

Continual Prototype Evolution: Learning Online from Non-Stationary Data Streams

Matthias De Lange & Tinne Tuytelaars
KU Leuven

{matthias.delange, tinne.tuytelaars}@kuleuven.be

Abstract

Attaining prototypical features to represent class distributions is well established in representation learning. However, learning prototypes online from streaming data proves a challenging endeavor as they rapidly become outdated, caused by an ever-changing parameter space during the learning process. Additionally, continual learning does not assume the data stream to be stationary, typically resulting in catastrophic forgetting of previous knowledge. As a first, we introduce a system addressing both problems, where prototypes evolve continually in a shared latent space, enabling learning and prediction at any point in time. To facilitate learning, a novel objective function synchronizes the latent space with the continually evolving prototypes. In contrast to the major body of work in continual learning, data streams are processed in an online fashion without task information and can be highly imbalanced, for which we propose an efficient memory scheme. As an additional contribution, we propose the learner-evaluator framework that i) generalizes existing paradigms in continual learning, ii) introduces data incremental learning, and iii) models the bridge between continual learning and concept drift. We obtain state-of-the-art performance by a significant margin on eight benchmarks, including three highly imbalanced data streams. Code is publicly available.¹

1. Introduction

The prevalence of data streams in contemporary applications urges systems to learn in a continual fashion. Autonomous vehicles, sensory robot data, and video streaming yield never-ending streams of data, with abrupt changes in the observed environment behind every vehicle turn, robot entering a new room, or camera cut to a subsequent scene. Alas, learning from streaming data is far from trivial due to these changes, as learning schemes for neural networks have been developed assuming identically and independently dis-

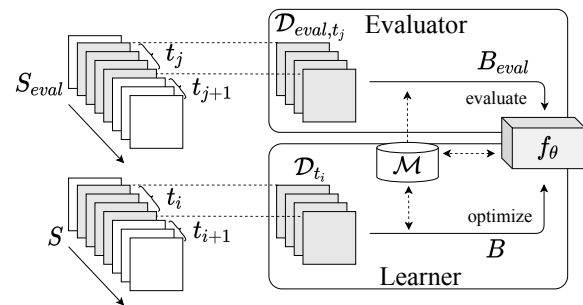


Figure 1: Overview of the *learner-evaluator* framework, overcoming the static training and testing paradigms by explicitly modelling continual optimization and evaluation from data streams in the *learner* and *evaluator*. The framework generalizes to both continual learning and concept drift with resources transparently defined as the horizon \mathcal{D} and operational memory \mathcal{M} .

tributed (iid) data. When this assumption is violated, catastrophic forgetting of previously acquired knowledge occurs [11], emanating a trade-off between neural stability to retain the current state of knowledge and neural plasticity to swiftly adopt the new knowledge [14]. The search for this delicate balance is considered the main hurdle for continual learning systems.

Although a lot of progress has been established in the literature, often strong assumptions apply [9], impeding applicability for real-world systems. Static training and testing paradigms prevail, whereas a true continual learner should enable both training and testing simultaneously and independently. Therefore, we propose the *learner-evaluator* framework, redefining perspective on existing paradigms in the field. Within this framework, we introduce *data incremental learning*, enabling completely task-free learning and evaluation.

Furthermore, we introduce *Continual Prototype Evolution* (CoPE), a new online data incremental learner with prototypes perpetually representing the most salient features of the class population, shifting the catastrophic forgetting

¹<https://github.com/matttdl/ContinualPrototypeEvolution>

problem from the full network parameter space to the lower-dimensional latent space. As a first, our prototypes evolve continually with the data stream, enabling learning and evaluation at any point in time. Similar to representativeness heuristics in human cognition [22], the class prototypes are the cornerstone for nearest neighbor classification. Additionally, our system is robust to highly imbalanced data streams by the combination of replay with a balancing memory population scheme. We find batch information in the latent space to be crucial in the challenging non-stationary and online processing regime, which we incorporate in our novel *pseudo-prototypical proxy loss*.

2. The learner-evaluator framework

To date, the paradigms of task, class, and domain incremental learning [46] dominate the continual learning literature. However, strong and differing assumptions often lead to confusion and overlap between implementations of these definitions. Furthermore, the concept of a static training and testing phase is still ubiquitous, whereas continual learning systems should enable both aspects continually and independently. Therefore, we propose a generalizing framework which disentangles the continual learning system into two entities: the *learner* and the *evaluator*. Figure 1 presents an overview of the framework.

The learner optimizes predicting function $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ parameterized by θ , mapping the input space \mathcal{X} to the target output space \mathcal{Y} . The learner receives data samples $(\mathbf{x}_i, \mathbf{y}_i)$ from stream S and has simultaneous access to the *horizon* \mathcal{D} , i.e. the observable subset of stream S which can be processed for multiple iterations. Data sample i is constituted by input feature $\mathbf{x}_i \in \mathcal{X}$ and corresponding (self-)supervision signal \mathbf{y}_i for which the output space for classification is defined as a discrete set of observed classes $\mathcal{Y}_i \leftarrow \mathcal{Y}_{i-1} \cup \{\mathbf{y}_i\}$. To manage memory usage and to enable multiple updates and stochasticity in the optimization process, updates for θ are typically performed based on a small-scale *processing batch* $B \subseteq \mathcal{D}$. The data and size of the horizon \mathcal{D} are determined by the specific setup or application, ranging from standard offline learning with $\mathcal{D} = S$ to online continual learning with $\mathcal{D} = B$. Furthermore, the learner might need additional resources after observing data from $B \subseteq \mathcal{D}$, such as stored samples or model copies, confined by the *operational memory* \mathcal{M} .

The evaluator acts independently from the learner by evaluating f_θ with horizon \mathcal{D}_{eval} from the evaluation stream S_{eval} , with small-scale processing batches $B_{eval} \subseteq \mathcal{D}_{eval}$. This stream can contain yet unobserved concepts by the learner in S to measure zero-shot performance. The framework provides leeway for the concept distributions in S_{eval} being either static or dynamically evolving, determining how performance of the learner is measured. On the one hand, static concept distributions can measure the degree to

which the knowledge of learned concepts is preserved, as commonly used in continual learning. On the other hand, evolving concept distributions measure performance for the current distribution in horizon \mathcal{D}_{eval} only, where concepts might drift from their original representation, also known as concept drift [38]. Evaluation can occur asynchronously on-demand or periodically with periodicity ρ determining the resolution of the evaluation samples.

