
Generalized Implicit Follow-The-Regularized-Leader

Keyi Chen¹ Francesco Orabona¹

Abstract

We propose a new class of online learning algorithms, generalized implicit Follow-The-Regularized-Leader (FTRL), that expands the scope of FTRL framework. Generalized implicit FTRL can recover known algorithms, such as FTRL with linearized losses and implicit FTRL, and it allows the design of new update rules, as extensions of aProx and Mirror-Prox to FTRL. Our theory is constructive in the sense that it provides a simple unifying framework to design updates that directly improve the worst-case upper bound on the regret. The key idea is substituting the linearization of the losses with a Fenchel-Young inequality. We show the flexibility of the framework by proving that some known algorithms, like the Mirror-Prox updates, are instantiations of the generalized implicit FTRL. Finally, the new framework allows us to recover the temporal variation bound of implicit OMD, with the same computational complexity.

1. Introduction

Online learning is a setting where the learner receives an arbitrary sequence of loss functions, selects points before knowing the loss functions, and is evaluated on the values of the loss functions on the points it selects (Cesa-Bianchi & Lugosi, 2006; Orabona, 2019; Cesa-Bianchi & Orabona, 2021). More in detail, at round t the learner outputs a point \mathbf{x}_t in a feasible set $V \subseteq \mathbb{R}^d$. Then, it receives a loss function $\ell_t : V \rightarrow \mathbb{R}$ and it pays the value $\ell_t(\mathbf{x}_t)$. Given the arbitrary nature of the losses, the learner cannot guarantee to have a small cumulative loss, $\sum_{t=1}^T \ell_t(\mathbf{x}_t)$. On the other hand, it is possible to minimize the *regret*, that is the difference between the cumulative loss of the algorithm and the one of

any arbitrary comparator $\mathbf{u} \in V$:

$$\text{Regret}_T(\mathbf{u}) \triangleq \sum_{t=1}^T \ell_t(\mathbf{x}_t) - \sum_{t=1}^T \ell_t(\mathbf{u}).$$

In particular, a successful online learning algorithm must guarantee a regret that grows sublinearly in time for any $\mathbf{u} \in V$. In this way, its average performance approaches the one of the best comparator in hindsight.

There are two families of online learning algorithms: Online Mirror Descent (OMD) (Nemirovskij & Yudin, 1983; Warmuth & Jagota, 1997) and Follow-the-Regularized-Leader (FTRL) (Shalev-Shwartz, 2007; Abernethy et al., 2008; Hazan & Kale, 2008). They stem from two similar but complementary approaches: the update of OMD aims at minimizing a linearization of the current loss without going too far from its previous prediction \mathbf{x}_t , while FTRL minimizes the sum of all the losses (or their linear approximation) plus a regularization term. On the contrary to the first approaches in online learning that focused on specific algorithms (e.g., the Winnow algorithm (Littlestone, 1988)), the theory of these two frameworks is particularly interesting because it allows *both the design and the analysis of generic online learning algorithms*.

While FTRL and OMD provide similar bounds in most situations, they are not completely equivalent. For example, FTRL has an advantage over OMD in unbounded domains, where it allows to use time-varying regularizers. In fact, OMD allows the use of time-varying stepsizes only in domains where its associated Bregman divergence is bounded. On the other hand, in the cases where we can use time-varying stepsizes, OMD can achieve a superior adaption to the gradients (see, e.g., Theorem 2 in Streeter & McMahan (2010) versus Theorem 2 in Orabona & Pál (2015)). In this view, these two frameworks are complementary.¹ Moreover, there exists another orthogonal axis on the use of the actual loss functions or a linear surrogate for both frameworks. We summarize all the variants of OMD and FTRL in Table 1.

Our motivation stems from the fact that in practical cases, all the variants that use full losses offer a big advantage in terms of empirical performance at the cost of a higher

¹Boston University, Boston, MA, USA. Correspondence to: Keyi Chen <keyichen@bu.edu>, Francesco Orabona <francesco@orabona.com>.

¹See also the blog post on this topic by Tim van Erven at <https://www.timvanerven.nl/blog/ftrl-vs-omd/>.

Table 1. Summary of implicit and linearized updates for FTRL and OMD. (The Bregman divergence $B_\psi(\mathbf{x}; \mathbf{y})$ is defined as $\psi(\mathbf{x}) - \psi(\mathbf{y}) - \langle \nabla \psi(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle$. The * denotes the Fenchel conjugate.)

Algorithm	Update
OMD (Warmuth & Jagota, 1997)	$\mathbf{x}_{t+1} = \operatorname{argmin}_{\mathbf{x} \in V} B_\psi(\mathbf{x}; \mathbf{x}_t) + \eta_t(\ell_t(\mathbf{x}_t) + \langle \mathbf{g}_t, \mathbf{x} - \mathbf{x}_t \rangle)$
Implicit OMD (Warmuth & Jagota, 1997)	$\mathbf{x}_{t+1} = \operatorname{argmin}_{\mathbf{x} \in V} B_\psi(\mathbf{x}; \mathbf{x}_t) + \eta_t \ell_t(\mathbf{x})$
FTRL (linearized) (Abernethy et al., 2008)	$\mathbf{x}_{t+1} = \operatorname{argmin}_{\mathbf{x} \in V} \psi_{t+1}(\mathbf{x}) + \sum_{i=1}^t (\ell_i(\mathbf{x}_i) + \langle \mathbf{g}_i, \mathbf{x} - \mathbf{x}_i \rangle)$
FTRL (full losses) (McMahan, 2017)	$\mathbf{x}_{t+1} = \operatorname{argmin}_{\mathbf{x} \in V} \psi_{t+1}(\mathbf{x}) + \sum_{i=1}^t \ell_i(\mathbf{x})$
Implicit FTRL (McMahan, 2010)	$\mathbf{x}_{t+1} = \operatorname{argmin}_{\mathbf{x} \in V} \psi_{t+1}(\mathbf{x}) + \ell_t(\mathbf{x}) + \sum_{i=1}^{t-1} (\ell_i(\mathbf{x}_i) + \langle \mathbf{g}_i, \mathbf{x} - \mathbf{x}_i \rangle)$
Generalized Implicit FTRL [This work]	$\mathbf{x}_{t+1} = \operatorname{argmin}_{\mathbf{x} \in V} \psi_{t+1}(\mathbf{x}) + \sum_{i=1}^t \langle \mathbf{z}_i, \mathbf{x} \rangle$ \mathbf{z}_i such that $\psi_{i+1,V}^*(\sum_{j=1}^i \mathbf{z}_j) + \ell_i^*(\mathbf{z}_i) \leq \psi_{i+1,V}^*(\sum_{j=1}^{i-1} \mathbf{z}_j - \mathbf{g}_i) + \ell_i^*(\mathbf{g}_i)$

computational complexity. On the theoretical side, the situation is not so clear given that in the worst case using the full losses can be equivalent to their linearized version, as it should be clear considering linear losses. In particular, the standard theoretical framework for FTRL does not allow a clear analysis of the implicit case. Moreover, while for implicit OMD it has been proven that one can achieve lower regret if the temporal variation of the losses is small, it is unclear if the same guarantee can be achieved for FTRL without the computational cost of using full losses.

In this paper, we aim at bridging this gap by proposing a *generalized* version of implicit FTRL. We go beyond implicit and linearized updates: *we directly construct the update rule in a way that minimizes an upper bound on the regret*. Our framework effectively expands the scope of the FTRL framework, fully retaining its coupling between design and analysis. Also, our updates come with a worst-case guarantee to never be worse than the standard linearized ones.

We show the flexibility of our framework recovering known update schemes, like the Mirror-Prox update (Nemirovski, 2004), or extending updates specifically designed for OMD to the FTRL case, like the aProx one (Asi & Duchi, 2019). Moreover, for the first time, we show an implicit version of FTRL that recovers the temporal variation bound of implicit OMD (Campolongo & Orabona, 2020), but with the same computational complexity of implicit OMD.

Related Work While there are many works on implicit mirror descent in both the online and offline setting (see, e.g., Moreau, 1965; Martinet, 1970; Rockafellar, 1976; Kivinen & Warmuth, 1997; Parikh & Boyd, 2014; Campolongo & Orabona, 2020; Shtoff, 2022), the number of works that deal with implicit updates for FTRL is quite limited. We are only aware of McMahan (2010), which quantifies a gain only for specific regularizers. However, the framework in McMahan (2010) is non-constructive in the sense that it is difficult to see how to generalize implicit updates. Joulani et al. (2017) extends this last result, but it does not provide a link with the maximization of the dual function that governs

the regret upper bound.

The closest approach to our framework is the one of Shalev-Shwartz & Singer (2007a;b), which develop a theory of FTRL updates as maximization of a dual function. However, their framework is limited to a specific shape of regularizers and it does not deal with implicit updates.

For implicit OMD, Campolongo & Orabona (2020) showed that implicit updates give rise to regret guarantees that depend on the temporal variability of the losses, so that constant regret is achievable if the variability of the losses is zero. They suggest that FTRL with full losses can achieve the same guarantee, but they also point out that given its computational complexity it would be “not worth pursuing.” Here, we show how to achieve the same bound of implicit OMD with our generalized implicit FTRL, while retaining the same computational complexity of implicit OMD.

