



HAL
open science

Contrastive Learning for Online Semi-Supervised General Continual Learning

Nicolas Michel, Romain Negrel, Giovanni Chierchia, Jean-François Bercher

► **To cite this version:**

Nicolas Michel, Romain Negrel, Giovanni Chierchia, Jean-François Bercher. Contrastive Learning for Online Semi-Supervised General Continual Learning. 2022 IEEE International Conference on Image Processing (ICIP), Oct 2022, Bordeaux, France. pp.1896-1900, 10.1109/ICIP46576.2022.9897290 . hal-03866436

HAL Id: hal-03866436

<https://hal.science/hal-03866436>

Submitted on 22 Nov 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

CONTRASTIVE LEARNING FOR ONLINE SEMI-SUPERVISED GENERAL CONTINUAL LEARNING

Nicolas Michel, Romain Negrel, Giovanni Chierchia, Jean-François Bercher

Univ Gustave Eiffel, CNRS, LIGM, F-77454 Marne-la-Vallée, France

ABSTRACT

We study Online Continual Learning with missing labels and propose SemiCon, a new contrastive loss designed for partly labeled data. We demonstrate its efficiency by devising a memory-based method trained on an unlabeled data stream, where every data added to memory is labeled using an oracle. Our approach outperforms existing semi-supervised methods when few labels are available, and obtain similar results to state-of-the-art supervised methods while using only 2.6% of labels on Split-CIFAR10 and 10% of labels on Split-CIFAR100.

Index Terms— Continual Learning, Contrastive Learning, Semi-Supervised Learning, Memory

1. INTRODUCTION

In the last decade, deep neural networks have demonstrated their efficiency achieving state-of-the-art results in several computer vision tasks, such as image classification or object detection. Efficiently training such networks relies on the following assumptions.

- A1.** Data is identically and independently distributed.
- A2.** Training data can be seen multiple times by the model during the learning process.
- A3.** Training data is fully labeled.

However, these assumptions are seldom true in real-world environments with a continuous data stream. In such scenarios, Assumption A1 cannot be verified, Assumption A2 is difficult to ensure as the amount of data grows indefinitely, and Assumption A3 would imply an oracle labeling every new incoming data. While numerous successful training strategies have been designed to leverage unlabeled data [1, 2] when A3 cannot be met, coping with the absence of A1 and A2 remains problematic [3]. A known consequence is Catastrophic Forgetting [4], and is likely to be observed if no specific measure is taken. Continual Learning (CL) aims to maintain performances in the absence of A1, and Online Continual Learning (OCL) addresses the lack of A1 and A2.

In this paper, we propose an OCL algorithm for image classification with few labels. Our method is designed to train a neural network in the absence of A1, A2 and A3. We achieve this by introducing a unified Semi-Supervised Contrastive Loss to leverage labeled and unlabeled data. Despite

This work has received support from Agence Nationale de la Recherche (ANR) for the project APY, with reference ANR-20-CE38-0011-02. This work was granted access to the HPC resources of IDRIS under the allocation 2022-AD011012603 made by GENCI. Code is available at <https://github.com/Nicolas1203/ossagl>

the lack of A3, our results show that the proposed approach is comparable to state-of-the-art supervised OCL methods, and performs better than existing semi-supervised OCL methods when labels are scarce.

The paper is organized as follows. Section 2 discusses the related work. Section 3 describes our approach: it sets up the problem, introduces a new Semi-Supervised Contrastive Loss, and describes the algorithm. Section 4 assesses the performance of our approach on a standard benchmark. Finally, Section 5 concludes the paper.

2. RELATED WORK

Our method stands at the junction of Contrastive Learning, Continual Learning and Semi-Supervised learning. In the following section we will briefly recall these fields and elaborate onto specific scenarios where label information is incomplete.

2.1. Contrastive Learning

Contrastive Learning has become especially popular in recent years for learning representation from data [1, 2, 5, 6, 7, 8]. The intuition is quite simple. Similar samples (called positives) should have close representations, while dissimilar samples (called negatives) should have their representations as far away as possible.

