



Active learning for data streams: a survey

Davide Cacciarelli^{1,2}  · Murat Kulahci^{1,3}

Received: 10 May 2023 / Revised: 13 September 2023 / Accepted: 10 October 2023 /
Published online: 20 November 2023
© The Author(s) 2023

Abstract

Online active learning is a paradigm in machine learning that aims to select the most informative data points to label from a data stream. The problem of minimizing the cost associated with collecting labeled observations has gained a lot of attention in recent years, particularly in real-world applications where data is only available in an unlabeled form. Annotating each observation can be time-consuming and costly, making it difficult to obtain large amounts of labeled data. To overcome this issue, many active learning strategies have been proposed in the last decades, aiming to select the most informative observations for labeling in order to improve the performance of machine learning models. These approaches can be broadly divided into two categories: static pool-based and stream-based active learning. Pool-based active learning involves selecting a subset of observations from a closed pool of unlabeled data, and it has been the focus of many surveys and literature reviews. However, the growing availability of data streams has led to an increase in the number of approaches that focus on online active learning, which involves continuously selecting and labeling observations as they arrive in a stream. This work aims to provide an overview of the most recently proposed approaches for selecting the most informative observations from data streams in real time. We review the various techniques that have been proposed and discuss their strengths and limitations, as well as the challenges and opportunities that exist in this area of research.

Keywords Stream-based active learning · Online active learning · Data streams · Online learning · Unlabeled data · Query strategies · Selective sampling · Concept drift · Experimental design · Bandits

Editor: Joao Gama.

✉ Davide Cacciarelli
dcac@dtu.dk

Murat Kulahci
muku@dtu.dk

¹ Department of Applied Mathematics and Computer Science, Technical University of Denmark, Kgs. Lyngby, Denmark

² Department of Mathematical Sciences, Norwegian University of Science and Technology, Trondheim, Norway

³ Department of Business Administration, Technology and Social Sciences, Luleå University of Technology, Luleå, Sweden

1 Introduction

The deployment of machine learning models in real-world applications is often reliant on the availability of significant amounts of annotated data. While recent advancements in sensor technology have facilitated the collection of larger amounts of data, this data is not always labeled and ready for use in training models. Indeed, the process of obtaining labeled observations for supervised learning models can be cost-prohibitive and time-consuming, as it often requires quality inspections or manual annotation. In such cases, active learning proves to be a valuable strategy to identify the most informative data points for use in training, thereby reducing the overall cost of labeling and improving the performance of the model. Over the years, a plethora of active learning approaches have been proposed in the literature, each with its own benefits and limitations. These approaches seek to strike a balance between the cost of labeling and the quality of the model by selectively choosing the most informative observations for querying. By carefully selecting the most informative observations, active learning helps to minimize the amount of labeled data required and streamlines the learning process, contributing to its overall efficiency.

While several surveys have been published on pool-based active learning (Aggarwal et al., 2014; Settles, 2009; Fu et al., 2013; Kumar & Gupta, 2020), which involves selecting a fixed set of observations from a pool of unlabeled data, the dynamic and sequential nature of many real-world problems often renders these approaches impractical. This has led to growing interest in the online variant of active learning, also referred to as stream-based active learning, which involves continuously selecting and labeling observations as they arrive in a stream, allowing for real-time adaptation to changing data distributions. Lughofer (2017) provided a review of online active learning approaches with a focus on fuzzy models. However, since its publication, numerous other online active learning approaches have been proposed, and to the best of our knowledge, no other surveys have been published to synthesize these developments. Moreover, surveys purely focusing on online learning from data streams (Lu et al., 2018; Tieppo et al., 2022; Lima et al., 2022; Hoi et al., 2021) discuss methods that assume a complete availability of labels, which is not the case in many real-world applications. The aim of this review is to fill this gap by providing a comprehensive overview¹ of the most recently developed query strategies for online active learning. It is worth noting that in certain cases, stream-based active learning is narrowly defined as the act of selecting the most informative observations from a data stream to fit a predictive model. Instead, the act of determining which observations to query while making predictions is referred to as online selective sampling (Hanneke & Yang, 2021). In this work, we cover and examine all the methods that address the crucial problem of selecting the most informative data points to label from a data stream in an online fashion. We will present the techniques that have been proposed so far, discussing their strengths and limitations, as well as the challenges and opportunities that exist in this field. In addition, we will provide an overview of evaluation strategies for online active learning algorithms and highlight some real-world applications. Finally, we will identify potential future research directions in this area.

¹ We conducted a search on SCOPUS and Google Scholar using the following keywords: “on-line active learning”, “online active learning”, “stream-based active learning”, “single pass active learning”, “online selective sampling”, “sequential selective sampling”, and “active learning” combined with “data stream”. Each paper was reviewed individually to determine its relevance to online active learning. We eliminated irrelevant papers and manually added some papers that did not contain these keywords but used online active learning methods or were relevant to our discussion. Additionally, we included related papers that were necessary to understand the bigger picture from the references of the reviewed strategies.

This survey comprehensively explores various facets of active learning, encompassing both theoretical foundations and practical challenges. By delving into this review, we aim to shed light on pertinent research questions, including:

1. *Query strategy.* What sampling strategy should be used to maximize learning efficiency in a streaming context?
2. *Timing of queries.* When and how often should data points be queried to balance learning and resource constraints?
3. *Model updates.* When should predictive models be updated and how can they adapt to changing data distributions and concept drift?
4. *Scalability.* How can active learning methods be made scalable and efficient for high-velocity data streams?
5. *Evaluation.* What are appropriate evaluation metrics for assessing the performance of stream-based active learning algorithms?

The structure of this paper is as follows. In Sect. 2, we provide an overview of active learning, including the main instance selection criteria, an overview of the main active learning scenarios, and the connection between active learning and semi-supervised learning. Section 3 represents the core of the review, with a brief overview of how online active learning approaches have been classified, followed by a detailed description of the state-of-the-art approaches. In Sect. 4, we examine evaluation strategies for online active learning algorithms. Section 5 highlights real-world applications and challenges. Section 6 provides a summary of the most common online active learning methods and highlights potential directions for future research. Finally, Sect. 7 provides conclusions and summarizes the key contributions of the review.

2 Preliminaries on active learning

In supervised learning, we seek to learn a function that can predict the output variable, also known as response, given a set of input variables, also known as covariates. This function is often learned by training a model on a labeled dataset that consists of a large number of input–output pairs. However, obtaining labeled examples is not always straightforward, and it may not be possible or practical to label all the available data. In these cases, active learning can be used to select a subset of the data for labeling in order to improve the performance of the model, when there is a budget constraint on the number of unlabeled observations that can be queried. Indeed, there are many examples of how a classification or regression model can achieve a performance that is similar to what can be achieved when all the labels are available, using only a small fraction of the available observations.

2.1 Instance selection criteria

The main challenge in active learning is deciding which data points to label. There are many strategies for selecting data points in active learning, and most of them can be associated with one of these groups:

- *Uncertainty-based query strategies.* These approaches focus on selecting data points that the model is least confident about, in order to reduce its uncertainty (Lu et al., 2016; Tong & Koller, 2002). When using classification models, the most widely used is the margin-

based query strategy, where data points close to the decision boundary are selected (Roth & Small, 2006; Balcan et al., 2007).

- *Expected error or variance minimization.* These strategies estimate the future error or variance, when a newly labeled example is made available, and try to minimize it directly (Cohn et al., 1996; Roy & McCallum, 2001).
- *Expected model change maximization.* This strategy involves selecting data points that would have the greatest impact on the estimate of the current model parameters if they were labeled and added to the training set (Cai et al., 2013).
- *Disagreement-based query strategies.* These approaches focus on selecting data points where there is disagreement among multiple models or experts (Hanneke, 2014; Wang, 2011; Steve & Liu, 2014; Sheng et al., 2008). One of the most common approaches that use an ensemble of models is query by committee (Seung et al., 1992; Freund et al., 1997; Burbidge et al., 2007), which uses an ensemble of models to identify instances where the models have conflicting predictions.
- *Diversity- and density-based approaches.* These methods exploit the structural information of the instances and try to select data points that are diverse and representative of the overall distribution of the data. One example of this approach is the use of Mahalanobis distance to seek observations that are far from the currently labeled data points (Ge, 2014; Cacciarelli et al., 2022a). Clustering may be applied to label representative data points (Nguyen & Smeulders, 2004; Min et al., 2020; Ienco et al., 2013), and graph-based methods can be employed to explore the structure information of labeled and unlabeled data points (Zhang et al., 2020b) or to build upon the semi-supervised label propagation strategy (Long et al., 2008).
- *Hybrid strategies.* These are active learning algorithms that combine multiple instance selection criteria (Donmez et al., 2007; Huang et al., 2014). For example, by combining margin-based sampling with clustering the learner can select the most uncertain observations within different areas of the input space.

By considering these different strategies, one can select the most appropriate approach for a given problem based on the characteristics of the data and the specific requirements of the application.

2.2 Active learning scenarios

Active learning can be broadly categorized into three macro scenarios, based on how the unlabeled instances are supplied to the learner and then selected to be labeled by an oracle. Regardless of the particular query strategy being employed, these macro scenarios provide a framework for understanding the flow of information and the decision-making steps involved in active learning. These scenarios serve as a high-level categorization of different methods for approaching the active learning problem, each with its own set of advantages and disadvantages depending on the specific use case. Understanding these macro scenarios is crucial for selecting the appropriate active learning technique for a particular problem and for comparing different active learning algorithms. In the next subsections, each of the three macro scenarios will be discussed.

2.2.1 Membership query synthesis active learning

This scenario represents the case when the learner is given complete freedom to ask for the label of any data point belonging to the input space or for a synthetically generated one.




Fig. 1 Membership query synthesis active learning

Some examples of membership query synthesis active learning include image classification, where the learner can generate modified versions of existing images to be labeled, or object detection, where the learner can generate new instances by combining and transforming existing instances. In natural language processing (NLP) tasks such as text classification or sentiment analysis, the learner might generate synthetic examples in the form of sentences or paragraphs that cover a wider range of variations in the language. Also, in speech recognition, the learner might generate synthetic speech samples in different accents, pronunciations, or speaking styles in order to improve the recognition accuracy. However, as highlighted by Baum and Lang (1992) and Settles (2009), the main drawback of this strategy is that it could generate unlabeled examples for which no labels can be associated by a human annotator (e.g., a mixture between a number and a letter). A general flowchart for this scenario is reported in Fig. 1, where the scheme is repeated until a budget constraint on the requested labels is met, or a stopping criterion on the achieved performance is satisfied.

In the context of deep active learning (Ren et al., 2022), the membership query synthesis scenario can be addressed by using generative models. For instance, generative adversarial networks (GANs) have been used to generate additional instances from the input space that may provide more informative labels for the learner (Goodfellow et al., 2014). This can be done by using GANs for data augmentation, as GANs are capable of generating diverse and high-quality instances (Zhu & Bento, 2017). Another approach is to combine the use of variational autoencoders (VAEs) (Kingma & Welling, 2013) and Bayesian data augmentation, as demonstrated by Tran et al. (Tran et al., 2019, 2017). The authors used VAEs to generate instances from the disagreement regions between multiple models, and Bayesian data augmentation to incorporate the uncertainty of the generated instances in the learning process.

2.2.2 Pool-based active learning

Pool-based active learning is one of the most widely studied scenarios in the machine learning literature. The goal is to select the most informative subset of observations from a closed, static set of unlabeled data points. The majority of the proposed pool-based active learning approaches have been developed for classification tasks (Cai et al., 2013), with image classification being a common application in computer vision (Li & Guo, 2013), as manually labeling large image datasets can be a challenging task.

The flowchart in Fig. 2 provides an overview of pool-based active learning sampling schemes, where k represents the number of unlabeled instances whose label is queried at each round. Traditional machine learning models that do not require substantial computational resources to train are typically associated with a choice of k equal to one (Vahdat et al.,

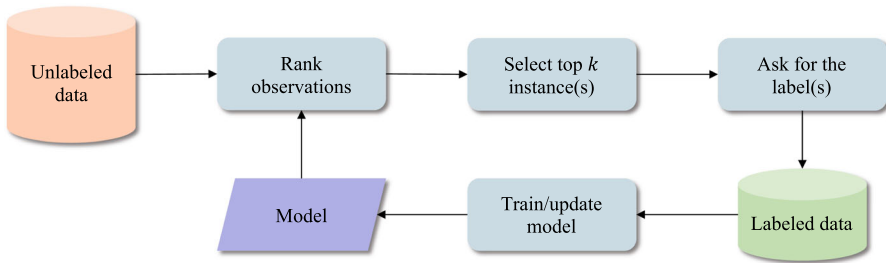


Fig. 2 Pool-based active learning

2019). This allows a timely update of the instance selection criteria, avoiding the redundant labeling of similar data points. However, larger values of k have also been used in practice, such as the analysis performed by Ge (2014) for values ranging from 5 to 30 or the approach used by Cai et al. (2013) to add 3% of the total number of observations to the training set each time. Using a higher k value may be more practical when working with large models, as repeated training can be computationally expensive and challenging. To this extent, batch mode active learning is generally considered to be a more efficient and effective option for image classification or detection tasks compared to the one-by-one query strategy, as the latter can be resource-intensive and time-consuming when working with large neural networks (Ren et al., 2022). This is because re-training the model with just one new data point with high input dimensionality may not result in significant improvement (Ren et al., 2022). In general, the choice of k may be problem- or model-specific, as it represents a trade-off between computational efficiency and the risk of querying redundant labels.

To enhance pool-based active learning, many approaches combine uncertainty-based instance selection criteria with acquisition functions such as entropy (Shannon, 1948; Wu et al., 2022), mutual information (Haussmann et al., 2020), or variation ratio (Schmidt et al., 2020). Entropy is commonly used as an acquisition function in active learning because it provides a way to measure the uncertainty of the model predictions for a given data point. The entropy of a probability distribution is a measure of the amount of disorder or randomness in the distribution. In the context of active learning, the entropy of a model's predicted class probabilities for a data point can be used as a measure of the model's uncertainty about the correct class label for that data point. Acquiring examples with the highest uncertainty is one way to select data points for annotation, but it is not the only way. Mutual information and variation ratio can also be used on the predictions obtained with the current model, in order to seek a diverse set of data points for which the predictions are the most uncertain. For a more comprehensive discussion on pool-based active learning, readers are referred to the surveys (Aggarwal et al., 2014; Settles, 2009; Fu et al., 2013; Kumar & Gupta, 2020).




Fig. 3 Single-pass online active learning

2.2.3 Online active learning

In this type of active learning, we cannot greedily select the most informative observations from a static pool, as the instances are generated in a continuous stream and cannot be stored in their entirety before a decision is made. This is similar to the famous statistical puzzle known as the secretary problem (Freeman, 1983), where a hiring manager must make a hiring decision for each applicant as they are interviewed, without the benefit of seeing all applicants first. In general, online active learning is a crucial scenario for various real-world applications where the ability to make a sampling decision in real-time is of utmost importance. A few examples are:

- *Chemical or manufacturing processes.* In these applications, a learner is tasked with predicting the quality of the final product but may only have a short timeframe to make the sampling decision, to avoid traceability issues, particularly in high-volume production (Schmitt et al., 2013; Lieber et al., 2012). Also, tasks like predictive maintenance and visual inspection might benefit from a real-time selection of new examples to be labeled and included in the training set (Rožanec et al., 2022).
- *Video streaming and clinical trials.* In these cases, a decision must be made on the fly, as users arrive or volunteers appear sequentially, and there may not be enough time to accumulate a pool of potential users or patients (Fowler et al., 2023; Riquelme, 2017).
- *Text classification:* In NLP, online active learning can be used for tasks such as sentiment analysis and spam detection, where the learner continuously learns from new incoming data points which need to be labeled to update the model in real-time and improve accuracy (Kranjc et al., 2015).
- *Fraud detection.* To effectively detect fraudulent activities, the learner must continuously select new examples to label so that it can continuously update its decision-making process (Carcillo et al., 2018, 2017).
- *Online customer service.* Online customer service agents can use online active learning to improve their performance by continuously learning from customer interactions. To do this, the learner must continuously select new examples to label or customer information to obtain, so that it can predict the best response based on past interactions and improve its accuracy over time (Zheng & Padmanabhan, 2006).
- *Marketing.* Online active learning could also be applied in the field of marketing to select informative examples in real-time and continuously optimize customer targeting and personalization (Carnein & Trautmann, 2019; Jamil & Khan, 2016).

One of the defining features of online active learning strategies is their data processing capabilities. Figures 3 and 4 provide a visual representation of the two main approaches;

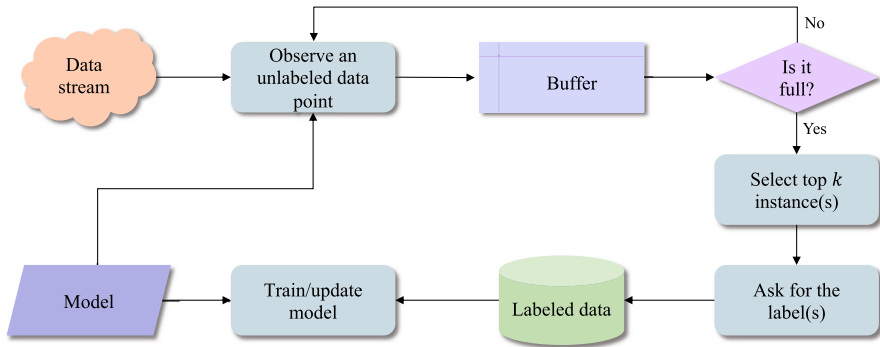


Fig. 4 Window- or batch-based online active learning

single-pass and window-based. Single-pass algorithms observe and evaluate each incoming data point on the fly, whereas window-based algorithms, also referred to as batch-based methods, observe a fixed-size chunk of data at a time. In this approach, the learner evaluates the entire batch of data and selects the top k observations as the most informative ones to be labeled. This approach is referred to as best-out-of-window sampling. The specific value of k and the dimensionality of the buffer can vary based on the storage capabilities of the system and the computational time required to update the model. Window-based methods are useful in situations where data is generated in large quantities and the algorithm does not have a tight constraint on the time available for decision-making. In contrast, single-pass methods are necessary when the algorithm needs to make a decision immediately after observing a specific data point.

Another critical property in the design of an effective online active learning strategy is the assumption made about the data stream distribution. One important difference to consider is whether the data stream is stationary or drifting. A stationary data stream is characterized by a stable data generating process where the statistical properties of the data distribution that remain constant over time. Conversely, a drifting data stream is marked by changing statistical properties of the data distribution over time, potentially due to alterations in the underlying data generating process. The distinction between stationary and drifting data streams is significant because it affects the performance of the active learning strategies. Online active learning strategies that have been developed for stationary data streams may lead to suboptimal performance when applied to drifting data streams. This is because concept drift can alter the scale of the informativeness measure of unlabeled data points or even urge a complete change of the model, with the acquisition of more observations to accommodate the new concept. Therefore, it is important to accurately assess the nature of the data stream distribution in the design of an active learning strategy. A failure to do so can result in a suboptimal performance and a reduced ability to effectively leverage the strengths of active learning. Another important property to consider when designing an active learning strategy is the label delay or verification latency. This refers to the time needed by the oracle to provide the label when it is requested by the learner. In some cases, there may be a delay L in the oracle providing the label after it has been requested. This property must be taken into account when designing a sampling strategy as there may be redundant label requests for similar instances if this issue is not properly addressed. Label delay can be classified into null latency, intermediate latency, or extreme latency (Souza et al., 2018). The case with null latency, or immediate availability of the label upon request, is commonly used in the stream

mining community, but may not be realistic for many practical applications. Extreme latency, where labels are never made available to the learner, is closer to an unsupervised learning task. Intermediate latency assumes a delay $0 < L < \infty$ in the availability of the labels from the oracle.

Finally, the training efficiency of the online active learning algorithms should also be taken into consideration. There are two main training approaches in active learning: incremental training and complete re-training. Incremental training involves updating model parameters with a small batch of new data, without starting the training process from scratch (Polikar et al., 2001; Wu et al., 2019; Shilton et al., 2005; Istrate et al., 2018). This approach allows the model to learn from new data while preserving its existing knowledge. This can be achieved through fine-tuning the model parameters with the new data, or by using techniques such as elastic weight consolidation, which prevent previous knowledge from being erased. Complete re-training, on the other hand, involves training a new model from scratch using the entire labeled data collected so far. This approach discards the previous knowledge of the model and starts anew, which may result in the loss of knowledge learned from previous data. Complete re-training is typically used when the amount of new data is substantial, the previous model is no longer relevant, or when the model architecture needs to be altered. It is important to note that the choice of training approach in online active learning algorithms can have a significant impact on the overall performance and effectiveness of the model.

2.3 Connection between active learning and semi-supervised learning

Semi-supervised learning is a field of research that is closely related to active learning, as both methods are developed to deal with limited labeled data. While active learning aims to minimize the amount of labeled data required to train a model, semi-supervised learning is a technique that trains a model using a combination of labeled and unlabeled data. Active learning can be considered a special case of semi-supervised learning, as it allows the model to actively select which data points it wants to be labeled, rather than relying on a fixed set of labeled data. In the context of online learning, Kulkarni et al. (2016) conducted a study that provided an overview of semi-supervised learning techniques for classifying data streams. These techniques do not address the primary question of active learning, which is *when to query*, but they are useful in exploiting the information contained in the unlabeled data points and in addressing issues related to model update and retraining in limited labeled data environments. It is also worth noting that semi-supervised learning can be used in combination with active learning to improve the data selection strategy. By leveraging the strengths of both methods, it is possible to achieve better performance and more efficient learning compared to using either method alone.

Semi-supervised learning approaches can be distinguished into three categories, unsupervised preprocessing, wrapper methods, and graph-based methods. Unsupervised preprocessing refers to the use of unsupervised learning techniques, such as dimensionality reduction (Cacciarelli & Kulahci, 2023), clustering, or feature extraction, to preprocess the entire dataset, labeled and unlabeled, before it is fed to the supervised model (Frumosu & Kulahci, 2018). The goal is to transform the data into a more useful representation that can be learned more easily by a supervised model and can support the sampling of more informative data points. This strategy can also help reduce the dimensionality of the learning problem, thus improving the model parameter estimation when only a few queries can be made. Related to the online active learning problem, Rožanec et al. (2022) used a pre-trained network to extract salient features from unlabeled images before starting the sampling routine. Simi-

larly, Cacciarelli et al. (2022a) used an autoencoder trained on all the available unlabeled data points to improve the performance of online active learning for linear regression models.

Wrapper methods, on the other hand, use one or more supervised learners that are trained on labeled data and pseudo-labeled unlabeled data. There are two main variants of wrapper methods, self-training and co-training. Self-training uses a single supervised model that is trained on labeled data, and pseudo-labels are used for the data points with confident predictions. Co-training, on the other hand, extends self-training to multiple supervised models, where two or more models exchange the most confident predictions to obtain pseudo-labels. Pseudo-labels can be very beneficial in label-scarce environments, but one must be mindful of the confirmation bias issue, where the model might rely on incorrect self-created labels. This problem has been extensively analyzed by Baykal et al. (2022) in the active distillation scenario, which is a strategy where a smaller model, known as the student model, is trained to mimic the behavior of a larger pre-trained model, known as the teacher model (Hoang et al., 2021; Kwak et al., 2022). In this context, confirmatory bias refers to the student model tendency to reproduce the predictions of the teacher model, even when the teacher predictions are incorrect. This can happen when the student model is trained to mimic the teacher model output too closely, without considering the underlying errors. To mitigate this, active distillation techniques use sample selection methods that encourage the student model to learn from data points where the teacher model makes errors, rather than just reproducing the teacher model predictions. In the more general active learning framework, confirmation bias might also refer to the tendency of an active learning algorithm to select examples that confirm its current hypothesis, rather than selecting examples that would challenge or improve it.

