

---

# On Generalization of Decentralized Learning with Separable Data

---

Hossein Taheri

University of California, Santa Barbara

Christos Thrampoulidis

University of British Columbia

## Abstract

Decentralized learning offers privacy and communication efficiency when data are naturally distributed among agents communicating over an underlying graph. Motivated by overparameterized learning settings, in which models are trained to zero training loss, we study algorithmic and generalization properties of decentralized learning with gradient descent on separable data. Specifically, for decentralized gradient descent (DGD) and a variety of loss functions that asymptote to zero at infinity (including exponential and logistic losses), we derive novel finite-time generalization bounds. This complements a long line of recent work that studies the generalization performance and the implicit bias of gradient descent over separable data, but has thus far been limited to centralized learning scenarios. Notably, our generalization bounds approximately match in order their centralized counterparts. Critical behind this, and of independent interest, is establishing novel bounds on the training loss and the rate-of-consensus of DGD for a class of self-bounded losses. Finally, on the algorithmic front, we design improved gradient-based routines for decentralized learning with separable data and empirically demonstrate orders-of-magnitude of speed-up in terms of both training and generalization performance.

## 1 INTRODUCTION

### 1.1 Motivation

Machine learning tasks often revolve around inference from data using empirical risk minimization (ERM):

$$\min_{w \in \mathbb{R}^d} \hat{F}(w) := \frac{1}{n} \sum_{i=1}^n f(w, x_i). \quad (1)$$

Here  $f : \mathbb{R}^d \times \mathbb{R}^{d'} \rightarrow \mathbb{R}$  is a loss function and  $x_i := y_i a_i$ , where  $(a_i, y_i)_{i=1}^n \stackrel{\text{iid}}{\sim} \mathcal{D}$  represent features and labels, sampled from a distribution  $\mathcal{D}$ . In large scale machine learning, due to privacy concerns and communication constraints, data points are often distributed on a set of local computing agents. Decentralized learning methods aim at minimizing the global loss function (1) while agents communicate their parameters on an underlying connected graph. The most ubiquitous of these algorithms is Decentralized Gradient Descent (DGD). Here the  $\ell$ th agent runs a step of gradient descent followed by an averaging step in which every agent replaces its parameter with the average of its neighbors [Nedic and Ozdaglar, 2009]:

$$w_\ell^{(t+1)} = \sum_{k \in \mathcal{N}_\ell} A_{\ell k} w_k^{(t)} - \eta_t \nabla \hat{F}_\ell(w_\ell^{(t)}). \quad (2)$$

The superscripts signify the iteration number and  $A_{\ell k}$  refers to the averaging weights used by agent  $\ell$  for the parameter of agent  $k \in \mathcal{N}_\ell$  where  $\mathcal{N}_\ell$  is the set of neighbors of agent  $\ell$ . The global loss  $\hat{F}$  is the average of local loss functions  $\hat{F}_\ell$ ,  $\ell \leq N$ , where each  $\hat{F}_\ell$  is formed as the average empirical risk evaluated on the local training dataset  $\mathcal{S}_\ell$  of the  $\ell$ th agent:

$$\hat{F}(w) = \frac{1}{N} \sum_{\ell=1}^N \hat{F}_\ell(w), \quad \hat{F}_\ell(w) = \frac{1}{n_\ell} \sum_{x_j \in \mathcal{S}_\ell} f(w, x_j), \quad (3)$$

where  $n_\ell$  denotes the dataset size of agent  $\ell$ . Convergence properties of the train loss  $\hat{F}(\cdot)$  in DGD have been studied extensively in literature, e.g., [Nedic and Ozdaglar, 2009, Nedić and Olshevsky, 2014, Yuan et al., 2016, Lian et al., 2017, Nedić and Olshevsky, 2016]. The bulk of these studies

build upon classical optimization theory [Nesterov, 2003] suited for studying the train loss per iteration. In particular, it is well-established in the literature that DGD converges at the rate  $\frac{1}{T} \sum_{t=1}^T \hat{F}(\bar{w}^{(t)}) - \hat{F}^* = O(\frac{1}{\sqrt{T}})$  for smooth convex functions [Nedić and Olshevsky, 2014]. Here  $\bar{w}^{(t)}$  is the average of local parameters  $w_\ell^{(t)}$ . Our results in Sections 2.1-2.2 show a rate of  $\hat{F}(\bar{w}^{(T)}) = O(\frac{(\log T)^2}{T})$  and  $\|W^{(T)} - \bar{W}^{(T)}\|_F^2 = O(\frac{(\log T)^4}{T^2})$  for the training loss and consensus error of DGD over separable data with “exponentially tailed” losses.

The study of generalization performance of DGD algorithms in the literature is mostly limited to empirical observations e.g., [Jiang et al., 2017, Wang et al., 2019, Koloskova et al., 2019], making the theory behind test error performance largely unexplored. Moreover, the traditional wisdom in convergence analysis of DGD algorithms assumes the existence of a finite norm minimizer, which is often the case for ERM with non-separable training data, e.g. [Koloskova et al., 2020]. However, modern machine learning models operate in over-parameterized settings where the model perfectly interpolates the training data, i.e., it achieves perfect accuracy on the training data [Zhang et al., 2021]. Understanding the challenges imposed by over-parameterization and the behavior of gradient descent on separable data has been the subject of several recent works [Soudry et al., 2018, Ji and Telgarsky, 2018, Arora et al., 2019, Nacson et al., 2019, Chizat and Bach, 2020, Shamir, 2021, Ji et al., 2021, Ji and Telgarsky, 2021, Schliserman and Koren, 2022]. Yet, they are all focused on centralized GD, while here we study the impact of the consensus error of DGD on both training and generalization errors.

Our first goal is to complement prior general results on the convergence of training loss in DGD by considering specific, but commonly encountered, settings in ERM over separable data. This includes the analysis of non-smooth objectives such as the exponential loss, analysis of logistic regression in the separable regime where the optimum is achieved at infinity, and analysis of objectives satisfying the PL condition. The second goal is to study, for the first time in these settings, convergence rates of the DGD test loss  $F(\bar{w}^{(t)}) := \mathbb{E}_{x \sim \mathcal{D}}[f(\bar{w}^{(t)}, x)]$ . Finally, we leverage recent advances in the study of centralized learning with separable data to design fast algorithms for decentralized learning. We discuss our contributions below.

**Contributions.** In Sections 2.1 and 2.3, we derive convergence rates for the training and test loss of DGD over separable data. Our results hold for convex losses satisfying realizability and self-boundedness, as well as, convex

losses satisfying self-boundedness and the PL condition. In Section 2.2, we prove under additional self-boundedness assumptions on the Hessian and gradient, which hold for exponentially tailed losses, that the test loss bound can be improved to approximately match the test loss bounds of centralized GD. When specialized to decentralized logistic regression on separable data, our results provide the first generalization guarantees of DGD. In Section 2.4, we propose two algorithms for speeding up the convergence of decentralized learning under separable data. Numerical experiments demonstrate that our proposed algorithms significantly improve both the train test error of decentralized logistic regression.

## 1.2 Further related works

**Decentralized learning.** Over the last few years there have been numerous research works which consider the convergence of first order methods for decentralized learning; an incomplete list includes [Nedic and Ozdaglar, 2009, Nedić and Olshevsky, 2014, Yuan et al., 2016, Lian et al., 2017, Jiang et al., 2017, Assran et al., 2018, Pu et al., 2020, Koloskova et al., 2020, Kovalev et al., 2020, Xin et al., 2021, Toghiani and Uribe, 2022a, Toghiani and Uribe, 2022b]. While DGD is suboptimal for strongly-convex objectives [Nedić and Olshevsky, 2014, Nedić and Olshevsky, 2016], alternative algorithms, namely EXTRA and Grading Tracking, for achieving exponential rate appeared in [Shi et al., 2015, Nedic et al., 2017] and were studied further in [Koloskova et al., 2021, Xin et al., 2021]. More recently, [Lin et al., 2021] proposes accelerated methods for improving generalization and training accuracy of decentralized algorithms; however, their study of generalization error is empirical. While this paper was nearing completion we became aware of the recent works [Sun et al., 2021, Richards et al., 2020a] which study the generalization bounds of decentralized methods for Lipschitz convex losses (see also [Richards et al., 2020b, Sun et al., 2022]). However, we consider exponentially tailed losses under the separable data regime and prove faster convergence and generalization rates under these conditions. Compared to these works, we also propose improved algorithms for learning with separable data. Finally, we highlight that our rates on the train loss are comparable to [Koloskova et al., 2020, Theorem 2]. While [Koloskova et al., 2020] also derives convergence of DGD train loss on separable data, their analysis is valid only for bounded optimizers. In contrast, we derive training loss bounds which are true for the case of unbounded optimizers as is the case for logistic regression over separable data.

**Implicit bias of GD.** An early work on the behavior of ERM with vanishing regularization on separable data appeared in [Rosset et al., 2003]. Closely related, a line of recent works [Soudry et al., 2018, Ji and Telgarsky, 2018, Nacson et al., 2019, Ji and Telgarsky, 2020, Ji and Telgarsky, 2021, Shamir, 2021, Schliserman and Koren, 2022] studies the parameter convergence, as well as training and test loss convergence, of gradient descent on separable data, showing that for (a class of) monotonic losses the solution to ERM and the max-margin solution are the same in direction, i.e.,  $\|\hat{w}^{(t)} - \hat{w}_{\text{MM}}\| \rightarrow 0$ . Here  $\hat{w}^{(t)} := w^{(t)}/\|w^{(t)}\|$  and  $\hat{w}_{\text{MM}} := w_{\text{MM}}/\|w_{\text{MM}}\|$ , where the vector  $w_{\text{MM}}$  is the solution to the hard-margin support vector machine problem,

$$w_{\text{MM}} := \arg \min_{w \in \mathbb{R}^d} \|w\| \quad \text{s.t.} \quad y_i w^\top a_i \geq 1, \quad \forall i \in [n].$$

Notably, [Ji and Telgarsky, 2018, Soudry et al., 2018] characterized the rate of directional convergence to be  $\|\hat{w}^{(T)} - \hat{w}_{\text{MM}}\| = O(1/\log(T))$  and for the training loss to be  $\hat{F}(w^{(T)}) = O(\frac{1}{\eta T})$ . Recently, Shamir [Shamir, 2021] and Schliserman and Koren [Schliserman and Koren, 2022] showed that the test loss of GD for logistic regression on linearly separable data satisfies  $F(w^{(T)}) = \tilde{O}(\frac{1}{\eta T} + \frac{1}{n})$  signifying that overfitting does not happen during the iterates of GD. In Section 2.2 (Remark 5), we show that the test loss of DGD with logistic regression on linearly separable data satisfies  $\mathbb{E}[F(\bar{w}^{(T)})] = \tilde{O}(\frac{1}{\eta T} + \frac{1}{n} + \eta^2)$ , where the expectation is taken over training samples chosen i.i.d. from the dataset. As we explain, the term  $\eta^2$  captures the impact of consensus error (i.e., decentralization) on the generalization rate.

While directional convergence is significantly slow for gradient descent, following the update rule  $w^{(t+1)} = w^{(t)} - \eta_t \frac{\nabla \hat{F}(w^{(t)})}{\|\nabla \hat{F}(w^{(t)})\|}$ , it can be improved to  $1/\sqrt{t}$  with decaying  $\eta_t$  at rate  $1/\sqrt{t}$  for linear models [Nacson et al., 2019]. Furthermore [Taheri and Thrampoulidis, 2023] proved improved training convergence of this algorithm for two-layer neural networks, suggesting the benefits extend to non-linear settings. These results apply to centralized optimization scenarios. However, in decentralized learning settings, the local loss functions are kept private and any information about the global loss, such as its gradient  $\|\nabla \hat{F}(\bar{w}^{(t)})\|$  is hidden from the agents. In Section 2.4, we propose algorithms which address these challenges and extend the normalized GD update rule to decentralized learning scenarios. Furthermore, we prove the asymptotic convergence of normalized local parameters  $w_i^{(t)}/\|w_i^{(t)}\|$  to the solution of centralized GD.

