
IMPLICIT REGULARIZATION OF BREGMAN PROXIMAL POINT ALGORITHM AND MIRROR DESCENT ON SEPARABLE DATA

Anonymous authors

Paper under double-blind review

ABSTRACT

Bregman proximal point algorithm (BPPA), as one of the centerpieces in the optimization toolbox, has been witnessing emerging applications. With simple and easy to implement update rule, the algorithm bears several compelling intuitions for empirical successes, yet rigorous justifications are still largely unexplored. We study the computational properties of BPPA through classification tasks with separable data, and demonstrate provable algorithmic regularization effects associated with BPPA. We show that BPPA attains non-trivial margin, which closely depends on the condition number of the distance generating function inducing the Bregman divergence. We further demonstrate that the dependence on the condition number is tight for a class of problems, thus showing the importance of divergence in affecting the quality of the obtained solutions. In addition, we extend our findings to mirror descent (MD), for which we establish similar connections between the margin and Bregman divergence. We demonstrate through a concrete example, and show BPPA/MD converges in direction to the maximal margin solution with respect to the Mahalanobis distance. Our theoretical findings are among the first to demonstrate the benign learning properties BPPA/MD, and also provide strong corroborations for a careful choice of divergence in the algorithmic design.

1 INTRODUCTION

The role of optimization algorithms has become arguably one of the most critical factors in the empirical successes of training deep models. As the go-to choice for modern machine learning, first-order algorithms, including (stochastic) gradient descent and their adaptive counterparts (Kingma and Ba, 2014; Duchi et al., 2011), have received tremendous attention, with detailed investigations dedicated to understanding the effect of batch size (Goyal et al., 2017; Smith et al., 2018; Keskar et al., 2016), learning rate (Li et al., 2019; He et al., 2019; Lewkowycz et al., 2020), and momentum (Sutskever et al., 2013; Smith, 2018) across a broad spectrum of applications.

Meanwhile, Bregman proximal point algorithm (Eckstein, 1993; Kiwiel, 1997) has been drawing substantial interests. The resounding successes of this classical algorithm are particularly evident for applications including knowledge distillation (Furlanello et al., 2018), mean-teacher learning paradigm (Tarvainen and Valpola, 2017), few-shot learning (Zhou et al., 2019), policy optimization (Green et al., 2019), and fine-tuning pre-trained models (Jiang et al., 2020), yielding competitive performance compared to its first-order counterparts. In the general form, Bregman proximal point algorithm updates parameters by minimizing a loss $L(\cdot)$, while regularizing the weighted distance to the previous iterate measured by some divergence function $\mathcal{D}(\cdot, \cdot)$,

$$\theta_{t+1} = \operatorname{argmin}_{\theta} L(\theta) + 1/(2\eta_t)\mathcal{D}(\theta, \theta_t). \quad (1.1)$$

Popular choices of divergence function used in practice include the squared ℓ_2 -norm distance $D_{LS}(\theta, \theta_t) = \mathbb{E}_{\mathcal{D}} \|f_{\theta}(x) - f_{\theta_t}(x)\|_2^2$ (Tarvainen and Valpola, 2017), and Kullback-Leibler based divergence $D_{KL}(\theta, \theta_t) = \mathbb{E}_{\mathcal{D}} \operatorname{KL}(f_{\theta'}(x) \| f_{\theta}(x))$ (Furlanello et al., 2018), where \mathcal{D} denotes the data distribution. Such a simple update is of great practical purposes, as it is easy to describe, and admits simple implementation by adopting suitable off-the-shelf black-box optimization algorithms (Solodov and Svaiter, 2000; Monteiro and Svaiter, 2010; Zaslavski, 2010). The updating form also suggests plausible intuitions for its empirical successes, including iteratively constraining the search space, alleviating aggressive updates, and preventing catastrophic forgetting (Schulman et al., 2015; Li and Hoiem, 2017). However, none of the intuitions have been rigorously justified, and

theoretical understandings for the empirical successes of Bregman proximal point algorithm remains underexplored.

A first, and a natural question is whether Bregman proximal point algorithm benefits from the same kind of mechanism that (stochastic) gradient descent (GD/SGD) enjoys for having the generalization properties. In particular, in many important applications, GD/SGD is widely believed as the “the algorithm that finds the right kind of solutions” for problems with non-unique solutions. Such a claim is supported with numerous provable examples: GD/SGD converges to the minimum-norm solution of under-determined linear systems (Gunasekar et al., 2018), converges to the max-margin solution for separable data (Soudry et al., 2018; Nacson et al., 2019), aligns layers of deep linear networks (Ji and Telgarsky, 2018), and converges to a generalizable solution for nonlinear networks (Brutzkus et al., 2017; Allen-Zhu et al., 2018) in the presence of infinitely many overfitting solutions. Given the emerging successes of Bregman proximal point algorithm, and the aforementioned evidences on its first-order counterparts (e.g. GD/SGD) finding generalizable solutions, one would naturally wonder

Does Bregman proximal point algorithm converge to a solution with favorable qualities?

Another important question with great practical implications for Bregman proximal point algorithm is how the divergence measure $\mathcal{D}(\cdot, \cdot)$ affects the solution. Instead of directly applying the Euclidean distance based divergence, it is widely observed that the successful application of Bregman proximal point algorithm is contingent on the careful design of divergence measure, based on the task at hand (Li and Hoiem, 2017; Hinton et al., 2015). Take the example of fine-tuning language model, the symmetrized Kullback-Leibler based divergence evaluated on the predictions of the updated model (i.e., θ_{t+1}) and previous model (i.e., θ_t) yields the state-of-the-art result (Jiang et al., 2020). Identifying the underlying mechanism for the success or failure of a given divergence choice is not only of theoretical interest, but also can significantly reduce human effort in searching/designing the suitable divergence for a given task. As an important addition, one may also ask whether the impact of divergence on the Bregman proximal point algorithm find natural counterparts in commonly adopted first-order algorithms (e.g. mirror descent, (Nemirovski and Yudin, 1983)). In such cases, better task-dependent algorithmic designs could be proposed in conjunction with the suitable divergence. To this end, we raise our second question.

How does divergence affect the qualities of the solution obtained by Bregman proximal point algorithm (and other first-order algorithms)?

In this paper, we initiate our study to address our previously proposed questions. We focus on a non-trivial example of an under-determined system – training linear classifiers using separable data. In particular, for exponential tail losses (e.g., exponential loss), the empirical loss function has infimum zero that is asymptotically attainable at infinity along infinitely many directions. The natural candidate for measuring the quality of the obtained classifier is its margin, i.e., the minimum distance between the samples and the decision hyperplane. For such a problem, we summarize our theoretical findings below as concrete answers to the previous questions.

- We show that Bregman proximal point algorithm (BPPA) obtains a solution with non-trivial margin lower-bound. As a concrete demonstration, we tailor our main theorem for Mahalanobis distance, and show that BPPA converges in direction to the maximal margin solution. We provide non-asymptotic analyses of the margin and empirical loss for constant stepsize BPPA, and propose a more aggressive stepsize rule for a provable exponential speed-up.
- We establish a dependence of such a margin lower-bound on the condition number of the distance generating function for defining the divergence. In addition, we provide a class of problems where the margin lower-bound is tight, demonstrating that the Bregman divergence is crucial in affecting the quality of the obtained solution.
- We extend our findings to first-order algorithms. Specifically, we show that mirror descent (MD) enjoys the same previously mentioned margin properties. We also provide non-asymptotic convergence analyses of the margin and empirical loss for constant stepsize MD, and its exponential speed-up using a varying stepsize scheme. Our findings for MD strictly complement prior works on under-determined regression problems (Gunasekar et al., 2018; Azizan and Hassibi, 2019).

Notations. We denote $[n] := \{1, \dots, n\}$; $\text{sgn}(z) = 1$ if $z \geq 0$ and -1 elsewhere. We use w.r.t in short for “with respect to”. For any $\|\cdot\|$ in Euclidean space \mathbb{R}^d , we use $\|\cdot\|_* = \max_{\|y\| \leq 1} \langle \cdot, y \rangle$ to denote its dual norm. Note that we have $(\|\cdot\|_*)_* = \|\cdot\|$.

2 PROBLEM SETUP

We study the binary classification on linearly separable data. Specifically, the dataset is $\mathcal{S} = \{(x_i, y_i)\}_{i=1}^n \subset \mathbb{R}^d \times \{+1, -1\}$, where x_i is the feature vector, and y_i is the label. In addition, there exists a linear classifier $u \in \mathbb{R}^d$ such that $y_i \langle u, x_i \rangle > 0$ for all $i \in [n]$. That is, the decision rule $f_u(\cdot) = \text{sgn}(\langle u, \cdot \rangle)$ achieves the perfect accuracy on the dataset, with $y_i = f_u(x_i)$ for all $i \in [n]$.

For each linear classifier $f_u(\cdot)$ with perfect accuracy, we define its $\|\cdot\|_*$ -norm margin as the minimum distance in $\|\cdot\|_*$ -norm from the feature vectors to the decision boundary $\mathcal{H}_u = \{x : \langle x, u \rangle = 0\}$. It is well known that the $\|\cdot\|_*$ -norm margin, denoted as γ_u , only depends on the direction of the classifier and satisfies $\gamma_u = \min_{i \in [n]} \left\langle x_i y_i, \frac{u}{\|u\|} \right\rangle$, where $\|\cdot\|$ is the dual norm of $\|\cdot\|_*$. The $\|\cdot\|_*$ -norm margin measures how well the data is separated by decision rule $f_u(\cdot)$, measured in $\|\cdot\|_*$ -norm, and is an important measure on the generalizability and robustness of the decision rule. Given a norm $\|\cdot\|_*$ on \mathbb{R}^d , we define the optimal linear classifier with the maximum $\|\cdot\|_*$ -margin below.

Definition 2.1 (Maximum $\|\cdot\|_*$ -norm Margin Classifier). *Given a linearly separable dataset $\{(x_i, y_i)\}_{i \in [n]}$, we define the maximum $\|\cdot\|_*$ -norm margin classifier $u_{\|\cdot\|_*}$, and its associated maximum $\|\cdot\|_*$ -norm margin $\gamma_{\|\cdot\|_*}$ as*

$$u_{\|\cdot\|_*} = \operatorname{argmax}_{\|u\| \leq 1} \min_{i \in [n]} \langle u, y_i x_i \rangle, \quad \gamma_{\|\cdot\|_*} = \max_{\|u\| \leq 1} \min_{i \in [n]} \langle u, y_i x_i \rangle.$$

For a separable dataset, we consider finding the classifier by minimizing the empirical loss

$$L_{\mathcal{S}}(\theta) = \frac{1}{n} \sum_{i=1}^n \ell(\langle \theta, y_i x_i \rangle). \quad (2.1)$$

Here we focus on the exponential loss $\ell(x) = \exp(-x)$, and our analyses can be readily extended to other losses with tight exponential tail (e.g., logistic loss).

Observation. One can readily verify that with a separable dataset \mathcal{S} , the empirical loss has infimum 0 but possesses no finite solution that attains the infimum. Thus any optimization algorithm minimizing the loss $L_{\mathcal{S}}(\cdot)$ will observe the explosion on the norm of iterate.

It has been shown that various optimization algorithms, including (stochastic) gradient descent and steepest descent, converge in direction to the maximum margin classifier in different norms (Soudry et al., 2018; Nacson et al., 2019; Gunasekar et al., 2018; Ji and Telgarsky, 2019; 2021). Connections between gradient descent and the regularization path of homotopy method have also been established (Ji et al., 2020). A striking feature behind such phenomena is that there is no explicit regularization in the loss function, and such effects have been termed as the *implicit (algorithmic) regularization*.

Up to date, most of the implicit regularization effects are attributed to (stochastic) gradient descent, given their prevalence in applications. However, as Bregman proximal point algorithm (BPPA) becomes increasingly popular in various domains, there exists considerable lack of understanding on the computational properties of BPPA. In addition, practitioners often find the choice of divergence function crucially important for the performance of BPPA (Jiang et al., 2020; Furlanello et al., 2018). This empirical evidence thus calls for a detailed characterization on the connection between computational properties and the divergence function of BPPA.

In what follows, we study the BPPA for solving problem (2.1) in detail. The BPPA (Algorithm 1) is an adaptation of the vanilla proximal point algorithm (Rockafellar, 1976a;b) to non-euclidean geometry, by using Bregman divergence as the divergence measure in (1.1). Specifically, given a distance generating function $w(\cdot)$ that is convex and differentiable, we define the Bregman divergence $D_w(\cdot, \cdot)$ associated with $w(\cdot)$ as $D_w(\theta, \theta') = w(\theta) - w(\theta') - \langle \nabla w(\theta'), \theta - \theta' \rangle$. Throughout our discussions, we only impose the following mild assumption on Bregman divergence function $D_w(\cdot, \cdot)$.

Algorithm 1 Bregman Proximal Point Algorithm (BPPA)

Input: Distance generating function $w(\cdot)$, stepsizes $\{\eta^t\}_{t \geq 0}$, samples $\{x_i, y_i\}_{i=1}^n$.

Initialize: $\theta^0 \leftarrow 0$.

for $t = 0, \dots$ **do**

$$\text{Update } \theta_{t+1} = \operatorname{argmin}_{\theta} L_{\mathcal{S}}(\theta) + \frac{1}{2\eta_t} D_w(\theta, \theta_t). \quad (2.2)$$

end for

Assumption 1. We assume that the distance generating function of Bregman divergence $D_w(\cdot, \cdot)$ is L_w -smooth and μ_w -strongly convex w.r.t. $\|\cdot\|$ -norm. That is,

$$\frac{\mu_w}{2} \|\theta - \theta'\|^2 \leq w(\theta) - w(\theta') - \langle \nabla w(\theta'), \theta - \theta' \rangle \leq \frac{L_w}{2} \|\theta - \theta'\|^2.$$

3 ALGORITHMIC REGULARIZATION OF BPPA

We show BPPA achieves a $\|\cdot\|_*$ -norm margin that is at least $\sqrt{\mu_w/L_w}$ -fraction of the maximal one.

Theorem 3.1 (Constant Stepsize BPPA). *Let $D_{\|\cdot\|_*} = \max_{i \in [n]} \|x_i\|_*$, where $\|\cdot\|_*$ denotes the dual norm of $\|\cdot\|$. Then under Assumption 1, for any constant stepsize $\eta_t = \eta > 0$, the following hold.*

(1) We have $\lim_{t \rightarrow \infty} L_S(\theta_t) = 0$. Specifically, we have that $L_S(\theta_t)$ diminishes at the following rate,

$$L_S(\theta_t) \leq \frac{1}{\gamma_{\|\cdot\|_*} \eta t} + \frac{L_w \log^2(\gamma_{\|\cdot\|_*} \eta t)}{4\gamma_{\|\cdot\|_*}^2 \eta t} = \mathcal{O}\left(\frac{L_w \log^2(\gamma_{\|\cdot\|_*} \eta t)}{\gamma_{\|\cdot\|_*}^2 \eta t}\right).$$

(2) We have that the margin is asymptotically lower bounded by

$$\lim_{t \rightarrow \infty} \min_{i \in [n]} \left\langle \frac{\theta_t}{\|\theta_t\|}, y_i x_i \right\rangle \geq \sqrt{\frac{\mu_w}{L_w}} \gamma_{\|\cdot\|_*}, \quad (3.1)$$

where $\gamma_{\|\cdot\|_*}$ is defined in Definition 2.1. In addition, for any given $\epsilon > 0$, there exists a t_0 satisfying

$$t_0 := \tilde{\mathcal{O}}\left(\max\left\{\frac{D_{\|\cdot\|_*}^2}{\epsilon^2 \gamma_{\|\cdot\|_*}^2}, \exp\left(\frac{D_{\|\cdot\|_*}^2}{\gamma_{\|\cdot\|_*}^2} \sqrt{\frac{L_w}{\mu_w}}\right) \frac{1}{\gamma_{\|\cdot\|_*}^2 \eta}\right\}\right),$$

such that for $t \geq t_0$ number of iterations, we have

$$\left\langle \frac{\theta_t}{\|\theta_t\|}, y_i x_i \right\rangle \geq (1 - \epsilon) \sqrt{\frac{\mu_w}{L_w}} \gamma_{\|\cdot\|_*}, \quad \forall i \in [n].$$

We highlight that (1) The choice of Bregman divergence in BPPA is flexible and can be *data dependent*. Properly chosen data-dependent divergence can adapt to data geometry much better than data-independent divergence, leading to better separation and margin. In Section 5 we demonstrate how BPPA can benefit significantly from such an adaptivity of carefully designed data-dependent divergence. (2) Our analysis on the convergence requires handling non-finite minimizers, which implies divergence of iterate $\|\theta_t\| \rightarrow \infty$. The optimization problem of our interest does not meet the standard assumptions in the classical analysis of BPPA in the literature, and requires a careful choice of reference point in order to derive non-trivial convergence results. (3) Our result is closely related, but should not be confused with the homotopy method in (Rosset et al., 2004), which can be viewed as performing only one proximal step at the origin, with an extremely large stepsize. (4) Finally, our result is a generalization of Telgarsky (2013); Gunasekar et al. (2018) to non-euclidean settings with Bregman divergence. Working with Bregman divergence poses unique challenges, as it is previously unclear how to relate the primal margin progress to the per-iteration progress over the dual space.

Theorem 3.1 shows that if the distance generating function $w(\cdot)$ is well-conditioned w.r.t. $\|\cdot\|$ -norm, then Bregman proximal point algorithm will output a solution with near optimal $\|\cdot\|_*$ -norm margin. As a concrete realization of Theorem 3.1, we consider the Mahalanobis distance $\|\cdot\|_A := \sqrt{\langle \cdot, A \cdot \rangle}$ induced by a positive definite matrix A .

Corollary 3.1. *Let $\|\cdot\| = \|\cdot\|_A$ for some positive definite matrix A . Under the same conditions as in Theorem 3.1, BPPA with distance generating function $w(\cdot) = \langle \cdot, A \cdot \rangle$ converges to the maximum $\|\cdot\|_*$ -margin solution, where $\|\cdot\|_* = \|\cdot\|_{A^{-1}}$. Specifically, we have*

$$L_S(\theta_t) \leq \frac{1}{\gamma_{\|\cdot\|_*} \eta t} + \frac{L_w \log^2(\gamma_{\|\cdot\|_*} \eta t)}{4\gamma_{\|\cdot\|_*}^2 \eta t} = \mathcal{O}\left(\frac{L_w \log^2(\gamma_{\|\cdot\|_*} \eta t)}{\gamma_{\|\cdot\|_*}^2 \eta t}\right).$$

In addition, we have $\lim_{t \rightarrow \infty} \min_{i \in [n]} \left\langle \frac{\theta_t}{\|\theta_t\|}, y_i x_i \right\rangle = \gamma_{\|\cdot\|_*}$. Specifically, for any given $\epsilon > 0$, there exists a t_0 satisfying $t_0 := \tilde{\mathcal{O}}\left(\max\left\{\frac{D_{\|\cdot\|_*}^2}{\epsilon^2 \gamma_{\|\cdot\|_*}^2}, \exp\left(\frac{D_{\|\cdot\|_*}^2}{\gamma_{\|\cdot\|_*}^2} \epsilon^2\right) \frac{1}{\gamma_{\|\cdot\|_*}^2 \eta}\right\}\right)$, such that for $t \geq t_0$ number of iterations, we have

$$\left\langle \frac{\theta_t}{\|\theta_t\|}, y_i x_i \right\rangle \geq (1 - \epsilon) \gamma_{\|\cdot\|_*}, \quad \forall i \in [n].$$

Finally, we have the direction convergence that $\lim_{t \rightarrow \infty} \frac{\theta_t}{\|\theta_t\|} = u_{\|\cdot\|_*}$.

Note that similar directional convergence results have been shown in Gunasekar et al. (2018) for steepest descent w.r.t. $\|\cdot\|_A$ norm. The directional convergence of BPPA obtained here, however, is not a simple corollary of known results, since existing analyses focus on first-order algorithms in euclidean setting (e.g., GD/SGD, steepest descent). Such existing analyses do not simply extend to non-first-order algorithms in non-euclidean setting, such as BPPA.

