



HAL
open science

A Survey on Semi-supervised Learning for Delayed Partially Labelled Data Streams

Heitor Murilo Gomes, Maciej Grzenda, Rodrigo Fernandes De Mello, Jesse Read, Minh-Huong Le Nguyen, Albert Bifet

► **To cite this version:**

Heitor Murilo Gomes, Maciej Grzenda, Rodrigo Fernandes De Mello, Jesse Read, Minh-Huong Le Nguyen, et al.. A Survey on Semi-supervised Learning for Delayed Partially Labelled Data Streams. ACM Computing Surveys, 2022, 55 (4), pp.75. 10.1145/3523055 . hal-04468366

HAL Id: hal-04468366

<https://hal.science/hal-04468366v1>

Submitted on 6 Jan 2025

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

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

A Survey on Semi-Supervised Learning for Delayed Partially Labelled Data Streams

HEITOR MURILO GOMES, AI Institute, University of Waikato

MACIEJ GRZENDA, Faculty of Mathematics and Information Science, Warsaw University of Technology

RODRIGO MELLO, ICMC, University of São Paulo

JESSE READ, LIX, École Polytechnique, Institut Polytechnique de Paris

MINH HUONG LE NGUYEN, Télécom Paris, Institut Polytechnique de Paris

ALBERT BIFET, AI Institute, University of Waikato

Unlabelled data appear in many domains and are particularly relevant to streaming applications, where even though data is abundant, labelled data is rare. To address the learning problems associated with such data, one can ignore the unlabelled data and focus only on the labelled data (supervised learning); use the labelled data and attempt to leverage the unlabelled data (semi-supervised learning); or assume some labels will be available on request (active learning). The first approach is the simplest, yet the amount of labelled data available will limit the predictive performance. The second relies on finding and exploiting the underlying characteristics of the data distribution. The third depends on an external agent to provide the required labels in a timely fashion. This survey pays special attention to methods that leverage unlabelled data in a semi-supervised setting. We also discuss the delayed labelling issue, which impacts both fully supervised and semi-supervised methods. We propose a unified problem setting, discuss the learning guarantees and existing methods, explain the differences between related problem settings. Finally, we review the current benchmarking practices and propose adaptations to enhance them.

Additional Key Words and Phrases: semi-supervised learning, data streams, concept drift, verification latency, delayed labeling

ACM Reference Format:

Heitor Murilo Gomes, Maciej Grzenda, Rodrigo Mello, Jesse Read, Minh Huong Le Nguyen, and Albert Bifet. 2021. A Survey on Semi-Supervised Learning for Delayed Partially Labelled Data Streams. 1, 1 (June 2021), 35 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Situations where all the data are appropriately labelled, which allow us to perform supervised learning, are ideal, but many important problems are either unlabelled or only partially labelled. When dealing with streaming data, it is reasonable to expect some non-negligible verification latency, i.e. the label of an instance will be available sometime in the future, but not immediately. We identify data streams that exhibit both unlabelled data and verification latency as *Delayed Partially Labelled Data Streams*. These characteristics refer to how (and if) labels are made available to the learning algorithm, as illustrated in Figure 1.

A simple approach to cope with such data streams is to ignore both the unlabelled data and the labelling delay. Several methods were proposed, and evaluated, assuming a streaming scenario

Authors' addresses: Heitor Murilo Gomes, hgomes@waikato.ac.nz, AI Institute, University of Waikato; Maciej Grzenda, m.grzenda@mini.pw.edu.pl, Faculty of Mathematics and Information Science, Warsaw University of Technology; Rodrigo Mello, mello@icmc.usp.br, ICMC, University of São Paulo; Jesse Read, jesse.read@polytechnique.edu, LIX, École Polytechnique, Institut Polytechnique de Paris; Minh Huong Le Nguyen, minh.lenguyen@telecom-paris.fr, Télécom Paris, Institut Polytechnique de Paris; Albert Bifet, abifet@waikato.ac.nz, AI Institute, University of Waikato.

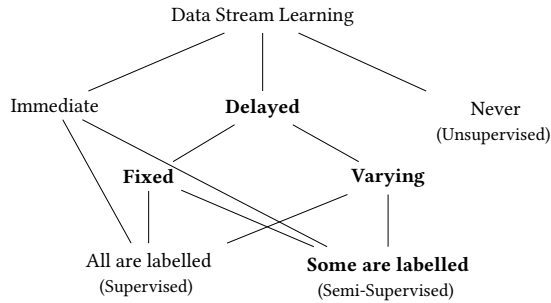


Fig. 1. Learning from data streams according to labels arrival time, based on [41]. Highlighted in **bold** the dimensions associated with delayed partially labelled data streams.

where all labels are immediately available [33, 72, 81]. More recently, some authors investigated how to leverage unlabelled data using semi-supervised learning (SSL) [50, 59], or active learning [106]. On top of that, significant advances were made in modelling and analysing the impact of delayed labelling in supervised learning and concept drift detection for data streams [45, 73, 105].

We focus the discussion on SSL methods for leveraging unlabelled data to enhance a supervised learning algorithm’s predictive performance. The basic assumption is that the algorithm has no influence over the labelling process, making active learning impractical. This work aims to organise the existing literature on SSL for data streams to facilitate new researchers and practitioners to navigate it. Concomitantly, we seek to elucidate the connections between the SSL and the delayed labelling literature to shed light on novel avenues for research. One challenging aspect of coping with delayed partially labelled data streams concerns the fair evaluation of algorithms. To assist in this perspective, we thoroughly discuss evaluation procedures for delayed partially labelled streams. This paper also aims to highlight the associations between related machine learning tasks, such as transductive learning, and to formalise the delayed partially labelled data streams.

This survey extends the existing literature by focusing on SSL and delayed labelling for data streams. It is complementary to the vast literature on semi-supervised learning for stationary data [21, 86, 103]; the evaluation of data streams [37], delayed labelling data streams [45] and SSL algorithms in general [71]; concept drift detection assuming immediate labelling [38, 91] or delayed labelling [105]; active learning for streaming data [104, 106]; and data stream mining in general [5, 11, 36, 43].

The rest of this work is organised as follows. We first introduce the problem statement, clearly identifying similarities and differences with related problems in Section 2. Next, in Section 3, we point out theoretical learning guarantees for SSL in both offline and online scenarios. Section 4 introduces existing SSL methods for streaming data. Section 5 includes a thorough discussion regarding the realistic assessment of SSL methods for data streams. The final Section 6 concludes the paper and discusses avenues for future research as envisioned by the authors.

2 PROBLEM DEFINITION

In this section, we introduce the definitions and explicitly state assumptions concerning the problem setting. Precisely, we begin with a general definition of **supervised learning** and then describe **verification latency**, and **partially labelled data** in the context of **evolving data streams**. We devote the end of this section to discuss the related problems to the setting we introduce.

Definition 2.1. Instance data: Let $X = \{x_0, \dots, x_1, x_\infty\}$ represent an open-ended sequence of observations collected over time, containing input examples in which $x_k \in \mathbb{R}^n$ and $n \geq 1$.

Definition 2.2. Labels: Let y be an open-ended sequence of target values collected over time, such that for every entry in y there is a corresponding entry in X , but the contrary may not be true, i.e., entries in X without a corresponding entry in y may exist. Furthermore, when y_k depicts a finite set of possible values, i.e., $y_k \in \{l_1, \dots, l_L\}$ for $L \geq 2$, it is said to be a classification task, while when $y_k \in \mathbb{R}$ it denotes a regression task.

Definition 2.3. Data Stream: Let Υ be a data stream i.e. a sequence of tuples $\mathcal{S}_1, \mathcal{S}_2, \dots$ which includes two types of tuples i.e.

$$\mathcal{S}_a = \begin{cases} \{(x_k, ?)\} & \text{if no true label is available yet} \\ \{(\cdot, y_k)\} & \text{when a true label for } x_k \text{ becomes available} \end{cases}$$

Hence, $\{(x_k, ?)\}$ is a tuple containing the observation, whereas $\{(\cdot, y_k)\}$ is a tuple containing the label corresponding to this observation.

Definition 2.4. Temporal-Mapping function: Let $T(\cdot)$ denote a function that extracts the precise discrete time unit t that x_k and y_k became available. It is relevant to mention that $T(x_k) \leq T(y_k)$ must always hold, indicating that the input data x_k becomes available at same the moment or before y_k .

Definition 2.5. Stream section: Let $\Psi[T_{\min}, T_{\max}]$ denote a stream section i.e. a sequence of instances and true labels that became available during a time window $[T_{\min}, T_{\max}]$. This means, $\forall x_k, y_k \in \Psi[T_{\min}, T_{\max}] : (T_{\min} \leq T(x_k) \leq T_{\max}) \wedge (T_{\min} \leq T(y_k) \leq T_{\max})$.

Definition 2.6. Verification latency: Let $V(x_k, y_k) = T(y_k) - T(x_k)$ represent the time difference a.k.a. ‘‘verification latency’’ of the labelled instance represented by the tuple (x_k, y_k) .

Definition 2.7. Infinitely delayed labels: Let $V(x_k, y_k) = \infty$ denote the verification latency of an infinitely delayed labelled instance a.k.a. unlabelled instance.

$T(x_k) = T(y_k)$, as seen in Definition 2.4, denotes a situation where both the input example and its label are provided at the same time instant, what is the same as receiving training instances from some batch learning task. Asymptotically, $V(x_k, y_k) \rightarrow \infty$ so that an observation x_k has no corresponding label y_k (see Definition 2.6).

Based on the aforementioned definitions:

- **(i) Immediate and fully labelled.** $\forall x \in X \wedge \forall y \in Y T(y) - T(x) = 1$, i.e., the verification latency between x and y corresponds exactly one time unit.
- **(ii) Delayed and fully labelled.** $\forall x \in X \wedge \forall y \in Y T(y) - T(x) = D$, where D is a random variable representing the discrete delay between x and y limited by the finite range $D \in \mathcal{Z}_+$, where $\max(D)$ denotes the maximum delay.
- **(iii) Immediate and partially labelled.** If we relax the constraint that every X has a corresponding entry on Y , we obtain a setting where X is only partially labelled. It is useful to emphasize the difference between entries in X which will be labelled as X_L and those that will not be labelled as X_U , and also to ascertain that often $|X_L| \ll |X_U|$ as the labelling process can be costly.
- **(iv) Delayed and partially labelled.** Similarly to (iii), we extend (ii) such that labels are delayed and some of them never arrived, i.e. they are infinitely delayed.

The majority of the literature with respect to semi-supervised learning for data streams has been devoted to (iii), while the intersection between delayed and partially labelled data, as in (iv),

is yet to be thoroughly explored. Besides the matters of label availability, another concept that is worth discussing in our definitions is whether the data distribution is **stationary** or **evolving**. In general, we assume evolving data distributions, thus **concept drifts** are deemed to occur, which may inadvertently influence the decision boundaries, and affect learned models. Note that if a concept drift is accurately detected (without false negatives) and dealt with by fully or partially resetting models as appropriate an independent and identically distributed (iid) assumption can be made (on a per-concept basis), since each concept can be treated as a separate iid stream, thus a series of iid streams to be dealt with¹. Nevertheless, the typical nature of a data-stream as being fast, dynamic and partially labelled encourages the in-depth study of methods for properly evaluate algorithms under these settings and semi-supervised algorithms to exploit unlabelled data.

2.1 Related Problems

In this section, we provide a short description of learning problems that are related to SSL for data streams, but that are not further scrutinized in this paper to avoid diverging from the delayed partially labelled problem setting.

Active learning. When dealing with an abundant amount of missing labels or a costly labelling procedure, *active learning* can be a viable alternative to SSL. Active learning approaches attempt to overcome the labelling bottleneck by querying the label of selected unlabelled instances to an oracle, such that the instances to be labelled are the most uncertain (e.g. a point lying close to the discriminative hyperplane) and that the answered labels can bring the highest value to the learning process. In this way, the active learner aims to achieve high accuracy using as few labelled instances as possible, thereby minimizing the cost of collecting labelled data [79]. Žliobaitė et al. [106] introduced a theoretical framework for active learning from drifting data streams. The authors stated three conditions for a successful active learning strategy in a drifting scenario: balancing the labelling budget over an infinite amount of time, perceiving changes anywhere in the instance space, and preserving the distribution of incoming data for detecting changes. Furthermore, in [106] three strategies were presented and empirically evaluated, assuming that an external adaptive learning algorithm is present.

Despite the advances in active learning for streaming data, it is sometimes hard to employ such strategies. The first reason is that the oracle’s response time may be too slow, as it often relies on a human expert. Second, still related to the labelling response time, if a concept drift occurs, the instances selected to be labelled may be outdated. The latter issue can be amended by using active learning strategies that take drift into account, as shown in [106]. Besides the issues involving the instability of the concepts, and delay to obtain the labels, Zhu et al. [104] also discusses the challenges related to the pool of candidates (instances to be labelled) being dynamic and issues related to the volume of data. To address these challenges, Zhu et al. [104] proposed an ensemble-based active learning classifier, where the minimization of the ensemble variance guides the labelling process.

Transductive learning. Transductive learning concerns a situation where the unlabelled test data set contains the whole of instances to be predicted, thus instead of producing a general model for predicting any future instance, the output is the predictions. This is a “closed world” assumption, where a successful solution is one where the algorithm can approximate the true labels of the instances solely for the finite test data set. This differs from inductive learning, where the goal is to yield a learning model capable of generalizing to previously unseen instances. Transduction is a powerful technique to leverage unlabelled data, but it is limited to situations where the goal is to

¹Not in every case a concept drifting stream can be decomposed into a sequence of iid streams. Theoretically, gradual (or incremental) drifts may occur where the distribution changes after every instance.

produce accurate predictions to a given set of instances and not devise a general rule. The majority of the algorithms for stream learning tend to focus on inductive learning. One possible reason is that traditional transductive methods require many computations. Thus frequently performing these may be prohibitive in a stream setting where predictions are often required to be fast. To circumvent this problem, Ho and Wechsler [49] proposed an incremental version of the transductive confidence machine (TCM) [39]. However, even though it is feasible to alleviate the computational aspects, another essential issue is that since data streams are unbounded, it is challenging to generate a closed set of instances.

Weakly multi-labelled data. Semi-supervised learning often stems from the case of having limited human labelling power to label all examples. Such a scenario is particularly inherent to data streams, where there are many instances, and they are arriving continuously. It is also aggravated when there are multiple label variables associated with each input – the so-called *multi-label* learning problem [84]. In this case, multiple labels $y_k \subseteq \{l_1, \dots, l_L\}$ (i.e., a *subset*) are associated with each instance. In this context, *weakly labelled* data (see, e.g., [82]) refers to instances where some, but not all of the relevant/true labels, have been applied to an instance. Specifically, the absence of a label in this subset does not necessarily imply that it is not relevant; and this is the challenge: to identify which of the non-relevant labels are missing in the labelled examples (it is not clear which ones are missing). A related concept of partial multi-label learning [95] is the generalization that additionally accounts for the possibility of false-positives (labels signalled as relevant, which are actually not). If we view a subset as binary indicator variables (as is typical in the literature), these problems become equivalently to L parallel (and possibly interdependent) noisily-labelled streams. Similar issues exist in the general multi-output case (extending to regression) [75, 89].

