



# Fairness-Enhancing Interventions in Stream Classification

Vasileios Iosifidis<sup>1,2(✉)</sup>, Thi Ngoc Han Tran<sup>1</sup>, and Eirini Ntoutsi<sup>1,2</sup>

<sup>1</sup> Leibniz University, Hannover, Germany  
tranthingochan.03@gmail.com

<sup>2</sup> L3S Research Center, Hannover, Germany  
{iosifidis, ntoutsi}@L3S.de

**Abstract.** The wide spread usage of automated data-driven decision support systems has raised a lot of concerns regarding accountability and fairness of the employed models in the absence of human supervision. Existing fairness-aware approaches tackle fairness as a batch learning problem and aim at learning a fair model which can then be applied to future instances of the problem. In many applications, however, the data comes sequentially and its characteristics might evolve with time. In such a setting, it is counter-intuitive to “fix” a (fair) model over the data stream as changes in the data might incur changes in the underlying model therefore, affecting its fairness. In this work, we propose fairness-enhancing interventions that modify the input data so that the outcome of any stream classifier applied to that data will be fair. Experiments on real and synthetic data show that our approach achieves good predictive performance and low discrimination scores over the course of the stream.

**Keywords:** Data mining · Fairness-aware learning ·  
Stream classification

## 1 Introduction

Despite the wide spread belief that data-driven decision making is objective in contrast to human-based decision making that is subject to biases and prejudices, several cases have been documented, e.g., [9, 13], in which data-driven decision making incurs discrimination. As a recent example, a Bloomberg report has suggested signs of racial discrimination in Amazon’s same-day delivery service [13]. The sensitive attribute race was not employed as a predictive attribute in Amazon’s model(s), however the location of the users might have acted as a proxy for race. As a result, predominantly black ZIP codes were excluded from services and amenities. Therefore, the wide spread usage of automated data-driven decision support systems has raised a lot of concerns regarding accountability and fairness of the employed models in the absence of human supervision [1, 24, 28]. Such issues result in societal and legal implications, therefore, recently the domain of discrimination-aware data mining [23] has attracted a lot of attention and several

methods have been proposed ranging from discrimination discovery to discrimination elimination and explanation of model decisions.


Most of these methods, however, tackle fairness as a batch learning problem aiming at learning a “fair” model which can be then used for predicting future instances of the population. In many modern applications, however, data is generated sequentially and its characteristics might change with time, i.e., the data is non-stationary. Such dynamic environments (or, data streams) call for model adaptation [12]. As an example, in the EU, the non-native population has significantly changed in the last years due to European refugee crisis and internal EU migration with a potential effect on the racial discrimination in the labor market. In such *non-stationary* environments, the main challenge for supervised learning is the so-called concept drifts, i.e., changes in the underlying data distribution which affect the learning model as the relationships between input and class variables might evolve with time [26]. Existing solutions from the data stream mining domain tackle this issue by adapting the learning models online. However, as the decision boundary of the classifier changes as a result of model adaptation, the fairness of the model might get hurt.

An example of an evolving stream with discrimination is shown in Fig. 1; one can see the deprived and favored communities (w.r.t. some sensitive attribute) over time as well as their class assignments. The favored community dominates the stream. The decision boundary of the classifier (solid line) changes in response to changes in the underlying data. As a result, the associated fairness of the model also changes, calling for “fairness-enhancing interventions” (dashed line). It is important, therefore, model adaptation to also consider fairness<sup>4</sup> to ensure that a valid fairness-aware classifier is maintained over the stream. In this work, we propose fairness-enhancing interventions that modify the input data before updating the classifier. Our method belongs to the category of pre-processing approaches to fairness, investigated that far only in the context of static learning [5, 6, 14–17]. Our contributions are: (a) we introduce the fairness-aware classification problem for streams (b) we propose pre-processing fairness-enhancing interventions for streams (c) we propose a synthetic generator for simulating different drift and fairness behaviors in a stream and (d) we present an extensive experimental evaluation with different stream learners and on different datasets.

The rest of the paper is as follows: In Sect. 2, we overview the related work. Our approach is presented in Sect. 3. Experimental results are discussed in Sect. 4. Finally, conclusions and outlook are presented in Sect. 5.

## 2 Related Work

Although more than twenty different notions of fairness have been proposed in the last few years [25, 29], still there is no agreement on which measure to apply in each situation. The most popular is that of statistical parity [29] that checks whether the favored and deprived communities have equal probability of being assigned to the positive class. This is the measure we also adopt in this work.



**Fig. 1.** An evolving stream with discrimination: At each timepoint, the current decision boundary (solid line) and the “fairness-corrected” boundary (dashed line) are displayed.

