



Semi-supervised Drifted Stream Learning with Short Lookback

Weijieying Ren
University of Central Florida
Orlando, USA
wjyren@knights.ucf.edu

Pengyang Wang
University of Macau
Macau, China
pywang@um.edu.mo

Xiaolin Li
Nanjing University
Nanjing, China
lixl@nju.edu.cn

Charles E. Hughes
University of Central Florida
Orlando, USA
charles.hughes@ucf.edu

Yanjie Fu*
University of Central Florida
Orlando, USA
yanjie.fu@ucf.edu

ABSTRACT

In many scenarios, 1) data streams are generated in real time; 2) labeled data are expensive and only limited labels are available in the beginning; 3) real-world data is not always i.i.d. and data drift over time gradually; 4) the storage of historical streams is limited. This learning setting limits the applicability and availability of many Machine Learning (ML) algorithms. We generalize the learning task under such setting as a semi-supervised drifted stream learning with short lookback problem (SDSL). SDSL imposes two under-addressed challenges on existing methods in semi-supervised learning and continuous learning: 1) robust pseudo-labeling under gradual shifts and 2) anti-forgetting adaptation with short lookback. To tackle these challenges, we propose a principled and generic generation-replay framework to solve SDSL. To achieve robust pseudo-labeling, we develop a novel pseudo-label classification model to leverage supervised knowledge of previously labeled data, unsupervised knowledge of new data, and structure knowledge of invariant label semantics. To achieve adaptive anti-forgetting model replay, we propose to view the anti-forgetting adaptation task as a flat region search problem. We propose a novel minimax game-based replay objective function to solve the flat region search problem and develop an effective optimization solver. Experimental results demonstrate the effectiveness of the proposed method.

CCS CONCEPTS

• **Computing methodologies** → **Lifelong machine learning**: *Transfer learning*; *Semi-supervised learning settings*.

KEYWORDS

Distribution shift; Semi-supervised learning; Continual learning

ACM Reference Format

Weijieying Ren, Pengyang Wang, Xiaolin Li, Charles E. Hughes, and Yanjie Fu[1]. 2022. Semi-supervised Drifted Stream Learning with Short Lookback. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
KDD '22, August 14–18, 2022, Washington, DC, USA
© 2022 Association for Computing Machinery.
ACM ISBN 978-1-4503-9385-0/22/08...\$15.00
<https://doi.org/10.1145/3534678.3539297>

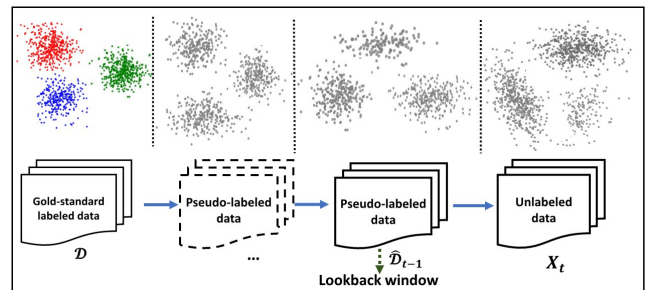


Figure 1: Learning in the semi-supervised, stream, gradually drifted data environment with short lookback.

and Data Mining (KDD '22), August 14–18, 2022, Washington, DC, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3534678.3539297>

1 INTRODUCTION

Many mobile Apps, such as Snapchat, generate unlabeled internet traffic streams in real time. In-App activities (e.g., share photos, videos, text, and drawings) could drift over time, resulting in distribution shifts. Due to mobile privacy concerns, many companies implement a very short data retention duration policy. We only have access to the most recent data (e.g., a short lookback window). Therefore, learning with short lookback in unlabeled drifted streams is critical for in-App activity classification. This scenario can be generalized as a new learning problem: *Semi-supervised Drifted Stream learning with short Lookback (SDSL)*, which is depicted in Figure 1. SDSL can enable a model to adaptively learn from an unlabeled stream of evolving distribution shifts, with limited initial labels and small lookback windows for training. Solving SDSL can address multiple critical issues to increase the availability and applicability of ML algorithms. For example, in many scenarios, 1) data are generated in real time; 2) and labeled data are expensive and only limited labels are available in the beginning of time; 3) real-world data is not always i.i.d. and data drift over time gradually; 4) the storage of historical streams is limited, so model updating can only be achieved based on a very short lookback window.

There are two major challenges in solving SDSL: 1) robust pseudo-labeling under distribution shifts, 2) anti-forgetting adaptation with short lookback. First, except for the initially given labels, all the incoming streams are unlabeled. The evolving distribution shifts further introduce bias into the task of labeling new data: training a classifier on the previously labeled data, but utilizing the classifier to

*Corresponding author

predict labels of forthcoming drifted data. Robust pseudo-labeling seeks to answer: how can we generate robust and quality pseudo-labels of unlabeled streams to augment old data? Second, SDSL can suffer from forgetting. This is because: 1) due to stream storage limitations and privacy-driven short data retention policies, less old data are stored for retraining; 2) and when a model adapts to new drifted data, the model parameters change to fit new data and forget old knowledge. A model that forgets old knowledge will perform poorly when new data with old distribution re-appear in the future stream. Therefore, anti-forgetting adaptation under short lookback is intended to answer: how can we learn new knowledge, while not forgetting old knowledge with limited historical data for replay?

Relevant works can only partially solve SDSL. First, SDSL is related to Semi-Supervised Learning (SSL) algorithms [1, 5], which combine a small amount of labeled data with a large amount of unlabeled data during training. However, in classic SSL, 1) the labeled and unlabeled data are assumed to sample from an i.i.d distribution; 2) and unlabeled data is static without evolving shift over streams [1]. Second, SDSL is related to continual learning that learns knowledge of new data without forgetting learned knowledge of old data. However, classic continual learning assumes newly generated data in streams are all labeled [15, 18]. Third, SDSL is related to domain adaptation [24, 25] that aims to transfer knowledge learned from one or multiple labeled source domains to an unlabeled target domain. However, in classic domain adaption, both source domains and target domains are static [15, 18]. Existing studies demonstrate the inability to jointly address both robust pseudo-labeling and anti-forgetting adaptation with short lookback in SDSL. As a result, we need a novel perspective to derive the novel formulation and solver of SDSL.

Our insights: an integrated robust and antiforgetting perspective. We formulate a generic learning problem of SDSL for semi-supervised, stream, limited lookback memory, evolving distribution shift environments. We show that semi-supervised learning in streams can be solved by iterating the label generation and the model replay process, where the label generation is to generate pseudo-labels for the newly coming unlabeled data stream, and the model replay is to retrain the learning model with old data of short lookback and pseudo-labeled new data. Robust pseudo-labels are important for effective SDSL. We find that leveraging invariant structure knowledge in streaming data can fight against the bias introduced by evolving distribution drifts for the robust pseudo-label generation. We demonstrate that to better learn patterns at the transition between old and new data, the model replay needs to be adaptive while anti-forgetting, even with a short lookback window. This requirement can be reformulated into a task of automated identification of flat regions. We highlight that the automated flat gradient region identification problem is indeed a minimax game. Solving the minimax game can effectively help the model to achieve both the anti-forgetting replay with limited old lookback data and adaption to new data.

