

<https://doi.org/10.1145/3593013.3593984>

This item is likely protected under Title 17 of the U.S. Copyright Law. Unless on a Creative Commons license, for uses protected by Copyright Law, contact the copyright holder or the author.

Access to this work was provided by the University of Maryland, Baltimore County (UMBC) ScholarWorks@UMBC digital repository on the Maryland Shared Open Access (MD-SOAR) platform.

Please provide feedback

Please support the ScholarWorks@UMBC repository by emailing scholarworks-group@umbc.edu and telling us what having access to this work means to you and why it's important to you. Thank you.

Preventing Discriminatory Decision-making in Evolving Data Streams

ZICHONG WANG, Michigan Technological University, USA

NRIPSUTA SAXENA, University of Southern California, USA

TONGJIA YU, Columbia University, USA

SNEHA KARKI, Michigan Technological University, USA

TYLER ZETTY, Michigan Technological University, USA

ISRAAT HAQUE, Dalhousie University, Canada

SHAN ZHOU, Michigan Technological University, USA

DUKKA KC, Michigan Technological University, USA

IAN STOCKWELL, University of Maryland, Baltimore County, USA

ALBERT BIFET, University of Waikato, New Zealand

WENBIN ZHANG, Michigan Technological University, USA

Bias in machine learning has rightly received significant attention over the last decade. However, most fair machine learning (fair-ML) work to address bias in decision-making systems has focused solely on the offline setting. Despite the wide prevalence of online systems in the real world, work on identifying and correcting bias in the online setting is severely lacking. The unique challenges of the online environment make addressing bias more difficult than in the offline setting. First, Streaming Machine Learning (SML) algorithms must deal with the constantly evolving real-time data stream. Second, they need to adapt to changing data distributions (concept drift) to make accurate predictions on new incoming data. Adding fairness constraints to this already complicated task is not straightforward. In this work, we focus on the challenges of achieving fairness in biased data streams while accounting for the presence of concept drift, accessing one sample at a time. We present *Fair Sampling over Stream (FS²)*, a novel fair rebalancing approach capable of being integrated with SML classification algorithms. Furthermore, we devise the first unified performance-fairness metric, *Fairness Bonded Utility (FBU)*, to evaluate and compare the trade-off between performance and fairness of different bias mitigation methods efficiently. FBU simplifies the comparison of fairness-performance trade-offs of multiple techniques through one unified and intuitive evaluation, allowing model designers to easily choose a technique. Overall, extensive evaluations show our measures surpass those of other fair online techniques previously reported in the literature.

Additional Key Words and Phrases: Fairness, Data Stream, Concept Drift, Fairness Drift

ACM Reference Format:

Zichong Wang, Nripsuta Saxena, Tongjia Yu, Sneha Karki, Tyler Zetty, Israat Haque, Shan Zhou, Dukka Kc, Ian Stockwell, Albert Bifet, and Wenbin Zhang. 2023. Preventing Discriminatory Decision-making in Evolving Data Streams. 1, 1 (February 2023), 17 pages. <https://doi.org/10.1145/xxxxx.xxxxxx>

Authors' addresses: Zichong Wang, zichongw@mtu.edu, Michigan Technological University, Houghton, Michigan, USA, 49931; Nripsuta Saxena, nsaxena@usc.edu, University of Southern California, Los Angeles, California, USA, 90007; Tongjia Yu, tongjia2428@columbia.edu, Columbia University, New York, New York, USA, 10027; Sneha Karki, snehak@mtu.edu, Michigan Technological University, Houghton, Michigan, USA, 49931; Tyler Zetty, tjzetty@mtu.edu, Michigan Technological University, Houghton, Michigan, USA, 49931; Israat Haque, israat@dal.ca, Dalhousie University, Halifax, Nova Scotia, Canada, B3H 4R2; Shan Zhou, shanzhou@mtu.edu, Michigan Technological University, Houghton, Michigan, USA, 49931; Dukka Kc, dbkc@mtu.edu, Michigan Technological University, Houghton, Michigan, USA, 49931; Ian Stockwell, istock1@umbc.edu, University of Maryland, Baltimore County, Baltimore, Maryland, USA, 21250; Albert Bifet, abifet@waikato.ac.nz, University of Waikato, New Zealand; Wenbin Zhang, wenbinzh@mtu.edu, Michigan Technological University, Houghton, Michigan, USA, 49931.

2023. Manuscript submitted to ACM

Manuscript submitted to ACM

1 INTRODUCTION

Machine learning-based decision-making systems are constantly increasing in importance due to their wide use in fields as diverse as criminal justice, job application screening, loan decisions, and resource allocation. However, the growing adoption of automated decision-making systems has led to heightened scrutiny of these models' issues of fairness and accountability [12, 36, 49, 52, 56, 58]. The example of Amazon scrapping an automated hiring tool after it was found to be biased against women [28] is one of many [11, 30]. The harmful impacts of such biases makes it imperative to develop fairness-conscious classifiers that lead to accurate predictions free of discrimination against marginalized subgroups in society. While various strategies have been proposed for fair decision-making systems, for example, to identify and eliminate discrimination [19, 40, 55, 57], and ameliorate the intrinsic bias in historical data [43, 47, 48, 50], they ignore a crucial space that is all around us in the real-world: the online setting.

In the real world, data is generated in real-time, and its characteristics and distribution may vary over time. Hence, it is necessary to design non-discriminatory decision systems that consider the unique challenges of the online setting. However, most work in fair-ML approaches fairness as a static issue, and makes the assumption that all data can be afforded for multiple scans and that the underlying population characteristics do not evolve over time [19, 37, 38, 46, 54]. Therefore they don't address the challenges faced in the online setting. The first significant challenge is **handling the constantly produced, real-time data stream**. In other words, the data is infinite. Since no memory can load infinite data, fairness-aware learning for the online setting should be able to process each incoming instance "on arrival," without requiring storage and reprocessing. Most fair-ML, however, assumes the entire dataset is available to be scanned multiple times. The second significant challenge is that **the target concepts may also evolve over time**. In other words, the instance at time $t + 1$ may have been generated by a different function than the one that generated the instance at time t . For example, consider the following medical diagnostics scenario. In early 2020, most patients presenting with symptoms such as fever, cough, etc would likely have received a diagnosis of the seasonal flu. However, over time we became aware of Covid-19, and while in the beginning Covid-19 would have represented a minority of the cases as compared to flu, as time passed this new disease became the majority class. What was once the minority class became the majority class, and therefore the class we would have to balance in our dataset changed over time. Addressing bias as well as a changing distribution concurrently is difficult. The third significant challenge is **an intuitive measure that is capable of quantifying the trade-off between fairness and performance**. So far, fairness evaluations have primarily reported fairness improvement and accuracy loss as two separate metrics without a reliable method for jointly measuring the inherent trade-off between them. However, in practical real-world applications, business analytics for example require a single metric jointly for both [46]. Therefore, a unified and efficient measurement system is essential for clear comprehension of the trade-off between fairness and accuracy of the model.

To address the aforementioned challenges, we present Fair Sampling over Stream (FS^2), a novel meta-strategy for online stream-based decision-making and Fair Bonded Utility (FBU), a novel unified fairness-performance trade-off measurement. Our approach eliminates the bias in the data stream prior to applying the SML classification algorithm and is thus model agnostic. The novel fairness improvement method takes into account continuous data streams, eliminates discrimination, and accounts for evolving data distribution. Additionally, it strikes an effective and robust balance between prediction accuracy and fairness performance. Our major contributions are:

- FS^2 , a meta-strategy that builds on the popular SMOTE technique (only applicable to batches), and can be integrated with any data stream classifier. Unlike current approaches to fair stream learning that rely on SMOTE, FS^2 is able to add a fairness guarantee without performance degradation, and is not restricted to batches.
- A novel fairness-performance metric that assesses the fairness-accuracy trade-off of ML bias mitigation methods in a unified and intuitive manner.
- Qualitative and quantitative experiments on five groups of biased data streams show the effectiveness of the proposed unified metrics and fairness-conscious online learners in streaming environments.