When the distance generating function $w(\cdot)$ is ill-conditioned w.r.t. $\|\cdot\|$ -norm (i.e., $\sqrt{\mu_w/L_w} \ll 1$), it might be tempting to suggest that the margin lower bound in (3.1) is loose, and what really happens is $\lim_{t \rightarrow \infty} \min_{i \in [n]} \left\langle \frac{\theta_t}{\|\theta_t\|}, y_i x_i \right\rangle = \gamma_{\|\cdot\|_*}$. However, as we show in the following proposition, there exists a class of problems, where the lower bound in (3.1) is in fact a tight upper bound (up to a factor of 2), demonstrating that the dependence on condition number of distance generating function $w(\cdot)$ w.r.t. $\|\cdot\|$ -norm is *not* a proof artifact.

Proposition 3.1 (Tight Dependence on Condition Number). *There exists a sequence of problems $\{\mathcal{P}^{(m)}\}_{m \geq 1}$, where each $\mathcal{P}^{(m)} = (\mathcal{S}^{(m)}, \|\cdot\|^{(m)}, w^{(m)})$ denotes the dataset, the norm, and the distance generating function of the m -th problem. For each m , the distance generating function $w^{(m)}(\cdot)$ is $\mu_w^{(m)}$ -strongly convex and $L_w^{(m)}$ -smooth w.r.t. $\|\cdot\|$ -norm. Then Bregman proximal point algorithm applied to each problem in $\{\mathcal{P}^{(m)}\}_{m \geq 1}$ yields*

$$\lim_{t \rightarrow \infty} \min_{(x,y) \in \mathcal{S}^{(m)}} \left\langle \frac{\theta_t}{\|\theta_t\|}, yx \right\rangle / \gamma_{\|\cdot\|_*^{(m)}} \leq 2 \sqrt{\frac{\mu_w^{(m)}}{L_w^{(m)}}}, \quad \forall m \geq 1, \quad (3.2)$$

In addition, for any $m \geq 4$, we have $\lim_{t \rightarrow \infty} \min_{(x,y) \in \mathcal{S}^{(m)}} \left\langle \frac{\theta_t}{\|\theta_t\|}, yx \right\rangle / \gamma_{\|\cdot\|_*^{(m)}} \leq 2 \sqrt{\frac{\mu_w^{(m)}}{L_w^{(m)}}} < 1$.

In fact,

$$\lim_{t \rightarrow \infty} \min_{(x,y) \in \mathcal{S}^{(m)}} \left\langle \frac{\theta_t}{\|\theta_t\|}, yx \right\rangle / \gamma_{\|\cdot\|_*^{(m)}} \rightarrow 0, \quad \text{as } m \rightarrow \infty. \quad (3.3)$$

Combine Theorem 3.1, Corollary 3.1 and Proposition 3.1, we conclude that the margin of the obtained solution by BPPA has non-trivial dependence on the condition number of the distance generating function $w(\cdot)$. This observation provides a strong evidence that the Bregman divergence $D_w(\cdot, \cdot)$ in BPPA is highly important to the quality of the obtained solution, and advocates a careful design of Bregman divergence when using the BPPA. Our theoretical findings also aligns the empirical evidences on the importance of divergence found in knowledge distillation and model fine-tuning (Jiang et al., 2020; Furlanello et al., 2018).

We have shown that BPPA with constant stepsize achieves a margin that is at least $\sqrt{\mu_w/L_w}$ -fraction of the maximal one. Meanwhile, our complexity bound in Theorem 3.1 shows that to obtain such a margin lower bound, it might take an exponential number of iterations. We next show by employing a more aggressive stepsize scheme, we can attain the same margin lower bound in a polynomial number of iterations, while speeding up the convergence of the empirical loss $\{L_S(\theta_t)\}_{t \geq 0}$ drastically.

Theorem 3.2 (Varying Stepsize BPPA). *Given any positive sequence $\{\alpha_t\}_{t \geq 0}$, letting the stepsizes $\{\eta_t\}_{t \geq 0}$ be $\eta_t = \frac{\alpha_t}{L_S(\theta_t)}$, then the following facts hold.*

(1) $\lim_{t \rightarrow \infty} L_S(\theta_t) = 0$. Specifically, for any $t \geq 0$, we have $L_S(\theta_{t+1}) \leq L_S(\theta_t) \beta(\alpha_t)$, where $\beta(\alpha) = \min_{\beta \in (0,1)} \max \left\{ \beta, \exp \left(-\frac{2\alpha\beta^2\gamma_{\|\cdot\|_*}^2}{L_w} \right) \right\} < 1$.

(2) Letting $\alpha_t = \frac{1}{\sqrt{t+1}}$, we have $\lim_{t \rightarrow \infty} \min_{i \in [n]} \left\langle \frac{\theta_t}{\|\theta_t\|}, y_i x_i \right\rangle \geq \sqrt{\frac{\mu_w}{L_w}} \gamma_{\|\cdot\|_*}$. In particular, for any $\epsilon \in (0, \frac{1}{2})$, there exists a t_0 satisfying

$$t_0 = \mathcal{O} \left(\left(\frac{L_w}{\gamma_{\|\cdot\|_*} \sqrt{\mu_w} \epsilon} \right)^8 \right), \quad (3.4)$$

such that in $t \geq t_0$ number of iterations, we have

$$\left\langle \frac{\theta_t}{\|\theta_t\|}, y_i x_i \right\rangle \geq (1 - \epsilon) \sqrt{\frac{L_w}{\mu_w}} \gamma_{\|\cdot\|_*}, \quad \forall i \in [n].$$

Additionally, the convergence rate of $\{L_S(\theta_t)\}_{t \geq 0}$ is given by $L_S(\theta_t) = \mathcal{O}\left(\exp\left(-\frac{\gamma_{\|\cdot\|_*}^2}{L_w} \sqrt{t}\right)\right)$.

We remark that (1) We do not optimize for the best polynomial dependence on $1/\epsilon$ in the iteration complexity (3.4), as our main goal is to show the exponential gap between the complexity presented in Theorem 3.1 and here. We refer interested readers to Appendix B, where we show that we can improve the polynomial dependence with more tailored analysis. (2) We also demonstrate that the empirical loss converges almost exponentially faster with our choice of stepsizes. (3) We reiterate that using the aggressive stepsizes does not change our established margin lower bound, and the exact convergence to the maximum margin solution demonstrated in Corollary 3.1 still holds for this scheme of stepsizes, which can be achieved with a polynomial number of iterations.

• **Inexact Implementation of BPPA.** The proximal update (2.2) requires solving a non-trivial optimization problem, and there has been fruitful results of inexact implementation of BPPA in optimization literature (Rockafellar, 1976b; Yang and Toh, 2021; Solodov and Svaiter, 2000; Monteiro and Svaiter, 2010). Here based on the *varying stepsize scheme* proposed in Theorem 3.2, we discuss the feasibility of a *gradient descent based inexact BPPA* that: (1) admits a simple implementation and achieves *polynomial complexity*, (2) retains the margin properties of exact BPPA. Specifically, at the t -th iteration, the gradient descent based inexact BPPA solves the proximal step

$$\hat{\theta}_{t+1} \approx \operatorname{argmin}_{\theta} \phi_t(\theta) := \frac{1}{n} \sum_{i=1}^n \exp(-\langle \theta, y_i x_i \rangle) + \frac{1}{2\eta_t} D_w(\theta, \hat{\theta}_t) \quad (3.5)$$

up to a pre-specified accuracy δ_t with gradient descent. Our key observation comes from the fact that when applying gradient descent to $\phi_t(\cdot)$ with small enough stepsizes, the iterate would stay in a region that has relative smoothness M_t and relative strong convexity μ_t bounded by

$$M_t \leq L_S(\hat{\theta}_t) + \frac{1}{\eta_t} = L_S(\hat{\theta}_t) \left(1 + \frac{1}{\alpha_t}\right), \quad \mu_t \geq \frac{1}{\eta_t} = \frac{L(\hat{\theta}_t)}{\alpha_t},$$

both measured w.r.t. Bregman divergence $D_h(\cdot, \cdot)$ (Lu et al., 2018). Note that the first inequality follows by our choice of stepsize η_t in Theorem 3.2. Thus the effective condition number $\kappa_t := M_t/\mu_t$ of $\phi_t(\cdot)$ is bounded by $\kappa_t = 1 + \alpha_t = \mathcal{O}(1)$, which implies that the t -th proximal step requires $\mathcal{O}\left(\kappa_t \log\left(\frac{1}{\delta_t}\right)\right) = \mathcal{O}\left(\log\left(\frac{1}{\delta_t}\right)\right)$ number of gradient descent steps. Summing up across t_0 iterations (3.4), we need up to $\mathcal{O}\left(\sum_{t=1}^{t_0} \log\left(\frac{1}{\delta_t}\right)\right)$ gradient descent steps, which depends polynomially on t_0 even if we choose an extremely high accuracy $\delta_t = \mathcal{O}(\exp(-t))$ for each inexact proximal step (3.5).

4 ALGORITHMIC REGULARIZATION OF MIRROR DESCENT

Inspired by the results in the previous section, we further show that mirror descent (MD, Algorithm 2), as a generalization of gradient descent to non-euclidean geometry, possesses similar connection between the margin and Bregman divergence. We remark that our results are the first to characterize the algorithmic regularization effect of MD for classification tasks, while previous literature exclusively focus on under-determined regression problems (Gunasekar et al., 2018; Azizan and Hassibi, 2019).

Algorithm 2 Mirror Descent Algorithm (MD)

Input: Distance generating function $w(\cdot)$, stepsizes $\{\eta^t\}_{t \geq 0}$, samples $\{x_i, y_i\}_{i=1}^n$.
Initialize: $\theta^0 \leftarrow 0$.
for $t = 0, \dots$ **do**
 Compute gradient $\nabla L_S(\theta_t) = \frac{1}{n} \sum_{i=1}^n \exp(-\langle \theta_t, y_i x_i \rangle) (-y_i x_i)$.
 Update $\theta_{t+1} = \operatorname{argmin}_{\theta} \langle \nabla L_S(\theta_t), \theta - \theta_t \rangle + \frac{1}{2\eta_t} D_w(\theta, \theta_t)$.
end for

Theorem 4.1 (Constant Stepsize MD). *Let $D_{\|\cdot\|_*} = \max_{i \in [n]} \|x_i\|_*$, where $\|\cdot\|_*$ denotes the dual norm of $\|\cdot\|$, and $D_{\|\cdot\|_2} = \max_{i \in [n]} \|x_i\|_2$. Under Assumption 1, let μ_2 be the strong convexity parameter of $w(\cdot)$ w.r.t. $\|\cdot\|_2$ -norm. Then for any constant stepsize $\eta_t = \eta \leq \frac{\mu_2}{2D_{\|\cdot\|_2}}$, we have that*

(1) $\lim_{t \rightarrow \infty} L_S(\theta_t) = 0$. Specifically, we have that $L_S(\theta_t)$ diminishes at the following rate,

$$L_S(\theta_t) \leq \frac{1}{\gamma_{\|\cdot\|_*} \eta t} + \frac{L_w \log^2(\gamma_{\|\cdot\|_*} \eta t)}{4\gamma_{\|\cdot\|_*}^2 \eta t} = \mathcal{O}\left(\frac{L_w \log^2(\gamma_{\|\cdot\|_*} \eta t)}{\gamma_{\|\cdot\|_*}^2 \eta t}\right).$$

(2) We have that the margin is asymptotically lower bounded by

$$\lim_{t \rightarrow \infty} \min_{i \in [n]} \left\langle \frac{\theta_t}{\|\theta_t\|}, y_i x_i \right\rangle \geq \sqrt{\frac{\mu_w}{L_w}} \gamma_{\|\cdot\|_*}.$$

In addition, for any $\epsilon > 0$, there exists a t_0 satisfying

$$t_0 = \mathcal{O} \left(\exp \left(\frac{D_{\|\cdot\|_*}^{3/2} D_{\|\cdot\|_2} L_w \eta}{\gamma_{\|\cdot\|_*}^2 \mu_w^{1/2} \mu_2^{3/2} \epsilon^{3/2}} \log \left(\frac{1}{\epsilon} \right) \right) \right), \quad (4.1)$$

such that any $t \geq t_0$, we have

$$\left\langle \frac{\theta_t}{\|\theta_t\|}, y_i x_i \right\rangle \geq (1 - \epsilon) \sqrt{\frac{\mu_w}{L_w}} \gamma_{\|\cdot\|_*},$$

Theorem 4.1 shows that mirror descent attains the same $\|\cdot\|_*$ -norm margin lower bound as BPPA, which is $\sqrt{\mu_w/L_w}$ -fraction of the maximal margin. Note that $\mu_2 > 0$ is a direct consequence of Assumption 1 and the equivalence of norm on finite-dimensional vector space.

Similar to Corollary 3.1, let $\|\cdot\| = \|\cdot\|_A$ be the Mahalanobis distance, then MD equipped with distance generating function $w(\cdot) = \langle \cdot, A \cdot \rangle$ converges to the maximum $\|\cdot\|_*$ -norm margin classifier.

Corollary 4.1. *Let $\|\cdot\| = \|\cdot\|_A$ for some positive definite matrix A . Then under the same conditions as in Theorem 4.1, the MD with distance generating function $w(\cdot) = \langle \cdot, A \cdot \rangle$ converges to the maximum $\|\cdot\|_*$ -margin solution, where $\|\cdot\|_* = \|\cdot\|_{A^{-1}}$. Specifically, we have $L_S(\theta_t) = \mathcal{O} \left(\frac{L_w \log^2(\gamma_{\|\cdot\|_*} \eta t)}{\gamma_{\|\cdot\|_*}^2 \eta t} \right)$.*

In addition, we have $\lim_{t \rightarrow \infty} \min_{i \in [n]} \left\langle \frac{\theta_t}{\|\theta_t\|}, y_i x_i \right\rangle = \gamma_{\|\cdot\|_*}$. Specifically, for any given $\epsilon > 0$, there exists a t_0 with $t_0 = \mathcal{O} \left(\exp \left(\frac{D_{\|\cdot\|_*}^{3/2} D_{\|\cdot\|_2} L_w \eta}{\gamma_{\|\cdot\|_*}^2 \mu_w^{1/2} \mu_2^{3/2} \epsilon^{3/2}} \log \left(\frac{1}{\epsilon} \right) \right) \right)$, such that for $t \geq t_0$ number of iterations, we have

$$\left\langle \frac{\theta_t}{\|\theta_t\|}, y_i x_i \right\rangle \geq (1 - \epsilon) \gamma_{\|\cdot\|_*}, \quad \forall i \in [n].$$

Finally, we have direction convergence that $\lim_{t \rightarrow \infty} \frac{\theta_t}{\|\theta_t\|} = u_{\|\cdot\|_*}$.

Note that Corollary 4.1 recovers the directional convergence of steepest descent in Gunasekar et al. (2018) w.r.t $\|\cdot\|_A$, which coincides with MD with distance generating function $w(\cdot) = \langle \cdot, A \cdot \rangle$.

Theorem 4.1 guarantees the near optimal $\|\cdot\|_*$ -norm margin when the distance generating function $w(\cdot)$ is well-conditioned w.r.t. $\|\cdot\|$ -norm. For cases when $w(\cdot)$ is ill-conditioned, we demonstrate that there exists a class of problem for which the margin lower bound is tight.

Proposition 4.1. *There exists a sequence of problems $\{\mathcal{P}^{(m)}\}_{m \geq 1}$ by the same construction as in Proposition 3.1, such that the margin lower bound in Theorem 4.1 is tight up to a non-trivial factor of 2. Specifically, we have (3.2) and (3.3) also hold for MD.*

Finally, we propose a more aggressive stepsize scheme for MD that achieves the same margin lower bound. In addition, instead of requiring an exponential number of iterations (4.1) as constant stepsize MD, such a stepsize scheme only needs a polynomial number of iterations, and achieves an almost exponential speedup for the empirical loss $\{L_S(\theta_t)\}_{t \geq 0}$.

Theorem 4.2 (Varying Stepsize MD). *Let the stepsizes $\{\eta_t\}_{t \geq 0}$ be given by $\eta_t = \frac{\alpha_t}{L_S(\theta_t)}$, where $\alpha_t = \min\left\{\frac{\mu_2}{2D_{\|\cdot\|_2}}, \frac{1}{\sqrt{t+1}}\right\}$. Then under the same conditions as in Theorem 4.1,*

(1) We have $\lim_{t \rightarrow \infty} \min_{i \in [n]} \left\langle \frac{\theta_t}{\|\theta_t\|}, y_i x_i \right\rangle \geq \sqrt{\frac{\mu_w}{L_w}} \gamma_{\|\cdot\|_*}$. In addition, for any $\epsilon > 0$, there exists a t_0 satisfying $t_0 = \mathcal{O} \left(\left(\frac{D_{\|\cdot\|_2} L_w}{\gamma_{\|\cdot\|_*} \mu_2 \sqrt{\mu_w \epsilon}} \right)^4 \right)$, such that for any $t \geq t_0$, we have

$$\left\langle \frac{\theta_t}{\|\theta_t\|}, y_i x_i \right\rangle \geq (1 - \epsilon) \sqrt{\frac{\mu_w}{L_w}} \gamma_{\|\cdot\|_*}, \quad \forall i \in [n].$$

(2) We have $\lim_{t \rightarrow \infty} L_S(\theta_t) = 0$. In addition, the convergence rate is given by

$$L_S(\theta_t) = \mathcal{O} \left(\exp \left(-\frac{\gamma_{\|\cdot\|_*}^2}{L_w} \sqrt{t} \right) \right).$$

5 EXPERIMENTS

Synthetic Data. We take $\mathcal{S} = \{((-0.5, 1), +1), ((-0.5, -1), -1), ((-0.75, -1), -1), ((2, 1), +1)\}$. One can readily verify that the maximum $\|\cdot\|_2$ -norm margin classifier is $u_{\|\cdot\|_2} = (0, 1)$. For both BPPA and MD, we take the Bregman divergence as $D_w(x, y) = \|x - y\|_2^2$, which corresponds to the vanilla proximal point algorithm and gradient descent algorithm. Note that both algorithms are guaranteed to converge in direction towards $u_{\|\cdot\|_2} = (0, 1)$, following Corollary 3.1 and 4.1.

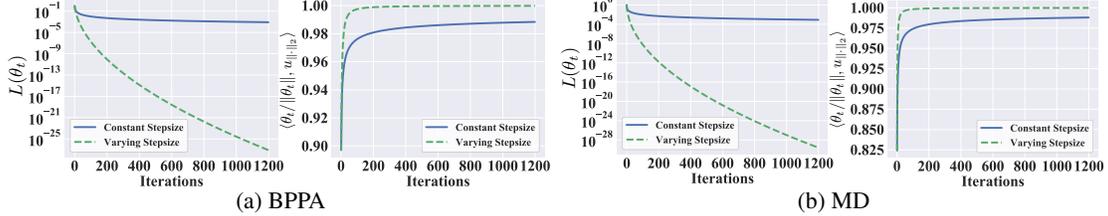


Figure 1: BPPA and MD run on the simple data set \mathcal{S} .

We take $\eta_t = \eta = 1$ for the constant stepsize BPPA/MD, and $\eta_t = \frac{1}{L(\theta_t)\sqrt{t+1}}$ for the varying stepsize BPPA/MD, following the stepsize choices in Theorem 3.1, 3.2, 4.1 and 4.2. To implement the proximal step in BPPA at the t -th iteration, we take 128 number of gradient descent steps with stepsize $0.2\eta_t$, following our discussion at the end of Section 3. We initialize all algorithms at the origin and run 1200 iterations. From Figure 1, we can clearly observe that both BPPA and MD converge in direction to the maximum $\|\cdot\|_2$ -norm margin classifier $u_{\|\cdot\|_2}$, which is consistent with our theoretical findings. In addition, by adopting the varying stepsize scheme proposed in Theorem 3.2 and 4.2, both BPPA and MD converge exponentially faster than their constant stepsize counterparts.

Data-dependent Bregman Divergence. We illustrate through an example on how properly chosen *data-dependent divergence can lead to much improved separation* compared to data-independent divergence, even on simple linear models.