Missing values. Weakly multi-labelled data is also related to having *missing values* in the output/label space, except in this latter case it is clear which values are missing. This can be illustrated with an example in vector notation: $y_k = \{l_1, l_3\} \Leftrightarrow [1, 0, 1, 0]$ (supposing $L = 4$) where in the missing-valued case we may have $[?, 0, 1, ?]$ (compared to the weakly-labelled case where, e.g., $[0, 0, 1, 0]$ where l_1 is a false negative in our label set). Of course, it is also common to have missing values in the input space (as this affects all kinds of machine learning). This context is not a main focus of this survey. However, we note that a common method to deal with missing values is imputation. And, by building classifier or regression models to carry out this imputation (according to the variable domain being imputed), it is possible to frame a missing value imputation as a weakly-labelled multi-label problem [69]; which in turn can be seen as L parallel partially-labelled streams.

Initially Labelled Streaming Environment Labelled data may only be available at the beginning of the learning process. Therefore, a supervised learning algorithm can be trained with the initial data, and another unsupervised mechanism used to update the model during execution. This is a challenging problem setting as new labelled data is not available throughout execution, therefore it is not possible to confidently verify the accuracy of the model during execution. This setting was explored by Kreml [58], where the APT algorithm was proposed to track concept drifts even in the absence of labelled data. Later, Dyer et al. [35] proposed the COMPOSE framework to tackle the same problem setting, which also featured a detection mechanism that was independent of labelled data.

Few-shot learning. Few-shot learning [90] refers to feeding the learning algorithm with very few amount of labelled data. This approach is very popular in fields such as computer vision, where a model for object categorization is able to discern between different objects even without having multiple instances of each object for training. The term few-shot is accompanied by low-shot, 1-shot and 0-shot, which refer to training with a low amount of instances per class, only one per

class and not even one labelled instance for each class, respectively. As expected, as the number of labelled instances shrinks, the harder to produce accurate models. Approaching few-shot learning (and its variants) using semi-supervised learning is a common technique, also, when possible it is usual to leverage pre-trained models from similar domains (transfer learning).

Concept evolution. In some problems, the number of labels may vary over time. This problem is known as concept evolution [67]. Concept evolution characterizes a challenging problem where some instances are not only unlabelled but belong to a class that has not yet been identified. This is true for scenarios where one want to characterize malware per family instead of the comparatively more manageable task of classifying applications into malware or “goodware” (binary classification). In this survey, we do not approach such a problem as it requires a different definition of the problem as not all class labels are known *a priori*. A practical approach to address concept evolution in data streams is to leverage the clustering assumption [86] or apply novelty (i.e. anomaly) detection techniques to identify novel classes. Masud et al. [64] introduced DXMiner, an algorithm capable of detecting novel classes by identifying novel clusters, while Masud et al. [65] used an outlier detector and a probabilistic approach to detect novel classes. Abdallah et al. [1] proposes a method to continuously monitor the flow of the streaming data to detect concept evolutions, whether they are normal or abnormal.

3 LEARNING GUARANTEES

Supervised learning relies on different theoretical frameworks to ensure the conditions under which learning is guaranteed, being the Statistical Learning Theory (SLT) the most prominent contribution [87]. According to SLT, supervised learning is defined as the process involved in converging to the best as possible classification or regression function $f : X \rightarrow \mathcal{Y}$ contained in the algorithm bias \mathcal{F} , a.k.a. its space of admissible functions, in which X corresponds to the input space and Y to the output space containing labels.

This convergence process is essentially focused on approaching some loss measurement $R_{\text{emp}}(f)$ (or empirical risk) computed on training examples $(x_i, y_i) \in X \times Y$ to its expected value $R(f)$ (or risk) which is only computable by having the joint probability distribution (JPD) $P(X, Y)$. The basic and most important concept behind this convergence is to make possible the use of the empirical risk $R_{\text{emp}}(f)$ as a good estimation for the risk $R(f)$, provided $f \in \mathcal{F}$. Observe that by making sure $R_{\text{emp}}(f) \rightarrow R(f)$ and the training sample size $n \rightarrow \infty$, one can use the empirical risk to select the best classification function f^* by using:

$$f^* = \arg \min_{f \in \mathcal{F}} R_{\text{emp}}(f),$$

assuming the impossibility of computing the risk $R(f)$ for real-world problems, because we would never have access to the JPD.

Based on the Law of Large Numbers [29], Vapnik [87] formulated the Empirical Risk Minimization Principle (ERMP) to represent $R_{\text{emp}}(f) \rightarrow R(f)$ as $n \rightarrow \infty$ in form:

$$P(\sup_{f \in \mathcal{F}} |R_{\text{emp}}(f) - R(f)| > \epsilon) \leq 2\mathcal{N}(\mathcal{F}, 2n) \exp(-2n\epsilon^2), \quad (1)$$

given f is selected from the algorithm bias \mathcal{F} , the supremum reinforces the worst possible classifier that most influences in the divergence between both risks, $\mathcal{N}(\mathcal{F}, 2n)$ is the shattering coefficient or growth function defining the number of distinct classifications built from \mathcal{F} , $R_{\text{emp}}(f), R(f) \in [0, 1]$ and $\epsilon \in \mathbb{R}_+$.

Given the use of the Law of Large Numbers, a set of assumptions must be ensured to prove learning, otherwise the ERMP becomes inconsistent. The first assumption is that the JPD $P(X, Y)$ is fixed, so it does not change along with the data sampling, otherwise the convergence could not

be ensured given samples would follow a different probability distribution. Second, all samplings obtained from $P(X, Y)$ must be independent of one another and identically distributed so that every possible event from JPD will have its probability of being chosen as defined by its corresponding density.

It is relevant to mention that SLT can be mapped into other theoretical frameworks such as PAC-Learning and regularization methods [88]. Thus, from a such theoretical point of view, the following sections assess learning guarantees for both semi-supervised offline and online scenarios.

3.1 Semi-supervised learning in offline scenarios

From the perspective of the semi-supervised learning on offline scenarios, the assumptions after the Law of Large Numbers can be still met depending on the target application, to mention: (i) the joint probability distribution (JPD) $P(X, Y)$ must be fixed, and (ii) samplings from such JPD must be independent from each other. From such theory, if the JPD changes over time, we could somehow manage to obtain as much guarantee as possible so that the Empirical Risk Minimization Principle (ERMP) becomes partially consistent, and thus we can come up with some learning bounds. Complementary, if instances are not independent from one another, one option is to restructure data spaces as discussed in [24].

In this section, we consider that our semi-supervised offline scenario is represented by a single, and thus fixed, JPD whose data instances are independent of each other, while the next section considers the opposite scenario common in online learning. Therefore, let us have some dataset $(x_i, y_i) \in X \times Y$, for $i = 1, \dots, n$, containing n input examples x_i and their corresponding class labels $y_i = \{-1, +1\}$ with three possibilities: a negative, a positive and an empty label information. Consider as the absence of a class label so that one has no information about such instance, consequently its relative misclassification cannot be computed using a loss function $\ell(x_i, y_i, f(x_i))$ provided a classifier f . The absence of class labels is what makes this scenario be defined as a semi-supervised learning task, otherwise it would be a typical supervised task.

If we had all class labels, so that $y_i = \{-1, +1\}$, the ERMP after the SLT would be sufficient to represent learning bounds which formulates the conditions for which the empirical risk approaches the expected risk, i.e., $R_{\text{emp}}(f) \rightarrow R(f)$, as the sample size $n \rightarrow \infty$:

$$P(\sup_{f \in \mathcal{F}} |R_{\text{emp}}(f) - R(f)| > \epsilon) \leq 2\mathcal{N}(\mathcal{F}, 2n) \exp(-2n\epsilon^2), \quad (2)$$

given the empirical risk is computed on a sample:

$$R_{\text{emp}}(f) = \frac{1}{n} \sum_{i=1}^n \ell(x_i, y_i, f(x_i)),$$

and the expected risk on the JPD:

$$R(f) = \mathbb{E}(\ell(X, Y, f(X))).$$

In the same supervised scenario, Vapnik [87] proved the Generalization bound from Inequation 2 as follows:

$$R(f) \leq R_{\text{emp}}(f) + \sqrt{\frac{4}{n} (\log(2\mathcal{N}(\mathcal{F}, 2n)) - \log \delta)}, \quad (3)$$

for $\delta = 2\mathcal{N}(\mathcal{F}, 2n) \exp(-2n\epsilon^2)$, so that one can estimate how far the expected risk is from the risk computed on some sample plus a divergence defined by the squared-root term.

Once we do not have access to a fully labeled dataset so that $y_i = \{-1, +1\}$, we relax the learning bounds provided by the ERMP by redefining the sample size $n = \sum_i \mathbf{I}(y_i)$ as the number of labeled

instances according to the indicator function:

$$\mathbf{I}(y_i) := \begin{cases} 1 & \text{if } y_i \in \{-1, +1\}, \\ 0 & \text{otherwise.} \end{cases},$$

considering all available examples in some dataset $(x_i, y_i) \in X \times Y$, for $i = 1, \dots, n_{\text{all}}$, from which the same SLT bounds (see Inequations 2 and 3) can be ensured provided the JPD $P(X, Y)$ is fixed and data instances are iid.

In attempt to show how this first and straight-forward conclusion is useful to define learning bounds for semi-supervised learning in offline scenarios, consider some dataset $(x_i, y_i) \in X \times Y$, for $i = 1, \dots, n_{\text{all}}$ instances sampled from $P(X, Y)$. Consider that a fraction ν of instances was pre-labeled in $\{-1, +1\}$ so that $n = \nu \times n_{\text{all}}$, therefore considering some pre-defined δ as the upper probability bound for (see Inequation 2):

$$P(\sup_{f \in \mathcal{F}} |R_{\text{emp}}(f) - R(f)| > \epsilon) \leq \delta, \quad (4)$$

we can study the minimal training sample size [25] to ensure such bound which relies on $\delta = 2\mathcal{N}(\mathcal{F}, 2n) \exp(-2n\epsilon^2)$ as proved by Vapnik [87]. In that sense, let us assume the Shattering coefficient function $\mathcal{N}(\mathcal{F}, 2n) = n^2$ for a specific semi-supervised algorithm working on some d -dimensional Hilbert input space, thus defining the maximal number of distinct classifications as the sample size n grows. Given the shattering coefficient represents the complexity of the algorithm bias, term δ reflects such complexity in terms of the pre-labeled sample size n available for computing the loss function $\ell(x_i, y_i, f(x_i))$ provided every classifier $f \in \mathcal{F}$.

Thus, after assuming $\mathcal{N}(\mathcal{F}, 2n) = n^2$, we compute δ as follows:

$$\delta = 2\mathcal{N}(\mathcal{F}, 2n) \exp(-2n\epsilon^2) = 2n^2 \exp(-2n\epsilon^2) = 2 \exp(2 \log n - 2n\epsilon^2), \quad (5)$$

from which we can analyze the minimal labeled sample size n , characterizing the acceptable divergence between the empirical risk $R_{\text{emp}}(f)$ and its expected value $R(f)$. As this scenario considers a polynomial shattering coefficient, this curve produced by δ as n varies will approach zero. For instance, if we decide to accept a divergence of 5% ($\epsilon = 0.05$) and set $\delta = 0.1$ to have a probability of getting both risks acceptably close with probability of 0.9 (90% of the cases), we would need $n = 3,908$ training labeled examples. Consequently, one can find the minimal training sample size to ensure a given divergence ϵ and some probabilistic upper bound δ . For the sake of comparisons, we suggest to consider at least $\epsilon \leq 0.05$ and $\delta \leq 0.05$.

It is still worth to discuss how an algorithm bias changes the minimal training sample size necessary to address some semi-supervised task. The more complex an algorithm is, the steeper will be its shattering coefficient curve thus directly requiring more data instances to ensure the same learning guarantees, provided ϵ and δ . As discussed in [25], such complexity is related to the number of hyperplanes used to devise a proper decision boundary, the number of dimensions the input space has, as well as other factors such as how the dataset is organized (e.g. graph or table of variables). The need for estimating the shattering coefficient function to proceed with further algorithmic analysis has been motivating several studies in the last years [25, 26]².

Alternatively, we may consider the Generalization bound (see Inequation 3) to study a model we wish to induce from data. In that circumstance, assume we have estimated the shattering coefficient

²We suggest the following R Package to estimate the Shattering coefficient function – <https://cran.r-project.org/package=shattering>.

$\mathcal{N}(\mathcal{F}, 2n) = n^2$ after setting $\delta = 0.05$, then we have:

$$R(f) \leq R_{\text{emp}}(f) + \sqrt{\frac{4}{n} (\log(2\mathcal{N}(\mathcal{F}, 2n)) - \log \delta)} = R_{\text{emp}}(f) + \sqrt{\frac{4}{n} (\log(2n^2) - \log 0.05)},$$

$$R(f) \leq R_{\text{emp}}(f) + \sqrt{\frac{8 \log(n)}{n} + \frac{14.7555}{n}},$$

from which we conclude the empirical risk $R_{\text{emp}}(f)$ diverges from its expected value $R(f)$ according to term $\sqrt{\frac{8 \log(n)}{n} + \frac{14.7555}{n}}$, which naturally converges to:

$$\lim_{n \rightarrow \infty} \sqrt{\frac{8 \log(n)}{n} + \frac{14.7555}{n}} = 0,$$

as the training sample size n tends to infinity, thus proving learning in such a semi-supervised scenario. However, we may consider what such square-root term brings in terms of information and comparison among different learning settings. Consequently, the greater such term is, the more complex the algorithm bias and the necessary training sample size to ensure learning bounds.

The square-root term represents the variance provided the space of admissible functions \mathcal{F} . It consequently relates to the cardinality of the classification functions enclosed in the algorithm bias and the acceptable upper bound for the ERMP (see Inequation 2) given $\delta = \mathcal{N}(\mathcal{F}, 2n) \exp(-2n\epsilon^2)$, being therefore a way of regularizing the learning process. Regularization strategies are used to reduce the error by fitting an appropriate set of functions given some training set, consequently avoiding overfitting [88].

3.2 Semi-supervised learning in online scenarios

From the perspective of the semi-supervised learning on online scenarios, the assumptions after the Law of Large Numbers (LLN) must be somehow dealt with, to mention once more: (i) the joint probability distribution (JPD) $P(X, Y)$ must be fixed, and (ii) samplings from such JPD must be independent from one another. We easily conclude that both assumptions limit online learning in which we certainly expect the JPD to change over time, a classical aspect known as concept drift in the data streams scenario, as well as data observations will most certainly present some degree of dependence. Therefore, some strategy must be employed to still make the Empirical Risk Minimization Principle (ERMP) consistent so learning can be theoretically ensured.

As proposed by Pagliosa and Mello [24], Dynamical system tools can be used to reconstruct the input space X so that all dependences are represented in terms of a new set of dimensions. They employ Takens' embedding theorem [83] to reconstruct some unidimensional time series $S = \{s_1, \dots, s_t\}$ into some high dimensional space referred to as phase space Φ whose points or states $\phi_t \in \Phi$ are in form:

$$\phi_t = (s_t, s_{t+\tau}, s_{t+2\tau}, \dots, s_{t+(m-1)\times\tau}),$$

given τ refers to the necessary time delay to unfold the temporal relationships, a.k.a. time delay, and m corresponds to the embedding dimension or simply the number of axes necessary to represent all dependencies.