The remainder of the paper is organized as follows. Section 2 and Section 3 review relevant work and the theoretical background knowledge regarding discrimination-aware learning. Next, we discuss the limitations of class balancing technique in Section 4. We present FS^2 and bias-performance measurement metric in Section 5. Section 6 analyzes the experimental results. Finally, we conclude the paper in Section 7.

2 RELATED WORK

2.1 Fairness-aware Learning

A number of approaches have been proposed to address the problem of bias and discrimination in machine learning, and can be broadly grouped under three categories depending on the part of the pipeline fairness interventions are applied to: pre-, in-, or post-processing. i) *Pre-processing techniques* focus on mitigating and correcting bias in the data used for training the model, and therefore are model-independent. The most popular ones include massaging [24] and reweighting [6], which have also been extended for addressing bias in the input data before updating the online model in [22]. ii) *In-processing techniques* transform an algorithm to improve fairness, typically by embedding fairness in the objective function by means of regularization or other constraints. [26] is among the first line of works in this category by integrating discrimination into the splitting criteria for fair tree induction. More recently, [51] further improves this splitting criteria and applies it in the online settings. iii) *Post-processing techniques* are based on either adapting the decision boundary or simply modifying the output of a model (e.g. the prediction labels). With supplemental prediction thresholds in place, [20] works against discrimination while in [7] the decision boundary of AdaBoost is adjusted with regards to fairness. We would like to stress on the fact that transferring such approaches to an online setup is not as simple since boundary/prediction might evolve due to the non-stationary spread in an online setting.

2.2 Stream Learning

The main challenges of learning in a stream setting are concept drift, where the joint data distribution changes over time and class imbalance [1, 15, 27]. To address the first challenge, learning methods need to adapt to changes incrementally by incorporating new information into the model [32, 41, 42], and by forgetting outdated information [16, 17, 53]. For the second challenge, various sampling methods have been proposed including the most representative (C-SMOTE) [3] to address class imbalance in a continuous data stream while ensuring the model can adapt to shifts in the data distribution. However, these techniques focus entirely on accuracy without considering fairness.

2.3 Fairness-aware Stream Learning

Despite the wide prevalence of the stream learning setting as well as bias and discrimination within the real world, surprisingly little work has explored this setting so far. In addition to the aforementioned fair online methods, [2] focuses on individual level and requires the existence of an auditor to detect fairness violations, while Online SMOTE Boost [39] is made cost-sensitive by adding various parameters of the Poisson distribution for different classes to address bias. In addition, FABBOO [23] adjusts the training distribution on-the-fly, considering both the imbalanced nature

of the stream and the model’s discriminatory behavior as evaluated from past data. One common drawback to all of these methods is that their balancing approaches are driven by the SMOTE which could further exacerbate bias (c.f., Section 4). Our method addresses this drawback by performing an additional clustering for fairness guarantee.

3 BASIC CONCEPTS

Let D be a sequence of instances d_1, d_2, \dots, d_n that arrive over time at timepoints t_1, t_2, \dots, t_n , where each instance $d \in \mathbb{R}$. Similarly, let Y be a sequence of corresponding class labels, such that for each $d \in D$ there exists a corresponding class label $y \in Y$. We assume each data instance in the continuous stream has the form $d = \{A, S, Y\}$, where A is a set of attributes, S denotes the sensitive attribute (e.g. gender/race), and Y is the true class label. In fair-ML, a sensitive attribute is a characteristic legally protected from discrimination (e.g., race, color, gender, etc). Typically, one sensitive attribute class has been historically marginalized, such as people of color or women, and another that has not. We refer to the former as the *unprivileged group* and the latter as the *privileged group*. Without loss of generality, we assume a binary classifier, $f() : D \rightarrow y$, such that $y \in \{-1, +1\}$. The class label predicted by this classifier shall be denoted by \hat{Y} .

Since we explore the *online learning environment*, each new data instance from the continuous data stream shall be processed one by one. A characteristic of the online learning setting is that the true class label of each instance is revealed to the learner before the arrival of the next instance. Therefore, the way things proceed in the online setting is as follows. For each new data instance d arriving at time point t , a class label is predicted using f_{t-1} , i.e. the classifier from time point $t - 1$. The true class label of data point d is revealed to the online learner before the arrival of the next data instance. The model is then updated using this true class label to get f_t . This sequence of events is referred to as the *first-test-then-train* or *prequential evaluation* setup [14].

Another characteristic typical to the online setting is the idea of *concept drift*. Concept drift refers to the changes in underlying distribution of the continuous data stream over time. In other words, it denotes the changes in the joint distribution. Mathematically, $P_{t_a}(D, Y) \neq P_{t_b}(D, Y)$ for two different timepoints t_a and t_b . Such changes often render the current classifier ineffective, since it is unable to deal with the changes in the data distribution, thereby necessitating a model update. In this work we focus on the scenario where the continuous data stream is also imbalanced, i.e. one class is much more represented than the other. Machine learning models typically neglect the minority class in such situations to prevent overfitting and loss of generalization [44]. Our technique does not require the minority class to be predefined in advance, nor does it assume that the same class will always be the minority class. In other words, we allow for the possibility that role of the minority class may alternate between the two classes. We analyze the extent of the imbalance in data by computing the imbalance ratio (IR) with Equation 1:

$$IR = \frac{S_{t_{min}}}{S_{t_{maj}} + S_{t_{min}}} \quad (1)$$

where $S_{t_{min}}$ is the minority class and $S_{t_{maj}}$ is the majority class in the data stream at time t .

The traditional approach to fairness-aware classification in the offline setting seeks to learn a mapping, $F : (A, S, Y) \Rightarrow Y$, where the goal is to assign a class label to each data instance accurately without discriminating against the unprivileged group. The desirable class label (e.g. receiving a loan in loan decisions; receiving bail in bail decisions) is known as the ‘*favorable*’ label, and the other, undesirable class label (e.g., being denied a loan or bail) is called the ‘*unfavorable*’ label. We focus on parity-based approaches that compare a fair-ML model’s behavior towards the unprivileged and privileged groups. The two measures we employ are statistical parity [25], and equal opportunity [20]. We extend them to the online setting.

Statistical parity quantifies the disparity between the probability of the privileged and unprivileged groups receiving a benefit. Cumulative Statistical Parity Difference (*CSPD*) in Equation 2 represents the cumulative difference in statistical parity between the privileged and unprivileged groups over time. $\lambda \in [0, 1]$ is a decay factor used to regulate the amount of influence of previous time points. In other words, the larger λ , the higher the contribution of historical information. It is specified by the user.

$$CSPD_t = (1 - \lambda) \times \left(\frac{\sum_{i=1}^t (F_i(d_i) \wedge Y = 1 | S = \bar{s})}{\sum_{i=1}^t (d_i | S = \bar{s})} - \frac{\sum_{i=1}^t (F_i(d_i) \wedge Y = 1 | S = s)}{\sum_{i=1}^t (d_i | S = s)} \right) + \lambda \times CSPD_{t-1} \quad (2)$$

However, statistical parity may result in the phenomenon of reverse discrimination in certain conditions (i.e. the scenario when data instances that are not deserving of a positive assignment receive one incorrectly), since it does not consider real class labels. Therefore, we also consider Equal Opportunity, which measures the difference in True Positive Rates (TPR) between the privileged and unprivileged groups. Cumulative Equal Opportunity Difference *CEOD* in Equation 3 is the cumulative difference in equal opportunity between the privileged and unprivileged groups over time.

$$CEOD_t = (1 - \lambda) \times \left(\frac{\sum_{i=1}^t (F_i(d_i) \wedge Y = 1 | S = \bar{s}, Y = 1)}{\sum_{i=1}^t (d_i | S = \bar{s}, Y = 1)} - \frac{\sum_{i=1}^t (F_i(d_i) \wedge Y = 1 | S = s, Y = 1)}{\sum_{i=1}^t (d_i | S = s, Y = 1)} \right) + \lambda \times CEOD_{t-1} \quad (3)$$

where \wedge is the logical symbol and λ is used for correction in the early stages of the stream (division by 0).

