

NAPA-VQ: Neighborhood Aware Prototype Augmentation with Vector Quantization for Continual Learning

Tamasha Malepathirana Damith Senanayake
Saman Halgamuge
Dept. of Mechanical Engineering
The University of Melbourne

{tamasha.malepathirana, damith.senanayake, saman.halgamuge}@unimelb.edu.au

Abstract

Catastrophic forgetting; the loss of old knowledge upon acquiring new knowledge, is a pitfall faced by deep neural networks in real-world applications. Many prevailing solutions to this problem rely on storing exemplars (previously encountered data), which may not be feasible in applications with memory limitations or privacy constraints. Therefore, the recent focus has been on Non-Exemplar based Class Incremental Learning (NECIL) where a model incrementally learns about new classes without using any past exemplars. However, due to the lack of old data, NECIL methods struggle to discriminate between old and new classes causing their feature representations to overlap. We propose NAPA-VQ: Neighborhood Aware Prototype Augmentation with Vector Quantization, a framework that reduces this class overlap in NECIL. We draw inspiration from Neural Gas to learn the topological relationships in the feature space, identifying the neighboring classes that are most likely to get confused with each other. This neighborhood information is utilized to enforce strong separation between the neighboring classes as well as to generate old class representative prototypes that can better aid in obtaining a discriminative decision boundary between old and new classes. Our comprehensive experiments on CIFAR-100, TinyImageNet, and ImageNet-Subset demonstrate that NAPA-VQ outperforms the State-of-the-art NECIL methods by an average improvement of 5%, 2%, and 4% in accuracy and 10%, 3%, and 9% in forgetting respectively. Our code can be found in <https://github.com/TamashaM/NAPA-VQ.git>.

1. Introduction

The achievements of deep neural networks over the years have grown significantly, and their efficiency and applicability have been demonstrated by numerous state-of-the-art

works [29, 23, 35, 20]. However, a major requirement for the optimal operation of gradient-based optimization – the ubiquitous learning paradigm – is the data samples being independently and identically distributed (IID) [21]. This assumption may not hold in real data owing to various factors such as the addition of new classes of data, the removal of old data due to memory or availability constraints, and the changes in the data-generating phenomena (concept drift). As a result, neural networks may experience catastrophic forgetting where the network forgets the previously learned knowledge upon acquiring new knowledge [25].

“Continual Learning” is a field of research pursuing mechanisms to mitigate this forgetting [13]. In this manuscript, we focus on one paradigm of continual learning, named Class Incremental Learning (CIL) [38]. In CIL, a neural network is trained over a series of tasks and at each task, the network learns a new set of classes. At any given time, the network should classify between all learned classes thus far. Among the techniques proposed for CIL, rehearsal-based methods have demonstrated promising results in mitigating forgetting by storing exemplars (old samples) and reusing them while learning new tasks [38, 5, 36]. However, such storage is not always possible due to memory limitations and privacy constraints [45]. Therefore, we focus on Non-Exemplar based CIL (NECIL), a more pragmatic yet challenging scenario, which attempts to preserve the old knowledge without storing any exemplars [57, 58].

NECIL methods often struggle with overlapping old and new class representations due to the unavailability of exemplars, resulting in catastrophic forgetting. [57]. While prototypes of old classes in the deep feature space are a viable alternative to reusing exemplars [57, 56], if not properly generated, the class boundaries refined using such prototypes tend to be muddled, causing confusion between the old and new classes, and in turn, leading to catastrophic forgetting. To overcome this limitation, we propose NAPA-VQ: Neighborhood-Aware Prototype Augmentation with

Vector Quantization framework. NAPA-VQ not only proposes a novel way to create prototypes of old classes by considering class neighborhoods but also incorporates a novel quantization mechanism to create clearer class boundaries by enforcing a strong separation between the neighboring classes. This reduction in representation overlap effectively mitigates catastrophic forgetting. NAPA-VQ builds on the principles of unsupervised Neural Gas (NG) [49] and supervised Learning Vector Quantization (LVQ) [26] to facilitate this neighborhood awareness and enforce more discriminative boundaries.

NAPA-VQ contains two components: (I) a Neighborhood-aware Vector Quantizer (NA-VQ) and (II) a Neighborhood-aware Prototype Augmenter (NA-PA). NA-VQ learns the topology of the feature space manifold \mathcal{Z} which is the output of a deep feature extractor (*e.g.* ResNet), identifying the neighboring classes that share similar features and are hence prone to get confused with each other. This knowledge of the neighboring classes and their class distributions is utilized by NA-VQ to increase their separability and by NA-PA for old class prototype augmentation, *i.e.*, to generate surrogate exemplars to facilitate the retention of old information when new classes are being learned.

To summarize,

- We propose an improved supervised vector quantization method to discretize the latent space and improve class separation.
- We propose a prototype augmentation method that uses the topological information of classes in the latent space to avoid confusion between classes and catastrophic forgetting.
- We demonstrate the utility of the above two contributions combined in NECIL, obtaining superior performance compared to the existing NECIL methods on CIFAR-100, TinyImageNet, ImageNet-Subset, and ImageNet-1k datasets.

2. Related work

2.1. Incremental learning

The techniques proposed to combat catastrophic forgetting can be broadly categorized into three [13]. (1) The regularisation-based methods that add an extra regularisation loss term either to penalize changes to the network parameters that are important for previous tasks [25, 33, 54, 1] or to distill knowledge from previous tasks to the current task [32, 16, 55]. (2) The parameter isolation-based methods that assign each task with an isolated set of parameters to prevent task interference either by dynamically increasing the network capacity [39, 52, 2] or by masking previous

task parameters in a fixed size network [17, 42, 37]. Although parameter isolation methods are effective in overcoming catastrophic forgetting, they experience either a linear increase in network parameters or a decrease in capacity per task as the number of tasks grows [44]. (3) The rehearsal-based methods that store a small subset of previous task data to either retrain [38, 5, 8] or constrain the optimisation [36, 7, 3] during the learning of new tasks in order to retain the discriminability between old and new classes. However, these methods also encounter pitfalls due to memory limitations, and other pragmatic concerns such as privacy or consent issues when storing samples. An alternative to rehearsal-based methods is “pseudo-rehearsal”, which involves training a generative model to mimic past task distributions [43, 41]. Despite the encouraging results, generative models are computationally expensive to train [11] and are also prone to catastrophic forgetting [48]. This motivated the development of NECIL strategies that neither depends on real nor fake past samples [57, 58, 34, 53].

