

Continual Distillation of Teachers from Different Domains

Nicolas Michel^{1,4} Maorong Wang² Jiangpeng He^{3,†} Toshihiko Yamasaki^{1,†}

¹The University of Tokyo ²National Institute of Informatics

³Indiana University Bloomington ⁴Japanese-French Laboratory of Informatics

{nicolas, yamasaki}@cvm.t.u-tokyo.ac.jp, maorong@nii.ac.jp, jhe2@iu.edu

Abstract

Deep learning models continue to scale, with some requiring more storage than many large-scale datasets. Thus, we introduce a new paradigm: Continual Distillation (CD), where a student learns sequentially from a stream of teacher models without retaining access to earlier teachers. CD faces two challenges: teacher training data is unavailable, and teachers have varying expertise. We show that external unlabeled data enables Unseen Knowledge Transfer (UKT), allowing the student to acquire information from domains not present in the training data, while known to the teacher. We also show that sequential distillation causes Unseen Knowledge Forgetting (UKF) when transferred knowledge is lost after training on later teachers. To better trade off between UKT and UKF, we propose Self External Data Distillation (SE2D), a method that preserves logits on external data to stabilize learning across heterogeneous teachers. Experiments on multiple benchmarks show that SE2D reduces UKF and improves cross-domain generalization. The code and implementation for this work are publicly available at: https://github.com/Nicolas1203/continual_distillation.

1. Introduction

Over the last decade, deep learning models have reached unprecedented scales, creating a growing need for computation-efficient strategies. Consequently, Continual Learning (CL) [8, 24, 35] has become a major branch of contemporary deep learning research. The main idea is straightforward: a model is trained on a sequence of datasets where previous data becomes unavailable over time. The rationale is that re-training on both past and current data as new data arrives can be computationally expensive or require substantial storage. In this context, data access is limited during training.

With the recent adoption of Foundation Models (FMs) [1, 9, 11, 21, 26] as the backbone of modern deep

learning, we propose a new paradigm, **Continual Distillation (CD)**, tailored to FMs, illustrated in Figure 1. In CD, instead of learning from a sequence of datasets, we propose to *learn from a sequence of models trained on different datasets*. Specifically, a single student model learns sequentially from multiple teacher models without retaining access to earlier ones. Similar to training data in CL, new FMs are regularly introduced over time, and storing them is cumbersome since they can rapidly require more storage than large-scale datasets. For instance, storing 10B parameters requires approximately 38GB [7], while FMs often exceed 100B parameters. Even accessing FMs through restricted APIs presents serious limitations, as previous versions may become unavailable after updates. Our CD setup reflects a realistic scenario in which one aims to leverage the ever-evolving stream of FMs to train smaller and more specialized models through distillation. In analogy to CL, where a single model benefits from a continuous stream of data, in CD, a single student model benefits from a continuous stream of teacher models.

However, CD introduces two main challenges. First, the original training data of a foundation model is typically unavailable, undisclosed, or prohibitively large to reuse [2]. Thus, choosing distillation data is critical, and such data can potentially emerge from a domain unknown to the teacher [36]. Second, each teacher generally exhibits distinct abilities and performance; for instance, one may excel at recognizing animals, while another may specialize in distinguishing insects.

In this work, we focus on the case of Continual Distillation, where, analogous to Domain Incremental Learning, teachers have been trained on a domain incremental set of datasets. However, we assume that teachers share partially overlapping domains. We believe this assumption to be realistic, as, for example, it is safe to assume that all FMs have been trained on ImageNet. Therefore, we make the following key assumptions: (1) each teacher is trained on data from different domains while sharing a specific domain; (2) the training data for distillation is fixed; and (3) the training

[†]Equal supervision.

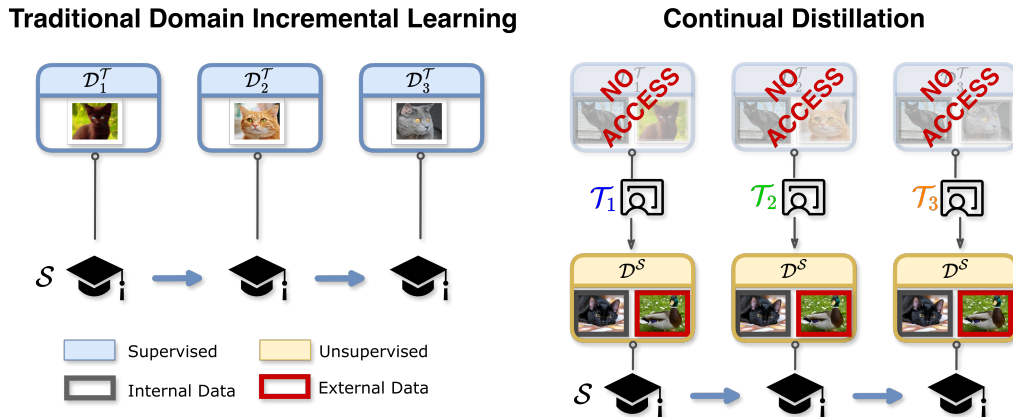


Figure 1. Overview of the Continual Distillation problem. A student model S learns through distillation from a sequence of teachers $\{\mathcal{T}_1, \mathcal{T}_2, \mathcal{T}_3\}$. During distillation, part of the teachers’ training data is unavailable. The objective is for the student to maintain high performance on **all domains known to the teachers**, but not necessarily introduced to the student.

data is unlabeled. The ultimate objective is to obtain a student model that achieves high performance on every domain known by at least one teacher, without having access to all teacher training data or labels. In this context, we find that training data can be decomposed into two categories: **External Data (ED)**, unknown to all teachers, and **Internal Data (ID)**, known to all teachers. Importantly, we discover that ED enables the transfer of knowledge from domains unseen by the student but known by the teacher during distillation, which we refer to as **Unseen Knowledge Transfer (UKT)**. Naturally, we observe that the student inevitably forgets unseen knowledge transferred by previous teachers when learning from future ones. We refer to this phenomenon as **Unseen Knowledge Forgetting (UKF)**. The central problem of CD becomes reaching an optimal UKT-UKF trade-off.