4 EFFECTS OF CLASS BALANCING ON FAIRNESS

Inequality in class distribution is a hallmark of biased data. As the instances in the underprivileged group are not as prevalent as the privileged group, the model may overlook patterns in the underprivileged group classification. Since machine learning is data-driven, it is easy for a model trained on a biased dataset to assign favorable labels to the privileged subgroup and unfavorable labels to the unprivileged subgroup in pursuit of higher accuracy. To address unequal representation, class balancing techniques such as oversampling and undersampling are used. Oversampling increases the minority class representation by creating synthetic instances, while undersampling balances representation by reducing majority class instances. We examine the drawbacks of these techniques and how they affect the bias and fairness in a model's output in detail. While certain limitations of popular class-balancing techniques are well known (e.g., oversampling may lead to overfitting if samples from a few classes are replicated multiple times [8, 13, 45]; undersampling may lead to hidden information being ignored due to the samples that were removed [9, 34, 53]), some other drawbacks are not as well studied, particularly in terms of fairness research when alleviating bias.

Traditional class-balancing techniques make the minority class samples statistically equal to the majority class samples. However, since they aim to solve the imbalance only in terms of the target variable (i.e. the class label), this process is blind to the inherent properties of the dataset. As a result, the bias in the dataset may actually increase after balancing the class labels, causing the model's fairness to decrease further. This is especially problematic in the case of multiple privileged and unprivileged subgroups. Not all privileged subgroups are privileged in the same way and/or to the same extent, nor are all unprivileged subgroups unprivileged in the same way, and/or to the same extent. Undersampling can lead to a disproportional increase

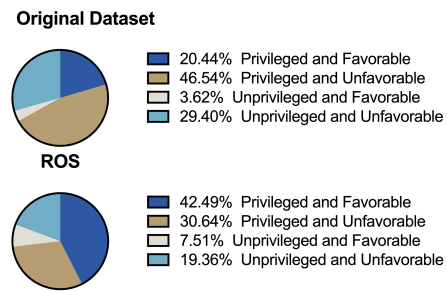


Fig. 1. The change of distribution bias in the Bank Marketing dataset before and after class balancing.

in the number of instances of the more favorable privileged group.

On the other hand, oversampling may result in a disproportionate increase in the number of instances of the more favored unprivileged group because they may be a bigger proportion of the minority class, leaving the worst favored unprivileged group severely underrepresented. After extensive experiments we see in Fig. 1 that using common class balancing techniques can indeed lead to more distribution bias in the dataset. In Fig. 1 the original Bank Marketing dataset, only 24% of the samples are positive (subscribe) while the rest are negative (non-subscribe). After class balancing with Random Over Sampling (ROS), 11,788 more positive samples were added, of which only 15% were unprivileged subgroup. This increases bias in the distribution, as more privileged subgroup samples were added during class balancing. In other words, the representation of unprivileged favorable groups in the data set declines further. Fig. 2 represents a limitation of another common class balancing technique, SMOTE. As the Figure 2 shows, the minority class samples synthesized by SMOTE may belong to the majority class in the feature space. Meanwhile, random selection of samples for synthesizing instances can lead to further in-class bias in the dataset, especially in datasets with multiple sensitive attributes.

5 METHODOLOGY

In this section, first we introduce our proposed performance-fairness metric, Fairness Bonded Utility (FBU). FBU is the first metric to unify the fairness-performance trade-off into one intuitive and versatile metric capable of comparing any combination of fairness and performance metrics. Then we present Fair Sampling over Stream (FS^2), a fair rebalancing technique for a biased data stream capable of being integrated with any stream-learning classification technique.

5.1 Fairness Bonded Utility

A long-standing problem in fair-ML has been the trade-off dance that occurs between model performance and fairness. Typically, a model's performance degrades as fairness guarantees become stronger. This complicates the decision about which fairness metric to use for which application since analyzing this trade-off between two different numbers (performance and fairness) is not always intuitive. Fairness Bonded Utility (FBU) solves this problem by assigning each fairness technique into one of five intuitive "effectiveness" levels ranging from "win-win" (fairness and performance both improve) to "lose-lose" (both degrade). At a high level, FBU does this by measuring how the performance of a model varies with different fairness techniques in a two-dimensional coordinate space and computing a trade-off baseline that categorizes the various fairness techniques into effectiveness levels. FBU can drastically simplify the agonizing decision-making process of which fairness technique to use for which application scenario.

Before going into details, we first construct the environment for the FBU to measure the performance of different fairness techniques. Mathematically, it is a two-dimensional coordinate system (see Figure

3). The x -axis represents the model's bias, and the y -axis depicts the model's performance. Our technique is versatile and can be used with any fairness metrics to measure bias (e.g., CSPD, CEOD, etc.) and any performance metrics (e.g.,

Manuscript submitted to ACM




Fig. 2. Illustration of oversampling by SMOTE and FS^2 . Circle and Triangular nodes symbolize samples with majority and minority labels, respectively. The hollow nodes are samples created based on the closest neighbors of minority samples within each cluster.

accuracy, f1-score, etc.) of the model designer’s choosing. Next, we outline the creation of the trade-off baseline and the evaluation of the outcomes.

Creating a Trade-off Baseline: The idea behind the trade-off baseline in FBU is centered around the zero-normalization principle presented by Speicher et al. [35]. It posits that bias is minimized when every individual receives the same label. In other words, bias would be 0 when everyone receives the same assignment as the output, although this may lead to decreased performance, and vice versa.

We use this principle to create multiple pseudo-models, F_p , where $p \in \{10\%, 20\%, \dots, 100\%\}$. In each pseudo-model, we substitute a proportion p of the original model’s predictions with the same label. We vary the proportion of substituted predictions in each pseudo model, with p ranging from 10% to 100%. In other words, we will replace 10% of the original model’s predictions with the same label in F_{10} . In F_{100} we replace 100% of the original model’s predictions with the same label. As the proportion of predictions substituted increases, the model’s fairness will increase but its accuracy will decrease.

Now we use the original model (with no fairness technique) and the created pseudo models to form the trade-off baseline. We derive the first point for the trade-off baseline from the (performance, bias) coordinate of the original model (point F_{ori} in Fig. 3(a)). Then we plot the (performance, bias) coordinates of the ten pseudo models (i.e., F_{10} , F_{30} , F_{60} , etc. in Fig. 3(a)). Finally we connect them all together to form the trade-off baseline.

