

Heterogeneous Continual Learning

Divyam Madaan^{1,2*}, Hongxu Yin¹, Wonmin Byeon¹, Jan Kautz¹, Pavlo Molchanov¹

¹NVIDIA, ²New York University

divyam.madaan@nyu.edu, {dannyy, wbyeon, jkautz, pmolchanov}@nvidia.com

Abstract

We propose a novel framework and a solution to tackle the continual learning (CL) problem with changing network architectures. Most CL methods focus on adapting a single architecture to a new task/class by modifying its weights. However, with rapid progress in architecture design, the problem of adapting existing solutions to novel architectures becomes relevant. To address this limitation, we propose Heterogeneous Continual Learning (HCL), where a wide range of evolving network architectures emerge continually together with novel data/tasks. As a solution, we build on top of the distillation family of techniques and modify it to a new setting where a weaker model takes the role of a teacher; meanwhile, a new stronger architecture acts as a student. Furthermore, we consider a setup of limited access to previous data and propose Quick Deep Inversion (QDI) to recover prior task visual features to support knowledge transfer. QDI significantly reduces computational costs compared to previous solutions and improves overall performance. In summary, we propose a new setup for CL with a modified knowledge distillation paradigm and design a quick data inversion method to enhance distillation. Our evaluation of various benchmarks shows a significant improvement on accuracy in comparison to state-of-the-art methods over various networks architectures.

1. Introduction

Over the past decade, we have seen a plethora of innovations in deep neural networks (DNNs) that have led to remarkable improvement in performance on several applications [34, 35, 40]. AlexNet [24] was the first breakthrough showing the potential of deep learning on the ImageNet benchmark [12], it was followed by various architectural advances such as VGG [46], Inception [48], ResNet and its variants [15, 53, 58], efficient architectures [18, 49], and applications of Transformer [52] in computer vision [14, 32]; however, all these architectures are trained from scratch and

compared on ImageNet classification [12]. In real-world scenarios, the dataset size is not constant, yet storage of billions of data instances and retraining the model representations from scratch is computationally expensive and infeasible. Further, access to old datasets is often not available due to data privacy. Thus, it is crucial to transfer the knowledge learned by the previous model representations to the recent state-of-the-advances without forgetting knowledge and accessing prior datasets.

Continual learning (CL) [50] is a common way to train representations on a data stream without forgetting the learned knowledge. However, all prior CL techniques [1, 4, 33, 41, 42, 43, 59] continually adapt a fixed representation structure – a fully-connected network with two hidden layers for MNIST [26] and standard ResNet-18 [15] for all the other datasets (see Figure 1). While recent works investigate the effect of depth and width [37] on continual learning, to our knowledge, no work has focused on the real problem of learning on sequentially arriving tasks with changing network architectures. Moreover, most works assume the weight transfer with the same architecture [43, 56], in case of the previous model being available only as black box, these methods are not applicable.

In this work, we argue that maintaining the same network structure during training is not a realistic assumption for practical applications such as autonomous driving, clinical applications, and recommendation systems. Instead, constant upgrades to stronger models with the best-known architecture are vital for the best customer experience and competitive advantage, and all streams of previous data are not only large but difficult to store, transfer, manage, and protect against security breach. As a motivating example, consider the tasks of clinical diagnosis and segmentation using medical images, where it is essential to continually update the model on arriving data to adapt to the changing environment. However, continually adapting the old model architectures can potentially hurt the current and future task performance. Also, it is often not allowed to retain the previous patients’ data to update to the new model due to privacy concerns. Based on this motivation, our goal is to continually update the state-of-the-art deep learning architectures with-

*Work done during an internship at NVIDIA.

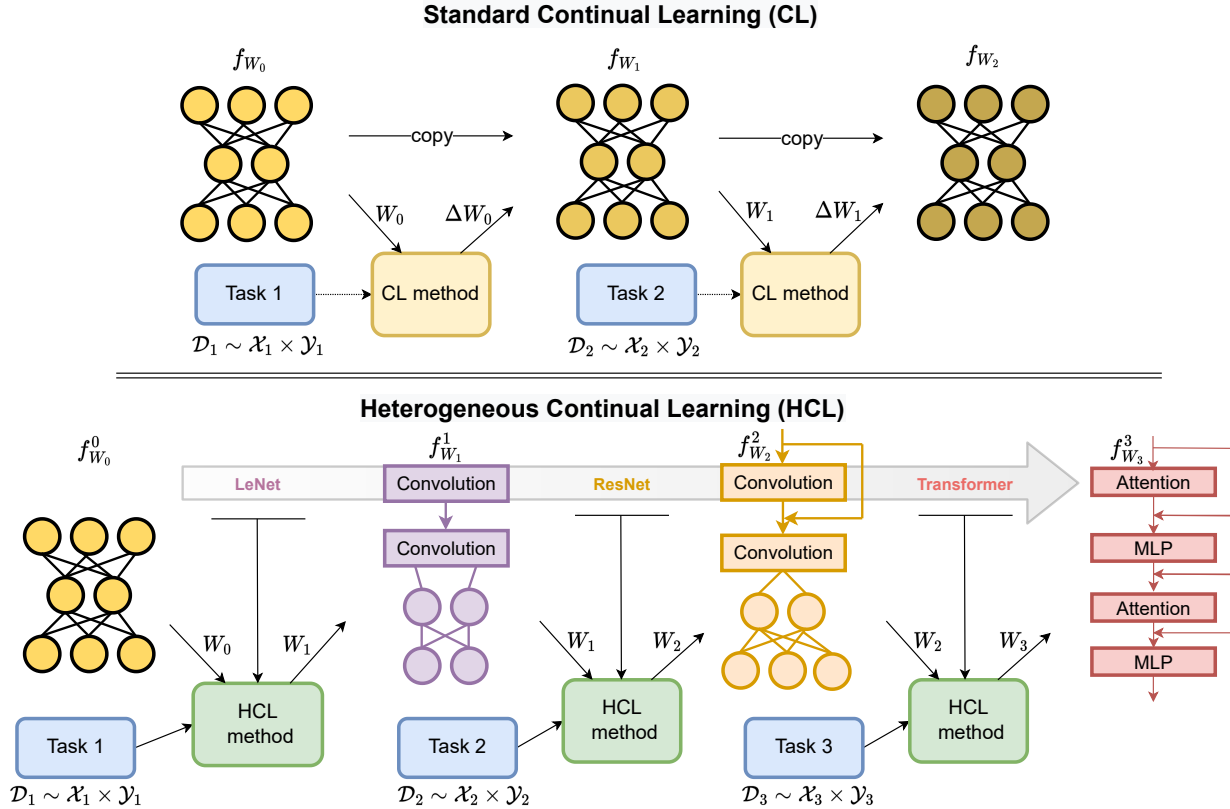


Figure 1. **Illustration of standard and heterogeneous continual learning.** Standard CL updates the same representational structure while preserving previous information with incoming tasks. On the contrary, the goal of HCL is to continuously evolve the network architecture while maintaining the accuracy of new data instances.

out storing the previous data while preserving the knowledge from previously learned representations.