Task, class, and domain incremental learning are based on the composition of the learner’s observable stream subset in horizon \mathcal{D}_t . In all three scenarios, the learner receives data of the form $(\mathbf{x}_i, \mathbf{y}_i, t_i)$ with t_i the indicator for horizon $\mathcal{D}_{t=t_i}$ of the current task, set of classes, or domain. In *task incremental learning* the learner’s horizon spans all data of a given task, i.e. $\mathcal{D}_t = \{(\mathbf{x}_i, \mathbf{y}_i, t_i) \in S \mid t = t_i\}$ [9, 47]. Moreover, this setup assumes the evaluator to receive data in the same form as the learner, hence having explicit access to t_i . This confines prediction to a task-specific subset of output nodes, also referred to as a head. In *class incremental learning*, the horizon is determined based on a subset of classes \mathcal{Y}_t , i.e. $\mathcal{D}_t = \{(\mathbf{x}_i, \mathbf{y}_i, t_i) \in S \mid \mathbf{y}_i \in \mathcal{Y}_{t=t_i}\}$, for which the learner requires t_i to identify the transition between subsequent class subsets \mathcal{Y}_{t_i} in S [37, 6, 43, 50]. However, the evaluator considers the entire output space without the need for identifier t . *Domain incremental learning* holds the same assumptions as class incremental learning, although with concepts drifting from one domain to the other for a typically fixed output space, exemplified by the widely used permuted-MNIST setup [13]. To measure to what extent knowledge is preserved, continual learning assumes in general that for new t_i ’s a fixed set of evaluation samples is added to the evaluator’s horizon $\mathcal{D}_{eval, t=t_i}$, drawn mutually exclusively from the same distribution as the learner’s horizon $\mathcal{D}_{t=t_i}$.

Data incremental learning is a more general paradigm we introduce to facilitate learning from any data stream, with no assumption on the horizon or the order of observing the data. In contrast to existing paradigms, when the *learner* observes horizon \mathcal{D} of data stream S , data incremental learning does not disclose an identifier t . Consequently, there is no explicit indication to which subset of the stream is being observed in the horizon \mathcal{D} . Therefore, the learner either processes observed data directly in an online fashion with processing batch $B = \mathcal{D}$, or infers an implicit identifier t from statistics in stream S . Similar to class and domain incremental learning, the *evaluator* operates without t on the full output space. The data incremental learning paradigm endows continual learning systems with increased practical use, as real-world streaming applications typically lack supervision signal t . Moreover, making this identifier available may depend on arbitrary choices, such as how to delineate tasks, introducing unwanted bias.

Data incremental learning resembles online learning in its

	evaluator		learner	
	sample	sample	horizon \mathcal{D}	iid
online learning	$(\mathbf{x}_i, \mathbf{y}_i)$	$(\mathbf{x}_i, \mathbf{y}_i)$	batch ($\mathcal{D} = B$)	✓
continual learning				
task incr.	$(\mathbf{x}_i, \mathbf{y}_i, t_i)$	$(\mathbf{x}_i, \mathbf{y}_i, t_i)$	task ($\mathcal{D}_{t=t_i}$)	✗
class incr.	$(\mathbf{x}_i, \mathbf{y}_i)$	$(\mathbf{x}_i, \mathbf{y}_i, t_i)$	class subset ($\mathcal{D}_{t=t_i}$)	✗
domain incr.	$(\mathbf{x}_i, \mathbf{y}_i)$	$(\mathbf{x}_i, \mathbf{y}_i, t_i)$	domain ($\mathcal{D}_{t=t_i}$)	✗
data incr.	$(\mathbf{x}_i, \mathbf{y}_i)$	$(\mathbf{x}_i, \mathbf{y}_i)$	any subset ($B \leq \mathcal{D} < S$)	✗

Table 1: The learner-evaluator framework disentangling online learning with four continual learning paradigms. Data incremental learning is fully independent of identifier t as in online learning, but evades the iid assumption.

independence of identifier t . However, online learning typically makes the rigorous assumption of learning from an iid data stream [41] or does not focus on preventing catastrophic forgetting [34, 26], whereas data incremental learning disregards the iid assumption and additionally generalizes beyond the online learning horizon ($\mathcal{D} = B$) by allowing to store larger subsets of the data stream ($\mathcal{D} \subset S$). The data incremental learning paradigm is also referred to as task-free [3] or task-agnostic [54] learning, but is here defined based on transition t of the horizon, which is more general than the task-based definition generating possible ambiguity of how a task is specifically defined. Note that online continual learning mostly refers to the four continual learning paradigms with horizon $\mathcal{D} = B$, which can also be viewed as learning for a single epoch for each task, class subset, or domain. Table 1 compares data incremental learning with online learning and the three main continual learning paradigms.

3. Prior work

Continual learning systems are able to learn with limited resources from data streams prone to severe distribution shifts. The main body of works presumes the presence of tasks, which divide the data streams into large discrete subsets, and are indicated to the learner with a task identifier [25, 30, 53, 1, 10, 52]. Replay methods retain representative data for observed data distributions, currently unavailable in the learner’s horizon \mathcal{D} . The replay data is either obtained directly from operational memory \mathcal{M} with stored samples [37, 31, 48, 18] or generated using generative models [42, 23, 39, 50]. **GEM** [31] uses replay in a constraint optimization perspective to project gradients towards a local joint task optimum. **iCaRL** [37] employs exemplars to distill knowledge [20] to the learner from a previous model version, with new class exemplars stored in a queue to optimally represent the class mean in feature space. The prototypes are then used for nearest neighbor prediction by the evaluator, in the same vein as [16]. Nonetheless, all three works strongly rely on task identifier t for the learner,

mostly unavailable for real-world data streams. Moreover, in both prototypical approaches [37, 16] the prototypes remain static between the given task transitions and become outdated. Consequently, before using the evaluator they exhaustively recalculate the prototypes based on all exemplars in memory. *In contrast, our prototypes evolve in an online fashion with the data stream and remain representative for the continual learner and evaluator at all times.*