Proximal updates on truncated linear models were introduced in Asi & Duchi (2019) for the OMD algorithm. Chen et al. (2022b) used gradient flow on the same truncated linear models with a coin-betting algorithm (Orabona & Pál, 2016), but their approach does not seem to satisfy a regret guarantee. Chen et al. (2022a) have used truncated linear models in an FTRL-based parameter-free algorithm (Orabona & Pál, 2021) with a novel decomposition of the regret. However, their approach is ad-hoc and seems difficult to generalize.

2. Definitions and Basic Tools

We define here some basic concepts and tools of convex analysis, we refer the reader to, e.g., Rockafellar (1970); Bauschke & Combettes (2011) for a complete introduction to this topic. We will consider extended value function that can assume infinity values too. A function f is *proper* if it is nowhere $-\infty$ and finite somewhere. A function $f : V \subseteq \mathbb{R}^d \rightarrow [-\infty, +\infty]$ is *closed* if $\{\mathbf{x} : f(\mathbf{x}) \leq \alpha\}$ is closed for every $\alpha \in \mathbb{R}$. For a proper function $f : \mathbb{R}^d \rightarrow (-\infty, +\infty]$, we define a *subgradient* of f in $\mathbf{x} \in \mathbb{R}^d$ as a vector $\mathbf{g} \in \mathbb{R}^d$ that satisfies $f(\mathbf{y}) \geq f(\mathbf{x}) + \langle \mathbf{g}, \mathbf{y} - \mathbf{x} \rangle$, $\forall \mathbf{y} \in \mathbb{R}^d$. We denote the set of subgradients of f in \mathbf{x} by $\partial f(\mathbf{x})$. The *indicator function of the set* V , $i_V : \mathbb{R}^d \rightarrow (-\infty, +\infty]$,

has value 0 for $\mathbf{x} \in V$ and $+\infty$ otherwise. We denote the *dual norm* of a norm $\|\cdot\|$ by $\|\cdot\|_*$. A proper function $f : \mathbb{R}^d \rightarrow (-\infty, +\infty]$ is μ -strongly convex over a convex set $V \subseteq \text{int dom } f$ w.r.t. $\|\cdot\|$ if $\forall \mathbf{x}, \mathbf{y} \in V$ and $\forall \mathbf{g} \in \partial f(\mathbf{x})$, we have $f(\mathbf{y}) \geq f(\mathbf{x}) + \langle \mathbf{g}, \mathbf{y} - \mathbf{x} \rangle + \frac{\mu}{2} \|\mathbf{x} - \mathbf{y}\|^2$. A function $f : V \rightarrow \mathbb{R}$, differentiable in an open set containing V , is L -smooth w.r.t. $\|\cdot\|$ if $f(\mathbf{y}) \leq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{L}{2} \|\mathbf{x} - \mathbf{y}\|^2$ for all $\mathbf{x}, \mathbf{y} \in V$. For a function $f : \mathbb{R}^d \rightarrow [-\infty, \infty]$, we define the *Fenchel conjugate* $f^* : \mathbb{R}^d \rightarrow [-\infty, \infty]$ as $f^*(\boldsymbol{\theta}) = \sup_{\mathbf{x} \in \mathbb{R}^d} \langle \boldsymbol{\theta}, \mathbf{x} \rangle - f(\mathbf{x})$. From this definition, we immediately have the Fenchel-Young inequality: $f(\mathbf{x}) + f^*(\boldsymbol{\theta}) \geq \langle \boldsymbol{\theta}, \mathbf{x} \rangle$, $\forall \mathbf{x}, \boldsymbol{\theta}$. We will also make use of the following properties of Fenchel conjugates.

Theorem 2.1 ((Orabona, 2019, Theorem 5.7)). *Let $f : \mathbb{R}^d \rightarrow (-\infty, +\infty]$ be proper. Then, the following conditions are equivalent:*

- $\boldsymbol{\theta} \in \partial f(\mathbf{x})$.
- $\langle \boldsymbol{\theta}, \mathbf{y} \rangle - f(\mathbf{y})$ achieves its supremum in \mathbf{y} at $\mathbf{y} = \mathbf{x}$.
- $f(\mathbf{x}) + f^*(\boldsymbol{\theta}) = \langle \boldsymbol{\theta}, \mathbf{x} \rangle$.

Moreover, if f is also convex and closed, we have an additional equivalent condition

- $\mathbf{x} \in \partial f^*(\boldsymbol{\theta})$.

Theorem 2.2 ((Orabona, 2019, Theorem 6.11)). *Let $\psi : \mathbb{R}^d \rightarrow (-\infty, +\infty]$ be a proper, closed, convex function, and $\text{dom } \partial \psi$ be non-empty. Then, ψ is $\lambda > 0$ strongly convex w.r.t. $\|\cdot\|$ iff ψ^* is $\frac{1}{\lambda}$ -smooth w.r.t. $\|\cdot\|_*$ on \mathbb{R}^d .*

3. Generalized Implicit FTRL

In this section, we introduce our novel generalized formulation of the implicit FTRL algorithm. The main idea is to depart from the implicit or linearized updates, and directly design updates that improve the upper bound on the regret. More in detail, the basic analysis of most of the online learning algorithms is based on the definition of subgradients:

$$\ell_t(\mathbf{x}_t) - \ell_t(\mathbf{u}) \leq \langle \mathbf{g}_t, \mathbf{x}_t - \mathbf{u} \rangle, \forall \mathbf{g}_t \in \partial \ell_t(\mathbf{x}_t). \quad (1)$$

This allows to study the regret on the linearized losses as a proxy for the regret on the losses ℓ_t . However, we can do better. We introduce a new fundamental and more general strategy: using the Fenchel-Young inequality, we have

$$\ell_t(\mathbf{x}_t) - \ell_t(\mathbf{u}) \leq \ell_t(\mathbf{x}_t) - \langle \mathbf{z}_t, \mathbf{u} \rangle + \ell_t^*(\mathbf{z}_t), \forall \mathbf{z}_t.$$

In particular, the algorithm will choose \mathbf{z}_t to make a certain upper bound involving this quantity to be tighter. This is a better inequality than (1) because when we select $\mathbf{z}_t = \mathbf{g}_t \in \partial \ell_t(\mathbf{x}_t)$, using Theorem 2.1, we recover (1). So, this inequality subsumes the standard one for subgradients, but, using $\mathbf{z}_t \in \partial \ell_t(\mathbf{x}_{t+1})$, it also subsumes the similar inequality used in the implicit case, as we show in Section 3.1. Moreover, we will see in Section 6 that it covers cases where \mathbf{z}_t is not a subgradient of ℓ_t .

Algorithm 1 Generalized Implicit FTRL

Require: Non-empty closed set $V \subseteq \mathbb{R}^d$, a sequence of regularizers $\psi_1, \dots, \psi_T : \mathbb{R}^d \rightarrow (-\infty, +\infty]$

- $\boldsymbol{\theta}_1 = \mathbf{0}$
- for** $t = 1$ **to** T **do**
- Output $\mathbf{x}_t \in \text{argmin}_{\mathbf{x} \in V} \psi_t(\mathbf{x}) - \langle \boldsymbol{\theta}_t, \mathbf{x} \rangle$
- Receive $\ell_t : V \rightarrow \mathbb{R}$ and pay $\ell_t(\mathbf{x}_t)$
- Set $\mathbf{g}_t \in \partial \ell_t(\mathbf{x}_t)$
- Set \mathbf{z}_t such that $H_t(\mathbf{z}_t) \leq H_t(\mathbf{g}_t)$ or $H'_t(\mathbf{z}_t) \leq H'_t(\mathbf{g}_t)$ where H_t and H'_t are defined in (2) and (3)
- Set $\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t - \mathbf{z}_t$
- end for**

The analysis shows that the optimal setting of \mathbf{z}_t is the one that minimizes the function

$$H_t(\mathbf{z}) \triangleq \psi_{t+1, V}^*(\boldsymbol{\theta}_t - \mathbf{z}) + \ell_t^*(\mathbf{z}) \quad (2)$$

or

$$H'_t(\mathbf{z}) \triangleq \psi_{t, V}^*(\boldsymbol{\theta}_t - \mathbf{z}) + \ell_t^*(\mathbf{z}), \quad (3)$$

where $\psi_{t, V}$ is the restriction of the regularizer used at time t on the feasible set V , i.e., $\psi_{t, V} \triangleq \psi_t + i_V$. However, we can show that any setting of \mathbf{z}_t that guarantees $H_t(\mathbf{z}_t) < H_t(\mathbf{g}_t)$ (or $H'_t(\mathbf{z}_t) < H'_t(\mathbf{g}_t)$) guarantee a strict improvement in the worst-case regret w.r.t. using the linearized losses.

One might wonder why the need for two different updates using H_t or H'_t . The reason is that when using time-varying regularizers that depend on the data, like in the FTRL version of AdaGrad (McMahan & Streeter, 2010; Duchi et al., 2011), if λ_{t+1} depends on \mathbf{z}_t it might make the calculation of the update particularly difficult. This can be avoided using the update involving H'_t .

