Supplementary Material of "A Simple Online Algorithm for Competing with Dynamic Comparators"

Yu-Jie Zhang, Peng Zhao, Zhi-Hua Zhou
National Key Laboratory for Novel Software Technology,
Nanjing University, Nanjing 210023, China
{zhangyj, zhaop, zhouzh}@lamda.nju.edu.cn

Abstract

This is the supplementary material of the paper "A Simple Online Algorithm for Competing with Dynamic Comparators". In Section A, we present the proofs of path-length and temporal variability bounds. Then, we provide the proof of Lemma 3 in Section B. Meanwhile, we provide additional experimental results on classification tasks in Section C. Section D presents the technical lemmas. Notice that the adaptive bound of Theorem 1 can be directly obtained from Theorem 2, since the vanilla Hedge is essentially a special case of the optimistic Hedge when no predictable sequences are provided.

A Proofs of Path-length and Temporal Variability Bounds

Before showing the proofs of path-length and temporal variability bounds, we present a general dynamic regret bound for OMD with the fixed step size, which supports the comparison between learner's decisions and any sequence $\{u_1, \dots, u_T\}$, where u_t can be any point in the feasible domain $\mathcal{X}$.

Lemma 4. *When $\mathcal{R}(\cdot)$ is a 1-strongly convex function, running OMD with any fixed step size $\eta > 0$ satisfies,*

$$\sum_{t=1}^{T} f_t(x_t) - \sum_{t=1}^{T} f_t(u_t) \leq \frac{\eta D_T}{2} + \frac{R_{\max}^2 + \gamma \sum_{t=2}^{T} \|u_t - u_{t-1}\|}{2\eta},$$

provided $\mathcal{D}_{\mathcal{R}}(x,z) - \mathcal{D}_{\mathcal{R}}(y,z) \leq \gamma\|x-y\|, \forall x,y,z \in \mathcal{X}$, where u_t can be any comparator in the feasible domain $\mathcal{X}$.

Proof of Lemma 1. According to the update procedure of OMD (5) with the fixed step size η and Lemma 8, we have

$$\langle x_t - \widehat{x}_{t+1}, \eta M_t \rangle \leq D_{\mathcal{R}}(\widehat{x}_{t+1}, \widehat{x}_t) - D_{\mathcal{R}}(\widehat{x}_{t+1}, x_t) - D_{\mathcal{R}}(x_t, \widehat{x}_t); \tag{16}$$

$$\langle \widehat{x}_{t+1} - u_t, \eta \nabla f_t(x_t) \rangle \leq D_{\mathcal{R}}(u_t, \widehat{x}_t) - D_{\mathcal{R}}(u_t, \widehat{x}_{t+1}) - D_{\mathcal{R}}(\widehat{x}_{t+1}, \widehat{x}_t). \tag{17}$$

Each iteration of the dynamic regret can be decomposed as,

$$\begin{aligned} f_t(x_t) - f_t(u_t) &\leq \langle \nabla f_t(x_t), x_t - u_t \rangle \\ &= \underbrace{\langle \nabla f_t(x_t) - M_t, x_t - \widehat{x}_{t+1} \rangle}_{\texttt{term A}} + \underbrace{\langle M_t, x_t - \widehat{x}_{t+1} \rangle}_{\texttt{term B}} + \underbrace{\langle \nabla f_t(x_t), \widehat{x}_{t+1} - u_t \rangle}_{\texttt{term C}}, \end{aligned} \tag{18}$$

where the first inequality holds for Jensen's inequality. Then, we proceed to bound the each term. First, we can see that

$$\texttt{term A} \leq \|x_t - \widehat{x}_{t+1}\| \cdot \|\nabla f_t(x_t) - M_t\|_* \leq \frac{\eta}{2}\|\nabla f_t(x_t) - M_t\|_*^2 + \frac{1}{2\eta}\|x_t - \widehat{x}_{t+1}\|^2, \tag{19}$$