We have n labeled data $\{(x_i, y_i)\}_{i=1}^m$ sampled from a mixture of sphere distribution: $y_i \sim \text{Bernoulli}(1/2)$, $x_i \sim \text{Unif}(\mathbb{S}_{y_i\mu}(r))$, where $\mathbb{S}_z(r)$ denotes the sphere centered at z with radius r in \mathbb{R}^d . In addition, we also have m unlabeled data $\{\tilde{x}_j\}_{j=1}^m$, following the same distribution as $\{x_i\}_{i=1}^m$, with no labels given. Clearly, the maximum $\|\cdot\|_2$ -margin classifier for the mixture of sphere distribution considered here is given by the linear classifier $f^*(\cdot) = \text{sign}(\langle \cdot, \mu \rangle)$.

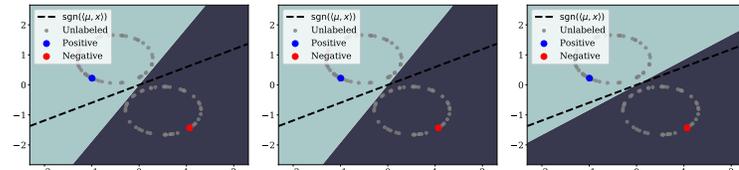


Figure 2: BPPA with Bregman divergence $D^{(3)}$ (right) significantly improves alignment with optimal classifier μ , compared to $D^{(1)}$ (left) and $D^{(2)}$ (middle).

Divergence	Alignment
$D^{(1)}(\cdot, \cdot)$	0.8703
$D^{(2)}(\cdot, \cdot)$	0.8175
$D^{(3)}(\cdot, \cdot)$	0.9754

Table 1: $\left\langle \frac{\theta^T}{\|\theta^T\|_2}, \mu \right\rangle$ averaged over 8 runs.

We choose $n = d = 2, m = 100, r = 0.8$, and generate $\mu \sim \text{Unif}(\mathbb{S}_0(1))$. We compare three types of Bregman divergence, given by $D^{(1)}(\theta, \theta') = \|\theta - \theta'\|_2^2$ (vanilla proximal point), $D^{(2)}(\theta, \theta') = (\theta - \theta')^\top \widehat{\Sigma}(\theta - \theta')$, and $D^{(3)}(\theta, \theta') = (\theta - \theta')^\top \widehat{\Sigma}^{-1}(\theta - \theta')$, where $\widehat{\Sigma} = \frac{1}{m} \sum_{j=1}^m x_j x_j^\top$ denotes the empirical covariance matrix. Note that $D^{(2)}$ and $D^{(3)}$ are data-dependent from their construction. For each divergence function, we run BPPA with 8 independent runs, the results are reported in Figure 2 and Table 1. We make two important remarks on the empirical results:

- Data-dependent divergence $D^{(3)}$ gives the best separation despite limited labeled data (in fact only 2!), much improved over data-independent squared ℓ_2 -distance $D^{(1)}$.
- Not all data-dependent divergence helps, $D^{(2)}$ shows degradation compared to $D^{(1)}$.

We further remark that by utilizing Corollary 3.1, one can completely characterize the solution obtained by BPPA for each of the divergence in closed form. Using such a characterization allows one to corroborate the empirical phenomenon with our developed theories, deferred in Appendix A.

CIFAR-100. We demonstrate the potential of extending our theoretical findings for linear models to practical networks, using ResNet-18 (He et al., 2016), ShuffleNetV2 (Ma et al., 2018), MobileNetV2 (Sandler et al., 2018), with CIFAR-100 dataset (Krizhevsky et al., 2009). At each iteration of BPPA, the updated model parameter θ_{t+1} is given by solving the proximal step $\theta_{t+1} = \operatorname{argmin}_{\theta} 1/n \sum_{i=1}^n \ell(f_{\theta}(x_i); y_i) + 1/(2\eta_t)D(\theta; \theta_t)$ for all $t \geq 0$, where D denotes divergence function, and θ_0 is obtained by standard training with SGD. We consider inexact implementation of the proximal step, discussed in (3.5). Specifically, each proximal step is solved by using SGD, with a batch size of 128, an initial learning rate of 0.1 which is subsequently divided by 5 at the 60th, 120th, and 160th epoch. We consider two divergence functions widely used in practice, defined by $D_{LS}(\theta', \theta) = 1/(2n) \sum_{i=1}^n \|f_{\theta}(x_i) - f_{\theta'}(x_i)\|_2^2$ (Tarvainen and Valpola, 2017), and $D_{KL}(\theta, \theta') = 1/(2n) \sum_{i=1}^n \operatorname{KL}(f_{\theta'}(x_i) \| f_{\theta}(x_i))$ (Furlanello et al., 2018). For each of the divergence, we run BPPA with 3 proximal steps, with the proximal stepsize $\eta_t = \eta = 0.025$ for D_{KL} , and $\eta_t = \eta = 0.2$ for D_{LS} ($\eta_t = 0.025$ gives significantly worse performance). For standard training with SGD, we use a batch size of 128, an initial learning rate of 0.1 further divided by 5 at the 60th, 120th, and 160th epoch. The results are reported in Figure 3.

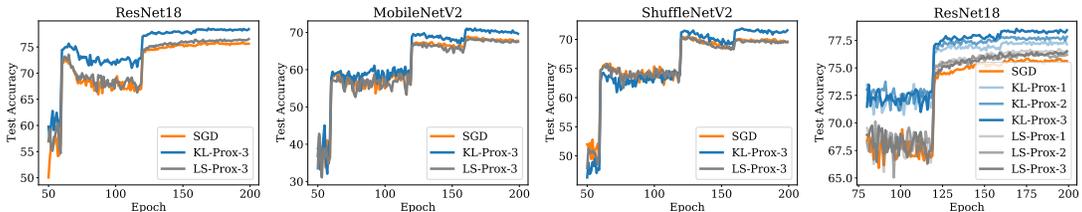


Figure 3: BPPA with divergences D_{KL} and D_{LS} on CIFAR-100 dataset. KL-Prox- k denotes learning curve of the k -th proximal step with D_{KL} ; LS-Prox- k denotes learning curve of the k -th proximal step with D_{LS} .

One can clearly see from Figure 3: (1) Across different model architectures, BPPA with D_{KL} outperforms standard training with SGD; (2) BPPA with D_{LS} yields negligible differences compared to SGD. The qualitative difference of D_{KL} and D_{LS} strongly indicates that the divergence function serves an important role in affecting the model performance learned by BPPA, which we view as an important evidence showing broader applicability of our developed divergence-dependent margin theories. In addition, the learned model with D_{KL} improves gradually w.r.t the total number of proximal steps. For ResNet-18, the accuracy increases from 75.83% (standard training) to 78.56% after 3 proximal steps – an additional 1.4% improvement over Tf-KD_{self} (see Table 2), which can be viewed as BPPA with 1 proximal step. We view such findings as the evidence suggesting the scope of algorithmic regularization associated with BPPA goes beyond simple linear models.

We make further remarks on the previously proposed method in Yuan et al. (2019), named Teacher-free Knowledge Distillation via self-training (Tf-KD_{self}), which is equivalent to BPPA with 1 proximal steps, using $D_{KL}(\theta, \theta')$ as the divergence function. Tf-KD_{self} was shown to improve over SGD for various network architectures on CIFAR-100 and Tiny-ImageNet. We include the reported results on CIFAR-100 therein in Table 2 for completeness.

Model	SGD	Tf-KD _{self}
MobileNetV2	68.38	70.96 (+2.58)
ShuffleNetV2	70.34	72.23 (+1.89)
ResNet18	75.87	77.10 (+1.23)
GoogLeNet	78.72	80.17 (+1.45)
DenseNet121	79.04	80.26 (+1.22)

Table 2: Comparison of Tf-KD_{self} (2-step BPPA) and SGD on CIFAR-100.

6 CONCLUSION AND FUTURE DIRECTION

To conclude, we have shown that for binary classification task with linearly separable data, the Bregman proximal point algorithm and mirror descent attain a $\|\cdot\|_*$ -norm margin that is closely related to the condition number of the distance generating function w.r.t. $\|\cdot\|$ -norm. We list two directions worthy of future investigations. (1) Our analyses exploit the fact that the Bregman divergence is defined over the *model parameters*, while many popular data-dependent divergences are defined over the *model output* (e.g. prediction confidence). Making this non-trivial extension to data-dependent divergence can also help demystify the mechanism of the data-dependent divergence. (2) Our current analyses focus on linear models, and the extension to nonlinear neural networks requires more delicate definitions of margin and divergence. We leave this direction as our long-term investigation plan.

REFERENCES

- ALLEN-ZHU, Z., LI, Y. and LIANG, Y. (2018). Learning and generalization in overparameterized neural networks, going beyond two layers. *arXiv preprint arXiv:1811.04918* .
- AZIZAN, N. and HASSIBI, B. (2019). Stochastic gradient/mirror descent: Minimax optimality and implicit regularization. In *International Conference on Learning Representations*.
- BRUTZKUS, A., GLOBERSON, A., MALACH, E. and SHALEV-SHWARTZ, S. (2017). Sgd learns over-parameterized networks that provably generalize on linearly separable data. *arXiv preprint arXiv:1710.10174* .
- DEVLIN, J., CHANG, M.-W., LEE, K. and TOUTANOVA, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* .
- DUCHI, J., HAZAN, E. and SINGER, Y. (2011). Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research* **12**.
- ECKSTEIN, J. (1993). Nonlinear proximal point algorithms using bregman functions, with applications to convex programming. *Mathematics of Operations Research* **18** 202–226.
- FRENCH, R. M. (1999). Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences* **3** 128–135.
- FURLANELLO, T., LIPTON, Z., TSCHANNEN, M., ITTI, L. and ANANDKUMAR, A. (2018). Born again neural networks. In *International Conference on Machine Learning*. PMLR.
- GOYAL, P., DOLLÁR, P., GIRSHICK, R., NOORDHUIS, P., WESOLOWSKI, L., KYROLA, A., TULLOCH, A., JIA, Y. and HE, K. (2017). Accurate, large minibatch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677* .
- GREEN, S., VINEYARD, C. M. and KOÇ, C. K. (2019). Distillation strategies for proximal policy optimization. *arXiv preprint arXiv:1901.08128* .
- GUNASEKAR, S., LEE, J., SOUDRY, D. and SREBRO, N. (2018). Characterizing implicit bias in terms of optimization geometry. In *International Conference on Machine Learning*. PMLR.
- HE, K., ZHANG, X., REN, S. and SUN, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- HE, T., ZHANG, Z., ZHANG, H., ZHANG, Z., XIE, J. and LI, M. (2019). Bag of tricks for image classification with convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- HINTON, G., VINYALS, O. and DEAN, J. (2015). Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531* .
- Ji, Z., DUDÍK, M., SCHAPIRE, R. E. and TELGARSKY, M. (2020). Gradient descent follows the regularization path for general losses. In *Proceedings of Thirty Third Conference on Learning Theory* (J. Abernethy and S. Agarwal, eds.), vol. 125 of *Proceedings of Machine Learning Research*. PMLR.
- Ji, Z. and TELGARSKY, M. (2018). Gradient descent aligns the layers of deep linear networks. *arXiv preprint arXiv:1810.02032* .
- Ji, Z. and TELGARSKY, M. (2019). The implicit bias of gradient descent on nonseparable data. In *Proceedings of the Thirty-Second Conference on Learning Theory* (A. Beygelzimer and D. Hsu, eds.), vol. 99 of *Proceedings of Machine Learning Research*. PMLR, Phoenix, USA.
- Ji, Z. and TELGARSKY, M. (2021). Characterizing the implicit bias via a primal-dual analysis. In *Algorithmic Learning Theory*. PMLR.

-
- JIANG, H., HE, P., CHEN, W., LIU, X., GAO, J. and ZHAO, T. (2020). SMART: Robust and efficient fine-tuning for pre-trained natural language models through principled regularized optimization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online.
- KAKADE, S., SHALEV-SHWARTZ, S., TEWARI, A. ET AL. (2009). On the duality of strong convexity and strong smoothness: Learning applications and matrix regularization. *Unpublished Manuscript*, <http://ttic.uchicago.edu/shai/papers/KakadeShalevTewari09.pdf> **2**.
- KEESKAR, N. S., MUDIGERE, D., NOCEDAL, J., SMELYANSKIY, M. and TANG, P. T. P. (2016). On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv preprint arXiv:1609.04836* .
- KINGMA, D. P. and BA, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* .
- KIRKPATRICK, J., PASCANU, R., RABINOWITZ, N., VENESS, J., DESJARDINS, G., RUSU, A. A., MILAN, K., QUAN, J., RAMALHO, T., GRABSKA-BARWINSKA, A. ET AL. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences* **114** 3521–3526.
- KIWIEL, K. C. (1997). Proximal minimization methods with generalized bregman functions. *SIAM journal on control and optimization* **35** 1142–1168.
- KRIZHEVSKY, A., HINTON, G. ET AL. (2009). Learning multiple layers of features from tiny images .
- LEWKOWYCZ, A., BAHRI, Y., DYER, E., SOHL-DICKSTEIN, J. and GUR-ARI, G. (2020). The large learning rate phase of deep learning: the catapult mechanism. *arXiv preprint arXiv:2003.02218* .
- LI, Y., WEI, C. and MA, T. (2019). Towards explaining the regularization effect of initial large learning rate in training neural networks. *arXiv preprint arXiv:1907.04595* .
- LI, Z. and HOIEM, D. (2017). Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence* **40** 2935–2947.
- LU, H., FREUND, R. M. and NESTEROV, Y. (2018). Relatively smooth convex optimization by first-order methods, and applications. *SIAM Journal on Optimization* **28** 333–354.
- MA, N., ZHANG, X., ZHENG, H.-T. and SUN, J. (2018). Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *Proceedings of the European conference on computer vision (ECCV)*.
- MCCLOSKEY, M. and COHEN, N. J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, vol. 24. Elsevier, 109–165.
- MONTEIRO, R. D. and SVAITER, B. F. (2010). Convergence rate of inexact proximal point methods with relative error criteria for convex optimization. *submitted to SIAM Journal on Optimization* .
- NACSON, M. S., LEE, J., GUNASEKAR, S., SAVARESE, P. H. P., SREBRO, N. and SOUDRY, D. (2019). Convergence of gradient descent on separable data. In *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics* (K. Chaudhuri and M. Sugiyama, eds.), vol. 89 of *Proceedings of Machine Learning Research*. PMLR.
- NEMIROVSKI, A. S. and YUDIN, D. B. (1983). Problem complexity and method efficiency in optimization .
- ROCKAFELLAR, R. T. (1976a). Augmented lagrangians and applications of the proximal point algorithm in convex programming. *Mathematics of operations research* **1** 97–116.
- ROCKAFELLAR, R. T. (1976b). Monotone operators and the proximal point algorithm. *SIAM journal on control and optimization* **14** 877–898.

-
- ROSSET, S., ZHU, J. and HASTIE, T. (2004). Boosting as a regularized path to a maximum margin classifier. *The Journal of Machine Learning Research* **5** 941–973.
- SANDLER, M., HOWARD, A., ZHU, M., ZHMOGINOV, A. and CHEN, L.-C. (2018). Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- SCHULMAN, J., LEVINE, S., ABBEEL, P., JORDAN, M. and MORITZ, P. (2015). Trust region policy optimization. In *International conference on machine learning*. PMLR.
- SMITH, L. N. (2018). A disciplined approach to neural network hyper-parameters: Part 1—learning rate, batch size, momentum, and weight decay. *arXiv preprint arXiv:1803.09820*.
- SMITH, S. L., KINDERMANS, P.-J. and LE, Q. V. (2018). Don’t decay the learning rate, increase the batch size. In *International Conference on Learning Representations*.
- SOLODOV, M. V. and SVAITER, B. F. (2000). Error bounds for proximal point subproblems and associated inexact proximal point algorithms. *Mathematical programming* **88** 371–389.
- SOUDRY, D., HOFFER, E. and SREBRO, N. (2018). The implicit bias of gradient descent on separable data. In *International Conference on Learning Representations*.
- SUTSKEVER, I., MARTENS, J., DAHL, G. and HINTON, G. (2013). On the importance of initialization and momentum in deep learning. In *Proceedings of the 30th International Conference on Machine Learning* (S. Dasgupta and D. McAllester, eds.), vol. 28 of *Proceedings of Machine Learning Research*. PMLR, Atlanta, Georgia, USA.
- TARVAINEN, A. and VALPOLA, H. (2017). Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *arXiv preprint arXiv:1703.01780*.
- TELGARSKY, M. (2013). Margins, shrinkage, and boosting. In *International Conference on Machine Learning*. PMLR.
- YANG, L. and TOH, K.-C. (2021). Bregman proximal point algorithm revisited: a new inexact version and its variant. *arXiv preprint arXiv:2105.10370*.
- YUAN, L., TAY, F. E., LI, G., WANG, T. and FENG, J. (2019). Revisit knowledge distillation: a teacher-free framework.
- ZASLAVSKI, A. J. (2010). Convergence of a proximal point method in the presence of computational errors in hilbert spaces. *SIAM Journal on Optimization* **20** 2413–2421.
- ZHANG, T., WU, F., KATIYAR, A., WEINBERGER, K. Q. and ARTZI, Y. (2021). Revisiting few-sample {bert} fine-tuning. In *International Conference on Learning Representations*.
- ZHOU, P., YUAN, X., XU, H., YAN, S. and FENG, J. (2019). Efficient meta learning via mini-batch proximal update. In *Advances in Neural Information Processing Systems* (H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox and R. Garnett, eds.), vol. 32. Curran Associates, Inc.

Supplementary Material for “Implicit Regularization of Bregman Proximal Point Algorithm and Mirror Descent on Separable Data”

A DISCUSSION ON EXPERIMENT SECTION

The observed phenomenon for our second experiment can be explained by Corollary 3.1. Recall that μ is the maximal ℓ_2 -norm max-margin solution of the mixture of sphere distribution. In Corollary 3.1, take $A = A^{(1)} := I_d$, $A = A^{(2)} := I_d r^2/d + \mu\mu^\top$ and $A = A^{(3)} := (I_d r^2/d + \mu\mu^\top)^{-1}$ respectively. Note $A^{(1)}, A^{(2)}, A^{(3)}$ by definition are positive definite, hence apply the result of Corollary 3.1, $A^{(3)}$ promotes a solution $\mu^{(3)}$ that has larger directional alignment with μ (hence achieving better separation), while $A^{(2)}$ promotes a solution $\mu^{(2)}$ that has smaller directional alignment with μ (hence has worse separation). It remains to note that $\widehat{\Sigma} \rightarrow A^{(2)}$, $\widehat{\Sigma}^{-1} \rightarrow A^{(3)}$ for m sufficiently large, hence the solutions of BPPA with $D^{(2)}$ and $D^{(3)}$ converge to $\mu^{(2)}$ and $\mu^{(3)}$ for large enough m , respectively. In conclusion, BPPA with $D^{(2)}$ converges to a classifier with smaller alignment with μ (worse separation), while BPPA with $D^{(3)}$ converges to a classifier with larger alignment with μ (better separation).

B PROOFS IN SECTION 3

Proof of Proposition 3.1. For each $m \in \mathbb{Z}^+$, we consider the following simple problem $\mathcal{P}^{(m)}$. First, the data set $\mathcal{S}^{(m)} = \{(x_i, y_i)\}_{i=1}^2 \subset \mathbb{R}^m \times \{+1, -1\}$ contains only two data-points, where $x_1 = z^{(m)}$, $y_1 = 1$ and $x_2 = -z^{(m)}$, $y_2 = 1$ for some vector $z^{(m)} \in \mathbb{R}^m$ to be chosen later.

