

Adaptive Fairness-Aware Online Meta-Learning for Changing Environments

Chen Zhao*
chen.zhao@utdallas.edu
University of Texas at Dallas
Richardson, Texas, USA

Kai Jiang
kai.jiang@utdallas.edu
University of Texas at Dallas
Richardson, Texas, USA

Feng Mi*
feng.mi@utdallas.edu
University of Texas at Dallas
Richardson, Texas, USA

Latifur Khan
lkhan@utdallas.edu
University of Texas at Dallas
Richardson, Texas, USA

Xintao Wu
xintaowu@uark.edu
University of Arkansas
Fayetteville, Arkansas, USA

Feng Chen
feng.chen@utdallas.edu
University of Texas at Dallas
Richardson, Texas, USA

ABSTRACT

The fairness-aware online learning framework has arisen as a powerful tool for the continual lifelong learning setting. The goal for the learner is to sequentially learn new tasks where they come one after another over time and the learner ensures the statistic parity of the new coming task across different protected sub-populations (e.g. race and gender). A major drawback of existing methods is that they make heavy use of the *i.i.d* assumption for data and hence provide static regret analysis for the framework. However, low static regret cannot imply a good performance in changing environments where tasks are sampled from heterogeneous distributions. To address the fairness-aware online learning problem in changing environments, in this paper we first construct a novel regret metric FairSAR by adding long-term fairness constraints onto a strongly adapted loss regret. Furthermore, to determine a good model parameter at each round, we propose a novel adaptive fairness-aware online meta-learning algorithm, namely FairSAOML, which is able to adapt to changing environments in both bias control and model precision. The problem is formulated in the form of a bi-level convex-concave optimization with respect to the model's primal and dual parameters that are associated with the model's accuracy and fairness, respectively. Theoretic analysis provides sub-linear upper bounds for both loss regret and violation of cumulative fairness constraints. Our experimental evaluation on different real-world datasets with settings of changing environments suggests that the proposed FairSAOML significantly outperforms alternatives based on the best prior online learning approaches.

CCS CONCEPTS

• **Computing methodologies** → **Artificial intelligence; Machine learning**; • **Applied computing** → *Law, social and behavioral sciences*; • **Social and professional topics** → User characteristics.

*Both authors contributed equally to this research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD'22, August 14–18, 2022, Washington, DC, USA

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9385-0/22/08...\$15.00

<https://doi.org/10.1145/3534678.3539420>

KEYWORDS

Fairness, online meta learning, changing environment, adaption

ACM Reference Format:

Chen Zhao, Feng Mi, Xintao Wu, Kai Jiang, Latifur Khan, and Feng Chen. 2022. Adaptive Fairness-Aware Online Meta-Learning for Changing Environments. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'22)*, August 14–18, 2022, Washington, DC, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3534678.3539420>

1 INTRODUCTION

In real world, data containing bias are likely collected sequentially over time and the distribution assumptions of them may vary at some critical time. For example, a recent news [18] by *New York Times* reports that systematic algorithms become increasingly more discriminative to African Americans in bank loan during COVID-19 than pre-pandemic. These algorithms are built up from a sequence of data streams collected one after another over time, where at each time decision-makings are biased on the protected race population. This reveals (1) online algorithms completely ignore the importance of learning with fairness, in which fairness is defined by the equality of a predictive utility across different sub-populations, and predictions of a model are statistically independent on protected characters (e.g. race); (2) machine learning models make heavy use of the *i.i.d* assumption but this does not hold when environment changes (e.g. before and after the pandemic).

To control bias over time and especially ensure group fairness across different protected sub-populations, fairness-aware online algorithms capture supervised learning problems for which fairness is a concern and they are to sequentially train predictive models free from biases. Specifically, the goal of such algorithms is to ensure that both static loss regret, which compares the cumulative loss of the learner to that of the best fixed action in hindsight, and the violation of the sum of fair notions sub-linearly increase in the total number of rounds [28]. Although these works achieve state-of-the-art theoretic guarantees, the metric of static regret is only meaningful for stationary environments, and low static regret does not necessarily imply a good performance in changing environments since the time-invariant comparators may behave badly [24].

To address the limitation of changing environments in online learning, strongly adaptive regret [3] and dynamic regret [29] attract people's attention. Dynamic regret handles changing environments from a global prospective and it compares the cumulative

loss of the learner against any sequence of comparators but allows the comparators change over time. By contrast, strongly adaptive regret takes a local perspective and it gives more attentions on short time intervals. This regret can be viewed as the maximum statistic regret over all intervals [3]. Although several works [13, 24, 25] achieve sub-linear loss regret in online learning with changing environments, they completely ignore the significance of learning with fairness, which is a crucial hallmark of human intelligence.

In this paper, we introduce a novel problem, that is *fairness-aware online meta-learning in changing environments*. In this problem setting, a sequence of data batches (*i.e.* tasks) are collected one after another over time and the domains of these tasks may vary. The first goal of this work is to generalize the predictive learning accuracy and model fairness to the new data domain. Secondly, both loss regret and violation of cumulative fairness constraints are minimized and sublinearly increase in time.

To this end, technically we propose a novel online learning algorithm, namely *fair strongly adaptive online meta-learner* (FairSAOML). In this algorithm, a set of geometric covering intervals with different lengths are carefully designed. We determine model parameters by formulating a problem composed by two main levels: online fair interval-level and meta-level learning. Problems in two levels interplay each other with two parts of parameters: primal parameters θ regarding model accuracy and dual parameters λ adjusting fairness notions. More concretely, at round $t \in [T]$, a subset of intervals are selected to activate a number of experts where each expert runs an interval-specific algorithm. An expert takes a meta-solution pair $(\theta_{t-1}, \lambda_{t-1})$ from the previous round as input and outputs an interval-level solution $(\theta_{t,I}, \lambda_{t,I})$ at interval I . Then a meta-algorithm combines the weighted actions of all experts so that a solution pair (θ_t, λ_t) at round t can make predictions for the next round. The main contributions are summarized:

- To the best of our knowledge, for the first time a fairness-aware online meta-learning framework in changing environment is proposed. We first introduce a novel adaptive fairness-aware regret FairSAR. Then a novel algorithm FairSAOML is proposed to find a good decision at each round. Specifically, at each round the problem is formulated as a constrained bi-level convex-concave optimization with respect to a primal-dual parameter pair.
- Theoretically grounded analysis justifies the efficiency and effectiveness of the proposed method by demonstrating tighter bounds $O((\tau \log T)^{1/2})$ for the loss regret and $O((\tau T \log T)^{1/4})$ for violation of fairness constraints.
- We validate the performance of our approach with state-of-the-art techniques on real-world datasets. Our results demonstrate FairSAOML can effectively adapt both accuracy and fairness in changing environments and it shows substantial improvements over the best prior works.

2 RELATED WORK

Changing environments in online learning. Since the pioneering work [29] in online learning, numerous subsequent researches [10, 22] have been developed under the assumption of stationary environment with static regret. Low static regret, however, cannot imply a good performance in changing environment due to time-invariant comparators. To address this limitation, two regret

metrics, dynamic regret [29] and adaptive regret [11], are devised to measure the learner's performance in changing environments. To bound the general dynamic regret, the path-length of comparators [25, 29] is introduced and further developed. Unlike dynamic regret, adaptive regret handles changing environments from a local perspective by focusing on comparators in short intervals. To reduce the time complexity of adaptive regret based online algorithms, geometric covering intervals [3, 13, 24] and data streaming techniques [8] are developed. Although existing methods achieve state-of-the-art performance, a major drawback is that they immerse in minimizing objective functions but ignore the model fairness of prediction.

Fairness-aware online learning problems assume individuals arrive one at a time and the goal of such algorithms is to train predictive models free from biases. From the perspective of optimization, group fairness notions are normally considered as constraints added on learning objectives. However, when the constraints are complex, the computational burden of the projection onto constraints may be too high. Several closely related works, including FairFML [28], FairGLC [23], FairAOGD [12], aim to improve the theoretic guarantees by relaxing the output through a simpler close-form projection. However, these methods are not ideal for continual lifelong learning with changing task distributions, as they assume that all samples come from the same data distribution.

Online meta-learning addresses the issue of learning with fast adaptation, where a meta-learner learns knowledge transfer from history tasks onto new coming ones. FTML [7] can be considered as an application of MAML [6] in the setting of online learning. FairFML [28] extends FTML by controlling bias in a online working paradigm with task-specific adaptation. Unfortunately, none of such techniques are devised to adapt changing environments.

In this paper, to bridge above mentioned areas, we study the problem of fairness-aware online meta-learning to deal with changing task environments. In particular, at each round, model parameters are determined by the proposed novel algorithm FairSAOML. This algorithm refers to ideas of dynamic programming and expert tracking techniques. Inspired by fairness-aware online learning and meta-learning, a bi-level adaptation strategy is used to accommodate changing environments and learn model with accuracy and fairness.

