LOCALLY ADAPTIVE MULTI-OBJECTIVE LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

We consider the general problem of learning a predictor that satisfies multiple objectives of interest simultaneously. We work in an online setting where the data distribution can change arbitrarily over time. Here, multi-objective learning captures many common targets such as online calibration, regret, and multiaccuracy. In the online setting, existing approaches to this problem that minimize the set of objectives over the entire time horizon can fail to adapt to distribution shifts. We correct this and propose algorithms that guarantee small error for all objectives over any local time interval of a given width. Empirical evaluations on datasets from energy forecasting and algorithmic fairness show that our methods can be used to guarantee unbiasedness of the predictions over subgroups of concern and ensure robustness under distribution shift.

1 INTRODUCTION

In an ever-changing world, real-time decision making necessitates coping with arbitrary distribution shifts and adversarial behavior. These shifts can arise from seasonality, change in data distribution induced by feedback loops or policy changes, and exogenous shocks such as pandemics or economic crises. Online learning is a powerful framework for analyzing sequential data that makes no assumptions on the data distribution.

Multi-objective learning is a generic framework that refers to any task in which a predictor must satisfy multiple objectives or criterion of interest simultaneously. This general framework has led to the development of online learning algorithms for numerous applications including multicalibration (Hebert-Johnson et al., 2018), multivalid conformal prediction (Gupta et al., 2022), and multi-group learning (Deng et al., 2024). Despite being a desirable and promising notion, methods from the online multi-objective learning literature have had little influence on the practice of machine learning.

We attribute this to two shortcomings. First, the majority of the algorithms proposed in the literature are not adaptive to abrupt changes in the data distribution: they learn a predictor that minimizes the objectives over the *entire time horizon*. In changing environments and in the presence of adversarial behavior, such algorithms will fail to cope with distribution shifts. Second, most prior work is purely theoretical with scant empirical evaluation. As a result, the practical aspects of multi-objective online algorithms have received limited consideration.

In this work, we aim to overcome the above shortcomings. We propose a locally adaptive multi-objective learning algorithm that outputs predictors which (approximately) satisfy a set of objectives over all local time intervals $I \subseteq [T]$. As we discussed above, multi-objective learning can be used to address many common prediction tasks. As a case study, in this work, we focus on the multiaccuracy problem in which the goal is to learn predictiors which are unbiased under covariate shift. We seek a small multiaccuracy error while preserving accuracy relative to a given sequence of baseline predictions. This is a problem of significant and broad interest across real-time decision-making and deployed machine learning systems.

To close the empirical gap in this literature, we perform experiments on electricity demand forecasting and predicting recidivism over time. We show that our proposed algorithm has low multiaccuracy error over all intervals while the baselines have poor adaptivity. An alternative objective to multiaccuracy that is popular in the literature is multicalibration (Haghtalab et al., 2023a; Garg et al., 2024). Despite being a stronger condition, we show that in practice existing online multicalibration algorithms only achieve multiaccuracy at relatively slow rates. While adaptive extensions of online multicalibration

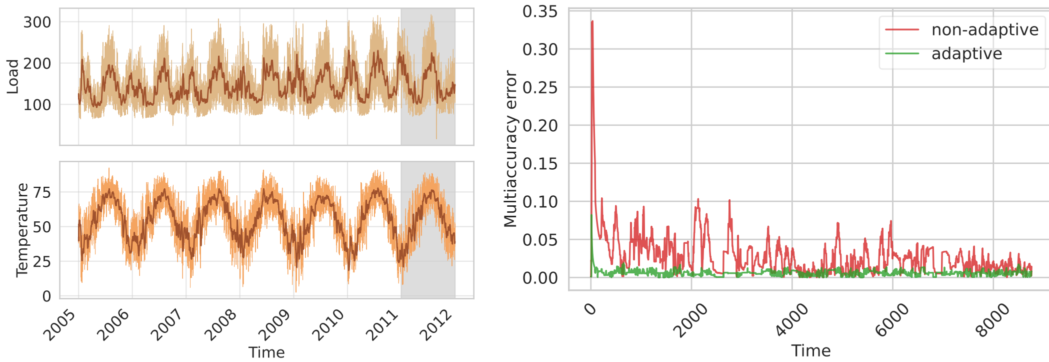

(a) (b)

Figure 1: GEFCom14-L electric load forecasting dataset. On the left hand side are the time series for load and temperature. The dark brown curves denote the 2-week moving average. The shaded region shows the competition duration. On the right-hand side, we plot local multiaccuracy error.

algorithms have been discussed (Lee et al., 2022), we show that they are less efficient and effective. We will release a codebase that implements our algorithm and all the baselines used in the paper.

We note that although we focus on multiaccuracy in this paper, our general algorithm extends to other multi-objective learning problems including multi-group learning (Tosh & Hsu, 2022) and omniprediction (Gopalan et al., 2022). We discuss these extensions toward the end.

1.1 Peek at results

To demonstrate the significance of local adaptivity in practice, we consider the probabilistic electricity load forecasting track of the Global Energy Forecasting Competition 2014 (GEFCom2014) (Hong et al., 2016). The aim in the load forecasting track GEFCom2014-L is to forecast month-ahead quantiles of hourly loads for a US utility from 01/2011 to 12/2011 based on historical load and temperature data (Figure 1a).

We consider the binary task of predicting whether the electricity demand exceeds 150MW at time step t and evaluate whether the predictions are multiaccurate with respect to discrete temperature groups $\{[0, 20), [20, 40), \ldots, [80, 100)\}$ (in °F). Informally, obtaining mutliaccuracy with respect to temperature ensures our predictions are accurate at any time-of-day and during all seasons. Figure 1b shows the multiaccuracy error of our proposed locally adaptive algorithm compared to a non-adaptive multiaccuracy algorithm, plotted as a weekly (168-hourly) rolling average. We can see that the multiacuracy error of the adaptive algorithm is close to zero across all time intervals, while the non-adaptive variant has high variance.

1.2 Related work

Our work is most closely related to the literature on multicalibration (Hebert-Johnson et al., 2018), multiaccuracy (Kim et al., 2019), and omniprediction (Gopalan et al., 2022). Each of these multiobjective criteria have been studied in both the online and batch settings. Most closely related to our work, Kim et al. (2019) and Globus-Harris et al. (2023) give algorithms for obtaining multi-accurate and multi-calibrated (resp.) predictors in the batch setting that are guaranteed to have accuracy no worse than that of a given base predictor.

In the online advesarial setting, a number of works develop algorithms for obtaining multiaccuraty, multicalibration, and/or omniprediction globally over all time steps (Lee et al., 2022; Garg et al., 2024; Okoroafor et al., 2025; Haghtalab et al., 2023a; Noarov et al., 2025). Our work will in particular build on the algorithmic framework developed in (Lee et al., 2022). This methodology has deep roots in the online learning literature and builds on ideas in blackwell approachability (Blackwell, 1956) and its connection to no-regret learning Abernethy et al. (2011).

To obtain time-local guarantees we will draw on the literature on adaptive and strongly-adaptive regret (Herbster & Warmuth, 1998; Daniely et al., 2015; Jun et al., 2017; Haghtalab et al., 2023b). Our work will most closely rely upon the work of Gradu et al. (2023) to obtain multiobjective error bounds over any local time interval. In the context of multiobjective learning, local guarantees have been discussed previously in Lee et al. (2022). However, the literature contains no empirical evaluations of these methods. We provide experiments evaluating the algorithms of Lee et al. (2022) in Section 4.3 and find that our approach achieves significantly lower error rates in practice.

