

# Meta-Learning with Less Forgetting on Large-Scale Non-Stationary Task Distributions

Zhenyi Wang<sup>1</sup>, Li Shen<sup>2\*</sup>, Le Fang<sup>1</sup>, Qiuling Suo<sup>1</sup>, Donglin Zhan<sup>3</sup>, Tiehang Duan<sup>4</sup>, and Mingchen Gao<sup>1\*</sup>

<sup>1</sup> State University of New York at Buffalo, USA {zhenyiwa, lefang, qiulings, mgao8}@buffalo.edu

<sup>2</sup> JD Explore Academy, Beijing, China mathshenli@gmail.com

<sup>3</sup> Columbia University, New York, NY, USA dz2478@columbia.edu

<sup>4</sup> Meta, Seattle, WA, USA tiehang.duan@gmail.com

**Abstract.** The paradigm of machine intelligence moves from purely supervised learning to a more practical scenario when many loosely related unlabeled data are available and labeled data is scarce. Most existing algorithms assume that the underlying task distribution is stationary. Here we consider a more realistic and challenging setting in that task distributions evolve over time. We name this problem as **Semi-supervised meta-learning with Evolving Task diStributions**, abbreviated as **SETS**. Two key challenges arise in this more realistic setting: (i) how to use unlabeled data in the presence of a large amount of unlabeled out-of-distribution (OOD) data; and (ii) how to prevent catastrophic forgetting on previously learned task distributions due to the task distribution shift. We propose an **OOD Robust and knowleDge presERved semi-supeRvised meta-learning approach (ORDER)** <sup>‡</sup>, to tackle these two major challenges. Specifically, our ORDER introduces a novel mutual information regularization to robustify the model with unlabeled OOD data and adopts an optimal transport regularization to remember previously learned knowledge in feature space. In addition, we test our method on a very challenging dataset: SETS on large-scale non-stationary semi-supervised task distributions consisting of (at least) 72K tasks. With extensive experiments, we demonstrate the proposed ORDER alleviates forgetting on evolving task distributions and is more robust to OOD data than related strong baselines.

## 1 Introduction

Meta-learning [51] focuses on learning lots of related tasks to acquire common knowledge adaptable to new unseen tasks, and has achieved great success in many computer vision problems [26,61,27]. Recently, it has been shown that additional unlabeled data is beneficial for improving the meta-learning model performance [47,35,36] especially when labeled data is scarce. One common assumption of

---

\* Co-correspondence author.

<sup>‡</sup> we use ORDER to denote the task distributions sequentially arrive with some ORDER

these works [47,35,36] is that the task distribution (i.e., a family of tasks) is stationary during meta training. However, real-world scenarios are more complex and often require learning across different environments.

One instance is in personalized self-driving systems [7]. Each task refers to learning a personal object recognition subsystem for an individual user. The labeled data (annotated objects) for each user is scarce and expensive, but unlabeled data (e.g., images captured by their in-car cameras) is abundant. Learning useful information from these unlabeled data would be very helpful for learning the system. Moreover, the task distributions are not stationary, as users’ driving behaviors are different across regions. Assume the learning system is first deployed in Rochester and later extended to New York. The original dataset for Rochester may be no longer available when training on New York users due to storage or business privacy constraints. Similar scenarios occur when Amazon extends the market for the personalized conversation AI product, Alexa, to different countries [19,29].

As can be seen from the above examples, the first significant challenge is that task distributions may evolve during the learning process; the learning system could easily forget the previously learned knowledge. Second, despite that the additional unlabeled data are potentially beneficial to improve model performance [47,35,36], blindly incorporating them may bring negative effects. This is due to the fact that unlabeled data could contain both useful (In-Distribution (ID)) and harmful (Out-Of-Distribution (OOD)) data. How to leverage such mixed (ID and OOD) unlabeled data in meta-learning is rarely explored in previous literature.

In general, these challenges can be summarized into a new realistic problem setting that the model learns on a sequence of distinct task distributions without prior information on task distribution shift. We term this problem setup as *semi-supervised meta-learning on evolving task distributions* (**SETS**). To simulate the task distribution shift in SETS, we construct a dataset by composing 6 datasets, including *CIFARFS* [12], *AIRCRAFT* [38], *CUB* [67], *Miniimagenet* [56], *Butterfly* [16] and *Plantae* [28]. A model will be sequentially trained on the datasets in a meta-learning fashion, simulating the real scenarios of evolving environment. However, at the end of training on the last domain, the meta learned model could easily *forget* the learned knowledge to solve *unseen* tasks sampled from the task distributions trained previously.

In this paper, we aim to tackle the SETS problem, which requires the meta-learning model to adapt to a new task distribution and retain the knowledge learned on previous task distributions. To better use unlabeled data of each task, we propose a regularization term driven by mutual information, which measures the mutual dependency between two variables. We first propose an OOD detection method to separate the unlabeled data for each task into ID and OOD data, respectively. Then, we minimize the mutual information between class prototypes and OOD data to reduce the effect of OOD on the class prototypes. At the same time, we maximize the mutual information between class prototypes and ID data to maximize the information transfer. We derive tractable sample-based variational bounds specialized for SETS for mutual information maximization

and minimization for practical computation since the closed-form of probability distribution in mutual information is not available.

To mitigate catastrophic forgetting during the learning process, we propose a novel and simple idea to reduce forgetting in feature space. Specifically, in addition to storing the raw data as used in standard continual learning, we also store the data features of memory tasks. To achieve this goal, we use optimal transport (OT) to compute the distance between the feature distribution of memory tasks at the current iteration and the feature distribution of memory tasks stored in the memory. This ensures minimal feature distribution shift during the learning process. The proposed OT framework can naturally incorporate the unlabeled data features of previous tasks to further improve the memorization of previous knowledge. Besides, as the feature size is much smaller than the raw data, this makes it scalable to large-scale problems with a very long task sequence.

Our contributions are summarized as three-fold:

- To our best knowledge, this is the first work considering the practical and challenging *semi-supervised meta-learning on evolving task distributions* (SETS) setting.
- We propose a novel algorithm for SETS setting to exploit the unlabeled data and a mutual information inspired regularization to tackle the OOD problem and an OT regularization to mitigate forgetting.
- Extensive experiments on the constructed dataset with at least 72K tasks demonstrate the effectiveness of our method compared to strong baselines.

## 2 Related Work

**Meta Learning** [51] focuses on rapidly adapting to new unseen tasks by learning prior knowledge through training on a large number of similar tasks. Various approaches have been proposed [56,54,21,22,24,50,5,40,4,39,34,66,69,65]. All of these methods work on the simplified setting where the model is meta trained on a stationary task distribution in a fully supervised fashion, which is completely different from our proposed SETS. On the other hand, online meta-learning [23] aims to achieve better performance on future tasks and stores all the previous tasks to avoid forgetting with a small number of tasks. Furthermore, [62,63] has proposed a meta learning method to solve large-scale task stream with sequential domain shift. However, our method is different from the previous work in several aspects. Specifically, (i) we additionally focus on leveraging unlabeled data and improving model robustness; (ii) we only store a small number of tasks to enable our method to work in our large-scale setting; (iii) we consider that the task distributions could be evolving and OOD-perturbed during meta training under the semi-supervised learning scenario. In summary, SETS is more general and realistic than online meta-learning and other settings.

**Continual Learning (CL)** aims to maintain previous knowledge when learning on a sequence of tasks with distribution shift. Many works mitigate catastrophic forgetting during the learning process [37,14,48,68,31,42,20,2,10,3,64].

