CONTINUAL COMPETITIVE MEMORY

Alexander Ororbia

Rochester Institute of Technology
Rochester, NY, USA
ago@cs.rit.edu

ABSTRACT

In this article, we propose a novel form of unsupervised learning that we call continual competitive memory (CCM) as well as a simple framework to unify related neural models that operate under the principles of competition. The resulting neural system, which takes inspiration from adaptive resonance theory, is shown to offer a rather simple yet effective approach for combating catastrophic forgetting in continual classification problems. We compare our approach to several other forms of competitive learning and find that: 1) competitive learning, in general, offers a promising pathway towards acquiring sparse representations that reduce neural cross-talk, and, 2) our proposed variant, the CCM, which is designed with task streams in mind, is needed to prevent the overriding of old information. CCM yields promising results on continual learning benchmarks including Split MNIST and Split NotMNIST.

Keywords lifelong learning, competitive learning, memory systems, unsupervised self-organization

1 Introduction

Lifelong machine learning, otherwise known as continual or never-ending learning [40, 27], stands as one of the greatest challenges facing statistical learning research, especially for models based on artificial neural networks (ANNs). In this problem context, the goal of an agent is much harder than just learning one single prediction task from one single dataset; it must, instead, learn multiple, different prediction tasks *across* several datasets, much as human agents do. In order to do so, the agent must be able to aggregate and transfer its knowledge as new pattern vectors from new problems are encountered. ANNs particularly struggle to continually learn due to the well-known fact that they tend to *catastrophically forget* [25, 32, 8, 21], or rather, they completely erase the knowledge acquired from earlier encountered tasks when processing samples from new tasks. In essence, the problem of catastrophic forgetting (or interference) becomes prominent when the data being presented to an ANN is no longer independently and identically distributed, a fundamental assumption that drives most modern-day machine learning machinery, particularly systems based on deep ANNs. This kind of forgetting not only imposes practical issues but also impedes the use of ANNs as potential models of mind.

While a great deal of research, including both classical and modern efforts, have attempted to address the problem of catastrophic forgetting (or interference), introducing regularization frameworks [18], new mechanisms that allow networks to grow/expand or prune their internal synaptic structures [38, 30], or forms of pattern rehearsal [34, 35, 33], most if not all of these efforts make the use of external human guidance in the form of task-descriptors, or special vectors that indicate the differences between task datasets. Task descriptors themselves are not necessarily neuro-cognitively unrealistic, given that when asked to solve high-level problems, humans are often given some indicator of the task to be completed (either through audio-visual cues, environmental context, or verbal communication). However, for the relatively low-level pattern recognition problems that deep neural networks are often evaluated on, e.g., object recognition, image categorization, generative modeling of low-level features like pixels, it seems rather unlikely that systems of real neurons would be provided, from some external arbiter, such explicit information in order to identify the differences between such tasks.

In this work, we tackle the far greater challenge imposed by continual learning – learning *without* external task descriptors at **both** training and test time phases. Some of the few, more recent efforts that have attempted to tackle such a problem [1] have labeled this scenario as “selfless” or task-free (sequential) learning, motivating the need for

agent systems that learn without task boundary information. Crafting continually learning agents that can would have strong implications for applications ranging from space exploration robotics to autonomous vehicles.

Under the problem context above, our contributions in this article are as follows:

- We craft a simple modeling framework – the neural competitive learning (NCL) framework – for unifying neural models that learn through competitive (Hebbian) learning and examine several important ones in the online streaming setting, such as classical adaptive resonance theory.
- We propose a novel, unsupervised learning system, i.e., *continual competitive memory*, that builds on the key ideas of competitive learning systems and is generalized for the purpose of continually learning from data streams without task descriptors.
- We measure the amount of forgetting in our CCM model and learning process using several metrics that test knowledge retention as well as class conformity across several low-level pattern classification tasks comparing to important NCL baselines as well as a backprop-based neural model.

Our experimental results indicate that the proposed CCM system not only outperforms the baselines tested and examined but also yields, when it is allowed to drive a feedforward neural predictor, competitive performance on two key benchmarks, Split MNIST and Split NotMNIST. Our general NCL framework as well as our proposed CCM neural model offer a foundation upon which future research in both growing and fixed-capacity neural systems for task-free continual learning may be built on top of.

2 The Neural Competitive Learning Framework

We start by first defining the variant of the continual learning problem this work aims to tackle. We then will present a general framework that serves as a modern-day generalization of the classical principles put forth in [37], which we call the neural competitive learning (NCL) framework, and unifies the types of neural models that we will design for the problem setting. The NCL framework will also be used to distinguish the key components required by adaptive artificial neural systems that could be classified as “competitive”.

2.1 Desiderata for Streaming Lifelong Learning Systems

The form of continual learning that is of interest to this study, and simultaneously one of the problem’s hardest forms, satisfies the following criteria:

1. **Data is processed sequentially online:** the agent receives samples, one at a time (or, at best, in small mini-batches), from a potentially infinite stream of data. Once a datum has been seen, it will not be stored and is not directly available ever again.
2. **Task descriptors are not provided:** the agent is not provided with any supervised signal from the experimenter as to what task it is operating on, not even an encoded integer identifier (as is common in many continual learning setups). This condition holds at both training and testing times.
3. **Raw sensory data is not be stored:** the agent should not physically store data patterns (as in the case of replay buffers common in rehearsal-based approaches) or, at worst, store a temporary, negligible quantity at any instant (to simulate working memory). This feasibility requirement ensures that the agent’s memory footprint does not grow uncontrollably with respect to the size of the stream (which could be infinite).
4. **Patterns from the stream might or not come with labels:** the agent should not rely on labels to self-organize as there is no guarantee that the stream will provide labels for each and every single data point it is to produce. Nevertheless, even if the agent is to use labels, the agent is still expected to learn and operate even when parts of the stream are unsupervised.

The NCL models that we examine, develop, or design in the rest of the paper satisfy all four of the above criterion. With respect to [42], the above setup means that NCL systems handle the class-incremental learning (Class-IL) scenario, where an agent must solve each task seen so far as well as infer which task it is being confronted with.

2.2 Preliminaries

Notation: In this paper, $\otimes$ indicates a Hadamard product while $\cdot$ indicates a matrix/vector multiplication (or dot product if the two objects it is applied to are vectors of the exact same shape) and $(\mathbf{v})^T$ denotes the transpose. $\|\mathbf{v}\|_p$ represents a p -norm (also used to represent a distance function, as in the case when the difference between two vectors is provided as the argument) – $p = 1$ chooses the Manhattan (L1) distance and $p = 2$ selects the Euclidean (L2) distance. $|\Theta|$ denotes the cardinality of the set Θ . $r(\mathbf{v})$ and $r(\mathbf{M})$ return the number of rows in vector $\mathbf{v}$ and in matrix $\mathbf{M}$, respectively, while $c(\mathbf{v})$ and $c(\mathbf{M})$ return the number of columns in $\mathbf{v}$ and $\mathbf{M}$, respectively.

$\mathbf{M}[:, i]$ is the slice operator, meant to extract the i th column vector of matrix $\mathbf{M}$ – note that $\mathbf{M}[i, :]$ means extract the i th row. $\mathbf{M}[j, i]$ means extract the scalar element/value in the position (j, i) in $\mathbf{M}$. In this paper, we will sometimes refer to a higher-order tensor (usually three-dimensional/3D) and retrieving the i th 2D slice is denoted as $\mathbf{M}[i, :, :]$.

$\mathbf{z} = \text{Enc}(w, d_z)$, the “one-hot encoding” function, means convert a list of integers w into a binary vector $\mathbf{z}$, where each value in $\mathbf{z}$ is one at each integer/index inside of w and zero elsewhere. Note that $\arg \min_i \mathbf{v}$ returns the index i of $\mathbf{v}$ with the minimum value and, conversely, $\arg \max_i \mathbf{v}$ returns the index of $\mathbf{v}$ with maximum value.