1.3 Preliminaries

We use $\mathcal{X}$ to denote our feature space and $\mathcal{Y} = [a, b]$ to denote our label space, which we assume to be bounded and convex. Our goal is to learn a sequence of predictors $p_t \in \mathcal{Y}, t = 1, 2, \dots, T$ that guarantee loss minimization simultaneously for every objective within a set $\mathcal{L}$ over time. Each objective, or criterion, is a function $\ell : [0, 1] \times \mathcal{X} \times \mathcal{Y} \to [-1, 1]$ that takes as input a predictor p_t, features $x_t \in \mathcal{X}$, and label $y_t \in \mathcal{Y}$ and returns an objective value. We assume the loss is bounded in $[-1, 1]$. We will use $[T]$ to denote the set $\{1, 2, \dots, T\}$.

Multi-objective learning aims to learn a predictor that simultaneously minimizes every objective in $\mathcal{L}$. We study the problem of multi-objective learning in an online, adversarial setting.

Definition 1 (Online multi-objective learning)**.** For a set of objectives $\mathcal{L} = \{\ell : [0, 1] \times \mathcal{X} \times \mathcal{Y} \to [0, 1]\}$ and sequence of data points $x_t, y_t, t \in [T]$, the online multi-objective problem is to learn a sequence of predictors p_t that minimize

$$\max_{\ell \in \mathcal{L}} \frac{1}{T} \sum_{t=1}^{T} \ell(p_t(x_t), x_t, y_t),$$

where (x_t, y_t) can be generated advesarially dependent on the entire history of data and predictions up to time t.

Next, we define two instantiations of multi-objective problems that are commonly studied in the literature—multiaccuracy and multicalibration. We parameterize the multiaccuracy criterion by a function class $\mathcal{F}$ and the goal is to be unbiased for all $f \in \mathcal{F}$, i.e., there is no systematic correlation between the prediction residuals and any $f \in \mathcal{F}$.

Definition 2 (Online multiaccuracy)**.** Let $\mathcal{F} = \{f : \mathcal{X} \to [0, 1]\}$ be a class of functions on $\mathcal{X}$. In online multiaccuracy, we instantiate $\ell_{\mathrm{MA}_{f,\sigma}}(p_t(x_t), x_t, y_t) = \sigma f(x_t) \cdot (y_t - p_t(x_t))$ for every sign $\sigma = \{\pm\}$ and $f \in \mathcal{F}$ and define the multiaccuracy error ℓ_{MA} in the ℓ_∞ norm as

$$\ell_{\mathrm{MA}} = \sup_{f \in \mathcal{F}, \sigma \in \{\pm\}} \frac{1}{T} \sum_{t=1}^{T} \sigma f(x_t) \cdot (y_t - p_t(x_t)). \tag{1}$$

Another popular online prediction target is calibration. In a binary classification task, calibration asks that among instances with predicted probability v, a fraction v of them are observed to be truly labeled as 1. To implement this in practice, we discretize the label interval $[0, 1]$ into m bins $V_m := \{[0, 1/m), [1/m, 2/m), \dots\}$ and ask for calibration in each bin $v \in V_m$. Multicalibration is a stronger and fine-grained notion that requires a predictor to be calibrated under all reweightings $f \in \mathcal{F}$. Standard calibration is then the special case where $\mathcal{F} = \{x \mapsto 1\}$ is a singleton function class containing only the identity.

Definition 3 (Online multicalibration)**.** Fix a set of functions $\mathcal{F}$ and $m \geq 1$. In online multicalibration we instantiate $\ell_{\mathrm{MC}_{f,\sigma,v}}(p_t(x_t), x_t, y_t) = \sigma f(x_t) \cdot \mathbb{1}\{p_t(x_t) \in v\} \cdot (y_t - p_t(x_t))$ for every sign $\sigma = \{\pm\}$, $f \in \mathcal{F}$, and $v \in V_m$ and define the multicalibration error ℓ_{MC} in the ℓ_∞ norm as

$$\ell_{\mathrm{MC}} = \sup_{f \in \mathcal{F}, \sigma \in \{\pm\}, v \in V_m} \frac{1}{T} \sum_{t=1}^{T} \sigma f(x_t) \cdot \mathbb{1}\{p_t(x_t) \in v\} \cdot (y_t - p_t(x_t)). \tag{2}$$

A direct calculation shows that the online multicalibration error always upperbounds the multiaccuracy error.

In this work, we give a multi-objective learning algorithm that achieves small multiaccuracy error while preserving accuracy relative to a base predictor sequence $\tilde{p}_t, t \in [T]$. We define the latter accuracy objective as regret

$$\ell_{\text{reg}} := \frac{1}{T} \sum_{t=1}^{T} \ell_{\text{acc}}(p_t(x_t), y_t) - \ell_{\text{acc}}(\tilde{p}_t(x_t), y_t), \tag{3}$$

where $\ell_{\text{acc}} \geq 0$ is any proper loss for the mean (i.e. any loss such that $\mathbb{E}_{y\sim P}[y] \in \text{argmin}_p \mathbb{E}_{y\sim P}[\ell(p, y)]$ for all distributions P on $\mathcal{Y}$.) (e.g., log loss or squared error/Brier score).

2 Methods

2.1 Online multi-objective learning

The online multi-objective learning problem can be viewed as a sequential prediction task over T rounds. A common approach is to imagine a two-player game between a learner, who observes $x_t \in \mathcal{X}$ and chooses a predictor p_t, and an adversary who maintains a probability distribution $q^{(t)} \in \Delta(\mathcal{L})$ over the mixture of objectives $\mathcal{L}$. After the learner acts, the label y_t is revealed and the realized objective values $\{\ell(p_t(x_t, x_t, y_t) : \ell \in \mathcal{L}\}$ are observed. The goal of the learner is to simultaneously minimize every objective in $\mathcal{L}$. In other words, the learner's goal is to minimize the maximum accumulated objective value after T rounds:

$$\max_{\ell \in \mathcal{L}} \sum_{t=1}^{T} \ell(p_t(x_t), x_t, y_t). \tag{4}$$

In parallel, the adversary updates its distribution $q^{(t)}$ to assign larger probability to objectives that incurred larger value in round t.

At its core, our algorithm is an exponential reweighting algorithm (Freund & Schapire, 1997). We uniformly initialize probabilities $q_\ell^{(1)}$ for each objective $\ell \in \mathcal{L}$ and update $q_\ell^{(t)}$ using multiplicative weights. We fix a function class $\mathcal{F}_{\text{MA}} \subseteq \{f : \mathcal{X} \to [0, 1]\}$ that we desire multiaccuracy with respect to and define $\mathcal{L} := \{\ell_{\text{MA}_{f,\sigma}} : f \in \mathcal{F}_{\text{MA}}, \sigma \in \{\pm\}\} \cup \{\ell_{\text{reg}}\}$ including the regret objective. Thus, we maintain $q^{(t)}$ over $2|\mathcal{F}_{\text{MA}}| + 1$ objectives with $\sum_{f,\sigma} q^{(t)}_{\text{MA}_{f,\sigma}} + q^{(t)}_{\text{reg}} = 1$. This algorithmic technique has been previously employed for multicalibration objectives (Lee et al., 2022; Haghtalab et al., 2023a); however, our empirical evaluation shows that it does not yield an effective online multiaccuracy algorithm.

The learner best-responds to the adversary by choosing a predictor p_t that minimizes the expected objective value over the mixture $q^{(t)}$ for the *worst case* label: $\underset{p_t}{\text{argmin}} \max_{y \in \mathcal{Y}} \sum_\ell q_\ell^{(t)} \ell(p_t(x_t), x_t, y_t)$. The key idea to upper bound (4) using this minimax optimal strategy was proposed in Lee et al. (2022). The algorithms we present in this work build on this idea. We show in Section 3 that this strategy yields predictors with low multiaccuracy error and regret.

