Regret Bounds for Optimistic Follow The Leader: Applications in Portfolio Selection and Linear Regression

Sudeep Raja Putta SP3794@COLUMBIA.EDU
Shipra Agrawal SA3305@COLUMBIA.EDU
Columbia University

Abstract

Follow-The-Leader (FTL) is a simple online learning algorithm that is often overlooked. This paper investigates the FTL algorithm and two of its variant. The Optimistic FTL (OFTL) algorithm in the context of online learning with hints and the Follow The Approximate Leader (FTAL) with adaptive curvature. We provide a general regret inequality for OFTL that explicitly captures the effect of the hints and the curvature of the cost functions. This directly leads to a regret bound for FTAL. We generalize prior regret bounds of FTAL by incorporating adaptive curvature and movement of the iterates. We demonstrate the applicability of our results by deriving regret bounds for the online portfolio selection problem using FTAL with adaptive curvature. We further show the applicability of OFTL by obtaining a uniform regret bound for online linear regression. Our analysis contributes to a better understanding of FTL and its variants in various online learning scenarios.

1. Introduction

We focus on *online convex optimization* (OCO). Here, a player interacts with an environment for T rounds. In each round, the player selects an action from a convex set, $w_t \in \mathcal{D} \subseteq \mathbb{R}^n$. The environment picks a convex function $f_t : \mathcal{D} \to \mathbb{R}$. The player incurs a $f_t(w_t)$ cost and observes f_t. The player's objective is to minimize $\sum_{t=1}^{T} f_t(w_t)$. The *regret* of the player compared to the cost of fixed point $w \in \mathcal{D}$ is $\mathcal{R}_T(w) = \sum_{t=1}^{T} f_t(w_t) - f_t(w)$. A straightforward strategy for the player is to select w_t using *Follow The Leader* (FTL). FTL can be succinctly expressed as $w_t \in \arg\min_{w\in\mathcal{D}} \sum_{s=1}^{t-1} f_s(w)$.

Unfortunately, FTL can have $O(T)$ regret even with linear functions [21, Example 2.10]. This occurs because FTL's iterates can be forced into alternating between opposite corners of $\mathcal{D}$ in every iteration, making it *"unstable"*. Nevertheless, FTL has $O(\log T)$ regret when the functions are strongly convex [21, Corollary 7.24]. Even for linear functions, FTL's regret is $O(\log T)$ if the decision set's boundary exhibits sufficient curvature [13, 14].

In the framework of optimistic online learning, at each round t the player is presented with a hint function, m_t, which can be utilized in the selection of w_t. Following this, the cost function, f_t, is disclosed, and the player incurs a cost of $f_t(w_t)$. If m_t serves as a beneficial hint for f_t, it could potentially lead to a reduction in regret. The Optimistic FTL (OFTL) can be succinctly expressed as $w_t \in \arg\min_{w\in\mathcal{D}} m_t(w) + \sum_{s=1}^{t-1} f_s(w)$. One can construct *surrogate* convex functions and run FTL/OFTL on them instead. This approach leads to the *Follow The Approximate Leader* (FTAL).

This paper examines OFTL and FTAL for problems with adaptive curvature on both bounded and unbounded domains. We illustrate the relevance of our findings by considering two specific problems: online portfolio selection, which exhibits adaptive curvature and a bounded domain, and online linear regression, which has constant curvature and an unbounded domain.

© S.R. Putta & S. Agrawal.

1.1. Our Contributions

In this paper, we introduce a general regret inequality for OFTL that explicitly considers the curvature of the cost functions and hints. When applied to quadratic surrogate functions, it leads to a regret bound for FTAL. By incorporating adaptive curvature, we extend the previous regret bounds for FTAL that were obtained by Hazan et al. [11]. We demonstrate the practicality of our findings by deriving new regret bounds for the online portfolio selection problem, and a uniform regret bound for online linear regression comparable to similar results in [8].

1.2. Notation

Let $\mathbb{R}^n_+$ be the non-negative orthant of $\mathbb{R}^n$ and $\Delta_n = \{w \in \mathbb{R}^n_+ : \sum_{i=1}^n w_i = 1\}$ be the probability simplex. Let $\mathcal{D}$ be a convex set. If $\mathcal{D}$ is bounded, its diameter is $D = \max_{w,w' \in \mathcal{D}} \|w - w'\|$. The Bregman Divergence of a function F is $\mathrm{B}_F(x\|y) = F(x) - F(y) - \nabla F(y)^\top (x - y)$. Let $\nabla_t = \nabla f_t(w_t)$ and $\beta_t = \beta_t(w_t)$. For a PSD matrix X, the minimum non-zero singular value of X is $\sigma_{\min}(X)$, and its Moore-Penrose pseudo inverse is X^+.

2. Optimistic Follow The Leader

Let $g_t(w) = \sum_{s=1}^t f_s(w)$. Algorithm 1 describes OFTL.

Algorithm 1: Optimistic Follow The Leader

for $t = 1$ *to* T **do**
  Receive the hint function m_t
  Compute $w_t = \arg\min \|w\|$ such that $w \in \arg\min_{w \in \mathcal{D}} m_t(w) + \sum_{s=1}^{t-1} f_s(w)$
  Receive f_t
end

2.1. New Regret Inequality

Let w'_t be the iterates obtained if we ran FTL instead of OFTL, i.e., w'_t are the iterates with $m_t(w) = 0$.

Theorem 1 *For any $w \in \mathcal{D}$, any sequence of convex cost functions $f_1, \dots, f_T$ and hint functions $m_1, \dots, m_T$ such that $\arg\min_{w \in \mathcal{D}} \sum_{s=1}^{t-1} f_s(w)$ and $\arg\min_{w \in \mathcal{D}} m_t(w) + \sum_{s=1}^{t-1} f_s(w)$ are non empty, the iterates of Optimistic FTL (Algorithm 1) satisfies the inequality :*

$$\mathcal{R}_T(w) \le \sum_{t=1}^T \Big((\nabla f_t(w_t) - \nabla m_t(w_t))^\top (w_t - w'_{t+1}) - B_{g_t}(w'_{t+1}\|w_t) - B_{g_{t-1}}(w_t\|w'_t) \Big)$$

If we set $m_t(w) = 0$, we immediately obtain a regret bound for FTL.

Corollary 2 *For any $w \in \mathcal{D}$ and any sequence of convex functions $f_1, \dots, f_T$ such that $\arg\min_{w \in \mathcal{D}} \sum_{s=1}^{t-1} f_s(w)$ is non empty, the iterates of FTL (Algorithm 1 with $m_t(w) = 0$) satisfy:*

$$\mathcal{R}_T(w) \le \sum_{t=1}^T \nabla_t^\top (w_t - w_{t+1}) - B_{g_t}(w_{t+1}\|w_t) = \sum_{t=1}^T f_t(w_t) - f_t(w_{t+1}) - B_{g_{t-1}}(w_{t+1}\|w_t)$$

A comparison of our regret bound with prior results appears in Appendix A.1.

3. Follow The Approximate Leader

In this section, we set the hint $m_t(w) = 0$. Consider cost functions that have a quadratic lower-bound with adaptive curvature:

Assumption 3 *Assume the cost functions f_t satisfy:*

$$f_t(w) \geq f_t(w') + \nabla f_t(w')^\top (w - w') + \frac{\beta_t(w')}{2} (\nabla f_t(w')^\top (w - w'))^2 \quad \forall w, w' \in \mathcal{D}$$

Define the function $\hat{f}_t(w) = f_t(w_t) + \nabla f_t(w_t)^\top (w - w_t) + \frac{\beta_t(w_t)}{2} (\nabla f_t(w_t)^\top (w - w_t))^2$.

Algorithm 2: Follow The Approximate Leader

---

for $t = 1$ *to* T **do**

  Compute the current iterate $w_t = \arg\min \|w\|$ such that $w \in \arg\min_{w \in \mathcal{D}} \sum_{s=1}^{t-1} \hat{f}_s(w)$

  Receive f_t. Let $\nabla_t = \nabla f_t(w_t)$ and $\beta_t = \beta_t(w_t)$.

  Construct lower-bound $\hat{f}_t(w) = f_t(w_t) + \nabla_t^\top (w - w_t) + \frac{\beta_t}{2} (\nabla_t^\top (w - w_t))^2$

end

3.1. New Regret Inequality

We have $\hat{f}_t(w) \leq f_t(w)$ for all $w \in \mathcal{D}$ and $\hat{f}_t(w_t) = f_t(w_t)$. This implies $\mathcal{R}_T(w) \leq \sum_{t=1}^T \hat{f}_t(w_t) - \hat{f}_t(w)$. Since FTAL runs FTL on $\hat{f}_t(w)$, we can apply Corollary 2 to obtain a regret bound for FTAL.

Theorem 4 *For any $w \in \mathcal{D}$ and any sequence of convex cost functions $f_1, \ldots, f_T$ that satisfy Assumption 3, the iterates of FTAL (Algorithm 2) satisfy:*

$$\mathcal{R}_T(w) \leq \frac{n}{2 \min_t \beta_t} \left(\log \left(\frac{\mathcal{G}_T^2 \mathcal{M}_T^2 \min_t \beta_t}{n^2} + 1 \right) + 1 \right)$$

Where $\mathcal{M}_T^2 = \sum_{t=1}^T \|w_t - w_{t+1}\|_2^2$ and $\mathcal{G}_T^2 = \sum_{t=1}^T \beta_t \|\nabla_t\|_2^2$

A comparison of our regret bound with prior results appears in Appendix A.2

4. Application: Online Portfolio Selection

The online portfolio selection problem is a repeated game of sequential investment between an investor (the player) and a market (the environment) consisting of n assets (stocks). In each round, the investor selects a $w_t \in \Delta_n$ and later observes the *returns* $r_t \in \mathbb{R}_+^n$ from the market. The investor's loss function is $f_t(w) = -\log(r_t^\top w)$. Let $\hat{r} = \min_{t,i} r_t[i]$, be the smallest return observed by the player. We make the *no junk bonds* assumption.

Assumption 5 *For all t, we have $\|r_t\|_\infty = 1$. There is an unknown constant $r > 0$ (called the market variability parameter) such that $\hat{r} \geq r$.*

We summarize our contributions along with relevant prior work using FTL style algorithms in Table 1. Here, Q_T represents the quadratic variation of r_t, defined as $Q_T = \sum_{t=1}^T \|r_t - \bar{r}_T\|_2^2$, where $\bar{r}_T = \frac{1}{T} \sum_{t=1}^T r_t$ and $L_T^\star = \min_{w \in \Delta_n} \sum_{t=1}^T f_t(w)$. A detailed comparison with prior work is in Appendix A.3.