We also introduce the following useful operators/macros. Given column vector $\mathbf{v} \in \mathcal{R}^{r \times 1}$ and row vector $\mathbf{u} \in \mathcal{R}^{1 \times c}$, we define the following:

- $\mathbf{o} = \text{BC}_c(\mathbf{v}) = \mathbf{v} \cdot \mathbf{1}_c$, where $\mathbf{1}_c = \{1\}^{1 \times c}$. The result is a matrix $\mathbf{o} \in \mathcal{R}^{r \times c}$ with $\mathbf{v}$ copied into each column ($\mathbf{v}$ is replicated column-wise c times),
- $\mathbf{o} = \text{BR}_r(\mathbf{u}) = \mathbf{1}_r \cdot \mathbf{u}$, where $\mathbf{1}_r = \{1\}^{r \times 1}$. The result is a matrix $\mathbf{o} \in \mathcal{R}^{r \times c}$ with $\mathbf{u}$ copied into each row ($\mathbf{u}$ is replicated row-wise r times).

Note that the above two operators are typically implemented internally in some programming languages as broadcasting/replication.

Problem Setting: Any of the neural systems we will investigate in this study are to be applied to streams of pattern vectors following the form $\{(\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2), \dots, (\mathbf{x}_N, \mathbf{y}_N)\}$ where any single input vector $\mathbf{x}_j$ contains d_x dimensions, i.e., $\mathbf{x}_j \in \mathcal{R}^{d_x \times 1}$, and any single label vector contains d_y dimensions, i.e., $\mathbf{y}_j \in \{0, 1\}^{d_y \times 1}$ (typically $\mathbf{y}_j$ is the one-hot encoding of a class index y_j). While N does indicate that the data stream is finite and terminates after N steps, it is more likely in real-world problem settings that the stream would be infinite, i.e., $N = \infty$. Each data point could come from one of T possible tasks/datasets (in the experiments presented later, we will investigate benchmarks with $T = 5$ tasks).

Note that the neural models we consider and design in this work are unsupervised in nature, i.e., they do not require or use class labels to conduct learning and/or inference (whereas models based on learning vector quantization [20] do require labels), easily satisfying the fifth criterion of Section 2.1. Nonetheless, later we will introduce simple extensions that will allow an NCL model or our proposed CCM to conduct classification given that the benchmarks we evaluated on contain ground truth labels. Such extensions will allow us to evaluate model quality and generalization in greater detail.

2.3 General Criterion and Organization

Our framework – the Neural Competitive Learning (NCL) framework – breaks down competitive neural systems into two key components: 1) a best matching unit function, and 2) a synaptic weight adjustment scheme. These are the two essential pieces that must be designed to create a minimally functional NCL model.

Our motivation behind investigating and generalizing competitive learning to the continual learning problem space is ultimately to reduce excessive cross-talk [45] (which refers to the information capacity of an ANN) – catastrophic interference is effectively cross-talk with a “vengeance” [7]. Classical work [7] argued that activation overlap – the average shared activation value over all neural units in a given layer – can create the cross-talk that either places the ANN beyond its information capacity (thus resulting in the destruction of old knowledge) or results in the situation where learning even a single new input can disrupt previously acquired information. One path for tackling the activation overlap problem would be to bias the ANN to acquire semi-distributed representations – a “sweet-spot” between fully distributed representations that generalize well (but lead to catastrophic forgetting) and local representations that are immune to forgetting (but do not generalize well). Our intuition is that, by using representations that are driven by neuronal units that fight for the right to activate in response to particular input stimuli, naturally emergent semi-distributed representations will emerge from the complex system. These representations will be far less likely to collide (provided that the right mechanisms are built into the memory system) yet still generalize well (since many possible input configurations will still match to a given memory unit).

Furthermore, competition across units creates sparse activity patterns. Since sparsity has been shown to prevent forgetting in classical systems such as sparse distributed memory [16] (before saturation), we argue that a carefully designed neural memory system that utilizes competition among its units can effectively be used to tackle the difficult problem of task-free continual learning described earlier.

Best Matching Unit (BMU) Activation Function: The most fundamental and necessary component for any neural system to be considered competitive is the design of a function for selecting a one or more units out of a

Algorithm 1 The algorithm for selecting the K best matching units.

Input: Pre-activation vector $\mathbf{h}$, K BMU desired, winner_type string flag
function FINDBMU($\mathbf{h}$, K , winner_type)
 $w = \emptyset$, $h_{min} = \min(\mathbf{h})$, $h_{max} = \max(\mathbf{h})$ ▷ Run K -winners selection process
▷ Initialize statistics
for $k = 1$ to K **do**
 if winner_type is “max” **then** ▷ winner_type is unit w/ maximum value
 $i = \arg \max_i \mathbf{h}$, $\mathbf{h}[1, i] = h_{max}$
 else ▷ winner_type is unit w/ minimum value (“min”)
 $i = \arg \min_i \mathbf{h}$, $\mathbf{h}[1, i] = h_{min}$
 $w \leftarrow w \cup \{i\}$ ▷ Record k th index of k th winning neuron
Return w ▷ Note, in this study, we treat $\{i\} = i$

pool of neurons, conditioned on sensory input or external activation. In basic competitive neural models, this is known as the *best matching unit* (BMU) and often refers to selecting one single neuron¹, which satisfies a particular criterion/constraint, out of many. However, a good deal of work in competitive learning often involves computing more than just one single neuron (at least the second BMU) (e.g., growing neural gas [9]) and our framework accounts for this by allowing more than one unit to potentially win the competition. Selecting a winning neuron can either be done by selecting the maximum, i.e., the $\max()$ and $\arg \max()$, or minimal value, i.e., $\min()$ and $\arg \min()$, out of a set of activation values and is often dependent on how the neural post-activities are computed in the first place. For simplicity, we focus on models that compute only a single hidden competitive layer of d_z neurons. Pre-activation values are denoted as $\mathbf{h} \in \mathcal{R}^{d_{z1} \times 1}$ and post-activation values are denoted as $\mathbf{z} \in \mathcal{R}^{d_z \times 1}$. However, we note that these models could be extended to L layers to create a deep competitive neural system.

The most prominent ways of computing pre-activation values are through subtraction and distance calculation or dot products and we focus on these two in this work. The parameters for an NCL model, at minimum, include $\Theta = \{\mathbf{M}\}$ where $\mathbf{M} \in \mathcal{R}^{d_z \times d_x}$. If the first approach is taken, then formally the pre-activity is computed as:

$$\mathbf{h}[i, 1] = \|\mathbf{x} - (\mathbf{M})^T[:, i]\|_p \quad (1)$$

where the i th element of the pre-activation $\mathbf{h}$ is a subtraction of the current pattern vector and the i th column of the transposed memory matrix, input as argument to the p -norm function. The second approach follows a form of computation more typical to feed-forward artificial neural networks (ANNs):

$$\mathbf{h} = \mathbf{M} \cdot \mathbf{x}, \text{ or, } \mathbf{h}[i, 1] = \mathbf{M}[i, :] \cdot \mathbf{x}. \quad (2)$$

Once the pre-activation value $\mathbf{h}$ has been computed for all neurons in layer 1, we may apply the best-matching post-activation function that conducts the actual required competition as shown in the simple procedure depicted in Algorithm 1. In this algorithm, we observe that the above algorithm can either compute the K minimal or maximal neurons (or K BMUs, stored in the list/set w). Furthermore, note that the above algorithm makes use of both $\max()/\min()$ and $\arg \max()/\arg \min()$ functions to obtain not only the values of winning neurons but also their indices within the activity vector $\mathbf{h}$.