Once we have the \mathbf{z}_t , we treat them as the subgradient of surrogate linear losses. So, putting it all together, Algorithm 1 shows the final algorithm. We now show a regret guarantee for this algorithm. First, we state a general Lemma and then instantiate it in a few interesting cases.

Theorem 3.1. *Let $V \subseteq \mathbb{R}^d$ be closed and non-empty and $\psi_t : V \rightarrow \mathbb{R}$. With the notation in Algorithm 1, define by $F_t(\mathbf{x}) = \psi_t(\mathbf{x}) + \sum_{i=1}^{t-1} \langle \mathbf{z}_i, \mathbf{x} \rangle$, so that $\mathbf{x}_t \in \text{argmin}_{\mathbf{x} \in V} F_t(\mathbf{x})$. Finally, assume that $\text{argmin}_{\mathbf{x} \in V} F_t(\mathbf{x})$ and $\partial \ell_t(\mathbf{x}_t)$ are not empty for all t .*

- For any $\mathbf{z}_t \in \mathbb{R}^d$ and any $\mathbf{u} \in \mathbb{R}^d$, we have

$$\begin{aligned} \text{Regret}_T(\mathbf{u}) &\leq \psi_{T+1}(\mathbf{u}) - \min_{\mathbf{x} \in V} \psi_1(\mathbf{x}) \\ &\quad + \sum_{t=1}^T [\psi_{t+1, V}^*(\boldsymbol{\theta}_t - \mathbf{g}_t) - \psi_{t, V}^*(\boldsymbol{\theta}_t) + \langle \mathbf{x}_t, \mathbf{g}_t \rangle - \delta_t] \\ &\quad + F_{T+1}(\mathbf{x}_{T+1}) - F_{T+1}(\mathbf{u}), \end{aligned}$$

where $\delta_t \triangleq H_t(\mathbf{g}_t) - H_t(\mathbf{z}_t)$.

- If $\psi_{t+1}(\mathbf{x}) \geq \psi_t(\mathbf{x})$ for any $\mathbf{x} \in V$, then, for any $\mathbf{z}_t \in \mathbb{R}^d$, we have

$$\begin{aligned} \text{Regret}_T(\mathbf{u}) &\leq \psi_{T+1}(\mathbf{u}) - \min_{\mathbf{x} \in V} \psi_1(\mathbf{x}) \\ &\quad + \sum_{t=1}^T [\psi_{t,V}^*(\boldsymbol{\theta}_t - \mathbf{g}_t) - \psi_{t,V}^*(\boldsymbol{\theta}_t) + \langle \mathbf{x}_t, \mathbf{g}_t \rangle - \delta'_t] \\ &\quad + F_{T+1}(\mathbf{x}_{T+1}) - F_{T+1}(\mathbf{u}), \end{aligned}$$

where $\delta'_t \triangleq H'_t(\mathbf{g}_t) - H'_t(\mathbf{z}_t)$.

Proof. The proof is composed of simple but not obvious steps. The first important observation is that the definition of \mathbf{x}_t in the algorithm corresponds exactly to the one of FTRL on the linear losses $\langle \mathbf{z}_t, \cdot \rangle$. Hence, we can use the FTRL equality in Orabona (2019, Lemma 7.1):

$$\begin{aligned} - \sum_{t=1}^T \langle \mathbf{z}_t, \mathbf{u} \rangle &= \sum_{t=1}^T [F_t(\mathbf{x}_t) - F_{t+1}(\mathbf{x}_{t+1})] + \psi_{T+1}(\mathbf{u}) \\ &\quad - \min_{\mathbf{x} \in V} \psi_1(\mathbf{x}) + F_{T+1}(\mathbf{x}_{T+1}) - F_{T+1}(\mathbf{u}), \end{aligned}$$

where we have cancelled out the terms $\langle \mathbf{z}_t, \mathbf{x}_t \rangle$ on both sides.

Now, use Fenchel-Young inequality, to have $\langle \mathbf{z}_t, \mathbf{u} \rangle \leq \ell_t(\mathbf{u}) + \ell_t^*(\mathbf{z}_t)$. Hence, we have

$$\begin{aligned} - \sum_{t=1}^T \ell_t(\mathbf{u}) &\leq \sum_{t=1}^T [F_t(\mathbf{x}_t) - F_{t+1}(\mathbf{x}_{t+1}) + \ell_t^*(\mathbf{z}_t)] \\ &\quad + \psi_{T+1}(\mathbf{u}) - \min_{\mathbf{x} \in V} \psi_1(\mathbf{x}) \\ &\quad + F_{T+1}(\mathbf{x}_{T+1}) - F_{T+1}(\mathbf{u}). \end{aligned}$$

Observe that

$$\begin{aligned} F_t(\mathbf{x}_t) &= \min_{\mathbf{x} \in V} \psi_t(\mathbf{x}) + \sum_{i=1}^{t-1} \langle \mathbf{z}_i, \mathbf{x} \rangle \\ &= - \max_{\mathbf{x} \in V} \langle \boldsymbol{\theta}_t, \mathbf{x} \rangle - \psi_t(\mathbf{x}) = -\psi_{t,V}^*(\boldsymbol{\theta}_t). \end{aligned}$$

In the same way, we have $-F_{t+1}(\mathbf{x}_{t+1}) = \psi_{t+1,V}^*(\boldsymbol{\theta}_{t+1})$. Also, for any $\mathbf{g}_t \in \partial \ell_t(\mathbf{x}_t)$, by Theorem 2.1 we have $\ell_t^*(\mathbf{g}_t) = \langle \mathbf{x}_t, \mathbf{g}_t \rangle - \ell_t(\mathbf{x}_t)$. Hence, each term in the sum can be written as

$$\begin{aligned} F_t(\mathbf{x}_t) - F_{t+1}(\mathbf{x}_{t+1}) + \ell_t^*(\mathbf{z}_t) &= \psi_{t+1,V}^*(\boldsymbol{\theta}_{t+1}) - \psi_{t,V}^*(\boldsymbol{\theta}_t) + \ell_t^*(\mathbf{z}_t) \\ &= H_t(\mathbf{z}_t) - \psi_{t,V}^*(\boldsymbol{\theta}_t). \end{aligned}$$

Now, we just add and subtract $H_t(\mathbf{g}_t) = \psi_{t+1,V}^*(\boldsymbol{\theta}_t - \mathbf{g}_t) + \langle \mathbf{g}_t, \mathbf{x}_t \rangle - \ell_t(\mathbf{x}_t)$ to obtain the stated bound.

The second case is similar. We just have to observe that if $\psi_{t+1,V} \geq \psi_{t,V}$, then $\psi_{t+1,V}^* \leq \psi_{t,V}^*$. Hence, each term in

the sum can be upper bounded as

$$\begin{aligned} F_t(\mathbf{x}_t) - F_{t+1}(\mathbf{x}_{t+1}) + \ell_t^*(\mathbf{z}_t) &\leq \psi_{t,V}^*(\boldsymbol{\theta}_{t+1}) - \psi_{t,V}^*(\boldsymbol{\theta}_t) + \ell_t^*(\mathbf{z}_t) \\ &= H'_t(\mathbf{z}_t) - \psi_{t,V}^*(\boldsymbol{\theta}_t). \end{aligned}$$

As before, adding and subtracting $H'_t(\mathbf{g}_t) = \psi_{t,V}^*(\boldsymbol{\theta}_t - \mathbf{g}_t) + \langle \mathbf{x}_t, \mathbf{g}_t \rangle - \ell_t(\mathbf{x}_t)$ gives the stated bound. \square

The Theorem is stated with very weak assumption to show its generality, but it is immediate to obtain concrete regret guarantees just assuming, for example, strongly convex regularizers and convex and Lipschitz losses and using well-known methods as Orabona (2019, Lemma 7.8)

However, we can already understand why this is an interesting guarantee. Let's first consider the case that $\mathbf{z}_t = \mathbf{g}_t$. In this case, we exactly recover the linearized FTRL algorithm. Even the guarantee in the Theorem exactly recovers the best known one (Orabona, 2019, Corollary 7.9), with $\delta_t = 0$ and $\delta'_t = 0$. Now, if we set \mathbf{z}_t such that $H_t(\mathbf{z}_t) < H_t(\mathbf{g}_t)$ or $H'_t(\mathbf{z}_t) < H'_t(\mathbf{g}_t)$ we will have that $\delta_t > 0$ or $\delta'_t > 0$. Hence, in each single term of the sum we have a negative factor that makes the regret bound smaller. While it might be difficult to give a lower bound to δ_t and δ'_t without additional assumptions, the main value of this analysis is in giving a *unifying way to design generalized implicit updates for FTRL*. In fact, in the next sections we will show a number of possibilities that this framework enables.

Next, we will gain more understanding on the updates in Algorithm 1, comparing them to implicit OMD.

3.1. Comparison with Implicit Online Mirror Descent

In this section, we show that when \mathbf{z}_t is set to minimize $H_t(\mathbf{z})$ or $H'_t(\mathbf{z})$, we recover different variants of implicit updates.