Finally, graph-based methods construct a graph on all available data and fit a supervised model, where the loss comprises a supervised loss and a regularization term that penalizes the difference between the labels predicted for connected data points. In the online active learning scenario, the graph structure can be used to model the similarity between data points, and the active learning algorithm can select the examples to label based on their position on the graph, such as selecting examples that are in low-density regions or are distant from other labeled examples.

3 Online active learning approaches

In this review, we present a taxonomy of online active learning strategies into four categories:

1. *Stationary data stream classification approaches.* These methods are designed to tackle online classification tasks, where the model is updated on the fly using newly labeled examples selected from a stream of data that does not change significantly over time. These methods are particularly useful in scenarios where the data distribution is relatively stable, such as quality control in industrial processes, where stationarity is often ensured by control actions taken at regular intervals and continuous maintenance of the components of the system (Bisgaard & Kulahci, 2011). Another example is represented by human activity recognition using wearable devices, where data is collected over time from wearable devices such as fitness trackers to identify patterns of activity like walking, running, or sleeping. This scenario would fall into this category because the data stream is relatively stable, and the model can be updated in real-time as new labeled examples become available (Miu et al., 2015).
2. *Drifting data stream classification approaches.* These online active learning strategies are specifically designed to handle classification tasks in dynamic environments where the

data distribution constantly changes. These approaches are designed to adapt to changes in the data distribution in order to maintain high classification accuracy. Some real-world applications might be fraud detection or intrusion detection. In financial fraud detection, fraudsters often change their methods to evade detection, so a classification model used for fraud detection must be able to adapt to new patterns of fraud as they emerge or to new customer habits (Zhang et al., 2022). In real-time intrusion detection, computer networks detection systems must be able to detect new forms of cyberattacks as they appear, so the classification models used must be able to adapt to changes in the data distribution over time (Nixon et al., 2021). This scenario would fall into this category because the data stream is constantly changing, and the model must be able to adapt to changes in the data distribution over time to maintain high accuracy.

3. *Evolving fuzzy system approaches.* These approaches are based on a type of fuzzy system that can adapt and change over time, in response to new data or changes in the environment (Gu et al., 2023). In traditional fuzzy systems, the rules and membership functions that define the system are fixed and do not change over time. Evolving fuzzy systems, on the other hand, are able to adapt their rules and membership functions based on new data or changes in the environment. This is particularly useful in applications where the data or the environment is non-stationary and evolves over time, such as in control systems for autonomous vehicles, where we must be able to adapt to changes in the environment, such as traffic patterns, road conditions, and weather (Naranjo et al., 2007; Wang et al., 2015).
4. *Experimental design and bandit approaches.* These methods, mostly related to regression models, actively select the most informative data points to improve model predictions. This category includes online active linear regression and sequential decision-making strategies like bandit algorithms or reinforcement learning. These methods adaptively select the most promising options in a given situation. An example is given by online advertising, where a model is used to select the most promising advertisements to display to users based on their browsing history and other factors (Avadhanula et al., 2021). This scenario would fall into this category because the model must adaptively select the most promising options in real-time based on the information available at that time. Also, in clinical trials, a model is used to select the most promising patients to enroll in a clinical trial based on their medical history and other personal information. Finally, in drug development studies (Réda et al., 2020), online active learning can be used to select the most promising compounds for further testing and development, based on their potential efficacy and safety.

This categorization provides a comprehensive overview of the different types of online active learning strategies and how they can be applied in various scenarios. While the simplest active learning strategy, random sampling, is available and involves selecting data points randomly from the stream for annotation, we will primarily focus on more specialized strategies designed to address scenarios where informed decisions are crucial due to resource constraints or where the data distribution is non-stationary.

Figure 5 depicts a general framework illustrating the essential components shared by the various categories of online active learning algorithms. The accompanying callouts highlight key options utilized by these methods. The following sections will provide an in-depth analysis of these strategies. For a more detailed flowchart regarding the drift detection and adaptation process, please refer to Lu et al. (2018), Lima et al. (2022).

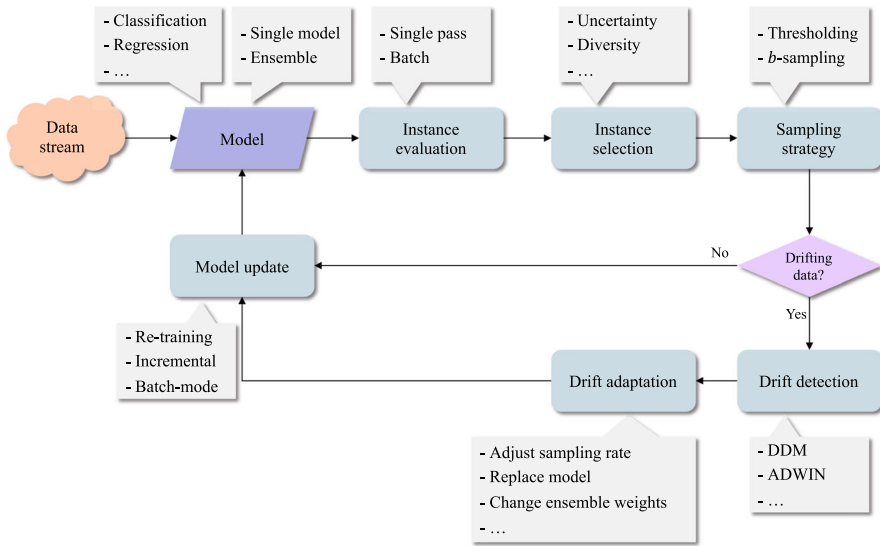


Fig. 5 Online active learning: general framework

3.1 Stationary data stream classification approaches

In online active learning, a commonly employed strategy is to request labels for data points that are considered to be informative enough based on a pre-determined threshold. This threshold can be established through a variety of techniques, depending on the instance selection criterion used to evaluate the informativeness of the unlabeled observations. Another method, sometimes referred to as *b*-sampling, is to calculate the probability that a data point will be queried by adjusting the parameter of a Bernoulli random variable, as proposed by Cesa-Bianchi et al. in one of the pioneering studies on online active learning (Cesa-Bianchi et al., 2004, 2006). They used a linear predictor characterized by the weight vector $\mathbf{w} \in \mathbb{R}^d$ and, at each time step t , after observing the current data point \mathbf{x}_t , the binary output $y \in \{-1, +1\}$ is predicted using

$$\hat{y}_t = \text{SGN} \left(\mathbf{w}_{t-1}^\top \mathbf{x}_t \right) \tag{1}$$

where \mathbf{w}_{t-1} is the weight vector estimated with the previously seen labeled examples $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_{t-1}, y_{t-1})$. The value $\mathbf{w}_{t-1}^\top \mathbf{x}_t$ is the margin, \hat{p}_t , of \mathbf{w}_{t-1} on the instance \mathbf{x}_t . If the learner queries the label y_t , a new weight vector is estimated using the newly added labeled example (\mathbf{x}_t, y_t) with the regular perceptron update rule (Rosenblatt, 1958) as in

$$\mathbf{w}_t = \mathbf{w}_{t-1} + M_t y_t \mathbf{x}_t \tag{2}$$

where M_t represents the indicator function of the event $\hat{y}_t \neq y_t$. If the label is not requested, the model remains unchanged, and we have $\mathbf{w}_t = \mathbf{w}_{t-1}$. At each time step t , the learner decides whether to query the label of a data point \mathbf{x}_t by drawing a Bernoulli random variable $Z_t \in \{0, 1\}$, whose parameter is given by

$$P_t = \frac{b}{b + |\hat{p}_t|} \tag{3}$$

where $b > 0$ is a positive smoothing constant that can be tuned to adjust the labeling rate. In general, as $\widehat{\rho}_t$ approaches 0, the sampling probability P_t converges to 1, suggesting that the labels are requested for highly uncertain observations. The sampling scheme introduced by Cesa-Bianchi et al. (2004) is referred to as selective sampling perceptron, and it is reported in Algorithm 1.

Algorithm 1 Selective sampling perceptron

Require: a data stream \mathbf{S} , an initial model $\mathbf{w}_0 = (0, \dots, 0)^\top$, a time horizon T , a sampling budget B , a parameter b .

$t \leftarrow 1$ ▷ Timestamp
 $c \leftarrow 0$ ▷ Labeling cost

while $c \leq B, t \leq T$ **do**

Observe an incoming data point $\mathbf{x}_t \in \mathbf{S}$ and set $\widehat{\rho}_t = \mathbf{w}_{t-1}^\top \mathbf{x}_t$

Predict the label $\widehat{y}_t = \text{SGN}(\widehat{\rho}_t)$

Draw a Bernoulli random variable Z_t of parameter $P_t = b / (b + |\widehat{\rho}_t|)$

if $Z_t = 1$ **then** ▷ Sampling decision

Ask for the true label y_t and update the model

$c \leftarrow c + 1$ ▷ Pay for the label

else

Discard \mathbf{x}_t

end if

$t \leftarrow t + 1$

end while

A similar approach to the one proposed by Cesa-Bianchi et al. (2004) was investigated by Dasgupta et al. (2005), who presented one of the first thresholding techniques for online active learning. They suggested setting a threshold on the margin, with the idea of sampling data points \mathbf{x}_t with a value of $|\widehat{\rho}_t|$ lower than a given threshold Γ . The threshold is initially set at a high value and iteratively divided by two until enough misclassifications occur among the queried points. The linear classifier is updated using the reflection concept [60] to give more focus to recent data points. Sculley (2007) built on the works of Cesa-Bianchi and Dasgupta to analyze the online active learning scenarios for real-time spam filtering. The author compares two models, a perceptron and a support vector machine (SVM), and tries three different instance selection criteria, the fixed thresholding approach by Dasgupta et al. (2005), the Bernoulli-based approach by Cesa-Bianchi et al. (2004), and a newly developed logistic margin sampling. The perceptron is updated as per Dasgupta et al. (2005), while the SVM is retrained on all available labeled observations each time a new data point is added. According to the logistic margin sampling strategy, the sampling decision is taken by drawing a Bernoulli random variable $Z_t \in \{0, 1\}$ with a parameter given by

$$P_t = e^{-\gamma|\widehat{\rho}_t|} \quad (4)$$

As in the traditional b -sampling approach introduced by Cesa-Bianchi et al. (2004), this sampling strategy depends on the uncertainty, meant as the distance from the prediction hyperplane. The main difference between the two strategies is the shape of the resulting sampling distribution, which can be observed in Fig. 6.

The selective sampling perceptron approach has also been investigated by Lu et al. (2016), who proposed an online passive-aggressive active learning variant of the algorithm. Similarly to the b -sampling approach, at each time step t , a Bernoulli random variable $Z_t \in \{0, 1\}$ is drawn to decide whether to query the label of the current data point \mathbf{x}_t or not. In this case, the parameter of Z_t is given by

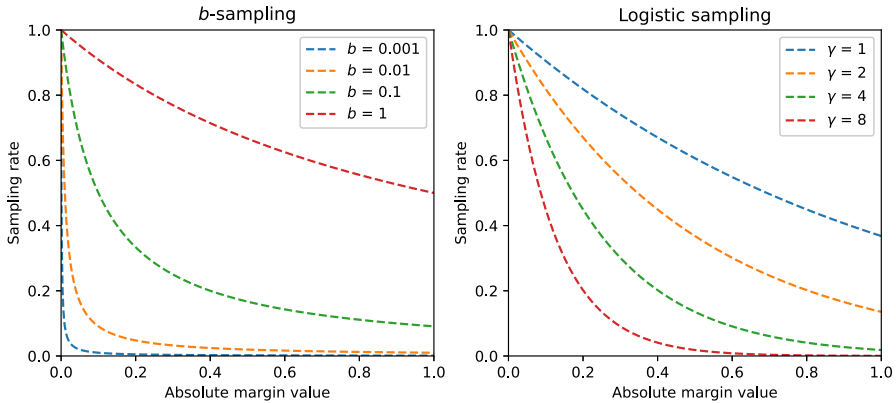


Fig. 6 Shape of the sampling distributions for b -sampling (a) and logistic sampling (b), for different values of b and γ

$$P_t = \frac{\delta}{\delta + |\hat{p}_t|} \tag{5}$$

where $\delta \geq 1$ is a smoothing parameter. Besides not allowing the smoothing parameter to assume a value lower than 1, the sampling distribution is the same as the one governed by the parameter in Eq. 3. The main difference lies in the passive-aggressive approach used for updating the weight vector. Indeed, while the traditional perceptron update, shown in Eq. 2, only uses misclassified examples to update the model, the passive-aggressive approach updates the weight vector $\mathbf{w} \in \mathbb{R}^d$ whenever the current loss $\ell_t(\mathbf{w}_{t-1}; (\mathbf{x}_t, y_t))$ is nonzero (Crammer et al., 2006). The new parameter \mathbf{w}_t is found using

$$\mathbf{w}_t = \mathbf{w}_{t-1} + \tau_t y_t \mathbf{x}_t \tag{6}$$

where τ_t represents the step size, and can be computed according to three different policies

$$\tau_t = \begin{cases} \ell_t(\mathbf{w}_{t-1}; (\mathbf{x}_t, y_t)) / \|\mathbf{x}_t\|^2 \\ \min(\kappa, \ell_t(\mathbf{w}_{t-1}; (\mathbf{x}_t, y_t)) / \|\mathbf{x}_t\|^2) \\ \ell_t(\mathbf{w}_{t-1}; (\mathbf{x}_t, y_t)) / (\|\mathbf{x}_t\|^2 + 1/2\kappa) \end{cases} \tag{7}$$

where κ is a penalty cost parameter. Passive-aggressive algorithms are known for their aggressive approach in updating the model, which is motivated by the fact that traditional perceptron updates might waste data points that have been correctly classified but with low prediction confidence.

A related issue to the update of the weight vector \mathbf{w}_t was emphasized by Bordes et al. (2005), who noted that always picking the most misclassified example is a reasonable sampling strategy only when the training examples are highly confident. When dealing with noisy labels, this strategy could lead to the selection of misclassified examples or examples lying on the wrong side of the optimal decision boundary. To address this, they suggested a more conservative approach that selects examples for updating \mathbf{w}_t based on a minimax gradient strategy.

In addition to confidence in the labels of the training examples, confidence in the model itself must be considered when the sampling strategy is based solely on model predictions. Hao et al. (2018b) pointed out that a margin-based sampling strategy may be suboptimal when

the classifier is not precise, especially in the early rounds of active learning when the model performance may be poor due to limited training feedback, leading to misleading sampling decisions. This issue is also referred to as cold-start active learning (Houlsby et al., 2014; Yuan et al., 2020; Jin et al., 2022). To address this, Hao et al. (2018b) propose considering second-order information in addition to margin value when deciding whether or not to query the label of a data point \mathbf{x}_t . In general, first-order online active learning strategies only consider the margin value, while second-order methods also take into account the confidence associated with it. To do this, they assume that the weight vector of the classifier $\mathbf{w} \in \mathbb{R}^d$ is distributed as

$$\mathbf{w} \sim \mathcal{N}(\boldsymbol{\mu}, \Sigma) \quad (8)$$

where the values μ_i and $\Sigma_{i,i}$ encode the model knowledge and confidence in the weight vector for the i th feature w_i . The covariance between the i th and j th features is captured by the term $\Sigma_{i,j}$. The smaller the variance associated with the coefficient w_i , the more confident the learner is about its mean value μ_i . The objective of the proposed method is to take into account the confidence of the model when updating the model and making the sampling decision. With regards to the model update, when the true label y_t of \mathbf{x}_t is queried, the Gaussian distribution in Eq. 8 is updated by minimizing an objective function based on the Kullback–Leibler divergence (Joyce, 2011) to ensure the updated model is not too different from the previous one. The sampling decision uses an additional parameter to the margin \hat{p}_t , which is defined as

$$c_t = \frac{-\eta}{2 \left(\frac{1}{v_t} + \frac{1}{\gamma} \right)} \quad (9)$$

where $\eta, \gamma > 0$ are two fixed hyper-parameters and v_t represents the variance of the margin related to the data point \mathbf{x}_t . The intuition is that, when the variance v_t is high, the model has not been sufficiently trained on instances similar to \mathbf{x}_t , and querying its label would lead to a model improvement. Then, a soft margin-based approach is employed by computing

$$\rho_t = |\hat{p}_t| + c_t \quad (10)$$

If $\rho_t \leq 0$, the label is always queried as the model is extremely uncertain about the margin. Instead, when $\rho_t > 0$, the model is more confident, and the labeling decision is taken by drawing a Bernoulli random variable of parameter

$$P_t = \frac{\delta}{\delta + \rho_t} \quad (11)$$

where $\delta > 0$ is a smoothing parameter. Finally, Hao et al. (2018b) also introduced a cost-sensitive variant of the loss function, for dealing with class-imbalanced applications. For a comprehensive discussion on imbalanced data stream analysis, please see Aguiar et al. (2023).

The cold-start issue related to the application of active learning to imbalanced datasets has also been highlighted by Qin et al. (2021), who used extreme learning machines (Huang et al., 2006) and extended the active learning framework initially proposed by Yu et al. (2015) to the multiclass classification scenario. They highlighted the challenge of the lack of instances for certain classes in imbalanced datasets, which can seriously impact the predictive ability of the model for those classes. To address this issue, they propose a sampling strategy that considers both diversity and uncertainty. The diversity is calculated by computing pairwise

Manhattan distance between the unlabeled observations. The uncertainty of a data point \mathbf{x}_t is computed by taking the difference between the largest two posterior probabilities as in

$$\text{margin}(\mathbf{x}_t) = p(y = c_b | \mathbf{x}_t) - p(y = c_{sb} | \mathbf{x}_t) \tag{12}$$

where c_b and c_{sb} are the classes with the highest posterior probabilities. This approach is also referred to as best-versus-second-best margin and, as highlighted by Joshi et al. (2009), is a good indicator of uncertainty when a large number of classes are present in the data. It should be noted that the sampling strategy introduced by Qin et al. (2021) is not suited for single-pass active learning as it requires computing similarity and uncertainty measures for all the unlabeled observations in the current batch.

Another approach to deal with class imbalance in active learning was proposed by Ferdowsi et al. (2013), who used linear SVMs and a sampling strategy that switches between multiple instance selection criteria online. This approach, however, is limited to a pool-based setting and requires predicting an unsupervised evaluation score for all available unlabeled instances. The impact of the last queried observations on the scores associated with the unlabeled data points is evaluated, and a greedy approach is used to decide which instance selection criterion to trust. SVMs have also been used by Ghassemi et al. (2016), who proposed a differentially private approach to online active learning. The privacy concerns are tackled both during the instance selection and the training phase, by randomizing the strategy introduced by Tong and Koller (2002). The informativeness of a data point \mathbf{x}_t is measured by its closeness to the current hyperplane \mathbf{w}_t as in

$$c(t) = \exp(-d(\mathbf{x}_t, \mathbf{w}_t)) \in [0, 1] \tag{13}$$

where the distance function $d(\mathbf{x}_t, \mathbf{w}_t)$ is defined as

$$d(\mathbf{x}_t, \mathbf{w}_t) \triangleq \frac{|\langle \mathbf{w}_t, \mathbf{x}_t \rangle|}{\|\mathbf{w}_t\|} \tag{14}$$

In the traditional framework, the label y_t is queried if we have $c(t) > \Gamma$, where Γ is a pre-defined threshold. It should be noted that $c(t) > \Gamma$ is equivalent to $d(\mathbf{x}_t, \mathbf{w}_t) \leq \log 1/\Gamma$, which means that the observation \mathbf{x}_t is in a sampling region of width $2 \log 1/\Gamma$ around \mathbf{w}_t . However, to avoid a deterministic decision process on the labeling and ensure privacy, some randomness needs to be introduced. This can be done in two ways. First, the labeling decision can be modeled as a Bernoulli random variable of parameter p if $c(t) < \Gamma$ or $(1 - p)$ if $c(t) \geq \Gamma$, where $p < 1/2$. Another approach is based on the exponential mechanism introduced by McSherry and Talwar (2007). According to this strategy, the algorithm sets a constant probability of labeling data points within a sampling region defined by α , and a decaying probability for points outside of it. The selection strategy is represented by a Bernoulli of parameter

$$q(t) = \begin{cases} e^{-\alpha\epsilon/\Delta} & d(\mathbf{x}_t, \mathbf{w}_t) \leq \alpha \\ e^{-d(\mathbf{x}_t, \mathbf{w}_t)\epsilon/\Delta} & d(\mathbf{x}_t, \mathbf{w}_t) > \alpha \end{cases} \tag{15}$$

where $\epsilon > 0$ and $\Delta = (1 - \alpha/M)M$. The authors assumed all data points belonging to the stream to be bounded in norm by M , $\|\mathbf{x}_t\| \leq M$ for $t = 1, \dots, T$. To tackle the privacy concerns while training, the authors propose two mini-batch strategies, to avoid the problem of slow convergence that may result from introducing noise according to the private stochastic gradient descent scheme (Bassily et al., 2014; Song et al., 2013; Duchi et al., 2013).

Two different approaches have been proposed by Ma et al. (2016) and Shah and Manwani (2020). Ma et al. (2016) proposed a query-while-learning strategy for decision tree classifiers. They used entropy intervals extracted from the evidential likelihood to determine the

dominant attributes, which are ordered based on the information gain ratio. When a new data point \mathbf{x}_t is observed, its label is queried only if there does not exist a dominant attribute. This will help to identify one and narrow the entropy interval. However, it should be noted that the authors consider a query while learning framework that only partially relates to online active learning. Shah and Manwani (2020) investigated the online active learning problem for reject option classifiers. Given the high cost that is sometimes associated with a misclassification error, these models are given the option of not predicting anything, for example when dealing with a highly ambiguous instance. A typical application of reject option classifiers is in the medical field, when making a diagnosis with ambiguous symptoms might be particularly difficult. In this case, it could be more beneficial not to provide a prediction but suggest further tests instead. They proposed an approach based on a non-convex double ramp loss function ℓ_{dr} (Manwani et al., 2013), where the label of the current example \mathbf{x}_t is queried only if it falls in the linear region of the loss given by $|f_t(\mathbf{x}_t)| \in [\rho_t - 1, \rho_t + 1]$, which is the region where the parameter would be updated. Here, ρ refers to the bandwidth parameter of the reject option classifier that determines the rejection region.

Fujii and Kashima (2016) investigated the problem of Bayesian online active learning. They provided a general framework based on policy-adaptive submodularity to handle data streams in an online setting. The authors distinguish between the stream setting, where the labeling decision can be made within a given timeframe, and the secretary setting, introduced in Sect. 2, where the labeling decision must be made immediately. The proposed framework can be applied in a variety of active learning scenarios, such as active classification, active clustering, and active feature selection. The framework is based on the concept of adaptive submodular maximization, which extends the idea of submodular maximization. A set function is considered to be submodular if it satisfies the property of diminishing returns, meaning that adding an element to a smaller set has a greater impact on the function value than adding the same element to a larger set. Adaptive submodular maximization allows the model to adapt to the changing distribution of data over time, by adjusting the set function to reflect the current state of knowledge. This leads to more efficient use of available data and improved performance.