Self-Supervised Contrastive Learning. Contrastive Learning was first designed for unlabeled data, where positives are formed by adding noise to the input. For images, noisy versions are augmentations or views of the original image. All views of the same inputs are positives, while any other image in the training batch is considered a negative. This self-supervision has several drawbacks: large batches are required for sampling enough negatives [1], and images from the same class might be considered as negatives [5, 6, 2].

Supervised Contrastive Learning. A supervised contrastive approach has been proposed by Khosla *et al.* [8] using label information to overcome some limitations of self-supervision. The authors consider every image from the same class as positives and show significant improvement.

2.2. Continual Learning

Continual Learning (CL) has gained popularity in the past years for image classification. The problem is as follows. Consider a sequential learning setup with a sequence $\{\mathcal{T}_1, \dots, \mathcal{T}_K\}$ of K tasks, and $\mathcal{D}_k = (X_k, Y_k)$ the corresponding data-label pairs. The number of tasks K is unknown and could potentially be infinitely large. For any values $k_1, k_2 \in \{1, \dots, K\}$, if $k_1 \neq k_2$ then we have $Y_{k_1} \cap Y_{k_2} = \emptyset$ and the number of

classes in each task is the same. Catastrophic Forgetting occurs when a model’s performance drastically drops on past tasks while learning the current task [4, 9].

General Continual Learning. Early CL studies, and still the majority of current methods, rely on settings where the task-id k is known [10, 11, 12], or at least the task-boundary is known (i.e. knowing when the task change occurs). However, such information is usually unavailable while training in a real environment [10, 13, 14, 15]. Buzzega *et al.* [16] introduced the General Continual Learning (GCL) scenario where task-ids and task-boundaries are unavailable.

Online GCL. Working with a data stream, the model should be able to learn without storing all incoming data. The most realistic scenario Online General Continual Learning (OGCL), described by Buzzega *et al.* [16], adds one more constraint to GCL: data are presented once in the incoming stream. This means that the model should adapt to current data without having access to all past data, even if we can afford to store a limited amount of stream data.

CL with missing labels. While most CL research focuses on the supervised setting, recent works consider CL where few or no labels are available [17, 18, 19]. However, none of them deals with the OGCL setting except the STAM architecture [20], where the authors present an online clustering suited for unsupervised OGCL. In this paper we focus on the OGCL with few labels, which we present in Section 3.

3. ONLINE SEMI-SUPERVISED GENERAL CONTINUAL LEARNING

In this section we formally define the Online Semi-Supervised General Continual Learning (OSSGCL) setting, and propose an approach to address the underlying problem.

3.1. Problem Definition

In the supervised case, each value in $Y = \cup_{k=1}^K Y_k$ is accessible. We consider the semi-supervised case, where we iterate over an incremental unlabeled data stream $\mathcal{S} = \{X_1, \dots, X_K\}$ and use an oracle to label specifically selected data. In this context, we have access to a subset $Y^l \subset Y$ with $p = \frac{|Y^l|}{|Y|}$ being the percentage of labels available. Usually in semi-supervised learning we want p as small as possible. In addition, we consider that we have labeled examples for every classes encountered. We define X_k^u and (X_k^l, Y_k^l) the sets of unlabeled and labeled data for task k . The problem then becomes learning sequentially from $\{\mathcal{T}_1, \dots, \mathcal{T}_K\}$ tasks on partially labeled datasets $\mathcal{D}_k = X_k^u \cup (X_k^l, Y_k^l)$.

3.2. Proposed Approach

To tackle the OSSGCL problem, we design a new Semi-supervised Contrastive Loss (SemiCon) combining a Supervised Contrastive Loss [8] and a Self-supervised Contrastive Loss [1].