Experimentally, we observe that while the usage of ED enables UKT, mainstream distillation strategies fail to address UKF, as their primary focus relies on maximizing knowledge transfer only. We therefore propose Self External Data Distillation (SE2D), a CL-inspired method tailored for the CD problem, focusing on preserving logits of ED to maintain performance on domains unseen by the student. SE2D allows positive UKT while maintaining performance in older domains. Our contributions are as follows:

- We introduce the paradigm of Continual Distillation, motivated by the practical challenges that arise when previous teacher models are no longer accessible.
- We demonstrate that the choice of distillation data is crucial for transferring knowledge to domains unseen by the student, and choose to take advantage of External Data, unknown to the teachers.
- We identify Unseen Knowledge Forgetting, mitigate it by preserving External Data logits, and validate our approach across various benchmarks.

2. Related Work

Continual Learning. Continual Learning (CL) [5, 24] focuses on enabling models to learn from a sequence of tasks while retaining previously acquired knowledge. Traditionally, each CL task is defined by different non-overlapping datasets with unique properties. In Class Incremental Learning (CIL) [10], each task is composed of a unique set of classes. In Domain Incremental Learning (DIL) [3], classes are shared, but input distributions vary. While access and change in data have been widely studied, we are, to the best of our knowledge, the first to consider an alternative scenario where data is fixed but accessible model change over time. In this case, we focus on a situation where the teachers’ training domain differs, not unlike a DIL setup, where each teacher would be trained on a specific task of a DIL setting. Since foundation models become ubiquitous [26], large, and ever-changing, a paradigm shift from ever-evolving data to ever-evolving teachers appears relevant. While Knowledge Distillation is often employed in CL to mitigate forgetting [22, 28, 32], we instead investigate the phenomenon of forgetting within distillation itself, treating it as the central challenge of our study.

Knowledge Distillation. Knowledge Distillation (KD) [15] is a technique that allows a smaller student model to replicate the behavior of a larger, more capable teacher model. Ideally, KD assumes that both the teacher and student are trained on the same dataset, and has been widely used for model compression [29, 39] and transfer learning [38]. When such data is not available, distillation is considered to be data-free [36]. In that context, the objective is usually to reproduce the teacher’s training domain as closely as possible. Thanks to the capability of knowledge transfer from KD, it has become one of the cornerstones for solving the forgetting problem in CL [12, 28]. Nevertheless, the ideal




Figure 2. Visualization \mathcal{D}_e and \mathcal{D}_i and different teacher training domains in our experimental setting. \mathcal{D}_i and \mathcal{D}_e are publicly available like ImageNet or Wikipedia. However, some private or restricted data might be included in training, such as medical data.

setting (i.e., the students are distilled with the original teachers’ training dataset) of KD rarely holds in CL due to the unavailability of historical data. Consequently, CL methods conduct distillation with new training samples [20], small memory buffer samples [22], generated samples [37], or adversarial samples [14].

While prior work has explored multi-teacher knowledge distillation and the use of distillation in continual learning, our setting differs in several key aspects. First, existing multi-teacher distillation methods typically assume simultaneous access to all teachers, whereas we consider a sequential setting where only one teacher is accessible at a time. Second, unlike continual learning approaches that focus on data streams, our formulation assumes a fixed dataset and instead models a stream of evolving teacher models. Finally, in contrast to standard distillation-based continual learning methods, we do not assume access to past data or replay buffers, nor to previously seen teachers. This combination of constraints defines a distinct and practically relevant scenario that, to the best of our knowledge, has not been explicitly studied before. We therefore frame this problem as Continual Distillation, emphasizing the shift from data-centric to model-centric CL.

3. Continual Distillation

3.1. Generic Definition

We define Continual Distillation (CD) as the process of distilling the knowledge from a **sequence of teacher models** continuously into **one student model**, on a **fixed dataset**. When distilling from one teacher to the student, other teachers are considered unavailable. The distillation process from a given teacher to the student is analogous to a task in standard CL. Formally, given a sequence of teachers $\{\mathcal{T}_0, \mathcal{T}_1, \dots, \mathcal{T}_N\}$, each trained on a dataset \mathcal{D}_t^T , the student \mathcal{S} is optimized to minimize distillation loss \mathcal{L}_{dist} , with respect to \mathcal{T}_t on a distillation dataset \mathcal{D}^S . Importantly, we only consider distillation and no label-dependent loss is considered. In this work, we focus on logits-based distillation as representations are architecture-dependent, computation-intensive, and require access to the entire teacher model. We present the overall procedure in Figure 1. We denote by $\mathbb{D}_{\mathcal{S}}(\mathcal{T}_t, \mathcal{D}^S)$ the operation of distilling

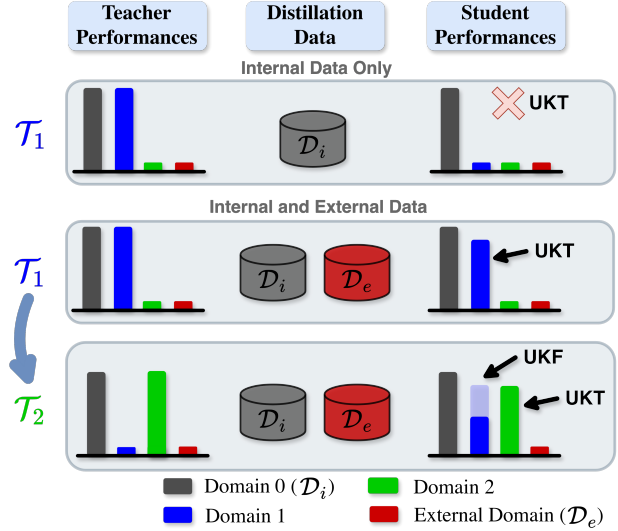


Figure 3. Illustration of Unseen Knowledge Transfer (UKT) and Unseen Knowledge Forgetting (UKF). When distilling only from \mathcal{T}_1 on Internal Data \mathcal{D}_i , only Domain 0 knowledge is transferred. When distilling from the same teacher \mathcal{T}_1 on $\mathcal{D}_e \cup \mathcal{D}_i$, Domains 0 and 1 knowledge is transferred, although the student has never seen Domain 1. However, when continuing the sequence and distilling from \mathcal{T}_2 , while the student acquires knowledge from Domain 2, it forgets part of the knowledge from Domain 1.

from teacher \mathcal{T} trained with \mathcal{D}_t^T to student \mathcal{S} on distillation dataset \mathcal{D}^S .