Consider norm $\|\cdot\|^{(m)} = \|\cdot\|_1$, and distance generating function $w^{(m)}(\theta) = \frac{\|\theta\|_2^2}{2}$, both defined on \mathbb{R}^m . From the simple identity $\frac{1}{2} \|\theta'\|_2^2 = \frac{1}{2} \|\theta\|_2^2 + \langle \theta, \theta' - \theta \rangle + \frac{1}{2} \|\theta' - \theta\|_2^2$, together with the fact that

$$\|\theta - \theta'\|_2^2 \leq \|\theta - \theta'\|_1^2, \quad \|\theta - \theta'\|_2^2 \geq \frac{1}{m} \|\theta - \theta'\|_1^2,$$

we conclude that $w^{(m)}(\cdot)$ is 1-smooth and $\frac{1}{m}$ -strongly convex w.r.t. $\|\cdot\|^{(m)}$ -norm, and $\sqrt{\mu_w^{(m)}/L_w^{(m)}} = \sqrt{1/m}$. By Corollary 3.1, taking $A = I_m$, we conclude that $\lim_{t \rightarrow \infty} \frac{\theta_t}{\|\theta_t\|_2} = u_2^{(m)}$, where

$$u_2^{(m)} = \operatorname{argmax}_{\|\theta\|_2 \leq 1} \min_{(x,y) \in \mathcal{S}^{(m)}} \langle \theta, yx \rangle = \frac{z^{(m)}}{\|z^{(m)}\|_2}$$

is the maximum $\|\cdot\|_2$ -norm margin SVM. Hence we have

$$\lim_{t \rightarrow \infty} \min_{(x,y) \in \mathcal{S}^{(m)}} \left\langle \frac{\theta_t}{\|\theta_t\|_1}, yx \right\rangle = \left\langle u_2^{(m)} \cdot \frac{\|u_2^{(m)}\|_2}{\|u_2^{(m)}\|_1}, z \right\rangle = \frac{\|z^{(m)}\|_2^2}{\|z^{(m)}\|_1}.$$

On the other hand, we can readily verify that $\gamma_{\|\cdot\|_*^{(m)}} := \max_{\|\theta\|^{(m)} \leq 1} \langle \theta, z^{(m)} \rangle = \|z^{(m)}\|_\infty$. Thus we conclude that

$$\lim_{t \rightarrow \infty} \min_{(x,y) \in \mathcal{S}^{(m)}} \left\langle \frac{\theta_t}{\|\theta_t\|_1}, yx \right\rangle / \gamma_{\|\cdot\|_*^{(m)}} = \frac{\|z^{(m)}\|_2^2}{\|z^{(m)}\|_1 \|z^{(m)}\|_\infty}.$$

Now taking $z^{(m)} = (1, \frac{1}{\sqrt{m}}, \dots, \frac{1}{\sqrt{m}})$, one can readily verify that

$$\frac{\|z^{(m)}\|_2^2}{\|z^{(m)}\|_1 \|z^{(m)}\|_\infty} = \frac{2 - 1/m}{\sqrt{m} - 1/\sqrt{m} + 1} \leq \frac{2}{\sqrt{m}}, \quad \forall m \geq 1.$$

Thus we conclude our proof with

$$\lim_{t \rightarrow \infty} \min_{(x,y) \in \mathcal{S}^{(m)}} \left\langle \frac{\theta_t}{\|\theta_t\|_1}, yx \right\rangle / \gamma_{\|\cdot\|_*^{(m)}} = \frac{\|z^{(m)}\|_2^2}{\|z^{(m)}\|_1 \|z^{(m)}\|_\infty} \leq \frac{2}{\sqrt{m}} = 2\sqrt{\frac{\mu_w^{(m)}}{L_w^{(m)}}}, \quad \forall m \geq 1. \quad \square$$

Lemma B.1. Let $f(\cdot)$ be convex, and $D_w(\theta', \theta) = w(\theta') - w(\theta) - \langle \nabla w(\theta), \theta' - \theta \rangle$ be the Bregman divergence associated with $w(\cdot)$. Letting

$$\theta_{t+1} = \operatorname{argmin}_{\theta} f(\theta) + \frac{1}{2\eta} D_w(\theta, \theta_t), \quad (\text{B.1})$$

then for all θ , we have

$$f(\theta_{t+1}) + \frac{1}{2\eta} D_w(\theta, \theta_{t+1}) \leq f(\theta) + \frac{1}{2\eta} D_w(\theta, \theta_t) - \frac{1}{2\eta} D_w(\theta_{t+1}, \theta_t).$$

Proof. By the optimality condition of (B.1), we have

$$\left\langle f'(\theta_{t+1}) + \frac{1}{2\eta} \nabla D_w(\theta_{t+1}, \theta_t), \theta - \theta_{t+1} \right\rangle \geq 0,$$

where $\nabla D_w(\theta_{t+1}, \theta_t)$ denotes the gradient of $D_w(\cdot, \theta_t)$ at θ_{t+1} , and $f'(\theta_{t+1}) \in \partial f(\theta_{t+1})$ denotes the subgradient of $f(\cdot)$ at θ_{t+1} . Now by the definition of the Bregman divergence, we have

$$D_w(\theta, \theta_t) = D_w(\theta_{t+1}, \theta_t) + \langle \nabla D_w(\theta_{t+1}, \theta_t), \theta - \theta_{t+1} \rangle + D_w(\theta, \theta_{t+1}).$$

Combining this with the fact the $\langle f'(\theta_{t+1}), \theta - \theta_{t+1} \rangle \leq f(\theta) - f(\theta_{t+1})$, we conclude our proof. \square

Lemma B.2. For proximal point algorithm with Bregman divergence $D_w(\cdot, \cdot)$, where $w(\cdot)$ is L_w -smooth w.r.t. to $\|\cdot\|$ -norm, we have

$$L(\theta_t) - L(\theta) \leq \frac{L_w}{4\eta t} \|\theta\|^2.$$

Proof. Given the update rule

$$\theta_{s+1} = \operatorname{argmin}_{\theta} L(\theta) + \frac{1}{2\eta} D_w(\theta, \theta_s),$$

together with Lemma B.1, wherein we take $f(\cdot) = L(\cdot)$, and $\theta = \theta_s$, we have that

$$L(\theta_{s+1}) \leq L(\theta_s) - \frac{1}{2\eta} D_w(\theta_s, \theta_{s+1}) - \frac{1}{2\eta} D_w(\theta_{s+1}, \theta_s). \quad (\text{B.2})$$

Since $w(\cdot)$ is convex, we have $D_w(\theta_s, \theta_{s+1}) \geq 0$ and $D_w(\theta_{s+1}, \theta_s) \geq 0$. Thus, we have monotone improvement $L(\theta_{s+1}) \leq L(\theta_s)$. On the other hand, by Lemma B.1, we also have

$$L(\theta_{s+1}) - L(\theta) \leq \frac{1}{2\eta} D_w(\theta, \theta_s) - \frac{1}{2\eta} D_w(\theta, \theta_{s+1}). \quad (\text{B.3})$$

Summing up the previous inequality from $s = 0$ to $s = t - 1$, we have

$$\sum_{s=1}^t L(\theta_s) - L(\theta) \leq \frac{1}{2\eta} D_w(\theta, \theta_0).$$

Additionally, since $w(\cdot)$ is L_w -smooth w.r.t. $\|\cdot\|$ norm, we have

$$D_w(\theta, \theta_0) \leq \frac{L_w}{2} \|\theta - \theta_0\|^2.$$

Combining with the fact $L(\theta_{s+1}) \leq L(\theta_s)$ for all $s \geq 0$ and $\theta_0 = 0$, we arrive at

$$L(\theta_t) - L(\theta) \leq \frac{L_w}{4\eta t} \|\theta\|^2. \quad \square$$

Finally, we show the convergence of loss for binary classification task with separable data. We remark that with separability, the empirical loss $L(\cdot)$ has infimum zero, but such an infimum *can not be attained*.

♠ **Proof of Theorem 3.1:**

◇ **Proof of Theorem 3.1-(I).** Define $u_R = R\mu_{\|\cdot\|_*}$, where

$$\begin{aligned} u_{\|\cdot\|_*} &= \operatorname{argmax}_{\|u\| \leq 1} \min_{i \in [n]} \langle u, y_i x_i \rangle, \\ \gamma_{\|\cdot\|_*} &= \max_{\|u\| \leq 1} \min_{i \in [n]} \langle u, y_i x_i \rangle. \end{aligned}$$

By Lemma B.2, we have $L(\theta_t) \leq L(u_R) + \frac{L_w}{4\eta t} \|u_R\|^2$. On the other hand, we know that

$$L(u_R) = \frac{1}{n} \sum_{i=1}^n \exp \{-R \langle u_{\|\cdot\|_*}, y_i x_i \rangle\} \leq \exp(-R\gamma_{\|\cdot\|_*}).$$

Thus, $L(\theta_t) \leq \exp(-R\gamma_{\|\cdot\|_*}) + \frac{R^2 L_w}{4\eta t}$. Finally, by taking $R = \frac{\log(t\gamma_{\|\cdot\|_*}\eta)}{\gamma_{\|\cdot\|_*}}$, we have

$$L(\theta_t) \leq \frac{1}{\gamma_{\|\cdot\|_*} \eta t} + \frac{L_w \log^2(\gamma_{\|\cdot\|_*} \eta t)}{4\gamma_{\|\cdot\|_*}^2 \eta t} = \mathcal{O} \left\{ \frac{L_w \log^2(\gamma_{\|\cdot\|_*} \eta t)}{\gamma_{\|\cdot\|_*}^2 \eta t} \right\}. \quad (\text{B.4})$$

□

We proceed to show the parameter convergence of the proximal point algorithm with constant stepsize. Before that, we provide some technical lemmas regarding Bregman divergence and Fenchel conjugate duality. The first lemma shows that the Bregman distance between primal variables (x, y) induced by $w(\cdot)$ is the same as the Bregman distance between the dual variables $(z_y = \nabla w(y), z_x = \nabla w(x))$ induced by $w^*(\cdot)$.

Lemma B.3. *Let $w(\cdot)$ be strictly convex and differentiable, and $w^*(\cdot) = \max_x \langle \cdot, x \rangle - w(x)$ be its convex conjugate function. Then we have*

$$D_w(x, y) = D_{w^*}(\nabla w(y), \nabla w(x)), \quad (\text{B.5})$$

where D_w and D_{w^*} denote the Bregman divergence induced by $w(\cdot)$ and $w^*(\cdot)$ respectively.

Proof. By the definition of conjugate function, we have $w^*(v) = \max_u \langle v, u \rangle - w(u)$. Since $w(\cdot)$ is strictly convex, we have $w^*(\cdot)$ is differentiable, and

$$\nabla w^*(v) = \operatorname{argmax}_u \langle v, u \rangle - w(u) \Rightarrow \langle v, \nabla w^*(v) \rangle - w(\nabla w^*(v)) = w^*(v), \quad \forall v.$$

Since $w(\cdot)$ is proper, we have $w = (w^*)^*$, which also gives

$$\begin{aligned} \nabla w(u) &= \operatorname{argmax}_v \langle u, v \rangle - w^*(v) \Rightarrow \langle u, \nabla w(u) \rangle - w^*(\nabla w(u)) = w(u), \\ &\Rightarrow u = \nabla w^*(\nabla w(u)), \quad \forall u. \end{aligned}$$

By the definition of D_{w^*} , we have

$$\begin{aligned} D_{w^*}(\nabla w(y), \nabla w(x)) &= w^*(\nabla w(y)) - w^*(\nabla w(x)) - \langle \nabla w^*(\nabla w(x)), \nabla w(y) - \nabla w(x) \rangle \\ &= w^*(\nabla w(y)) - w^*(\nabla w(x)) - \langle x, \nabla w(y) - \nabla w(x) \rangle \\ &= w^*(\nabla w(y)) - \langle \nabla w(y), y \rangle - [w^*(\nabla w(x)) - \langle \nabla w(x), x \rangle] - \langle \nabla w(y), x - y \rangle \\ &= w(x) - w(y) - \langle \nabla w(y), x - y \rangle \\ &= D_w(x, y). \end{aligned}$$

□

The second lemma establishes the duality between smoothness and strong convexity w.r.t. to general $\|\cdot\|$ -norm. We refer interested readers to Kakade et al. (2009) for the detailed proof.

Lemma B.4. *[Theorem 6, Kakade et al. (2009)] If $f(\cdot)$ is L -smooth and μ -strongly convex ($\mu > 0$) with respect to $\|\cdot\|$ -norm, then $f^*(\cdot)$ is $\frac{1}{\mu}$ -smooth and $\frac{1}{L}$ strongly convex with respect to $\|\cdot\|_*$ -norm. Here $f^*(\cdot) = \max_x \langle x, \cdot \rangle - f(x)$ denotes the convex conjugate of $f(\cdot)$, and $(\|\cdot\|, \|\cdot\|_*)$ are a pair of dual norm.*

◇ **Proof of Theorem 3.1-(2).** From (B.2), we have

$$\begin{aligned}
& \sum_{i=1}^n \exp(-\langle \theta_{t+1}, y_i x_i \rangle) = L(\theta_{t+1}) \\
& \leq L(\theta_t) - \frac{1}{2\eta} D_w(\theta_t, \theta_{t+1}) - \frac{1}{2\eta} D_w(\theta_{t+1}, \theta_t) \\
& = L(\theta_t) \left\{ 1 - \frac{1}{2\eta L(\theta_t)} D_w(\theta_t, \theta_{t+1}) - \frac{1}{2\eta L(\theta_t)} D_w(\theta_{t+1}, \theta_t) \right\} \\
& \leq L(\theta_t) \exp \left\{ -\frac{1}{2\eta L(\theta_t)} D_w(\theta_t, \theta_{t+1}) - \frac{1}{2\eta L(\theta_t)} D_w(\theta_{t+1}, \theta_t) \right\}.
\end{aligned}$$

Hence for each $i \in [n]$, the normalized margin can be lower bounded by

$$\left\langle \frac{\theta_{t+1}}{\|\theta_{t+1}\|}, y_i x_i \right\rangle \geq \frac{-\log n - \log L(\theta_0)}{\|\theta_{t+1}\|} + \frac{1}{2\eta \|\theta_{t+1}\|} \sum_{s=0}^t \frac{D_w(\theta_t, \theta_{t+1}) + D_w(\theta_{t+1}, \theta_t)}{L(\theta_s)}. \quad (\text{B.6})$$

By optimality condition of the proximal update $\theta_{t+1} \in \operatorname{argmin}_{\theta} L(\theta) + \frac{1}{2\eta} D_w(\theta, \theta_t)$, we have that

$$\frac{1}{2\eta} (\nabla w(\theta_{t+1}) - \nabla w(\theta_t)) + \nabla L(\theta_{t+1}) = 0, \quad (\text{B.7})$$

Thus we have

$$\begin{aligned}
& \|\nabla w(\theta_{t+1}) - \nabla w(\theta_t)\|_* = \|\nabla w(\theta_{t+1}) - \nabla w(\theta_t)\|_* \|u_{\|\cdot\|_*}\| \\
& \geq \langle \nabla w(\theta_{t+1}) - \nabla w(\theta_t), u_{\|\cdot\|_*} \rangle \\
& = -2\eta \langle \nabla L(\theta_{t+1}), u_{\|\cdot\|_*} \rangle = \frac{2\eta}{n} \sum_{i=1}^n \langle u_{\|\cdot\|_*}, y_i x_i \rangle \exp(-\langle \theta_{t+1}, y_i x_i \rangle) \\
& \geq 2\eta L(\theta_{t+1}) \gamma_{\|\cdot\|_*},
\end{aligned} \quad (\text{B.8})$$

where the first inequality follows from $\|u_{\|\cdot\|_*}\| = 1$, the second inequality follows from the Fenchel-Young Inequality, the second equality uses the definition of $\nabla L(\theta_{t+1})$, and the final inequality follows from $\langle u_{\|\cdot\|_*}, y_i x_i \rangle \geq \gamma_{\|\cdot\|_*}$.

Now by Lemma B.3, we have $D_w(\theta_t, \theta_{t+1}) = D_{w^*}(\nabla w(\theta_{t+1}), \nabla w(\theta_t))$. Since $w(\cdot)$ is L_w -smooth w.r.t. $\|\cdot\|$ -norm, by Lemma B.4 we have that $w^*(\cdot)$ is $\frac{1}{L_w}$ -strongly convex w.r.t. $\|\cdot\|_*$ -norm, which gives us

$$\begin{aligned}
& D_{w^*}(\nabla w(\theta_{t+1}), \nabla w(\theta_t)) \\
& = w^*(\nabla w(\theta_{t+1})) - w^*(\nabla w(\theta_t)) - \langle \nabla w^*(\nabla w(\theta_t)), \nabla w(\theta_{t+1}) - \nabla w(\theta_t) \rangle \\
& \geq \frac{1}{2L_w} \|\nabla w(\theta_{t+1}) - \nabla w(\theta_t)\|_*^2.
\end{aligned}$$

Thus we have

$$\begin{aligned}
& \sqrt{D_w(\theta_t, \theta_{t+1})} = \sqrt{D_{w^*}(\nabla w(\theta_{t+1}), \nabla w(\theta_t))} \\
& \geq \sqrt{\frac{1}{2L_w}} \|\nabla w(\theta_{t+1}) - \nabla w(\theta_t)\|_* \geq \sqrt{\frac{2}{L_w}} \eta L(\theta_{t+1}) \gamma_{\|\cdot\|_*},
\end{aligned}$$

where the last inequality follows from (B.8).

Together with the fact that $w(\cdot)$ is μ_w -strongly convex w.r.t. $\|\cdot\|$ -norm, we have

$$D_w(\theta_t, \theta_{t+1}) \geq \sqrt{\frac{2}{L_w}} \eta L(\theta_{t+1}) \gamma_{\|\cdot\|_*} \sqrt{D_w(\theta_t, \theta_{t+1})} \geq \sqrt{\frac{\mu_w}{L_w}} \eta L(\theta_{t+1}) \gamma_{\|\cdot\|_*} \|\theta_{t+1} - \theta_t\|. \quad (\text{B.9})$$

By the same argument, we also have that

$$D_w(\theta_{t+1}, \theta_t) \geq \sqrt{\frac{\mu_w}{L_w}} \eta L(\theta_{t+1}) \gamma_{\|\cdot\|_*} \|\theta_{t+1} - \theta_t\|.$$

Together with (B.6), we have

$$\begin{aligned}
& \left\langle \frac{\theta_{t+1}}{\|\theta_{t+1}\|}, y_i x_i \right\rangle \\
& \geq \frac{-\log n - \log L(\theta_0)}{\|\theta_{t+1}\|} + \frac{1}{2\eta \|\theta_{t+1}\|} \sum_{s=0}^t \frac{D_w(\theta_s, \theta_{s+1}) + D_w(\theta_{s+1}, \theta_s)}{L(\theta_s)} \\
& \geq \frac{-\log n - \log L(\theta_0)}{\|\theta_{t+1}\|} + \frac{1}{\eta \|\theta_{t+1}\|} \sum_{s=0}^n \sqrt{\frac{\mu_w}{L_w}} \eta \frac{L(\theta_{s+1})}{L(\theta_s)} \gamma_{\|\cdot\|_*} \|\theta_{s+1} - \theta_s\| \\
& \geq \frac{-\log n - \log L(\theta_0)}{\|\theta_{t+1}\|} + \frac{1}{\|\theta_{t+1}\|} \sum_{s=0}^n \sqrt{\frac{\mu_w}{L_w}} \frac{L(\theta_{s+1})}{L(\theta_s)} \gamma_{\|\cdot\|_*} \|\theta_{s+1} - \theta_s\| \tag{B.10} \\
& \geq \frac{-\log n - \log L(\theta_0)}{\|\theta_{t+1}\|} + \frac{\gamma_{\|\cdot\|_*}}{\|\theta_{t+1}\|} \sqrt{\frac{\mu_w}{L_w}} \sum_{s=0}^t \frac{L(\theta_{s+1})}{L(\theta_s)} \|\theta_{s+1} - \theta_s\| \\
& \geq \frac{-\log n - \log L(\theta_0)}{\|\theta_{t+1}\|} + \gamma_{\|\cdot\|_*} \sqrt{\frac{\mu_w}{L_w}} \left\{ \sum_{s=0}^t \frac{L(\theta_{s+1})}{L(\theta_s)} \|\theta_{s+1} - \theta_s\| \right\} / \left\{ \sum_{s=0}^t \|\theta_{s+1} - \theta_s\| \right\}.
\end{aligned}$$

Next, we provide two technical lemmas regarding properties of the iterates produced by proximal point algorithms.