where the first inequality satisfies due to Hölder inequality and the second holds due to the fact that $ab \leq \frac{\eta}{2}a^2 + \frac{1}{2\eta}b^2$ for $a, b \in \mathbb{R}$ and any $\eta > 0$. According to (16) and (17), term B and term C can be bounded as

$$\texttt{term B} \leq \frac{1}{\eta}\left(D_{\mathcal{R}}(\widehat{x}_{t+1}, \widehat{x}_t) - D_{\mathcal{R}}(\widehat{x}_{t+1}, x_t) - D_{\mathcal{R}}(x_t, \widehat{x}_t)\right); \tag{20}$$

$$\texttt{term C} \leq \frac{1}{\eta}\left(D_{\mathcal{R}}(u_t, \widehat{x}_t) - D_{\mathcal{R}}(u_t, \widehat{x}_{t+1}) - D_{\mathcal{R}}(\widehat{x}_{t+1}, \widehat{x}_t)\right). \tag{21}$$

Due to the strong convexity of regularizer $\mathcal{R}$, we have $D_{\mathcal{R}}(x, y) \geq \frac{1}{2}\|x - y\|^2$ for any $x, y \in \mathcal{X}$ [Mohri et al., 2018]. Therefore, by plugging (19), (21) and (20) into (18), we obtain that

$$\begin{aligned} f_t(x_t) - f_t(u_t) \leq{} & \frac{\eta}{2}\|\nabla f_t(x_t) - M_t\|_*^2 + \frac{1}{2\eta}\|x_t - \widehat{x}_{t+1}\|^2 \\ & + \frac{1}{\eta}\left(D_{\mathcal{R}}(u_t, \widehat{x}_t) - D_{\mathcal{R}}(u_t, \widehat{x}_{t+1}) - \frac{1}{2}\|\widehat{x}_{t+1} - x_t\|^2 - \frac{1}{2}\|x_t - \widehat{x}_t\|^2\right). \end{aligned}$$

Summing over the index from $t = 1$ to T, we have

$$\begin{aligned} \sum_{t=1}^T f_t(x_t) - \sum_{t=1}^T f_t(u_t) & \leq \frac{\eta}{2} D_T + \sum_{t=1}^T \frac{D_{\mathcal{R}}(u_t, \widehat{x}_t) - D_{\mathcal{R}}(u_t, \widehat{x}_{t+1})}{\eta} \\ & \leq \frac{\eta}{2} D_T + \sum_{t=1}^{T-1} \frac{D_{\mathcal{R}}(u_{t+1}, \widehat{x}_{t+1}) - D_{\mathcal{R}}(u_t, \widehat{x}_{t+1})}{\eta} + \frac{R_{\max}^2}{\eta}. \end{aligned}$$

Under the condition that $\mathcal{D}_{\mathcal{R}}(x, z) - \mathcal{D}_{\mathcal{R}}(y, z) \leq \gamma\|x - y\|, \forall x, y, z \in \mathcal{X}$, we can further bound the above inequality as

$$\sum_{t=1}^T f_t(x_t) - \sum_{t=1}^T f_t(u_t) \leq \frac{\eta}{2} D_T + \sum_{t=1}^{T-1} \frac{\gamma\|u_{t+1} - u_t\|}{\eta} + \frac{R_{\max}^2}{\eta} \leq \frac{\eta}{2} D_T + \frac{R_{\max}^2 + \gamma\sum_{t=2}^T \|u_t - u_{t-1}\|}{\eta},$$

which completes the proof. □

Since the base algorithm that we run is Algorithm 2, which is the OMD incorporating surrogate loss function, we offer a counterpart of Lemma 4 for these base algorithms.