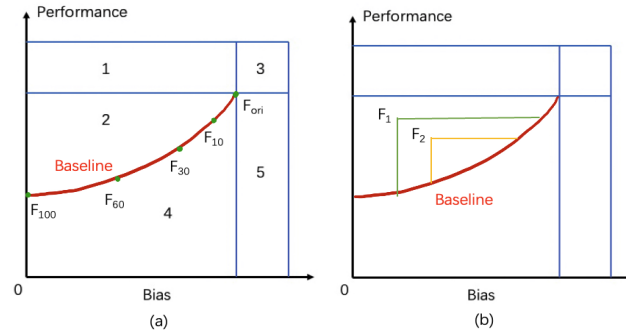


Fig. 3. The FBU fairness-accuracy trade-off baseline is depicted by the original trade-off point (F_{ori}) and the points generated by the pseudo models (F_{10}, \dots, F_{100}). A bias reduction method is considered effective if it shows a superior trade-off compared to the FBU baseline (i.e., it lies above the red line).

Five effectiveness levels: The trade-off baseline categorizes the different bias mitigation techniques being compared into five intuitive effectiveness levels, each with a different bias-performance trade-off. The first level, region 1 in Fig. 3(a), is a “win-win”: A technique belongs in this category if it improves both model performance and fairness relative to the trade-off baseline. Region 2 in Fig. 3(a) denotes a “good” trade-off: a technique in region 2 enhances either model performance or fairness compared to the trade-off baseline, making it overall better than the trade-off baseline. If a technique improves model performance but leads to a decrease in fairness, it is considered an “inverted” trade-off and falls into region 3 (Fig. 3(a)). A technique in region 4 in Fig. 3(a) is a “poor” trade-off: it leads to a decline in either model performance or fairness compared to the baseline, making it overall worse than the trade-off baseline. A region 5 technique (in Fig. 3(a)) represents a “lose-lose” trade-off since it decreases both model performance and fairness compared to the original model.

Quantitative Evaluation of Trade-offs: The “win-win” (region 1), “lose-lose” (region 5), and “poor” (region 4) trade-off regions offer clear indications of the effectiveness of the bias mitigation technique. Therefore, we concentrate on quantifying the trade-off quality of bias mitigation techniques in the “good” trade-off region to enable a more detailed comparison. Specifically, we will measure the strengths and weaknesses of different bias mitigation methods that fall in the “good” region (region 2) based on the area enclosed by the bias-performance points and baseline (see Fig. 3(b)). Techniques with larger areas have better, more desirable bias-performance trade-offs. Using area as a measure of trade-off, instead of other criteria such as distance from the trade-off baseline, guarantees a fair comparison even in cases where the trade-off baseline is curved.

The final output of FBU for a technique will be five percentages: one percentage for each region representing how many cases for that technique fell in that region. The number of cases per region is computed by: $n_r \times n_t \times n_f \times p_m$, where n_r is the number of run times, n_t number of techniques compared against, n_f number of fairness metrics used, and p_m is the number of performance metrics used.

5.2 Fair Sampling over Stream

There are two fundamental challenges facing Fair Sampling over Stream (FS^2): the first is the lack of the entire dataset in the rebalancing phase. The second is synthesizing fair data. We first discuss the online monitoring of the continuous data stream and model updates. Then we elaborate on the challenge of synthesizing fair data. We follow with the FS^2 algorithm, and end the section with explaining FS^2 classification.

5.2.1 Online Monitoring and Model Update. In continuous data streams, the status of minority and majority classes can evolve. In other words, a class currently identified as a minority may turn into the majority class, or vice versa [41]. Additionally, approaches focusing solely on fairness in recent outcomes (i.e., in the immediate future) are inadequate in combating discrimination over time. Although individual instances of discrimination may seem minor when evaluated alone, they can accumulate and result in substantial bias in the long term. Hence, monitoring the subgroup distribution in the data stream over time is vital for both the efficiency and fairness of the model. However, the primary challenge faced by FS^2 is its inability to consistently monitor data to address both concept drift and fairness concerns. In streaming machine learning, we assume the data stream to be infinite. Firstly, no memory can physically store infinite data; secondly, this would be against the stream paradigm approach. Therefore, storing every new sample in memory is impossible until the data stream ends.

To address this problem, we employ a different approach from previous non-stationary studies [10, 18]. FS^2 employs ADWIN [5] to detect changes in both accuracy and fairness, reflecting both concept and fairness drifts. In other words, drift is detected when either fairness or the data distribution evolves. Specifically, we will use ADWIN to save the most recently seen samples and apply the synthesis algorithm (explained below in Section 5.2.2) using the samples stored in ADWIN. ADWIN maintains a variable-length window of recently seen items and can automatically detect and adapt the window size to the current rate of change in the data. As shown in Equation 4, ADWIN employs a threshold known as δ (i.e., performance threshold and fairness threshold) to configure the error to have two levels automatically. These levels are known as the warning level and the change level.

$$levelError = \log\left(\frac{2 \times \log n}{\delta}\right) \quad (4)$$

We identify the warning level by multiplying δ by 10, and δ itself helps determine the change level. Since δ is in the denominator, multiplying it by 10 will result in a lower value than the one obtained by using δ alone. Therefore, we will

reach the warning level before the change level and can detect changes early to take appropriate action. The width of the window at any given moment is represented by n .

ADWIN monitors the error that occurs over the data in the window. If the error reaches a threshold of the warning level, ADWIN assumes that either a concept drift or fairness drift is starting to occur, and it starts collecting new samples in a new window. If the error exceeds the change level, ADWIN assumes either a concept drift or a fairness drift has occurred, and it substitutes the old window with the new one.

5.2.2 Synthetic fair data. Here we propose a new strategy for synthesizing fair data to address the challenges posed by class balancing and its impact on model fairness. Our proposed approach eliminates the in-class and between-class imbalance in the oversampling part. Unlike traditional class balancing techniques such as SMOTE, our method identifies the true minority classes in the dataset (see Figure 2), and adds fairness constraints to the generation process for synthetic samples. At a high level, rather than directly running SMOTE, our approach first determines similar samples via clustering, then filters the clusters using silhouette score to ensure a good clustering. Finally, it uses SMOTE within each cluster to generate synthetic minority samples. Notably, our approach not only enhances the fairness of the balancing method but also preserves the accuracy of the prediction. Below we discuss our approach in detail.

First, we use a clustering algorithm to identify homogeneous subgroups of feature similarity in the feature space. In our experiments, we use the KNN algorithm. Note that any clustering algorithm can be applied, but different clustering algorithms may have different clustering results and time complexity. By partitioning the dataset according to the selected number of clustering families M , samples with similarities in the feature space will be classified into the same clusters. We will measure the quality of clustering by silhouette analysis. The silhouette score measures how similar an object is to its own cluster compared to others, and the value is between $[-1, 1]$. A value of 1 indicates perfect separation, and a value of -1 suggests that the object is in the wrong cluster. Values between 0 and 1 show the degree of separation, with values closer to 1 indicating a stronger separation.

After determining the optimal number of clusters M , the algorithm performs filtering detection for each cluster. The algorithm will calculate the silhouette score of each sample and then remove samples with the lowest 20% silhouette score from the dataset. The intuition for this is as follows. Let's imagine a realistic data distribution situation where a sample of a minority group exists near the cluster boundaries. In this case, we can assume that SMOTE selects the minority group sample near the clustering boundary as the template. Then the simulated data generated based on this sample will also be close to the clustering boundary. In addition, there is a risk that the generated simulated data are even closer to the clustering boundary than the original data. This increases the number of samples near the cluster boundaries, thus blurring the boundaries between different clusters, making it difficult to separate these samples in the feature space and a higher probability of being assigned to incorrect clusters. As a result, new instances generated based on these samples may further increase the distribution bias.