Summary of Proposed Approach. Inspired by these findings, this paper presents the first attempt to develop a principled and generic generation-replay framework for the SDSL problem by iterating the robust pseudo-label generation and the adaptive anti-forgetting model replay. The framework has two goals: 1) robust pseudo-label generation against evolving distribution shifts in the




Figure 2: The training and testing stage of SDSL.

generation step; 2) balancing anti-forgetting replay and effective adaptation in the replay step. To achieve Goal 1, we develop a three-step robust label generation method to leverage multi-level knowledge. In particular, we find that robustness of pseudo-labels can be improved by modeling supervised knowledge from previously labeled data, unsupervised knowledge from new unlabeled drifted data, and structure knowledge from invariant label class semantics. Step 1 develops a supervised neural encoder-based classifier trained on previously labeled data. Then, we adopt a center-based clustering method to adjust labels on new drifted data, and leverage the invariance of label class semantics to regularize the neural encoder-based classifier. To achieve Goal 2, we develop an adaptive anti-forgetting model replay technique. In particular, we reformulate the adaptive anti-forgetting model replay into a flat region search problem. We propose a novel minimax game-based replay objective to automatically find the flat region to minimize the predictive loss on both previous data and new drifted data. And we develop an effective optimization method to solve the minimax game. Finally, we present extensive empirical results to demonstrate the effectiveness of our method for learning in the semi-supervised, streaming, gradually drifted, and short lookback setting.

2 PROBLEM STATEMENT

The SDSL Problem. Let \mathcal{D} to denote gold standard label data at $t=0$ and $\hat{\mathcal{D}}$ to denote pseudo-labeled data at $t>0$. Considering the existence of the initial labeled data that includes a feature matrix \mathbf{X}_0 and a gold standard label vector \mathbf{y}_0 in the very beginning (denoted by $\mathcal{D} = \{\mathbf{X}_0, \mathbf{y}_0\}$), and a drifted unlabeled data stream (denoted by a feature matrix list of stream segmentations $[\mathbf{X}_t]_{t=1}^T$), we aim to train an adaptive model to classify all the data points of the unlabeled data stream into a fixed number of classes. At time t , the model generates a list of pseudo-labeled sets $[\mathbf{y}_1, \dots, \mathbf{y}_{t-1}]$ for the list of unlabeled stream segmentations $[\mathbf{X}_1, \dots, \mathbf{X}_{t-1}]$.

During the training phase (Figure 2), we iteratively learn the adaptive model h that takes only the initial gold standard labeled data \mathcal{D} and the short lookback of the $(t-1)$ -th pseudo-labeled stream segmentation (denoted by $\hat{\mathcal{D}}_{t-1} = \{\mathbf{X}_{t-1}, \hat{\mathbf{y}}_{t-1}\}$) as inputs, and predicts the pseudo-labels (denoted by $\hat{\mathbf{y}}_t$) of the t -th stream segmentation (denoted by \mathbf{X}_t) at the time t . Formally, the approximation function h with learning parameters θ is given by:

$$h_{\theta}(\mathcal{D}, \hat{\mathcal{D}}_{t-1}, \mathbf{X}_t) \rightarrow \hat{\mathbf{y}}_t. \quad (1)$$

The optimization objective is to learn the model that can 1) generate robust pseudo-labels, and 2) prevent forgetting old data while adapting well to new data.

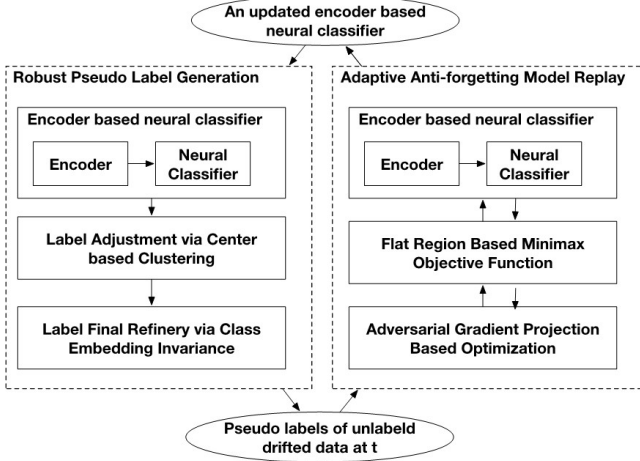


Figure 3: Overview of the proposed framework.

3 THE GENERATION-REPLAY FRAMEWORK

3.1 Overview

Figure 3 shows our proposed generation-replay framework including two components: 1) pseudo label generation; 2) adaptive anti-forgetting model replay. To achieve robust pseudo-label generation, we propose a three-step approach to improve the robustness of generated pseudo-labels by leveraging the knowledge from previously labeled data, new unlabeled data, and invariant label classes semantics, i.e., class correlations. In particular, in Step 1, we first train an integrated encoder and a classifier by minimizing the predictive loss on previous gold-labeled data and pseudo-labeled data at $t-1$. We then use the trained classifier to predict and obtain the initial pseudo-labels of forthcoming unlabeled drifted data at t . Since the trained neural classifier learns previous knowledge that is partially overlapped with the knowledge of forthcoming drifted data, the labels of new drifted data that are non-overlapped with previous knowledge could be biased and inaccurate. In Step 2, we then propose to use an unsupervised centroid-based clustering method to adjust the labels of new drifted data by repeatedly assigning the new drifted data to the nearest class centroid and updating class centroids until converged. The adjusted pseudo-labels of the forthcoming unlabeled drifted data are compared with model predicted labels to create loss signals as feedback to improve the neural integrated encoder-classifier method. Besides, we find that the semantic meanings of label classes remain invariant during the SDSL. We propose to leverage the invariance property of label class embeddings to further refine the class centroids. Specifically, we optimize the label class centroid via fixing label latent embedding learned from initial gold-labeled data. After obtaining the refined class centroids, we reassign forthcoming drifted data to the nearest class cluster centroid to generate debiased and robust labels. To achieve adaptive anti-forgetting model replay, we aim to retrain the integrated encoder-classifier to prevent forgetting previous knowledge while adapting well to new drifted data. We formulate this joint objective into a problem of searching the flat region. It is challenging to not just search the flat region but also identify the width of the flat region. We found that the challenge can be converted into a formulation of solving a minimax game. We develop an effective

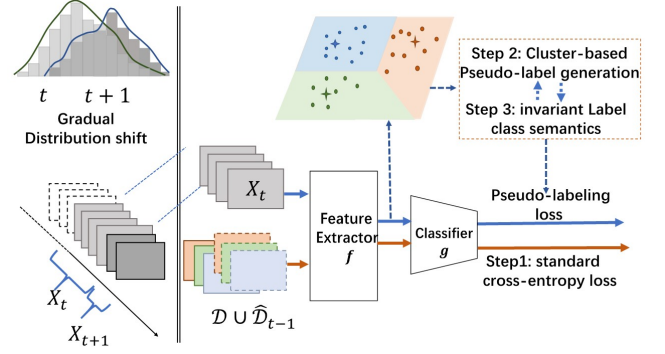


Figure 4: The Robust Pseudo-Label Generation Process.

optimization method to solve the minimax game to find the flat region and identify its width.

3.2 Robust Pseudo Label Generation

Why Robust Pseudo Label Generation Matters. Conventional semi-supervised learning generates pseudo-labels based on the predictions of a supervised model on labeled data (i.e., pseudo-label generation stage) and then integrates the pseudo-labeled new data to retrain an updated model (i.e., replay stage). In this way, the network gradually generalizes to unlabeled data in a self-paced curriculum learning manner [7]. However, such a strategy is not applicable to the SDSL setting because the unlabeled stream data drift over time. Note that the supervised model is trained based on previously labeled data. When the supervised model predicts on forthcoming drifted unlabeled data, the generated pseudo-labels are likely to be biased and inaccurate, which will propagate errors to the replay stage.