**Pre-processing Fairness-Enhancing Interventions:** Methods in this category work under the assumption that in order to learn a fair classifier, the training data should be discrimination-free. To this end, they try to balance the representation of the different groups in the population. Massaging [15] modifies the data distribution by re-labeling some of the instances which reside close to the decision boundary in order to neutralize discriminatory effects. Re-weighting [5] assigns different weights to the different group, e.g., the deprived group will receive a higher score comparing to the favored group. These methods are typically model-agnostic and therefore, any classifier is applicable after the pre-processing phase.

**In-processing Fairness-Enhancing Interventions:** Methods in this category directly modify the learning algorithm to ensure that it will produce fair results. As such, they are algorithm-specific; e.g., [18] proposes a decision tree that encodes fairness by employing a modified entropy-based attribute splitting criterion and [8] includes sensitive attributes in the learning process by utilizing a joint loss function that makes explicit trade-off between fairness and accuracy.

**Post-processing Fairness-Enhancing Interventions:** Post-processing methods modify the results of a trained classifier to ensure the chosen fairness criterion is met; e.g., [18] modifies the leaf labels of a decision tree, [22] changes the confidence values of classification rules and [10] shifts the decision boundary of an AdaBoost learner until the fairness criterion is fulfilled.

**Stream Classification:** Data stream algorithms must be able to adapt to concept drifts in order to maintain a good performance over the stream [12]. Model adaptation is typically enabled by: (i) incorporating new instances from the stream into the model and (ii) forgetting or downgrading outdated information from the model. The former calls for online/incremental algorithms, whereas the latter calls for methods that are able to forget e.g., [11, 19]. We discuss several stream classifiers in the experiments (Sect. 4).

**Sequential Fairness:** When a sequence of decisions has to be taken, the notion of sequential fairness is relevant. For example, [27], studies fair online item ranking for groups and [20] how fairness criteria interact with temporal indicators of well-being and affect discriminated populations on the long-term.

Our work lies in the intersection of pre-processing methods for fairness and stream classification methods. The former, however, focus solely on the static case, i.e., they assume that the data is stationary, whereas the latter focus solely on predictive accuracy and ignore fairness. To the best of our knowledge, this is the first work trying to bridge the two domains.

### 3 Fairness-Aware Stream Classification

A data stream  $S$  is a potentially infinite sequence of instances arriving over time, each instance described in a feature space  $A = (A_1, A_2 \cdots A_d)$ . One of the attributes is the *sensitive attribute*, denoted by  $SA$ , with values  $SA = \{s, \bar{s}\}$ ; we refer to  $s$  and  $\bar{s}$  as “deprived” and “favored”, respectively. We also assume a binary class attribute  $C = \{rejected, granted\}$ . We refer to “granted” class value as *target class*. We process the stream in chunks of fixed size,  $S_1, \dots, S_t$  with  $S_t$  being the most recent chunk. We assume the fully supervised learning setting, where the labels of the instances are available shortly after their arrival. Therefore, the goal is to make a prediction for the instances based on the current classifier and use the labels later on for update (the so-called prequential evaluation [12]). The underlying stream population is subject to changes, which might incur concept drifts, i.e., the decision boundary might change overtime (c.f., solid line in Fig. 1) and therefore, fairness implications may take place (c.f., dashed line in Fig. 1). A stream classifier typically takes care of concept drifts, but does not consider fairness.

The *discrimination aware stream classification problem* therefore is to maintain a classifier that performs well (i.e., the predictive accuracy is high) and does not discriminate (i.e., the discrimination score is low, c.f. Eq. 1) over the course of the stream. In this work, we follow the pre-processing approaches to fairness-aware learning that intervene at the input data to ensure a fair representation of the different communities. In particular, we monitor the discrimination in each incoming chunk from the stream (Sect. 3.1) and if the discrimination score exceeds a user defined threshold  $\epsilon$ , we “correct” the chunk for fairness (Sect. 3.2) before feeding it into the learner (Sect. 3.3). We assume an initialization phase at the beginning of the stream for which an initial fairness-aware classifier  $F_0$  is trained upon an initial dataset  $S_0$  from the stream. An overview of our approach is depicted in Fig. 2, where M1–M4 are the adaptation strategies introduced in Sect. 3.3.