So far, we discussed several single model approaches to active learning, which have shown promising results in various applications. However, it is important to note that single models have their limitations and can sometimes struggle to capture complex patterns and diverse representations present in the data. To address these limitations, researchers have proposed the use of ensembles or committees as an alternative (Krawczyk et al., 2017). An ensemble or committee refers to a group of multiple models that collaborate to produce a more robust and accurate prediction by combining their individual predictions. The models in an ensemble or committee can be trained on different subsets of the data or with varying hyperparameters, and the final prediction is typically made through either voting or weighted averaging. Ensembles or committees can also be regarded as a collection of models that work together to make a prediction, either by exchanging information or learning from one another. Among this class of methods, a common sampling strategy is represented by disagreement-based active learning. A framework to perform disagreement-based active learning in online settings was recently introduced by Huang et al. (2022). They characterized the learner by a hypothesis space \mathcal{H} of Vapnik–Chervonenkis (VC) dimension d , which is composed of all the classifiers currently under consideration, and a Tsybakov noise model (Mammen & Tsybakov, 1999; Tsybakov, 2004). Each classifier $h \in \mathcal{H}$ is a measurable function mapping the observation \mathbf{x}_t to binary output $y_t = \{0, 1\}$. The disagreement among two classifiers is given by $d(h_1, h_2) = \mathbb{P}[h_1(\mathbf{x}) \neq h_2(\mathbf{x})]$ and the disagreement region is defined as

$$D(h_1, h_2) = \{\mathbf{x} \in \mathcal{X} : h_1(\mathbf{x}) \neq h_2(\mathbf{x})\} \tag{16}$$

The online active learning strategy is represented by the policy $\pi = (\{v_t\}, \{\lambda_t\})$, where $\{v_t\}_{t \geq 1}$ is the map of the queried data points, and $\{\lambda_t\}_{t \geq 1}$ is the sequence of prediction rules. Over the time horizon T , the performance of the policy π is evaluated using the label complexity and the regret. The label complexity measures the expected number of labels queried, with respect to the stochastic process induced by π , and it is given by

$$\mathbb{E}[Q(T)] = \mathbb{E}\left[\sum_{t=1}^T \mathbb{1}[v_t = 1]\right] \tag{17}$$

The regret, on the other hand, represents the expected number of excess classification errors with respect to h^* , and it is obtained as

$$\mathbb{E}[R(T)] = \mathbb{E}\left[\sum_{t \leq T: v_t=0} \mathbb{1}[\lambda_t \neq y_t] - \mathbb{1}[h^*(\mathbf{x}_t) \neq y_t]\right] \tag{18}$$

The objective of the algorithm is to minimize the label complexity with a constraint on the regret. At the first round, the initial version space is the entire hypothesis space \mathcal{H} , while the initial region of disagreement is the whole input space \mathcal{X} . Then, at time step t , the learner updates the version space \mathcal{H}_t using the M collected labels, and computes a new region of disagreement as

$$\mathcal{D}(\mathcal{H}_t) = \{\mathbf{x} \in \mathcal{X} : \exists h_1, h_2 \in \mathcal{H}_t, h_1(\mathbf{x}) \neq h_2(\mathbf{x})\} \tag{19}$$

If $\mathbf{x}_t \in \mathcal{D}(\mathcal{H}_t)$, then the label of the current data point is queried, otherwise a prediction is produced using an arbitrary classifier in \mathcal{H}_t . At the end of the iteration t , the set \mathcal{Z}_t of M collected labeled examples is used to estimate the empirical error $\epsilon_{\mathcal{Z}_t}(h)$ of the classifiers in \mathcal{H} and identify the best currently available classifier. Then, the version space is updated by removing all the suboptimal hypotheses whose empirical error exceeds the one obtained with h_t^* by a threshold $\Delta_{\mathcal{Z}_t}(h, h_t^*)$. The threshold regulates the trade-off between reducing label complexity by narrowing the region of disagreement and increasing the regret by eliminating good classifiers.

The disagreement concept was also used by Desalvo et al. (2021), while proposing an approach to online active learning for binary classification tasks based on surrogate losses. The overall framework is similar to the disagreement-based one used by Huang et al. (2022), with the main difference being the use of weak-labels to optimize the sampling strategy. At each time step t , the learner observes the unlabeled data point \mathbf{x}_t and either decides to request its label or assigns a pseudo-label \hat{y}_t . Then, the pseudo labels \hat{y}_t and the true labels y_t processed so far are used together to obtain an estimate of the empirical risk $\epsilon_{\mathcal{S}_t}(h)$, where \mathcal{S}_t is obtained by combining the collected labeled examples \mathcal{Z}_t with the pseudo-labeled ones $\hat{\mathcal{Z}}_t$. This represents an example of combining active learning and semi-supervised learning, as highlighted in Sect. 2.3.

Loy et al. (2012) presented a Bayesian framework that leverages the principle of committee consensus to balance exploration and exploitation in online active learning. The aim of exploration is to discover new, previously unknown classes, while exploitation focuses on refining the decision boundary for known classes. To address the issue of unknown classes, the framework uses a Pitman-Yor Processes (PYP) prior model (Pitman & Yor, 1997) with a Dirichlet process mixture model (DPMM). A DPMM is a non-parametric clustering and classification model that models the data generating process using a mixture of probability

distributions. Each data point is assigned to a cluster, which is associated with a probability distribution over the classes. The number of clusters is modeled using a Dirichlet process, which is a distribution over distributions that allows for an infinite number of clusters but ensures that the number of actual clusters is always finite. At each time step t , the learner samples two random hypotheses h_1 and h_2 from the model. Then, it computes the posterior probability of the current class c corresponding to k , $p(c = k | \mathbf{x}_t)$, for each of the two hypotheses. Finally, $h_i(\mathbf{x}_t) = \arg \max p(c | \mathbf{x}_t)$ is calculated for $i = 1, 2$. The label of the current data point is queried in two cases: first, if $h_1(\mathbf{x}_t) \neq h_2(\mathbf{x}_t)$, meaning the two hypotheses disagree, and second, if $h_i(\mathbf{x}_t) = K + 1 \forall i$, where K is the number of currently known classes, meaning the current data point belongs to a new class.

The DPMM has also been used by Mohamad et al. (2020), who proposed a semi-supervised strategy for performing active learning in online human activity recognition with sensory data. To account for the possibility of dealing with different sensor network layouts, the authors proposed pre-training a conditional restricted Boltzmann machine (Taylor & Hinton, 2009; Taylor et al., 2006) and used it to extract generic features from the sensory input. The instance selection strategy follows a Bayesian approach, in trying to minimize the uncertainty about the model parameters. To assess the usefulness of labeling the data point \mathbf{x}_t , they measure the discrepancy between the model uncertainty computed from the data observed until the time step t and the expected risk associated with y_t . This gives a hint of how the current label would impact the current model uncertainty. A dynamically adaptive threshold Γ is finally used to determine whether the current expected risk is greater than the current risk.

A different kind of committee has been considered by Hao et al. (2018a). They proposed a framework for minimizing the number of queries made by an online learner that is trying to make the best possible forecast, given the advice received from a pool of experts. To do so, they adapted the exponentially weighted average forecaster (EWA) and the greedy forecaster (GF) to the online active learning scenario. A comprehensive analysis of forecasters to perform prediction with expert advice can be found in the book by Cesa-Bianchi and Lugosi (2006). In general, at each time step t , the learner or forecaster has access to the predictions for the data point \mathbf{x}_t made by the N experts, $f_{i,t}(\mathbf{x}_t) : \mathbb{R}^d \rightarrow [0, 1]$ with $i = 1, \dots, N$. Based on these predictions, it outputs its own prediction p_t for the outcome y_t . Then, if the label is revealed, the predictions made by the forecaster and the experts are scored using a nonnegative loss function ℓ . The objective of the learner is to minimize the cumulative regret over the time horizon T , which can be seen as the difference between its loss and the one obtained with each expert i as in

$$R_{i,T} = \sum_{t=1}^T (\ell(p_t, y_t) - \ell(f_{i,t}(\mathbf{x}_t), y_t)) = \widehat{L}_T - L_{i,T} \quad (20)$$

The most simple approach to obtain a prediction p_t from the learner is to compute a weighted average of the experts predictions as in

$$p_t = \frac{\sum_{i=1}^N \omega_{i,t} f_{i,t}(\mathbf{x}_t)}{\sum_{i=1}^N \omega_{i,t}} \quad (21)$$

where $\omega_{i,t} \geq 0$ is the weight assigned at time t to the i th expert. With the EWA, the weight for the i th expert are obtained using

$$\omega_{i,t} = \frac{e^{\eta R_{i,t-1}}}{\sum_{i=1}^N e^{\eta R_{i,t-1}}} \quad (22)$$

where η is a positive decay factor and $R_{i,t-1}$ is the cumulative loss of expert i observed until step t . The exponential decay factor η determines the weight given to the past losses, with more recent losses having a higher weight and older losses having a lower weight. Instead, the GF works by minimizing, at each time step, the largest possible increase of the potential function for all the possible outcomes of y_t . The potential function is the function that assigns a potential value to each expert, which captures the quality of an expert advice based on its past performance. Hao et al. (2018a) extended the EWF and GF by proposing the active EWF (AEWF) and active GF (AGF). The key idea is that, while the standard EWF and GF assume the availability of the true label y_t after each prediction, in the online active learning framework the loss ℓ can only be measured a limited number of times. To factor this in, a binary variable $Z_t \in \{0, 1\}$ is introduced to decide whether or not at round t the label is requested. Consequently, the cumulative loss suffered by the i th expert on the instances queried by the active forecaster is given by

$$\widehat{L}_{i,T} = \sum_{t=1}^T \ell(f_{i,t}(\mathbf{x}_t), y_t) \cdot Z_t \quad (23)$$

The sampling strategy is based on the determination of a confidence condition on the difference between the prediction p_t of the fully supervised forecaster and the prediction \widehat{p}_t made by the active forecaster. For the active forecaster we have that $\widehat{p}_t = \pi_{[0,1]}(\overline{p}_t)$, where \overline{p}_t depends on the chosen model. The AEWF is based upon the observation that if we have

$$\max_{1 \leq i, j \leq N} |f_{i,t}(\mathbf{x}_t) - f_{j,t}(\mathbf{x}_t)| \leq \delta \quad (24)$$

then $|p_t - \widehat{p}_t| \leq \delta$, where δ is a tolerance threshold. This means that the prediction of the forecaster is close to the one obtained in the fully supervised setting if the maximum difference of advice between any two experts is not too large. This assumption might not hold in the presence of noisy or bad experts and, to tackle this problem, the authors proposed a robust variant of the AEWF. The AGF uses instead a confidence condition based on the fact that if

$$\max_{1 \leq i, j \leq N} |f_{i,t}(\mathbf{x}_t) - \overline{p}_t| \leq \delta \quad (25)$$

then $|p_t - \widehat{p}_t| \leq \delta$. The general scheme for performing online active learning with expert advice is reported in Algorithm 2.

A similar framework, in conjunction with multiple kernel learning (MKL), has been investigated by Chae and Hong (2021). They propose an active MKL (AMKL) algorithm based on random feature approximation. In general, online MKL based on random feature approximation is a method for online learning and prediction that combines multiple kernel functions to improve the performance of a learning algorithm (Jin et al., 2010; Hoi et al., 2013). In MKL, multiple kernel functions are used to capture different aspects of the data, and the optimal combination of kernels is learned from the data. The online version of MKL based on random feature approximation is designed to handle data that arrives sequentially, and the learning algorithm is updated after each new data point. In kernel-based learning, the target function $f(\mathbf{x})$ is assumed to belong to a reproducing Hilbert kernel space (RKHS). In the proposed AMKL the learner uses an ensemble of N kernel functions. At each time step t , two main steps are implemented. First, each kernel function $\widehat{f}_{i,t}(\mathbf{x}_t)$, with $i = 1, \dots, n$, is optimized independently of the other kernel functions. This is referred to as local step. Then, in the global step, the learner seeks the best function approximation $\widehat{f}_t(\mathbf{x}_t)$ by combining the N kernel functions as in

Algorithm 2 Online active learning with expert advice

Require: a data stream \mathbf{S} , a loss function ℓ , a time horizon T , a set of N experts, a tolerance threshold δ , a sampling budget B .

```

t ← 1
c ← 0
while c ≤ B, t ≤ T do
    Observe an incoming data point  $\mathbf{x}_t \in \mathcal{S}$ 
    Receive advice by experts  $\{f_{i,t}(\mathbf{x}_t) : i = 1, \dots, N\}$ 
    Generate prediction  $\bar{p}_t$  for the label  $y_t$  and set  $\hat{p}_t = \pi_{[0,1]}(\bar{p}_t)$ 
    Draw a Bernoulli random variable  $Z_t$  of parameter  $P_t = b/(b + |\hat{p}_t|)$ 
    if Equation 24 or 25 is satisfied then
        Discard  $\mathbf{x}_t$ 
    else
        Ask for the true label  $y_t$ 
        c ← c + 1
    end if
    t ← t + 1
end while

```

▷ Timestamp
▷ Labeling cost
▷ Sampling decision
▷ Pay for the label

$$\hat{f}_t(\mathbf{x}_t) = \sum_{i=1}^N \hat{v}_{i,t} \hat{f}_{i,t}(\mathbf{x}_t) \tag{26}$$

where $\hat{v}_{i,t}$ refers to the weight for the i th kernel function at round t . Similarly to the case with expert advice, the weights are determined by minimizing the regret over the time horizon T , which is defined as the difference between the loss of the learner and the one obtained with the best kernel function $f_{i,t}^*$. To do so, the weights are computed based on the past losses ℓ as

$$\hat{\omega}_{i,t} = \exp \left(-\eta_g \sum_{\tau \in \mathcal{A}_{t-1}} \ell \left(\hat{f}_{i,\tau}(\mathbf{x}_\tau), y_\tau \right) \right) \tag{27}$$

where $\eta_g > 0$ is a tunable parameter and \mathcal{A}_{t-1} is an index of time stamps t indicating the instances for which has label has been requested, thus permitting to measure the loss. Then, the weights $\hat{v}_{i,t}$ are obtained from $\hat{\omega}_{i,t}$ as follows

$$\hat{v}_{i,t} = \frac{\hat{\omega}_{i,t}}{\sum_{i=1}^N \hat{\omega}_{i,t}} \tag{28}$$

Finally, the instance selection criterion is based on a confidence condition, denoted by with $\delta > 0$, on the similarity of the learned kernel function, which is a similar to the condition used by Hao et al. (2018a) in the formulation of the AEWAF

$$\max_{1 \leq j \leq N} \sum_{i=1}^N \hat{v}_{i,t} \ell \left(\hat{f}_{i,t}(\mathbf{x}_t), \hat{f}_{j,t}(\mathbf{x}_t) \right) \leq \delta \tag{29}$$

3.2 Drifting data stream classification approaches

Active learning strategies belonging to this category aim to tackle online classification tasks in time-varying data streams affected by distribution shifts. We can classify distribution shifts into three main categories, depending on whether they concern the feature space \mathbf{x} or the output dimension y . A shift that only affects the input distribution $p(\mathbf{x})$, and not the conditional distribution $p(y | \mathbf{x})$, is referred to as covariate shift (Zhou et al., 2021; Wu et al., 2021; Li et




Fig. 7 Different types of drifts that can affect the data stream: abrupt drift (a), gradual drift (b), incremental drift (c), recurring concepts (d). C_1 and C_2 indicate the two concepts that might characterize the data distribution

al., 2021) or virtual drift (Baier et al., 2021). In these circumstances, for two different time steps, t_i and $t_{i+\Delta}$, we have that $p_{t_i}(\mathbf{x}) \neq p_{t_{i+\Delta}}(\mathbf{x})$ and $p_{t_i}(y | \mathbf{x}) = p_{t_{i+\Delta}}(y | \mathbf{x})$, meaning that the underlying model is not being altered by phenomena like class swaps or coefficient changes. Conversely, in the presence of a real concept drift (Baier et al., 2021; Suárez-Cetrulo et al., 2023), the conditional distribution changes, and we have $p_{t_i}(y | \mathbf{x}) \neq p_{t_{i+\Delta}}(y | \mathbf{x})$. In this scenario, the predictive performance of the fitted model dramatically deteriorates, and a model update or replacement becomes necessary. An example of this kind of distribution shift can be identified in the changes of the consumer behaviors over time, or following a major event as the COVID-19 pandemic (Zwanka & Buff, 2021). However, it should be noted that virtual drifts and real concept drifts often occur together (Tsymbol et al., 2008), leading to a situation where we have both $p_{t_i}(\mathbf{x}) \neq p_{t_{i+\Delta}}(\mathbf{x})$ and $p_{t_i}(y | \mathbf{x}) \neq p_{t_{i+\Delta}}(y | \mathbf{x})$ (Lu et al., 2018). Lastly, we can incur in a label distribution shift (Wu et al., 2021) when the shift only affects $p(y)$, leading to $p_{t_i}(y) \neq p_{t_{i+\Delta}}(y)$. This situation can be observed in many real-world scenarios where the target distribution changes over time. A typical example is the prediction of diseases like influenza, whose distribution can dramatically change depending on the season, or in the presence of sudden outbreaks.

Another key characteristic of distribution shifts is represented by the change rate, namely how fast the new concept or distribution is introduced into the data stream. To this extent, we can identify four kinds of drifts (Lu et al., 2018; Lima et al., 2022), which are illustrated in Fig. 7. A sudden or abrupt drift is a drift that can be immediately detected from two consecutive time steps, t_i and t_{i+1} . It refers to a sudden and clearly identifiable change in the data distribution. An example of this would be a sudden change in the weather, which would affect the behavior of customers at a retail store. The change is noticeable, and the model needs to be updated immediately. A gradual drift exhibits a transition phase, where a mixture or overlap between the two distributions p_{t_i} and $p_{t_{i+\Delta}}$ exists. In this case, the change is slower and more difficult to detect, making it challenging to update the model. An example would be a change in consumer behavior over time, which is hard to detect but can have a significant impact on a business. Another type of drift is the incremental drift, which has an extremely low transition rate, which makes it very difficult to detect changes between the data points observed in the transition period. This type of drift is often caused by changes in the data generating process that happen gradually over time, in small steps rather than all at once. An example would be changes in the types of products that are popular among customers, which happen gradually and are hard to detect. Finally, a data stream can also be affected by recurring concepts, which sequentially alternate over time. An example would be a retail store where the same types of products are popular at different times of the year, such as winter coats and summer dresses. The model needs to be able to detect and adapt to these recurring concepts in order to maintain good performance.

In online active learning for drifting data streams, some approaches address the presence of concept drifts by combining active learning strategies with drift detectors (Zhang et al., 2020a; Krawczyk et al., 2018). Drift detectors are algorithms that try to detect distribution shifts and identify when the context is changing. They can be divided into three macro-categories (Lu et al., 2018). The first group of methods is represented by the error-based drift detectors, which try to detect online changes in the error rate of a base classifier. Among these, one of the most commonly employed strategies is the drift detection method (DDM) proposed by Gama et al. (2004). Another popular approach is the adaptive window (ADWIN) strategy proposed by Bifet and Gavaldà (2007). The second class of drift detectors is called data distribution-based drift detection, and the third class is represented by multiple hypothesis testing strategies. While the first class contains the majority of the proposed approaches, it assumes that we are able to observe the labels of all the incoming data points to assess the error rate. Instead, the last two classes could be implemented even in an unsupervised manner. An exhaustive overview on unsupervised drift detection methods has been proposed by Gemaque et al. (2020). While the unsupervised nature of the data distribution-based and multiple hypothesis testing strategies make them ideal for the active learning scenario, it should be noted that real concept drifts can hardly be detected in a completely unsupervised fashion. Indeed, in a circumstance when the input distribution $p(\mathbf{x})$ remains unaltered while the underlying model relating the input variables \mathbf{x} to the label y changes, it would not be possible to detect the change of concept without collecting labels. This is why Krawczyk et al. (2018) propose to apply an error-based drift detector to the few labels collected during the online active learning routine. To this extent, they use the ADWIN (Bifet & Gavaldà, 2007) method to detect drifts and decide when the current model needs to be updated or replaced. The proposed general framework for dealing with online active learning with drifting data streams is reported in Algorithm 3.

Moreover, the authors proposed the use of a time-variable threshold to balance the budget use over time. Their approach is based on the intuition that, when a new concept is introduced, more labeling effort will be required to quickly collect representative observations belonging to the new concept and replace the outdated model. This is obtained by adjusting a time-variable threshold to balance the budget use over time. Given a threshold Γ on the uncertainty of the classifier and a labeling rate adjustment $r \in [0, 1]$, the threshold is reduced to $\Gamma - r$ when ADWIN raises a warning and to $\Gamma - 2r$ when a real drift is detected. Thus, when allocating the labeling budget, the key requirement is that the labeling rate employed when a drift is detected should be strictly larger than the one used in static conditions. A similar thresholding idea has also been used by Castellani et al. (2022), who proposed an active learning strategy for non-stationary data streams in the presence of verification latency. They used a piece-wise constant budget function, where the labeling rate α is increased to α_{high} when a drift is detected and, after a while, reduced to α_{low} . Finally, the labeling rate is restored to its nominal value α . A visual representation of the labeling approach is shown in Fig. 8. The length of the time segments where the labeling rate is altered depends on the desired values for α_{high} and α_{low} , constraining the overall labeling rate to be equal to α .

The authors also tackled the verification latency issue by considering the spatial information of a queried point for which the label has not been made available yet by the oracle. In this way, it is possible to avoid oversampling from regions where many close points have a high utility, namely a low classification confidence. While assessing the utility of the incoming data points the authors use real and pseudo-labels by propagating the information contained in the already labeled observations, as suggested by Pham et al. (2022). The idea is to use the spatial information of the queried labels by estimating the still missing labels with a weighted majority vote of the label of its k -nearest neighbors labels, where the weight for

Algorithm 3 Online active learning with drifting data streams

Require: a data stream S , a classifier Θ , a drift detector Θ , a sampling strategy Υ , a labeling rate α , a sampling budget B .

```

 $t \leftarrow 1$  ▷ Timestamp
 $c \leftarrow 0$  ▷ Labeling cost
while  $c \leq B$  and  $t \leq |S|$  do
  Observe incoming data point  $x_t \in S$ 
  if  $\Upsilon(x_t) = \text{True}$  then ▷ Sampling decision
    Ask for the true label  $y_t$ 
     $c \leftarrow c + 1$  ▷ Pay for the label
    Update classifier  $\Psi$  with the labeled example  $(x_t, y_t)$ 
    Update drift detector  $\Theta$  with the labeled example  $(x_t, y_t)$ 
    if drift warning = True then
      Start to train a new classifier  $\Psi_{\text{new}}$ 
      Increase labeling rate  $\alpha$ 
    else
      if drift detected = True then ▷ A detection is always preceded by a warning
        Replace  $C$  with  $C_{\text{new}}$ 
        Further increase  $\alpha$ 
      else
        Return to initial labeling rate  $\alpha$ 
      end if
    end if
    if  $C_{\text{new}}$  exists then ▷ Keeps being updated in the background until replacement
      Update classifier  $C_{\text{new}}$  with the labeled example  $(x_t, y_t)$ 
    end if
  end if
   $t \leftarrow t + 1$ 
end while

```




Fig. 8 Piece-wise constant budget function introduced by Castellani et al. (2022). The sampling rate α is increased to α_{high} when a drift is detected (t_{drift}), then reduced to α_{low} between t_{r1} and t_{r2} , before being restored to its nominal value

each nearest neighbor depends on the arrival time of the labels. The verification latency issue in online active learning with drifting data streams was also extensively analyzed by Pham et al. (2022). Consider the general case where at time t_i^x we draw an instance \mathbf{x}_i , and find it interesting enough to send it to the oracle, which will send back the label y_i only at time t_i^y , where $t_i^y > t_i^x$. Before the requested label arrives, we might encounter another instance similar to \mathbf{x}_i and ask again for its label, since the learner could not update its utility function or threshold. Similarly, we might use outdated information when updating the policy in a future

window. To tackle these issues, the authors propose a forgetting and simulating strategy to avoid using soon-to-be outdated observations and prevent redundant labeling. The instance selection is based upon the variable uncertainty strategy proposed by Zliobaite et al. (2014) and the balanced incremental quantile filter by Kottke et al. (2015). If we denote the current sliding window at time t_n^x as $\mathcal{W}_n = [t_n^x - \Delta, t_n^x)$ and use windows of fixed size Δ , we know that the sliding window that would be used for training when the label y_n related to \mathbf{x}_n arrives will be given by $\mathcal{D}_n = [t_n^y - \Delta, t_n^y)$. The forgetting step is then implemented by discarding outdated labeled examples that are included in \mathcal{W}_n but will not be included in \mathcal{D}_n . If a_i is a Boolean variable indicating whether the i th observation has been labeled, the set of instances selected to be forgotten is given by

$$O_n = [(\mathbf{x}_i, y_i) \forall i < n : a_i = 1 \wedge t_i^x, t_i^y \in \mathcal{W}_n \setminus \mathcal{D}_n]. \quad (30)$$

Similarly, there is a second set of observations, with time stamps $\mathcal{D}_n^+ = \mathcal{D}_n \setminus \mathcal{W}_n = [t_n^x, t_n^y)$, where there might be instances that have been queried but whose label is not currently available. To avoid losing such information and redundantly asking for the label of similar instances, the algorithm simulates incoming labels with a bagging approach by averaging across multiple utility estimations. They also consider an alternative simulation approach based on fuzzy labeling.