NECIL methods benefit from powerful feature extractors learning transferable features across tasks, as demonstrated by SDC [53], which showed that embedding networks suffer significantly less from catastrophic forgetting. PASS [57] also showed that self-supervised learning alleviates task-level overfitting. Furthermore, to maintain the decision boundaries of previously learned classes, PASS introduced a class-mean prototype augmentation technique based on Gaussian noise. While this technique aids in the retention of old information, it can be further improved by leveraging the knowledge of the distribution of classes in the feature space. Accordingly, IL2A [56] proposed storing covariance matrices to retain class variations, but this approach can be memory intensive. SSRE [58] proposed a dynamic structure reorganization strategy to retain and transfer knowledge between tasks along with a prototype selection mechanism that utilizes an up-sampling technique of non-augmented class-mean prototypes. Similar to these approaches, we also store the mean prototype, while proposing a new method to augment them. To this end, we use the topological connections derived from an NG-like vector quantization to generate prototypes that lie within the shared feature regions of the confusing classes which aid in establishing better class discrimination.

2.2. Vector quantization

Vector Quantization (VQ), a technique used to discretize a continuous data space into a finite set of “coding vectors” (CVs) was popularised with the advents of Self-organizing Maps (SOMs) [27]. In addition to quantizing the data manifold, a SOM captures a topological mapping from data to the CVs. Neural Gas (NG) networks [49, 19], on the other hand, were introduced to address a shortcoming of the original SOM by allowing a generic graph structure rather than

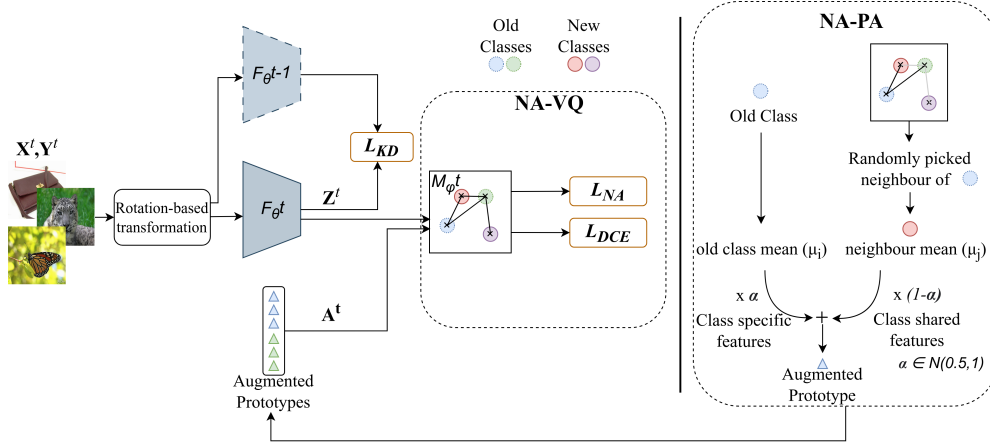


Figure 1. Illustration of NAPA-VQ. Data from the current task t are augmented using a rotation-based technique [57] and are fed to the feature extractor. The obtained feature representations (\mathbf{Z}^t) and the NA-PA generated old class representative prototypes (\mathbf{A}^t) are sent to the vector quantizer (NA-VQ) to identify and repel confusing classes, establishing better discrimination in the feature space. Knowledge Distillation (L_{KD}) is used to minimize the feature drift across tasks.

a fixed lattice structure. In NG, the CVs are adjusted to capture the data-dense regions, and the edges between these CVs are formed based on their proximity. These edges and the CVs form a graph that approximates the topology of the data manifold.

Coding vector-based learning can be traced back to the K-nearest neighbor (K-NN) algorithm [18]. For instance, Learning Vector Quantization (LVQ) was proposed to derive the CVs used in a 1-NN classifier [26, 40]. Despite common roots, LVQ algorithms and unsupervised VQ algorithms such as SOM differ in their primary usages of CVs; the unsupervised algorithms attempt to obtain a set of CVs to best represent the data while LVQ algorithms attempt to reduce the misclassification rate by focusing on the decision boundaries between classes. These complementary properties allow us to combine unsupervised and supervised VQ methods [22, 12] to obtain CVs to both reduce the misclassification rate and represent the data distribution [27].

Multiple studies explored the integration of the hierarchical feature-extracting capability of deep feature extractors with VQ [12, 50, 4] which were also later adapted to Continual Learning. TPCIL [46] proposed to retain the topology of the feature space to preserve old knowledge over the increments. IDLVQ [9] proposed to adapt a margin-based loss for the task of few-shot class incremental learning (FSCIL) – a special case of CIL, therefore not directly transferrable to CIL/NECIL – to create a large margin between classes to mitigate overlap. TOPIC [47] was also proposed for the FSCIL setting with the aim of preserving old knowledge by stabilizing a NG network. We highlight that changes to the topology are possible due to the inevitable feature drift occurring over incremental steps thus a method that uses both augmented prototypes and new data to update the topological graph between CVs is war-

rented.

3. Methodology

3.1. Preliminaries and notations

In CIL, a model is continually trained over a series of tasks, where at each task the model learns a set of new classes that are distinct from the previously learned classes. At any given time, the model should classify between samples from all classes seen thus far. The training data at task t is denoted as $\mathbf{D}^t = \{\mathbf{X}^t, \mathbf{Y}^t\}$ where $\mathbf{X}^t = \{x_i^t\}_{i=1}^{N^t}$ is the set of input images and $\mathbf{Y}^t = \{y_i^t | y_i^t \in C^t\}_{i=1}^{N^t}$ are their target labels. N^t and C^t correspond to the number of samples and the set of classes at task t . We further define $P^t = \sum_{j=1}^t |C^j|$ as the total number of classes seen at the end of task t .

3.2. Overview of the framework

An illustration of our framework is shown in Fig. 1 which consists of a feature extractor (F_θ), a Vector Quantizer (NA-VQ), and a prototype augments (NA-PA). F_θ is used to obtain the feature space of the input data. NA-VQ quantizes this feature space by learning a set of CVs named M_ϕ such that m_i is associated with class i . These CVs are trained to effectively improve the inter-class variance and reduce intra-class variance in the feature space reducing the representational overlap in classes. The module parameters θ and ϕ are shared across all tasks but are updated continually with the data at the current task, thus t denotes the states of the parameters at each task. Consequently, F_{θ^t} and M_{ϕ^t} refer to the states of the feature extractor and the set of CVs at task t . The goal of task t is to jointly update $F_{\theta^{t-1}}$ and $M_{\phi^{t-1}}$ using \mathbf{D}^t to obtain F_{θ^t} and M_{ϕ^t} . At incremental tasks (when $t > 0$), a set of old class prototypes (\mathbf{A}^t)

