

TARGET LAYER REGULARIZATION FOR CONTINUAL LEARNING USING CRAMER-WOLD GENERATOR

Anonymous authors

Paper under double-blind review

ABSTRACT

We propose an effective regularization strategy (CW-TaLaR) for solving continual learning problems. It uses a penalizing term expressed by the Cramer-Wold distance between two probability distributions defined on a target layer of an underlying neural network that is shared by all tasks, and the simple architecture of the Cramer-Wold generator for modeling output data representation. Our strategy preserves target layer distribution while learning a new task but does not require remembering previous tasks' datasets. We perform experiments involving several common supervised frameworks, which prove the competitiveness of the CW-TaLaR method in comparison to a few existing state-of-the-art continual learning models.

1 INTRODUCTION

The concept of *continual learning (CL)*, which aims to reduce the distance between human and artificial intelligence, seems to be considered recently by deep learning community as one of the main challenges. Generally speaking, it means the ability of the neural network to effectively learn consecutive tasks (in either supervised or unsupervised scenarios) while trying to prevent forgetting already learned information. Therefore, when designing an appropriate strategy, it needs to be ensured that the network weights are updated in such a way that they correspond to both the current and all previous tasks. However, in practice, it is quite likely that constructed CL model will suffer from either *intransigence* (hard acquiring new knowledge, see Chaudhry et al. (2018)) or *catastrophic forgetting (CF)* phenomenon (tendency to lose past knowledge, see McCloskey & Cohen (1989)).

In recent years, methods of overcoming the above-mentioned problems are subject to wide and intensive investigation. Following Maltoni & Lomonaco (2019), the most popular CL strategies may be assigned into the following fuzzy categories:

- (i) *architectural strategies*, involving specific (eventually growing) architectures and/or weight freezing/pruning,
- (ii) *(pseudo) rehearsal strategies*, in which past information is remembered, preferably (to avoid increasing memory consumption) exploiting a generative model, and then replayed in future training,
- (iii) *regularization strategies*, introducing a penalization term into the loss function, which promotes selective consolidation of past information or slows training on new tasks.

As respective pure examples, we can recall here: (i) Progressive Neural Network (PNN, see Rusu et al. (2016)), (ii) Experience Replay for Streaming Learning (ExStream, see Hayes et al. (2019)), and (iii) Learning without Forgetting (LwF, see Li & Hoiem (2016)). For a more complete and detailed overview of the existing CL models and their classification according to the above distinction rules, we refer the reader to Maltoni & Lomonaco (2019), especially to the very informative Venn diagram provided as Figure 1.

This paper aims to show that it is achievable to construct an effective regularization strategy with a penalizing term expressed by the Cramer-Wold distance (introduced by Knop et al. (2020)) between two probability distributions designed to represent current and past information, both defined on a target layer of a neural network that is shared by all models dedicated to solve individual tasks. To memorize the past knowledge, an additional simple generator architecture is learned to retrieve the

network output data from Gaussian noise, which can be considered as a step towards the pseudo-rehearsal model. Following Knop et al. (2020) we call it the Cramer-Wold (or CW) generator. It is worth noting that such strategy preserves the network output distribution while learning a new task, but does not require remembering/replaying (usually high dimensional) previous tasks' datasets. A general description of our approach, involving theoretical (mathematical) background, the reader can find in Section 3.