We propose a novel CL framework called *Heterogeneous Continual Learning (HCL)* to learn continual representations on a data stream without any restrictions on the training configuration or network architecture of the representations (see Figure 1). In particular, we consider the continual learner as a stream of learners, where we train each learner with the same or different backbone architecture. To this end, we consider a diverse range of heterogeneous architectures including LeNet [27], ResNet and its variants [15, 53, 58], and recent architectures such as RegNet [54] and vision transformers [32]. However, due to the assumption of the same architectural structure across different tasks, conventional regularization [1, 30, 44, 59] and architectural methods [29, 43, 56] are not directly applicable for HCL.

To continually learn heterogeneous representations, we revisit knowledge distillation [16] – a method to transfer the knowledge learned across different networks. While knowledge distillation has been formulated for continual learning [4, 30, 41], prior methods do not outperform the experience replay based methods [2, 4, 41, 42]. Moreover, the distillation pipeline has not been considered for network

evolution. This work shows that knowledge distillation can outperform state-of-the-art methods with and without replay examples in standard and heterogeneous continual learning settings with proper modifications. Specifically, we incorporate *consistency* and *augmentation* [3] into the knowledge distillation paradigm for CL and low-temperature transfer with label-smoothing [6] for learning the network representations. Our solution comes at the cost of storing and reusing only the most recent model.

The continual learner might not have access to the prior data. Therefore, motivated by DeepInversion (DI) [55], we consider *data-free continual learning*. In particular, DI optimizes random noise to generate class-conditioned samples; however, it requires many steps to optimize a single batch of instances. This increases its computational cost and slows CL. In response to this, we propose *Quick Deep Inversion (QDI)*, which utilizes current task examples as the starting point for synthetic examples and optimizes them to minimize the prediction error on the previous tasks. Specifically, this leads to an interpolation between the current task and previous task instances, which promotes current task adaptation while minimizing catastrophic forgetting. We compare our proposed method on various CL benchmarks against

state-of-the-art methods, where it outperforms them across all settings with a wide variety of architectures. In summary, the contributions of our work are:

- We propose a novel CL framework called *Heterogeneous Continual Learning (HCL)* to learn a stream of different architectures on a sequence of tasks while transferring the knowledge from past representations.
- We revisit knowledge distillation and propose *Quick Deep Inversion (QDI)*, which inverts the previous task parameters while interpolating the current task examples with minimal additional cost.
- We benchmark existing state-of-the-art solutions in the new setting and outperform them with our proposed method across a diverse stream of architectures for both task-incremental and class-incremental CL.

2. Related work

Replay-based approaches. These approaches formulate different criterion [2, 4, 41, 42] to select a representative subset that is revisited in future tasks. While most approaches [2, 41, 42] store the input examples, Buzzega et al. [4] retain the logits to preserve the knowledge from the past. However, all these methods are limited in practical scenarios due to privacy concerns and excessive memory requirements to store the data from previous tasks. Our approach does not strictly require a buffer to store data or representation, and we also do not assume that the learners’ architecture is fixed.

Regularization methods. These methods penalize the change in the representations by adding a regularization term in the loss [1, 30, 44, 59]. Li & Hoiem [30] minimizes the change in the output representations, and Zenke et al. [59] estimates the important parameters and restricts their change during training. However, the regularization cannot be enforced when the architectures are different across tasks.

Architectural methods. They consider the inclusion of task-dependent representations to prevent catastrophic forgetting during continual learning [10, 29, 43, 56, 57]. Rusu et al. [43] instantiates a new column for the new task and require inputs from both columns and Yoon et al. [57] decomposes each layer parameters to task-specific and task-shared parameters. The objective of these methods is to reduce forgetting and the network expansion is restricted to the same structure. In contrast, we do not make any assumption on the network structure and incorporate the recent advances in deep learning to train the continual learner without any knowledge of prior data or its architecture. To the best of our knowledge, this is the first work that explores continual learning with different representational structures while adapting to the current task and minimizing forgetting.

Optimal architectures for continual learning. Prior works also investigate most optimal architecture for continual learn-

ing. Lee et al. [28] formulated a data-driven layer-based transfer function to learn the layers that should be transferred to mitigate forgetting. Mirzadeh et al. [36] showed that increasing the width mitigates forgetting, whereas increasing the depth negatively affects forgetting, and Mirzadeh et al. [38] investigated the effect of various components in an architecture. Carta et al. [5] proposed Ex-Model Continual Learning to learn the continual learner from a stream of trained models; however, their proposed scheme is not directly applicable to the HCL setting as the architecture of their learner is fixed and it requires an extra buffer to store the past expert and CL representations.

3. Heterogeneous Continual Learning

In real-world scenarios, the data distribution continuously evolves with time; thus, adapting the network representations to the distribution shifts is essential while preserving the knowledge on the past sequence of data. Meanwhile, deep learning techniques are improving rapidly, providing better model architectures and techniques with higher performance on various tasks. However, prior CL methods [1, 4, 41, 42, 43, 59] either optimize the performance on a single architecture [15] or focus on finding the optimal architecture for CL [36, 38].

To motivate further, we provide a simple experiment on CIFAR-10 split into two tasks demonstrating the benefit of adapting to novel neural architectures in

Figure 2. With the new ability to switch the network architecture, (e.g., from LeNet to ResNet) the performance surpasses standard CL setup that is limited by a rigid but weak architecture (HCL vs. CL LeNet). More appealingly, it even approximates the stronger model deployed throughout the CL process (HCL vs. CL ResNet), while alleviating the requirement of replay buffer to store prior data experiences (HCL is data-free vs. CL with buffer), as we will show next in methodology and experiment sections.