3 PRELIMINARIES

3.1 Notations

An index set of a sequence of tasks is defined as $[T] = \{1, \dots, T\}$. Vectors are denoted by lower case bold face letters, *e.g.* the primal variables $\theta \in \Theta$ and the dual variables $\lambda \in \mathbb{R}_+^m$ where their i -th entries are θ_i, λ_i . Vectors with subscripts of task indices, such as θ_t, λ_t where $t \in [T]$, indicate model parameters for the task at round t . The Euclidean ℓ_2 -norm of θ is denoted as $\|\theta\|$. Given a differentiable function $\mathcal{L}(\theta, \lambda) : \Theta \times \mathbb{R}_+^m \rightarrow \mathbb{R}$, the gradient at θ and λ is denoted as $\nabla_\theta \mathcal{L}(\theta, \lambda)$ and $\nabla_\lambda \mathcal{L}(\theta, \lambda)$, respectively. Scalars are denoted by lower case italic letters, *e.g.* $\eta > 0$. $\Pi_{\mathcal{B}}$ is the projection operation to the set \mathcal{B} . $[u]_+$ denotes the projection of the vector u on the nonnegative orthant in \mathbb{R}_+^m , namely $[u]_+ = (\max\{0, u_1\}, \dots, \max\{0, u_m\})$. Some important notations are listed in Table 2 of Appendix A.1.

3.2 Fairness-Aware Constraints

In general, group fairness criteria used for evaluating and designing machine learning models focus on the relationships between the protected attribute and the system output [20, 26, 27]. The problem of group unfairness prevention can be seen as a constrained optimization problem. For simplicity, we consider one binary protected attribute (e.g. gender) in this work. However, our ideas can be easily extended to many protected attributes with multiple levels.

Let $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$ be the data space, where $\mathcal{X} = \mathcal{E} \cup \mathcal{S}$. Here $\mathcal{E} \subset \mathbb{R}^d$ is an input space, $\mathcal{S} = \{0, 1\}$ is a protected space, and $\mathcal{Y} = \{-1, 1\}$ is an output space for binary classification. Given a task (batch) of samples $\{\mathbf{e}_i, y_i, s_i\}_{i=1}^n \in (\mathcal{E} \times \mathcal{Y} \times \mathcal{S})$ where n is the number of datapoints, a fine-grained measurement to ensure fairness in class label prediction is to design fair classifiers by controlling the notions of fairness between protected subgroups, such as demographic parity and equality of opportunity [15, 21].

DEFINITION 1 (NOTIONS OF FAIRNESS [15, 21]). A classifier $h : \Theta \times \mathbb{R}^d \rightarrow \mathbb{R}$ is fair when its predictions are independent of the protected attribute $\mathbf{s} = \{s_i\}_{i=1}^n$. To get rid of the indicator function and relax the exact values, a linear approximated form of the difference between protected subgroups is defined [15],

$$g(\theta) = \left| \mathbb{E}_{(\mathbf{e}, y, s) \in \mathcal{Z}} \left[\frac{1}{\hat{p}_1(1 - \hat{p}_1)} \left(\frac{s+1}{2} - \hat{p}_1 \right) h(\theta, \mathbf{e}) \right] \right| - \epsilon \quad (1)$$

where $|\cdot|$ is the absolute function and $\epsilon > 0$ is the fairness relaxation determined by empirical analysis. \hat{p}_1 is an empirical estimate of p_{r1} . p_{r1} is the proportion of samples in group $s = 1$ and correspondingly $1 - p_{r1}$ is the proportion of samples in group $s = 0$.

Notice that, in Definition 1, when $\hat{p}_1 = \mathbb{P}_{(\mathbf{e}, y, s) \in \mathcal{Z}}(s = 1)$, the fairness notion $g(\theta)$ is defined as the difference of demographic parity (DDP). Similarly, when $\hat{p}_1 = \mathbb{P}_{(\mathbf{e}, y, s) \in \mathcal{Z}}(y = 1, s = 1)$, $g(\theta)$ is defined as the difference of equality of opportunity (DEO) [15]. Therefore, parameters θ in the domain of a task is feasible if it satisfies the fairness constraint $g(\theta) \leq 0$.

3.3 Fairness-Aware Online Learning

The protocol of fairness-aware online convex optimization can be viewed as a repeated game between a learner and an adversary, where the learner is faced with tasks $\{\mathcal{D}_t\}_{t=1}^T$ one after another. At each round $t \in [T]$,

- **Step 1:** The learner selects a model parameter θ_t in the fair domain Θ .
- **Step 2:** The adversary reveals a loss function $f_t : \Theta \times \mathbb{R}^d \rightarrow \mathbb{R}$ and m fairness functions $g_i : \Theta \times \mathbb{R}^d \rightarrow \mathbb{R}, \forall i \in [m]$.
- **Step 3:** The learner incurs an instantaneous loss $f_t(\theta_t, \mathcal{D}_t)$ and m fairness notions $g_i(\theta_t, \mathcal{D}_t), \forall i \in [m]$.
- **Step 4:** Advance to round $t + 1$.

The goal of fairness-aware online learning [23, 28] is to (1) minimize the loss regret over the rounds which is to compare to the cumulative loss of the best fixed model in hindsight, and (2) ensure the total violation of fair constraints sublinearly increase in T . The loss regret is typically referred as *static regret* since the comparator is time-invariant. To control bias and especially ensure group fairness across different protected sub-populations, fairness notions

are considered as constraints added on optimization problems.

$$\begin{aligned} \min_{\theta_1, \dots, \theta_T \in \Theta} \quad & \text{Regret}(T) = \sum_{t=1}^T f_t(\theta_t, \mathcal{D}_t) - \min_{\theta \in \Theta} \sum_{t=1}^T f_t(\theta, \mathcal{D}_t) \quad (2) \\ \text{subject to} \quad & \sum_{t=1}^T g_i(\theta_t, \mathcal{D}_t) \leq O(T^\gamma), \forall i \in [m], \gamma \in (0, 1) \end{aligned}$$

where the summation of fair constraints is defined as *long-term constraints* in [17]. The big O notation in the constraint is to bound the total violation of fairness sublinear in T . The main drawback of using the metric of static regret is that it is only meaningful for stationary environments, and low static regret cannot imply a good performance in changing environments since the time-invariant comparator in Eq.(2) may behave badly [24].

4 METHODOLOGY

4.1 Settings and Problem Formulation

To address the limitation of changing environments in online learning, *adaptive regret* (AR) based on [11] is defined as the maximum static regret over any contiguous intervals. However, AR does not respect short intervals well. To this end, *strongly adaptive regret* (SAR) [4] is proposed to improve AR, which emphasizes the dependence on lengths of intervals and it takes the form that

$$\text{SAR}(T, \tau) = \max_{[s, s+\tau-1] \subseteq [T]} \left(\sum_{t=s}^{s+\tau-1} f_t(\theta_t, \mathcal{D}_t) - \min_{\theta \in \Theta} \sum_{t=s}^{s+\tau-1} f_t(\theta, \mathcal{D}_t) \right) \quad (3)$$

where τ indicates the length of time interval. In SAR, the learner is competing with changing comparators, as θ varies with s over $[s, s + \tau - 1]$.

In this paper, we consider the online meta-learning setting that is similar in [7, 28] but tasks are sampled from heterogeneous distributions. Instead of static regret, we define a novel regret FairSAR in Eq.(4). Let $\{\theta_t\}_{t=1}^T$ be the sequence of model parameters generated in the **Step 1** of the learning protocol (see Sec. 3.3). The goal of our problem is to minimize FairSAR under the long-term fair constraints:

$$\begin{aligned} \text{FairSAR}(T, \tau) = \max_{[s, s+\tau-1] \subseteq [T]} & \left(\sum_{t=s}^{s+\tau-1} f_t(\mathcal{G}_t(\theta_t, \mathcal{D}_t^S), \mathcal{D}_t^V) \right. \\ & \left. - \min_{\theta \in \Theta} \sum_{t=s}^{s+\tau-1} f_t(\mathcal{G}_t(\theta, \mathcal{D}_t^S), \mathcal{D}_t^V) \right) \quad (4) \\ \text{subject to} \quad & \max_{[s, s+\tau-1] \subseteq [T]} \left(\sum_{t=s}^{s+\tau-1} g_i(\mathcal{G}_t(\theta_t, \mathcal{D}_t^S), \mathcal{D}_t^V) \right) \leq O(T^\gamma), \forall i \in [m] \end{aligned}$$

where $\gamma \in (0, 1)$. $\mathcal{D}_t^S, \mathcal{D}_t^V \subset \mathcal{D}_t$ are the support and validation set. $\mathcal{G}_t(\cdot)$ is the base learner which corresponds to one or multiple gradient steps [6]. Different from traditional online learning settings, the long-term constraint violation $g(\cdot) : \mathcal{B} \times \mathbb{R}^d \rightarrow \mathbb{R}$ is satisfied. To facilitate our analysis, θ_t is originally chosen from its domain $\Theta = \{\theta \in \mathbb{R}^d : g_i(\theta, \mathcal{D}_t) \leq 0, \forall i \in [m]\}$. A projection operator is hence typically applied to the updated variables in order to make them feasible [12, 17, 23]. In order to lower the computational complexity and accelerate the online processing speed, we relax the domain Θ to \mathcal{B} , where $\Theta \subseteq \mathcal{B} = \mathbb{S}\mathbb{K}$ with \mathbb{K} being the unit ℓ_2 ball centered at the origin, and $S = \max\{r > 0 : r = \|\theta_1 - \theta_2\|, \forall \theta_1, \theta_2 \in \Theta\}$.