Leveraging Multi-level Knowledge to Robustify Pseudo-Label Generation. We find that label generation of unlabeled drifted data under SDSL can be made more robust by integrating supervised knowledge from previously labeled data, unsupervised knowledge from new unlabeled drifted data, and structure knowledge from invariant label class semantic meanings and relationships. Based on our unique insight, we propose a step-by-step testable method that includes three steps. Figure 4 illustrates the framework of the robust pseudo-label generation process.

Step 1: Leveraging Supervised Knowledge In a stream, data distributions drift gradually. In other words, the distribution of forthcoming unlabeled drifted data partially overlaps with the distribution of previous data. As shown in Figure 4, historical gold-labeled data and pseudo-labeled data still contain useful knowledge that can be used to predict the overlapped part of forthcoming drifted unlabeled data.

To this end, we develop an encoder-based neural classification model. This model jointly includes a multi-layer neural encoder f and a neural classifier g as an approximation function of the classification model: $h = g(f(\cdot))$. The neural encoder takes a data point as input and outputs the embedding of the data point. The neural classifier takes the embedding of the data point as features and outputs the predicted labels. Aside from the approximation function, we use the Cross-Entropy (CE) loss $\ell(\cdot, \cdot)$ to measure classification errors. Formally, the optimization objective function

can be formulated as:

$$\min_{\theta} \mathcal{L}_{CE}(\mathcal{D}, \hat{\mathcal{D}}_{t-1}; \theta) = \sum_{(\mathbf{x}_i, y_i) \in \{\mathcal{D} \cup \hat{\mathcal{D}}_{t-1}\}} \ell(h_{\theta}(\mathbf{x}_i), y_i). \quad (2)$$

In this way, we leverage the gold labels and previous pseudo-labels of historical data as supervision signals to learn the encoder-based neural classifier to further generate initial pseudo-labels for current unlabeled data \mathbf{X}_t , which will be introduced next.

Step 2: Leveraging Unsupervised Knowledge. Since the trained encoder-based neural classifier learns previous knowledge that is partially overlapped with the knowledge of forthcoming drifted data, the labels of new drifted data that are non-overlapped with previous data distributions could be biased and inaccurate. How can we improve the label quality of new drifted data whose patterns are not seen in the knowledge of previous data?

We find that unsupervised knowledge in new drifted unlabeled data is helpful for improving and refining the labels of such data themselves. Centroid-based clustering is an unsupervised learning method to exploit unsupervised information to discover data grouping patterns. Different from using the centroid-based clustering for data grouping, we propose to use such a method for data label adjustment. The high-level idea is to exploit the centroid-based clustering to adjust the labels of new drifted data by repeatedly assigning the new drifted data to the nearest class centroid and updating class centroids until converged. The underlying insight is that label class centroid-based clustering reassigns labels based on the global pattern structure of new unlabeled drifted data.

Specifically, in Step 2, we firstly use the trained encoder-based neural classifier to classify the labels of new drifted unlabeled data. We then use these classified labels to compute the centroids of all the label classes. In the initialization of the centroid-based clustering, we exploit all the class centroids as the initial cluster centroids. Formally, the centroid embedding in class c are initialized as via:

$$\mathbf{u}_c^{(0)} = \frac{\sum_{\mathbf{x}_i \in \mathbf{X}_t} \delta(h(\mathbf{x}_i)) f(\mathbf{x}_i)}{\sum_{\mathbf{x}_i \in \mathbf{X}_t} \delta(h(\mathbf{x}_i))}, \quad (3)$$

where $\delta(\cdot)$ is the softmax function [10]. We then repeat two tasks: assigning data points to the nearest class and updating class centroids, until converged (e.g., maximum number of iterations). Particularly, in the data reassignment task, given the class centroid \mathbf{u}_c , we construct the nearest centroid classifier to assign each unlabeled data point $\mathbf{x}_i \in \mathbf{X}_t$ to a class cluster as:

$$\hat{y}_i = \arg \min_{c \in C} d(f(\mathbf{x}_i), \mathbf{u}_c), \quad (4)$$

where $d(\cdot, \cdot)$ measures the cosine distance between data \mathbf{x}_i and the class centroid \mathbf{u}_c . And C denotes class numbers. Given the new pseudo-labels, we update the class centroids at each iteration k by:

$$\mathbf{u}_c^{(k)} = \frac{\sum_{\mathbf{x}_i \in \mathbf{X}_t} f(\mathbf{x}_i) * \mathbb{1}(y_i == 1)}{\sum_{\mathbf{x}_i \in \mathbf{X}_t} \mathbb{1}(y_i == c)}, \quad (5)$$

where $\mathbb{1}(\cdot)$ is the indicator function. The centroid-based clustering can reduce the error caused by supervised prediction under drifted data, and generate adjusted pseudo-labels.

Finally, the adjusted pseudo-labels of the forthcoming unlabeled drifted data are compared with model predictions to create loss signals as feedback to improve the training of the encoder-based

neural classifier. The objective function can be formulated as:

$$\min_{\theta} \mathcal{L}_{PL}(\mathbf{X}_t; \theta) = \sum_{\mathbf{x}_i \in \mathbf{X}_t} \ell(h_{\theta}(\mathbf{x}_i), \hat{y}_i). \quad (6)$$

Step 3: Leveraging Structure Knowledge of Invariant Label semantics Intuitively, the semantic meanings of label classes remain invariant during the SDSL. We show that leveraging the invariance property of label class semantics can further refine the quality of the generated pseudo-labels of new drifted data. Our insight is based on the invariance property of label classes. The embeddings of label classes should remain invariant over timelines.

The underlying idea is that label quality can be improved by reassigning data points based on improved label class centroids. To improve the accuracy of label class centroids, we can treat label class centroids as a matrix and exploit factorization-based matrix reconstruction to reconstruct improved label class centroids. Notably, factorizing the class centroid matrix can be factorized into the embeddings of features and the embeddings of label classes, which links to the invariance regularization of label class embedding.

Specifically, in Step 3, we first obtain the gold standard label class embedding $\hat{\mathbf{V}}$ by factorizing the class centroid matrix \mathbf{U} of the gold standard labeled data (\mathcal{D}) at $t=0$. We then perform factorization-based class centroid matrix reconstruction, by fixing the embedding of label classes $\hat{\mathbf{V}}$ to that learned from initial gold standard label data at $t=0$. Formally, the regularization term is defined as:

$$\mathcal{R}(\mathbf{U}_t) = \min_{\mathbf{U}_t, \mathbf{H}_t} \|\mathbf{U}_t - \mathbf{H}_t^T \hat{\mathbf{V}}\|_2. \quad (7)$$

With the given label class embedding $\hat{\mathbf{V}}$, we update class centroids \mathbf{U}_t and feature latent embeddings \mathbf{H}_t iteratively. By solving the optimization problem, we obtain the refined class centroids. Finally, we reassign forthcoming drifted data to the nearest class cluster centroid to generate debiased and robust pseudo-labels.

Final loss function The final objective function of robust pseudo-label generation can be represented as :

$$\mathcal{L}_{total} = \mathcal{L}_{CE}(\mathcal{D}, \hat{\mathcal{D}}_{t-1}; \theta) + \mathcal{L}_{PL}(\mathbf{X}_t; \theta) + \mathcal{R}(\mathbf{U}_t), \quad (8)$$

where $\mathcal{L}_{CE}(\mathcal{D}, \hat{\mathcal{D}}_{t-1}; \theta)$ represents cross-entropy loss on gold-labeled data \mathcal{D} and pseudo-labeled data $\hat{\mathcal{D}}_{t-1}$. Besides, $\mathcal{R}(\mathbf{U}_t)$ term stabilizes the updating of centroids which contains invariant label class semantics. A detailed optimization process can be referred to Algorithm 1 in Appendix A.1.