Lemma 5. *When $\mathcal{R}(\cdot)$ is a 1-strongly convex function, the i-th base algorithm running Algorithm 2 with a fixed step size η_i enjoys*

$$\sum_{t=1}^T \ell_t(x_t^i) - \sum_{t=1}^T \ell_t(u_t) \leq \frac{\eta_i D_T}{2} + \frac{R_{\max}^2 + \gamma\sum_{t=2}^T \|u_t - u_{t-1}\|}{2\eta_i}, \tag{22}$$

provided $\mathcal{D}_{\mathcal{R}}(x, z) - \mathcal{D}_{\mathcal{R}}(y, z) \leq \gamma\|x - y\|, \forall x, y, z \in \mathcal{X}$, where u_t can be any comparator in the feasible domain $\mathcal{X}$.

Proof of Lemma 5. The proof of Lemma 5 is almost the same with that of Lemma 4 as we can just substitute the prediction x_t with x_t^i and $\widehat{x}_t$ with $\widehat{x}_t^i$. First, according to the update procedure of OMD with the surrogate loss, we can decompose each iteration of the dynamic regret following (18),

$$\ell_t(x_t^i) - \ell_t(u_t) \leq \underbrace{\langle \nabla f_t(x_t) - M_t, x_t^i - \widehat{x}_{t+1}^i \rangle}_{\texttt{term A}} + \underbrace{\langle M_t, x_t^i - \widehat{x}_{t+1}^i \rangle}_{\texttt{term B}} + \underbrace{\langle \nabla f_t(x_t), \widehat{x}_{t+1}^i - u_t \rangle}_{\texttt{term C}}, \tag{23}$$

where term A, term B and term C can be bounded by,

$$\texttt{term A} \leq \frac{\eta_i}{2}\|\nabla f_t(x_t) - M_t\|_*^2 + \frac{1}{2\eta}\|x_t^i - \widehat{x}_{t+1}^i\|^2; \tag{24}$$

$$\texttt{term B} \leq \frac{1}{\eta_i}\left(D_{\mathcal{R}}(\widehat{x}_{t+1}^i, \widehat{x}_t^i) - D_{\mathcal{R}}(\widehat{x}_{t+1}^i, x_t^i) - D_{\mathcal{R}}(x_t^i, \widehat{x}_t^i)\right); \tag{25}$$

$$\texttt{term C} \leq \frac{1}{\eta_i}\left(D_{\mathcal{R}}(u_t, \widehat{x}_t^i) - D_{\mathcal{R}}(u_t, \widehat{x}_{t+1}^i) - D_{\mathcal{R}}(\widehat{x}_{t+1}^i, \widehat{x}_t^i)\right). \tag{26}$$

By plugging (24), (26) and (25) into (23) and the fact that $D_{\mathcal{R}}(x,y) \geq \frac{1}{2}\|x-y\|^2$ for any $x, y \in \mathcal{X}$, we obtain

$$\begin{aligned}\ell_t(x_t) - \ell_t(x_t^*) \leq{}& \frac{\eta_i}{2}\|\nabla f_t(x_t) - M_t\|_*^2 + \frac{1}{2\eta_i}\|x_t^i - \widehat{x}_{t+1}^i\|^2 \\ &+ \frac{1}{\eta_i}\left(D_{\mathcal{R}}(u_t, \widehat{x}_t^i) - D_{\mathcal{R}}(u_t, \widehat{x}_{t+1}^i) - \frac{1}{2}\|\widehat{x}_{t+1}^i - x_t^i\|^2 - \frac{1}{2}\|x_t^i - \widehat{x}_t^i\|^2\right).\end{aligned}$$

Summing over T iterations and following the same argument in the proof of Lemma 4, we can obtain the dynamic regret bound

$$\sum_{t=1}^{T} \ell_t(x_t^i) - \sum_{t=1}^{T} \ell_t(u_t) \leq \frac{\eta_i D_T}{2} + \frac{R_{\max}^2 + \gamma\sum_{t=2}^{T}\|u_t - u_{t-1}\|}{2\eta_i}, \tag{27}$$

which completes the proof. □

A.1 Proof of the Path-length Bound (Lemma 1)

Proof of Lemma 1. Lemma 1 can be seen as a corollary of Lemma 4. We can complete the proof by setting the comparator $u_t = x_t^*$ for $t \in [T]$. □

A.2 Proof of Temporal Variability Bound (Lemma 2 and Lemma 6)

In this section, we first derive the temporal variability bound for OMD with fixed step size (Lemm 2), followed by the analysis on the our algorithm where optimistic Hedge is employed to aggregate multiple OMDs running on surrogate loss (Lemma 6).