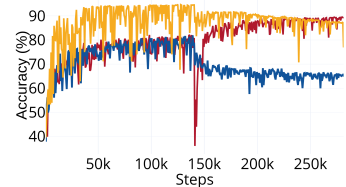


Figure 2. Average accuracy comparison between standard CL – ResNet18 and LeNet trained with buffer, ours with adaptive architectures without buffer on CIFAR-10 dataset with two tasks containing five classes each.

3.1. Problem setup

The problem of standard CL considers the learning objective where the continual learner $f_{W_t} : \mathcal{X}_t \rightarrow \mathcal{Y}_t$ learns a sequence of T tasks by updating the **fixed representation structure** for each task. Each task $t \in T$ contains training data $\mathcal{D}_t = \{x_i, y_i\}_{i=1}^{N_t} \sim \mathcal{X}_t \times \mathcal{Y}_t$, which is composed of N_t identically and independently distributed examples. More

formally, we minimize the following objective:

$$\underset{W_t}{\text{minimize}} \mathbb{E}_{(x_i, y_i) \sim \mathcal{D}_t} [\ell(f_{W_t}(x_i), y_i)], \quad (1)$$

where $\ell : \mathcal{Y}_t \times \mathcal{Y}_t \rightarrow \mathbb{R}_{\geq 0}$ is the task-specific loss function. In this work, we consider the continual learner as a **stream of architectures** $\{f_{W_1}^1, \dots, f_{W_t}^t\}$, where the learner can completely change the backbone architecture to incorporate recent architectural developments [14, 15, 32, 51] in order to improve model performance.

However, when the architectures are different, there is no natural knowledge transfer mechanism, and the network parameters are initialized randomly, which makes this problem challenging compared to standard continual learning.

Particularly, each architectural representation $f_{W_t}^t : \mathcal{X}_t \rightarrow \mathcal{Y}_t, \forall t \in \{1, \dots, T\}$ is trained on task distribution \mathcal{D}_t and our objective is to train these stream of networks on a sequence of tasks without forgetting the knowledge of the previous set of tasks. Additionally, when the network structure remains the same, we want to transfer the learned representations sequentially to train incoming tasks. Overall, the learning objective of the HCL is:

$$\underset{W_t}{\text{minimize}} \mathbb{E}_{(x_i, y_i) \sim \mathcal{D}_t} [\ell(f_{W_t}^t(x_i), y_i)], \quad (2)$$

For notation simplicity, we discard W_t in $f_{W_t}^t$ for the rest of the paper. In this work, we focus on the continual learner that does not rely on task identifiers during training and uses constant memory following prior works [4, 11]. We consider task-incremental (task-IL) and class-incremental learning (class-IL) during inference but do not use task labels during training. In task-IL, we provide the task identity to select the classification head, whereas class-IL uses a shared head across all classes. Consequently, class-IL is more challenging due to an equal weight for all the tasks and prone to higher forgetting and lower accuracy than task-IL; however, recent works [9] also raise concerns regarding the practicality of class-IL in real-world scenarios.

3.2. Proposed method

We address the problem of HCL with our method summarized in **Figure 3**. As in the standard CL, new task data and an optional buffer are used to directly train a new model (stronger student) initialized from scratch with a task loss. Additionally, a knowledge transfer data (KDA) is collected given new task data, (optional) buffer and (optional) synthesised samples by quick deep inversion. The KDA supports network transfer via knowledge distillation between old and new models. To describe our method, we first revisit knowledge distillation [16] in **Section 3.2.1** and propose a training paradigm that outperforms state-of-the-art methods. Then, we introduce quick deep inversion in **Section 3.2.2** to recover past visual features promptly. It inverts the previous task model to generate synthetic examples

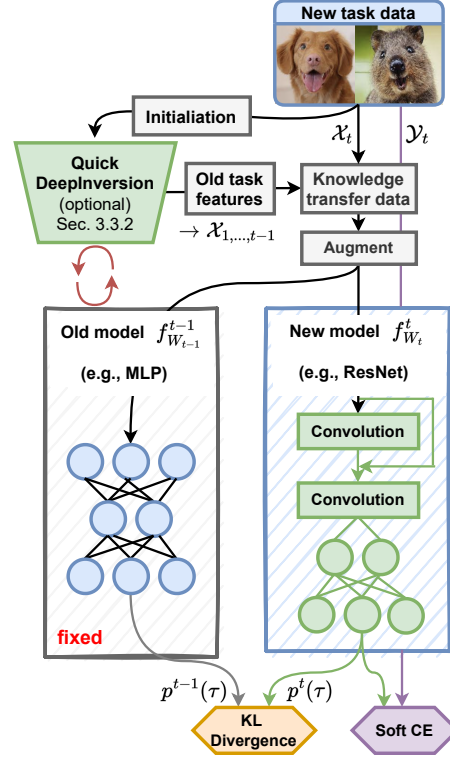


Figure 3. Proposed method to perform continual learning with neural network architecture change.

3.2.1 Revisiting knowledge distillation

Due to a difference in model representations and the absence of the training data used to train previous models, we concentrate on knowledge distillation (KD) [16] for HCL. However, our goal is orthogonal to standard KD; instead of distilling the knowledge from a large model (or ensemble) to a smaller (weaker) model for efficient deployment we distill from the latter to the former. Our motivation is to improve model performance by transferring to a state-of-the-art model.

While KD has been adopted in CL [4, 13, 17, 30, 41], it is (i) limited to homogeneous architectures across different tasks, and (ii) considered a weak baseline in comparison to replay-based methods. Contrarily, we focus on HCL and propose a modified paradigm inspired by the recent advances in KD to improve the performance and lift both limitations.

Label smoothing. Prior works [39, 45] investigating the compatibility between label smoothing (LS) and knowledge distillation have derived opposite conclusions. Recently, Chandrasegaran et al. [6] highlighted that LS-trained teacher with a low-temperature transfer is essential to achieve students with higher performance. This inspired us to adapt label smoothing to produce better soft targets for KD, which improves the knowledge transfer from a smaller to a larger model architecture reducing the forgetting for both the standard CL and HCL scenarios.