3.3 Adaptive Anti-forgetting Model Replay

After generating robust pseudo-labels for the newly coming unlabeled drifted data at time t , we treat such pseudo-labeled data as part of training data, and retrain the encoder-based neural classifier using the gold label data \mathcal{D} and the newly generated pseudo-labeled data $\hat{\mathcal{D}}_t$ at time t . To simplify our description, we ignore \mathcal{D} in the following sections, since \mathcal{D} is available across timelines.

Why Anti-forgetting Adaptation Matters? The key challenge of retraining the model is that, due to the privacy concerned short data retention policies in mobile and social applications and limited memory storage capacity of big stream data, a continuous learning model adapts to new data, new patterns, and new knowledge, while forgetting old data and knowledge at the same time. Figure 5 (a) shows that after adapting to new data, the new model (θ_2) shifts to

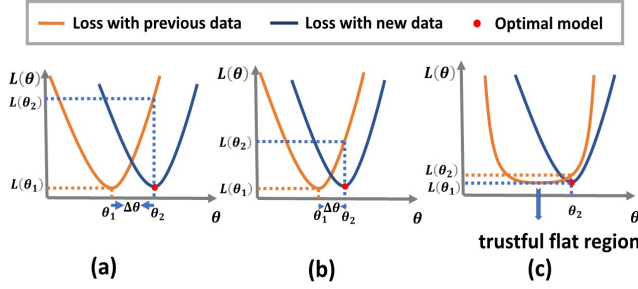


Figure 5: Motivation of the flat region.

the right side of the old model (θ_1). There are two key observations of adapting to drifted data: 1) the overlapped area between the new model and the old model becomes smaller and smaller. 2) the red inflection point of loss minimal in the new model results in a higher loss in the old model. Both observations indicate the new adapted model loses previous knowledge. How can we strive for a balance between anti-forgetting and adaptation? Many studies develop various technical solutions [9] to impose a regularization: the parameters of the new model should not deviate too much from previous parameters, i.e., $\min_{\theta_2} \|\theta_1 - \theta_2\|_2$. Such a regularization term, however, can still cause a significant performance drop on previous data. Figure 5 (b) shows the new model parameters are forced to be similar to the old model parameters, so that the overlapped area of the two models grows larger. However, the red-color loss minimal point of the new model still obtains a high loss in the old model, i.e., $\mathcal{L}(\theta_2) \gg \mathcal{L}(\theta_1)$.

Anti-forgetting Adaptation as A Minimax Game. Inspired by [20], we leverage the concept of the flat region to strive a balance between anti-forgetting and adaptation in semi-supervised stream learning. Figure 5 (c) shows why the flat region (denoted by $\theta^* - a \leq \theta \leq \theta^* + a$, where a is the width of the flat region) concept works. Fundamentally, the flat region represents a set of optimal or near optimal candidate model instantiations in the model space on old data. Finding such a flat region can effectively increase the overlapped area of the new model and the old model and maintain a low loss on previous data, while at the same time allowing the new model to adapt and shift to new data. This flat region searching and optimization relaxation process to strive for a balance between anti-forgetting and adaptation can be described by finding a model parameter (θ) that satisfies:

$$\begin{aligned} \min_{\theta} \quad & \sum_{x_i \in \mathcal{D}_t} \mathcal{L}(x_i; \theta) \\ \text{s.t.} \quad & \theta^* - a \leq \theta \leq \theta^* + a. \end{aligned} \quad (9)$$

where the width a of the flat region is manually specified based on empirical and domain experiences [20, 21].

However, in a dynamic learning environment of SDSL, the best flat region width will dynamically vary when both old data and new data change at different times. Therefore, it is impractical to directly integrate the above formulation of the flat region. The key research question is: can we find the flat region while automatically identify the best width of the flat region? We find that searching the flat region while identifying the best width can be reformulated into a computationally tangible minimax game. Assuming the parameters of a model can vary in a certain flat region defined by $\theta^* - \xi \leq$

$\theta \leq \theta^* + \xi$. We identify the upper bound of the model loss in the region of ξ and find the worst case by maximizing the training loss of the network on the new data \mathcal{D}_t , which is described by $\max_{\xi} \sum_{x_i \in \mathcal{D}_t} \mathcal{L}_{\mathcal{D}_t}(x_i; \theta + \xi)$. After the loss upper bound (worst case) over the region is measured, we minimize and lower the loss upper bound over the region to find the feasible parameters that can minimize the current loss. Besides, as depicted in Figure 5.(c), ξ lies in the neighborhoods around parameters trained on previous data \mathcal{D}_{t-1} , the optimization of ξ should follow $\xi \in \mathcal{M}$ where \mathcal{M} represents the space span by previous parameters trained on \mathcal{D}_{t-1} . Formally, the objective function can be formulated as:

$$\begin{aligned} \min_{\theta} \max_{\xi} \quad & \sum_{x_i \in \mathcal{D}_t} \mathcal{L}_{\mathcal{D}_t}(x_i; \theta + \xi) \\ \text{s.t.} \quad & \xi \in \mathcal{M}. \end{aligned} \quad (10)$$

Solving the Optimization Problem. Based on the gradient projection method [19], the adversarial weight perturbation ξ can be updated with the projection on space \mathcal{M} via the step size η_1 ,

$$\xi \leftarrow \xi + \eta_1 \text{proj}_{\mathcal{M}}(\nabla_{(\theta_t)} \mathcal{L}_{\mathcal{D}_t}(\theta_t + \xi)). \quad (11)$$

Notably, θ at time $t-1$ preserves the previously acquired knowledge which is spanned by \mathcal{M} . When coming to \mathcal{D}_t , we only update θ along the orthogonal direction of \mathcal{M} , leading to the least change (or locally no change) to the learned θ . Especially, the parameter θ can be adaptively updated with each \mathcal{D}_t as:

$$\theta \leftarrow \theta - \eta_2 (I - \text{proj}_{\mathcal{M}})(\nabla_{(\theta_t)} \mathcal{L}_{\mathcal{D}_t}(\theta_t + \xi)). \quad (12)$$

where I is the identity matrix and η_2 is the step size. Concretely, we present the Algorithm overview of the model replay stage in Appendix A.2. Additionally, we show a theoretical analysis on why the flat region can help mitigate forgetting when adapting on shifted streaming data with short lookback and why our method works in Appendix A.3.

4 EXPERIMENTS

We conduct extensive experiments on various datasets to evaluate the performance of our method. Specifically, our experiments aim to answer the following questions: **Q1:** Can our method outperform baselines on the semi-supervised drifted stream learning problem? **Q2:** Can our method generate robust pseudo-labels? **Q3:** Can our method effectively alleviate the forgetting problem? **Q4:** Can the flat region theory be supported by empirical investigations?

4.1 Experimental Setup

4.1.1 Data Description. We conducted experiments on eight datasets, including four widely-used synthetic benchmark datasets of stream learning research (i.e., UG_2C_2D, UG_2C_3D, UG_2C_5D and MG_2C_2D) and four real-world stationary classification datasets (i.e., Optdigits, Spam, Satimage and Twonorm). Table 2 in Appendix A.4 shows the statistics of datasets. Specifically, we exploited the same setting in [5, 22] to simulate the distribution shift manually by regrouping the instances for the four classification datasets. At each time, 1,000 instances arrived for UG_2C_2D and 2,000 instances arrived for UG_2C_3D, UG_2C_5D, MG_2C_2D, which were split into test and unlabeled datasets in a 30% and 70% ratio. For the Optdigits, Twonorm and Satimage datasets, 200 instances arrived each time and were split into 160 as unlabeled data and 40 as test

data. For the Spam dataset, 400 instances arrived at every time step and were split into 280 as unlabeled data and 120 as test data.