According to their approach, a single-dimensional data stream could be reconstructed into some phase space so that their temporal dependencies would be represented; therefore, all states $\phi_t \in \Phi$ would be independent of one another, thus solving Assumption (ii) of the LLN if one needs to perform some regression on unidimensional data. However, it leaves some important open questions: (a) how to deal with multidimensional data streams?; and (b) how to deal with the classification task of semi-supervised data streams?

Question (a) associated with Assumption (ii) was previously answered by Serra et al. [78] who used the same concepts from dynamical systems to reconstruct multivariate time series S for $s_t \in \mathbb{R}^d$, given $d > 1$, as follows (the upper index corresponds to each variable composing the multivariate time series S):

$$\phi_t = (s_t^1, s_{t+\tau}^1, \dots, s_{t+(m-1)\times\tau}^1, s_t^2, s_{t+\tau}^2, \dots, s_{t+(m-1)\times\tau}^2, \dots, s_t^d, s_{t+\tau}^d, \dots, s_{t+(m-1)\times\tau}^d),$$

so that, in addition to represent the temporal relationships of a single variable with itself, it also unfolds the dependencies that each variable of the time series (upper index) has with the others. Therefore, assuming that a multidimensional data stream has some time index as data observations arrive, one can extend Serra's framework to solve Assumption (ii) of the LLN, thus answering the first question. Observe it is not an unreasonable assumption to require data observations are indexed over time.

Observe that someone may doubt the presence of data dependence among stream observations. However, it is not difficult to mention several real-world phenomena illustrating such scenarios, such as in the context of air temperatures of a given world region, climatic variables, interaction of chemicals in reactions and the growth of populations along time [54]. Several researchers associated with the area of dynamical systems have been applying the same tools to obtain iid spaces [55, 77, 85].

Now, we get back to the Assumption (i) of the LLN, which requires the joint probability distribution (JPD) $P(X, Y)$ to be fixed in order to ensure learning. In that specific scenario, we suggest modelling the current JPD using past data observations and, as soon as some relevant data distribution change or drift is identified, past data must be discarded and the learning algorithm must start buffering new observations to induce a new model. Such approach is used in [27] to build up new models as data arrives, using McDiarmid's inequality to detect data drifts and indicate the best moment to retrain learning models using recent collected observations.

At last, all theoretical concepts addressed in this section intend to support other researchers to analyze their learning algorithms in an attempt of obtaining as much as possible guarantees while tackling partially labelled real-world streaming problems.

4 METHODS

Supervised machine learning is defined by using labelled data to train algorithms to predict unseen and unlabelled instances. These unlabelled examples do not influence algorithm anyhow. In most applications, obtaining labelled data is time-consuming and expensive, as labelling often depends upon human annotators. On the other hand, acquiring unlabelled data is an easier task, but these data cannot update supervised models directly. Semi-supervised learning is a paradigm of learning that exploit unlabelled data to leverage models trained with labelled data.

The caveat is that semi-supervised learning methods make strong assumptions about the data or the model [86, 103]. For example, one can assume a common underlying density across points belonging to a single class, or a common manifold underlying the points of each class. Figure 2 illustrates two such examples. Deciding which class to assign the test data point is relatively intuitive looking at all data points, but is not clear when considering labelled data points. This highlights precisely the advantages of using the unlabelled points.

Zhou and Li [102] organize techniques that leverage unlabelled data in roughly three categories: semi-supervised learning, active learning and transductive learning. This high-level organization does not take into account the constraints and objectives of the learning problem, for example, active learning is not applicable if an oracle is not feasible. More recently, Engelen and Hoos [86] organized techniques first in two classes, transductive and inductive. The majority of the techniques fall under the inductive category, since similarly to active learning, transductive learning

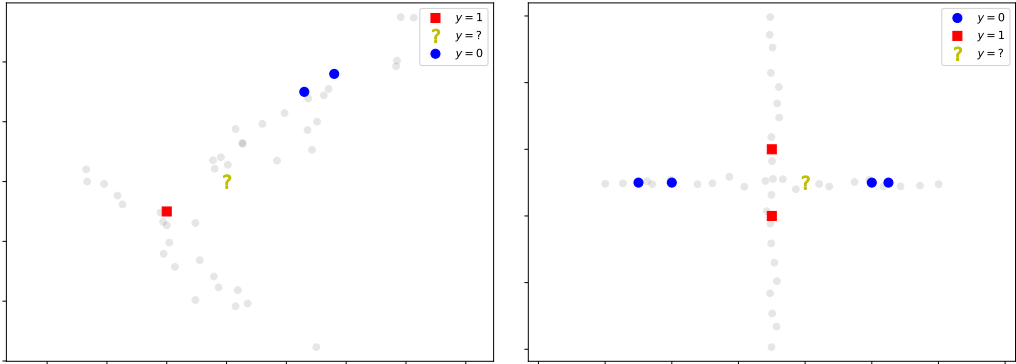


Fig. 2. Two illustrations of the utility of unlabelled data points in semi-supervised learning: A model is asked to produce a decision for a particular test point (shown in yellow), having observed many points, only a small number of which are labelled (shown in class 1 and 0 shown in red and blue, respectively). A semi-supervised method makes use of the unlabelled points to deduce a dense area per class (as would be appropriate in the example on the left), or a manifold (as appropriate on the right; a linear manifold in this particular example).

assumes specific characteristics of the learning problem as discussed in Section 2.1. Engelen and Hoos further divided inductive techniques into wrapper methods, unsupervised preprocessing, and intrinsically SSL. Wrapper methods includes those that leverage existing supervised learning algorithms, such as co-training, self-training and boosting algorithms. Unsupervised preprocessing denotes techniques that seek to improve the performance by extracting useful features from the input data without relying on labelled data. Finally, intrinsically SSL techniques include methods that are direct extensions of supervised learning methods, i.e., they extend the objective function of the supervised method to account for unlabelled data.

Even though it makes sense to leverage unlabelled data for the application of machine learning to streaming sources, this practice is relatively recent compared to similar approaches applied to static data. Therefore, some methodologies explored in the previously mentioned taxonomies are under-represented. For example, only a few works explore transductive learning for data streams, noticeable [30]. This section focuses on inductive methods, further categorizing such methods as: Intrinsically SSL, Self-Training, Learning by Disagreement, Representation Learning, and Unsupervised and SSL Drift Detection. All these methods categories can be found in the batch literature, except for drift detection. Drift detection methods are of extreme importance when dealing with streaming data as unsupervised or semi-supervised drift detection can serve several purposes, such as indicating when to acquire new labels, signal relevant changes to the domain that might have not yet influenced the decision boundary, and others.

4.1 Intrinsically SSL

We start our characterization by discussing streaming learners that exploit the unlabelled instances directly as part of objective function or optimization procedure.

Maximum-margin. Support Vector Machines (SVMs) are a popular method for supervised machine learning, based on the hypothetical maximum-margin classifier. The maximum-margin classifier's goal is to separate a binary classification problem, such that the hyperplane that splits the input space maximizes the margin. The margin is the perpendicular distance from the line to the closest points, namely the support vectors. In the construction of the classifier, only the support vectors are relevant. When learning a fully supervised SVM, the hyperplane is learned

from fully labelled training data through an optimization procedure that maximizes the margin. In complex data scenarios, it is unlikely that there is any hyperplane that entirely separates the data. Thus the constrain of obtaining the maximum-margin is relaxed, and the goal is to find a soft-margin classifier, i.e. one that allows some of the training examples to violate the maximum-margin assumption.

In the streaming setting, SVMs were explored for both supervised and semi-supervised problems. SVMs require considerable computational resources when trained on large volumes of data as the training problem involves the solution of a quadratic programming problem. Domeniconi and Gunopulos [32] compared several incremental techniques for constructing SVMs and shown that they achieve predictive performance closer to the batch version while requiring less training time. Differently, from prior incremental strategies for constructing SVMs, the techniques presented in [32] inspect data only once, which is more suitable for data stream processing than other techniques. Zhang et al. [99] introduces a relational k-means based transfer semi-supervised SVM learning framework (RK-TS3VM), where instances are organized into four types: labelled (type I) and unlabelled (type III) from the same distribution as the data arriving shortly; and labelled (type II) and unlabelled (type IV) from a similar distribution to the data arriving shortly. Learning from type I instances follows the traditional approach to maximize the margin given labelled instances, thus solving a constrained convex optimization problem. To learn from type II, III and IV, the authors had to modify the objective function (type II and III) and rely on a relational k-means (RK) model to build new features for the labelled examples using information extracted from type IV instances. Empirical results presented in [99] shows that RK-TS3VM outperform fully supervised SVM models trained only on the labelled data as well as S3VMs (semi-supervised support vector machines) [7].

Generative models. Generative models hypothesizes a model $p(x, y) = p(y) \times p(x|y)$ where $p(x|y)$ is an identifiable mixture distribution. Using the Expectation-Maximization (EM) algorithm [28], the mixture components can be identified given a large amount of unlabelled data. However, generative models require approximately correct modelling of the joint probability $p(x, y)$. If the modelling is incorrect, the unlabelled data may hurt performance. In contrast to discriminative models that only aim to estimate the conditional probability $p(y|x)$, $p(x, y)$ is more complicated to capture and too much/too little effort may lead to an incorrect model. Moreover, even if the modelling is correct, unlabelled data may hurt learning if a local maximum is far from the global maximum while using the EM algorithm.

Besides these limitations, the EM can be very slow to compute, especially when the data's dimensionality is high. Therefore, it is unlikely to apply the original EM algorithm to streaming data. Cappé and Moulines [17] introduced an online version of the EM algorithm with provably optimal convergence behaviour. The online EM algorithm [17] is analogous to the standard batch EM algorithm, which facilitates its implementation.

Nigam et al. [70] investigated the application of EM to text classification, such that the text documents are represented using a bag-of-words (BoW) model. Even though a BoW representation may conceal much of the complexities of written text, the authors show that there is a positive correlation between the generative model probability and the classification accuracy for some domains. In these cases, the application of EM alongside Naive Bayes suffices to leverage predictive performance. Such an approach could be adapted to data streams, given the combination of the online EM method [17] and the natural adaptation of Naive Bayes to perform incremental updates.

To compensate for the drawbacks of generative models, Grandvalet and Bengio [44] proposed a method called *Entropy Regularization*, aiming to only learn from unlabelled data that are informative, that is, when classes are well apart to favour the low-density separation assumption. Grandvalet and Bengio [44] argue that unlabelled data are not always beneficial, mostly when class overlap occurs,

and its informativeness should be encoded as a prior to modify the estimation process. The strength of the prior is controlled by a tuning parameter λ . A deterministic annealing process helps driving the decision boundary away from unlabelled data, thus avoiding poor local minima. Ultimately, the method they propose estimates the posterior distribution by maximising the likelihood of the parameters based on labelled data and at the same time being regularised by unlabelled data via λ .

4.2 Self-training

Self-training figures as another commonly used technique for semi-supervised learning. The idea is to let a classifier learn from its previous mistakes and try to reinforce itself. Self-training acts as a wrapper algorithm that takes any arbitrary classifier. Therefore, if we have an existing, fully-supervised learner that is complicated and hard to modify, self-training is an approach worth considering. Self-training has seen its application in natural language processing tasks such as word sense disambiguation [98] and sentiment analysis [62].

In an offline scenario, self-training works as follows. Given a dataset S that consists of a set of labelled data X_L and unlabelled data X_U such that $S = X_L \cup X_U$, a classifier C is trained on X_L and after that used to predict the labels in X_U . The predictions with a high confidence score are assumed true and added to X_L as new labelled data. The process repeats until convergence. When implementing a self-training algorithm, we must ponder the following issues: (i) how to evaluate the confidence of a prediction, and (ii) what the threshold for a "high" confidence score is? These issues remain relevant in an online scenario. Additionally, the learner must be adapted to learn incrementally from labelled and unlabelled instances coming from the stream.

Wei and Keogh [92] introduced experiments using a self-training (i.e. self-labelling) approach for time series classification. Special considerations were taken into account to leverage a one-nearest-neighbour classifier by using unlabelled data. The main challenge in adopting such a strategy to a streaming scenario is that it requires multiple passes over the input data. More recently, Jawed et al. [53] proposed a semi-supervised time series classification algorithm that leverages features learned from the self-supervised task on unlabelled data. It exploits the unlabelled training data with a forecasting task which provides a strong surrogate supervision signal for feature learning.

Le Nguyen et al. [59] proposed a self-training learner designed to receive as input either a single instance or a batch of instances at a time. A distance-based score was proposed to overcome the fact that some classifiers are unable to produce confidence scores. The confidence threshold that determines whether instances are used for self-training could be fixed or adaptive concerning the average confidence scores observed in a window. Le Nguyen et al. [59] observed that the variant using a windowed input, distance-based scoring, and fixed confidence threshold achieves the best performance. Similarly to [59], Khezri et al. [56] uses a the self-training approach which uses streaming classifiers predictions along with distance-based methods to select a set of high-confidence predictions for the unlabeled data.

4.3 Learning by disagreement

Learning by disagreement incorporates several strategies, which takes the form of learners "teaching" other learners. The canonical example is co-training [14], in which two models are trained on two different views of the same data. Multi-view learning [97] generalizes co-training to more situations where more than two views are available. Also, if multiple views are not available, one approach is to enforce the disagreement among the learners' predictions [102]. This artificially simulates multiple views from single-view data. The disagreement among learners can be achieved through many diversity-inducing techniques, such as bootstrapping aggregation [15].

Co-training. Blum and Mitchell [14] introduced co-training, relying on the intuition of using two separate learners to "guide" each other with the predicted labels they are most confident of

during the learning process. To achieve good predictive performance, co-training relies on two assumptions [14, 61]: the *consistency assumption* and the *independence assumption*. *Consistency* implies that the instance distribution is compatible with the target function, i.e. for most instances, the target functions over each view yield the same class label. Furthermore, the two views must be *conditionally independent*, given the class label. Given two views X^A and X^B and class label y , X^A and X^B are conditionally independent given class y iff, given any value of y , the probability distribution of X^A is the same for all values of X^B , and the probability distribution of X^B is the same for all values of X^A . The batch co-training procedure is relatively simple. First, learners A and B are trained on labelled views X^A and X^B , respectively. Second, alternately A and B yield predictions for the unlabelled data. Third, the most confident predictions produced by A are added to the training set of B and vice-versa. This process repeats until reaching a stopping criterion. The first apparent challenge in applying co-training to streaming data is that it is impractical to repeat the process iteratively. Another issue that arises in an evolving streaming setting is that each view's underlying data may change over time. During periods of change, learners will likely yield incorrect predictions with high confidence contributing to their counterparts' predictive performance degradation.

Learning by disagreement. The key ideas behind learning by disagreement is to generate multiple learners; let them collaborate to exploit the unlabelled data; and maintain a large disagreement between them. Learning by disagreement comprehends methods such as Tri-training [101] and Co-Forest [60], and it can be considered a generalization of co-training [14]. In these strategies, an set of learners (or ensemble) is trained on single view data, and the different views are simulated by enforcing diversity with respect to predictions through known techniques, such as bootstrapping [15] or random subsets of features, as in Random Forest [16]. One attractive idea behind using ensembles is that the majority of the methods do not restrict which base learner should be used, thus it becomes a fairly general technique for leveraging unlabelled data. The downside, as with similar techniques, is that if not properly regularized the base learners may converge to the same decisions, as members of the ensemble train each other.