Assume that the ℓ_t are closed and convex. Also, assume that $\psi_{t,V}^*$ is differentiable, that is true, for example, when ψ_t is strongly convex by Theorem 2.2. Then, observe that by the first-order optimality condition and Theorem 2.1, we have

$$\begin{aligned} \mathbf{z}_t &= \underset{\mathbf{z}}{\operatorname{argmin}} H_t(\mathbf{z}) \\ &\Leftrightarrow \nabla \psi_{t+1,V}^*(\boldsymbol{\theta}_t - \mathbf{z}_t) \in \partial \ell_t^*(\mathbf{z}_t) \\ &\Leftrightarrow \mathbf{z}_t \in \partial \ell_t(\nabla \psi_{t+1,V}^*(\boldsymbol{\theta}_t - \mathbf{z}_t)) = \partial \ell_t(\mathbf{x}_{t+1}). \end{aligned} \quad (4)$$

Hence, in this case, we have that the optimal \mathbf{z}_t is the gradient at the *next* point \mathbf{x}_{t+1} . This is exactly what happens in the implicit updates. This condition that depends on the gradient in the next iterate is similar to other implicit updates.

Under the same assumptions, we also have

$$\begin{aligned} \mathbf{z}_t = \underset{\mathbf{z}}{\operatorname{argmin}} H'_t(\mathbf{z}) &\Leftrightarrow \nabla \psi_{t,V}^*(\boldsymbol{\theta}_t - \mathbf{z}_t) \in \partial \ell_t^*(\mathbf{z}_t) \\ &\Leftrightarrow \mathbf{z}_t \in \partial \ell_t(\nabla \psi_{t,V}^*(\boldsymbol{\theta}_{t+1})). \end{aligned} \quad (5)$$

In this other case, the update also has an implicit flavor but the subgradient is queried on a point different from the next point, where the difference depends on how much $\nabla \psi_{t,V}^*$ differs from $\nabla \psi_{t+1,V}^*$.

Let's see this connection even more precisely, considering *proximal updates*. Hence, for simplicity, let's consider the case that $V = \mathbb{R}^d$, similar considerations hold in the constrained case. Consider the case that $\psi_t(\mathbf{x}) = \frac{\lambda_t}{2} \|\mathbf{x}\|_2^2$. In this case, the update can be written with the *proximal operator* of the loss functions. In particular, the proximal operator of ηf , is defined as

$$\operatorname{Prox}_{\eta f}(\mathbf{y}) \triangleq \underset{\mathbf{x} \in \mathbb{R}^d}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{x} - \mathbf{y}\|_2^2 + \eta f(\mathbf{x}).$$

If the function f is differentiable we have that $\operatorname{Prox}_{\eta f}(\mathbf{y}) = \mathbf{y} - \eta \nabla f(\operatorname{Prox}_{\eta f}(\mathbf{y}))$. In words, the proximal update moves by a quantity that depends on the gradient on the updated point. The implicit nature of these updates justifies the name ‘‘implicit updates’’ used in the online learning literature. More generally, we have that $\operatorname{Prox}_{\eta f}(\mathbf{y}) \in \mathbf{y} - \eta \partial f(\operatorname{Prox}_{\eta f}(\mathbf{y}))$. We list some common proximal operators in Appendix A.

Assuming λ_{t+1} does not depend on \mathbf{z}_t , using the proximal operator we can rewrite the update in (4) as

$$\begin{aligned} \mathbf{x}_{t+1} &= \frac{\boldsymbol{\theta}_{t+1}}{\lambda_{t+1}} = \operatorname{Prox}_{\frac{\ell_t}{\lambda_{t+1}}} \left(\frac{\boldsymbol{\theta}_t}{\lambda_{t+1}} \right) \\ &= \operatorname{Prox}_{\frac{\ell_t}{\lambda_{t+1}}} \left(\frac{\lambda_t \mathbf{x}_t}{\lambda_{t+1}} \right). \end{aligned} \quad (6)$$

Similarly, we can rewrite the update in (5) as

$$\begin{aligned} \frac{\boldsymbol{\theta}_{t+1}}{\lambda_t} &= \frac{\boldsymbol{\theta}_t}{\lambda_t} - \frac{\mathbf{z}_t}{\lambda_t} = \mathbf{x}_t - \frac{\mathbf{z}_t}{\lambda_t} \in \mathbf{x}_t - \frac{1}{\lambda_t} \partial \ell_t(\nabla \psi_{t,V}^*(\boldsymbol{\theta}_{t+1})) \\ &= \mathbf{x}_t - \frac{1}{\lambda_t} \partial \ell_t \left(\frac{\boldsymbol{\theta}_{t+1}}{\lambda_t} \right). \end{aligned}$$

Hence, we have that $\frac{\boldsymbol{\theta}_{t+1}}{\lambda_t} = \operatorname{Prox}_{\frac{\ell_t}{\lambda_t}}(\mathbf{x}_t)$ and we get

$$\mathbf{x}_{t+1} = \frac{\boldsymbol{\theta}_{t+1}}{\lambda_{t+1}} = \frac{\lambda_t}{\lambda_{t+1}} \operatorname{Prox}_{\frac{\ell_t}{\lambda_t}}(\mathbf{x}_t). \quad (7)$$

It is instructive to compare both updates with the one of Implicit Online Mirror Descent using $\psi(\mathbf{x}) = \frac{1}{2} \|\mathbf{x}\|_2^2$ as distance generating function and stepsizes $\frac{1}{\lambda_t}$. In this case, we would update with

$$\begin{aligned} \mathbf{x}_{t+1} &= \underset{\mathbf{x}}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{x}_t - \mathbf{x}\|_2^2 + \frac{1}{\lambda_t} \ell_t(\mathbf{x}) \\ &= \operatorname{Prox}_{\frac{\ell_t}{\lambda_t}}(\mathbf{x}_t). \end{aligned} \quad (8)$$

Comparing (4) and (5) to (8), we see, when $\lambda_t \leq \lambda_{t+1}$ as it is usual, the two updates above shrink a bit towards the zero vector, that is the initial point \mathbf{x}_1 , before or after the proximal operator. This shrinking is given by the FTRL update and it is the key difference with Implicit OMD update. The different update also corresponds to a different guarantee: the regret of the generalized implicit FTRL holds for unbounded domains too, while in Implicit OMD with time-varying stepsizes can have linear regret on unbounded domains (Orabona & Pál, 2018). Interestingly, a similar shrinking has been proposed in Fang et al. (2020) to fix the unbounded issue in OMD. Clearly, the updates (4) and (5) become equivalent to (8) for λ_t constant in t , that is exactly the only case when implicit/proximal online mirror descent works for unbounded domains.

4. Temporal Variability Bound

In this section, we quantify the advantage of the generalized implicit FTRL updates in the case of slow temporal variability of the loss functions.

It was observed in Campolongo & Orabona (2020) that implicit OMD satisfies regret guarantees that depend on the temporal Variability V_T :

$$V_T \triangleq \sum_{t=2}^T \max_{\mathbf{x} \in V} \ell_t(\mathbf{x}) - \ell_{t-1}(\mathbf{x}).$$

In Campolongo & Orabona (2020, Appendix E) they also show that FTRL with full losses guarantees a similar guarantee, but at a much higher computational price. Indeed, FTRL with full losses requires solving a finite sum optimization problem at each step, whose size increases with the number of iterations. Such computational burden induced Campolongo & Orabona (2020) to say that such approach is ‘‘not worth of pursuing.’’

Here, we show that the Algorithm 1 can satisfy the same guarantee of implicit OMD with the same computational complexity too. First, we show the following Lemma.

Lemma 4.1. *Under the assumptions of Theorem 3.1, further assume V to be convex, $\psi_t : V \rightarrow \mathbb{R}$ closed, λ_t -strongly convex w.r.t. $\|\cdot\|$, and subdifferentiable in V , ℓ_t closed, convex, and subdifferentiable in V , and $\lambda_{t+1} \geq \lambda_t$. Set $\mathbf{z}_t \in \underset{\mathbf{z}}{\operatorname{argmin}} H_t(\mathbf{z})$. Then, we have*

$$\begin{aligned} \operatorname{Regret}_T(\mathbf{u}) &\leq \psi_{T+1}(\mathbf{u}) - \min_{\mathbf{x} \in V} \psi_1(\mathbf{x}) \\ &+ \sum_{t=1}^T \left(\ell_t(\mathbf{x}_t) - \ell_t(\mathbf{x}_{t+1}) - \frac{\lambda_t}{2} \|\mathbf{x}_{t+1} - \mathbf{x}_t\|_2^2 \right), \forall \mathbf{u} \in V. \end{aligned}$$

Proof. First of all, the existence and unicity of \mathbf{x}_t is guaranteed by ψ_t being closed and strongly convex (see, e.g., Orabona, 2019, Theorem 6.8).