Table 1: Prior work and our contributions($\star$) in Online Portfolio Selection

Algorithm	Regret	Run-time (per round)
FTL[11]	$\hat{r}^{-2} n \log(T)$	$n^3 T$
ONS/FTAL [1, 11]	$r^{-1} n \log(T)$	n^3
FTAL+Adaptive Curvature*	$\hat{r}^{-1} n \log(T)$	n^3
Exp-Concave FTL [9, 10]	$\hat{r}^{-2} n \log(Q_T + n)$	$n^3 T$
FTRL [9, 10]	$r^{-3} n \log(Q_T + n)$	n^3
FTRL+Adaptive Curvature*	$\hat{r}^{-2} n \log(Q_T + n)$	n^3
ONS [22]	$r^{-1} n \log(r^{-2} L_T^\star + r^{-3} n)$	n^3
FTRL+Adaptive Curvature*	$\hat{r}^{-1} n \log(\hat{r}^{-2} L_T^\star + \hat{r}^{-3} n)$	n^3

4.1. New Adaptive Curvature Lower-bound and Regret bounds

Under Assumption 5, we have that $0 < r \leq \hat{r} \leq r_t^\top w \leq 1$ for all $w \in \Delta_n$. So, we apply the result:

Lemma 6 *[24] For all $x, y \in (0, 1]$, we have:* $\frac{y}{x} - 1 - \log\left(\frac{y}{x}\right) \geq \frac{1}{2}\frac{(x-y)^2}{x}$

This leads to an adaptive curvature quadratic lower-bound with $\beta_t(w) = r_t^\top w$:

Lemma 7 *Under Assumption 5, for all $w, w' \in \Delta_n$, we have for $f_t(w) = -\ln(r_t^\top w)$:*

$$f_t(w) \geq f_t(w') + \nabla f_t(w')^\top (w - w') + \frac{(r_t^\top w')}{2}(\nabla f_t(w')^\top (w - w'))^2$$

Using the FTAL algorithm (Algorithm 2) with this lower-bound, we have:

Theorem 8 *For any $w \in \Delta_n$ and any sequence of returns $r_1, \ldots, r_T$ such that Assumption 5 holds, the iterates of FTAL (Algorithm 2), with $\nabla_t = -r_t/(r_t^\top w_t)$ and $\beta_t = r_t^\top w_t$ satisfy the inequality:*

$$\mathcal{R}_T(w) \leq \frac{n}{2\hat{r}}\left(\log\left(\frac{2T^2}{n} + 1\right) + 1\right)$$

Consider the following FTRL with ℓ_2 regularization:

$$w_t = \arg\min_{w \in \Delta_n} \frac{1}{2}\|w\|_2^2 + \sum_{s=1}^{t-1}\left(f_s(w_s) - \frac{r_s^\top (w - w_s)}{(r_s^\top w_s)^\top} + \frac{(r_s^\top (w - w_s))^2}{2(r_s^\top w_s)}\right) \tag{1}$$

Theorem 9 *For any $w \in \Delta_n$ and any sequence of returns $r_1, \ldots, r_T$ such that Assumption 5 holds, the iterates of Equation 1 satisfy* $\mathcal{R}_T(w) = O\left(\frac{n}{\hat{r}^2}\log(Q_T + n)\right)$

This matches Exp-Concave FTL [10] while having $O(n^3)$ run time.

Lemma 10 *Under Assumption 5, the cost functions $f_t(w) = -\log(r_t^\top w)$ are $n/\hat{r}^2$-smooth on Δ_n.*

Due to cost functions having the smoothness property, we can show that the iterates of Equation 1 have a regret bounded by $O(\log L_T^\star)$.

Theorem 11 *For any $w \in \Delta_n$ and any sequence of returns $r_1, \dots, r_T$ such that Assumption 5 holds, the iterates of Equation 1 satisfy $\mathcal{R}_T(w) = O\left(\frac{n}{\hat{r}} \log\left(\frac{L_T^\star}{\hat{r}^2} + \frac{n}{\hat{r}^3}\right)\right)$, where $L_T^\star = \min_{w\in\Delta_n} \sum_{t=1}^T f_t(w)$*

5. Application: Online Linear Regression

In the online linear regression problem, at each round t, a feature vector $x_t \in \mathbb{R}^n$ is revealed to the player. The player then picks an estimator $w_t \in \mathbb{R}^n$ and predicts the value $\hat{y}_t = x_t^\top w_t$. The environment then reveals the value y_t, and the player pays the cost $f_t(w) = \frac{1}{2}(x_t^\top w_t - y_t)^2$. We can analyze the regret using optimistic OCO, as the feature vector x_t could be used to create a hint.

5.1. New regret bound

We consider the same forecaster as [8], that runs OFTL (algorithm 1) on the costs $f_t(w) = \frac{1}{2}(x_t^\top w - y_t)^2$ and uses hint $m_t(w) = \frac{1}{2}(x_t^\top w)^2$. Let $X_t = \sum_{s=1}^t x_s x_s^\top$ and $Y_t = \sum_{s=1}^t y_s x_s$. We have:

Theorem 12 *For any $w \in \mathbb{R}^n$ and any sequence of pairs $(x_1, y_1), \dots, (x_T, y_T)$, the iterates of OFTL (Algorithm 1) with $f_t(w) = \frac{1}{2}(x_t^\top w - y_t)^2$, $m_t(w) = \frac{1}{2}(x_t^\top w)^2$ and $\mathcal{D} = \mathbb{R}^n$, satisfy the regret the inequality $\sum_{t=1}^T \frac{1}{2}(x_t^\top w_t - y_t)^2 - \frac{1}{2}(x_t^\top w - y_t)^2 \le \frac{(\max_t y_t^2)n}{2}\left(\log\left(\frac{\mathcal{X}_T^2 \mathcal{M}'^2_T}{n^2(\max_t y_t^2)} + 1\right) + 1\right)$, Where $\mathcal{X}_T^2 = \sum_{t=1}^T \|x_t\|_2^2$ and $\mathcal{M}'^2_T = \sum_{t=1}^T \|w_t - w'_{t+1}\|_2^2$. This implies the regret bound:*

$$\le \frac{(\max_t y_t^2)n}{2}\left(\log\left(\frac{\sum_{t=1}^T \|x_t\|_2^2}{n^2}\left(\sum_{t=1}^T \frac{\|x_t\|_2^2}{\sigma_{\min}(X_t)^2}\right) + 1\right) + 1\right)$$

Since this holds for all $w \in \mathbb{R}^n$, it is a uniform regret bound. Moreover, it is also scale-invariant to multiplying features by a scalar constant. However, the $\sigma_{\min}(X_t)$ term in the regret bound may be small, leading to large regret. However, for many sequences of feature vectors x_t this term could be reasonable. Moreover, this bound has the optimal leading term $(\max_t y_t^2)n$ in front of the logarithm.

6. Conclusions

In this paper, we studied bounds for the Optimistic FTL algorithm that explicitly shows the effect of the curvature of the cost functions and hints on regret. We derive a regret bound for the FTAL algorithm for cost functions with an adaptive quadratic lower-bound. In contrast to prior work, this regret bound holds for both bounded and unbounded domains. For bounded domains, FTAL has a better data-dependant leading constant.

For the online portfolio selection problem, we first show a new adaptive curvature quadratic lower-bound. Using this lower-bound with the FTAL algorithm, we show a $O(\hat{r}^{-1} n \log(T))$ regret bound. Using the same lower-bound with a ℓ_2 regularized FTRL algorithm, we show a $O(\hat{r}^{-2} n \log(Q_T + n))$ regret and a first-order $O(\hat{r}^{-1} n \log(\hat{r}^{-2} L_T^\star + \hat{r}^{-3} n))$ regret bound, where $\hat{r} = \min_{t,i} r_t[i]$.

As a future research direction, it is interesting to explore if the adaptive curvature lower-bound could be used to obtain an $O(n \log(T)))$ regret algorithm that runs in $O(n^3)$ time per iteration. A natural definition of variation in the returns of the online portfolio selection problem would be

$V_T = \sum_{t=2}^{T} \|r_t - r_{t-1}\|$. Applying the variation regret bounds from [5] or Orabona [21, Corollary 7.38] for instance, does not give us regret bounds depending on V_T. Thus, obtaining an $O(\log(V_T))$ is also an open problem.

For the online linear regression problem, we use OFTL to obtain a uniform regret bound that holds for all $w \in \mathbb{R}^n$. While this bound has the optimal leading term, the term inside the logarithm is not uniformly bounded for all sequences of feature vectors and may be very large. An open problem posed in Gaillard et al. [8] asks if it is possible to obtain a doubly uniform regret bound that holds for all $w \in \mathbb{R}^n$ and all features x_t such that $\|x_t\| \leq X$.

References

[1] Amit Agarwal, Elad Hazan, Satyen Kale, and Robert E. Schapire. Algorithms for portfolio management based on the newton method. In William W. Cohen and Andrew W. Moore, editors, *Machine Learning, Proceedings of the Twenty-Third International Conference (ICML 2006), Pittsburgh, Pennsylvania, USA, June 25-29, 2006*, volume 148 of *ACM International Conference Proceeding Series*, pages 9–16. ACM, 2006. doi: 10.1145/1143844.1143846. URL https://doi.org/10.1145/1143844.1143846.

[2] Katy S. Azoury and Manfred K. Warmuth. Relative loss bounds for on-line density estimation with the exponential family of distributions. *Mach. Learn.*, 43(3):211–246, 2001. doi: 10.1023/A:1010896012157. URL https://doi.org/10.1023/A:1010896012157.

[3] Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006. ISBN 978-0-521-84108-5. doi: 10.1017/CBO9780511546921.

[4] Nicolò Cesa-Bianchi and Francesco Orabona. Online learning algorithms. *Annual Review of Statistics and Its Application*, 8(1):165–190, 2021. doi: 10.1146/annurev-statistics-040620-035329.