Knowledge distillation. To improve the KD paradigm for HCL, we incorporate *consistency* and *augmentation* [3]. Specifically, fixed teaching and distillation without augmentation leads to overfitting of the larger model on the current task performance. On the contrary, consistent teaching using the same input views for the student and teacher and augmentation improves generalization while reducing forgetting. In contrast to standard knowledge distillation, we use the current task instances for distillation, as the previous task data is often unavailable in real-world applications due to data privacy and legal restrictions.

Additionally, our method is applicable to task-free continual learning during training as it does not require task identity. In addition, we do not use warmup before the current task as we find that it is not essential to improve model performance. We conduct ablation experiments to justify the design choices of every component in Table 2. Our overall objective for the lifelong learner is:

$$\underset{W_t}{\text{minimize}} \mathbb{E}_{(x_i, y_i) \sim \mathcal{D}_t} [\ell(f^t(x_i), y_i^t(\psi))] + \alpha \cdot \text{KL}(p_i^t(\tau), p_i^{t-1}(\tau)), \quad (3)$$

where $y_i^t(\psi) = y(1 - \psi) + \psi/C$, ψ denotes the mixture parameter to interpolate the hard targets to uniform distribution defined using ψ , C is the number of classes, $p^t(\tau)$ and $p^{t-1}(\tau)$ denote the temperature-scaled output probabilities for current task model and past-task model, respectively, with τ as the corresponding temperature. α is the hyper-parameter that controls the strength of KD loss.

3.2.2 Enhancing data efficiency

To further improve our data-free HCL paradigm, we extend DeepInversion (DI) [55], which optimizes images \tilde{x} initialized with Gaussian noise to improve knowledge transfer from the prior tasks given only the trained models and no data. The objective of the optimization is to excite particular features, or, as in our work, classes from previous tasks. The optimization is guided by proxy image priors: (i) total variation \mathcal{L}_{tv} , (ii) \mathcal{L}_{ℓ_2} of the generated samples and (iii) feature distribution regularization $\mathcal{L}_{\text{feature}}$. DI requires many steps to generate examples prior to every task, which increases the computational cost of CL method and, therefore, becomes less appealing.

To tackle this limitation, we propose *Quick Deep Inversion (QDI)* that initializes the synthetic examples with the current task data prior to the optimization. There is no reason to believe that randomly sampled Gaussian noise is an optimal starting point for generating past task examples. We optimize the current task examples such that features will fall to the manifold learned by the previous model and the domain shift is minimized. As a result such images are classified as past task classes. QDI generates examples $\tilde{x}_{\text{prior}}\{f^{t-1}, k\}$ to approximate features from all prior tasks $\{1, \dots, t-1\}$ by inverting the last model f^{t-1} with k optimization steps:

$$\tilde{x}_{\text{prior}}\{f^{t-1}, k\} = \underset{\tilde{x}}{\text{argmin}} \mathcal{L}(f^{t-1}(\tilde{x}), \tilde{y}) + \alpha_{tv} \mathcal{L}_{tv}(\tilde{x}) + \alpha_{\ell_2} \mathcal{L}_{\ell_2}(\tilde{x}) + \alpha_{\text{feature}} \mathcal{L}_{\text{feature}}, \quad (4)$$

where the synthesized examples are optimized towards prior classes $\tilde{y} \sim \mathcal{Y}_{\{1, \dots, t-1\}}$ to minimize forgetting. $\alpha_{tv}, \alpha_{\ell_2}, \alpha_{\text{feature}}$ denote the hyper-parameters that determine the strength of individual losses. To improve synthesis speed, we initialize \tilde{x} with the current task input image x_t , which provides a $4\times$ speed-up and leads to more natural images:

$$\tilde{x}_{\text{prior}}\{f^{t-1}, 0\} = \tilde{x}_{\text{prior}, k=0} = x_t.$$

The initialized \tilde{x}_{prior} can then be optimized using Equation 5, regularized for realism by the target model f^{t-1} , hence quickly unveiling previous task visual features (see Figure 4) on top of the current task image through

$$\mathcal{L}_{\text{feature}} = \sum_{l \in L} \left[d(\mu_l(\tilde{x}_{\text{prior}}), \mathbb{E}[\mu_l(x) | x \sim \mathcal{X}_{\{1, \dots, t-1\}}]) + d(\sigma_l(\tilde{x}_{\text{prior}}), \mathbb{E}[\sigma_l(x) | x \sim \mathcal{X}_{\{1, \dots, t-1\}}]) \right]. \quad (5)$$

$d(\cdot, \cdot)$ denotes the distance metric for feature regularization. In this paper, we use the mean-squared distance. Akin to Yin et al. [55], we assume Gaussian distribution for feature maps and focus on batch-wise mean $\mu_l(x)$ and variance $\sigma_l(x)$ for the layer l . We note that these statistics are implicitly captured through the batch normalization in f^{t-1} without storing input data for all previous tasks $\{1, \dots, t-1\}$: $\mathbb{E}[\mu_l(x) | x \sim \mathcal{X}_{\{1, \dots, t-1\}}] \simeq \text{BN}_l(\text{running mean})$, $\mathbb{E}[\sigma_l(x) | x \sim \mathcal{X}_{\{1, \dots, t-1\}}] \simeq \text{BN}_l(\text{running var})$. Otherwise, we approximate them using values calculated with post-convolution feature maps given a current task batch to the target model, f^{t-1} , leveraging its feature extraction capability for previous tasks. This efficient QDI allows us to use it for continual learning with minimal additional cost to existing methods, and the learning objective in Equation 3 can be updated as:

$$\underset{W_t}{\text{minimize}} \mathbb{E}_{(x_i, y_i) \sim \mathcal{D}_t} [\ell(f^t(x_i), y_i^t(\psi))] + \alpha \cdot \text{KL}(p_i^t(\tau), p_i^{t-1}(\tau)) + \beta \cdot \text{KL}(\tilde{p}_i^t(\tau), \tilde{p}_i^{t-1}(\tau)), \quad (6)$$

where $\tilde{p}^t(\tau)$ and $\tilde{p}^{t-1}(\tau)$ are the output probabilities of the generated examples scaled with temperature τ using the current and past task-models, β is the hyper-parameter to control the strength of QDI knowledge distillation loss. QDI utilizes the generated examples to enhance knowledge distillation with a four times speedup to prior data-inversion methods [55] due to the current task being a better prior than pixel-wise Gaussian to learn the generated data.