2.2 Locally adaptive multi-objective learning

The minimization goal defined in (4) ceases to be useful when environments are changing and the data distribution shifts arbitrarily over time. As a simple example, fix the singleton function class $\mathcal{F}_{\text{MA}} = \{x \mapsto t\}$ and consider targeting just the multiaccuracy error (i.e., set $\mathcal{L} = \{\ell_{\text{MA}_{f,\sigma}} : f \in \mathcal{F}_{\text{MA}}, \sigma \in \{\pm\}\}$). Let the labels be given as $y_t = 1$ for the first $T/2$ rounds and $y_t = 0$ for the last $T/2$ rounds. Here, the constant predictor $p_t = 1/2$ minimizes the multiaccuracy error in (1). Nevertheless, this predictor performs poorly in the individual intervals $1 \leq t \leq T/2$ and $t > T/2$ compared to the optimal predictor that switches from $p_t = 1$ to $p_t = 0$ after $t = T/2$.

To account for distribution shifts in changing environments, we will now strengthen out objective by requiring the learner to choose the "locally best" predictor in all intervals $I = [r, s] \subset [T]$. The locally adaptive multi-objective learning problem is then to learn a sequence of predictors p_t that

Algorithm 1 Locally adaptive multiaccurate mean estimation

Input: Function class $\mathcal{F}_{\text{MA}} \subseteq \{f : \mathcal{X} \to [0,1]\}$; base predictor sequence $\tilde{p}_t, t \in [T]$; hyperparameters η, γ.
Input: Sequence of samples $\{(x_1, y_1), \ldots, (x_T, y_T)\}$
1: $q^{(1)}_{\text{MA}_{f,\sigma}} = \frac{1}{2|\mathcal{F}_{\text{MA}}|+1}, \quad \forall f \in \mathcal{F}_{\text{MA}}, \sigma \in \{\pm 1\}$.
2: $q^{(1)}_{\text{reg}} = \frac{1}{2|\mathcal{F}_{\text{MA}}|+1}$
3: **for** each $t \in [T]$ **do**
4: $\quad p_t(x_t) := \underset{p}{\operatorname{argmin}} \max_{y \in \mathcal{Y}} \sum_{f,\sigma} q^{(t)}_{\text{MA}_{f,\sigma}} \sigma f(x_t)(y - p(x_t)) + q^{(t)}_{\text{reg}}(\ell_{\text{acc}}(p(x_t), y) - \ell_{\text{acc}}(\tilde{p}_t(x_t), y))$
5: $\quad \tilde{q}^{(t+1)}_{\text{MA}_{f,\sigma}} \propto q^{(t)}_{\text{MA}_{f,\sigma}} \exp\left(\eta \cdot \sigma f(x_t)(y_t - p_t(x_t))\right)$ for all $f \in \mathcal{F}_{\text{MA}}, \sigma \in \{\pm 1\}$
6: $\quad \tilde{q}^{(t+1)}_{\text{reg}} \propto q^{(t)}_{\text{reg}} \exp\left(\eta \cdot (\ell_{\text{acc}}(p_t(x_t), y_t) - \ell_{\text{acc}}(\tilde{p}_t(x_t), y_t))\right)$
7: $\quad q^{(t+1)}_{\text{MA}_{f,\sigma}} = (1-\gamma)\, \tilde{q}^{(t+1)}_{\text{MA}_{f,\sigma}} + \frac{\gamma}{2|\mathcal{F}_{\text{MA}}|+1}$
8: $\quad q^{(t+1)}_{\text{reg}} = (1-\gamma)\, \tilde{q}^{(t+1)}_{\text{reg}} + \frac{\gamma}{2|\mathcal{F}_{\text{MA}}|+1}$
Output: Sequence of (randomized) predictors $p_1, \ldots, p_T$

minimize

$$\sup_{I=[r,s]\subseteq[T]} \left[\max_{\ell \in \mathcal{L}} \sum_{t=r}^{s} \ell(p_t(x_t), x_t, y_t) \right]. \tag{5}$$

We develop a locally adaptive multiaccuracy algorithm (Algorithm 1) that minimizes (5). We lend our procedure adaptivity by modifying an algorithm proposed in the online control setting by Gradu et al. (2023), which is inspired by the Fixed Share algorithm (Herbster & Warmuth, 1998). At every update step, we add a uniform exploration term to the exponential weights update. This prevents any objective's weight from becoming too small and provably allows the learner to compete with the best predictor in every interval.

Strong adaptivity. We note that there exist online learning algorithms that guarantee *strongly adaptive regret* (Daniely et al., 2015), i.e., low error for intervals of any length $|I|$. This is a stronger guarantee compared to the guarantee over fixed interval widths and comes with a $\log T$ overhead in runtime and memory. It is straightforward to apply this to our algorithm—in fact, strong adaptivity extensions have been discussed for multicalibration (Lee et al., 2022; Haghtalab et al., 2023a). However, we find that they do not perform well in practice (Section 4.4).

2.3 Significance of the regret objective

In our applications, we will start with a base forecaster, $\tilde{p}_t$ that was constructed in advance for that application. Our goal will be to improve $\tilde{p}_t$ to be multiaccurate. While doing this, it is important that we do not degrade the accuracy of $\tilde{p}_t$, thereby rendering its predictions less useful. Our algorithm achieves small multiaccuracy error while preserving the accuracy relative to a base predictor by including an additional regret objective. In our empirical evaluation we will show that in the absence of this objective, a version of our method that targets multiaccuracy alone incurs high accuracy loss relative to the base predictors $\tilde{p}_t, t \in [T]$.

In general, even in the absence of a base predictor, it is not advisable to solely target multiaccuracy. Indeed, if we exclude the regret objective in Algorithm 1 one can show that the best response of the learner yields the predictor: $p_t = \mathbb{1}\{\sum_{f,\sigma} q^{(t)}_{\text{MA}_{f,\sigma}} \sigma f(x_t) > 0\}$. This solution has a pathological behavior where the predictor will only take values 0 or 1 at every step. This makes the predictions useless for real-time decision-making in an online setting. Our regret objective recovers the predictor from this problem by enforcing solutions that do not lie in the extremes. In practical settings where $\tilde{p}_t$ is not available in advance, we recommend combining our procedure with a standard online learning algorithm (e.g. online gradient or mirror descent) that provides an appropriate baseline (Algorithm 2).

3 THEORY

We now show that our proposed algorithm achieves small multiaccuracy error and regret over any local time interval of fixed width. From here on, we use shorthands $\ell^{(t)}_{\mathrm{MA}_{f,\sigma}} := \sigma f(x_t)(y_t - p_t(x_t))$ and $\ell^{(t)}_{\mathrm{reg}} := \ell_{\mathrm{acc}}(p_t(x_t), y_t) - \ell_{\mathrm{acc}}(\tilde{p}_t(x_t), y_t)$. $\ell^{(t)}_{\mathrm{MA}}$ and $q^{(t)}_{\mathrm{MA}}$ denote vectors of size $2|\mathcal{F}_{\mathrm{MA}}|$ that comprise, respectively, the objective values and probabilities for all multiaccuracy objectives. All proofs are deferred to Appendix A.

We first show that the maximum objective value over any time interval I is upper bounded by the cumulative expected value of the individual objectives.

Lemma 1. *Let $c = \sup_{\ell \in \{\ell^{(t)}_{\mathrm{MA}}, \ell^{(t)}_{\mathrm{reg}}\}} \frac{e^{\eta \ell}}{2}$. Assume that $\gamma \leq 1/2$. Then, for any interval $I = [r, s] \subseteq [T]$,*

$$\sum_{t=r}^{s} q^{(t)\top}_{\mathrm{MA}} \ell^{(t)}_{\mathrm{MA}} + q^{(t)}_{\mathrm{reg}} \ell^{(t)}_{\mathrm{reg}} \geq \max\left\{ \max_{f,\sigma} \sum_{t=r}^{s} \ell^{(t)}_{\mathrm{MA}_{f,\sigma}}, \sum_{t=r}^{s} \ell^{(t)}_{\mathrm{reg}} \right\} - c\eta \left(\sum_{t=r}^{s} q^{(t)\top}_{\mathrm{MA}} {\ell^{(t)}_{\mathrm{MA}}}^2 + q^{(t)}_{\mathrm{reg}} {\ell^{(t)}_{\mathrm{reg}}}^2 \right) - \frac{1}{\eta}\left(\log\left(\frac{2|\mathcal{F}_{\mathrm{MA}}| + 1}{\gamma} \right) + |I| 2\gamma \right), \tag{6}$$

where the data $(x_t, y_t), t \in [T]$ are fixed and the randomness is from $q^{(t)}_{\ell}, \ell \in \mathcal{L}$ in Algorithm 1.