[5] Chao-Kai Chiang, Tianbao Yang, Chia-Jung Lee, Mehrdad Mahdavi, Chi-Jen Lu, Rong Jin, and Shenghuo Zhu. Online optimization with gradual variations. In Shie Mannor, Nathan Srebro, and Robert C. Williamson, editors, *COLT 2012 - The 25th Annual Conference on Learning Theory, June 25-27, 2012, Edinburgh, Scotland*, volume 23 of *JMLR Proceedings*, pages 6.1–6.20. JMLR.org, 2012. URL http://proceedings.mlr.press/v23/chiang12/chiang12.pdf.

[6] Thomas M. Cover. Universal portfolios. *Mathematical Finance*, 1(1):1–29, 1991. doi: https://doi.org/10.1111/j.1467-9965.1991.tb00002.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-9965.1991.tb00002.x.

[7] Steven de Rooij, Tim van Erven, Peter D. Grünwald, and Wouter M. Koolen. Follow the leader if you can, hedge if you must. *J. Mach. Learn. Res.*, 15(1):1281–1316, 2014. URL http://dl.acm.org/citation.cfm?id=2638576.

[8] Pierre Gaillard, Sébastien Gerchinovitz, Malo Huard, and Gilles Stoltz. Uniform regret bounds over $\mathbb{R}^d$ for the sequential linear regression problem with the square loss. In Aurélien Garivier and Satyen Kale, editors, *Algorithmic Learning Theory, ALT 2019, 22-24 March 2019, Chicago,*

Illinois, USA, volume 98 of *Proceedings of Machine Learning Research*, pages 404–432. PMLR, 2019. URL http://proceedings.mlr.press/v98/gaillard19a.html.

[9] Elad Hazan and Satyen Kale. On stochastic and worst-case models for investing. In Yoshua Bengio, Dale Schuurmans, John D. Lafferty, Christopher K. I. Williams, and Aron Culotta, editors, *Advances in Neural Information Processing Systems 22: 23rd Annual Conference on Neural Information Processing Systems 2009. Proceedings of a meeting held 7-10 December 2009, Vancouver, British Columbia, Canada*, pages 709–717. Curran Associates, Inc., 2009. URL https://proceedings.neurips.cc/paper/2009/hash/26337353b7962f533d78c762373b3318-Abstract.html.

[10] Elad Hazan and Satyen Kale. An online portfolio selection algorithm with regret logarithmic in price variation. *Mathematical Finance*, 25(2):288–310, 2015. doi: https://doi.org/10.1111/mafi.12006. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/mafi.12006.

[11] Elad Hazan, Amit Agarwal, and Satyen Kale. Logarithmic regret algorithms for online convex optimization. *Mach. Learn.*, 69(2-3):169–192, 2007. doi: 10.1007/s10994-007-5016-8. URL https://doi.org/10.1007/s10994-007-5016-8.

[12] David P. Helmbold, Robert E. Schapire, Yoram Singer, and Manfred K. Warmuth. On-line portfolio selection using multiplicative updates. *Mathematical Finance*, 8(4):325–347, 1998. doi: https://doi.org/10.1111/1467-9965.00058. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/1467-9965.00058.

[13] Ruitong Huang, Tor Lattimore, András György, and Csaba Szepesvári. Following the leader and fast rates in linear prediction: Curved constraint sets and other regularities. In Daniel D. Lee, Masashi Sugiyama, Ulrike von Luxburg, Isabelle Guyon, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 4970–4978, 2016.

[14] Ruitong Huang, Tor Lattimore, András György, and Csaba Szepesvári. Following the leader and fast rates in online linear prediction: Curved constraint sets and other regularities. *Journal of Machine Learning Research*, 18(145):1–31, 2017. URL http://jmlr.org/papers/v18/17-079.html.

[15] Rémi Jézéquel, Dmitrii M Ostrovskii, and Pierre Gaillard. Efficient and near-optimal online portfolio selection. *arXiv preprint arXiv:2209.13932*, 2022.

[16] Adam Kalai and Santosh S. Vempala. Efficient algorithms for universal portfolios. *J. Mach. Learn. Res.*, 3:423–440, 2002. URL http://jmlr.org/papers/v3/kalai02a.html.

[17] Adam Tauman Kalai and Santosh S. Vempala. Efficient algorithms for online decision problems. *J. Comput. Syst. Sci.*, 71(3):291–307, 2005. doi: 10.1016/j.jcss.2004.10.016. URL https://doi.org/10.1016/j.jcss.2004.10.016.

[18] Bin Li and Steven Chu Hong Hoi. *Online portfolio selection: principles and algorithms*. Crc Press, 2018.

[19] Haipeng Luo, Chen-Yu Wei, and Kai Zheng. Efficient online portfolio with logarithmic regret. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 8245–8255, 2018. URL https://proceedings.neurips.cc/paper/2018/hash/91c77393975889bd08f301c9e13a44b7-Abstract.html.

[20] Zakaria Mhammedi and Alexander Rakhlin. Damped online newton step for portfolio selection. In Po-Ling Loh and Maxim Raginsky, editors, *Conference on Learning Theory, 2-5 July 2022, London, UK*, volume 178 of *Proceedings of Machine Learning Research*, pages 5561–5595. PMLR, 2022. URL https://proceedings.mlr.press/v178/mhammedi22b.html.

[21] Francesco Orabona. A modern introduction to online learning. *CoRR*, abs/1912.13213, 2019.

[22] Francesco Orabona, Nicolò Cesa-Bianchi, and Claudio Gentile. Beyond logarithmic bounds in online learning. In Neil D. Lawrence and Mark A. Girolami, editors, *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2012, La Palma, Canary Islands, Spain, April 21-23, 2012*, volume 22 of *JMLR Proceedings*, pages 823–831. JMLR.org, 2012. URL http://proceedings.mlr.press/v22/orabona12.html.

[23] Laurent Orseau, Tor Lattimore, and Shane Legg. Soft-bayes: Prod for mixtures of experts with log-loss. In Steve Hanneke and Lev Reyzin, editors, *International Conference on Algorithmic Learning Theory, ALT 2017, 15-17 October 2017, Kyoto University, Kyoto, Japan*, volume 76 of *Proceedings of Machine Learning Research*, pages 372–399. PMLR, 2017. URL http://proceedings.mlr.press/v76/orseau17a.html.

[24] Sudeep Raja Putta and Shipra Agrawal. Scale-free adversarial multi armed bandits. In Sanjoy Dasgupta and Nika Haghtalab, editors, *International Conference on Algorithmic Learning Theory, 29 March - 1 April 2022, Paris, France*, volume 167 of *Proceedings of Machine Learning Research*, pages 910–930. PMLR, 2022. URL https://proceedings.mlr.press/v167/putta22a.html.

[25] Alexander Rakhlin and Karthik Sridharan. Optimization, learning, and games with predictable sequences. In Christopher J. C. Burges, Léon Bottou, Zoubin Ghahramani, and Kilian Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, pages 3066–3074, 2013. URL https://proceedings.neurips.cc/paper/2013/hash/f0dd4a99fba6075a9494772b58f95280-Abstract.html.

[26] Nathan Srebro, Karthik Sridharan, and Ambuj Tewari. Smoothness, low noise and fast rates. In John D. Lafferty, Christopher K. I. Williams, John Shawe-Taylor, Richard S. Zemel, and Aron Culotta, editors, *Advances in Neural Information Processing Systems 23: 24th Annual Conference on Neural Information Processing Systems 2010. Proceedings of a meeting held*

6-9 December 2010, Vancouver, British Columbia, Canada, pages 2199–2207. Curran Associates, Inc., 2010. URL https://proceedings.neurips.cc/paper/2010/hash/76cf99d3614e23eabab16fb27e944bf9-Abstract.html.

[27] Chung-En Tsai, Hao-Chung Cheng, and Yen-Huan Li. Online self-concordant and relatively smooth minimization, with applications to online portfolio selection and learning quantum states. In Shipra Agrawal and Francesco Orabona, editors, *International Conference on Algorithmic Learning Theory, February 20-23, 2023, Singapore*, volume 201 of *Proceedings of Machine Learning Research*, pages 1481–1483. PMLR, 2023. URL https://proceedings.mlr.press/v201/tsai23a.html.

[28] Tim van Erven, Wouter M. Koolen, and Dirk van der Hoeven. Metagrad: Adaptation using multiple learning rates in online learning. *J. Mach. Learn. Res.*, 22:161:1–161:61, 2021. URL http://jmlr.org/papers/v22/20-1444.html.

[29] Volodya Vovk. Competitive on-line statistics. *International Statistical Review*, 69 (2):213–248, 2001. doi: https://doi.org/10.1111/j.1751-5823.2001.tb00457.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1751-5823.2001.tb00457.x.

[30] Julian Zimmert, Naman Agarwal, and Satyen Kale. Pushing the efficiency-regret pareto frontier for online learning of portfolios and quantum states. In Po-Ling Loh and Maxim Raginsky, editors, *Conference on Learning Theory, 2-5 July 2022, London, UK*, volume 178 of *Proceedings of Machine Learning Research*, pages 182–226. PMLR, 2022. URL https://proceedings.mlr.press/v178/zimmert22a.html.

Appendix A. Prior Work

A.1. Optimistic FTL

Regret bounds in the optimistic online learning literature have mostly been limited to optimistic FTRL with linear costs [25], [21, Theorem 7.35]. Our regret bound in Theorem 1 generalizes these bounds to convex cost functions and hints. Regret bounds for FTL were initially derived using the Be-The-Leader (BTL) lemma Kalai and Vempala [17]. Their FTL regret bound is $\mathcal{R}_T(w) \leq \sum_{t=1}^{T} f_t(w_t) - f_t(w_{t+1})$. Our regret bound includes an additional term, $-\mathrm{B}_{g_{t-1}}(w_{t+1} \| w_t)$.

Orabona [21, Lemma 7.1] provides a general regret equality that can derive regret bounds for various online algorithms. It implies the FTL bound: $\mathcal{R}_T(w) \leq \sum_{t=1}^{T} f_t(w_t) - f_t(w_{t+1}) + g_{t-1}(w_t) - g_{t-1}(w_{t+1})$. Comparing this with Corollary 2, we have the extra term $\nabla g_{t-1}(w_t)^\top(w_{t+1} - w_t)$. If the domain is unconstrained, this term is 0 as $\nabla g_{t-1}(w_t) = 0$. In the constrained case, $\nabla g_{t-1}(w_t)^\top(w_{t+1} - w_t) \geq 0$. This term is necessary to obtain regret inequalities with the Bregman divergence term. Even in Orabona [21, Lemma 7.4], this term is added to obtain regret inequalities.