Recent works focus on online data incremental learning (Section 2) in which the learner operates completely task-free. **Reservoir** [49] is a replay baseline with strong potential to outperform continual learning methods [7]. Samples are stored in replay memory \mathcal{M}_r with probability $|\mathcal{M}_r|/n$, with n the number of observed samples and fixed replay buffer size $|\mathcal{M}_r|$. **MIR** [2] extends Reservoir sampling with a loss-based retrieval strategy, with the cost of additional forward passes and a model copy to attain the losses for a subset of samples. The Reservoir buffer population approximately follows the data stream distribution, severely deteriorating the performance of underrepresented tasks in imbalanced data streams, as shown in Section 6.2. An alternative memory population scheme is used in **GSS** [4] by extending the GEM constraint optimization perspective to an instance-based level. GSS adds samples to the buffer based on their gradients, whereas GEM requires the number of tasks and the task transitions to divide memory equally over all tasks a priori. Another memory population is used in iCaRL, incrementally subdividing over all classes after learning a task by iteratively adding observed samples from the horizon \mathcal{D}_t to optimally approximate the class mean in feature space. As this is computationally expensive, recent works explore other balancing schemes [24, 8], where we propose a simple but effective class-based Reservoir scheme with uniform retrieval.

Another branch of works are parameter isolation methods [9], allocating parameters to subsets of the data. Several task incremental works assign parameters based on the task identifier [32, 40]. A new line of work instead focuses on data incremental model expansion. **CURL** [36] enables task-free and unsupervised adaptation using a multi-component variational auto-encoder, with generative replay from a model copy avoiding forgetting in the current model. **CN-DPM** [29] allocates data subsets to expert networks following a Dirichlet process mixture. In contrast to these capacity expansion based methods, CoPE evades unbound allocation of resources, as the memory and network capacity are fixed with the replay memory dynamically subdivided over categories occurring in the data stream.

Besides the focus on continual learning in this work, our learner-evaluator framework generalizes to concept drift as well [38], for which we refer to an overview in [45, 12]. Further, in deep embedding learning most commonly pairs [15] and triplets [17] of samples are considered in contrastive




Figure 2: Main setup. The learner updates network f_θ and prototypes $\mathbf{p}^y, \forall y \in \mathcal{Y}$ continually. The PPP-loss encourages inter-class variance (red arrows) and reduces intra-class variance (green arrows).

losses, whereas other works use batch information in lifted structure embeddings [35] or instance-wise softmax embeddings [51]. These approaches fully depend on the batch size, whereas our pseudo-prototypical proxy loss aggregates both decoupled prototypes and the additional batch pseudo-prototypes to defy class interference in the latent space. Learning prototypical representations also shows promising results in few-shot learning [44].

4. Continual Prototype Evolution

The online data incremental learning setup of the CoPE learner is described in Figure 2. Embedding network f_θ maps processing batch B , composed of samples in horizon \mathcal{D} from the non-iid data stream S and operational memory \mathcal{M} , to low-dimensional \mathbb{R}^d latent space, followed by a nearest neighbor classifier. We enforce $\|f_\theta(\mathbf{x}_i)\| = 1$ with an L2 normalization layer. \mathcal{M} is subdivided in replay memory \mathcal{M}_r and prototypical memory \mathcal{M}_p . CoPE comprises three main components: continually evolving representations, balanced replay and the pseudo-prototypical proxy (PPP) loss. In the following, we discuss these components and formalize the optimal choice of prototype, with \mathbf{f}_i^c denoting latent space projection $f_\theta(\mathbf{x}_i^c)$ for an instance \mathbf{x}_i of class c . For the full algorithm, we refer to Appendix A.

4.1. Evolving representations

Each observed class $c \in \mathcal{Y}$ is represented by a slowly progressing prototype \mathbf{p}^c in operational memory \mathcal{M}_p . The nearest neighbor classifier finds the most similar prototype for the given query \mathbf{x}_i , predicting $c^* = \arg \max_{c \in \mathcal{Y}} \mathbf{f}_i^c \mathbf{p}^c$. Similar to [33, 37], the class-prototype approximates the center of mass in the latent space, which we formally justify in Section 4.4. The main crux with storing representations is to prevent them from becoming obsolete as the embedding network evolves. This is further complicated by the shifting data distributions in the non-stationary regime, incurring catastrophic forgetting. Experience replay from a buffer \mathcal{M}_r is a well known approach to address this forgetting. In our setup the replayed exemplars provide additional information about the current state of the embedding space, enabling rehearsal to rectify approximation \mathbf{p}^c to the true center of

mass. Concretely, the sampled batch B_n equals the horizon \mathcal{D} from data stream S and joins batch $B_{\mathcal{M}}$ of equal size from memory \mathcal{M}_r , constituting B as $B_n \cup B_{\mathcal{M}}$. However, updating the prototypes by fully relying on features extracted from B incurs an unstable optimization process as the representative prototypes depend on stochastic sampling of the class distributions. Therefore, we design the prototypes to evolve continually with a high momentum based update for each observed batch, aiming to stabilize the impetuous changes in the data stream:

$$\begin{aligned} \mathbf{p}^c &\leftarrow \alpha \mathbf{p}^c + (1 - \alpha) \bar{\mathbf{p}}^c, \\ \text{s.t. } \bar{\mathbf{p}}^c &= \frac{1}{|B^c|} \sum_{\mathbf{x}^c \in B^c} f_\theta(\mathbf{x}^c) \end{aligned} \quad (1)$$

with momentum parameter $\alpha \in [0, 1]$, the batch subset $B^c = \{(\mathbf{x}_i, y_i = c) \in B\}$ of class c , and $\bar{\mathbf{p}}^c$ the corresponding center of mass in latent space for the current batch. Due to triangle inequality $\bar{\mathbf{p}}^c$ is no longer unit length and requires to be L2-normalized after the update in Eq.(1). We empirically validate the effectiveness of high momentum with $\alpha \approx 1$ in the ablation study in Appendix D.