Note that c is typically small as we rescale the objectives to be bounded by 1. Next, we show that the cumulative expected value of the objectives is non-positive over any interval I.

Lemma 2. *For any interval $I = [r, s] \subseteq [T]$,*

$$\sum_{t=r}^{s} q^{(t)\top}_{\mathrm{MA}} \ell^{(t)}_{\mathrm{MA}} + q^{(t)}_{\mathrm{reg}} \ell^{(t)}_{\mathrm{reg}} \leq 0.$$

This lemma follows from the minimax-optimal strategy of the learner and has been shown to hold in Lee et al. (2022). We combine the previous two lemmas to get our main result.

Theorem 1. *Assume that $\gamma \leq 1/2$. Then, for any interval $I = [r, s] \subseteq [T]$,*

$$\max\left\{ \max_{f,\sigma} \frac{1}{|I|} \sum_{t=r}^{s} \ell^{(t)}_{\mathrm{MA}_{f,\sigma}}, \frac{1}{|I|} \sum_{t=r}^{s} \ell^{(t)}_{\mathrm{reg}} \right\} \leq \frac{c\eta}{|I|} \left(\sum_{t=r}^{s} q^{(t)\top}_{\mathrm{MA}} {\ell^{(t)}_{\mathrm{MA}}}^2 + \ q^{(t)}_{\mathrm{reg}} {\ell^{(t)}_{\mathrm{reg}}}^2 \right) + \frac{1}{\eta |I|}\left(\log\left(\frac{2|\mathcal{F}_{\mathrm{MA}}| + 1}{\gamma} \right) + |I| 2\gamma \right), \tag{7}$$

where the data $(x_t, y_t), t \in [T]$ are fixed and the randomness is from $q^{(t)}_{\ell}, \ell \in \mathcal{L}$ in Algorithm 1.

If we substitute the optimal values $\gamma = \frac{1}{2|I|}$ and $\eta = \sqrt{\frac{\log((2|\mathcal{F}_{\mathrm{MA}}|+1) \cdot 2|I|) + 1}{c \sum_{t=r}^{s} q^{(t)\top}_{\mathrm{MA}} {\ell^{(t)}_{\mathrm{MA}}}^2 + q^{(t)}_{\mathrm{reg}} {\ell^{(t)}_{\mathrm{reg}}}^2}}$ in (7), we obtain

$$\max\left\{ \max_{f,\sigma} \frac{1}{|I|} \sum_{t=r}^{s} \ell^{(t)}_{\mathrm{MA}_{f,\sigma}}, \frac{1}{|I|} \sum_{t=r}^{s} \ell^{(t)}_{\mathrm{reg}} \right\} \leq \frac{(c+1)}{|I|} \sqrt{\left(\log((2|\mathcal{F}_{\mathrm{MA}}| + 1) \cdot 2|I|) + 1 \right)} \cdot \sqrt{\sum_{t=r}^{s} q^{(t)\top}_{\mathrm{MA}} {\ell^{(t)}_{\mathrm{MA}}}^2 + \ q^{(t)}_{\mathrm{reg}} {\ell^{(t)}_{\mathrm{reg}}}^2}$$

$$= O\left(\sqrt{\frac{\log(|\mathcal{L}| \cdot |I|)}{|I|}} \right)$$

We note that the optimal values for γ and η used above depend on quantities that are unknown a priori in practice: the interval width $|I|$, the constant c, and the expected squared objective

$\sqrt{\sum_{t=r}^{s} q_{\text{MA}}^{(t)\top} \ell_{\text{MA}}^{(t)}{}^2 + q_{\text{reg}}^{(t)} \ell_{\text{reg}}^{(t)}{}^2}$. We let the user pick a fixed interval width τ, noting that a smaller choice of τ gives stronger locally adaptive guarantees at the cost of a looser upper bound. We follow Gibbs & Candès (2024) in selecting an adaptive value of η that updates online as

$$\eta = \eta_t := \sqrt{\frac{\log((2|\mathcal{F}_{\text{MA}}| + 1) \cdot 2\tau) + 1}{\sum_{s=t-\tau+1}^{t} q_{\text{MA}}^{(s)\top} \ell_{\text{MA}}^{(s)}{}^2 + q_{\text{reg}}^{(s)} \ell_{\text{reg}}^{(s)}{}^2}},$$

where we fix $c = 1$. This choice lets the algorithm adaptively track changes in the moving average of the expected squared objective over the most recent τ time steps.

4 Experiments

4.1 Datasets

GEFCom2014 electric load forecasting. Global Energy Forecasting Competition 2014 (GEFCom2014) (Hong et al., 2016) was a probabilistic energy forecasting competition conducted with four tracks on load, price, wind and solar forecasting. In this work, we study the electricity demand forecasting track GEFCom2014-L. In Section 1.1, we shared details regarding the task and Figure 1a displays the load and temperature trends over time. We consider a binary load prediction task for our empirical evaluation where $y_t = \mathbb{1}\{\text{load}_t \geq 150\text{MW}\}$. We construct our baseline predictions $\tilde{p}_t$ using the quantiles forecasted by Ziel & Liu (2016) who outperform top entries of the competition.

COMPAS dataset. Larson et al. (2016) analyzed the COMPAS tool used to predict recidivism for criminal defendants in Broward County, Florida and found that certain groups of defendants are more likely to be incorrectly judged as high risk of recidivism. In Figure 2, we plot the true recidivism rate over time for different racial groups. We consider the recidivism prediction task and evaluate the local multiaccuracy of predictors with respect to the African-American, Caucasian, and Hispanic subgroups. We use the COMPAS risk scores provided in the dataset as our baseline predictions. Following the analysis of Barenstein (2019), we drop the data points beyond the two year cutoff.

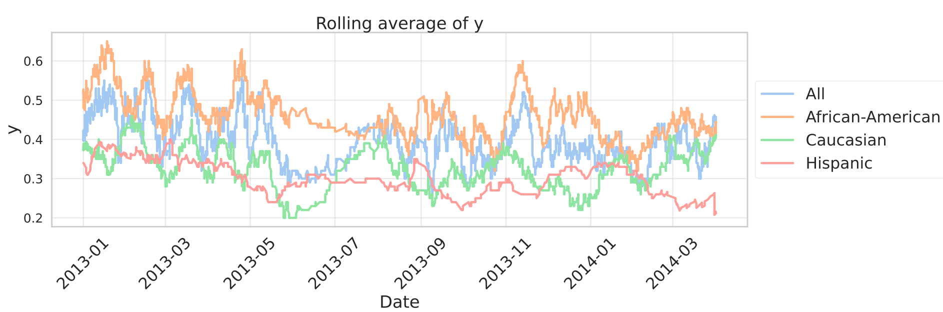

Figure 2: COMPAS dataset. Rolling average of true recidivism over time.

4.2 Baselines

We consider baselines that differ in their adaptivity and the set of objectives we consider in $\mathcal{L}$. We refer to our locally adaptive algorithm with the multiaccuracy and regret objectives (Algorithm 1) as MA+reg and its non-adpative variant as MA+reg (non-adaptive). The non-adaptive variant is a special case of Algorithm 1 with $\gamma = 0$ and $\eta = O(\sqrt{\frac{|\mathcal{L}|}{T}})$.