Hazan et al. [11, Theorem 5] provide a regret bound for the special case of *ridge functions*, where $f_t(w) = h_t(w^\top v_t)$ for some scalar convex function h_t. Given that $\mathcal{D}$ is bounded, $h_t'(v_t^\top w) \leq b$, $h_t''(v_t^\top w) \geq a$ for all $w \in \mathcal{D}$, and $|v_t| \leq R$, they demonstrate that the regret of FTL is $O\left(\frac{nb^2}{a}\log\left(1 + \frac{DRaT}{bn}\right)\right)$. We can derive this bound by applying these assumptions to Corollary 2.

A.2. FTAL

In the particular case where $\mathcal{D}$ is bounded, $\beta_t(w) = \beta$ and $\|\nabla f_t(w)\| \le G$ for all $w \in \mathcal{D}$, the bound in Theorem 4 becomes $O\left(n/\beta \log(DGT\beta/n)\right)$. For functions that are α exp-concave, the above bound becomes $O\left(\left(\alpha^{-1} + GD\right) n \log\left(T/n\right)\right)$. These bounds recover the FTAL bounds appearing in Hazan et al. [11]. ONS [11] uses weaker constant-curvature lower-bounds that have the worst case curvature constant $\beta_{\min} = \inf_{w\in\mathcal{D}} \beta(w)$ instead of adaptive curvature. If $\mathcal{D}$ is bounded, then picking the learning rate, leads to the regret bound $O\left(n/\beta_{\min} \log\left(D^2\beta_{\min}^2(\sum_{t=1}^T \|\nabla_t\|_2^2)/n^2 + 1\right)\right)$.

The quantity preceding FTAL's regret scales as $(\min_t \beta_t)^{-1}$, whereas for ONS, it scales as $(\beta_{\min})^{-1}$. Since $(\min_t \beta_t)^{-1}$ is a data-dependent quantity, it could be significantly smaller than $(\beta_{\min})^{-1}$. Second, ONS can only be applied when $\mathcal{D}$ is bounded, whereas FTAL can be applied in bounded and unbounded settings. Note that FTAL needs to know the gradient $\nabla f_t(w_t)$ and curvature $\beta_t(w_t)$ at each round to obtain the regret bound of Theorem 4. If $\mathcal{D}$ is bounded, $\|\nabla f_t(w)\| \le G$ and $\beta_t(w)$ is unavailable, the MetaGrad algorithm of [28] has regret bound of $O\left(n/\beta_{\min} \log\left(DGT/n\right)\right)$.

A.3. Online Portfolio Selection

Regret bounds for this problem can be divided into two categories:

Regret independent of r: Cover [6] showed that the Universal Portfolio(UP) has a regret bound of $O(n \log(T))$, but computing it requires $O(n^4 T^{14})$ run time per round[16]. There have been several works that explore the trade-off between run-time and regret [19, 20, 23, 27, 30]. Recently, [15] gave an algorithm that uses the *volumetric barrier* along with FTRL that has $O(n \log(T))$ regret and has a run time of $O(n^2 T)$. See [15] for a detailed explanation of prior work in this area.

Regret dependent on r: Helmbold et al. [12] showed that the Exponentiated Gradient(EG) algorithm has a regret of $O(r^{-1}\sqrt{T \log(n})$. However, r needs to be known to achieve this rate. Using the AdaHedge algorithm [7] that automatically tunes learning rates, it is possible to get $O(\hat{r}^{-1}\sqrt{T \log(n)})$. This is better than EG as $\hat{r} \ge r$ and r does not need to be known.

The FTL regret bound from [11, Theorem 5] can be applied directly to the portfolio selection problem, which yields a regret of $O(\hat{r}^{-2} n \log(T))$. If r is known, then using a constant curvature lower-bound, [1, 11] show that ONS and FTAL can have a $O(r^{-1} n \log(T))$ regret. In this paper, we show that using an adaptive curvature lower-bound, FTAL can have a regret of $O(\hat{r}^{-1} n \log(T))$ without knowing r. Note that MetaGrad can obtain the regret bound $O(r^{-1} n \log(T))$ without knowing r. However, the adaptive curvature FTAL is simpler compared to MetaGrad.

While algorithms with regret independent of r offer superior worst-case performance guarantees, they tend to be overly conservative. Empirical studies, such as those conducted by [1, 12, 18] demonstrate that EG and ONS outperform Cover's UP. Furthermore, Hazan and Kale [9, 10] propose algorithms akin to FTL/FTAL that exhibit regret of order $O(\log Q_T)$. Here, Q_T represents the quadratic variation of r_t, defined as $Q_T = \sum_{t=1}^T \|r_t - \bar{r}_T\|_2^2$, where $\bar{r}_T = \frac{1}{T}\sum_{t=1}^T r_t$. Using our adaptive curvature quadratic lower-bound, we can show a regret of $O(\hat{r}^{-2} n \log(Q_T + n)$.

Additionally, [22, Theorem 1] show that if the cost functions are *smooth*, ONS can achieve an $O(\log L_T^\star)$ regret, where $L_T^\star = \min_{w\in\Delta_n} \sum_{t=1}^T f_t(w)$. Using our adaptive curvature lower-bound we show a regret of $\hat{r}^{-1} n \log(\hat{r}^{-2} L_T^\star + \hat{r}^{-3} n)$.

A.4. Online Linear Regression

A detailed introduction to this topic can be found in [3, Chapter 11] and [4]. Of particular interest is the Vovk-Azoury-Warmuth (VAW) forecaster introduced by [29] and [2]. It can be interpreted as an optimistic FTRL that uses $\frac{1}{2}(x_t^\top w)$ as the hint function. The iterates are computed as: $w_t = \arg\min_{w\in\mathbb{R}^d} \frac{1}{2}\|w\|_2^2 + \frac{1}{2}(x_t^\top w)^2 + \sum_{s=1}^{t-1}\frac{1}{2}(x_s^\top w - y_s)^2$. The regret for this algorithm is $\mathcal{R}_T(w) \le \frac{\|w\|_2^2}{2} + \frac{n}{2}\log\left(1+\frac{T}{n}\max_t \|x_t\|_2^2\right)(\max_t y_t^2) \quad \forall w \in \mathbb{R}^n$.

[8] study the un-regularized version of the VAW forecaster, that is an optimistic FTL, obtained by running algorithm 1 with $m_t(w) = \frac{1}{2}(x_t^\top w)^2$ and $f_t(w) = \frac{1}{2}(x_t^\top w - y_t)^2$. They give a uniform regret [8, Theorem 11] bound that depends on $\sigma_{\min}(\sum_{s=1}^t x_s x_s^\top)$. We give a similar regret bound that is much simpler to obtain than theirs.

The regret bound of [8] is:

$$\mathcal{R}_T(w) \le n(\max_t y_t^2)(1+\log(T)) + n(\max_t y_t^2)\underbrace{\left(\frac{1}{n}\sum_{\tau\in\mathcal{T}}\log\left(\frac{\mathcal{X}^2}{\sigma_{\min}(X_\tau)}\right)\right)}_{O_T(1)}$$

Here $\mathcal{X} = \max_t \|x_t\|_2$, $X_t = \sum_{s=1}^t x_s x_s^\top$ is the t'th gram matrix and $\mathcal{T}$ is the set of indices τ such that $\text{rank}(X_{\tau-1}) \ne \text{rank}(X_\tau)$. Thus, once the gram matrices are full rank, the $O_T(1)$ term stops growing, however, it could be quite large for some sequences of x_t.

If x_{t+1} is in the span of the eigenvectors of X_t, then $\text{rank}(X_t) = \text{rank}(X_{t+1})$ and $\sigma_{\min}(X_t) \le \sigma_{\min}(X_{t+1})$. On the other hand, if x_{t+1} is not in the span of the eigenvectors of X_t, then $\text{rank}(X_t) + 1 = \text{rank}(X_{t+1})$. Thus, we can bound the sum $\sum_{t=1}^T \frac{\|x_t\|_2^2}{\sigma_{\min}(X_t)^2}$ in the regret bound of Theorem 12:

$$\sum_{t=1}^T \frac{\|x_t\|_2^2}{\sigma_{\min}(X_t)^2} \le \mathcal{X}^2 \sum_{t=1}^T \frac{1}{\sigma_{\min}(X_t)^2} \le \mathcal{X}^2 \sum_{\tau\in\mathcal{T}} \frac{t_\tau}{\sigma_{\min}(X_\tau)^2} \le \mathcal{X}^2 T \max_{\tau\in\mathcal{T}} \frac{1}{\sigma_{\min}(X_\tau)^2}$$

Here t_τ is the number of gram matrices that have the same rank as X_τ, where $\tau \in \mathcal{T}$. Applying this bound and simplifying our result, we get:

$$\mathcal{R}_T(w) \le n(\max_t y_t^2)(1+\log(T)) + n(\max_t y_t^2)\underbrace{\left(\max_{\tau\in\mathcal{T}}\log\left(\frac{\mathcal{X}^2}{n\sigma_{\min}(X_\tau)}\right)\right)}_{O_T(1)}$$

Thus, in the same vein as the regret bound of [8], our regret bound also has an $O_T(1)$ term stops growing once the gram matrices are full rank.

[8] study the un-regularized version of the VAW forecaster. They give a uniform regret bound that depends on $\sigma_{\min}(\sum_{s=1}^t x_s x_s^\top)$. We give a similar regret bound that is much simpler to obtain.

Appendix B. Auxiliary Results

Let $\text{dom}(F)$ be the domain of function F. Recall the definition of Bregman divergence $\mathrm{B}_F(x\|y) = F(x) - F(y) - \nabla F(y)^\top(x-y)$.

Lemma 13 *For any $v, w \in dom(\nabla F)$ and $u \in dom(F)$ we have:*

$$B_F(u\|w) - B_F(u\|v) - B_F(v\|w) = (\nabla F(w) - \nabla F(v))^\top(v-u)$$

Proof We can obtain the right-hand side by a straightforward expansion of the Bargeman divergences on the left-hand side. ■

Let I be the $n \times n$ identity matrix. We state the following lemma, which is a tighter version of Lemma 11 in Hazan et al. [11]

Lemma 14 *[11, Lemma 11] Let $x_1, \dots, x_t$ be a sequence of vectors in $\mathbb{R}^n$. Define $H_t = \epsilon I + \sum_{s=1}^{t} x_s x_s^\top$. Then, the following holds:*

$$\sum_{t=1}^{T} x_t^\top H_t^{-1} x_t \leq n \log\left(1 + \frac{\sum_{t=1}^{T} \|x_t\|_2^2}{n\epsilon}\right)$$

In Lemma 11 of [11], they give the bound $n \log(1 + T \sup_t \|x_t\|_2^2/\epsilon)$

We state a lemma from Putta and Agrawal [24].