#### 3.1 Detecting Classifier Discrimination in Data Streams

Let  $F$  be the current (non-discriminating) stream classifier and  $S_t$  be the current chunk received from the stream. We evaluate the discriminative behavior of  $F$  over  $S_t$ , i.e.,  $disc_S(F, S_t)$  by evaluating the predictions of  $F$  over instances of  $S_t$ . First, we define four communities in each chunk  $S_t$  by combining the sensitive attribute  $SA$  with the (predicted) class attribute  $C$  (both binary attributes):




Fig. 2. Fairness-aware stream classification overview

 Table 1. (Chunk  $S_t$ ) Communities

Sensitive Attribute SA	(predicted) class	
	Rejected	Granted
$s$ (Female)	$DR_t$ ( <i>deprived rejected</i> )	$DG_t$ ( <i>deprived granted</i> )
$\bar{s}$ (Male)	$FR_t$ ( <i>favoured rejected</i> )	$FG_t$ ( <i>favoured granted</i> )

As discrimination measure, we employ statistical parity that evaluates whether the favored and deprived groups have equal probabilities of being granted [25]:

$$disc_S(F, S_t) = \frac{FG_t}{FG_t + FR_t} - \frac{DG_t}{DG_t + DR_t} \quad (1)$$

If the discrimination value exceeds the threshold  $\epsilon$ , i.e.,  $disc_S(F, S_t) > \epsilon$ , the discrimination performance of the model degrades, due to, e.g., changes in the distribution that reside in the newly arrived chunk  $S_t$ . A typical stream classifier would update  $F$  based on  $S_t$  to adapt to the incoming stream. However, to also account for fairness, we first “correct”  $S_t$  for fairness (c.f., Sect. 3.2), before employing its instances for updating the model (c.f., Sect. 3.3).

### 3.2 Fairness-Enhancing Data Interventions in Data Streams

If discrimination is detected,  $S_t$  is “corrected” for fairness before being used for model update (Sect. 3.3). To this end, we employ two different data intervention techniques: massaging and re-weighting.

**Chunk-Based Massaging.** Massaging [15] modifies the data distribution by swapping the class labels of certain instances (from “granted” into “rejected” or vice versa) from each of the deprived rejected (DR) and favoured granted (FG) communities. The amount of affected instances,  $M_t$ , from each community is derived by Eq. 1 and is as follows:

$$M_t = \frac{FG_{S_t} * (DG_{S_t} + DR_{S_t}) - DG_{S_t} * (FG_{S_t} + FR_{S_t})}{|S_t|} \quad (2)$$

The best candidate instances for label swapping are those close to the decision boundary, as intuitively their alternation will have the least impact on the model while it will fulfill the discrimination requirement (Eq. 1). To this end, we employ a ranker  $R_t$  trained on  $S_t$  that estimates the class probabilities of the instances in  $S_t$ . Then,  $M_t$  instances assigned with high probability to  $DR$  and  $M_t$  instances assigned with low probability to  $FG$  are selected for label swapping.

**Chunk-Based Re-weighting.** Re-weighting [5] modifies the data distribution by assigning different weights to each community (c.f., Table 1) to “enforce” a fair allocation of deprived and favored instances w.r.t the target class. Similarly to massaging, the exact weights depend on the  $S_t$  distribution in the different communities. Below we provide the weight for the favoured granted community, same rationale holds for the other communities:

$$W_t^{FG} = \frac{|\bar{s}_{S_t}| * |\{x \in S_t | (x.C = \text{“granted”})\}|}{|S_t| * |FG_{S_t}|} \quad (3)$$

Each instance  $x \in S_t$  is weighted by “inheriting” the weight of its community.

**Massaging vs Re-weighting.** Both methods modify the data distribution to equalize the number of deprived and favored communities in the target class. However, there are fundamental differences between the two approaches: massaging interferes at the instance level by altering single instances, whereas re-weighting affects a whole community by lowering/increasing its weight. Moreover, massaging is more intrusive than re-weighting as it alters the class labels. Both interventions result in a “corrected” chunk  $S'_t$  ( $|S_t| = |S'_t|$ ) used for updating the classifier (c.f., Sect. 3.3).

### 3.3 Fairness-Aware Classifier Adaptation in Data Streams

The update of a classifier should take into account both concept drifts and fairness. For the former, we work with stream classifiers, like, Hoeffding Trees, Accuracy Updated Ensembles and Naive Bayes that already adapt to concept drifts. In that sense, the concept drift problem is directly tackled by the learner. For the latter, we “correct” the input stream per chunk, using either massaging or re-weighting, to ensure that learners are trained on “fair” data (c.f., Sect. 3.2). In particular, we propose update strategies for fairness-aware stream classification:

- **ACCUM&FULLTRAIN** (shortly  $M1$ ):  $F$  is continuously updated over the stream using the original current chunk  $S_t$ , if no discrimination is detected, or its “corrected” counterpart  $S'_t$ , if discrimination is detected.
- **RESET&FULLTRAIN** (shortly  $M2$ ): Similar to  $M1$ , but if discrimination is detected,  $F$  is reset and a new model is created from the “corrected”  $S'_t$ .