**Notation** We use  $\|\cdot\|$  to denote the  $\ell_2$ -norm of vectors and the operator norm of matrices. The Frobenius norm of a matrix  $W$  is shown by  $\|W\|_F$ . The set  $\{i \in \mathbb{N} : i \leq N\}$  is

denoted by  $[N]$ . The gradient and hessian of a function  $F : \mathbb{R}^d \rightarrow \mathbb{R}$  are denoted by  $\nabla F(\cdot)$  and  $\nabla^2 F(\cdot)$ , respectively. For functions  $f, g : \mathbb{R} \rightarrow \mathbb{R}$ , we write  $f(t) = O(g(t))$  when  $|f(t)| \leq Mg(t)$  after  $t \geq t_0$  for positive constants  $M, t_0$ . Finally, we write  $f(t) = \tilde{O}(g(t))$  when  $f(t) = O(g(t)h(t))$  for a polylogarithmic function  $h$ .

## 2 MAIN RESULTS

Throughout the paper we make the following standard assumption on the mixing matrix  $A = [A_{ij}]_{N \times N}$  corresponding to the underlying connected network.

**Assumption 1** (Mixing matrix). *The mixing matrix  $A \in \mathbb{R}^{N \times N}$  is symmetric, doubly stochastic with bounded spectrum i.e.,  $|\lambda_i(A)| \in (0, 1]$  and  $\lambda_2(A) < 1$ .*

First, we state a lemma which relates the generalization loss of DGD at iteration  $t$  to its train loss and consensus error up to iteration  $t$ . The lemma is derived based on a stability analysis [Bousquet and Elisseeff, 2002, Hardt et al., 2016, Lei and Ying, 2020]. Specifically we use a self-boundedness and a realizability assumption [Schliserman and Koren, 2022] which makes the stability analysis feasible for settings such as logistic regression on separable data. Additionally, we assume convexity and  $L$ -smoothness of the loss function. Formally, we assume the following, where for simplicity, we use the short-hand  $f_x(w) := f(w, x)$  for the loss incurred at a generic  $x \in \mathcal{D}$  in the data distribution  $\mathcal{D}$ .

**Assumption 2** (Convexity). *The loss functions  $f_x : \mathbb{R}^d \rightarrow \mathbb{R}$  are convex and differentiable, satisfying,  $f_x(w) \leq f_x(v) + \langle \nabla f_x(w), w - v \rangle$ .*

**Assumption 3** (Smoothness). *The loss functions  $f_x : \mathbb{R}^d \rightarrow \mathbb{R}$  are  $L$ -smooth and differentiable, i.e.  $f_x(w) \leq f_x(v) + \langle \nabla f_x(v), w - v \rangle + \frac{L}{2} \|w - v\|^2$ .*

**Assumption 4** (Self-boundedness of the gradient). *The loss functions  $f_x : \mathbb{R}^d \rightarrow \mathbb{R}$  satisfy the self-boundedness property with the parameters  $c > 0$  and  $\alpha \in [\frac{1}{2}, 1]$ , i.e.,*

$$\|\nabla f_x(w)\| \leq c(f_x(w))^\alpha.$$

Assumption 4 is weaker than Assumption 3, since an  $L$ -smooth non-negative function  $f$  satisfies  $\|\nabla f(w)\|^2 \leq 2L(f(w) - f^*) \leq 2Lf(w)$ , where  $f^* := \inf_w f(w) \geq 0$ . However, we make use of the smoothness property whenever it suits the analysis, particularly to bound training loss.

Additionally, we make the following assumptions: All local parameters are initiated at zero i.e,  $w_\ell^{(1)} = 0$  for all  $\ell \leq N$ . We assume for simplicity of exposition, that each agent has access to  $n/N$  ( $n_\ell = n/N$ ) samples from the dataset. The general case can be treated with minor modifications. We

also assume that  $f_x(w) \geq 0$  for all  $w$  and the minimum of each loss is zero i.e.,  $f_i^* = 0$ .

Before our key lemma, we introduce a few necessary notations. We define matrix  $W^{(t)} \in \mathbb{R}^{N \times d}$  as the concatenation of all agents' parameters at iteration  $t$ , i.e.,  $W = [w_1^{(t)}, \dots, w_N^{(t)}]^\top$ . We also denote by  $\bar{w}^{(t)} := \frac{1}{N} \sum_{\ell=1}^N w_\ell^{(t)}$  the average of local parameters, and denote by  $\bar{W}^{(t)} = [\bar{w}^{(t)}, \dots, \bar{w}^{(t)}] \in \mathbb{R}^{N \times d}$  its concatenated matrix.

**Lemma 1** (Key lemma, Informal version). *Let Assumptions 1-4 hold. Consider the iterates of decentralized gradient descent in Eq.(2) with a fixed positive step-size  $\eta \leq \frac{2}{L}$ . Then, for the test loss  $F$  at iteration  $T \geq 1$ , it holds that*

$$\begin{aligned} \mathbb{E} [F(\bar{w}^{(T)})] &\lesssim \mathbb{E} [\hat{F}(\bar{w}^{(T)})] \\ &+ \frac{\eta^2 L^2 c^2 T^2}{n^{3-2\alpha}} \mathbb{E} \left[ \left( \frac{1}{T} \sum_{t=1}^T \hat{F}(\bar{w}^{(t)}) \right)^{2\alpha} \right] \\ &+ \frac{\eta^2 L^4}{N} \mathbb{E} \left[ \left( \sum_{t=1}^T \|W^{(t)} - \bar{W}^{(t)}\|_F \right)^2 \right], \quad (4) \end{aligned}$$

where the expectation is over the training set of  $n$  i.i.d samples.

The precise statement and the proof of Lemma 1 are deferred to Appendix A. Lemma 1 bounds the test loss with respect to the train loss and the consensus error. In the following sections, we show how Lemma 1 yields test loss bounds on DGD by establishing bounds on the train loss and consensus errors under different assumptions on the loss function.

It is worth remarking that Eq. (4) is in fact valid not only for DGD, but also for Decentralized Gradient Tracking (DGT). DGT is another popular algorithm for distributed learning that can accelerate train error convergence over DGD by modifying the update in Eq. (2) such that each agent keeps a running estimate of the global gradient [Nedic et al., 2017]. The reason why (4) continues to hold for DGD is that the proof of Lemma 1 only relies on the updates of the ‘‘averaged’’ parameter  $\bar{w}^{(t)} := \frac{1}{N} \sum_{\ell=1}^N w_\ell$  and that the update rule of  $\bar{w}^{(t)}$  for both DGD and DGT is derived as  $\bar{w}^{(t)} = \bar{w}^{(t-1)} - \frac{\eta}{N} \sum_{\ell=1}^N \nabla \hat{F}_\ell(w_\ell^{(t-1)})$ . Thus, starting with Eq.(4) one can also obtain test loss bounds of DGT after replacing appropriate bounds of DGT for the training loss and consensus error. We leave this to future work.

## 2.1 Convergence with general convex losses

The upper-bound in Eq.(4) shows how the consensus error and train loss of DGD affect the test loss.

The next lemma bounds the training loss and consensus error

of DGD for general convex losses. The proof is deferred to Appendix B.1

**Lemma 2** (Training bounds for convex losses). *Under Assumptions 1-3, for any  $w \in \mathbb{R}^d$  and for a fixed step-size*

$$\eta < \frac{1}{L} \min \left\{ 1 - \alpha_1, \sqrt{\frac{1 - \alpha_1}{2\alpha_2}} \right\},$$

where  $\alpha_1 \in (3/4, 1)$ ,  $\alpha_2 > 4$  are parameters that depend only on the mixing matrix, the train loss and consensus error of DGD (2) satisfy:

$$\frac{1}{T} \sum_{t=1}^T \hat{F}(\bar{w}^{(t)}) \leq \frac{2\|w\|^2}{\eta T} + 4\hat{F}(w), \quad (5)$$

$$\frac{1}{NT} \sum_{t=1}^T \|W^{(t)} - \bar{W}^{(t)}\|_F^2 \leq \frac{\alpha_2 \eta^2 L^2}{1 - \alpha_1} \left( \frac{2\|w\|^2}{\eta T} + 4\hat{F}(w) \right).$$

To bound the training loss for functions  $f(\cdot)$  where the optimum is attained at infinity we need a realizability assumption. In particular, we choose  $w \in \mathbb{R}^d$  (in Lemma 2) using the following.

**Assumption 5** (Realizability). *The loss functions  $f_x : \mathbb{R}^d \rightarrow \mathbb{R}$  satisfy the realizability condition, i.e.  $\exists$  decreasing function  $\rho : \mathbb{R}_+ \rightarrow \mathbb{R}_+$  such that for every  $\varepsilon > 0$  there exists  $\hat{w} \in \mathbb{R}^d$  with  $\|\hat{w}\| \leq \rho(\varepsilon)$  that satisfies  $f_x(\hat{w}) \leq \varepsilon$ .*

The set of Assumptions 2-5 covers classification over linearly separable data with logistic loss, in addition to losses with other exponential-type tails  $\exp(-w^r)$  and polynomial tail  $w^{-r}$ , for  $r > 0$ .

*Remark 1* (Training loss of DGD on separable data). The realizability assumption as stated appeared recently in [Schliserman and Koren, 2022] (and was implicitly used in [Ji and Telgarsky, 2018, Shamir, 2021]). It can be checked that for linearly separable training data with margin  $\gamma$ , loss functions with an exponential tail such as logistic loss satisfy this assumption with  $\rho(\varepsilon) = \frac{1}{\gamma} \log(\frac{1}{\varepsilon})$  (e.g., see Proposition 26 and [Schliserman and Koren, 2022, Lemma 4]). Based on Lemma 2, this leads to the following bound for DGD training loss for all  $\varepsilon > 0$ ,

$$\frac{1}{T} \sum_{t=1}^T \hat{F}(\bar{w}^{(t)}) \leq \frac{2 \log(1/\varepsilon)^2}{\gamma^2 \eta T} + 4\varepsilon. \quad (6)$$

In particular, choosing  $\varepsilon = 1/T$ , gives a rate of  $O(\frac{(\log T)^2}{\eta T})$ , surprisingly matching up to logarithmic factors the corresponding rate for centralized GD in [Ji and Telgarsky, 2018, Theorem 1.1].

*Remark 2.* The bounds of Lemma 2 are true for any dataset  $\{x_i\}_{i \in [n]}$  provided that Assumptions 2 and 3 hold for all  $f_x = f_{x_i} = f(w, x_i) := f_i(w)$ ,  $i \in [n]$ . Similarly, (6)



holds provided Assumption 5 is true over the training set (i.e. provided the training dataset is separable). However, bounding the test loss in Lemma 1, requires bounding the *expectation over all datasets* of the train/consensus errors. This is guaranteed by Assumptions 2-5 as they hold for any point  $x$  in the distribution.

**Theorem 3** (Test loss with convex losses). *Under Assumptions 1-5, by choosing*

$$\eta < \frac{1}{L\sqrt{T}} \min \left\{ 1 - \alpha_1, \sqrt{\frac{1 - \alpha_1}{2\alpha_2}} \right\}$$

where  $\alpha_1 \in (3/4, 1)$ ,  $\alpha_2 > 4$  are parameters that depend only on the mixing matrix and assuming  $\varepsilon \leq \frac{\rho(\varepsilon)^2}{\eta T}$ , the test error of DGD for iteration  $T \geq 1$  satisfies:

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E} \left[ F(\bar{w}^{(t)}) \right] = O \left( \frac{\rho(\varepsilon)^2}{\sqrt{T}} + \frac{L^2 c^2 \rho(\varepsilon)^{4\alpha}}{n^{3-2\alpha}} T^{1-\alpha} + \frac{L^4 \rho(\varepsilon)^2}{\sqrt{T}} \right), \quad (7)$$

where the expectation is over the training set of  $n$  i.i.d samples.