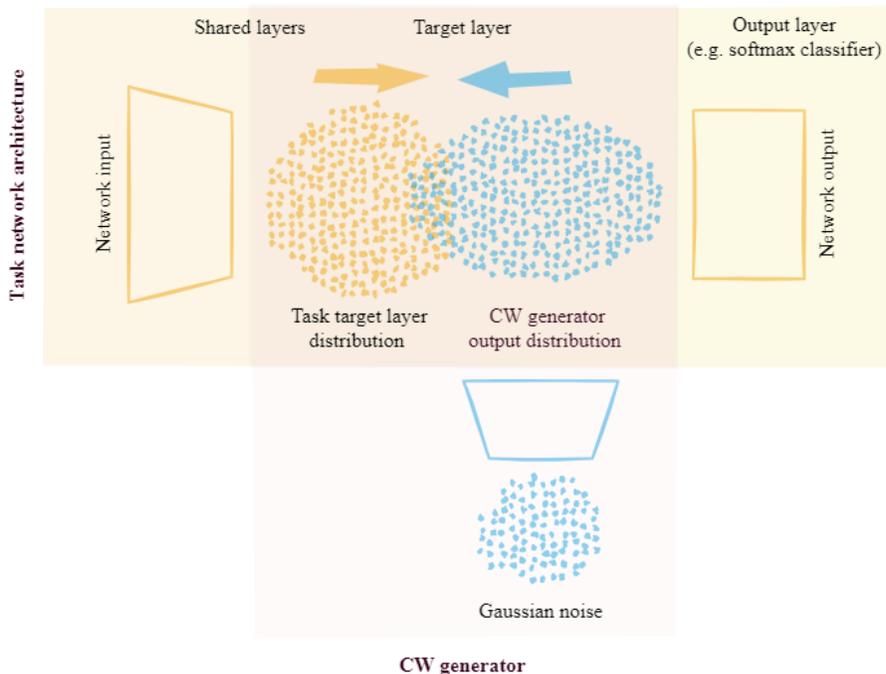


Figure 1: In each step of the proposed algorithm, we regularize our model by forcing closeness of task target layer distribution (representing current information) and CW generator output distribution (representing past information). This applies to both solving individual tasks as well as training the CW generator. Note that our construction does not require extra memory to remember previous tasks' datasets. (The figure was produced using `diagrams.net` software.)

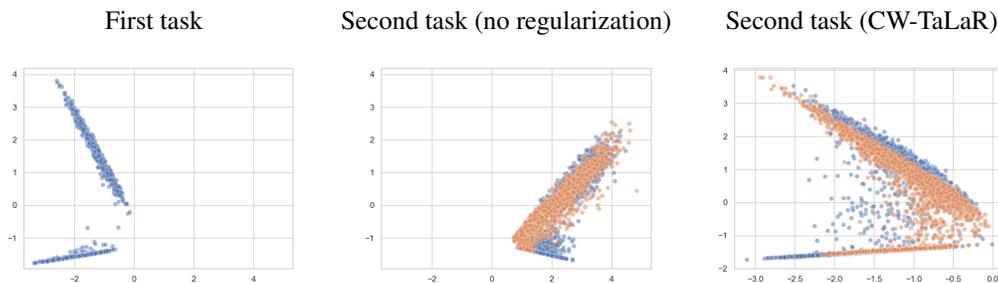


Figure 2: The figures demonstrate how CW-TaLaR works in the real world CL scenario. Blue dots represent output produced on a target layer for samples from the first task while orange dots represent output for samples from the second task. In the first column, we present how output data look like after training the first task. In the second column, we show how output data behave when there is no regularization incorporated. Note that in such a scenario output for the first task samples completely changed their distribution. In the third column, we show how our method prevents output data from diverging from the previous tasks' distribution which eventually improves model continual learning capabilities. The above results were obtained in the Split MNIST experimental setting (see Section 5). Output data dimensionality has been reduced with PCA for the purpose of illustration.

We apply our method (see Section 4), which we call henceforth *Cramer-Wold Target Layer Regularization (CW-TaLaR)* strategy, for different categories of supervised benchmark CL learning scenarios proposed by Yen-Chang Hsu, involving *Incremental Task Learning*, *Incremental Domain Learning* and *Incremental Class Learning* on Split MNIST, Split CIFAR-10 and Permuted MNIST datasets. The conducted experiments show (see Section 5) that the results of the CW-TaLaR method are comparable or better than those obtained by the other existing state-of-the-art regularization-based CL models, including (online) Elastic Weights Consolidation (EWC, see Kirkpatrick et al. (2017)), Synaptic Intelligence (SI, see Zenke et al. (2017)) and Memory Aware Synapses (MAS, see Aljundi et al. (2018)).

A concept of our regularization strategy and an intuitive motivation (involving the real world CL scenario) for the CW-TaLaR model are presented in Figures 1 and 2.

Our contribution can be summarized as follows: (i) we introduce a novel CL strategy (CW-TaLaR), which is based on the Cramer-Wold distance, (ii) we compare the CW-TaLaR method with the other regularization-based models, including EWC, SI and MAS.

2 CRAMER-WOLD DISTANCE AND CW GENERATOR

The key idea behind the Cramer-Wold distance (d_{cw}) between two multi-dimensional distributions, which was firstly defined by Knop et al. (2020), is to consider squared L_2 distance computed across multiple single-dimensional smoothed (with a Gaussian kernel) projections of their density functions. In a practical context, where we deal with two D -dimensional samples $X = (x_i)$ and $Y = (y_i)$, calculation of d_{cw} boils down to the following general formula:

$$d_{\text{cw}}^2(X, Y) = \int_{S^{D-1}} \|\text{sm}_\gamma(v^T X) - \text{sm}_\gamma(v^T Y)\|_2^2 d\sigma_{D-1}(v), \quad (1)$$