3.2. Specific Problem Scenario

Traditional KD assumes that the teacher training datasets are available for distillation. In CD, not only are such datasets considered unavailable, but dataset domains might differ from one teacher training to another. In other words, \mathcal{D}_t^T , $\mathcal{D}_{t'}^T$ and \mathcal{D}^S may cover partially or totally different domains. Therefore, various scenarios can be defined in CD depending on the domain overlap between teachers training data \mathcal{D}_t^T and distillation data \mathcal{D}^S . Realistically, when considering FMs, it is safe to assume that part of the training data is shared. Typically, publicly available datasets such as Wikipedia or ImageNet are commonly used for training FMs. However, additional data, exclusive to a specific model, might also be included during training. A visualisation of such a scenario is presented in Figure 2. Formally, we consider the case of partially exclusive teacher domains such that all teachers share a specific domain:

$$\mathcal{D}_t^T \cap \mathcal{D}_{t'}^T = \mathcal{D}_i, \forall t, t', \quad (1)$$

with \mathcal{D}_i the **Internal Data (ID)**. The remaining distillation data is considered unknown to all teachers, hence:

$$\mathcal{D}^S = \mathcal{D}_e \cup \mathcal{D}_i, \quad (2)$$

where \mathcal{D}_e denotes the **External Data (ED)** such that for any teacher of index t , $\mathcal{D}_e \cap \mathcal{D}_t = \emptyset$.

3.3. Unseen Knowledge Transfer and Forgetting

We define ‘unseen’ as domains not present in the student’s training data but present in the teacher’s knowledge. In CD, \mathcal{D}_e is unknown to the teachers. Such an ED can either be introduced or can appear unknowingly when generating data, a standard procedure of data-free distillation. While this could appear to be a limitation, we observe that leveraging ED allows the student to acquire knowledge about domains that have never been explicitly seen during training. We refer to this phenomenon as **Unseen Knowledge Transfer (UKT)**. Intuitively, when integrating ED, generic knowledge is transferred because of the teacher’s uncertainty. Conversely, when the teacher is confident, specific knowledge only is transferred. In CD, we propose to take advantage of UKT by purposefully integrating ED during distillation to extract additional knowledge from the teacher.

However, in CD, the student sequentially learns from multiple teachers, each providing distinct unseen domain knowledge. While UKT enables the student to acquire information about domains not directly represented in the teacher data, this transferred knowledge is often fragile. As the student learns from subsequent teachers, they tend to lose information previously transferred from earlier ones. We refer to this phenomenon as **Unseen Knowledge Forgetting (UKF)**. UKF differs fundamentally from the catastrophic forgetting traditionally studied in Domain Incremental Learning, as the forgotten knowledge does not originate from the student’s own training data but from the teacher’s knowledge. Since the student is never directly exposed to such knowledge, we call it *unseen knowledge*. An intuitive illustration of UKF and UKT is given in Figure 3.

3.4. Self External Data Distillation (SE2D)

We introduce Self-External Data Distillation (SE2D), a method designed to mitigate UKF in Continual Distillation. In SE2D, the student model is trained not only from the current teacher but also from its own checkpoint saved after the previous task. Such a strategy is quite common in Continual Learning [20, 27]; however, we propose to adapt it to the specific problem at hand. Therefore, the distillation process from the checkpoint is performed exclusively on *external data* \mathcal{D}_e , which is unknown to all teachers. An overview of our proposed approach is given in Figure 4.

Prior observations indicate that performance on domains unseen by the student yet known by past teachers depends heavily on these external samples. We restrict self-distillation to external data \mathcal{D}_e to specifically preserve knowledge that is not directly supported by the shared internal domain \mathcal{D}_i . Applying self-distillation on \mathcal{D}_i would mainly reinforce already stable knowledge, while our goal

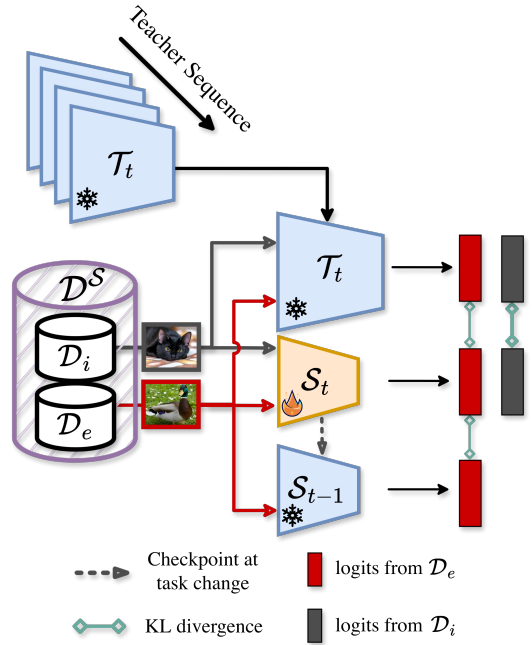


Figure 4. Overview of Self External Data Distillation (SE2D). While the distillation from the teacher sequence is done on the entire distillation data, distillation from the checkpoint of the student at the end of the previous task is done only on \mathcal{D}_e .

is to maintain transferred knowledge from unseen domains, which is primarily captured through \mathcal{D}_e .