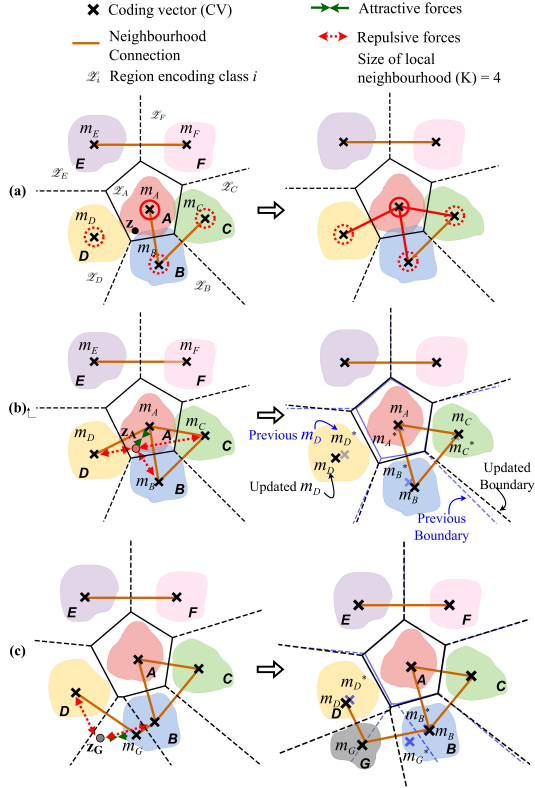


Figure 2. Conceptual illustration of NA-VQ. Each color represents a single class. (a) Given z , edges are created between the closest CV (solid circle) and the next $K - 1$ closest CVs (dotted circles) to approximate the local topology around z . (b) Class A is overlapped with classes B, C, and D thus a sample from class A (z_A) could be misclassified into B, C, or D. NA-VQ pulls z_A and m_A together and pushes z_A and direct neighbors of m_A (m_B, m_C , and m_D) away from each other, reducing the class overlap. Closer the CV to z_A , the higher these forces. The edge between m_A and m_D is pruned due to edge weakening over time. (c) When a new class G is introduced, m_G is inserted randomly and refined through attractive and repulsive forces over the iterations, identifying a less-overlapping space in the feature space.

generated using our NA-PA technique is used alongside \mathbf{D}^t to retain the discrimination between old and new classes.

3.3. Neighborhood-Aware Vector Quantizer (NA-VQ)

At task t , we extract the features of \mathbf{X}^t by computing $\mathbf{Z}^t \subset \mathbb{R}^n$ using F_{θ^t} , where n is the dimensionality of the feature space (Eq. 1).

$$\mathbf{Z}^t = F_{\theta^t}(\mathbf{X}^t) \quad (1)$$

We assume \mathbf{Z}^t lies in a feature space manifold \mathcal{Z} which captures the topological properties between the classes. NA-VQ partitions \mathcal{Z} into disjoint regions such that \mathcal{Z}_i contains the feature distribution of class i . We encode \mathcal{Z}_i using m_i , i.e., $M_{\phi^t} = \{m_i \in \mathbb{R}^n | i = 1, \dots, P^t\}$ contains

one CV per class. The algorithm can be summarised into two iterative steps (1) the topology approximation and (2) the CV adaptation. These steps are performed concurrently, i.e., the topology learned depends on the CV adaptation and vice versa. A conceptual illustration of NA-VQ is shown in Fig. 2.

(1) Topology approximation. Inspired by NG [49], the topology of \mathcal{Z} is approximated by learning an undirected graph $G^t = \langle M_{\phi^t}, E^t \rangle$, where CVs are the nodes and the topological connections between the CVs are the edges represented by the adjacency matrix E . At the beginning of task t , $M_{\phi^{t-1}}$ and E^{t-1} are extended to M_{ϕ^t} and E^t to accommodate task t . Specifically, C^t new randomly initialized CVs are inserted without any edges that link to, from, or between them. During the learning of task t , E^t is modified by adding, decaying, and removing edges between the CVs.

Concretely, given $(x, y) \in (\mathbf{X}^t, \mathbf{Y}^t)$ and its feature representation $z \in \mathbf{Z}^t$, we calculate the distance from z to each CV in M_{ϕ^t} : $D = \{d(z, m_i) | i = 1, \dots, P^t\}$ where d is the Euclidean distance. D is then sorted to assign a rank to each CV: $R = \{r_i | i = 1, \dots, P^t\}$ such that $\forall i, d(z, m_{r_i}) < d(z, m_{r_{i+1}})$. Next, the edges between the CVs are formed based on a connectivity factor denoted as K . A higher value of K leads to a denser graph, whereas a lower value of K results in a sparser graph. Specifically, edges are created between the closest CV and the next $K - 1$ closest CVs, i.e., $e_{r_1, j} \leftarrow 1$ for $j = r_2, \dots, r_K$. Consequently, CVs that lie within the high-density regions of \mathcal{Z} develop edges, allowing to identify classes that share similar features in \mathcal{Z} . Since the CVs get updated over time, the edges created in a previous iteration may become obsolete as their endpoints may have moved. To remove such edges we employ an edge decaying mechanism. All edges from the closest CV are decayed by a constant multiplier $\epsilon \in (0, 1)$ so that the edges created at a previous iteration that no longer fall within a high-density region are weakened, i.e., $e_{r_1, j} \leftarrow e_{r_1, j} * \epsilon$ for $j = 1, \dots, P^t$. If the edge strength goes below a predefined e_{\min} , such edges would be pruned, i.e., if $e_{i, j} < e_{\min}$, $e_{i, j} \leftarrow 0$.

(2) CV adaptation. During the learning of task t , the class label y in $(x, y) \in (\mathbf{X}^t, \mathbf{Y}^t)$ is employed to adapt CVs in M_{ϕ^t} to improve the discriminability in the feature space. Concretely, we create attractive forces between z and m_y and repulsive forces between z and any other “confusing CVs” of m_y . The confusing CVs of m_y (referred to as N_y^-) are the direct neighbors of m_y as determined by graph G^t in Step (1).

$$N_y^- = \{m_i | e_{y, i} > 0 \text{ and } y \neq i \text{ and } i = 1, \dots, P^t\} \quad (2)$$

Any $m_i \in N_y^-$ is considered to be close to m_y , thus, there is a high likelihood that m_i will be mistaken for m_y as the winner CV when presented with sample x . To counteract this, we propose a Neighborhood-Adaptation loss (L_{NA})

(Eq. 3) to bring z and m_y closer together while pushing z and N_y^- farther apart.

$$L_{NA}^t(\mathbf{X}^t, \mathbf{Y}^t; M_{\phi^t}, F_{\theta^t}) = \sum_{\substack{x, y \in \\ \mathbf{X}^t, \mathbf{Y}^t}} ReLU(d(z, m_y) - d_{\text{neigh-}})$$

$$\text{where } d_{\text{neigh-}} = \sum_{m_i \in N_y^-} W_{m_i} \times d(z, m_i) \text{ and } z = F_{\theta^t}(x) \quad (3)$$