We explain the baselines below:

1. Baseline predictor $\tilde{p}_t$: These are the predictions that were constructed in advance for the application and are our input to Algorithm 1.
2. Multiaccuracy algorithm (MA) with $\mathcal{L} := \{\ell_{\text{MA}_{f,\sigma}} : f \in \mathcal{F}_{\text{MA}}, \sigma \in \{\pm\}\}$: This is a specific case of our algorithm where the set $\mathcal{L}$ does not include the regret objective.

3. Multicalibration (MC): We implement the online multicalibration algorithm from Lee et al. (2022). This is a competitive algorithm as multicalibration is a stronger condition than multiaccuracy. Lee et al. (2022) show that their algorithm can guarantee *multicalibeating* in addition to multicalibration, i.e., the predictions satisfy an accuracy objective (specifically, low squared error) on subgroups along with calibration. Hence, we consider the squared error objective as ℓ_{acc} in our regret objective for fair comparison.

4.3 Local multiaccuracy and regret evaluation

In this section, we evaluate the local multiaccuracy error ℓ_{MA} and regret ℓ_{reg} incurred by the algorithms we defined above.

First, we look at the results on GEFCom2014-L dataset (Figure 3). We take the interval width $|I| = 336$ hours (2 weeks) for this set of experiments. We compute empirical local multiaccuracy and regret rates over this moving 2 week-window. It can be seen that the constructed baseline predictor $\tilde{p}_t$ has high local multiaccuracy error and all algorithms improve over this baseline. Overall, the adaptive algorithms (MA and MA+reg) have close to zero multiaccuracy over all local intervals. On the other hand, the non-adaptive algorithms have high variance. It is interesting to note that the multicalibration algorithm (MC) has significantly slower multiaccuracy rates in practice.

Next, we turn to study the empirical local regret of these algorithms. As expected, the MA baseline has non-zero local regret and we lose accuracy with respect to the predictor $\tilde{p}_t$ in the absence of the regret objective. MA+reg consistently preserves or improves accuracy over $\tilde{p}_t$. As promised by the multicalibration+multicalibeating algorithm in Lee et al. (2022), we observe that MC generally has negative regret, although with poorer adaptivity compared to MA+reg.

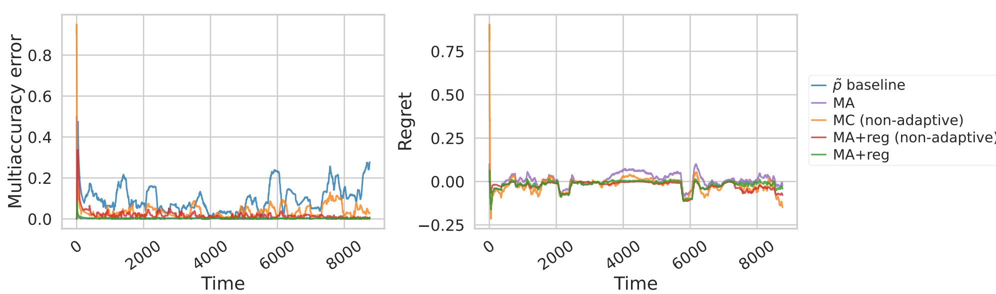

Figure 3: Local multiaccuracy error and regret on GEFCom2014-L dataset.

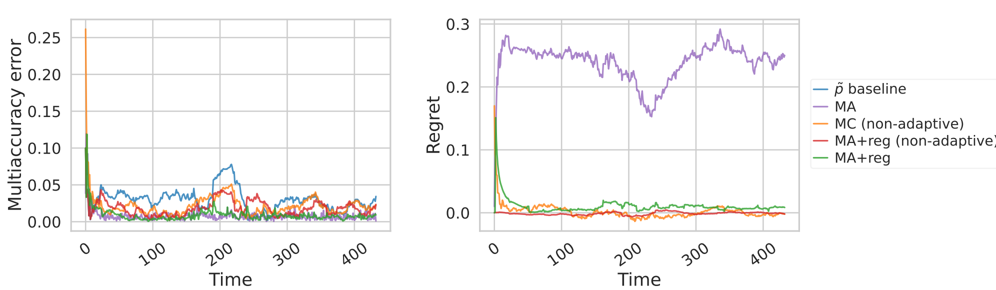

Figure 4: Local multiaccuracy error and regret on COMPAS dataset.

Now, we examine our results on the COMPAS dataset (Figure 4). Here, we fix $|I| = 50$ days. Our findings from above are seen to generalize here. $\tilde{p}_t$, MC, and MA+reg (non-adaptive) show minimal adaptivity to the underlying shifts and are expected to perform poorly across some subgroups over local intervals; whereas, our proposed algorithm has significantly better local multiaccuracy. We note that while MA performs slightly better in terms of multiaccuracy compared to MA+reg, it suffers from significantly higher regret as can be seen from the right plot.

4.4 Comparison with strongly adaptive baseline

Finally, we compare our algorithm with a strongly adaptive extension of the online multicalibration algorithm proposed in Lee et al. (2022) (Figure 5). This algorithm guarantees low multicalibration error on all subintervals in T at the expense of higher runtime and memory. While we use the fixed width value $|I| = 336$ in our algorithm, we perform a general evaluation here over different interval widths. We find that while adaptivity slightly improves the performance of the multicalibration algorithm, MA+reg still has significantly better local multiaccuracy across all interval widths. We note that there may be practical modifications that improve the empirical performance of the adaptive multicalibration algorithm.

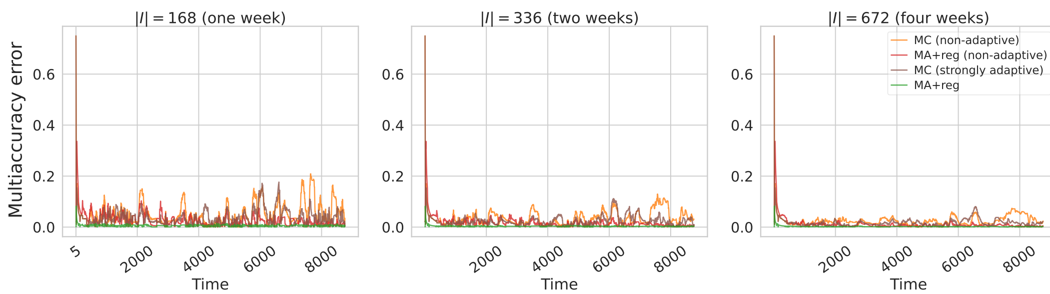

Figure 5: Local multiaccuracy error on GEFCom2014-L under different interval widths.

5 Extensions

We now discuss the several extensions of our general algorithm.

Quantile estimation. Analogous to mean estimation, our algorithm can be used to update predicted quantiles to satisfy group-conditional coverage guarantees while preserving the quantile loss ℓ_α (also referred to as pinball loss) with respect to baseline quantile predictions. We provide the full algorithm in Algorithm 3. Note that we have to allow θ_t to be random in this algorithm.

Multi-group learning. Our algorithm can be extended to multi-group learning with the set of objectives of the form $\mathbb{1}\{x_t \in g\}(\ell(p_t(x_t), y_t) - \ell(f(x_t), y_t))$ for groups $g \in \mathcal{G}$ and functions $f \in \mathcal{F}$.

Omniprediction. Omniprediction is a straightforward extension of our algorithm where the set of objectives are of the form $\ell(p_t(x_t), y_t) - \ell(f(x_t), y_t)$ for functions $f \in \mathcal{F}$ and losses $\ell \in \mathcal{L}$.

6 Discussion

We present adaptive multi-objective learning algorithms that guarantee small error for all objectives over local time intervals. In this growing literature, we hope our work serves as an initial step toward bridging the empirical gap. We note two limitations of our work: firstly, our guarantees hold over intervals of fixed width. Developing efficient algorithms that provide stronger guarantees in this setting remains an important problem. Second, our empirical evaluation focuses on a subset of objectives and validation on broader problems is interesting future work.