Contrastive Learning framework. Following [1], consider the following elements: a data augmentation process $\text{Aug}(\cdot)$ that transforms any given input according to some random procedure with $a \neq b \Leftrightarrow \text{Aug}_a(\cdot) \neq \text{Aug}_b(\cdot)$; an encoder $\text{Enc}_\theta(\cdot)$ that maps the input to the latent space; a projection head $\text{Proj}_\phi(\cdot)$ that maps the latent representation to another space where the loss is applied. The encoder can be

any function and the obtained latent representation will be used for downstream tasks. The projection head is discarded when training is complete to keep only the encoder.

Multiview batch with missing labels. We define $\mathcal{B} = \{x_i^l, y_i^l\}_{i=1..b_l} \cup \{x_j^u\}_{j=1..b_u}$ as the incoming batch of $b = b_l + b_u$ inputs with labeled and unlabeled data. Considering a, b random numbers, we then work on $\mathcal{B}_I = \text{Aug}_a(\mathcal{B}) \cup \text{Aug}_b(\mathcal{B})$, the "multiview batch" [8] of $2b$ samples over indices $i \in I$ and $I = I_l \cup I_u$ with I_l the indices over labeled samples and I_u the indexes over unlabeled samples. Moreover, $h_i = \text{Enc}(x_i)$ is the latent representation of x_i with $H_I = \{h_i\}_{i \in I}$, and $z_i = \text{Proj}(h_i)$ is the projection of h_i with $Z_I = \{z_i\}_{i \in I}$. We also define $P(i) = \{j \in I \setminus \{i\} \mid y_j = y_i\}$ the indices over the positives of i (similar samples), and τ the temperature.

SemiCon Loss. We introduce SemiCon, a unified loss designed to train a contrastive model in the context of missing labels. We construct this loss by combining two terms. The first one, \mathcal{L}_m (1), corresponds to a supervised contrastive loss on labeled memory data with unlabeled streaming data being considered as negatives (dissimilar samples).

$$\mathcal{L}_m = - \sum_{i \in I_l} \frac{1}{|P(i)|} \sum_{p \in P(i)} \log \frac{e^{z_i \cdot z_p / \tau}}{\sum_{a \in I \setminus \{i\}} e^{z_i \cdot z_a / \tau}} \quad (1)$$

The second term, \mathcal{L}_u (2), stands for an unsupervised contrastive loss on unlabeled stream data with labeled steaming data being considered as negatives with $j(i)$ the index such that i and $j(i)$ are indices of augmented samples having the same input source.

$$\mathcal{L}_u = - \sum_{i \in I_u} \log \frac{e^{z_i \cdot z_{j(i)} / \tau}}{\sum_{a \in I \setminus \{i\}} e^{z_i \cdot z_a / \tau}} \quad (2)$$

We define a unified loss $\mathcal{L}_{\text{SemiCon}} = \mathcal{L}_m + \alpha \mathcal{L}_u$ where $\alpha \in [0, +\infty[$ is a weighting hyper-parameter representing the importance of unlabeled data during training. We can see that $(\forall i \in I_u) P(i) = \{j(i)\}$ so the loss can be expressed as

$$\mathcal{L}_{\text{SemiCon}} = - \sum_{i \in I} \frac{g_\alpha(i)}{|P(i)|} \sum_{p \in P(i)} \log \frac{e^{z_i \cdot z_p / \tau}}{\sum_{a \in I \setminus \{i\}} e^{z_i \cdot z_a / \tau}} \quad (3)$$

with $g_\alpha(i) = \alpha$ if $i \in I_l$, and $g_\alpha(i) = 1$ otherwise. Methods replaying past data can suffer from overfitting on memory data [21] and while SemiCon handles missing labels, it also gives control over how a model should balance learning from memory and streaming data separately.

Training Procedure. We propose an approach inspired by the work of Mai *et al.* [22]. They defined Supervised Contrastive Replay (SCR) which combines a Supervised Contrastive Loss [8] and a memory-based strategy [23]. Their method achieves state-of-the-art results in online CL when every data is labeled. We adapt SCR to work with a limited amount of labeled data using SemiCon as objective. Our approach relies on two points: (a) each data added to the memory buffer is labeled, (b) we leverage labeled and unlabeled data in a unified contrastive objective using SemiCon.