L_{NA} computes the difference between $d(z, m_y)$ and $d_{\text{neigh-}}$ when $d(z, m_y) > d_{\text{neigh-}}$, the scenario where x is most likely to be misclassified. $d(z, m_y)$ is the distance from z to its correct CV m_y and $d_{\text{neigh-}}$ is a linear combination of the distances from z to N_y^- . The weight given to $m_i \in N_y^-$ (W_{m_i}) is calculated using a monotonic decaying function (Eq. 4) to reduce the impact of m_i on L_{NA} as the distance between m_i and z increases.

$$W_{m_i} = \frac{e^{-\beta \times d(z, m_i)}}{\sum_{m_j \in N_y^-} e^{-\beta \times d(z, m_j)}} \text{ where } \beta > 0 \quad (4)$$

To avoid the repelling forces in L_{NA} from diverging any $m_i \in N_y^-$ from their respective class distributions, we take three steps. First, the repelling forces are distributed over multiple confusing CVs. Second, $ReLU$ prevents any adjustments to the CVs when $d(z, m_y) < d_{\text{neigh-}}$, since it is unlikely that this scenario would lead to misclassification. Finally, Distance-based cross-entropy loss (L_{DCE}) [51], as described next, is used to encourage m_i to more accurately represent the distribution of class i which aids in minimizing the distance between the feature representations of class i and m_i .

The probability of sample x , belonging to class y can be measured by the distance between $z = F_{\theta^t}(x)$ and m_y [51], i.e., $p(y|x) \propto -d(z, m_y)$. Considering the non-negative and sum-to-one properties of the probability, we can define $p(y|x)$ using a softmax function as shown in Eq. 5 where τ is a temperature parameter.

$$p(y|x) = \frac{e^{-\frac{d(z, m_y)}{\tau}}}{\sum_{j=1}^{P^t} e^{-\frac{d(z, m_j)}{\tau}}} \quad (5)$$

Since the true probability distribution is a one hot encoded vector, we can define L_{DCE} for \mathbf{X}^t as Eq. 6.

$$L_{DCE}^t(\mathbf{X}^t, \mathbf{Y}^t; M_{\phi^t}, F_{\theta^t}) = \sum_{x, y \in \mathbf{X}^t, \mathbf{Y}^t} -\log p(y|x) \quad (6)$$

These two steps are conducted concurrently at each mini-batch gradient update during the optimization of θ and ϕ . By backpropagating the gradients calculated for L_{NA} and L_{DCE} through F_{θ} and M_{ϕ} , we effectively reduce the class

overlapping in the feature space and establish a more discriminative decision boundary between classes. During the incremental steps ($t > 0$), we freeze the old class representative CVs but update their topological connections using both the augmented prototypes (Sec. 3.4) and \mathbf{D}^t .

3.4. Neighborhood-Aware Prototype Augmenter (NA-PA)

In NECIL, we cannot directly compute L_{NA} or L_{DCE} for the previous task samples. While the class mean in the feature space serves as a central representative point for each old class, it does not capture the class variance. Thus, if the model relies solely on class-means without any augmentation, it may overfit to the class mean and forget to distinguish samples of the same class that are close to but not exactly equal to the class mean. Augmentation mitigates this issue by generating multiple representative points for each old class, considering the underlying class variance. Therefore, we store representative mean prototypes and augment them on the fly with NA-PA during the learning of each new task. Specifically, we generate prototypes in the regions where the model is uncertain about the class labels. We consider pairs of classes that the model may confuse between using the topological graph G^t established in Sec. 3.3, and generate prototypes to distinguish between these pairs of classes, increasing the quality of the prototypes used for retaining old knowledge.

Inspired by the work [10], we identify that the features of each class can be decomposed into a class-specific component (features placed closer to the class mean) and a class-shared component (features that lie between itself and another class). For a given old class, the samples near the boundaries shared with its confusing classes have the biggest impact on recovering good decision boundaries because they are closer to the regions where the model is uncertain about the class labels. Thus, NA-PA generates augmented representations of the old classes by fusing the class-specific features from the old classes with the class-generic/shared features from their confusing classes to create such high-impact prototypes (Fig. 1). At task $t (> 0)$, an augmented prototype a_i of old class i is generated as shown in Eq. 7 by fusing the mean prototype of class i (μ_i) with the mean prototype of class j (μ_j), a randomly picked neighbor of class i . By varying α , we ensure the augmented prototypes are composed with varying degrees of uncertainty, i.e., the lower the α , the higher the uncertainty. Here, we are using the mean prototypes instead of CVs as these CVs are adjusted continually to reduce misclassification and may potentially be positioned near class boundaries (away from class means) for improved discrimination. We determined that class-means serve as better representatives of the class-specific features than the CVs.

$$a_i = \alpha \mu_i + (1 - \alpha) \mu_j \text{ where } m_j \in N_i^- \text{ \& } \alpha \sim \mathcal{N}(0.5, 1) \quad (7)$$

\mathbf{A}^t is the collection of augmented prototypes and their class labels representing all old classes up to task t . \mathbf{A}^t is used alongside \mathbf{D}^t in NA-VQ for topology approximation as well CV adaptation. Specifically, we calculate \hat{L}_{NA}^t and \hat{L}_{DCE}^t using \mathbf{A}^t .

$$\hat{L}_{DCE}^t(\mathbf{A}^t; M_{\phi^t}) = \sum_{a_i \in \mathbf{A}^t} -\log p(i|a_i) \quad (8)$$

$$\hat{L}_{NA}^t(\mathbf{A}^t; M_{\phi^t}) = \sum_{a_i \in \mathbf{A}^t} \text{ReLU}(d(a_i, m_i) - d_{\text{neigh}^-}) \quad (9)$$

3.5. Knowledge distillation

As F_{θ}^t gets updated continually, the actual feature distributions of old classes drift away from their original distributions. To mitigate this drift we incorporate a feature-level knowledge distillation (L_{KD}) [57, 56] that attempts to align the feature spaces of the current and the previous models.

$$L_{KD}^t(\mathbf{X}^t | F_{\theta}^{t-1}, F_{\theta}^t) = d(F_{\theta}^{t-1}(\mathbf{X}^t), F_{\theta}^t(\mathbf{X}^t)) \quad (10)$$

The total loss used in our framework is shown below and λ_1 and λ_2 are loss weights (See Supp. Materials for explanations on loss weights)

$$L_{total}^t = L_{DCE}^t + L_{NA}^t \text{ when } t = 0 \quad (11)$$

$$L_{total}^t = L_{DCE}^t + \lambda_1 \hat{L}_{DCE}^t + L_{NA}^t + \hat{L}_{NA}^t + \lambda_2 L_{KD}^t \text{ when } t > 0$$

3.6. Rotation-based data augmentation

In order to learn richer features, we transform the training data using the same rotation-based approach used in [57]. Concretely, the training samples are rotated by 90, 180 and 270 degrees to generate 3 new pseudo-classes, learning $4P^t$ classes instead of P^t classes at the training stage. However, the classification occurs only between the original classes during the evaluation stage.