After the filtering is complete, the algorithm checks each cluster. Specifically, the algorithm detects the proportion of minority class samples in each cluster to determine its synthetic weight, which is directly related to the proportion of minority samples in the cluster. If the percentage of minority samples in a cluster is higher, the more representative the cluster is for minority samples. In other words, clusters with a high percentage of minority class samples are assigned higher synthetic weights. For each cluster, the synthetic weight is calculated according to Equation 5:

$$ClusterWeight = \begin{cases} W_i = \frac{N_i}{N_{Total}} \\ W_1 + W_2 + \dots + W_n = 1 \end{cases} \quad (5)$$

where N_i is the number of minority class samples in each cluster, N_{Total} is the total number of minority class samples, and W_i is the synthetic weight of cluster i . The sum of all weights is one, and the number of samples selected for each cluster is calculated according to Equation 6:

$$G_i = N_{num} \times W_i \quad (6)$$

where N_{num} is the total number of samples to be synthesized. After calculating the number of samples selected for each cluster, the model will randomly select samples in the clusters, and a new minority sample is generated according to the SMOTE [9] generation algorithm based on the k-nearest neighbors of the minority samples.

5.2.3 FS^2 Algorithm. The algorithm for FS^2 is outlined in Algorithm 1 using pseudocode.

The FS^2 algorithm utilizes two sliding windows of samples, W and W_{label} (line 1). W contains all samples related to the current concept. W_{label} is an array indicating whether the corresponding samples in W belong to the privileged favorable subgroup, privileged unfavorable subgroup, unprivileged favorable subgroup, or unprivileged unfavorable subgroup. The FS^2 algorithm utilizes nine counters (lines 2-4): C_0 , C_1 , C_2 and C_3 are used to count the number of instances in each subgroup, while $\overline{C_0}$, $\overline{C_1}$, $\overline{C_2}$ and $\overline{C_3}$ count the number of instances of different subgroups generated by the synthesis algorithm, respectively. The C_{gen} counter for each sample d_i in the sliding window W keeps track of the number of times d_i is used to generate synthetic samples.

FS^2 tracks instances of different classes. As we mention above, after new samples are continuously collected, the number of samples of different classes in W can deviate significantly. In this case, the representation of different subgroups in W may change, FS^2 will synthesize new samples to balance the representation of different subgroups. In line 5, the variable $adwin$ is a change detector that ensures the two windows (W , W_{label}) remain consistent with any concept drift by using the class value of the incoming sample. $adwin$ is used later for adapting the model to concept and/or fairness drift as needed. In line 6 the variables $balR$, $fairR$ are initialized to 0.

For each new instance d_i (line 7), we first train the pipeline learner l using the new sample (line 8). After this process, the new sample d_i is saved in the sliding window W (line 9). On line 10, the function `updateWindows` adds a new bit into W_{label} depending on the new sample's class label Y (i.e., 0 if $Y = 0$). On line 11, the function `updateCounters` increments the relative counter C_0 , C_1 , C_2 , or C_3 depending on the sample's class label Y and sensitive attribute S .

After that, the $adwin$ is updated with the sensitive attribute S and class value Y of the incoming sample d_i (line 12) and the `checkDrift` function uses $adwin$ to check for any concept or fairness drift (line 13). The function adapts by reducing W in accordance with the window maintained by $adwin$ to the occurrence of concept drift or fairness drift. Additionally, it updates W_{label} and the counters C_0 , C_1 , C_2 and C_3 based on this reduction. Furthermore, if the instances removed from W have been used to create synthetic samples, the function also updates the counter of instances generated for that class ($\overline{C_0}$, $\overline{C_1}$, $\overline{C_2}$ or $\overline{C_3}$). Line 14 determines the actual representation of each subgroup and then saves the associated window and counters into W_{syn} , C_{syn} , $\overline{C_{syn}}$.

Then we evaluate whether the instances in $W_{synthetic}$ satisfy or surpass the hyperparameter $minSize$ set by the user (line 15). $minSize$ denotes the minimum number of data samples already seen before rebalancing can occur. For example, if $minSize = 5$, that means we must have seen at least 5 data samples before rebalancing can occur. If both criteria are met, the rebalancing stage begins. If not, the algorithm waits for another data sample to arrive. The imbalance ratio between the number of minority and majority instances is calculated according to Equation 1, and Equation 2 is used to compute the statistical parity difference. The results are stored in $balR$ and $fairR$, respectively (line 16).

Now we present the rebalancing part of the algorithm. If $balR$ is less than p_1 it means that W is unbalanced, and if $fairR$ is less than the hyperparameter f_1 , then W is unfair. Therefore, new synthetic samples \hat{d}_i are generated using the

function *FairGenerate* until W is fair and balanced (line 18). In other words, the condition to stop synthesizing samples is that $balR$ equals p_1 and $fairR$ equals f_1 . As described in the previous section, our method of generating samples will be fairer than the SMOTE.

Before generating new data points, FS^2 will calculate the number of instances introduced for each batch subgroup. Then it will cluster the instances in sliding window W , identify similar individuals through the feature space, filter the boundary samples, and calculate the synthetic weights for the different clusters. FS^2 will randomly select samples in each cluster to synthesize new instances d_i and update the counter S_{gen} until the synthesis number of that cluster is reached. In addition, to avoid overfitting, we will keep the one-time principle, meaning that a sample will be used only once to synthesize instances during the rebalancing phase. The algorithm violates the one-time principle only in one case: when the number of synthesized instances needed to rebalance W is larger than the number of samples. Only in this case will the function reuse the same samples multiple times during the same rebalancing phase. Finally, we will use the new synthesized instances d_i to train the learner l (line 19) and calculate the new $balR$ and $fairR$ (line 21).

Algorithm 1: FS^2 Leaning Algorithm

Input: a discriminated data stream D , the pipeline learner l , optional sampling ratio p_1 and f_1

```

1  $W, W_{label} \leftarrow \emptyset;$ 
2  $C_0, C_1, C_2, C_3 \leftarrow 0;$ 
3  $\overline{C_0}, \overline{C_1}, \overline{C_2}, \overline{C_3} \leftarrow 0;$ 
4  $C_{gen} \leftarrow \emptyset;$ 
5  $adwin \leftarrow \emptyset;$ 
6  $balR, fairR \leftarrow 0;$ 
7 for each instance  $d_i$  in  $D$  do
8    $Train(d_i, l);$ 
9    $W \leftarrow Add(d_i);$ 
10   $UpdateWindows(d_i, W_{label});$ 
11   $C_0, C_1, C_2, C_3 \leftarrow updateCounters(d_i);$ 
12   $adwin \leftarrow Add(d_i);$ 
13   $CheckDrift(adwin, W, W_{label}, C_1, C_2, C_3, C_4, \overline{C_0}, \overline{C_1}, \overline{C_2}, \overline{C_3}, C_{gen});$ 
14   $W_{syn}, C_{syn}, \overline{C_{syn}} \leftarrow CheckRep(C_1, C_2, C_3, C_4, \overline{C_0}, \overline{C_1}, \overline{C_2}, \overline{C_3});$ 
15  if  $CheckSize(minSize, C_{syn})$  then
16     $balR, fairR \leftarrow Calculate(C_1, C_2, C_3, C_4, \overline{C_0}, \overline{C_1}, \overline{C_2}, \overline{C_3})$  according to Equation (1) and Equation (2);
17    while  $f_1 > fairR$  or  $p_1 > balR$  do
18       $\hat{d}_i \leftarrow FairGenerate(W_{syn}, S_{gen});$ 
19       $Train(\hat{d}_i, l);$ 
20       $\overline{C_{syn}} += 1;$ 
21       $balR, fairR \leftarrow Calculate(C_1, C_2, C_3, C_4, \overline{C_0}, \overline{C_1}, \overline{C_2}, \overline{C_3})$  according to Equation (1) and Equation (2);
22    end
23  end
24 end

```

5.2.4 FS^2 Classification. FS^2 is a meta-strategy that solves class imbalance and cumulative discriminatory results in data streams. In particular, we have integrated the Adaptive Hoeffding Trees (AHT) [4] classifier. AHT is a stream-based decision tree induction algorithm that ensures the tree's adaptation to changes in the underlying data distribution by

updating the tree with new instances from the stream and replacing underperforming sub-trees. Therefore, for the online setting, AHT is superior to the majority of other classifiers. FS^2 can be pipelined with any data stream classifier.