4.1.2 Baseline Algorithms. Since there are limited studies working on the SDSL setting, we compared our method with the semi-supervised learning methods [5], domain adaptation methods via evolving shifted data [9, 24] and other competitive baseline methods: (1) **Supervised Training (ST)**: simply trains the model on the gold standard labeled data once and tests the model on stream data without adaptation. (2) **Joint Training (JT)**: assumes that the gold standard labels are available for all the streaming data at each time. The model is jointly trained on all the labeled data ever seen. (3) **Pseudo-labeling with high confidence (PL-Conf)** [7]: stores examples with high softmax probability at each time. (4) **Evolution Adaptive Meta-Learning (EAML)** [9]: is a strong baseline that adapts to gradually shifted data without forgetting. EAML penalizes model parameters with a L_2 Regularizer to alleviate forgetting. (5) **Resource Constrained SSL under Distribution Shift (Record)** [5]: exploits a generation-detection-restoring pipeline. Differently, Record needs to generate the pseudo-labeled set based on the previously trained model and restore influential samples ever seen with a memory buffer. (6) **Domain-adversarial training of neural networks (DANN)** [24]: is a representative domain adaptation method. We applied DANN to the evolving distribution shift setting by training the model with labeled data, and learning an invariant embedding on the evolving shifted data sequentially.

We reported the average results with standard deviations of 5 runs for all experiments. Following the setting in [5], we chose the mean teacher (MT) [23] as the base model of SSL classifier in our framework. We used a two-layer multi-layer perceptron as a feature extractor. The memory reply buffer is set as 100 for our framework (the look back size) and baselines, i.e., only 100 unlabeled examples can be stored in memory. For PL_Conf, we chose the 100 most confident samples to retrain the model. The parameters of all the baseline models are defined in accordance with their respective publications. The step size η_1 and η_2 are set as 0.01.

4.1.3 Evaluation Metric. Following the setting in [5, 9], we evaluated the classification performance by averaging classification accuracy through each time as Acc_t and we evaluate the memorization ability by averaging classification accuracy on the final model as Acc_T :

$$Acc_t = \frac{1}{T} \sum_{i=1}^T R_{t,i}, \quad Acc_T = \frac{1}{T} \sum_{i=1}^T R_{T,i}. \quad (13)$$

where T is the total number of data sequences. $R_{i,j}$ is the test classification accuracy of the model at time j after learning the last sample from i -th data.

4.2 Q1: Overall Comparison

To answer Q1, **we compared our method with baselines that leverage unlabeled data through different strategies.** Figure 6 shows the mean classification accuracy and standard deviation on the test data for five runs. The Y-axis represents accuracy (Acc_t here), while the shaded regions show standard error computed using various random seeds. Figure 6 shows that our method achieves significant improvements over the baselines and is even comparable with the JT method (upper bound of the setting). The observation validates the effectiveness of leveraging multi-level knowledge

(supervised, unsupervised, and structure) for robust labeling and minimax-based flat region solver for anti-forgetting adaptation.

One interesting observation arises from the results that unlabeled data matter for streaming adaptation in the semi-supervised setting. Specifically, both DANN and EAML perform poorly in streaming shifted data. Our method and Record that utilize pseudo labels show significantly better performance. This implies pseudo labels can introduce auxiliary information about shifted data and improve generalization ability for streaming data. Moreover, while Record benefits from the influential shifted data detection mechanism, the pseudo labels are generated by the classifier trained on previous data. The performance gap between our method and Record validates the necessity to mitigate classifier bias and verified our motivation to leverage the structure knowledge of invariant label class semantics.

4.3 Q2: Study of Robust Pseudo Labeling

4.3.1 Effectiveness of Robust Pseudo Labeling. To validate the effectiveness of the proposed robust pseudo labeling method, we compare our method with two variants. Since our proposed solution is a generation-replay pipeline, we replace the generation part (pseudo label generation) with two widely used methods to construct the two variants: (1) the variant that takes PL_conf as the generation part, denoted by “Ours_PL”; (2) the variant that takes the Record’s generation method as the generation part, denoted by “Ours_Record”. The difference is the variant “Ours_PL” selects the most confident samples based on the softmax-based predictive confidence, while the variant “Ours_Record” selects the influential samples for the pseudo-labels generation.

Figure 7 shows our method uses the robust pseudo labeling method and outperforms the two variants. Specifically, the variant “Ours_Record” performs better than the variant “Ours_PL” since the Record method selects the most influential samples for the distribution change. This observation verified our motivation that mining unlabeled shifted data could boost adaptation performance. However, the selected pseudo labels still lie in the overlap region with previous data. Thus, it is insufficient to incorporate the non-overlapped information. In contrast, our method achieves consistently promising performance. The boost in performance verifies our motivation that class centroid-based clustering can exploit the global pattern structure and assign accurate labels for the new unlabeled drifted data that are non-overlapped with previous knowledge.

4.3.2 Ablation study of Robust Pseudo labeling. Our proposed robust pseudo labeling method generates promising pseudo labels for unlabeled data by preserving the invariant label semantics. To validate the contribution of the invariant label semantics, we conduct the ablation study, in which we compare our method with the variant that omits the Invariant Label Semantics (ILS) constraint in Equation 7. We denote the variant as “Ours w/o ILS”.

Figure 8 suggests that the ILS constraint contributes a stable classifier learning and can generate high-quality pseudo-labeled pairs. so, the model is easy to adapt well and mitigate forgetting to some extent due to the denoise ability. The observation verifies our motivation that the utilization of unsupervised knowledge of new data could provide extra information than previous labeled data and benefits the pseudo label generation of shifted data.