Practically, at each distillation step t , the student \mathcal{S}_t learns from both the current teacher \mathcal{T}_t and the previous student checkpoint \mathcal{S}_{t-1} :

$$\mathcal{L}_{\text{SE2D}} = \mathcal{L}_{\text{KD}}(\mathcal{S}_t, \mathcal{T}_t; \mathcal{D}^S) + \mathcal{L}_{\text{KD}}(\mathcal{S}_t, \mathcal{S}_{t-1}; \mathcal{D}_e), \quad (3)$$

where \mathcal{L}_{KD} denotes the temperature-scaled Kullback–Leibler divergence between the softmax distributions of the student logits $z_S(x)$ and teacher logits $z_T(x)$:

$$\mathcal{L}_{\text{KD}}(\mathcal{S}, \mathcal{T}; \mathcal{D}^S) = T^2 \mathbb{E}_{x \sim \mathcal{D}^S} \left[\text{KL} \left(\sigma \left(\frac{z_T(x)}{T} \right) \parallel \sigma \left(\frac{z_S(x)}{T} \right) \right) \right], \quad (4)$$

where T is the distillation temperature and $\sigma(\cdot)$ denotes the softmax function. This simple yet effective mechanism enables the student to accumulate knowledge over time while retaining transferable information from previous distillation stages. In Section 5, we present experimental results of SE2D across various benchmarks.

4. Experimental Setup

In the following, we describe the experimental setup for Continual Distillation. More details regarding the implementation are given in the appendix.

4.1. Domain Selection

Teachers Domains. To reproduce a Continual Distillation context, we work on datasets containing data of the same classes but coming from different domains. For example, in Figure 1, each domain contains the class “cat”; however, the color of the cat is different depending on the considered domain. Another example would be the condition in which the picture was taken (inside or outside). Such a domain shift is typically considered in DIL [16, 33].

External Data Selection. In CD, the choice of ED is crucial. In realistic scenarios, ED could be introduced accidentally, as knowing FMs exact training domain is unlikely. In this work, we deliberately exploit such external data, specifically selected to fall outside all teacher training domains, to study its influence on knowledge transfer and forgetting. We consider two scenarios: (1) **related external domains**, where the data share the same semantic classes as the teacher domains (e.g., if teachers are trained on sea animals such as dolphins, sharks, and whales, the external data may include jellyfish); and (2) **unrelated external domains**, where the data are semantically distinct (e.g., using images of trucks or digits instead of marine animals). In both cases, teachers have minimal performance on ED. Additional discussion is included in appendix.

4.2. Teacher Selection

For the experimental setup, we hypothesize that most contemporary FMs share a substantial amount of common knowledge but differ primarily in their performance on specific tasks. Based on this assumption, we train a sequence of teacher models that all share a common domain while each possesses a unique domain not shared by any other teacher. Consequently, all teachers in the sequence provide a shared knowledge base together with a teacher-specific domain that the student must learn and retain through the use of external data. An example is given in Table 1 where each teacher is trained on pairs (0, 1), (0, 2), and so on.

4.3. Datasets

For simulating CD, we build upon Domain Incremental Learning datasets. We consider each domain separately and pre-train teachers on subsets of such domains. **CIFAR20** [30], a variation of CIFAR-100 using the 20 superclasses instead of the 100 fine-grained classes. Since in CIFAR-100, each superclass is composed of 5 subclasses, using the superclasses allows for defining 5 different domains where each domain is images from different subclasses but identical super-class. CIFAR20 contains 10,000 train images and 2,000 test images per domain. All such domains are considered related. Additionally, we experiment with CUB [34] and MNIST as unrelated

datasets. **Digits**, that we define as the combination of various digit datasets, each one representing a different domain. Namely, we mix MNIST [19], MNIST-M [13], USPS [17] and SVHN [23]. As a related domain, we consider KM-NIST [6], a dataset composed of Japanese Hiragana. We reckon this dataset as having similar difficulty as MNIST, even though of different classes. **DomainNet**, an adapted version of the DomainNet dataset [25], containing six visual domains: Real, Clipart, Painting, Infograph, Sketch, and Quickdraw. Each domain shares the same set of 345 classes but differs significantly in style and texture statistics. This dataset contains around 600,000 images, and the domains are unbalanced, increasing the difficulty. All domains are considered related.

4.4. Considered Methods

We selected mainstream and recent state-of-the-art baselines for logits-based distillation methods. We considered the following. **KL-divergence**. This method consists of a standard distillation loss and serves as a baseline [18]. **Logits Standardization (LS)** [31] is a method recently proposed that improves upon logits distillation by standardizing student and teacher logits. **Medium Difficulty Samples (MDS)** [4] is a data-pruning strategy that considers distilling on samples of medium difficulty only. The original method was proposed in a supervised scenario where the teacher’s cross-entropy was considered as the criterion for assessing sample difficulty. In our setup, we adapted this method with the same principle, using teacher entropy as a sample difficulty estimator. **Decoupled Knowledge Distillation (DKD)** [39], a standard distillation method that decomposes the conventional distillation objective into two components: a target class term and a non-target class term. While most efficient in supervised scenarios, DKD can easily be used in unsupervised scenarios by considering the teacher’s maximum prediction as the target. **Self-Distillation**. Distillation is regularly used in CL [12, 22, 28]. A common strategy is to save a checkpoint of the current model at the end of the task and use such a checkpoint for distillation when training on the subsequent task. This methods uses both \mathcal{D}_i and \mathcal{D}_e for distillation.

5. Experimental Results

5.1. External Data Impact on UKT and UKF

External Data Improves Knowledge Transfer. A first observation that can be made in the CD setting is the necessity of training with ED to fully distill knowledge from the teacher. To showcase this effect, we train a sequence of teachers on CIFAR20 on domain pairs $\{0, 1\}$, $\{0, 2\}$, $\{0, 3\}$, while distilling to the student model on ID only (domain 0), using the KL-divergence. The results are displayed in Table 1, where the accuracy of the student after each task is

Table 1. Accuracy (%) of the student model on test sets, domain-wise, on CIFAR20, for various domain overlaps, after distilling for 1 epoch. The larger the ratio $|\mathcal{D}_e|/|\mathcal{D}^S|$, the more the student performs on domains unseen during training (underlined values).