Lemma B.5. *The iterates produced by proximal point algorithm $\{\theta_s\}_{s \geq 0}$ satisfy*

$$\sum_{s=0}^t \|\theta_{s+1} - \theta_s\| \geq \|\theta_{t+1}\| \geq \frac{1}{D_{\|\cdot\|_*}} \left\{ \log \left(2\gamma_{\|\cdot\|_*}^2 \eta t \right) - 2 \log \log \left(\gamma_{\|\cdot\|_*} \eta t \right) \right\}, \tag{B.11}$$

$$1 > \frac{L(\theta_{s+1})}{L(\theta_s)} \geq \exp \left\{ -D_{\|\cdot\|_*} \sqrt{\frac{2}{s\mu_w \gamma_{\|\cdot\|_*}^2} + \frac{L_w \log^2(\gamma_{\|\cdot\|_*} \eta s)}{2\mu_w \gamma_{\|\cdot\|_*}^2 s}} \right\}. \tag{B.12}$$

Proof. For any $t \geq 0$, since $\theta_0 = 0$, we have

$$\|\theta_{t+1}\| \leq \sum_{s=0}^t \|\theta_{s+1} - \theta_s\|.$$

Now by the convergence of empirical loss (B.4), together with the assumption that $\|x_i\|_* \leq D_{\|\cdot\|_*}$,

$$\exp(-D \|\theta_{t+1}\|) \leq L(\theta_t) \leq \frac{1}{\gamma_{\|\cdot\|_*} \eta t} + \frac{L_w^2 \log^2(\gamma_{\|\cdot\|_*} \eta t)}{4\gamma_{\|\cdot\|_*}^2 \eta t} = \mathcal{O} \left\{ \frac{L_w \log^2(\gamma_{\|\cdot\|_*} \eta t)}{4\gamma_{\|\cdot\|_*}^2 \eta t} \right\},$$

Thus we have the following lower bound,

$$\sum_{s=0}^t \|\theta_{s+1} - \theta_s\| \geq \|\theta_{t+1}\| \geq \frac{1}{D_{\|\cdot\|_*}} \left\{ \log \left(4\gamma_{\|\cdot\|_*}^2 \eta t / L_w \right) - 2 \log \log \left(\gamma_{\|\cdot\|_*} \eta t \right) \right\}.$$

On the other hand, we have

$$\frac{L(\theta_{s+1})}{L(\theta_t)} = \frac{\sum_{i=1}^n \exp(-\langle \theta_s, y_i x_i \rangle) \exp(-\langle \theta_{s+1} - \theta_s, y_i x_i \rangle)}{\sum_{i=1}^n \exp(-\langle \theta_s, y_i x_i \rangle)} \geq \exp(-\|\theta_{s+1} - \theta_s\| D_{\|\cdot\|_*}).$$

We then establish an upper bound on $\|\theta_{s+1} - \theta_s\|$. By (B.2) and the fact the $L(\theta) \geq 0$ for all θ , we have

$$D_w(\theta_s, \theta_{s+1}) + D_w(\theta_{s+1}, \theta_s) \leq 2\eta (L(\theta_s) - L(\theta_{s+1})).$$

On the other hand, since $w(\cdot)$ is μ_w -strongly convex, we have

$$\frac{\mu_w}{2} \|\theta_{s+1} - \theta_s\|^2 \leq D_w(\theta_s, \theta_{s+1}), \quad \frac{\mu_w}{2} \|\theta_{s+1} - \theta_s\|^2 \leq D_w(\theta_{s+1}, \theta_s).$$

Thus we obtain

$$\|\theta_{s+1} - \theta_s\| \leq \sqrt{\eta(L(\theta_s) - L(\theta_{s+1}))} \leq \sqrt{\frac{2\eta}{\mu_w} L(\theta_s)}, \quad \frac{L(\theta_{s+1})}{L(\theta_t)} < 1.$$

Together with the convergence of empirical loss in (B.4), we have

$$\|\theta_{s+1} - \theta_s\| \leq \sqrt{\frac{2}{s\mu_w\gamma_{\|\cdot\|_*}} + \frac{L_w \log^2(\gamma_{\|\cdot\|_*}\eta s)}{2\mu_w\gamma_{\|\cdot\|_*}^2 s}} = \mathcal{O} \left\{ \sqrt{\frac{L_w \log^2(\gamma_{\|\cdot\|_*}\eta s)}{\mu_w\gamma_{\|\cdot\|_*}^2 s}} \right\}, \quad (\text{B.13})$$

Thus we have

$$\frac{L(\theta_{s+1})}{L(\theta_t)} \geq \exp(-\|\theta_{s+1} - \theta_s\| D_{\|\cdot\|_*}) = \Omega \left\{ \exp \left(-D_{\|\cdot\|_*} \sqrt{\frac{L_w \log^2(\gamma_{\|\cdot\|_*}\eta s)}{\mu_w\gamma_{\|\cdot\|_*}^2 s}} \right) \right\}. \quad \square$$

Now by (B.10), we have

$$\begin{aligned} \left\langle \frac{\theta_{t+1}}{\|\theta_{t+1}\|}, y_i x_i \right\rangle &\geq \underbrace{\frac{-\log n - \log L(\theta_0)}{\|\theta_{t+1}\|}}_{(a)} \\ &\quad + \underbrace{\gamma_{\|\cdot\|_*} \sqrt{\frac{\mu_w}{L_w}} \left\{ \sum_{s=0}^t \frac{L(\theta_{s+1})}{L(\theta_s)} \|\theta_{s+1} - \theta_s\| \right\}}_{(b)} \bigg/ \left\{ \sum_{s=0}^t \|\theta_{s+1} - \theta_s\| \right\}. \end{aligned} \quad (\text{B.14})$$

We then derive a large t to ensure $\left\langle \frac{\theta_{t+1}}{\|\theta_{t+1}\|}, y_i x_i \right\rangle \geq (1 - \epsilon) \sqrt{\frac{\mu_w}{L_w} \gamma_{\|\cdot\|_*}}$ for all $i \in [n]$.

We first address the term (a) in (B.14). By (B.11) in Lemma B.5 we have

$$\|\theta_{t+1}\| \geq \frac{1}{D_{\|\cdot\|_*}} \left\{ \log \left(4\gamma_{\|\cdot\|_*}^2 \eta t / L_w \right) - 2 \log \log \left(\gamma_{\|\cdot\|_*} \eta t \right) \right\}.$$

Thus combining with $\theta_0 = 0$, (i.e., $L(\theta_0) = 1$), we need

$$\begin{aligned} \left| \frac{-\log n - \log L(\theta_0)}{\|\theta_{t+1}\|} \right| &= \frac{\log n}{\|\theta_{t+1}\|} \leq \frac{D_{\|\cdot\|_*} \log n}{\log \left(4\gamma_{\|\cdot\|_*}^2 \eta t / L_w \right) - 2 \log \log \left(\gamma_{\|\cdot\|_*} \eta t \right)} \\ &\leq \sqrt{\frac{\mu_w}{L_w} \gamma_{\|\cdot\|_*} \frac{\epsilon}{2}}. \end{aligned} \quad (\text{B.15})$$

One can readily verify that by taking

$$t = \Theta \left\{ \exp \left(\frac{D_{\|\cdot\|_*} \log n}{\epsilon} \gamma_{\|\cdot\|_*} \right) \right\},$$

the previous inequality (B.15) holds.

It remains to bound term (b) in (B.14). We need

$$\left\{ \sum_{s=0}^t \frac{L(\theta_{s+1})}{L(\theta_s)} \|\theta_{s+1} - \theta_s\| \right\} \bigg/ \left\{ \sum_{s=0}^t \|\theta_{s+1} - \theta_s\| \right\} \geq 1 - \frac{\epsilon}{2}. \quad (\text{B.16})$$

For any \underline{t} such that $0 \leq \underline{t} \leq t - 1$,

$$\begin{aligned} &\left\{ \sum_{s=0}^t \frac{L(\theta_{s+1})}{L(\theta_s)} \|\theta_{s+1} - \theta_s\| \right\} \bigg/ \left\{ \sum_{s=0}^t \|\theta_{s+1} - \theta_s\| \right\} \\ &= \frac{1}{\sum_{k=0}^t \|\theta_{k+1} - \theta_k\|} \left\{ \sum_{s=0}^{\underline{t}} \|\theta_{s+1} - \theta_s\| \frac{L(\theta_{s+1})}{L(\theta_s)} + \sum_{s=\underline{t}+1}^t \|\theta_{s+1} - \theta_s\| \frac{L(\theta_{s+1})}{L(\theta_s)} \right\} \\ &> \frac{1}{\sum_{k=0}^t \|\theta_{k+1} - \theta_k\|} \left\{ \sum_{s=\underline{t}+1}^t \|\theta_{s+1} - \theta_s\| \frac{L(\theta_{s+1})}{L(\theta_s)} \right\}. \end{aligned}$$

By (B.12) in Lemma B.5, we have

$$\frac{L(\theta_{s+1})}{L(\theta_s)} \geq \exp \left\{ -D_{\|\cdot\|_*} \sqrt{\frac{2}{s\mu_w\gamma_{\|\cdot\|_*}} + \frac{L_w \log^2(\gamma_{\|\cdot\|_*}\eta s)}{2\mu_w\gamma_{\|\cdot\|_*}^2 s}} \right\}.$$

One can readily verify that, for $\underline{t} = \tilde{\Theta} \left(\frac{L_w D_{\|\cdot\|_*}^2}{\mu_w \epsilon^2 \gamma_{\|\cdot\|_*}^2} \right)$, and any $s \geq \underline{t}$

$$\frac{L(\theta_{s+1})}{L(\theta_s)} \geq 1 - \frac{\epsilon}{4}. \quad (\text{B.17})$$

Thus we have

$$\frac{1}{\sum_{k=0}^{\underline{t}} \|\theta_{k+1} - \theta_k\|} \left\{ \sum_{s=\underline{t}+1}^{\underline{t}} \|\theta_{s+1} - \theta_s\| \frac{L(\theta_{s+1})}{L(\theta_s)} \right\} \geq \left[1 - \frac{\sum_{s=0}^{\underline{t}} \|\theta_{s+1} - \theta_s\|}{\sum_{k=0}^{\underline{t}} \|\theta_{k+1} - \theta_k\|} \right] \left(1 - \frac{\epsilon}{4} \right). \quad (\text{B.18})$$

By (B.13) in the proof of Lemma B.5, we have

$$\|\theta_{s+1} - \theta_s\| \leq \sqrt{\frac{2}{s\mu_w\gamma_{\|\cdot\|_*}} + \frac{L_w \log^2(\gamma_{\|\cdot\|_*}\eta s)}{2\mu_w\gamma_{\|\cdot\|_*}^2 s}},$$

Thus we obtain the following upper bound on $\sum_{s=0}^{\underline{t}} \|\theta_{s+1} - \theta_s\|$ that

$$\begin{aligned} \sum_{s=0}^{\underline{t}} \|\theta_{s+1} - \theta_s\| &\leq \sqrt{\frac{2\underline{t}}{\mu_w\gamma_{\|\cdot\|_*}} + \frac{1}{\gamma_{\|\cdot\|_*}}} \sqrt{\frac{L_w}{\mu_w}} \left[\sqrt{2\underline{t}_o \log^2(\gamma_{\|\cdot\|_*}\eta)} + \sqrt{\log^2(\underline{t})\underline{t}} \right] \\ &= \mathcal{O} \left(\frac{\sqrt{\underline{t} \log^2(\underline{t})}}{\gamma_{\|\cdot\|_*}} \sqrt{\frac{L_w}{\mu_w}} \right). \end{aligned}$$

Together with (B.11), we obtain

$$\begin{aligned} \frac{\sum_{s=0}^{\underline{t}} \|\theta_{s+1} - \theta_s\|}{\sum_{k=0}^{\underline{t}} \|\theta_{k+1} - \theta_k\|} &\leq D_{\|\cdot\|_*} \frac{\sqrt{\frac{2\underline{t}}{\mu_w\gamma_{\|\cdot\|_*}} + \frac{1}{\gamma_{\|\cdot\|_*}}} \sqrt{\frac{L_w}{\mu_w}} \left[\sqrt{2\underline{t}_o \log^2(\gamma_{\|\cdot\|_*}\eta)} + \sqrt{\log^2(\underline{t})\underline{t}} \right]}{\log(2\gamma_{\|\cdot\|_*}^2 \eta \underline{t}) - 2 \log \log(\gamma_{\|\cdot\|_*} \eta \underline{t})} \\ &= \mathcal{O} \left(\frac{D_{\|\cdot\|_*} \sqrt{\underline{t} \log^2(\underline{t})}}{\gamma_{\|\cdot\|_*} \log(\gamma_{\|\cdot\|_*}^2 \eta \underline{t})} \sqrt{\frac{L_w}{\mu_w}} \right). \end{aligned}$$

Thus, letting $t = \Theta \left\{ \exp \left(\frac{D_{\|\cdot\|_*} \sqrt{\underline{t} \log \underline{t}}}{\gamma_{\|\cdot\|_*} \epsilon} \sqrt{\frac{L_w}{\mu_w}} \right) / \gamma_{\|\cdot\|_*}^2 \eta \right\}$, we have

$$\frac{\sum_{s=0}^{\underline{t}} \|\theta_{s+1} - \theta_s\|}{\sum_{k=0}^{\underline{t}} \|\theta_{k+1} - \theta_k\|} \leq \frac{\epsilon}{4}. \quad (\text{B.19})$$

In summary, for $\underline{t} = \tilde{\Theta} \left(\frac{D_{\|\cdot\|_*}^2 L_w}{\epsilon^2 \gamma_{\|\cdot\|_*}^2 \mu_w} \right)$, and

$$\begin{aligned} t &= \max \left\{ \underline{t}, \Theta \left\{ \exp \left(\frac{D_{\|\cdot\|_*} \sqrt{\underline{t} \log \underline{t}}}{\gamma_{\|\cdot\|_*} \epsilon} \sqrt{\frac{L_w}{\mu_w}} \right) / \gamma_{\|\cdot\|_*}^2 \eta \right\} \right\} \\ &= \max \left\{ \tilde{\Theta} \left(\frac{D_{\|\cdot\|_*}^2 L_w}{\epsilon^2 \gamma_{\|\cdot\|_*}^2 \mu_w} \right), \Theta \left\{ \exp \left(\frac{L_w D_{\|\cdot\|_*}^2}{\mu_w \gamma_{\|\cdot\|_*}^2 \epsilon^2} \log \left(\frac{1}{\epsilon} \right) \right) / \gamma_{\|\cdot\|_*}^2 \eta \right\} \right\}, \end{aligned}$$

by combining (B.17), (B.18), and (B.19), we have

$$\left\{ \sum_{s=0}^{\underline{t}} \frac{L(\theta_{s+1})}{L(\theta_s)} \|\theta_{s+1} - \theta_s\| \right\} / \left\{ \sum_{s=0}^{\underline{t}} \|\theta_{s+1} - \theta_s\| \right\} \geq \left(1 - \frac{\epsilon}{4} \right)^2 > 1 - \frac{\epsilon}{2}.$$

Or equivalently, we have (B.16) holds.

Finally, combine (B.14), (B.15), and (B.16), we conclude that there exists $t_0 = \max \left\{ \tilde{\mathcal{O}} \left(\frac{L_w D_{\|\cdot\|_*}^2}{\mu_w \epsilon^2 \gamma_{\|\cdot\|_*}^2} \right), \mathcal{O} \left\{ \exp \left(\frac{L_w D_{\|\cdot\|_*}^2}{\mu_w \gamma_{\|\cdot\|_*}^2} \log \left(\frac{1}{\epsilon} \right) \right) / \gamma_{\|\cdot\|_*}^2 \eta \right\} \right\}$, such that

$$\left\langle \frac{\theta_{t+1}}{\|\theta_{t+1}\|}, y_i x_i \right\rangle \geq (1 - \epsilon) \sqrt{\frac{\mu_w}{L_w}} \gamma_{\|\cdot\|_*}, \quad \forall i \in [n],$$

for all $t \geq t_0$. Hence we conclude our proof. \square

\square

\square

♠ **Proof of Corollary 3.1.** Define $\phi(u) = \min_{i \in [n]} \langle u, y_i x_i \rangle$. We observe that $\phi(\cdot)$ is concave, and $\phi(\cdot)$ is positive homogenous.

By the previous observations, we claim two properties of $\mathcal{U}^* = \operatorname{argmax}_{\|u\| \leq 1} \phi(u)$:

(a) \mathcal{U}^* is singleton, i.e., $\mathcal{U}^* = \{u_{\|\cdot\|_*}\}$.

(b) $\|u_{\|\cdot\|_*}\| = 1$.

To show (a), suppose that there exists $u_1, u_2 \in \mathcal{U}^*$ and $u_1 \neq u_2$. Let $u_\alpha = \alpha u_1 + (1 - \alpha) u_2$ for any $\alpha \in (0, 1)$. Since $\|\cdot\|$ is a strictly convex function, $\{u : \|u\| \leq 1\}$ is a strictly convex set. Hence $u_\alpha \in \operatorname{int}(\{u : \|u\| \leq 1\})$ and $\|u_\alpha\| < 1$. In addition $\phi(u_\alpha) \geq \alpha \phi(u_1) + (1 - \alpha) \phi(u_2) = \phi(u_1)$ given the fact that $\phi(\cdot)$ is concave and $\phi(u_1) = \phi(u_2)$.

Then we have

$$\phi \left(\frac{u_\alpha}{\|u_\alpha\|} \right) = \frac{1}{\|u_\alpha\|} \phi(u_\alpha) \geq \frac{1}{\|u_\alpha\|} \phi(u_1) > \phi(u_1),$$

which contradicts the fact that $u_1 \in \mathcal{U}^*$. Thus we have (a) holds.

On the other hand, using positive homogeneity, for any $u_* \in \mathcal{U}^*$ such that $\|u_*\| < 1$, we have $\bar{u}_* = u_* / \|u_*\|$ satisfy $\phi(\bar{u}_*) > \phi(u_*)$ and $\|\bar{u}_*\| = 1$, which again yields a contradiction. Thus we have (b) hold.

To show asymptotic convergence, since we have shown in Theorem 3.1-(2) that $\lim_{t \rightarrow \infty} \phi \left(\frac{\theta_t}{\|\theta_t\|} \right) = \max_{\|u\| \leq 1} \phi(u)$, given that $\mathcal{U}^* = \{u_{\|\cdot\|_*}\}$ is a singleton and $\phi(\cdot)$ is continuous, we must have $\lim_{t \rightarrow \infty} \frac{\theta_t}{\|\theta_t\|} = u_{\|\cdot\|_*}$. \square

♠ **Proof of Theorem 3.2:**

◇ **Proof of Theorem 3.2-(I).** By (B.2), we have $L(\theta_{s+1}) \leq L(\theta_s) - \frac{1}{2\eta_s} D_w(\theta_s, \theta_{s+1}) - \frac{1}{2\eta_s} D_w(\theta_{s+1}, \theta_s)$. Recall that our stepsize is given by $\eta_s = \frac{\alpha_s}{L(\theta_s)}$, which gives us

$$L(\theta_{s+1}) \leq L(\theta_s) - \frac{1}{2\eta_s} D_w(\theta_s, \theta_{s+1}) - \frac{1}{2\eta_s} D_w(\theta_{s+1}, \theta_s) \tag{B.20}$$

$$\begin{aligned} &\leq L(\theta_s) \exp \left\{ -\frac{1}{2\eta_s L(\theta_s)} D_w(\theta_s, \theta_{s+1}) - \frac{1}{2\eta_s L(\theta_s)} D_w(\theta_{s+1}, \theta_s) \right\} \\ &= L(\theta_s) \exp \left\{ -\frac{1}{2\alpha_s} D_w(\theta_s, \theta_{s+1}) - \frac{1}{2\alpha_s} D_w(\theta_{s+1}, \theta_s) \right\}. \end{aligned} \tag{B.21}$$

In addition, given the fact that $w(\cdot)$ is L_w -smooth w.r.t. $\|\cdot\|$ -norm, By Lemma B.4, we have that $w^*(\cdot)$ is $\frac{1}{L_w}$ -strongly convex w.r.t. $\|\cdot\|_*$ -norm, and hence

$$\begin{aligned} D_w(\theta_s, \theta_{s+1}) &= D_{w^*}(\nabla w(\theta_{s+1}), \nabla w(\theta_s)) \\ &\geq \frac{1}{2L_w} \|\nabla w(\theta_{s+1}) - \nabla w(\theta_s)\|_*^2 \\ &\geq \frac{2\eta_s^2}{L_w} L^2(\theta_{s+1}) \gamma_{\|\cdot\|_*}^2 \\ &= \frac{2\alpha_s^2}{L_w} \left(\frac{L(\theta_{s+1})}{L(\theta_s)} \right)^2 \gamma_{\|\cdot\|_*}^2, \end{aligned} \tag{B.22}$$

where the first equality holds by Lemma B.5, and the second inequality holds by (B.8). Following the same arguments, we also have $D_w(\theta_{s+1}, \theta_s) \geq \frac{2\alpha_s^2}{L_w} \left(\frac{L(\theta_{s+1})}{L(\theta_s)} \right)^2 \gamma_{\|\cdot\|_*}^2$. Hence we obtain

$$\frac{L(\theta_{s+1})}{L(\theta_s)} \leq \exp \left\{ - \left(\frac{L(\theta_{s+1})}{L(\theta_s)} \right)^2 \frac{2\alpha_s \gamma_{\|\cdot\|_*}^2}{L_w} \right\}.$$

We conclude that $\beta_s := L(\theta_{s+1})/L(\theta_s) < 1$. Now for any $\beta \in (0, 1)$, if $\beta_s \geq \beta$, then

$$\beta_s = \frac{L(\theta_{s+1})}{L(\theta_s)} \leq \exp \left(-\beta^2 \frac{2\alpha_s \gamma_{\|\cdot\|_*}^2}{L_w} \right).$$

Hence we have

$$\frac{L(\theta_{s+1})}{L(\theta_s)} \leq \beta_s = \max \left\{ \beta, \exp \left(-\beta^2 \frac{2\alpha_s \gamma_{\|\cdot\|_*}^2}{L_w} \right) \right\}, \quad \forall \beta \in (0, 1).$$

Since the the previous inequality holds for any $\beta \in (0, 1)$, we minimize the right hand side with respect to β , and obtain

$$\frac{L(\theta_{s+1})}{L(\theta_s)} \leq \min_{\beta \in (0, 1)} \max \left\{ \beta, \exp \left(-\beta^2 \frac{2\alpha_s \gamma_{\|\cdot\|_*}^2}{L_w} \right) \right\} := \beta(\alpha_s). \tag{B.23}$$

Hence loss $\{L(\theta_s)\}_{s \geq 0}$ contracts with contraction coefficient $\beta(\alpha_s)$. \square

To proceed, we first establish a lemma stating that $L(\theta_{s+t})/L(\theta_s)$ has a lower bound.