Synaptic Weight Adjustment: The other fundamental function needed to fully specify an NCL system is its method for adjusting synaptic weight values given the results of the competition function (described in the last section). Typically, the form of the update used in competitive learning is a local Hebbian or anti-Hebbian rule [5]. The goal of this routine, in any of the models/algorithms we examine in this paper, is to produce an update matrix $\Delta \mathbf{M}$ (of the same shape as $\mathbf{M}$) which is then used to actually alter the values within $\mathbf{M}$ as follows:

$$\mathbf{M} \leftarrow \mathbf{M} + \alpha \Delta \mathbf{M} \quad (3)$$

where $0 < \alpha < 1$ controls the magnitude of the update applied to parameters $\mathbf{M}$. Note that other rules for changing the values of $\mathbf{M}$ are possible, such as an adaptive learning rate, e.g., Adam [17] or RMSprop [41].

3 Continual Competitive Memory

In this section, we describe and formally define our proposed NCL system for task-free lifelong learning, the continual competitive memory (CCM) model.

¹We will, throughout the paper, also refer to these as “prototypes” or memory units/slots.

Figure 1: Continual competitive memory (CCM) shown processing the j th pattern from a data continuum.

3.1 Learning and Inference

Fundamentally, the architecture of our proposed CCM is organized into what we call “task-memory blocks”, i.e., $\mathbf{M}_m \in \mathcal{R}^{d_{mem} \times d_x}$ is associated with a hypothetical task m and d_{mem} is the number of memory slots to be allocated for a given task. These blocks are all generally stored in a pool of M blocks which we refer to as the “long-term memory” region of the system. However, at any one instant in time, one out of all these task-memory blocks is placed in what we refer to as the “short-term memory” region of the system. The block currently placed in short-term storage is denoted $\mathbf{M}_t$ where t is known as the “task-pointer”, an integer variable maintained by the CCM (and initialized to $t = 0$). This task pointer represents the CCM’s knowledge of what current task it is operating on (out of M ones currently known by the CCM).

The CCM also maintains two other important variables – the alarm a , which, when exceeding integer threshold a_θ , indicates a “change of task” and signals to the CCM to create a brand new task memory block $\mathbf{M}_{M+1}$. Once the new memory block is created, the task pointer is moved to position $M + 1$. Furthermore, a recall vector $\mathbf{r} \in \mathbb{Z}^{M \times 1}$, a vector of recall integers, one per current memory block, is maintained by the CCM. Whenever the $\max(\mathbf{r})$ exceeds a threshold r_θ , the CCM is triggered to enter into “recall” mode and ultimately changes its task pointer t to point to the long-term memory block most frequently referred to (within a window of samples).

In terms of parameters, in addition to the counter/tracking variables described above, the CCM is formally specified by the set $\Theta = \{(\Upsilon_1, \mathbf{M}_1, \mathcal{C}_1), \dots, (\Upsilon_M, \mathbf{M}_M, \mathcal{C}_M)\}$ where, at initialization, $M = 1$. Coupled with each actual memory block $\mathbf{M}_m$ are two lists: Υ_m is a list of length d_{mem} containing the per-memory learning rates (that will be decayed over after pattern is generated by the stream) and $\mathcal{C}_m$ is a list of length d_{mem} where each element tracks the number of patterns readily matched to its assigned memory slot in $\mathbf{M}_m$. Note that M , while unbounded in this paper’s implementation of the CCM, will only grow if the CCM detects a shift in the data stream’s distribution (which would be a strong enough shift to indicate the presence of a new task) and only if the CCM cannot recall seeing the data it is currently processing at all (when peeking into its long-term memory using $\mathbf{r}$).

From a mechanics standpoint, when presented with an input pattern $\mathbf{x}$ (such as the j th pattern $\mathbf{x}_j$ from a stream), CCM first computes the dot product between the normalized of $\mathbf{x}$ ($\bar{\mathbf{x}}$) with each memory unit in the short-term task memory block $\mathbf{M}_t$. If the dot product of any of these potential matches exceeds ρ , then the CCM simply updates the matched prototype with the highest dot product/similarity score (and updates the count for that unit in $\mathcal{C}_m$ and decays its learning rate in Υ_m). If none of these dot products exceed ρ , CCM then increments the alarm a by one and then immediately checks if the current input matches any task memory units in its long-term memory. Note that if $\mathbf{x}$ matches a slot q in the long term memory, exceeding the recall vigilance meta-parameter ρ_r^2 , then $\mathbf{r}[q, 1]$ is incremented by one. If $\max(\mathbf{r}) > r_\theta$, then the CCM will enter into recall mode and retrieve the task memory block $q = \arg \max(\mathbf{r})$ that has received more than r_θ consecutive hits, changing its task pointer $t = q$. If $a > a_\theta$, however, then a fresh, empty memory block $\mathbf{M}_{M+1}$ will be created and the task pointer becomes t to $M + 1$.

The full process (as well as additional implementation details) of the CCM’s learning and inference process is presented in Algorithm 2. Graphically, in Figure 1, we depict the CCM (at a high-level) processing the j th pattern vector from the data stream/continuum. One might observe that the CCM builds upon useful concepts originally

²Higher values for ρ and ρ_r , i.e., 0.8, are used since the CCM should be sure that an input belongs to any task memory $\mathbf{M}_m$.

Algorithm 2 Inference and learning for the continual competitive memory (CCM) model.

Input: sample $\mathbf{x}$, $t = 1$, $\Theta = \{(\Upsilon_1, \mathbf{M}_1, \mathcal{C}_1)\}$, $a = 0$, $\mathbf{r} = \mathbf{0}$
Constants: $a_\theta > 0$, $r_\theta > 0$, $C_\theta > 0$, $0 < \rho < 1$, $0 < \rho_r < 1$, $0 < \epsilon \ll 1$, $0 \ll \gamma_v < 1$
// Adjust synapses given pattern vector
function UPDATE($\mathbf{x}$, ρ , ϵ , γ_v , Θ , t , a , $\mathbf{r}$, a_θ , r_θ , C_θ)
 $\Upsilon_t, \mathbf{M}_t, \mathcal{C}_t \leftarrow \Theta[t]$
 $\bar{\mathbf{x}} = \mathbf{x} / (||\mathbf{x}|| + \epsilon)$, $\mathbf{h}_t = \mathbf{M}_t \cdot \bar{\mathbf{x}}$
 $h_i = -1$, $i = -1$, $C_i = \infty$, $k = 0$
while $h_i \leq \rho \wedge C_i > C_\theta \wedge k < r(\mathbf{h}_t)$ **do** ▷ Enter search mode
 $i = \text{FINDBMU}(\mathbf{h}_t, K = 1, \text{"max"})$, $h_i = \mathbf{h}_t[i, 1]$, $C_i = \mathcal{C}_t[i]$,
 $\mathbf{h}_t[i, 1] = -1$, $k \leftarrow k + 1$
if $h_i > \rho \vee C_i \leq C_\theta$ **then**
// Update i th slot in short-term task memory block $\mathbf{M}_t$
 $v_i = \Upsilon_t[i]$, $\mathbf{M}_t[i, :] = \frac{\mathbf{M}_t[i, :] + v_i((\bar{\mathbf{x}})^T - \mathbf{M}_t[i, :])}{||\mathbf{M}_t[i, :] + v_i((\bar{\mathbf{x}})^T - \mathbf{M}_t[i, :])||_2 + \epsilon}$, $\Upsilon_t[i] = \gamma_v v_i$, $\mathcal{C}_t[i] \leftarrow \mathcal{C}_t[i] + 1$
else
// Check for recall in long-term memory, otherwise create new task memory block
 $a \leftarrow a + 1$, $q = -1$
while $m < |\Theta| \wedge h_i \leq \rho_r$ **do**
if $j \neq t$ **then**
 $\Upsilon_m, \mathbf{M}_m, \mathcal{C}_m \leftarrow \Theta[m]$, $\mathbf{h}_m = \mathbf{M}_m \cdot \bar{\mathbf{x}}$
 $i = \text{FINDBMU}(\mathbf{h}_m, K = 1, \text{"max"})$, $h_m = \mathbf{h}_m[i, 0]$, $q = m$
 $m \leftarrow m + 1$
if $h_m > \rho_r$ **then**
 $\mathbf{r}[m, 1] \leftarrow \mathbf{r}[m, 1] + 1$
if $\max(\mathbf{r}) > r_\theta$ **then** ▷ Point back to recalled memory block
 $t \leftarrow q$, $a \leftarrow 0$, $\mathbf{r} \leftarrow \mathbf{0}$
else if $a > a_\theta$ **then**
 $\Upsilon_{|\Theta|+1}, \mathbf{M}_{|\Theta|+1}, \mathcal{C}_{|\Theta|+1} \leftarrow \text{INITPARAMS}()$ ▷ Create new memory block $|\Theta| + 1$
 $\Theta \leftarrow \Theta \cup \{(\Upsilon_{|\Theta|+1}, \mathbf{M}_{|\Theta|+1}, \mathcal{C}_{|\Theta|+1})\}$, $t \leftarrow |\Theta| + 1$, $a \leftarrow 0$, $\mathbf{r} \leftarrow \mathbf{0}$ ▷ Point to new $\mathbf{M}_{|\Theta|+1}$
// Determine which memory block pattern $\mathbf{x}$ belongs to
function GETTASKPOINTER($\mathbf{x}$, Θ , ϵ)
 $\mathbf{M} = \emptyset$
for $m = 1$ to $|\Theta|$ **do** ▷ Concatenate all memory blocks into one single matrix
 $\Upsilon_m, \mathbf{M}_m, \mathcal{C}_m \leftarrow \Theta[m]$, $d_{mem} = r(\mathbf{M}_m)$
 $(m > 1 \rightarrow \mathbf{M} = [\mathbf{M}, \mathbf{M}_m]_r) \wedge (m \leq 1 \rightarrow \mathbf{M} = \mathbf{M}_m)$
 $\bar{\mathbf{x}} = \mathbf{x} / (||\mathbf{x}|| + \epsilon)$, $\mathbf{h} = \mathbf{M} \cdot \bar{\mathbf{x}}$
 $i = \text{FINDBMU}(\mathbf{h}, K = 1, \text{"max"})$, $t_i = \lfloor i / d_{mem} \rfloor$
Return t_i ▷ Output task memory block pointer