Similarly to Krawczyk et al. (2018), the ADWIN drift detector has also been used by Zhang et al. (2020a) while proposing a method for dealing with online active learning in environments characterized by concept drifts and class imbalance. The instance selection criterion is based on the predictive uncertainty, which they estimate using the best-versus-second-best margin value (Eq. 12), as they tackle a multi-class classification problem. An initial pool of n observations is passively collected from the stream to initialize the active learning strategy. Then, a threshold Γ_i is estimated for each class as in

$$\Gamma_i = \begin{cases} \frac{nm}{n_i L} & \frac{n}{n_i L} \geq 1 \\ m & \frac{n}{n_i L} < 1 \end{cases} \quad (31)$$

where $i = 1, \dots, L$ is the number of classes and m is a pre-defined constant used to control the size of the threshold. The model is represented by an ensemble of N classifiers and, when ADWIN detects a concept drift, the classifier with the higher error is replaced with a newly trained one. Finally, the class imbalance issue is also taken into account in two ways, during the training of the ensemble with the use of class-specific weights, and during the active learning routine, by dynamically adjusting the threshold to select more observations belonging to the minority class.

Recently, Cheng et al. (2023) presented another approach to combine online active learning with drift detection. Their method involves segmenting the data stream \mathbf{S} into fixed-length chunks and then detecting drifts by comparing the distributions of adjacent chunks. After a drift is detected, a multi-objective optimization problem is formulated in order to identify the most relevant and diverse data points within the current batch. For a data point \mathbf{x}_t , relevance is defined as its contribution to the new concept, and diversity as the Pearson correlation coefficient with other instances in the same region. Instead, Martins et al. (2023) proposed to sample the most uncertain data points from each chunk, using a meta-learning framework to fine-tune the threshold used for each window. This allows to reduce the need for labels while maintaining a steady adaptation to the new concepts.

Another window-based approach to perform active learning from data streams has been proposed by Zhu et al. (2007). The authors developed an ensemble E by partitioning the data stream \mathbf{S} into chunks and then training each of the k models composing the ensemble E on

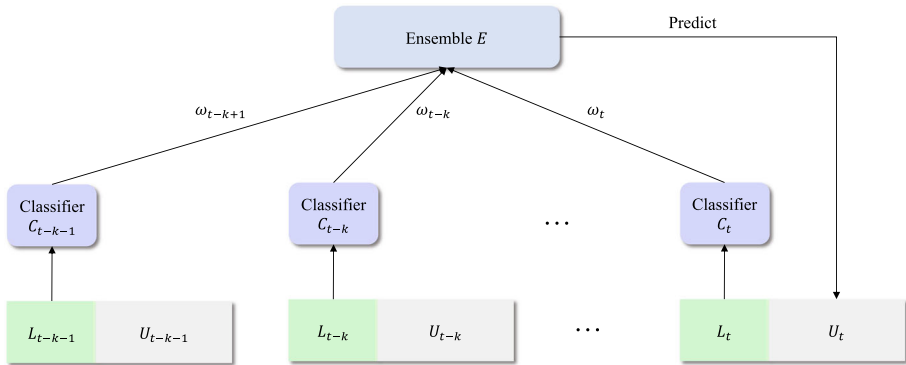


Fig. 9 Ensemble-based active learning framework for data streams proposed by Zhu et al. (2007)

a different chunk of data. In this way, even if the previous observations become unavailable, the models can be used when taking the sampling decision in order to take into account a global uncertainty measure, which is a more robust approach than treating each chunk as a static dataset. At time step t , the learner receives a data chunk \mathbf{S}_t , which is used to build the current classifier C_t . At this point, the ensemble is composed by C_t , together with the most recent $k - 1$ classifiers, $C_{t-k+1}, \dots, C_{t-1}$, trained on the labeled examples sampled from the previously observed data chunks, $L_{t-k+1}, \dots, L_{t-1}$. At each iteration, the objective is to predict the remaining unlabeled data points from the current chunk, U_t . The ensemble-based active learning framework is depicted in Fig. 9. The instances selected to be queried are the ones with the largest ensemble variance, and the predictions are obtained by combining the predictions of the single classifiers using the weights $\omega_{t-k+1}, \dots, \omega_{t-1}$. Finally, a weight updating rule is used to adapt to dynamic data streams.

Shan et al. (2019) and Zhang et al. (2018) developed online active learning strategies by building upon the pairwise classifiers strategy introduced by Xu et al. (2016). The pairwise strategy makes use of two models, a stable classifier C_s and a dynamic classifier C_d , and divides the data stream into batches as in Zhu et al. (2007). The prediction for an incoming data point \mathbf{x}_t is obtained with a weighted average of the predictions obtained from the two classifiers as in

$$f_E(\mathbf{x}_t) = \omega_s f_{C_s}(\mathbf{x}_t) + \omega_d f_{C_d}(\mathbf{x}_t) \tag{32}$$

where ω_s and ω_d are the weights associated with the stable and the dynamic classifier, respectively. At time t , the stable classifier C_s is trained on the labeled portions of all the batches processed so far, L_1, \dots, L_{t-1} . Conversely, C_d is trained exclusively on L_{t-1} . The key idea is that whenever the reactive classifier starts to outperform the stable classifier, the stable classifier is replaced by the reactive one, which is eventually reset. This replacement allows the learner to adapt to the drift and focus on the most recent instances, forgetting the seemingly obsolete data points. The main drawback of this approach is that it cannot effectively address gradual drifts as the replacement with the classifier trained on the most recent observations makes the learner forget about observations away from the current window. Hence, similarly to the approach of Zhu et al. (2007), Shan et al. (2019) proposed an extension of this approach, based on an ensemble of classifiers in trying to contemporarily address gradual drifts and abrupt drifts. In their strategy, the stable classifier learns from all the labeled instances and the reactive classifier is replaced by an ensemble of dynamic classifiers, trained on multilevel

sliding windows to capture changes in the data stream at different time intervals. The instance selection approach combines random sampling and uncertainty sampling, where the latter is based on the margin value of the predictions obtained by the ensemble. It should be noted that the prediction f_E for the data point \mathbf{x}_t is obtained as a weighted combination of the predictions obtained with the stable and dynamic classifiers as in

$$f_E(\mathbf{x}_t) = \omega_s f_{C_s}(\mathbf{x}_t) + \sum_{d=1}^D \omega_d f_{C_d}(\mathbf{x}_t) \quad (33)$$

The stable classifier has a constant weight $\omega_s = 0.5$ and plays a crucial role in trying to learn the overall trend and direction of concept drift. Conversely, the dynamic classifiers have gradually decaying weights, according to a damped sliding windowing approach where each weight is initialized at $\frac{1}{D}$ and then reduced according to its creation time

$$\omega_d = \begin{cases} \omega_d \left(1 - \frac{1}{D}\right) & d = 1, \dots, D - 1 \\ \frac{1}{D} & d = D \end{cases} \quad (34)$$

The most recent classifiers are useful in detecting sudden concept drifts and have highest weights while the old dynamic classifiers have lower weights and can help to identify gradual drifts. The same pairwise strategy based on an ensemble composed by a stable classifier and D dynamic classifiers was used by Zhang et al. (2022). They modified the original strategy by introducing a reinforcement mechanism to adjust the weights ω_d according to the prediction performance and the class imbalance issue. The weights adjustment strategy is described by Algorithm 4. It should be noted that this procedure is only implemented after the true label y_t has been revealed by the oracle. The damped class imbalance ratio (DCIR) value is obtained by taking into account the number of observations for each class collected so far. This is expected to be useful when dealing with imbalanced classes. With regards to the instance selection criterion, the authors consider a hybrid strategy combining uncertainty sampling and random sampling, since approaches solely based on uncertainty could ignore a concept change that is not close to the boundary. Woźniak et al. (2023) recently proposed another ensemble-based active learning strategy where the data points to be labeled are selected from the current chunk using the budget labeling active learning strategy introduced by Zybiewski et al. (2020). According to this approach, the learner selects both random and informative data points, where the informativeness is determined using the support function threshold, which in the case of binary classification problems can be interpreted as a distance from the decision boundary.

Algorithm 4 Weight adjustment for dynamic classifiers

Require: a labeled observation (\mathbf{x}_t, y_t) , number of classes K , number of dynamic classifiers D , current weights ω_d with $d = 1, \dots, D$, DCIR $_{\kappa}$ for $\kappa \in K$.

```

if DCIR $[y_t] < \frac{1}{K}$  then                                ▷ Check if it belongs to the minority class
  for  $d$  in  $(1, D)$  do
    if  $C_d(\mathbf{x}_t) = y_t$  then                                ▷ Check if the prediction made by  $C_d$  is correct
       $\omega_d \leftarrow \omega_d \left(1 + \frac{1}{D}\right)$                 ▷ Increase weight of classifier  $C_d$ 
    else
       $\omega_d \leftarrow \omega_d \left(1 - \frac{1}{D}\right)$                 ▷ Decrease weight of classifier  $C_d$ 
    end if
  end for
end if

```

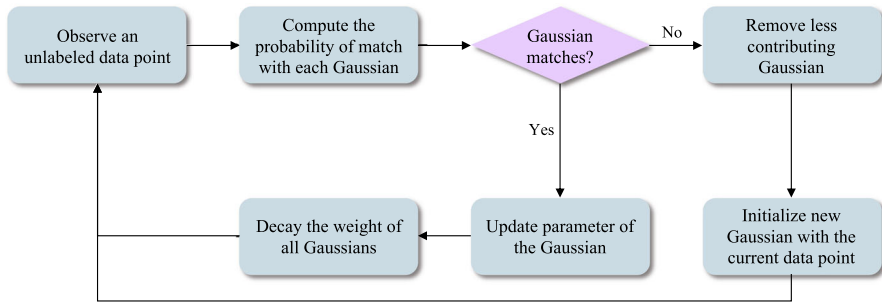


Fig. 10 Main steps of the growing Gaussian mixture model used by Mohamad et al. (2018)

Another way to perform online active learning in time-varying data streams is to use clustering-based approaches. Halder et al. (2023) extended the framework based on stable and dynamic classifiers by introducing a clustering step that aims to train the new stable classifier C_s on the most informative and representative instances from each data block. Similarly, Ienco et al. (2013) investigated a clustering-based approach in a batch-based scenario, where only a fraction of the incoming block of observations can be labeled. They extend the pre-clustering approach (Nguyen & Smeulders, 2004), which had been previously studied in the pool-based scenario, to the stream-based case. The sampling strategy takes into account an extra-cluster metric, to sort the clusters, and an intra-cluster one, to sort the observations within each cluster. When a new batch arrives, observations are clustered, and clusters are sorted based on the homogeneity of the clusters, which is measured taking into account the number of (predicted) classes within each cluster. If a cluster is balanced in the number of expected classes, it is regarded to as an uncertain cluster that covers a more difficult area of the input space. Within each cluster, the certainty of an observations is determined by its representativeness, namely the distance from the centroid, and the uncertainty, meant as the maximum a posteriori probability among all the predicted classes for \mathbf{x}_t . When the clusters and observations are ranked, the learner starts to iteratively ask the observations label in an alternate fashion. To sample the most representative data points from each batch, Zhang et al. (2023) suggested the use of density-peak clustering and recognize the incomplete clusters in the dynamic feature space through the altitude of these data points. This allows to query the observations belonging to those regions in the following iterations.

Recently, Yin et al. (2023) proposed an adaptive data stream classification method based on microclustering. After initializing micro-clusters from the initial training data, they collected new labels using a mixed strategy that combines random sampling with a class-weighted margin score. Then, the micro-cluster learning model is dynamically updated to adapt to the presence of concept drifts.

Another approach that tries to exploit the clustering nature of the incoming observations has been proposed by Mohamad et al. (2018), with the use of bi-criteria active learning algorithm that considers both density in the input space and label uncertainty. The density-based criterion makes use of the growing Gaussian mixture model proposed (GGMM) by Bouchachia and Vanaret (2014), which is used to find clusters in the data and estimate its density. This model creates a new cluster when a new data point \mathbf{x}_t has a Mahalanobis distance greater than a given closeness threshold from the nearest cluster, among the currently available ones. A flowchart describing the main steps of the GGMM is depicted in Fig. 10.

A Bayesian logistic regression model is used for addressing the label uncertainty criterion and the concept drift. As the classifier parameters \mathbf{w}_t are assumed to evolve over time, the

model is incrementally updated using a discrepancy measure, which is computed as the difference between the uncertainty of the model in \mathbf{x}_t before and after the true label y_t is added to the training set. The query strategy follows the b-sampling approach, in trying to sample, with high probability, the observations that contribute the most to the current error. The combination of density and uncertainty is also employed by Liu et al. (2021), who proposed a cognitive dual query strategy for online active learning in the presence of concept drifts and noise. The local density measure is used to obtain representative instances and the uncertainty criterion aim to select data points where the classifier is less confident. The cognitive aspect takes into account Ebbinghaus's law of human memory (Ebbinghaus, 2013) to determine an optimal replacement policy. The proposed strategy tries to tackle both gradual and abrupt drifts. The drift is generally considered as a change in the underlying joint probability distribution from one time step t to another, namely $p_t(\mathbf{x}, y) \neq p_{t+1}(\mathbf{x}, y)$. The local density of an observation \mathbf{x}_t is defined by the number of times that \mathbf{x}_t is the nearest neighbor of other instances (Ienco et al., 2014). Since we are in an online framework, the authors proposed to measure the local density using a sliding window model, referred to as a cognition window. Based on the concept of memory strength, the model determines when the current window is full and needs to be updated. Finally, the labeling decision is taken by using two thresholds, one for the local density and one for the classifier uncertainty.

A different sliding window-based online active learning strategy is the one proposed by Kurlej and Woźniak (2011). The authors proposed a sliding window approach based on a nearest neighbors classifier. The reference set for the k-nearest neighbors model is a window, and it is updated in two ways: in a first-in-first-out manner or using the examples selected by the active learning strategy. Since the reference set is updated over time, this method can effectively deal with concept drift and time-varying data streams. The sampling strategy is also based on two criteria. The first one is similar to the margin-based approaches, an instance is queried if it has a low distance from two observations belonging to different classes. The second criterion, similar to the greedy sampling strategy, seeks observations that have a large minimum distance from the observations in the current reference set. Both criteria are implemented by setting a threshold on the distances.

A simpler approach for taking into account the time-varying aspect of evolving data stream is to force the model to focus on the most recent observations. Along these lines, Chu et al. (2011) propose a framework based on a Bayesian probit model and a time-decay variant. Online Bayesian learning is used to maintain a posterior distribution of the weight vector of a linear classifier over time \mathbf{w}_t , and the time-decay strategies are employed to tackle the concept drift and give more importance to recent observations. They also propose an online approximation technique that can handle weighted examples, which is based upon Minka (2001). They tested different sampling strategies, built upon an online probit classifier. The instance selection criteria are based on entropy, function-value, and random sampling.

3.3 Evolving fuzzy systems approaches

An alternative way to take into account the time-varying nature of evolving data streams is the use of evolving fuzzy systems (EFS) (Lughofer, 2011), which are soft computing techniques that can efficiently deal with novelty and knowledge expansion. EFS are self-developing, self-learning fuzzy rule-based or neuro-fuzzy systems that self-adapt both their parameters and their structure on-the-fly. They try to mimic human-like reasoning by modeling it with a dynamically developing fuzzy rule-based structure and implementing it utilizing data streams using a formal learning process. The basic rule structure of a fuzzy model is given by




Fig. 11 Learning modules of an EFS (Ge & Zeng, 2020)

$$\begin{aligned} \text{Rule}_i : & \text{if } (x_1 \text{ is } X_{i1}) \text{ and } \dots \text{ and } (x_n \text{ is } X_{in}) \\ & \text{then } (y_i = a_{i0} + a_{i1}x_1 + \dots + a_{in}x_n) \end{aligned} \quad (35)$$

where Rule_i with $i = 1, 2, \dots, R$ is one of several fuzzy rules in the current rule base; x_j ($j = 1, 2, \dots, n$) are input variables; y_i denotes the output of the i th fuzzy rule; X_{ij} denotes the j th prototype (focal point) of the i th fuzzy rule; a_{ij} denotes the j th parameter of the i th fuzzy rule. For a more thorough discussion on EFS and their use in online learning, please see (Lughofer, 2017, 2011; Ge & Zeng, 2020; Gu et al., 2022). The main components of an EFS are shown in Fig. 11. The two key components of an EFS are the structure evolving scheme, which contains the rule generation and simplification modules, and the parameters updating scheme. The rule generation module defines when a new rule needs to be added to the current model. The rule merging and pruning steps simplify the models by removing redundant rules and combining two rules when their similarity is larger than a given threshold. The parameter updating modules are used to keep track of the model evolution. These learning modules are used to update the EFS every time a new labeled example (\mathbf{x}_t, y_t) is made available.

The first single-pass active learning approach based on the use of evolving classification models has been proposed by Lughofer (2012). The proposed algorithm is based on two key concepts, conflict and ignorance. The former is related to an incoming data point lying close to the boundary between any two classes; the latter considers the distance of the incoming observation from the currently labeled training set, in the feature space. This suggests that the data point falls within a region that has not been thoroughly explored by the learner. Later on, Lughofer and Pratama (2018) also proposed the first online active learning approach for evolving regression models. Similarly to their previous work (Lughofer, 2012), the authors consider the ignorance about the input space in the instance selection criterion. Moreover, they also consider the uncertainty in the model outputs and in the model parameters. The predictive uncertainty is assessed in terms of confidence intervals using locally adaptive error bars. The error bars are inspired by Škrjanc (2009) and the authors propose a new merging approach for dealing with the case of overlapping fuzzy rules. The uncertainty in the model parameters is instead evaluated using the A-optimality criterion, which will be discussed in Sect. 3.4 together with other alphabetic optimality criteria. Instead of leveraging the uncertainty about the output, Pratama et al. (2015) set a dynamic threshold based on the variable uncertainty strategy introduced by Zliobaite et al. (2014) while trying to address the what-to-learn question in the training of a recurrent fuzzy classifier. The key idea is that the model is iteratively retrained using data points that fall within rules with low support, which were formed using the smallest amount of observations. Recently, Lughofer and Škrjanc (2023) proposed an online active learning strategy for fuzzy models based on three criteria.

- D-optimality in the consequent space to reduce parameter uncertainty, as in Cacciarelli et al. (2022b).

- Overlap degree in the antecedent space to reduce the number of data points lying in the overlap regions of two different rules.
- Novelty content in the antecedent space, indicating the required knowledge expansion through rule evolution.

A different kind of threshold, based on the spherical potential theory, has been suggested by Subramanian et al. (2014), with the proposal of a meta-cognitive component that evaluates the novelty content of incoming data points. This is done using a knowledge measure represented by the spherical potential, which has been thoroughly investigated in kernel-based approaches (Hoffmann, 2007). The spherical potential is used to set a threshold and decide whether to add a new rule to capture the knowledge in the current sample. It should be noted that the authors also used a threshold based on the prediction error, which could not be used with scarcity of labels. The prediction error is assessed using the hinge loss error function (Suresh et al., 2008; Zhang, 2004).

Fuzzy models have also been used to solve computer vision tasks. Weigl et al. (2016) analyze the visual inspection quality control case, which is also considered by Rožanec et al. (2022). They assess the usefulness of the images in a single-pass manner, but the instances that are selected to be queried are accumulated in a buffer, which is later on assigned to an oracle for labeling. Choosing the size of the buffer represents a trade-off problem between timely updating the classifier and requiring continuous interventions from a human annotator. The active learning strategy works by setting a threshold on the certainty of the model with regards to the incoming data points. The authors take into account two model classes, a random forest classifier and an evolving fuzzy classifier. When using random forest, certainty is computed using the best-versus-second-best margin score. Instead, when using evolving fuzzy classifiers, the sample selection criterion takes into account the conflict and ignorance concepts as in Lughofer (2012).

Finally, Cernuda et al. (2014) combine the use of fuzzy models with a sampling approach inspired by the multivariate statistical process control literature. Indeed, using a latent structure model, they propose a query strategy based on the Hotelling T^2 and the squared prediction error (SPE) statistics, which have been extensively used in anomaly detection problems (Cacciarelli & Kulahci, 2022; Gajjar et al., 2018; Vanhatalo & Kulahci, 2016; Vanhatalo et al., 2017). Ge (2014) used these statistics for pool-based active learning in conjunction with a principal component regression model. The key idea is to use the Hotelling T^2 and the SPE statistics to measure the distance between the currently labeled training set and a new unlabeled data point. A high value in one of the two statistics would most likely suggest that the new observation is violating the current model, and thus its inclusion in the training set could bring some valuable information. Similarly, Cernuda et al. (2014) use the Hotelling T^2 and the SPE statistics with a partial least squares model. Then, when a new observation is added to the training set, they retrain a TS fuzzy model using a sliding window approach.