During the training phase, we iterate over an unlabeled data stream \mathcal{S} . For each incoming stream batch $\mathcal{B}_s \in \mathcal{S}$, we randomly sample a labeled data batch \mathcal{B}_m from memory \mathcal{M}

Algorithm 1 Proposed Training Method

Input: Unlabeled data stream \mathcal{S} ; Memory \mathcal{M} ; $\text{Aug}(\cdot)$; $\text{Enc}_\theta(\cdot)$; $\text{Proj}_\phi(\cdot)$; Oracle \mathcal{O}
Output: Encoder Enc_θ ; Memory \mathcal{M}
for $\mathcal{B}_s \in \mathcal{S}$ **do**
 $\mathcal{B}_m \leftarrow \text{MemoryRetrieval}(\mathcal{M})$ ▷ Unlabeled data
 $\mathcal{B} \leftarrow \mathcal{B}_s \cup \mathcal{B}_m$ ▷ Labeled data
 $\mathcal{B}_I \leftarrow \text{Aug}_a(\mathcal{B}) \cup \text{Aug}_b(\mathcal{B})$ ▷ Combined Batch
 $Z_I \leftarrow \text{Proj}_\phi(\text{Enc}_\theta(\mathcal{B}_I))$ ▷ a,b randoms
 $\theta, \phi \leftarrow \text{SGD}(\mathcal{L}_{\text{SemiCon}}(Z_I), \theta, \phi)$ ▷ eq.(3)
 $\mathcal{M} \leftarrow \text{MemoryUpdate}(\mathcal{B}_s, \mathcal{O}, \mathcal{M})$
return: $\theta; \mathcal{M}$

and work on $\mathcal{B} = \mathcal{B}_s \cup \mathcal{B}_m$. Each data batch \mathcal{B} is augmented, and the obtained multiview batch \mathcal{B}_I is fed to the network to compute image projections Z_I . The objective function $\mathcal{L}_{\text{SemiCon}}$ is computed on Z_I , and the model parameters are updated using vanilla Steepest Gradient Descent (SGD). After each SGD step, memory data is updated using Reservoir Sampling [24]. Each selected stream data is labeled using the Oracle \mathcal{O} before memory storage.

During the testing phase, memory data representations $H_{\mathcal{M}} = \{\text{Enc}_\theta(x)\}_{x \in \mathcal{M}}$ are computed and a classifier is trained on $H_{\mathcal{M}}$. Our approach is detailed in Algorithm 1. Similarly to Mai *et al.*, we use the Nearest Class Mean (NCM) classifier for the testing phase. Any other classifier can be used on top of the obtained representations; however, we found no significant difference in performance.

4. EXPERIMENTAL RESULTS

In this section, we describe two CL datasets, introduce baselines and present our results compared to state-of-the-art.

4.1. Datasets

We use modified versions of standard image classification datasets [25] to build an incremental learning environment. These datasets are built on CIFAR10 and CIFAR100 by splitting them into several tasks of non-overlapping classes. Specifically, we work on Split-CIFAR10 and Split-CIFAR100. We divide CIFAR10 in 5 tasks with 2 classes per task and CIFAR100 in 10 tasks with 10 classes per task. Each dataset contains 50,000 train images and 10,000 test images.

4.2. Baselines

To assess our results, we compare them to several baselines which respect the OGCL setting and are listed below:

- **offline:** Supervised upper bound. The model is trained without any CL specific constraints.
- **fine-tuned:** Supervised lower bound that trains the model in a CL setting without precautions to avoid forgetting.
- **SCR:** Current state-of-the-art on fully supervised OGCL and closest method to our work.
- **SCR - Memory Only** (SCR-MO): SCR method, but trained using memory data only.
- **Experience Replay** (ER) [23]: ER is a simple baseline which applies reservoir sampling [24] for memory update,

just as SCR, but trained with a cross entropy loss rather than a contrastive loss.