Figure 1: (Upper) An graphical illustration of AGC intervals (base=2) when $T = 18$. The total interval set \mathcal{I} consists of 4 subsets $\{I_0, I_1, I_2, I_3\}$ and each contains different numbers of intervals with fixed length. Intervals covered by shadow is an example of target subset C_5 when $t = 5$. **(Lower)** A demonstration of active experts activated by C_t and their corresponding sleeping experts.

In the protocol stated in Sec. 3.3, the key step (**Step 1**) is to find a good parameter θ_t at each round. Let \mathcal{K} be an online learning algorithm determining θ_t . A trick commonly used in designing a meta algorithm for changing environments is to initiate a new instance of \mathcal{K} at every round [11]. Therefore, we run \mathcal{K} independently for each interval $I \in \{[t, \infty) | t = 1, 2, \dots\}$ and denote \mathcal{K}_I as the run of \mathcal{K} on interval I . At round t , we combine the decisions from the runs by weighted average. The key idea is that at time t , some of the outputs in $\mathcal{K} \in \{\mathcal{K}_I\}_{I \ni t}$ are not based on any data prior to time $t' < t$, so that if the environment changes at t' , those outputs may be given a larger weight by the meta algorithm, allowing it to adapt more quickly to the change. A main drawback, however, is a factor of t increase in the time complexity. To avoid this, we reduce it to $O(\log t)$ by restarting algorithms on a designed set of adaptive geometric covering (AGC) intervals.

4.2 AGC Intervals and Experts

In Eq.(4), FairSAR evaluates the learner's performance on each time interval and it is the maximum regret over any contiguous intervals. Similar to the seminal work of SAR [3], we construct a number of interval sets, namely adaptive geometric covering (AGC) intervals, where each set contains various intervals with same length. AGC intervals can be considered as a special case of a more general set of intervals and they hence efficiently reduce time complexity to $O(\log t)$. A set of contiguous AGC intervals \mathcal{I} are defined as

$$\mathcal{I} = \bigcup_{k \in [\lfloor \log_2^T \rfloor - 1] \cup \{0\}} \mathcal{I}_k \quad (5)$$

where $\forall k, \mathcal{I}_k = \{I_k^i | [(i-1) \cdot 2^k + 1, \min\{T, i \cdot 2^k\}] : i \in \mathbb{N}\}$. An example with $T = 18$ is given in Figure 1 to illustrate the composition of AGC intervals. Notice that the log base equals to 2 is not required, but larger base number leads to less interval levels.

Furthermore, inspired by *learning with expert advice* problems [13], we construct a set \mathcal{U} of interval-level learning processes, defined as experts, where $|\mathcal{U}| = \lfloor \log_2^T \rfloor$. To better adapt changing environment, only one expert is assigned to the corresponding interval subset \mathcal{I}_k . At each round $t \in [T]$, we introduce a target set

$C_t \subset \mathcal{I}$ which includes a set of intervals starting from t

$$C_t = \{I | I \in \mathcal{I}, t \in I, (t-1) \notin I\} \quad (6)$$

C_t aims to dynamically activate a subset of experts in \mathcal{U} (an example of $C_{t=5}$ is given in Figure 1). A set of active experts at round t is denoted as $\mathcal{A}_t \subseteq \mathcal{U}$ and the corresponding inactive ones are called sleeping experts. The set of sleeping experts is denoted $\mathcal{S}_t = \mathcal{U} \setminus \mathcal{A}_t$. Note that only active experts are used to learn and further update interval-level model parameters through a base learner $\mathcal{G}_t(\cdot)$ that corresponds to one or multiple gradient steps [6]. The *interval-level* parameter update for an active expert E_I in \mathcal{A}_t at round t is defined

$$\begin{aligned} \theta_{t,I} &:= \mathcal{G}_t(\theta, \mathcal{D}_{t,I}^S) = \arg \min_{\theta} f_t(\theta, \mathcal{D}_{t,I}^S) \\ \text{subject to } g_i(\theta, \mathcal{D}_{t,I}^S) &\leq 0, \forall i \in [m] \end{aligned} \quad (7)$$

where the loss function $f_t(\cdot)$ and the fairness function $g_i(\cdot)$ are defined based on the support set $\mathcal{D}_{t,I}^S \subset \mathcal{D}_{t,I}$ associated with E_I .

4.3 Learning Dynamically for Bi-Level Adaptation

Recall that in the protocol of fairness-aware online learning (Sec. 3.3), the main goal for the learner is to sequentially decide on model parameter θ_t that performs well on the loss sequence and the long-term fair constraints. Crucially, inspired by [6], we consider a setting where at each round t the learner can perform a number of expert-specific updates at an interval level in the active set \mathcal{A}_t .

As specified in Eq.(4), model parameters at each round t are determined by formulating problems with a nested bi-level adaptation process: interval-level and meta-level. Each level corresponds to a sub-learner, i.e. base and meta learner respectively, described in Figure 2. The problem of learning a meta-level parameter θ_t is embedded with the optimization problem of finding interval-level parameters $\theta_{t,I}$ in Eq.(7). For experts in the sleeping set \mathcal{S}_t , the base learner is not applied. The *meta-level* problem takes the form

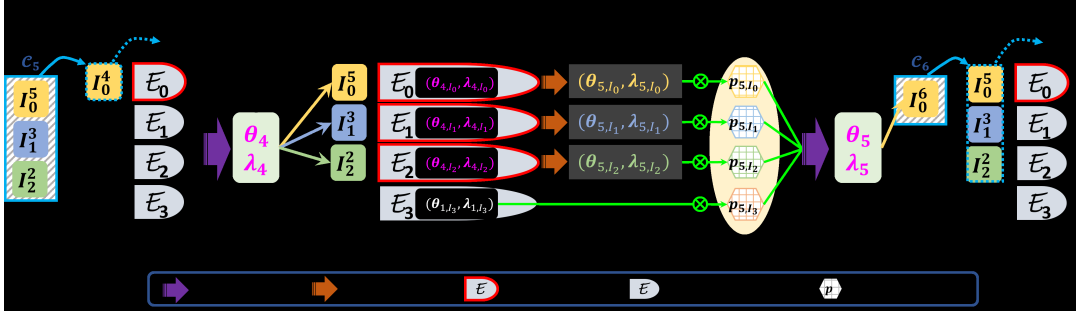


Figure 2: An overview of our FairSAOML to determine model parameter pair at each round. A target set (shadowed) of intervals are initially selected and are later used to activate corresponding experts. Each active expert runs through a base learner for the interval-level parameter-pair adaption and its weight is updated. The meta-level parameter-pair is finally attained through the meta learner by combining weighted actions of all experts.

in Eq.(8). For simplicity, we abuse the subscription E_I with I .

$$\min_{\theta \in \mathcal{B}} \sum_{E_I \in \mathcal{A}_t} p_{t,I} \cdot f_t(\mathcal{G}_t(\theta, \mathcal{D}_{t,I}^S), \mathcal{D}_{t,I}^Q) + \sum_{E_I \in \mathcal{S}_t} p_{t,I} \cdot f_t(\theta_{t',I}, \mathcal{D}_{t,I}^Q) \quad (8)$$

$$\text{s.t.} \sum_{E_I \in \mathcal{A}_t} p_{t,I} \cdot g_t(\mathcal{G}_t(\theta, \mathcal{D}_{t,I}^S), \mathcal{D}_{t,I}^Q) + \sum_{E_I \in \mathcal{S}_t} p_{t,I} \cdot g_t(\theta_{t',I}, \mathcal{D}_{t,I}^Q) \leq 0$$

where $p_{t,I} \geq 0$ is the expert weight of E_I at t . $\mathcal{D}_{t,I}^Q \subset \mathcal{D}_{t,I}$ is the query set where $\mathcal{D}_{t,I}^Q \cap \mathcal{D}_{t,I}^S = \emptyset$. $\theta_{t',I}$ is the interval-level model parameter for a sleeping expert $E_I \in \mathcal{S}_t$ where the round index $t' < t$ represents the last time this expert was activated.

In the following section, we introduce our proposed algorithm FairSAOML. In stead of optimizing primal parameters only, it efficiently deals with the bi-level optimization problem of Eq.(7)(8) by approximating a sequence of pairs of primal-dual meta parameters $\{(\theta_t, \lambda_t)\}_{t=1}^T$ where the pair respectively responds for adjusting accuracy and fairness level.

4.4 An Efficient Algorithm: FairSAOML

To find a good model parameter pair (θ_t, λ_t) at each round, an efficient working flow is proposed in Algorithm 1 and an overview of our proposed FairSAOML is given in Figure 2. Inspired by dynamic programming and expert-tracking [16] techniques, experts at each round are recursively divided into active and sleeping sets. Model parameters in active experts are locally updated, but for the ones in sleeping experts, they are directly inherited from the previous round. Specifically, in the beginning of round t , a target set C_t containing intervals is used to activate a subset of experts in \mathcal{U} . For each active expert E_I in \mathcal{A}_t , an interval-level algorithm takes the meta-level solution $(\theta_{t-1}, \lambda_{t-1})$ and outputs an expert-specific solution pair $(\theta_{t,I}, \lambda_{t,I})$. Finally, through the meta-learner, we combine the weighted solutions of all experts and move to the next round.