Proof of Lemma 2. First, we can decompose the dynamic regret in the following way,

$$\mathbf{Reg}_T^d = \sum_{t=1}^{T} f_t(x_t) - \sum_{t=1}^{T} f_t(x_t^*) = \underbrace{\sum_{t=1}^{T} f_t(x_t) - \sum_{t=1}^{T} f_t(u_t)}_{\texttt{term A}} + \underbrace{\sum_{t=1}^{T} f_t(u_t) - \sum_{t=1}^{T} f_t(x_t^*)}_{\texttt{term B}}. \tag{28}$$

Here, we insert a term of $\sum_{t=1}^{T} f_t(u_t)$, where u_t can be any comparator in the feasible set. The flexibility of the sequence of $\{u_1, \ldots, u_T\}$ is of great importance. In this proof, we will specify $\{u_1, \ldots, u_T\}$ as a piece-wise stationary sequence that changes every $\Delta \in [T]$ iterations. Concretely, denoting by $\mathcal{I}_i = [s_i, e_i]$ the i-th interval with the length Δ, where $s_i = (i-1)\cdot\Delta + 1$ and $e_i = i \cdot \Delta$, for all $t \in \mathcal{I}_i$, we specify u_t as the best fixed decision $x_{\mathcal{I}_i}^* = \arg\min_{x \in \mathcal{X}} \sum_{t \in \mathcal{I}_i} f_t(x)$ of the corresponding interval $\mathcal{I}_i$.

In such a case, according to Lemma 4, term A is bounded by

$$\texttt{term A} \leq \frac{\eta D_T}{2} + \frac{R_{\max}^2 + \gamma\sqrt{2}R_{\max}\lceil T/\Delta\rceil}{2\eta}. \tag{29}$$

Because the sequence $\{u_1, \ldots, u_T\}$ changes at most $\lceil T/\Delta\rceil$ times and the deviation is bounded due to the boundedness of $\mathcal{X}$ in terms of Bergman divergence that $\|x - y\| \leq \sqrt{2}R_{\max}$, for all $x, y \in \mathcal{X}$.

Then we proceed to analyze the term B. By the construction of the comparator sequence $\{u_1, \ldots, u_T\}$, term B is the difference between the cumulative loss of best decisions within each interval and that of the worst case comparators, which can be bounded following the argument in Besbes et al. [2015].

For contentedness, we present the proof here. At first, we first bound this difference on the interval $\mathcal{I}_i$,

$$\begin{aligned}&\sum_{t \in \mathcal{I}_i} f_t(x_{\mathcal{I}_i}^*) - \sum_{t \in \mathcal{I}_i} f_t(x_t^*) \\ ={}& \sum_{t \in \mathcal{I}_i} f_t(x_{\mathcal{I}_i}^*) - \sum_{t \in \mathcal{I}_i} f_{s_i}(x_{s_i}^*) + \sum_{t \in \mathcal{I}_i} f_{s_i}(x_{s_i}^*) - \sum_{t \in \mathcal{I}_i} f_t(x_t^*)\end{aligned}$$

$$
\begin{aligned}
&\leq \sum_{t\in\mathcal{I}_i} f_t(x^*_{s_i}) - \sum_{t\in\mathcal{I}_i} f_{s_i}(x^*_{s_i}) + \sum_{t\in\mathcal{I}_i} f_{s_i}(x^*_t) - \sum_{t\in\mathcal{I}_i} f_t(x^*_t) \\
&\leq 2\Delta \sum_{t\in\mathcal{I}_i} \sup_{x\in\mathcal{X}} |f_{t+1}(x) - f_t(x)|, \qquad (30)
\end{aligned}
$$

