

A BATCHED COMPUTATIONS

In Algorithm 2 we report the same operations of Algorithm 1 when a mini-batch of size b is provided by the stream at each time instant. The case of batched computations is pretty straightforward, with a major difference in the key scrambling routine. As a matter of fact, scrambling would require to serialize the processing of the element in the mini-batch, that is not a desirable feature. For this reason, we restrict to 1 the number of weak keys that can be potentially replaced for each mini-batch, selecting the one associated to the batch element that is returning the smallest similarity score with respect to K .

Algorithm 2 Learning with a Continual Memory Neuron when a batch of b samples is provided at each time instant by the stream. Notice that X is the mini-batch matrix ($b \times d$), function α is intended to compute attention scores for each element of the batch, \dagger and o are now arrays of length b . Scrambling involves up to 1 key for each mini-batch (function SCRAMBLEONE).

Require: Stream \mathcal{S} ; generic loss function $\text{loss}(\dots)$, learning rates ρ, β ; $K \leftarrow \text{rand}$, $M \leftarrow \text{rand}$, $\mu \leftarrow 0$'s, $\eta \leftarrow \tau^\eta$'s.

while true do

- $X \leftarrow \text{next_neuron_inputs}(\mathcal{S})$ ▷ X is a $b \times d$ matrix
- $A, S \leftarrow \alpha(X, K, \delta)$ ▷ A and S are $b \times m$ matrices
- $\dagger_h \leftarrow \arg \max_{j \in \{1, \dots, m\}} \{A_{hj}\}$, $h = 1, \dots, b$ ▷ Indices of the winning keys
- $K, M, \dagger, \mu \leftarrow \text{SCRAMBLEONE}(X, K, M, S, \dagger, \mu, \eta)$ ▷ Possibly replace a weak key
- $K_{\dagger_h} \leftarrow K_{\dagger_h} + \beta \nabla \text{sim}(\psi(X_h), K_{\dagger_h})$, $h = 1, \dots, b$, ▷ Upd. winning keys, Eq. 6
- $A \leftarrow \alpha(X, K, \delta)$ ▷ Refresh attention (from scratch)
- $\mu_{\dagger_h} \leftarrow \mu_{\dagger_h} + 1$, $\eta_{\dagger_h} \leftarrow 0$, $h = 1, \dots, b$ ▷ Increase winning keys usages, reset ages
- $\eta = \eta + b$ ▷ Increase all ages
- $W \leftarrow AW$ ▷ Generate weights, Eq. 3, W is a $b \times d$ matrix
- $o_h \leftarrow W_h X'_h$, $h = 1, \dots, b$ ▷ Compute output, Eq. 2
- $M_{\dagger} \leftarrow M_{\dagger} - \rho \nabla_{M_{\dagger}} \text{loss}(o)$ ▷ Upd. winning memory unit

end while

function SCRAMBLEONE($X, K, M, S, \dagger, \mu, \eta$)

- $\mathcal{W} \leftarrow \{z: \mu_z < \tau^\mu \wedge \eta_z \geq \tau^\eta\}$ ▷ Set of weak keys (if any)
- $k \leftarrow \arg \min_{h \in \{1, \dots, b\}} \{S_{h\dagger_h}\}$ ▷ Idx of the sample with lowest similarity to its winning key
- if** $S_{k\dagger_k} < \tau^\alpha \wedge \mathcal{W} \neq \emptyset$ **then**
 - $j \leftarrow \arg \max_{z \in \mathcal{W}} \{\eta_z\}$ ▷ If the current match is loose
 - $K_j \leftarrow \psi(X_k)$ ▷ The weakest key is the oldest
 - $M_j \leftarrow M_{\dagger}$ ▷ Replace weakest key with data sample
 - $\mu_j \leftarrow 0$ ▷ Warm-start for replaced memory unit
 - $\dagger_k \leftarrow j$ ▷ Reset usage counter
 - $\dagger_k \leftarrow j$ ▷ Update winning key index
- end if**
- return** K, M, \dagger, μ

end function

B METRICS

Following the notation of Section 4, we are given a data stream \mathcal{S} partitioned into non-overlapping time intervals, and we indicate with t_j the last time instant of the j -th interval I_j , with $t_N = T$. In the following description, we assume class indices to be ordered with respect to the time in which they become available, to keep the notation simple. We indicate with $p(x|I_i)$ the data distribution in the i -th interval, with θ_j is the model developed up to t_j (being it a CMN-based network or the network of another competitor), while \mathcal{D}_i is an held out test set with data sampled from $p(x|I_i)$. Then, we indicate with