where sm_γ means a smoothing function using a Gaussian kernel with the variance γ , and S^{D-1} denotes the unit sphere in \mathbb{R}^D with the normalized surface measure σ_{D-1} . Then, in light of Theorem 1 from Knop et al. (2020) and assuming sufficiently large dimension of data ($D \geq 20$), in all numerical experiments we can use an approximate version of (1), namely:

$$2\sqrt{\pi} d_{\text{cw}}^2(X, Y) \approx \frac{1}{n^2} \sum_{ij} (\gamma_n + \frac{\|x_i - x_j\|^2}{2D-3})^{-\frac{1}{2}} + \frac{1}{n^2} \sum_{ij} (\gamma_n + \frac{\|y_i - y_j\|^2}{2D-3})^{-\frac{1}{2}} - \frac{2}{n} \sum_{ij} (\gamma_n + \frac{\|x_i - y_j\|^2}{2D-3})^{-\frac{1}{2}}, \quad (2)$$

where γ_n is a bandwidth chosen using Silverman’s rule of thumb¹.

Since the Cramer-Wold distance turned out to have a closed-form for spherical Gaussians, it was used by Knop et al. (2020) to construct the Cramer-Wold autoencoder (CWAE). Moreover, the authors of CWAE indicated that it is also possible to apply $d_{\text{cw}}^2(X, G(Z))$ as an objective function for a data generator G (hence called the CW generator), mapping a latent Gaussian noise Z into the data space. We want to emphasize that this simple concept showed promising results when trained on MNIST and Fashion MNIST datasets (see Section 8 of Knop et al. (2020)), which motivated us to consider it as a part of our model.

3 PROPOSED ALGORITHM

We deal with a general CL scenario, in which a sequence of tasks T_1, \dots, T_n , using data samples $(X_1, Y_1) = (x_i^1, y_i^1), \dots, (X_n, Y_n) = (x_i^n, y_i^n)$ drawn from some (unknown) distributions on data spaces $\mathcal{X} \times \mathcal{Y}_1, \dots, \mathcal{X} \times \mathcal{Y}_n$, is given. Let $f_\theta: \mathcal{X} \rightarrow \mathcal{S} \subset \mathbb{R}^D$ be a neural network shared by all architectures that we need to optimize in order to solve all the tasks one by one (i.e. a new task is collected only when the current task is over). Assume that a solution of each task T_j for $j = 1, \dots, n$ may be obtained by minimizing an appropriate objective $\mathcal{I}_j(\theta) = \mathcal{L}_j(f_\theta(X_j), Y_j)$, where $\mathcal{L}_j: \mathcal{S} \times \mathcal{Y}_j \rightarrow \mathbb{R}$ means a respective loss.

In our model, we regularize each cost function $\mathcal{I}_j(\theta)$ for $j = 2, \dots, n$, by tending to add tuned squared Cramer-Wold distance calculated between two distributions for tasks T_{j-1} and T_j , which

¹Specifically, $\gamma_n = \hat{\sigma}(\frac{4}{3n})^{2/5}$, where $\hat{\sigma}$ denotes a standard deviation calculated using joined X and Y samples.

are defined on the target layer \mathcal{S} and represented by the output samples $f_{\theta_{j-1}^*}(X_{j-1})$ and $f_{\theta}(X_j)$, where θ_{j-1}^* denotes the optimal parameter value achieved when T_{j-1} is over. However, since we do not already have access to previous tasks' datasets when we solve task T_j , we render a missing sample $f_{\theta_{j-1}^*}(X_{j-1})$ using the CW generator network $G_{\gamma} : \mathcal{Z} \rightarrow \mathcal{S}$, which is trained independently to generate the output of $f_{\theta_{j-1}^*}$. More precisely, our general algorithm can be presented in the following steps.

Solving T_1 . We solve T_1 by minimizing $\mathcal{I}_1(\theta)$ and obtain $\theta_1^* = \arg \min_{\theta} \mathcal{I}_1(\theta)$.

Training G_{γ} . We train G_{γ} by minimizing $d_{\text{cw}}^2(f_{\theta_1^*}(X_1), G_{\gamma}(Z))$, where $Z = (z_i)$ means a sample from $N(0, I)$ (a Gaussian noise on the latent \mathcal{Z}), and obtain $\gamma_1^* = \arg \min_{\gamma} d_{\text{cw}}^2(f_{\theta_1^*}(X_1), G_{\gamma}(Z))$.

Solving T_2 . We solve T_2 by minimizing

$$\tilde{\mathcal{I}}_2(\theta) = \mathcal{I}_2(\theta) + \lambda d_{\text{cw}}^2(G_{\gamma_1^*}(Z), f_{\theta}(X_2)), \quad (3)$$

where λ is a hyperparameter², and obtain $\theta_2^* = \arg \min_{\theta} \tilde{\mathcal{I}}_2(\theta)$.