Lemma 15 *[24] For all $x, y \in (0, 1]$, we have: $\frac{y}{x} - 1 - \log\left(\frac{y}{x}\right) \geq \frac{1}{2}\frac{(x-y)^2}{x}$*

In order to obtain the $O(\log Q_T)$ regret bound, we state a slightly modified version of a theorem from Hazan and Kale [10].

Theorem 16 *[10, Theorem 1.1] Let the cost functions be $f_t(w) = h_t(w^\top v_t)$ for a scalar function h_t. Consider the iterates:*

$$w_t = \arg\min_{w \in \mathcal{D}} \frac{1}{2}\|w\|_2^2 + \sum_{s=1}^{t-1} h_s(w^\top v_s)$$

If $\|v_t\| \leq R$, $\|w\| \leq D$ for all $w \in \mathcal{D}$, $h_t'(w_t^\top v_t) \in [-a, 0]$ and $h_t''(w^\top v_t) \geq b$ for all $w \in \mathcal{D}$, then:

$$\mathcal{R}_T(w) \leq O\left(\frac{a^2 n}{b}\log(1 + bQ_T + bR^2) + aRD\log(1 + Q_T/R^2) + D^2\right)$$

Here $Q_T = \min_\mu \sum_{t=1}^{T} \|v_t - \mu\|$

In the statement of the theorem in [10], they assume that $h_t = h$ for all t and $h'(w^\top v_t) \in [-a, 0]$ for all $w \in \mathcal{D}$. However, they later note that the proof of the theorem is flexible enough to handle different functions h_t for different t. Furthermore, the proof only requires the bound a on the magnitude of the first derivatives at the points w_t, which the algorithm produces, and not the entire domain $\mathcal{D}$.

Srebro et al. [26] prove the following lemma about smooth functions.

Lemma 17 *[26, Lemma 3.1] If a non-negative function f is H-smooth on the domain $\mathcal{D}$, then $\|\nabla f(w)\| \leq \sqrt{4Hf(w)}$ for all $w \in \mathcal{D}$*

Orabona et al. [22] show the following useful self-bounding result:

Lemma 18 *[22, . Corollary 5] Let $a, b, c, d, x > 0$ satisfy $x \leq a\log(bx + c) + d$, then:*

$$x \leq a\log\left(2\left(ab\log\left(\frac{2ab}{e}\right) + db + c\right)\right) + d$$

Here e is the base of the natural logarithm.

Appendix C. Proof of Theorem 1

Theorem 1 *For any $w \in \mathcal{D}$, any sequence of convex cost functions $f_1, \ldots, f_T$ and hint functions $m_1, \ldots, m_T$ such that $\arg\min_{w\in\mathcal{D}} \sum_{s=1}^{t-1} f_s(w)$ and $\arg\min_{w\in\mathcal{D}} m_t(w) + \sum_{s=1}^{t-1} f_s(w)$ are non empty, the iterates of Optimistic FTL (Algorithm 1) satisfies the inequality :*

$$\mathcal{R}_T(w) \le \sum_{t=1}^{T} \Big((\nabla f_t(w_t) - \nabla m_t(w_t))^\top (w_t - w'_{t+1}) - B_{g_t}(w'_{t+1}\|w_t) - B_{g_{t-1}}(w_t\|w'_t) \Big)$$

Proof w'_t are the iterates obtained if we had used FTL and w_t are the actual iterates of OFTL. Recall that $g_t(w) = \sum_{s=1}^{t} f_s(w)$. Consider the term $f_t(w_t) - f_t(w)$. We add and subtract $f_t(w'_{t+1})$ and use the definition of Bregman divergence to obtain:

$$\begin{aligned}
&f_t(w_t) - f_t(w) = f_t(w_t) - f_t(w'_{t+1}) + f_t(w'_{t+1}) - f(w)\\
&= \nabla f_t(w_t)^\top (w_t - w'_{t+1}) - \mathrm{B}_{f_t}(w'_{t+1}\|w_t) + \underbrace{\nabla f_t(w'_{t+1})^\top (w'_{t+1} - w)}_{(1)} - \mathrm{B}_{f_t}(w\|w'_{t+1})
\end{aligned}$$

Since $g_t(w) - g_{t-1}(w) = f_t(w)$, we have:

$$\begin{aligned}
&(1) = \nabla f_t(w'_{t+1})^\top (w'_{t+1} - w) = (\nabla g_t(w'_{t+1}) - \nabla g_{t-1}(w'_{t+1}))^\top (w_{t+1} - w)\\
&= (\nabla g_t(w'_{t+1}) - \nabla g_{t-1}(w_t))^\top (w_{t+1} - w) + \underbrace{(\nabla g_{t-1}(w_t) - \nabla g_{t-1}(w'_{t+1}))^\top (w'_{t+1} - w)}_{(2)}
\end{aligned}$$

Due to Lemma 13, the term (2) is:

$$(\nabla g_{t-1}(w_t) - \nabla g_{t-1}(w'_{t+1}))^\top (w'_{t+1} - w) = \mathrm{B}_{g_{t-1}}(w\|w_t) - \mathrm{B}_{g_{t-1}}(w\|w'_{t+1}) - \mathrm{B}_{g_{t-1}}(w'_{t+1}\|w_t)$$

Substituting this back in the expression for $f_t(w_t) - f_t(w)$, we have:

$$\begin{aligned}
f_t(w_t) - f_t(w) &= \nabla f_t(w_t)^\top (w_t - w'_{t+1}) - \mathrm{B}_{f_t}(w'_{t+1}\|w_t) - \mathrm{B}_{f_t}(w\|w'_{t+1})\\
&\quad + \mathrm{B}_{g_{t-1}}(w\|w_t) - \mathrm{B}_{g_{t-1}}(w\|w'_{t+1}) - \mathrm{B}_{g_{t-1}}(w'_{t+1}\|w_t)\\
&\quad + (\nabla g_t(w'_{t+1}) - \nabla g_{t-1}(w_t))^\top (w_t - w)\\
&= \nabla f_t(w_t)^\top (w_t - w'_{t+1}) - \mathrm{B}_{g_t}(w'_{t+1}\|w_t)\\
&\quad + \mathrm{B}_{g_{t-1}}(w\|w_t) - \mathrm{B}_{g_t}(w\|w'_{t+1})\\
&\quad + (\nabla g_t(w'_{t+1}) - \nabla g_{t-1}(w_t))^\top (w'_{t+1} - w)\\
&= (\nabla f_t(w_t) - \nabla m_t(w_t))^\top (w_t - w'_{t+1}) - \mathrm{B}_{g_t}(w'_{t+1}\|w_t)\\
&\quad + \underbrace{\mathrm{B}_{g_{t-1}}(w\|w_t) - \nabla m_t(w_t)^\top (w - w_t)}_{(3)} - \mathrm{B}_{g_t}(w\|w'_{t+1})\\
&\quad + (\nabla g_t(w'_{t+1}) - \nabla g_{t-1}(w_t) - \nabla m_t(w_t))^\top (w'_{t+1} - w)
\end{aligned}$$

In the last expression, note that all the $\nabla m_t(w_t)$ terms add to 0. Consider term (3):

$$\begin{aligned}
\mathrm{B}_{g_{t-1}}(w\|w_t) - \nabla m_t(w_t)^\top (w - w_t) &= g_{t-1}(w) - g_{t-1}(w_t)\\
&\quad - (\nabla g_{t-1}(w_t) + \nabla m_t(w_t))^\top (w - w_t)\\
&= \mathrm{B}_{g_{t-1}}(w\|w'_t) - \mathrm{B}_{g_{t-1}}(w_t\|w'_t) + \nabla g_{t-1}(w'_t)(w - w_t)\\
&\quad - (\nabla g_{t-1}(w_t) + \nabla m_t(w_t))^\top (w - w_t)
\end{aligned}$$

Substituting this back in the expression for $f_t(w_t) - f_t(w)$, we have:

$$\begin{aligned}
f_t(w_t) - f_t(w) &= (\nabla f_t(w_t) - \nabla m_t(w_t))^\top (w_t - w'_{t+1}) - \mathrm{B}_{g_t}(w'_{t+1} \| w_t) - \mathrm{B}_{g_{t-1}}(w_t \| w'_t) \\
&\quad + \mathrm{B}_{g_{t-1}}(w \| w'_t) - \mathrm{B}_{g_t}(w \| w'_{t+1}) \\
&\quad + (\nabla g_t(w'_{t+1}) - \nabla g_{t-1}(w_t) - \nabla m_t(w_t))^\top (w'_{t+1} - w) \\
&\quad + (\nabla g_{t-1}(w_t) + \nabla m_t(w_t) - \nabla g_{t-1}(w'_t))^\top (w_t - w) \\
&= (\nabla f_t(w_t) - \nabla m_t(w_t))^\top (w_t - w'_{t+1}) - \mathrm{B}_{g_t}(w'_{t+1} \| w_t) - \mathrm{B}_{g_{t-1}}(w_t \| w'_t) \\
&\quad + \mathrm{B}_{g_{t-1}}(w \| w'_t) - \mathrm{B}_{g_t}(w \| w'_{t+1}) \\
&\quad + (\nabla g_{t-1}(w'_t) - \nabla g_t(w'_{t+1}))^\top w \\
&\quad + (\nabla g_{t-1}(w_t) + \nabla m_t(w_t))^\top (w_t - w'_{t+1}) \\
&\quad + \nabla g_t(w'_{t+1})^\top w'_{t+1} - \nabla g_{t-1}(w'_t)^\top w_t
\end{aligned}$$

Since w_t minimizes $g_{t-1}(w) + m_t(w)$, we have $(\nabla g_{t-1}(w_t) + \nabla m_t(w_t))^\top (w_t - w'_{t+1}) \leq 0$

$$\begin{aligned}
&\leq (\nabla f_t(w_t) - \nabla m_t(w_t))^\top (w_t - w'_{t+1}) - \mathrm{B}_{g_t}(w'_{t+1} \| w_t) - \mathrm{B}_{g_{t-1}}(w_t \| w'_t) \\
&\quad + \mathrm{B}_{g_{t-1}}(w \| w'_t) - \mathrm{B}_{g_t}(w \| w'_{t+1}) \\
&\quad + (\nabla g_{t-1}(w'_t) - \nabla g_t(w'_{t+1}))^\top w \\
&\quad + \nabla g_t(w'_{t+1})^\top w'_{t+1} - \nabla g_{t-1}(w'_t)^\top w_t
\end{aligned}$$