**Remark 3** (DGD with logistic regression never overfits). The proof of Theorem 3 is delayed to Appendix B.2. As in Remark 1, we take logistic regression on separable data with margin  $\gamma > 0$  as our case study. For logistic regression (as well as other loss functions with an exponential tail), it can be verified that the self-boundedness assumption holds with  $\alpha = 1$ . Similar to Remark 1 it holds that  $\rho(\varepsilon) = \frac{1}{\gamma} \log(\frac{1}{\varepsilon})$ , thus choosing  $\varepsilon = 1/\sqrt{T}$  results in a test loss rate  $\tilde{O}(\frac{1}{\sqrt{T}} + \frac{1}{n})$  by Eq.(7). This indicates that the upper-bound decreases at a rate of  $\tilde{O}(\frac{1}{\sqrt{T}})$  until after  $T = n^2 \cdot (\max(\frac{1}{Lc}, \frac{L}{c}))^4$  iterations where the upper bound essentially reduces to  $\tilde{O}(\frac{L^2 c^2}{n})$ . Additionally, the fact that the upper-bound is decreasing proves that with appropriate choice of step-size, overfitting never happens along the path of DGD at any iteration.

**Remark 4** (Log factors). The attentive reader will have recognized in Remarks 1 and 3 that due to the “ $\rho(\varepsilon) = \mathcal{O}(\log(T))$ ” factor, the upper bound on the test loss in Eq. (7) increases (very) slowly with  $\log^4(T)$ . Note that this term becomes dominant only when  $T$  is exponentially large with respect to the sample size  $n$  and the margin  $\gamma$ . Our experiments in Sec. 3.2 confirm this slow logarithmic increase late in the training phase. Analogous behavior, but for centralized GD training, are discussed in [Soudry et al., 2018, Schliserman and Koren, 2022].

## 2.2 On the convergence of DGD with exponentially-tailed losses

In this section, we show that our guarantees can be improved for exponentially tailed losses. First, we note that

the bounds in Lemma 2 and Theorem 3 hold for the average loss across iterations  $t \leq T$ . It is straight-forward to see that if DGD is a descent algorithm i.e.,  $\hat{F}(\bar{w}^{(t+1)}) \leq \hat{F}(\bar{w}^{(t)})$  for all  $t \leq T$ , then  $\hat{F}(\bar{w}^{(T)}) \leq \frac{1}{T} \sum_{t=1}^T \hat{F}(\bar{w}^{(t)})$ ; thus implying that the upper-bounds on training and test loss hold for the last iterate of DGD. We will prove that DGD is indeed a “descent algorithm” for a class of convex losses which include popular choices such as the logistic loss and even non-smooth choices including the exponential loss. Moreover, we show that the consensus error of Lemma 2 as well as the test loss bounds of Theorem 3 can be improved compared to the results of the previous section.

In particular, we use the following assumptions together with the self-boundedness gradient assumption (Assumption 4) with  $\alpha = 1$  as well as the convexity assumption.

**Assumption 6** (Self-bounded Hessian). *The local losses  $\hat{F}_\ell : \mathbb{R}^d \rightarrow \mathbb{R}$  satisfy the following for the Hessian matrices  $\nabla^2 \hat{F}_\ell$  and a positive constant  $h$ ,*

$$\|\nabla^2 \hat{F}_\ell(w)\| \leq h \hat{F}_\ell(w).$$

**Assumption 7** (Self-lowerbounded gradient). *The global loss satisfies for a constant  $\tau$  that*

$$\|\nabla \hat{F}(w)\| \geq \tau \hat{F}(w).$$

Assumptions 2, 4, 6 and 7 include linear classification with non-smooth losses such as the exponential loss, losses with super-exponential tails ( $\exp(-x^r)$ ,  $r > 1$ ) and the logistic loss; e.g., see Proposition 24 in the appendix.

**Theorem 4** (Last iterate convergence of DGD). *Consider DGD with the loss functions and mixing matrix satisfying Assumptions 1,2,6,7 and Assumption 4 with  $\alpha = 1$  and  $c = h$ . Assume that the step-size satisfies  $\eta < \frac{\delta}{\hat{F}(1)}$ , for a constant  $\delta$  depending only on the mixing matrix and on  $\tau, h$ , then DGD is a descent algorithm i.e, for all  $t \geq 1$  it holds that  $\hat{F}(\bar{w}^{(t+1)}) \leq \hat{F}(\bar{w}^{(t)})$ . Moreover, the train loss and the consensus error of DGD at iteration  $T$  satisfy the following for all  $w \in \mathbb{R}^d$ ,*

$$\hat{F}(\bar{w}^{(T)}) \leq 4\hat{F}(w) + \frac{2\|w\|^2}{\eta T},$$

$$\|W^{(T)} - \bar{W}^{(T)}\|_F^2 = O \left( h^2 \eta^2 \hat{F}^2(w) + \frac{h^2 \|w\|^4}{T^2} \right). \quad (8)$$

The proof of Theorem 4 is included in Appendix C.1. In the following remark, we discuss the implications of this result. **Remark 5** (Improved rates). While similar to Lemma 2, for logistic regression we have  $\hat{F}(\bar{w}^{(T)}) = \tilde{O}(\frac{1}{\eta T} + \frac{1}{T})$ , for the consensus error rate we have by applying Theorem 4 and noting that  $\rho(\varepsilon) = \log(1/\varepsilon)/\gamma$ ,

$$\|W^{(T)} - \bar{W}^{(T)}\|_F^2 = O \left( h^2 \eta^2 \varepsilon^2 + \frac{h^2 (\log(1/\varepsilon))^4}{\gamma^4 T^2} \right).$$

After choosing  $\varepsilon = 1/T$ , we have the improved rate  $\|W^{(T)} - \bar{W}^{(T)}\|_F^2 = \tilde{O}(\frac{1}{T^2})$ , which is superior over the rate  $\tilde{O}(\frac{1}{T})$  for general convex losses with constant  $\eta$  (Lemma 2). For the test loss, employing Lemma 1 with the new rates for the consensus error leads to the following rate for DGD with logistic regression,

$$\mathbb{E} \left[ F(\bar{w}^{(T)}) \right] = \tilde{O} \left( \frac{1}{\eta T} + \frac{1}{n} + \eta^2 \right). \quad (9)$$

In accordance to Remark 2, we can conclude the above from Lemma 1 provided Assumptions 4 and 7. Thus, the bounds of Theorem 4 remain true for all training sets within the data distribution. We note that the resulting bound in (9) is a superior rate for the test loss of logistic regression, compared to the rate of Remark 3. Concretely, setting  $\eta = 1/T^{1/3}$  gives a rate of  $\tilde{O}(1/T^{2/3} + 1/n)$ , faster than the  $\tilde{O}(1/\sqrt{T} + 1/n)$  rate in Remark 3. On the other hand, it is slightly slower compared to its centralized counterpart  $\tilde{O}(1/T + 1/n)$  in [Shamir, 2021, Schliserman and Koren, 2022]. As revealed by Lemma 1, the additional  $\eta^2$  factor in (9) captures impact of the consensus term, which is unavoidable in decentralized learning.

### 2.3 Convergence under the PL condition

Next, we show how our previous results change when the global loss satisfies the  $\mu$ -PL condition. Formally, the PL condition [Polyak, 1963, Lojasiewicz, 1963] is defined as follows.

**Assumption 8** (PL condition). *The loss function  $\hat{F} : \mathbb{R}^d \rightarrow \mathbb{R}$  satisfies the Polyak-Lojasiewicz (PL) condition with parameter  $\mu > 0$ :  $\|\nabla \hat{F}(w)\|^2 \geq 2\mu(\hat{F}(w) - \hat{F}^*)$ .*

The next lemma shows that DGD enjoys an exponential rate under the PL condition and smoothness. and data separability (i.e.,  $\hat{F}^* = 0$ ). See Appendix D.1 for a proof.

**Lemma 5** (Train loss under the PL condition). *Let Assumptions 1,3 and 8 hold and let the step-size  $\eta \leq \min\{\frac{1-\alpha_1}{\mu}, \frac{1}{2L^2} \sqrt{\frac{(1-\alpha_1)\mu}{\alpha_2}}, \frac{1}{L}\}$ , where the constants  $\alpha_1 \in (3/4, 1)$  and  $\alpha_2 > 4$  depend only on the mixing matrix. Define  $\zeta := 1 - \frac{\eta\mu}{2}$ , then under the data separability assumption, the iterates of DGD satisfy for all  $t \geq 1$ ,*

$$\begin{aligned} \hat{F}(\bar{w}^{(t)}) &\leq \zeta^{t-1} \hat{F}(\bar{w}^{(1)}), \\ \frac{1}{N} \left\| W^{(t)} - \bar{W}^{(t)} \right\|_F^2 &\leq \frac{2\alpha_2 \eta^2 L^2 \hat{F}(\bar{w}^{(1)})}{1 - \alpha_1} \zeta^{t-1}. \end{aligned}$$

We use this lemma combined with our key lemma 1 to obtain the test loss bound in the next theorem. The proof is provided in Appendix D.2.

**Theorem 6** (Test loss under the PL condition). *Let Assumptions 1-4 hold. Further assume 8 holds for all*

---

#### Algorithm 1: FDLR

---

**Input:** Initial values  $w_\ell^{(1)}, v_\ell^{(1)} \in \mathbb{R}^d$  for all agents  $\ell \in [N]$ , step size  $\eta_t$  and mixing matrix

$$A = [A_{\ell k}]_{N \times N}$$

```

1 for  $t = 1, \dots, T$  all agents  $\ell \in [N]$  in parallel do
2    $w_\ell^{(t+\frac{1}{2})} = w_\ell^{(t)} - \eta_t \frac{v_\ell^{(t)}}{\|v_\ell^{(t)}\|}$ 
3   send and receive local variables  $w_\ell^{(t+\frac{1}{2})}$  and  $v_\ell^{(t)}$ 
4    $w_\ell^{(t+1)} = \sum_{k \in \mathcal{N}_\ell} A_{\ell k} w_k^{(t+\frac{1}{2})}$ 
5    $v_\ell^{(t+1)} = \sum_{k \in \mathcal{N}_\ell} A_{\ell k} v_k^{(t)} + \nabla \hat{F}_\ell(w_\ell^{(t+1)}) - \nabla \hat{F}_\ell(w_\ell^{(t)})$ 

```

---

*training sets in the distribution. Let  $\eta$  and  $\zeta$  be as in Lemma 5. Then the iterates of DGD satisfy for all  $T \geq 1$ ,*

$$\mathbb{E} \left[ F(\bar{w}^{(T)}) \right] = O \left( \zeta^T + \frac{L^2 c^2}{n^{3-2\alpha} \mu^{2\alpha}} (\eta T)^{2-2\alpha} + \frac{\eta^2 L^4}{\mu^2} \right).$$

**Remark 6.** The bound above involves  $(\eta T)^{2-2\alpha}$ . When  $\alpha < 1$ , as in the case of smooth functions such as highly over-parameterized Least-squares  $f(w, x) = (1 - w^\top x)^2$  where  $d \gg n$ , the bound becomes vacuous as it is increasing with  $T$ . This suggests the existence of overfitting in DGD under such scenarios; with the optimal value of  $T$  achieved at the very early steps of training. See Appendix G.1 for experiments that confirm this behavior.

### 2.4 Improved algorithms for decentralized learning with separable data

In this section, we consider decentralized learning with exponentially tail losses on separable data and propose modifications to the DGD algorithm for improving the convergence rates based on the normalized GD mechanism.