Semi-supervised continual learning (SSCL) [33,60,53] is a recent extension of CL to the semi-supervised learning setting. SETS is completely different from CL and SSCL in several aspects: (i) our SETS focuses on a large-scale task sequence (more than 72K), while CL and SSCL are not scalable to this large-scale setting; (ii) our goal is to enable the model to generalize to the *unseen* tasks from all the previous task distributions, while SSCL works on the generalization to the testing data of previous tasks.

**Semi-supervised Few-shot Learning (SSFSL)** [47,35,36] aims to generalize to unseen semi-supervised tasks when meta training on semi-supervised few-shot learning tasks in an *offline* learning manner. It assumes that the few-shot learning task distributions are stationary during meta training. Unlike these works, we work on a more challenging and realistic setting of evolving task distributions in an *online* learning manner. Meanwhile, SETS also aims at mitigating catastrophic forgetting for the previous learned tasks during online learning process as well, which makes it more challenging and practical to real-world scenarios.

We summarize and compare different benchmarks in Table 1, including continual learning (CL), semi-supervised learning (SSL), continual few-shot learning (CFSL), semi-supervised continual learning (SSCL), semi-supervised few-shot learning (SSFSL), and SETS (Ours) in terms of whether they consider catastrophic forgetting (CF), use unlabeled data, few-shot learning, and generalize to unseen tasks. In summary, our setting is more comprehensive and practical and covers more learning aspects in real applications.

Table 1: Comparisons among different benchmarks.

Settings	CF	Unlabeled	Few-Shot	Unseen	Tasks
CL [31]	✓	✗	✗	✗	✗
SSL [11]	✗	✓	✗	✗	✗
CFSL [6]	✓	✗	✓	✗	✗
SSCL [60]	✓	✓	✗	✗	✗
SSFSL [47]	✗	✓	✓	✓	✓
SETS (Ours)	✓	✓	✓	✓	✓

### 3 Problem Setting

**General setup.** SETS (Figure 1 for illustration), online meta-learns on a sequence of task distributions,  $\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_L$ . For time  $t = 1, \dots, \tau_1$ , we randomly sample mini-batch semi-supervised tasks  $\mathcal{T}_t$  at each time  $t$  from task distribution  $P(\mathcal{D}_1)$ ; for time  $t = \tau_1 + 1, \dots, \tau_2$ , we randomly sample mini-batch semi-supervised tasks  $\mathcal{T}_t$  at each time  $t$  from task distribution  $P(\mathcal{D}_2)$ ; for time  $t = \tau_i + 1, \dots, \tau_{i+1}$ , we randomly sample mini-batch semi-supervised tasks  $\mathcal{T}_t$  at each time  $t$  from task distribution  $P(\mathcal{D}_i)$ , where  $P(\mathcal{D}_i)$  is the task distribution (a collection of a large number of tasks) in  $\mathcal{D}_i$ . Thus, in SETS, each task from each domain sequentially arrives. Time  $t = \tau_1, \tau_2, \dots, \tau_i, \dots$  is the time when task distribution shift happens, but we do not assume any prior knowledge about them. This is a




Fig. 1: Illustration of SETS on a large-scale non-stationary semi-supervised task distributions  $\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_L$  (top) with at least 72K tasks. Each task (bottom) contains labeled support data (S), unlabeled data (U) (in-distribution and out-of-distribution data) and query data (Q).

more practical and general setup. Each time interval  $|\tau_i - \tau_{i-1}|$  is large enough to learn a large number of tasks from each task distribution adequately. Each task  $\mathcal{T}$  is divided into support, unlabeled and query data  $\{\mathcal{S}, \mathcal{U}, \mathcal{Q}\}$ . The support set  $\mathcal{S}$  consists of a collection of labeled examples,  $\{(\mathbf{x}^k, y^k)\}$ , where  $\mathbf{x}^k$  is the data and  $y^k$  is the label. The unlabeled set  $\mathcal{U}$  contains only inputs:  $\mathcal{U} = \{\tilde{\mathbf{x}}^1, \tilde{\mathbf{x}}^2, \dots\}$ . The goal is to predict labels for the examples in the task’s query set  $\mathcal{Q}$ . Our proposed learning system maintains a memory buffer  $\mathcal{M}$  to store a small number of training tasks from previous task distributions. We use reservoir sampling [57] (RS) to maintain tasks in the memory. RS assigns equal probability for each incoming task without needing to know the total number of training tasks in advance. The total number of meta training tasks is much larger than the capacity of memory buffer, making it infeasible to store all the tasks. More details about memory buffer update and maintenance are provided in Appendix B. We evaluate model performance on the unseen tasks sampled from both current and all previous task distributions.

**Evolving semi-supervised few-shot episode distributions construction.** In most existing works, the meta-learning model is trained in a stationary setting. By contrast, in SETS, the model is meta trained on an evolving sequence of datasets. A large number of semi-supervised tasks sampled from each dataset form task distribution  $\mathcal{D}_i$ . We thus have a *sequence of semi-supervised task distributions*  $\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_L$ . In episodic meta-learning, a task is usually denoted as an episode. When training on task distribution  $\mathcal{D}_i$ , to sample an  $N$ -way  $K$ -shot training episode, we first uniformly sample  $N$  classes from the set of training classes  $\mathcal{C}_i^{train}$ . For the labeled support set  $\mathcal{S}$  (training data), we then sample  $K$  images from the labeled split of each of these classes. For the unlabeled set  $\mathcal{U}$ , we sample  $Z$  images from the unlabeled split of each of these classes as ID data. When including OOD in  $\mathcal{U}$ , we additionally sample  $R$  images from external datasets as OOD. In practical scenarios, the OOD data could dominate the unlabeled portion  $\mathcal{U}$ . To be more realistic, we construct the unlabeled set with different proportions of OOD to simulate real-world scenarios. The query data (testing data)  $\mathcal{Q}$  of the episode is comprised of a fixed number of images from the labeled split of each of the  $N$  chosen classes. The support set  $\mathcal{S}$  and query set  $\mathcal{Q}$  contain different inputs but with the same class set.

At the end of meta training on task distribution  $\mathcal{D}_L$ , the meta trained model is expected to generalize to the *unseen* tasks of each dataset in the sequence. The evaluation on *unseen* classes of previous datasets is to measure the forgetting. We adopt the  $N$ -way  $K$ -shot classification on *unseen* classes of each dataset. The semi-supervised testing episodes are constructed similarly as above. The accuracy is averaged on the query set  $\mathcal{Q}$  of many meta-testing episodes from all the trained datasets.

**Constructed dataset.** To simulate realistic evolving semi-supervised task distributions, we construct a new large-scale dataset and collect 6 datasets, including **CIFARFS** [12], **AIRCRAFT** [38], **CUB** [67], **Miniimagenet** [56], **Butterfly** [16] and **Plantae** [28]. The OOD data for each task are sampled from another three datasets, including **Omniglot** [32], **Vggflower** [43] and **Fish** [70]. For each dataset, we randomly sample 12K tasks, we thus have 72K tasks in total. This task sequence is much larger than existing continual learning models considered. We perform standard 5-way-1-shot and 5-way-5-shot learning on this new dataset. All the datasets are publicly available with more details provided in Appendix A due to the space limitation.

## 4 Methodology

In the following, we focus on the  $N$ -way  $K$ -shot setting of SETS as described above. We first describe the standard SSFSL task in Section 4.1, then present our method for handling unlabeled data, especially OOD data in Section 4.2. Next, we present our method for mitigating catastrophic forgetting in Section 4.3. At last, we summarize our proposed methodology in Section 4.4.