Taking the summation over the t terms $\sum_{t=1}^T f_t(w_t) - f_t(w)$, we have:

$$\begin{aligned}
\mathcal{R}_T(w) \leq \sum_{t=1}^T &(\nabla f_t(w_t) - \nabla m_t(w_t))^\top (w_t - w'_{t+1}) - \mathrm{B}_{g_t}(w'_{t+1} \| w_t) - \mathrm{B}_{g_{t-1}}(w_t \| w'_t) \\
&+ \underbrace{\sum_{t=1}^T \mathrm{B}_{g_{t-1}}(w \| w'_t) - \mathrm{B}_{g_t}(w \| w'_{t+1})}_{(4)} \\
&+ \underbrace{\sum_{t=1}^T (\nabla g_{t-1}(w'_t) - \nabla g_t(w'_{t+1}))^\top w}_{(5)} \\
&+ \underbrace{\sum_{t=1}^T \nabla g_t(w'_{t+1})^\top w'_{t+1} - \nabla g_{t-1}(w'_t)^\top w_t}_{(6)}
\end{aligned}$$

Note that $g_0(w) = 0$ and the hint after round T, i.e., $m_{T+1}(w) = 0$. We can telescope term (4) to get:

$$\sum_{t=1}^T \mathrm{B}_{g_{t-1}}(w \| w'_t) - \mathrm{B}_{g_t}(w \| w'_{t+1}) = \mathrm{B}_{g_0}(w \| w'_1) - \mathrm{B}_{g_T}(w \| w'_{T+1}) = 0 - \mathrm{B}_{g_T}(w \| w'_{T+1}) \leq 0$$

Term (5) can be telescoped as:

$$\sum_{t=1}^{T}(\nabla g_{t-1}(w'_t) - \nabla g_t(w'_{t+1}))^\top w = (\nabla g_0(w'_1) - \nabla g_T(w'_{T+1}))^\top w = -\nabla g_T(w'_{T+1})^\top w$$

Since $m_{T+1}(w) = 0$, we have $w_{T+1} = w'_{T+1}$.Finally for term (6):

$$\begin{aligned}\sum_{t=1}^{T}\nabla g_t(w'_{t+1})^\top w'_{t+1} - \nabla g_{t-1}(w'_t)^\top w_t &= \sum_{t=1}^{T-1}\nabla g_t(w'_{t+1})^\top (w'_{t+1} - w_{t+1}) + \nabla g_T(w'_{T+1})^\top w_{T+1} \\ &\leq \nabla g_T(w'_{T+1})^\top w'_{T+1}\end{aligned}$$

Here, we used the fact that w'_{t+1} minimizes $g_t(w)$. So $\nabla g_t(w'_{t+1})^\top (w'_{t+1} - w) \leq 0$ for all $w \in \mathcal{D}$. Combining the upper bounds for terms (5) and (6):

$$(5) + (6) \leq g_T(w'_{T+1})^\top (w'_{T+1} - w) \leq 0$$

Thus, we have the result:

$$\mathcal{R}_T(w) \leq \sum_{t=1}^{T}(\nabla f_t(w_t) - \nabla m_t(w_t))^\top (w_t - w'_{t+1}) - \mathrm{B}_{g_t}(w'_{t+1}\|w_t) - \mathrm{B}_{g_{t-1}}(w_t\|w'_t)$$

■

Appendix D. Proof of Theorem 4

Theorem 4 *For any $w \in \mathcal{D}$ and any sequence of convex cost functions $f_1, \ldots, f_T$ that satisfy Assumption 3, the iterates of FTAL (Algorithm 2) satisfy:*

$$\mathcal{R}_T(w) \leq \frac{n}{2\min_t \beta_t}\left(\log\left(\frac{\mathcal{G}_T^2 \mathcal{M}_T^2 \min_t \beta_t}{n^2} + 1\right) + 1\right)$$

Where $\mathcal{M}_T^2 = \sum_{t=1}^{T} \|w_t - w_{t+1}\|_2^2$ and $\mathcal{G}_T^2 = \sum_{t=1}^{T} \beta_t \|\nabla_t\|_2^2$

Proof Algorithm 2 runs FTL on on $\hat{f}_t(w)$ where $\hat{f}_t(w)$ is defined as:

$$\hat{f}_t(w) = f_t(w_t) + \nabla f_t(w_t)^\top (w - w_t) + \frac{\beta_t(w_t)}{2}(\nabla f_t(w_t)^\top (w - w_t))^2$$

As $\hat{f}_t(w) \leq f_t(w)$ for all $w \in \mathcal{D}$ and $\hat{f}_t(w_t) = f_t(w_t)$, we have:

$$\mathcal{R}_T(w) = \sum_{t=1}^{T} f_t(w_t) - f_t(w) \leq \sum_{t=1}^{T} \hat{f}_t(w_t) - \hat{f}_t(w)$$

Applying Corollary 2, we have:

$$\sum_{t=1}^{T} \hat{f}_t(w_t) - \hat{f}_t(w) \leq \sum_{t=1}^{T} \nabla \hat{f}_t(w_t)^\top (w_t - w_{t+1}) - \mathrm{B}_{\hat{g}_t}(w_{t+1}\|w_t)$$

Here $\hat{g}_t(w) = \sum_{s=1}^{t} \hat{f}_s(w)$. Observe that $\nabla \hat{f}_t(w_t) = \nabla f_t(w_t) = \nabla_t$. Moreover, we have:

$$\begin{aligned}
\mathrm{B}_{\hat{g}_t}(w_{t+1} \| w_t) &= \frac{1}{2}(w_{t+1} - w_t)^\top \left(\sum_{s=1}^{t} \beta_s(w_s) \nabla f_s(w_s) \nabla f_s(w_s)^\top \right) (w_{t+1} - w_t) \\
&= \frac{1}{2}(w_{t+1} - w_t)^\top \left(\sum_{s=1}^{t} \beta_s \nabla_s \nabla_s^\top \right) (w_{t+1} - w_t) \\
&= \frac{1}{2}(w_{t+1} - w_t)^\top \left(\sum_{s=1}^{t} \beta_s \nabla_s \nabla_s^\top + \epsilon I \right) (w_{t+1} - w_t) - \frac{1}{2}\epsilon \|w_{t+1} - w_t\|_2^2
\end{aligned}$$

Thus, we have:

$$\begin{aligned}
\mathcal{R}_T(w) \leq \sum_{t=1}^{T} \nabla_t^\top (w_t - w_{t+1}) - \frac{1}{2}(w_{t+1} - w_t)^\top \left(\sum_{s=1}^{t} \beta_s \nabla_s \nabla_s^\top + \epsilon I \right) (w_{t+1} - w_t) \\
+ \frac{1}{2}\epsilon \sum_{t=1}^{T} \|w_{t+1} - w_t\|_2^2
\end{aligned}$$

Using the fact that $a^\top x - \frac{1}{2} x^\top A x \leq \frac{1}{2} a^\top A^{-1} a$ when A is a positive definite matrix.

$$\begin{aligned}
&\leq \frac{1}{2} \sum_{t=1}^{T} \nabla_t^\top \left(\sum_{s=1}^{t} \beta_s \nabla_s \nabla_s^\top + \epsilon I \right)^{-1} \nabla_t + \frac{\epsilon}{2} \mathcal{M}_T^2 \\
&\leq \frac{1}{2 \min_t \beta_t} \sum_{t=1}^{T} \beta_t \nabla_t^\top \left(\sum_{s=1}^{t} \beta_s \nabla_s \nabla_s^\top + \epsilon I \right)^{-1} \nabla_t + \frac{\epsilon}{2} \mathcal{M}_T^2
\end{aligned}$$

Using Lemma 14 with $x_t = \sqrt{\beta_t} \nabla_t$

$$\leq \frac{n}{2 \min_t \beta_t} \log \left(1 + \frac{\sum_{t=1}^{T} \beta_t \|\nabla_t\|_2^2}{n\epsilon} \right) + \frac{\epsilon}{2} \mathcal{M}_T^2$$

Picking $\epsilon = \frac{n}{\mathcal{M}_T^2 \min_t \beta_t}$

$$\leq \frac{n}{2 \min_t \beta_t} \log \left(1 + \frac{\mathcal{G}_T^2 \mathcal{M}_T^2 \min_t \beta_t}{n^2} \right) + \frac{n}{2 \min_t \beta_t}$$

This completes the proof. ■

Remark Note that ϵ is not a parameter of the algorithm. It only appears in the analysis. It can be chosen to be any positive constant dependent on problem parameters. So, we can state an alternate regret bound for FTAL:

$$\mathcal{R}_T(w) \leq \inf_{\epsilon>0} \left(\frac{n}{2\min_t \beta_t} \log\left(1 + \frac{\sum_{t=1}^T \beta_t \|\nabla_t\|_2^2}{n\epsilon}\right) + \frac{\epsilon}{2}\mathcal{M}_T^2 \right) \tag{2}$$

Appendix E. Results on FTRL

We show a general regret bound for FTRL similar to Corollary 2.

Corollary 19 *Let $F(w)$ be a strongly convex regularizer. Consider the iterates of the following FTRL:*

$$w_t = \arg\min_{w\in\mathcal{D}} F(w) + \sum_{s=1}^{t-1} f_s(w)$$

For any $w \in \mathcal{D}$ and any sequence of convex functions $f_1, \ldots, f_T$, we have:

$$\mathcal{R}_T(w) \leq F(w) - F(w_1) + \sum_{t=1}^T \nabla f_t(w_t)^\top (w_t - w_{t+1}) - B_{g_t+F}(w_{t+1}|w_t)$$

Proof Since FTRL can be thought of as an FTL on the sequence of functions, $f_0, f_1, \ldots, f_T$ where $f_0 = F$, we apply Corollary 2:

$$\begin{aligned}
&\sum_{t=0}^T f_t(w_t) - f_t(w) \leq \sum_{t=0}^T \left(f_t(w_t) - f_t(w_{t+1}) - \sum_{s=0}^{t-1} \mathrm{B}_{f_s}(w_{t+1}\|w_t) \right) \\
\implies F(w_0) - F(w) + &\sum_{t=1}^T f_t(w_t) - f_t(w) \leq F(w_0) - F(w_1) \\
&\qquad + \sum_{t=1}^T \left(f_t(w_t) - f_t(w_{t+1}) - \mathrm{B}_{g_{t-1}+F}(w_{t+1}\|w_t) \right) \\
\implies &\sum_{t=1}^T f_t(w_t) - f_t(w) \leq F(w) - F(w_1) \\
&\qquad + \sum_{t=1}^T \nabla f_t(w_t)^\top (w_t - w_{t+1}) - \mathrm{B}_{g_t+F}(w_{t+1}|w_t)
\end{aligned}$$

This completes the proof. ∎

We can get a regret bound for FTRL similar to Theorem 4 where we add a regularizer to FTAL.