4.2. Balanced replay

Similar to [37, 8], the total buffer size $|\mathcal{M}_r|$ is equally divided over the number of observed classes $|\mathcal{Y}|$ in a dynamic fashion. This scheme ensures consistent buffer capacity for all classes, making memory allocation independent of the data stream characteristics. As S is typically highly imbalanced in real-world scenarios, this memory scheme prevents classes to be eradicated from the buffer and assumes equal importance to represent each class at all times. Consequently, random retrieval from the buffer resembles sampling an iid replay batch, assuming sufficient samples per class are observed in S . Furthermore, each class-specific replay memory \mathcal{M}_r^c can simply capture a random subset of its parent class distribution to approximate its center of mass. This avoids computationally expensive herding techniques as in iCaRL [37], which would require recalculation of the feature means for each change of the memory size or network parameters.

4.3. Pseudo-Prototypical Proxy loss

The learner optimizes f_θ to project an instance $\mathbf{f}_i^c \in \mathbb{R}^d$ of class c close to its corresponding prototype \mathbf{p}^c in the latent space. As the prototype acts as a surrogate for the class mean in latent space, the cluster population has a common reference point to reduce intra-class variance, and we can enforce inter-class variance by remaining distant from the other class prototypes. Additionally, due to the embedding architecture we can use intrinsic information of the batch samples in the latent space. Therefore, we exploit the supervision signal \mathbf{y}_i in a sample $(\mathbf{x}_i, \mathbf{y}_i) \in B$ not only to indicate which class \mathbf{x}_i belongs to, but also to make the distinction between positive and negative pairs in B . Consequently, we can define one-against-all subsets for an instance of class c , with positives from the same class in $B^c = \{(\mathbf{x}_i, \mathbf{y}_i = c) \in B\}$ and negatives in B^k . Starting from these sets, the prototypical attractor and repeller sets for an instance \mathbf{x}_i^c are constituted with the class prototype \mathbf{p}^c and the other instances in B . First, the other instances of class c act as pseudo-prototypes $\hat{\mathbf{p}}^c$ in attractor set $\mathbb{P}_i^c = \{\mathbf{p}^c\} \cup \{\hat{\mathbf{p}}_j^c = f_\theta(\mathbf{x}_j^c) \mid \forall \mathbf{x}_j^c \in B^c, i \neq j\}$. To avoid gradient redundancy of the pseudo-prototypes, attractor sets remain static and are excluded from gradient computation. Second, the samples of other classes $\mathbf{x}_j^k \in B^k$ should instead avoid both \mathbf{x}_i^c in latent space and the class representative \mathbf{p}^c , defined by repeller set $\mathbb{U}_i^c = \{\mathbf{p}^c, \hat{\mathbf{p}}_i^c = f_\theta(\mathbf{x}_i^c)\}$. The attractor set \mathbb{P}_i^c for \mathbf{x}_i^c facilitates a decrease in intra-class variance with \mathbf{p}^c safeguarding the absence of positive batch pairs with $1 \leq |\mathbb{P}_i^c| \leq |B^c|$, whereas the repeller \mathbb{U}_i^c exploits \mathbf{x}_i^c and corresponding prototype as a reference point to increase inter-class variance. To incorporate the attractor and repeller sets, we formulate a binary classification problem similar to [51], with the joint probability that instance \mathbf{x}_i^c is predicted as class c and instances $\mathbf{x}_j^k \in B^k$ not being predicted as class c in

$$P_i = P(c|\mathbf{x}_i^c) \prod_{\mathbf{x}_j^k} (1 - P_i(c|\mathbf{x}_j^k)) \quad (2)$$

with the assumption of independence between \mathbf{x}_i^c and \mathbf{x}_j^k being recognized as c . We define the expected posterior probabilities for the attractor and repeller sets of instance \mathbf{x}_i^c respectively as

$$P(c|\mathbf{x}_i^c) = \mathbb{E}_{\hat{\mathbf{p}}^c \in \mathbb{P}_i^c} [P(c|\mathbf{f}_i^c, \hat{\mathbf{p}}^c)] \quad (3)$$

$$P_i(c|\mathbf{x}_j^k) = \mathbb{E}_{\hat{\mathbf{p}}^c \in \mathbb{U}_i^c} [P(c|\mathbf{f}_j^k, \hat{\mathbf{p}}^c)] \quad (4)$$

with $\hat{\mathbf{p}}^c$ a proxy for the latent mean of class c in

$$P(c|\mathbf{f}, \hat{\mathbf{p}}^c) = \frac{\exp(\mathbf{f}^T \hat{\mathbf{p}}^c / \tau)}{\exp(\mathbf{f}^T \hat{\mathbf{p}}^c / \tau) + \sum_{k \neq c} \exp(\mathbf{f}^T \mathbf{p}^k / \tau)} \quad (5)$$

where temperature τ controls the concentration level of the distribution [20] and assuming a cosine similarity metric $\mathbf{f}_i^T \mathbf{f}_j$ with vectors normalized to unit length. We reformulate the objective in Eq.(2) as loss function \mathcal{L} by taking the negative log-likelihood and summing over all the instances in B , which approximates the true joint probability with assumed independent pairs in the batch:

$$\mathcal{L} = -\frac{1}{|B|} \left[\sum_i \log P(c|\mathbf{x}_i^c) + \sum_i \sum_{\mathbf{x}_j^k} \log(1 - P_i(c|\mathbf{x}_j^k)) \right]. \quad (6)$$