Ensemble methods are popular approaches for supervised tasks involving data streams [40], for several reasons: (1) ensembles can leverage predictive performance, often surpassing what is achievable with a single (complex) learner; (2) ensemble-based methods can be coupled with concept drift detection [12, 41, 42]. However, ensemble methods are costly in terms of computational resources, which can be a major concern when dealing with data streams. Even though, several ensemble methods are embarrassingly parallel, streaming and parallel implementations of such algorithms requires extra efforts to better exploit the distributed setting [18, 41, 63].

The extension of co-training and learning by disagreement for data streams is not trivial as the algorithms that implement such techniques relies on multiple passes over the training data. Soemers et al. [80] leverages SSL in an unusual way while training an incremental regression tree, i.e. FIMT-DD [51]). In [80] the goal is to use FIMT-DD to cluster credit card transactions, and then apply a contextual multi-armed bandit (CMAB) algorithm that makes use of the structure of the FIMT-DD model. SSL is used to assist in the training of the FIMT-DD model. Since the FIMT-DD does not split a leaf node until a sufficiently large number of instances have been observed, the number of instances at each leaf can be large, which adversely affects the CMAB algorithm. To circumvent such problem, in regular intervals, logistic regression models are used to predict the labels of the transactions at the leaves, such predictions are then presented to the tree as true labels, which then can lead to further tree splits. This approach can be viewed, even though not explicitly mentioned by the authors, as a particular case of learning by disagreement.

4.4 Representation Learning

A general strategy for semi-supervised learning is to use unlabelled examples to build a representation of the input data, and then use this representation as input to a model for obtaining predictions. This technique is sometimes referred to as feature learning [6]. The idea is that an improved representation will lead to improved predictions; and since representation learning can naturally be an unsupervised task, training labels are not required. Figure 4 shows an illustration of this strategy.

Restricted Boltzmann machines (RBMs) are an example of kind of model that has been used in semi-supervised data stream contexts [76]. Trained using contrastive divergence, a single iteration can be carried out per instance, thus making them suitable for streams. As in the general strategy of representation learning, it is assumed that this representation improves the learning and prediction process whenever training labels are available, or predictions required, respectively.

One can use the incrementally-learned representation \mathbf{z} as input to any off-the-shelf data-streams classifier (naive Bayes, Hoeffding tree, etc.). A second option is to use the RBM's weights as the first layer of a neural network, to then be fine-tuned with back propagation [48] whenever a training label is available, with some form of stochastic gradient descent; a natural incremental algorithm. Predictions are carried out via a forward pass as in any multi-layer neural network.

In RBMs the variables are binary, $z_j \in \{0, 1\}$ but one may also use the probabilistic interpretations $[P(z_1|\mathbf{x}), \dots, P(z_k|\mathbf{x})]$ as the representation for an instance \mathbf{x} .

In a multi-label context, one may also obtain a representation of the label vector \mathbf{y} in a related manner [23] although to our knowledge streaming variations have not yet been developed.

Auto-encoders are another suitable (and related) approach. An auto-encoder is a neural network that learns to predict its own input. However, usually only the inner layer representation (\mathbf{z}) is of interest (hence, one can view Figure 4 (left) as an auto-encoder with the top part of the network removed). Again, as a neural network, gradient-descent based method, learning can be an inherently incremental process. This, as well as their non-linearities, make them more suitable and powerful for streams than linear methods such as principal components analysis [6].

It can easily be argued that RBMs are a particular kind of auto encoder. In both cases, it can be emphasised that many-layer (i.e., deep) models can be used (deep representation learning). In the case of RBMs, this is typically (but not always) done greedily. In a standard auto-encoder, it is simply a deep neural network where a single layer \mathbf{z} is taken. Again: the layer can be taken and given to any off-the-shelf data-stream learner (i.e., as a meta method), or turned into an instance- or batch-incremental neural network allowing back-propagation whenever labelled examples are provided by the stream.

Cluster representations are useful to identify cohesive sets of input instances, which in turn can be exploited by an SSL algorithm. The cluster-then-label technique assumes that instances belonging to the same cluster may share the same label. Applying classic clustering algorithms, such as k-means, to streaming data is challenging as such algorithms repeatedly iterate over the data. The majority of the stream clustering methods incrementally update micro-clusters (summarised representations of the input data). The actual clustering algorithm is only occasionally executed in an offline step using the micro-clusters as input. Fig. 3 illustrates a situation with three clusters summarising several micro-clusters and their respective instances. The instances in Fig. 3 are just for illustration purposes; the whole meaning of using micro-clusters is not to store the actual instances after the micro-cluster is updated.

One such clustering algorithm to follow this approach is CluStream [2]. CluStream takes a fixed number of micro-clusters m , which are updated whenever a new instance arrives. The offline phase of CluStream employs k-means to the micro-clusters. Recently, Le Nguyen et al. [59] proposed

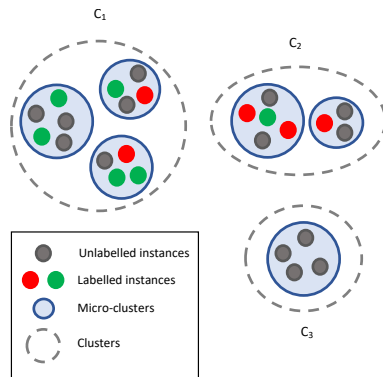


Fig. 3. An example of clusters, micro-clusters and instances. C_1 and C_2 clusters contain a majority of instances belonging to the green and red class labels, respectively. C_3 has no labelled instances; thus no inference about the label of new instances assigned to it can be made.

a cluster-then-label approach utilizing CluStream, such that each cluster had an associated class label frequency counter. Pseudo-labels were assigned to arriving unlabelled data according to the most frequent label associated with its closest cluster. A similar strategy was earlier explored by Masud et al. [66], where an individual model created K micro-clusters from a chunk of data. The prediction was given after determining the closest k nearest clusters from it. The predicted class label was the one with the highest frequency of labelled data across all of the closest k clusters.

Realistically, any unsupervised method that can produce a useful representation of the (unlabelled) data can be considered potentially useful in the semi-supervised settings. And any algorithm for such a representation that may be suitable for a data stream is thus suitable for semi-supervised learning in a data-stream setting. Mixture models are typically trained using the EM algorithm, which is an iterative algorithm requiring several sweeps over the data, however it can be adapted to streams [17]. In fact, the EM algorithm and k-means are special cases of self-training (see Section 4.2).

4.5 Unsupervised and SSL Drift Detection

Real-world problems tend to be very dynamic. For example, consumer behaviour may change as time goes by, a group of people can change their opinion about a product or a political party, the attacks a network receives change as new barriers are created, and so on. Learning from data that distribution may change over time is challenging for conventional batch machine learning algorithms. These algorithms assume that the data distribution is static. Conventionally, data streams that contain drifts are identified as evolving streams.

There are many aspects to consider when discussing concept drift, including its cause, rate, and data distribution. Generally, a drift can be characterized either as “real” or “virtual” [38]. A real concept drift happens when changes affect the class labels’ posterior probabilities, $p(y|X)$, i.e., the output variable distribution changes affecting the upcoming predictions. Virtual concept drift is said to occur if the distribution of the incoming data $p(X)$ changes without affecting $p(y|X)$. Usually, there is not much interest in virtual drifts because they do not change the output’s conditional distribution. A sizeable amount of research has been dedicated to discuss different aspects concerning concept drift [38, 91]. This section focuses on discussing concept drift in scenarios where labels are delayed and often partially available.

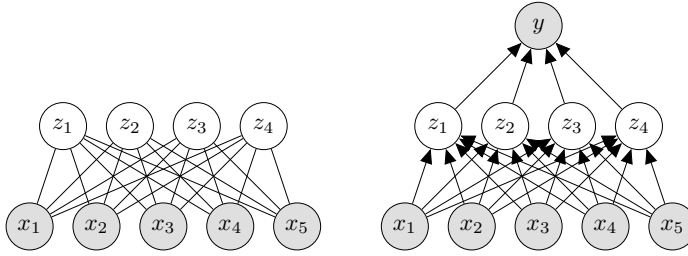


Fig. 4. An unsupervised model (left), able to form a representation of data points as z_1, \dots, z_k . In this figure an undirected graphical representation is depicted, but representations of generative models (where arrows point from z to x) are also possible, depending on the learning algorithm chosen for this step. In a second step, the representation can then be used directly as *input* to the supervised learning model (along with training label y , whenever it is available; i.e., learns to map $z \mapsto y$). A second specific option is to consider the representation *part of* a neural network (shown here, right – where arrows show the direction of the forward pass) and use a backward pass through all layers whenever a training label y is available – thereby fine-tuning the representation for discriminative power. Shaded nodes are those observed in the data stream and white nodes are the latent/hidden representation that is learned.

Most concept drift detection algorithms are applied to the univariate stream of correct/incorrect classifier predictions. Such strategies require that labelled data is available as soon as possible to respond to concept drifts in a timely fashion. From a practical standpoint, despite their intrinsic differences, most drift detectors trigger when the observed model’s predictive performance starts to degrade. Algorithms such as ADWIN [10] and EDDM [4], were tested under the assumption that labelled data is available almost immediately. However, if the ground-truth is not immediately available, then these algorithms’ ability to timely detect drifts is severely decreased.

To illustrate the impact of delayed labelling on a drifting stream, we present a small experiment with data generated using the AGRAWAL generator with 3 abrupt concept drifts (at instances 25, 000, 50, 000, and 75, 000). In these experiments, we used an ensemble algorithm [42] capable of detecting and adapting to changes by resetting base models whenever changes are detected (by ADWIN [10]) on their univariate stream of correct/incorrect predictions. Figure 5 depicts the amount of concept drifts detected (y-axis) over the processing of 100, 000 instances with and without delayed labelling. In Figure 5, the detections for the “No delay” experiment shows a high rate of detection immediately after the concept drifts, except for a few arbitrary drift signals in-between the concept drifts. A labelling delay of 10, 000 instances severely impacts the ability to detect the changes, as shown in the Delayed labelling variant, where the detections still occur but with a shift of approximately 10, 000 instances. The impact on the predictive performance in the delayed and not delayed scenarios is clearly observable in Figure 6. The experiment depicts the prequential accuracy³ as the 100, 000 instances are processed. The delayed labelling causes longer periods of poor accuracy (below 50% on a balanced binary problem).

The occurrence of a concept drift indicates that something is off concerning the learning model. Several actions can be taken after drift is detected, such as raise the alarm to the user, trigger automatic changes to the underlying model, or selectively request new labelled data (active learning). Research is often devoted to automatic detection and recovery [41]; however, these can be frowned upon because they give less control over the model. Raising the alarm and signalling the need for new labelled data is a less drastic approach that gives the data scientists behind the model more control. This control is especially useful in scenarios where prediction stability and fairness

³A window of 1, 000 instances was used in this prequential accuracy evaluation.

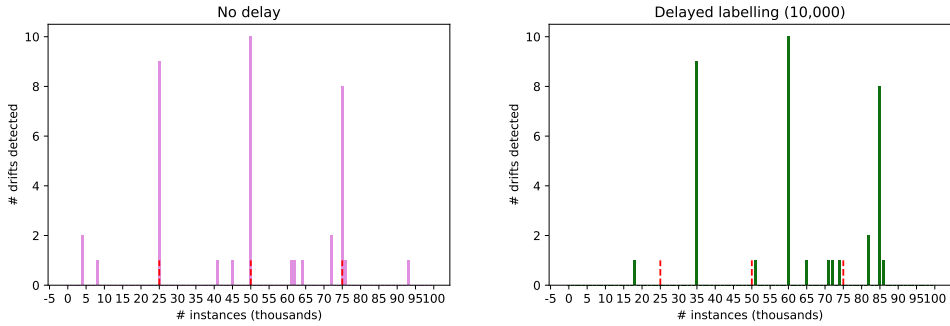


Fig. 5. Drifts detected by a 10 learner SRP model using ADWIN on AGRAWAL with and without labelling delay. Red dotted vertical lines indicate the location of concept drifts.

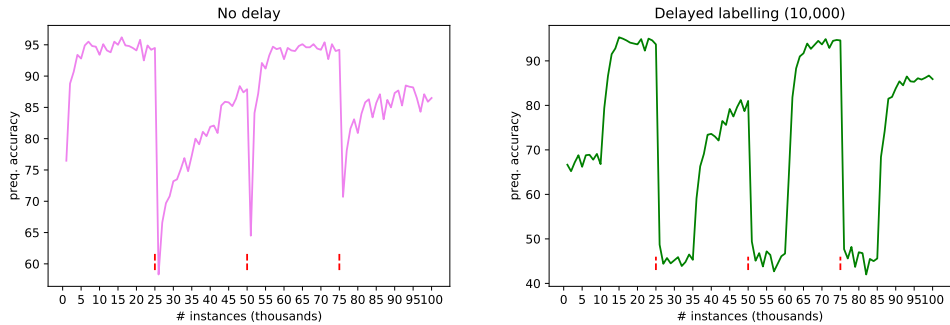


Fig. 6. Accuracy by a 10 learner SRP model using ADWIN on AGRAWAL with and without labelling delay. Red dotted vertical lines indicate the location of concept drifts.

are equally (or more) important than accuracy. A brand new model might be accurate for the immediately new data, yet it may be unstable to produce predictions given that it was only trained on a small amount of data. On top of that, automatically re-training the model leaves room for unfair treatment of underrepresented groups in the data, e.g., it is harder to prevent gender and race discrimination if the model is being updated on the fly without being reviewed. Therefore, one interesting application is to detect concept drift and only notify the user. Depending on the machine learning pipeline, a new model might be created on the fly to replace the old model, but it will not be deployed immediately without user interference. Another option is to notify the user to suggest instances for labelling using active learning strategies for data streams [106].

Žliobaite [105] presented an analytical view of the conditions that must be met to allow concept drift detection in a delayed labelled setting. These conditions depends on characteristics of the concept drift, i.e. how the change affected the input probability $P(X)$ and the posterior probabilities $P(w_i|X)$ of the class labels w_i . Žliobaite [105] thoroughly discussed when it can be expected that changes are observable or not as shown in Table 1. One interesting connection between semi-supervised learning and unsupervised concept drift detection is that if the underlying marginal data distribution $P(X)$ over the input does not contain information about $P(y|X)$ or indicates changes on $P(y|X)$, then it is impossible to exploit unlabelled data to improve a supervised learner (SSL) or detect a change ((3) in Table 1).

		$P(w_i X)$	
		change	no change
$P(X)$	change	(1) important & observable	(2) observable
	no change	(3) important & unobservable	(4) no drift

Table 1. When concept drift detection is observable and important (affects the decision rules), adapted from [105]

A handful of algorithms focus on drift detection on delayed, partially labelled or unlabelled data streams. Examples include SUN [94] and the method from Klinkenberg [57] based on support vector machines. The former uses a clustering algorithm to produce ‘concept clusters’ at the leaves of an incremental decision tree, and drifts are identified according to the deviation between history concept clusters and the current clusters. Haque et al. [47] proposed an approach that dynamically determines the boundaries of windows by detecting significant changes in classifier confidence scores using a limited number of labelled instances. This approach is integrated with a classifier to form a complete SSL framework that utilises dynamic chunk boundaries to address concept drift and concept evolution. Cerqueira et al. [19] presents an unsupervised drift detector based on a teacher-student approach, where a predictive model (teacher) is built using an initial batch of labelled training data. The teacher’s predictions are used as class labels to train a surrogate model (student), which will learn to mimic the teacher. A drift detection algorithm is used to identify variations in the mimicking error of the student. The hypothesis is that if the mimicking error increases, then it means that a concept drift has occurred.