Theorem 20 *Assume cost function $f_1, \ldots, f_T$ satisfy Assumption 3.Consider the iterates of the following FTRL:*

$$w_t = \arg\min_{w\in\mathcal{D}} \frac{\epsilon}{2}\|w\|_2^2 + \sum_{s=1}^{t-1} \hat{f}_s(w)$$

Here $\hat{f}_t(w) = f_t(w_t) + \nabla f_t(w_t)^\top (w - w_t) + \frac{\beta_t(w_t)}{2}(\nabla f_t(w_t)^\top (w - w_t))^2$. *For any* $w \in \mathcal{D}$, *the iterates satisfy:*

$$\mathcal{R}_T(w) \leq \frac{\epsilon}{2}\|w\|_2^2 + \frac{n}{2\min_t \beta_t}\log\left(1 + \frac{\sum_{t=1}^T \beta_t \|\nabla_t\|_2^2}{n\epsilon}\right)$$

Proof As $\hat{f}_t(w) \leq f_t(w)$ for all $w \in \mathcal{D}$ and $\hat{f}_t(w_t) = f_t(w_t)$, we have:

$$\mathcal{R}_T(w) = \sum_{t=1}^T f_t(w_t) - f_t(w) \leq \sum_{t=1}^T \hat{f}_t(w_t) - \hat{f}_t(w)$$

Applying Corollary 19 with $F(w) = \frac{\epsilon}{2}\|w\|_2^2$, we have:

$$\sum_{t=1}^T \hat{f}_t(w_t) - \hat{f}_t(w) \leq F(w) - F(w_1) + \sum_{t=1}^T \nabla \hat{f}_t(w_t)^\top (w_t - w_{t+1}) - \mathrm{B}_{\hat{g}_t + F}(w_{t+1}\|w_t)$$

Here $\hat{g}_t(w) = \sum_{s=1}^t \hat{f}_s(w)$. Observe that $\nabla \hat{f}_t(w_t) = \nabla f_t(w_t) = \nabla_t$. Moreover, we have:

$$\begin{aligned}
\mathrm{B}_{\hat{g}_t+F}(w_{t+1}\|w_t) &= \frac{1}{2}(w_{t+1} - w_t)^\top \left(\sum_{s=1}^t \beta_s(w_s)\nabla f_s(w_s)\nabla f_s(w_s)^\top + \epsilon I\right)(w_{t+1} - w_t) \\
&= \frac{1}{2}(w_{t+1} - w_t)^\top \left(\sum_{s=1}^t \beta_s \nabla_s \nabla_s^\top + \epsilon I\right)(w_{t+1} - w_t)
\end{aligned}$$

Thus, we have:

$$\mathcal{R}_T(w) \leq \frac{\epsilon}{2}\|w\|_2^2 + \sum_{t=1}^T \nabla_t^\top (w_t - w_{t+1}) - \frac{1}{2}(w_{t+1} - w_t)^\top \left(\sum_{s=1}^t \beta_s \nabla_s \nabla_s^\top + \epsilon I\right)(w_{t+1} - w_t)$$

Using the fact that $a^\top x - \frac{1}{2}x^\top A x \leq \frac{1}{2}a^\top A^{-1} a$ when A is a positive definite matrix.

$$\begin{aligned}
&\leq \frac{\epsilon}{2}\|w\|_2^2 + \frac{1}{2}\sum_{t=1}^T \nabla_t^\top \left(\sum_{s=1}^t \beta_s \nabla_s \nabla_s^\top + \epsilon I\right)^{-1} \nabla_t \\
&\leq \frac{\epsilon}{2}\|w\|_2^2 + \frac{1}{2\min_t \beta_t}\sum_{t=1}^T \beta_t \nabla_t^\top \left(\sum_{s=1}^t \beta_s \nabla_s \nabla_s^\top + \epsilon I\right)^{-1} \nabla_t
\end{aligned}$$

Using Lemma 14 with $x_t = \sqrt{\beta_t}\nabla_t$

$$\leq \frac{\epsilon}{2}\|w\|_2^2 + \frac{n}{2\min_t \beta_t}\log\left(1 + \frac{\sum_{t=1}^T \beta_t \|\nabla_t\|_2^2}{n\epsilon}\right)$$

This completes the proof. ■

Appendix F. Proofs in Section 4

Lemma 7 *Under Assumption 5, for all $w, w' \in \Delta_n$, we have for $f_t(w) = -\ln(r_t^\top w)$:*

$$f_t(w) \geq f_t(w') + \nabla f_t(w')^\top (w - w') + \frac{(r_t^\top w')}{2} (\nabla f_t(w')^\top (w - w'))^2$$

Proof Under Assumption 5, we have $0 < r_t^\top w \leq 1$ for all $w \in \Delta_n$. So, we can apply Lemma 6 with $y = r_t^\top w$ and $x = r_t^\top w'$. Note that $\nabla f_t(w) = -r_t/(r_t^\top w)$. This gives:

$$\begin{aligned}
&\frac{r_t^\top w}{r_t^\top w'} - 1 - \log\left(\frac{r_t^\top w}{r_t^\top w'}\right) \geq \frac{1}{2}\frac{(r_t^\top w' - r_t^\top w)^2}{r_t^\top w'} \\
\implies& -\log(r_t^\top w) \geq -\log(r_t^\top w') - \frac{r_t^\top (w - w')}{r_t^\top w'} + \frac{1}{2}\frac{(r_t^\top (w - w'))^2}{r_t^\top w'} \\
\implies& f_t(w) \geq f_t(w') + \nabla f_t(w')^\top (w - w') + \frac{(r_t^\top w')}{2} (\nabla f_t(w')^\top (w - w'))^2
\end{aligned}$$

Completing the proof. ∎

Theorem 8 *For any $w \in \Delta_n$ and any sequence of returns $r_1, \ldots, r_T$ such that Assumption 5 holds, the iterates of FTAL (Algorithm 2), with $\nabla_t = -r_t/(r_t^\top w_t)$ and $\beta_t = r_t^\top w_t$ satisfy the inequality:*

$$\mathcal{R}_T(w) \leq \frac{n}{2\hat{r}}\left(\log\left(\frac{2T^2}{n} + 1\right) + 1\right)$$

Proof Applying Theorem 4 with $\nabla_t = -r_t/(r_t^\top w_t)$ and $\beta_t = r_t^\top w_t$, we have:

$$\begin{aligned}
\mathcal{M}_T^2 &\leq 2T \\
\mathcal{G}_T^2 &= \sum_{t=1}^{T} \frac{\|r_t\|_2^2}{r_t^\top w_t} \leq \frac{nT}{\min_t \beta_t} \implies \mathcal{G}_T^2 \min_t \beta_t \leq nT \\
\min_t \beta_t &\geq \hat{r}
\end{aligned}$$

Substituting these quantities in the result of Theorem 4, we have the bound:

$$\mathcal{R}_T(w) \leq \frac{n}{2\hat{r}}\left(\log\left(\frac{2T^2}{n} + 1\right) + 1\right)$$

∎

Theorem 9 *For any $w \in \Delta_n$ and any sequence of returns $r_1, \ldots, r_T$ such that Assumption 5 holds, the iterates of Equation 1 satisfy $\mathcal{R}_T(w) = O\left(\frac{n}{\hat{r}^2}\log(Q_T + n)\right)$*

Proof The iterates are computed as:

$$w_t = \arg\min_{w \in \Delta_n} \frac{1}{2}\|w\|_2^2 + \sum_{s=1}^{t-1} \left(f_s(w_s) - \frac{r_s^\top (w - w_s)}{(r_s^\top w_s)^\top} + \frac{(r_s^\top (w - w_s))^2}{2(r_s^\top w_s)} \right)$$

We can apply Theorem 16 with $v_t = r_t$. The function $h_t(x) = f_t(w_t) - \frac{x - r_s^\top w_s}{(r_s^\top w_s)^\top} + \frac{(x - r_s^\top w_s)^2}{2(r_s^\top w_s)}$. This gives $h_t'(r_t^\top w_t) = \frac{-1}{r_t^\top w_t} \in [\frac{-1}{\hat{r}}, 0]$ and $h_t''(r_t^\top w) = \frac{-1}{r_t^\top w_t} \geq 1$. Thus we have $\|r_t\| \leq \sqrt{n} = R$, $D = 1$, $a = \frac{1}{\hat{r}}$ and $b = 1$. So, we have the regret bound:

$$\mathcal{R}_T(w) \leq O\left(\frac{n}{\hat{r}^2} \log(1 + Q_T + n) + \frac{\sqrt{n}}{\hat{r}} \log(1 + Q_T/n) + 1 \right)$$

Here $Q_T = \min_\mu \sum_{t=1}^T \|r_t - \mu\| = \sum_{t=1}^T \|r_t - \bar{r}_T\|$, where $\bar{r}_T = \frac{1}{T} \sum_{t=1}^T r_t$. ■

Lemma 10 *Under Assumption 5, the cost functions $f_t(w) = -\log(r_t^\top w)$ are $n/\hat{r}^2$-smooth on Δ_n.*

Proof For any $w, w' \in \Delta_n$, we have:

$$\|\nabla f_t(w) - \nabla f_t(w')\|_2 = \left\| \frac{r_t}{r_t^\top w} - \frac{r_t}{r_t^\top w'} \right\|_2 = \frac{\|r_t\|_2^2 \|w - w'\|_2}{(r_t^\top w)(r_t^\top w')} \leq \frac{n}{\hat{r}^2} \|w - w'\|$$