Our first algorithm –Fast Distributed Logistic Regression (FDLR)– is summarized in Algorithm 1. Each agent keeps two local variables  $w_\ell, v_\ell \in \mathbb{R}^d$  which are also communicated to neighbor agents at each round. In matrix notation, Algorithm 1 has the following updates:

$$\begin{aligned} W^{(t+1)} &= A(W^{(t)} - \eta \tilde{V}^{(t)}), \\ V^{(t+1)} &= A V^{(t)} + \nabla \hat{F}(W^{(t+1)}) - \nabla \hat{F}(W^{(t)}). \end{aligned}$$

As in (2),  $A \in \mathbb{R}^{N \times N}$  is the mixing matrix of the undirected network of agents, which satisfies the regularity conditions in Assumption 1. Furthermore  $W^{(t)}, V^{(t)}, \nabla \hat{F}(W^{(t)}) \in \mathbb{R}^{N \times d}$  are formed by stacking  $w_\ell^{(t)}, v_\ell^{(t)}$  and local gradients  $\nabla \hat{F}_\ell(w_\ell^{(t)})$  for all  $\ell \in [N]$  as their rows. The matrix  $\tilde{V}^{(t)} \in \mathbb{R}^{N \times d}$  is formed by concatenation of the vectors  $v_\ell^{(t)} / \|v_\ell^{(t)}\|$  as its rows. In step (2) of Algorithm 1, every agent  $\ell$  runs in parallel an update

---

**Algorithm 2:** FDLR with Nesterov momentum
 

---

**Input:** Initial values  $w_\ell^{(1)}, v_\ell^{(1)}, z^{(1)} \in \mathbb{R}^d$  for all agents  $\ell \in [N]$ , hyper-parameters  $\eta_t, \gamma_t$  and mixing matrix  $A = [A_{\ell k}]_{N \times N}$

```

1 for  $t = 1, \dots, T$  all agents  $i \in [N]$  in parallel do
2    $z_\ell^{(t+1)} = \gamma_t(z_\ell^{(t)} + \frac{v_\ell^{(t)}}{\|v_\ell^{(t)}\|})$ 
3    $w_\ell^{(t+\frac{1}{2})} = w_\ell^{(t)} - \eta_t(z_\ell^{(t+1)} + \frac{v_\ell^{(t)}}{\|v_\ell^{(t)}\|})$ 
4   send and receive local variables  $w_\ell^{(t+\frac{1}{2})}$  and  $v_\ell^{(t)}$ 
5    $w_\ell^{(t+1)} = \sum_{k \in \mathcal{N}_\ell} A_{\ell k} w_k^{(t+\frac{1}{2})}$ 
6    $v_\ell^{(t+1)} = \sum_{k \in \mathcal{N}_\ell} A_{\ell k} v_k^{(t)} + \nabla \hat{F}_\ell(w_\ell^{(t+1)}) - \nabla \hat{F}_\ell(w_\ell^{(t)})$ 
    
```

---

rule which resembles the distributed gradient descent update rule (aka Eq. (2)), with the difference that the local gradient  $\nabla \hat{F}_\ell(w_\ell^{(t)})$  is replaced by  $v_\ell^{(t)} / \|v_\ell^{(t)}\|$ . In the next step, agents send their local parameters  $w_\ell^{(t)}, v_\ell^{(t)}$  to their neighbors. Step (4) is the consensus step at which agent  $\ell$  computes a weighted average of  $w_k^{(t+\frac{1}{2})}$  sent from neighbor agents  $k$ , in order to update  $w_\ell^{(t)}$ . Step (5) uses the newly computed local gradient  $\nabla \hat{F}_\ell(w_\ell^{(t+\frac{1}{2})})$  and the gradient computed in the previous step to update the local parameter  $w_\ell^{(t)}$ . The purpose behind introducing the variable  $v_\ell^{(t)}$  is to estimate the global gradient. This idea is previously used in the gradient tracking algorithm (e.g. see [Nedic et al., 2017]) and the idea also relates to stochastic variance reduced gradient (SVRG) [Johnson and Zhang, 2013]. The following theorem proves that for exponentially decaying loss functions and separable data, FDLR with time-decaying step-size  $\eta_t = 1/\sqrt{t}$  converges successfully in direction to the solution of centralized gradient descent. The proof is provided in Appendix E.

**Theorem 7** (Asymptotic convergence of FDLR). *Let the sequence  $\{w_\ell^{(t)}\}$  be generated by FDLR (Algorithm 1) trained with logistic or exponential loss on a separable dataset with  $\eta_t = O(1/\sqrt{t})$ . Then, for all  $\ell \in [N]$ ,  $\lim_{t \rightarrow \infty} w_\ell^{(t)} / \|w_\ell^{(t)}\| = w_{\text{MM}} / \|w_{\text{MM}}\|$ , where  $w_{\text{MM}}$  is the solution to max-margin problem.*

Based on the above result, we anticipate that FDLR has good test performance. In fact, we will show in Section 3 that FDLR achieves good test performance orders of magnitude faster than DGD. To get some insight on this and also on the nature of the FDLR updates consider the infinite time limit. In this limit, when the matrix  $A$  satisfies the mixing Assumption 1, it can be checked that  $V^{(\infty)} = \frac{1}{n} \mathbf{1} \mathbf{1}^\top \nabla \hat{F}(W^{(\infty)})$ . Hence, as  $t \rightarrow \infty$  the variables  $v_\ell^{(t)}$  for all agents converge to the same global

gradient  $\sum_{\ell=1}^N \nabla \hat{F}_\ell(w_\ell^{(t)})$ . Realizing this, we can see that Step (2) of FDLR is asymptotically approximating a normalized GD update, i.e., for large  $t$ , each agent performs an update  $w_\ell^{(t+1/2)} \approx w_\ell^{(t)} - \eta_t \frac{\sum_{\ell=1}^N \nabla \hat{F}_\ell(w_\ell^{(t)})}{\|\sum_{\ell=1}^N \nabla \hat{F}_\ell(w_\ell^{(t)})\|_2}$ . Previously, normalized gradient descent has been used to speed up convergence in centralized logistic regression over separable data [Nacson et al., 2019]. Here, we essentially extend this idea to a decentralized setting and argue that FDLR is the canonical way to do so. In particular, the idea of introducing additional variables  $v_\ell$  that keep track of the global gradient is critical for the algorithm's success. That is, a naive implementation with updates  $w_\ell^{(t+1/2)} = w_\ell^{(t)} - \eta_t \nabla \hat{F}_\ell(w_\ell^{(t)}) / \|\nabla \hat{F}_\ell(w_\ell^{(t)})\|_2$  based only on the local gradients would fail. At the other end, just introducing variables  $v_\ell$  without performing a normalized gradient update (i.e. implementing gradient tracking) also fails to give significant speed ups over DGD. See Section 3 for experiments in support of this claim.

We also present a yet improved Algorithm 2, which combines FDLR with Nesterov Momentum. The key innovation of Algorithm 2 compared to FDLR is its step (3), where now the local parameter  $w_\ell^{(t)}$  is updated by a weighted average ( $z_\ell^{(t+1)}$ ) of normalized gradients from previous iterations. Similar to our previous remarks regarding FDLR, extending the Nesterov accelerated variant of normalized GD for centralized logistic regression [Ji et al., 2021] to the distributed setting is more subtle as now each agent has access only to local gradients. Our experiments in Section 3 verify the correctness of the proposed implementation of Algorithm 2 as it achieves significant speed ups over both DGD and FDLR.

### 3 NUMERICAL EXPERIMENTS

In this section, we present numerical experiments to verify our theoretical results and demonstrate the benefits of our proposed algorithms. We begin with a numerical study of the performance of FDLR.

#### 3.1 Experiments on FDLR

In Fig. 1 (Left), we compare the performance of FDLR and its momentum variant to DGD and gradient tracking (GT) for exponential loss with signed measurements (i.e.,  $y = \text{sign}(a^\top w^*)$  for samples  $a$ , labels  $y$  and the true vector of regressors  $w^*$ ) with  $n = 100$ ,  $d = 25$ . The underlying graph is selected as an Erdos-Rényi graph with  $N = 50$  agents and connectivity probability  $p_c = 0.3$ . On the  $y$ -axis, directional convergence represents the distance of normalized  $w_\ell^{(t)}$  to the normalized final solution for agent  $\ell = 1$ , i.e.,  $\| \frac{w_1^{(t)}}{\|w_1^{(t)}\|} - \frac{w_{\text{MM}}}{\|w_{\text{MM}}\|} \|$  (see Theorem 7). The




Figure 1: Directional parameter convergence of our proposed Algorithms 1-2 compared to the vanilla distributed gradient descent and gradient tracking algorithms on (Left) synthetic data  $y = \text{sign}(a^\top w^*)$  and (Right) on two classes from the UCI WINE dataset.

hyper-parameters  $\eta_t, \gamma_t$  are fine-tuned for each algorithm to represent the best of each algorithm and the final values are  $\eta_t = 0.1, 0.05, 0.5$  and  $0.2$  for Distributed GD, GT, Alg. 1 and Alg. 2, respectively and  $\gamma_t = 0.8$  for Alg. 2. Our algorithms significantly outperform the well-known distributed learning algorithms in directional convergence to the final solution. Regardless, in this case we noticed that the gain obtained by including the momentum is small. In Fig. 1(Right), we consider a binary classification task on a real-world dataset (two classes of the UCI WINE dataset [win, ]) where  $d = 13$  and  $n = 107$ . We compare the performance of FDLR (blue line) and its momentum variant (red line) with DGD and DGT on an Erdos-Rényi graph with  $N = 10$  and  $p_c = 0.4$ . The hyper-parameters are fine-tuned to  $\eta_t = 12, 1, 0.9, 2$  for DGD, DGT, Alg. 1 and Alg. 2 respectively, and  $\gamma_t = 0.88$  for Alg. 2. Notably, while Alg. 1 significantly outperforms both DGD and DGT, the benefits of adding the momentum are also significant in this case as Alg. 2 demonstrates a faster rate of convergence than Alg. 1.

The two plots in Fig. 2 (Left) illustrate the train and test errors of DGD/DGT and our proposed algorithms for the same setting as Fig. 1(Left) with  $n = 800$  and  $d = 50$ . Here the hyper-parameters are fine-tuned to be  $\eta_t = 0.01, 0.01, 0.4, 0.5$  for DGD, DGT, Alg. 1 and Alg. 2, respectively and  $\gamma_t = 0.5$  for Alg. 2. Fig. 2(Right) shows the training and test losses. Here, we use the same dataset with  $N = 10, p_c = 0.4$  and an exponential loss. The hyper-parameters are fine-tuned to  $\eta_t = 0.01, 0.012, 0.4, 0.6$  and  $\gamma_t = 0.9$ . Both of our algorithms outperform the commonly used DGD and DT, in both training error and test error performance. Also, the gains of adding the momentum are significant, since FDLR with Nesterov momentum (Algorithm 2) reaches an approximation of its final test accuracy in 50 iterations, while the same happens for FDLR with approximately 300 iterations. An interesting phenomenon in Fig. 2 (see also

Fig. 3(Right)) is the behavior of test loss: while during the starting phase the test loss is monotonically decreasing, after sufficient iterations the test loss starts increasing. This behavior of test loss is indeed captured by the bounds on the test loss of DGD in Theorems 3-4 and Remarks 3-5. In particular, the increase in test loss is observed in the bound for the test loss  $O(\frac{(\log T)^2}{\sqrt{T}} + \frac{(\log T)^2}{n})$  in Remark 3, where the presence of the term  $\frac{(\log T)^2}{n}$  suggests that the bound after sufficient iterations starts to slowly increase. See also Remark 4.

### 3.2 Experiments on convergence of DGD

Next, we investigate the convergence behavior of DGD for the train loss and the consensus error. We consider the same network topology, mixing matrix and data setup as in the last figure. Left and middle plots in Fig. 3 show the consensus error  $\frac{1}{N} \|W^{(t)} - \bar{W}^{(t)}\|_F^2$  and the train loss  $\hat{F}(\bar{w}^{(t)})$  in solid lines, for two over-parameterization ratios  $d/n$ . We recall that  $d$  and  $n$  represent the dimension of the ambient space and the dataset size, respectively. The dashed lines help to show for each solid line, the approximate rate of convergence after sufficient number of iterations. Notably, we observe that the convergence rates on the consensus error ( $\tilde{O}(1/t^2)$ ) and on the train error ( $\tilde{O}(1/t)$ ) stated in Remark 5 are attained in both cases (recall that  $\tilde{O}(\cdot)$  hides logarithmic factors). Fig. 3 (Right) depicts the Test loss of DGD for  $d/n = 0.05$ . For comparison, the corresponding curve for centralized GD is also shown. Here step-sizes are fine-tuned to represent the best of each algorithm. In agreement with our findings in Remark 5, we observe approximately similar behavior between the convergence behavior of two algorithms. As before, the slight increase in the curves of test loss are due to the logarithmic factors in test loss upperbounds.