where the first inequality comes from the optimality of $x^*_{\mathcal{I}_i}$ over $\mathcal{I}_i$ and $x^*_{s_i} = \arg\min_{x\in\mathcal{X}} f_{s_i}(x)$. The last inequality holds as $f_{s_i}(x_t) - f_t(x_t) \leq \sum_{t\in\mathcal{I}_i} \sup_{x\in\mathcal{X}} |f_{t+1}(x) - f_t(x)|$ for all $t \in \mathcal{I}_i$ for all $x \in \mathcal{X}$. Thus, by summing over all intervals, we have

$$\texttt{term B} \leq 2\Delta V_T.$$

We emphasize that it is crucial to specify the comparator u_t be the minimizer of the interval with respect to the *original function* f_t instead of the surrogate loss function ℓ_t, which is introduced in the master-base aggregation.

By plugging (29) and (30) into (28), we obtain

$$\mathbf{Reg}^d_T \leq \frac{\eta D_T}{2} + \frac{R^2_{\max} + \gamma\sqrt{2}R_{\max}\lceil T/\Delta\rceil}{2\eta} + 2\Delta V_T.$$

By further setting $\eta = \sqrt{(R^2_{\max} + \gamma\sqrt{2}R_{\max}\lceil T/\Delta\rceil)/(1+D_T)}$, we have

$$\mathbf{Reg}^d_T \leq \sqrt{(1+D_T)(R^2_{\max} + \gamma\sqrt{2}R_{\max}\lceil T/\Delta\rceil)} + 2\Delta V_T,$$

which completes the proof. □

However, the base algorithm we run is Algorithm 2, which incorporates the surrogate loss function in the learning process. Lemma 2 can not be expanded for analyzing the regret for these base algorithms, as their entemporal variability term becomes $V^\ell_T = \sum_{t=2}^T \sup_{x\in\mathcal{X}} |\ell_t(x) - \ell_{t-1}(x)|$, which is hard to be converted to the desired V_T term in terms of original function f_t.

Thus, instead of analyzing the dynamic regret of the single base algorithm, as shown in Appendix A of the main paper, we bound the overall dynamic regret of the whole algorithm, where term A of (11) is bounded as Lemma 6.

Lemma 6. *Under the same condition of Theorem 1, running the master algorithm (Algorithm 1) with N base OMDs (Algorithm 2), we have*

$$\sum_{i=1}^{\lceil T/\Delta\rceil} \sum_{t\in\mathcal{I}_i} \ell_t(x_t) - \sum_{i=1}^{\lceil T/\Delta\rceil} \sum_{t\in\mathcal{I}_i} \ell_t(x^*_{\mathcal{I}_i}) \leq 2R_{\max}\sqrt{(4+2\ln N)D_T} + 2\sqrt{(1+D_T)(C_1 + C_2\lceil T/\Delta\rceil)},$$

holds for any $\Delta \in [T]$where $\mathcal{I}_i = [s_i, e_i]$ is the i-th interval with the length Δ, and $s_i = (i-1)\cdot\Delta + 1$, $e_i = i\cdot\Delta$. Notation $x^*_{\mathcal{I}_i} = \arg\min_{x\in\mathcal{X}} \sum_{t\in\mathcal{I}_i} f_t(x)$ refers to the best fixed decision in interval $\mathcal{I}_i$.

Proof of Lemma 6. First we can decompose the dynamic regret as

$$
\begin{aligned}
&\sum_{i=1}^{\lceil T/\Delta\rceil} \sum_{t\in\mathcal{I}_i} \ell_t(x_t) - \sum_{i=1}^{\lceil T/\Delta\rceil} \sum_{t\in\mathcal{I}_i} \ell_t(x^*_{\mathcal{I}_i}) \\
\leq\ & \underbrace{\sum_{i=1}^{\lceil T/\Delta\rceil} \sum_{t\in\mathcal{I}_i} \ell_t(x_t) - \sum_{i=1}^{\lceil T/\Delta\rceil} \sum_{t\in\mathcal{I}_i} \ell_t(x^i_t)}_{\texttt{Reg-Mas(i)}} + \underbrace{\sum_{i=1}^{\lceil T/\Delta\rceil} \sum_{t\in\mathcal{I}_i} \ell_t(x^i_t) - \sum_{i=1}^{\lceil T/\Delta\rceil} \sum_{t\in\mathcal{I}_i} \ell_t(x^*_{\mathcal{I}_i})}_{\texttt{Reg-Base(i)}}
\end{aligned}
$$