proposed in adaptive resonance theory (ART) [12, 13, 14]³, particularly its use of a “vigilance” meta-parameter to drive neuronal unit competition entering during the BMU search phase. In that sense, one could consider the CCM to have generalized ART to the task-free online continual learning setting with the critical distinction that the CCM focuses on different functionality with respect to sequential learning. Furthermore, CCM is structured in terms of a short and long-term memory parameter organization and utilizes a memory update rule that is rather different than the one used in ART. As we will see in our experiments later, while ART is in of itself quite powerful, the CCM’s design to self-organize its memory contents with respect to short and long-term storage regions proves to be important when learning from longer, high-dimensional task streams.

3.2 Detecting Tasks through Modality Pointers at Test-Time Inference

One of the key uses of the CCM, at various points of the stream, is to, when presented with a pool of held-out test data patterns, infer which of its own current task memory blocks that the data pool belongs to. In Algorithm 2, we present a simple procedure (GETTASKPOINTER()) that does this for a single input pattern, leveraging the

³ART originally proposed with non-stationary data streams in mind, where catastrophic interference is a prominent issue.

current state of the CCM’s long-term memory. In essence, the CCM will compare the normalized input pattern $\bar{\mathbf{x}}$ to each possible memory slot in $\mathbf{M}$, which is a row-wise concatenation of all memory blocks in long-term storage. To compute which task memory block the resultant BMU index i refers to, one may simply divide i by the total number of unique memory slots in any $\mathbf{M}_m$, i.e., d_{mem} and take the floor of the result. Finally, when provided with a data pool, the CCM infers the pool’s task by applying GETTASKPOINTER() to each pattern, storing the resulting integer t_i for each in a list $\mathcal{P}$, and finally computes the modal value t_M of this list of integers, i.e., $t_M = \text{mode}(\mathcal{P})$.

Coupling the CCM to a Neural Predictor: Armed with the above process for inferring the modal task pointer at test-time, we may now briefly describe a powerful extension of the CCM when processing data from a stream. In short, imagine that we have access to a prediction function $f_{\Theta_f}(\mathbf{x})$. The function takes on the form of an ANN where a stack of nonlinear transformations, or $\{f_\ell(\mathbf{z}_{\ell-1}; \theta_\ell)\}_{\ell=1}^L$, is applied to the input $\mathbf{x}$. As an example, if the network is a multilayer perceptron (MLP), each transformation $\mathbf{z}_\ell = f_\ell(\mathbf{z}_{\ell-1})$ produces an output $\mathbf{z}_\ell$ from the value $\mathbf{z}_{\ell-1}$ of the previous layer with the help of a weight matrix $\theta_\ell = \{\mathbf{W}_{(\ell-1) \rightarrow \ell}\}$. f_ℓ is decomposed into two operations (biases omitted for clarity):

$$\mathbf{z}_\ell = \phi_\ell(\mathbf{h}_\ell), \quad \mathbf{h}_\ell = \mathbf{W}_{(\ell-1) \rightarrow \ell} \cdot \mathbf{z}_{\ell-1} \quad (4)$$

where ϕ_ℓ is an activation function, $\mathbf{z}_\ell \in \mathcal{R}^{H_\ell}$ is the post-activation of layer ℓ while $\mathbf{h}_\ell \in \mathcal{R}^{H_\ell}$ is the pre-activation vector of layer ℓ (H_ℓ is the number of neurons in layer ℓ). Note that $\mathbf{z}^0 = \mathbf{x}$.

To leverage the CCM’s task pointer and couple it to the MLP predictor to prevent catastrophic interference when updating MLP parameters via backprop, we take inspiration from the biologically-inspired complementary neural system proposed in [29]. Specifically, we introduce one final set of parameters for each layer of the MLP (except the bottom-most and top-most layers), i.e., a binary matrix $\mathbf{Q}^\ell \in \{0, 1\}^{|\Theta| \times H_\ell}$ where each row contains a unique set of M_H ones ($M_H \ll H_\ell$). The CCM’s task pointer, either t during training within the stream or t_M during test-time inference, is encoded via $\mathbf{t} = \text{ENC}(t, |\Theta|)$ which is then multiplied with each $\mathbf{Q}^\ell$ to obtain a gating vector $\mathbf{g}^\ell$ for each hidden layer of the MLP. The gate vector is formally calculated as $\mathbf{g}^\ell = \mathbf{Q}^\ell \cdot \mathbf{t}$ which is multiplied element-wise with each layer of the MLP as follows: $\mathbf{z}^\ell \leftarrow \mathbf{z}^\ell \otimes \mathbf{g}^\ell$.

The above coupling of the CCM with the MLP (we refer to this as CCM-MLP in our experiments), allows the CCM to drive portions or sub-networks of the MLP predictor when processing data and facilitate stronger label predictive generalization. Furthermore, this demonstrates that the CCM is compatible with popular ANNs like the MLP.

4 Neural Competitive Learning Baselines

In this section, we describe several important NCL baseline models that our CCM will be compared against, including our adaptation of adaptive resonance theory (ART).