The underlying assumption for **ACCUM&FULLTRAIN** is that if trained with “fair” chunks, the classifier  $F$  should be fair. In practice, though and due to the complex interaction between input data and learning algorithms, this might not be true (c.f., [17]); therefore, we also propose the **RESET&FULLTRAIN** model that resets the learner once its predictions incur discrimination.

In addition, we propose two variations that focus more on fairness. The rationale is similar to the previous approaches but the model is updated only if discrimination is detected and only via “corrected” data. Therefore, these two models adapt slower to concept drifts comparing to the first two models, as their adaptation occurs only if discrimination is detected. In particular:

- ACCUM&CORRECTEDTRAIN (shortly *M3*):  $F$  is updated over the stream only if discrimination is detected. Update is based on “corrected” chunks  $S'_t$ .
- RESET&CORRECTEDTRAIN (shortly *M4*): Similar to *M3*, but once discrimination is detected  $F$  is reset and a new model is created from the corrected chunk  $S'_t$ . Thus, *M4* adapts only via reset, when discrimination is detected.

## 4 Experiments

We evaluate the performance of our methods for discrimination elimination in data streams using both real and synthetic datasets (Sect. 4.1). As evaluation measures, we use the performance of the model, in terms of accuracy and discrimination, over the new coming chunk from the stream. We report on the performance of the different methods over the stream but also on the overall accuracy-vs-fairness behavior of the different methods. We experiment with a variety of stream classifiers such as Naive Bayes (NB), Hoeffding Tree (HT), Accuracy Updated Ensemble (AUE) and k-Nearest Neighbors (KNN). The aforementioned models are updated based on the new incoming chunk from the stream, however they differ w.r.t how they handle historical information. NB and HT classifiers do not forget, whereas AUE forgets by replacing old learners with new ones. kNNs on the other hand, rely solely on the last chunk for the predictions, due to its internal buffer. An overview of each classifier is given below:

- Naive Bayes (NB): A probabilistic classifier that makes a simplistic assumption on the class-conditional independence of the attributes. The stream version of NBs [2] is an online algorithm, i.e., the model is updated based on new instances from the stream, but does not forget historical information.
- Hoeffding Tree (HT): A decision tree classifier for streams that uses the Hoeffding bound to make a reliable decision on the best splitting attribute from a small data sample [7]. HT is an online algorithm (so, it is updated based on new instances from the stream) but does not forget.
- Accuracy Updated Ensemble (AUE): An ensemble model that adapts to concept drifts by updating its base-learners based on the current data distribution, tuning their weights according to their predictive power on the most recent chunk [4]. The model replaces old learners with newer ones trained upon more recent chunks. We used HTs as base learners for the ensemble and we set the maximum number of base learners to 10.
- KNN: A lazy learner which predicts based on the class labels of the neighboring instances [3]. In particular, the previous chunk instances and their labels are used to make predictions for the instances of the current chunk. The neighborhood is set to  $k = 10$ .

We evaluate our strategies M1–M4 (c.f., Sect. 3.3) for the different classifiers, as well as against the following baselines that do not explicitly handle discrimination:

- B1 **B.NoSA** (Baseline NoSensitiveAttribute): The classifier  $F$  does not employ  $SA$  neither in training nor in testing. The model is continuously updated over the stream from the original chunks  $S_t$ . Intuitively, the model tackles discrimination by omitting  $SA$ .
- B2 **B.RESET** (Baseline Reset): If discrimination is detected, the old model  $F$  is deleted and a new model is learned on  $S_t$ . The model is updated over the stream, but without any correction. Discrimination is being monitored and if it is detected again, the whole procedure starts over. Intuitively, this approach tackles discrimination by resetting the model when discrimination is detected.

For the massaging techniques, we use NB as a ranker which according to [5] is the best ranker. We implemented our methods<sup>1</sup> in MOA [3]. For all of our reported experiments, we consider a discrimination threshold of  $\epsilon = 0.0$ , that is, we do not tolerate any discrimination, and a chunk size of  $|S| = 1,000$  instances. The effect of these parameters is discussed in Sect. 4.2.

#### 4.1 Datasets

As real dataset we employ the census-income (or adult-census) dataset, which comprises one of the most popular datasets in this area; we simulate the stream using the file order. Due to lack of stream data for fairness, we extend an existing stream generator to simulate different discrimination scenarios in data streams.

