (Li and Hoiem, 2017), MetaSGD-CL

also alleviated catastrophic forgetting via modifications in the backward pass. Thus, it can likely be modified to prevent forgetting via regularisation.

Alternative meta-learning approaches have been employed in continual learning. Riemer et al. (2018) combined GEM with Reptile (Nichol et al., 2018). While Javed and White (2019) continually learnt via representation learning with *Model-Agnostic Meta-Learning* (MAML) (Finn et al., 2017)-like techniques. A large-scale comparison for such hybrid approaches would thus be an interesting future study.

7. Conclusion

This paper introduced **MetaSGD-CL** as an extension to Chaudhry et al. (2019)’s *ER* technique to remedy the overfitting on old tasks when memory-based continual learning was allocated with small memory sizes. MetaSGD-CL leveraged meta-learning by *learning a learning rate per parameter per past task*. It alleviated catastrophic forgetting and also prevented base learners from overfitting and achieved high performances for all tasks. In addition, MetaSGD-CL optimised base learners in fewer iterations and showed higher robustness against noises in training data.

Acknowledgments

This research was supported by the Australian Government Research Training Program (AGRTP) Scholarship. We also thank our reviewers for the constructive feedback.

References

Hongjoon Ahn, Sungmin Cha, Donggyu Lee, and Taesup Moon. Uncertainty-based continual learning with adaptive regularisation. In *Neural Information Processing Systems*, 2019.

- Marcin Andrychowicz, Misha Denil, Sergio Gomez, Matthew W Hoffman, David Pfau, Tom Schaul, Brendan Shillingford, and Nando De Freitas. Learning to learn by gradient descent by gradient descent. In *Neural Information Processing Systems*, 2016.
- Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with a-gem. In *International Conference on Learning Representations*, 2018.
- Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and Marc’Aurelio Ranzato. On tiny episodic memories in continual learning. In *International Conference on Machine Learning Workshop*, 2019.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, 2017.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Computer Vision and Pattern Recognition*, 2016.
- Khurram Javed and Martha White. Meta-learning representations for continual learning. In *Neural Information Processing Systems*, 2019.
- Ronald Kemker, Marc McClure, Angelina Abitino, Tyler Hayes, and Christopher Kanan. Measuring catastrophic forgetting in neural networks. In *AAAI Conference on Artificial Intelligence*, 2018.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimisation. In *International Conference on Learning Representations*, 2015.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. In *Proceedings of the National Academy of Sciences*. National Acad Sciences, 2017.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. In *University of Toronto Tech Report*, 2009.
- Nicholas Kuo, Mehrtash Harandi, Nicolas Fourrier, Christian Walder, Gabriela Ferraro, and Hanna Suominen. M2sgd: Learning to learn important weights. In *Computer Vision and Pattern Recognition Workshop on Continual Learning in Vision*, 2020.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. In *Proceedings of the IEEE*, 1998.
- Zhenguang Li, Fengwei Zhou, Fei Chen, and Hang Li. Meta-sgd: Learning to learn quickly for few-shot learning. *arXiv preprint arXiv:1707.09835*, 2017.
- Zhizhong Li and Derek Hoiem. Learning without forgetting. In *Pattern Analysis and Machine Intelligence*, 2017.
- David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. In *Neural Information Processing Systems*, 2017.
- Zheda Mai, Hyunwoo Kim, Jihwan Jeong, and Scott Sanner. Batch-level experience replay with review for continual learning. In *Computer Vision and Pattern Recognition Workshop*, 2020.
- Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist

- networks: The sequential learning problem. In *Psychology of Learning and Motivation*, 1989.
- Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *arXiv preprint*, 2018.
- Sachin Ravi and Hugo Larochelle. Optimisation as a model for few-shot learning. In *International Conference on Learning Representations*, 2017.
- Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *Computer Vision and Pattern Recognition*, 2017.
- Matthew Riemer, Ignacio Cases, Robert Ajemian, Miao Liu, Irina Rish, Yuhai Tu, and Gerald Tesauro. Learning to learn without forgetting by maximizing transfer and minimizing interference. In *International Conference on Learning Representations*, 2018.
- Anthony Robins. Catastrophic forgetting, rehearsal and pseudorehearsal. In *Connection Science*, 1995.
- Joan Serra, Didac Suris, Marius Miron, and Alexandros Karatzoglou. Overcoming catastrophic forgetting with hard attention to the task. In *International Conference on Machine Learning*, 2018.
- Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. In *Neural Information Processing Systems*, 2017.
- Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Computer Vision and Pattern Recognition*, 2015.