### 4.1 Standard Semi-Supervised Few-Shot Learning

Prototypical network [54] has been extended to semi-supervised few-shot learning by [47]. We denote the feature embedding function as  $h_{\theta}(\mathbf{x})$  with parameter  $\theta$ . The class-wise prototype for support set  $\mathcal{S}$  is calculated as  $\mathbf{p}_c = \sum_{\mathbf{x} \in \mathcal{S}_c} h_{\theta}(\mathbf{x}) / \sum_{\mathbf{x} \in \mathcal{S}_c} 1$ , where  $\mathcal{S}_c$  is the set of data in  $\mathcal{S}$  belonging to category  $c$ . Each class prototype is refined by unlabeled data as:

$$\mathbf{p}'_c = \frac{\sum_{\mathbf{x} \in \mathcal{S}_c} h_{\theta}(\mathbf{x}) + \sum_{\tilde{\mathbf{x}} \in \mathcal{U}} h_{\theta}(\tilde{\mathbf{x}}) \mu_c(\tilde{\mathbf{x}})}{\sum_{\mathbf{x} \in \mathcal{S}_c} 1 + \sum_{\tilde{\mathbf{x}} \in \mathcal{U}} \mu_c(\tilde{\mathbf{x}})}$$

by assigning each unlabeled data  $\tilde{\mathbf{x}} \in \mathcal{U}$  with a soft class probability  $\mu_c(\tilde{\mathbf{x}}) = \frac{\exp(-\|h_{\theta}(\tilde{\mathbf{x}}) - \mathbf{p}_c\|_2^2)}{\sum_{c'} \exp(-\|h_{\theta}(\tilde{\mathbf{x}}) - \mathbf{p}'_{c'}\|_2^2)}$ . To work with OOD data, they use an additional class with zero mean in addition to the normal  $c$  classes in the support set. The learning objective is to simply maximize the average log-probability of the correct class assignments over query examples  $\mathcal{Q}$  and rigorously defined as following:

$$\mathcal{L}_{meta} = P(\mathcal{Q}|\theta, \mathcal{S}; \mathcal{U}) = P(\mathcal{Q}|\theta, \{\mathbf{p}'_c\}) = \sum_{(\mathbf{x}_i, y_i) \in \mathcal{Q}} P(y_i | \mathbf{x}_i, \theta, \{\mathbf{p}'_c\}). \quad (1)$$

## 4.2 Mutual-Information For Unlabeled Data Handling

To be more realistic, the unlabeled data contains OOD data for each task, whose data categories are not present in the support set. In practice, the unlabeled OOD data could potentially negatively impact the performance. It has been shown that existing works [47,35,36] are particularly sensitive to OOD data. How to properly use unlabeled data is challenging when OOD data dominates the unlabeled data, which is common in practical scenarios. Instead of sidestepping this problem by simply discarding the detected OOD data, we propose to exploit discriminative information implied by unlabeled OOD samples. To explicitly address this problem, we propose to learn the representations that are robust to OOD data from an information-theoretical perspective. We first give the detailed definitions for in-distribution (ID) and out-of-distribution (OOD) below.

**Definition 1.** Suppose a task  $\mathcal{T}$  with data  $\{\mathcal{S}, \mathcal{U}, \mathcal{Q}\}$ , contains labeled categories  $\mathcal{C}$  in  $\mathcal{S}$ . For unlabeled data  $\mathcal{U}$ , in-distribution unlabeled data is defined as the set of data  $\mathbf{x}$  whose categories  $c \in \mathcal{C}$ , i.e.,  $\mathcal{U}_{id} = \{\mathbf{x} \mid \mathbf{x} \in \mathcal{U}, c \in \mathcal{C}\}$ ; similarly, the out-of-distribution data is defined as the set of data  $\mathbf{x}$  whose categories  $c \notin \mathcal{C}$ , i.e.,  $\mathcal{U}_{ood} = \{\mathbf{x} \mid \mathbf{x} \in \mathcal{U}, c \notin \mathcal{C}\}$ .

**OOD sample detection.** We propose to automatically divide the unlabeled data for each task into ID and OOD. The existing confidence-based approach is highly unreliable for detecting OOD data in the few-shot learning setting since the network can produce incorrect high-confidence predictions for OOD samples [41]. We instead propose a metric-based OOD sample detection method in feature embedding space. First, we calculate the class-wise prototypes  $\mathbf{p}_c$  for each class  $c$ . We define the distance of data  $\mathbf{x}$  to the prototypes set as:  $d(h_{\theta}(\mathbf{x}), \{\mathbf{p}_c\}_{c=1}^N) \stackrel{def}{=} \min_c d(h_{\theta}(\mathbf{x}), \mathbf{p}_c)$ , i.e., the smallest distance to all the prototypes in the embedding space; where we use Euclidean distance as the

---

### Algorithm 1 OOD sample detection.

---

**Require:** A task data  $\{\mathcal{S}, \mathcal{U}, \mathcal{Q}\}$ ; unlabeled query  $\mathcal{Q}$  serves as ID-calibration set;

- 1: calculate prototypes  $\mathbf{p}_c$  for each class  $c$  with support data  $\mathcal{S}$ .
- 2: **for**  $\mathbf{x}_i \in \mathcal{Q}$  **do**
- 3:    $\mathbf{w}_i = h_{\theta}(\mathbf{x}_i) / \|h_{\theta}(\mathbf{x}_i)\|_2$  (normalization)
- 4:    $d_i = d(\mathbf{w}_i, \{\mathbf{p}_c\}_{c=1}^N) \stackrel{def}{=} \min_c d(\mathbf{w}_i, \mathbf{p}_c)$
- 5: **end for**
- 6: thresh =  $\mu(\{d_i\}) + \sigma(\{d_i\})$  [ $\mu(\{d_i\})$  is the mean of  $\{d_i\}$  and  $\sigma(\{d_i\})$  is the standard deviation of  $\{d_i\}$ ]
- 7: **for**  $\mathbf{x}_i \in \mathcal{U}$  **do**
- 8:    $\mathbf{w}_i = h_{\theta}(\mathbf{x}_i) / \|h_{\theta}(\mathbf{x}_i)\|_2$  (normalization)
- 9:    $d_i = d(\mathbf{w}_i, \{\mathbf{p}_c\}_{c=1}^N) \stackrel{def}{=} \min_c d(\mathbf{w}_i, \mathbf{p}_c)$
- 10:   **if**  $d_i > \text{thresh}$  **then**
- 11:      $\mathcal{U}_{ood} = \mathcal{U}_{ood} \cup \mathbf{x}_i$
- 12:   **else**
- 13:      $\mathcal{U}_{id} = \mathcal{U}_{id} \cup \mathbf{x}_i$
- 14:   **end if**
- 15: **end for**
- 16: return ID data  $\mathcal{U}_{id}$  and OOD data  $\mathcal{U}_{ood}$

---

distance metric  $d$  for simplicity. Intuitively, the data closer to the class prototypes are more likely to be ID data, the data farther to the class prototypes are more likely to be OOD. We adopt a method similar to the state-of-the-art few-shot OOD detection algorithm, named cluster-conditioned detection [52]. We use the unlabeled query data  $\mathcal{Q}$  for each task as an ID calibration set to determine the threshold between OOD and ID data. The OOD sample detection is presented in Algorithm 1. The threshold is based on the Gaussian assumption of  $d(h_{\theta}(\mathbf{x}), \{\mathbf{p}_c\}_{c=1}^N)$  (3-Sigma rule).

We adopt mutual information (MI) to separately handle ID data and OOD data, where MI describes the mutual dependence between the two variables and is defined as follows.