5 FAIR COMPARATIVE ANALYSIS

Like other machine learning methods, SSL methods should be evaluated in a realistic process to verify their capabilities while considering other applicable methods. Van Engelen and Hoos [86] observed that additional factors have to be considered during evaluation compared to fully supervised learning scenarios.

First and foremost, the question arises of whether the use of a semi-supervised approach yields performance gains compared to supervised methods [71]. Furthermore, a comparison of an SSL method of interest with other SSL methods is required. Similarly to other machine learning methods, the selection of the data for which predictions are evaluated has to be followed by calculating performance measures. Interestingly, due to the latency of ground truth labels, multiple predictions made for a single instance at different times before the arrival of its true label may be considered in the evaluation [45]. The objective of this section is to address the unique aspects of the evaluation of semi-supervised stream mining methods while taking into account non-negligible delays in the availability of ground truth labels.

5.1 Evaluation of machine learning methods

Evaluation of machine learning methods necessitates applying an appropriate combination of error estimation methods, which include but are not limited to the selection of data used for model development and testing. In addition, performance measures matching domain needs have to be selected. An extensive study on evaluating learning algorithms by Japkowicz and Shah [52] has already provided in-depth coverage of performance measures for classification, how they should be calculated and how different candidate methods should be compared. However, the latter work is focused on the fully supervised batch learning paradigm, i.e. the use of fully labelled data sets rather than partially labelled data streams for both learning and evaluation purposes.

In recent years, there has been an increasing number of studies on evaluating learning algorithms under other settings than fully supervised batch learning. In particular, Oliver et al. [71] examined the impact of various factors such as hyperparameter tuning, class imbalance, and volume of unlabelled data on the evaluation of semi-supervised deep learning methods and revealed how important careful consideration and documenting of these and other factors could be. However, in this major study on evaluating SSL methods, the focus was limited to batch learning. As far as the evaluation of machine learning methods applied to data streams is concerned, new procedures aimed at how data streams should be used for online learning and evaluation purposes were developed. A summary of the most popular of these methods was made in [11]. Importantly, these methods do not consider label latency. Besides new performance measures such as measures reflecting temporal dependencies in the data [13], intermediate performance measures capturing the performance of multiple predictions made over time for a single instance [46] were proposed. A particular limitation of these works from the perspective of our study is that these measures assume a fully supervised setting.

While each of these studies refers to at least one of the aspects of the evaluation of SSL methods for delayed partially labelled data streams, to the best of our knowledge, only selected aspects of the evaluation of SSL methods were covered in them. There are relatively few studies on SSL under delayed labelling settings (as defined in Section 2), which we refer to below.

Taking into account the complexity of the evaluation of machine learning methods [52], in this section we aim to focus on the unique aspects of the evaluation of SSL methods for partly labelled delayed data streams, rather than provide a holistic view of all aspects of the evaluation of stream mining methods. Our intention is not to repeat these aspects which are common to the evaluation of other stream mining techniques, but to pay particular attention to how the evaluation relevant for this study differs from the evaluation suggested for related tasks. Furthermore, we conclude this section with a unified view of the recommendations we make.

All data streams considered in this section are assumed to be delayed data streams. Furthermore, let us observe that while data streams are infinite by definition, the evaluation of learning methods inevitably relies on stream sections $\Psi[T_{\min}, T_{\max}]$. In particular, the period $[T_{\min}, T_{\max}]$ may be the period of a single concept or may span multiple concepts.

The remainder of this section is organised as follows. First, an overview of the evaluation processes and measures applicable to stream mining methods is presented. This is followed by a discussion of the way standard evaluation practices can be tailored to enable a comparison of SSL methods with fully supervised methods. Next, the role of data streams used in the evaluation and other key factors influencing the evaluation outcomes are discussed. This provides the basis for the unified evaluation process proposed in this study, which concludes this section.

5.2 Evaluation of stream mining models

5.2.1 Evaluation process. The evaluation of machine learning methods for data streams is often focused on “**how**” past predictions influence the current model’s predictive performance measure. The most straightforward approach is to evaluate predictions in a **test-then-train** fashion; as the name implies, each instance is first used for testing and then for training. The predictive performance of the learning algorithm in test-then-train represents the average value of all the instances assessed up to that point in time [11, 45]. In test-then-train evaluation, the latest predictive performance estimation is influenced by all previous predictions. This characteristic is desirable when the goal is to understand how well the model performs up to a given point in time. However, it may not give a clear view of how well the model is performing at a given period of the stream. For example, several recent incorrect predictions (perhaps caused by a concept drift) may be shadowed by thousands of old correct predictions.

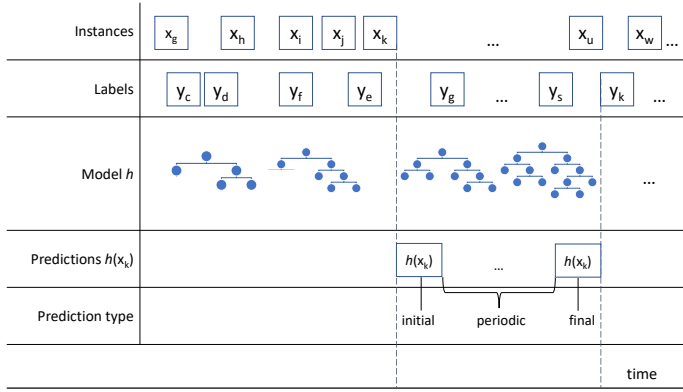


Fig. 7. The types of predicted labels under delayed labelling setting. Based on Grzenda et al. [45]

One approach to avoid undesired influence from previous predictions is to perform a **periodic holdout** evaluation, where training and testing are interleaved using predefined windows, such that one window is used for training and the following used only for testing. This approach can be considered wasteful as labelled data that could be used for training (after testing) is discarded. Thus, an alternative approach is to use a **prequential evaluation** [11], where test-then-train is applied to a sliding window, or a fading factor is used, to forget old predictions. A more in-depth discussion of both approaches to prequential evaluation can be found in Gama et al. [37]. Importantly, when data are only partly labelled, the prequential evaluation is still applicable, as the loss function can be calculated only for those observations for which labels y_k are known [37]. A well-accepted approach for increasing confidence in the evaluation results of batch evaluations is cross-validation. The challenge associated with cross-validation on a streaming setting is that it is infeasible to reprocess the whole stream. To cope with this constraint Bifet et al. [9] introduced an approach for cross-validation in an online setting, where the models are trained and tested in parallel on different folds of data. Furthermore, in delayed data streams, predictions are typically made at $T(x_k)$ and can be evaluated only after the corresponding true labels arrive. This means **verification latency** [31] occurs.

Grzenda et al. [45] claim that besides **“how”** predictions affect the predictive performance, it is also key to consider **“when”** labels are made available as part of the evaluation. This leads to the concept of **continuous re-evaluation** introduced in [45] and further explored in [46]. The goal of continuous re-evaluation is to observe if, and how fast models can transform an initial, possibly incorrect prediction made at $T(x_k)$ into a correct prediction before the true label arrives at $T(y_k)$. While waiting for a label y_k to determine whether a prediction was correct or not, the model is incrementally trained with labels from other instances that arrived before $T(y_k)$. These updates to the model may change the initial prediction yield for x_k .

This evaluation is essential to scenarios where evolving predictions are relevant, such as continuously re-assessing whether a recently released application is “malware” or “goodware” until the ground-truth is available [20]. Continuous re-evaluation is the generalisation of the test-then-train approach, as it provides a way of calculating and assessing **initial predictions** made for x_k at $T(x_k)$, possible **periodic predictions** made for the instance before $T(y_k)$, and **final predictions** made at $T(y_k)$ i.e. immediately before using the true label to possibly update a model [45]. Fig. 7 illustrates the way all these types of predictions are produced for the x_k instance. Continuous re-evaluation assumes that for every instance, its true label is available with non-infinite delay [45].

Even in scenarios where evolving predictions are not essential, it might still be useful to also consider intermediary predictive performance as it indicates how fast the model can adapt to predictions given that other labelled data were made available. Finally, one could potentially adapt continuous re-evaluations to partially delayed labelled streamas, the main difference being that only instances where y arrives will be used for assessment.

5.2.2 Performance measures. A fundamental aspect of every evaluation of stream mining methods is the selection of performance measures to be calculated. Similarly to other stream mining settings, apart from measures such as accuracy, measures developed to address the unique aspects of streaming data, including possible temporal dependencies, should be considered. Kappa temporal (κ^+) [13] is of particular importance for many data streams as it compares the performance of a model against the performance of a simple “No Change” model. The “No Change” model always predicts the next y using the previous true label, which causes it to be consistently correct when temporal dependencies are expected. A summary of all the aforementioned measures can be found in Bifet et al. [11]. In addition, when multiple predictions are made for a single instance, the values of measures such as accuracy and kappa can be aggregated into intermediate performance measures [46]. In this way, the performance of both initial and periodic predictions made for an instance x_k before its true label arrival $T(y_k)$ can be assessed [45].

Furthermore, it is essential to note that for the evaluation to be complete, the memory and computational overhead should be reported. In [68], running time and memory allocation were reported for both supervised and semi-supervised technique. Importantly, the evaluation of stream classification methods in [41] revealed that some of the tested methods failed to process data streams even when 200GB of operating memory were made available for these methods. Hence, ideally, CPU and memory consumption should be reported for all evaluated methods, including SSL methods, along with the measures focused on assessing the ability of individual methods to minimise the loss of prediction.

Streaming algorithms are expected to be efficient, so it is reasonable to also assess them in terms of computational resources [11]. On top of that, other key features of the data stream should be reported, such as the proportion of labelled and unlabelled instances, the number of true labels in the latest window, the proportion of individual classes, and others. All these factors may have a substantial impact on the performance of a stream mining method.

5.3 Comparison of semi-supervised vs. supervised methods

5.3.1 Evaluation based on removing some labels. When an SSL method is considered, its merits should be verified through comparison against supervised methods, including methods of possibly lower computational complexity. This should be done by using an appropriate combination of evaluation process and performance measures. Whether other periodic predictions are justified by the domain problem or not, continuous re-evaluation adapted to a partially labelled setting can be used to analyse the performance of just initial predictions, or possibly also final and periodic predictions. As previously mentioned, computational resources are paramount to stream mining algorithms. Hence, when a supervised method yields the same performance as an SSL method, which is achieved at a lower computational overhead, it is natural that the supervised method will be preferred. SSL methods may require more computational resources as they potentially use all incoming instances for training.

Comparison of an SSL method against a fully supervised method can be made in two ways. First of all, some labels can be removed from an initial data stream to provide a delayed and partially labelled data stream processed by an SSL method [47, 59, 68]. We will refer to such a data stream as a *reduced partially labelled data stream*, which we denote by $F_U(\Psi[T_{\min}, T_{\max}], p_u)$. We propose

to generate such a data stream, by removing with probability p_u individual true labels $\{(\cdot, y_k)\}$ from $\Psi[T_{\min}, T_{\max}]$. In this way, in every run of an evaluation process fed with $\Psi[T_{\min}, T_{\max}]$ data, a possibly different reduced partially labelled data stream $F_U(\Psi[T_{\min}, T_{\max}], p_u)$ will be generated and used to evaluate the impact of the reduced number of true labels on the evaluation process. Importantly, this means that each of the originally labelled instances $\{(\mathbf{x}_k, ?)\}$ is converted with probability p_u into an unlabelled one.

Furthermore, let us observe that any fully supervised method will ignore the existence of unlabelled instances. Hence, it will operate on what we call a *reduced fully labelled data stream* $F_L(\Psi[T_{\min}, T_{\max}])$. This stream refers to the one created from the initial stream after removing instances for which no labels have arrived until T_{\max} . In other words, $F_L(\Psi[T_{\min}, T_{\max}])$ neglects the existence of unlabelled instances. Hence, it provides input for fully supervised learning methods.

The practice of removing labels to create partly labelled data sets is frequently present in studies on SSL methods. Van Engelen and Hoos [86] observed that data sets used in research are usually obtained by removing several labels from existing supervised learning data sets. In line with these practices, a comparison of the performance measures attained by a fully supervised method operated on a $F_L(\Psi[T_{\min}, T_{\max}])$ data stream and an SSL method operated on a $F_U(F_L(\Psi[T_{\min}, T_{\max}]), p_u)$ data stream can be made. Such a comparison of SSL methods operating on partly labelled data with supervised methods using fully labelled data streams has been made *inter alia* in [35, 47, 59, 68].

Moreover, the impact of the p_u value on the performance measures of an SSL method should be analysed to provide insight into the way the method responds to varying volumes of labelled and unlabelled data. In particular, a supervised method's performance can be compared with an SSL method's performance operating on a reduced number of labelled instances. It is now well established that some batch SSL algorithms may work well or not depending on the volume of labelled and unlabelled data [71, 86]. Analysing individual methods' performance under varied ratios of labelled and unlabelled instances produced from the same set [71, 94] and diverse data sets with different quantities of labelled and unlabelled data [86] was already recommended to address this phenomenon. Analysing the impact of p_u on individual methods' performance is a way to adapt these findings to the needs of SSL evaluation under streaming scenarios. Le Nguyen et al. [59] presented a summarised analysis considering different ratios (from 90% to 99%) of unlabelled instances for streaming evaluation.

Comparison of an SSL method against a fully supervised method based on removing some true labels is particularly challenging for the SSL method, as it makes the latter method rely on a lower number of labelled instances than the fully supervised method. Still, as shown in Haque et al. [47] and Masud et al. [68], an SSL stream mining method may provide accuracy comparable to or even competitive with a fully supervised technique under such circumstances. Further examples of works reporting that SSL approaches, even in such cases, can yield accuracy comparable to purely supervised learning are provided in the study Oliver et al. [71], which is focused on the evaluation of deep SSL methods in a batch setting.

It is important to observe that $F_U(\Psi, p_u)$ can be created from an initially available data stream, which could be either fully or partially labelled data stream. While we propose that an SSL method executed on $F_U(F_L(\Psi[T_{\min}, T_{\max}]), p_u)$ is compared with a fully supervised method $F_L(\Psi[T_{\min}, T_{\max}])$, this does not exclude the use of a partially labelled original stream Ψ . In particular, the SSL method can use both originally unlabelled instances and unlabelled instances caused by the use of F_U . Let us note that constraining SSL methods to make them use only those unlabelled instances which were originally labelled, would not reflect the real needs and opportunities provided by SSL techniques.

5.3.2 Evaluation based on removing unlabelled instances. Another way of comparing the performance of a fully supervised method with the performance of an SSL method is based on removing

unlabelled instances. Unlike the former approach, under this scenario, the initial data stream has to be a partially labelled data stream Ψ , rather than fully labelled. The objective of the evaluation is to verify the merits of using unlabelled instances by comparing results attained on the partially labelled Ψ with the results provided by a fully supervised method on a fully labelled stream $F_L(\Psi)$ i.e. the Ψ stream constrained to fully labelled instances. Indeed, the interest in semi-supervised learning is partly driven by the abundance of unlabelled data combined with scarce labelled data. In such cases removing unlabelled instances is acceptable rather than removing already scarce labels.