4. Experiments

4.1. Experimental setup

Datasets and architectures. In this subsection, we provide the details about the datasets and the architectures, where we consider various architectural innovations in the ImageNet challenge [12] for HCL.

1. **Split CIFAR-10 [23].** It consists of CIFAR-10 sized 32×32 split into five tasks, where each task contains two different classes. We consider a different architecture for each task in HCL: LeNet [25], ResNet [15], DenseNet [20], SENet [19], and RegNet [54].
2. **Split CIFAR-100 [23].** It is a split of CIFAR-100, where the 100 object classes are divided into 20 tasks. We consider nine architectures in the HCL: LeNet [25], AlexNet [24], VGG [46] with batch normalization [21], InceptionNet [48], ResNet [15], ResNeXt [53], DenseNet [20], SENet [19] for two tasks each and RegNet [54] for last four tasks.
3. **Split Tiny-ImageNet.** It is a variant of ImageNet [12] that consists of images sized 64×64 from 200 classes. It consists of 20 tasks, where each task contains 10 classes. The set of architectures in HCL include LeNet [25], ResNet [15], ResNeXt [53], SENet [19], and RegNet [54] architectures, where each architecture is used for two consecutive tasks.

Baselines. We compare our designed KD paradigm (Equation 3), with QDI (Equation 6) against standard fine-tuning (FINETUNE), SI [59], LWF [30], ICARL [41] A-GEM [7], ER [8], DER and DER++ [4], DI [55] and multi task training (MULTITASK), where all tasks are trained jointly with ResNet18 – most prominent architecture in CL literature. For a fair comparison, we extend our method with the classification loss using BUFFER in Equation 3. The buffer-size for replay-based method is fixed to 200 for all the methods that use exemplar or generative replay. We use the same set of augmentations for training all the methods.

Evaluation. To evaluate and compare our method in both standard CL and HCL setting, we use the average accuracy and forgetting as two metrics.

1. **Average accuracy.** It is the average of the accuracy of all tasks after the completion of training with the sequence of T tasks. More formally, average accuracy $\mathcal{A}_T = \frac{1}{T} \sum_{t=1}^T a_{T,t}$, where $a_{i,j}$ is the accuracy of task j after completion of task i .
2. **Average forgetting.** It is the averaged difference in accuracy between the accuracy after training and the maximum accuracy for each task. More formally, average forgetting $\mathcal{F}_T = \frac{1}{T-1} \sum_{t=1}^T \max_{i \in \{1, \dots, T\}} (a_{i,t} - a_{T,t})$

4.2. Quantitative results

Task-incremental continual learning. First, we compare the performance of our proposed method with other baselines for task-IL on standard and heterogeneous continual learning settings. In Table 1, we find that the performance of all the methods drop significantly by going from the SCL to the HCL setting. Contrary, our proposed methods achieve the lowest drop across all methods; moreover, it obtains the best performance with and without buffer across all considered scenarios. Additionally, we show that in standard CL and HCL, QDI leads to a $\sim 40\%$ and $\sim 22\%$ absolute improvement in average accuracy over state-of-the-art data-free methods for Split CIFAR-100 and Tiny-ImageNet. Interestingly, we find that distillation with synthetic examples generated by QDI obtains comparable or better performance to the replay buffer highlighting the flexibility of data-free continual learning.

Moreover, QDI improves synthesis speed by $4\times$ for generating the synthetic examples in comparison to DI [55] for each task. All variants of our method achieve the least forgetting in comparison to the baselines across all datasets for both the standard CL and HCL scenarios. Our proposed training scheme is especially beneficial when past-task data is unavailable due to privacy concerns. It allows the adaptation to the new network structure for the new task by utilizing the saved model checkpoints. It enables training with different network structures without storing any past-task instances.

Class-incremental continual learning. We also show the validation of our proposed method on class-IL in the supplementary material, where task-labels are not provided during training or inference. We observe that all our methods outperform the baselines with and without buffer in average accuracy and forgetting. Specifically, we show that ours with QDI and BUFFER examples show a five to ten percent improvement in average accuracy for Split CIFAR-100 and Tiny-ImageNet datasets. However, we find that buffer is an important component of our method for class-IL; it is consistent to the observations in prior works [5, 55]. Further, we observe that the synthetic examples generated by QDI lead to unstable training for CIFAR-10, but improve the performance for other datasets. We believe this is an outcome of the model’s dependence on domain over semantics [47] and further investigation of data-free continual learning techniques for class-IL would be an interesting direction for future work.

4.3. Additional analysis

Ablation studies. To dissect the contribution of different components in our proposed method for HCL, we isolate and study them on Split CIFAR-10 in Table 2. First, we find that knowledge distillation with augmented instances, in the con-