Figure 2: Training/test misclassification errors and train/test losses for our proposed algorithms compared to the decentralized gradient descent and gradient tracking algorithms on synthetic data  $y = \text{sign}(a^\top w^*)$  with  $n = 800, d = 50$ .




Figure 3: Consensus error, train loss and test loss for DGD with exponential loss and linearly separable data. Left and middle plots verify the rates  $\tilde{O}(1/t^2)$  and  $\tilde{O}(1/t)$  (Theorem 4 and Remark 5) for *consensus error* and *training loss* of DGD. Right plot shows test loss for DGD and GD show approximately similar convergence behavior under separable data.

## 4 CONCLUSIONS

We studied the behavior of train loss and test loss of decentralized gradient descent (DGD) methods when training dataset is separable. To the best of our knowledge, this yields the first rigorous guarantees for the generalization error of DGD in such a setting. For the same setting, we also proposed fast algorithms and empirically verified that they accelerate both training and test accuracy. We believe our work opens several directions, with perhaps the most exciting one being the analysis of non-convex objectives. We are also interested in extending our results to other distributed settings such as federated learning [Kairouz et al., 2021] and Gradient Tracking e.g., [Nedic et al., 2017].

## Acknowledgements

This work is partially supported by NSF Grant CCF-2009030.

## References

[win,] Uci wine data set, web address : <https://archive.ics.uci.edu/ml/datasets/wine>.

[Arora et al., 2019] Arora, S., Du, S., Hu, W., Li, Z., and Wang, R. (2019). Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *International Conference on Machine Learning*, pages 322–332. PMLR.

[Assran et al., 2018] Assran, M., Loizou, N., Ballas, N., and Rabbat, M. (2018). Stochastic gradient push for distributed deep learning. *arXiv preprint arXiv:1811.10792*.

[Bousquet and Elisseeff, 2002] Bousquet, O. and Elisseeff, A. (2002). Stability and generalization. *The Journal of Machine Learning Research*, 2:499–526.

[Chizat and Bach, 2020] Chizat, L. and Bach, F. (2020). Implicit bias of gradient descent for wide two-layer neural networks trained with the logistic loss. In *Conference on Learning Theory*, pages 1305–1338. PMLR.

[Hardt et al., 2016] Hardt, M., Recht, B., and Singer, Y. (2016). Train faster, generalize better: Stability of stochastic gradient descent. In *International conference on machine learning*, pages 1225–1234. PMLR.

[Ji et al., 2021] Ji, Z., Srebro, N., and Telgarsky, M. (2021). Fast margin maximization via dual acceleration. In

- International Conference on Machine Learning*, pages 4860–4869. PMLR.
- [Ji and Telgarsky, 2018] Ji, Z. and Telgarsky, M. (2018). Risk and parameter convergence of logistic regression. *arXiv preprint arXiv:1803.07300*.
- [Ji and Telgarsky, 2020] Ji, Z. and Telgarsky, M. (2020). Directional convergence and alignment in deep learning. *Advances in Neural Information Processing Systems*, 33:17176–17186.
- [Ji and Telgarsky, 2021] Ji, Z. and Telgarsky, M. (2021). Characterizing the implicit bias via a primal-dual analysis. In *Algorithmic Learning Theory*, pages 772–804. PMLR.
- [Jiang et al., 2017] Jiang, Z., Balu, A., Hegde, C., and Sarkar, S. (2017). Collaborative deep learning in fixed topology networks. *Advances in Neural Information Processing Systems*, 30.
- [Johnson and Zhang, 2013] Johnson, R. and Zhang, T. (2013). Accelerating stochastic gradient descent using predictive variance reduction. *Advances in neural information processing systems*, 26.
- [Kairouz et al., 2021] Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., et al. (2021). Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2):1–210.
- [Koloskova et al., 2021] Koloskova, A., Lin, T., and Stich, S. U. (2021). An improved analysis of gradient tracking for decentralized machine learning. *Advances in Neural Information Processing Systems*, 34.
- [Koloskova et al., 2019] Koloskova, A., Lin, T., Stich, S. U., and Jaggi, M. (2019). Decentralized deep learning with arbitrary communication compression. *arXiv preprint arXiv:1907.09356*.
- [Koloskova et al., 2020] Koloskova, A., Loizou, N., Boreiri, S., Jaggi, M., and Stich, S. (2020). A unified theory of decentralized sgd with changing topology and local updates. In *International Conference on Machine Learning*, pages 5381–5393. PMLR.
- [Kovalev et al., 2020] Kovalev, D., Salim, A., and Richtárik, P. (2020). Optimal and practical algorithms for smooth and strongly convex decentralized optimization. *Advances in Neural Information Processing Systems*, 33:18342–18352.
- [Lei and Ying, 2020] Lei, Y. and Ying, Y. (2020). Fine-grained analysis of stability and generalization for stochastic gradient descent. In *International Conference on Machine Learning*, pages 5809–5819. PMLR.
- [Lian et al., 2017] Lian, X., Zhang, C., Zhang, H., Hsieh, C.-J., Zhang, W., and Liu, J. (2017). Can decentralized algorithms outperform centralized algorithms? a case study for decentralized parallel stochastic gradient descent. In *Advances in Neural Information Processing Systems*, pages 5330–5340.
- [Lin et al., 2021] Lin, T., Karimireddy, S. P., Stich, S., and Jaggi, M. (2021). Quasi-global momentum: Accelerating decentralized deep learning on heterogeneous data. In *International Conference on Machine Learning*, pages 6654–6665. PMLR.
- [Lojasiewicz, 1963] Lojasiewicz, S. (1963). A topological property of real analytic subsets. *Coll. du CNRS, Les equations aux derive es partielles*.
- [Nacson et al., 2019] 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 *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 3420–3428. PMLR.
- [Nedić and Olshevsky, 2014] Nedić, A. and Olshevsky, A. (2014). Distributed optimization over time-varying directed graphs. *IEEE Transactions on Automatic Control*, 60(3):601–615.
- [Nedić and Olshevsky, 2016] Nedić, A. and Olshevsky, A. (2016). Stochastic gradient-push for strongly convex functions on time-varying directed graphs. *IEEE Transactions on Automatic Control*, 61(12):3936–3947.
- [Nedic et al., 2017] Nedic, A., Olshevsky, A., and Shi, W. (2017). Achieving geometric convergence for distributed optimization over time-varying graphs. *SIAM Journal on Optimization*, 27(4):2597–2633.
- [Nedic and Ozdaglar, 2009] Nedic, A. and Ozdaglar, A. (2009). Distributed subgradient methods for multi-agent optimization. *IEEE Transactions on Automatic Control*, 54(1):48–61.
- [Nesterov, 2003] Nesterov, Y. (2003). *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media.
- [Polyak, 1963] Polyak, B. (1963). Gradient methods for the minimisation of functionals. *Ussr Computational Mathematics and Mathematical Physics*, 3:864–878.

- [Pu et al., 2020] Pu, S., Shi, W., Xu, J., and Nedic, A. (2020). Push-pull gradient methods for distributed optimization in networks. *IEEE Transactions on Automatic Control*.
- [Richards et al., 2020a] Richards, D. et al. (2020a). Graph-dependent implicit regularisation for distributed stochastic subgradient descent. *Journal of Machine Learning Research*, 21(2020).
- [Richards et al., 2020b] Richards, D., Rebeschini, P., and Rosasco, L. (2020b). Decentralised learning with random features and distributed gradient descent. In *International Conference on Machine Learning*, pages 8105–8115. PMLR.
- [Rosset et al., 2003] Rosset, S., Zhu, J., and Hastie, T. J. (2003). Margin maximizing loss functions. In *NIPS*.
- [Schliserman and Koren, 2022] Schliserman, M. and Koren, T. (2022). Stability vs implicit bias of gradient methods on separable data and beyond. *arXiv preprint arXiv:2202.13441*.
- [Shamir, 2021] Shamir, O. (2021). Gradient methods never overfit on separable data. *Journal of Machine Learning Research*, 22(85):1–20.
- [Shi et al., 2015] Shi, W., Ling, Q., Wu, G., and Yin, W. (2015). Extra: An exact first-order algorithm for decentralized consensus optimization. *SIAM Journal on Optimization*, 25(2):944–966.
- [Soudry et al., 2018] Soudry, D., Hoffer, E., Nacson, M. S., Gunasekar, S., and Srebro, N. (2018). The implicit bias of gradient descent on separable data. *The Journal of Machine Learning Research*, 19(1):2822–2878.
- [Sun et al., 2021] Sun, T., Li, D., and Wang, B. (2021). Stability and generalization of decentralized stochastic gradient descent. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 9756–9764.
- [Sun et al., 2022] Sun, Y., Maros, M., Scutari, G., and Cheng, G. (2022). High-dimensional inference over networks: Linear convergence and statistical guarantees. *arXiv preprint arXiv:2201.08507*.
- [Taheri and Thrampoulidis, 2023] Taheri, H. and Thrampoulidis, C. (2023). Fast convergence in learning two-layer neural networks with separable data. In *AAAI Conference on Artificial Intelligence*.
- [Toghani and Uribe, 2022a] Toghani, M. T. and Uribe, C. A. (2022a). Communication-efficient distributed cooperative learning with compressed beliefs. *IEEE Transactions on Control of Network Systems*.
- [Toghani and Uribe, 2022b] Toghani, M. T. and Uribe, C. A. (2022b). Scalable average consensus with compressed communications. In *2022 American Control Conference (ACC)*, pages 3412–3417. IEEE.
- [Wang et al., 2019] Wang, J., Tantia, V., Ballas, N., and Rabbat, M. (2019). Slowmo: Improving communication-efficient distributed sgd with slow momentum. *arXiv preprint arXiv:1910.00643*.
- [Xin et al., 2021] Xin, R., Khan, U. A., and Kar, S. (2021). An improved convergence analysis for decentralized online stochastic non-convex optimization. *IEEE Transactions on Signal Processing*, 69:1842–1858.
- [Yuan et al., 2016] Yuan, K., Ling, Q., and Yin, W. (2016). On the convergence of decentralized gradient descent. *SIAM Journal on Optimization*, 26(3):1835–1854.
- [Zhang et al., 2021] Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. (2021). Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM*, 64(3):107–115.

## APPENDIX

In this section, we present the proofs of all theorems and lemmas stated in the main body. We organize the appendix as follows,

The formal statement and proof of Lemma 1 are included in Appendix A.

The proofs for Section 2.1 are included in Appendix B.

The proofs for Section 2.2 are included in Appendix C.

The proofs for Section 2.3 are included in Appendix D.

The proof of Theorem 7 is included in Appendix E.

Auxiliary results on our assumptions are included in Appendix F.

Finally, we conduct complementary experiments in Appendix G.

### Notation

Throughout the appendix we use the following notations:

$$\begin{aligned}\bar{w} &:= \frac{1}{N} \sum_{\ell=1}^N w_\ell, \quad \bar{W} := [\bar{w}, \bar{w}, \dots, \bar{w}]^\top \in \mathbb{R}^{N \times d}, \\ W &:= [w_1, \dots, w_N]^\top \in \mathbb{R}^{N \times d}, \\ \hat{F}(W) &:= \frac{1}{N} \sum_{\ell=1}^N \hat{F}_\ell(w_\ell), \\ \nabla \hat{F}(w) &:= \frac{1}{n} \sum_{i=1}^n \nabla f(w, x_i), \\ \nabla \hat{F}(W) &:= [\nabla \hat{F}_1(w_1), \nabla \hat{F}_2(w_2), \dots, \nabla \hat{F}_N(w_N)]^\top \in \mathbb{R}^{N \times d}, \\ \bar{\nabla} \hat{F}(W) &:= \frac{1}{N} \sum_{\ell=1}^N \nabla \hat{F}_\ell(w_\ell), \\ \nabla \hat{F}_\ell(w_\ell) &:= \frac{N}{n} \sum_{x_j \in \mathcal{S}_\ell} \nabla f(w_\ell, x_j).\end{aligned}$$

where recall that  $d$  is the dimension of ambient space,  $n$  is the total sample size,  $N$  is the number of agents and each agent has access to  $n/N$  samples.