4.1 The Incremental Winner-Take-All (iWTA) Model

The simplest competitive learning model is the winner-take-all (WTA) system. In this model, it is typical to combine the subtractive pre-activation rule (Equation 1) with a distance function and Algorithm 1. The resulting model, with parameter matrix $\mathbf{M} \in \mathcal{R}^{d_z \times d_x}$ and generalized for finding the K BMU given input $\mathbf{x}$, is:

$$\mathbf{h}[i, 1] = \|\mathbf{x} - (\mathbf{M})^T[:, i]\|_p, \quad \forall i = 0, 1, \dots, r(\mathbf{M}) \quad (5)$$

$$w = \text{FINDBMU}(\mathbf{h}, K, \text{“min”}), \quad \text{and}, \quad \mathbf{z} = \text{Enc}(w, d_z) \quad (6)$$

where $K = 1$ recovers the classical unsupervised vector quantization model [11]. To update the synaptic weights of the WTA model, we employ the following Hebbian rule:

$$\Delta \mathbf{M}[i, :] = \mathbf{z}[i, 1] \left((\mathbf{x})^T - \mathbf{M}[i, :] \right). \quad (7)$$

If we observe that each subtraction vector in the rule above is simply being multiplied by the i th value in the binary column vector $\mathbf{z}$, we can write a vectorized form of the rule as follows:

$$\Delta \mathbf{M} = (\text{BR}_{d_z}((\mathbf{x})^T) - \mathbf{M}) \otimes \text{BC}_{d_x}(\mathbf{z}). \quad (8)$$

Since this NCL model will adapt to patterns online, we shall further refer to it as incremental WTA (iWTA).

4.2 The Incremental Gaussian Mixture Model (iGMM)

The iWTA model presented above can be generalized further by adapting aspects of the venerable mixture of Gaussians model, otherwise known as the Gaussian mixture model (GMM), often used in density estimation and clustering [31, 26]. The first adaptation of the GMM to the NCL setup, which we call the incremental GMM (iGMM), assumes isotropic Gaussian components with a fixed scalar variance $\gamma = 1/\sigma^2$ (a hyper-parameter

Algorithm 3 The inference and learning processes for ART-C 2A.

Input: sample $\mathbf{x}$, $\mathbf{M} = \emptyset$, $\Upsilon = \emptyset$, $0 < \rho < 1$
Constants: $0 < v_0 < 1$, $0 < \epsilon \ll 1$, $C_\theta > 0$, $0 \ll \gamma_v < 1$
function UPDATE($\mathbf{x}$, ρ , ϵ , $\mathbf{M}$, C_θ)
 $\bar{\mathbf{x}} = \mathbf{x} / (||\mathbf{x}|| + \epsilon)$, $\mathbf{h} = \mathbf{M} \cdot \bar{\mathbf{x}}$
if $\mathbf{M} \neq \emptyset$ **then**
 $h_i = -1, i = -1$
while $h_i \leq \rho \wedge k < r(\mathbf{h})$ **do** ▷ Enter search mode
 $i = \text{FINDBMU}(\mathbf{h}, K = 1, \text{"max"})$, $h_i = \mathbf{h}[i, 1]$ ▷ Find current BMU
 $\mathbf{h}[i, 1] = -1$, $k \leftarrow k + 1$
if $h_i > \rho$ **then** ▷ Enter resonance mode
 $v_i = \Upsilon[i]$, $\mathbf{M}[i, :] = \frac{v_i(\bar{\mathbf{x}})^T + (1-v_i)\mathbf{M}[i, :]}{||v_i(\bar{\mathbf{x}})^T + (1-v_i)\mathbf{M}[i, :]|_2}$, $\Upsilon[i] = \gamma_v v_i$ ▷ Update weight vector i
else ▷ Conduct mismatch reset
 $\mathbf{M} = [\mathbf{M}, (\bar{\mathbf{x}})^T]_r$
if $r(\mathbf{M}) > C_\theta$ **then**
Apply constraint reset satisfaction according to [14], yielding corrected $\mathbf{M}$
Compute new vigilance ρ according to [14]
else
 $\mathbf{M} = (\bar{\mathbf{x}})^T$, $\Upsilon = \{v_0\}$ ▷ Set very first weight vector to be input pattern

decided by the user). This yields a set of Gaussian kernels with which the input pattern $\mathbf{x}$ is compared against:

$$\mathbf{h}[i, 1] = ||\mathbf{x} - (\mathbf{M})^T[:, i]||_p, \forall i = 0, 1, \dots, r(\mathbf{M}) \quad (9)$$

$$\mathbf{z} = \frac{\exp(-\gamma \mathbf{h}/T)}{\sum_j -\gamma \exp(\mathbf{h}/T)[j, 1]} \quad (\text{Calculate posterior probabilities over units}) \quad (10)$$

where we have introduced the temperature meta-parameter T (we set $T = 0.0285$ in this study). Notice that the above formulation produces a probabilistic weight per unit/prototype rather than a hard quantization like in WTA.

The update for the iGMM also follows a simple subtractive form:

$$\Delta \mathbf{M} = \gamma \left((\text{BR}_{d_z}((\mathbf{x})^T) - \mathbf{M}) \otimes \text{BC}_{d_x}(\mathbf{z}) \right). \quad (11)$$

but the prototype comparisons are weighted by the posterior probabilities over prototypes as well as the shared variance. This means that all units will be moved to some degree by the presence of an input $\mathbf{x}$.

4.3 Adaptive Resonance Theory (ART)

In this section, we present the fast variant of adaptive resonance theory (ART) [12] designed to tackle real-valued/continuous inputs, known as ART 2A. Specifically, we implement and explore the ART-C 2A model [14].

In standard ART, a vigilance meta-parameter ρ is used to detect novelty – if a match computed between an input pattern and a memory unit exceeds ρ then the memory is adjusted (using a moving-average like rule), otherwise, a new category/memory is created (initialized to the current pattern). In Algorithm 3, our implementation of the ART-C 2A process as an NCL model also notably includes an extension known as constraint reset, meant to prevent an unbounded growth of memory units [14]. ART-C 2A involves a bit more than just calculation of the BMU post-activation before updating weights – it, as does CCM, conducts a search through its activities to find a BMU that satisfies a “vigilance” threshold check. If none of the currently available activities pass the vigilance check, then a new category/weight vector is instantiated. The constraint reset mechanism, again, ensures that this NCL model never has more than C_θ units. To satisfy this bound on units, once the number of units exceed C_θ , a merging operation is applied that replaces the two nearest neighbor memories (in terms of dot-product similarity) and replaces them both with a single vector containing their average. Furthermore, the vigilance parameter is dynamically re-computed upon each input presentation [14] (based on the maximum vigilance computed given the current input $\mathbf{x}$). Note that we modified the underlying ART-C 2A process further by incorporating a per-unit learning rate that was decayed as a function of the number of input samples similar to that used in our CCM (we found this improved generalization a bit when learning online).

We comment that the literature related to ART is vast and a plethora of model variants exist. Given that it was our interest to only examine one of the simplest/fastest variants of ART, i.e., ARTC-2A, when making comparisons, the

reader should read the results reported for our ART model noting that potentially one of the more complex/intricate ART systems (such as those based on differential equations [12]) might yield different and different results.

4.4 Modal Classification

All of the NCL models presented above, including our CCM, can further be extended to offer some basic discriminative functionality. While some models like ART have special generalizations such as ARTMAP [2], we will focus, in this study, on generalizing all of the unsupervised memory systems above in the same way. Specifically, an NCL model can be made to conduct classification of its input through the following two alterations:

- During learning/synaptic adjustment, the NCL model must maintain and update a single frequency vector $\mathbf{y}_i^c \in \mathbb{Z}^{d_y \times 1}$ for each memory unit. meaning there are d_z (or $d_{mem} \times d_z$ for the CCM) total frequency vectors yielding an extra tensor parameter $\mathbf{Y}^c \in \mathbb{Z}^{d_z \times d_y \times 1}$. This means that, upon presentation of $(\mathbf{x}_j, \mathbf{y}_j)$, the model computes $\mathbf{h}$ and $i = \text{FINDBMU}(\mathbf{h}, K, \text{winner_type})$ and then updates i th slice of $\mathbf{Y}^c$ through simple addition, i.e., $\mathbf{Y}^c[i, :, :] = \mathbf{Y}^c[i, :, :] + \mathbf{y}_j$ (since $\mathbf{y}_j$ is one-hot encoded, it already contains an increment of one at the correct class index for sample j). This means $\mathbf{y}_j$ is bound to its relevant unit.
- During (test-time) inference, upon the presentation of $\mathbf{x}_j$, the NCL model classifies the pattern by first calculating pre-activity $\mathbf{h}$ as well as the winning neuron $i = \text{FINDBMU}(\mathbf{h}, K, \text{winner_type})$ and then retrieving the i th slice of $\mathbf{Y}^c$. The final label assignment is made via the following rule: $y_j = \arg \max_y \mathbf{Y}^c[i, :, :]$ where y_j is an integer representing the NCL model's predicted class for $\mathbf{x}_j$.