METHOD	\mathcal{B}	SPLIT CIFAR-10		SPLIT CIFAR-100		SPLIT TINY-IMAGENET	
		\mathcal{A}_T (\uparrow)	\mathcal{F}_T (\downarrow)	\mathcal{A}_T (\uparrow)	\mathcal{F}_T (\downarrow)	\mathcal{A}_T (\uparrow)	\mathcal{F}_T (\downarrow)
STANDARD CONTINUAL LEARNING							
FINETUNE	–	62.93 (\pm 1.41)	41.45 (\pm 2.39)	40.13 (\pm 3.60)	53.64 (\pm 3.65)	19.31 (\pm 0.71)	59.50 (\pm 0.81)
SI [59]*	–	68.05 (\pm 5.91)	38.76 (\pm 0.89)	–	–	36.32 (\pm 0.13)	–
LWF [30]*	–	63.29 (\pm 2.35)	32.56 (\pm 0.56)	–	–	15.85 (\pm 0.58)	–
DI [55]	–	89.11 (\pm 3.70)	7.12 (\pm 2.43)	49.64 (\pm 0.59)	35.29 (\pm 0.75)	44.29 (\pm 1.88)	22.42 (\pm 3.23)
KD (OURS)	–	95.01 (\pm 0.79)	0.87 (\pm 0.35)	87.13 (\pm 0.45)	4.26 (\pm 0.41)	65.48 (\pm 0.12)	7.33 (\pm 2.03)
KD w/ QDI (OURS)	–	95.46 (\pm 0.15)	0.37 (\pm 0.09)	88.30 (\pm 0.17)	2.00 (\pm 0.15)	66.79 (\pm 0.45)	7.28 (\pm 0.49)
ICARL [41]*	✓	88.99 (\pm 2.13)	2.63 (\pm 3.48)	–	–	28.19 (\pm 1.47)	–
A-GEM [7]	✓	87.09 (\pm 0.84)	12.54 (\pm 1.23)	63.88 (\pm 1.23)	29.98 (\pm 1.40)	22.44 (\pm 0.24)	55.43 (\pm 0.71)
ER [42]	✓	91.39 (\pm 0.41)	5.99 (\pm 0.43)	67.06 (\pm 1.01)	25.31 (\pm 1.27)	21.03 (\pm 2.41)	57.88 (\pm 2.63)
DER [4]	✓	92.45 (\pm 0.26)	5.79 (\pm 0.20)	67.74 (\pm 0.88)	24.77 (\pm 0.94)	30.97 (\pm 0.14)	45.25 (\pm 0.48)
DER++ [4]	✓	92.16 (\pm 0.68)	5.96 (\pm 0.74)	69.03 (\pm 0.74)	23.06 (\pm 0.69)	33.63 (\pm 0.10)	40.66 (\pm 0.83)
KD w/ BUFFER (OURS)	✓	95.81 (\pm 0.03)	0.73 (\pm 0.34)	80.35 (\pm 0.53)	10.76 (\pm 0.65)	69.07 (\pm 0.26)	4.37 (\pm 0.17)
MULTITASK*	✓	98.31 (\pm 0.12)	N/A	93.89 (\pm 0.78)	N/A	82.04 (\pm 0.10)	N/A
HETEROGENEOUS CONTINUAL LEARNING							
FINETUNE	–	61.99 (\pm 1.98)	42.15 (\pm 2.94)	27.08 (\pm 1.20)	64.54 (\pm 1.45)	14.83 (\pm 0.53)	56.31 (\pm 0.60)
DI [55]	–	81.16 (\pm 2.77)	15.84 (\pm 3.01)	38.33 (\pm 1.44)	47.08 (\pm 1.37)	31.69 (\pm 1.99)	32.42 (\pm 3.63)
KD (OURS)	–	93.13 (\pm 0.35)	3.97 (\pm 0.72)	68.97 (\pm 0.72)	21.45 (\pm 0.91)	52.86 (\pm 2.78)	11.73 (\pm 2.59)
KD w/ QDI (OURS)	–	90.30 (\pm 1.27)	6.34 (\pm 0.15)	76.44 (\pm 1.36)	12.31 (\pm 0.99)	53.39 (\pm 1.67)	13.66 (\pm 1.18)
A-GEM [7]	✓	76.70 (\pm 0.67)	24.31 (\pm 1.03)	35.37 (\pm 1.29)	56.41 (\pm 2.19)	15.79 (\pm 0.60)	54.10 (\pm 1.00)
ER [42]	✓	81.78 (\pm 2.15)	17.57 (\pm 2.86)	52.59 (\pm 1.03)	36.41 (\pm 0.54)	27.82 (\pm 1.43)	39.66 (\pm 1.24)
DER [4]	✓	79.92 (\pm 0.51)	19.99 (\pm 0.64)	53.01 (\pm 0.61)	37.33 (\pm 0.58)	25.27 (\pm 1.34)	40.81 (\pm 0.57)
DER++ [4]	✓	81.33 (\pm 0.76)	18.79 (\pm 0.75)	54.81 (\pm 1.65)	35.20 (\pm 1.63)	29.74 (\pm 1.40)	36.79 (\pm 2.16)
KD w/ BUFFER (OURS)	✓	93.95 (\pm 0.22)	2.49 (\pm 0.44)	73.98 (\pm 0.13)	16.40 (\pm 0.18)	54.84 (\pm 0.87)	10.82 (\pm 0.80)

Table 1. **Accuracy and forgetting** with task-IL on standard CL and HCL. The best results are highlighted in **bold**. \mathcal{B} denotes replay-buffer, \mathcal{A}_T , \mathcal{F}_T denote average accuracy and forgetting after the completion of training. * denotes the methods whose numbers were used from Buzzega et al. [4] and – indicates the unavailability of results. All other experiments are over three independent runs.

trast to a common setting of precomputed logits, improves the performance by $\sim 8\%$ and $\sim 13\%$ for task-IL and class-IL respectively due to the better generalization of current task representations. Next, we observe that label-smoothing improves the task-IL performance while hurting the class-IL average accuracy highlighting the presence of a trade-off between both the settings. Additionally, we measure the effect of warmup and task-identity for HCL following prior works [22, 30] for standard CL, where we find that warm-up of the linear head before current-task training and inclusion of task-identity hurt the performance for HCL.

Further, we conduct an ablation on different distance metrics used for KD in Equation 3 for HCL. While, a modified version of cross entropy (CE) or mean-square distance (MSE) are commonly used for knowledge distillation in

CL [4, 30, 41], we find the KL-divergence (KL) is more stable and often obtains better performance in Table 3 for both class-IL and task-IL. We believe that the ablation of these design choices will provide critical insights towards building future methods for standard and heterogeneous CL. **Transfer to vision transformer.** To show the flexibility of our framework beyond convolutional architectures, we design a two task experiment with CIFAR-10 consisting of five classes each. The first task is trained with ResNet18 [15] followed by Swin-T [32] on the second task. Table 4 shows the comparison of our proposed method with only KD against FINETUNE, DER++ trained with HCL and MULTITASK learned with CIFAR-10 on Swin-T [32]. First, we remark that Swin-T (MULTITASK: 88.71 ± 0.40) achieves lower accuracy in comparison to ResNet18 (MULTITASK:

MODEL	TASK-IL	CLASS-IL
KD	83.43 (± 0.42)	21.11 (± 0.19)
+ AUG-KD	91.76 (± 2.03)	34.11 (± 1.07)
+ LS	93.13 (± 0.35)	30.21 (± 0.11)
+ WARMUP	91.64 (± 1.37)	28.26 (± 0.28)
+ TI	90.14 (± 1.33)	22.04 (± 0.27)

Table 2. **Ablation** showing \mathcal{A}_t to measure the effect of various components with Split-CIFAR10. The inclusion of distillation with augmented images (AUG-KD) and label-smoothing (LS) significantly improve performance, whereas WARMUP and task-identity (TI) degrade performance. **Blue** and **red** colors denote the components that were included and excluded from the method respectively.