- **Experience Replay - Memory Only** (ER-MO): ER-MO is essentially ER, but trained using memory data only.

Even though other semi-supervised CL methods exist, none respects the OSSGCL setting and thus cannot be used in the comparison [14, 18, 26].

4.3. Implementation Detail

For every experiment we train a reduced ResNet-18 [27] from scratch following previous works [13, 22] and the projection layer Proj_ϕ is a multi layer perceptron [1] with one hidden layer, a ReLU activation and an output size of 128. For memory based methods, we use a memory batch size $|\mathcal{B}_m|$ of 100 on Split-CIFAR10 and of 500 on Split-CIFAR100. More details on the impact of memory batch size can be found in section 4. For online methods, we use a stream batch size $|\mathcal{B}_s|$ of 10, which ensures 5,000 SGD steps on both datasets. Each compared method is trained using SGD optimizer with a learning rate of 0.1, no regularization, and a temperature $\tau = 0.07$ for contrastive loss. For every experiment we used the same augmentation procedure as in [1] and for contrastive methods we construct the multiview batch using one augmentation for each view, while the original SCR implementation used the original image as one view and an augmentation as the other view. We obtained experimentally better results using one augmentation for each view. All experiments are performed 10 times. The average results and their standard deviations are shown in the next section. For the offline baseline, we use the same optimizer and network as other methods and train for 50 epochs.

4.4. Results

In the following we briefly recall a standard CL metric, describe the impact of two hyper-parameters and analyze the obtained results.

Metrics. We use the accuracy averaged across all tasks after training on the last task. This metric is referred to as the final average accuracy [11, 10].

Memory Batch Size selection. We study the impact of the memory batch size $|\mathcal{B}_m|$ on the performance. We use SCR, SCR-MO and our proposed approach as the case of study. As shown in figure 1, every method follows the same trend and benefits from larger $|\mathcal{B}_m|$. We select $|\mathcal{B}_m| = 100$ for every training on Split-CIFAR10 and likewise $|\mathcal{B}_m| = 500$ on Split-CIFAR100. Moreover, even with $\alpha = 0$, our method consistently outperforms SCR-MO on smaller $|\mathcal{B}_m|$ values. This demonstrate that using unlabeled negated can significantly enhance performance when few labels are available.

Impact of α . We evaluate the impact of α on our method’s performance. We keep every other parameter fixed when experimenting on α values. Intuitively, α corresponds to the importance we want to give to unlabeled streaming data against labeled memory data. We observe on figure 2 that the optimal value for α depends on the memory size and tends to be close to one. Figure 3 confirms this observation by comparing performance for $\alpha = 1$ to best performance for any α and implies that $\alpha = 1$ is an acceptable default parameter for our method. Also, our method performance becomes




Fig. 1. Impact of increasing Memory Batch Size $|\mathcal{B}_m|$. We set $\alpha = 0$ while augmenting $|\mathcal{B}_m|$. The memory size is set to 200 for Split-CIFAR10 and 2k for Split-CIFAR100

comparable to SCR-MO for larger memory sizes. We interpret previous observation as the consequence of α playing the role of a regularization parameter. When the memory size is small, the model tends to overfit on memory data and performs better when α is larger. Likewise when the memory size is large, the model has enough information in memory and performs better when α is smaller. Looking at the results on figure 2, best performances are obtained when we use the information from both stream and memory.




Fig. 2. Average Accuracy for $\alpha \in [10^{-1.5}, 10^{0.75}]$ for our method. Every other parameter is the same for SCR, SCR-MO and our method. Left figure corresponds to Split-CIFAR10 with a memory size M of 200 and right figure corresponds to $M=500$.