Retraining G_{γ} . We retrain G_{γ} by minimizing $d_{\text{cw}}^2(f_{\theta_2^*}(X_2), G_{\gamma}(Z))$ and obtain $\gamma_2^* = \arg \min_{\gamma} d_{\text{cw}}^2(f_{\theta_2^*}(X_2), G_{\gamma}(Z))$.

...

Solving T_j . We solve T_j by minimizing

$$\tilde{\mathcal{I}}_j(\theta) = \mathcal{I}_j(\theta) + \lambda d_{\text{cw}}^2(G_{\gamma_{j-1}^*}(Z), f_{\theta}(X_j)) \quad (4)$$

and obtain $\theta_j^* = \arg \min_{\theta} \tilde{\mathcal{I}}_j(\theta)$.

Retraining G_{γ} . We retrain G_{γ} by minimizing $d_{\text{cw}}^2(f_{\theta_j^*}(X_j), G_{\gamma}(Z))$ and obtain $\gamma_j^* = \arg \min_{\gamma} d_{\text{cw}}^2(f_{\theta_j^*}(X_j), G_{\gamma}(Z))$.

...

Solving T_n . We solve T_n by minimizing

$$\tilde{\mathcal{I}}_n(\theta) = \mathcal{I}_n(\theta) + \lambda d_{\text{cw}}^2(G_{\gamma_{n-1}^*}(Z), f_{\theta}(X_n)) \quad (5)$$

and obtain $\theta_n^* = \arg \min_{\theta} \tilde{\mathcal{I}}_n(\theta)$.

Note that due to flexibility of the above construction and various ways of interpretation of its particular elements, the proposed method is enough general to be applied to many typical supervised and unsupervised learning frameworks. However, since in further study we concentrate on application of our model in solving classification problems, in the next section we describe a more precise version of the proposed algorithm designed for a supervised setting.

4 ADJUSTMENT TO CLASSIFICATION PROBLEMS

Under notation from the previous section, assume that T_1, \dots, T_n is a sequence of classification tasks. In this case data space \mathcal{Y}_j is a collection of labels l_1, \dots, l_{m_j} (that may differ between tasks), and the last shared layer \mathcal{S} (containing the output of f_{θ}) is mapped into a logit layer \mathbb{R}^{m_j} by some single layer network ϕ_{δ}^j . Moreover, each loss \mathcal{L}_j is expressed by the total cross-entropy evaluated over the label sample $Y = (y_i)$ relative to the output of a classifier (a softmax function) applied for a given logit sample $\phi_{\delta}^j(S) = \phi_{\delta}^j(s_i)$, i.e.:

$$\mathcal{L}_j(S, Y) = - \sum_i \sum_k \mathbb{I}_{\{y_i=l_k\}} \log \text{softmax}(\phi_{\delta}^j(s_i)), \quad (6)$$

²The hyperparameter λ (which may be different in each step) aims to find a balance between intransigence and forgetting properties.

where $\mathbb{I}_{\{.\}}$ means a characteristic function.

The above approach covers all three categories of CL scenarios, i.e. Incremental Task Learning (ITL), Incremental Class Learning (ICL) and Incremental Domain Learning (IDL), which were introduced by Yen-Chang Hsu in relation to the differences between tasks, taking into account the marginal distributions of input data and target labels. Hence, we can apply our method to many variants of experimental setup that were proposed in Yen-Chang Hsu. In the next section we describe considered experimental scenarios as well as present the obtained results, comparing them to those provided by the (online) EWC, SI and MAS models.

5 EXPERIMENTS

To evaluate effectiveness of the CW-TaLaR model in a CL setting, we adhere to experimental setups proposed by Yen-Chang Hsu. We compare our strategy with a selection of regularization-based methods, i.e. (online) EWC, SI and MAS, as well as with some vanilla baselines. We use (mentioned above) three classic scenarios: ITL, IDL and ICL, which are applied for a few datasets commonly used in various CL settings: Split MNIST, Split CIFAR-10 and Permuted MNIST.

In ITL the output layer is split into 5 separate heads that serve as binary classifiers (only one active per task), each of which discriminates between 2 different classes. On the other hand, ICL uses a single head that grows accordingly, from 2 to 10 every 2 outputs. In turn, IDL consists of a single head that is shared by all tasks, which serve as a classifier that distinguishes between 2 (for Split MNIST and Split CIFAR-10) or 10 (for Permuted MNIST) different classes.