Domain	0	1	2	3	$ \mathcal{D}_e / \mathcal{D}^S $
After $\mathbb{D}(\mathcal{T}_{01}, \mathcal{D}_0^S)$	93.25	37.80	45.30	36.90	0%
After $\mathbb{D}(\mathcal{T}_{02}, \mathcal{D}_0^S)$	93.95	31.60	<u>42.20</u>	35.25	0%
After $\mathbb{D}(\mathcal{T}_{03}, \mathcal{D}_0^S)$	92.10	32.00	39.05	<u>33.00</u>	0%
After $\mathbb{D}(\mathcal{T}_{012}, \mathcal{D}_{014}^S)$	92.70	93.15	<u>68.35</u>	53.80	33%
After $\mathbb{D}(\mathcal{T}_{013}, \mathcal{D}_{014}^S)$	92.45	92.95	48.95	<u>72.55</u>	33%
After $\mathbb{D}(\mathcal{T}_{01}, \mathcal{D}_{04}^S)$	96.35	<u>77.15</u>	48.95	48.45	50%
After $\mathbb{D}(\mathcal{T}_{02}, \mathcal{D}_{04}^S)$	96.35	43.30	<u>80.10</u>	46.55	50%
After $\mathbb{D}(\mathcal{T}_{03}, \mathcal{D}_{04}^S)$	95.70	49.40	57.60	<u>77.80</u>	50%
After $\mathbb{D}(\mathcal{T}_{01}, \mathcal{D}_{034}^S)$	94.60	<u>85.20</u>	51.85	57.15	66%
After $\mathbb{D}(\mathcal{T}_{02}, \mathcal{D}_{034}^S)$	94.60	44.00	<u>83.55</u>	51.55	66%

reported. It is important to note that initially, considered teachers achieve above 95% accuracy on their respective domains. Eventually, distilling only on the domain 0 yields competitive performance on this domain; however, performances on other domains remain extremely limited at any training step. The student performs only on domains that have been encountered during training. In Table 1, we maintain the same teacher sequence but include ED for distillation. Such data are from the domain 4 of CIFAR20, which is unknown to the teacher. In this case, we can observe that the student maintains performance on domain 0 while achieving much stronger performances on other domains, despite never being encountered during training. UKT is observed only when distilling with external data.

External Data Accentuates UKF. While ED typically facilitates knowledge transfer, it can bias the model toward the most recent teacher at the expense of earlier knowledge. Tables 2 and 3 show that ED does not guarantee uniform gains and may even accentuate forgetting. For instance, on Digits, DKD and LS suffer significant drops on *SVHN* and *MNIST-M*; notably, DKD’s *MNIST-M* performance falls from 54.50% to 33.84%. However, this effect is not universal, as DKD shows slight gains on DomainNet’s *Infograph* (Table 4). This suggests that while ED can trigger substantial UKF, the impact depends heavily on the distillation method and domain distribution.

5.2. UKF Mitigation

Traditional Methods and UKF. In Tables 2, 3, and 4, we present the results for all methods in CD setups with related ED. Firstly, it can easily be seen that, as expected, performances on earlier domains tend to be the lowest for all methods, which is a direct demonstration of UKF where the student forgot transferred knowledge on unseen domains. In all scenarios, all traditional distillation methods suffer from UKF. Secondly, Self-Distillation can partially miti-




Figure 5. Average Accuracy of the student across all domains known by teachers at the end of training on CIFAR20.

gate UKF and surpass distillation approaches due to the regularization effect of self-distillation. However, a clear trade-off emerges between UKT and UKF, not unlike the stability-plasticity trade-off in Continual Learning, where a model has to balance learning and remembering capabilities. Therefore, while self-distillation surpasses most baselines on earlier domains, it struggles to achieve high performances on more recent domains.

Performances of SE2D. SE2D allows for better knowledge retention when compared to most baselines, especially on older tasks where it can surpass baselines by more than 9% on domain 1 of CIFAR20, as reported in Table 2. However, such results hold only for sufficiently related external data where the student can adequately learn from the teacher from one task to another. In Table 2, when using MNIST in combination, the advantages of SE2D become largely limited. This behavior is even more pronounced when training on DomainNet, where SE2D falls behind Self Distillation. We identify the main reason to be the low teacher quality, giving poor supervision on ED for SE2D. Additionally, results in Table 4 show that even when using domains from DomainNet (supposedly related), performance gain is inconsistent for all methods. We believe the vast discrepancy between domains in this dataset hinders UKT, highlighting the importance of the origin of ED for efficient CD. This is discussed in more detail in appendix.

5.3. Discussion

External Data Ratio Matters. Another phenomenon that can be observed in Table 1 is that increasing $\frac{|\mathcal{D}_e|}{|\mathcal{D}^S|}$, the proportion of ED compared to ID, enhances performances on domains unseen by the student. Therefore, a clear trend emerges: the more data unknown to the teacher is used, the stronger the UKT is observed.

External Data Origin Matters. Naturally, the origin of the external data has an impact on the intensity of UKT.

Table 2. Performances (% , higher is better) of the student at the end of training on CIFAR20 for 4 scenarios. Internal Data Only (D0), Related External Data (D4), CUB as ED, and MNIST as ED. The number of runs is set to 3. The gain columns shows the gain over using Internal Data Only. Grey : Internal Data; Blue : Domain known by the teacher; Red : External Data (ED); White: Ignored.