Note that the above two modifications to our NCL models means that we are conducting classification by estimating the target label as the mode of the known training sample labels bound over time to the relevant memory $\mathbf{h}[:, i]$ that responds to $\mathbf{x}_j$, i.e., assigning the majority vote label. This choice of mechanism allows us to separate out the classification process from the pattern matching process of the memory system, providing us with a general family of models that can handle the case of data streams that either have no labels at all (unsupervised) or present some samples with labels and others without (semi-supervised).

We note that there are other mechanisms that could be used instead to conduct classification (and potentially yield better discriminative performance), including modifying the learning process of each NCL to work similar to learning vector quantization [20] or to learn a classifier jointly with the memory system (as in the case of ARTMAP [2]). Nonetheless, we shall leave investigation of such generalizations for future work given that our interest in this paper is to investigate memory retention without joint reliance on the label.

5 Experiments

5.1 Metrics

In this section, we describe the set of evaluation metrics used to compare our CCM against the competitive learning baselines. Specifically, we measure performance with respect to two properties: 1) knowledge retention and generalization across tasks in a stream, and 2) cluster validity through class conformity.

Generalization and Knowledge Retention: In order to quantify each model's ability to generalize and retain previous knowledge, we measure average accuracy (ACC), or mean performance across tasks, and backwards transfer (BWT), two key metrics often used to quantify how well a continual learning model retains knowledge.

To use these metrics, we must first compose a task matrix R (as in [22]), which is an $T \times T$ matrix of task accuracy scores (normalized to $[0, 1]$). With respect to R , the two key metrics we measure are defined as:

$$\text{ACC} = \frac{1}{T} \sum_{i=1}^T R_{T,i}, \quad \text{BWT} = \frac{1}{T-1} \sum_{i=1}^{T-1} R_{T,i} - R_{i,i}. \quad (12)$$

BWT measures the influence that learning a task i has on the performance of task $k < i$. A positive BWT indicates that a learning task i increases performance on a preceding task k . As such, a higher BWT is better and a strongly negative BWT means there is stronger (more catastrophic) forgetting.

Class Conformity: We next measure how well the clusters acquired by each model conform to the target class distribution associated with each benchmark. Since our NCL models are applied to supervised learning streams, we are interested in examining explicit clustering metrics that determine how similar each model's clustering it to the input data's (empirical) ground truth label distribution. In this paper, we record three such explicit metrics

MNIST	FMI	ARI	V-M
iWTA	0.7464 ± 0.004	0.7179 ± 0.0044	0.8234 ± 0.0023
iGMM	0.7526 ± 0.0292	0.7249 ± 0.0323	0.8256 ± 0.0175
ART-C 2A	0.6635 ± 0.0245	0.6237 ± 0.0285	0.7817 ± 0.0129
CCM (Ours)	0.8990 ± 0.0003	0.8877 ± 0.0004	0.9132 ± 0.0002
NotMNIST	FMI	ARI	V-M
iWTA	0.6896 ± 0.0156	0.6551 ± 0.0174	0.7887 ± 0.0080
iGMM	0.7014 ± 0.0213	0.6680 ± 0.0236	0.7967 ± 0.0122
ART-C 2A	0.6495 ± 0.0205	0.6096 ± 0.0221	0.7716 ± 0.0115
CCM (Ours)	0.7319 ± 0.0114	0.7021 ± 0.0126	0.8111 ± 0.0060

Table 1: Class conformity measurements for each NCL model. Mean & standard deviation reported over 10 trials.

	MNIST		NotMNIST	
	ACC	BWT	ACC	BWT
ICarl [29]	0.9400 ± 0.4100	-0.1000 ± 0.0040	0.8870 ± 0.1020	-0.1090 ± 0.0070
Mnemonics [29]	0.9600 ± 0.3200	-0.9910 ± 0.0050	0.9500 ± 0.0710	-0.0110 ± 0.0070
MLP	0.6132 ± 0.0133	-0.4671 ± 0.0165	0.6222 ± 0.0243	-0.4389 ± 0.0336
iWTA	0.8285 ± 0.0022	-0.0824 ± 0.0122	0.8386 ± 0.0068	-0.0468 ± 0.0109
iGMM	0.8329 ± 0.0145	-0.0746 ± 0.0231	0.8437 ± 0.0128	-0.0350 ± 0.0048
ART-C 2A	0.7220 ± 0.0216	-0.1829 ± 0.0164	0.7852 ± 0.0164	-0.0755 ± 0.0028
CCM (Ours)	0.9418 ± 0.0011	-0.0211 ± 0.0001	0.8784 ± 0.0110	-0.0074 ± 0.0081
CCM-MLP (Ours)	0.9853 ± 0.0002	-0.0006 ± 0.0001	0.9553 ± 0.0006	-0.0024 ± 0.0019

Table 2: Knowledge retention metrics for each NCL model. Mean & standard deviation reported over 10 trials.

(bounded in the range of $[0, 1]$): adjusted random index (ARI) [15], the Fowlkes-Mallows index (FMI) [6], and the V-measure (V-M) [36] (we refer the reader to the references for implementation details of each metric).

For all three of the above conformity metrics, a higher score closer to 1.0 is desired. Values closer to zero generally indicate that a model’s clustering and the ground truth labels are largely independent while values closer to one indicate significant agreement. Furthermore, we selected these three metrics given that all of them make no assumption as to what kind of cluster structure that any particular algorithm should yield. Notably, the V-measure was used in [14] (which proposed ARTC-2A) and is derived from conditional entropy analysis, offering a measurement that accounts for two desirable properties one would want in a model’s resultant clustering: homogeneity, where each cluster contains solely members of a single class/label, and completeness (class entropy), where all members of a particular class are assigned to the same cluster by a given model.

5.2 Data Benchmarks for Online Lifelong Learning

Split Image Benchmarks: The MNIST dataset⁴ contains 28×28 images with gray-scale pixel feature values in the range of $[0, 255]$. The only preprocessing applied to this data is to normalize the pixel values to the range of $[0, 1]$ by dividing them by 255. NotMNIST⁵ is a more difficult variation of MNIST created by replacing the digits with characters of varying fonts/glyphs (letters A-J). The preprocessing and training, validation, and testing splits were created to match the setup of MNIST with the exception that there is more data in NotMNIST, i.e., 100,000 data points are in the training split we used for our experiments (adapted from the smaller variant in [28]).

To create Split MNIST and Split NotMNIST, the samples in each are organized by class. A series of five tasks ($t = \{1, 2, 3, 4, 5\}$) are created as follows: 1) **Task 1** – categorize the first two classes, e.g., 0 and 1 in MNIST, 2) **Task 2** – categorize the next two classes, e.g., 2 and 3 in MNIST, 3) **Task 3** – categorize the next two classes, e.g., 4 and 5 in MNIST, 4) **Task 4** – categorize the next two classes, e.g., 6 and 7 in MNIST, and, 5) **Task 5** – categorize the last two classes, e.g., 8 and 9 in MNIST. Note that for each task above, the label would be one-hot encoded into a binary vector of size 2, i.e., $\mathbf{y}_j \in \{0, 1\}^{2 \times 1}$. Mini-batches of samples produced by each data stream were constrained to be 1 to simulate online learning (data in each task was shuffled only once prior to simulation).