We repeated each experiment 10 times and reported the average accuracy over all tasks and runs for all applied strategies. (See the following three subsections for the detailed results.) It should be noted that for the CW-TaLaR method we added an additional penalizing term to the objective function (i.e. L_2 norm taken from network output sample on a target layer) that prevents large magnitudes which may occur due to potentially unlimited ReLU values.

5.1 SPLIT MNIST

We summarize the results of all experiments performed on the Split MNIST dataset in Table 1. Note that CW-TaLaR significantly improves performance in single-head CL scenarios (IDL and ICL).

Method	ITL	IDL	ICL
Adam	93.50 ± 3.40	54.48 ± 0.80	19.88 ± 0.02
SGD	94.45 ± 2.24	75.83 ± 0.40	18.70 ± 0.19
Adagrad	97.06 ± 0.53	63.78 ± 3.38	19.55 ± 0.03
L2	95.39 ± 3.62	65.52 ± 2.89	19.44 ± 2.89
EWC	94.18 ± 3.52	55.74 ± 1.25	19.83 ± 0.03
Online EWC	94.18 ± 4.98	58.66 ± 1.37	19.86 ± 0.02
SI	97.80 ± 1.43	65.66 ± 1.49	20.54 ± 0.78
MAS	97.97 ± 1.15	72.24 ± 5.37	20.24 ± 0.95
CW-TaLaR (ours)	97.53 ± 0.53	80.64 ± 1.25	38.65 ± 3.02
Offline (upper bound)	99.59 ± 0.09	98.69 ± 0.13	97.79 ± 0.20

Table 1: Average accuracy (% , higher is better) with standard deviation achieved on validation dataset by each method applied for Split MNIST in three CL scenarios, calculated over 5 tasks and 10 runs (at the end of the training process).

Figure 3 shows that our strategy effectively struggles with catastrophic forgetting, especially in the case of IDL setting, where all other methods fail to overcome this issue since their accuracy fall drastically from time to time, as well as in the case of ICL setting, where significantly wins throughout the entire learning process.

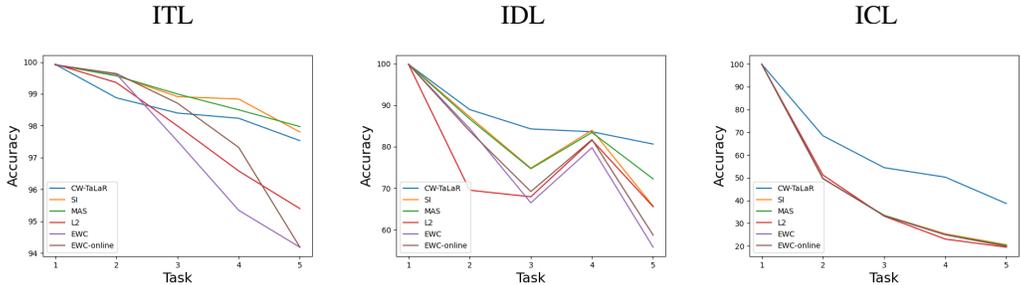


Figure 3: Comparison of average accuracy over currently completed tasks achieved by each method applied for Split MNIST in three CL scenarios, calculated on validation dataset during the training process. (All results were additionally averaged over 10 runs.) Note that CW-TaLaR is superior to overcome forgetting in single-head scenarios.

5.2 SPLIT CIFAR-10

We summarize the results of all experiments performed on the Split CIFAR-10 dataset in Table 2. Note that CW-TaLaR in all scenarios gives scores comparable to SI and MAS, although works slightly worse than (online) EWC. This observation seems to be confirmed by Figure 4, which shows comparison of efficiency of all considered CL strategies in overcoming catastrophic forgetting during the training process. We suppose that such loss of dominance when compared to the results for Split MNIST is due to significant change in applied architecture (here we use CNN instead of MPL).

Method	ITL	IDL	ICL
Adam	72.65 ± 3.39	73.13 ± 1.39	19.31 ± 0.05
SGD	68.93 ± 5.68	73.13 ± 1.35	17.51 ± 2.42
Adagrad	70.90 ± 6.50	72.57 ± 1.54	16.02 ± 0.54
L2	76.30 ± 0.96	73.40 ± 1.71	20.36 ± 2.13
EWC	86.07 ± 4.06	78.16 ± 0.68	19.24 ± 0.11
Online EWC	87.57 ± 3.89	77.46 ± 0.89	19.22 ± 0.12
SI	84.41 ± 2.60	75.28 ± 1.26	19.27 ± 0.10
MAS	82.13 ± 1.80	73.50 ± 1.54	19.30 ± 0.04
CW-TaLaR (ours)	82.09 ± 2.04	74.89 ± 0.61	19.20 ± 0.08
Offline (upper bound)	95.25 ± 0.64	91.35 ± 0.32	84.83 ± 1.60

Table 2: Average accuracy (% , higher is better) with standard deviation, achieved on validation dataset by each method applied for Split CIFAR-10 in three CL scenarios, calculated over 5 tasks and 10 runs (at the end of the training process).