for any $i \in [N]$, where the first term is regret of the master algorithm while the second term is dynamic regret of the i-th base algorithm. According to Lemma 3, the master regret w.r.t any base algorithm indexed by i can be bounded as

$$\texttt{Reg-Mas(i)} \leq 2R_{\max}\sqrt{(4+2\ln D_T)}.$$

As for the regret of base algorithm, according to Lemma 4, by setting u_t as the best fixed decision $x^*_{\mathcal{I}_i} = \arg\min_{x\in\mathcal{X}} \sum_{t\in\mathcal{I}_i} f_t(x)$ of the corresponding interval $\mathcal{I}_i$, we have

$$\texttt{Reg-Base(i)} \le \frac{\eta_i D_T}{2} + \frac{R^2_{\max} + \gamma \sum_{i=2}^{\lceil T/\Delta\rceil} \|x^*_{\mathcal{I}_i} - x^*_{\mathcal{I}_{i-1}}\|}{2\eta_i},$$

holds for any $i \in [N]$. Due to the boundedness of $\mathcal{X}$ in terms of Bergman divergence, we have $\|x - y\| \le \sqrt{2}R_{\max}$, for all $x, y \in \mathcal{X}$. Thus, the base regret can be further bounded as

$$\texttt{Reg-Base(i)} \le \frac{\eta_i D_T}{2} + \frac{R^2_{\max} + \gamma\sqrt{2}R_{\max}\lceil T/\Delta\rceil}{2\eta_i},$$

. Combining $\texttt{Reg-Mas(i)}$ with $\texttt{Reg-Base(i)}$, we confirm that

$$\begin{aligned}&\sum_{i=1}^{\lceil T/\Delta\rceil}\sum_{t\in\mathcal{I}_i} \ell_t(x_t) - \sum_{i=1}^{\lceil T/\Delta\rceil}\sum_{t\in\mathcal{I}_i} \ell_t(x^*_{\mathcal{I}_i}) \\ &\le 2R_{\max}\sqrt{4 + 2\ln D_T} + \frac{\eta_i D_T}{2} + \frac{R^2_{\max} + \gamma\sqrt{2}R_{\max}\lceil T/\Delta\rceil}{2\eta_i},\end{aligned}$$

which holds for any $i \in [N]$. Since $\Delta \in [T]$, there must exist a step size $\eta' \in \mathcal{P}$ satisfies $\eta^*_{\texttt{var}}/2 \le \eta' \le \eta^*_{\texttt{var}}$, where $\eta^*_{\texttt{var}} = \sqrt{(R^2_{\max} + \gamma\sqrt{2}R_{\max}\lceil T/\Delta\rceil)/(1 + D_T)}$ is the optimal tunning. Thus, we can choose the base algorithm running with η' as the intermediate term, where the dynamic regret can be further bounded by,

$$\begin{aligned}&\sum_{i=1}^{\lceil T/\Delta\rceil}\sum_{t\in\mathcal{I}_i} \ell_t(x_t) - \sum_{i=1}^{\lceil T/\Delta\rceil}\sum_{t\in\mathcal{I}_i} \ell_t(x^*_{\mathcal{I}_i}) \\ &\le 2R_{\max}\sqrt{4 + 2\ln D_T} + \frac{\eta^*_{\texttt{var}} D_T}{2} + \frac{R^2_{\max} + \gamma\sqrt{2}R_{\max}\lceil T/\Delta\rceil}{\eta^*_{\texttt{var}}} \\ &\le 2R_{\max}\sqrt{4 + 2\ln D_T} + 2\sqrt{(1 + D_T)(R^2_{\max} + \gamma\sqrt{2}R_{\max}\lceil T/\Delta\rceil)},\end{aligned}$$

which completes the proof by setting $C_1 = R^2_{\max}$ and $C_2 = \gamma\sqrt{2}R_{\max}$. □