CIFAR20 - Internal Data Only							
Method	D0	D1	D2	D3	D4 ✗	Avg. (0-3)	Gain (↑)
\mathcal{T}_{best} (upper bound)	97.75	95.80	96.70	95.75	-	96.5	0.00
KL-divergence	98.10 ± 0.05	41.40 ± 0.84	53.38 ± 0.35	54.87 ± 3.70	-	61.94 ± 1.24	0.00
DKD [CVPR'22]	95.70 ± 1.55	35.53 ± 1.94	45.42 ± 2.05	41.02 ± 2.55	-	54.42 ± 2.02	0.00
LS [CVPR'24]	96.50 ± 1.52	39.13 ± 2.90	49.77 ± 4.40	49.58 ± 5.78	-	58.75 ± 3.65	0.00
MDS [ICLR'25]	94.20 ± 0.73	34.30 ± 2.80	44.70 ± 1.40	41.00 ± 1.90	-	53.55 ± 1.72	0.00
Self-Distillation	97.45 ± 0.43	37.58 ± 2.06	50.42 ± 1.41	45.83 ± 2.15	-	57.82 ± 1.51	0.00
CIFAR20 + Related External Data							
Method	D0	D1	D2	D3	D4	Avg. (0-3)	Gain (↑)
\mathcal{T}_{best} (upper bound)	97.75	95.80	96.70	95.75	-	96.5	0.00
KL-divergence	97.05 ± 0.09	48.55 ± 1.15	55.08 ± 0.70	84.77 ± 0.87	-	71.36 ± 0.70	9.42
DKD [CVPR'22]	96.05 ± 0.50	44.13 ± 0.98	51.67 ± 0.92	68.55 ± 2.60	-	65.10 ± 1.25	10.68
LS [CVPR'24]	96.85 ± 0.15	47.25 ± 0.69	54.25 ± 0.46	83.20 ± 1.87	-	70.39 ± 0.79	11.64
MDS [ICLR'25]	96.55 ± 0.07	45.26 ± 2.10	54.90 ± 0.71	73.51 ± 0.56	-	67.56 ± 0.86	14.01
Self-Distillation	97.71 ± 0.18	61.23 ± 0.83	64.21 ± 0.51	76.58 ± 0.94	-	74.93 ± 0.61	17.11
SE2D (ours)	97.46 ± 0.19	70.71 ± 1.05	62.85 ± 0.50	73.65 ± 1.67	-	76.17 ± 0.85	n/a
CIFAR20 + CUB							
Method	D0	D1	D2	D3	CUB	Avg. (0-3)	Gain (↑)
\mathcal{T}_{best} (upper bound)	97.75	95.80	96.70	95.75	-	96.5	0.00
KL-divergence	97.24 ± 0.37	43.89 ± 1.02	55.13 ± 0.76	71.80 ± 1.73	-	67.02 ± 0.97	5.08
DKD [CVPR'22]	93.39 ± 0.88	33.46 ± 2.51	43.10 ± 2.93	40.70 ± 2.20	-	52.66 ± 2.13	-1.76
LS [CVPR'24]	94.79 ± 3.17	38.53 ± 4.87	50.21 ± 5.74	59.32 ± 11.60	-	60.71 ± 6.34	1.96
MDS [ICLR'25]	97.15 ± 0.50	41.12 ± 0.96	52.38 ± 1.76	60.38 ± 2.48	-	62.76 ± 0.89	9.21
Self-Distillation	97.47 ± 0.12	47.97 ± 1.07	58.40 ± 1.34	61.97 ± 1.82	-	66.45 ± 1.09	8.63
SE2D (ours)	97.74 ± 0.10	53.93 ± 0.43	58.02 ± 0.45	64.54 ± 1.81	-	68.56 ± 0.70	n/a
CIFAR20 + MNIST							
Method	D0 (ID)	D1	D2	D3	MNIST	Avg. (0-3)	Gain (↑)
\mathcal{T}_{best} (upper bound)	97.75	95.80	96.70	95.75	-	96.5	0.00
KL-divergence	94.45 ± 4.20	38.24 ± 6.06	48.44 ± 8.56	57.97 ± 15.71	-	59.78 ± 8.63	-2.16
DKD [CVPR'22]	91.00 ± 3.35	30.71 ± 3.60	41.36 ± 3.61	37.96 ± 4.56	-	50.26 ± 3.78	-4.16
LS [CVPR'24]	92.04 ± 4.43	35.93 ± 6.37	45.79 ± 7.50	52.42 ± 13.24	-	56.54 ± 7.88	-2.21
MDS [ICLR'25]	86.83 ± 1.04	28.53 ± 2.20	37.92 ± 2.41	37.38 ± 3.24	-	47.17 ± 1.57	-6.38
Self-Distillation	91.90 ± 4.94	36.35 ± 9.46	45.23 ± 9.18	43.58 ± 12.43	-	54.26 ± 9.00	-3.56
SE2D (ours)	92.64 ± 4.66	39.94 ± 9.73	47.55 ± 9.55	48.88 ± 13.23	-	57.25 ± 9.29	n/a

Table 3. Performances (% , higher is better) of the student at the end of training on Digits for 2 scenarios. Internal Data Only (MNIST) and Related External Data (KMNIST). The number of runs is set to 3. Average and standard deviations are reported.

Digits - Internal Data Only							
Method	MNIST	SVHN	MNIST-M	USPS	KMNIST ✗	Avg.	Gain (↑)
\mathcal{T}_{best} (upper bound)	99.2	97.84	99.25	98.8	-	98.77	0.00
KL-divergence	99.17 ± 0.04	35.80 ± 1.49	62.80 ± 2.69	95.81 ± 0.46	-	73.40 ± 1.17	0.00
DKD [CVPR'22]	98.70 ± 0.11	33.35 ± 1.28	54.50 ± 2.45	95.07 ± 0.43	-	70.40 ± 1.07	0.00
LS [CVPR'24]	99.13 ± 0.09	37.60 ± 1.40	64.10 ± 1.73	96.00 ± 0.25	-	74.21 ± 0.87	0.00
MDS [ICLR'25]	98.15 ± 1.13	34.51 ± 4.80	54.97 ± 9.49	93.21 ± 3.01	-	70.21 ± 3.80	0.00
Self-Distillation	99.23 ± 0.10	35.08 ± 1.03	65.87 ± 1.38	95.28 ± 0.53	-	73.87 ± 0.76	0.00
Digits + Related External Data							
Method	MNIST	SVHN	MNIST-M	USPS	KMNIST	Avg.	Gain (↑)
KL-divergence	99.13 ± 0.05	31.53 ± 1.55	59.84 ± 2.57	96.51 ± 0.10	-	71.75 ± 1.07	-1.65
DKD [CVPR'22]	98.35 ± 0.32	25.21 ± 1.62	33.84 ± 4.28	92.87 ± 1.63	-	62.57 ± 1.96	-7.83
LS [CVPR'24]	99.13 ± 0.08	32.28 ± 0.37	61.47 ± 2.15	96.33 ± 0.33	-	72.30 ± 0.73	-1.91
MDS [ICLR'25]	99.12 ± 0.04	33.03 ± 1.79	60.74 ± 0.97	96.13 ± 0.05	-	62.50 ± 0.90	-7.71
Self-Distillation	99.38 ± 0.01	55.86 ± 1.60	90.76 ± 0.35	96.33 ± 0.15	-	85.58 ± 0.53	11.71
SE2D (ours)	99.33 ± 0.04	61.84 ± 2.05	90.44 ± 0.18	96.33 ± 0.10	-	87.00 ± 0.60	n/a

Table 4. Performances (% , higher is better) of the student at the end of training on DomainNet for 2 scenarios. Internal Data Only (Clipart) and Related External Data (Sketch). The number of runs is set to 3. Average and standard deviations are reported.