Thus $f_t(w)$ is $n/\hat{r}^2$ smooth on Δ_n. ■

Theorem 11 *For any $w \in \Delta_n$ and any sequence of returns $r_1, \ldots, r_T$ such that Assumption 5 holds, the iterates of Equation 1 satisfy $\mathcal{R}_T(w) = O\left(\frac{n}{\hat{r}} \log\left(\frac{L_T^\star}{\hat{r}^2} + \frac{n}{\hat{r}^3} \right) \right)$, where $L_T^\star = \min_{w \in \Delta_n} \sum_{t=1}^T f_t(w)$*

Proof The iterates are computed as:

$$w_t = \arg\min_{w \in \Delta_n} \frac{1}{2}\|w\|_2^2 + \sum_{s=1}^{t-1} \left(f_s(w_s) - \frac{r_s^\top (w - w_s)}{(r_s^\top w_s)^\top} + \frac{(r_s^\top (w - w_s))^2}{2(r_s^\top w_s)} \right)$$

We can apply Theorem 20 with $\epsilon = 1$, $\beta_t = r_t^\top w_t$:

$$\mathcal{R}_T(w) \leq \frac{1}{2}\|w\|_2^2 + \frac{n}{2 \min_t \beta_t} \log\left(1 + \frac{\sum_{t=1}^T (r_t^\top w_t) \|\nabla_t\|_2^2}{n} \right) \tag{3}$$

Note that $\|w\| \leq 1$, $\min_t \beta_t \geq \hat{r}$ and $(r_t^\top w_t) \leq 1$:

$$\leq \frac{1}{2} + \frac{n}{2\hat{r}} \log\left(1 + \frac{\sum_{t=1}^T \|\nabla_t\|_2^2}{n} \right)$$

From Lemma 10, we know f_t is $n/\hat{r}^2$ smooth on Δ_n. Under Assumption 5, f_t are non-negative. So we apply Lemma 17, which gives us the bound:

$$\|\nabla f_t(w_t)\|_2^2 \leq \frac{4n}{\hat{r}^2} f_t(w_t)$$

So, we have:

$$\sum_{t=1}^{T} f_t(w_t) - \sum_{t=1}^{T} f_t(w) \leq \frac{1}{2} + \frac{n}{2\hat{r}} \log\left(1 + \frac{4}{\hat{r}^2}\sum_{t=1}^{T} f_t(w_t)\right)$$

Now, we apply Lemma 18 with $a = n/2\hat{r}$, $b = 4/\hat{r}^2$, $c = 1$, $d = 1/2 + \sum_{t=1}^{T} f_t(w)$ and $x = \sum_{t=1}^{T} f_t(w_t)$. So, we have for all $w \in \Delta_n$:

$$\sum_{t=1}^{T} f_t(w_t) - \sum_{t=1}^{T} f_t(w) \leq \frac{1}{2} + \frac{n}{2\hat{r}} \log\left(2\left(\frac{2n}{\hat{r}^3}\log\left(\frac{4n}{e\hat{r}^3}\right) + \frac{2}{\hat{r}^2} + \frac{4}{\hat{r}^2}\sum_{t=1}^{T} f_t(w) + 1\right)\right)$$

Let $w^\star \in \arg\min_{w\in\Delta_n} \sum_{t=1}^{T} f_t(w)$ and $L_T^\star = \min_{w\in\Delta_n} \sum_{t=1}^{T} f_t(w) = \sum_{t=1}^{T} f_t(w^\star)$. Using $w^\star$ in the above bound, we have:

$$\mathcal{R}_T(w) \leq \mathcal{R}(w^\star) \leq \frac{1}{2} + \frac{n}{2\hat{r}} \log\left(2\left(\frac{2n}{\hat{r}^3}\log\left(\frac{4n}{e\hat{r}^3}\right) + \frac{2}{\hat{r}^2} + \frac{4}{\hat{r}^2}L_T^\star + 1\right)\right)$$

■

Remark From Equation (3), we can obtain the regret bound $\mathcal{R}_T(w) \leq \frac{1}{2} + \frac{n}{2\hat{r}}\log\left(1 + \frac{T}{\hat{r}}\right)$. We can obtain a similar regret bound for FTAL using Equation (2) and plugging in $\epsilon = (2T)^{-1}$. This gives: $\mathcal{R}_T(w) \leq \frac{1}{2} + \frac{n}{2\hat{r}}\log\left(1 + \frac{T^2}{\hat{r}}\right)$.

Appendix G. Proof of Theorem 12

Let $X_t = \sum_{s=1}^{t} x_s x_s$ and $Y_t = \sum_{s=1}^{t} y_s x_s$.

Theorem 12 *For any $w \in \mathbb{R}^n$ and any sequence of pairs $(x_1, y_1), \ldots, (x_T, y_T)$, the iterates of OFTL (Algorithm 1) with $f_t(w) = \frac{1}{2}(x_t^\top w - y_t)^2$, $m_t(w) = \frac{1}{2}(x_t^\top w)^2$ and $\mathcal{D} = \mathbb{R}^n$, satisfy the regret the inequality $\sum_{t=1}^{T} \frac{1}{2}(x_t^\top w_t - y_t)^2 - \frac{1}{2}(x_t^\top w - y_t)^2 \leq \frac{(\max_t y_t^2)n}{2}\left(\log\left(\frac{\mathcal{X}_T^2 \mathcal{M}'^2_T}{n^2(\max_t y_t^2)} + 1\right) + 1\right)$, Where $\mathcal{X}_T^2 = \sum_{t=1}^{T}\|x_t\|_2^2$ and $\mathcal{M}'^2_T = \sum_{t=1}^{T}\|w_t - w'_{t+1}\|_2^2$. This implies the regret bound:*

$$\leq \frac{(\max_t y_t^2)n}{2}\left(\log\left(\frac{\sum_{t=1}^{T}\|x_t\|_2^2}{n^2}\left(\sum_{t=1}^{T}\frac{\|x_t\|_2^2}{\sigma_{\min}(X_t)^2}\right) + 1\right) + 1\right)$$

Proof We use the regret bound from Theorem 1 and ignore the last term as it is non-negative:

$$\begin{aligned}
\mathcal{R}_T(w) &\leq \sum_{t=1}^{T}(\nabla f_t(w_t) - \nabla m_t(w_t))^\top (w_t - w'_{t+1}) - \mathrm{B}_{g_t}(w'_{t+1}\|w_t) - \mathrm{B}_{g_{t-1}}(w_t\|w'_t) \\
&\leq \sum_{t=1}^{T}(\nabla f_t(w_t) - \nabla m_t(w_t))^\top (w_t - w'_{t+1}) - \mathrm{B}_{g_t}(w'_{t+1}\|w_t)
\end{aligned}$$

Since $\nabla f_t(w_t) = (x_t^\top w_t - y_t)x_t$ and $\nabla m_t(w_t) = (x_t^\top w_t)x_t$, the term $(\nabla f_t(w_t) - \nabla m_t(w_t)) = -y_t x_t$. The Bregmen term is:

$$\begin{aligned}
\mathrm{B}_{g_t}(w'_{t+1} \| w_t) &= \frac{1}{2}(w_t - w'_{t+1})^\top X_t (w_t - w'_{t+1}) \\
&= \frac{1}{2}(w_t - w'_{t+1})^\top (X_t + \epsilon I)(w_t - w'_{t+1}) - \frac{1}{2}\|w_t - w'_{t+1}\|_2^2
\end{aligned}$$

Thus, we have:

$$\begin{aligned}
\mathcal{R}_T(w) \leq & \sum_{t=1}^{T} (-y_t x_t)^\top (w_t - w'_{t+1}) - \frac{1}{2}(w_t - w'_{t+1})^\top (X_t + \epsilon I)(w_t - w'_{t+1}) \\
& + \frac{1}{2}\epsilon \sum_{t=1}^{T} \|w_t - w'_{t+1}\|_2^2
\end{aligned}$$

Using the fact that $a^\top x - \frac{1}{2}x^\top A x \leq \frac{1}{2}a^\top A^{-1} a$ when A is a positive definite matrix.

$$\begin{aligned}
& \leq \frac{1}{2}\sum_{t=1}^{T} y_t^2 x_t^\top (X_t + \epsilon I)^{-1} x_t + \frac{1}{2}\epsilon {\mathcal{M}'_T}^2 \\
& \leq \frac{(\max_t y_t^2)}{2} \sum_{t=1}^{T} x_t (X_t + \epsilon I)^{-1} x_t + \frac{1}{2}\epsilon {\mathcal{M}'_T}^2
\end{aligned}$$

Using Lemma 14

$$\leq \frac{(\max_t y_t^2) n}{2} \log\left(1 + \frac{\sum_{t=1}^{T}\|x_t\|_2^2}{n\epsilon}\right) + \frac{1}{2}\epsilon {\mathcal{M}'_T}^2$$

Thus, we have the regret bound:

$$\mathcal{R}_T(w) \leq \inf_{\epsilon > 0}\left(\frac{(\max_t y_t^2) n}{2} \log\left(1 + \frac{\sum_{t=1}^{T}\|x_t\|_2^2}{n\epsilon}\right) + \frac{1}{2}\epsilon {\mathcal{M}'_T}^2\right)$$

Picking $\epsilon = \frac{n(\max_t y_t^2)}{{\mathcal{M}'_T}^2}$

$$\leq \frac{(\max_t y_t^2) n}{2} \log\left(1 + \frac{\mathcal{X}_T^2 {\mathcal{M}'_T}^2}{n^2 (\max_t y_t^2)}\right) + \frac{(\max_t y_t^2) n}{2}$$

We can bound ${\mathcal{M}'_T}^2$ as:

$$\begin{aligned}
{\mathcal{M}'_T}^2 &= \sum_{t=1}^{T}\|w'_{t+1} - w_t\|_2^2 = \sum_{t=1}^{T}\|X_t^+ Y_t - X_t^+ Y_{t-1}\|_2^2 = \sum_{t=1}^{T} y_t^2 \|X_t^+ x_t\|_2^2 \\
&\leq (\max_t y_t^2) \sum_{t=1}^{T} \frac{\|x_t\|_2^2}{\sigma_{\min}(X_t)^2}
\end{aligned}$$

So, we have the regret bound:

$$\mathcal{R}_T(w) \leq \frac{(\max_t y_t^2)n}{2}\left(\log\left(\frac{\sum_{t=1}^T \|x_t\|_2^2}{n^2}\left(\sum_{t=1}^T \frac{\|x_t\|_2^2}{\sigma_{\min}(X_t)^2}\right)+1\right)+1\right)$$

■