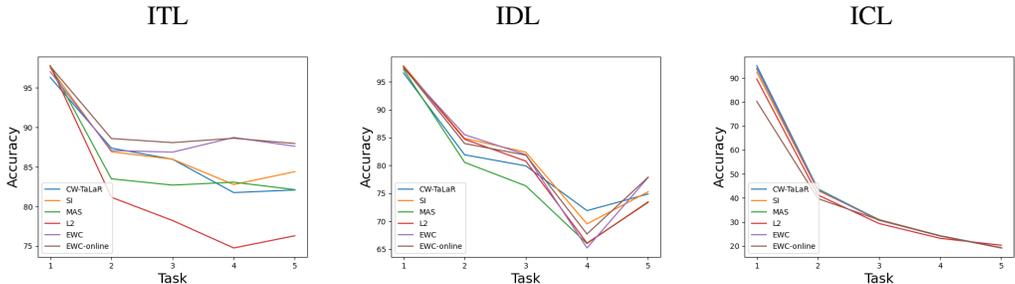


Figure 4: Comparison of average accuracy over currently completed tasks achieved by each method applied for Split CIFAR-10 in three CL scenarios, calculated on validation dataset during the training process. (All results were additionally averaged over 10 runs.) Note that scores for CW-TaLaR are comparable with those for SI and MAS.

5.3 PERMUTED MNIST

We summarize the results of IDL experiments performed on the Permuted MNIST dataset in Table 3. Note that almost all obtained scores (excluding those for EWC) are comparable. However, closer inspection of learning process, which is shown in Figure 5, allows to guess that the CW-TaLaR method is superior in learning new tasks, although may slightly suffer from forgetting in the case when the number of tasks is relatively large.

Method	IDL
Adam	76.19 \pm 2.17
SGD	70.35 \pm 1.53
Adagrad	85.24 \pm 0.75
L2	86.92 \pm 1.68
EWC	77.81 \pm 1.32
Online EWC	87.63 \pm 0.99
SI	83.92 \pm 1.22
MAS	87.82 \pm 1.47
CW-TaLaR (ours)	86.38 \pm 1.08
Offline (upper bound)	97.83 \pm 0.05

Table 3: Average accuracy (% , higher is better) with standard deviation, achieved on validation dataset by each method applied for Permuted MNIST in IDL scenario, calculated over 5 tasks and 10 runs (at the end of the training process).

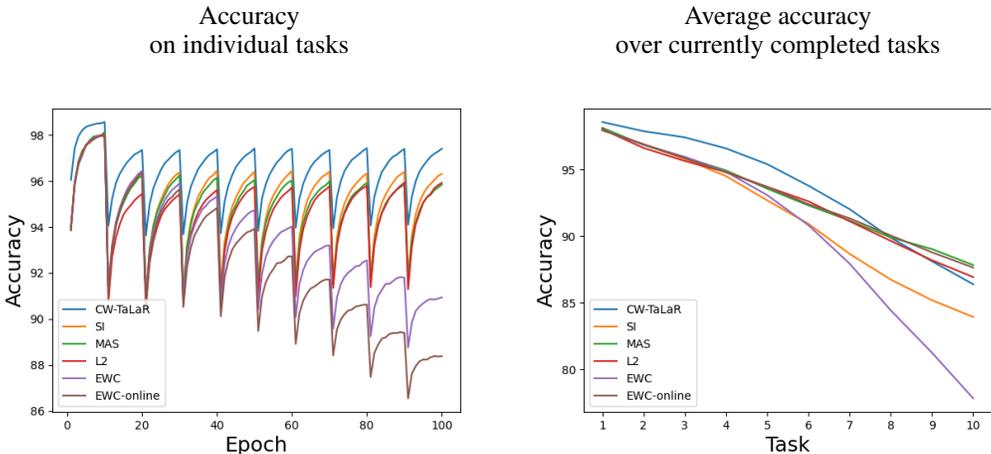


Figure 5: Comparison of average accuracy over currently completed tasks achieved by each method applied for Permuted MNIST in IDL scenario, calculated on validation dataset during the training process. (All results were additionally averaged over 10 runs.) Note that CW-TaLaR is superior in learning new tasks or for relatively short term scenarios.