DomainNet - Internal Data Only								
Method	Clipart	Infograph	Painting	Quickdraw	Real	Sketch \times	Avg.	Gain (\uparrow)
\mathcal{T}_{best} (upper bound)	74.35	35.83	66.66	66.19	78.90	-	64.39	0.00
KL-divergence	77.08 \pm 0.45	18.16 \pm 0.24	44.17 \pm 0.50	17.19 \pm 0.73	67.29 \pm 0.24	-	44.78 \pm 0.29	0.00
DKD [CVPR'22]	77.32 \pm 0.22	18.04 \pm 0.09	42.76 \pm 0.40	17.43 \pm 0.31	63.99 \pm 0.14	-	43.91 \pm 0.17	0.00
LS [CVPR'24]	77.21 \pm 0.37	17.21 \pm 0.91	40.95 \pm 0.14	16.42 \pm 1.01	60.39 \pm 0.20	-	42.43 \pm 0.44	0.00
MDS [ICLR'25]	76.92 \pm 0.42	18.08 \pm 0.23	42.14 \pm 0.38	17.44 \pm 1.03	63.29 \pm 0.20	-	43.57 \pm 0.34	0.00
Self-Distillation	80.57 \pm 0.08	20.85 \pm 0.33	46.23 \pm 0.11	23.53 \pm 0.25	65.38 \pm 0.04	-	47.31 \pm 0.11	0.00
DomainNet + Related External Data								
Method	Clipart	Infograph	Painting	Quickdraw	Real	Sketch	Avg.	Gain (\uparrow)
KL-divergence	76.00 \pm 0.18	18.89 \pm 0.09	44.77 \pm 0.79	15.53 \pm 0.15	70.65 \pm 0.46	-	45.17 \pm 0.31	0.39
DKD [CVPR'22]	76.53 \pm 0.14	18.70 \pm 0.44	44.24 \pm 0.32	16.24 \pm 0.40	68.29 \pm 0.34	-	44.80 \pm 0.12	0.89
LS [CVPR'24]	75.28 \pm 0.39	16.88 \pm 0.71	40.78 \pm 0.49	15.47 \pm 0.43	62.76 \pm 0.64	-	42.24 \pm 0.31	-0.20
MDS [ICLR'25]	75.32 \pm 0.39	17.53 \pm 0.07	42.28 \pm 0.11	16.12 \pm 0.61	67.29 \pm 0.11	-	43.71 \pm 0.19	0.14
Self-Distillation	80.10 \pm 0.18	21.53 \pm 0.09	48.15 \pm 0.23	25.28 \pm 0.19	68.75 \pm 0.01	-	48.76 \pm 0.07	1.45
SE2D (ours)	78.05 \pm 0.11	21.98 \pm 0.29	47.76 \pm 0.44	23.81 \pm 0.34	68.43 \pm 0.14	-	48.01 \pm 0.20	n/a

Intuitively, a large domain difference between external data and teacher domains might make distillation more challenging. To showcase the impact of the domain gap, we experimented with **related and unrelated external domains** on CIFAR20. Table 2 presents the results with various ED scenarios. Notably, it can be observed that leveraging ED leads to consistent improvement in average on all domains when related enough to ID. For example, when using D4 and CUB as ED, the performances of KL-divergence distillation increase from 61.94% to 71.36% and 67.02%, respectively. However, when using MNIST as ED, the performances slightly drop to 59.78%. Such a trend is observed for all considered methods, as it can be observed in Figure 5. Semantically, CUB consists of images of birds and is more similar to CIFAR20 than MNIST, even though the number of classes does not align. Interestingly, we observe that leveraging ED is essential to promote UKT. However, as the domain shift between ED and ID increases (e.g., in CUB), performance tends to degrade. When the domain gap becomes too large (as in MNIST), using ED can even result in lower performance than without it.

Limitations of SE2D. While SE2D reduces UKF, its impact relies on (1) the domain gap between teacher and external data and (2) teacher performance on domains unseen by the student. SE2D also requires data-origin knowledge; the student must distinguish between the teacher’s known and unknown domains. This is particularly complex when data are generated to imitate training sets, making it non-trivial to identify data outside the teacher’s domain.

6. Conclusions and Future Work

In this work, we introduced a new paradigm titled Continual Distillation, where a single model learns on a fixed dataset from a sequence of teachers. Such a new setup is relevant

in the context of ever-evolving Foundation Models, which are costly to train, expensive to store, demanding to run inference on, and in many cases only accessible via restricted APIs. In such a context, we observe that the domain of origin of distillation data is crucial for controlling which knowledge is indeed transferred to the student. Notably, we unveil two characteristics: Unknown Knowledge Transfer and Unknown Knowledge Forgetting, which represent the ability of the student to modify its knowledge on domains that they have never encountered. Such knowledge control depends only on the teacher and the data used for distillation. The objective then becomes reaching the best UKT-UKF trade-off. In that sense, we proposed Self External Data Distillation (SE2D), which allows us to reduce UKF and maintain strong average performance on all domains, including domains unseen during training. However, we uncover that the domain gap between external data and the teacher domain must be carefully considered in order to foster UKT. Similarly, performant teachers are required for SE2D performance to be ensured.

Eventually, UKT comprises opportunities and risks, as uncontrolled or undesired knowledge could be involuntarily embedded in a student model through distillation depending on the considered data. Such a vulnerability could be easily exploited and introduce unknown bias to model training. Such an aspect of UKT should be explored in future work. Potential future directions include working with larger models, such as language or multimodal models.