Lemma B.6. *The iterates produced by Proximal Point Algorithm with varying stepsize $\eta_t = \frac{\alpha_t}{L(\theta_t)}$ satisfy*

$$\frac{L(\theta_{t+1})}{L(\theta_t)} \geq \underline{\beta}(\alpha_t) := \max_{\beta \in (0, 1)} \min \left\{ \beta, \exp \left(-\alpha_t D_{\|\cdot\|_*}^2 \beta \right) \right\}, \quad \forall t \geq 0. \tag{B.24}$$

Proof. Since $w(\cdot)$ is μ_w -strongly convex w.r.t. $\|\cdot\|$ -norm, by Lemma B.4, we have that $w^*(\cdot)$ is $\frac{1}{\mu_w}$ -smooth w.r.t. to $\|\cdot\|_*$ norm. Thus,

$$\begin{aligned} D_w(\theta_s, \theta_{s+1}) &= D_{w^*}(\nabla w(\theta_{s+1}), \nabla w(\theta_s)) \\ &\leq \frac{1}{2\mu_w} \|\nabla w(\theta_{s+1}) - \nabla w(\theta_s)\|_*^2 \\ &= \frac{2\eta_s^2}{\mu_w} \|\nabla L(\theta_{s+1})\|_*^2 \\ &\leq \frac{2\eta_s^2}{\mu_w} L^2(\theta_{s+1}) D_{\|\cdot\|_*}^2 = \frac{2\alpha_s^2}{\mu_w} \left(\frac{L(\theta_{s+1})}{L(\theta_s)} \right)^2 D_{\|\cdot\|_*}^2, \end{aligned}$$

where the first equality holds by Lemma B.3, the second equality holds by the optimality condition of the proximal update (B.7), and the last inequality holds by the fact that $\|\nabla L(\theta_{s+1})\| = \left\| \frac{1}{n} \sum_{i=1}^n \exp(-\langle \theta_{s+1}, y_i x_i \rangle) y_i x_i \right\| \leq L(\theta_{s+1}) D_{\|\cdot\|_*}$ and $\|x_i\|_* \leq D_{\|\cdot\|_*}$. Thus by $w(\cdot)$ is μ_w -strongly convex w.r.t. $\|\cdot\|$ -norm, i.e., $D_w(\theta_s, \theta_{s+1}) \geq \frac{\mu_w}{2} \|\theta_s - \theta_{s+1}\|^2$, we obtain $\|\theta_{s+1} - \theta_s\| \leq \frac{2\alpha_s}{\mu_w} \frac{L(\theta_{s+1})}{L(\theta_s)} D_{\|\cdot\|_*}$.

Thus we have

$$\begin{aligned}
L(\theta_{s+1}) &= \frac{1}{n} \sum_{i=1}^n \exp(-\langle \theta_s, y_i x_i \rangle) \exp(-\langle \theta_{s+1} - \theta_s, y_i x_i \rangle) \\
&\geq \frac{1}{n} \sum_{i=1}^n \exp(-\langle \theta_s, y_i x_i \rangle) \exp(-\|\theta_{s+1} - \theta_s\| D_{\|\cdot\|_*}) \\
&\geq L(\theta_s) \exp\left(-\frac{2\alpha_s D_{\|\cdot\|_*}^2 L(\theta_{s+1})}{\mu_w L(\theta_s)}\right).
\end{aligned}$$

Letting $\beta_s := \frac{L(\theta_{s+1})}{L(\theta_s)}$, then for any $\beta \in (0, 1)$, if $\beta_s \geq \beta$, we have $\frac{L(\theta_{s+1})}{L(\theta_s)} \geq \exp\left(-\frac{2\alpha_s D_{\|\cdot\|_*}^2}{\mu_w} \beta\right)$.

Thus we conclude that

$$\frac{L(\theta_{s+1})}{L(\theta_s)} = \beta_s \geq \min\left\{\beta, \exp\left(-\frac{2\alpha_s D_{\|\cdot\|_*}^2}{\mu_w} \beta\right)\right\}, \quad \forall \beta \in (0, 1).$$

Maximizing the right hand side w.r.t. β , we have

$$\frac{L(\theta_{s+1})}{L(\theta_s)} \geq \underline{\beta}(\alpha_s) := \max_{\beta \in (0, 1)} \min\left\{\beta, \exp\left(-\frac{2\alpha_s D_{\|\cdot\|_*}^2}{\mu_w} \beta\right)\right\}. \quad (\text{B.25})$$

□

◇ **Proof of Theorem 3.2-(2).** We start by (B.21) and obtain

$$\begin{aligned}
L(\theta_t) &= L(\theta_{t-1}) \exp\left\{-\frac{1}{2\alpha_s} D_w(\theta_{t-1}, \theta_t) - \frac{1}{2\alpha_s} D_w(\theta_t, \theta_{t-1})\right\} \\
&\leq L(\theta_0) \exp\left\{-\frac{1}{2\alpha_s} \sum_{s=0}^{t-1} D_w(\theta_s, \theta_{s+1}) - \frac{1}{2\alpha_s} \sum_{s=0}^{t-1} D_w(\theta_{s+1}, \theta_s)\right\}.
\end{aligned}$$

From which we obtain that for any $i \in [n]$,

$$\begin{aligned}
\left\langle \frac{\theta_t}{\|\theta_t\|}, y_i x_i \right\rangle &\geq \frac{-\log n - \log L(\theta_0)}{\|\theta_t\|} + \frac{\sum_{s=0}^{t-1} D_w(\theta_s, \theta_{s+1})}{2\alpha_s \|\theta_t\|} + \frac{\sum_{s=0}^{t-1} D_w(\theta_{s+1}, \theta_s)}{2\alpha_s \|\theta_t\|} \\
&\geq \frac{\sum_{s=0}^{t-1} D_w(\theta_s, \theta_{s+1})}{2\alpha_s \sum_{s=0}^{t-1} \|\theta_{s+1} - \theta_s\|} + \frac{\sum_{s=0}^{t-1} D_w(\theta_{s+1}, \theta_s)}{2\alpha_s \sum_{s=0}^{t-1} \|\theta_{s+1} - \theta_s\|} - \frac{\log n + \log L(\theta_0)}{\|\theta_t\|}
\end{aligned} \quad (\text{B.26})$$

$$\geq \underbrace{\gamma_{\|\cdot\|_*} \sqrt{\frac{\mu_w}{L_w}} \frac{\sum_{s=0}^{t-1} \frac{L(\theta_{s+1})}{L(\theta_s)} \|\theta_{s+1} - \theta_s\|}{\sum_{s=0}^{t-1} \|\theta_{s+1} - \theta_s\|}}_{(a)} - \underbrace{\frac{\log n + \log L(\theta_0)}{\sum_{s=0}^{t-1} \|\theta_{s+1} - \theta_s\|}}_{(b)}, \quad (\text{B.27})$$

where the second inequality follows from triangle inequality, and the third inequality follows from (B.9) where we have

$$D_w(\theta_s, \theta_{s+1}) \geq \sqrt{\frac{\mu_w}{L_w}} \eta_s L(\theta_{s+1}) \gamma_{\|\cdot\|_*} \|\theta_{s+1} - \theta_s\| = \alpha_s \sqrt{\frac{\mu_w}{L_w}} \frac{L(\theta_{s+1})}{L(\theta_s)} \gamma_{\|\cdot\|_*},$$

and similarly $D_w(\theta_{s+1}, \theta_s) \geq \alpha_s \sqrt{\frac{\mu_w}{L_w}} \underline{\beta}(\alpha_s) \gamma_{\|\cdot\|_*} \|\theta_{s+1} - \theta_s\|$. We proceed to bound the two terms in (B.27) separately.

We first give two technical observations. By (B.22) we readily obtain the first one that

$$\frac{L_w}{2} \|\theta_{s+1} - \theta_s\|^2 \geq D_w(\theta_s, \theta_{s+1}) \geq \frac{2\alpha_s^2}{L_w} \left(\frac{L(\theta_{s+1})}{L(\theta_s)}\right)^2 \gamma_{\|\cdot\|_*}^2,$$

or equivalently, $\|\theta_{s+1} - \theta_s\| \geq \frac{2\alpha_s}{L_w} \frac{L(\theta_{s+1})}{L(\theta_s)} \gamma_{\|\cdot\|_*}$.

For the second observation, we start from (B.20), which leads to

$$\frac{L(\theta_s)}{2\alpha_s} \mu_w \|\theta_s - \theta_{s+1}\|^2 \leq \frac{1}{2\eta_s} (D_w(\theta_s, \theta_{s+1}) + D_w(\theta_{s+1}, \theta_s)) \leq L(\theta_s), \quad (\text{B.28})$$

where the first inequality follows from the definition of η_s , together with the fact that $w(\cdot)$ is μ_w -strongly convex w.r.t. $\|\cdot\|$ -norm. Hence we have $\|\theta_s - \theta_{s+1}\| \leq \sqrt{\frac{2\alpha_s}{\mu_w}}$.

Claim: For any $\epsilon > 0$, there exists \underline{t} , such that for all $s \geq \underline{t}$, we have $\frac{L(\theta_{s+1})}{L(\theta_s)} \geq 1 - \epsilon$.

Given the pervious claim, for any $t > \underline{t}$, we can lower bound (a)

$$\begin{aligned} \frac{\sum_{s=0}^{t-1} \frac{L(\theta_{s+1})}{L(\theta_s)} \|\theta_{s+1} - \theta_s\|}{\sum_{s=0}^{t-1} \|\theta_{s+1} - \theta_s\|} &\geq (1 - \epsilon) \frac{\sum_{s=\underline{t}}^{t-1} \|\theta_{s+1} - \theta_s\|}{\sum_{s=0}^{t-1} \|\theta_{s+1} - \theta_s\|} \\ &\geq (1 - \epsilon) \left(1 - \frac{\sum_{s=0}^{\underline{t}} \|\theta_{s+1} - \theta_s\|}{\sum_{s=0}^{t-1} \|\theta_{s+1} - \theta_s\|} \right) \end{aligned}$$

Now by our two observations, we have

$$\begin{aligned} \sum_{s=0}^{\underline{t}-1} \|\theta_{s+1} - \theta_s\| &\leq \sqrt{\frac{2}{\mu_w}} \sum_{s=0}^{\underline{t}-1} \sqrt{\alpha_s}, \\ \sum_{s=0}^{t-1} \|\theta_{s+1} - \theta_s\| &\geq \frac{2\gamma_{\|\cdot\|_*}}{L_w} \sum_{s=0}^{t-1} \alpha_s \frac{L(\theta_s)}{L(\theta_s)} \geq \frac{2\gamma_{\|\cdot\|_*}}{L_w} \sum_{s=\underline{t}}^{t-1} \alpha_s \frac{L(\theta_s)}{L(\theta_s)} \geq \frac{2\gamma_{\|\cdot\|_*}(1-\epsilon)}{L_w} \sum_{s=\underline{t}}^{t-1} \alpha_s. \end{aligned}$$

Thus we obtain

$$\begin{aligned} \frac{\sum_{s=0}^{t-1} \frac{L(\theta_{s+1})}{L(\theta_s)} \|\theta_{s+1} - \theta_s\|}{\sum_{s=0}^{t-1} \|\theta_{s+1} - \theta_s\|} &\geq (1 - \epsilon) \left(1 - \sqrt{\frac{2}{\mu_w}} \frac{L_w}{2\gamma_{\|\cdot\|_*}(1-\epsilon)} \frac{\sum_{s=0}^{\underline{t}-1} \sqrt{\alpha_s}}{\sum_{s=0}^{t-1} \alpha_s} \right) \\ &\geq (1 - \epsilon) \left(1 - \sqrt{\frac{2}{\mu_w}} \frac{L_w}{\gamma_{\|\cdot\|_*}} \frac{\sum_{s=0}^{\underline{t}-1} \sqrt{\alpha_s}}{\sum_{s=0}^{t-1} \alpha_s} \right), \end{aligned}$$

where the last inequality follows from $\epsilon < \frac{1}{2}$. Now take $\alpha_s = (s+1)^{-1/2}$, we have

$$\begin{aligned} \frac{\sum_{s=0}^{t-1} \frac{L(\theta_{s+1})}{L(\theta_s)} \|\theta_{s+1} - \theta_s\|}{\sum_{s=0}^{t-1} \|\theta_{s+1} - \theta_s\|} &\geq (1 - \epsilon) \left(1 - \sqrt{\frac{2}{\mu_w}} \frac{L_w}{\gamma_{\|\cdot\|_*}} \mathcal{O} \left(\frac{t^{3/4}}{\sqrt{t} - \sqrt{\underline{t}}} \right) \right) \\ &= (1 - \epsilon) \left(1 - \sqrt{\frac{2}{\mu_w}} \frac{L_w}{\gamma_{\|\cdot\|_*}} \mathcal{O} \left(\frac{t^{3/8}}{\sqrt{t} - t^{1/4}} \right) \right) \\ &= (1 - \epsilon) \left(1 - \sqrt{\frac{2}{\mu_w}} \frac{L_w}{\gamma_{\|\cdot\|_*}} \mathcal{O} \left(\frac{1}{t^{1/8}} \right) \right), \end{aligned} \quad (\text{B.29})$$

where the last two equalities hold by taking $t = \underline{t}^2$.

It remains to determine the size of \underline{t} . Note that for this we need

$$\frac{L(\theta_{s+1})}{L(\theta_s)} \geq 1 - \epsilon, \quad \forall s \geq \underline{t}.$$

By Lemma B.6, we have $\frac{L(\theta_{s+1})}{L(\theta_s)} \geq \underline{\beta}(\alpha_s) := \max_{\beta \in (0,1)} \min \{ \beta, \exp(-\alpha_s D_{\|\cdot\|_*}^2 \beta) \}$. To make $\underline{\beta}(\alpha_s) = (1 - \epsilon)$, we must have

$$\alpha_s = \frac{-\log(1 - \epsilon)}{D_{\|\cdot\|_*}^2 \underline{\beta}(\alpha_s)} = \frac{-\log(1 - \epsilon)}{D_{\|\cdot\|_*}^2 (1 - \epsilon)}.$$

Thus by choosing $\underline{t} \geq \frac{\log^2(1-\epsilon)}{D_{\|\cdot\|_*}^2 (1-\epsilon)^2} = \Theta \left(\frac{1}{D_{\|\cdot\|_*}^2 \epsilon^2} \right)$, we have $\underline{\beta}(\alpha_{\underline{t}}) \geq 1 - \epsilon$. Since $\underline{\beta}(\cdot)$ is a decreasing function, and $\{\alpha_s\}_{s \geq 0}$ is a decreasing sequence, we have $\underline{\beta}(\alpha_s) \geq 1 - \epsilon$ for all $s \geq \underline{t}$.

Thus in summary, by taking

$$t \geq \max \left\{ \underline{t}^2, \left(\frac{L_w}{\gamma_{\|\cdot\|_*} \sqrt{\mu_w} \epsilon} \right)^8 \right\} = \Theta \left(\left(\frac{L_w}{\gamma_{\|\cdot\|_*} \sqrt{\mu_w} \epsilon} \right)^8 \right),$$

from (B.29), we have control over term (a) as

$$\frac{\sum_{s=0}^{t-1} \frac{L(\theta_{s+1})}{L(\theta_s)} \|\theta_{s+1} - \theta_s\|}{\sum_{s=0}^{t-1} \|\theta_{s+1} - \theta_s\|} \geq (1 - \epsilon)^2 \geq 1 - 2\epsilon.$$

To control term (b), recall

$$\sum_{s=0}^{t-1} \|\theta_{s+1} - \theta_s\| \geq \frac{2\gamma_{\|\cdot\|_*} (1 - \epsilon)}{L_w} \sum_{s=\underline{t}}^{t-1} \alpha_s \geq \frac{\gamma_{\|\cdot\|_*}}{L_w} \sum_{s=\underline{t}}^{t-1} \alpha_s = \Theta \left(\frac{\gamma_{\|\cdot\|_*}}{L_w} \sqrt{t} \right),$$

where the last inequality follows from $\epsilon < \frac{1}{2}$ and $t = \underline{t}^2$. Thus by taking

$$t \geq \max \left\{ \underline{t}^2, \frac{\log^2 n L_w^3}{\gamma_{\|\cdot\|_*}^4 \mu_w \epsilon^2} \right\} = \Theta \left(\frac{1}{D_{\|\cdot\|_*}^4 \epsilon^4} \right),$$

we have a bound for term (b) that $\frac{\log n + \log L(\theta_0)}{\sum_{s=0}^{t-1} \|\theta_{s+1} - \theta_s\|} \leq \sqrt{\frac{\mu_w}{L_w}} \gamma_{\|\cdot\|_*} \epsilon$.

Putting everything together, there exists $t_0 = \Theta \left(\left(\frac{L_w}{\gamma_{\|\cdot\|_*} \sqrt{\mu_w} \epsilon} \right)^8 \right)$, such that for all $t \geq t_0$, we have

$$\left\langle \frac{\theta_t}{\|\theta_t\|}, y_i x_i \right\rangle \geq (1 - 3\epsilon) \sqrt{\frac{\mu_w}{L_w}} \gamma_{\|\cdot\|_*}, \quad \forall i \in [n].$$

□

◇ **Proof of Theorem 3.2-(2).** To show loss convergence, by (B.23), we have

$$L(\theta_t) \leq \exp \left\{ \sum_{s=0}^{t-1} \log \beta(\alpha_s) \right\}.$$

Now by the definition of $\beta(\alpha_s)$ we have

$$\beta(\alpha_s) = \exp \left\{ -\beta^2(\alpha_s) \frac{2\alpha_s \gamma_{\|\cdot\|_*}^2}{L_w} \right\} \leq \exp \left(-\underline{\beta}^2(\alpha_s) \frac{2\alpha_s \gamma_{\|\cdot\|_*}^2}{L_w} \right).$$

By choosing $\epsilon = \frac{1}{2}$ in the proof of Theorem 3.2-(2), we have that for $t \geq \underline{t} \geq \frac{4 \log^2(2)}{D_{\|\cdot\|_*}}$, $\underline{\beta}(\alpha_t) \geq \frac{1}{2}$.