### A Proof of Lemma 1

**Lemma 8** (Formal statement of Lemma 1). *Consider the iterates of decentralized gradient descent in Eq.(2) with a fixed positive step-size  $\eta \leq \frac{2}{L}$ . Let Assumptions 1-4 hold. Then for the test loss  $F$  at iteration  $T \geq 1$ , it holds that*

$$\begin{aligned}\mathbb{E} [F(\bar{w}^{(T)})] &\leq 4\mathbb{E} [\hat{F}(\bar{w}^{(T)})] + \frac{9L^2 c^2 \eta^2 T^2}{n^{3-2\alpha}} \mathbb{E} \left[ \left( \frac{1}{T} \sum_{t=1}^T \hat{F}(\bar{w}^{(t)}) \right)^{2\alpha} \right] \\ &\quad + \frac{9L^4 \eta^2}{N} \mathbb{E} \left[ \left( \sum_{t=1}^T \|W^{(t)} - \bar{W}^{(t)}\|_F \right)^2 \right] + \frac{9L^4 \eta^2}{N} \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[ \left( \sum_{t=1}^T \|W_{\neg i}^{(t)} - \bar{W}_{\neg i}^{(t)}\|_F \right)^2 \right].\end{aligned}\tag{10}$$

where the expectation is over training samples and  $W_{\neg i}^{(t)}, \bar{W}_{\neg i}^{(t)}$  denote the parameter matrix and averaged parameter matrix at iteration  $t$  for the DGD algorithm when the  $i$ -th data sample is left out.



*Proof.* The proof relies on algorithmic stability [Bousquet and Elisseeff, 2002, Hardt et al., 2016]. Specifically, we build on the framework introduced by [Lei and Ying, 2020] (and also used recently by [Schliserman and Koren, 2022]). Unlike these works, our analysis is for decentralized gradient descent.

We define  $w_{\ell, \neg i}^{(t)}$  as the parameter of agent  $\ell$  resulting from decentralized gradient descent at iteration  $t$ , when the  $i^{\text{th}}$  training sample  $i \leq n$  is left out during training. We emphasize that the  $i^{\text{th}}$  sample may or may not belong to the dataset of agent  $\ell$ .

We define  $\bar{w}_{\neg i}^{(t)} \in \mathbb{R}^d$

$$\bar{w}_{\neg i}^{(t)} := \frac{1}{N} \sum_{j=1}^N w_{j, \neg i}^{(t)},$$

as the average of all agents' parameters at iteration  $t$ , when the  $i^{\text{th}}$  sample is left out of the algorithm. Thus, the parameter matrices  $W_{\neg i}^{(t)}, \bar{W}_{\neg i}^{(t)} \in \mathbb{R}^{N \times d}$  are defined as follows,

$$\begin{aligned} W_{\neg i}^{(t)} &:= [w_{1, \neg i}^{(t)}, w_{2, \neg i}^{(t)}, \dots, w_{N, \neg i}^{(t)}], \\ \bar{W}_{\neg i}^{(t)} &:= [\bar{w}_{\neg i}^{(t)}, \bar{w}_{\neg i}^{(t)}, \dots, \bar{w}_{\neg i}^{(t)}]. \end{aligned}$$

The first step in the proof is to bound the term  $\frac{1}{n} \sum_{i=1}^n \|\bar{w}^{(t)} - \bar{w}_{\neg i}^{(t)}\|^2$ . By definition of DGD in Eq.(2), we have the following update rule for the averaged parameter,

$$\bar{w}^{(t+1)} = \bar{w}^{(t)} - \eta \bar{\nabla} \hat{F}(W^{(t)}).$$

Analogously,

$$\bar{w}_{\neg i}^{(t+1)} = \bar{w}_{\neg i}^{(t)} - \eta \bar{\nabla} \hat{F}(W_{\neg i}^{(t)}) = \bar{w}_{\neg i}^{(t)} - \frac{\eta}{n} \sum_{\ell=1}^N \sum_{x_j \in S_\ell, x_j \neq x_i} \nabla f(w_{\ell, \neg i}^{(t)}, x_j).$$

Thus by adding and subtracting  $\bar{\nabla} \hat{F}(\bar{W}^{(t)})$  and  $\bar{\nabla} \hat{F}(\bar{W}_{\neg i}^{(t)})$ , we have

$$\begin{aligned} &\|\bar{w}^{(t+1)} - \bar{w}_{\neg i}^{(t+1)}\| \\ &= \left\| \bar{w}^{(t)} - \eta \bar{\nabla} \hat{F}(W^{(t)}) - (\bar{w}_{\neg i}^{(t)} - \eta \bar{\nabla} \hat{F}(W_{\neg i}^{(t)})) \right\| \\ &= \left\| \bar{w}^{(t)} - \eta \bar{\nabla} \hat{F}(\bar{W}^{(t)}) + \eta (\bar{\nabla} \hat{F}(W^{(t)}) - \bar{\nabla} \hat{F}(\bar{W}^{(t)})) \right. \\ &\quad \left. - (\bar{w}_{\neg i}^{(t)} - \eta \nabla \hat{F}(\bar{W}_{\neg i}^{(t)})) + \eta \bar{\nabla} \hat{F}(W_{\neg i}^{(t)}) - \eta \nabla \hat{F}(\bar{W}_{\neg i}^{(t)}) \right\| \\ &\leq \left\| \bar{w}^{(t)} - \eta \bar{\nabla} \hat{F}(\bar{W}^{(t)}) - (\bar{w}_{\neg i}^{(t)} - \eta \bar{\nabla} \hat{F}(\bar{W}_{\neg i}^{(t)})) \right\| \\ &\quad + \eta \left\| \bar{\nabla} \hat{F}(W_{\neg i}^{(t)}) - \nabla \hat{F}(\bar{W}_{\neg i}^{(t)}) \right\| + \eta \left\| \bar{\nabla} \hat{F}(W^{(t)}) - \nabla \hat{F}(\bar{W}^{(t)}) \right\|. \end{aligned}$$

For the last term, using smoothness, we can write

$$\begin{aligned} \left\| \bar{\nabla} \hat{F}(W^{(t)}) - \nabla \hat{F}(\bar{W}^{(t)}) \right\| &= \frac{1}{N} \left\| \sum_{\ell=1}^N \nabla \hat{F}_\ell(w_\ell^{(t)}) - \nabla \hat{F}_\ell(\bar{w}^{(t)}) \right\| \\ &\leq \frac{1}{N} \sum_{\ell=1}^N \left\| \nabla \hat{F}_\ell(w_\ell^{(t)}) - \nabla \hat{F}_\ell(\bar{w}^{(t)}) \right\| \\ &\leq \frac{L}{N} \sum_{\ell=1}^N \|w_\ell^{(t)} - \bar{w}^{(t)}\| \\ &\leq \frac{L}{\sqrt{N}} \left( \sum_{\ell=1}^N \|w_\ell^{(t)} - \bar{w}^{(t)}\|^2 \right)^{1/2} = \frac{L}{\sqrt{N}} \|W^{(t)} - \bar{W}^{(t)}\|_F. \end{aligned}$$

The second term is upper-bounded similarly. Using these bounds, splitting the gradient  $\nabla \hat{F}(\bar{W}^{(t)}) = \frac{1}{n} \sum_{i' \neq i} f(\bar{w}^{(t)}, x_{i'}) + \frac{1}{n} f(\bar{w}^{(t)}, x_i)$ , using smoothness and convexity  $\|w + \eta \nabla f(w) - v - \eta \nabla f(v)\| \leq \|w - v\|$  for  $\eta \leq 2/L$  [Nesterov, 2003] and employing Assumption 4 we can write

$$\begin{aligned}
 \|\bar{w}^{(t+1)} - \bar{w}_{-i}^{(t+1)}\| &\leq \|\bar{w}^{(t)} - \eta \nabla \hat{F}(\bar{W}^{(t)}) - (\bar{w}_{-i}^{(t)} - \eta \nabla \hat{F}(\bar{W}_{-i}^{(t)}))\| \\
 &\quad + \frac{\eta L}{\sqrt{N}} \|W_{-i}^{(t)} - \bar{W}_{-i}^{(t)}\|_F + \frac{\eta L}{\sqrt{N}} \|W^{(t)} - \bar{W}^{(t)}\|_F \\
 &\leq \frac{1}{n} \sum_{i' \neq i} \|\bar{w}^{(t)} - \eta \nabla f(\bar{w}^{(t)}, x_{i'}) - \bar{w}_{-i}^{(t)} + \eta \nabla f(\bar{w}_{-i}^{(t)}, x_{i'})\| + \frac{1}{n} \|\bar{w}^{(t)} - \eta \nabla f(\bar{w}^{(t)}, x_i) - \bar{w}_{-i}^{(t)}\| \\
 &\quad + \frac{\eta L}{\sqrt{N}} \|W_{-i}^{(t)} - \bar{W}_{-i}^{(t)}\|_F + \frac{\eta L}{\sqrt{N}} \|W^{(t)} - \bar{W}^{(t)}\|_F \\
 &\leq \|\bar{w}^{(t)} - \bar{w}_{-i}^{(t)}\| + \frac{\eta}{n} \|\nabla f(\bar{w}^{(t)}, x_i)\| + \frac{\eta L}{\sqrt{N}} \|W_{-i}^{(t)} - \bar{W}_{-i}^{(t)}\|_F + \frac{\eta L}{\sqrt{N}} \|W^{(t)} - \bar{W}^{(t)}\|_F \\
 &\leq \|\bar{w}^{(t)} - \bar{w}_{-i}^{(t)}\| + \frac{c\eta}{n} (f(\bar{w}^{(t)}, x_i))^\alpha + \frac{\eta L}{\sqrt{N}} \|W_{-i}^{(t)} - \bar{W}_{-i}^{(t)}\|_F + \frac{\eta L}{\sqrt{N}} \|W^{(t)} - \bar{W}^{(t)}\|_F.
 \end{aligned}$$

By summing over  $t \in [T]$ ,

$$\|\bar{w}^{(T+1)} - \bar{w}_{-i}^{(T+1)}\| \leq \frac{c\eta}{n} \sum_{t=1}^T (f(\bar{w}^{(t)}, x_i))^\alpha + \frac{\eta L}{\sqrt{N}} \sum_{t=1}^T \|W_{-i}^{(t)} - \bar{W}_{-i}^{(t)}\|_F + \frac{\eta L}{\sqrt{N}} \sum_{t=1}^T \|W^{(t)} - \bar{W}^{(t)}\|_F.$$

We define for the ease of notation the following two consensus terms,

$$e^{(T)} := \left( \sum_{t=1}^T \|W^{(t)} - \bar{W}^{(t)}\|_F \right)^2 \quad \text{and} \quad e_{-i}^{(T)} := \left( \sum_{t=1}^T \|W_{-i}^{(t)} - \bar{W}_{-i}^{(t)}\|_F \right)^2. \quad (11)$$

Thus the bound for the squared term can be written as follows

$$\|\bar{w}^{(T+1)} - \bar{w}_{-i}^{(T+1)}\|^2 \leq \frac{3c^2\eta^2}{n^2} \left( \sum_{t=1}^T (f(\bar{w}^{(t)}, x_i))^\alpha \right)^2 + \frac{3\eta^2 L^2}{N} e^{(T)} + \frac{3\eta^2 L^2}{N} e_{-i}^{(T)}.$$