Among others, Oliver et al. [71] removed unlabelled instances to verify whether the performance obtained by training a model on $X_L \cup X_U$ (i.e. union of labelled data and unlabelled data) is better than the performance observed on labelled instances X_L alone. Oliver et al. [71] observed that such a baseline is also frequently reported in other studies.

The comparison of SSL methods exploiting both labelled and unlabelled parts of data streams to fully supervised methods which discard unlabelled data was performed in the study proposing a semi-supervised SVM learning framework [99]. The SSL methods proposed in the study outperformed the methods discarding unlabelled data. A related aspect of the impact the growing number of unlabelled training instances used by an SSL method on the overall accuracy was addressed in [68]. The growth in the number of unlabelled training instances used by an SSL method resulted in accuracy improvements. In [59], the cluster-and-label method with pseudo-labeling was compared with its version without pseudo-labelling and found to outperform it for a number of synthetic and real data streams. This kind of comparison is one more example of investigating the benefits arising from including unlabelled training instances. Interestingly, the original data streams were fully labelled. Hence, Le Nguyen et al. [59] illustrate the case of removing some labels from a fully labelled data stream first and considering or not unlabelled instances in pseudo-labelling next.

5.3.3 The benefits of both evaluation scenarios. It follows from the related works that by a “fully supervised baseline” two different baselines are meant in different studies. The first of them, which we refer to as the *label removal scenario* relies on removing some labels from an initially fully labelled data stream. The other approach i.e. the *unlabelled instance removal scenario* is to use a partially labelled data stream to remove unlabelled instances from it. Both scenarios can be used to develop a fully labelled data stream and use it for supervised learning. Using two different approaches can in fact be justified by the merits of each of them, which we summarise in Table 2.

In the *label removal scenario*, the fully supervised approach is assumed to yield an upper bound for the SSL performance, as the SSL model relies on a lower number of labelled instances. Moreover, the SSL model can access the same number of instances as its fully supervised counterpart, but true labels of some of these instances are not available for the SSL method. Hence, the benefits of using unlabelled instances are not expected to surpass the benefits of using the same instances provided with true labels. On the other hand, the *unlabelled instance removal scenario* means that all the labelled instances are available for both the fully supervised and SSL methods. In addition, an SSL method may benefit from unlabelled instances available for this method only. This means that in the case of the unlabelled instance removal scenario, we intuitively expect the performance of fully supervised models to be worse than the performance of SSL models. Let us observe that in the case of both expected upper and lower bound, the performance of the supervised method can be in some cases unexpectedly, respectively, worse and better than the performance of an SSL method. These dependencies are not rigid due to multiple factors, some of which are: a) possible superiority of the supervised part of an SSL method over the fully supervised method, b) noisy true labels disturbing fully supervised learning, c) the use of data augmentation techniques by some of the SSL methods increasing their chances of surpassing supervised methods despite the lower number of labelled instances.

Table 2. Two scenarios of obtaining reference fully supervised baselines for the comparison of the performance of SSL methods (SSM) with fully supervised methods (FSM)

Scenario	Original method	Modified stream and method	Key features	Sample works
Label removal	FSM	$F_U(\Psi[T_{\min}, T_{\max}], p_u)$ SSM	<ul style="list-style-type: none"> • suitable when the source data stream is fully labelled • the number of labelled instances available for SSL methods lower than for fully supervised setting • the same number of instances available for both fully supervised and SSL methods • fully supervised setting provides an expected upper bound for SSL model performance 	[35, 47, 59, 68, 86, 94]
Unlabelled instance removal	SSM	$F_L(\Psi[T_{\min}, T_{\max}])$ FSM	<ul style="list-style-type: none"> • suitable for evaluation made for real partially labelled data streams, for which complete labelling is not feasible • the same number of labelled instances available for both fully supervised and SSL methods • fully supervised setting provides an expected lower bound for SSL model performance 	[59, 71, 99]

To sum up, we propose both scenarios described above to be applied simultaneously to benchmark SSL methods against fully supervised methods. This is because comparison under a lower number of labelled instances and the same number of labelled instances accompanied by unlabelled instances are inherently different, and each provides additional insight into the functioning of SSL methods.

5.4 Reference data streams

5.4.1 The selection of data used for comparative analysis. The evaluation of individual stream mining methods under consideration should be made on a benchmark set of data streams. Similarly to other stream mining studies, we propose that evaluation performed with real data streams should be accompanied by evaluation performed with synthetic data streams including the streams for which predefined concept drift events, including the periods affected by gradual concept drift, can be defined. The evaluation of both synthetic and real data streams is a common practice in works proposing new stream mining methods [12, 41, 42, 99, 99].

By definition, the evaluation requires multiple partly labelled delayed data streams to be included. However, synthetic data streams are typically fully labelled and rely on immediately available labels. This includes synthetic data streams frequently used in the evaluation of stream mining

methods such as Agrawal [45], Hyperplane [45], LED [45, 59], and Random Tree [59]. As proposed in Le Nguyen et al. [59], labels from their instances can be removed with probability p to provide a partly labelled data stream.

In their recent study, Le Nguyen et al. [59] proposed establishing a baseline to evaluate semi-supervised learning methods in data streams. Importantly, this includes the extension of the MOA framework, which enables such evaluation. Even though this proposal does not consider the delayed labelling, but immediate labelling only, it can serve as a starting point for defining a baseline set of delayed data streams and developing software serving evaluation needs. Data streams for which no natural delay exists, including all the synthetic data streams listed above, can be converted to delayed ones by adding a fixed delay [45].

To sum up, some reference data streams can be developed under the label removal scenario from their fully labelled versions, but also from real data streams. In this way, partially labelled data streams can be developed. Next, fully labelled data streams can be developed by applying the unlabelled instance removal scenario to the former streams. As a consequence, the results of such studies can be compared to the studies already made under a fully labelled setting for the original data streams. Such data streams ideally should be accompanied by real partially labelled delayed data streams illustrating the abundance of unlabelled data.

5.4.2 Key aspects of the evaluation process. Let us observe that for the evaluation of SSL methods to be fair, it is important to document all the assumptions and limitations it relies on, but also alternative approaches. Let us first discuss some of the assumptions which may have a potentially significant impact on the interpretation of evaluation results and on the evaluation process needed.

First of all, in some studies, an assumption can be made that the number and distribution of classes in the labelled and unlabelled parts of a data stream are the same. In some tasks, such as binary classification, in which the probability that an instance has no label depends neither on the instance data nor the true label, this approach can be justified. However, as pointed out in Oliver et al. [71], the predictive performance of SSL techniques can degrade drastically when the assumption of equal distribution of classes in the labelled and unlabelled parts of a data set is not met.

Another important aspect of the evaluation is whether hyperparameter tuning aimed at finding the best settings of the individual stream mining methods under comparison has been performed or not. In the case of stream mining methods, unlike in the case of batch methods such as the methods analysed in Oliver et al. [71], a single pass over the data is expected i.e. every instance should be inspected at most once [11]. Hence, in contrast to batch learning, hyperparameter tuning should not rely on multiple runs with the same data. Whether such multiple runs were applied or not, we suggest it should be one of the solutions explicitly documented in the performance analysis of SSL methods applied to data streams. Not surprisingly, comparing a method for which large-scale hyperparameter tuning was carried out before its ultimate performance assessment to methods not tuned in such a way may be misleading.

One more important aspect of the evaluation is sensitivity analysis i.e. the analysis of the impact of varied hyperparameter settings on the performance of the methods under comparison, which is frequently illustrated with the plots showing the impact of parameter settings on the accuracy of the models. The sensitivity analysis of this kind was performed in [35] and [68]. When the method relies on fixed parameter settings not evolved during stream processing, such sensitivity analysis provides an insight into the resilience of the methods to varied, including possibly suboptimal parameter settings.

Furthermore, as noted by Oliver et al. [71], a rarely reported baseline against which to compare SSL approaches is transfer learning. More precisely, the authors suggested comparing results

attained with the SSL method to the results observed when a model was trained with a different (but similar and large) data set to be tuned with the small data set representing the ultimate task in the next stage. Not surprisingly, whether this approach can be applied, depends on the availability of similar data making the training of the initial model possible.

In a stream setting, a similar data stream may be a data stream containing instances collected from a similar sensor. Hence, to fully document possible alternatives, the availability of related data streams which could be used to provide initial stream mining models in a transfer learning setting could be checked. Nevertheless, transfer learning is a challenging approach for online scenarios [100].

5.4.3 Key features of data streams. The impact of label latency. Much of the research on supervised learning for delayed data streams had focused on evaluating predictions made for the instances when they were received from a data stream Ψ i.e. initial predictions. In contrast to the works considering label latency, immediate labelling studies assume that an instance's true label is available immediately after this instance. In such cases, the test-then-train approach is frequently applied. This approach, when adopted to delayed labelling, suggests evaluating final predictions i.e. predictions made for the instances immediately before the arrival of the true labels of these instances.

Considering the needs of the evaluation of SSL methods for delayed data streams, let us observe that it should be focused on initial predictions. However, final predictions should also be included to reveal to what extent models evolve and predictions change in turn whilst waiting for true labels [45, 46].

Class label distribution. One more factor that may have a substantial impact on the performance of SSL methods is the presence of instances related to new classes (concept evolution [11]) in the unlabelled data. Oliver et al. [71] showed that in such cases performance can even be degraded compared to not using any unlabelled data at all. In the case of stream mining methods, the ability to deal with novel classes is frequently assumed. Still, whether and if so how many examples of different classes were present in the form of both labelled and unlabelled instances should be documented in the evaluation of individual methods, both fully supervised and semi-supervised.

Volume of unlabelled data. A good match between the data characteristics and the SSL method biases is required to allow improvement by leveraging unlabelled data. On top of that, SSL methods' merits depend on the volume of unlabelled data available to them. A larger volume of unlabelled data can raise the confidence of such methods, for example, those that rely on ensuring that the decision boundary must pass through low-density areas in the input space [7]. Hence, for the evaluation of an SSL method to be complete, an analysis of the impact of the volume of unlabelled data on the performance of the method should be included. Not surprisingly, when the number of unlabelled instances available for the method is minimal, the performance advantage of SSL over fully supervised methods is unlikely to be significant, if any. Some of the works made in the immediate labelling setting have already analysed the impact of a growing number of unlabelled instances on the performance of the methods they proposed [68].

Oliver et al. [71] showed that the sensitivity of different SSL approaches to the amount of labelled and unlabelled data significantly varies. This finding, while observed for batch learning techniques, indicates that accuracy gains arising from the use of unlabelled data must be interpreted with caution. In particular, the superiority of a SSL method following from experiments not including an evaluation of the impact of the proportion of unlabelled data on the performance of this method should be considered to be true for the quantity of labelled and unlabelled data used in the experiment, rather than for other ratios of unlabelled data.

The influence of the number of unlabelled instances can be observed for both real and synthetic streams. Even though the number of unlabelled instances for individual concepts may be strictly constrained by the data, it is possible to use only a subset of the available unlabelled instances in some of the runs of the method under consideration. Hence, individual method runs may rely on a different number of unlabelled instances and possibly different proportions of labelled vs. unlabelled instances.

5.4.4 Additional evaluation of active learning methods. Some SSL methods rely on active learning (AL). Active learning can be used not only to increase the availability of labelled data, but also to contribute to model adaptation to concept drift. A method relying on active learning to obtain additional labelled instances when concept drift is detected was proposed *inter alia* in [34]. When active learning becomes a part of an SSL method, additional evaluation of the method has to be considered. It is important to note that comparison of active learning models vs. models trained on initially available labelled data should only take into account the cost of obtaining additional labels by the active learning method. In particular, in the case of the active learning method, the superiority of the method in terms of the performance of its predictions is not sufficient to confirm the actual improvement offered by the method over a method not requesting extra labels from an oracle such as a human expert.

The evaluation of the cost of obtaining extra labels from an oracle can be made a) in an on-line manner to control the number of requests for additional labels, in order not to jeopardise the benefits of the SSL/AL method and b) in an off-line manner i.e., calculated after the active learning method has been executed. When active learning methods are considered, which is problem-dependent, the problem of selecting the best method for a data stream or set of data streams can be defined as a multi-objective optimisation problem, as both the performance of the method and the cost of obtaining extra labels have to be considered.

Recent studies on the evaluation of stream mining methods for delayed data streams [45, 46] reported that the accuracy of initial predictions is typically lower than the accuracy of final predictions. This phenomenon was observed both for synthetic and real data streams. This difference between the accuracy of initial and final predictions was even more significant for concept drifting data streams. This is because a more recent model benefits from a larger number of labelled instances, possibly reflecting recent changes in the underlying process [45]. Therefore, for the evaluation of active learning approaches to be realistic, not only the cost of obtaining additional labels from an oracle but also the latency with which these labels are available should be considered. This latency cannot be neglected, especially when a human expert is assumed to be involved in the labelling process. If this latency of obtaining additional labels were neglected, the evolution of a model benefiting from these labels would be assumed to be faster than actually possible. As a consequence, taking into account the results of the aforementioned studies, the value of performance measures reported for active learning while not considering labelling latency could potentially be unrealistically superior to the measure values reported for other methods.

To sum up, when a SSL method relies on the active learning paradigm, a recommendation can be made to both report the cost of obtaining extra labels and consider in the evaluation of the method the latency of obtaining additional labels from the method.

5.5 Unified fair evaluation

Taking into account all the aforementioned aspects of the evaluation of SSL methods applied to delayed partially labelled data streams, let us propose Alg. 1 for such evaluation. Importantly, the algorithm aims to show logical data flow rather than its physical implementation. Similarly to the seminal work on Hoeffding trees [33], additional measures can be applied at the implementation

stage to reduce the computational load and storage needs of Alg. 1, some of which are outlined below. The input for the algorithm is the set of reference data streams, which are expected to include both real and synthetic data streams. The algorithm starts by determining dependent data streams. As discussed above, any evaluation is constrained to a certain time period and the set of instances and labels from this period.

In Alg. 1, two categories of Ψ_q streams can be used i.e. fully labelled data stream as used in the label removal scenario, or a partially labelled data stream, used as an input in the unlabelled instance removal scenario. We suggest that both cases can be unified i.e. in both cases, the evaluation can include the input data stream and its labelled part only. Moreover, a particular disadvantage of comparing the performance of an SSL method observed on partially labelled Ψ_s with the performance of a fully supervised method applied to $F_L(\Psi_s)$ is the fact that the performance of both methods is analysed for only one proportion of labelled and unlabelled instances - already present in the input Ψ_s stream. Hence, we propose testing the impact of removing some of the labels on both methods, including the unlabelled instance removal scenario. As a consequence, for every Ψ_q data stream, two categories of fully labelled data streams i.e. Ψ_{UFS} and Ψ_{LFS} , providing an expected upper bound and lower bound respectively for the SSL method are developed. In the case of Ψ_{LFS} streams, the number of such streams matches the number of different p_u settings controlling the number of removed true instances. The partially labelled data streams are used to evaluate SSL methods M_{SSL} , while the fully labelled data streams are used to evaluate M_{FS} methods. By a method a combination of stream mining method and its hyperparameter settings is meant. In this way, sensitivity analysis of individual methods can be performed. Ideally, both real and synthetic data streams including the streams with known presence of concept drift should be represented in the reference stream sections.