B Proof of Lemma 3

For self-containedness, we present the proof of Lemma 3 here, which is a counterpart of Lemma 4 in Rakhlin and Sridharan [2013] without exploiting the local norm. Lemma 4 in Rakhlin and Sridharan [2013] requires the learning rate satisfying $\epsilon\|f_t - m_t\|_\infty \le 1/4$ at any iteration. This limitation can be eliminated based on Theorem 19 in Syrgkanis et al. [2015] by taking optimistic Hedge as a special case of Optimistic Follow the Regularized Leader (OFTRL).

Proof of Lemma 3. Note that the update procedure of optimistic Hedge is actually the solution of

$$w_{t+1} = \operatorname*{arg\,min}_{w\in\Delta_N} \epsilon \left\langle \sum_{t=1}^{T} f_t + m_{t+1}, w \right\rangle + \mathcal{R}_H(w), \tag{31}$$

where $\Delta_N \subseteq \mathbb{R}^N$ is the probability simplex and $\mathcal{R}_H(w_t) = \sum_{i=1}^N w^i_t \log w^i_t$ is a 1-strongly convex function with respect to $\|\cdot\|_1$. The update procedure (31) is known as optimistic FTRL [Rakhlin and Sridharan, 2013, Syrgkanis et al., 2015], whose static regret can be bounded by the following lemma.

Lemma 7 (Theorem 19 in Syrgkanis et al. [2015])**.** *The regret of a player under optimistic FTRL and with respect to any $w \in \Delta_N$ is upper bounded by:*

$$\sum_{t=1}^{T}\langle f_t, w_t - w\rangle \le \frac{\ln N + 2}{\epsilon} + \sum_{t=1}^{T}\|f_t - m_t\|_\infty \|w_t - \widehat{w}_{t+1}\|_1, \tag{32}$$

where $\widehat{w}_{t+1} = \arg\min_{w\in\Delta_N} \epsilon\langle\sum_{t=1}^{T+1} f_t, w\rangle + \mathcal{R}_H(w)$.

The right hand side of (32) can be further bounded by the stability of the FTRL type algorithm . Since $\mathcal{R}_H$ is 1-strongly convex with respect to $\|\cdot\|_1$, according to Lemma 9, we have

$$\|w_t - \widehat{w}_{t+1}\|_1 \le \epsilon \|f_t - m_t\|_\infty. \quad (33)$$

Plugging (33) into (32) and setting w as the vector with i-th entry being 1 and the rest being 0, we complete the proof with the learning rate $\epsilon = \sqrt{(2 + \ln N)/D_\infty}$. □

C Additional Experiments

In this section, we offer additional experimental results to further justify the effectiveness and efficiency of our proposal, where binary classification tasks with squared hinge loss are considered. Specifically, at iteration t, a sample (z_t, y_t) arrives with predictable sequences M_t, where $z_t \in \mathbb{R}^d$ is the feature vector and $y_t \in \{+1, -1\}$ is the label. Then the learner makes prediction x_t and suffers the squared loss $f_t(x_t) = (1 - y_t \cdot \langle z_t, x_t \rangle)^2$.

Setting. The experimental setup is similar to that in the main paper, where OMD, AOMD and our proposal with $\mathcal{R}(x) = \frac{1}{2}\|x\|_2^2$ are compared. We use three classification datasets [1]: IJCNN1, mushroom and svmguide1 to simulate the non-stationary environments, whose labels flip 9 times over 1200 iterations. All configurations are repeated 5 times and we normalize the features of the datasets to range from $[0, 1]$ at the beginning of the learning process.