**Definition 2.** *The mutual information between variables  $\mathbf{a}$  and  $\mathbf{b}$  is defined as*

$$\mathcal{I}(\mathbf{a}, \mathbf{b}) = \mathbb{E}_{p(\mathbf{a}, \mathbf{b})} \left[ \log \frac{p(\mathbf{a}, \mathbf{b})}{p(\mathbf{a})p(\mathbf{b})} \right], \quad (2)$$

where  $p(\mathbf{a})$  and  $p(\mathbf{b})$  are marginal distribution of  $\mathbf{a}$  and  $\mathbf{b}$ ,  $p(\mathbf{a}, \mathbf{b})$  is the joint distribution.

In information theory,  $\mathcal{I}(\mathbf{a}, \mathbf{b})$  is a measure of the mutual dependence between the two variables  $\mathbf{a}$  and  $\mathbf{b}$ . More specifically, it quantifies the "amount of information" about the variable  $\mathbf{a}$  (or  $\mathbf{b}$ ) captured via observing the other variable  $\mathbf{b}$  (or  $\mathbf{a}$ ). The higher dependency between two variables, the higher MI is. Intuitively, the higher dependency between embedded OOD data and class prototypes, the higher the negative impact on the model performance. On the contrary, more dependency between embedded ID data and class prototypes could improve the performance. To enhance the effect of ID data on the class prototypes, we maximize the MI between class prototypes and ID data. To reduce the negative effect of OOD data on the class prototypes, we further consider minimizing the MI between class prototypes and OOD data:

$$\mathcal{L}_{total} = \mathcal{L}_{meta} + \lambda (\mathcal{I}(\mathbf{e}_{ood}, \mathbf{p}_{c'}) - \mathcal{I}(\mathbf{e}_{id}, \mathbf{p}_{c'})), \quad (3)$$

where  $\mathcal{L}_{meta}$  is defined in Eq. (1);  $\mathbf{e}_{ood} = h_{\theta}(\mathbf{x})$ ,  $\mathbf{x} \in \mathcal{U}_{ood}$  and  $\mathbf{e}_{id} = h_{\theta}(\mathbf{x})$ ,  $\mathbf{x} \in \mathcal{U}_{id}$ .  $c' = \operatorname{argmin}_c \|\mathbf{e} - \mathbf{p}_c\|$ ;  $\mathcal{I}(\mathbf{e}_{ood}, \mathbf{p}_{c'})$  is the mutual information between embedded OOD data and class prototypes;  $\mathcal{I}(\mathbf{e}_{id}, \mathbf{p}_{c'})$  is the mutual information between embedded ID data and class prototypes;  $\lambda$  is a hyperparameter to control the relative importance of different terms. Since this objective is to minimize  $-\mathcal{I}(\mathbf{e}_{id}, \mathbf{p}_{c'})$ , it is equivalent to maximize  $\mathcal{I}(\mathbf{e}_{id}, \mathbf{p}_{c'})$ . However, the computation of MI values is intractable [44,18], since the form of joint probability distribution  $p(\mathbf{a}, \mathbf{b})$  in Eq. (2) is not available in our case and only samples from the joint distribution are accessible. To solve this problem, we derive the variational lower bound of MI for MI maximization and variational upper bound of MI for MI minimization specialized for SETS, respectively. Due to the limited space, all proofs for Lemmas and Theorems are provided in Appendix D.

**Theorem 1.** *For a task  $\mathcal{T} = \{\mathcal{S}, \mathcal{U}, \mathcal{Q}\}$ , suppose the unlabeled OOD data embedding  $\mathbf{e}_{ood} = h_{\theta}(\mathbf{x})$ ,  $\mathbf{x} \in \mathcal{U}_{ood}$ . The nearest prototype corresponds to  $\mathbf{e}_{ood}$  is  $\mathbf{p}_{c'}$ ,*



where  $c' = \operatorname{argmin}_c \|\mathbf{e}_{ood} - \mathbf{p}_c\|$ . Given a collection of samples  $\{(\mathbf{e}_{ood}^i, \mathbf{p}_{c'}^i)\}_{i=1}^{i=L} \sim P(\mathbf{e}_{ood}, \mathbf{p}_{c'})$ , the variational upper bound of mutual information  $\mathcal{I}(\mathbf{e}_{ood}, \mathbf{p}_{c'})$  is:

$$\mathcal{I}(\mathbf{e}_{ood}, \mathbf{p}_{c'}) \leq \sum_{i=1}^{i=L} \log P(\mathbf{p}_{c'}^i | \mathbf{e}_{ood}^i) - \sum_{i=1}^{i=L} \sum_{j=1}^{j=L} \log P(\mathbf{p}_{c'}^i | \mathbf{e}_{ood}^j) = \mathcal{I}_{ood}. \quad (4)$$

The bound is tight (equality holds) when  $\mathbf{p}_{c'}$  and  $\mathbf{e}_{ood}$  are independent.

This upper bound requires the unknown conditional probability distribution  $P(\mathbf{p}_{c'} | \mathbf{e}_{ood})$  and thus it is intractable. In this work, we approximate it with variational distribution  $Z_\phi(\mathbf{p}_{c'} | \mathbf{e}_{ood})$  parameterized by a neural network.

**Lemma 1.** For a task  $\mathcal{T} = \{\mathcal{S}, \mathcal{U}, \mathcal{Q}\}$ , suppose the unlabeled ID data embedding  $\mathbf{e}_{id} = h_\theta(\mathbf{x})$ ,  $\mathbf{x} \in \mathcal{U}_{id}$ . The nearest prototype corresponds to  $\mathbf{e}_{id}$  is  $\mathbf{p}_{c'}$ , where  $c' = \operatorname{argmin}_c \|\mathbf{e}_{id} - \mathbf{p}_c\|$ . Given a collection of samples  $\{(\mathbf{e}_{id}^i, \mathbf{p}_{c'}^i)\}_{i=1}^{i=L} \sim P(\mathbf{e}_{id}, \mathbf{p}_{c'})$ , the variational lower bound of  $\mathcal{I}(\mathbf{e}_{id}, \mathbf{p}_{c'})$  is:

$$\mathcal{I}(\mathbf{e}_{id}, \mathbf{p}_{c'}) \geq \sum_{i=1}^{i=L} f(\mathbf{e}_{id}^i, \mathbf{p}_{c'}^i) - \sum_{j=1}^{j=L} \log \sum_{i=1}^{i=L} e^{f(\mathbf{e}_{id}^i, \mathbf{p}_{c'}^j)}. \quad (5)$$

This bound is tight if  $f(\mathbf{e}_{id}, \mathbf{p}_{c'}) = \log P(\mathbf{p}_{c'} | \mathbf{e}_{id}) + c(\mathbf{p}_{c'})$ .

Based on Lemma 1, the variational lower bound is derived as in Theorem 2.