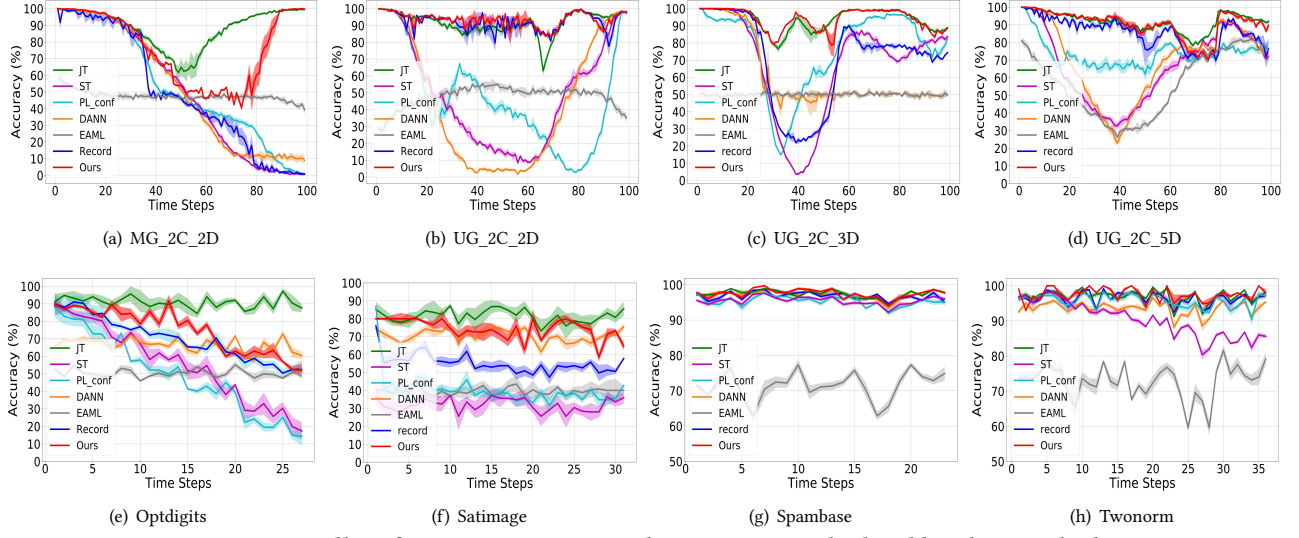
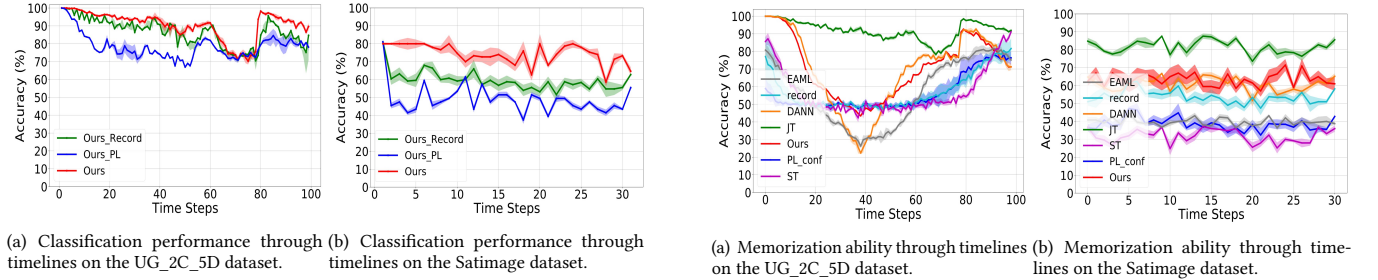


Figure 6: Overall performance comparisons between our method and baseline methods.

Table 1: Comparison of effectiveness on alleviating forgetting. Noted that Joint Training (JT) shows the ideal performance of the SDSL setting. The closer to JT, the better the performance.


Method	MG_2C_2D	UG_2C_2D	UG_2C_3D	UG_2C_5D	Optdigits	Satimage	Spam	Twonorm
ST	0.479±0.020	0.367±0.020	0.438 ±0.019	0.570±0.015	0.439±0.032	0.318±0.019	0.965±0.018	0.881±0.009
JT	0.593±0.012	0.907±0.0252	0.837±0.0103	0.919±0.007	0.921±0.019	0.816 ±0.018	0.971±0.016	0.966±0.012
PL_conf	0.513±0.020	0.531±0.021	0.567±0.219	0.614±0.019	0.406±0.025	0.385±0.026	0.966±0.021	0.956±0.011
DANN	0.532±0.039	0.349±0.011	0.578±0.018	0.701±0.014	0.415±0.031	0.607±0.019	0.969±0.013	0.948±0.018
EAML	0.508±0.014	0.499±0.016	0.499±0.012	0.569±0.016	0.527±0.018	0.397±0.015	0.725±0.014	0.729±0.021
Record	0.499±0.021	0.883±0.010	0.599±0.035	0.725±0.020	0.613±0.021	0.536±0.026	0.965±0.018	0.961±0.010
Ours	0.549±0.021	0.894±0.017	0.717±0.015	0.768±0.057	0.657±0.024	0.635±0.034	0.973±0.012	0.964±0.015



(a) Classification performance through timelines on the UG_2C_5D dataset. (b) Classification performance through timelines on the Satimage dataset.

(a) Memorization ability through timelines on the UG_2C_5D dataset. (b) Memorization ability through timelines on the Satimage dataset.

Figure 7: Validation of robust pseudo labeling.



(a) Classification performance (Acc_7) through timelines on Satimage data. (b) Memorization ability (Acc_7) through timelines on Satimage data.

Figure 8: Ablation study in Pseudo labeling.


4.4 Q3: Study of Alleviating Forgetting

4.4.1 Effectiveness of alleviating forgetting. To answer Q3, we validated the competence of our method on remembering previous

Figure 9: Memorization ability through timelines. Noted that JT shows the ideal performance of the SDSL setting. The closer to JT (in green curve), the better the performance.

knowledge. Noted that JT is the upper bound of the performance since JT assumes the availability of ground-truth labels across streaming data and restores all these data for training the model. In our experiments, we considered a sequential setting where shifted data come one after another. To validate the effectiveness of the proposed algorithm, we trained the model until the final task, and then tested the model performance on all previously seen tasks. All methods stop at $t = T$ after seeing all tasks.

Table 1 shows the average accuracy across all the tasks (Acc_T) on all the datasets. Specifically, we show the detailed results for each task along the timeline on *UG_2C_5D* and Satimage datasets,



(a) Classification performance through timelines on the Satimage data. (b) Memorization ability through timelines on the Satimage data.


Figure 10: Ablation study on flat region searching term.

which is illustrated in Figure 9. The results indicate that our method achieves better performance than other baseline methods on alleviating forgetting. Specifically, Record takes a reply method and stores influential samples with a memory buffer on each sequential task. However, the restricted storing buffer determines the generalization ability and is incapable of recovering all the learned information. On the contrary, EAML, DANN and our method relax the storage requirements. But EAML uses the meta-learning strategies and is inferior to other techniques due to the lack of high-quality validation set on shifted data. Moreover, DANN mitigates forgetting via learning an invariant representation, which is insufficient to store distinctive knowledge about previous data. Compared to these baseline algorithms, the performance of our method is closest to the upper bound (JT). The reason is that the generated high-quality pseudo labels provide exact supervisions to adjust model decision boundary, meanwhile, seek a flat minimal region to further enhance generalization ability.

4.4.2 Ablation study of alleviating forgetting. We introduce a Flat Region (FR) constraint to better alleviate the forgetting issue. In the experiment, we also conduct an ablation study to investigate the contribution of the flat region constraint. Specifically, we compare our method with a variant that omits the FR constraint in Equation 10, denoted as “Ours w/o FR”.

Figure 10 shows our method benefits from finding the flat minimal region. The flatness of the optimal minima makes the model insensitive to parameter change. Besides, the model updates the parameters in the orthogonal direction of previous parameters space, which preserves previous knowledge when adapting to new data.

4.5 Empirical Validation for Flat Region Theory



(a) Memorization ability through timelines on the UG_2C_5D dataset. (b) Memorization ability through timelines on the Satimage dataset.

Figure 11: Flat region validation.

In the Appendix A.2, we introduced the theoretical analysis of the flat region concept for boosting the model generalization ability. The flat region enables the model insensitive to the drift, resulting in promising generalization ability. Our question is: whether the

flat region exists? If yes, can our method find a wider flat region? Therefore, we conducted an empirical validation to answer the question. Suppose the flat region exists, the deviation of the model parameters within the region will not cause significant fluctuation of the model performance. Therefore, in the experiment, we validated the flat region theory based on a noise-sampling method. Specifically, to find the local optimal θ^* , we measure its flatness as follows. We first sampled noise from a pre-defined region $[0, b]$, then injected the noise to the trained model parameters only in the testing phase, and reported Acc_T . We changed the value of b , and averaged the model classification accuracy to measure its sensitiveness to noise. Intuitively, if the flat region exists, the model will be more robust against (less sensitive to) the injected noise. We compare our method with one variant that does not take flat region search (denoted as Record) in the same setting. Due to the page limitation, we only present the result on UG_2C_5D and Satimage.