From Theorem 2.1, for any $\mathbf{g}'_t \in \partial \ell_t(\mathbf{x}_{t+1})$, we have $\ell_t^*(\mathbf{g}'_t) = \langle \mathbf{x}_{t+1}, \mathbf{g}'_t \rangle - \ell_t(\mathbf{x}_{t+1})$. Hence, from (4), we have

$$\begin{aligned} & \psi_{t+1,V}^*(\boldsymbol{\theta}_{t+1}) - \psi_{t,V}^*(\boldsymbol{\theta}_t) + \ell_t^*(\mathbf{z}_t) \\ &= \psi_{t+1,V}^*(\boldsymbol{\theta}_t - \mathbf{z}_t) - \psi_{t,V}^*(\boldsymbol{\theta}_t) + \langle \mathbf{x}_{t+1}, \mathbf{z}_t \rangle - \ell_t(\mathbf{x}_{t+1}). \end{aligned}$$

Using this identity, we have

$$\begin{aligned} & \psi_{t+1,V}^*(\boldsymbol{\theta}_t - \mathbf{z}_t) - \psi_{t,V}^*(\boldsymbol{\theta}_t) + \langle \mathbf{x}_{t+1}, \mathbf{z}_t \rangle \\ &= \langle \boldsymbol{\theta}_t - \mathbf{z}_t, \mathbf{x}_{t+1} \rangle - \psi_{t+1}(\mathbf{x}_{t+1}) - \langle \boldsymbol{\theta}_t, \mathbf{x}_t \rangle + \psi_t(\mathbf{x}_t) \\ & \quad + \langle \mathbf{x}_{t+1}, \mathbf{z}_t \rangle \\ & \leq \psi_t(\mathbf{x}_t) - \psi_t(\mathbf{x}_{t+1}) + \langle \boldsymbol{\theta}_t, \mathbf{x}_{t+1} - \mathbf{x}_t \rangle. \end{aligned}$$

From the first-order optimality condition of \mathbf{x}_t , we have that $\boldsymbol{\theta}_t \in \partial \psi_t(\mathbf{x}_t) + \partial i_V(\mathbf{x}_t)$. Moreover, for all $\mathbf{g}''_t \in \partial i_V(\mathbf{x}_t)$, by definition we have $\langle \mathbf{g}''_t, \mathbf{y} - \mathbf{x}_t \rangle \leq 0$ for all $\mathbf{y} \in V$. Hence, for $\mathbf{g}'_t \in \partial \psi_t(\mathbf{x}_t)$ and $\mathbf{g}''_t \in \partial i_V(\mathbf{x}_t)$ such that $\boldsymbol{\theta}_t = \mathbf{g}'_t + \mathbf{g}''_t$, we have

$$\begin{aligned} & \psi_t(\mathbf{x}_t) - \psi_t(\mathbf{x}_{t+1}) + \langle \boldsymbol{\theta}_t, \mathbf{x}_{t+1} - \mathbf{x}_t \rangle \\ &= \psi_t(\mathbf{x}_t) - \psi_t(\mathbf{x}_{t+1}) + \langle \mathbf{g}'_t + \mathbf{g}''_t, \mathbf{x}_{t+1} - \mathbf{x}_t \rangle \\ & \leq -\frac{\lambda_t}{2} \|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2, \end{aligned}$$

where in the inequality we also used the strong convexity of ψ_t . Using this inequality in Theorem 3.1 and summing over time, we have

$$\begin{aligned} & \sum_{t=1}^T \ell_t(\mathbf{x}_{t+1}) - \sum_{t=1}^T \ell_t(\mathbf{u}) \\ & \leq \psi_{T+1}(\mathbf{u}) - \min_{\mathbf{x} \in V} \psi_1(\mathbf{x}) - \sum_{t=1}^T \frac{\lambda_t}{2} \|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2. \end{aligned}$$

By adding and subtracting $\sum_{t=1}^T \ell_t(\mathbf{x}_t)$ to both sides and reordering the terms, we have the stated bound. \square

This Lemma mirrors Theorem 5.2 in Campolongo & Orabona (2020), with the important difference that here we do not need the Bregman divergence to be bounded on the feasible set V , thanks to the use of FTRL instead of OMD. We can now state the immediate corollary on a regret bound that depends on the temporal variation.

Corollary 4.2. *Under the assumptions of Lemma 4.1, for any $\mathbf{u} \in V$, we have*

$$\begin{aligned} \text{Regret}_T(\mathbf{u}) & \leq \psi_{T+1}(\mathbf{u}) - \min_{\mathbf{x} \in V} \psi_1(\mathbf{x}) \\ & \quad + \ell_1(\mathbf{x}_1) - \ell_T(\mathbf{x}_{T+1}) + V_T. \end{aligned}$$

From this result, following (Campolongo & Orabona, 2020), it is relatively easy to obtain the following adaptive regret guarantee. The only difficulty is the fact that we need ψ_{t+1} to be independent of \mathbf{z}_t to have a simpler update rule. We

solve this problem using an increasing regularizer that is “behind of two steps”. In this way, we have that λ_{t+1} depends on quantities that are all known at the beginning of round t . The proof is in Appendix B.

Corollary 4.3. *Under the assumptions of Lemma 4.1, further assume $\|\mathbf{g}_t\|_* \leq G$ for all t . Define $\gamma_t = \ell_t(\mathbf{x}_t) - \ell_t(\mathbf{x}_{t+1}) - \frac{\lambda_t}{2} \|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2$ and $\lambda_t = \frac{1}{\beta^2} \left(G\beta + \sum_{i=1}^{t-2} \gamma_i \right)$. Assume that ψ is closed and 1-strongly convex w.r.t. $\|\cdot\|$ and set $\psi_t = \lambda_t \psi$. Then, for any $\mathbf{u} \in V$, we have*

$$\begin{aligned} \text{Regret}_T(\mathbf{u}) & \leq \min \left(\frac{1}{\beta} (\ell_1(\mathbf{x}_1) - \ell_T(\mathbf{x}_{T+1}) + V_T), \right. \\ & \quad \left. G + \sqrt{\frac{5}{4} \sum_{t=1}^T \|\mathbf{g}_t\|_*^2} \right) \left(\frac{\psi(\mathbf{u})}{\beta} + \beta \right). \end{aligned}$$

5. Two-step Updates

The choice of \mathbf{z}_t that minimizes the regret upper bound requires solving the optimization problem $\min_{\mathbf{z}} H_t(\mathbf{z})$ or $\min_{\mathbf{z}} H'_t(\mathbf{z})$. We have seen in Section 3.1 that this corresponds to (some variant) of a implicit/proximal update and, depending on ℓ_t , it can be of difficult calculation. However, as we said, any choice better than \mathbf{g}_t will cause a provable gain. Hence, a viable solution is to *approximately* solve for the optimal \mathbf{z}_t .

Here, we propose a simple approximation: set \mathbf{z}_t as

$$\mathbf{z}_t \in \partial \ell_t(\nabla \psi_{t+1,V}^*(\boldsymbol{\theta}_t - \mathbf{g}_t)) \quad (9)$$

or as

$$\mathbf{z}_t \in \partial \ell_t(\nabla \psi_{t,V}^*(\boldsymbol{\theta}_t - \mathbf{g}_t)). \quad (10)$$

In words, we set \mathbf{z}_t to be a subgradient after one fake update. This is exactly the approach used in the Mirror-Prox algorithm (Nemirovski, 2004), an offline optimization algorithm. In the next theorem, when the loss functions ℓ_t are smooth and the regularizer is chosen appropriately, we show that this choice can be used in the generalized implicit FTRL too and it cannot be worse than using \mathbf{g}_t .

Theorem 5.1. *Assume $\psi_t(\mathbf{x})$ proper, closed, and λ_t -strongly convex with respect to $\|\cdot\|$. Assume $\ell_t(\mathbf{x})$ closed and λ_t -smooth w.r.t. $\|\cdot\|_*$ for all t . Then, using (9) and assuming $\lambda_{t+1} \geq L_t$, we have $H_t(\mathbf{z}_t) \leq H_t(\mathbf{g}_t)$. On the other hand, when using (10) and assuming $\lambda_t \geq L_t$ we have $H'_t(\mathbf{z}_t) \leq H'_t(\mathbf{g}_t)$.*

Proof. We only prove that statement for (9), the other one is similar. We would like to prove that

$$\begin{aligned} H_t(\mathbf{z}_t) &= \psi_{t+1,V}^*(\boldsymbol{\theta}_t - \mathbf{z}_t) + \ell_t^*(\mathbf{z}_t) \\ & \leq \psi_{t+1,V}^*(\boldsymbol{\theta}_t - \mathbf{g}_t) + \ell_t^*(\mathbf{g}_t) = H_t(\mathbf{g}_t). \end{aligned}$$

This is equivalent to prove

$$\psi_{t+1,V}^*(\boldsymbol{\theta}_t - \mathbf{z}_t) - \psi_{t+1,V}^*(\boldsymbol{\theta}_t - \mathbf{g}_t) \leq \ell_t^*(\mathbf{g}_t) - \ell_t^*(\mathbf{z}_t).$$