**Census-Income** [21]: The learning task is to predict whether a person earns more than 50K/year using demographic features. We consider gender as the sensitive attribute with females being the deprived community and males being the favored community. In addition, we consider an annual income of more than 50K as the target class. The dataset consists of 48,842 records and has an overall discrimination of 19.45%. In Fig. 3a, the discrimination score and the different community volumes ( $DR_t, DG_t, FR_t, FG_t$ ) are shown over time using a chunk size of  $|S| = 1,000$  instances. The discrimination score ranges between 15%–25% overtime.

**Synthetic Generator:** Our generator comprises an extension of the static fairness generator of [30] that represents each community using a Gaussian distribution. It forces the DG community to be closer to the negative class whereas the FG community is placed further away from the negative class. An example can be already seen in Fig. 1. We extend this idea in a stream setting by varying the amount of discrimination over the stream while introducing concept drifts.

In particular, we initialize four Gaussians, as follows, similarly to the static generator:  $p(DG) = N([2; 2], [3, 1; 1, 3])$ ,  $p(FG) = N([2.5; 2.5], [3, 1; 1, 3])$ ,

<sup>1</sup> Code will be made available online.



$p(DR) = N([0.5; 0.5], [3, 3; 1, 3])$  and  $p(FR) = N([-2; -2], [3, 1; 1, 3])$ . In the initialization phase, all Gaussians contribute equally to each community with  $n$  instances each, giving a total of  $N = 4n$  instances for the initial chunk. With respect to discrimination, we introduce a parameter  $SPP$  that controls the statistical parity by controlling the number of generated instances  $x$  in the DG community over the stream. The exact amount of instances  $x$  can be derived from Eq. 1 as follows:

$$SPP = \frac{n}{2 * n} - \frac{x}{x + n} \Rightarrow x = n * \frac{1 - 2 * SPP}{1 + 2 * SPP} \quad (4)$$

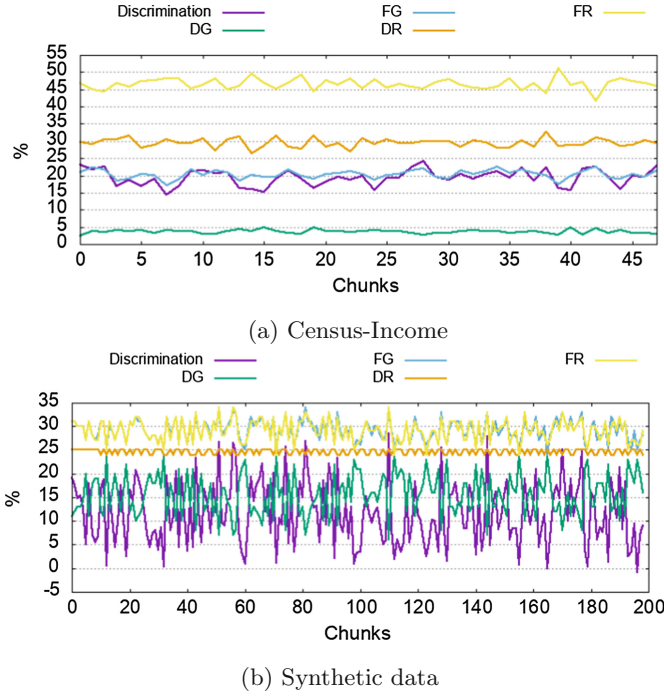
where  $n$  is the amount of instances generated by each Gaussian in a chunk and  $x$  is the amount of instances for the DG community based on the desired  $SPP$  value; the rest  $n - x$  instances generated originally by its corresponding Gaussian are evenly distributed to the FG and FR communities. This way, the ratio of positive instances in the favored community remains the same. To simulate concept drifts in the population, we change the means of the Gaussians at random points over the stream. To maintain the initial problem (unfair treatment of one community), we shift the means all together at a random direction *up*, *down*, *left* or *right* by a random value  $k \in [0, 2]$ .

For evaluation purposes, we generate a synthetic dataset of 200,000 instances (200 chunks,  $N = 1,000$  instances per chunk), 4 numerical attributes, 1 binary sensitive attribute and 1 binary class. We inserted 20 concept drift at random points and vary  $SPP$  randomly over time from 0% to 30%. The dataset characteristics are shown in Fig. 3b.