As far as the main instance loop operating on the instances of a single stream is concerned, let us emphasise that we take into account the initial predictions i.e. predictions made at the time of receiving instance data, final predictions i.e. predictions made at the time of receiving a true label, and periodic predictions. Furthermore, if active learning methods are included in the evaluation, additional labels received on request from an oracle can be included and possibly used to update a model at the time of actually having them available. Last, but not least, we propose these labels to be used for updating a model, but not for updating performance indicators, as the performance indicators should rely entirely on the performance observed on input instances. This is because the distribution of instances for which additional labels are requested is not likely to match the distribution of x examples. Furthermore, by performance indicators both the indicators aggregating the similarity of predicted and true labels, including accuracy, κ , κ^+ , and intermediate performance measures [46] and indicators revealing resource consumption, such as computation time and memory use are meant. In the former case, the assessment of initial, periodic and final prediction may reveal varied abilities of individual methods to evolve the models before true label arrival. In parallel, stream statistics including label latency histograms, distribution of classes, and the volume of labelled and unlabelled data can be collected.

As far as implementation aspects are concerned, all the streams present in Ω_q sets can be developed in parallel. In particular, multiple runs for every p_u and stream mining method combination can also be executed in parallel. Furthermore, all dependent streams can be gradually produced and processed in parallel without the need to store all their instances and labels. In contrast, every time a new instance or a label of an instance arrives, it may or may not, depending on whether it is included in a dependent stream, become a part of the dependent stream and be processed in the instance-based loop of the algorithm for this dependent stream.

Input: $\{\Psi_q[T_{\min}(q), T_{\max}(q)] : q = 1, \dots, Q\}$ - reference stream sections, $\{p_i : i = 1, \dots, P\}$ - label removal probabilities, R - the number of runs, M_{SSL} and M_{FS} - semi-supervised and fully supervised stream mining methods to evaluate, respectively

Data: $\mathcal{S}_1, \mathcal{S}_2, \dots$ - data stream; L - list of examples $(\{(x_k, ?)\}, \dots)$, containing the data of instances awaiting their true labels, $P(k)$ - list of timestamped predictions made for $\mathcal{S}_i = \{(x_k, ?)\}$, h_i - the prediction model after processing i instances, $M(\Psi)$ - applicable methods i.e. M_{SSL} and M_{FS} for partly and fully labelled Ψ , respectively

begin

```

for  $q = 1, \dots, Q$  do
   $\Psi_s = \Psi_q[T_{\min}(q), T_{\max}(q)]; \Psi_{\text{UFS}} = F_L(\Psi_s);$ 
   $\Omega_q = \{\Psi_s, \Psi_{\text{UFS}}\};$ 
  for  $p_u \in \{p_i : i = 1, \dots, P\}$  do
    for  $r=1, \dots, R$  do
       $\Psi_{\text{SSL}} = F_U(\Psi_{\text{UFS}}, p_u); \Psi_{\text{LFS}} = F_L(\Psi_{\text{SSL}});$ 
       $\Omega_q = \Omega_q \cup \{\Psi_{\text{SSL}}, \Psi_{\text{LFS}}\};$ 
    end
  end
  for  $\Psi \in \Omega_q$  do
    for  $M \in M(\Psi)$  do
       $h_1 = \phi; L = \phi; P = \phi;$ 
      /* Instance loop */
      for  $i = 1, \dots$  do
         $\mathcal{S}_i = \text{fetchNext}(\Psi);$ 
        /* New unlabelled instance arrived */
        if  $\mathcal{S}_i = \{(x_k, ?)\}$  then
           $L.\text{add}(\{(x_k, ?)\});$ 
          /* obtain first time prediction */
           $P(k).\text{addFirst}(h_i(x_k), t(\mathcal{S}_i));$ 
           $h_{i+1} = \text{trainSSL}(M, h_i, \{(x_k, ?)\});$ 
        else
          /*  $\mathcal{S}_i = \{(x_k, y_k)\}$ , i.e. a true label arrived */
          /* if the label is a delayed label */
          if  $c(\mathcal{S}_i) = \text{DELAYED}$  then
            /* obtain final prediction */
             $P(k).\text{addFinal}(h_i(x_k), t(\mathcal{S}_i));$ 
            /* Calculate performance measures and stream statistics */
             $\text{updPerformance}(P(k), y_k, t(\mathcal{S}_i)); L.\text{remove}(k);$ 
            /* generate new periodic predictions for instances awaiting true labels */
             $L.\text{generateNewPredictions}(h_i);$ 
          end
          /* Update the model with SSL method for SSL streams or FSM method for
            UFS and LFS streams, based on delayed and active learning labels */
           $h_{i+1} = \text{train}(M, h_i, \{(x_k, y_k)\});$ 
        end
      end
    end
  end
end

```

end

Algorithm 1: Evaluation of semi-supervised methods under delayed labelling setting

6 CONCLUSION AND PERSPECTIVES

In this paper, we discussed semi-supervised learning from the perspective of delayed partially labelled data streams. This setting is a realistic representation of several real-world problems involving the application of machine learning for data streams. We present several aspects of SSL in this context, including: related problems; learning guarantees; classic batch methods and their online counter parts; and fair evaluation of SSL methods.

SSL methods have been explored (and continue to be explored) with varying levels of success in the batch setting, therefore it is worth reflecting upon how this success can be reflected in the settings of partially-labelled data streams. An explosive area to look at is that of deep learning where the interest in SSL techniques has increased rapidly in the community [86]. SSL is particularly relevant to these models which are notoriously data hungry, and often not enough labeled data is available. Data augmentation can also be useful in this setting: augmenting the training data with new examples artificially created from existing ones. A promising, albeit challenging, venue of research is the application of such techniques to unlabelled examples; some have already been proposed [96]. Extending such methods to data streams poses several challenges, namely that of computational complexity. Augmentation strategies defined for batch data already require costly operations, such as SMOTE [22] which was originally used to address class imbalance problems, but can also be used to augment the training data. Recently, Bernardo et al. [8] proposed a meta-strategy named continuous-SMOTE that perform the oversampling step only on a recent subset of instances.

The application of deep neural networks to streaming is controversial. Although gradient descent is naturally instance-incremental (or, more commonly, minibatch-incremental), a large amount of labeled data is required, and often the convergence of gradient descent algorithms require many passes over the data, which is not possible a streaming setting. There are indeed approaches to apply deep learning models to streams [3, 74], however the advancements in the field are lagging behind in comparison to the batch community.

Transfer learning is an increasingly popular and powerful technique to improve the performance of learning one concept, given that an earlier concept has already been learned. There is an obvious connection here to learning in a concept-drifting data stream, but also more generally when trying to stream one concept making use of another existing already-learned task. Zhao et al. [100] introduced a framework for online transfer learning, including algorithms to tackle domains of common and different feature spaces, as well as an algorithm to address concept-drifting streams. More recently, Wu et al. [93] explored multiple homogeneous and heterogeneous sources for online transfer learning, however concept drift was not taken into account for the scenarios studied. We remark upon this topic because SSL can be leveraged for the application of transfer learning, since more accurate models can be produced even when labelled data is scarce in the target domain. In Section 5, we mentioned that comparing the predictive performance of SSL methods against fully supervised methods that benefit from transfer learning leads to a realistic benchmark. However, this practice is yet to become popular as transfer learning is not widely used in streaming in comparison to batch settings.

Also most SSL tasks are focussed on classification, several SSL techniques can be adapted to regression [86], such as self-training, co-training, and learning by disagreement. One possible pathway for future research is the application of such SSL techniques in regression problems.

ACKNOWLEDGMENTS

Maciej Grzenda: The project was funded by POB Research Centre for Artificial Intelligence and Robotics of Warsaw University of Technology within the Excellence Initiative Program - Research University (ID-UB).

REFERENCES

- [1] Zahraa S Abdallah, Mohamed Medhat Gaber, Bala Srinivasan, and Shonali Krishnaswamy. 2016. Anynovel: detection of novel concepts in evolving data streams. *Evolving Systems* 7, 2 (2016), 73–93.
- [2] Charu C Aggarwal, Jiawei Han, Jianyong Wang, and Philip S Yu. 2003. A framework for clustering evolving data streams. In *International Conference on Very Large Data Bases (VLDB)*. 81–92.
- [3] Andri Ashfahani and Mahardhika Pratama. 2019. Autonomous deep learning: Continual learning approach for dynamic environments. In *Proceedings of the 2019 SIAM International Conference on Data Mining*. SIAM, 666–674.
- [4] Manuel Baena-García, José del Campo-Ávila, Raúl Fidalgo, Albert Bifet, Ricard Gavaldà, and Rafael Morales-Bueno. 2006. Early drift detection method. (2006).
- [5] Maroua Bahri, Albert Bifet, João Gama, Heitor Murilo Gomes, and Silviu Maniu. 2021. Data stream analysis: Foundations, major tasks and tools. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* (2021).
- [6] Yoshua Bengio, Aaron Courville, and Pascal Vincent. 2013. Representation Learning: A Review and New Perspectives. *IEEE Trans. Pattern Anal. Mach. Intell.* 35, 8 (2013), 1798–1828.
- [7] Kristin P Bennett and Ayhan Demiriz. 1999. Semi-supervised support vector machines. In *Advances in Neural Information processing systems*. 368–374.
- [8] Alessio Bernardo, Heitor Murilo Gomes, Jacob Montiel, Bernhard Pfahringer, Albert Bifet, and Emanuele Della Valle. 2020. C-SMOTE: Continuous Synthetic Minority Oversampling for Evolving Data Streams. In *IEEE International Conference on Big Data*.
- [9] Albert Bifet, Gianmarco de Francisci Morales, Jesse Read, Geoff Holmes, and Bernhard Pfahringer. 2015. Efficient Online Evaluation of Big Data Stream Classifiers. In *ACM SIGKDD*. 59–68.
- [10] Albert Bifet and Ricard Gavaldà. 2007. Learning from time-changing data with adaptive windowing. In *SIAM international conference on data mining*. 443–448.
- [11] Albert Bifet, Ricard Gavaldà, Geoff Holmes, and Bernhard Pfahringer. 2018. *Machine Learning for Data Streams: with Practical Examples in MOA*. MIT Press.
- [12] Albert Bifet, Geoff Holmes, and Bernhard Pfahringer. 2010. Leveraging Bagging for Evolving Data Streams. In *PKDD*. 135–150.
- [13] Albert Bifet, Jesse Read, Indrė Žliobaitė, Bernhard Pfahringer, and Geoff Holmes. 2013. Pitfalls in Benchmarking Data Stream Classification and How to Avoid Them. In *Machine Learning and Knowledge Discovery in Databases*. Springer, 465–479.
- [14] Avrim Blum and Tom Mitchell. 1998. Combining labeled and unlabeled data with co-training. In *Conference on Computational learning theory*. 92–100.
- [15] Leo Breiman. 1996. Bagging predictors. *Machine learning* 24, 2 (1996), 123–140.
- [16] Leo Breiman. 2001. Random Forests. *Machine Learning* 45, 1 (2001), 5–32.
- [17] Olivier Cappé and Eric Moulines. 2009. On-line expectation–maximization algorithm for latent data models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 71, 3 (2009), 593–613.
- [18] G. Cassales, H. Gomes, A. Bifet, B. Pfahringer, and H. Senger. 2020. Improving Parallel Performance of Ensemble Learners for Streaming Data Through Data Locality with Mini-Batching. In *IEEE International Conference on High Performance Computing and Communications (HPCC)*.
- [19] Vitor Cerqueira, Heitor Murilo Gomes, and Albert Bifet. 2020. Unsupervised Concept Drift Detection using a Student–Teacher Approach. In *International Conference on Discovery Science*. Springer, 190–204.
- [20] Fabrício Ceschin, Heitor Murilo Gomes, Marcus Botacin, Albert Bifet, Bernhard Pfahringer, Luiz S Oliveira, and André Grégio. 2020. Machine Learning (In) Security: A Stream of Problems. *arXiv preprint arXiv:2010.16045* (2020).
- [21] Olivier Chapelle, Bernhard Scholkopf, and Alexander Zien. 2006. *Semi-Supervised Learning*. MIT Press.
- [22] Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. 2002. SMOTE: synthetic minority over-sampling technique. *Journal of artificial intelligence research* 16 (2002), 321–357.
- [23] Moustapha Cisse, Maruan Al-Shedivat, and Samy Bengio. 2016. ADIOS: Architectures Deep In Output Space. In *Proceedings of The 33rd International Conference on Machine Learning*, Vol. 48. PMLR, 2770–2779.
- [24] Lucas de Carvalho Pagliosa and Rodrigo Fernandes de Mello. 2017. Applying a kernel function on time-dependent data to provide supervised-learning guarantees. *Expert Syst. Appl.* 71 (2017), 216–229.
- [25] Rodrigo Fernandes de Mello. 2019. On the Shattering Coefficient of Supervised Learning Algorithms. *arXiv preprint arXiv:1911.05461* (2019).
- [26] Rodrigo F de Mello, Chaitanya Manapragada, and Albert Bifet. 2019. Measuring the Shattering coefficient of Decision Tree models. *Expert Systems with Applications* 137 (2019), 443–452.
- [27] Rodrigo Fernandes de Mello, Yule Vaz, Carlos Henrique Grossi Ferreira, and Albert Bifet. 2019. On learning guarantees to unsupervised concept drift detection on data streams. *Expert Syst. Appl.* 117 (2019), 90–102.
- [28] Arthur P Dempster, Nan M Laird, and Donald B Rubin. 1977. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)* 39, 1 (1977), 1–22.