Given that $\psi_{t+1}(\mathbf{x}_t)$ is λ_{t+1} -strongly convex, by Theorem 2.2, we have $\psi_t^*(\boldsymbol{\theta})$ is $1/\lambda_{t+1}$ -smooth with respect to $\|\cdot\|_*$. By the definition of smoothness, we have

$$\begin{aligned} & \psi_{t+1,V}^*(\boldsymbol{\theta}_t - \mathbf{z}_t) - \psi_{t+1,V}^*(\boldsymbol{\theta}_t - \mathbf{g}_t) \\ & \leq \langle \nabla \psi_{t+1}^*(\boldsymbol{\theta}_t - \mathbf{g}_t), \mathbf{g}_t - \mathbf{z}_t \rangle + \frac{1}{2\lambda_{t+1}} \|\mathbf{g}_t - \mathbf{z}_t\|_*^2. \end{aligned}$$

Given that $\ell_t(\mathbf{x}_t)$ is L_t -smooth w.r.t $\|\cdot\|_*$, by Theorem 2.2 $\ell_t^*(\mathbf{g})$ is $1/L_t$ strongly convex w.r.t $\|\cdot\|$. So, by the definition of the strong convexity, we have

$$\ell_t^*(\mathbf{g}_t) - \ell_t^*(\mathbf{z}_t) \geq \langle \mathbf{q}_t, \mathbf{g}_t - \mathbf{z}_t \rangle + \frac{1}{2L_t} \|\mathbf{g}_t - \mathbf{z}_t\|_*^2,$$

for all $\mathbf{q}_t \in \partial \ell_t^*(\mathbf{z}_t)$. Defining $\mathbf{x}'_{t+1} \triangleq \nabla \psi_{t+1}^*(\boldsymbol{\theta}_t - \mathbf{g}_t)$, by Theorem 2.1, we have $\mathbf{x}'_{t+1} \in \partial \ell_t^*(\mathbf{z}_t)$. Hence, we can select \mathbf{q}_t such that $\mathbf{x}'_{t+1} = \mathbf{q}_t$. Finally, using the assumption on $\lambda_{t+1} \geq L_t$, we have the stated bound. \square

6. Going Beyond Subgradients with aProx

Till now, in all the updates we have considered \mathbf{z}_t was set to be a subgradient of ℓ_t in a specific point. In this section, we show that we can go beyond this idea.

Asi & Duchi (2019) introduced aProx updates, that is proximal updates on surrogate loss functions. In particular, they used truncated linear lower bounds to the loss functions as surrogate functions. These simple surrogates are motivated by the fact that they are strictly better than linear approximation and at the same time they allow writing the proximal update in a closed form. Moreover, they showed empirically that in certain situations the performance of the algorithms becomes much more resistant to the tuning of the stepsizes.

One might just use the same truncated lower bounds in implicit FTRL, but it would not be clear why this should give any advantage in the theoretical bound. Indeed, even in Asi & Duchi (2019) it is not completely clear what part of the theory tells us that we should expect a better performance from these updates.

Here, we show how the *updates in the generalized implicit FTRL are actually a generalization of the aProx ones*. In particular, we generalize the aProx updates to arbitrary regularizers and show that all of them satisfy $H_t(\mathbf{z}_t) \leq H_t(\mathbf{g}_t)$ and $H'_t(\mathbf{z}_t) \leq H'_t(\mathbf{g}_t)$. In words, the aProx updates are guaranteed to be at least as good as the subgradient \mathbf{g}_t in minimizing the worst-case regret.

In order to consider truncated linear lower bounds to the functions ℓ_t , in this section we will assume that the loss

functions ℓ_t are lower bounded. Given that the regret is invariant to additive constants in the losses, without loss of generality we can assume the lower bound to be 0 for all the loss functions. Hence, define the truncated linear model $\hat{\ell}_t : V \rightarrow \mathbb{R}$ around \mathbf{x}_t to be

$$\hat{\ell}_t(\mathbf{x}) \triangleq \max(\ell_t(\mathbf{x}_t) + \langle \mathbf{g}_t, \mathbf{x} - \mathbf{x}_t \rangle, 0),$$

where $\mathbf{g}_t \in \partial \ell_t(\mathbf{x}_t)$. For brevity of notation, our notation does not stress the fact that the truncated linear model depends on \mathbf{x}_t and the specific subgradient \mathbf{g}_t .

The idea to extend aProx to the case of generalized implicit FTRL, we use the truncated linear lower bound in the update of \mathbf{z}_t . So, we define

$$\mathbf{z}_t = \underset{\mathbf{z}}{\operatorname{argmin}} \psi_{t+1,V}^*(\boldsymbol{\theta}_t - \mathbf{z}_t) + \hat{\ell}_t^*(\mathbf{z}_t) \quad (11)$$

or

$$\mathbf{z}_t = \underset{\mathbf{z}}{\operatorname{argmin}} \psi_{t,V}^*(\boldsymbol{\theta}_t - \mathbf{z}_t) + \hat{\ell}_t^*(\mathbf{z}_t). \quad (12)$$

Theorem 6.1. *Assume the loss functions $\ell_t : V \rightarrow \mathbb{R}$ to be convex, closed, and subdifferentiable in V for all t . Set \mathbf{z}_t using (11) or (12). Then, we have that $H_t(\mathbf{z}_t) \leq H_t(\mathbf{g}_t)$ or $H'_t(\mathbf{z}_t) \leq H'_t(\mathbf{g}_t)$ respectively.*

Proof. We consider the update (11), the other case is very similar and we omit it.

First, we derive some inequalities on the quantities of interest. From Theorem 2.1, given that $\mathbf{g}_t \in \partial \ell_t(\mathbf{x}_t)$ and $\mathbf{g}_t \in \partial \ell_t(\mathbf{x}_t)$ we have both $\ell_t(\mathbf{x}_t) + \ell^*(\mathbf{g}_t) = \langle \mathbf{g}_t, \mathbf{x}_t \rangle$ and $\hat{\ell}_t(\mathbf{x}_t) + \hat{\ell}^*(\mathbf{g}_t) = \langle \mathbf{g}_t, \mathbf{x}_t \rangle$. Moreover, given that $\hat{\ell}_t(\mathbf{x}) \leq \ell_t(\mathbf{x})$ for any \mathbf{x} , we have $\hat{\ell}_t^*(\mathbf{z}) \geq \ell_t^*(\mathbf{z})$ for any \mathbf{z} . Finally, by the definition of truncated linear lower bound, we have $\ell_t(\mathbf{x}_t) = \hat{\ell}_t(\mathbf{x}_t)$.


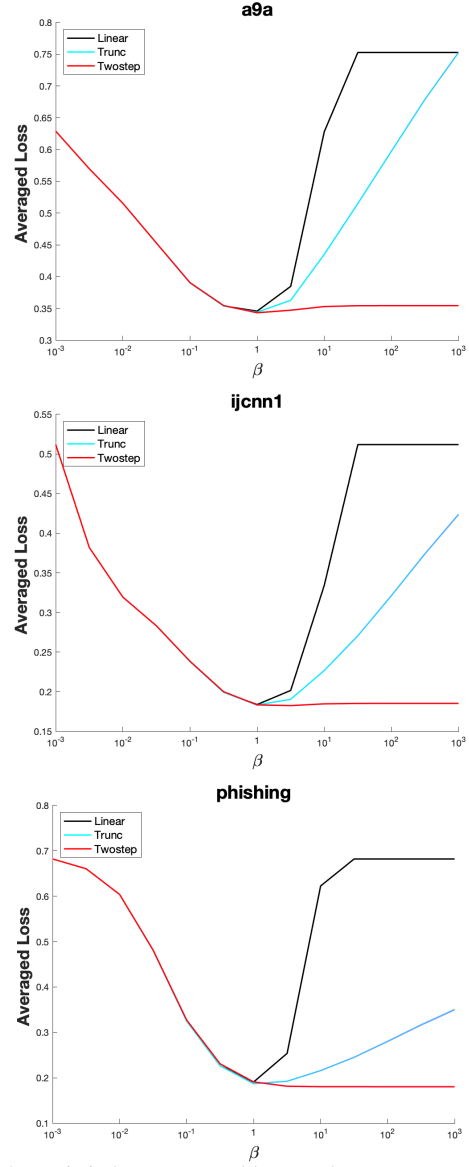
Hence, we have

$$\begin{aligned} & \psi_{t+1,V}^*(\boldsymbol{\theta}_t - \mathbf{z}_t) + \ell_t^*(\mathbf{z}_t) \\ & \leq \psi_{t+1,V}^*(\boldsymbol{\theta}_t - \mathbf{z}_t) + \hat{\ell}_t^*(\mathbf{z}_t) \\ & = \min_{\mathbf{z}} \psi_{t+1,V}^*(\boldsymbol{\theta}_t - \mathbf{z}) + \hat{\ell}_t^*(\mathbf{z}) \\ & \leq \psi_{t+1,V}^*(\boldsymbol{\theta}_t - \mathbf{g}_t) + \hat{\ell}_t^*(\mathbf{g}_t) \\ & = \psi_{t+1,V}^*(\boldsymbol{\theta}_t - \mathbf{g}_t) + \langle \mathbf{g}_t, \mathbf{x}_t \rangle - \hat{\ell}_t(\mathbf{x}_t) \\ & = \psi_{t+1,V}^*(\boldsymbol{\theta}_t - \mathbf{g}_t) + \langle \mathbf{g}_t, \mathbf{x}_t \rangle - \ell_t(\mathbf{x}_t) \\ & = \psi_{t+1,V}^*(\boldsymbol{\theta}_t - \mathbf{g}_t) + \ell_t^*(\mathbf{g}_t) = H_t(\mathbf{g}_t). \quad \square \end{aligned}$$

We can also immediately write closed form updates for generalized implicit FTRL with regularizer $\psi_t(\mathbf{x}) = \frac{\lambda_t}{2} \|\mathbf{x}\|^2$, that mirror the ones of aProx. The proof is in Section C.