## 4.2 Evaluation Results

For each dataset, we report on the discrimination-vs-accuracy behavior of the different classifiers under the different adaptation strategies. The discrimination-vs-accuracy plot (an example is shown in Fig. 4a) allows for a quick evaluation of the different behaviors. Values close to 0 in the  $x$ -axis mean fair models, whereas as the values increase the corresponding classifier become more discriminating. w.r.t accuracy ( $y$ -axis), good predictive power models reside close to 100%, whereas low  $y$  values indicate poor performing models. The ideal models are located on the up left region which indicates high accuracy and low discrimination performance models. The worst models are located in the bottom right region where low accuracy and high discriminating behavior take place. Up right and bottom left regions indicate unfair but accurate models and fair but inaccurate models, respectively.


**Census-Income.** For the massaging, c.f., Fig. 4a, our strategies achieve lower discrimination comparing to the baselines (our values are closer to 0 in the  $x$ -axis). As expected, the improvement w.r.t discrimination incurs a drop in accuracy, i.e., baselines have better accuracy comparing to our strategies (baseline values are closer to 100% in the  $y$ -axis). We also observe that some strategies depict very similar performance, e.g.,  $M2$  and  $M4$  when combined with  $HT$ .




**Fig. 3.** Discrimination and community size over the stream ( $|S| = 1,000$  instances/chunk)

The reason is that since  $\epsilon = 0$ , our discrimination detector is activated on almost every chunk from the stream and therefore strategies like  $M2$  and  $M4$  will both reset the model on each chunk. Accumulative strategies,  $M1$  and  $M3$ , perform better than reset strategies,  $M2$  and  $M4$ ; the reason is probably that the latter ones forget too fast. Regarding the different classifiers employed by our strategies, we can see that the best performing ones in terms of both accuracy and discrimination are  $KNN$  and  $AUE$ .  $AUE$  and  $HT$  models yield better accuracy and less discrimination when they do not discard previous knowledge. Although  $KNN$  is not the best performing model in terms of accuracy, it yields the lowest discrimination score with the smallest drop in accuracy when compared to its baselines, namely  $B.RESET$  and  $B.NoSA$ . In particular, discrimination drops from 19% to 4% while accuracy drops by almost 1%, when  $KNN$  is employed by  $M3$  and  $M4$  strategies.

In Fig. 4b, we compare the discrimination-vs-accuracy behavior of the different classifiers under re-weighting. Same as in massaging, our strategies reduce discrimination in predictions. Classifiers such as  $HT$  behave similarly under different strategies since the detector detects discrimination in almost every chunk.  $KNN$  on the other hand, doesn't take into consideration weights, hence all the strategies perform identically.



(a) Massaging




(b) Re-weighting

**Fig. 4.** Census-income (small): Discrimination-vs-accuracy of the different strategies

We also compare models overtime in Fig. 5. We have selected one model for each method (massaging/re-weighting) based on the discrimination-accuracy trade off (points which are closer to (0, 1) based on Euclidean distance) and the best baseline of those two models. Although HT’s baseline has the best accuracy overtime, its discrimination score is close to stream’s discrimination. On the other hand, re-weighting and massaging methods result in a significant drop in discrimination. The number of massaged instances  $M$  varies over the stream, based on how discriminative the KNN’s predictions are.

**Synthetic Data.** The dataset contains 20 concept drifts while its discrimination score varies overtime, as seen in Fig. 3b. The majority of baselines, in Figs. 6a and b, are able to adapt to concept drifts (i.e., they achieve high accuracy), however they cannot handle discrimination, which in some cases is even amplified comparing to the original stream overall discrimination. The vast majority of baselines