Reproducibility Statement

We provide the code to run our algorithms and reproduce our experiments as part of the supplementary material. We state the assumptions in the main text and include proofs for our theoretical results in Appendix A. We describe all experimental details in Section 4.

REFERENCES

Jacob Abernethy, Peter L. Bartlett, and Elad Hazan. Blackwell approachability and low-regret learning are equivalent. In *Proceedings of the Annual Conference on Learning Theory*, 2011.

Matias Barenstein. Propublica's compas data revisited. *arXiv preprint arXiv:1906.04711*, 2019.

David Blackwell. An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics*, 6(1):1 – 8, 1956.

Amit Daniely, Alon Gonen, and Shai Shalev-Shwartz. Strongly adaptive online learning. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 1405–1411. PMLR, 2015.

Samuel Deng, Jingwen Liu, and Daniel Hsu. Group-wise oracle-efficient algorithms for online multi-group learning. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.

Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139, 1997.

Sumegha Garg, Christopher Jung, Omer Reingold, and Aaron Roth. Oracle efficient online multicalibration and omniprediction. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 2725–2792, 2024.

Isaac Gibbs and Emmanuel J. Candès. Conformal inference for online prediction with arbitrary distribution shifts. *Journal of Machine Learning Research*, 25(162):1–36, 2024.

Ira Globus-Harris, Declan Harrison, Michael Kearns, Aaron Roth, and Jessica Sorrell. Multicalibration as boosting for regression. In *Proceedings of the 40th International Conference on Machine Learning*, Proceedings of Machine Learning Research, pp. 11459–11492, 2023.

Parikshit Gopalan, Adam Tauman Kalai, Omer Reingold, Vatsal Sharan, and Udi Wieder. Omnipredictors. In *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*, volume 215 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pp. 79:1–79:21, 2022.

Paula Gradu, Elad Hazan, and Edgar Minasyan. Adaptive regret for control of time-varying dynamics. In *Proceedings of The 5th Annual Learning for Dynamics and Control Conference*, volume 211, pp. 560–572. PMLR, 15–16 Jun 2023.

Varun Gupta, Christopher Jung, Georgy Noarov, Mallesh M. Pai, and Aaron Roth. Online Multivalid Learning: Means, Moments, and Prediction Intervals. In *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*, pp. 82:1–82:24, 2022.

Nika Haghtalab, Michael Jordan, and Eric Zhao. A unifying perspective on multi-calibration: Game dynamics for multi-objective learning. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023a.

Nika Haghtalab, Chara Podimata, and Kunhe Yang. Calibrated stackelberg games: Learning optimal commitments against calibrated agents. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023b.

Ursula Hebert-Johnson, Michael Kim, Omer Reingold, and Guy Rothblum. Multicalibration: Calibration for the (Computationally-identifiable) masses. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 1939–1948, 2018.

Mark Herbster and Manfred K. Warmuth. Tracking the best expert. *Machine Learning*, 32(2): 151–178, 1998.

Tao Hong, Pierre Pinson, Shu Fan, Hamidreza Zareipour, Alberto Troccoli, and Rob J. Hyndman. Probabilistic energy forecasting: Global energy forecasting competition 2014 and beyond. *International Journal of Forecasting*, 32(3):896–913, 2016.

Kwang-Sung Jun, Francesco Orabona, Stephen Wright, and Rebecca Willett. Improved Strongly Adaptive Online Learning using Coin Betting. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pp. 943–951. PMLR, 2017.

Michael P. Kim, Amirata Ghorbani, and James Zou. Multiaccuracy: Black-box post-processing for fairness in classification. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 247–254, 2019.

Jeff Larson, Surya Mattu, Lauren Kirchner, and Julia Angwin. How we analyzed the compas recidivism algorithm. *ProPublica*, 2016. URL https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm.

Daniel Lee, Georgy Noarov, Mallesh Pai, and Aaron Roth. Online minimax multiobjective optimization: Multicalibeating and other applications. In *Advances in Neural Information Processing Systems*, 2022.

Georgy Noarov, Ramya Ramalingam, Aaron Roth, and Stephan Xie. High-dimensional prediction for sequential decision making. In *Forty-second International Conference on Machine Learning*, 2025.

Princewill Okoroafor, Robert Kleinberg, and Michael P. Kim. Near-optimal algorithms for omniprediction. *arXiv preprint arXiv:2501.17205*, 2025.

Christopher J Tosh and Daniel Hsu. Simple and near-optimal algorithms for hidden stratification and multi-group learning. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 21633–21657. PMLR, 2022.

Florian Ziel and Bidong Liu. Lasso estimation for GEFCom2014 probabilistic electric load forecasting. *International Journal of Forecasting*, 32(3):1029–1037, 2016.

A Proofs

A.1 Proof of Lemma 1

We follow the calculations of Gradu et al. (2023) and Gibbs & Candès (2024). The primary difference between our work and these articles is that our losses may take on negative values. The constant c introduced below accounts for this.

Let $W^{(t+1)} := \sum_{f,\sigma} w^{(t)}_{\mathrm{MA}_{f,\sigma}} \exp\left(\eta \cdot \sigma f(x_t)(y_t - p_t(x_t))\right) + w^{(t)}_{\mathrm{reg}} \exp\left(\eta \cdot (\ell_{\mathrm{acc}}(p_t(x_t), y) - \ell_{\mathrm{acc}}(\tilde{p}_t(x_t), y))\right)$. We initialize $w^{(t)}_{\mathrm{MA}_{f,\sigma}} = 1$ for all $f \in \mathcal{F}, \sigma \in \{\pm 1\}$ and $w^{(t)}_{\mathrm{reg}} = 1$. By construction, the probabilities $q^{(t)}_{\mathrm{MA}_{f,\sigma}} := w^{(t)}_{\mathrm{MA}_{f,\sigma}} / (\sum_{f,\sigma} w^{(t)}_{\mathrm{MA}_{f,\sigma}} + w^{(t)}_{\mathrm{reg}})$ and $q^{(t)}_{\mathrm{reg}} := w^{(t)}_{\mathrm{reg}} / (\sum_{f,\sigma} w^{(t)}_{\mathrm{MA}_{f,\sigma}} + w^{(t)}_{\mathrm{reg}})$. Thus,

$$\frac{W^{(t+1)}}{W^{(t)}} = \sum_{f,\sigma} q^{(t)}_{\mathrm{MA}_{f,\sigma}} \exp\left(\eta \sigma f(x_t)(y_t - p_t(x_t))\right) + q^{(t)}_{\mathrm{reg}} \exp\left(\eta(\ell_{\mathrm{acc}}(p_t(x_t), y_t) - \ell_{\mathrm{acc}}(\tilde{p}_t(x_t), y_t)\right).$$

We use the inequalities $1 - a \leq \exp(-a)$ and $\exp(a) \leq 1 + a + \frac{e^{|a|}}{2} a^2$ to get

$$\frac{W^{(t+1)}}{W^{(t)}} \leq \exp\left(\eta q^{(t)\top}_{\mathrm{MA}} \ell^{(t)}_{\mathrm{MA}} + c\eta^2 q^{(t)\top}_{\mathrm{MA}} {\ell^{(t)}_{\mathrm{MA}}}^2 + \eta q^{(t)}_{\mathrm{reg}} \ell^{(t)}_{\mathrm{reg}} + c\eta^2 q^{(t)}_{\mathrm{reg}} {\ell^{(t)}_{\mathrm{reg}}}^2\right),$$

where $c = \sup_{\ell \in \{\ell^{(t)}_{\mathrm{MA}}, \ell^{(t)}_{\mathrm{reg}}\}} \frac{e^{\eta \ell}}{2}$. This inductively implies, for interval $I = [r, s]$

$$\frac{W^{(s+1)}}{W^{(r)}} \leq \exp\left(\sum_{t=r}^{s} \eta q^{(t)\top}_{\mathrm{MA}} \ell^{(t)}_{\mathrm{MA}} + c\eta^2 q^{(t)\top}_{\mathrm{MA}} {\ell^{(t)}_{\mathrm{MA}}}^2 + \eta q^{(t)}_{\mathrm{reg}} \ell^{(t)}_{\mathrm{reg}} + c\eta^2 q^{(t)}_{\mathrm{reg}} {\ell^{(t)}_{\mathrm{reg}}}^2\right).$$