3.4 Experimental design and bandit approaches

Optimal experimental design (Karlin & Studden, 1966) is a research field that is closely related to active learning. It deals with the design of experiments that allow for efficient estimation of model parameters or improved prediction performance while minimizing the number of required labeled examples, also referred to as the number of runs N . Many optimality criteria have been developed in striving to strike a balance between efficient use of resources and ensuring good performance of the model. The traditional framework of optimal experimental designs focuses on linear regression models of the form

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon} \tag{36}$$

where, given d input variables, \mathbf{y} is a $N \times 1$ vector of response variables, \mathbf{X} is a $N \times d$ model matrix, $\boldsymbol{\beta}$ is a $d \times 1$ vector of regression coefficients, and $\boldsymbol{\varepsilon}$ is a $N \times 1$ vector representing the noise, with covariance matrix $\sigma^2\mathbf{I}$. If the matrix $\mathbf{X}^\top\mathbf{X}$ is of full rank, an ordinary least square (OLS) estimator for $\boldsymbol{\beta}$ can be obtained using

$$\hat{\boldsymbol{\beta}} = \left(\mathbf{X}^\top\mathbf{X}\right)^{-1}\mathbf{X}^\top\mathbf{y} \tag{37}$$

In general, design optimality criteria leverage the information contained in the moment matrix, which is defined as $\mathbf{M} = \mathbf{X}^\top\mathbf{X}/N$. The matrix $\mathbf{X}^\top\mathbf{X}$ plays a crucial role in the estimation of the model coefficients $\boldsymbol{\beta}$, and it is important to perceive information about the design geometry. Indeed, with Gaussian noise characterized by $\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \sigma^2\mathbf{I})$, we know that

$$\hat{\boldsymbol{\beta}} \mid \mathbf{X} \sim \mathcal{N}\left(\boldsymbol{\beta}, \left(\mathbf{X}^\top\mathbf{X}\right)^{-1}\sigma^2\right) \tag{38}$$

and we can define a $100(1 - \alpha)\%$ confidence ellipsoid related to the solutions of $\boldsymbol{\beta}$ using

$$\frac{(\mathbf{b} - \hat{\boldsymbol{\beta}})^\top (\mathbf{X}^\top\mathbf{X}) (\mathbf{b} - \hat{\boldsymbol{\beta}})}{ds^2} \leq F_{\alpha,d,N-d} \tag{39}$$

where s^2 represents the residual mean square, $F_{\alpha,d,N-d}$ is the $100(1 - \alpha)$ percentile derived from the Fisher distribution, and \mathbf{b} indicates all the possible vectors that could be the true model parameter $\boldsymbol{\beta}$. The ellipsoid can also be expressed as $(\mathbf{b} - \hat{\boldsymbol{\beta}})^\top (\mathbf{X}^\top\mathbf{X}) (\mathbf{b} - \hat{\boldsymbol{\beta}}) \leq C$, where $C = ds^2 F_{\alpha,d,N-d}$. The volume of this ellipsoid is inversely proportional to the square root of the determinant of $\mathbf{X}^\top\mathbf{X}$, and the length of its axes is proportional to $1/\lambda_i$, where λ_i represents the i th eigenvalue of $\mathbf{X}^\top\mathbf{X}$, with $i = 1, \dots, d$. The so-called alphabetic optimality criteria pursue efficient designs by exploiting these properties (Kiefer, 1959). The most commonly employed optimality criteria for good parameter estimation are A-, D- and E-optimality:

- *A-optimality*. This criterion pursues good model parameter estimation by minimizing the sum of the variances of the regression coefficients. Knowing that the coefficients variances appear on the diagonal of the matrix $(\mathbf{X}^\top\mathbf{X})^{-1}$, it can be shown that an A-optimal design is given by a design \mathcal{D}^* that satisfies $\min_{\mathcal{D}} \text{tr}[\mathbf{M}(\mathcal{D})]^{-1} = \text{tr}[\mathbf{M}(\mathcal{D}^*)]^{-1}$.
- *D-optimality*. This criterion takes into account both the variance and covariance of the regression coefficients, directly minimizing the total volume of the confidence ellipsoid (Myers et al., 2016). A D-optimal design is given by a design \mathcal{D}^* that satisfies $\max_{\mathcal{D}} |\mathbf{M}(\mathcal{D})| = |\mathbf{M}(\mathcal{D}^*)|$ (John & Draper, 1975).
- *E-optimality*. This strategy tries to shrink the ellipsoid by minimizing the maximum eigenvalue of the covariance matrix.

The geometrical intuition behind these criteria is illustrated, in the two-dimensional case, in Fig. 12.

Finally, there are also optimality criteria that focus on developing models with good predictive properties. Within this class, *G-optimality* represents a criterion that is used to seek protection against the worst-case prediction variance in a region of interest \mathcal{R} . This is achieved by solving

$$\min_{\mathcal{D}} \left[\max_{\mathbf{x} \in \mathcal{R}} v(\mathbf{x}) \right] \tag{40}$$




Fig. 12 Confidence ellipsoid around the model parameters and optimality criteria: A-optimality (a) shrinks the hyperrectangular enclosing the confidence ellipsoid (Asprey & Macchietto, 2002; Galvanin, 2010), D-optimality (b) aim to shrink the total volume of the ellipsoid, and E-optimality (c) tries to reduce the length of the longest axis (Jamieson, 2018)

where $v(\mathbf{x})$ represents the scaled prediction variance of the current model in the data point \mathbf{x} , which can be computed as

$$v(\mathbf{x}) = N\mathbf{x}^{(m)\top} (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{x}^{(m)} \quad (41)$$

where $\mathbf{x}^{(m)}$ represents the data point where the variance is being estimated, expanded to the model form. It should be noted that G-optimality can be highly influenced by anomalous observations, as it protects against the highest possible variance over all the region \mathcal{R} . This issue can be tackled by using *I-* or *V-optimality*, which estimate the overall prediction variance over \mathcal{R} by integrating or averaging, respectively. For a more extensive discussion on optimal designs, please see Montgomery (2012) or Myers et al. (2016).

The use of optimality criteria has proven to be highly beneficial in offline experimental design, allowing practitioners to pre-determine the location of each design point with ease. However, these methods require modification to be applied in a stream-based scenario where data points arrive sequentially. A common approach for obtaining a near-optimal design with streaming observational data is represented by thresholding. Riquelme (2017) proposed a thresholding algorithm for online active linear regression, which is related to the A-optimality criterion. Their approach uses a norm-thresholding algorithm, where only observations with large, scaled norms are selected. The design is augmented with the observations \mathbf{x} whose norm exceeds a threshold Γ given by

$$\mathbb{P}(\|\mathbf{x}\| \geq \Gamma) = \alpha \quad (42)$$

where α is the ratio of observations we are willing to label out of the incoming data stream. Another approach related to the A-optimality criterion was proposed by Fontaine et al. (2021), who studied online optimal design under heteroskedasticity assumptions, with the objective of optimally allocating the total labeling budget between covariates in order to balance the variance of each estimated coefficient. Cacciarelli et al. (2022b) further extended the thresholding approach introduced by Riquelme (2017) by proposing a conditional D-optimality (CDO) algorithm. The terms conditional refers to the fact the design is marginally optimal, given an initial set of labeled observations to be augmented. The main steps of the CDO approach are reported in Algorithm 5. The authors exploited the connection between D-optimality and prediction variance previously highlighted by Myers et al. (2016). The sampling strategy selects observations by setting a threshold Γ given by

$$\mathbb{P} \left(\mathbf{x}_t^\top \left(\mathbf{X}^\top \mathbf{X} \right)^{-1} \mathbf{x}_t \geq \Gamma \right) = \alpha \tag{43}$$

where \mathbf{X} is the current set of labeled observations and \mathbf{x}_t is the data point that is currently under evaluation. The threshold is estimated using kernel density estimation (KDE) on a set of j unlabeled observations, which are taken passively from the data stream without querying any label. This provides an initial set of data, referred to as warm-up set, that can be used to estimate the covariance matrix and the threshold.

Algorithm 5 Online active learning using CDO

Require: an initial random design \mathbf{X} , a data stream \mathbf{S} , a warm-up length j , a sampling rate α , a budget B

```

 $t \leftarrow 1$  ▷ Timestamp
 $c \leftarrow 0$  ▷ Labeling cost
Set  $\mathbf{W} = \emptyset$  ▷ Warm-up set to estimate  $\Sigma$  and  $\Gamma$ 
while  $t \leq j$  do
  Observe incoming data point  $\mathbf{x}_t \in \mathbf{S}$ 
  Select  $\mathbf{x}_t : \mathbf{W} = \mathbf{W} \cup \mathbf{x}_t$ 
   $t \leftarrow t + 1$ 
end while
Estimate the covariance matrix  $\Sigma$  of  $\mathbf{W}$  and perform eigendecomposition  $\Sigma = \mathbf{U}\Lambda\mathbf{U}^\top$ 
Whiten the initial design by computing  $\mathbf{Z} = \Lambda^{-1/2}\mathbf{U}^\top\mathbf{X}$ 
Whiten the warm-up observations by computing  $\mathbf{V} = \Lambda^{-1/2}\mathbf{U}^\top\mathbf{W}$ 
Estimate  $\Gamma$  using KDE on  $\mathbf{V}$  with the desired sampling rate  $\alpha$  using Equation 43 with  $\mathbf{Z}$  and  $\mathbf{V}$ 
while  $c \leq B$  and  $t \leq |\mathbf{S}|$  do
  Observe incoming data point  $\mathbf{x}_t \in \mathbf{S}$ 
  Whiten  $\mathbf{x}_t$  by computing  $\mathbf{z}_t = \Lambda^{-1/2}\mathbf{U}^\top\mathbf{x}_t$ 
  if  $\mathbf{z}_t^\top (\mathbf{Z}^\top \mathbf{Z})^{-1} \mathbf{z}_t \geq \Gamma$  then
    Ask for the label  $y_t$  and augment the labeled dataset:  $\mathbf{Z} \leftarrow \mathbf{Z} \cup \mathbf{z}_t$ 
     $c \leftarrow c + 1$  ▷ Pay for the label
    Update threshold  $\Gamma$  to measure the prediction variance of the enlarged design
  else
    Discard  $\mathbf{x}_t$ 
  end if
   $t \leftarrow t + 1$ 
end while

```

Cacciarelli et al. (2023) also investigated how the presence of outliers affect the performance of online active linear regression strategies. They showed how the design optimality-based sampling strategies might be attracted to outliers, whose inclusion in the design eventually degrades the predictive performance of the model. This issue can be tackled by bounding the search area of the learner with two thresholds, as in

$$\mathbb{P} \left(\Gamma_1 \leq \mathbf{x}_t^\top \left(\mathbf{X}^\top \mathbf{X} \right)^{-1} \mathbf{x}_t \leq \Gamma_2 \right) = \alpha \tag{44}$$

where the choice of Γ_2 represents a trade-off between seeking protection against outliers and exploring uncertain regions of the input space.

The norm-thresholding approach was also extended by Riquelme et al. (2017a) to the case where the learner tries to estimate uniformly well a set of models, given a shared budget. This scenario is similar to a multi-armed bandit (MAB) problem where the learner wants to estimate the mean of a finite set of arms by setting a budget on the number of allowed pulls (Ruan et al., 2020; Audibert & Munos, 2010; Jamieson & Nowak, 2014; Soare et al., 2013). The authors propose a trace upper confidence bound (UCB) algorithm to simultaneously estimate the difficulty of each model and allocate the shared labeling budget

proportionally to these estimates. UCB is a common algorithm used in MAB problems to balance exploration and exploitation (Carpentier et al., 2015; Garivier & Moulines, 2008), which takes into account the predicted mean value and the predicted standard deviation, weighted by an adjustable parameter (Thompson et al., 2022). This allows to balance the exploitation of data points with a high predicted value and the exploration of areas with high uncertainty.

In general, MAB problems can be seen as a special case of sequential experimental design, where the goal is to sequentially choose experiments to perform with the aim of maximizing some outcome. The typical framework of a MAB problem can be regarded as an optimization problem where the learner must identify the option or arm with the highest reward, among a set of available arms characterized by different reward distributions. Both MAB and active learning paradigms involve a sequential decision-making process where the learner aims to maximize a reward or improve model accuracy by selecting an arm to pull or a data point to label, respectively, and receiving feedback (in the form of a reward or label request) for each selection. There are two main approaches to tackle MAB problems:

- *Regret minimization.* This approach is coherent with the objective of maximizing the cumulative reward observed over many trials. In this case, the learner must balance exploration, namely trying out different arms to learn more about the reward distributions, with exploitation, i.e., using current knowledge to choose the most promising arm. These kinds of algorithms strike a balance between learning a good model and obtaining high rewards. A few examples might be treatment design, online advertising and recommender systems.
- *Pure exploration.* In this case, we are interested in finding the most promising arm, with a certain confidence or given a fixed budget on the number of pulls. To do so, the objective is to learn a good model while minimizing the number of measurements or labels required. This scenario is suggested in circumstances where, due to safety constraints, we are not given complete freedom to change the variable levels and we are mostly interested in understanding the underlying model governing the system. Possible examples include drug discovery or soft sensor development (Fortuna et al., 2007; Shi & Xiong, 2018; Chan et al., 2018; Tang et al., 2018).

The pure exploration approach is particularly useful when coupled with the study of linear bandits, which are a type of contextual bandit algorithms that assume a linear relationship between the features of the context and the expected reward of each arm. In this type of problem, when an arm $\mathbf{x} \in \mathcal{X}$ is pulled, the learner observes a reward $r(\mathbf{x})$ that depends on an unknown parameter $\boldsymbol{\theta}^* \in \mathbb{R}^d$ according to the linear model

$$r(\mathbf{x}) = \mathbf{x}^\top \boldsymbol{\theta}^* + \varepsilon \quad (45)$$

where ε is a zero-mean i.i.d. noise. This is similar to active linear regression in that, in both cases, the learner aims to select the most informative data points to learn about the underlying model or system (Audibert & Munos, 2010; Jamieson & Nowak, 2014). Soare et al. (2014), investigated this problem, in the offline setting, using the G-optimality criterion and a newly proposed \mathcal{XY} -allocation algorithm. Jedra and Proutiere (2020) proposed a fixed-confidence algorithm for the same problem, while Azizi et al. (2022) analyzed the fixed-budget case, extending the framework to the case where the underlying model is represented by a generalized linear model (Filippi et al., 2010). An interesting variant of this problem is presented in the study of transductive experimental designs. A transductive design is a problem where we can pull arms from a set $\mathcal{X} \in \mathbb{R}^d$, with the objective of identifying the best arm or improve the predictions over a separate set of observations $\mathcal{Z} \in \mathbb{R}^d$, which is

given, in an unlabeled form, beforehand. A practical example of this case is when we are trying to infer the user preferences over a set of products, but we can only do that by pulling arms from a limited set of free trials. Alternatively, we might be interested in estimating the efficacy of a drug over a certain population, while doing experiments on a population with different characteristics. This problem has been tackled with an active learning approach by Yu et al. (2006), with the idea of exploiting unlabeled data points in \mathcal{Z} while evaluating the informativeness of the data points in \mathcal{X} . The transductive case of sequential experimental design has been explored by Fiez et al. (2019), but instead of performing active learning, they were interested in inferring the best reward over \mathcal{Z} , only pulling the arms in \mathcal{X} . Finally, this has been extended to the online scenario by Camilleri et al. (2021), balancing the trade-off between time complexity and label complexity, namely between the number of unlabeled observations spanned and the number of labels queried in order to stop the learning procedure and declare the best-arm.

In addition to MAB, reinforcement learning-based approaches can also be applied to active learning in order to optimize a decision-making policy that balances the exploration of uncertain data with the exploitation of information learned from previous observations. This can be particularly useful in applications where the goal is to maximize the expected cumulative reward over time, such as in robotics or game playing. Compared to MAB, reinforcement learning-based approaches offer a more general and flexible framework for active learning, allowing for a wider range of problem formulations and feedback signals (Menard et al., 2021; Fang et al., 2017; Rudovic et al., 2019). One approach to combining active learning and reinforcement learning is through modeling the sampling routine as a contextual-bandit problem, as proposed by Wassermann et al. (2019). In this approach, the rewards are based on the usefulness of the query behavior of the learner. The key intuition behind the use of reinforcement learning in online active learning is that the learner gets feedback after the requested label, based on how useful the request actually was. In contrast to the traditional active learning view, where most of the effort is dedicated to the instance selection phase, the learner is penalized ex-post for querying useless instances. The learner gets a positive reward ρ^+ if it asks for the label when it would have otherwise predicted the wrong class, and a negative reward ρ^- when querying was unnecessary as the model would have predicted the right label. The contextual bandit problem is implemented by building an ensemble of different models, with each expert suggesting whether to query or not based on whether its prediction certainty exceeds a threshold Γ . The models are assigned a decision power based on how past suggestions were rewarded and how coherent they were with the other experts' suggestions. When an observation is sent to the oracle for labeling, the reward is computed, and the objective function of the learner is to maximize the total reward over a time horizon T .

Another reinforcement learning-based approach has been proposed by Woodward and Finn (2017). They considered the case where at each time step t the learner needs to decide whether to predict the label of the unlabeled data point \mathbf{x}_t or pay to request its label y_t . The reinforcement learning framework is used to find an optimal policy $\pi^*(s_t)$ that takes into account the cost of asking for a label and the cost of making an incorrect prediction, where s_t represents the state that is given in input at the time t to a policy $\pi(s_t)$ that outputs the suggested action a_t . The authors approximate the action-value function using a long short-term memory (LSTM) neural network with a linear output layer. The optimal policy is determined by maximizing the long-term reward, after assigning a reward to a label request R_{req} , a correct prediction R_{corr} , and an incorrect prediction R_{inc} . It should be noted that R_{corr} and R_{inc} should be negative rewards, as they are associated with costly actions.

4 Evaluation strategies

The use of active learning approaches is becoming increasingly common in machine learning, allowing models to be trained more efficiently by selecting the most informative examples for labeling. To evaluate the performance of these approaches, it is typical to compare them to a passive random sampling strategy by generating learning curves that plot the model performance (e.g., accuracy, F1 score, or root mean square error) on a holdout test set over the number of labeled examples used for training. Learning curves are a useful tool for comparing the asymptotic performance of different strategies and their sample efficiency, with the slope of the curve reflecting the rate at which the model performance improves with additional labeled examples. A steeper slope indicates a more sample-efficient strategy. When multiple sampling strategies are being compared, a visual inspection of the learning curves may not be sufficient, and more rigorous statistical tests may be necessary. Reyes et al. (2018) recommend the use of non-parametric statistical tests to analyze the effectiveness of active learning strategies for classification tasks. The sign test (Steel, 1959) or the Wilcoxon signed-ranks test (Wilcoxon, 1945) can be used to compare two strategies, while the Friedman test (Friedman, 1940), the Friedman aligned-ranks test (Hodges & Lehmann, 1962), the Friedman test with Iman-Davenport correction (Iman & Davenport, 1980), or the Quade test (Quade, 1979) can be used when evaluating more than two strategies. These statistical tests can provide insight into whether the difference in performance between the active learning and passive random sampling strategies is statistically significant.

Algorithm 6 Prequential evaluation for online active learning

Require: an initial model w_0 , a data stream S , a budget B , an active learning strategy Q .

```

 $t \leftarrow 1$  ▷ Timestamp
 $P \leftarrow \emptyset$  ▷ Storing predictions
while  $c \leq B$  and  $i \leq |S|$  do
  Observe the data point  $x_t \in S$ 
  Predict the label  $\hat{y}_t$  and store it:  $P \leftarrow P \cup \hat{y}_t$ 
  if  $Q(x_t) = \text{True}$  then ▷ Sampling decision
    Ask for the true label  $y_t$  and update the model
     $c \leftarrow c + 1$  ▷ Pay for the label
  else
    Discard  $x_t$ 
  end if
   $t \leftarrow t + 1$ 
end while

```

Overall, the use of learning curves and statistical tests can provide valuable insights into the effectiveness and efficiency of different active learning strategies. By understanding the statistical significance of differences in performance between these strategies, researchers can make informed decisions about which approaches are more effective for a particular task or dataset. Furthermore, the choice of the evaluation scheme is crucial when assessing the performance of active learning approaches. If we use an evaluation scheme based on a holdout test set, at each learning step t the performance of the model is assessed using the same test set. This can be a reasonable approach if we are dealing with a stationary data stream, which does not evolve over time. Under these assumptions, using the same test set we might be able to better assess the prediction improvement as more labeled examples are included in the design. However, this approach might not be ideal when dealing with drifting data streams. In these circumstances, a prequential evaluation scheme can be more

Table 1 Evaluation strategies

Evaluation strategy	Works
Holdout test set	Desalvo et al. (2021), Wassermann et al. (2019), Rožanec et al. (2022), Narr et al. (2016), Ferdowski et al. (2013), Bordes et al. (2005), Suzuki et al. (2021), Ghassemi et al. (2016), Qin et al. (2021), Woodward and Finn (2017), Riquelme et al. (2017b), Cacciarelli et al. (2022b), Cacciarelli et al. (2023), Manjah et al. (2023)
Prequential/test-then-train	Zhang et al. (2022), Pham et al. (2022), Castellani et al. (2022), Chu et al. (2011), Zhang et al. (2018), Krawczyk et al. (2018), Xu et al. (2016), Mohamad et al. (2020), Weigl et al. (2016), Ienco et al. (2013), Zhang et al. (2020a)

useful to monitor the evolution of the prediction error over time (Suárez-Cetrulo et al., 2021; Cerqueira et al., 2020; Tieppo et al., 2022; Cacciarelli & Boresta, 2021). In online learning, prequential evaluation is also referred to as test-then-train approach, and it involves using each incoming instance first to measure the prediction error, and then to be included in the training set (Suárez-Cetrulo et al., 2023). The main steps of the test-then-train approach are reported in Algorithm 6. The key idea is that at each time step t , we first test the model by making a prediction, then we decide whether to query the true labels and finally we update our model.

An in-depth analysis and discussion between the use of a holdout test set and the prequential evaluation scheme for streaming data has been provided by Gama et al. (2009, 2013), who suggested the use of a prequential evaluation scheme with forgetting mechanisms. For scenarios with imbalanced data streams, a specialized prequential variant of the area under the curve metric has been proposed by Brzezinski and Stefanowski (2015, 2017). From an implementation perspective, Bifet et al. (2010) developed an open source software suite called MOA for data stream mining, which includes both the holdout and prequential strategies. This framework has found widespread application in the evaluation of online active learning strategies, as evidenced by the studies conducted by Liu et al. (2021), Shan et al. (2019), Weigl et al. (2016), Zhang et al. (2020a), Alabdulrahman et al. (2016).

In Table 1, we categorize the studies based on the experimental protocols they employed to evaluate the sampling strategies. The table exclusively includes approaches where the evaluation strategy was explicitly defined. In most cases, when assessing active learning methods in the context of drifting data streams, a prequential approach is favored. Conversely, for scenarios where the methods are ill-suited to handle concept drifts, holdout test sets tend to be the preferred choice. In approaches not featured in the table, the evaluation strategies exhibited some variations or lacked explicit specification. For instance, in the work by Fujii and Kashima (2016), their evaluation strategy involved training models on the queried data and subsequently testing them with the entire dataset. This approach differs from the conventional test-then-train paradigm since, in this case, models are tested on data they encountered during training, at least in part. Another example is found in Zhu et al. (2007), who utilized a window-based approach, assessing prediction accuracy across all observations in the current batch. On a different note, Hao et al. (2018a) employed the per-round regret metric, which quantifies the loss difference between the forecaster and the best expert at each iteration of the active learning process. In some instances, none of the previously mentioned methods were employed, as the analysis took a more theoretical perspective. This is exemplified by the works of Dasgupta et al. (2005), Chae and Hong (2021), Huang et al. (2022). Lastly,




Fig. 13 Low-cost active spam filtering (Sculley, 2007)

bandit algorithms employed a distinct evaluation approach, often aiming to identify the most promising arm with a fixed confidence or budget. In the fixed confidence setting, performance typically hinges on comparing label complexity to problem dimensionality or the number of arms pulled, as observed in Fiez et al. (2019). Alternatively, regret or error metrics were evaluated against the required number of trials, as demonstrated in the studies by Riquelme et al. (2017a), Sudarsanam and Ravindran (2018), Fontaine et al. (2021).

5 Real-world applications and challenges

5.1 Applications

Online active learning has been recognized as a powerful technique in scenarios where data is arriving at a high velocity, labeling data is expensive, and it is infeasible to store all the unlabeled data before making a decision about which observations to query to update the model. In particular, these techniques have proven particularly useful in dynamic and ever-evolving environments, where models need to adapt to new data in real-time, by selectively querying the most informative instances. One of the first real-world applications of online active learning has been presented by Sculley (2007), who investigated the scenario of low-cost active spam filtering (Fig. 13) where a filter is updated online by selecting the most informative emails in real time. Another application of online active learning in the field of IT has been recently presented by Zhang et al. (2020a). They analyzed the scenario of network protocol identification and proposed a method (presented in Sect. 3.2) to select the most representative instances on the fly and adapt the model to dynamic data distributions.