For baseline models, we not only compared to a tuned implementation of each of the NCL baseline models (described in Section 4) but to also an MLP with two hidden layers (first layer had 110 units and second layer had 104 units – these numbers were selected to ensure the MLP had the same number of total synaptic parameters as

⁴Available at the URL: <http://yann.lecun.com/exdb/mnist/>.

⁵URL: <http://yaroslavvb.blogspot.com/2011/09/notmnist-dataset.html>

Figure 2: Sample memories for each neural model – top row are the memories for MNIST and the bottom row are the memories for NotMNIST. Note that ART-C 2A and CCM memories are normalized such that they have unit Euclidean norms whereas the iGMM memories are not.

the NCL models). For iWTA and iGMM, the total number of memory units were set to 125 (which would be the maximum number of units allowed that both ART-C 2A and the CCM could grow by the upon reaching the end of the stream – for ART-C 2A, we set $C_\theta = 125$). For both CCM and ART-C 2A, $\epsilon = 0.00001$. For ART-C 2A, we set $\gamma_v = 1$ and $v_0 = 0.02$ for MNIST and $\gamma_v = 0.998$ and we set $v_0 = 0.1$ for NotMNIST. For the CCM, for MNIST, we set $a_\theta = 20$ and $\rho = \rho_r = 0.8$ while, for NotMNIST, we set $a_\theta = 30$ and $\rho = \rho_r = 0.755$. For both datasets, the CCM also used $\gamma_v = 0.998$, $v_0 = 0.35$, $r_\theta = 30$, and $C_\theta = 60$.

5.3 Results

In Tables 1 and 2, we present our main results for the various NCL models evaluated for this study. We observe that, for all metrics (both for class conformity and knowledge retention/generalization) measured, our CCM not only outperforms the MLP baselines and other NCL models but offers strong resistance to forgetting as indicated by its high ACC and low BWT scores (for both the CCM and coupled CCM-MLP model). Another promising result is that, with respect to the ACC and BWT, all of the NCL models improve on memory retention and generalization compared to the MLP trained with backprop. This promising result offers empirical evidence that competitive learning can serve as a means of learning semi-distributed representations that combat excessive neural cross-talk and, ultimately, catastrophic interference. Notice that at the top of Table 2, we have included two modern-day,

state-of-the-art MLP baselines trained via backprop – for our CCM using modal classification, we most nearly match the ICarl model and for the CCM-MLP integrated model, we outperform even these competitive baselines.

In Figure 2, we visualize the prototypes acquired by the iGMM, ART-C 2A, and the proposed CCM. While all of the models have appeared to acquire memory slots specialized for each of the digits (MNIST) or letters (NotMNIST), we notice that the iGMM does not fully utilize most of its capacity to represent data from each task, especially in the case of Split MNIST. We hypothesize that this might be due to the fact that, since all of the iGMM’s memory slots are adjusted in the presence of sensory input (proportional to their estimated posterior probabilities), there is less diversity among the different memory units (since constantly moving all units might prevent many from specializing). Furthermore, we observe that the CCM’s memories, as expected since its design embodies the inductive bias of a dual short-term and long-term memory organization of parameter blocks, offer a more balanced pool of digits from each class per task. This is likely a driver behind its high classification performance in general.

6 Related Work

Artificial neural systems that conduct inference and learning based on competitive learning has a fairly long history [37, 4, 43]. In essence, competitive learning centers around the idea that a pool of neurons/units compete for the right to activate – leading to different units that activate for different bundles or clusters of patterns, i.e., yielding a form of neural template matching. As more data is presented to this pool of neurons, each unit in the pool will ultimately converge to the center of the cluster it has come to model. This means that each neuron will activate more strongly for sensory inputs strongly related to its cluster and more weakly for patterns related to other units/clusters. Furthermore, there exists a vast array of neural models and their learning algorithms based on competition-oriented neural computation. These models range from specialized supervised systems driven by learning vector quantization [11] to explicit topology clustering systems based on the esteemed self-organizing map (SOM) [19]. Other models have generalized competitive learning in other directions, such as those that make use of the top $K > 1$ winning neurons as in the case of models that classify under competitive Hebbian learning [44, 23], including even compression systems based on neural gas and its variants [24, 9]. Many if not all of these models can be reformulated as NCL models. In addition, the NCL framework could be extended to account for unique features such as explicit constraints placed on the topology of competing neurons in order to include SOMs.

Adaptive resonance theory (ART) [12], from which our CCM certainly draws inspiration from, is a large and ever-growing family of powerful NCL models that favorably offer connections to real underlying neuro-biological processes [13]. While we only presented how one of the more classical (yet fast and efficient) forms of ART fall under the NCL framework in this paper, i.e., ART-C 2A, other models such as Fuzzy ART [3] and ARTMAP [2] could be similarly included, provided that sufficient modifications were made to the framework in order to account for special features unique to each variant.

In terms of investigating and generalizing competitive learning effectively to the grand challenge of task-free lifelong machine learning, far less work exists. Of the potential methods developed in prior efforts, the self-organizing incremental neural network (SOINN) [10, 39] represents one particular system that is also meant to operate in the face of data streams (including data that comes from tasks presented sequentially). SOINN is, in some sense, an interesting generalization of growing neural gas (GNG) [9], where neural prototypes are generated incrementally based on several criterion and are, crucially, guided by a relation graph between the units that is incrementally constructed and constantly updated. However, while this competitive learning model is promising for processing data streams, it is difficult to scale it to large problem spaces given that an entire graph of the units/centroids must be maintained and a vast array of key meta-parameters are introduced that must be carefully tuned to the problem space at hand. Furthermore, the experiments shown to evaluate the efficacy of the SOINN, while again yielding positive results, tend to focus on rather small-scale, low-dimensional problems/datasets.

7 Conclusions

In this paper, we examined several powerful, biologically-plausible competitive neural systems, unifying them under a common computational framework, i.e., the neural competitive learning (NCL) framework, and applied them to the challenging scenario of online continual learning, without using explicitly provided/set task descriptors. Furthermore, we proposed a novel neural memory based on competitive learning that we call continual competitive memory (CCM). Our results demonstrate that not only does our proposed CCM model generalize well to the task of online continual learning, offering performance competitive with modern-day continual learning baselines, but that competitive learning, in general, provides some robustness to catastrophic interference as compared to an artificial neural network trained via back-propagation of errors. While powerful, the CCM does not come

without its limitations: 1) it rejects data samples with a fixed threshold (a scheme to dynamically adjust the CCM’s thresholds would be highly desirable), and, 2) it does not offer an upper bound/constraint on the number of units it grows as it processes data from a stream. This last issue could be resolved by developing a constraint reset schema much like that employed in the adaptive resonance theory model we experimented with in this paper, however, great care would need to be taken to ensure that merging of units is done in such a way to not disrupt the long-term memory too much and risk accidental deletion of important task-specific knowledge. Finally, given that we have shown that the CCM readily integrates with other predictors to improve generalization, future work should entail investigating replacing the task selection basal ganglia model explored in [29] with our CCM, since the CCM overcomes many of the limitations faced by the original task selector, yielding a much more robust, neuro-cognitively plausible continual learning system.