- [29] Luc Devroye, László Györfi, and Gábor Lugosi. 1996. *A Probabilistic Theory of Pattern Recognition*. Stochastic Modelling and Applied Probability, Vol. 31. Springer. 1–638 pages.
- [30] Gregory Ditzler, Gail Rosen, and Robi Polikar. 2012. Transductive learning algorithms for nonstationary environments. In *The 2012 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 1–8.
- [31] G. Ditzler, M. Roveri, C. Alippi, and R. Polikar. 2015. Learning in Nonstationary Environments: A Survey. *IEEE Computational Intelligence Magazine* 10, 4 (2015), 12–25.
- [32] Carlotta Domeniconi and Dimitrios Gunopulos. 2001. Incremental support vector machine construction. In *Proceedings 2001 IEEE International Conference on Data Mining*. IEEE, 589–592.
- [33] Pedro Domingos and Geoff Hulten. 2000. Mining high-speed data streams. In *ACM SIGKDD*. 71–80.
- [34] Denis Moreira dos Reis, Peter Flach, Stan Matwin, and Gustavo Batista. 2016. Fast Unsupervised Online Drift Detection Using Incremental Kolmogorov-Smirnov Test. In *ACM SIGKDD*. 1545–1554.
- [35] K. B. Dyer, R. Capo, and R. Polikar. 2014. COMPOSE: A Semisupervised Learning Framework for Initially Labeled Nonstationary Streaming Data. *IEEE Transactions on Neural Networks and Learning Systems* 25, 1 (2014), 12–26.
- [36] Mohamed Medhat Gaber, Arkady Zaslavsky, and Shonali Krishnaswamy. 2005. Mining data streams: a review. *ACM Sigmod Record* 34, 2 (2005), 18–26.
- [37] João Gama, Raquel Sebastião, and Pedro Pereira Rodrigues. 2013. On evaluating stream learning algorithms. *Machine Learning* 90, 3 (01 Mar 2013), 317–346.
- [38] João Gama, Indrè Žliobaitė, Albert Bifet, Mykola Pechenizkiy, and Abdelhamid Bouchachia. 2014. A survey on concept drift adaptation. *ACM computing surveys (CSUR)* 46, 4 (2014), 44.
- [39] Alex Gammerman, Volodya Vovk, and Vladimir Vapnik. 1998. Learning by transduction. *Procs of the Fourteenth Conference Uncertainty in Artificial Intelligence* (1998).
- [40] Heitor Murilo Gomes, Jean Paul Barddal, Fabrício Enembreck, and Albert Bifet. 2017. A Survey on Ensemble Learning for Data Stream Classification. *Comput. Surveys* 50, 2, Article 23 (2017), 36 pages. <https://doi.org/10.1145/3054925>
- [41] Heitor Murilo Gomes, Albert Bifet, Jesse Read, Jean Paul Barddal, Fabrício Enembreck, Bernhard Pfahringer, Geoff Holmes, and Talel Abdesslem. 2017. Adaptive random forests for evolving data stream classification. *Machine Learning* 106, 9-10 (2017), 1469–1495.
- [42] Heitor Murilo Gomes, Jesse Read, and Albert Bifet. 2019. Streaming Random Patches for Evolving Data Stream Classification. In *IEEE International Conference on Data Mining*. IEEE.
- [43] Heitor Murilo Gomes, Jesse Read, Albert Bifet, Jean Paul Barddal, and João Gama. 2019. Machine learning for streaming data: state of the art, challenges, and opportunities. *ACM SIGKDD Explorations Newsletter* 21, 2 (2019), 6–22.
- [44] Yves Grandvalet and Y Bengio. 2006. Entropy Regularization. *MIT Press Scholarship Online: August 2013* (09 2006). <https://doi.org/10.7551/mitpress/9780262033589.003.0009>
- [45] Maciej Grzenda, Heitor Murilo Gomes, and Albert Bifet. 2020. Delayed labelling evaluation for data streams. *Data Mining and Knowledge Discovery* 34, 5 (2020), 1237–1266.
- [46] M. Grzenda, H. M. Gomes, and A. Bifet. 2020. Performance measures for evolving predictions under delayed labelling classification. In *International Joint Conference on Neural Networks (IJCNN)*. 1–8.
- [47] Ahsanul Haque, Latifur Khan, and Michael Baron. 2015. Semi Supervised Adaptive Framework for Classifying Evolving Data Stream. In *Advances in Knowledge Discovery and Data Mining*. Springer, 383–394.
- [48] Geoffrey Hinton and Ruslan Salakhutdinov. 2006. Reducing the Dimensionality of Data with Neural Networks. *Science* 313, 5786 (2006), 504 – 507.
- [49] Shen-Shyang Ho and Harry Wechsler. 2004. Learning from data streams via online transduction. *ICDM Workshop* (2004), 45–52.
- [50] Mohammad Javad Hosseini, Ameneh Gholipour, and Hamid Beigy. 2016. An ensemble of cluster-based classifiers for semi-supervised classification of non-stationary data streams. *Knowledge and Information Systems* 46, 3 (2016), 567–597.
- [51] Elena Ikonomovska, João Gama, and Sašo Džeroski. 2011. Learning model trees from evolving data streams. *Data mining and knowledge discovery* 23, 1 (2011), 128–168.
- [52] Nathalie Japkowicz and Mohak Shah. 2011. *Evaluating Learning Algorithms: A Classification Perspective*. Cambridge University Press, USA.
- [53] Shayan Jawed, Josif Grabocka, and Lars Schmidt-Thieme. 2020. Self-supervised Learning for Semi-supervised Time Series Classification. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, 499–511.
- [54] Holger Kantz and Thomas Schreiber. 2003. *Nonlinear Time Series Analysis*. Cambridge University Press, USA.
- [55] Artur Karimov, Erivelton G. Nepomuceno, Aleksandra Tutueva, and Denis Butusov. 2020. Algebraic Method for the Reconstruction of Partially Observed Nonlinear Systems Using Differential and Integral Embedding. *Mathematics* 8, 2 (2020), 300.

- [56] Shirin Khezri, Jafar Tanha, Ali Ahmadi, and Arash Sharifi. 2020. STDS: self-training data streams for mining limited labeled data in non-stationary environment. *Applied Intelligence* (2020), 1–20.
- [57] Ralf Klinkenberg. 2001. Using labeled and unlabeled data to learn drifting concepts. In *IJCAI Workshop on Learning from Temporal and Spatial Data*. 16–24.
- [58] Georg Kreml. 2011. The algorithm APT to classify in concurrence of latency and drift. In *International Symposium on Intelligent Data Analysis*. Springer, 222–233.
- [59] Minh Huong Le Nguyen, Heitor Murilo Gomes, and Albert Bifet. 2019. Semi-supervised Learning over Streaming Data using MOA. In *2019 IEEE International Conference on Big Data (Big Data)*. IEEE, 553–562.
- [60] Ming Li and Zhi-Hua Zhou. 2007. Improve computer-aided diagnosis with machine learning techniques using undiagnosed samples. *IEEE Transactions on Systems, Man, and Cybernetics* 37, 6 (2007), 1088–1098.
- [61] Charles X Ling, Jun Du, and Zhi-Hua Zhou. 2009. When does co-training work in real data?. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, 596–603.
- [62] Beatriz Maceirezo, Diane Litman, and Rebecca Hwa. 2004. Co-training for Predicting Emotions with Spoken Dialogue Data. In *Proceedings of the ACL 2004 on Interactive Poster and Demonstration Sessions (ACLdemo '04)*. Association for Computational Linguistics, Stroudsburg, PA, USA, Article 28. <https://doi.org/10.3115/1219044.1219072>
- [63] Diego Marrón, Eduard Ayguadé, José R. Herrero, Jesse Read, and Albert Bifet. 2017. Low-latency Multi-threaded Ensemble Learning for Dynamic Big Data Streams. In *IEEE International Conference on Big Data*. 223–232.
- [64] Mohammad M Masud, Qing Chen, Jing Gao, Latifur Khan, Jiawei Han, and Bhavani Thuraisingham. 2010. Classification and novel class detection of data streams in a dynamic feature space. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 337–352.
- [65] Mohammad M Masud, Qing Chen, Latifur Khan, Charu Aggarwal, Jing Gao, Jiawei Han, and Bhavani Thuraisingham. 2010. Addressing concept-evolution in concept-drifting data streams. In *2010 IEEE International Conference on Data Mining*. IEEE, 929–934.
- [66] Mohammad M Masud, Jing Gao, Latifur Khan, Jiawei Han, and Bhavani Thuraisingham. 2008. A practical approach to classify evolving data streams: Training with limited amount of labeled data. In *ICDM*. IEEE, 929–934.
- [67] Mohammad M Masud, Jing Gao, Latifur Khan, Jiawei Han, and Bhavani Thuraisingham. 2010. Classification and novel class detection in data streams with active mining. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. 311–324.
- [68] Mohammad M. Masud, Clay Woolam, Jing Gao, Latifur Khan, Jiawei Han, Kevin W. Hamlen, and Nikunj C. Oza. 2012. Facing the reality of data stream classification: coping with scarcity of labeled data. *Knowledge and Information Systems* 33, 1 (01 Oct 2012), 213–244. <https://doi.org/10.1007/s10115-011-0447-8>
- [69] Jacob Montiel, Jesse Read, Albert Bifet, and Talel Abdesslem. 2018. Scalable Model-based Cascaded Imputation of Missing Data. In *PAKDD*. Springer, 64–76.
- [70] Kamal Nigam, Andrew McCallum, and Tom M Mitchell. 2006. Semi-Supervised Text Classification Using EM.
- [71] Avital Oliver, Augustus Odena, Colin A Raffel, Ekin Dogus Cubuk, and Ian Goodfellow. 2018. Realistic Evaluation of Deep Semi-Supervised Learning Algorithms. In *Advances in Neural Information Processing Systems*. 3238–3249.
- [72] N.C. Oza. 2005. Online bagging and boosting. In *IEEE International Conference on Systems, Man and Cybernetics*, Vol. 3. 2340–2345 Vol. 3. <https://doi.org/10.1109/ICSMC.2005.1571498>
- [73] Joshua Plasse and Niall Adams. 2016. Handling delayed labels in temporally evolving data streams. In *IEEE ICDL*. 2416–2424.
- [74] Mahardhika Pratama, Andri Ashfahani, and Abdul Hady. 2019. Weakly supervised deep learning approach in streaming environments. In *2019 IEEE International Conference on Big Data (Big Data)*. IEEE, 1195–1202.
- [75] Jesse Read and Luca Martino. 2020. Probabilistic Regressor Chains with Monte-Carlo Methods. *Neurocomputing* 413 (2020), 471–486.
- [76] Jesse Read, Fernando Perez-Cruz, and Albert Bifet. 2015. Deep Learning in Partially-Labelled Data-Streams. In *SAC 2015: 30th ACM Symposium on Applied Computing*. ACM, 954–959.
- [77] Jesse Read, Ricardo A. Rios, Tatiane Nogueira, and Rodrigo Fernandes de Mello. 2020. Data Streams Are Time Series: Challenging Assumptions. In *Intelligent Systems - 9th Brazilian Conference, BRACIS Proceedings (Lecture Notes in Computer Science)*, Vol. 12320. Springer, 529–543.
- [78] Joan Serrà, Xavier Serra, and Ralph G Andrzejak. 2009. Cross recurrence quantification for cover song identification. *New Journal of Physics* 11, 9 (sep 2009).
- [79] Burr Settles. 2009. *Active Learning Literature Survey*. Computer Sciences Technical Report 1648. University of Wisconsin–Madison.
- [80] Dennis JNJ Soemers, Tim Brys, Kurt Driessens, Mark HM Winands, and Ann Nowé. 2018. Adapting to Concept Drift in Credit Card Transaction Data Streams Using Contextual Bandits and Decision Trees.. In *AAAI*.
- [81] W Nick Street and YongSeog Kim. 2001. A streaming ensemble algorithm (SEA) for large-scale classification. In *ACM SIGKDD international conference on Knowledge discovery and data mining*. 377–382.

- [82] Yu-Yin Sun, Yin Zhang, and Zhi-Hua Zhou. 2010. Multi-label learning with weak label. In *Proceedings of the twenty-fourth AAAI conference on artificial intelligence*. 593–598.
- [83] Floris Takens. 1981. Detecting strange attractors in turbulence. In *Dynamical Systems and Turbulence, Warwick 1980*. Springer, 366–381.
- [84] Grigorios Tsoumakas and Ioannis Katakis. 2007. Multi Label Classification: An Overview. *International Journal of Data Warehousing and Mining* 3, 3 (2007), 1–13.
- [85] Masayuki Ushio and Kazutaka Kawatsu. 2020. *Forecasting Ecological Time Series Using Empirical Dynamic Modeling: A Tutorial for Simplex Projection and S-map*. Springer, 193–213.
- [86] Jesper E Van Engelen and Holger H Hoos. 2020. A survey on semi-supervised learning. *Machine Learning* 109, 2 (2020), 373–440.
- [87] Vladimir Vapnik. 2013. *The nature of statistical learning theory*. Springer science & business media.
- [88] Ulrike von Luxburg and Bernhard Schölkopf. 2011. Statistical Learning Theory: Models, Concepts, and Results. In *Inductive Logic. Handbook of the History of Logic, Vol. 10*. Elsevier, 651–706.
- [89] Willem Waegeman, Krzysztof Dembczyński, and Eyke Hüllermeier. 2019. Multi-target prediction: a unifying view on problems and methods. *Data Mining and Knowledge Discovery* 33, 2 (2019), 293–324.
- [90] Yaqing Wang, Quanming Yao, James T Kwok, and Lionel M Ni. 2020. Generalizing from a few examples: A survey on few-shot learning. *ACM Computing Surveys (CSUR)* 53, 3 (2020), 1–34.
- [91] Geoffrey I Webb, Roy Hyde, Hong Cao, Hai Long Nguyen, and Francois Petitjean. 2016. Characterizing concept drift. *Data Mining and Knowledge Discovery* 30, 4 (2016), 964–994.
- [92] Li Wei and Eamonn Keogh. 2006. Semi-supervised time series classification. In *ACM SIGKDD*. 748–753.
- [93] Qingyao Wu, Hanrui Wu, Xiaoming Zhou, Mingkui Tan, Yonghui Xu, Yuguang Yan, and Tianyong Hao. 2017. Online transfer learning with multiple homogeneous or heterogeneous sources. *IEEE TKDE* 29, 7 (2017), 1494–1507.
- [94] Xindong Wu, Peipei Li, and Xuegang Hu. 2012. Learning from concept drifting data streams with unlabeled data. *Neurocomputing* 92 (2012), 145–155.
- [95] Ming-Kun Xie and Sheng-Jun Huang. 2018. Partial multi-label learning. In *AAAI*.
- [96] Qizhe Xie, Zihang Dai, Eduard Hovy, Minh-Thang Luong, and Quoc V Le. 2019. Unsupervised data augmentation for consistency training. *arXiv preprint arXiv:1904.12848* (2019).
- [97] Chang Xu, Dacheng Tao, and Chao Xu. 2013. A survey on multi-view learning. *arXiv preprint arXiv:1304.5634* (2013).
- [98] David Yarowsky. 1995. Unsupervised Word Sense Disambiguation Rivaling Supervised Methods. In *Annual Meeting on Association for Computational Linguistics*. Association for Computational Linguistics, 189–196.
- [99] Peng Zhang, Xingquan Zhu, and Li Guo. 2009. Mining data streams with labeled and unlabeled training examples. In *ICDM*. IEEE, 627–636.
- [100] Peilin Zhao, Steven CH Hoi, Jialei Wang, and Bin Li. 2014. Online transfer learning. *Artificial Intelligence* 216 (2014), 76–102.
- [101] Zhi-Hua Zhou and Ming Li. 2005. Tri-training: Exploiting unlabeled data using three classifiers. *IEEE Transactions on knowledge and Data Engineering* 17, 11 (2005), 1529–1541.
- [102] Zhi-Hua Zhou and Ming Li. 2010. Semi-supervised learning by disagreement. *Knowledge and Information Systems* 24, 3 (2010), 415–439.
- [103] Xiaojin Zhu. 2008. Semi-Supervised Learning Literature Survey. *Comput Sci, University of Wisconsin-Madison* 2 (07 2008).
- [104] Xingquan Zhu, Peng Zhang, Xiaodong Lin, and Yong Shi. 2007. Active learning from data streams. In *ICDM*. 757–762.
- [105] Indrè Žliobaitė. 2010. Change with delayed labeling: When is it detectable?. In *IEEE ICDMW*. 843–850.
- [106] Indrè Žliobaitė, Albert Bifet, Bernhard Pfahringer, and Geoffrey Holmes. 2013. Active learning with drifting streaming data. *IEEE transactions on neural networks and learning systems* 25, 1 (2013), 27–39.