Computer vision is another interesting area where online active learning can be applied. Deep learning models require a large amount of annotated data, making manual annotation of thousands of images one of the most challenging aspects of model development. However, it is important to note that the most effective deep active learning methods proposed so far are not easily adaptable to a stream-based setting. Many of these methods involve clustering or measuring pairwise similarity among image embeddings (Sener & Savarese, 2017; Agarwal et al., 2020; Ash et al., 2019; Citovsky et al., 2021; Prabhu et al., 2020), which cannot be easily done in a single-pass manner. As a result, most online applications of active learning in computer vision rely on the use of traditional models with uncertainty-based sampling. Narr et al. (2016) analyze the stream-based active learning problem for the classification of 3D objects. They used a mondrian forest classifier (Lakshminarayanan et al., 2014), which is an efficient alternative of random forest for the online learning scenario, and selected images with high classification uncertainty to be labeled. Rožanec et al. (2022) used online active learning to reduce the data labeling effort while performing vision-based process monitoring. Initially, features are extracted from the images using a pre-trained ResNet-18 model (He et

al., 2015) and then, using the mutual information criterion (Kraskov et al., 2004), only \sqrt{n} features (Hua et al., 2005) are retained to fit an online classifier, where n is the total number of observations in the training set. The authors combine a simple active learning strategy based on model uncertainty with five streaming classification algorithms, including Hoeffding tree (Hulten et al., 2001), Hoeffding adaptive tree (Bifet & Gavaldà, 2009), stochastic gradient tree (Gouk et al., 2019), streaming logistic regression, and streaming k-nearest neighbors. Recently, Saran et al. (2023) proposed a novel approach to streaming active learning with deep neural networks. Given a neural network with f with parameters θ , last-layer parameters θ_L , and the cross-entropy function ℓ , they compute the gradient representation of the data point \mathbf{x}_t , which is given by

$$g(\mathbf{x}_t) = \frac{\partial}{\partial \theta_L} \ell(f(\mathbf{x}_t; \theta), \hat{y}_t) \quad (46)$$

where $\hat{y}_t = \operatorname{argmax} f(\mathbf{x}_t; \theta)$. Then, the data points to be included in the batch for training the model are chosen by using a probability p_t proportional to the contribution of the current example to the covariance matrix of the examples collected so far, as in

$$p_t \propto \det(\hat{\Sigma}_t + g(\mathbf{x}_t)g(\mathbf{x}_t)^\top) \quad (47)$$

where $\hat{\Sigma}_t$ is the covariance matrix of the data points that have been selected to be included in the current batch, up to the time step t .

Online active learning has also been explored for object detection tasks. Manjah et al. (2023) proposed a stream-based active distillation (SBAD) framework by combining the concepts of active learning and self-supervision as described in Sect. 2.3. The SBAD framework enables the deployment of scalable deep-learning models as it does not rely on human annotators and takes into account the imperfection of the oracle when distilling knowledge from a large teacher model to a lightweight student. Indeed, the authors suggest setting a threshold on the confidence of the images and only querying images with high confidence in trying to avoid confirmation bias. The threshold is determined using a warm-up phase, similarly to the approach proposed by Cacciarelli et al. (2022b) presented in Algorithm 5. The SBAD pipeline for model development and evaluation is reported in Fig. 14.

The problem of performing active learning for object detection with streaming data has also been explored by Beck et al. (2023). In the case of a camera placed on an autonomous vehicle, the collected data encompasses various scenarios, including clear weather, foggy conditions, and rainy weather, all of which require the model to perform effectively. However, the frequency of these scenarios can vary significantly. In situations where one scenario is prevalent, a passive sampling strategy could tend to sample very few examples from the most rare slices. Instead, the proposed streamline approach by attempts to smartly allocate the budget to obtain more observations from the slices where the model is under-performing. The case of autonomous cars was also considered by Yan et al. (2023), who used a diversity-based online active learning strategy to reduce false alarm rate and learn unseen faults.

Another interesting industrial application has been recently presented by Ghiasi et al. (2023). They proposed a deployable framework that combines a thermodynamics-based compressor model and a Gaussian Process-based surrogate model with an online active learning module. The objective of the study was to minimize the power absorbed by the machine during the boil off process of centrifugal compressor. In the proposed framework, the simulator, the surrogate model, and the optimizer interact in real time to determine the new experimental points.

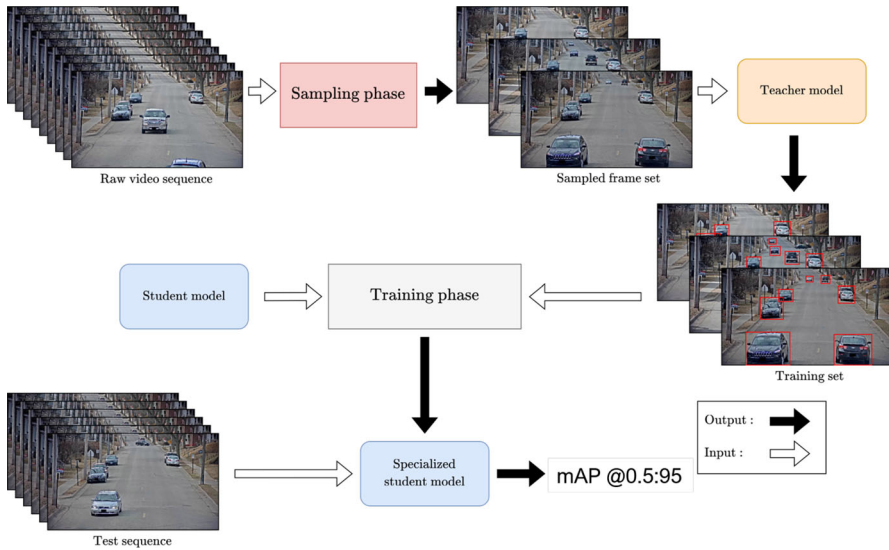


Fig. 14 SBAD framework (Manjah et al., 2023): sampling, fine-tuning and evaluation. The sampling is performed in a single-pass manner via thresholding

5.2 Challenges

When applying online active learning strategies to real-world problems, there are several potential issues to consider, including:

- **Algorithm scalability.** Online active learning algorithms need to be efficient and scalable to handle large datasets and high-velocity data streams. As the amount of data grows, the computational demands of active learning can become prohibitive, making it difficult to deploy in practice. The time required to make the sampling decision needs to be lower than the feed rate of the process being analyzed. If the algorithm is too slow, it may require a buffer, which reduces the benefits of online active learning.
- **Labeling quality.** Most online active learning strategies rely heavily on the quality of labeled data, which can be challenging to ensure in real-world scenarios. Human annotators may make errors, introduce biases, or interpret labeling instructions differently. For this reason, in real-life situations, it may be necessary to consider oracle imperfections like in the knowledge distillation case (Baykal et al., 2022). Another difficult aspect related to labeling quality is the delay or latency, which has been described in Sect. 2.2.3.
- **Data drift.** In real-world settings, data distributions may shift over time, making it challenging for models to adapt and continue providing accurate predictions. Changes in the data distribution may also affect the quality of the labeled data, as the criteria for selecting informative instances may become less effective. Methods from Sects. 3.2 and 3.3 should be used when dynamic and ever-changing behaviors are expected.
- **Model interpretability.** Besides simply asking for the most informative instances from a modeling perspective, it might be useful to provide additional information on why a particular instance is beneficial for improving the performance of the current model. In fields like healthcare and manufacturing this might help practitioners to improve their understanding of the underlying problem.

- *Evaluation.* When developing active learning methods from a research perspective, the different query strategies are evaluated assuming the ground-truth labels to be available for a held-out test set, or for the data stream being analyzed. However, in real life, the key motivation behind active learning is label scarcity and thus it might be difficult to thoroughly assess the effectiveness of the deployed sampling strategy.
- *Human-computer interaction.* In the context of active learning for data streams, the synergy between human labelers and computer systems plays a pivotal role in the labeling process. While the majority of online active learning methods focus on querying the most informative data points in real-time, we can distinguish between two distinct labeling scenarios:
 1. *Real-time annotation.* In most of the presented works, it is assumed that labels are immediately available when a data point is queried from the stream. This immediate access to true labels enables an optimized active learning routine, as the model can be promptly updated and can recommend exploration of new regions based on up-to-date information. However, this approach poses some implementation challenges that need to be addressed with the use of advanced data annotation tools (Feuz & Cook, 2013).
 2. *Postponed annotation.* There are cases where we must allow for a delay between data querying and labeling. For instance, methods that consider verification latency (Castellani et al., 2022; Pham et al., 2022) take into account the possibility of delayed labels. This is particularly relevant in situations where a physical quality inspection or medical treatment must occur before the label is revealed. Another example is in the training of deep neural networks, where real-time sampling from a data stream is necessary due to memory constraints (Manjah et al., 2023), but the labeling and model update phase may occur when a batch is collected, following a batch-mode active learning strategy (Ren et al., 2022).

6 Summary and future directions

This survey outlines the challenge of conducting active learning with data streams and investigates different approaches for selecting the most informative data points in real-time.

Table 2 provides a summary of the relevant state-of-the-art approaches, highlighting their main properties and settings. Our examination reveals that existing research has predominantly concentrated on creating online classification models, which can operate with both stationary and drifting data streams. However, there has been comparatively limited effort devoted to online active linear regression or dedicated to constructing online regression models in general.

We believe that there are several promising directions for future research in this field. First, we recommend further investigation into online active learning strategies specifically designed for regression models. Given the limited work in this area, there is a need for more advanced methods that can be applied to nonlinear models, beyond linear models or linear bandits. For example, there has been a recent spark of interest toward the use of Bayesian optimization for active learning in nonlinear regression problems (Mohamadi & Amindavar, 2020; Riis et al., 2022). Additionally, model-agnostic methods that can be applied to a variety of regression models could be valuable as they would provide a more

Table 2 Online active learning strategies: summary based on data processing capabilities, assumptions about the data stream, task of the model and model characteristics

Data processing	Data stream	Task	Model	Work(s)
Single-pass	Stationary	Classification	Single model	Cesa-Bianchi et al. (2004), Cesa-Bianchi et al. (2006), Dasgupta et al. (2005), Sculley (2007), Lu et al. (2016), Hao et al. (2018b), Ghassemi et al. (2016), Shah and Manwani (2020), Mohamad et al. (2020), Saran et al. (2023), Rožanec et al. (2022), Woodward and Finn (2017)
			Ensemble	Huang et al. (2022), Desalvo et al. (2021), Loy et al. (2012), Hao et al. (2018a), Chae and Hong (2021)
		Regression	Single model	Riquelme (2017), Fontaine et al. (2021), Cacciarelli et al. (2022b), Cacciarelli et al. (2023), Cacciarelli et al. (2022a)
	Drifting	Object detection Classification	Single model Single model	Manjah et al. (2023) Krawczyk et al. (2018), Castellani et al. (2022), Pham et al. (2022), Yin et al. (2023), Mohamad et al. (2018), Liu et al. (2021), Kurlej and Woźniak (2011), Chu et al. (2011)
Batch	Evolving	Classification	Ensemble	Zhang et al. (2020a), Shan et al. (2019), Zhang et al. (2018), Zhang et al. (2022)
		Regression	Single model	Lughofer (2012), Pratama et al. (2015)
	Stationary	Classification	Single model	Lughofer and Pratama (2018), Lughofer and Škrjanc (2023)
		Object detection	Single model	Bordes et al. (2005), Qin et al. (2021), Fujii and Kashima (2016)
Drifting	Classification	Single model	Beck et al. (2023) Cheng et al. (2023), Martins et al. (2023), Ienco et al. (2013), Zhang et al. (2023), Yan et al. (2023)	
Evolving	Classification	Ensemble	Zhu et al. (2007), Woźniak et al. (2023), Halder et al. (2023)	
			Single model	Subramanian et al. (2014), Weigl et al. (2016), Cernuda et al. (2014)

general solution to the problem. Second, we believe that there is potential for research into single-pass online sampling strategies for dynamic data streams. Ensemble models and batch-based approaches have been the dominant methods in online classification, but some of their assumptions or requirements may not hold in many real-world applications. For instance, in some applications, data may arrive in a continuous stream, and it may not be possible to divide it into batches due to time or memory constraints. In such cases, single-pass online sampling strategies that do not require the use or update of multiple models would be more practical. Moreover, it could be beneficial to develop online active learning strategies that are able to tackle all the types of distribution shifts introduced in Sect. 3.2. Finally, the combination of reinforcement learning and active learning in pool-based scenarios is an area of ongoing research. We believe that the study of online reinforcement learning to optimize sampling strategies could provide valuable insights into how to best perform active learning in dynamic environments.

7 Conclusion

The field of online active learning with data streams is a rapidly evolving and highly relevant area of research in machine learning. The ability to effectively learn from data streams in real-time is becoming increasingly important, as the amount of data generated by modern applications continues to grow at an exponential rate. However, obtaining annotated data to train complex prediction and decision-making models presents a major roadblock. This hinders the proper integration of artificial intelligence models with real-world applications such as healthcare, autonomous driving and industrial production. Our survey provides a comprehensive overview of the current state of the art in this field and highlights the challenges and opportunities that researchers face when developing methods for online active learning. We reviewed a wide range of strategies for selecting the most informative data points in online active learning, including methods based on uncertainty sampling, diversity sampling, query by committee, and reinforcement learning, among others. Our analysis has shown that these strategies have been applied in a variety of contexts, including online classification, online regression, and online semi-supervised learning. We hope that this survey will inspire further research in the field of online active learning with data streams and encourage the development of new and advanced methods for handling this type of data. In particular, we believe that there is significant potential for the development of model-agnostic and single-pass online active learning strategies that can be applied in practical settings.

Acknowledgements The authors gratefully acknowledge the support of the DTU Strategic Alliances Fund, which made this research possible. We would also like to extend our sincere thanks to John Sølve Tyssedal for his invaluable help and support throughout the project.

Author Contributions DC conceived and designed the analysis and wrote the paper. MK provided guidance and supervision throughout the work, reviewed the paper, and provided critical feedback to improve the manuscript.

Funding Open access funding provided by Technical University of Denmark The research reported in this paper was supported by the DTU Strategic Alliances Fund.

Data availability Not applicable.

Code Availability Not applicable.

Declarations

Conflict of interest The authors declare that they have no known competing financial interests or personal relationships that could have influenced the work reported in this paper.

Ethical approval Not applicable.

Consent to participate Not applicable.

Consent for publication Not applicable.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Agarwal, S., Arora, H., Anand, S., et al. (2020). Contextual diversity for active learning. In *European conference on computer vision 2020*. https://doi.org/10.1007/978-3-030-58517-4_9. arXiv:2008.05723.
- Aggarwal, C. C., Kong, X., Gu, Q., et al. (2014). *Data classification (Chapter: "Active learning: A survey")*. Taylor & Francis. <http://charuaggarwal.net/active-survey.pdf>.
- Aguiar, G., Krawczyk, B., & Cano, A. (2023). A survey on learning from imbalanced data streams: Taxonomy, challenges, empirical study, and reproducible experimental framework. *Machine Learning*, 1–79.
- Alabdulrahman, R., Viktor, H., & Paquet, E. (2016). An active learning approach for ensemble-based data stream mining. In *International conference on knowledge discovery and information retrieval, SCITEPRESS* (pp. 275–282).
- Ash, J. T., Zhang, C., Krishnamurthy, A., et al. (2019). Deep batch active learning by diverse, uncertain gradient lower bounds. In *2020 international conference on learning representations*. arXiv:1906.03671.
- Asprey, S., & Macchietto, S. (2002). Designing robust optimal dynamic experiments. *Journal of Process Control*, 12, 545–556. [https://doi.org/10.1016/S0959-1524\(01\)00020-8](https://doi.org/10.1016/S0959-1524(01)00020-8)
- Audibert, J. Y., & Munos, R. (2010). Best arm identification in multi-armed bandits. In *COLT—23th conference on learning theory*. <http://certis.enpc.fr/audibert/Mes%20articles/COLT10.pdf>.
- Avadhanula, V., Colini Baldeschi, R., Leonardi, S., et al. (2021). Stochastic bandits for multi-platform budget optimization in online advertising. *Proceedings of the Web Conference, 2021*, 2805–2817.
- Azizi, M. J., Kveton, B., & Ghavamzadeh, M. (2022). Fixed-budget best-arm identification in structured bandits. In *Proceedings of the thirty-first international joint conference on artificial intelligence (IJCAI-22)*. <https://www.ijcai.org/proceedings/2022/0388.pdf>.
- Baier, L., Schlör, T., Schöffner, J., et al. (2021). Detecting concept drift with neural network model uncertainty. In *Hawaii international conference on system sciences (HICSS) 2023*. arXiv:2107.01873.
- Balcan, M. F., Broder, A., & Zhang, T. (2007). Margin based active learning. In *COLT—23th conference on learning theory* 4739. https://doi.org/10.1007/978-3-540-72927-3_5.
- Bassily, R., Smith, A., & Thakurta, A. (2014). Private empirical risk minimization: Efficient algorithms and tight error bounds. In *2014 IEEE 55th annual symposium on foundations of computer science* (pp. 464–473). <https://doi.org/10.1109/FOCS.2014.56>.
- Baum, E., & Lang, K. (1992). Query learning can work poorly when a human oracle is used. In *Proceedings of the IEEE international joint conference on neural networks*.
- Baykal, C., Trinh, K., Iliopoulos, F., et al. (2022). *Robust active distillation*. arXiv:2210.01213.
- Beck, N., Kothawade, S., Shenoy, P., et al. (2023). *Streamline: Streaming active learning for realistic multi-distributional settings*. arXiv preprint arXiv:2305.10643.
- Bifet, A., & Gavaldà, R. (2007). Learning from time-changing data with adaptive windowing. In *Proceedings of the 2007 SIAM international conference on data mining* (pp. 443–448). <https://doi.org/10.1137/1.9781611972771.42>.

- Bifet, A., & Gavaldà, R. (2009). Adaptive learning from evolving data streams. In *IDA 2009: Advances in intelligent data analysis VIII* (pp. 249–260). https://doi.org/10.1007/978-3-642-03915-7_22.
- Bifet, A., Holmes, G., Pfahringer, B., et al. (2010). Moa: massive online analysis, a framework for stream classification and clustering. In *Proceedings of the first workshop on applications of pattern analysis, PMLR* (pp. 44–50).
- Bisgaard, S., & Kulahci, M. (2011). *Time series analysis and forecasting by example*. New York: Wiley.
- Bordes, A., Ertekin, S., Weston, J., et al. (2005). Fast kernel classifiers with online and active learning. *The Journal of Machine Learning Research*, 6. <https://jmlr.csail.mit.edu/papers/v6/bordes05a.html>.
- Bouchachia, A., & Vanaret, C. (2014). Gt2fc: An online growing interval type-2 self-learning fuzzy classifier. *IEEE Transactions on Fuzzy Systems*, 22, 999–1018. <https://doi.org/10.1109/TFUZZ.2013.2279554>
- Brzezinski, D., & Stefanowski, J. (2015). Prequential auc for classifier evaluation and drift detection in evolving data streams. In *3rd International workshop on new frontiers in mining complex patterns, (NFMCP 2014)* (pp. 87–101). https://doi.org/10.1007/978-3-319-17876-9_6.
- Brzezinski, D., & Stefanowski, J. (2017). Prequential auc: Properties of the area under the roc curve for data streams with concept drift. *Knowledge and Information Systems*, 52, 531–562. <https://doi.org/10.1007/s10115-017-1022-8>
- Burbidge, R., Rowland, J. J., & King, R.D. (2007). Active learning for regression based on query by committee. In *8th International conference on intelligent data engineering and automated learning, IDEAL 2007*. https://doi.org/10.1007/978-3-540-77226-2_22.
- Cacciarelli, D., & Boresta, M. (2021). What drives a donor? A machine learning-based approach for predicting responses of nonprofit direct marketing campaigns. *International Journal of Nonprofit and Voluntary Sector Marketing*. <https://doi.org/10.1002/nvsm.1724>
- Cacciarelli, D., & Kulahci, M. (2022). A novel fault detection and diagnosis approach based on orthogonal autoencoders. *Computers & Chemical Engineering*, 163, 107853. <https://doi.org/10.1016/j.compchemeng.2022.107853>
- Cacciarelli, D., & Kulahci, M. (2023). Hidden dimensions of the data: PCA vs autoencoders. *Quality Engineering*, 35, 741–750. <https://doi.org/10.1080/08982112.2023.2231064>
- Cacciarelli, D., Kulahci, M., & Tyssedal, J. (2022a). Online active learning for soft sensor development using semi-supervised autoencoders. In *ICML 2022 workshop on adaptive experimental design and active learning in the real world*. [arXiv:2212.13067](https://arxiv.org/abs/2212.13067).
- Cacciarelli, D., Kulahci, M., & Tyssedal, J. S. (2022b). Stream-based active learning with linear models. *Knowledge-Based Systems*, 254, 109664. <https://doi.org/10.1016/j.knsys.2022.109664>.
- Cacciarelli, D., Kulahci, M., & Tyssedal, J. S. (2023). Robust online active learning. *Quality and Reliability Engineering International*. <https://doi.org/10.1002/qre.3392>
- Cai, W., Zhang, Y., & Zhou, J. (2013). Maximizing expected model change for active learning in regression. In *Proceedings—IEEE international conference on data mining, ICDM* (pp. 51–60). <https://doi.org/10.1109/ICDM.2013.104>.
- Camilleri, R., Xiong, Z., Fazel, M., et al. (2021). Selective sampling for online best-arm identification. In *35th conference on neural information processing systems (NeurIPS 2021)*. [arXiv:2110.14864](https://arxiv.org/abs/2110.14864).
- Carcillo, F., Le Borgne, Y. A., Caelen, O., et al. (2017). An assessment of streaming active learning strategies for real-life credit card fraud detection. In *2017 IEEE international conference on data science and advanced analytics (DSAA)* (pp. 631–639). IEEE.
- Carcillo, F., Le Borgne, Y. A., Caelen, O., et al. (2018). Streaming active learning strategies for real-life credit card fraud detection: Assessment and visualization. *International Journal of Data Science and Analytics*, 5, 285–300.
- Carnein, M., & Trautmann, H. (2019). Customer segmentation based on transactional data using stream clustering. In *Advances in knowledge discovery and data mining: 23rd Pacific-Asia Conference, PAKDD 2019, Macau, China, April 14–17, 2019, proceedings, part I* (Vol. 23, pp. 280–292). Springer.
- Carpentier, A., Lazaric, A., Ghavamzadeh, M., et al. (2015). Upper-confidence-bound algorithms for active learning in multi-armed bandits.
- Castellani, A., Schmitt, S., & Hammer, B. (2022). Stream-based active learning with verification latency in non-stationary environments. https://doi.org/10.1007/978-3-031-15937-4_22. [arXiv:2204.06822](https://arxiv.org/abs/2204.06822).
- Cernuda, C., Lughofer, E., Mayr, G., et al. (2014). Incremental and decremental active learning for optimized self-adaptive calibration in viscose production. *Chemometrics and Intelligent Laboratory Systems*, 138, 14–29. <https://doi.org/10.1016/j.chemolab.2014.07.008>
- Cerqueira, V., Torgo, L., & Mozetič, I. (2020). Evaluating time series forecasting models: An empirical study on performance estimation methods. *Machine Learning*, 109, 1997–2028. <https://doi.org/10.1007/s10994-020-05910-7>
- Cesa-Bianchi, N., & Lugosi, G. (2006). *Prediction, learning, and games*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511546921>