6 EXPERIMENTAL EVALUATION

6.1 Datasets

Online fairness research is hindered by the lack of access to good datasets [31]. The large amount of data required and the necessity for concept drift to be present limit availability. We use three real-world datasets and one synthetic

dataset (see Table 1 for details). i) Bank Telemarketing Data (BNK) [29]: The objective of the classification task here is to determine whether a client will subscribe to a term deposit. ii) Give Me Some Credit Dataset (GMSC) [21]: The objective here is to determine whether to approve a loan application. iii) NYPD [33]: We use this dataset to predict whether a suspect was convicted of a felony or not. iv) Synthetic dataset: we followed the initialization process by the authors [22], which involves generating each attribute as a different Gaussian distribution. To simulate class imbalance and concept drift, we introduced these elements into the data stream by shifting the mean of average of each Gaussian distribution.

Table 1. Summary of datasets used in experiments.

Dataset \ CHAR	Sample#	Features#	Sensitive Attribute
BNK	52944	150	Age
GMSC	150,000	29	Age
NYPD	311,367	16	Gender
SYN	1,000,000	21	Synth

6.2 Baselines and Metrics

To evaluate the performance of our method, we compare FS^2 against four state-of-the-art online fairness streaming methods based on four metrics. The four baselines are: i) Fairness-aware Hoeffding Tree [51] (FAHT): FAHT is a fairness-focused variation of the Hoeffding tree algorithm that considers both performance and fairness when selecting split attributes. ii) Online Smooth Boosting [39] (OSBoost): OSBoost is made cost-sensitive by adding various parameters of the Poisson distribution for different classes. iii) Online fairness and class imbalance-aware boosting [23] (FABBOO): FABOO considers long-term class imbalance to fight class imbalance and discrimination from a boosting approach. iv) Continuous SMOTE [3] (C-SMOTE): C-SMOTE is an approach that addresses class imbalance by resampling the minority class using a sliding window, but does not consider fairness.

In terms of the four metrics, two performance metrics and two fairness metrics are employed. i) Balanced Accuracy: It is the arithmetic mean of Specificity (True Negative Rate) and Sensitivity (True Positive Rate). ii) Recall: measures the proportion of actual positive cases that the model correctly identified. iii) Cumulative Statistical Parity Difference (CumSPD): It is a measure of the cumulative statistical parity difference between the protected and deprived groups (Equation 2). This measure must be equal to 0 to be fair. iv) Cumulative Equal Opportunity Difference (CumEOD): It is a measure of cumulative equal opportunity difference between the privileged and unprivileged groups (Equation 3).

6.3 Experimental Results

6.3.1 FS^2 compared to state-of-the-art methods. We compare the performance-fairness metric of our proposed method (FS^2) against the four baselines. The results are demonstrated in Table 2, the darker hue of red denotes the best model performance followed by a lighter red/ pink shade that signifies the second-best performance. FS^2 performed the best in 12 of the 16 performance-fairness evaluations. It was a close second (with a margin of 1-3%) in the remaining 4 measurements. In the BNK dataset, FS^2 achieved a maximum Balanced Accuracy of 81%, surpassing the powerful C-SMOTE. The cumulative EOD and SPD were the best for FS^2 with a 0.02 and 0.01 score respectively, while its recall score of 0.78 was the second highest. FS^2 outperformed these models on the GMSC dataset in terms of Balanced Accuracy, Cumulative EOD and SPD with respective scores of 0.83, 0.01 and 0.01. It was second in terms of Recall with a score of 0.78. Similarly, for the NYPD dataset FS^2 surpassed all four baseline methods as it got the best score for all the fairness-performance metrics i.e. Balanced Accuracy (0.68), Recall (0.54), Cumulative SPD (0.06) and Cumulative EOD (0.03). For the SYN dataset, FS^2 performed fairly well—it was the second best in terms of Balanced Accuracy and Recall. But it outshined in terms of the fairness metrics as it scored a cumulative SPD (0.04) and OPD (0.06).

To summarize, FS^2 demonstrates superior performance compared to other methods in terms of both accuracy and fairness.




Fig. 4. FS^2 and other methods' effectiveness distribution in benchmark tasks; FS^2 shows the best balance, with 79% of mitigation cases having good or win-win results.

6.3.2 The Trade-Off of Effectiveness between FS^2 and other State-of-the-Art Methods.

With the proposed FBU, we can now evaluate the trade-off between fairness and performance based on one single illustrative metric. Figure 4 shows the overall results. As we can see, FS^2 achieves a good or win-win trade-off (i.e., it beats the trade-off baseline constructed by FBU) in most cases, i.e., 79% of the time.

In comparison, the corresponding percentages for FATH, OSBoost, FABBOO, and C-SMOTE were 32%, 25%, 58%, 31%, respectively. In addition, FS^2 has significantly fewer lose-lose trade-off cases (just 1%) than other existing methods. For example, FATH has a lose-lose trade-off rate of 15%, 15 times higher than that of FS^2 . Overall, FS^2 achieves the best trade-off, and it is a significant improvement over all existing methods. The obtained results in Figure 4 are also consistent with the results in Table 2, but are presented in a unified manner that is easy to interpret, verifying the theoretical design of FBU.

6.3.3 Influence of different decay factors. In this section, we analyze the effect of varying decay factor settings on model fairness and performance in the presence of diverse data flow variations. i) Fixed Class-imbalance: In this situation the class ratio is fixed over the data stream. We noted that low decay factor values negatively impact model performance, with Recall dropping below 60% for decay factor < 0.4 . However, for decay factor ≥ 0.5 , Balanced Accuracy and Recall were not significantly impacted. Furthermore, the model's capacity to alleviate unfair outcomes was unchanged for all decay factor values. ii) Increasing and decreasing class imbalance: we observed a steady rise in Recall as the decay factor increased. iii) Fluctuating Class-imbalance: low values of decay factor decrease the performance as in previous cases, while high values also hurt the minority class. In this case, the positive class alternates between minority and

Table 2. Overall predictive and fairness performance for the (darker cells show top rank and the lighter cells show the second rank).