Results interpretation. Table 1 shows the comparison of our approach to current supervised state-of-the-art methods and their memory-only counterparts on Split-CIFAR10 and Split-CIFAR100 with varying memory sizes M . Our approach achieves best results compared to other semi-supervised methods and performs comparably to supervised state-of-the-art while leveraging only a small portion of labels. This is especially remarkable for small memory sizes where only 2.6% labels are given to the model. Figure 3 indicates that our method outperforms SCR-MO when less than 20% labels are available. Also, our method is on par




Fig. 3. Relative performance compared to SCR on Split-CIFAR100 for SCR-MO and our method while increasing the percentage of available labels.

with SCR using only 10% labels while SCR-MO needs 20% labels to obtain comparable results. Our experiments also show that baselines using the information only available in memory perform competitively, with results close to their supervised counterparts for large memory sizes. This can be explained by the observation that the larger the memory, the closer the problem is to an offline supervised problem.

Method		M=200	M=500	M=2k	M=5k
Supervised	offline fine-tuned	80.0±1.2	16.4±2.0	43.2±2.3	3.6±0.7
	ER	42.6±1.9	52.3±4.7	23.8±1.3	28.5±1.0
	SCR	49.2±2.2	58.7±1.3	31.3±0.8	39.3±0.8
	% labels	2.6%	5.6%	16.9%	33%
Semi Supervised	ER-MO	41.0±3.5	49.9±3.3	23.7±0.9	27.9±1.0
	SCR-MO	46.0±2.0	56.4±1.4	30.9±0.6	38.9±0.8
	Ours	48.8±1.1	57.9±1.1	31.0±0.9	38.9±0.5
		Split-CIFAR10		Split-CIFAR100	

Table 1. Average accuracy on split-CIFAR10 and split-CIFAR100. We use $\alpha = 1.78$ and $\alpha = 0.18$ for split-CIFAR10 and split-CIFAR100 respectively. Best results for each scenario are displayed in bold. Displayed values correspond to the average and standard deviation over 10 experiments.

5. CONCLUSION

In this paper, we defined a novel OSSGCL setting, and introduced a new Semi-supervised Contrastive Loss (SemiCon). We demonstrated experimentally that semi-supervised approaches trained using memory data only can perform competitively to their supervised counterparts, while leveraging as few as 2.6% labels on split-CIFAR10. We proposed a new memory-based approach for the OSSGCL setting which successfully combines labeled and unlabeled data using the novel SemiCon loss. This criterion allows user-controlled balance between labeled and unlabeled data during training. We showed that our method can take advantage of unlabeled data, surpassing other semi-supervised baselines on Split-CIFAR datasets, and achieving similar performance to state-of-the-art supervised methods.