$$\text{acc}_i^{\theta_j} = \text{accuracy}(\mathcal{D}_i, \theta_j)$$

the accuracy on data sampled from $p(x|I_i)$ computed using the model parameters at t_j , i.e., θ_j . We collect the following matrix of accuracies during the learning procedure,

$$\begin{bmatrix} \text{Model/Test Data} & \mathcal{D}_1 & \mathcal{D}_2 & \dots & \mathcal{D}_j & \dots & \mathcal{D}_N \\ \theta_1 & \text{acc}_1^{\theta_1} & \text{acc}_2^{\theta_1} & \dots & \text{acc}_j^{\theta_1} & \dots & \text{acc}_N^{\theta_1} \\ \theta_2 & \text{acc}_1^{\theta_2} & \text{acc}_2^{\theta_2} & \dots & \text{acc}_j^{\theta_2} & \dots & \text{acc}_N^{\theta_2} \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ \theta_j & \text{acc}_1^{\theta_j} & \text{acc}_2^{\theta_j} & \dots & \text{acc}_j^{\theta_j} & \dots & \text{acc}_N^{\theta_j} \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ \theta_N & \text{acc}_1^{\theta_N} & \text{acc}_2^{\theta_N} & \dots & \text{acc}_j^{\theta_N} & \dots & \text{acc}_N^{\theta_N} \end{bmatrix}$$

that we indicate as continual confusion matrix (CCM, being CCM_j the matrix up to t_j), and we exploit it to compute the following measures. Notice that it is a square matrix.

- The average accuracy at t_z is defined as the average of the z -th row of the CCM, up to the z -th column (included),

$$\text{avg_accuracy}(\text{CCM}_z) = \frac{1}{z} \sum_{i=1}^z \text{acc}_i^{\theta_z},$$

and we commonly measure the average accuracy (ACC of Section 4) at the end of training, $t_z = t_N$.

- The average forgetting at t_z can be defined as

$$\text{avg_forgetting}(\text{CCM}_z) = \frac{1}{z-1} \sum_{i=1}^{z-1} \left(\text{acc}_i^* - \text{acc}_i^{\theta_z} \right),$$

where acc_i^* is the best accuracy obtained on data \mathcal{D}_i so far, i.e., $\max_{\theta_k \in \{\theta_1, \dots, \theta_{z-1}\}} \text{acc}_i^{\theta_k}$ (maximum of the i -th column up to row $z-1$). We commonly measure the average forgetting (FORG of Section 4) at the end of training, $t_z = t_N$.

- Forward transfer measures how learning at checkpoint t_z (positively) influences predictions on data introduced in future intervals,

$$\text{forward}(\text{CCM}_z) = \frac{2}{z(z-1)} \sum_{i=1}^{z-1} \sum_{j=i+1}^z \text{acc}_i^{\theta_j},$$

i.e., it is the average of the upper-triangular portion of the CCM (excluding the diagonal). Similarly to the previous cases, we commonly measure it at $t_z = t_N$, yielding FWD of Section 4.

C COMPUTATIONAL COST

A classic neuron model in a neural network requires u products to compute the output score, being u the size of the input, Eq. 1. We compare this cost in terms of the operations performed by CMNs, still using products as a reference. Computing the output of a CMN involves three main operations: (1) evaluating $\psi(x)$, whose cost is $C(\psi)$, that transforms the neuron input in a customizable manner, being \tilde{u} the size of the ψ -output; (2) computing the attention scores by α , Eq. 4, with cost $m\tilde{u}$ plus the cost of the softmax^δ operation, that is δ ; (3) blending memories, δu products, due to the sparsity of the attention scores; (4) computing the usual output function as in a classic neuron, u products. In total, we have $C(\psi) + m\tilde{u} + \delta + \delta u + u$. The cost of a layer of n classic neurons trivially becomes un , while the cost of a layer of CMNs that share the same K is

$$C(\psi) + m\tilde{u} + \delta + \delta un + un, \quad (7)$$

where only the last two terms depends on n , since the first three ones are about operations that are performed only once, being K shared. In order to reasonable relate the cost of classic and CM neurons, some basic considerations must be introduced. First of all, the cost $C(\psi)$ is expected to be way smaller than the cost of the whole layer. For example, when ψ is just limited to the