By averaging over  $i \in [n]$  and noting that  $\alpha \in [1/2, 1]$  so that  $x^\alpha$  is concave, we conclude that

$$\begin{aligned}
 \frac{1}{n} \sum_{i=1}^n \|\bar{w}^{(T+1)} - \bar{w}_{-i}^{(T+1)}\|^2 &\leq \frac{3c^2\eta^2}{n^3} \sum_{i=1}^n \left( \sum_{t=1}^T (f(\bar{w}^{(t)}, x_i))^\alpha \right)^2 + \frac{3\eta^2 L^2}{Nn} \sum_{i=1}^n e_{-i}^{(T)} + \frac{3\eta^2 L^2}{N} e^{(T)} \\
 &\leq \frac{3c^2\eta^2 T^2}{n^3} \sum_{i=1}^n \left( \frac{1}{T} \sum_{t=1}^T f(\bar{w}^{(t)}, x_i) \right)^{2\alpha} + \frac{3\eta^2 L^2}{Nn} \sum_{i=1}^n e_{-i}^{(T)} + \frac{3\eta^2 L^2}{N} e^{(T)} \\
 &\leq \frac{3c^2\eta^2 T^{2(1-\alpha)}}{n^{3-2\alpha}} \left( \sum_{t=1}^T \hat{F}(\bar{w}^{(t)}) \right)^{2\alpha} + \frac{3\eta^2 L^2}{Nn} \sum_{i=1}^n e_{-i}^{(T)} + \frac{3\eta^2 L^2}{N} e^{(T)}.
 \end{aligned}$$

Thus we have for iteration  $T$ :

$$\frac{1}{n} \sum_{i=1}^n \|\bar{w}^{(T)} - \bar{w}_{-i}^{(T)}\|^2 \leq \frac{3c^2\eta^2 T^2}{n^{3-2\alpha}} \left( \frac{1}{T} \sum_{t=1}^T \hat{F}(\bar{w}^{(t)}) \right)^{2\alpha} + \frac{3\eta^2 L^2}{Nn} \sum_{i=1}^n e_{-i}^{(T)} + \frac{3\eta^2 L^2}{N} e^{(T)}. \quad (12)$$

Next we use [Schliserman and Koren, 2022, Lemma 7] (see also [Lei and Ying, 2020, Theorem 2]), which states that for the  $L$ -smooth loss  $f$ , the test error of the output  $w$  of an algorithm taking as input a dataset  $(x_1, \dots, x_n)$  size  $n$ , satisfies the following,

$$\mathbb{E}[F(w)] \leq 4\mathbb{E}[\hat{F}(w)] + \frac{3L^2}{n} \sum_{i=1}^n \mathbb{E}[\|w - w_{-i}\|^2],$$

where expectations are taken over the training set  $(x_1, x_2, \dots, x_n)$ . We replace  $w$  with  $\bar{w}^{(T)}$  and by using (12) (which we can do because it holds true for all datasets since Assumptions 1-4 hold for every sample  $x$  in the distribution),

$$\begin{aligned} \mathbb{E}[F(\bar{w}^{(T)})] &\leq 4\mathbb{E}[\hat{F}(\bar{w}^{(T)})] + \frac{3L^2}{n} \sum_{i=1}^n \mathbb{E}[\|\bar{w}^{(T)} - \bar{w}_{-i}^{(T)}\|^2] \\ &\leq 4\mathbb{E}[\hat{F}(\bar{w}^{(T)})] + \frac{9L^2 c^2 \eta^2 T^2}{n^{3-2\alpha}} \mathbb{E}\left[\left(\frac{1}{T} \sum_{t=1}^T \hat{F}(\bar{w}^{(t)})\right)^{2\alpha}\right] \\ &\quad + \frac{9L^4 \eta^2}{Nn} \sum_{i=1}^n \mathbb{E}[e_{-i}^{(T)}] + \frac{9L^4 \eta^2}{N} \mathbb{E}[e^{(T)}]. \end{aligned}$$

This leads to (10) and completes the proof.  $\square$

Finally, we explain the informal version of the lemma presented in the main body (Lemma 1). Compared to the bound in Eq. (10), the informal Lemma 1 combines the consensus-error term  $e^{(T)}$  with the average leave-one-out consensus-error term  $\frac{1}{n} \sum_{i \in [n]} e_{-i}^{(T)}$  (recall the definitions in (11)). It is convenient doing that for the following reason. To apply Lemma 8, we need upper bounds on  $e^{(T)}$  and  $e_{-i}^{(T)}$  (for specific assumptions on the function class that is optimized). We do this in the section that follows. It turns out that the bounds we obtain for the consensus-error term  $e^{(T)}$  also holds for the leave-one-out consensus error terms  $e_{-i}^{(T)}$ ,  $i \leq [n]$ . The reason for that is that our bounds are not affected by the sample-size, but rather they depend crucially only on the smoothness parameter of the train loss. It is easy to see that the smoothness parameter of the leave-one-out train loss  $\frac{1}{n} \sum_{i' \neq i} f(w, x_{i'})$  is upper-bounded by the smoothness parameter of  $\frac{1}{n} \sum_{i'} f(w, x_{i'})$ . See Remark 7 for more details.

## B Proofs for Section 2.1

**Lemma 9** (Recursions for the consensus error). *Let the step-size  $\eta \leq (1 - \lambda)/4L$  where  $\lambda := \max(|\lambda_2(A)|, |\lambda_N(A)|)^2$ . The consensus error of DGD under Assumptions 1,3 satisfies the following:*

$$\|W^{(t)} - \bar{W}^{(t)}\|_F^2 < \alpha_1 \|W^{(t-1)} - \bar{W}^{(t-1)}\|_F^2 + \alpha_2 N \eta^2 L^2 \hat{F}(\bar{w}^{(t-1)}), \quad (13)$$

where  $\alpha_1 := \frac{3+\lambda}{4}$ ,  $\alpha_2 := 4(\frac{2}{1-\lambda} - 1)$ .

*Proof.* Denoting  $A^\infty := \lim_{t \rightarrow \infty} A^t = \frac{1}{N} \mathbf{1}\mathbf{1}^T$ , it holds by Assumption 1,

$$\begin{aligned} \|AW - \bar{W}\|_F^2 &= \|(A - A^\infty)(W - \bar{W})\|_F^2 = \sum_{i=1}^N \|(A - A^\infty)(W_i - \bar{W}_i)\|^2 \\ &\leq \sum_{i=1}^N \|A - A^\infty\|^2 \|W_i - \bar{W}_i\|^2 \leq \max(\lambda_2^2(A), \lambda_N^2(A)) \cdot \|W - \bar{W}\|_F^2, \end{aligned} \quad (14)$$

where  $W_i$  is the  $i$  th column of  $W$ . By Assumption 1,  $\lambda = \max(|\lambda_2(A)|, |\lambda_N(A)|)^2 < 1$ . For the consensus error, we can write,

$$\begin{aligned} \|W^{(t)} - \bar{W}^{(t)}\|_F^2 &= \|W^{(t)} - \bar{W}^{(t-1)} - \bar{W}^{(t)} + \bar{W}^{(t-1)}\|_F^2 \\ &\leq \|W^{(t)} - \bar{W}^{(t-1)}\|_F^2 \\ &= \|AW^{(t-1)} - \eta \nabla \hat{F}(W^{(t-1)}) - \bar{W}^{(t-1)}\|_F^2 \\ &\leq (1 + \beta) \|AW^{(t-1)} - \bar{W}^{(t-1)}\|_F^2 + (1 + \beta^{-1}) \eta^2 \|\nabla \hat{F}(W^{(t-1)})\|_F^2, \end{aligned}$$

where the second step is due to  $\|X - \bar{X}\|_F \leq \|X\|_F$  [Koloskova et al., 2019, Koloskova et al., 2020]. The last line holds for any  $\beta > 0$ , due to  $\|a + b\|^2 \leq (1 + \beta^{-1})\|a\|^2 + (1 + \beta)\|b\|^2$ .

Based on this inequality and by noting (14) and using the  $L$ -smoothness assumption, we can deduce that,

$$\begin{aligned}
 \|W^{(t)} - \bar{W}^{(t)}\|_F^2 &\leq (1 + \beta)\lambda \|W^{(t-1)} - \bar{W}^{(t-1)}\|_F^2 + (1 + \beta^{-1})\eta^2 \|\nabla \hat{F}(W^{(t-1)})\|_F^2 \\
 &\leq (1 + \beta)\lambda \|W^{(t-1)} - \bar{W}^{(t-1)}\|_F^2 + 2(1 + \beta^{-1})\eta^2 \|\nabla \hat{F}(W^{(t-1)}) - \nabla \hat{F}(\bar{W}^{(t-1)})\|_F^2 \\
 &\quad + 2(1 + \beta^{-1})\eta^2 \|\nabla \hat{F}(\bar{W}^{(t-1)})\|_F^2 \\
 &\leq (1 + \beta)\lambda \|W^{(t-1)} - \bar{W}^{(t-1)}\|_F^2 + 2(1 + \beta^{-1})\eta^2 L^2 \|W^{(t-1)} - \bar{W}^{(t-1)}\|_F^2 \\
 &\quad + 4(1 + \beta^{-1})\eta^2 LN \hat{F}(\bar{w}_{t-1}),
 \end{aligned} \tag{15}$$

where the last step is due to  $L$ -smoothness and the non-negativity of  $\hat{F}_\ell$ , i.e.

$$\|\nabla \hat{F}(\bar{W}^{(t-1)})\|_F^2 = \sum_{\ell=1}^N \|\nabla \hat{F}_\ell(\bar{w}^{(t-1)})\|^2 \leq 2L \sum_{\ell=1}^N (\hat{F}_\ell(\bar{w}^{(t-1)}) - \hat{F}_\ell^*) \leq 2LN \hat{F}(\bar{w}^{(t-1)}).$$

Thus,

$$\|W^{(t)} - \bar{W}^{(t)}\|_F^2 < ((1 + \beta)\lambda + 2(1 + \beta^{-1})\eta^2 L^2) \|W^{(t-1)} - \bar{W}^{(t-1)}\|_F^2 + 4(1 + \beta^{-1})\eta^2 LN \hat{F}(\bar{w}^{(t-1)}).$$

Next, choose  $\beta = (1 - \lambda)/(2\lambda)$ . Then, it follows from the assumption  $\eta \leq (1 - \lambda)/4L$  that

$$\begin{aligned}
 (1 + \beta)\lambda + 2(1 + \beta^{-1})\eta^2 L^2 &< (3 + \lambda)/4 = \alpha_1, \\
 4(1 + \beta^{-1}) &< 4(2/(1 - \lambda) - 1) = \alpha_2.
 \end{aligned}$$

This concludes the lemma.  $\square$

By telescoping summation over the iterates  $t = 1, \dots, T$  of the consensus error in Eq.(13), we end up with the consensus error at iteration  $T$ . The final expression is stated in the next lemma.