Results. Table 1 and Table 2 report contenders' cumulative loss and running time in seconds, when predictable sequences of middle quality are provided ($\lambda = 0.4$). The results show the same tendency as those in the main paper, where our proposal achieves comparable performance with AOMD and reduces the computational burden dramatically.

Table 1: Mean and standard deviation of contenders' cumulative loss, where the best algorithms are emphasized in bold (paired t-test at 95% significance level)

Dataset	OMD	AOMD	Ours
IJCNN1	778 ± 31	$\mathbf{583 \pm 15}$	$\mathbf{590 \pm 17}$
svmguide1	918 ± 11	$\mathbf{720 \pm 11}$	763 ± 11
mushroom	2234 ± 63	1968 ± 73	$\mathbf{1850 \pm 99}$

Table 2: Mean and standard deviation of contenders' running time in seconds, where the best algorithms are emphasized in bold (paired t-test at 95% significance level)

Dataset	OMD	AOMD	Ours
IJCNN1	$\mathbf{0.02 \pm 0.011}$	5541 ± 173	0.43 ± 0.200
svmguide1	$\mathbf{0.02 \pm 0.011}$	6258 ± 167	0.27 ± 0.064
mushroom	$\mathbf{0.02 \pm 0.013}$	5001 ± 171	0.37 ± 0.070

D Technical Lemmas

In this section, we first introduce the lemma on the property of Bregman divergence.

Lemma 8 (Beck and Teboulle [2003]). *Let $\mathcal{X}$ be a convex set in a Banach space $\mathcal{B}$, and regularizer $\mathcal{R} : \mathcal{X} \mapsto \mathbb{R}$ be a 1-strongly convex function on $\mathcal{X}$ with respect to a norm $\|\cdot\|$, and let $D_{\mathcal{R}}(\cdot, \cdot)$ be the Bregman divergence induced by $\mathcal{R}$. Then, any update of the form*

$$x^* = \arg\min_{x \in \mathcal{X}} \{\langle a, x \rangle + D_{\mathcal{R}}(x, c)\}$$

satisfies the following inequality

$$\langle x^* - d, a \rangle \le D_{\mathcal{R}}(d, c) - D_{\mathcal{R}}(d, x^*) - D_{\mathcal{R}}(x^*, c),$$

for any $d \in \mathcal{X}$.

Next, the following lemma shows the stability of the FTRL algorithm.

Lemma 9 (Lemma 20 in Syrgkanis et al. [2015]). *If $w_* = \arg\min_{w \in \mathcal{X}} \epsilon \langle w, F \rangle + \mathcal{R}_H(w)$ and $w'_* = \arg\min_{w \in \mathcal{X}} \epsilon \langle w, F' \rangle + \mathcal{R}(w)$ for a λ-strongly convex regularizer $\mathcal{R} : \mathcal{W} \mapsto \mathbb{R}$ with respect to a norm $\|\cdot\|$ and some $F \in \mathbb{R}^d$ and $F' \in \mathbb{R}^d$. Then we have*

$$\lambda \|w_* - w'_*\| \le \epsilon \|F - F'\|_*.$$

[1] The datasets are downloaded from `https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/`

References

A. Beck and M. Teboulle. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operation Research Letter*, 31(3):167–175, 2003.

O. Besbes, Y. Gur, and A. J. Zeevi. Non-stationary stochastic optimization. *Operations Research*, 63(5):1227–1244, 2015.

M. Mohri, A. Rostamizadeh, and A. Talwalkar. *Foundations of Machine Learning*. The MIT Press, second edition, 2018.

A. Rakhlin and K. Sridharan. Online learning with predictable sequences. In *Proceedings of the 26th Conference On Learning Theory*, pages 993–1019, 2013.

V. Syrgkanis, A. Agarwal, H. Luo, and R. E. Schapire. Fast convergence of regularized learning in games. In *Advances in Neural Information Processing Systems 28*, pages 2989–2997, 2015.