Corollary 6.2. *Set $\psi_t = \frac{\lambda_t}{2} \|\mathbf{x}\|_2^2$ and $\mathbf{g}_t \in \partial \ell_t(\mathbf{x}_t)$. Setting \mathbf{z}_t as in (11), we have that the update of generalized implicit FTRL is*

$$\mathbf{x}_{t+1} = \frac{\lambda_t}{\lambda_{t+1}} \mathbf{x}_t - \min \left(\frac{1}{\lambda_{t+1}}, \frac{\ell_t(\mathbf{x}_t)}{\|\mathbf{g}_t\|^2} \right) \mathbf{g}_t.$$


 Figure 1. Hinge loss, averaged loss vs. hyperparameter β .

 Figure 2. Logistic loss, averaged loss vs. hyperparameter β .

On the other hand, setting z_t as in (12), the update is

$$\mathbf{x}_{t+1} = \frac{\lambda_t}{\lambda_{t+1}} \mathbf{x}_t - \min \left(\frac{1}{\lambda_{t+1}}, \frac{\lambda_t}{\lambda_{t+1}} \frac{\ell_t(\mathbf{x}_t)}{\|\mathbf{g}_t\|^2} \right) \mathbf{g}_t.$$

7. Empirical Evaluation

As we said, in the worst case scenario any kind of implicit update cannot give any advantage over the usual updates. However, in practice it is well-known that things are vastly different. Hence, in this section, we compare the performance of different choices of z_t in Algorithm 1 when $\psi_t(\mathbf{x}) = \frac{\lambda_t}{2} \|\mathbf{x}\|_2^2$. In particular, we consider:

- FTRL with linearized losses (Linear): $z_t = \mathbf{g}_t$;
- Implicit FTRL with aProx updates (Trunc): $z_t =$

$$\min \left\{ 1, \frac{\lambda_t \ell_t(\mathbf{x}_t)}{\|\mathbf{g}_t\|^2} \right\} \mathbf{g}_t;$$

- Implicit FTRL with two-step updates (Twostep): $z_t = \partial \ell_t(\mathbf{x}_t - \mathbf{g}_t/\lambda_t)$;
- Implicit FTRL with (6) when the proximal operator has a closed form (Proximal).

We adopt the choice of λ_t from Corollary 4.3.

We conduct linear prediction experiments on datasets from LibSVM (Chang & Lin, 2011). We show here experiments on classification tasks using the hinge loss and the logistic loss, and regression tasks with absolute loss. We normalize the datasets and added a constant bias term to the features. Given that in the online learning setting, we do not have the training data and validation data to tune the β , we will plot the averaged loss, $\frac{1}{t} \sum_{i=1}^t \ell_i(\mathbf{x}_i)$, versus different choice of




Figure 3. Absolute loss, averaged loss vs. hyperparameter β .

β , that at the same time show the algorithms’ sensitivity to the hyperparameter β and their best achievable performance. We consider $\beta \in [10^{-3}, 10^3]$ for hinge loss and logistic loss, and $\beta \in [10^{-3}, 10^8]$ for the absolute loss. Each algorithm is run 15 times, we plot the average of the averaged losses and the 95% confidence interval. Note that the confidence intervals so small to be invisible, but for the larger values of the β for the Linear updates.

Figure 1 and Figure 2 show the averaged loss versus different selections of hyperparameter β for classification tasks with hinge loss and logistic loss respectively. Note that with the hinge loss aProx updates and proximal updates are completely equivalent. In all experiments, FTRL with linearized updates is more sensitive to the setting of β , and

its performance is almost uniformly worse than all the other generalized implicit updates. This is in line with previous results in Asi & Duchi (2019) in the offline setting. With the logistic loss, the proximal operator does not have a closed-form solution. In all the classification experiments, the performance of generalized implicit FTRL with two-step updates seems remarkable and a possible viable alternative to aProx. The confidence intervals for all implicit updates have a width smaller than 0.01, making them too narrow to be visible in the figures. In contrast, when using hinge loss, the performance of FTRL with linear models exhibits significant fluctuations across different repetitions when a large learning rate is used. This observation provides evidence supporting our assertion that the selection of hyperparameter β greatly affects the performance of FTRL with linear models, while implicit updates demonstrate robustness.

Figure 3 shows that FTRL with linearized updates is very sensitive to the choice of the hyperparameter β , while the implicit FTRL updates are robust. Again, Implicit FTRL with two-step updates achieves essentially the best performance. The confidence intervals in the regression tasks lead to a similar conclusion as in the classification tasks.

8. Conclusion and Future Work

In this work, we propose a new framework: generalized implicit Follow-the-Regularized-Leader. We show that generalized implicit FTRL can not only recover known algorithms, e.g., implicit FTRL and FTRL with linearized losses, but it also provides a theoretical guideline to design new algorithms, such as the extensions of aProx and Mirror-Prox. Indeed, we believe that the main contribution of our work lies precisely in the fact that it provides a unifying framework that is general, flexible, and theoretically grounded.

In the future, we plan to explore further this framework designing new z_t with low computational complexity. This is a promising direction because the two-step update seems to be already a valid alternative to the aProx updates, even if it comes at the computational expense of querying an additional gradient in each round.

Acknowledgements

We thank Alex Shtoff for discussion and feedback on a preliminary version of this paper. Francesco Orabona is supported by the National Science Foundation under the grants no. 2022446 “Foundations of Data Science Institute” and no. 2046096 “CAREER: Parameter-free Optimization Algorithms for Machine Learning”.

References

- Abernethy, J. D., Hazan, E., and Rakhlin, A. Competing in the dark: An efficient algorithm for bandit linear optimization. In Servedio, R. A. and Zhang, T. (eds.), *Proc. of Conference on Learning Theory (COLT)*, pp. 263–274. Omnipress, 2008.
- Asi, H. and Duchi, J. C. Stochastic (approximate) proximal point methods: Convergence, optimality, and adaptivity. *SIAM Journal on Optimization*, 29(3):2257–2290, 2019.
- Bauschke, H. H. and Combettes, P. L. *Convex analysis and monotone operator theory in Hilbert spaces*, volume 408. Springer, 2011.
- Campolongo, N. and Orabona, F. Temporal variability in implicit online learning. In *Advances in Neural Information Processing Systems*, volume 33. Curran Associates, Inc., 2020.
- Cesa-Bianchi, N. and Lugosi, G. *Prediction, learning, and games*. Cambridge University Press, 2006.
- Cesa-Bianchi, N. and Orabona, F. Online learning algorithms. *Annual Review of Statistics and Its Application*, 8:165–190, 2021.
- Chang, C.-C. and Lin, C.-J. LIBSVM: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2(3):1–27, 2011. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- Chen, K., Cutkosky, A., and Orabona, F. Implicit parameter-free online learning with truncated linear models. In *International Conference on Algorithmic Learning Theory*, pp. 148–175. PMLR, 2022a.
- Chen, K., Langford, J., and Orabona, F. Better parameter-free stochastic optimization with ODE updates for coin-betting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2022b.
- Crammer, K., Dekel, O., Keshet, J., Shalev-Shwartz, S., and Singer, Y. Online passive-aggressive algorithms. *Journal of Machine Learning Research*, 7:551–585, 2006.
- Duchi, J. C., Hazan, E., and Singer, Y. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12:2121–2159, 2011.
- Fang, H., Harvey, N., Portella, V., and Friedlander, M. Online mirror descent and dual averaging: keeping pace in the dynamic case. In Daum, III, H. and Singh, A. (eds.), *Proc. of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 3008–3017. PMLR, 13–18 Jul 2020.
- Hazan, E. and Kale, S. Extracting certainty from uncertainty: Regret bounded by variation in costs. In *Proc. of the 21st Conference on Learning Theory*, 2008.
- Joulani, P., Györfgy, A., and Szepesvári, C. A modular analysis of adaptive (non-)convex optimization: Optimism, composite objectives, and variational bounds. In *Proc. of the International Conference on Algorithmic Learning Theory (ALT)*, volume 76, pp. 681–720, 2017.
- Kivinen, J. and Warmuth, M. Exponentiated gradient versus gradient descent for linear predictors. *Information and Computation*, 132(1):1–63, January 1997.
- Kulis, B. and Bartlett, P. L. Implicit online learning. In *International Conference on Machine Learning*, pp. 575–582, 2010.
- Littlestone, N. Learning quickly when irrelevant attributes abound: a new linear-threshold algorithm. *Machine Learning*, 2(4):285–318, 1988.
- Martinet, B. Régularisation d'inéquations variationnelles par approximations successives. *rev. française informat. Recherche Opérationnelle*, 4:154–158, 1970.
- McMahan, H. B. A unified view of regularized dual averaging and mirror descent with implicit updates. *arXiv preprint arXiv:1009.3240*, 2010.
- McMahan, H. B. A survey of algorithms and analysis for adaptive online learning. *The Journal of Machine Learning Research*, 18(1):3117–3166, 2017.
- McMahan, H. B. and Streeter, M. J. Adaptive bound optimization for online convex optimization. In *COLT*, 2010.
- Moreau, J.-J. Proximité et dualité dans un espace hilbertien. *Bulletin de la Société mathématique de France*, 93:273–299, 1965.
- Nemirovski, A. Prox-method with rate of convergence $O(1/t)$ for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251, 2004.
- Nemirovskij, A. S. and Yudin, D. *Problem complexity and method efficiency in optimization*. Wiley, New York, NY, USA, 1983.
- Orabona, F. A modern introduction to online learning. *arXiv preprint arXiv:1912.13213*, 2019. Version 6.
- Orabona, F. and Pál, D. Scale-free algorithms for online linear optimization. In *International Conference on Algorithmic Learning Theory*, pp. 287–301. Springer, 2015.