**Theorem 2.** For a task  $\mathcal{T} = \{\mathcal{S}, \mathcal{U}, \mathcal{Q}\}$ , suppose the unlabeled ID data embedding  $\mathbf{e}_{id} = h_\theta(\mathbf{x})$ ,  $\mathbf{x} \in \mathcal{U}_{id}$ . The nearest prototype corresponds to  $\mathbf{e}_{id}$  is  $\mathbf{p}_{c'}$ , where  $c' = \operatorname{argmin}_c \|\mathbf{e}_{id} - \mathbf{p}_c\|$ . Given a collection of samples  $\{(\mathbf{e}_{id}^i, \mathbf{p}_{c'}^i)\}_{i=1}^{i=L} \sim P(\mathbf{e}_{id}, \mathbf{p}_{c'})$ , the variational lower bound of  $\mathcal{I}(\mathbf{e}_{id}, \mathbf{p}_{c'})$  is:

$$\mathcal{I}(\mathbf{e}_{id}, \mathbf{p}_{c'}) \geq \mathcal{I}_{id} := \sum_{i=1}^{i=L} f(\mathbf{e}_{id}^i, \mathbf{p}_{c'}^i) - \sum_{j=1}^{j=L} \left[ \frac{\sum_{i=1}^{i=L} e^{f(\mathbf{e}_{id}^i, \mathbf{p}_{c'}^j)}}{a(\mathbf{p}_{c'}^j)} + \log(a(\mathbf{p}_{c'}^j)) - 1 \right]. \quad (6)$$

The bound is tight (equality holds) when  $f(\mathbf{e}_{id}, \mathbf{p}_{c'}) = \log P(\mathbf{p}_{c'} | \mathbf{e}_{id}) + c(\mathbf{p}_{c'})$  and  $a(\mathbf{p}_{c'}) = \mathbb{E}_{p(\mathbf{e}_{id})} e^{f(\mathbf{e}_{id}, \mathbf{p}_{c'})}$ .  $a(\mathbf{p}_{c'})$  is any function that  $a(\mathbf{p}_{c'}) > 0$ .

Hence, by combining Theorems 1 and 2, we obtain an surrogate for the loss function Eq. (3) as:

$$\mathcal{L}_{total} = \mathcal{L}_{meta} + \lambda(\mathcal{I}(\mathbf{e}_{ood}, \mathbf{p}_{c'}) - \mathcal{I}(\mathbf{e}_{id}, \mathbf{p}_{c'})) \leq \mathcal{L}_{meta} + \lambda(\mathcal{I}_{ood} - \mathcal{I}_{id}). \quad (7)$$

### 4.3 Mitigate CF by Optimal Transport

A straightforward idea of mitigating catastrophic forgetting on previous task distributions is to replay the memory tasks when training on current task  $\mathcal{T}_t$ , i.e., we randomly sample a small batch tasks  $\mathcal{B}$  from the memory buffer  $\mathcal{M}$ . Thus the loss function becomes  $\mathcal{H}(\theta) = \mathcal{L}(\mathcal{T}_t) + \sum_{\mathcal{T}_i \in \mathcal{B}} \mathcal{L}(\mathcal{T}_i)$ . However, simple experience

replay only forces the model to remember a limited number of labeled data, i.e., a much larger number of previous unlabeled data is totally neglected, which causes the model to sub-optimally remember the previous knowledge. To further avoid forgetting, we directly regularize the network outputs on both labeled and unlabeled data since what ultimately matters is the network output, in addition to remembering output labels on memory tasks. To achieve this goal, we propose a novel constrained loss via optimal transport to reduce the *output (feature) distribution shift* on both labeled and unlabeled data.

**Optimal Transport.** A discrete distribution  $\boldsymbol{\mu} = \sum_{i=1}^n u_i \delta_{\mathbf{e}_i}(\cdot)$ , where  $u_i \geq 0$  and  $\sum_i u_i = 1$ ,  $\delta_{\mathbf{e}}(\cdot)$  denotes a spike distribution located at  $\mathbf{e}$ . Given two discrete distributions,  $\boldsymbol{\mu} = \sum_{i=1}^n u_i \delta_{\mathbf{e}_i}$  and  $\boldsymbol{\nu} = \sum_{j=1}^m v_j \delta_{\mathbf{g}_j}$ , respectively, the OT distance [8] between  $\boldsymbol{\mu}$  and  $\boldsymbol{\nu}$  is defined as the optimal value of the following problem:

$$\mathcal{L}_{\text{OT}} = \min_{\mathbf{w} \in \Pi(\boldsymbol{\mu}, \boldsymbol{\nu})} \sum_{i=1}^n \sum_{j=1}^m \mathbf{W}_{ij} \cdot d(\mathbf{e}_i, \mathbf{g}_j), \quad (8)$$

where  $\mathbf{W}_{ij}$  is defined as the joint probability mass function on  $(\mathbf{e}_i, \mathbf{g}_j)$  and  $d(\mathbf{e}, \mathbf{g})$  is the cost of moving  $\mathbf{e}$  to  $\mathbf{g}$  (matching  $\mathbf{e}$  and  $\mathbf{g}$ ).  $\Pi(\boldsymbol{\mu}, \boldsymbol{\nu})$  is the set of joint probability distributions with two marginal distributions equal to  $\boldsymbol{\mu}$  and  $\boldsymbol{\nu}$ , respectively. The exact calculation of the OT loss is generally difficult [25]. Therefore, for simplicity, we use CVXPY [1] for efficient approximation.

When storing a semi-supervised episode into the memory, we store the feature embedding for each data point in addition to the raw data. Memory features are the collection of  $\{\mathbf{h}_{\boldsymbol{\theta}}(\mathbf{x}), \mathbf{x} \in \bigcup_{\{\mathcal{S}, \mathcal{U}_{id}\} \in \mathcal{M}} \{\mathcal{S}, \mathcal{U}_{id}\}\}$ , the concatenation of labeled and unlabeled ID data embedding. Suppose the previously stored data features are  $\{\mathbf{g}_i\}_{i=1}^G$  with distribution  $\boldsymbol{\nu}$ , where  $G$  is the number of data examples stored in memory buffer. The data features generated with model parameters  $\boldsymbol{\theta}_t$  of current iteration are  $\{\mathbf{e}_i\}_{i=1}^G$  with distribution  $\boldsymbol{\mu}$ . Then, the loss function becomes minimizing on the optimal transport between  $\{\mathbf{e}_i\}_{i=1}^G$  and  $\{\mathbf{g}_i\}_{i=1}^G$ .

---

#### Algorithm 2 ORDER Algorithm

---

**Require:** evolving episodes  $\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_N$ , with  $\{\tau_1, \dots, \tau_i, \dots, \tau_{L-1}\}$  the time steps when task distribution shift; memory buffer  $\mathcal{M} = \{\}$ ; model parameters  $\boldsymbol{\theta}$ ; memory task features  $\mathcal{E} = \{\}$ .

- 1: **for**  $t = 1$  to  $N$  **do**
  - 2:   sample tasks  $\mathcal{B}$  from  $\mathcal{M}$  and  $\mathcal{B} = \mathcal{B} \cup \mathcal{T}_t$ .
  - 3:   **for**  $\mathcal{T} \in \mathcal{B}$  **do**
  - 4:      $\mathcal{T} = \{\mathcal{S}, \mathcal{U}, \mathcal{Q}\}$ , OOD detection to divide the unlabeled data  $\mathcal{U}$  into  $\mathcal{U}_{ood}$  and  $\mathcal{U}_{id}$
  - 5:     update parameters  $\boldsymbol{\theta}$  by minimizing  $\mathcal{H}(\boldsymbol{\theta})$
  - 6:   **end for**
  - 7:   **if** reservoir sampling: **then**
  - 8:      $\mathcal{M} = \mathcal{M} \cup \mathcal{T}_t$  // store raw task data
  - 9:      $\mathcal{E} = \mathcal{E} \cup \mathcal{E}(\mathcal{T}_t)$  // store task data features
  - 10:   **end if**
  - 11: **end for**
-

#### 4.4 Overall Learning Objective

By unifying the above-derived novel mutual information driven unlabeled data handling technique and optimal transport driven forgetting mitigation, we propose an OOD robust and knowledge preserved semi-supervised meta-learning algorithm (ORDER) to tackle the challenges in the SETS setting. Combining the OT loss from Eq. (8) with Eq. (7) and assigning  $\lambda$  and  $\beta$  as two hyperparameters for controlling the relative importance of the corresponding terms, we obtain the following final loss function:

$$\mathcal{H}(\theta) = \mathcal{L}_{meta}(\mathcal{T}_t) + \lambda(\mathcal{I}_{ood} - \mathcal{I}_{id}) + \sum_{\mathcal{T}_i \in \mathcal{B}} \mathcal{L}(\mathcal{T}_i) + \beta \mathcal{L}_{OT}. \quad (9)$$

The benefits of the second term are two-fold: (i) promote the effect of ID data on class prototypes by maximizing the MI between class prototypes and ID data; (ii) reduce the negativity of OOD data on class prototypes by minimizing the MI between class prototypes and OOD data. The last term is to force the model to remember in feature space. The detailed procedure of ORDER is shown in Algorithm 2.