- Cesa-Bianchi, N., Gentile, C., Zaniboni, L. (2004). Worst-case analysis of selective sampling for linear-threshold algorithms. In *Advances in neural information processing systems*. https://proceedings.neurips.cc/paper_files/paper/2004/hash/92426b262d11b0ade77387cf8416e153-Abstract.html.
- Cesa-Bianchi, N., Gentile, C., Zaniboni, L. (2006). Worst-case analysis of selective sampling for linear classification. *The Journal of Machine Learning Research*, 7. <https://www.jmlr.org/papers/volume7/cesa-bianchi06b/cesa-bianchi06b.pdf>.
- Chae, J., & Hong, S. (2021). Stream-based active learning with multiple kernels. In *2021 international conference on information networking (ICOIN)* (pp. 718–722). <https://doi.org/10.1109/ICOIN50884.2021.9333940>.
- Chan, L. L. T., Wu, Q. Y., & Chen, J. (2018). Dynamic soft sensors with active forward-update learning for selection of useful data from historical big database. *Chemometrics and Intelligent Laboratory Systems*, 175, 87–103. <https://doi.org/10.1016/j.chemolab.2018.01.015>
- Cheng, J., Zheng, Z., Guo, Y., et al. (2023). Active broad learning with multi-objective evolution for data stream classification. *Complex & Intelligent Systems*, 12, 1–18.
- Chu, W., Zinkevich, M., Li, L., et al. (2011). Unbiased online active learning in data streams. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining—KDD '11* (p. 195). <https://doi.org/10.1145/2020408.2020444>.
- Citovsky, G., DeSalvo, G., Gentile, C., et al. (2021). Batch active learning at scale. In *35th Conference on neural information processing systems, NeurIPS 2021*. [arXiv:2107.14263](https://arxiv.org/abs/2107.14263).
- Cohn, D. A., Ghahramani, Z., & Jordan, M. I. (1996). Active learning with statistical models. *Journal of Artificial Intelligence Research*, 4, 129–145. <https://doi.org/10.1613/jair.295>
- Crammer, K., Dekel, O., Keshet, J., et al. (2006). Online passive-aggressive algorithms. *The Journal of Machine Learning Research*. <https://jmlr.csail.mit.edu/papers/volume7/crammer06a/crammer06a.pdf>.
- Dasgupta, S., Kalai, A. T., & Monteleoni, C. (2005). Analysis of perceptron-based active learning. In *COLT '05—international conference on computational learning theory* (pp. 249–263). https://doi.org/10.1007/11503415_17.
- Desalvo, G., Gentile, C., & Thune, T. S. (2021). Online active learning with surrogate loss functions. In *Advances in neural information processing systems 34 (NeurIPS 2021)*. <https://proceedings.neurips.cc/paper/2021/hash/c1619d2ad66f7629c12c87fe21d32a58-Abstract.html>.
- Donmez, P., Carbonell, J., & Bennet, P. (2007). Dual strategy active learning. In *18th European conference on machine learning, ECML 2007*, 4701. https://doi.org/10.1007/978-3-540-74958-5_14.
- Duchi, J. C., Jordan, M. I., & Wainwright, M. J. (2013). Local privacy and statistical minimax rates. In *2013 IEEE 54th annual symposium on foundations of computer science* (pp. 429–438). <https://doi.org/10.1109/FOCS.2013.53>.
- Ebbinghaus, H. (2013). Memory: A contribution to experimental psychology. *Annals of Neurosciences*. <https://doi.org/10.5214/ans.0972.7531.200408>
- Fang, M., Li, Y., & Cohn, T. (2017). Learning how to active learn: A deep reinforcement learning approach. [arXiv:1708.02383](https://arxiv.org/abs/1708.02383).
- Ferdowsi, Z., Ghani, R., & Settini, R. (2013). Online active learning with imbalanced classes. In *2013 IEEE 13th international conference on data mining* (pp. 1043–1048). <https://doi.org/10.1109/ICDM.2013.12>.
- Feuz, K. D., & Cook, D. J. (2013). Real-time annotation tool (rat). In *Workshops at the twenty-seventh AAAI conference on artificial intelligence*.
- Fiez, T., Jain, L., Jamieson, K., et al. (2019). Sequential experimental design for transductive linear bandits. In *33rd conference on neural information processing systems (NeurIPS 2019)*. https://proceedings.neurips.cc/paper_files/paper/2019/file/8ba6c657b03fc7c8dd4dff8e45defcd2-Paper.pdf.
- Filippi, S., Cappe, O., Garivier, A., et al. (2010). Parametric bandits: The generalized linear case. In *Advances in neural information processing systems 23 (NIPS 2010)*. https://papers.nips.cc/paper_files/paper/2010/hash/c2626d850c80ea07e7511bbae4c76f4b-Abstract.html.
- Fontaine, X., Perrault, P., Valko, M., et al. (2021). Online α -optimal design and active linear regression. <http://proceedings.mlr.press/v139/fontaine21a/fontaine21a.pdf>.
- Fortuna, L., Graziani, S., Rizzo, A., et al. (2007). *Soft sensors for monitoring and control of industrial processes* (Vol. 22). Berlin: Springer. <https://doi.org/10.1007/978-1-84628-480-9>
- Fowler, K., Kokilepersaud, K., Prabhushankar, M., et al. (2023). Clinical trial active learning. In *The 14th ACM conference on bioinformatics, computational biology and health informatics (ACM-BCB)*.
- Freeman, P. R. (1983). The secretary problem and its extensions: A review. *International Statistical Review*, 51, 189–206.
- Freund, Y., Seung, H. S., Shamir, E., et al. (1997). Selective sampling using the query by committee algorithm. *Machine Learning*, 28, 133–168. <https://doi.org/10.1023/a:1007330508534>
- Friedman, M. (1940). A comparison of alternative tests of significance for the problem of m rankings. *The Annals of Mathematical Statistics*, 11, 86–92. <https://doi.org/10.1214/aoms/1177731944>

- Frumosu, F. D., & Kulahci, M. (2018). Big data analytics using semi-supervised learning methods. *Quality and Reliability Engineering International*, 34, 1413–1423. <https://doi.org/10.1002/qre.2338>
- Fu, Y., Zhu, X., & Li, B. (2013). A survey on instance selection for active learning. *Knowledge and Information Systems*, 35, 249–283. <https://doi.org/10.1007/s10115-012-0507-8>
- Fujii, K., & Kashima, H. (2016). Budgeted stream-based active learning via adaptive submodular maximization. In *30th annual conference on neural information processing systems, NIPS 2016*. <https://proceedings.neurips.cc/paper/2016/hash/07cdfd23373b17c6b337251c22b7ea57-Abstract.html>.
- Gajjar, S., Kulahci, M., & Palazoglu, A. (2018). Real-time fault detection and diagnosis using sparse principal component analysis. *Journal of Process Control*, 67, 112–128. <https://doi.org/10.1016/j.jprocont.2017.03.005>
- Galvanin, F. (2010). Optimal model-based design of experiments in dynamic systems: Novel techniques and unconventional applications. Thesis. <https://hdl.handle.net/11577/3427095>.
- Gama, J., Medas, P., Castillo, G., et al. (2004). Learning with drift detection. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 3171, 286–295. https://doi.org/10.1007/978-3-540-28645-5_29
- Gama, J., Sebastiao, R., & Rodrigues, P. P. (2009). Issues in evaluation of stream learning algorithms. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining* (pp. 329–338).
- Gama, J., Sebastiao, R., & Rodrigues, P. P. (2013). On evaluating stream learning algorithms. *Machine Learning*, 90, 317–346.
- Garivier, A., & Moulines, E. (2008). On upper-confidence bound policies for non-stationary bandit problems. [arXiv:0805.3415](https://arxiv.org/abs/0805.3415).
- Ge, D., & Zeng, X. J. (2020). Learning data streams online—An evolving fuzzy system approach with self-learning/adaptive thresholds. *Information Sciences*, 507, 172–184. <https://doi.org/10.1016/j.ins.2019.08.036>
- Ge, Z. (2014). Active learning strategy for smart soft sensor development under a small number of labeled data samples. *Journal of Process Control*, 24, 1454–1461. <https://doi.org/10.1016/j.jprocont.2014.06.015>
- Gemaque, R. N., Costa, A. F. J., Giusti, R., et al. (2020). An overview of unsupervised drift detection methods. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*. <https://doi.org/10.1002/widm.1381>
- Ghassemi, M., Sarwate, A. D., & Wright, R. N. (2016). Differentially private online active learning with applications to anomaly detection. In *AISeC 2016—Proceedings of the 2016 ACM workshop on artificial intelligence and security, co-located with CCS 2016* (pp. 117–128). <https://doi.org/10.1145/2996758.2996766>.
- Ghiassi, S., Pazzi, G., Del Grosso, C., et al. (2023). Combining thermodynamics-based model of the centrifugal compressors and active machine learning for enhanced industrial design optimization. In *1st workshop on the synergy of scientific and machine learning modeling@ ICML2023*.
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., et al. (2014). Generative adversarial networks. [arXiv:1406.2661](https://arxiv.org/abs/1406.2661).
- Gouk, H., Pfahringer, B., & Frank, E. (2019). Stochastic gradient trees. <http://proceedings.mlr.press/v101/gouk19a/gouk19a.pdf>.
- Gu, X., Han, J., Shen, Q., et al. (2022). Autonomous learning for fuzzy systems: A review. *Artificial Intelligence Review*. <https://doi.org/10.1007/s10462-022-10355-6>
- Gu, X., Han, J., Shen, Q., et al. (2023). Autonomous learning for fuzzy systems: A review. *Artificial Intelligence Review*, 56(8), 7549–7595.
- Halder, B., Hasan, K. A., Amagasa, T., et al. (2023). Autonomic active learning strategy using cluster-based ensemble classifier for concept drifts in imbalanced data stream. *Expert Systems with Applications*, 120578.
- Hanneke, S. (2014). Theory of disagreement-based active learning. *Foundations and Trends in Machine Learning*, 7, 131–309. <https://doi.org/10.1561/22000000037>
- Hanneke, S., & Yang, L. (2021). Toward a general theory of online selective sampling: Trading off mistakes and queries. In *proceedings of the 24th international conference on artificial intelligence and statistics*. <https://proceedings.mlr.press/v130/hanneke21a.html>.
- Hao, S., Hu, P., Zhao, P., et al. (2018). Online active learning with expert advice. *ACM Transactions on Knowledge Discovery from Data*. <https://doi.org/10.1145/3201604>
- Hao, S., Lu, J., Zhao, P., et al. (2018). Second-order online active learning and its applications. *IEEE Transactions on Knowledge and Data Engineering*, 30, 1338–1351. <https://doi.org/10.1109/TKDE.2017.2778097>
- Hausmann, E., Fenzi, M., Chitta, K., et al. (2020). Scalable active learning for object detection. In *Proceedings 31st IEEE intelligent vehicles symposium (IV)*. <https://doi.org/10.1109/IV47402.2020.9304793>.

- He, K., Zhang, X., Ren, S., et al. (2015). Deep residual learning for image recognition. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. <https://doi.org/10.1109/CVPR.2016.90>
- Hoang, T. N., Hong, S., Xiao, C., et al. (2021). Aid: Active distillation machine to leverage pre-trained black-box models in private data settings. *Proceedings of the Web Conference, 2021*, 3569–3581. <https://doi.org/10.1145/3442381.3449944>
- Hodges, J., & Lehmann, E. (1962). Rank methods for combination of independent experiments in analysis of variance. *The Annals of Mathematical Statistics*, 33, 482–497.
- Hoffmann, H. (2007). Kernel PCA for novelty detection. *Pattern Recognition*, 40, 863–874. <https://doi.org/10.1016/j.patcog.2006.07.009>
- Hoi, S. C., Sahoo, D., Lu, J., et al. (2021). Online learning: A comprehensive survey. *Neurocomputing*, 459, 249–289. <https://doi.org/10.1016/j.neucom.2021.04.112>
- Hoi, S. C. H., Jin, R., Zhao, P., et al. (2013). Online multiple kernel classification. *Machine Learning*, 90, 289–316. <https://doi.org/10.1007/s10994-012-5319-2>
- Houlsby, N., Hernandez-Lobato, J. M., & Ghahramani, Z. (2014). Cold-start active learning with robust ordinal matrix factorization. In *31st international conference on machine learning*. <https://proceedings.mlr.press/v32/houlsby14.html>.
- Hua, J., Xiong, Z., Lowey, J., et al. (2005). Optimal number of features as a function of sample size for various classification rules. *Bioinformatics*, 21, 1509–1515. <https://doi.org/10.1093/bioinformatics/bti171>
- Huang, B., Salgia, S., & Zhao, Q. (2022). Disagreement-based active learning in online settings. *IEEE Transactions on Signal Processing*, 70, 1947–1958. <https://doi.org/10.1109/TSP.2022.3159388>
- Huang, G. B., Zhu, Q. Y., & Siew, C. K. (2006). Extreme learning machine: Theory and applications. *Neurocomputing*, 70, 489–501. <https://doi.org/10.1016/j.neucom.2005.12.126>
- Huang, S. J., Jin, R., & Zhou, Z. H. (2014). Active learning by querying informative and representative examples. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36, 1936–1949. <https://doi.org/10.1109/TPAMI.2014.2307881>
- Hulten, G., Spencer, L., & Domingos, P. (2001). Mining time-changing data streams. In *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining—KDD '01* (pp. 97–106). <https://doi.org/10.1145/502512.502529>.
- Ienco, D., Bifet, A., Zliobaite, et al. (2013). Clustering based active learning for evolving data streams. In *16th international conference on discovery science*. https://doi.org/10.1007/978-3-642-40897-7_6.
- Ienco, D., Pfahringer, B., & Zliobaite, I. (2014). High density-focused uncertainty sampling for active learning over evolving stream data. In *BIGMINE'14: Proceedings of the 3rd international conference on big data, streams and heterogeneous source mining: Algorithms, systems, programming models and applications*. <https://proceedings.mlr.press/v36/ienco14.html>.
- Iman, R. L., & Davenport, J. M. (1980). Approximations of the critical region of the fbietkan statistic. *Communications in Statistics - Theory and Methods*, 9, 571–595. <https://doi.org/10.1080/03610928008827904>
- Istrate, R., Malossi, A. C. I., Bekas, C., et al. (2018). Incremental training of deep convolutional neural networks. [arXiv:1803.10232](https://arxiv.org/abs/1803.10232).
- Jamieson, K. (2018). Online and adaptive machine learning. *Regression (Part 7)*. <https://courses.cs.washington.edu/courses/cse599i/18wi/>.
- Jamieson, K., & Nowak, R. (2014). Best-arm identification algorithms for multi-armed bandits in the fixed confidence setting. In *2014 48th annual conference on information sciences and systems (CISS)* (pp. 1–6). <https://doi.org/10.1109/CISS.2014.6814096>.
- Jamil, S., & Khan, A. (2016). Churn comprehension analysis for telecommunication industry using alba. In *2016 international conference on emerging technologies (ICET)* (pp. 1–5). IEEE.
- Jedra, Y., & Proutiere, A. (2020). Optimal best-arm identification in linear bandits. In *34th conference on neural information processing systems (NeurIPS 2020)*. <https://proceedings.neurips.cc/paper/2020/file/7212a6567c8a6c513f33b858d868ff80-Paper.pdf>.
- Jin, Q., Yuan, M., Li, S., et al. (2022). Cold-start active learning for image classification. *Information Sciences*, 616, 16–36. <https://doi.org/10.1016/j.ins.2022.10.066>
- Jin, R., Hoi, S., & Yang, T. (2010). Online multiple kernel learning: Algorithms and mistake bounds. In *Proceedings of the 21st international conference on algorithmic learning theory*. https://doi.org/10.1007/978-3-642-16108-7_31.
- John, R. C. S., & Draper, N. R. (1975). D-optimality for regression designs: A review. *Technometrics*, 17, 15–23. <https://doi.org/10.1080/00401706.1975.10489266>
- Joshi, A. J., Porikli, F., & Papanikolopoulos, N. (2009). Multi-class active learning for image classification. In *2009 IEEE conference on computer vision and pattern recognition* (pp. 2372–2379). <https://doi.org/10.1109/CVPR.2009.5206627>.
- Joyce, J. M. (2011). Kullback–Leibler divergence. https://doi.org/10.1007/978-3-642-04898-2_327.

- Karlin, S., & Studden, W. J. (1966). Optimal experimental designs. *The Annals of Mathematical Statistics*, 37, 783–815.
- Kiefer, J. (1959). Optimum experimental designs. *Journal of the Royal Statistical Society Series B (Methodological)*. <https://www.jstor.org/stable/2983802>.
- Kingma, D. P., & Welling, M. (2013). Auto-encoding variational Bayes. In *2nd international conference on learning representations, ICLR*. [arXiv:1312.6114](https://arxiv.org/abs/1312.6114).
- Kottke, D., Kreml, G., & Spiliopoulou, M. (2015). *Probabilistic active learning in datastreams*. https://doi.org/10.1007/978-3-319-24465-5_13.
- Kranjc, J., Smailović, J., Podpečan, V., et al. (2015). Active learning for sentiment analysis on data streams: Methodology and workflow implementation in the clowflows platform. *Information Processing & Management*, 51(2), 187–203.
- Kraskov, A., Stögbauer, H., & Grassberger, P. (2004). Estimating mutual information. *Physical Review E*, 69, 066138. <https://doi.org/10.1103/PhysRevE.69.066138>
- Krawczyk, B., Minku, L. L., Gama, J., et al. (2017). Ensemble learning for data stream analysis: A survey. *Information Fusion*, 37, 132–156.
- Krawczyk, B., Pfahringer, B., & Wozniak, M. (2018). Combining active learning with concept drift detection for data stream mining. In *2018 IEEE international conference on big data (big data)* (pp. 2239–2244). <https://doi.org/10.1109/BigData.2018.8622549>.
- Škrjanc, I. (2009). Confidence interval of fuzzy models: An example using a waste-water treatment plant. *Chemometrics and Intelligent Laboratory Systems*, 96, 182–187. <https://doi.org/10.1016/j.chemolab.2009.01.009>
- Kulkarni, R. V., Patil, S. H., & Subhashini, R. (2016). An overview of learning in data streams with label scarcity. In *Proceedings of the international conference on inventive computation technologies, ICICT, 2016* (Vol. 2). <https://doi.org/10.1109/INVENTIVE.2016.7824874>.
- Kumar, P., & Gupta, A. (2020). Active learning query strategies for classification, regression, and clustering: A survey. *Journal of Computer Science and Technology*, 35, 913–945. <https://doi.org/10.1007/s11390-020-9487-4>
- Kurlej, B., & Woźniak, M. (2011). *Learning curve in concept drift while using active learning paradigm*. https://doi.org/10.1007/978-3-642-23857-4_13.
- Kwak, B., Kim, Y., & Kim, Y. J., et al. (2022). Trustal: Trustworthy active learning using knowledge distillation. In *The thirty-sixth AAAI conference on artificial intelligence (AAAI-22)*. [arXiv:2201.11661](https://arxiv.org/abs/2201.11661).
- Lakshminarayanan, B., Roy, D., & Teh, Y. W. (2014). Mondrian forests: Efficient online random forests. In *Advances in neural information processing systems (NIPS)*. https://proceedings.neurips.cc/paper_files/paper/2014/file/d1dc3a8270a6f9394f88847d7f0050cf-Paper.pdf.
- Li, A., Boyd, A., Smyth, P., et al. (2021). Detecting and adapting to irregular distribution shifts in Bayesian online learning. In *35th conference on neural information processing systems (NeurIPS 2021)*. <https://papers.nips.cc/paper/2021/file/362387494f6be6613daea643a7706a42-Paper.pdf>.
- Li, X., & Guo, Y. (2013). Adaptive active learning for image classification. In *2013 IEEE conference on computer vision and pattern recognition* (pp. 859–866). <https://doi.org/10.1109/CVPR.2013.116>.
- Lieber, D., Konrad, B., Deuse, J., et al. (2012). Sustainable interlinked manufacturing processes through real-time quality prediction. In *Leveraging technology for a sustainable world: Proceedings of the 19th CIRP conference on life cycle engineering*, University of California at Berkeley, Berkeley, USA, May 23–25, 2012 (pp. 393–398). Springer.
- Lima, M., Neto, M., Filho, T. S., et al. (2022). Learning under concept drift for regression—a systematic literature review. *IEEE Access*, 10, 45410–45429. <https://doi.org/10.1109/ACCESS.2022.3169785>
- Liu, S., Xue, S., Wu, J., et al. (2021). Online active learning for drifting data streams. *IEEE Transactions on Neural Networks and Learning Systems*. <https://doi.org/10.1109/TNNLS.2021.3091681>
- Long, J., Yin, J., Zhao, W., et al. (2008). Graph-based active learning based on label propagation. In *MDAI 2008: Modeling decisions for artificial intelligence* (pp. 179–190). https://doi.org/10.1007/978-3-540-88269-5_17.
- Loy, C. C., Hospedales, T. M., Xiang, T., et al. (2012). Stream-based joint exploration-exploitation active learning. In *Proceedings of the IEEE computer society conference on computer vision and pattern recognition* (pp. 1560–1567). <https://doi.org/10.1109/CVPR.2012.6247847>.
- Lu, J., Zhao, P., & Hoi, S. C. H. (2016). Online passive-aggressive active learning. *Machine Learning*, 103, 141–183. <https://doi.org/10.1007/s10994-016-5555-y>
- Lu, J., Liu, A., Dong, F., et al. (2018). Learning under concept drift: A review. *IEEE Transactions on Knowledge and Data Engineering*. <https://doi.org/10.1109/TKDE.2018.2876857>
- Lughofer, E. (2011). *Evolving fuzzy systems—Methodologies, advanced concepts and applications* (Vol. 266). Berlin: Springer. <https://doi.org/10.1007/978-3-642-18087-3>