We explain the main steps in Algorithm 1 below. In step 4, when a new task arrives at round t , a batch of data \mathcal{D}_t^V is randomly sampled from \mathcal{D}_t for validation purpose and the performance on θ_{t-1} achieved is recorded. A target set of intervals C_t is selected from \mathcal{I} in step 5. For each interval $I \in C_t$ (step 6-11), the corresponding expert $E_{t,I}$ is activated. We set adaptive stepsizes $\eta_{t,I} = S/(G\sqrt{|I|})$ where $|I| = 2^k$ denotes the interval length in

\mathcal{I}_k . S is the radius of the Euclidean ball \mathcal{B} and there exists a constant $G > 0$ that bounds the (sub)gradients of f_t and g_t . Following the setting used in [12], empirically we set $S = \sqrt{1+2\epsilon} - 1$ and $G = \max\{\sqrt{d} + S, \max_t\{\|\mathbf{e}_I\|_2, \mathbf{e}_I \in \mathcal{P}_t\}\}$, where \mathbf{e}_I is the non-protected features lied in the interval I and d is its feature dimension. \mathcal{P}_t is a set which includes all past intervals until time t . In steps 12-14, for all experts in \mathcal{U} , a following weight $p_{t,I}$ is assigned:

$$p_{t,I} = \frac{w(R_{t,I}, C_{t,I})}{\sum_{E_I \in \mathcal{U}} w(R_{t,I}, C_{t,I})} \quad (9)$$

Here, a weight function [16] is defined as $w(R, C) = \frac{1}{2}(\Phi(R+1, C+1) - \Phi(R-1, C-1))$, where $\Phi(R, C) = \exp([R]_+^2/3C)$ and $[r]_+ = \max(0, r)$ and $\Phi(0, 0) = 1$. In steps 15-24, our FairSAOML responds to the bi-level adaptation stated in Eq.(7) and (8). Specifically, to solve the interval-level problem in Eq.(7), for each active expert E_I in \mathcal{A}_t , we consider following Lagrangian function

$$\mathcal{F}_{t,I}(\theta_{t-1}, \lambda_{t-1}) = f_t(\theta_{t-1}, \mathcal{D}_{t,I}^S) + \sum_{i=1}^m \lambda_{t-1,i} \cdot g_t(\theta_{t-1}, \mathcal{D}_{t,I}^S) \quad (10)$$

where the interval-level parameter pair for an active expert E_I are initialized with the meta-level parameter $(\theta_{t-1}, \lambda_{t-1})$. For optimization with simplicity, cumulative constraints in Eq.(7) are approximated with the summarized regularization. Interval-level parameters are updated through a base learner $\mathcal{G}_t(\cdot)$. One example for the learner is updating with one gradient step [6] using the pre-determined adaptive stepsize $\eta_{t,I}$. Notice that for multiple gradient steps, $\theta_{t,I}$ and $\lambda_{t,I}$ interplay each other for updating.

$$\theta_{t,I} = \theta_{t-1} - \eta_{t,I} \nabla_{\theta} \mathcal{F}_{t,I}(\theta_{t-1}, \lambda_{t-1}) \quad (11)$$

$$\lambda_{t,I} = \lambda_{t-1} + \eta_{t,I} \nabla_{\lambda} \mathcal{F}_{t,I}(\theta_{t,I}, \lambda_{t-1}) \quad (12)$$

Next, to solve the meta-level problem in Eq.(8), we combine the actions of active experts together with sleeping experts. We consider the following augmented Lagrangian function and abuse the symbol t' with t in Eq.(8):

$$\mathcal{L}_t(\theta_{t,I}, \lambda_{t,I}) = \sum_{E_I \in \mathcal{U}} p_{t,I} \left(f_t(\theta_{t,I}, \mathcal{D}_{t,I}^Q) + \sum_{i=1}^m \left(\lambda_{i,t,I} \cdot g_t(\theta_{t,I}, \mathcal{D}_{t,I}^Q) - \frac{\delta(\eta_1 + \eta_2)}{2} \lambda_{i,t,I}^2 \right) \right) \quad (13)$$

where $\delta > 0$ is a constant determined by analysis. Note that the last augmented term on the dual variable is devised to prevent λ from

Algorithm 1 FairSAOML

```

1: Initialize meta-parameters pair  $(\theta_0, \lambda_0)$ , where  $\theta_0$  is the center
   of  $\mathcal{B}$  and  $\lambda_0 \in \mathbb{R}_+^m$  is randomly chosen
2: Initialize  $k = \lfloor \log_2^T \rfloor$  experts in set  $\mathcal{U}$  and for each expert
    $E_I \in \mathcal{U}$ , set  $R_{0,I} = 0, C_{0,I} = 0, p_{0,I} = 1$ .
3: for each  $t \in [T]$  do
4:   Sample  $\mathcal{D}_t^V \subset \mathcal{D}_t$  and record the performance of  $\theta_{t-1}$ 
5:   Subset  $\mathcal{C}_t$  from  $\mathcal{I}$  using Eq.(6)
6:   for each interval  $I \in \mathcal{C}_t$  do
7:     Activate expert  $E_{t,I}$  and set  $\eta_{t,I} = S/(G\sqrt{|I|})$ ,
        $(\theta_{t,I}, \lambda_{t,I}) \leftarrow (\theta_{t-1}, \lambda_{t-1})$ 
8:     Identify an  $E_J \in \mathcal{U}$ , such that  $|J| = |I|$ 
9:     Update  $R_{t,I} \leftarrow R_{t-1,I}$  and  $C_{t,I} \leftarrow C_{t-1,I}$ 
10:    Replace  $E_J$  with  $E_I$  in  $\mathcal{U}$ 
11:   end for
12:   for each expert in  $\mathcal{U}$  do
13:     Update  $p_{t,I}$  using  $R_{t,I}, C_{t,I}$  in Eq.(9)
14:   end for
15:   for  $n = 1, \dots, N_{meta}$  steps do
16:     for each active expert  $E_I$  in  $\mathcal{A}_t$  do
17:       Sample support set  $\mathcal{D}_{t,I}^S \subset \mathcal{D}_{t,I}$ 
18:       Adapt interval-level primal and dual variables with
        $\mathcal{D}_{t,I}^S$  using Eq.(11)(12)
19:     end for
20:     for each expert  $E_I$  in  $\mathcal{U}$  do
21:       Sample query set  $\mathcal{D}_{t,I}^Q \subset \mathcal{D}_{t,I}$ 
22:     end for
23:     Update meta-level primal and dual variables with  $\mathcal{D}_{t,I}^Q$ 
       using Eq.(14)(15)
24:   end for
25:   for each expert  $E_I \in \mathcal{U}$  do
26:      $R_{t+1,I} = R_{t,I} + \mathcal{F}_{t,I}(\theta_{t,I}, \lambda_{t,I}) - \mathcal{F}_{t,I}(\theta_{t-1,I}, \lambda_{t-1,I})$ 
27:      $C_{t+1,I} = C_{t,I} + \left| \mathcal{F}_{t,I}(\theta_{t,I}, \lambda_{t,I}) - \mathcal{F}_{t,I}(\theta_{t-1,I}, \lambda_{t-1,I}) \right|$ 
28:   end for
29: end for

```

being too large. The update rule for meta-level parameters follows:

$$\theta_t = \prod_{\mathcal{B}} \left(\theta_{t-1} - \eta_1 \nabla_{\theta} \mathcal{L}_t(\theta_{t,I}, \lambda_{t,I}) \right) \quad (14)$$

$$\lambda_t = \left[\lambda_{t-1} + \eta_2 \nabla_{\lambda} \mathcal{L}_t(\theta_{t,I}, \lambda_{t,I}) \right]_+ \quad (15)$$

where $\prod_{\mathcal{B}}$ is the projection operation to the relaxed domain \mathcal{B} that is introduced in Sec.4.1. This approximates the true desired projection with a simpler closed-form. Finally, in steps 25-28, for each expert we update its R and C values which determine the expert weight for the next round. The intuition of weight update is to re-adjust difference between the meta-solution and the interval-level solution given by the expert.

5 ANALYSIS

To analysis, we first make following assumptions as in [17, 24]. Examples where these assumptions hold include logistic regression and L_2 regression over a bounded domain. As for constraints, a family of fairness notions, such as DDP stated in Definition 1, are

applicable as discussed in [15]. For simplicity, in this section we omit \mathcal{D} used in $f_t(\cdot), \forall t$ and $g_i(\cdot), \forall i$.