References

- [1] ALJUNDI, R., ROHRBACH, M., AND TUYTELAARS, T. Selfless sequential learning. *arXiv preprint arXiv:1806.05421* (2018).
- [2] CARPENTER, G. A., GROSSBERG, S., AND REYNOLDS, J. H. Artmap: Supervised real-time learning and classification of nonstationary data by a self-organizing neural network. *Neural networks* 4, 5 (1991), 565–588.
- [3] CARPENTER, G. A., GROSSBERG, S., AND ROSEN, D. B. Fuzzy art: Fast stable learning and categorization of analog patterns by an adaptive resonance system. *Neural networks* 4, 6 (1991), 759–771.
- [4] DESIENO, D. Adding a conscience to competitive learning. In *ICNN* (1988), vol. 1.
- [5] FÖLDIAK, P. Forming sparse representations by local anti-hebbian learning. *Biological cybernetics* 64, 2 (1990), 165–170.
- [6] FOWLKES, E. B., AND MALLOWS, C. L. A method for comparing two hierarchical clusterings. *Journal of the American statistical association* 78, 383 (1983), 553–569.
- [7] FRENCH, R. M. Using semi-distributed representations to overcome catastrophic forgetting in connectionist networks. In *Proceedings of the 13th annual cognitive science society conference* (1991), vol. 1, pp. 173–178.
- [8] FRENCH, R. M. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences* 3, 4 (1999), 128–135.
- [9] FRITZKE, B., ET AL. A growing neural gas network learns topologies. *Advances in neural information processing systems* 7 (1995), 625–632.
- [10] FURAO, S., OGURA, T., AND HASEGAWA, O. An enhanced self-organizing incremental neural network for online unsupervised learning. *Neural Networks* 20, 8 (2007), 893–903.
- [11] GRAY, R. Vector quantization. *IEEE Assp Magazine* 1, 2 (1984), 4–29.
- [12] GROSSBERG, S. Competitive learning: From interactive activation to adaptive resonance. *Cognitive science* 11, 1 (1987), 23–63.
- [13] GROSSBERG, S. Adaptive resonance theory: How a brain learns to consciously attend, learn, and recognize a changing world. *Neural networks* 37 (2013), 1–47.
- [14] HE, J., TAN, A.-H., AND TAN, C.-L. Modified art 2a growing network capable of generating a fixed number of nodes. *IEEE Transactions on Neural Networks* 15, 3 (2004), 728–737.
- [15] HUBERT, L., AND ARABIE, P. Comparing partitions. *Journal of classification* 2, 1 (1985), 193–218.
- [16] KANERVA, P. *Sparse distributed memory*. MIT press, 1988.
- [17] KINGMA, D. P., AND BA, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [18] KIRKPATRICK, J., PASCANU, R., RABINOWITZ, N., VENESS, J., DESJARDINS, G., RUSU, A. A., MILAN, K., QUAN, J., RAMALHO, T., GRABSKA-BARWINSKA, A., ET AL. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences* 114, 13 (2017), 3521–3526.
- [19] KOHONEN, T. The self-organizing map. *Proceedings of the IEEE* 78, 9 (1990), 1464–1480.
- [20] KOHONEN, T. Learning vector quantization. In *Self-organizing maps*. Springer, 1995, pp. 175–189.
- [21] LEWANDOWSKY, S. On the relation between catastrophic interference and generalization in connectionist networks. *Journal of Biological Systems* 2, 03 (1994), 307–333.

- [22] LOPEZ-PAZ, D., ET AL. Gradient episodic memory for continual learning. In *Advances in Neural Information Processing Systems* (2017), pp. 6470–6479.
- [23] MARTINETZ, T. Competitive hebbian learning rule forms perfectly topology preserving maps. In *International conference on artificial neural networks* (1993), Springer, pp. 427–434.
- [24] MARTINETZ, T. M., BERKOVICH, S. G., AND SCHULTEN, K. J. 'neural-gas' network for vector quantization and its application to time-series prediction. *IEEE transactions on neural networks* 4, 4 (1993), 558–569.
- [25] MCCLOSKEY, M., AND COHEN, N. J. Catastrophic interference in connectionist networks: The sequential learning problem. *The psychology of learning and motivation* 24, 109 (1989), 92.
- [26] MCLACHLAN, G. J., LEE, S. X., AND RATHNAYAKE, S. I. Finite mixture models. *Annual review of statistics and its application* 6 (2019), 355–378.
- [27] MITCHELL, T., COHEN, W., HRUSCHKA, E., TALUKDAR, P., YANG, B., BETTERIDGE, J., CARLSON, A., DALVI, B., GARDNER, M., KISIEL, B., ET AL. Never-ending learning. *Communications of the ACM* 61, 5 (2018), 103–115.
- [28] ORORBIA, A., MALI, A., GILES, C. L., AND KIFER, D. Continual learning of recurrent neural networks by locally aligning distributed representations. *IEEE Transactions on Neural Networks and Learning Systems* (2020).
- [29] ORORBIA, A., MALI, A., KIFER, D., AND GILES, C. L. Lifelong neural predictive coding: Learning cumulatively online without forgetting. *arXiv preprint arXiv:1905.10696* (2019).
- [30] PARISI, G. I., AND WERMTER, S. Lifelong learning of action representations with deep neural self-organization. In *The AAAI 2017 Spring Symposium on Science of Intelligence: Computational Principles of Natural and Artificial Intelligence, Stanford, US* (2017), pp. 608–612.
- [31] PINTO, R. C., AND ENGEL, P. M. A fast incremental gaussian mixture model. *PloS one* 10, 10 (2015), e0139931.
- [32] RATCLIFF, R. Connectionist models of recognition memory: constraints imposed by learning and forgetting functions. *Psychological review* 97, 2 (1990), 285.
- [33] REBUFFI, S.-A., KOLESNIKOV, A., AND LAMPERT, C. H. icarl: Incremental classifier and representation learning. In *Proc. CVPR* (2017).
- [34] ROBINS, A. Catastrophic forgetting in neural networks: the role of rehearsal mechanisms. In *Artificial Neural Networks and Expert Systems, 1993. Proceedings., First New Zealand International Two-Stream Conference on* (1993), IEEE, pp. 65–68.
- [35] ROBINS, A. Catastrophic forgetting, rehearsal and pseudorehearsal. *Connection Science* 7, 2 (1995), 123–146.
- [36] ROSENBERG, A., AND HIRSCHBERG, J. V-measure: A conditional entropy-based external cluster evaluation measure. In *Proceedings of the 2007 joint conference on empirical methods in natural language processing and computational natural language learning (EMNLP-CoNLL)* (2007), pp. 410–420.
- [37] RUMELHART, D. E., AND ZIPSER, D. Feature discovery by competitive learning. *Cognitive Science* 9, 1 (1985), 75–112.
- [38] RUSU, A. A., RABINOWITZ, N. C., DESJARDINS, G., SOYER, H., KIRKPATRICK, J., KAVUKCUOGLU, K., PASCANU, R., AND HADSELL, R. Progressive neural networks. *arXiv preprint arXiv:1606.04671* (2016).
- [39] SHEN, F., AND HASEGAWA, O. A fast nearest neighbor classifier based on self-organizing incremental neural network. *Neural Networks* 21, 10 (2008), 1537–1547.
- [40] THRUN, S., AND MITCHELL, T. M. Lifelong robot learning. *Robotics and autonomous systems* 15, 1-2 (1995), 25–46.
- [41] TIELEMAN, T., AND HINTON, G. Lecture 6.5—RmsProp: Divide the gradient by a running average of its recent magnitude. COURSERA: Neural Networks for Machine Learning, 2012.
- [42] VAN DE VEN, G. M., AND TOLIAS, A. S. Three scenarios for continual learning. *arXiv preprint arXiv:1904.07734* (2019).
- [43] WANG, L. On competitive learning. *IEEE Transactions on neural networks* 8, 5 (1997), 1214–1217.
- [44] WHITE, R. H. Competitive hebbian learning. In *IJCNN-91-Seattle International Joint Conference on Neural Networks* (1991), vol. 2, IEEE, pp. 949–vol.
- [45] WILLSHAW, D. Holography, associative memory, and inductive generalization. In *Parallel models of associative memory*. Psychology Press, 2014, pp. 113–138.