## 5 Experiments

To show the benefit of using unlabeled data, we first compare semi-supervised meta-learning using unlabeled data with meta-learning without unlabeled data in Section 5.1. To evaluate the effectiveness of the proposed method, ORDER, for learning useful knowledge in SETS, we compare to SOTA meta-learning models in Section 5.2. Next, we compare ORDER to SOTA continual learning baselines to show the effectiveness of OR-

Table 2: 5-way, 1-shot and 5-shot accuracy compared to meta-learning baselines.

Algorithm	1-Shot	5-Shot
ProtoNet	31.96 ± 0.93	42.62 ± 0.73
ANIL	30.58 ± 0.81	42.24 ± 0.86
MSKM	33.79 ± 1.05	45.41 ± 0.58
LST	34.81 ± 0.81	46.09 ± 0.69
TPN	34.52 ± 0.83	46.23 ± 0.60
ORDER (Ours with only MI)	36.93 ± 0.90	48.85 ± 0.53
ORDER (Ours)	<b>41.25 ± 0.64</b>	<b>53.96 ± 0.71</b>

DER for mitigating forgetting in SETS in Section 5.3. For the evaluation, we calculate the average testing accuracy on previously unseen 600 testing tasks sampled from the unseen categories of each dataset in the dataset sequence to evaluate the effectiveness of the tested methods. Moreover, we conduct extensive ablation studies in Section 5.4. Due to the space limitation, we put detailed implementation details in Appendix B.

### 5.1 Benefit of Using Unlabeled Data in SETS

To verify the benefit of using unlabeled data in SETS, we compare with various meta-learning methods including supervised and semi-supervised ones. All of these

methods are adapted to SETS setting directly. Supervised meta-learning methods *without unlabeled data*, including (1) gradient-based meta-learning **ANIL** [45], which is a simplified model of MAML [22]; (2) metric-based meta-learning **Prototypical Network (ProtoNet)** [54]. In addition, we compare to representative SSFSL models *using unlabeled data*: (3) Masked Soft k-Means (**MSKM**) [47]; (4) Transductive Propagation Network (**TPN**) [36]; (5) Learning to Self-Train (**LST**) [35]. LST and MSKM can be viewed as extension of MAML(ANIL) and ProtoNet to SETS by using additional unlabeled data respectively. Table 2 shows that with additional unlabeled data, semi-supervised meta-learning methods substantially outperform corresponding meta-learning methods without using unlabeled data, demonstrating that unlabeled data could further improve performance in SETS. We also include a reduced version of our method, which uses Eq. (7) as the objective function, and keep other components the same as baselines without considering mitigating CF. We can see that with only the guidance of MI, our method can outperform baselines by more than 2.6% for 5-shot learning and 2.1% for 1-shot learning, which verifies the effectiveness of the proposed MI for handling unlabeled data.

## 5.2 Comparison to Meta-Learning

To compare to SOTA meta-learning methods, similar to the baselines in Section 5.1, Table 2 also shows the advantage of ORDER. We observe that our method significantly outperforms baselines for 5-shot learning by 7.7% and for 1-shot learning by 6.4%. This improvement shows the effectiveness of our method in this challenging setting. LST [35] performs relatively worse in SETS, probably due to the challenging OOD data and noise of pseudo-labels.

Table 3: 5-way, 1-shot and 5-shot classification accuracy compared to continual learning baselines.

Algorithm	1-Shot	5-Shot
Semi-Seq	33.79 ± 1.05	45.41 ± 0.58
Semi-ER	37.62 ± 0.97	50.35 ± 0.76
Semi-AGEM	36.97 ± 1.16	50.29 ± 0.69
Semi-MER	37.90 ± 0.83	50.46 ± 0.74
Semi-GPM	36.53 ± 0.81	49.78 ± 0.65
Semi-DEGCL	36.78 ± 0.89	50.07 ± 0.61
ORDER (Ours OT only)	39.21 ± 0.61	51.72 ± 0.61
ORDER (Ours)	<b>41.25 ± 0.64</b>	<b>53.96 ± 0.71</b>
Joint-training	49.91 ± 0.79	61.78 ± 0.75

## 5.3 Comparison to Continual Learning

To evaluate the effectiveness of ORDER for mitigating forgetting in SETS, we compare to SOTA CL baselines, including Experience replay (ER) [15], A-GEM[14], MER [48], GPM [49] and DEGCL [13]. Note that all of these methods are originally designed for mitigating forgetting for *a small number of tasks*, we adapt these methods to SETS by combing them with meta-learning methods for mitigating forgetting. Here, we combine CL methods with MSKM for illustration since it is more efficient and

performs comparably in SETS than the other two meta-learning methods. The combination baselines are named as Semi-, etc. We also compare to (1) **sequential training (Seq)**, which trains the latent task distributions sequentially without any external mechanism and helps us understand the model forgetting behavior; (2) **joint offline training**, which learns all the task distributions jointly in a stationary setting and provides the performance upper bound.

The memory buffer is updated by reservoir sampling mentioned in Section 3. More details about how to update the memory buffer is provided in Appendix B. All the baselines and our method maintain 200 tasks in memory and the number of OOD data  $R = 50$  for each task. Table 3 shows the results. We also include a version of our method with memory replay and only OT regularization, without MI regularization, and other components are kept as the same as baselines. Our method can outperform baselines by 1.3% for 1-shot learning and 1.2% for 5-shot learning. Our full method significantly outperforms baselines for 5-shot learning by 3.1% and for 1-shot learning by 3.3%. This improvement shows the effectiveness of ORDER for mitigating forgetting in SETS.

#### 5.4 Ablation Study

**Robustness to OOD.** As for the OOD data of each episode, we randomly sample OOD from the mixture of the OOD dataset as described in Section 3. Furthermore, it is more natural for the unlabeled set to include more OOD samples than ID samples. To investigate the sensitivity of the baselines and our method to OOD data, we fix the number of ID data for each task and gradually increase the amount of OOD data in the unlabeled set. We set  $R$  (number of OOD data) = 50, 100, 150, 200, 250, respectively. The results are shown in Figure 1 in Appendix C. Interestingly, we can observe that our method is much more robust to OOD data than baselines, even when the number of OOD data is large.

To study the individual effect of mutual information regularization and optimal transport, we conduct ablation studies to justify (i) whether adding mutual information regularization helps to be robust to OOD data; (ii) whether adding OT helps mitigate catastrophic forgetting. The results of both ablation studies are shown in Table 8 in Appendix C. The  $\checkmark$  means the corresponding column component is used and  $\times$  means do not use the corresponding column component. Specifically, “MI” means using mutual information regularization, “OT” means using optimal transport constraint loss, “ $\mathcal{U}_{id}$ ” means using additional ID unlabeled data features in addition to labeled data for OT loss. We observe that the improvement of robustness to OOD becomes more significant, especially with a larger amount of OOD data. When the number of OOD is 250 for each task, the improvement is nearly 4%. This demonstrates the effectiveness of mutual information regularization. The optimal transport (OT) component improves model performance by more than 2% with 250 OOD data, demonstrating the effectiveness of remembering in feature space with both labeled and unlabeled data. Additionally, we compare OT regularization with only labeled data features to baseline. We find the performance is moderately improved in this case. This shows the benefits of using unlabeled data for remembering.