ASSUMPTION 1 (CONVEX DOMAIN). *The convex set Θ is non-empty, closed, bounded, and it is described by m convex functions as $\Theta = \{\theta : g_i(\theta) \leq 0, \forall i \in [m]\}$. The relaxed domain \mathcal{B} (where $\Theta \subseteq \mathcal{B}$) contains the origin $\mathbf{0}$ and its diameter is bounded by S .*

ASSUMPTION 2. *Both the loss functions $f_t(\cdot), \forall t$ and constraint functions $g_i(\cdot), \forall i \in [m]$ satisfy the following assumptions*

- (1) (Lipschitz Continuous) $\forall \theta_1, \theta_2 \in \mathcal{B}$, $\|f_t(\theta_1) - f_t(\theta_2)\| \leq L_f \|\theta_1 - \theta_2\|, \|g_i(\theta_1) - g_i(\theta_2)\| \leq L_g \|\theta_1 - \theta_2\|$. Let $G = \max\{L_f, L_g\}$, $F = \max_{t \in [T]} \max_{\theta_1, \theta_2 \in \mathcal{B}} f_t(\theta_1) - f_t(\theta_2) \leq 2L_f S$, and $D = \max_{i \in [m]} \max_{\theta \in \mathcal{B}} g_i(\theta) \leq L_g S$
- (2) (Lipschitz Gradient) $f_t(\theta), \forall t$ are β_f -smooth and $g_i(\theta), \forall i$ are β_g -smooth, that is, $\forall \theta_1, \theta_2 \in \mathcal{B}$, $\|\nabla f_t(\theta_1) - \nabla f_t(\theta_2)\| \leq \beta_f \|\theta_1 - \theta_2\|, \|\nabla g_i(\theta_1) - \nabla g_i(\theta_2)\| \leq \beta_g \|\theta_1 - \theta_2\|$.
- (3) (Lipschitz Hessian) Twice-differentiable functions $f_t(\theta), \forall t$ and $g_i(\theta), \forall i$ have ρ_f and ρ_g -Lipschitz Hessian, respectively. That is, $\forall \theta_1 - \theta_2 \in \mathcal{B}$, $\|\nabla^2 f_t(\theta_1) - \nabla^2 f_t(\theta_2)\| \leq \rho_f \|\theta - \phi\|, \|\nabla^2 g_i(\theta_1) - \nabla^2 g_i(\theta_2)\| \leq \rho_g \|\theta_1 - \theta_2\|$.

ASSUMPTION 3 (STRONGLY CONVEXITY). *Suppose $f_t(\theta), \forall t$ and $g_i(\theta), \forall i$ have strong convexity, that is, $\forall \theta_1, \theta_2 \in \mathcal{B}$, $\|\nabla f_t(\theta_1) - \nabla f_t(\theta_2)\| \geq \mu_f \|\theta_1 - \theta_2\|, \|\nabla g_i(\theta_1) - \nabla g_i(\theta_2)\| \geq \mu_g \|\theta_1 - \theta_2\|$.*

Under above assumptions, we state the key Theorem 1 that the proposed FairSAOML enjoys sub-linear guarantee for both regret and long-term fairness constraints in the long run for Algorithm 1. Proof is given in Appendix B.

THEOREM 1. *Set $\theta^* = \arg \min_{\theta \in \Theta} \sum_{t=s}^{s+\tau-1} f_t(\mathcal{G}_t(\theta))$ where $[s, s+\tau-1] \subseteq [T]$. Under Assumptions 1, 2 and 3, the regret FairSAR proposed in Eq.(4) of FairSAOML in Algorithm 1 satisfies*

$$\max_{[s, s+\tau-1] \subseteq [T]} \left(\sum_{t=s}^{s+\tau-1} f_t(\mathcal{G}_t(\theta_t)) - f_t(\mathcal{G}_t(\theta^*)) \right) \leq O\left((\tau \log T)^{1/2}\right)$$

$$\max_{[s, s+\tau-1] \subseteq [T]} \left(\sum_{t=s}^{s+\tau-1} g_i(\mathcal{G}_t(\theta_t)) \right) \leq O\left((\tau T \log T)^{1/4}\right), \quad \forall i \in [m]$$

Discussion for Upper Bounds. Under aforementioned assumptions and provable convexity of Eq.(16) in θ (see Lemma 2 in Appendix B), the proposed FairSAOML in Algorithm 1 achieves sub-linear bounds in FairSAR for both loss regret and violation of fairness constraints. Although such bounds are comparable with the strongly adapted loss regret in [13, 24] (see Table 1) in terms of online learning in changing environment paradigms, for the first time we bound loss regret and cumulative fairness constraints simultaneously. On the other hand, in terms of fairness-aware online learning, our proposed method outperforms [12, 23, 28] by giving a tighter bound of fair constraint violations.

Complexity. The computational complexity of FairSAOML in Algorithm 1 at each round $t \in [T]$ is $O(N_{meta} \cdot |\mathcal{U}|)$ where N_{meta} is the number of meta-level iterations and $|\mathcal{U}| = O(\log T)$ is the total number of experts that needs to be maintained, and the complexity of each expert is $O(1)$.

Table 1: Comparison of upper bounds in loss regret and constraint violations across various methods.

Algorithms	Static Environment				Changing Environment		
	FTML[7]	FairFML[28]	FairAOGD[12]	FairGLC[23]	AOD[24]	CBCE[13]	FairSAOML(Ours)
Loss Regret	$O(\log T)$	$O(\log T)$	$O(T^{2/3})$	$O(\log T)$	$O((\tau \log T)^{1/2})$	$O((\tau \log T)^{1/2})$	$O((\tau \log T)^{1/2})$
Constraint Violations	-	$O((T \log T)^{1/2})$	$O(T^{2/3})$	$O((T \log T)^{1/2})$	-	-	$O((\tau T \log T)^{1/4})$

6 EXPERIMENTAL SETTINGS

Datasets. We use the following publicly available datasets. (1) *New York Stop-and-Frisk* [14] is a prominent dataset of a real-world application on policing in the new york city from 2009 to 2010. It documents whether a pedestrian who was stopped on suspicion of weapon possession would in fact possess a weapon. As this data had a pronounced racial bias on African Americans, for each frisked record, we consider race as the binary protected attribute, that is black and non-black. Besides, this dataset consists of records collected in five different sub-districts, Manhattan (**M**), Brooklyn (**B**), Queens (**Q**), Bronx (**R**) and Staten (**S**). Since there are large performance disparities across districts and race groups, each district is viewed as an independent domain. To adapt the online learning setting, data in each domain is further split into 32 tasks and each task corresponds to ten days of a month with 111 non-protected features. According to *DDP* values in Definition 1, the fairness levels from low to high are Bronx (0.74), Queens (0.68), Staten (0.65), Manhattan (0.53) and Brooklyn (0.44). The larger *DDP* values indicate lower fairness level. We hence consider two settings for domain adaptation where each setting contains 96 tasks in total: (i) fairness level from high to low: Brooklyn to Manhattan to Staten (**B**→**M**→**S**); and (ii) fairness level from low to high: Bronx to Queens then Staten (**R**→**Q**→**S**). (2) *MovieLens*¹ contains 100k ratings by 943 users on 1682 movies and each rating is given a binary label ("recommending" if rating greater than 3, "not recommending" otherwise). We consider gender as the protected attribute. To generate dynamic environments, following [19] we construct a larger dataset by combining three copies of the original data and flipping the original values of non-protected attributes by multiplying -1 for the middle copy. Therefore, each copy is considered as a data domain. Furthermore, each data copy is split into 30 tasks by timestamps and there are 90 tasks in total.

Evaluation Metrics. Two popular evaluation metrics are introduced that each allows quantifying the extent of bias taking into account the protected attribute. *Demographic Parity* (DP) [5] and *Equalized Odds* (EO) [9] can be formalized as

$$\text{DP} = \frac{P(\hat{Y} = 1|S = 0)}{P(\hat{Y} = 1|S = 1)}; \quad \text{EO} = \frac{P(\hat{Y} = 1|S = 0, Y = y)}{P(\hat{Y} = 1|S = 1, Y = y)}$$

where $y \in \{-1, 1\}$. Equalized odds requires that \hat{Y} have equal true positive rates and false positive rates between sub-groups. For both metrics, a value closer to 1 indicate fairness.

Competing Methods. We compare the performance of our algorithm FairSAOML with six baseline methods. These baselines are chosen from three perspectives: online meta learning (MaskFTML, FairFML), online fairness learning (FairFML, FairAOGD, FairGLC) and online learning in changing environment (AOD, CBCE). (1) MaskFTML [7]: the original FTML finds a sequence of meta parameters by simply applying MAML [6] at each round. To focus

on fairness learning, this approach is applied to modified datasets by simply removing protected attributes. (2) FairFML [28] controls bias in an online working paradigm and aims to attain zero-shot generalization with task-specific adaptation. Different from our FairSAOML, FairFML focuses on static environment and assumes tasks sampled from an unchangeable distribution. (3) FairAOGD [12] is proposed for online learning with long-term constraints. In order to fit bias-prevention and compare them to FairSAOML, we specify such constraints as DDP stated in Definition 1. (4) FairGLC [23] rectifies FairAOGD by square-clipping the constraints in place of $g_i(\cdot)$, $\forall i$. (5) AOD [24] minimizes the strongly adaptive regret by running multiple online gradient descent algorithms over a set of dense geometric covering intervals. (6) CBCE [13] adapts changing environment in an online learning paradigm by combining the sleeping bandits idea with the coin betting algorithm.

Settings. As discussed in Sec.5, the performance of our proposed method has been well justified theoretically for machine learning models whose objectives are strongly convex and smooth. However, in machine learning and fairness studies, due to the non-linearity of neural networks, many problems have a non-convex landscape where theoretical analysis is challenging. Nevertheless, algorithms originally developed for convex optimization problems like gradient descent have shown promising results in practical non-convex settings [7]. Taking inspiration from these successes, we describe practical instantiations for the proposed online algorithm, and empirically evaluate the performance in Sec.7.