L_2 normalization of x . Moreover, depending on the considered problem, there could be room for designing ψ such that \tilde{u} is smaller than u . Of course, this does not always hold. It is reasonable to assume the term δ in Eq. 7 to be way smaller than the other ones (being it a strong sparsity index, always $< m$), thus we discard it. As a result, we can compute the ratio R between the cost of a CMN layer and the corresponding classic layer,

$$R_C = \frac{m\tilde{u} + (\delta + 1)un}{un} = \frac{m\tilde{u}}{nu} + \delta + 1. \quad (8)$$

In case of multi-layer nets, with ℓ layers, we have

$$R_C = \frac{\sum_{i=1}^{\ell} m_{(i)}\tilde{u}_{(i)} + (\delta_{(i)} + 1)d_{(i)}n_{(i)}}{\sum_{i=1}^{\ell} d_{(i)}n_{(i)}}, \quad (9)$$

being i the layer index. In terms of memory consumption, a layer of n CMNs stores matrix K and n matrices of memory units (M), that is a total of $m\tilde{u} + mun$ floating point numbers, while in a classic layer only the weight matrix is stored (un floating point values). The ratio R_M for ℓ layer is then,

$$R_M = \frac{\sum_{i=1}^{\ell} m_{(i)}\tilde{u}_{(i)} + m_{(i)}u_{(i)}n_{(i)}}{u_{(i)}n_{(i)}} \quad (10)$$

while the additional memory usage introduced by ℓ CMN layers is

$$U_M = m_{(i)}\tilde{u}_{(i)} + (m_{(i)} - 1)u_{(i)}n_{(i)}. \quad (11)$$

A candidate way to compare CMN-based net with models that replay data from memory buffers, is to use the exact same network architecture, using classic neurons in place of CMNs. Then, we allow replay-based methods to sample $R_C - 1$ examples from the buffer at each time step. In fact, these buffer-based models make a prediction on a mini-batch of buffer data in addition to the currently streamed sample, according to the continual online learning setting experimented in this paper. Of course, when comparing with models that have more layers than the CMN-net, it is harder to keep a perfect balance in term of computational cost, so we allowed competitors to have a cost that is slightly larger than the one of the CMN-net, making the comparison more challenging. Moreover, we recall that the replay-based methods learn by exploiting the label-related information they store on the replay-buffer, while no-label-information is stored by the CMN-net (this the comparison becomes unfair when using very large replay buffers).

D HYPER-PARAMETERS

We evaluated multiple combinations of values for the main hyper-parameters of CMNs and competitors, that we summarize in the following, in addition to the already described parameter values of the main paper. In the case of CMN-based nets, in MODES and MOONS, we selected $m = 8$ memory units with $\delta = 2$, while we tested $\beta \in \{10^{-4}, 10^{-3}, 10^{-2}, 1\}$, $\tau_{\mu} \in \{50, 200\}$, $\tau_{\eta} \in \{50, 200\}$, $\tau_{\alpha} \in \{0.85, 0.95\}$, $\gamma \in \{1, 5, 25\}$. In NS-IMAGENET we considered $m \in \{10, 25, 50, 100\}$, $\delta \in \{2, 5\}$, $\beta \in \{10^{-3}, 10^{-2}\}$, $\tau_{\mu} \in \{50, 500, 5000\}$, $\tau_{\eta} \in \{50, 500, 5000\}$, $\tau_{\alpha} \in \{0.7, 0.85, 0.95\}$, $\gamma \in \{1, 5, 25\}$. The hidden layer size has been evaluated in $h \in \{5, 25\}$ for 2D data, while in $h \in \{50, 100\}$ for NS-IMAGENET. In all the models, we considered a learning rate $\rho \in \{10^{-4}, 10^{-3}, 10^{-2}, 1\}$, and trained with fixed-step-size gradient descent. We also evaluated the case of Adam, which yielded lower results on average. Indeed, adopting optimizers with memory such as Adam may be tricky: at every step, the model might select a different set of weights to be updated, making the statistics of the optimizer invalid. We leave the investigation about the effect of such optimizers for future work, restricting our analysis to memoryless optimizers, which do not suffer from this issue. We also considered a weight decay factor `DECAY` for the optimizers $\in \{10^{-4}, 10^{-3}, 0\}$. We trained GDumb for 10 epochs on the buffer data. Other minor internal parameters of the competitors were set to the values suggested in the respective papers. The results reported in the main paper are averaged over three runs with different seeds in $\{1234, 123456, 12346578\}$. For all the experiments in this work we used PyTorch, running on a Linux machine–NVIDIA GeForce RTX 3090 GPU (24 GB).