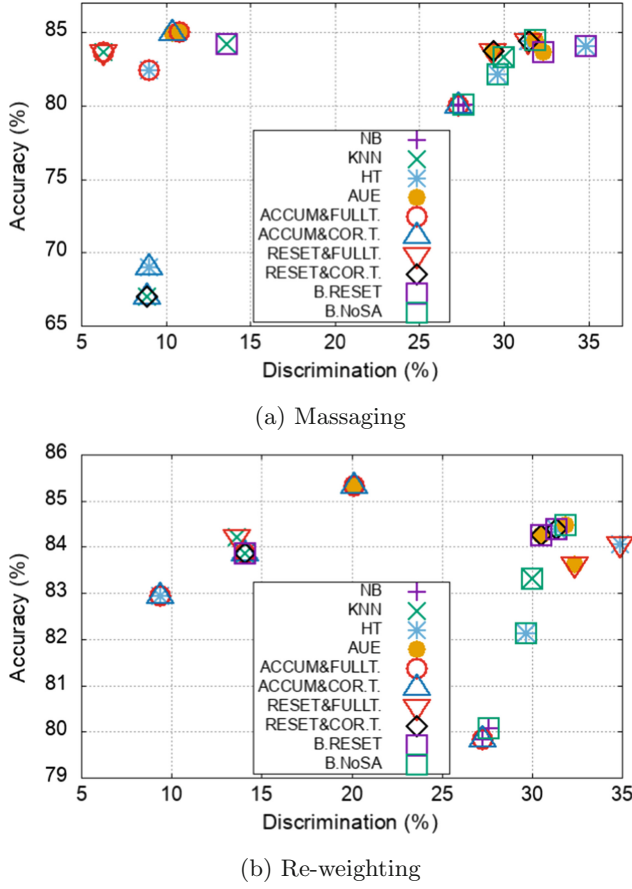


**Fig. 5.** Census-income: Accuracy (top), discrimination (middle) and # massaged instances (bottom) over the stream

occupy the up right region which means that models are able to adapt to concept drifts even though they are highly discriminating. By inspecting Fig. 6a, we can observe once again that accumulative models are less discriminating in contrast to reset models. KNN achieves high reduction in discrimination (up to 6%), while maintaining high accuracy. Classifiers such as AUE and HT perform well when combined with accumulative strategies while reset strategies incur higher discrimination. A possible reason for this behavior is that when trained on more data, a model can generalize better, especially in re-occurring concepts, comparing to reset strategies that rely solely on recent chunks. KNN is an exception as it performs well despite relying on an internal sliding window for its predictions. A possible reason is that a kNN learner is an instance-based learned and does not perform explicit generalization like HT and AUE. Similarly to census-income dataset, NB is not able to tackle discrimination.


In Fig. 6b, we observe that almost all baselines, same as in massaging, cover the up right region area. AUE’s performance is increasing while it becomes more discriminating in contrast to AUE in massaging. Same as before, HT and KNN have the least discriminating behavior while NB performs poorly. Again, reset strategies produce good accuracy models but fail to reduce discrimination.

In Fig. 7, we compare the best “correction” methods and the best baseline. M1 combined with KNN has the lowest discrimination score overtime. Its accuracy is slightly worse than its baseline. However, discrimination is lower than stream’s and baseline’s discrimination overtime. HT’s overall performance w.r.t accuracy is relatively good except the interval between 77th and 90th chunk where four concept drifts occurred incurring accuracy loss. Despite the accuracy degradation, HT achieved lower discrimination compared to other classifiers.



**Fig. 6.** Synthetic stream: Discrimination-vs-accuracy of the different strategies

**Parameter effect:** Due to lack of space we omit the time execution charts. A derived conclusion is that our strategies are executed slightly slower compared to the baselines and moreover, that the reset strategies are faster than the accumulative strategies. We have also experimented with different chunk sizes  $|S|$  and discrimination thresholds  $\epsilon$ . Based on our experiments, increasing  $\epsilon$  results in better accuracy models but their discrimination also increases. With respect to the chunk size effect, there was no clear effect on the performance except for the execution time that decreases with chunk size as less operations take place.



**Fig. 7.** Synthetic stream: Accuracy (top), discrimination (middle) and # massaged instances (bottom) over the stream

## 5 Conclusions and Future Work

In this work, we proposed an approach for fairness-aware stream classification, which is able to maintain good predictive performance models with low discrimination scores overtime. Our approach tackles discrimination by “correcting” the input stream w.r.t fairness and therefore, can be coupled with any stream classifier. Our experiments show that such a correction over the stream can reduce discrimination in model predictions, while the maintenance of the model over the stream allows for adaptation to underlying concept drifts. Comparing the different fairness-intervention methods, our experiments show that massaging performs better than re-weighting. A possible explanation is that massaging works at the individual instance level by swapping its class label, whereas re-weighting works at a group level by applying different weights to different communities. Moreover, massaging affects selected instances, which are closer to the boundary.

Our approach is model-agnostic, however our experiments show that the effect of “data correction for discrimination” depends on the classifier and therefore, how to “best correct” for specific classifiers is an interesting research direction. Moreover, we want to investigate in-processing fairness-aware stream classifiers that incorporate fairness notion directly in the classification algorithm.

**Acknowledgements.** The work is inspired by the German Research Foundation (DFG) project OSCAR (Opinion Stream Classification with Ensembles and Active learners) for which the last author is Co-Principal Investigator.