5.4 ARCHITECTURE

For all models we performed a grid search over hyperparameters. Our classification networks were trained on 128-element mini-batches using ADAM optimizer (see Kingma & Ba (2017)) with a starting learning rate $1e-4$, which decreases to $1e-5$ after the first task. CW generator was also trained on 128-element mini-batches using ADAM optimizer with learning rate set to $1e-4$. We applied ReLU as the activation function and the loss function (in all methods and all scenarios) is a standard cross-entropy for classification frameworks. Below we include architecture details for networks used in various experimental settings.

5.4.1 SPLIT MNIST

Each classification task was trained for 4 epochs. Generator was trained for 10 epochs.

classifier

fully connected networks with ReLU activation functions
layer sizes: $784 \rightarrow 1024 \rightarrow 512 \rightarrow 256 \rightarrow 2/10$ neurons depends on setting
softmax output layer

generator

noise dimension: 8
fully connected networks with ReLU activation functions
layer sizes: $8 \rightarrow 16 \rightarrow 32 \rightarrow 64 \rightarrow 128 \rightarrow 256$
ReLU output layer

5.4.2 PERMUTED MNIST

Each classification task was trained for 10 epochs. Generator was trained for 10 epochs. Architectures used for generator and classifiers were the same as those described for Split MNIST.

5.4.3 SPLIT CIFAR-10

Each classification task was trained for 12 epochs. Generator was trained for 10 epochs.

classifier

all convolutional layers use 3×3 kernel and are followed by batch normalization and ReLU activation
two convolutional layers with 32 channels
convolutional layer with 64 channels
max pooling layer
two convolutional layers with 64 channels
convolutional layer with 128 channels
max pooling layer
two convolutional layers with 128 channels
convolutional layer with 256 channels
max pooling layer
two dense ReLU layers with $256 \times 4 \times 4$ and 256 neurons
dense layer with 2 or 10 outputs depending on scenario
softmax output layer

generator

noise dimension: 64
fully connected networks with ReLU activation functions
layer sizes: $64 \rightarrow 128 \rightarrow 256 \rightarrow 512 \rightarrow 256$
ReLU output layer

6 RELATED WORK

As we indicated in Introduction the CL strategies can be assigned into three fuzzy collections: architectural strategies, (pseudo) rehearsal strategies and regularization strategies.

Such methods as Progressive Neural Networks (PNN, see Rusu et al. (2016)) or CopyWeights with Re-init (CWR, see Lomonaco & Maltoni (2017)) prevent catastrophic forgetting by changing their architectures. CWR trains in each task the same model with new parameters and adds them to the old ones. A similar idea was used in PNN, where in the next task a bigger neural network is trained, but a part of its weights is taken from the previous neural network. However, since our method

cannot be considered as an architectural strategy, we do not feel obligated to compare it to such CL models.

Full rehearsal method (mixing all older examples with new ones) is rather a theoretical solution. An example of effective rehearsal strategy was proposed by Hayes et al. (2019). It creates a special buffer for each label and each buffer aggregates the same number of examples. The old examples make a place for new ones by compression or removal. Up to some extent, CW-TaLaR algorithm can be considered as a pseudo-rehearsal method. However, instead of aggregating old examples, a generative model is trained so we need to only remember all its parameters.

Another approach to solving CL problems (Learning without Forgetting, LwF) was made by Li & Hoiem (2016). LwF relies on controlling a change of output data. It was achieved by regularization but also by adding a small number of nodes in each task. Therefore, it could be partially considered as an architectural strategy.

A pure regularization strategies usually assume that the optimal parameters for a new task can be found in the neighborhood of those obtained for older tasks. For example, this idea is underlined by Elastic Weight Consolidation (EWC, see Kirkpatrick et al. (2017)). EWC algorithm controls the change of parameters only by introducing regularization term in a loss function. The importance of parameters is weighted by the diagonal of the Fisher information matrix. An alternative for this method was presented by Kolouri et al. (2020). There the respective diagonal matrix in regularization term originated from the Cramer distance between output densities for old and new tasks. Another important modification of EWC is Synaptic Intelligence (SI, see Zenke et al. (2017)). The novel of this method lies in detecting the change of parameters in the previous tasks and weighting a regularization term by their importance.

For the other significant CL models we refer the reader to Lomonaco & Maltoni (2017); Hinton et al. (2015); Lopez-Paz & Ranzato (2017); Rebuffi et al. (2017); Parisi et al. (2018); Mishra & Narayanan (2021); Nguyen et al. (2018); van de Ven & Tolias (2019); Liu et al. (2018); Hadsell et al. (2020).