D.1 OPTIMAL HYPER-PARAMETERS

We report in Table 2 the best selected hyperparameters for the CMN model in all the considered datasets and settings described in the main paper.

Table 2: Optimal parameters. The best selected hyperparameters for the proposed CMN model, drawn from the grids described in the text, for the datasets. See the code for further details.

| Parameters | MODES | | | MOONS | | | NS-IMAGENET |
|---------------|-----------|-----------|-----------|-----------|-----------|-----------|-------------|
| | CI | CDI | CDID | CI | CDI | CDID | - |
| δ | 2 | 2 | 2 | 2 | 2 | 2 | 5 |
| β | 10^{-2} | 10^{-2} | 10^{-2} | 10^{-2} | 10^{-2} | 10^{-1} | 10^{-3} |
| ρ | 10^{-2} | 10^{-1} | 10^{-2} | 10^{-1} | 10^{-2} | 10^{-2} | 10^{-4} |
| γ | 25 | 25 | 25 | 25 | 25 | 5 | 1 |
| m | 8 | 8 | 8 | 8 | 8 | 8 | 100 |
| τ_α | 0.95 | 0.95 | 0.95 | 0.95 | 0.95 | 0.85 | 0.7 |
| τ_η | 50 | 50 | 50 | 50 | 200 | 200 | 5000 |
| τ_μ | 50 | 50 | 50 | 50 | 50 | 50 | 500 |
| DECAY | 0. | 0. | 0. | 0. | 0. | 10^{-3} | 10^{-3} |




Figure 6: MODES and MOONS data, **best** test accuracy (reference only) and std in the three setting we analyzed (CI, CDI, CDID).

E NETWORKS OF CMNS

We report in Figure 7 a more detailed view of the process of input projection, memory blending and WTA update occurring in each CMN. We now briefly discuss the specific case of convolutional layers and on networks with multiple stacked CMN-based layers.

In a convolutional layer, each spatial coordinate is associated with a neuron⁴ whose input is differently shifted with respect to every other neuron of the layer. At a first glance, this makes less obvious how to share keys (if need) among the neurons of the layer. However, neurons are still expected to share the same weights/filter, thus the key-attention scores should be the same for all of them, in order to coherently blend memories and generate the same filter. In this case, it is convenient to go back to the standard definition of convolution operation, where the whole input map of the layer is one of its operands, thus $\psi(x)$ is a projection of the whole input of the layer, and it is the same for all the neurons. As a result, the attention scores are computed only once per layer, independently of the resolution of the input map.

Layers of CMNs can be stacked into multi-layer networks, as usual. However, while the input of the first layer is not affected by CMN dynamics, the input of any other layer comes from the CMNs of the layer below. The proposed WTA key update scheme is not gradient-based, so that we also avoid gradients to propagate through the key-matching process, i.e., we consider $\hat{w}(x, K, M)$ of Eq. 2 to not depend on x for gradient computation purposes. As a result, gradients flow from the

⁴We consider a single filter/output-feature-map in this description, for simplicity. For the same reason, and without any loss of generality, we describe infinitely supported filters.

output layer down to the output of the layer below, as usual, and not through \hat{w} . Another intuition that we followed is that CMNs belonging to the lower layers in a deep convolutional network should have a smaller number of memory units than CMNs on the top layers, since lower-level features are likely to be more shared across inputs than the higher-level ones, where the semantics emerge (e.g., edge-like filters in the lowest levels vs. objects/object-parts related filters in the highest ones). Of course, in non-convolutional nets this is harder to say in advance, but we still follow the same intuition motivated by the need of reducing the variability in the outputs of the lower layers, to favor stability in the learning process of the upper layers. Investigating the interaction among layer requires specific studies that go beyond the scope of this paper.

F ADDITIONAL EXPERIMENTAL RESULTS

In Fig. 6 we report the upper-bound of the results presented in the main paper, obtained by selecting the best-performing model on the test data. Of course, these results are only intended to be a reference to understand the maximum performance each model could achieve, and not a way to make comparisons across different approaches. We remark that this is different from what we did in the main paper (Fig. 3), where we cross-validated the hyper-parameter values on the validation part of the stream and evaluated performance on the out-of-sample test sets.

Comparing Fig. 6 with Fig. 3, we notice that the CMN-based net is actually able to reach similar performance, thus being able to make the most out of the validation procedure. Since the validation set is limited to the first part of the streamed data, this results is very promising in terms of what can be achieved when working on longer streams with limited time-span dedicated to the validation of the parameters.




Figure 7: Larger instance of the contents of Fig. 1, for better readability.