4.4. Optimal prototypes

We update prototypes to approximate the mean of the parent distribution in Eq.(1). This assumption is optimal for Bregman divergences for which the cluster mean is shown to have minimal distance to its population [5]. This Bregman divergence is defined for a differentiable, strictly convex function φ as

$$d_\varphi(\mathbf{f}_i, \mathbf{f}_j) = \varphi(\mathbf{f}_i) - \varphi(\mathbf{f}_j) - (\mathbf{f}_i - \mathbf{f}_j)^T \nabla \varphi(\mathbf{f}_j), \quad (7)$$

for which the squared Euclidean distance with $\varphi(\mathbf{f}) = \|\mathbf{f}\|^2$ is a canonical example. The squared Euclidean distance is proportional to the cosine distance with vectors normalized to unit length: $\frac{1}{2}\|\mathbf{f}_i - \mathbf{f}_j\|^2 = 1 - \cos \angle(\mathbf{f}_i, \mathbf{f}_j)$. As the PPP-loss for Eq.(5) requires a similarity measure instead of a distance measure, we employ the complementary normalized cosine similarity $\cos \angle(\mathbf{f}_i, \mathbf{f}_j) = \mathbf{f}_i^T \mathbf{f}_j$ with $\|\mathbf{f}_i\| = \|\mathbf{f}_j\| = 1$. Besides the desirable cluster-mean property of its complement, this metric is also efficient for calculating the full batch similarity matrix using matrix multiplication libraries.

5. Experiments

This work examines five balanced data streams and 15 highly imbalanced variants based on Split-MNIST, Split-CIFAR10 and Split-CIFAR100, from which two low-capacity balanced setups are discussed in Appendix E. The learner is presented a data stream S , constituted by a sequence of tasks, each delineated by a subset of classes from the original dataset. Although the learner in CoPE is completely ignorant to the notion of task, this setup enables comparing to methods requiring task boundaries such as GEM and iCaRL. The evaluator uses a held-out dataset of static concepts in S_{eval} , evaluating with the subset of seen concepts \mathcal{Y} in \mathcal{D}_{eval} using the accuracy metric. The learner of CoPE processes data online with $B_n = \mathcal{D}$ in the data incremental setup. To enable fair comparison, this allows per-task processing of 1 epoch for methods requiring task boundaries with $B \subset \mathcal{D}_t$. We use vanilla stochastic gradient descent with a limited processing batch size $|B_n|$ of 10

as in [31, 4, 29]. All results are averaged over 5 different network initializations. Appendix details the full setup with additional experiments.

Balanced data streams contain a similar amount of data per task. We consider three benchmarks. First, **Split-MNIST** constitutes the MNIST [28] handwritten digit recognition dataset with 60k training samples, split into 5 tasks according to pairs of incrementing digits. Second, **Split-CIFAR10** considers the CIFAR10 [27] dataset, subdivided into 5 tasks with 2 labels each, where each task entails 10k training samples. Third, **Split-CIFAR100** is a variant of the CIFAR dataset with 100 different classes. The 50k training samples are subdivided in 20 tasks of 2.5k samples as in [31, 29]. For all datasets the evaluator considers the entire original test subset for S_{eval} .

Imbalanced data streams introduce a more realistic scenario without equality assumptions on the task durations in S . This addresses a common weakness in the literature mostly balancing the data streams artificially. Besides the imbalanced Split-MNIST setup [4], we introduce two novel and more challenging benchmarks based on Split-CIFAR10 and Split-CIFAR100, where data stream S comprises significantly more data in task T_i , denoted by $S(T_i)$. Split-MNIST and Split-CIFAR10 have respectively 2k and 4k samples in T_i , whereas tasks T_j for $j \neq i$ contain factor 10 less data for five variants $S(T_i)$, $\forall i \in \{1, \dots, 5\}$. Split-CIFAR100 defines T_i with 2.5k samples and 1k for the remaining tasks, with variants $i \in \{1, 5, 10, 15, 20\}$.

Architectures. MNIST setups use an MLP with 2 hidden layers of 400 units with 2k memories for the balanced setup as in [21, 29, 47], and 100 units with $|\mathcal{M}_r| = 0.3k$ for the imbalanced setup as in [4]. CIFAR setups use a slim version of Resnet18 [19] with a 1k memory size for CIFAR10 [4, 29], and 5k for CIFAR100 [31].

Methods compared to CoPE entail 11 baselines, with details on prior work discussed in Section 3. The upper reference point for performance when relaxing the challenging non-iid property of continual learning is set by **iid-online & iid-offline**. The learner shuffles the full data stream S to ensure the iid property, for which *iid-online* trains a single epoch and *iid-offline* multiple epochs. In contrast, the **finetune** learner considers non-iid data stream S sequentially, but optimizes solely for the new batch which typically results in worst-case catastrophic forgetting. **CoPE-CE** is a reference point for the merits of a prototypical approach by solely using the CoPE memory and sampling scheme, but with a typical cross-entropy loss and softmax classifier. **GEM** and **iCaRL** are standard replay methods considered in a class incremental setup, with the learner requiring knowledge about task boundaries. For online data incremental learning, we consider the **reservoir**, **MIR** and greedy **GSS** replay baselines, with **CURL** and **CN-DPM** instead relying on model expansion.




Figure 3: Accuracies over buffer sizes $|\mathcal{M}_r|$ for balanced Split-MNIST (*top*), Split-CIFAR10 (*center*) and Split-CIFAR100 (*bottom*) sequences. The legend reports averages over all observed buffer sizes. '*' indicates learner using task information.

6. Results and discussion

6.1. Balanced data streams

The results for the three balanced data streams in Table 2 consistently report state-of-the-art results for CoPE. The difficulty for learning online is reflected in the discrepancy of performance between iid-offline and iid-online, indicating increasing difficulty for a minimal 2% discrepancy for Split-




Figure 4: Accuracy (%) for imbalanced Split-MNIST (left), Split-CIFAR10 (center) and Split-CIFAR100 (right) sequences. The legend reports average accuracies over all the sequence variations.

MNIST, raised to 20% for Split-CIFAR10, and culminating to 30% in Split-CIFAR100. For Split-MNIST CoPE closes the gap with iid-online performance by 0.7% compared to main competitors GEM and DN-CPM. Furthermore, in the more challenging Split-CIFAR10 setup CoPE significantly increases the gained margin to 3.7%. In the most challenging Split-CIFAR100, CN-DPM, Reservoir and MIR are able to perform on par with the iid-online baseline, however, CoPE establishes an improvement of at least 1.5% over all four baselines.