Hence for $t \geq \frac{4 \log^2(2)}{D_{\|\cdot\|_*}}$, we have $\beta(\alpha_s) \leq \exp \left(-\frac{\alpha_s \gamma_{\|\cdot\|_*}^2}{2L_w} \right)$. Thus we conclude that

$$L(\theta_t) \leq \exp \left\{ \sum_{s=0}^{t-1} \log \beta(\alpha_s) \right\} \leq \exp \left\{ -\sum_{s=\underline{t}}^{t-1} \frac{\alpha_s \gamma_{\|\cdot\|_*}^2}{2L_w} \right\} = \mathcal{O} \left(\exp \left(-\frac{\gamma_{\|\cdot\|_*}^2}{L_w} \sqrt{t} \right) \right).$$

□

□

C PROOFS IN SECTION 4

Proof of Proposition 4.1. The proof is a direct consequence of Corollary 4.1 and the proof of Proposition 3.1. \square

♠ *Proof of Theorem 4.1:*

◇ *Proof of Theorem 4.1-(1):* We first show that for small enough stepsize η_t , we have monotonic improvement, i.e., $L(\theta_{t+1}) \leq L(\theta_t)$. Suppose not, and we have $L(\theta_{t+1}) > L(\theta_t)$, then given the second-order Taylor expansion, there exists $\tilde{\theta}_t = \alpha_t \theta_t + (1 - \alpha_t) \theta_{t+1}$ for some $\alpha_t \in [0, 1]$, such that

$$L(\theta_{t+1}) \leq L(\theta_t) + \langle \nabla L(\theta_t), \theta_{t+1} - \theta_t \rangle + \frac{1}{2} (\theta_{t+1} - \theta_t)^\top \nabla^2 L(\tilde{\theta}_t) (\theta_{t+1} - \theta_t).$$

By some calculation, we have

$$\nabla^2 L(\tilde{\theta}_t) = \sum_{i=1}^n \exp\left(-\langle \tilde{\theta}_t, y_i x_i \rangle\right) x_i x_i^\top \preceq D_{\|\cdot\|_2} L(\tilde{\theta}_t) I_d.$$

Since $L(\cdot)$ is a convex function, by the definition of $\tilde{\theta}_t$ we have

$$L(\tilde{\theta}_t) \leq \alpha_t L(\theta_t) + (1 - \alpha_t) L(\theta_{t+1}) \leq \max\{L(\theta_t), L(\theta_{t+1})\} = L(\theta_{t+1}),$$

where the last inequality holds by our assumption $L(\theta_{t+1}) > L(\theta_t)$. Hence we have

$$\begin{aligned} L(\theta_{t+1}) &\leq L(\theta_t) + \langle \nabla L(\theta_t), \theta_{t+1} - \theta_t \rangle + \frac{\max\{L(\theta_{t+1}), L(\theta_t)\}}{2} D_{\|\cdot\|_2} \|\theta_{t+1} - \theta_t\|_2^2 \quad (\text{C.1}) \\ &= L(\theta_t) + \langle \nabla L(\theta_t), \theta_{t+1} - \theta_t \rangle + \frac{L(\theta_{t+1})}{2} D_{\|\cdot\|_2} \|\theta_{t+1} - \theta_t\|_2^2. \end{aligned}$$

To proceed, we employ the well known three-point lemma in mirror descent analysis, which states the following.

Lemma C.1. *Given Bregman divergence $D_w(\cdot, \cdot)$ and mirror descent update*

$$\theta_{t+1} = \operatorname{argmin}_{\theta} \langle \nabla L(\theta_t), \theta - \theta_t \rangle + \frac{1}{2\eta_t} D_w(\theta, \theta_t),$$

we have

$$\langle \nabla L(\theta_{t+1}), \theta_{t+1} - \theta_t \rangle + \frac{1}{2\eta_t} D_w(\theta_{t+1}, \theta_t) \leq \frac{1}{2\eta_t} D_w(\theta, \theta_t) - \frac{1}{2\eta_t} D_w(\theta, \theta_{t+1}), \quad \forall \theta.$$

Proof. The proof follows exactly the same lines as in Lemma B.1. \square

Since distance generating function $w(\cdot)$ is μ_w -strongly convex w.r.t. $\|\cdot\|$ -norm, by the equivalence of norm, we know that there exists $\mu_2 > 0$, such that $w(\cdot)$ is μ_2 -strongly convex w.r.t. $\|\cdot\|_2$ -norm. That is $D_w(\theta', \theta) \geq \frac{\mu_2}{2} \|\theta' - \theta\|_2^2$ for any (θ', θ) . Thus combined with Lemma C.1,

$$\begin{aligned} L(\theta_{t+1}) &\leq L(\theta_t) + \langle \nabla L(\theta_t), \theta_{t+1} - \theta_t \rangle + \frac{L(\theta_{t+1})}{2} D_{\|\cdot\|_2} \|\theta_{t+1} - \theta_t\|_2^2 \\ &\leq L(\theta_t) + \langle \nabla L(\theta_t), \theta_{t+1} - \theta_t \rangle + \frac{L(\theta_{t+1})}{\mu_2} D_{\|\cdot\|_2} D_w(\theta_{t+1}, \theta_t) \\ &\leq L(\theta_t) + \langle \nabla L(\theta_t), \theta_{t+1} - \theta_t \rangle + \frac{1}{2\eta_t} D_w(\theta_{t+1}, \theta_t) - \left(\frac{1}{2\eta_t} - \frac{L(\theta_{t+1}) D_{\|\cdot\|_2}}{\mu_2} \right) D_w(\theta_{t+1}, \theta_t) \\ &\leq L(\theta_t) + \langle L(\theta_t), \theta - \theta_t \rangle + \frac{1}{2\eta_t} D_w(\theta, \theta_t) \\ &\quad - \frac{1}{2\eta_t} D_w(\theta, \theta_{t+1}) - \left(\frac{1}{2\eta_t} - \frac{L(\theta_{t+1}) D_{\|\cdot\|_2}}{\mu_2} \right) D_w(\theta_{t+1}, \theta_t). \end{aligned}$$

Taking $\theta = \theta_t$ in the previous inequality, we obtain

$$L(\theta_{t+1}) \leq L(\theta_t) - \frac{1}{2\eta_t} D_w(\theta_t, \theta_{t+1}) - \left(\frac{1}{2\eta_t} - \frac{L(\theta_{t+1}) D_{\|\cdot\|_2}}{\mu_2} \right) D_w(\theta_{t+1}, \theta_t).$$

Simple rearrangement gives us

$$L(\theta_{t+1}) \left(1 - \frac{D_w(\theta_{t+1}, \theta_t) D_{\|\cdot\|_2}}{\mu_2} \right) \leq L(\theta_t) \left(1 - \frac{1}{2\eta_t L(\theta_t)} D_w(\theta_t, \theta_{t+1}) - \frac{1}{2\eta_t L(\theta_t)} D_w(\theta_{t+1}, \theta_t) \right).$$

One can readily check that whenever $\frac{D_w(\theta_{t+1}, \theta_t) D_{\|\cdot\|_2}}{\mu_2} \leq \frac{1}{2\eta_t L(\theta_t)} \{D_w(\theta_t, \theta_{t+1}) + D_w(\theta_{t+1}, \theta_t)\}$, we have $L(\theta_{t+1}) \leq L(\theta_t)$ and hence a contradiction. Equivalently, we need

$$\eta_t \leq \frac{\mu_2}{2L(\theta_t)} \cdot \frac{D_w(\theta_t, \theta_{t+1}) + D_w(\theta_{t+1}, \theta_t)}{D_w(\theta_{t+1}, \theta_t) D_{\|\cdot\|_2}} \leq \frac{\mu_2 L_w}{L(\theta_t) \mu_w D_{\|\cdot\|_2}},$$

where the last inequality comes from $w(\cdot)$ is L_w -smooth and μ_w -strongly convex w.r.t. $\|\cdot\|$ -norm, which gives us $D_w(\theta_t, \theta_{t+1}) \leq \frac{L_w}{2} \|\theta_{t+1} - \theta_t\|^2$, and $\frac{\mu_w}{2} \|\theta_{t+1} - \theta_t\|^2 \leq D_w(\theta_{t+1}, \theta_t) \leq \frac{L_w}{2} \|\theta_{t+1} - \theta_t\|^2$.

We have shown that $L(\theta_{t+1}) \leq L(\theta_t)$ whenever $\eta_t \leq \frac{\mu_2 L_w}{L(\theta_t) \mu_w D_{\|\cdot\|_2}}$. In addition, by (C.1), whenever $\eta_t = \eta \leq \frac{\mu_2}{2L(\theta_t) D_{\|\cdot\|_2}}$, we have for any θ ,

$$\begin{aligned} L(\theta_{t+1}) &\leq L(\theta_t) + \langle \nabla L(\theta_t), \theta_{t+1} - \theta_t \rangle + \frac{L(\theta_t)}{2} D_{\|\cdot\|_2} \|\theta_{t+1} - \theta_t\|^2 \\ &\leq L(\theta_t) + \langle \nabla L(\theta_t), \theta_{t+1} - \theta_t \rangle + \frac{L(\theta_t)}{\mu_2} D_{\|\cdot\|_2} D_w(\theta_{t+1}, \theta_t) \\ &\leq L(\theta_t) + \langle \nabla L(\theta_t), \theta_{t+1} - \theta_t \rangle + \frac{1}{2\eta} D_w(\theta_{t+1}, \theta_t) \tag{C.2} \\ &\leq L(\theta_t) + \langle \nabla L(\theta_t), \theta - \theta_t \rangle + \frac{1}{2\eta} D_w(\theta, \theta_t) - \frac{1}{2\eta} D_w(\theta, \theta_{t+1}) \\ &\leq L(\theta) + \frac{1}{2\eta} D_w(\theta, \theta_t) - \frac{1}{2\eta} D_w(\theta, \theta_{t+1}). \end{aligned}$$

Thus $L(\theta_{t+1}) - L(\theta) \leq \frac{1}{2\eta} D_w(\theta, \theta_t) - \frac{1}{2\eta} D_w(\theta, \theta_{t+1})$, which is the same as the recursion (B.3) obtained in Lemma B.2. Finally, observe that whenever $\eta_t = \eta \leq \frac{\mu_2}{2L(\theta_0) D_{\|\cdot\|_2}} = \frac{\mu_2}{2D_{\|\cdot\|_2}}$, we always have $L(\theta_{t+1}) \leq L(\theta_t) \leq L(\theta_0)$ for all $t \geq 0$. Thus by simple induction argument, $\eta_t = \eta \leq \frac{\mu_2}{2L(\theta_t) D_{\|\cdot\|_2}}$ for all $t \geq 0$ and all the previous stated statements holds for all $t \geq 0$.

Follow the exact same lines as in Lemma B.2 starting from recursion (B.3), and the proof of Theorem 3.1-(1), we obtain

$$L(\theta_t) \leq \frac{1}{\gamma_{\|\cdot\|_*} \eta t} + \frac{L_w \log^2(\gamma_{\|\cdot\|_*} \eta t)}{4\gamma_{\|\cdot\|_*}^2 \eta t} = \mathcal{O} \left\{ \frac{L_w \log^2(\gamma_{\|\cdot\|_*} \eta t)}{\gamma_{\|\cdot\|_*}^2 \eta t} \right\}. \tag{C.3}$$

□

◇ *Proof of Theorem 4.1-(2):* Starting from (C.1), whenever $\eta_t = \eta \leq \frac{\mu_2}{2D_{\|\cdot\|_2}}$, we have $L(\theta_{t+1}) \leq L(\theta_t)$ for all $t \geq 0$ given previous proof of Theorem 4.1-(1), and hence for any θ ,

$$\begin{aligned} L(\theta_{t+1}) &\leq L(\theta_t) + \langle \nabla L(\theta_t), \theta_{t+1} - \theta_t \rangle + \frac{L(\theta_t)}{2} D_{\|\cdot\|_2} \|\theta_{t+1} - \theta_t\|^2 \\ &\leq L(\theta_t) + \langle \nabla L(\theta_t), \theta_{t+1} - \theta_t \rangle + \frac{L(\theta_t)}{\mu_2} D_{\|\cdot\|_2} D_w(\theta_{t+1}, \theta_t) \\ &= L(\theta_t) + \langle \nabla L(\theta_t), \theta - \theta_t \rangle + \frac{1}{2\eta} D_w(\theta_{t+1}, \theta_t) - \left(\frac{1}{2\eta} - \frac{L(\theta_t)}{\mu_2} D_{\|\cdot\|_2} \right) D_w(\theta_{t+1}, \theta_t) \\ &\leq L(\theta_t) + \langle \nabla L(\theta_t), \theta - \theta_t \rangle + \frac{1}{2\eta} D_w(\theta, \theta_t) - \frac{1}{2\eta} D_w(\theta, \theta_{t+1}) \\ &\quad - \left(\frac{1}{2\eta} - \frac{L(\theta_t)}{\mu_2} D_{\|\cdot\|_2} \right) D_w(\theta_{t+1}, \theta_t), \end{aligned}$$

where the last inequality follows from Lemma C.1. Taking $\theta = \theta_t$, we have

$$L(\theta_{t+1}) \leq L(\theta_t) - \frac{1}{2\eta} D_w(\theta_t, \theta_{t+1}) - \left(\frac{1}{2\eta} - \frac{L(\theta_t)}{\mu_2} D_{\|\cdot\|_2} \right) D_w(\theta_{t+1}, \theta_t) \quad (\text{C.4})$$

$$\begin{aligned} &\leq L(\theta_t) \exp \left\{ -\frac{1}{2\eta L(\theta_t)} D_w(\theta_t, \theta_{t+1}) - \frac{1}{L(\theta_t)} \left(\frac{1}{2\eta} - \frac{L(\theta_t)}{\mu_2} D_{\|\cdot\|_2} \right) D_w(\theta_{t+1}, \theta_t) \right\} \\ &\leq L(\theta_0) \exp \left\{ -\sum_{s=0}^t \frac{1}{2\eta L(\theta_s)} D_w(\theta_s, \theta_{s+1}) - \sum_{s=0}^t \frac{1}{L(\theta_s)} \left(\frac{1}{2\eta} - \frac{L(\theta_s)}{\mu_2} D_{\|\cdot\|_2} \right) D_w(\theta_{s+1}, \theta_s) \right\} \end{aligned} \quad (\text{C.5})$$

Now we can lower bound $D_w(\theta_s, \theta_{s+1})$ by

$$\begin{aligned} D_w(\theta_s, \theta_{s+1}) &= \sqrt{D_w(\theta_s, \theta_{s+1})} \sqrt{D_w(\theta_s, \theta_{s+1})} \\ &= \sqrt{D_{w^*}(\nabla w(\theta_{s+1}), \nabla w(\theta_s))} \sqrt{D_w(\theta_s, \theta_{s+1})} \\ &\geq \sqrt{\frac{1}{2L_w}} \|\nabla w(\theta_{s+1}) - \nabla w(\theta_s)\|_* \sqrt{\frac{\mu_w}{2}} \|\theta_{s+1} - \theta_s\|, \end{aligned} \quad (\text{C.6})$$

where the second equality holds by Lemma B.3, and the final inequality holds by Lemma B.4, together with Assumption 1. Note that the optimality condition of mirror descent update gives us

$$\nabla L(\theta_s) + \frac{1}{2\eta} (\nabla w(\theta_{s+1}) - \nabla w(\theta_s)) = 0.$$

By the definition of $(u_{\|\cdot\|_*}, \gamma_{\|\cdot\|_*})$, we have

$$\begin{aligned} \|\nabla w(\theta_{s+1}) - \nabla w(\theta_s)\|_* &= \|\nabla w(\theta_{s+1}) - \nabla w(\theta_s)\|_* \|u_{\|\cdot\|_*}\| \\ &\geq \langle \nabla w(\theta_{s+1}) - \nabla w(\theta_s), u_{\|\cdot\|_*} \rangle \\ &= 2\eta \langle -\nabla L(\theta_s), u_{\|\cdot\|_*} \rangle \\ &= \eta \frac{2}{n} \sum_{i=1}^n \exp(-\langle \theta_s, y_i x_i \rangle) \langle y_i x_i, u_{\|\cdot\|_*} \rangle \\ &\geq 2\eta L(\theta_s) \gamma_{\|\cdot\|_*}. \end{aligned} \quad (\text{C.7})$$

Thus combining (C.6) and (C.7), we conclude that

$$D_w(\theta_s, \theta_{s+1}) \geq \eta \sqrt{\frac{\mu_w}{L_w}} L(\theta_s) \gamma_{\|\cdot\|_*} \|\theta_{s+1} - \theta_s\|. \quad (\text{C.8})$$

Follow the exact same lines of argument, we can also show

$$D_w(\theta_{s+1}, \theta_s) \geq \eta \sqrt{\frac{\mu_w}{L_w}} L(\theta_s) \gamma_{\|\cdot\|_*} \|\theta_{s+1} - \theta_s\|. \quad (\text{C.9})$$

Using definition $L(\theta_{t+1}) = \frac{1}{n} \sum_{i=1}^n \exp(-\langle \theta_{t+1}, y_i x_i \rangle)$, together with (C.5), (C.8) and (C.9), we obtain

$$\begin{aligned} &\frac{1}{n} \sum_{i=1}^n \exp(-\langle \theta_{t+1}, y_i x_i \rangle) \\ &\leq L(\theta_0) \exp \left\{ -\sum_{s=0}^t \frac{1}{2\eta L(\theta_s)} D_w(\theta_s, \theta_{s+1}) - \sum_{s=0}^t \frac{1}{L(\theta_s)} \left(\frac{1}{2\eta} - \frac{L(\theta_s)}{\mu_2} D_{\|\cdot\|_2} \right) D_w(\theta_{s+1}, \theta_s) \right\} \\ &\leq L(\theta_0) \left\{ -\sum_{s=0}^t \frac{1}{2} \sqrt{\frac{\mu_w}{L_w}} \|\theta_{s+1} - \theta_s\| \gamma_{\|\cdot\|_*} - \sum_{s=0}^t \left(\frac{1}{2} - \frac{L(\theta_s)}{\mu_2} \eta D_{\|\cdot\|_2} \right) \sqrt{\frac{\mu_w}{L_w}} \gamma_{\|\cdot\|_*} \|\theta_{s+1} - \theta_s\| \right\}. \end{aligned}$$

Taking logarithm and dividing by $\|\theta_{t+1}\|$ on both sides, we have for any $i \in [n]$,

$$\begin{aligned}
& \left\langle \frac{\theta_{t+1}}{\|\theta_{t+1}\|}, y_i x_i \right\rangle \\
& \geq \frac{-\log n - \log L(\theta_0)}{\|\theta_{t+1}\|} + \frac{1}{2} \sqrt{\frac{\mu_w}{L_w}} \frac{\sum_{s=0}^t \|\theta_{s+1} - \theta_s\| \gamma_{\|\cdot\|_*}}{\|\theta_{t+1}\|} \\
& \quad + \sum_{s=0}^t \left(\frac{1}{2} - \frac{L(\theta_s)}{\mu_2} \eta D_{\|\cdot\|_2} \right) \sqrt{\frac{\mu_w}{L_w}} \gamma_{\|\cdot\|_*} \|\theta_{s+1} - \theta_s\| \\
& \geq \underbrace{\frac{-\log n - \log L(\theta_0)}{\|\theta_{t+1}\|}}_{(a)} + \underbrace{\frac{1}{2} \sqrt{\frac{\mu_w}{L_w}} \frac{\sum_{s=0}^t \|\theta_{s+1} - \theta_s\| \gamma_{\|\cdot\|_*}}{\sum_{s=0}^t \|\theta_{s+1} - \theta_s\|}}_{(b)} \\
& \quad + \underbrace{\sum_{s=0}^t \left(\frac{1}{2} - \frac{L(\theta_s)}{\mu_2} \eta D_{\|\cdot\|_2} \right) \sqrt{\frac{\mu_w}{L_w}} \gamma_{\|\cdot\|_*} \|\theta_{s+1} - \theta_s\|}_{(c)} \bigg/ \sum_{s=0}^t \|\theta_{s+1} - \theta_s\|. \tag{C.10}
\end{aligned}$$