To investigate whether using OOD data is helpful, we performed comparisons to cases where we only use ID-data for training and prediction with OOD data filtered out. The results are shown in Table 9 in Appendix C. The results show that filtering out OOD data is already helpful, but with additional OOD data and MI regularization, the performance is improved further. Results on other ablation settings are also available in Table 9.

**Sensitivity to dataset ordering.** To investigate the sensitivity of baselines and our method to dataset order, we also performed comparisons on two other dataset sequences, including (i) Butterfly, CUB, CIFARFS, Plantae, MiniImagenet, Aircraft and (ii) CUB, CIFARFS, Plantae, MiniImagenet, Butterfly, Aircraft. The results are shown in Appendix C (see details in Table 1 and Table 2 for order (i), Table 3 and Table 4 for order (ii)). In all cases, our method substantially outperforms the baselines and demonstrates its superiority.

**Sensitivity of hyperparameter.** The sensitivity analysis of hyperparameters  $\lambda$  and  $\beta$  of ORDER are provided in Appendix C (see details in Table 5).  $\beta$  controls the magnitude of optimal transport regularization. Results indicate the model performance is positively correlated with  $\beta$  when it increases until optimal trade-off is reached after which performance deteriorates with over regularization when  $\beta$  reaches 0.01. A similar trend is observed in mutual information regularization  $\lambda$ .

**More Results.** To investigate the effect of memory buffer size, we provide more experimental results with a memory buffer size of 50, 100, 200 tasks in Appendix C (see details in Table 6). With increasing memory buffer size, the performance improves further. The current size of the memory buffer is negligible compared to the total amount of 72K tasks.

## 6 Conclusion

In this paper, we step towards a challenging and realistic problem scenario named SETS. The proposed method, i.e., ORDER, is designed to leverage unlabeled data to alleviate the forgetting of previously learned knowledge and improve the robustness to OOD. Our method first detects unlabeled OOD data from ID data. Then with the guidance of mutual information regularization, it further improves the accuracy. The method is also shown to be less sensitive to OOD. The OT constraint is adopted to mitigate the forgetting in feature space. We compare different methods on the constructed dataset. For the future work direction, we are working on adapting the method to the unsupervised scenario that no labeled data in new coming tasks, which requires the model with stronger robustness and capability of learning feature representations.

**Acknowledgement** We thank all the anonymous reviewers for their thoughtful and insightful comments. This research was supported in part by NSF through grant IIS-1910492.

## References

1. Agrawal, A., Amos, B., Barratt, S., Boyd, S., Diamond, S., Kolter, Z.: Differentiable convex optimization layers. In: *Advances in Neural Information Processing Systems* (2019)
2. Aljundi, R., Babiloni, F., Elhoseiny, M., Rohrbach, M., Tuytelaars, T.: Memory aware synapses: Learning what (not) to forget. *The 2018 European Conference on Computer Vision (ECCV)* (2018)
3. Aljundi, R., Kelchtermans, K., Tuytelaars, T.: Task-free continual learning. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2019)
4. Andrychowicz, M., Denil, M., Gomez, S., Hoffman, M.W., Pfau, D., Schaul, T., Shillingford, B., de Freitas, N.: Learning to learn by gradient descent by gradient descent. *Advances in Neural Information Processing Systems* (2016)
5. Antoniou, A., Edwards, H., Storkey, A.: How to train your maml. *International Conference on Learning Representations* (2019)
6. Antoniou, A., Patacchiola, M., Ochal, M., Storkey, A.: Defining benchmarks for continual few-shot learning. <https://arxiv.org/abs/2004.11967> (2020)
7. Bae, I., Moon, J., Jhung, J., Suk, H., Kim, T., Park, H., Cha, J., Kim, J., Kim, D., Kim, S.: Self-driving like a human driver instead of a robocar: Personalized comfortable driving experience for autonomous vehicles. *NeurIPS Workshop* (2019)
8. Balaji, Y., Chellappa, R., Feizi, S.: Robust optimal transport with applications in generative modeling and domain adaptation. In: *Advances in Neural Information Processing Systems*. pp. 12934–12944 (2020)
9. Barber, D., Agakov, F.: The im algorithm: A variational approach to information maximization (2003)
10. Belouadah, E., Popescu, A.: Il2m: Class incremental learning with dual memory. In: *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 583–592 (2019)
11. Berthelot, D., Carlini, N., Goodfellow, I., Papernot, N., Oliver, A., Raffel, C.: Mixmatch: A holistic approach to semi-supervised learning (2019)
12. Bertinetto, L., Henriques, J.F., Torr, P.H.S., Vedaldi, A.: Meta-learning with differentiable closed-form solvers. *International Conference on Learning Representations* (2019)
13. Buzzega, P., Boschini, M., Porrello, A., Abati, D., Calderara, S.: Dark experience for general continual learning: a strong, simple baseline. *34th Conference on Neural Information Processing Systems* (2020)
14. Chaudhry, A., Ranzato, M., Rohrbach, M., Elhoseiny, M.: Efficient lifelong learning with a-gem. *Proceedings of the International Conference on Learning Representations* (2019)
15. Chaudhry, A., Rohrbach, M., Elhoseiny, M., Ajanthan, T., Dokania, P.K., Torr, P.H.S., Ranzato, M.: Continual learning with tiny episodic memories. <https://arxiv.org/abs/1902.10486> (2019)
16. Chen, T., Wu, W., Gao, Y., Dong, L., Luo, X., Lin, L.: Fine-grained representation learning and recognition by exploiting hierarchical semantic embedding. *ACM International Conference on Multimedia* (2018)
17. Chen, W.Y., Liu, Y.C., Kira, Z., Wang, Y.C., Huang, J.B.: A closer look at few-shot classification. In: *International Conference on Learning Representations* (2019)
18. Cheng, P., Hao, W., Dai, S., Liu, J., Gan, Z., Carin, L.: Club: A contrastive log-ratio upper bound of mutual information. *Proceedings of the 37th International Conference on Machine Learning* (2020)