3.7. Classification

We perform the nearest CV-based classification. At the end of task t , given a test sample x , we obtain $z = F_{\theta^t}(x)$, calculate the distance from normalized z to each normalized CV and assign the class label of the closest CV to x .

$$y_{pred} = \arg \min_{i=1, \dots, P^t} d\left(\frac{z}{\|z\|}, \frac{m_i}{\|m_i\|}\right) \quad (12)$$

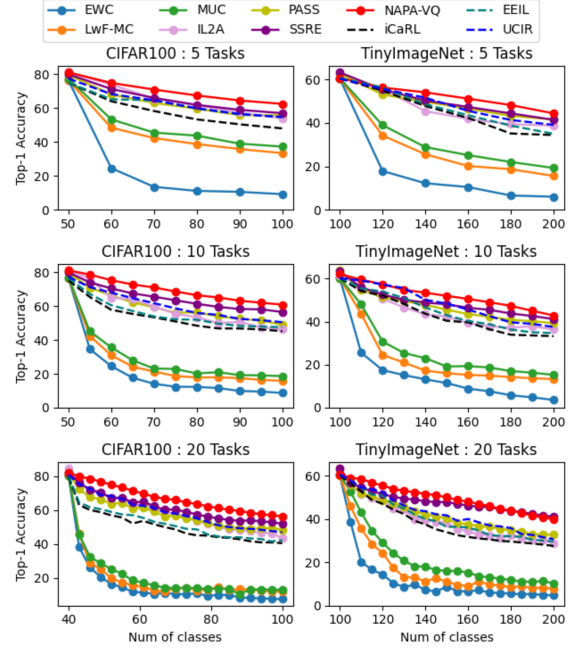


Figure 3. Detailed Accuracy curves showing the Top-1 Accuracy at each incremental step.

4. Experiments

4.1. Datasets

We perform comprehensive experiments using four datasets; CIFAR-100 [28], TinyImageNet [31], ImageNet-Subset [14], ImageNet-1K [14] in three incremental scenarios; $T = 5, 10$ and 20 where T is the number incremental tasks. For comparability, the classes are arranged into tasks using the same fixed random order and division settings as [57, 58] for the first three datasets. For the ImageNet-1K dataset, we train the model on 400 classes for the first task, and equal classes in the rest of the tasks.

4.2. Implementation details

For a fair comparison, we adapted the same backbone architecture, ResNet-18 [23] from [57]. The concrete details of our implementation can be found in the Supp. materials and the [publicly available codebase](#). All the experiments were conducted on the University of Melbourne’s high-performance computing system, Spartan [30].

4.3. Evaluation metrics

In line with previous works [57, 58], we report the standard metrics used to evaluate CIL strategies: Average Accuracy and Average Forgetting. Accuracy at task t is the average accuracy of all the classes that have been learned up to and during task t . **Average accuracy** [6] is the mean accuracy across all the tasks, including the initial task. **Forgetting** at any given time for a task previously encountered,

Table 1. Average Accuracy (%) of NAPA-VQ compared to the top three SOTA using four datasets. The higher the values, the better. T is the number of incremental tasks. Values for the methods with * were extracted from [58]. Our improvement is shown in red.

Method	CIFAR-100			TinyImageNet			ImageNet-Subset			ImageNet-1K		
	T=5	T=10	T=20	T=5	T=10	T=20	T=5	T=10	T=20	T=5	T=10	T=20
PASS *	63.47	61.84	58.09	49.55	47.29	42.07	66.84	61.80	54.46	-	-	-
IL2A	65.61	59.09	58.82	47.02	44.48	39.68	-	-	-	-	-	-
SSRE *	65.88	65.04	61.70	50.39	48.93	48.17	-	67.69	-	-	-	-
NAPA-VQ	70.44 (+4.56)	69.04 (+4)	67.42 (+5.72)	52.77 (+2.38)	51.78 (+2.85)	49.51 (+1.34)	69.15 (+2.31)	68.83 (+1.14)	63.09 (+8.63)	55.11	53.04	45.46

Table 2. Average Forgetting (%) of NAPA-VQ compared to the top three SOTA using four datasets. The lower the values, the better. T is the number of incremental tasks. Values for the methods with * were extracted from [58]. Our improvement is shown in red.

Method	CIFAR-100			TinyImageNet			ImageNet-Subset			ImageNet-1K		
	T=5	T=10	T=20	T=5	T=10	T=20	T=5	T=10	T=20	T=5	T=10	T=20
PASS *	25.20	30.25	30.61	18.04	23.11	30.55	19.66	25.85	30.98	-	-	-
IL2A	28.72	39.86	40.70	19.74	29.90	39.99	-	-	-	-	-	-
SSRE *	18.37	19.48	19.00	9.17	14.06	14.20	-	8.30	-	-	-	-
NAPA-VQ	6.90 (-11.47)	9.65 (-9.83)	9.08 (-9.92)	9.08 (-0.09)	10.81 (-3.25)	9.31 (-4.89)	7.17 (-12.49)	9.67 (+1.37)	14.49 (-16.49)	10.45	10.94	18.23

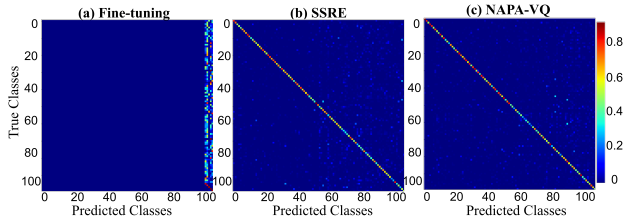


Figure 4. Confusion matrices for Fine-tuning, SSRE and NAPA-VQ for CIFAR-100. Both SSRE and NAPA-VQ reduce the task reccency bias observed in fine-tuning. Along the diagonal, NAPA-VQ has more red patches than SSRE.

is measured by the difference between the maximum accuracy for the task during the learning process and the current accuracy for the same task. **Average Forgetting** at the end of task t is therefore defined as the average of forgetting values for all the tasks learned up to task t [6]. We report the average forgetting at the end of the final task. Additional explanations related to evaluation metrics can be found in Supp. materials.