## References

1. Bhandari, E.: Big data can be used to violate civil rights laws, and the FTC agrees (2016)
2. Bifet, A., Frank, E.: Sentiment knowledge discovery in twitter streaming data. In: Pfahringer, B., Holmes, G., Hoffmann, A. (eds.) DS 2010. LNCS (LNAI), vol. 6332, pp. 1–15. Springer, Heidelberg (2010). [https://doi.org/10.1007/978-3-642-16184-1\\_1](https://doi.org/10.1007/978-3-642-16184-1_1)
3. Bifet, A., Holmes, G., Kirkby, R., Pfahringer, B.: MOA: massive online analysis. *JMLR* **11**, 1601–1604 (2010)
4. Brzeziński, D., Stefanowski, J.: Accuracy updated ensemble for data streams with concept drift. In: Corchado, E., Kurzyński, M., Woźniak, M. (eds.) HAIS 2011. LNCS (LNAI), vol. 6679, pp. 155–163. Springer, Heidelberg (2011). [https://doi.org/10.1007/978-3-642-21222-2\\_19](https://doi.org/10.1007/978-3-642-21222-2_19)
5. Calders, T., Kamiran, F., Pechenizkiy, M.: Building classifiers with independency constraints. In: ICDMW 2009, pp. 13–18. IEEE (2009)
6. Calmon, F., Wei, D., Vinzamuri, B., Ramamurthy, K.N., Varshney, K.R.: Optimized pre-processing for discrimination prevention. In: NIPS, pp. 3992–4001 (2017)
7. Domingos, P., Hulten, G.: Mining high-speed data streams. In: SIGKDD, pp. 71–80. ACM (2000)
8. Dwork, C., Immorlica, N., Kalai, A.T., Leiserson, M.D.: Decoupled classifiers for group-fair and efficient machine learning. In: FAT, pp. 119–133 (2018)
9. Edelman, B.G., Luca, M.: Digital discrimination: the case of airbnb.com (2014)
10. Fish, B., Kun, J., Lelkes, Á.D.: A confidence-based approach for balancing fairness and accuracy. In: SIAM, pp. 144–152 (2016)
11. Forman, G.: Tackling concept drift by temporal inductive transfer. In: SIGIR, pp. 252–259. ACM (2006)
12. Gama, J.: Knowledge Discovery from Data Streams. CRC Press, Boca Raton (2010)
13. Ingold, D., Soper, S.: Amazon doesn’t consider the race of its customers. Should it. Bloomberg, April 2016
14. Iosifidis, V., Ntoutsi, E.: Dealing with bias via data augmentation in supervised learning scenarios. Jo Bates Paul D. Clough Robert Jäschke, p. 24 (2018)
15. Kamiran, F., Calders, T.: Classifying without discriminating. In: IC4. IEEE (2009)
16. Kamiran, F., Calders, T.: Classification with no discrimination by preferential sampling. In: BeneLearn, pp. 1–6. Citeseer (2010)
17. Kamiran, F., Calders, T.: Data preprocessing techniques for classification without discrimination. *KAIS* **33**(1), 1–33 (2012)
18. Kamiran, F., Calders, T., Pechenizkiy, M.: Discrimination aware decision tree learning. In: ICDM, pp. 869–874. IEEE (2010)
19. Klinkenberg, R.: Learning drifting concepts: example selection vs. example weighting. *IDA* **8**, 281–300 (2004)
20. Liu, L.T., Dean, S., Rolf, E., Simchowitz, M., Hardt, M.: Delayed impact of fair machine learning. arXiv preprint [arXiv:1803.04383](https://arxiv.org/abs/1803.04383) (2018)
21. Merz, C.J., Murphy, P.M.: {UCI} repository of machine learning databases (1998)
22. Pedreschi, D., Ruggieri, S., Turini, F.: Measuring discrimination in socially-sensitive decision records. In: SDM, pp. 581–592. SIAM (2009)
23. Pedreshi, D., Ruggieri, S., Turini, F.: Discrimination-aware data mining. In: SIGKDD, pp. 560–568. ACM (2008)

24. USA: Executive Office of the President, Podesta, J.: Big data: seizing opportunities, preserving values. White House, Executive Office of the President (2014)
25. Romei, A., Ruggieri, S.: A multidisciplinary survey on discrimination analysis. *TKER* **29**(5), 582–638 (2014)
26. Schlimmer, J.C., Granger, R.H.: Beyond incremental processing: tracking concept drift. In: *AAAI*, pp. 502–507 (1986)
27. Stoyanovich, J., Yang, K., Jagadish, H.: Online set selection with fairness and diversity constraints. In: *EDBT* (2018)
28. Sweeney, L.: Discrimination in online ad delivery. *Queue* **11**(3), 10 (2013)
29. Verma, S., Rubin, J.: Fairness definitions explained (2018)
30. Zafar, M.B., Valera, I., Gomez Rodriguez, M., Gummadi, K.P.: Fairness constraints: mechanisms for fair classification. arXiv preprint [arXiv:1507.05259](https://arxiv.org/abs/1507.05259) (2017)