- Orabona, F. and Pál, D. Coin betting and parameter-free online learning. In Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 29*, pp. 577–585. Curran Associates, Inc., 2016.
- Orabona, F. and Pál, D. Scale-free online learning. *Theoretical Computer Science*, 716:50–69, 2018. Special Issue on ALT 2015.
- Orabona, F. and Pál, D. Parameter-free stochastic optimization of variationally coherent functions. *arXiv preprint arXiv:2102.00236*, 2021.
- Parikh, N. and Boyd, S. Proximal algorithms. *Foundations and Trends in optimization*, 1(3):127–239, 2014.
- Rockafellar, R. T. *Convex Analysis*. Princeton University Press, 1970.
- Rockafellar, R. T. Monotone operators and the proximal point algorithm. *SIAM journal on control and optimization*, 14(5):877–898, 1976.
- Shalev-Shwartz, S. *Online Learning: Theory, Algorithms, and Applications*. PhD thesis, The Hebrew University, 2007.
- Shalev-Shwartz, S. and Singer, Y. A primal-dual perspective of online learning algorithms. *Machine Learning*, 69:115–142, 2007a.
- Shalev-Shwartz, S. and Singer, Y. Convex repeated games and Fenchel duality. In *Advances in neural information processing systems*, pp. 1265–1272, 2007b.
- Shtoff, A. Efficient implementation of incremental proximal-point methods. *arXiv preprint arXiv:2205.01457*, 2022.
- Streeter, M. and McMahan, H. B. Less regret via online conditioning. *arXiv preprint arXiv:1002.4862*, 2010.
- Warmuth, M. K. and Jagota, A. K. Continuous and discrete-time nonlinear gradient descent: Relative loss bounds and convergence. In *Electronic proceedings of the 5th International Symposium on Artificial Intelligence and Mathematics*, volume 326, 1997.

A. Update Rule for Common Losses

In this section, we report the proximal operator of common losses for easy referencing. These formulas are well-known and they can be found, for example, in [Crammer et al. \(2006\)](#); [Kulis & Bartlett \(2010\)](#).

$$\begin{aligned}\ell_t(\mathbf{x}) = \max(1 - y_t \langle \mathbf{s}_t, \mathbf{x} \rangle, 0) &\Rightarrow \text{Prox}_{\frac{\ell_t}{\lambda}}(\mathbf{x}) = \mathbf{x} + \min\left(\frac{1}{\lambda}, \frac{\max(1 - y_t \langle \mathbf{s}_t, \mathbf{x} \rangle, 0)}{\|\mathbf{s}_t\|^2}\right) y_t \mathbf{s}_t \\ \ell_t(\mathbf{x}) = |\langle \mathbf{s}_t, \mathbf{x} \rangle - y_t| &\Rightarrow \text{Prox}_{\frac{\ell_t}{\lambda}}(\mathbf{x}) = \mathbf{x} - \min\left(\frac{1}{\lambda}, \frac{|\langle \mathbf{s}_t, \mathbf{x} \rangle - y_t|}{\|\mathbf{s}_t\|^2}\right) \mathbf{s}_t \\ \ell_t(\mathbf{x}) = \frac{1}{2}(\langle \mathbf{s}_t, \mathbf{x} \rangle - y_t)^2 &\Rightarrow \text{Prox}_{\frac{\ell_t}{\lambda}}(\mathbf{x}) = \mathbf{x} - \frac{(\langle \mathbf{s}_t, \mathbf{x} \rangle - y_t) \mathbf{s}_t}{\lambda + \|\mathbf{s}_t\|_2^2}.\end{aligned}$$

B. Proof of Corollary 4.3

Proof. From the regret guarantee in Lemma 4.1, we have that

$$\text{Regret}_T(\mathbf{u}) \leq \psi_{T+1}(\mathbf{u}) + \beta^2 \sum_{t=1}^T \gamma_t \leq \lambda_{T+1}(\psi(\mathbf{u}) + \beta^2), \quad \forall \mathbf{u} \in V.$$

Now, we upper bound $\sum_{t=1}^T \gamma_t$ in two different ways. In the first upper bound, we have

$$\begin{aligned}\sum_{t=1}^T \gamma_t &= \sum_{t=1}^T \left(\ell_t(\mathbf{x}_t) - \ell(\mathbf{x}_{t+1}) - \frac{\lambda_t}{2} \|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2 \right) = \ell_1(\mathbf{x}_1) - \ell_T(\mathbf{x}_{T+1}) + \sum_{t=2}^T (\ell_t(\mathbf{x}_t) - \ell_{t-1}(\mathbf{x}_t)) \\ &\leq \ell_1(\mathbf{x}_1) - \ell_T(\mathbf{x}_{T+1}) + V_T.\end{aligned}$$

For the second upper bound, we have

$$\begin{aligned}\gamma_t &= \ell_t(\mathbf{x}_t) - \ell(\mathbf{x}_{t+1}) - \frac{\lambda_t}{2} \|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2 \leq \langle \mathbf{g}_t, \mathbf{x}_t - \mathbf{x}_{t+1} \rangle - \frac{\lambda_t}{2} \|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2 \\ &\leq \frac{\|\mathbf{g}_t\|_*^2}{2\lambda_t} + \frac{\lambda_t}{2} \|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2 - \frac{\lambda_t}{2} \|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2 = \frac{\|\mathbf{g}_t\|_*^2}{2\lambda_t} \leq \beta \frac{\|\mathbf{g}_t\|_*}{2},\end{aligned}$$

where we used Fenchel-Young inequality and the second lower bound is obtained by using the fact that $\lambda_t \geq \lambda_1 = \frac{G}{\beta}$. Hence, we have

$$\lambda_{t+1} = \lambda_t + \frac{\gamma_t}{\beta^2} \leq \lambda_t + \min\left(\frac{\|\mathbf{g}_t\|_*}{2\beta}, \frac{\|\mathbf{g}_t\|_*^2}{2\beta^2 \lambda_t}\right).$$

Using Lemma 6.1 in [Campolongo & Orabona \(2020\)](#) and taking into account the fact that $\lambda_1 = \frac{G}{\beta}$, we have

$$\lambda_{T+1} \leq \frac{G}{\beta} + \sqrt{\frac{5}{4\beta^2} \sum_{t=1}^T \|\mathbf{g}_t\|_*^2}.$$

Putting all together, we have the stated bound. □

C. Proof of Corollary 6.2

Proof. The proximal operator of $\frac{\hat{\ell}_t}{\lambda}$ is

$$\text{Prox}_{\frac{\hat{\ell}_t}{\lambda}}(\mathbf{x}) = \mathbf{x} - \min\left(\frac{1}{\lambda}, \frac{\ell_t(\mathbf{x}_t)}{\|\mathbf{g}_t\|^2}\right) \mathbf{g}_t.$$

Hence, from (7), we have

$$\mathbf{x}_{t+1} = \frac{\lambda_t}{\lambda_{t+1}} \left(\mathbf{x}_t - \min \left(\frac{1}{\lambda_t}, \frac{\ell_t(\mathbf{x}_t)}{\|\mathbf{g}_t\|^2} \right) \mathbf{g}_t \right) = \frac{\lambda_t}{\lambda_{t+1}} \mathbf{x}_t - \min \left(\frac{1}{\lambda_{t+1}}, \frac{\lambda_t}{\lambda_{t+1}} \frac{\ell_t(\mathbf{x}_t)}{\|\mathbf{g}_t\|^2} \right) \mathbf{g}_t.$$

Instead, from (6), we have

$$\mathbf{x}_{t+1} = \frac{\lambda_t}{\lambda_{t+1}} \mathbf{x}_t - \min \left(\frac{1}{\lambda_{t+1}}, \frac{\ell_t(\mathbf{x}_t)}{\|\mathbf{g}_t\|^2} \right) \mathbf{g}_t.$$

□