19. Diethe, T.: Practical considerations for continual learning (Amazon) (2020)
20. Ebrahimi, S., Meier, F., Calandra, R., Darrell, T., Rohrbach, M.: Adversarial continual learning. The 2020 European Conference on Computer Vision (ECCV) (2020)
21. Edwards, H., Storkey, A.: Towards a neural statistician. ArXiv [abs/1606.02185](https://arxiv.org/abs/1606.02185) (2017)
22. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. International Conference on Machine Learning (2017)
23. Finn, C., Rajeswaran, A., Kakade, S., Levine, S.: Online meta-learning. In: Proceedings of International Conference on Machine Learning (2019)
24. Finn, C., Xu, K., Levine, S.: Probabilistic model-agnostic meta-learning. Advances in Neural Information Processing Systems (2018)
25. Genevay, A., Peyré, G., Cuturi, M.: Learning generative models with sinkhorn divergences (2018)
26. Guo, J., Zhu, X., Zhao, C., Cao, D., Lei, Z., Li, S.Z.: Learning meta face recognition in unseen domains. Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
27. Guo, X., Yang, C., Li, B., Yuan, Y.: Metacorrection: Domain-aware meta loss correction for unsupervised domain adaptation in semantic segmentation. Conference on Computer Vision and Pattern Recognition (CVPR) (2021)
28. Horn, G.V., Aodha, O.M., Song, Y., Cui, Y., Sun, C., Shepard, A., Adam, H., Perona, P., Belongie, S.: The inaturalist species classification and detection dataset. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
29. Huang, X., Qi, J., Sun, Y., Zhang, R.: Semi-supervised dialogue policy learning via stochastic reward estimation. Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (2020)
30. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. International Conference on Learning Representations (2014)
31. Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A.A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., Hassabis, D., Clopath, C., Kumaran, D., Hadsell, R.: Overcoming catastrophic forgetting in neural networks. Proceedings of the national academy of sciences (2017)
32. Lake, B., Salakhutdinov, R., Gross, J., Tenenbaum, J.: One shot learning of simple visual concepts. Conference of the Cognitive Science Society (2011)
33. Lee, K., Lee, K., Shin, J., Lee, H.: Overcoming catastrophic forgetting with unlabeled data in the wild. 2019 IEEE/CVF International Conference on Computer Vision (ICCV) (2019)
34. Lee, K., Maji, S., Ravichandran, A., Soatto, S.: Meta-learning with differentiable convex optimization. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
35. Li, X., Sun, Q., Liu, Y., Zheng, S., Zhou, Q., Chua, T.S., Schiele, B.: Learning to self-train for semi-supervised few-shot classification. Proceedings of the Advances in Neural Information Processing Systems (2019)
36. Liu, Y., Lee, J., Park, M., Kim, S., Yang, E., Hwang, S.J., Yang, Y.: Learning to propagate labels: Transductive propagation network for few-shot learning. International Conference on Learning Representations (2019)
37. Lopez-Paz, D., Ranzato, M.: Gradient episodic memory for continual learning. Advances in Neural Information Processing Systems (2017)
38. Maji, S., Rahtu, E., Kannala, J., Blaschko, M., Vedaldi, A.: Fine-grained visual classification of aircraft. <https://arxiv.org/abs/1306.5151> (2013)



39. Mishra, N., Rohaninejad, M., Chen, X., Abbeel, P.: A simple neural attentive meta-learner. *International Conference on Learning Representations* (2018)
40. Munkhdalai, T., Yu, H.: Meta networks. *Proceedings of the 34th International Conference on Machine Learning* (2017)
41. Nguyen, A., Yosinski, J., Clune, J.: Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2015)
42. Nguyen, C.V., Li, Y., Bui, T.D., Turner, R.E.: Variational continual learning. *Proceedings of the International Conference on Learning Representations* (2018)
43. Nilsback, M.E., Zisserman, A.: Automated flower classification over a large number of classes (2008)
44. Poole, B., Ozair, S., van den Oord, A., Alemi, A.A., Tucker, G.: On variational bounds of mutual information (2019)
45. Raghu, A., Raghu, M., Bengio, S., Vinyals, O.: Rapid learning or feature reuse? towards understanding the effectiveness of maml. *International Conference on Learning Representations* (2020)
46. Ravi, S., Larochelle, H.: Optimization as a model for few-shot learning. In: *International Conference on Learning Representations* (2017)
47. Ren, M., Triantafillou, E., Ravi, S., Snell, J., Swersky, K., Tenenbaum, J.B., Larochelle, H., Zemel, R.S.: Meta-learning for semi-supervised few-shot classification. *International Conference on Learning Representations* (2018)
48. Riemer, M., Cases, I., Ajemian, R., Liu, M., Rish, I., Tu, Y., Tesauro, G.: Learning to learn without forgetting by maximizing transfer and minimizing interference. *International Conference on Learning Representations* (2019)
49. Saha, G., Garg, I., Roy, K.: Gradient projection memory for continual learning. *International Conference on Learning Representations* (2021)
50. Santoro, A., Bartunov, S., Botvinick, M., Wierstra, D., Lillicrap, T.: Meta-learning with memory-augmented neural networks. *Proceedings of the 34th International Conference on Machine Learning* (2016)
51. Schmidhuber, J.: A neural network that embeds its own meta-levels. *IEEE International Conference on Neural Networks* (1993)
52. Sehwag, V., Chiang, M., Mittal, P.: Ssd: A unified framework for self-supervised outlier detection. *International Conference on Learning Representations* (2021)
53. Smith, J., Balloch, J., Hsu, Y.C., Kira, Z.: Memory-efficient semi-supervised continual learning: The world is its own replay buffer. <https://arxiv.org/abs/2101.09536> (2021)
54. Snell, J., Swersky, K., Zemel, R.S.: Prototypical networks for few-shot learning. *Advances in Neural Information Processing Systems* (2017)
55. Tseng, H.Y., Lee, H.Y., Huang, J.B., Yang, M.H.: Cross-domain few-shot classification via learned feature-wise transformation. *Proceedings of the International Conference on Learning Representations* (2020)
56. Vinyals, O., Blundell, C., Lillicrap, T., Kavukcuoglu, K., Wierstra, D.: Matching networks for one shot learning. <https://arxiv.org/pdf/1606.04080.pdf> (2016)
57. Vitter, J.S.: Random sampling with a reservoir. *Association for Computing Machinery* (1985)
58. Vitter, J.S.: Random sampling with a reservoir. *ACM Transactions on Mathematical Software* (1985)
59. Vuorio, R., Sun, S.H., Hu, H., Lim, J.J.: Multimodal model-agnostic meta-learning via task-aware modulation. *Proceedings of the Advances in Neural Information Processing Systems* (2019)

60. Wang, L., Yang, K., Li, C., Hong, L., Li, Z., Zhu, J.: Ordisco: Effective and efficient usage of incremental unlabeled data for semi-supervised continual learning. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2021)
61. Wang, Y.X., Ramanan, D., Hebert, M.: Meta-learning to detect rare objects. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (October 2019)
62. Wang, Z., Duan, T., Fang, L., Suo, Q., Gao, M.: Meta learning on a sequence of imbalanced domains with difficulty awareness. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 8947–8957 (2021)
63. Wang, Z., Shen, L., Duan, T., Zhan, D., Fang, L., Gao, M.: Learning to learn and remember super long multi-domain task sequence. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 7982–7992 (2022)
64. Wang, Z., Shen, L., Fang, L., Suo, Q., Duan, T., Gao, M.: Improving task-free continual learning by distributionally robust memory evolution. In: *International Conference on Machine Learning*. pp. 22985–22998 (2022)
65. Wang, Z., Wang, X., Shen, L., Suo, Q., Song, K., Yu, D., Shen, Y., Gao, M.: Meta-learning without data via wasserstein distributionally-robust model fusion. In: *The Conference on Uncertainty in Artificial Intelligence* (2022)
66. Wang, Z., Zhao, Y., Yu, P., Zhang, R., Chen, C.: Bayesian meta sampling for fast uncertainty adaptation. *International Conference on Learning Representations* (2020)
67. Welinder, P., Branson, S., Mita, T., Wah, C., Schroff, F., Belongie, S., Perona, P.: *Caltech-UCSD Birds 200* (2010)
68. Yoon, J., Yang, E., Lee, J., Hwang, S.J.: Lifelong learning with dynamically expandable networks. *International Conference on Learning Representations* (2018)
69. Zhou, Y., Wang, Z., Xian, J., Chen, C., Xu, J.: Meta-learning with neural tangent kernels. *International Conference on Learning Representations* (2021)
70. Zhuang, P., Wang, Y., Qiao, Y.: Wildfish: A large benchmark for fish recognition in the wild (2018)