7 CONCLUSION

In this paper, we presented the new CL learning strategy (CW-TaLaR) based on regularization involving the Cramer-Wold distance between target layer distributions representing current and past information, which does not suffer from growing memory consumption. We applied our method for a few supervised CL scenarios using various versions of experimental setup proposed in Yen-Chang Hsu.

Our experiments show that CW-TaLaR gives superior average accuracy scores (when compared to EWC, SI and MAS methods) for IDL and ICL scenarios on Split MNIST dataset, and very competitive results for the other considered settings (including ITL) on Split MNIST, Split CIFAR-10 and Permuted MNIST datasets.

In conclusion, the proposed method works very well in general but seems to be especially efficient for CL scenarios with a single-head classifiers (i.e IDL and ICL). However, the closer inspection of the results for Permuted MNIST and Split CIFAR-10 suggests that it may slightly suffer from forgetting, when either the number of tasks is relatively large or the applied architecture consists of convolutional (rather than fully connected) layers, although in these cases the average scores still remain competitive with respect to the other state-of-the-art CL models.

REFERENCES

- Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- Arslan Chaudhry, Puneet K. Dokania, Thalaiyasingam Ajanthan, and Philip H. S. Torr. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- Raia Hadsell, Dushyant Rao, Andrei A. Rusu, and Razvan Pascanu. Embracing change: Continual learning in deep neural networks. *Trends in Cognitive Sciences*, 24:1028–1040, 2020.

- Tyler L. Hayes, Nathan D. Cahill, and Christopher Kanan. Memory efficient experience replay for streaming learning. In *2019 International Conference on Robotics and Automation (ICRA)*, pp. 9769–9776, 2019.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network, 2015.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114:3521–3526, 2017.
- Szymon Knop, Przemysław Spurek, Jacek Tabor, Igor Podolak, Marcin Mazur, and Stanisław Jastrzebski. Cramer-wold auto-encoder. *Journal of Machine Learning Research*, 21:1–28, 2020.
- Soheil Kolouri, Nicholas A. Ketz, Andrea Soltoggio, and Praveen K. Pilly. Sliced cramer synaptic consolidation for preserving deeply learned representations. In *International Conference on Learning Representations*, 2020.
- Zhizhong Li and Derek Hoiem. Learning without forgetting. In *Computer Vision – ECCV 2016*, pp. 614–629, Cham, 2016. Springer International Publishing.
- Xialei Liu, Marc Masana, Luis Herranz, Joost Van de Weijer, Antonio M. López, and Andrew D. Bagdanov. Rotate your networks: Better weight consolidation and less catastrophic forgetting. In *2018 24th International Conference on Pattern Recognition (ICPR)*, pp. 2262–2268, 2018.
- Vincenzo Lomonaco and Davide Maltoni. Core50: a new dataset and benchmark for continuous object recognition. In Sergey Levine, Vincent Vanhoucke, and Ken Goldberg (eds.), *Proceedings of the 1st Annual Conference on Robot Learning*, volume 78 of *Proceedings of Machine Learning Research*, pp. 17–26. PMLR, 2017.
- David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. NIPS’17, pp. 6470–6479, Red Hook, NY, USA, 2017. Curran Associates Inc.
- Davide Maltoni and Vincenzo Lomonaco. Continuous learning in single-incremental-task scenarios. *Neural Networks*, 116:56–73, 2019.
- Michael McCloskey and Neal J. Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. volume 24 of *Psychology of Learning and Motivation*, pp. 109–165. Academic Press, 1989.
- Poonam Mishra and Rishikesh Narayanan. Stable continual learning through structured multiscale plasticity manifolds. *Current Opinion in Neurobiology*, 70:51–63, 2021.
- Cuong V. Nguyen, Yingzhen Li, Thang D. Bui, and Richard E. Turner. Variational continual learning, 2018.
- German I. Parisi, Jun Tani, Cornelius Weber, and Stefan Wermter. Lifelong learning of spatiotemporal representations with dual-memory recurrent self-organization. *Frontiers in Neurobotics*, 12, 2018.
- Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, G. Sperl, and Christoph H. Lampert. icarl: Incremental classifier and representation learning. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5533–5542, 2017.
- Andrei A. Rusu, Neil C. Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks, 2016.
- Gido M. van de Ven and Andreas S. Tolias. Three scenarios for continual learning, 2019.
- Anita Ramasamy Zsolt Kira Yen-Chang Hsu, Yen-Cheng Liu. Re-evaluating continual learning scenarios: A categorization and case for strong baselines.

Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 3987–3995. PMLR, 2017.