- Lughofer, E. (2012). Single-pass active learning with conflict and ignorance. *Evolving Systems*, 3, 251–271. <https://doi.org/10.1007/s12530-012-9060-7>
- Lughofer, E. (2017). On-line active learning: A new paradigm to improve practical useability of data stream modeling methods. *Information Sciences*, 415–416, 356–376. <https://doi.org/10.1016/j.ins.2017.06.038>
- Lughofer, E., & Pratama, M. (2018). Online active learning in data stream regression using uncertainty sampling based on evolving generalized fuzzy models. *IEEE Transactions on Fuzzy Systems*, 26, 292–309. <https://doi.org/10.1109/TFUZZ.2017.2654504>
- Lughofer, E., & Škrjanc, I. (2023). Online active learning for evolving error feedback fuzzy models within a multi-innovation context. *IEEE Transactions on Fuzzy Systems*. <https://doi.org/10.1109/TFUZZ.2023.3302403>
- Ma, L., Destercke, S., & Wang, Y. (2016). Online active learning of decision trees with evidential data. *Pattern Recognition*, 52, 33–45. <https://doi.org/10.1016/j.patcog.2015.10.014>
- Mammen, E., & Tsybakov, A. B. (1999). Smooth discrimination analysis. *The Annals of Statistics*. <https://doi.org/10.1214/aos/1017939240>
- Manjah, D., Cacciarelli, D., Standaert, B., et al. (2023). Stream-based active distillation for scalable model deployment. In *Proceedings of the IEEE/CVF computer vision and pattern recognition (CVPR) workshops*.
- Manwani, N., Desai, K., Sasidharan, S., et al. (2013). Double ramp loss based reject option classifier. In *19th Pacific-Asia conference on advances in knowledge discovery and data mining (PAKDD)*. https://doi.org/10.1007/978-3-319-57454-7_53.
- Martins, V. E., Cano, A., & Junior, S. B. (2023). Meta-learning for dynamic tuning of active learning on stream classification. *Pattern Recognition*, 138, 109359.
- McSherry, F., & Talwar, K. (2007). Mechanism design via differential privacy. In *48th annual IEEE symposium on foundations of computer science (FOCS'07)* (pp. 94–103). <https://doi.org/10.1109/FOCS.2007.41>.
- Menard, P., Domingues, O. D., Jonsson, A., et al. (2021). Fast active learning for pure exploration in reinforcement learning. In *Proceedings of the 38th international conference on machine learning*. <http://proceedings.mlr.press/v139/menard21a/menard21a-supp.pdf>.
- Min, F., Zhang, S. M., Ciucci, D., et al. (2020). Three-way active learning through clustering selection. *International Journal of Machine Learning and Cybernetics*, 11, 1033–1046. <https://doi.org/10.1007/s13042-020-01099-2>
- Minka, T. P. (2001). A family of algorithms for approximate Bayesian inference. Thesis. <https://hd.media.mit.edu/tech-reports/TR-533.pdf>.
- Miu, T., Missier, P., & Plötz, T. (2015). Bootstrapping personalised human activity recognition models using online active learning. *2015 IEEE international conference on computer and information technology; ubiquitous computing and communications; dependable, autonomic and secure computing; pervasive intelligence and computing* (pp. 1138–1147). IEEE.
- Mohamad, S., Bouchachia, A., & Sayed-Mouchaweh, M. (2018). A bi-criteria active learning algorithm for dynamic data streams. *IEEE Transactions on Neural Networks and Learning Systems*, 29, 74–86. <https://doi.org/10.1109/TNNLS.2016.2614393>
- Mohamad, S., Sayed-Mouchaweh, M., & Bouchachia, A. (2020). Online active learning for human activity recognition from sensory data streams. *Neurocomputing*, 390, 341–358. <https://doi.org/10.1016/j.neucom.2019.08.092>
- Mohamadi, S., & Amindavar, H. (2020). Deep Bayesian active learning, a brief survey on recent advances. [arXiv:2012.08044](https://arxiv.org/abs/2012.08044).
- Montgomery, D. C. (2012). *Design and analysis of experiments*. New York: Wiley. <https://doi.org/10.1002/9781118147634>
- Myers, R. H., Montgomery, D., & Anderson-Cook, C. M. (2016). Response surface methodology: Process and product optimization using designed experiments. Wiley. <https://www.wiley.com/en-au/Response+Surface+Methodology:+Process+and+Product+Optimization+Using+Designed+Experiments,+4th+Edition-p-9781118916018>.
- Naranjo, J. E., Sotelo, M. A., Gonzalez, C., et al. (2007). Using fuzzy logic in automated vehicle control. *IEEE Intelligent Systems*, 22(1), 36–45.
- Narr, A., Triebel, R., & Cremers, D. (2016). Stream-based active learning for efficient and adaptive classification of 3d objects. In *Proceedings—IEEE international conference on robotics and automation 2016-June* (pp. 227–233). <https://doi.org/10.1109/ICRA.2016.7487138>.
- Nguyen, H. T., & Smeulders, A. (2004). Active learning using pre-clustering. In *Proceedings of the twenty-first international conference on machine learning*. <https://doi.org/10.1145/1015330.1015349>.
- Nixon, C., Sedky, M., & Hassan, M. (2021). Reviews in online data stream and active learning for cyber intrusion detection—a systematic literature review. In *2021 Sixth international conference on fog and mobile edge computing (FMEC)* (pp. 1–6). IEEE.

- Pham, T., Kottke, D., Krempel, G., et al. (2022). Stream-based active learning for sliding windows under the influence of verification latency. *Machine Learning*, 111, 2011–2036. <https://doi.org/10.1007/s10994-021-06099-z>
- Pitman, J., & Yor, M. (1997). The two-parameter Poisson–Dirichlet distribution derived from a stable subordinator. *The Annals of Probability*, 25. <https://www.jstor.org/stable/20680193>.
- Polikar, R., Upda, L., Upda, S., et al. (2001). Learn++: An incremental learning algorithm for supervised neural networks. *IEEE Transactions on Systems, Man and Cybernetics. Part C (Applications and Reviews)*, 31, 497–508. <https://doi.org/10.1109/5326.983933>
- Prabhu, V., Chandrasekaran, A., Saenko, K., et al. (2020). Active domain adaptation via clustering uncertainty-weighted embeddings. <https://github.com/virajprabhu/CLUE>.
- Pratama, M., Anavatti, S. G., & Lu, J. (2015). Recurrent classifier based on an incremental metacognitive-based scaffolding algorithm. *IEEE Transactions on Fuzzy Systems*, 23, 2048–2066. <https://doi.org/10.1109/TFUZZ.2015.2402683>
- Qin, J., Wang, C., Zou, Q., et al. (2021). Active learning with extreme learning machine for online imbalanced multiclass classification. *Knowledge-Based Systems*, 231, 107385. <https://doi.org/10.1016/j.knsys.2021.107385>
- Quade, D. (1979). Using weighted rankings in the analysis of complete blocks with additive block effects. *Journal of the American Statistical Association*, 74, 680. <https://doi.org/10.2307/2286991>
- Réda, C., Kaufmann, E., & Delahaye-Duriez, A. (2020). Machine learning applications in drug development. *Computational and Structural Biotechnology Journal*, 18, 241–252.
- Ren, P., Xiao, Y., Chang, X., et al. (2022). A survey of deep active learning. *ACM Computing Surveys*, 54, 1–40. <https://doi.org/10.1145/3472291>
- Reyes, O., Altalhi, A. H., & Ventura, S. (2018). Statistical comparisons of active learning strategies over multiple datasets. *Knowledge-Based Systems*, 145, 274–288. <https://doi.org/10.1016/j.knsys.2018.01.033>
- Riis, C., Antunes, F., Hüttel, F. B., et al. (2022). Bayesian active learning with fully Bayesian Gaussian processes. In *Proceedings of advances in neural information processing systems 35 (NeurIPS 2022)*. https://proceedings.neurips.cc/paper_files/paper/2022/file/4f1fba885f266d87653900fd3045e8af-Paper-Conference.pdf.
- Riquelme, C. (2017). Online active learning with linear models. Thesis. <http://purl.stanford.edu/rp382fv8012>.
- Riquelme, C., Ghavamzadeh, M., & Lazaric, A. (2017a). Active learning for accurate estimation of linear models. In *Proceedings of the 34th international conference on machine learning*. <http://proceedings.mlr.press/v70/riquelme17a/riquelme17a.pdf>.
- Riquelme, C., Johari, R., & Zhang, B. (2017b). Online active linear regression via thresholding. In *Thirty-first AAAI conference on artificial intelligence*. www.aaai.org.
- Rožanec, J. M., Trajkova, E., Dam, P., et al. (2022). Streaming machine learning and online active learning for automated visual inspection. *IFAC-PapersOnLine*, 55, 277–282. <https://doi.org/10.1016/j.ifacol.2022.04.206>
- Rosenblatt, F. (1958). The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65, 386–408. <https://doi.org/10.1037/h0042519>
- Roth, D., & Small, K. (2006). Margin-based active learning for structured output spaces. *Machine Learning: ECML, 2006*, 413–424. https://doi.org/10.1007/11871842_40
- Roy, N., & McCallum, A. (2001). Toward optimal active learning through sampling estimation of error reduction. In *Proceedings of the eighteenth international conference on machine learning*. <https://dl.acm.org/doi/10.5555/645530.655646>.
- Ruan, Y., Yang, J., & Zhou, Y. (2020). Linear bandits with limited adaptivity and learning distributional optimal design. In *STOC 2021: Proceedings of the 53rd annual ACM SIGACT symposium on theory of computing*. <https://doi.org/10.1145/3406325.3451004>.
- Rudovic, O., Zhang, M., Schuller, B., et al. (2019). Multi-modal active learning from human data: A deep reinforcement learning approach. In *2019 international conference on multimodal interaction* (pp. 6–15). <https://doi.org/10.1145/3340555.3353742>.
- Saran, A., Yousefi, S., Krishnamurthy, A., et al. (2023). Streaming active learning with deep neural networks. In Krause, A., Brunkskill, E., Cho, K., et al. (Eds.), *Proceedings of the 40th international conference on machine learning, proceedings of machine learning research*. PMLR (Vol. 202, pp. 30005–30021). <https://proceedings.mlr.press/v202/saran23a.html>.
- Schmidt, S., Rao, Q., Tatsch, J., et al. (2020). Advanced active learning strategies for object detection. In *2020 IEEE intelligent vehicles symposium (IV)* (pp. 871–876). <https://doi.org/10.1109/IV47402.2020.9304565>.

- Schmitt, R., Jatzkowski, P., & Peterrek, M. (2013). Traceable measurements using machine tools. In *Laser metrology and machine performance X: 10th international conference and exhibition on laser metrology, machine tool, CMM & robotic performance, Lamdamap* (pp. 20–21).
- Sculley, D. (2007). Online active learning methods for fast label efficient spam filtering. In *Proceedings of the fourth conference on email and antispan*.
- Sener, O., & Savarese, S. (2017). Active learning for convolutional neural networks: A core-set approach. In *ICLR*.
- Settles, B. (2009). *Active learning literature survey*. Technical report 1648, University of Wisconsin-Madison Department of Computer Sciences. <https://burrsettles.com/pub/settles.activelearning.pdf>.
- Seung, H. S., Opper, M., & Sompolinsky, H. (1992). Query by committee. In *Proceedings of the fifth annual workshop on computational learning theory—COLT '92* (pp. 287–294). <https://doi.org/10.1145/130385.130417>.
- Shah, K., & Manwani, N. (2020). Online active learning of reject option classifiers. In *Proceedings of the AAAI conference on artificial intelligence* (Vol. 34, pp. 5652–5659). <https://doi.org/10.1609/aaai.v34i04.6019>.
- Shan, J., Zhang, H., Liu, W., et al. (2019). Online active learning ensemble framework for drifted data streams. *IEEE Transactions on Neural Networks and Learning Systems*, 30, 486–498. <https://doi.org/10.1109/TNNLS.2018.2844332>
- Shannon, E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, 27, 379–423.
- Sheng, V. S., Provost, F., & Ipeirotis, P. G. (2008). Get another label? Improving data quality and data mining using multiple, noisy labelers. In *Proceeding of the 14th ACM SIGKDD international conference on knowledge discovery and data mining—KDD 08* (p. 614). <https://doi.org/10.1145/1401890.1401965>.
- Shi, X., & Xiong, W. (2018). Approximate linear dependence criteria with active learning for smart soft sensor design. *Chemometrics and Intelligent Laboratory Systems*, 180, 88–95. <https://doi.org/10.1016/j.chemolab.2018.07.009>
- Shilton, A., Palaniswami, M., Ralph, D., et al. (2005). Incremental training of support vector machines. *IEEE Transactions on Neural Networks*, 16, 114–131. <https://doi.org/10.1109/TNN.2004.836201>
- Soare, M., Lazaric, A., & Munos, R. (2013). Active learning in linear stochastic bandits. Bayesian Optimization in Theory and Practice https://www.univ-orleans.fr/lifo/Members/soare/files/active_learning_linear_bandit.pdf.
- Soare, M., Lazaric, A., & Munos, R. (2014). Best-arm identification in linear bandits. In *27th Conference on neural information processing systems (NeurIPS 2014)*.
- Song, S., Chaudhuri, K., & Sarwate, A. D. (2013). Stochastic gradient descent with differentially private updates. In *2013 IEEE global conference on signal and information processing* (pp. 245–248). <https://doi.org/10.1109/GlobalSIP.2013.6736861>.
- Souza, V., Pinho, T., & Batista, G. (2018). Evaluating stream classifiers with delayed labels information. In *2018 7th Brazilian conference on intelligent systems (BRACIS)* (pp. 408–413). <https://doi.org/10.1109/BRACIS.2018.00077>.
- Steel, R. G. D. (1959). A multiple comparison sign test: Treatments versus control. *Journal of the American Statistical Association*, 54, 767. <https://doi.org/10.2307/2282500>
- Steve, H., & Liu, Y. (2014). Minimax analysis of active learning. *Journal of Machine Learning Research*. <https://www.jmlr.org/papers/volume16/hanneke15a/hanneke15a.pdf>.
- Subramanian, K., Das, A. K., Sundaram, S., et al. (2014). A meta-cognitive interval type-2 fuzzy inference system and its projection based learning algorithm. *Evolving Systems*, 5, 219–230. <https://doi.org/10.1007/s12530-013-9102-9>
- Sudarsanam, N., & Ravindran, B. (2018). Using linear stochastic bandits to extend traditional offline designed experiments to online settings. *Computers & Industrial Engineering*, 115, 471–485.
- Suresh, S., Sundararajan, N., & Saratchandran, P. (2008). Risk-sensitive loss functions for sparse multi-category classification problems. *Information Sciences*, 178, 2621–2638. <https://doi.org/10.1016/j.ins.2008.02.009>
- Suárez-Cetrulo, A. L., Kumar, A., & Miralles-Pechuán, L. (2021). Modelling the covid-19 virus evolution with incremental machine learning. In *29th Irish conference on artificial intelligence and cognitive science, AICS 2021*. <https://ceur-ws.org/Vol-3105/paper1.pdf>.
- Suárez-Cetrulo, A. L., Quintana, D., & Cervantes, A. (2023). A survey on machine learning for recurring concept drifting data streams. *Expert Systems with Applications*, 213, 118934. <https://doi.org/10.1016/j.eswa.2022.118934>
- Suzuki, K., Sunagawa, T., Sasaki, T., et al. (2021). Annotation cost reduction of stream-based active learning by automated weak labeling using a robot arm. In *2021 IEEE/RSJ international conference on intelligent robots and systems (IROS)* (pp. 9000–9007). <https://doi.org/10.1109/IROS51168.2021.9636355>.

- Tang, Q., Li, D., & Xi, Y. (2018). A new active learning strategy for soft sensor modeling based on feature reconstruction and uncertainty evaluation. *Chemometrics and Intelligent Laboratory Systems*, 172, 43–51. <https://doi.org/10.1016/j.chemolab.2017.11.001>
- Taylor, G., & Hinton, G. (2009). Factored conditional restricted Boltzmann machines for modeling motion style. In *Proceedings of the 26th international conference on machine learning, Montreal, Canada, 2009*. <https://doi.org/10.1145/1553374.1553505>.
- Taylor, G., Hinton, G., & Roweis, S. (2006). Modeling human motion using binary latent variables. In *Advances in neural information processing systems 19 (NIPS 2006)*. https://papers.nips.cc/paper_files/paper/2006/hash/1091660f3dff84fd648efe31391c5524-Abstract.html.
- Thompson, J., Walters, W. P., Feng, J. A., et al. (2022). Optimizing active learning for free energy calculations. *Artificial Intelligence in the Life Sciences*, 2, 100050. <https://doi.org/10.1016/j.ails.2022.100050>
- Tieppo, E., dos Santos, R. R., Barddal, J. P., et al. (2022). Hierarchical classification of data streams: A systematic literature review. *Artificial Intelligence Review*, 55, 3243–3282. <https://doi.org/10.1007/s10462-021-10087-z>
- Tong, S., & Koller, D. (2002). Support vector machine active learning with applications to text classification. *The Journal of Machine Learning Research*. <https://doi.org/10.1162/153244302760185243>
- Tran, T., Pham, T., Carneiro, G., et al. (2017). A Bayesian data augmentation approach for learning deep models. In *31st conference on neural information processing systems (NIPS 2017)*. https://proceedings.neurips.cc/paper_files/paper/2017/file/076023edc9187cflac1f1163470e479a-Paper.pdf.
- Tran, T., Do, T. T., Reid, I., et al. (2019). Bayesian generative active deep learning. In *Proceedings of the 36th international conference on machine learning*. [arXiv:1904.11643](https://arxiv.org/abs/1904.11643).
- Tsybakov, A. B. (2004). Optimal aggregation of classifiers in statistical learning. *The Annals of Statistics*. <https://doi.org/10.1214/aos/1079120131>
- Tsymbal, A., Pechenizkiy, M., Cunningham, P., et al. (2008). Dynamic integration of classifiers for handling concept drift. *Information Fusion*, 9, 56–68. <https://doi.org/10.1016/j.inffus.2006.11.002>
- Vahdat, A., Belbahri, M., & Nia, V. P. (2019). Active learning for high-dimensional binary features. In *15th international conference on network and service management (CNSM)*. <https://www.computer.org/csdl/proceedings-article/cnsm/2019/09012676/1hQr3hscsJG>.
- Vanhatalo, E., & Kulahci, M. (2016). Impact of autocorrelation on principal components and their use in statistical process control. *Quality and Reliability Engineering International*, 32, 1483–1500. <https://doi.org/10.1002/qre.1858>
- Vanhatalo, E., Kulahci, M., & Bergquist, B. (2017). On the structure of dynamic principal component analysis used in statistical process monitoring. *Chemometrics and Intelligent Laboratory Systems*, 167, 1–11. <https://doi.org/10.1016/j.chemolab.2017.05.016>
- Wang, L. (2011). Smoothness, disagreement coefficient, and the label complexity of agnostic active learning. *The Journal of Machine Learning Research*. <https://www.jmlr.org/papers/volume12/wang11b/wang11b.pdf>.
- Wang, X., Fu, M., Ma, H., et al. (2015). Lateral control of autonomous vehicles based on fuzzy logic. *Control Engineering Practice*, 34, 1–17.
- Wassermann, S., Cuvelier, T., & Casas, P. (2019). Ral-improving stream-based active learning by reinforcement learning. <https://hal.archives-ouvertes.fr/hal-02265426>.
- Weigl, E., Heidl, W., Lughofer, E., et al. (2016). On improving performance of surface inspection systems by online active learning and flexible classifier updates. *Machine Vision and Applications*, 27, 103–127. <https://doi.org/10.1007/s00138-015-0731-9>
- Wilcoxon, F. (1945). Individual comparisons by ranking methods. *Biometrics Bulletin*, 1, 80. <https://doi.org/10.2307/3001968>
- Woodward, M., & Finn, C. (2017). Active one-shot learning. In *NIPS 2016, deep reinforcement learning workshop*. [arXiv:1702.06559](https://arxiv.org/abs/1702.06559).
- Woźniak, M., Zybiewski, P., & Ksieniewicz, P. (2023). Active weighted aging ensemble for drifted data stream classification. *Information Sciences*, 630, 286–304.
- Wu, J., Chen, J., & Huang, D. (2022). Entropy-based active learning for object detection with progressive diversity constraint. In *2022 IEEE/CVF conference on computer vision and pattern recognition (CVPR)*. <https://doi.org/10.1109/CVPR52688.2022.00918>.
- Wu, R., Guo, C., Su, Y., et al. (2021). Online adaptation to label distribution shift. In *35th conference on neural information processing systems (NeurIPS 2021)*. <https://www.kaggle.com/Cornell-University/arxiv>.
- Wu, Y., Chen, Y., Wang, L., et al. (2019). Large scale incremental learning. In *2019 IEEE/CVF conference on computer vision and pattern recognition (CVPR)*. <https://doi.org/10.1109/CVPR.2019.00046>.
- Xu, W., Zhao, F., & Lu, Z. (2016). Active learning over evolving data streams using paired ensemble framework. In *2016 eighth international conference on advanced computational intelligence (ICACI)* (pp. 180–185). <https://doi.org/10.1109/ICACI.2016.7449823>.

- Yan, X., Sarkar, M., Lartey, B., et al. (2023). An online learning framework for sensor fault diagnosis analysis in autonomous cars. *IEEE Transactions on Intelligent Transportation Systems*. <https://doi.org/10.1109/TITS.2023.3305620>
- Yin, C., Chen, S., & Yin, Z. (2023). Clustering-based active learning classification towards data stream. *ACM Transactions on Intelligent Systems and Technology*, 14(2), 1–18.
- Yu, H., Sun, C., Yang, W., et al. (2015). AI-elm: One uncertainty-based active learning algorithm using extreme learning machine. *Neurocomputing*, 166, 140–150. <https://doi.org/10.1016/j.neucom.2015.04.019>
- Yu, K., Bi, J., & Tresp, V. (2006). Active learning via transductive experimental design. In *Proceedings of the 23rd international conference on machine learning*. <https://doi.org/10.1145/1143844.1143980>.
- Yuan, M., Lin, H. T., & Boyd-Graber, J. (2020). Cold-start active learning through self-supervised language modeling. In *Proceedings of the 2020 conference on empirical methods in natural language processing (EMNLP)*. <https://doi.org/10.18653/v1/2020.emnlp-main.637>.
- Zhang, H., Liu, W., Shan, J., et al. (2018). Online active learning paired ensemble for concept drift and class imbalance. *IEEE Access*, 6, 73815–73828. <https://doi.org/10.1109/ACCESS.2018.2882872>
- Zhang, H., Liu, W., Sun, L., et al. (2020a). Analyzing network traffic for protocol identification: An ensemble online active learning method. In *Proceedings—2020 6th international conference on big data and information analytics, BigDIA 2020* (pp. 167–172). <https://doi.org/10.1109/BigDIA51454.2020.00035>.
- Zhang, H., Ravi, S. S., & Davidson, I. (2020b). A graph-based approach for active learning in regression. In *Proceedings of the 2020 SIAM international conference on data mining (SDM)*. <https://doi.org/10.1137/1.9781611976236.32>.
- Zhang, H., Liu, W., & Liu, Q. (2022). Reinforcement online active learning ensemble for drifting imbalanced data streams. *IEEE Transactions on Knowledge and Data Engineering*, 34, 3971–3983. <https://doi.org/10.1109/TKDE.2020.3026196>
- Zhang, K., Liu, S., & Chen, Y. (2023). Online active learning framework for data stream classification with density-peaks recognition. *IEEE Access*, 11, 27853–27864.
- Zhang, T. (2004). Statistical behavior and consistency of classification methods based on convex risk minimization. *The Annals of Statistics*. <https://doi.org/10.1214/aos/1079120130>
- Zheng, Z., & Padmanabhan, B. (2006). Selectively acquiring customer information: A new data acquisition problem and an active learning-based solution. *Management Science*, 52(5), 697–712.
- Zhou, C., Ma, X., Michel, P., et al. (2021). Examining and combating spurious features under distribution shift. In *Proceedings of the 38th international conference on machine learning*. <https://github.com/violet-zct/>.
- Zhu, J. J., & Bento, J. (2017). Generative adversarial active learning. [arXiv:1702.07956](https://arxiv.org/abs/1702.07956).
- Zhu, X., Zhang, P., Lin, X., et al. (2007). Active learning from data streams. In *Proceedings—IEEE international conference on data mining, ICDM* (pp. 757–762). <https://doi.org/10.1109/ICDM.2007.101>.
- Zliobaite, I., Bifet, A., Pfahringer, B., et al. (2014). Active learning with drifting streaming data. *IEEE Transactions on Neural Networks and Learning Systems*, 25, 27–39. <https://doi.org/10.1109/TNNLS.2012.2236570>
- Zwanka, R. J., & Buff, C. (2021). Covid-19 generation: A conceptual framework of the consumer behavioral shifts to be caused by the covid-19 pandemic. *Journal of International Consumer Marketing*, 33, 58–67. <https://doi.org/10.1080/08961530.2020.1771646>
- Zyblewski, P., Ksieniewicz, P., & Woźniak, M. (2020). Combination of active and random labeling strategy in the non-stationary data stream classification. In *International conference on artificial intelligence and soft computing* (pp. 576–585). Springer.