4.4. Comparison with SOTA

We compare NAPA-VQ with the existing state-of-the-art (SOTA) methods in NECIL, including EWC [25], LwF_MC [38], MUC [34], PASS [57], IL2A [56], and SSRE [58]. We report the average accuracy and average forgetting of NAPA-VQ against the top three performing methods in Ta-

bles 1 and 2. Reported values are the average of three separate runs. As illustrated in Table 1, NAPA-VQ demonstrates an average improvement of 5%, 2%, and 4% in accuracy for CIFAR-100, TinyImageNet, and ImageNet-Subset, respectively, over the best existing NECIL-SOTA technique. The detailed accuracy curves for the compared methods for CIFAR-100, TinyImageNet (in Fig. 3) and ImageNet-Subset (in Supp. Fig. 1), show that NAPA-VQ maintains higher accuracies over the incremental tasks. Moreover, NAPA-VQ exhibits a significant reduction in forgetting by an average of 10%, 3% and 9% for CIFAR-100, TinyImageNet, and ImageNet-Subset, respectively. This reduction in forgetting is prominent when dealing with a larger number of tasks (Table 2). In addition, we provide results on ImageNet-1K, demonstrating the effectiveness of NAPA-VQ on large-scale datasets. Furthermore, we compare NAPA-VQ to traditional exemplar-based methods, iCARL [38], EEIL [5], and UCIR [24] trained using a limited number of exemplars (20) (Fig. 3) and show that NAPA-VQ obtains competitive performance.

4.5. Ablation study

The ablation study is conducted using the CIFAR-100 dataset to demonstrate the impact of NA-VQ and NA-PA (Table 3). We first train a baseline model using the Categorical Cross Entropy loss (L_{CCE}) and L_{KD} (Ablation-1) [58]. We then incorporate NA-VQ (Ablation-2 and 3) and NA-PA (Ablation-5) into the baseline model in sequence.

Table 3. Average Accuracy (%) and Average Forgetting (%) obtained for the ablation study conducted using CIFAR-100 to evaluate the effectiveness of NA-VQ and NA-PA. We refer to the combination of DCE and NA as NA-VQ.

Ablation	Average Accuracy \uparrow			Average Forgetting \downarrow		
	T=5	T=10	T=20	T=5	T=10	T=20
1) KD + CCE (Baseline)	26.21	16.70	10.93	85.29	89.77	93.49
2) KD + DCE	49.83	32.90	17.01	48.84	75.37	81.71
3) KD + DCE + NA \rightarrow KD + NA-VQ	49.31	42.16	42.13	45.03	41.76	30.09
4) KD + NA-VQ + Gaussian-PA	68.84	65.44	62.39	10.78	16.20	18.30
5) KD + NA-VQ + NA-PA \rightarrow NAPA-VQ	70.44	69.04	67.42	6.90	9.65	9.08

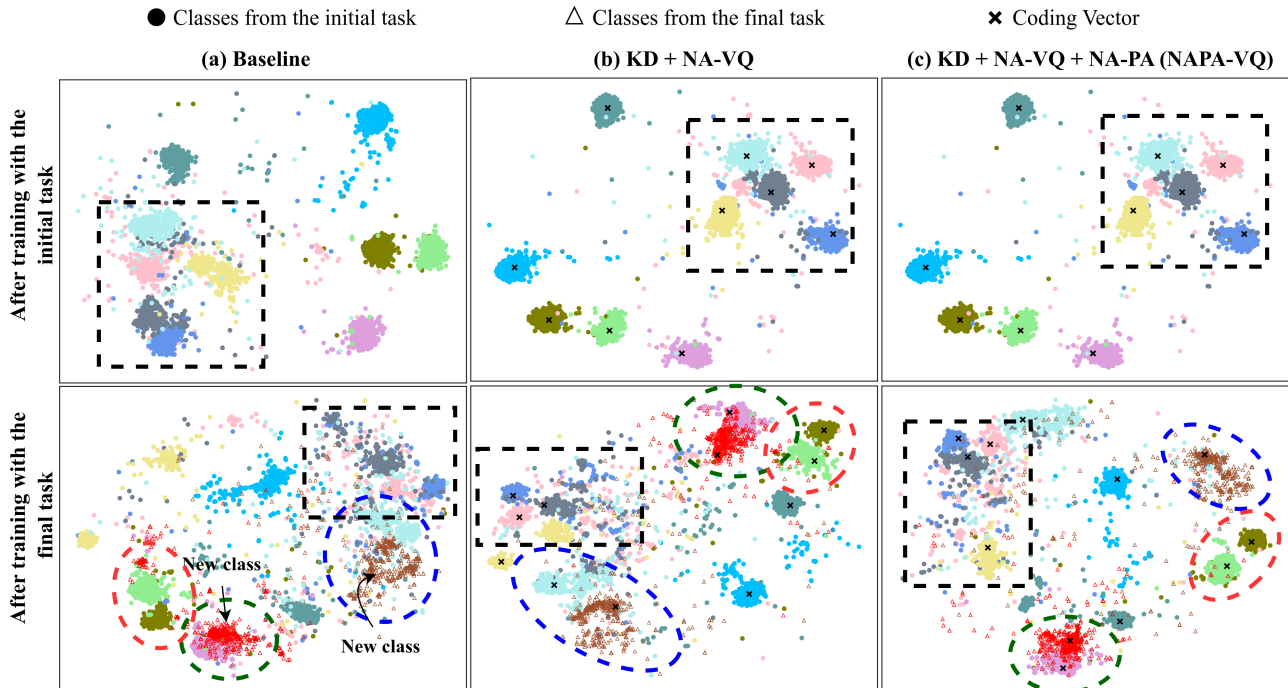


Figure 5. Visualization of the impact of NA-VQ and NA-PA on the feature representations. Each colour represents a single class. The areas highlighted depict the observable differences between experiments. After learning the initial task (first row), NA-VQ integrated models (b and c) reduces much of the class overlap seen in the baseline model (a). The feature space for (b) and (c) are identical at this stage since NA-PA is only applied to the incremental tasks. After learning the final task (second row), NA-VQ integrated model (b) shows better discrimination between the old classes as well as between old and new classes compared to the baseline (a). When NA-PA is integrated, the discrimination between old and new classes improves further.

For comparability, all the models use the same ResNet-18 architecture and the rotation-based data transformation.

NA-VQ combines two losses; L_{DCE} and L_{NA} thus these losses are separately incorporated into the baseline model to comprehend their individual impact. First we substitute L_{DCE} for L_{CCE} (Ablation-2). Although this replacement improves accuracy across all three incremental scenarios compared to the baseline, the forgetting is still significant when $T = 10$ and 20 . In Ablation-3, we add L_{NA} to L_{DCE} and L_{KD} and observe comparable or better accuracy in all three scenarios while significantly reducing forgetting compared to Ablation-2. This show-

cases the effectiveness of L_{NA} in mitigating the interference of feature representations over the incremental steps. Finally, we integrate NA-PA in Ablation-5 and compare results to Ablation-3 which does not employ any prototypes and Ablation-4 which employs Gaussian-augmented prototypes [57]. Our findings show that NA-PA enhances accuracy and reduces forgetting in all incremental scenarios. This improvement can be attributed to the prototypes generated closer to the boundaries which aid in identifying optimal decision boundaries, subsequently reducing the misclassification rate. Moreover, the advantages of NA-VQ and NA-PA are prominent when a larger number of tasks are in-

volved, highlighting the method’s importance.