	Split-MNIST	Split-CIFAR10	Split-CIFAR100
iid-offline	98.44 ± 0.02	83.02 ± 0.60	50.28 ± 0.66
iid-online	96.57 ± 0.14	62.31 ± 1.67	20.10 ± 0.90
finetune	19.75 ± 0.05	18.55 ± 0.34	3.53 ± 0.04
GEM	93.25 ± 0.36	24.13 ± 2.46	11.12 ± 2.48
iCaRL	83.95 ± 0.21	37.32 ± 2.66	10.80 ± 0.37
CURL [36]	92.59 ± 0.66	–	–
DN-CPM [29]	93.23 ± 0.09	45.21 ± 0.18	20.10 ± 0.12
reservoir	92.16 ± 0.75	42.48 ± 3.04	19.57 ± 1.79
MIR	93.20 ± 0.36	42.80 ± 2.22	20.00 ± 0.57
GSS	92.47 ± 0.92	38.45 ± 1.41	13.10 ± 0.94
CoPE-CE	91.77 ± 0.87	39.73 ± 2.26	18.33 ± 1.52
CoPE (ours)	93.94 ± 0.20	48.92 ± 1.32	21.62 ± 0.69

Table 2: The three balanced data stream accuracies (%) with standard deviation over 5 initializations. For expansion-based methods CURL and DN-CPM we report results from their original work.

Compared to balanced replay with standard cross-entropy (CoPE-CE), the prototypical approach (CoPE) proves effective with significant gains of 2.2%, 9.2% and 3.3% respectively over the three benchmarks. Except for GEM in Split-MNIST, class incremental learning methods GEM and iCaRL are not competitive in the online setting and additionally require from the setup to reveal an identifier t to the learner. From the expansion-based methods DN-CPM is competitive, whereas CURL is more suited for unsupervised learning and lags behind. Although Reservoir and extension

MIR perform on par with iid-online for Split-CIFAR100, the imbalanced experiments in Section 6.2 show that full reservoir-based population of the buffer strongly relies on the assumption of equally sized tasks, which is unlikely to occur in real-world data streams.

Buffer size ablation study in Figure 3 shows CoPE to prevail over all sizes of replay buffer \mathcal{M}_r compared to other replay methods, extending robustness to low capacity regimes. Although iCaRL shows competitive results for low capacity, CoPE scales with growing capacity leading to significantly outperforming iCaRL with 11% in Split-MNIST ($2k$) and Split-CIFAR100 ($5k$), and 17% in Split-CIFAR10 ($2k$).

6.2. Imbalanced data streams

Results for the highly imbalanced data stream benchmarks are reported in Figure 4. CoPE significantly outperforms all baselines in the three scenarios, with low standard deviation for the 15 variants indicating robustness over a wide spectrum of imbalanced sequences. GSS outperforms Reservoir and MIR for Split-MNIST, in correspondence with results in [4], whereas loss-based retrieval in MIR has significant gains for the challenging Split-CIFAR100 setting. However, CoPE surpasses both GSS and MIR for all three benchmarks, and on top of that operates profusely more resource efficient as discussed in Appendix C. The balancing memory scheme in CoPE-CE highly improves Reservoir over imbalanced Split-MNIST and Split-CIFAR10 variants with 10.8% and 3.4% respectively, and performs on par for Split-CIFAR100 where balancing over 100 classes with limited batch size proves more difficult. Although CoPE and CoPE-CE share memory and retrieval schemes, the prototypical CoPE surpasses the cross-entropy based CoPE-CE with 4.0%, 2.9% and 6.7% respectively on the three benchmarks, indicating the merits of the PPP-loss and continually evolving prototypes. Figure 5 compares the CoPE and CoPE-CE confusion matrices at the end of learning, with CoPE-CE exhibiting high plasticity for the later observed classes, whereas CoPE better preserves the recall over early learned classes, hence effectively alleviating catastrophic forgetting.

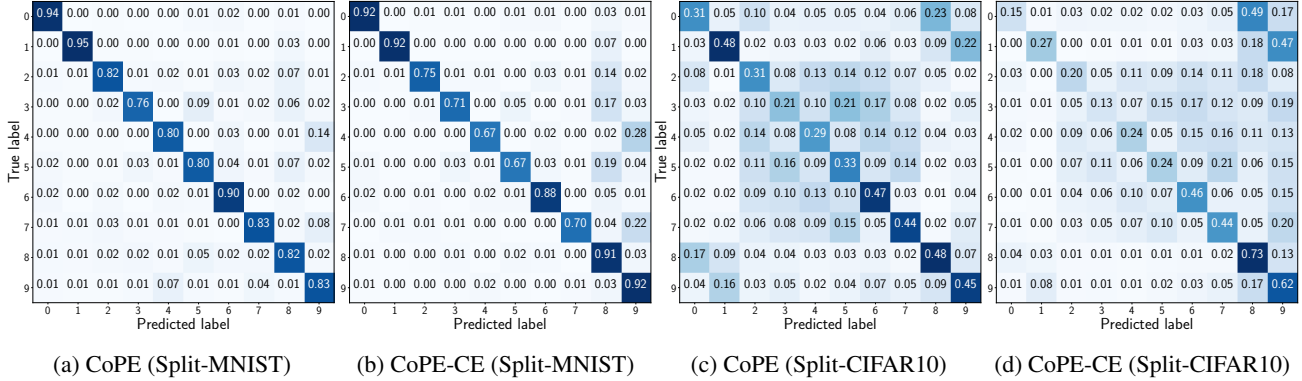


Figure 5: CoPE and CoPE-CE confusion matrices at the end of learning averaged over all variations $S(T_i)$ for the imbalanced Split-MNIST setup in (a) and (b), and Split-CIFAR10 in (c) and (d). Results are normalized per row.