Figure 11 shows that all the models have a respective flat region, but the range of the region is different. For example, for the UG_2C_5D dataset, when the upper bound b changes from 0 to 0.1, the performance of our method barely changes; when b is larger than 0.1, the performance begins to drop significantly. The observation reveals the fact that there exists a flat region $[0, 0.1]$, in which the model is insensitive to the drift. Similarly, we can also observe a flat region $[0, 0.15]$ for our method on the Satimage dataset. Among the comparisons, our method has the widest flat region, which means our method has the best generalization ability. Therefore, the empirical results validate the effectiveness of the seeking for a flat region method.

5 RELATED WORKS

Our work lies at the intersection of semi-supervised learning, continual learning (life-long learning), and domain adaptation. Next, we provide an overview of the related research efforts and briefly discuss the connections with our work.

Semi-Supervised Learning. Our work is related to semi-supervised learning (SSL) [7]. SSL is a special case of machine learning that leverages a large amount of unlabeled data with a small portion of labeled data to enhance the learning performance. SSL methods can be categorized into consistency-based [1], temporal ensembling [6], virtual adversarial training [13], pseudo labeling [7]. Most of SSL studies are designed both for offline and “identical and independent distribution” (i.i.d.) data but ignore the evolving nature of unlabeled samples. There are some emerging works designed for streaming data in the SSL setting by integrating local consistency propagation on graph [26, 27]. However, these methods assume the streaming data are i.i.d. with the labeled data which is not ideal for a realistic scenario. Recently, [5] considers learning from streaming data with a distribution shift in a semi-supervised way. However, they still generate pseudo labels via classifiers trained on previous data and maintain a memory buffer to store pseudo labeled data sequentially. Differently, We improve the pseudo label generation process considering shifted data, and only replay data with short lookback.

Continual Learning. Our method also connects to continual learning (CL). CL studies the problem of learning new information throughout their lifespan, without forgetting previously learned data or tasks. Generally speaking, there are three typical scenarios for CL: (1) class-incremental [11], (2) task-incremental [12], and

(3) data-incremental scenario (with class label set fixed) [2, 17]. Our work is related to the third category, but the incoming data is shifted without any supervision. In contrast, we release the requirement of the availability of labeled incoming data. Different from traditional regularization-based methods that penalize any changes to previous important parameters, we modify the objective function by introducing a flat region [20, 21] into SDSL setting. Different from previous works in CL that learn the flat region with prior knowledge [20, 21], we formulate the automated flat region identification problem as a minimax game into SDSL, which can ease the forgetting issue and adapt well to the new timeline.

Unsupervised Domain Adaptation. Our work is relevant to unsupervised domain adaptation [24]. Unsupervised domain adaptation aims to transfer the knowledge learned from labeled source domains to the unlabeled target domain. Existing works mainly focus on minimizing the discrepancy between the source and target distributions for learning domain-invariant features [4, 24]. Yet, recent theoretical analysis and empirical findings suggest that distinct marginal label distributions across domains provably undermine the target generalization. Therefore, to minimize the distance between labeling functions, [8, 16] accesses a small amount of labeled data in the target domain as extra supervision. Our work is enlightened by this line of works, however, we consider label class semantic constraints. Besides, we aim to alleviate the catastrophic forgetting problem, which is ignored by this pipeline.

6 CONCLUSION REMARKS

In this work, we provide a systematic analysis of semi-supervised drifted streaming learning with short lookback (SDSL), which is a realistic yet challenging setting without extensive study. To address this, we propose a novel method that follows the 'generation-replay' pipeline. To generate accurate pseudo labels for incoming shifted data, we leverage supervised knowledge of previously labeled data to label overlapped data, unsupervised knowledge of new data to refine non-overlapped data, as well as structure knowledge of invariant label semantic embedding to regularize the classifier. To achieve adaptive anti-forgetting model replay, we introduce the flat region notion and search the feasible region with a minimax game. Comprehensive experimental results verified our motivations and demonstrated the effectiveness of our method. In the future, we aim to explore more properties of unlabeled data to further improve the robustness of SDSL setting.

7 ACKNOWLEDGEMENT

This research was partially supported by the National Science Foundation (NSF) via the grant numbers: 2040950, 2006889, 2045567. This research was also supported in part by grants from the National Science Foundation Grants 2120240, 2114808, 1725554, and from the U.S. Department of Education Grants P116S210001, H327S210005, H327S200009. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsors.

REFERENCES

- [1] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin Raffel. 2019. Mixmatch: A holistic approach to semi-supervised learning. *arXiv preprint arXiv:1905.02249* (2019).

- [2] Aristotelis Chrysakis and Marie-Francine Moens. 2020. Online continual learning from imbalanced data. In *International Conference on Machine Learning*. PMLR.
- [3] Danruo Deng, Guangyong Chen, Jianye Hao, Qiong Wang, and Pheng-Ann Heng. 2021. Flattening Sharpness for Dynamic Gradient Projection Memory Benefits Continual Learning. *Advances in Neural Information Processing Systems* 34 (2021).
- [4] Basura Fernando, Amaury Habrard, Marc Sebban, and Tinne Tuytelaars. 2013. Unsupervised visual domain adaptation using subspace alignment. In *Proceedings of the IEEE international conference on computer vision*. 2960–2967.
- [5] Lan-Zhe Guo, Zhi Zhou, and Yu-Feng Li. 2020. Record: Resource constrained semi-supervised learning under distribution shift. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*.
- [6] Samuli Laine and Timo Aila. 2017. Temporal ensembling for semi-supervised learning. *International Conference on Learning Representations* (2017).
- [7] Dong-Hyun Lee et al. 2013. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, Vol. 3. 896.
- [8] Bo Li, Yezhen Wang, Shanghang Zhang, Dongsheng Li, Kurt Keutzer, Trevor Darrell, and Han Zhao. 2021. Learning invariant representations and risks for semi-supervised domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1104–1113.
- [9] Hong Liu, Mingsheng Long, Jianmin Wang, and Yu Wang. 2020. Learning to adapt to evolving domains. *Advances in Neural Information Processing Systems* 33 (2020), 22338–22348.
- [10] Weiyang Liu, Yandong Wen, Zhiding Yu, and Meng Yang. 2016. Large-margin softmax loss for convolutional neural networks. In *International Conference on Machine Learning*, Vol. 2. 7.
- [11] Yaoyao Liu, Yuting Su, An-An Liu, Bernt Schiele, and Qianru Sun. 2020. Mnemonics training: Multi-class incremental learning without forgetting. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*.
- [12] Davide Maltoni and Vincenzo Lomonaco. 2019. Continuous learning in single-incremental-task scenarios. *Neural Networks* 116 (2019), 56–73.
- [13] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. 2018. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence* (2018).
- [14] Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, and Nati Srebro. 2017. Exploring generalization in deep learning. *Advances in neural information processing systems* 30 (2017).
- [15] Simo Jialin Pan, Tsang, James Kwok, and Qiang Yang. 2010. Domain adaptation via transfer component analysis. *IEEE transactions on neural networks* (2010).
- [16] Xiaoying Ren, Jing Jiang, Ling Min Serena Khoo, and Hai Leong Chieu. 2021. Cross-Topic Rumor Detection using Topic-Mixtures. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*.
- [17] Xiaoying Ren, Linli Xu, Tianxiang Zhao, Chen Zhu, Junliang Guo, and Enhong Chen. 2018. Tracking and forecasting dynamics in crowdfunding: A basis-synthesis approach. In *2018 IEEE International Conference on Data Mining*. IEEE.
- [18] David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. 2019. Experience replay for continual learning. *Advances in Neural Information Processing Systems* 32 (2019).
- [19] Jo Bo Rosen. 1960. The gradient projection method for nonlinear programming. Part I. Linear constraints. *Journal of the society for industrial and applied mathematics* 8, 1 (1960), 181–217.
- [20] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. 2015. Trust region policy optimization. In *International conference on machine learning*. PMLR, 1889–1897.
- [21] Guangyuan Shi, Jiaxin Chen, Wenlong Zhang, Li-Ming Zhan, and Xiao-Ming Wu. 2021. Overcoming Catastrophic Forgetting in Incremental Few-Shot Learning by Finding Flat Minima. *Advances in Neural Information Processing Systems* (2021).
- [22] Vinicius MA Souza, Diego F Silva, João Gama, and Gustavo EAPA Batista. 2015. Data stream classification guided by clustering on nonstationary environments and extreme verification latency. In *Proceedings of the 2015 SIAM International Conference on Data Mining*. SIAM, 873–881.
- [23] Antti Tarvainen and Harri Valpola. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems* 30 (2017).
- [24] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. 2017. Adversarial discriminative domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 7167–7176.
- [25] Tianxiang Zhao, Xianfeng Tang, Xiang Zhang, and Suhang Wang. 2020. Semi-Supervised Graph-to-Graph Translation. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*.
- [26] Tianxiang Zhao, Xiang Zhang, and Suhang Wang. 2021. Graphsmote: Imbalanced node classification on graphs with graph neural networks. In *Proceedings of the 14th ACM international conference on web search and data mining*.
- [27] Tianxiang Zhao, Xiang Zhang, and Suhang Wang. 2022. Exploring edge disentanglement for node classification. In *Proceedings of the ACM Web Conference*.