4.6. Reduced overlapping in the feature space

We visualize the feature space of $T = 10$ scenario in Ablation-1 (Baseline), Ablation-3 (KD + NA-VQ) and Ablation-5 (NAPA-VQ) in Fig. 5 using t-SNE [15] to show the impact of NA-VQ and NA-PA. Specifically, we visualize the feature representations of a randomly selected subset of classes learned during the initial task at two-time points: (1) after the initial task training, and (2) after the final task training with a subset of new classes from the final task. Once the first task is learned, the class-representative features of the baseline model overlap, whereas those of the NA-VQ integrated models are more compact and distinct, reducing the misclassification rate. Once the final task is learned, the overlapping in the baseline model increases creating further confusion between classes. NA-VQ integrated model reduces this overlap due to the more discretized feature space of the old classes and the repulsive forces between old and new classes. When NA-PA is integrated on top of NA-VQ the discrimination between the old classes as well as the discrimination between the old and new classes improve further showing the positive impact of prototypes.

4.7. Comparison of confusion matrices

Fig. 4 shows a comparison between the confusion matrices generated for (1) simple Fine-tuning where a model is trained using CCE loss incrementally without using any strategies to mitigate forgetting, (2) SSRE and (3) NAPA-VQ. The diagonal entries in the matrices represent correct predictions, while off-diagonal entries denote misclassifications. The predictions of Fine-tuning are heavily biased towards the most recent tasks due to the forgetting of old classes. SSRE and NAPA-VQ eliminate much of this bias by correctly classifying both old and new classes. Although quite similar, more red patches are visible along the diagonal in NAPA-VQ compared to SSRE, which explains the higher average accuracies in NAPA-VQ compared to SSRE.

4.8. Impact of the connectivity factor K

To determine the impact of the connectivity factor K on performance, we conducted an experiment on the CIFAR-100 dataset by varying the value of K between 2 and 50, with $K = 2$ being the commonly used heuristic [49, 19]. The results show that as K increases, the performance improves but so does the running time (Supp. Fig. 3). The improved performance can be attributed to a wider neighborhood being considered to improve both decision boundary learning and prototype augmentation. A value of $K = 15$ was found to provide desirable performance without compromising algorithm efficiency.

5. Conclusion

In this manuscript, we proposed NAPA-VQ, a novel method for CIL that does not rely on previous task exemplars to retain old knowledge. Instead, we increase the discriminability of the feature space by using class neighborhood information captured by a topological approximation of the feature space. Furthermore, we show that generating representative prototypes for old classes by borrowing the shared features of their neighboring classes helps to establish good decision boundaries between the areas where the classes tend to overlap. Comprehensive experiments on four benchmarking datasets demonstrate the superiority of our method over existing NECIL methods. While the proposed method exerts no limit on how many CVs to be used per class, we used one CV per class in our experiments. A future study may explore the effect of using a larger number of CVs per class on both the running time and the incremental learning performance.

Acknowledgements T.M. acknowledges Melbourne Graduate Research Scholarship and GCI Women in STEM Student Award support scheme. D.S. and S.H. acknowledge Australian Research Council grant DP210101135. The authors thank Sachith Seneviratne, Maneesha Perera, Rashindrie Perera, and Nisal Ranasinghe for proofreading.

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. 2
- [2] Rahaf Aljundi, Punarjay Chakravarty, and Tinne Tuytelaars. Expert gate: Lifelong learning with a network of experts. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3366–3375, 2017. 2
- [3] Rahaf Aljundi, Min Lin, Baptiste Goujaud, and Yoshua Bengio. Gradient based sample selection for online continual learning. *Advances in neural information processing systems*, 32, 2019. 2
- [4] Sebastian Blaes and Thomas Burwick. Few-shot learning in deep networks through global prototyping. *Neural Networks*, 94:159–172, 10 2017. 3
- [5] 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)*, volume CS, pages 233–248, 2018. 1, 2, 7
- [6] Arslan Chaudhry, Puneet K. Dokania, Thalaiyasingam Ajanthan, and Philip H.S. Torr. Riemannian Walk for Incremental Learning: Understanding Forgetting and Intransigence. In *Proceedings of the European conference on computer vision (ECCV)*, pages 556–572, 2018. 6, 7
- [7] Arslan Chaudhry, Marc Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with a