For each task we set the number of fairness constraints to one, i.e. $m = 1$. For the rest, we follow the same settings as used in online meta learning [7, 28]. In particular, we meta-train with support size of 400 for each class and 800 for a query set, whereas 90% (hundreds of datapoints) of task samples for evaluation. Besides, for *New York Stop-and-Frisk* dataset, we choose base = 2 and the total number of experts is $\lfloor \log_2^9 \rfloor = 6$. Similarly, we choose base = 3 for the *MovieLens* dataset and hence it has $\lfloor \log_3^{90} \rfloor = 4$ experts in total. All the baseline models that are used to compare with our proposed approach share the same neural network architecture and parameter settings. All the experiments are repeated 10 times with the same settings and the mean and standard deviation results are reported. Details on the settings and hyperparameter tuning are given in Appendix A.2 and A.3.

7 RESULTS

7.1 Overall Performance

As for all methods showcased in in Figure 3, higher is better for all plots, while shaded regions indicate standard errors. The learning curves show that the proposed FairSAOML effectively controls bias as the learner sees more tasks, and it eventually meets the fair condition of "80%-rule" [1] where DP and EO at last several rounds are beyond 0.8. Furthermore, FairSAOML substantially outperforms most of the alternative approaches in achieving the best model precision represented by high accuracies. Besides, in terms of learning

¹<https://grouplens.org/datasets/movielens/100k/>

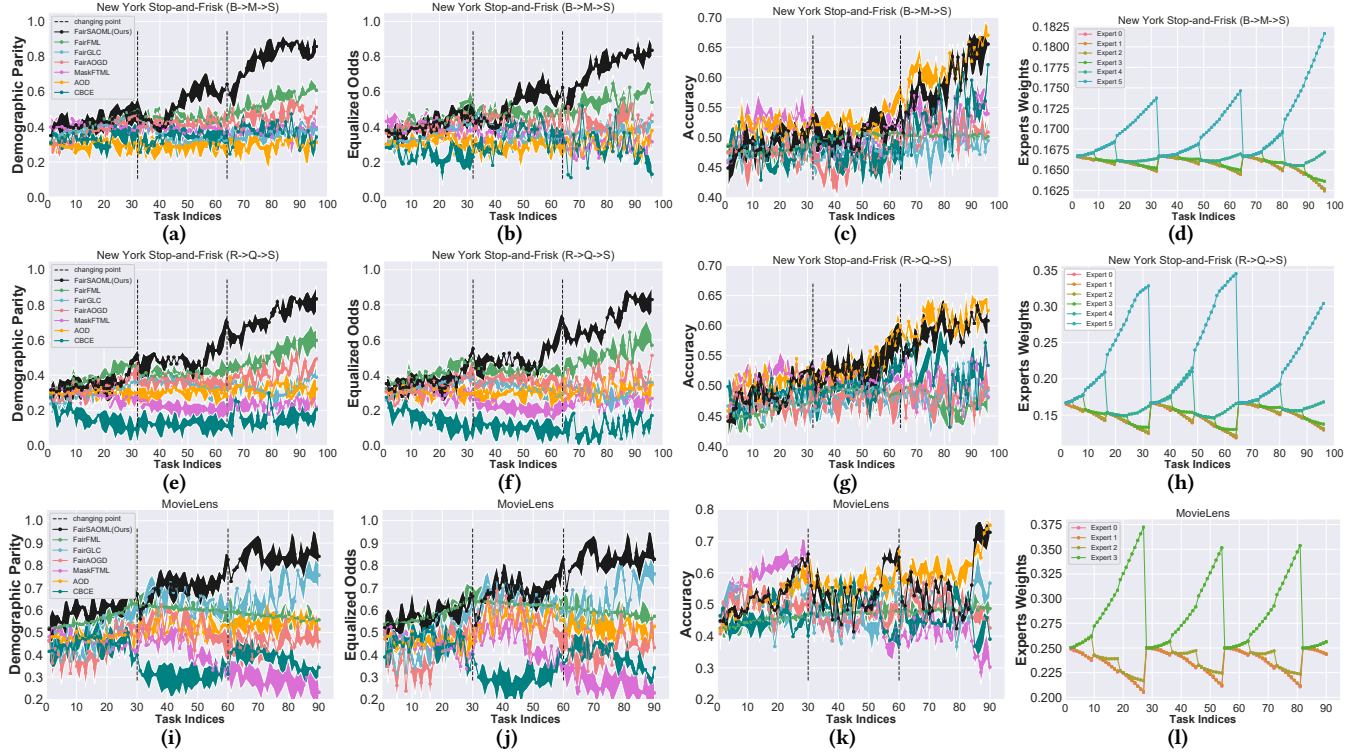


Figure 3: Model performance over real-world datasets through each round. New York Stop-and-Frisk (a-d) B→M→S, (e-h) R→Q→S; (i-l) MovieLens.




Figure 4: Running time comparison over baseline methods




Figure 5: Ablation studies on the MovieLens dataset.

efficiency, our FairSAOML shows the least running time compared with baselines in Figure 4. This results from that (1) at each time the total number of experts are fixed (*i.e.* 4 in *MovieLens*), and (2) instead of using the entire data task, only a subset (support) of data is used for parameters update of active experts.

7.2 Adaptability to Changing Environments

The main focus of the designed experiments is to test the adaptability of our FairSAOML with respect to both fairness and model accuracy from one data domain to another. To better visualize changing environments, we manually add vertical dotted lines in Figure 3 to distinguish different domains at task indices of changing data domains. Our experimental results demonstrate that although FairSAOML may not dominate the performance over other baseline methods in the first data domain, it can effectively adapt to changes in the environment and so that the performance is constantly improving in both bias control and predictive accuracy.

In Sec. 4.2, we introduce that experts serve as key components in FairSAOML and model parameters pair (θ_t, λ_t) at round t are

determined by combining weighted expert advises (see Figure 2). Figure 3 (d,h,l) show weights change of all experts in FairSAOML. We observe that (1) expert-weight changes periodically; (2) experts associated with longer intervals are assigned with larger weights and these weights keep growing as the learner sees more tasks; (3) on the contrary, smaller weights are given to experts which carry short intervals and become less valued. No surprisingly, heavier weights on experts with long intervals enable our FairSAOML to adapt to the instability of model performance caused by changing environments.

As for baseline methods, MaskFTML outperforms other methods in accuracy in the first domain (see Figure 3 (c,g,k)), but it is not comparable with respect to model fairness. This indicates that attempting to make decision-makers blind to the protected attribute is not able to ameliorate prediction fairness. Although FairFML, FairAOGD, and FairGLC are able to control bias in the first domain, they fail to adapt fairness or predictive accuracy when data domain changes. Since AOD and CBCE are initially designed




Figure 6: Sensitive analysis on the MovieLens dataset.

for online learning in changing environments, they merely focus on learning accuracy but ignore control model fairness when data domain shifted. Besides, higher accuracy in AOD leads to a trade-off performance on fairness.

7.3 Ablation Studies and Sensitive Analysis

We conduct ablation studies on the *MovieLens* dataset to demonstrate the contributions of two key components in FairSAOML: expert weights $p_{t,l}$ and the base learner highlighted in Figure 2. Specifically, at each round, meta-level parameters are derived by combining weighted expert decisions. By removing expert weights, all experts will contribute equally. Furthermore, in active experts, base learners stated in Eq.(7) are used to further update model parameters at an interval level. Without base learners, all active experts share the same model parameters inherited from the previous round, and hence they are equally weighted. The key findings in Figure 5 are (1) expert weights play an important role in FairSAOML, and (2) base learners effectively enhance model performance on bias control and predictive accuracy.

Some sensitive analysis on the *MovieLens* dataset are shown in Figure 6 where intervals are considered using different bases switched from 2 to 5. According to Eq.(5), the setting with the smallest base (*i.e.* 2) indicates the most experts (*i.e.* 6) and hence its largest expert carries the longest intervals (*i.e.* 32). We observe that in terms of model fairness, settings with smaller base slightly outperform the ones with larger base in the first domain, but opposite findings are given in the last domain. This is because (1) in the first domain, the largest experts carry more information in smaller base setting than larger base setting; (2) while in the last domain, the largest experts in smaller base settings become impure that take data across two domains. More results are given in Appendix A.4.

8 CONCLUSION

To address fairness-aware online learning in changing environments where the learner is faced with data tasks sampled from heterogeneous distributions one after another, we first introduce a novel regret FairSAR which extends strongly adaptive regret by adding fairness-aware long-term constraints. We claim that for the first time a fairness-aware online meta-learning in changing environments framework is proposed. Then, to determine model parameters at each round, we propose a novel learning algorithm, namely FairSAOML. In this algorithm, we dynamically activate a subset of experts at each round and update their parameters at an interval-level. Model parameters at the meta-level are further achieved by combining weighted actions of all experts. Detailed theoretic analysis and corresponding proofs justify the efficiency and effectiveness of the proposed algorithm by demonstrating upper bounds for the loss regret and violation of fair constraints. Empirical studies based on real-world datasets show that our method

outperforms state-of-the-art online learning techniques in both model accuracy and fairness.