Dataset	Sensitive Attribute	Methods	Balanced Accuracy	Recall	Cumulative SPD	Cumulative EOD
BNK	Age	FAHT	0.73	0.52	0.16	0.18
		OSBoost	0.73	0.71	0.18	0.19
		FABBOO	0.76	0.54	0.05	0.08
		C-SMOTE	0.80	0.81	0.34	0.14
		FS^2	0.81	0.78	0.02	0.01
GMSC	Age	FAHT	0.62	0.28	0.02	0.04
		OSBoost	0.65	0.57	0.04	0.06
		FABBOO	0.81	0.76	0.02	0.01
		C-SMOTE	0.80	0.80	0.07	0.02
		FS^2	0.83	0.79	0.01	0.01
NYPD	Gender	FAHT	0.55	0.28	0.17	0.28
		OSBoost	0.52	0.07	0.25	0.15
		FABBOO	0.62	0.50	0.07	0.06
		C-SMOTE	0.56	0.42	0.34	0.16
		FS^2	0.68	0.54	0.06	0.03
SYN	synth.	FAHT	0.62	0.28	0.08	0.16
		OSBoost	0.57	0.31	0.08	0.11
		FABBOO	0.67	0.58	0.09	0.07
		C-SMOTE	0.67	0.62	0.26	0.16
		FS^2	0.65	0.61	0.04	0.05

majority. High decay factor values allow the model to consider a greater amount of history, leading to class imbalance weights assigned by FS^2 not being adjusted to recent changes. Overall, the value of the decay factor directly affects the performance of FS^2 . Low values of the decay factor produce fluctuating class imbalance weights that decrease the model's performance, and in some cases, values close to 1 are also not appropriate.




Fig. 5. Impact of decay factor on the synthetic dataset of varying class imbalance

6.3.4 Is FS^2 capable of handling concept drift? We tracked the variation of the model performance on the simulated dataset using FBU, and the results are shown in Figure 6. Our approach (the orange curve, towards the top of the graph) is designed for streaming data, and we do observe that the model has the capability to adapt and regain its performance after a concept drift takes place (the curve goes down and then goes up). And since we can handle fairness drifts, we observe that our model can adapt to both fairness drift and conceptual drift. For example, C-SMOTE has a significant

decrease in the overall performance of the model because of the lack of detection of fairness drift. At the same time, the intuition and validity of FBU are well represented. Before FBU, we would need to show the changes in performance and fairness separately, and it would be difficult to analyze the specific impact tradeoff between them. FBU makes this analysis far simpler.

6.3.5 Effect of varying Window size. In our study, we evaluated the impact of various window sizes on our method, specifically using window sizes of 500, 1000, 2000, 5000, and 10000. For predictive performance, we present results on Balanced Accuracy, and for fairness we report CumSPD. As Figure 7 shows, we observe that when the window size surpasses 2000, Balanced Accuracy and CumSPD remain constant. This is due to the occurrence of concept drift, the large sliding window (window size > 2000) is not able to discard outdated instances from before the concept drift, while the smaller window (size = 500 or 1000) is able to accommodate the newer instances after concept drift and discards the older ones. Generally, smaller window settings yield more fair results, but at the cost of decreased prediction performance. Conversely, excessively large window settings can impede the model’s ability to adjust decision boundaries during concept drift, thus negatively affecting model fairness.

7 CONCLUSION

Ensuring models make fair predictions in the presence of a continuous infinite stream of data is an open challenge in real-world streaming machine learning applications. In this work, we proposed Fair Bonded Utility (FBU), the first fairness metric to unify the comparison of the fairness-performance trade-offs of multiple fairness techniques into one intuitive evaluation that simplifies the task of choosing a fairness technique. Then we presented Fair Sampling over Stream (FC²), a novel fair class-balancing technique capable of handling continuous data streams with concept drift. Extensive experiments show our approach gives superior fairness results according to multiple widely-used fairness metrics without sacrificing performance. The pre-processing nature of our method can help improve performance and fairness in many varied applications due to ease of integration.

REFERENCES

- [1] Charu C Aggarwal. 2007. *Data streams: models and algorithms*. Vol. 31. Springer Science & Business Media.
- [2] Yahav Bechavod, Christopher Jung, and Steven Z Wu. 2020. Metric-free individual fairness in online learning. *Advances in neural information processing systems* 33 (2020), 11214–11225.
- [3] Alessio Bernardo, Heitor Murilo Gomes, Jacob Montiel, Bernhard Pfahringer, Albert Bifet, and Emanuele Della Valle. 2020. C-smote: Continuous synthetic minority oversampling for evolving data streams. In *2020 IEEE International Conference on Big Data (Big Data)*. IEEE, 483–492.
- [4] Albert Bifet and Ricard Gavalda. 2009. Adaptive learning from evolving data streams. In *Advances in Intelligent Data Analysis VIII: 8th International Symposium on Intelligent Data Analysis, IDA 2009, Lyon, France, August 31-September 2, 2009. Proceedings 8*. Springer, 249–260.




Fig. 6. Proportion of mitigation cases that beat the trade-off baseline of FS² and existing methods in synthesis dataset with concept drift.




Fig. 7. Effect of window size on model performance and fairness.

- [5] Albert Bifet and Ricard Gavaldà. 2007. Learning from Time-Changing Data with Adaptive Windowing. *Proceedings of the 7th SIAM International Conference on Data Mining* 7. <https://doi.org/10.1137/1.9781611972771.42>
- [6] Toon Calders, Faisal Kamiran, and Mykola Pechenizkiy. 2009. Building classifiers with independency constraints. In *2009 IEEE international conference on data mining workshops*. IEEE, 13–18.
- [7] Toon Calders and Sicco Verwer. 2010. Three naive bayes approaches for discrimination-free classification. *Data mining and knowledge discovery* 21 (2010), 277–292.
- [8] Joydallya Chakraborty, Suvodeep Majumder, and Tim Menzies. 2021. Bias in machine learning software: Why? how? what to do?. In *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 429–440.
- [9] Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. 2002. SMOTE: synthetic minority over-sampling technique. *Journal of artificial intelligence research* 16 (2002), 321–357.
- [10] Shang-Tse Chen, Hsuan-Tien Lin, and Chi-Jen Lu. 2012. An online boosting algorithm with theoretical justifications. *arXiv preprint arXiv:1206.6422* (2012).
- [11] Kate Crawford. 2016. Artificial intelligence’s white guy problem. *The New York Times* 25, 06 (2016), 5.
- [12] Amit Datta, Michael Carl Tschantz, and Anupam Datta. 2014. Automated experiments on ad privacy settings: A tale of opacity, choice, and discrimination. *arXiv preprint arXiv:1408.6491* (2014).
- [13] Georgios Douzas, Fernando Bacao, and Felix Last. 2018. Improving imbalanced learning through a heuristic oversampling method based on k-means and SMOTE. *Information Sciences* 465 (2018), 1–20.
- [14] Joao Gama. 2010. *Knowledge discovery from data streams*. CRC Press.
- [15] João Gama, Indrė Žliobaitė, Albert Bifet, Mykola Pechenizkiy, and Abdelhamid Bouchachia. 2014. A survey on concept drift adaptation. *ACM computing surveys (CSUR)* 46, 4 (2014), 44.
- [16] Adel Ghazikhani, Reza Monsefi, and Hadi Sadoghi Yazdi. 2014. Online neural network model for non-stationary and imbalanced data stream classification. *International Journal of Machine Learning and Cybernetics* 5 (2014), 51–62.
- [17] Adel Ghazikhani, Reza Monsefi, and Hadi Sadoghi Yazdi. 2013. Ensemble of online neural networks for non-stationary and imbalanced data streams. *Neurocomputing* 122 (2013), 535–544.
- [18] Heitor M Gomes, Albert Bifet, Jesse Read, Jean Paul Barddal, Fabricio Enembreck, Bernhard Pfahringer, Geoff Holmes, and Talel Abdesslem. 2017. Adaptive random forests for evolving data stream classification. *Machine Learning* 106 (2017), 1469–1495.
- [19] Sara Hajian, Francesco Bonchi, and Carlos Castillo. 2016. Algorithmic bias: From discrimination discovery to fairness-aware data mining. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. 2125–2126.
- [20] Moritz Hardt, Eric Price, and Nati Srebro. 2016. Equality of opportunity in supervised learning. *Advances in neural information processing systems* 29 (2016).
- [21] <https://www.kaggle.com/c/GiveMeSomeCredit/overview>. 2021. Give Me Some Credit. <https://doi.org/10.34740/KAGGLE/DSV/2242482>
- [22] Vasileios Iosifidis, Thi Ngoc Han Tran, and Eirini Ntoutsi. 2019. Fairness-enhancing interventions in stream classification. In *Database and Expert Systems Applications: 30th International Conference, DEXA 2019, Linz, Austria, August 26–29, 2019, Proceedings, Part I* 30. Springer, 261–276.
- [23] Vasileios Iosifidis, Wenbin Zhang, and Eirini Ntoutsi. 2021. Online fairness-aware learning with imbalanced data streams. *arXiv preprint arXiv:2108.06231* (2021).
- [24] Faisal Kamiran and Toon Calders. 2009. Classifying without discriminating. In *2009 2nd international conference on computer, control and communication*. IEEE, 1–6.
- [25] Faisal Kamiran and Toon Calders. 2012. Data preprocessing techniques for classification without discrimination. *Knowledge and information systems* 33, 1 (2012), 1–33.
- [26] Faisal Kamiran, Toon Calders, and Mykola Pechenizkiy. 2010. Discrimination aware decision tree learning. In *2010 IEEE international conference on data mining*. IEEE, 869–874.
- [27] Bartosz Krawczyk, Leandro L Minku, João Gama, Jerzy Stefanowski, and Michał Woźniak. 2017. Ensemble learning for data stream analysis: A survey. *Information Fusion* 37 (2017), 132–156.
- [28] Nicol Turner Lee, Paul Resnick, and Genie Barton. 2019. Algorithmic bias detection and mitigation: Best practices and policies to reduce consumer harms. *Brookings Institute: Washington, DC, USA* 2 (2019).
- [29] Sérgio Moro, Paulo Cortez, and Paulo Rita. 2014. A data-driven approach to predict the success of bank telemarketing. *Decision Support Systems* 62 (2014), 22–31. <https://doi.org/10.1016/j.dss.2014.03.001>
- [30] Marcelo OR Prates, Pedro H Avelar, and Luís C Lamb. 2020. Assessing gender bias in machine translation: a case study with google translate. *Neural Computing and Applications* 32 (2020), 6363–6381.
- [31] Tai Le Quy, Arjun Roy, Vasileios Iosifidis, Wenbin Zhang, and Eirini Ntoutsi. 2022. A survey on datasets for fairness-aware machine learning. *Data Mining and Knowledge Discovery* (2022).
- [32] Jesse Read, Nikolaos Tziortziotis, and Michalis Vazirgiannis. 2019. Error-space representations for multi-dimensional data streams with temporal dependence. *Pattern Analysis and Applications* 22, 3 (2019), 1211–1220.
- [33] John F Sargent and Dana A Shea. 2017. *Office of Science and Technology Policy (OSTP): History and Overview*. Congressional Research Service.
- [34] Gaganpreet Sharma. 2017. Pros and cons of different sampling techniques. *International journal of applied research* 3, 7 (2017), 749–752.