- gem. In *7th International Conference on Learning Representations, ICLR 2019*, pages 1–20, 2019. 2
- [8] Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip H S Torr, and Marc’Aurelio Ranzato. On tiny episodic memories in continual learning. *arXiv preprint arXiv:1902.10486*, 2019. 2
- [9] Kuilin Chen and Chi-guhn Lee. Incremental Few-shot Learning via Vector Quantization in deep embedded space. In *International Conference on Learning Representations*, 2021. 3
- [10] Peng Chu, Xiao Bian, Shaopeng Liu, and Haibin Ling. Feature Space Augmentation for Long-Tailed Data. In *Computer Vision–ECCV*, pages 694–710, 2020. 5
- [11] Gido M 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
- [12] Harm De Vries, Roland Memisevic, and Aaron Courville. Deep learning vector quantization. *ESANN 2016 - 24th European Symposium on Artificial Neural Networks*, (April):503–508, 2016. 3
- [13] Matthias Delange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Ales Leonardis, Greg Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8828(c):1–20, 2021. 1, 2
- [14] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255, 2009. 6
- [15] Laurens der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11), 2008. 9
- [16] Prithviraj Dhar, Rajat Vikram Singh, Kuan Chuan Peng, Ziyang Wu, and Rama Chellappa. Learning without memorizing. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2019-June:5133–5141, 2019. 2
- [17] Chrisantha Fernando, Dylan Banarse, Charles Blundell, Yori Zwols, David Ha, Andrei A Rusu, Alexander Pritzel, and Daan Wierstra. Pathnet: Evolution channels gradient descent in super neural networks. *arXiv preprint arXiv:1701.08734*, 2017. 2
- [18] Evelyn Fix and J. L. Hodges. Discriminatory Analysis. Non-parametric Discrimination: Consistency Properties. *International Statistical Review / Revue Internationale de Statistique*, 57(3):238, 12 1989. 3
- [19] Bernd Fritzke. A growing neural gas network learns topologies. *Advances in Neural Information Processing Systems*, 1994. 2, 9
- [20] Ross Girshick. Fast R-CNN. In *Proceedings of the IEEE international conference on computer vision.*, pages 1440–1448, 2015. 1
- [21] Raia Hadsell, Dushyant Rao, Andrei A. Rusu, and Razvan Pascanu. Embracing Change: Continual Learning in Deep Neural Networks. *Trends in Cognitive Sciences*, 24(12):1028–1040, 12 2020. 1
- [22] Barbara Hammer, Marc Strickert, and Thomas Villmann. Supervised neural gas with general similarity measure. *Neural Processing Letters*, 21(1):21–44, 2005. 3
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2016-December:770–778, 12 2016. 1, 6
- [24] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2019-June:831–839, 2019. 7
- [25] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences of the United States of America*, 114(13):3521–3526, 2017. 1, 2, 7
- [26] Teuvo Kohonen. Improved versions of learning vector quantization. In *ijcnn international joint conference on Neural networks*, pages 545–550. Publ by IEEE, 1990. 2, 3
- [27] Teuvo Kohonen. The Self-Organizing Map. *Proceedings of the IEEE*, 78(9):1464–1480, 1990. 2, 3
- [28] Alex Krizhevsky, Geoffrey Hinton, and others. Learning multiple layers of features from tiny images. 2009. 6
- [29] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. ImageNet Classification with Deep Convolutional Neural Networks. *Advances in Neural Information Processing Systems*, 25, 2012. 1
- [30] Lev Lafayette, Greg Sauter, Linh Vu, and Bernard Meade. Spartan performance and flexibility: An hpc-cloud chimera. *OpenStack Summit, Barcelona, 27*, 2016. 6
- [31] Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015. 6
- [32] Zhizhong Li and Derek Hoiem. Learning without Forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12):2935–2947, 2018. 2
- [33] Xialei Liu, Marc Masana, Luis Herranz, Joost Van De Weijer, Antonio M. Lopez, and Andrew D. Bagdanov. Rotate your Networks: Better Weight Consolidation and Less Catastrophic Forgetting. *Proceedings - International Conference on Pattern Recognition*, 2018-August:2262–2268, 11 2018. 2
- [34] Yu Liu, Sarah Parisot, Gregory Slabaugh, Xu Jia, Ales Leonardis, and Tinne Tuytelaars. More Classifiers, Less Forgetting: A Generic Multi-classifier Paradigm for Incremental Learning. In *Computer Vision–ECCV 2020: 16th European Conference*, pages 699–716, 2020. 2, 7
- [35] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully Convolutional Networks for Semantic Segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. 1
- [36] David Lopez-Paz and Marc’ Aurelio Ranzato. Gradient Episodic Memory for Continual Learning. In I Guyon,

- U Von Luxburg, S Bengio, H Wallach, R Fergus, S Vishwanathan, and R Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. 1, 2
- [37] 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. 2
- [38] Sylvestre Alvisé Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H. Lampert. iCaRL: Incremental classifier and representation learning. *Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017*, 2017-Janua:5533–5542, 2017. 1, 2, 7
- [39] Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016. 2
- [40] Atsushi Sato and Keiji Yamada. Generalized Learning Vector Quantization. *Advances in neural information processing systems*, pages 423–429, 1996. 3
- [41] Ari Seff, Alex Beatson, Daniel Suo, and Han Liu. Continual learning in generative adversarial nets. *arXiv preprint arXiv:1705.08395*, 2017. 2
- [42] Joan Serra, Didac Suris, Marius Miron, and Alexandros Karatzoglou. Overcoming catastrophic forgetting with hard attention to the task. In *International Conference on Machine Learning*, pages 4548–4557, 2018. 2
- [43] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. In *Advances in Neural Information Processing Systems*, volume 30, 2017. 2
- [44] James Smith, Yen-Chang Hsu, Jonathan Balloch, Yilin Shen, Hongxia Jin, and Zsolt Kira. Always be dreaming: A new approach for data-free class-incremental learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9374–9384, 2021. 2
- [45] James Seale Smith, Junjiao Tian, Yen-Chang Hsu, and Zsolt Kira. A Closer Look at Rehearsal-Free Continual Learning. (2019):1–12, 2022. 1
- [46] Xiaoyu Tao, Xinyuan Chang, Xiaopeng Hong, and Xing Wei. Topology-Preserving Class-Incremental Learning for Image Classification. *European Conference on Computer Vision, (Cil)*, 2020. 3
- [47] Xiaoyu Tao, Xiaopeng Hong, Xinyuan Chang, Songlin Dong, Xing Wei, and Yihong Gong. Few-Shot Class-Incremental Learning. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 12180–12189, 2020. 3
- [48] Hoang Thanh-Tung and Truyen Tran. Catastrophic forgetting and mode collapse in gans. In *2020 international joint conference on neural networks (ijcnn)*, pages 1–10, 2020. 2
- [49] Thomas Martinetz and Klaus Schulten. A "Neural-Gas" Network Learns Topologies. *Artificial Neural Networks*, page 397–402, 1991. 2, 4, 9
- [50] T. Villmann, M. Biehl, A. Villmann, and S. Saralajew. Fusion of deep learning architectures, multilayer feedforward networks and learning vector quantizers for deep classification learning. *12th International Workshop on Self-Organizing Maps and Learning Vector Quantization, Clustering and Data Visualization, WSOM 2017 - Proceedings*, 8 2017. 3
- [51] Hong-ming Yang, Xu-yao Zhang, Fei Yin, and Cheng-lin Liu. Robust Classification with Convolutional Prototype Learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 3474–3482, 2018. 5
- [52] Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. Lifelong learning with dynamically expandable networks. *arXiv preprint arXiv:1708.01547*, 2017. 2
- [53] 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. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 6980–6989, 2020. 2
- [54] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *34th International Conference on Machine Learning, ICML 2017*, volume 8, pages 6072–6082, 2017. 2
- [55] Juntao Zhang, Jie Zhang, Shalini Ghosh, Dawei Li, Serafetin Tasci, Larry Heck, Heming Zhang, and C.-C. Jay Kuo. Class-incremental Learning via Deep Model Consolidation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 3 2020. 2
- [56] Fei Zhu, Zhen Cheng, Xu-Yao Zhang, and Cheng-Lin Liu. Class-Incremental Learning via Dual Augmentation. In *Advances in Neural Information Processing Systems*, number NeurIPS, pages 14306–14318, 2021. 1, 2, 6, 7
- [57] Fei Zhu, Xu-yao Zhang, Chuang Wang, Fei Yin, and Cheng-lin Liu. Prototype Augmentation and Self-Supervision for Incremental Learning. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 5267–5276, 2021. 1, 2, 3, 6, 7, 8
- [58] Kai Zhu, Wei Zhai, Yang Cao, Jiebo Luo, and Zheng-Jun Zha. Self-Sustaining Representation Expansion for Non-Exemplar Class-Incremental Learning. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 9296–9305, 2022. 1, 2, 6, 7