ACKNOWLEDGMENTS

The research reported herein was supported by the National Science Foundation (NSF) under grant number OAC-1931541, 2147375, 2137335, IIS-1954376, IIS-1815696, Army Research Office Contract W911NF2110032, and IBM faculty award (Research).

REFERENCES

- [1] Dan Biddle. 2005. Adverse Impact and Test Validation: A Practitioner's Guide to Valid and Defensible Employment Testing. Gower (2005).
- [2] Nicolò Cesa-Bianchi and Gábor Lugosi. 2006. Prediction, Learning, and Games. Cambridge University Press (2006).
- [3] Amit Daniely, Alon Gonen, and Shai Shalev-Shwartz. 2015. Strongly Adaptive Online Learning. In *ICML*.
- [4] Amit Daniely, Alon Gonen, and Shai Shalev-Shwartz. 2015. Strongly Adaptive Online Learning. In *ICML*.
- [5] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Rich Zemel. 2011. Fairness Through Awareness. *CoRR* (2011).
- [6] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. *ICML* (2017).
- [7] Chelsea Finn, Aravind Rajeswaran, Sham Kakade, and Sergey Levine. 2019. Online Meta-Learning. *ICML* (2019).
- [8] András György, Tamás Linder, and Gábor Lugosi. 2012. Efficient tracking of large classes of experts. *IEEE Transactions on Information Theory* (2012).
- [9] Moritz Hardt, Eric Price, and Nathan Srebro. 2016. Equality of opportunity in supervised learning. *NeurIPS* (2016).
- [10] Elad Hazan and Edgar Minasyan. 2020. Faster Projection-free Online Learning. In *Proceedings of Thirty Third Conference on Learning Theory*.
- [11] Elad Hazan and C. Seshadhri. 2007. Adaptive Algorithms for Online Decision Problems. *Electronic Colloquium on Computational Complexity (ECCC)* (2007).
- [12] Rodolphe Jenatton, Jim Huang, and Cedric Archambeau. 2016. Adaptive Algorithms for Online Convex Optimization with Long-term Constraints. *ICML* (2016).
- [13] Kwang-Sung Jun, Francesco Orabona, Stephen Wright, and Rebecca Willett. 2017. Improved Strongly Adaptive Online Learning using Coin Betting. In *AISTATS*.
- [14] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, Tony Lee, Etienne David, Ian Stavness, Wei Guo, Berton Earnshaw, Imran Haque, Sara M Beery, Jure Leskovec, Anshul Kundaje, Emma Pierson, Sergey Levine, Chelsea Finn, and Percy Liang. 2021. WILDS: A Benchmark of in-the-Wild Distribution Shifts. In *ICML*.
- [15] Michael Lohaus, Michael Perrot, and Ulrike Von Luxburg. 2020. Too Relaxed to Be Fair. In *ICML*.
- [16] Haipeng Luo and Robert E Schapire. 2015. Achieving all with no parameters: Adanormalhedge. In *Conference on Learning Theory*. PMLR, 1286–1304.
- [17] Mehrdad Mahdavi, Rong Jin, and Tianbao Yang. 2012. Trading regret for efficiency: online convex optimization with long term constraints. *JMLR* (2012).
- [18] Jennifer Miller. 2020. Is an Algorithm Less Racist Than a Loan Officer? www.nytimes.com/2020/09/18/business/digital-mortgages.html (2020).
- [19] Yuanyu Wan, Bo Xue, and Lijun Zhang. 2021. Projection-free Online Learning in Dynamic Environments. *AAAI* (2021).
- [20] Zhuoyi Wang, Yuqiao Chen, Chen Zhao, Yu Lin, Xujiang Zhao, Hemeng Tao, Yigong Wang, and Latifur Khan. 2021. CLEAR: Contrastive-Prototype Learning with Drift Estimation for Resource Constrained Stream Mining. In *WWW*.
- [21] Yongkai Wu, Lu Zhang, and Xintao Wu. 2019. On Convexity and Bounds of Fairness-aware Classification. *WWW*.
- [22] Jiahao Xie, Zebang Shen, Chao Zhang, Boyu Wang, and Hui Qian. 2020. Efficient projection-free online methods with stochastic recursive gradient. In *AAAI*.
- [23] Jianjun Yuan and Andrew Lamperski. 2018. Online convex optimization for cumulative constraints. *NeurIPS* (2018).
- [24] Lijun Zhang, Shiyin Lu, and Tianbao Yang. 2020. Minimizing Dynamic Regret and Adaptive Regret Simultaneously. *AISTATS* (2020).
- [25] Lijun Zhang, Shiyin Lu, and Zhi-Hua Zhou. 2018. Adaptive Online Learning in Dynamic Environments. In *International Conference on Neural Information Processing Systems*. *NeurIPS* 2018.
- [26] Chen Zhao and Feng Chen. 2019. Rank-Based Multi-task Learning For Fair Regression. *IEEE International Conference on Data Mining (ICDM)* (2019).
- [27] Chen Zhao and Feng Chen. 2020. Unfairness Discovery and Prevention For Few-Shot Regression. *ICKG* (2020).
- [28] Chen Zhao, Feng Chen, and Bhavani Thuraisingham. 2021. Fairness-Aware Online Meta-learning. *ACM SIGKDD* (2021).
- [29] Martin Zinkevich. 2003. Online Convex Programming and Generalized Infinitesimal Gradient Ascent. *ICML* (2003).

A ADDITIONAL EXPERIMENT DETAILS

A.1 Notations

Table 2: Important notations and corresponding descriptions.

Notations	Descriptions
T	Total number of learning tasks
t	Indices of tasks
τ	Length of time intervals in general
$\mathcal{D}_t^S, \mathcal{D}_t^V, \mathcal{D}_t^Q$	Support/Validation/Query set of data \mathcal{D}_t
θ_t, λ_t	Meta-level primal/dual parameters at round t
$\theta_{t,I}, \lambda_{t,I}$	Interval-level primal/dual parameters for an expert E_I at round t
$f_t(\cdot)$	Loss function at round t
$g_i(\cdot)$	Fairness function
m	Total number of fairness notions
i	Indices of fairness notions
$\mathcal{G}(\cdot)$	Base learner
\mathcal{U}	Expert set
$\mathcal{A}_t, \mathcal{S}_t$	Active/Sleeping expert set at round t
\mathcal{I}	AGC interval set
C_t	Target set of intervals at round t
\mathcal{B}	Relaxed primal domain
$\prod_{\mathcal{B}}$	Projection operation onto domain \mathcal{B}
η_1, η_2	Learning rates
$p_{t,I}$	Expert weight of E_I at round t
δ	Augmented constant

A.2 Data Pre-Processing

The *New York Stop-and-Frisk* dataset [14] is provided by the New York Police Department and contains information about the stop, question, and frisk policy implemented by the NYPD. The full dataset contains records of this policy dating from 2003 to 2019. It includes a record for each stop made, with information about the time of the stop, the location of the stop, information about the officer, the suspect, and other various features. The data used in this paper is from the year 2011, which saw the highest number of stops of any year. This data is used to attempt to predict whether someone will get frisked based on the circumstances of the stop and the features of the individual suspect. Specifically, we extract relevant features (predictor and response) that are utilized for analysis and remove those that may result in data leakage. Missing entries are imputed using precinct (*pct*) data and further we remove the *pct* column due to difficulty in binning them into separate precincts. Data records containing unknown or inconsistent values are dropped.

The *MovieLens* dataset was collected through the MovieLens web site (movielens.umn.edu) during the seven-month period from September 19th, 1997 through April 22nd, 1998. This data has been cleaned up - users who had less than 20 ratings or did not have complete demographic information were removed from this data set. We first merge the main three spreadsheets by its primary keys: *user-id* and *item-id*. Then some unused features are removed including *user-id*, *movie-title*, *release-date*, *video-release-date*, *IMDb-URL* and *zip-code*.

A.3 Implementation Details and Hyperparameter Tuning

Our neural network trained follows the same architecture used in [6], which contains 2 hidden layers of size of 40 with ReLU activation functions. In the training process of the *MovieLens* (*New York Stop-and-Frisk*) data, each gradient is computed using a batch size of 200 (800) examples where each binary class contains 100 (400) examples. For each dataset, we tune the following hyperparameters: (1) the initial dual meta parameter λ_0 is chosen from $\{0.00001, 0.0001, 0.001, 0.01, 0.1, 1, 10, 100, 1000, 10000\}$; (2) the interval-level gradient steps are chosen from 1 to 10; (3) the number of iterations N_{meta} are chosen from $\{20, 25, 30, 35, 40, 45, 50, 55, 60, 65, 70, 75, 80, 85, 90, 95, 100\}$; (4) learning rates η_1 and η_2 for updating meta-level parameters in Eq.(14) and (15) are chosen from $\{0.0001, 0.0005, 0.001, 0.005, 0.01, 0.05, 0.1, 0.5, 1, 5, 10, 50, 100, 500, 1000\}$; (5) the positive constant δ used in the augmented term are chosen from $\{10, 25, 50, 75, 100\}$.