On the other hand, for any fixed $\ell \in \mathcal{L}$, $w^{(t+1)}_{\ell} \geq w^{(t)}_{\ell} (1 - \gamma) \exp\left(\eta \ell^{(t)}\right)$. Without loss of generality, we proceed with a fixed f and σ, noting that the same calculations will follow for ℓ_{reg}. This gives

$$\begin{aligned}\frac{W^{(s+1)}}{W^{(r)}} \geq \frac{w^{(s+1)}_{\mathrm{MA}_{f,\sigma}}}{W^{(r)}} &\geq (1-\gamma)^{|I|} q^{(t)}_{\mathrm{MA}_{f,\sigma}} \exp\left(\sum_{t=r}^{s} \eta \ell^{(t)}_{\mathrm{MA}_{f,\sigma}}\right) \\ &\geq (1-\gamma)^{|I|} \frac{\gamma}{2|\mathcal{F}_{\mathrm{MA}}| + 1} \exp\left(\sum_{t=r}^{s} \eta \ell^{(t)}_{\mathrm{MA}_{f,\sigma}}\right)\end{aligned}$$

Combining the two inequalities and taking logarithm on both sides yields

$$\begin{aligned}|I|(1-\gamma) + \log\left(\frac{\gamma}{2|\mathcal{F}_{\mathrm{MA}}| + 1}\right) + \sum_{t=r}^{s} \eta \ell^{(t)}_{\mathrm{MA}_{f,\sigma}} \leq \sum_{t=r}^{s} \eta q^{(t)\top}_{\mathrm{MA}} \ell^{(t)}_{\mathrm{MA}} + c\eta^2 q^{(t)\top}_{\mathrm{MA}} {\ell^{(t)}_{\mathrm{MA}}}^2 + \\ \eta q^{(t)}_{\mathrm{reg}} \ell^{(t)}_{\mathrm{reg}} + c\eta^2 q^{(t)}_{\mathrm{reg}} {\ell^{(t)}_{\mathrm{reg}}}^2.\end{aligned}$$

We rearrange to get the following inequality

$$\begin{aligned}\sum_{t=r}^{s} q^{(t)\top}_{\mathrm{MA}} \ell^{(t)}_{\mathrm{MA}} + q^{(t)}_{\mathrm{reg}} \ell^{(t)}_{\mathrm{reg}} \geq \sum_{t=r}^{s} \ell^{(t)}_{\mathrm{MA}_{f,\sigma}} - c\eta\left(\sum_{t=r}^{s} q^{(t)\top}_{\mathrm{MA}} {\ell^{(t)}_{\mathrm{MA}}}^2 + q^{(t)}_{\mathrm{reg}} {\ell^{(t)}_{\mathrm{reg}}}^2\right) + \\ \frac{1}{\eta}|I|(1-\gamma) + \log\left(\frac{\gamma}{2|\mathcal{F}_{\mathrm{MA}}| + 1}\right).\end{aligned}$$

As $\gamma \leq 1/2$, we can use the inequality $\log(1-\gamma) \geq -2\gamma$ to get the final inequality

$$\sum_{t=r}^{s} q_{\text{MA}}^{(t)\top} \ell_{\text{MA}}^{(t)} + q_{\text{reg}}^{(t)} \ell_{\text{reg}}^{(t)} \geq \sum_{t=r}^{s} \ell_{\text{MA}_{f,\sigma}}^{(t)} - c\eta \left(\sum_{t=r}^{s} q_{\text{MA}}^{(t)\top} {\ell_{\text{MA}}^{(t)}}^2 + q_{\text{reg}}^{(t)} {\ell_{\text{reg}}^{(t)}}^2 \right) - \frac{1}{\eta} \left(\log \left(\frac{2|\mathcal{F}_{\text{MA}}|+1}{\gamma} \right) + |I| 2\gamma \right).$$

As the same calculation holds for any objective $\ell \in \mathcal{L}$, we get the final result

$$\sum_{t=r}^{s} q_{\text{MA}}^{(t)\top} \ell_{\text{MA}}^{(t)} + q_{\text{reg}}^{(t)} \ell_{\text{reg}}^{(t)} \geq \max \left\{ \max_{f,\sigma} \sum_{t=r}^{s} \ell_{\text{MA}_{f,\sigma}}^{(t)}, \sum_{t=r}^{s} \ell_{\text{reg}}^{(t)} \right\} - c\eta \left(\sum_{t=r}^{s} q_{\text{MA}}^{(t)\top} {\ell_{\text{MA}}^{(t)}}^2 + q_{\text{reg}}^{(t)} {\ell_{\text{reg}}^{(t)}}^2 \right) - \frac{1}{\eta} \left(\log \left(\frac{2|\mathcal{F}_{\text{MA}}|+1}{\gamma} \right) + |I| 2\gamma \right),$$

A.2 Proof of Lemma 2

This result was shown in Lee et al. (2022) and we follow the same calculations.

Let $u^{(t)}(p,y) := \sum_{f,\sigma} q_{\text{MA}_{f,\sigma}}^{(t)} \sigma f(x_t)(y - p(x_t)) + q_{\text{reg}}^{(t)}(\ell_{\text{acc}}(p(x_t),y) - \ell_{\text{acc}}(\tilde{p}_t(x_t),y))$. Let $\Delta(\mathcal{Y})$ denote the space of distributions over $\mathcal{Y}$. Applying Sion's Minimax Theorem, we get $\min_p \max_{y \in \mathcal{Y}} u^{(t)}(p,y) = \min_p \max_{P \in \Delta(\mathcal{Y})} \mathbb{E}_{y \sim P}[u^{(t)}(p,y)] = \max_{P \in \Delta(\mathcal{Y})} \min_p \mathbb{E}_{y \sim P}[u^{(t)}(p,y)]$. This conveys that the minimax-optimal strategy p_t of the learner can achieve $u^{(t)}(p,y)$ as low as if the adversary moved first and the learner could best-respond. In this latter case, $p = \mathbb{E}_{y \sim P}[y]$ gives $u^{(t)}(p,y) \leq 0$.

Thus, the minimax optimal strategy guarantees that $\min_p \max_{y \in \mathcal{Y}} u^{(t)}(p,y) \leq 0$ for all $t \in [T]$. This yields the desired inequality

$$\sum_{t=r}^{s} q_{\text{MA}}^{(t)\top} \ell_{\text{MA}}^{(t)} + q_{\text{reg}}^{(t)} \ell_{\text{reg}}^{(t)} \leq 0.$$

A.3 Proof of Theorem 1

Applying Lemma 2 to the inequality (6) in Lemma 1 gives

$$\max \left\{ \max_{f,\sigma} \sum_{t=r}^{s} \ell_{\text{MA}_{f,\sigma}}^{(t)}, \sum_{t=r}^{s} \ell_{\text{reg}}^{(t)} \right\} - c\eta \left(\sum_{t=r}^{s} q_{\text{MA}}^{(t)\top} {\ell_{\text{MA}}^{(t)}}^2 + q_{\text{reg}}^{(t)} {\ell_{\text{reg}}^{(t)}}^2 \right) - \frac{1}{\eta} \left(\log \left(\frac{2|\mathcal{F}_{\text{MA}}|+1}{\gamma} \right) + |I| 2\gamma \right) \leq 0.$$