	PPP-loss		Batch Size $ B_n $				
	<i>incl.</i> $\hat{\mathbf{p}}$	<i>excl.</i> $\hat{\mathbf{p}}$	10 (Online)	20	50	100	200
Split-MNIST	93.9 ± 0.2	92.4 ± 0.6	93.9 ± 0.2	93.9 ± 0.6	93.7 ± 0.3	93.1 ± 0.6	89.3 ± 0.5
Split-CIFAR10	48.9 ± 1.3	41.3 ± 2.0	48.9 ± 1.3	48.4 ± 1.9	43.4 ± 2.7	37.4 ± 3.0	37.0 ± 1.3
Split-CIFAR100	21.6 ± 0.7	16.3 ± 0.7	21.6 ± 0.7	21.7 ± 0.7	16.5 ± 0.4	13.8 ± 0.5	11.2 ± 0.4

Table 3: Accuracies (%) for ablating pseudo-prototypes $\hat{\mathbf{p}}$ in the PPP-loss and varying batch size.

6.3. PPP-loss analysis

The PPP-loss exploits relational information in the small processing batch B by introducing pseudo-prototypes $\hat{\mathbf{p}}$. This leads to questioning to what extent the pseudo-prototypes actually contribute to the quality of the embedding and how this relates to the batch size. We examine both inquiries in Table 3 for the three balanced data streams by comparing inclusion and exclusion of the pseudo-prototypes $\hat{\mathbf{p}}$ in the PPP-loss, and extending the batch size $|B_n|$. First, including the pseudo-prototypes significantly improves overall performance, especially for the harder CIFAR-based data streams. Although both setups use batch information to update the prototypes following Eq.(1), it seems crucial to use additional pseudo-prototypes in the PPP-loss to improve latent space quality. Second, results for smaller batch sizes of 10 and 20 are very similar, and deteriorate towards increasing sizes. The PPP-loss implements the expectation over the prototype and the pseudo-prototypes, assuming uniform distribution in Eq.(3) and Eq.(4). Although this assumption impedes significance of the prototype for increasingly higher batch sizes, it results in ideal robustness for small online processing batches. Small batches maintain the additional benefit of more frequent prototype and network parameter updates for the same amount of processed data.

7. Conclusion

The novel learner-evaluator framework introduced a new perspective on current paradigms in continual learning. To

overcome the standard paradigm of static training and testing phases, we explicitly model continual optimization and evaluation in the *learner* and *evaluator* respectively. We formalized the required resources as the *horizon* \mathcal{D} , containing the simultaneously available data of the data stream, and the *operational memory* \mathcal{M} for operation of the learning algorithm. Transitions in the horizon $\mathcal{D}_t \rightarrow \mathcal{D}_{t+1}$ enable a uniform differentiation between existing incremental learning paradigms and the horizon size encloses the range from online ($\mathcal{D} = B$) to offline ($\mathcal{D} = S$) learning. Using the framework, we defined the task-free *data incremental learning* paradigm, requiring no additional information on the horizon’s identifier t for both the learner and evaluator. In this challenging setup, we proposed Continual Prototype Evolution (CoPE) as a prototypical solution to learn online from non-stationary data streams. As a first, CoPE prevents the prototypes becoming obsolete in an ever evolving representation space, while using the prototypes to combat catastrophic forgetting. The three main components are continually evolving prototypes, a novel pseudo-prototypical proxy loss, and an efficient balancing replay scheme, proving remarkably effective over 11 baselines in both balanced and highly imbalanced benchmarks. We hope to encourage further research in online data incremental learning.

Acknowledgements

The first author was funded by the KUL C1 Macchina project and Flanders AI Research Program.

References

- [1] 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)*, pages 139–154, 2018. 3
- [2] Rahaf Aljundi, Eugene Belilovsky, Tinne Tuytelaars, Laurent Charlin, Massimo Caccia, Min Lin, and Lucas Page-Caccia. Online continual learning with maximal interfered retrieval. In *Advances in Neural Information Processing Systems*, pages 11849–11860, 2019. 3
- [3] Rahaf Aljundi, Klaas Kelchtermans, and Tinne Tuytelaars. Task-free continual learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11254–11263, 2019. 3
- [4] Rahaf Aljundi, Min Lin, Baptiste Goujaud, and Yoshua Bengio. Gradient based sample selection for online continual learning. In *Advances in Neural Information Processing Systems*, pages 11816–11825, 2019. 3, 6, 7
- [5] Arindam Banerjee, Srujana Merugu, Inderjit S Dhillon, and Joydeep Ghosh. Clustering with bregman divergences. *Journal of machine learning research*, 6(Oct):1705–1749, 2005. 5
- [6] Francisco M Castro, Manuel J Marín-Jiménez, Nicolás Guil, Cordelia Schmid, and Karteek Alahari. End-to-end incremental learning. In *Proceedings of the European conference on computer vision (ECCV)*, pages 233–248, 2018. 2
- [7] Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and Marc’Aurelio Ranzato. Continual learning with tiny episodic memories. *arXiv preprint arXiv:1902.10486*, 2019. 3
- [8] Aristotelis Chrysakis and Marie-Francine Moens. Online continual learning from imbalanced data. *Proceedings of Machine Learning Research*, 2020. 3, 4
- [9] Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Ales Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *arXiv preprint arXiv:1909.08383*, 2019. 1, 2, 3
- [10] Matthias De Lange, Xu Jia, Sarah Parisot, Ales Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. Unsupervised model personalization while preserving privacy and scalability: An open problem. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14463–14472, 2020. 3
- [11] Robert M French. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 3(4):128–135, 1999. 1
- [12] João Gama, Indrè Žliobaitė, Albert Bifet, Mykola Pechenizkiy, and Abdelhamid Bouchachia. A survey on concept drift adaptation. *ACM computing surveys (CSUR)*, 46(4):1–37, 2014. 3
- [13] Ian J Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211*, 2013. 2
- [14] Stephen Grossberg. *Studies of mind and brain : neural principles of learning, perception, development, cognition, and motor control*. Boston studies in the philosophy of science 70. Reidel, Dordrecht, 1982. 1
- [15] Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’06)*, volume 2, pages 1735–1742. IEEE, 2006. 3
- [16] Xu Han, Yi Dai, Tianyu Gao, Yankai Lin, Zhiyuan Liu, Peng Li, Maosong Sun, and Jie Zhou. Continual relation learning via episodic memory activation and reconsolidation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6429–6440, 2020. 3
- [17] Ben Harwood, BG Kumar, Gustavo Carneiro, Ian Reid, Tom Drummond, et al. Smart mining for deep metric learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2821–2829, 2017. 3
- [18] Tyler L Hayes, Kushal Kafle, Robik Shrestha, Manoj Acharya, and Christopher Kanan. Remind your neural network to prevent catastrophic forgetting. In *European Conference on Computer Vision*, pages 466–483. Springer, 2020. 3
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 6
- [20] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 3, 5
- [21] Yen-Chang Hsu, Yen-Cheng Liu, Anita Ramasamy, and Zsolt Kira. Re-evaluating continual learning scenarios: A categorization and case for strong baselines. *arXiv preprint arXiv:1810.12488*, 2018. 6
- [22] Daniel Kahneman and Amos Tversky. Subjective probability: A judgment of representativeness. *Cognitive psychology*, 3(3):430–454, 1972. 2
- [23] Nitin Kamra, Umang Gupta, and Yan Liu. Deep generative dual memory network for continual learning. *arXiv preprint arXiv:1710.10368*, 2017. 3
- [24] Chris Dongjoo Kim, Jinseo Jeong, and Gunhee Kim. Imbalanced continual learning with partitioning reservoir sampling. *arXiv preprint arXiv:2009.03632*, 2020. 3
- [25] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017. 3
- [26] J Zico Kolter and Marcus A Maloof. Dynamic weighted majority: An ensemble method for drifting concepts. *The Journal of Machine Learning Research*, 8:2755–2790, 2007. 3
- [27] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 6
- [28] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 6