- [35] Zeyu Sun, Jie M Zhang, Mark Harman, Mike Papadakis, and Lu Zhang. 2020. Automatic testing and improvement of machine translation. In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering*. 974–985.
- [36] Latanya Sweeney. 2013. Discrimination in online ad delivery. *Commun. ACM* 56, 5 (2013), 44–54.
- [37] Xuejiao Tang, Lihua Zhang, et al. 2020. Using machine learning to automate mammogram images analysis. In *IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*. 757–764.
- [38] Sahil Verma and Julia Rubin. 2018. Fairness definitions explained. In *Proceedings of the international workshop on software fairness*. 1–7.
- [39] Boyu Wang and Joelle Pineau. 2016. Online bagging and boosting for imbalanced data streams. *IEEE Transactions on Knowledge and Data Engineering* 28, 12 (2016), 3353–3366.
- [40] Qianwen Wang, Zhenhua Xu, Zhutian Chen, Yong Wang, Shixia Liu, and Huamin Qu. 2020. Visual analysis of discrimination in machine learning. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2020), 1470–1480.
- [41] Shuo Wang, Leandro L Minku, and Xin Yao. 2013. A learning framework for online class imbalance learning. In *2013 IEEE Symposium on Computational Intelligence and Ensemble Learning (CIEL)*. IEEE, 36–45.
- [42] Shuo Wang, Leandro L Minku, and Xin Yao. 2014. Resampling-based ensemble methods for online class imbalance learning. *IEEE Transactions on Knowledge and Data Engineering* 27, 5 (2014), 1356–1368.
- [43] Zichong Wang, Yang Zhou, et al. 2023. Towards Fair Machine Learning Software: Understanding and Addressing Model Bias Through Counterfactual Thinking. (2023).
- [44] Gary M Weiss. 2004. Mining with rarity: a unifying framework. *ACM Sigkdd Explorations Newsletter* 6, 1 (2004), 7–19.
- [45] Shen Yan, Hsien-te Kao, and Emilio Ferrara. 2020. Fair class balancing: Enhancing model fairness without observing sensitive attributes. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 1715–1724.
- [46] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez-Rodriguez, and Krishna P Gummadi. 2019. Fairness constraints: A flexible approach for fair classification. *The Journal of Machine Learning Research* 20, 1 (2019), 2737–2778.
- [47] Mingli Zhang, Xin Zhao, et al. 2020. Deep discriminative learning for autism spectrum disorder classification. In *International Conference on Database and Expert Systems Applications*. Springer, 435–443.
- [48] Wenbin Zhang, Albert Bifet, Xiangliang Zhang, Jeremy C Weiss, and Wolfgang Nejdl. 2021. FARF: A Fair and Adaptive Random Forests Classifier. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, 245–256.
- [49] Wenbin Zhang, Tina Hernandez-Boussard, and Jeremy C Weiss. 2023. Censored Fairness through Awareness. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- [50] Wenbin Zhang, Juyong Kim, Zichong Wang, Pradeep Ravikumar, and Jeremy Weiss. 2023. Individual Fairness Guarantee in Learning with Censorship. (2023).
- [51] Wenbin Zhang and Eirini Ntoutsi. 2019. FAHT: an adaptive fairness-aware decision tree classifier. *arXiv preprint arXiv:1907.07237* (2019).
- [52] Wenbin Zhang, Shimei Pan, Shuigeng Zhou, Toby Walsh, and Jeremy C Weiss. 2022. Fairness Amidst Non-IID Graph Data: Current Achievements and Future Directions. *arXiv preprint arXiv:2202.07170* (2022).
- [53] Wenbin Zhang and Jianwu Wang. 2017. A hybrid learning framework for imbalanced stream classification. In *2017 IEEE International Congress on Big Data (BigData Congress)*. IEEE, 480–487.
- [54] Wenbin Zhang and Jianwu Wang. 2018. Content-bootstrapped collaborative filtering for medical article recommendations. In *IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*.
- [55] Wenbin Zhang and Jeremy Weiss. 2021. Fair Decision-making Under Uncertainty. In *2021 IEEE International Conference on Data Mining (ICDM)*. IEEE.
- [56] Wenbin Zhang and Jeremy Weiss. 2022. Fairness with Censorship and Group Constraints. *Knowledge and Information Systems* (2022).
- [57] Wenbin Zhang and Jeremy C Weiss. 2022. Longitudinal fairness with censorship. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 12235–12243.
- [58] Indre Zliobaite. 2015. A survey on measuring indirect discrimination in machine learning. *arXiv preprint arXiv:1511.00148* (2015).