Acknowledgments

This work was partially financially supported by JST ASPIRE Program, Japan, Grant Number JPMJAP2303. This work was partially supported by the JSPS Postdoctoral Fellowship for Research in Japan (Fellowship ID: P24752).

References

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. 1
- [2] Muhammad Awais, Muzammal Naseer, Salman Khan, Rao Muhammad Anwer, Hisham Cholakkal, Mubarak Shah, Ming-Hsuan Yang, and Fahad Shahbaz Khan. Foundation models defining a new era in vision: a survey and outlook. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(4):2245–2264, 2025. 1
- [3] 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*, pages 15920–15930, 2020. 2
- [4] Yudong Chen, Xuwei Xu, Frank de Hoog, Jiajun Liu, and Sen Wang. Medium-difficulty samples constitute smoothed decision boundary for knowledge distillation on pruned datasets. In *The Thirteenth International Conference on Learning Representations*, 2025. 5
- [5] Zhiyuan Chen and Bing Liu. Lifelong machine learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 12(3):1–207, 2018. 2
- [6] Tarin Clanuwat, Mikel Bober-Irizar, Asanobu Kitamoto, Alex Lamb, Kazuaki Yamamoto, and David Ha. Deep learning for classical japanese literature. *arXiv preprint arXiv:1812.01718*, 2018. 5
- [7] Ankita De, Edward Wang, Rohan Varma, Anjali Sridhar, and Kartikay Khandelwal. Scaling multimodal foundation models in torchmultimodal with pytorch distributed. <https://pytorch.org/blog/scaling-multimodal-foundation-models-in-torchmultimodal-with-pytorch-distributed>, 2025. 1
- [8] 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*, 44(7):3366–3385, 2021. 1
- [9] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pages 4171–4186, 2019. 1
- [10] Na Dong, Yongqiang Zhang, Mingli Ding, and Yancheng Bai. Class-incremental object detection. *Pattern Recognition*, 139:109488, 2023. 2
- [11] 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 *International Conference on Learning Representations*, 2021. 1
- [12] Arthur Douillard, Matthieu Cord, Charles Ollion, Thomas Robert, and Eduardo Valle. Podnet: Pooled outputs distillation for small-tasks incremental learning. In *16th European Conference on Computer Vision (ECCV)*, pages 86–102, 2020. 2, 5
- [13] Yaroslav Ganin, Evgeniya Ustinova, Hubert Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. In *International Conference on Machine Learning*, 2016. 5
- [14] Dipam Goswami, Albin Soutif-Cormerais, Yuyang Liu, Sandesh Kamath, Bart Twardowski, Joost Van De Weijer, et al. Resurrecting old classes with new data for exemplar-free continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 28525–28534, 2024. 3
- [15] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 2
- [16] Yen-Chang Hsu, Yen-Cheng Liu, Anita Ramasamy, and Zsolt Kira. Re-evaluating continual learning scenarios: A categorization and case for strong baselines. *arXiv preprint arXiv:1810.12488*, 2018. 5
- [17] Jonathan J. Hull. A database for handwritten text recognition research. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(5):550–554, 2002. 5
- [18] Solomon Kullback. *Information theory and statistics*. Courier Corporation, 1997. 5
- [19] Yann LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998. 5
- [20] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12):2935–2947, 2017. 3, 4
- [21] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019. 1
- [22] Nicolas Michel, Maorong Wang, Ling Xiao, and Toshihiko Yamasaki. Rethinking momentum knowledge distillation in online continual learning. In *Proceedings of the 41st International Conference on Machine Learning*, pages 35607–35622. PMLR, 2024. 2, 3, 5
- [23] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bisacco, Baolin Wu, Andrew Y Ng, et al. Reading digits in natural images with unsupervised feature learning. In *NIPS workshop on deep learning and unsupervised feature learning*, page 7, 2011. 5
- [24] German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113:54–71, 2019. 1, 2
- [25] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source

- domain adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1406–1415, 2019. 5
- [26] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763, 2021. 1, 2
- [27] Amal Rannen, Rahaf Aljundi, Matthew B Blaschko, and Tinne Tuytelaars. Encoder based lifelong learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1320–1328, 2017. 4
- [28] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2001–2010, 2017. 2, 5
- [29] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*, 2014. 2
- [30] Stefan Stojanov, Samarth Mishra, Ngoc Anh Thai, Nikhil Dhanda, Ahmad Humayun, Chen Yu, Linda B Smith, and James M Rehg. Incremental object learning from contiguous views. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8777–8786, 2019. 5
- [31] Shangquan Sun, Wenqi Ren, Jingzhi Li, Rui Wang, and Xiaochun Cao. Logit standardization in knowledge distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15731–15740, 2024. 5
- [32] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. 2
- [33] Guido M van de Ven, Tinne Tuytelaars, and Andreas S Tolias. Three types of incremental learning. *Nature Machine Intelligence*, 4(12):1185–1197, 2022. 5
- [34] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The caltech-ucsd birds-200-2011 dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011. 5
- [35] Liyuan Wang, Xingxing Zhang, Hang Su, and Jun Zhu. A comprehensive survey of continual learning: Theory, method and application. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. 1
- [36] Yuzheng Wang, Ding kang Yang, Zhaoyu Chen, Yang Liu, Siao Liu, Wenqiang Zhang, Lihua Zhang, and Lizhe Qi. De-confounded data-free knowledge distillation for handling distribution shifts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12615–12625, 2024. 1, 2
- [37] Chenshen Wu, Luis Herranz, Xialei Liu, Joost Van De Weijer, Bogdan Raducanu, et al. Memory replay gans: Learning to generate new categories without forgetting. *Advances in Neural Information Processing Systems*, 31, 2018. 3
- [38] Junho Yim, Donggyu Joo, Jihoon Bae, and Junmo Kim. A gift from knowledge distillation: Fast optimization, network minimization and transfer learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4133–4141, 2017. 2
- [39] Borui Zhao, Quan Cui, Renjie Song, Yiyu Qiu, and Jiajun Liang. Decoupled knowledge distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11953–11962, 2022. 2, 5