6. REFERENCES

- [1] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton, “A simple framework for contrastive learning of visual representations,” *PMLR*, pp. 1597–1607, 2020.
- [2] Aaron van den Oord, Yazhe Li, and Oriol Vinyals, “Representation learning with contrastive predictive coding,” *NIPS*, 2018.
- [3] Zheda Mai, Ruiwen Li, Jihwan Jeong, David Quispe, Hyunwoo Kim, and Scott Sanner, “Online continual learning in image classification: An empirical survey,” *Neurocomputing*, vol. 469, pp. 28–51, 2022.
- [4] Michael McCloskey and Neal J Cohen, “Catastrophic interference in connectionist networks: The sequential learning problem,” in *Psychology of learning and motivation*, vol. 24, pp. 109–165. Elsevier, 1989.
- [5] Michael Gutmann and Aapo Hyvärinen, “Noise-contrastive estimation: A new estimation principle for unnormalized statistical models,” in *Proceedings of the thirteenth international conference on artificial intelligence and statistics*. JMLR Workshop and Conference Proceedings, 2010, pp. 297–304.
- [6] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin, “Unsupervised feature learning via non-parametric instance discrimination,” *CVPR*, pp. 3733–3742, 2018.
- [7] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He, “Improved baselines with momentum contrastive learning,” *NIPS*, 2018.
- [8] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan, “Supervised contrastive learning,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 18661–18673, 2020.
- [9] Robert M French, “Catastrophic forgetting in connectionist networks,” *Trends in cognitive sciences*, vol. 3, no. 4, pp. 128–135, 1999.
- [10] Yen-Chang Hsu, Yen-Cheng Liu, Anita Ramasamy, and Zsolt Kira, “Re-evaluating continual learning scenarios: A categorization and case for strong baselines,” in *Continual Learning Workshop of 32nd Conference on Neural Information Processing Systems (NeurIPS 2018)*, Montréal, Québec, dec 2018.
- [11] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, . A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska, D. Hassabis, C. Clopath, D. Kumaran, and R. Hadsell, “Overcoming catastrophic forgetting in neural networks,” vol. 114, no. 13, pp. 3521–3526, Publisher: National Academy of Sciences Section: Biological Sciences.
- [12] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert, “iCaRL: Incremental classifier and representation learning,” pp. 2001–2010, 2017.
- [13] David Lopez-Paz and Marc’Aurelio Ranzato, “Gradient episodic memory for continual learning,” *Advances in neural information processing systems*, vol. 30, 2017.
- [14] Jiangpeng He and Fengqing Zhu, “Unsupervised continual learning via pseudo labels,” *arXiv:2104.07164*, 2021, version: 2.
- [15] Pranshu Ranjan Singh, Saisubramaniam Gopalakrishnan, Qiao ZhongZheng, Ponnuthurai N. Suganthan, Savitha Ramasamy, and ArulMurugan Ambikapathi, “Task-Agnostic Continual Learning Using Base-Child Classifiers,” in *2021 IEEE International Conference on Image Processing (ICIP)*, Sept. 2021, pp. 794–798, ISSN: 2381-8549.
- [16] Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara, “Dark experience for general continual learning: a strong, simple baseline,” *Advances in neural information processing systems*, vol. 33, pp. 15920–15930, 2020.
- [17] Dushyant Rao, Francesco Visin, Andrei Rusu, Razvan Pascanu, Yee Whye Teh, and Raia Hadsell, “Continual unsupervised representation learning,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [18] James Smith, Jonathan Balloch, Yen-Chang Hsu, and Zsolt Kira, “Memory-efficient semi-supervised continual learning: The world is its own replay buffer,” *IJCNN*, pp. 1–8, 2021.
- [19] Matteo Boschini, Pietro Buzzega, Lorenzo Bonicelli, Angelo Porrello, and Simone Calderara, “Continual semi-supervised learning through contrastive interpolation consistency,” *arXiv preprint arXiv:2108.06552*, 2021.
- [20] James Smith, Cameron Taylor, Seth Baer, and Constantine Dovrolis, “Unsupervised progressive learning and the stam architecture,” *International Joint Conferences on Artificial Intelligence Organization*, 2021.
- [21] 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*, June 2019.
- [22] Zheda Mai, Ruiwen Li, Hyunwoo Kim, and Scott Sanner, “Supervised contrastive replay: Revisiting the nearest class mean classifier in online class-incremental continual learning,” *CVPR*, pp. 3589–3599, 2021.
- [23] David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne, “Experience replay for continual learning,” in *Advances in Neural Information Processing Systems*. vol. 32, Curran Associates, Inc.
- [24] Jeffrey S Vitter, “Random sampling with a reservoir,” *ACM Transactions on Mathematical Software (TOMS)*, vol. 11, no. 1, pp. 37–57, 1985.
- [25] Alex Krizhevsky, “Learning multiple layers of features from tiny images,” M.S. thesis, Department of Computer Science, University of Toronto, 2009.
- [26] Dhanajit Brahma, Vinay Kumar Verma, and Piyush Rai, “Hypernetworks for continual semi-supervised learning,” *arXiv preprint arXiv:2110.01856*, 2021.
- [27] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Deep residual learning for image recognition,” *CVPR*, pp. 770–778, 2016.