MODEL	TASK-IL	CLASS-IL
CE	91.06 (± 0.63)	20.99 (± 0.27)
MSE	89.68 (± 0.99)	20.09 (± 0.26)
KL	93.13 (± 0.35)	30.21 (± 0.11)

Table 3. **Ablation** showing \mathcal{A}_t using validation set to measure the effect of different distance metrics. We measure the average accuracy and forgetting for HCL on Split-CIFAR10 dataset.

98.31 \pm 0.12) because ViTs are data-hungry [31]. Second, we show that our proposed method with evolving architectures obtains performance similar to the multitask baseline. Lastly, our proposed method obtains better average accuracy and lower forgetting over the evaluated baselines showing the benefit of training with multiple architectures.

Visualization of generated examples. We visualize the generated examples by DI and QDI in the HCL setting for ResNet18 architecture with CIFAR-10 dataset in Figure 4. To explain the interplay of the generated examples initialized with the current task instances and the past task examples, we consider the example of a current task (*e.g.*, a cat or bird) to invert the visual features of the past task (car). We observe that the input-features swiftly generated by QDI have a combination of both classes and superior in realism; in contrast, DI only contains the past-task features and requires more optimization steps from random noise for full convergence. We believe this interpolation also helps in current task-training while also alleviating catastrophic forgetting.

5. Conclusion

In this work, we propose a novel continual learning setting named *Heterogeneous Continual Learning (HCL)*, where the lifelong learner learns on a sequence of tasks while adopting state-of-the-art deep-learning techniques and architectures. We benchmark the state-of-the-art CL solutions in this new setting and observe a large degradation in performance. To

MODEL	\mathcal{A}_T (\uparrow)	\mathcal{F}_T (\downarrow)
FINETUNE	55.47 (± 2.06)	77.68 (± 3.68)
DER++	69.05 (± 2.56)	51.99 (± 1.91)
KD (OURS)	88.77 (± 0.55)	10.95 (± 1.34)
MULTITASK	88.71 (± 0.40)	–

Table 4. **Transformer analysis.** Average accuracy and forgetting for HCL with ResNet18 to Swin-T with two tasks on CIFAR-10.

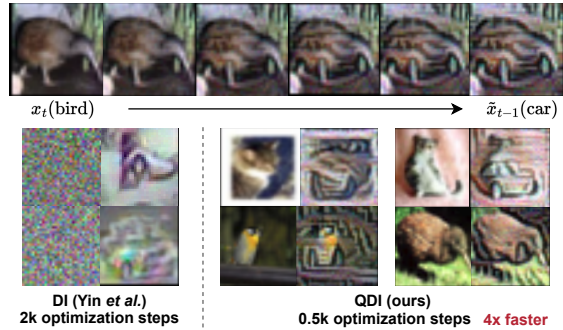


Figure 4. Visualization of QDI and comparison to prior art. **Top:** prior feature synthesis from current task data. **Bottom:** samples for DI that start from $\mathcal{N}(0, 1)$ and QDI from x_t in pairs (left – optimization start point; right – optimization convergence point) for ResNet18 on CIFAR-10 dataset.

tackle this limitation, we revisit knowledge distillation and propose a modified paradigm inspired by the recent advances in knowledge distillation. Additionally, we propose *Quick Deep Inversion (QDI)* that generates synthetic examples using current task instances to enhance knowledge distillation performance in the data-free continual learning setup at a negligible additional cost. Our experimental evaluation shows that without tuning the training configurations, we achieve a significant improvement for both the standard and heterogeneous continual learning scenarios across all datasets, highlighting our proposed approach’s efficacy in real-world applications.

Limitations. Although the proposed approach is applicable for general deep-learning architectures, we found some limitations with the current approach/setup. First, in some cases, the recent architectures can be sub-optimal and degrade the model performance due to their training configuration, model size, or hyper-parameter mismatch. We note that in this work, we did not tune the training configuration or hyper-parameters for each model for fair comparisons. However, model specific hyper-parameter tuning can further improve model performance. Second, our work is not yet applicable to unsupervised CL with heterogeneous architectures, where the learner is expected to learn on unlabelled data with changing architectures because model inversion and KD with unsupervised CL is not a trivial extension due to the absence of class-dependent information to generate class-conditioned examples.

References

- [1] Hongjoon Ahn, Sungmin Cha, Donggyu Lee, and Taesup Moon. Uncertainty-based continual learning with adaptive regularization. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. 1, 2, 3
- [2] Rahaf Aljundi, Min Lin, Baptiste Goujaud, and Yoshua Bengio. Gradient based sample selection for online continual learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. 2, 3
- [3] Lucas Beyer, Xiaohua Zhai, Amélie Royer, Larisa Markeeva, Rohan Anil, and Alexander Kolesnikov. Knowledge distillation: A good teacher is patient and consistent. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2, 5
- [4] Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. Dark experience for general continual learning: a strong, simple baseline. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 1, 2, 3, 4, 6, 7
- [5] Antonio Carta, Andrea Cossu, Vincenzo Lomonaco, and Davide Bacciu. Ex-model: Continual learning from a stream of trained models. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2022. 3, 6
- [6] Keshigeyan Chandrasegaran, Ngoc-Trung Tran, Yunqing Zhao, and Ngai-Man Cheung. Revisiting label smoothing and knowledge distillation compatibility: What was missing? In *Proceedings of the International Conference on Machine Learning (ICML)*, 2022. 2, 4
- [7] Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with a-gem. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019. 6, 7
- [8] Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaisyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and M Ranzato. Continual learning with tiny episodic memories. *arXiv preprint arXiv:1902.10486*, 2019. 6
- [9] Andrea Cossu, Gabriele Graffieti, Lorenzo Pellegrini, Davide Maltoni, Davide Bacciu, Antonio Carta, and Vincenzo Lomonaco. Is class-incremental enough for continual learning? *Frontiers in Artificial Intelligence*, 2022. 4
- [10] Xiaoliang Dai, Hongxu Yin, and Niraj K Jha. Incremental learning using a grow-and-prune paradigm with efficient neural networks. *IEEE Transactions on Emerging Topics in Computing*, 2020. 3
- [11] Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Aleš Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 4
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. 1, 6
- [13] Prithviraj Dhar, Rajat Vikram Singh, Kuan-Chuan Peng, Ziyang Wu, and Rama Chellappa. Learning without memorizing. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 4
- [14] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021. 1, 4
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 1, 2, 3, 4, 6, 7
- [16] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 2, 4
- [17] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 4
- [18] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017. 1
- [19] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE International*

- Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 6
- [20] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 6
- [21] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2015. 6
- [22] Heechul Jung, Jeongwoo Ju, Minju Jung, and Junmo Kim. Less-forgetting learning in deep neural networks. *arXiv preprint arXiv:1607.00122*, 2016. 7
- [23] Alex Krizhevsky. Learning multiple layers of features from tiny images. *University of Toronto*, 05 2012. 6
- [24] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2012. 1, 6
- [25] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. Backpropagation applied to handwritten zip code recognition. *Neural Computation*, 1989. 6
- [26] Yann LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998. 1
- [27] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 1998. 2
- [28] Seungwon Lee, Sima Behpour, and Eric Eaton. Sharing less is more: Lifelong learning in deep networks with selective layer transfer. In *Proceedings of the 38th International Conference on Machine Learning*, 2021. 3
- [29] Xilai Li, Yingbo Zhou, Tianfu Wu, Richard Socher, and Caiming Xiong. Learn to grow: A continual structure learning framework for overcoming catastrophic forgetting. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2019. 2, 3
- [30] Zhizhong Li and Derek Hoiem. Learning without forgetting. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016. 2, 3, 4, 6, 7
- [31] Yahui Liu, Enver Sangineto, Wei Bi, Nicu Sebe, Bruno Lepri, and Marco De Nadai. Efficient training of visual transformers with small-size datasets. *arXiv preprint arXiv:2106.03746*, 2021. 8
- [32] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021. 1, 2, 4, 7
- [33] Divyam Madaan, Jaehong Yoon, Yuanchun Li, Yunxin Liu, and Sung Ju Hwang. Representational continuity for unsupervised continual learning. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2022. 1
- [34] Thomas McGrath, Andrei Kapishnikov, Nenad Tomašev, Adam Pearce, Demis Hassabis, Been Kim, Ulrich Paquet, and Vladimir Kramnik. Acquisition of chess knowledge in alphazero. *arXiv preprint arXiv:2111.09259*, 2021. 1
- [35] Sachin Mehta, Marjan Ghazvininejad, Srinivasan Iyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. Delight: Deep and light-weight transformer. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021. 1
- [36] Seyed Iman Mirzadeh, Arslan Chaudhry, Huiyi Hu, Razvan Pascanu, Dilan Gorur, and Mehrdad Farajtabar. Wide neural networks forget less catastrophically. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2022. 3
- [37] Seyed Iman Mirzadeh, Arslan Chaudhry, Dong Yin, Huiyi Hu, Razvan Pascanu, Dilan Gorur, and Mehrdad Farajtabar. Wide neural networks forget less catastrophically. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2022. 1
- [38] Seyed Iman Mirzadeh, Arslan Chaudhry, Dong Yin, Timothy Nguyen, Razvan Pascanu, Dilan Gorur, and Mehrdad Farajtabar. Architecture matters in continual learning. *arXiv preprint arXiv:2202.00275*, 2022. 3
- [39] Rafael Müller, Simon Kornblith, and Geoffrey E Hinton. When does label smoothing help? *Neurips*, 2019. 4
- [40] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2021. 1
- [41] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 2, 3, 4, 6, 7

- [42] David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience replay for continual learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. 1, 2, 3, 7
- [43] 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. 1, 2, 3
- [44] Jonathan Schwarz, Jelena Luketina, Wojciech M Czarnecki, Agnieszka Grabska-Barwinska, Yee Whye Teh, Razvan Pascanu, and Raia Hadsell. Progress & compress: A scalable framework for continual learning. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2018. 2, 3
- [45] Zhiqiang Shen, Zechun Liu, Dejia Xu, Zitian Chen, Kwang-Ting Cheng, and Marios Savvides. Is label smoothing truly incompatible with knowledge distillation: An empirical study. In *iclr*, 2021. 4
- [46] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2015. 1, 6
- [47] 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 International Conference on Computer Vision (ICCV)*, 2021. 6
- [48] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Proceedings of the AAAI National Conference on Artificial Intelligence (AAAI)*, 2017. 1, 6
- [49] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2019. 1
- [50] Sebastian Thrun. Is learning the n-th thing any easier than learning the first? In *Neurips*, 1996. 1
- [51] Ilya Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Daniel Keysers, Jakob Uszkoreit, Mario Lucic, and Alexey Dosovitskiy. Mlp-mixer: An all-mlp architecture for vision, 2021. 4
- [52] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems (NeurIPS)*, 30, 2017. 1
- [53] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 2, 6
- [54] Jing Xu, Yu Pan, Xinglin Pan, Steven Hoi, Zhang Yi, and Zenglin Xu. Regnet: self-regulated network for image classification. *IEEE Transactions on Neural Networks and Learning Systems*, 2022. 2, 6
- [55] Hongxu Yin, Pavlo Molchanov, Jose M Alvarez, Zhizhong Li, Arun Mallya, Derek Hoiem, Niraj K Jha, and Jan Kautz. Dreaming to distill: Data-free knowledge transfer via DeepInversion. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2, 5, 6, 7
- [56] Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. Lifelong learning with dynamically expandable networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018. 1, 2, 3
- [57] Jaehong Yoon, Saehoon Kim, Eunho Yang, and Sung Ju Hwang. Scalable and order-robust continual learning with additive parameter decomposition. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2020. 3
- [58] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016. 1, 2
- [59] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2017. 1, 2, 3, 6, 7