- [29] Soochan Lee, Junsoo Ha, Dongsu Zhang, and Gunhee Kim. A neural dirichlet process mixture model for task-free continual learning. In *International Conference on Learning Representations*, 2019. 3, 6, 7
- [30] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017. 3
- [31] David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. In *Advances in Neural Information Processing Systems*, pages 6467–6476, 2017. 3, 6
- [32] Arun Mallya and Svetlana Lazebnik. Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7765–7773, 2018. 3
- [33] Thomas Mensink, Jakob Verbeek, Florent Perronnin, and Gabriela Csurka. Distance-based image classification: Generalizing to new classes at near-zero cost. *IEEE transactions on pattern analysis and machine intelligence*, 35(11):2624–2637, 2013. 4
- [34] Claire Monteleoni and Tommi Jaakkola. Online learning of non-stationary sequences. In S. Thrun, L. Saul, and B. Schölkopf, editors, *Advances in Neural Information Processing Systems*, volume 16. MIT Press, 2004. 3
- [35] Hyun Oh Song, Yu Xiang, Stefanie Jegelka, and Silvio Savarese. Deep metric learning via lifted structured feature embedding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4004–4012, 2016. 4
- [36] Dushyant Rao, Francesco Visin, Andrei Rusu, Razvan Pascanu, Yee Whye Teh, and Raia Hadsell. Continual unsupervised representation learning. In *Advances in Neural Information Processing Systems*, pages 7645–7655, 2019. 3, 7
- [37] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 2001–2010, 2017. 2, 3, 4
- [38] Jeffrey C Schlimmer and Richard H Granger. Beyond incremental processing: Tracking concept drift. In *AAAI*, pages 502–507, 1986. 2, 3
- [39] Ari Seff, Alex Beatson, Daniel Suo, and Han Liu. Continual learning in generative adversarial nets. *arXiv preprint arXiv:1705.08395*, 2017. 3
- [40] Joan Serra, Didac Suris, Marius Miron, and Alexandros Karatzoglou. Overcoming catastrophic forgetting with hard attention to the task. *arXiv preprint arXiv:1801.01423*, 2018. 3
- [41] Shai Shalev-Shwartz et al. Online learning and online convex optimization. *Foundations and trends in Machine Learning*, 4(2):107–194, 2011. 3
- [42] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. In *Advances in Neural Information Processing Systems*, pages 2990–2999, 2017. 3
- [43] Konstantin Shmelkov, Cordelia Schmid, and Karteek Alahari. Incremental learning of object detectors without catastrophic forgetting. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3400–3409, 2017. 2
- [44] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in neural information processing systems*, pages 4077–4087, 2017. 4
- [45] Alexey Tsymbal. The problem of concept drift: definitions and related work. *Computer Science Department, Trinity College Dublin*, 106(2):58, 2004. 3
- [46] Gido M van de Ven and Andreas S Tolias. Generative replay with feedback connections as a general strategy for continual learning. *arXiv preprint arXiv:1809.10635*, 2018. 2
- [47] Gido M van de Ven and Andreas S Tolias. Three scenarios for continual learning. *arXiv preprint arXiv:1904.07734*, 2019. 2, 6
- [48] Eli Verwimp, Matthias De Lange, and Tinne Tuytelaars. Rehearsal revealed: The limits and merits of revisiting samples in continual learning. In *Proceedings of the IEEE International Conference on Computer Vision*, 2021. 3
- [49] Jeffrey S Vitter. Random sampling with a reservoir. *ACM Transactions on Mathematical Software (TOMS)*, 11(1):37–57, 1985. 3
- [50] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, Zhengyou Zhang, and Yun Fu. Incremental classifier learning with generative adversarial networks. *arXiv preprint arXiv:1802.00853*, 2018. 2, 3
- [51] Mang Ye, Xu Zhang, Pong C Yuen, and Shih-Fu Chang. Unsupervised embedding learning via invariant and spreading instance feature. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6210–6219, 2019. 4, 5
- [52] Lu Yu, Bartłomiej Twardowski, Xialei Liu, Luis Herranz, Kai Wang, Yongmei Cheng, Shangling Jui, and Joost van de Weijer. Semantic drift compensation for class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6982–6991, 2020. 3
- [53] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3987–3995. JMLR. org, 2017. 3
- [54] Chen Zeno, Itay Golan, Elad Hoffer, and Daniel Soudry. Task agnostic continual learning using online variational bayes. *arXiv preprint arXiv:1803.10123*, 2018. 3