A APPENDIX

A.1 Algorithm of Robust Pseudo-Label Generation

Algorithm 1 Robust Pseudo-Label Generation

Require: gold-label data \mathcal{D} , Pseudo-labeled set $\hat{\mathcal{D}}_{t-1}$, Randomly initialization on model parameter θ , Randomly initialization on latent feature matrix \mathbf{H}_t .

- 1: Pretraining the model on golden label set \mathcal{D} with feature extractor and classifier.
 - 2: Generating latent label vector \mathbf{V} with classifier parameters via SVD decomposition.
 - 3: **for** t in T **do**
 - 4: Training the model with \mathcal{D} and \mathcal{D}_{t-1} with θ .
 - 5: Applying the model with θ on \mathbf{X}_t to initialize prototypes \mathbf{U}_t in Eq.3.
 - 6: **while** Not Converged **do**
 - 7: Calculating $\hat{\mathbf{H}}_t$ with current \mathbf{U}_t and $\hat{\mathbf{V}}$ in Eq.7.
 - 8: Generating \hat{y}_t in Eq.4.
 - 9: Updating prototypes \mathbf{U}_t with Eq.5 and Eq.7.
 - 10: Update θ with gradients to minimize Loss \mathcal{L}_{full} .
 - 11: **end while**
 - 12: **return** Pseudo-labeled set \mathcal{D}_t and model parameter θ .
 - 13: **end for**
-

A.2 Algorithm of Adaptive Anti-forgetting Model Replay

After the robust pseudo-label generation stage, we obtain pseudo-labeled pairs \mathcal{D}_t (also collected it with memory buffer as shown in the paper.). Then we replay the data \mathcal{D} and \mathcal{D}_t , and then move to the next robust pseudo-label generation stage ($t+1$). Here, we represent the adaptive anti-forgetting model replay stage.

Algorithm 2 Adaptive Anti-forgetting Model Replay

Require: gold-label data \mathcal{D} , Pseudo-labeled set $\hat{\mathcal{D}}_t$, trained model parameters after time $t-1$, step size η_1 and η_2 , the model parameters θ train on $t-1$

- 1: **while** Not Converged **do**
 - 2: update θ on \mathcal{D} and $\hat{\mathcal{D}}_t$ via Eq.12
 - 3: update ξ via Eq.11
 - 4: **end while**
 - 5: **return** model parameter θ .
-

A.3 Theoretical Analysis

We present a theoretical analysis on why the flat region can characterize the continual learning property on streaming data and why our method works. Without loss of generalization, we simplify the drifted stream data with \mathcal{D}_{t-1} and \mathcal{D}_t , which are sampled from data distribution Q_{t-1} and Q_t , respectively. Based on previous works on PAC-Bayes bound [3, 14], given a ‘prior’ distribution \mathbf{P} (a common assumption is zero mean, σ^2 variance Gaussian distribution) over

the weights, the expected error can be bounded with probability at least $1 - \delta$:

$$\begin{aligned} & \min_{\Delta\theta} \mathbb{E}_{\xi} [\mathcal{L}_{Q_{t-1} \cup Q_t}(\theta + \Delta\theta + \xi)] \\ & \leq \underbrace{\min_{\Delta\theta \in M^c} \mathbb{E}_{\xi} [\mathcal{L}_{\mathcal{D}_{t-1}}(\theta + \Delta\theta + \xi)] + 4\sqrt{\frac{1}{n} KL(\theta + \xi || P) + \ln \frac{2N}{\delta}}}_{\text{Generalization Gap}} \\ & \quad + \max_{\xi \in M} [\mathcal{L}_{\mathcal{D}_t}(\theta + \Delta\theta + \xi)] - \mathcal{L}_{\mathcal{D}_{t-1}}(\theta + \Delta\theta) + \mathcal{L}_{\mathcal{D}_{t-1}}(\theta + \Delta\theta) \end{aligned} \quad (14)$$

where $\xi \in M$ and $\Delta\theta$ is updated along the orthogonal direction of previous optimal solution θ learned on previous data \mathcal{D}_{t-1} , i.e., $\theta \in M^c$. So that $\min_{\Delta\theta \in M^c} \mathbb{E}_{\xi} [\mathcal{L}_{\mathcal{D}_{t-1}}(\theta + \Delta\theta + \xi)]$ does not change too much compared with the previously minimized $\mathcal{L}_{\mathcal{D}_{t-1}}(\theta)$. Similarly, the updated parameters will not increase the training loss on \mathcal{D}_{t-1} . The second term depicts the Kullback Leibler (KL) divergence to the ‘prior’ \mathbf{P} [14]. Our method exactly optimizes the worst-case of the flatness of weight loss landscape $\max_{\xi \in M} \mathcal{L}_{\mathcal{D}_t}(\theta + \Delta\theta + \xi) - \mathcal{L}_{\mathcal{D}_{t-1}}(\theta + \Delta\theta)$ to control the above PAC-Bayes bound, which theoretically justifies why our method works.

A.4 Dataset Description

Table 2: Statistic Analysis of Datasets.

Dataset	Instances	Features	Classes
UG_2C_2D	100,000	2	2
UG_2C_3D	200,000	3	2
UG_2C_5D	200,000	5	2
MG_2C_2D	200,000	2	2
Optdigits	5620	64	10
Satimage	6435	36	7
Spambase	9324	500	2
Twonorm	7400	2	2

A.5 Additional Implementation Details

All experiments were conducted on the Ubuntu 18.04.5 LTS operating system, Intel(R) Core(TM) i9-10900X CPU@ 3.70GHz, and 1 way SLI RTX 3090 and 128GB of RAM, with the framework of Python 3.8.5 and PyTorch 1.8.1.