A.4 Additional Results

Due to space limit, more experimental results on the *New York Stop-and-Frisk* dataset are given in Figure 7 and 8.

B PROOF SKETCH OF THEOREM 1

Under Assumptions 1, 2 and 3, we first target Eq.(13) and have

LEMMA 2 (THEOREM 1 IN [28]). *Suppose f and $g : \Theta \times \mathbb{R}_+^m \rightarrow \mathbb{R}$ satisfy Assumptions 1, 2 and 3. The interval-level update and the augmented Lagrangian function $\mathcal{L}_t(\theta, \lambda)$ are defined in Eq.(11)(12) and Eq.(13). Then, the function $\mathcal{L}_t(\theta, \lambda)$ is convex-concave with respect to the arguments θ and λ , respectively. Furthermore, as for $\mathcal{L}_t(\cdot, \lambda)$, if stepsize $\eta_{t,I}$ for each active expert E_I is selected as $\eta_{t,I} \leq \min\{\frac{\mu_f + \lambda m \mu_g}{8(L_f + \lambda m L_g)(\rho_f + \lambda m \rho_g)}, \frac{1}{2(\beta_f + \lambda m \beta_g)}\}$, then $\mathcal{L}_t(\cdot, \lambda)$ enjoys $\frac{2}{8}(\beta_f + \lambda m \beta_g)$ -smooth and $\frac{1}{8}(\mu_f + \lambda m \mu_g)$ -strongly convex, where $\bar{\lambda} \geq 0$ is the mean value of λ .*

According to Theorems 1 and 3 in [16] and the Lemma 1 in [24], we have the following lemma with respect to Eq.(13) that

LEMMA 3. *Under Assumption 2, for any interval $I = [i, j] \in \mathcal{I}$, FairSAOML satisfies*

$$\sum_{u=i}^t \mathcal{L}_u(\theta_u, \lambda_u) - \sum_{u=i}^t \mathcal{L}_u(\theta_{u,I}, \lambda_{u,I}) \leq S \sqrt{6L_f L_g (t-i-1)c(t)}$$

where $c(t) \leq 1 + \ln t + \ln(1 + \log_2^T) + \ln \frac{5+3 \ln(1+t)}{2}$.

By applying the Lemma 3 with the Theorem 2 in [24], we have

LEMMA 4. *Under Assumption 1 and 2, for any interval $I = [i, j] \in \mathcal{I}$, for any $(\theta, \lambda) \in \Theta \times \mathbb{R}_+^m$ FairSAOML satisfies*

$$\sum_{t \in I} \mathcal{L}_t(\mathcal{G}_t(\theta_t), \lambda) - \sum_{t \in I} \mathcal{L}_t(\theta, \lambda_{t,I}) \leq S \sqrt{|I|(\sqrt{6L_f L_g c(t)} + G)}$$

To extend our Lemma 4 to any interval $I = [r, s] \subseteq [T]$, we refer the following lemma

LEMMA 5 (LEMMA 3 IN [24]). *For any interval $[r, s] \subseteq [T]$, it can be partitioned into two sequences of disjoint and consecutive intervals,*




Figure 7: Running time comparison and ablation studies on the New York Stop-and-Frisk (B→M→S) dataset.




Figure 8: Running time comparison and ablation studies on the New York Stop-and-Frisk (R→Q→S) dataset.

denoted by $I_{-p}, \dots, I_0 \in \mathcal{I}$ and $I_1, \dots, I_q \in \mathcal{I}$, such that

$$|I_{-i}|/|I_{-i-1}| \leq 1/2, \forall i \geq 1 \quad \text{and} \quad |I_i|/|I_{i-1}| \leq 1/2, \forall i \geq 2$$

Next, we prove Theorem 1.

PROOF. By applying Lemma 4 onto Lemma 5 and set θ^* being the optimal solution for $\min_{\theta \in \Theta} \sum_{t=r}^s f_t(\mathcal{G}_t(\theta))$ where $[r, s] \subseteq [T]$, we have

$$\begin{aligned} & \sum_{t=r}^s \mathcal{L}_t(\mathcal{G}_t(\theta_t), \lambda) - \sum_{t=r}^s \mathcal{L}_t(\mathcal{G}_t(\theta^*), \lambda_{t,I}) \\ &= \sum_{i=-p}^q \left(\sum_{t \in I_i} \mathcal{L}_t(\mathcal{G}_t(\theta_t), \lambda) - \sum_{t \in I_i} \mathcal{L}_t(\mathcal{G}_t(\theta^*), \lambda_{t,I}) \right) \\ &\leq \sum_{i=-p}^q S\sqrt{|I_i|}(\sqrt{6L_f L_g c(s)} + G) \\ &\leq 2S(\sqrt{6L_f L_g c(s)} + G) \sum_{i=0}^{\infty} \sqrt{2^{-i}|I|} \\ &\leq 8S(\sqrt{6L_f L_g c(s)} + G)\sqrt{|I|} \end{aligned} \quad (16)$$

By expanding Eq.(16) using Eq.(13) and following the Theorem 3.1 in [2], we have

$$\begin{aligned} & \sum_{t=r}^s \left\{ f_t(\mathcal{G}_t(\theta_t)) - f_t(\mathcal{G}_t(\theta^*)) \right\} \\ &+ \sum_{i=1}^m \left\{ \lambda_i \sum_{t=r}^s g_i(\mathcal{G}_t(\theta_t)) - \sum_{t=r}^s \lambda_{t,i} g_i(\mathcal{G}_t(\theta^*)) \right\} \\ &- \frac{\delta(\eta_1 + \eta_2)(s-r+1)}{2} \|\lambda\|^2 + \frac{\delta(\eta_1 + \eta_2)}{2} \sum_{t=r}^s \|\lambda\|^2 \\ &\leq 8S(\sqrt{6L_f L_g c(s)} + G)\sqrt{|I|} \end{aligned}$$

Here, we approximately average $p_{t,I}$ for all experts E_I in \mathcal{U} and hence the subscription I is omitted. Inspired by the proof of Theorem 4 in [17], we take maximization for λ over $(0, +\infty)$ and get

$$\begin{aligned} & \sum_{t=r}^s \left\{ f_t(\mathcal{G}_t(\theta_t)) - f_t(\mathcal{G}_t(\theta^*)) \right\} \\ &+ \sum_{i=1}^m \left\{ \frac{[\sum_{t=r}^s g_i(\mathcal{G}_t(\theta_t))]^2}{2(\delta(\eta_1 + \eta_2)(s-r+1) + \frac{m}{\eta_1 + \eta_2})} - \sum_{t=r}^s \lambda_{t,i} g_i(\mathcal{G}_t(\theta^*)) \right\} \\ &\leq 8S(\sqrt{6L_f L_g c(s)} + G)\sqrt{|I|} \end{aligned}$$

Since $g_i(\mathcal{G}_t(\theta^*)) \leq 0$ and $\lambda_{t,i} \geq 0, \forall i \in [m]$, the resulting inequality becomes

$$\begin{aligned} & \sum_{t=r}^s \left\{ f_t(\mathcal{G}_t(\theta_t)) - f_t(\mathcal{G}_t(\theta^*)) \right\} \\ &+ \sum_{i=1}^m \frac{[\sum_{t=r}^s g_i(\mathcal{G}_t(\theta_t))]^2}{2(\delta(\eta_1 + \eta_2)(s-r+1) + \frac{m}{\eta_1 + \eta_2})} \leq 8S(\sqrt{6L_f L_g c(s)} + G)\sqrt{|I|} \end{aligned}$$

Due to non-negative of $\frac{[\sum_{t=r}^s g_i(\mathcal{G}_t(\theta_t))]^2}{2(\delta(\eta_1 + \eta_2)(s-r+1) + \frac{m}{\eta_1 + \eta_2})}$, we have

$$\begin{aligned} & \sum_{t=r}^s \left\{ f_t(\mathcal{G}_t(\theta_t)) - f_t(\mathcal{G}_t(\theta^*)) \right\} \leq 8S(\sqrt{6L_f L_g c(s)} + G)\sqrt{|I|} \\ &= O((|I| \log s)^{1/2}) \end{aligned}$$

Furthermore, we have $\sum_{t=r}^s \left\{ f_t(\mathcal{G}_t(\theta_t)) - f_t(\mathcal{G}_t(\theta^*)) \right\} \geq -F(s-r+1)$ according to the assumption and set $\eta_1 = \eta_2 = O(1/\sqrt{s})$. We have

$$\sum_{t=r}^s g_i(\mathcal{G}_t(\theta_t)) \leq O((|I|s \log s)^{1/4})$$

Therefore, as for FairSAR proposed in Eq.(4) we complete the proof.

$$\begin{aligned} & \max_{[s, s+\tau-1] \subseteq [T]} \left(\sum_{t=s}^{s+\tau-1} f_t(\mathcal{G}_t(\theta_t)) - f_t(\mathcal{G}_t(\theta^*)) \right) \leq O((\tau \log T)^{1/2}) \\ & \max_{[s, s+\tau-1] \subseteq [T]} \left(\sum_{t=s}^{s+\tau-1} g_i(\mathcal{G}_t(\theta_t)) \right) \leq O((\tau \log T)^{1/4}), \quad \forall i \in [m] \end{aligned}$$

□