To get asymptotic convergence, observe that by (C.3), $L(\theta_t) \rightarrow 0$. Thus we have

$$\|\theta_{t+1}\| \rightarrow \infty, \quad \sum_{s=0}^t \|\theta_{s+1} - \theta_s\| \geq \|\theta_{t+1}\| \rightarrow \infty, \quad \frac{1}{2} - \frac{L(\theta_s)}{\mu_2} \eta D_{\|\cdot\|_2} \rightarrow \frac{1}{2},$$

from which we conclude

$$(a) \rightarrow 0, \quad (b) = \frac{1}{2} \sqrt{\frac{\mu_w}{L_w}} \gamma_{\|\cdot\|_*}, \quad (c) \rightarrow \frac{1}{2} \sqrt{\frac{\mu_w}{L_w}} \gamma_{\|\cdot\|_*},$$

where we use the simple fact that $\lim_{t \rightarrow \infty} \frac{\sum_{s=0}^t a_s b_s}{\sum_{s=0}^t b_s} = a$ if $\lim_{t \rightarrow \infty} \sum_{s=0}^t b_s = \infty$ and $\lim_{t \rightarrow \infty} a_s = a$. Thus we obtain

$$\lim_{t \rightarrow \infty} \left\langle \frac{\theta_{t+1}}{\|\theta_{t+1}\|}, y_i x_i \right\rangle \geq \sqrt{\frac{\mu_w}{L_w}} \gamma_{\|\cdot\|_*}.$$

To get detailed convergence rate, we proceed to treat term (a) and (c) in (C.10) separately. In order to make $\left\langle \frac{\theta_{t+1}}{\|\theta_{t+1}\|}, y_i x_i \right\rangle \geq (1 - \epsilon) \sqrt{\frac{\mu_w}{L_w}} \gamma_{\|\cdot\|_*}$, it suffices to make

$$(a) \leq \sqrt{\frac{\mu_w}{L_w}} \frac{\epsilon \gamma_{\|\cdot\|_*}}{2}, \quad (c) \geq \left(1 - \frac{\epsilon}{2}\right) \sqrt{\frac{\mu_w}{L_w}} \gamma_{\|\cdot\|_*}. \tag{C.11}$$

To bound term (a), observe that from the loss convergence (C.3), together with assumption that $\|x_i\|_* \leq D_{\|\cdot\|_*}$,

$$\exp(-D \|\theta_{t+1}\|) \leq L(\theta_t) \leq \frac{1}{\gamma_{\|\cdot\|_*} \eta t} + \frac{L_w^2 \log^2(\gamma_{\|\cdot\|_*} \eta t)}{4\gamma_{\|\cdot\|_*}^2 \eta t} = \mathcal{O} \left\{ \frac{L_w \log^2(\gamma_{\|\cdot\|_*} \eta t)}{4\gamma_{\|\cdot\|_*}^2 \eta t} \right\}.$$

Thus we have the following lower bound,

$$\sum_{s=0}^t \|\theta_{s+1} - \theta_s\| \geq \|\theta_{t+1}\| \geq \frac{1}{D_{\|\cdot\|_*}} \left\{ \log \left(4\gamma_{\|\cdot\|_*}^2 \eta t / L_w \right) - 2 \log \log(\gamma_{\|\cdot\|_*} \eta t) \right\}. \tag{C.12}$$

One can readily verify that by taking

$$t = \Theta \left\{ \exp \left(\frac{D_{\|\cdot\|_*} \log n}{\epsilon} \gamma_{\|\cdot\|_*} \right) \right\},$$

we have $(a) \leq \sqrt{\frac{\mu_w}{L_w}} \frac{\epsilon \gamma_{\|\cdot\|_*}}{2}$.

To bound (c), it suffices to make

$$\frac{\eta D_{\|\cdot\|_2}}{\mu_2} \frac{\sum_{s=0}^t \|\theta_{s+1} - \theta_s\| L(\theta_s)}{\sum_{s=0}^t \|\theta_{s+1} - \theta_s\|} \leq \frac{\epsilon}{2}.$$

Note that for any $0 \leq \underline{t} \leq t-1$, we have

$$\begin{aligned} \frac{\sum_{s=0}^{\underline{t}} \|\theta_{s+1} - \theta_s\| L(\theta_s)}{\sum_{s=0}^{\underline{t}} \|\theta_{s+1} - \theta_s\|} &\leq \frac{\sum_{s=0}^{\underline{t}} \|\theta_{s+1} - \theta_s\| L(\theta_s)}{\sum_{s=0}^{\underline{t}} \|\theta_{s+1} - \theta_s\|} + \frac{\sum_{s=\underline{t}+1}^t \|\theta_{s+1} - \theta_s\| L(\theta_s)}{\sum_{s=0}^{\underline{t}} \|\theta_{s+1} - \theta_s\|} \\ &\leq \frac{\sum_{s=0}^{\underline{t}} \|\theta_{s+1} - \theta_s\|}{\sum_{s=0}^{\underline{t}} \|\theta_{s+1} - \theta_s\|} + \frac{\sum_{s=\underline{t}+1}^t \|\theta_{s+1} - \theta_s\| L(\theta_s)}{\sum_{s=0}^{\underline{t}} \|\theta_{s+1} - \theta_s\|} \\ &\leq \frac{\sum_{s=0}^{\underline{t}} \|\theta_{s+1} - \theta_s\|}{\sum_{s=0}^{\underline{t}} \|\theta_{s+1} - \theta_s\|} + L(\theta_{\underline{t}}), \end{aligned}$$

where we use the fact the loss is monotonically improving, thus $L(\theta_s) \leq L(\theta_0) = 1$ for all $s \geq 0$, and $L(\theta_s) \leq L(\theta_{\underline{t}})$ for all $s \geq \underline{t}$.

On the other hand, taking $\theta = \theta_t$ in (C.2), we have $L(\theta_{t+1}) \leq L(\theta_t) - \frac{1}{2\eta} D_w(\theta_t, \theta_{t+1})$. Since $w(\cdot)$ is μ_w -strongly convex w.r.t. $\|\cdot\|$ -norm, and $L(\theta_t) \geq 0$, we obtain

$$\|\theta_{t+1} - \theta_t\| \leq \sqrt{\frac{2}{\mu_w} (L(\theta_t) - L(\theta_{t+1}))} \leq \sqrt{\frac{2}{\mu_w} L(\theta_t)}.$$

Combining with loss convergence in (C.3), we have

$$\|\theta_{s+1} - \theta_s\| \leq \sqrt{\frac{2}{s\mu_w\gamma_{\|\cdot\|_*}} + \frac{L_w \log^2(\gamma_{\|\cdot\|_*} \eta s)}{2\mu_w\gamma_{\|\cdot\|_*}^2 s}},$$

Thus we obtain the following upper bound on $\sum_{s=0}^{\underline{t}} \|\theta_{s+1} - \theta_s\|$ as

$$\begin{aligned} \sum_{s=0}^{\underline{t}} \|\theta_{s+1} - \theta_s\| &\leq \sqrt{\frac{2\underline{t}}{\mu_w\gamma_{\|\cdot\|_*}}} + \frac{1}{\gamma_{\|\cdot\|_*}} \sqrt{\frac{L_w}{\mu_w}} \left[\sqrt{t_o 2 \log^2(\gamma_{\|\cdot\|_*} \eta)} + \sqrt{\log^2(\underline{t})\underline{t}} \right] \\ &= \mathcal{O} \left(\frac{\sqrt{\underline{t} \log^2(\underline{t})}}{\gamma_{\|\cdot\|_*}} \sqrt{\frac{L_w}{\mu_w}} \right). \end{aligned}$$

Together with (C.12), we obtain

$$\begin{aligned} \frac{\sum_{s=0}^{\underline{t}} \|\theta_{s+1} - \theta_s\|}{\sum_{k=0}^{\underline{t}} \|\theta_{k+1} - \theta_k\|} &\leq D_{\|\cdot\|_*} \frac{\sqrt{\frac{2\underline{t}}{\mu_w\gamma_{\|\cdot\|_*}}} + \frac{1}{\gamma_{\|\cdot\|_*}} \sqrt{\frac{L_w}{\mu_w}} \left[\sqrt{t_o 2 \log^2(\gamma_{\|\cdot\|_*} \eta)} + \sqrt{\log^2(\underline{t})\underline{t}} \right]}{\log(2\gamma_{\|\cdot\|_*}^2 \eta \underline{t}) - 2 \log \log(\gamma_{\|\cdot\|_*} \eta \underline{t})} \\ &= \mathcal{O} \left(\frac{D_{\|\cdot\|_*} \sqrt{\underline{t} \log^2(\underline{t})}}{\gamma_{\|\cdot\|_*} \log(\gamma_{\|\cdot\|_*}^2 \eta \underline{t})} \sqrt{\frac{L_w}{\mu_w}} \right). \end{aligned}$$

Thus we can choose $t = \Theta \left\{ \exp \left(\frac{D_{\|\cdot\|_*} \sqrt{\underline{t} \log \underline{t} D_{\|\cdot\|_2} \eta}{\gamma_{\|\cdot\|_*} \epsilon \mu_2} \sqrt{\frac{L_w}{\mu_w}} \right) / \gamma_{\|\cdot\|_*}^2 \eta \right\}$, so that

$$\frac{\eta D_{\|\cdot\|_2} \sum_{s=0}^{\underline{t}} \|\theta_{s+1} - \theta_s\|}{\mu_2 \sum_{k=0}^{\underline{t}} \|\theta_{k+1} - \theta_k\|} \leq \frac{\epsilon}{4}. \quad (\text{C.13})$$

In addition, by the loss convergence (C.3), we can choose $\underline{t} = \tilde{\Theta} \left(\frac{L_w D_{\|\cdot\|_2}}{\mu_2 \gamma_{\|\cdot\|_*}^2 \epsilon} \right)$, so that $\frac{\eta D_{\|\cdot\|_2}}{\mu_2} L(\underline{t}) \leq \frac{\epsilon}{4}$.

Thus in conclusion, we can choose

$$t = \tilde{\Theta} \left(\exp \left(\frac{D_{\|\cdot\|_*}^{3/2} D_{\|\cdot\|_2} L_w \eta}{\gamma_{\|\cdot\|_*}^2 \mu_w^{1/2} \mu_2^{3/2} \epsilon^{3/2}} \log \left(\frac{1}{\epsilon} \right) \right) \right),$$

so that $\frac{\eta D_{\|\cdot\|_2}}{\mu_2} \frac{\sum_{s=0}^{\underline{t}} \|\theta_{s+1} - \theta_s\|}{\sum_{k=0}^{\underline{t}} \|\theta_{k+1} - \theta_k\|} \leq \frac{\epsilon}{4}$.

Finally, we conclude that there exists

$$t_0 = \max \left\{ \Theta \left(\exp \left(\frac{D_{\|\cdot\|_*} \log n}{\epsilon} \gamma_{\|\cdot\|_*} \right) \right), \Theta \left(\exp \left(\frac{D_{\|\cdot\|_*}^{3/2} D_{\|\cdot\|_2} L_w \eta}{\gamma_{\|\cdot\|_*}^2 \mu_w^{1/2} \mu_2^{3/2} \epsilon^{3/2}} \log \left(\frac{1}{\epsilon} \right) \right) \right) \right\},$$

such that for any $t \geq t_0$, we have (C.11) holds, which implies

$$\left\langle \frac{\theta_{t+1}}{\|\theta_{t+1}\|}, y_i x_i \right\rangle \geq (1 - \epsilon) \sqrt{\frac{\mu_w}{L_w}} \gamma_{\|\cdot\|_*}, \quad \forall i \in [n].$$

□

□

♠ *Proof of Theorem 4.2.*

◇ *Proof of Theorem 4.2-(1):* Following similar lines as we obtain (C.10), we can show that for any $i \in [n]$,

$$\begin{aligned} \left\langle \frac{\theta_{t+1}}{\|\theta_{t+1}\|}, y_i x_i \right\rangle &\geq \frac{-\log n - \log L(\theta_0)}{\|\theta_{t+1}\|} + \frac{1}{2} \sqrt{\frac{\mu_w}{L_w}} \frac{\sum_{s=0}^t \|\theta_{s+1} - \theta_s\| \gamma_{\|\cdot\|_*}}{\|\theta_{t+1}\|} \\ &\quad + \sum_{s=0}^t \left(\frac{1}{2} - \frac{L(\theta_s)}{\mu_2} \eta D_{\|\cdot\|_2} \right) \sqrt{\frac{\mu_w}{L_w}} \gamma_{\|\cdot\|_*} \|\theta_{s+1} - \theta_s\| \\ &\geq \frac{-\log n - \log L(\theta_0)}{\|\theta_{t+1}\|} + \frac{1}{2} \sqrt{\frac{\mu_w}{L_w}} \frac{\sum_{s=0}^t \|\theta_{s+1} - \theta_s\| \gamma_{\|\cdot\|_*}}{\sum_{s=0}^t \|\theta_{s+1} - \theta_s\|} \\ &\quad + \sum_{s=0}^t \left(\frac{1}{2} - \frac{L(\theta_s)}{\mu_2} \eta_s D_{\|\cdot\|_2} \right) \sqrt{\frac{\mu_w}{L_w}} \gamma_{\|\cdot\|_*} \|\theta_{s+1} - \theta_s\| \Big/ \sum_{s=0}^t \|\theta_{s+1} - \theta_s\| \\ &\geq \sqrt{\frac{\mu_w}{L_w}} \gamma_{\|\cdot\|_*} + \frac{-\log n - \log L(\theta_0)}{\sum_{s=0}^t \|\theta_{s+1} - \theta_s\|} \\ &\quad - \sqrt{\frac{\mu_w}{L_w}} \gamma_{\|\cdot\|_*} \sum_{s=0}^t \frac{L(\theta_s)}{\mu_2} \eta_s D_{\|\cdot\|_2} \|\theta_{s+1} - \theta_s\| \Big/ \sum_{s=0}^t \|\theta_{s+1} - \theta_s\|. \quad (\text{C.14}) \end{aligned}$$

In addition, we have

$$\begin{aligned} \frac{L_w}{2} \|\theta_{s+1} - \theta_s\|^2 &\geq D_w(\theta_s, \theta_{s+1}) = D_{w^*}(\nabla w(\theta_{s+1}), \nabla w(\theta_s)) \\ &\geq \frac{2\eta_s^2}{L_w} L^2(\theta_s) \gamma_{\|\cdot\|_*}^2 = \frac{2\alpha_s^2}{L_w} \gamma_{\|\cdot\|_*}^2, \quad (\text{C.15}) \end{aligned}$$

where the first inequality follows from $w(\cdot)$ being L_w -smooth w.r.t. $\|\cdot\|$ -norm, the second equality follows from Lemma B.3, and the second inequality follows from (C.7), and the final equality follows from the definition of stepsize $\eta_s = \frac{\alpha_s}{L(\theta_s)}$. Thus we obtain $\|\theta_{s+1} - \theta_s\| \geq \frac{2\alpha_s \gamma_{\|\cdot\|_*}}{L_w}$. From which we conclude

$$\sum_{s=0}^t \|\theta_{s+1} - \theta_s\| \geq \sum_{s=0}^t \frac{2\gamma_{\|\cdot\|_*}}{L_w} \alpha_s = \Omega \left(\frac{2\gamma_{\|\cdot\|_*}}{L_w} \sqrt{t} \right). \quad (\text{C.16})$$

On the other hand, whenever $\eta_s \leq \frac{\mu_2}{2D_{\|\cdot\|_2} L(\theta_s)}$, or equivalently, $\alpha_s \leq \frac{\mu_2}{2D_{\|\cdot\|_2}}$, we have

$$\frac{\mu_w}{2} \|\theta_{s+1} - \theta_s\|^2 \leq D_w(\theta_s, \theta_{s+1}) \leq 2\eta_s \{L(\theta_s) - L(\theta_{s+1})\} \leq 2\eta_s L(\theta_s) = 2\alpha_s,$$

where the first inequality follows from $w(\cdot)$ being μ_w -strongly convex w.r.t. $\|\cdot\|$ -norm, the second inequality comes from (C.4), and the last inequality follows from $L(\theta_s) \geq 0$ and the definition of α_s .

Hence we have $\|\theta_s - \theta_{s+1}\| \leq 2\sqrt{\frac{\alpha_s}{\mu_w}}$.

Thus we have

$$\begin{aligned} & \sum_{s=0}^t \frac{L(\theta_s)}{\mu_2} \eta_s D_{\|\cdot\|_2} \|\theta_{s+1} - \theta_s\| \bigg/ \sum_{s=0}^t \|\theta_{s+1} - \theta_s\| \\ & \leq \frac{D_{\|\cdot\|_2} L}{\gamma_{\|\cdot\|_*} \mu_2 \sqrt{\mu_w}} \frac{\sum_{s=0}^t \alpha_s^{3/2}}{\sum_{s=0}^t \alpha_s} = \Omega \left(\frac{D_{\|\cdot\|_2} L}{\gamma_{\|\cdot\|_*} \mu_2 \sqrt{\mu_w}} t^{-1/4} \right). \end{aligned} \quad (\text{C.17})$$

Thus combining (C.14), (C.16), and (C.17), we conclude that there exists t_0 with $t_0 = \max \left\{ \Theta \left(\frac{\log n L}{\gamma_{\|\cdot\|_*} \epsilon} \right)^2, \Theta \left(\frac{D_{\|\cdot\|_2} L_w}{\gamma_{\|\cdot\|_*} \mu_2 \sqrt{\mu_w} \epsilon} \right)^4 \right\}$, such that for $t \geq t_0$, we have

$$\begin{aligned} \left\langle \frac{\theta_{t+1}}{\|\theta_{t+1}\|}, y_i x_i \right\rangle & \geq \sqrt{\frac{\mu_w}{L_w}} \gamma_{\|\cdot\|_*} + \frac{-\log n - \log L(\theta_0)}{\sum_{s=0}^t \|\theta_{s+1} - \theta_s\|} \\ & \quad - \sqrt{\frac{\mu_w}{L_w}} \gamma_{\|\cdot\|_*} \sum_{s=0}^t \frac{L(\theta_s)}{\mu_2} \eta_s D_{\|\cdot\|_2} \|\theta_{s+1} - \theta_s\| \bigg/ \sum_{s=0}^t \|\theta_{s+1} - \theta_s\| \\ & \geq (1 - \epsilon) \sqrt{\frac{\mu_w}{L_w}} \gamma_{\|\cdot\|_*}, \quad \forall i \in [n]. \end{aligned}$$

□

◇ *Proof of Theorem 4.2-(2)*: Follow the exact same line as we show (C.5), we conclude that with stepsize $\eta_s \leq \frac{\alpha_s}{L(\theta_s)} \leq \frac{\mu_2}{2D_{\|\cdot\|_2} L(\theta_s)}$, we have

$$\begin{aligned} L(\theta_{t+1}) & \leq L(\theta_0) \exp \left\{ - \sum_{s=0}^t \frac{1}{2\eta_s L(\theta_s)} D_w(\theta_s, \theta_{s+1}) - \sum_{s=0}^t \frac{1}{L(\theta_s)} \left(\frac{1}{2\eta_s} - \frac{L(\theta_s)}{\mu_2} D_{\|\cdot\|_2} \right) D_w(\theta_{s+1}, \theta_s) \right\} \\ & \leq L(\theta_0) \exp \left\{ - \sum_{s=0}^t \frac{1}{2\eta_s L(\theta_s)} D_w(\theta_s, \theta_{s+1}) \right\} = L(\theta_0) \exp \left\{ - \sum_{s=0}^t \frac{1}{2\alpha_s} D_w(\theta_s, \theta_{s+1}) \right\}. \end{aligned}$$

In addition, from (C.15), we have

$$D_w(\theta_s, \theta_{s+1}) \geq \frac{2\alpha_s^2}{L_w} \gamma_{\|\cdot\|_*}^2.$$

Thus we conclude that

$$L(\theta_{t+1}) \leq L(\theta_0) \exp \left\{ - \frac{\gamma_{\|\cdot\|_*}^2}{L_w} \sum_{s=0}^t \alpha_s \right\} = \mathcal{O} \left(\exp \left(- \frac{\gamma_{\|\cdot\|_*}^2}{L_w} \sqrt{t} \right) \right).$$

□

□