Rearranging and dividing both sides by $|I|$ yields the desired inequality,

$$\max \left\{ \max_{f,\sigma} \frac{1}{|I|} \sum_{t=r}^{s} \ell_{\text{MA}_{f,\sigma}}^{(t)}, \frac{1}{|I|} \sum_{t=r}^{s} \ell_{\text{reg}}^{(t)} \right\} \leq \frac{c\eta}{|I|} \left(\sum_{t=r}^{s} q_{\text{MA}}^{(t)\top} {\ell_{\text{MA}}^{(t)}}^2 + q_{\text{reg}}^{(t)} {\ell_{\text{reg}}^{(t)}}^2 \right) + \frac{1}{\eta |I|} \left(\log \left(\frac{2|\mathcal{F}_{\text{MA}}|+1}{\gamma} \right) + |I| 2\gamma \right).$$

B Deferred algorithms

Algorithm 2 Locally adaptive multiaccurate mean estimation (learning $\tilde{p}_t$ online)

Input: Function class $\mathcal{F}_{\text{MA}} \subseteq \{f : \mathcal{X} \to [0,1]\}$; $\mathcal{F}_{\text{pred}} = \{f_\beta : \beta \in \mathbb{R}\}$; hyperparameters η, γ.
Input: Sequence of samples $\{(x_1, y_1), \ldots, (x_T, y_T)\}$
1: $q^{(1)}_{\text{MA}_{f,\sigma}} = \frac{1}{2|\mathcal{F}_{\text{MA}}|+1}, \quad \forall f \in \mathcal{F}_{\text{MA}}, \sigma \in \{\pm 1\}$.
2: $q^{(1)}_{\text{reg}} = \frac{1}{2|\mathcal{F}_{\text{MA}}|+1}$
3: $\beta_1 = 0$
4: **for** each $t \in [T]$ **do**
5: $\quad \beta_{t+1} = \beta_t - \gamma \nabla_\beta \ell\left(f_{\beta_t}(x_t), y_t\right)$
6: $\quad \tilde{p}_t(x_t) := f_{\beta_t}(x_t)$
7: $\quad p_t(x_t) := \operatorname*{argmin}_p \max_{y \in \{0,1\}} \sum_{f,\sigma} q^{(t)}_{\text{MA}_{f,\sigma}} \sigma f(x_t)(y - p(x_t)) + q^{(t)}_{\text{reg}}(\ell_{\text{acc}}(p(x_t), y) - \ell_{\text{acc}}(\tilde{p}_t(x_t), y))$
8: $\quad \tilde{q}^{(t+1)}_{\text{MA}_{f,\sigma}} \propto q^{(t)}_{\text{MA}_{f,\sigma}} \exp\left(\eta \cdot \sigma f(x_t)(y_t - p_t(x_t))\right)$ for all $f \in \mathcal{F}, \sigma \in \{+,-\}$
9: $\quad \tilde{q}^{(t+1)}_{\text{reg}} \propto q^{(t)}_{\text{reg}} \exp\left(\eta \cdot (\ell(p_t(x_t), y_t) - \ell(\tilde{p}_t(x_t), y_t)\right)$
10: $\quad q^{(t+1)}_{\text{MA}_{f,\sigma}} = (1-\gamma)\, \tilde{q}^{(t+1)}_{\text{MA}_{f,\sigma}} + \frac{\gamma}{2|\mathcal{F}_{\text{MA}}|+1}$
11: $\quad q^{(t+1)}_{\text{reg}} = (1-\gamma)\, \tilde{q}^{(t+1)}_{\text{reg}} + \frac{\gamma}{2|\mathcal{F}_{\text{MA}}|+1}$
Output: Sequence of (randomized) predictors $p_1, \ldots, p_T$

Algorithm 3 Locally adaptive multiaccurate quantile estimation

Input: Function class $\mathcal{F}_{\text{MA}} \subseteq \{f : \mathcal{X} \to [0,1]\}$; quantile level α; quantile predictions $\tilde{\theta}_t, t \in [T]$; hyperparameters η, γ.
Input: Sequence of samples $\{(x_1, y_1), \ldots, (x_T, y_T)\}$
1: $q^{(1)}_{\text{MA}_{f,\sigma}} = \frac{1}{2|\mathcal{F}_{\text{MA}}|+1}, \quad \forall f \in \mathcal{F}_{\text{MA}}, \sigma \in \{\pm 1\}$.
2: $q^{(1)}_{\text{reg}} = \frac{1}{2|\mathcal{F}_{\text{MA}}|+1}$
3: **for** each $t \in [T]$ **do**
4: $\quad \theta_t(x_t) := \operatorname*{argmin}_{\theta \in \Delta(\Theta)} \max_{y \in \mathcal{Y}} \sum_{f,\sigma} q^{(t)}_{\text{MA}_{f,\sigma}} \sigma f(x_t)(\mathbb{1}\{y \leq \theta\} - \alpha) + q^{(t)}_{\text{reg}} \left(\ell_\alpha(\theta, y) - \ell_\alpha(\tilde{\theta}_t, y_t)\right)$
5: $\quad \tilde{q}^{(t+1)}_{\text{MA}_{f,\sigma}} \propto q^{(t)}_{\text{MA}_{f,\sigma}} \exp\left(\eta \cdot \sigma \cdot f(x_t)(\mathbb{1}\{y_t \leq \theta_t(x_t)\} - \alpha)\right)$ for all $f \in \mathcal{F}_{\text{MA}}, \sigma \in \{\pm 1\}$
6: $\quad \tilde{q}^{(t+1)}_{\text{reg}} \propto q^{(t)}_{\text{reg}} \exp\left(\eta \cdot (\ell_\alpha(\theta_t(x_t), y_t) - \ell_\alpha(\tilde{\theta}_t, y_t)\right)$
7: $\quad q^{(t+1)}_{\text{MA}_{f,\sigma}} = (1-\gamma)\, \tilde{q}^{(t+1)}_{\text{MA}_{f,\sigma}} + \frac{\gamma}{2|\mathcal{F}_{\text{MA}}|+1}$
8: $\quad q^{(t+1)}_{\text{reg}} = (1-\gamma)\, \tilde{q}^{(t+1)}_{\text{reg}} + \frac{\gamma}{2|\mathcal{F}_{\text{MA}}|+1}$
Output: Sequence of (randomized) quantile predictors $\theta_1, \ldots, \theta_T$

C ADDITIONAL PLOTS

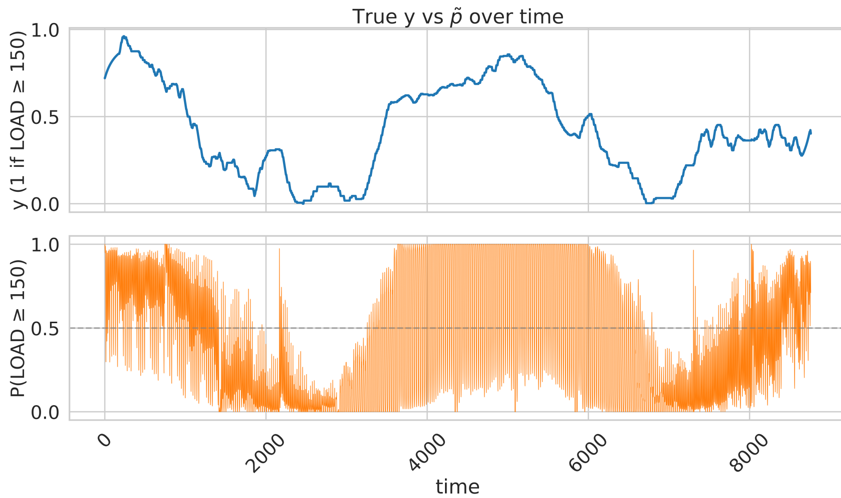

Figure 6: GEFCom2014-L: True load (y) vs $\tilde{p}_t$
. We plot the moving average of the binary y over a window size 336 (2 weeks) and compare with the baselines predictions $\tilde{p}_t$ constructed using quantile forecasts.