**Lemma 10.** *Under the assumptions of Lemma 9, it holds for  $T > 1$  that,*

$$\|W^{(T)} - \bar{W}^{(T)}\|_F^2 < \alpha_1^{T-1} \|W_1 - \bar{W}_1\|_F^2 + (\alpha_2 \eta^2 LN) \sum_{t=1}^{T-1} \alpha_1^{t-1} \hat{F}(\bar{w}^{(T-t)}).$$

**Lemma 11.** *Under the assumptions of Lemma 9 and the zero initialization assumption for all agents, the average consensus error satisfies,*

$$\frac{1}{NT} \sum_{t=1}^T \|W^{(t)} - \bar{W}^{(t)}\|_F^2 \leq \frac{\alpha_2 \eta^2 L}{(1 - \alpha_1)T} \sum_{t=1}^{T-1} \hat{F}(\bar{w}^{(t)}).$$

*Proof.* By Lemma 10 and the zero initialization and non-negativity assumptions, we have

$$\begin{aligned}
 \frac{1}{NT} \sum_{t=1}^T \|W^{(t)} - \bar{W}^{(t)}\|_F^2 &\leq \frac{\alpha_2 \eta^2 L}{T} \sum_{t=2}^T \sum_{\tau=1}^{t-1} \alpha_1^{\tau-1} \hat{F}(\bar{w}^{(t-\tau)}) \leq \frac{\alpha_2 \eta^2 L}{T} \sum_{\tau=1}^{T-1} \alpha_1^{\tau-1} \sum_{t=1}^{T-\tau} \hat{F}(\bar{w}^{(t)}) \\
 &\leq \frac{\alpha_2 \eta^2 L}{T} \sum_{\tau=1}^{T-1} \alpha_1^{\tau-1} \sum_{t=1}^{T-1} \hat{F}(\bar{w}^{(t)}) \leq \frac{\alpha_2 \eta^2 L}{(1 - \alpha_1)T} \sum_{t=1}^{T-1} \hat{F}(\bar{w}^{(t)}).
 \end{aligned}$$

$\square$

**Lemma 12.** *Under Assumptions 2,3 and for all  $w \in \mathbb{R}^d$ , the DGD updates satisfy the following recursions:*

$$\frac{2\eta - 4L\eta^2}{T} \sum_{t=1}^{T-1} \hat{F}(\bar{w}^{(t)}) \leq \frac{\|\bar{w}^{(1)} - w\|^2}{T} + 2\eta \hat{F}(w) + \frac{2L^2\eta^2 + \eta L}{NT} \sum_{t=1}^{T-1} \|W^{(t)} - \bar{W}^{(t)}\|_F^2$$



*Proof.* We start by upper bounding the following quantity:

$$\|\bar{w}^{(t+1)} - w\|^2 = \|\bar{w}^{(t)} - \eta \bar{\nabla} \hat{F}(W^{(t)}) - w\|^2 = \|\bar{w}^{(t)} - w\|^2 + \eta^2 \|\bar{\nabla} \hat{F}(W^{(t)})\|^2 - 2\eta \langle \bar{w}^{(t)} - w, \bar{\nabla} \hat{F}(W^{(t)}) \rangle$$

For the second term above, using  $L$ -smoothness and non-negativity of the loss, we obtain:

$$\begin{aligned} \|\bar{\nabla} \hat{F}(W^{(t)})\|^2 &= \|\bar{\nabla} \hat{F}(W^{(t)}) - \nabla \hat{F}(\bar{w}^{(t)}) + \nabla \hat{F}(\bar{w}^{(t)})\|^2 \\ &\leq 2\|\bar{\nabla} \hat{F}(W^{(t)}) - \nabla \hat{F}(\bar{w}^{(t)})\|^2 + 2\|\nabla \hat{F}(\bar{w}^{(t)})\|^2 \\ &\leq \frac{2L^2}{N} \sum_{\ell=1}^N \|w_\ell^{(t)} - \bar{w}^{(t)}\|^2 + 4L\hat{F}(\bar{w}^{(t)}). \end{aligned}$$

For the third term, by using  $L$ -smoothness and convexity properties we can write,

$$\begin{aligned} \langle \bar{w}^{(t)} - w, \bar{\nabla} \hat{F}(W^{(t)}) \rangle &= \frac{1}{N} \sum_{\ell=1}^N \langle \bar{w}^{(t)} - w, \nabla \hat{F}_\ell(w_\ell^{(t)}) \rangle \\ &= \frac{1}{N} \sum_{\ell=1}^N \langle \bar{w}^{(t)} - w_\ell^{(t)}, \nabla \hat{F}_\ell(w_\ell^{(t)}) \rangle + \frac{1}{n} \sum_{\ell=1}^N \langle w_\ell^{(t)} - w, \nabla \hat{F}_\ell(w_\ell^{(t)}) \rangle \\ &\geq \frac{1}{N} \sum_{\ell=1}^N (\hat{F}_\ell(\bar{w}^{(t)}) - \hat{F}_\ell(w_\ell^{(t)})) - \frac{L}{2} \|w_\ell^{(t)} - \bar{w}^{(t)}\|^2 + \frac{1}{N} \sum_{\ell=1}^N (\hat{F}_\ell(w_\ell^{(t)}) - \hat{F}_\ell(w)) \\ &= \hat{F}(\bar{w}^{(t)}) - \hat{F}(w) - \frac{L}{2N} \|W^{(t)} - \bar{W}^{(t)}\|_F^2. \end{aligned}$$

Combining these inequalities we derive the following:

$$\begin{aligned} \|\bar{w}^{(t+1)} - w\|^2 &\leq \|\bar{w}^{(t)} - w\|^2 + \eta^2 \left( 2L^2 \|W^{(t)} - \bar{W}^{(t)}\|_F^2 / N + 4L\hat{F}(\bar{w}^{(t)}) \right) \\ &\quad - 2\eta \left( \hat{F}(\bar{w}^{(t)}) - \hat{F}(w) - \frac{L}{2} \|W^{(t)} - \bar{W}^{(t)}\|_F^2 / N \right). \end{aligned}$$

Summing these equations for  $t = 1, 2, \dots, T-1$ ,

$$\begin{aligned} \|\bar{w}^{(T)} - w\|^2 &\leq \|\bar{w}^{(1)} - w\|^2 + \sum_{t=1}^{T-1} \frac{2L^2\eta^2 + \eta L}{N} \|W^{(t)} - \bar{W}^{(t)}\|_F^2 + \sum_{t=1}^{T-1} (4L\eta^2 - 2\eta) \hat{F}(\bar{w}^{(t)}) + 2 \sum_{t=1}^{T-1} \eta \hat{F}(w) \\ &\implies \frac{2\eta - 4L\eta^2}{T} \sum_{t=1}^{T-1} \hat{F}(\bar{w}^{(t)}) \leq \frac{\|\bar{w}^{(1)} - w\|^2}{T} + \frac{2L^2\eta^2 + \eta L}{NT} \sum_{t=1}^{T-1} \|W^{(t)} - \bar{W}^{(t)}\|_F^2 + 2\eta \hat{F}(w). \end{aligned}$$

□

## B.1 Proof of Lemma 2

**Lemma 13** (Restatement of Lemma 2). *Under Assumptions 1-4 and zero initialization, for any  $w \in \mathbb{R}^d$  and for a fixed step-size  $\eta < \min\{\frac{1-\alpha_1}{L}, \frac{1}{L}\sqrt{\frac{1-\alpha_1}{2\alpha_2}}\}$ , where  $\alpha_1 \in (3/4, 1)$ ,  $\alpha_2 > 4$  are parameters that depend only on the mixing matrix, the following holds for the train loss and consensus error of DGD:*

$$\frac{1}{T} \sum_{t=1}^T \hat{F}(\bar{w}^{(t)}) \leq \frac{2\|w\|^2}{\eta T} + 4\hat{F}(w), \quad (16)$$

$$\frac{1}{NT} \sum_{t=1}^T \|W^{(t)} - \bar{W}^{(t)}\|_F^2 \leq \frac{\alpha_2 \eta^2 L}{(1-\alpha_1)} \left( \frac{2\|w\|^2}{\eta T} + 4\hat{F}(w) \right). \quad (17)$$

*Proof.* Recalling the initialization  $w_\ell^{(1)} = 0 \Rightarrow \bar{w}^{(1)} = 0$  and using  $\eta < 1/(4L)$ , we deduce from Lemma 12 that,

$$\frac{1}{T} \sum_{t=1}^{T-1} \hat{F}(\bar{w}^{(t)}) \leq \frac{\|w\|^2}{\eta T} + 2\hat{F}(w) + \frac{2L^2\eta + L}{NT} \sum_{t=1}^{T-1} \|W^{(t)} - \bar{W}^{(t)}\|_F^2$$

By Lemma 11,

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^{T-1} \hat{F}(\bar{w}^{(t)}) &\leq \frac{\|w\|^2}{\eta T} + 2\hat{F}(w) + \frac{(2L^2\eta + L)\alpha_2\eta^2 L}{T(1-\alpha_1)} \sum_{t=1}^{T-1} \hat{F}(\bar{w}^{(t)}) \\ &\leq \frac{\|w\|^2}{\eta T} + 2\hat{F}(w) + \frac{1}{2T} \sum_{t=1}^{T-1} \hat{F}(\bar{w}^{(t)}). \end{aligned} \quad (18)$$

where the condition on  $\eta$  on the lemma's statement ensures that  $(2L^2\eta + L)\alpha_2\eta^2 L/(1-\alpha_1) < 1/2$ . This gives the statement of the lemma for the training loss in Eq.(16). Appealing again to Lemma 11 for the consensus error yields (17).  $\square$

*Remark 7* (Bounds for leave-one-out consensus error). The bound in Eq. (17) also applies to the leave-one-out consensus-error term  $\frac{1}{T} \sum_{t=1}^T \|W_{-i}^{(t)} - \bar{W}_{-i}^{(t)}\|_F^2$ . To see this starting from Lemma 11 note that we still have

$$\frac{1}{T} \sum_{t=1}^T \|W_{-i}^{(t)} - \bar{W}_{-i}^{(t)}\|_F^2 \leq \frac{\alpha_2\eta^2 LN}{(1-\alpha_1)T} \sum_{t=1}^{T-1} \hat{F}_{-i}(\bar{w}_{-i,t}), \quad (19)$$

where we denote the leave-one-out train loss  $\hat{F}_{-i}(w) := \frac{1}{n} \sum_{i' \neq i} f(w, x_{i'})$ . This is true because the smoothness parameter of  $\hat{F}_{-i}(w)$  is  $(1-1/n)L \leq L$ . Moreover, applying Lemma 12 to the leave-one-out loss (and using again that it's smoothness parameter is upper bounded by  $L$ ), we have for all  $w$  that

$$\frac{2\eta - 4L\eta^2}{T} \sum_{t=1}^{T-1} \hat{F}_{-i}(\bar{w}_{-i}^{(t)}) \leq \frac{\|\bar{w}_{-i}^{(1)} - w\|^2}{T} + 2\eta\hat{F}_{-i}(w) + \frac{2L^2\eta^2 + \eta L}{NT} \sum_{t=1}^{T-1} \|W_{-i}^{(t)} - \bar{W}_{-i}^{(t)}\|_F^2$$

But, from the initialization assumption  $\bar{w}_{-i}^{(1)} = 0$  and also  $\hat{F}_{-i}(w) \leq \hat{F}(w)$  since the functions are assumed non-negative. Hence, and also using (19), shows that

$$\frac{2\eta - 4L\eta^2}{T} \sum_{t=1}^{T-1} \hat{F}_{-i}(\bar{w}_{-i}^{(t)}) \leq \frac{\|w\|^2}{T} + 2\eta\hat{F}(w) + \frac{(2L^2\eta^2 + \eta L)\alpha_2\eta^2 L}{T(1-\alpha_1)} \sum_{t=1}^{T-1} \hat{F}_{-i}(\bar{w}_{-i,t}).$$

Note that after using  $\eta < 1/(4L)$  this is exactly analogous to Eq. (18) for the train loss, which leads to the same bound  $\hat{F}_{-i}(\bar{w}_{-i}^{(t)}) \leq \frac{2\|w\|^2}{\eta T} + 4\hat{F}(w)$  for the leave-one-out loss. Plugging this back to Eq. (19) shows that the bound in (17) also holds for the leave-one-out consensus term.

## B.2 Proof of Theorem 3

We are ready to prove Theorem 3 by combining our results from Lemmas 2 and 1. We state the proof for general choice of step-size  $\eta$ . In particular, Theorem 3 follows by the next theorem after choosing  $\eta = O(1/\sqrt{T})$ .

**Theorem 14** (Theorem 3 for general  $\eta$ ). *Consider DGD under Assumptions 1-5, and choose  $\eta < \min\{\frac{1-\alpha_1}{L}, \frac{1}{L}\sqrt{\frac{1-\alpha_1}{2\alpha_2}}\}$ . The following bound holds for the averaged test error of DGD with separable data up to iteration  $T$ , assuming  $\varepsilon \leq \rho(\varepsilon)^2/\eta T$ ,*

$$\frac{1}{T} \sum_{t=1}^T F(\bar{w}^{(t)}) = O\left(\frac{\rho(\varepsilon)^2}{\eta T} + \frac{L^2 c^2 \rho(\varepsilon)^{4\alpha}}{n^{3-2\alpha}} (\eta T)^{2-2\alpha} + L^4 \rho(\varepsilon)^4 \eta^3 T\right).$$

*Proof.* By Lemma 1,

$$\mathbb{E}\left[F(\bar{w}^{(t)})\right] = O\left(\mathbb{E}\left[\hat{F}(\bar{w}^{(t)})\right] + \frac{L^2 c^2 \eta^2 t^2}{n^{3-2\alpha}} \mathbb{E}\left[\left(\frac{1}{t} \sum_{\tau=1}^t \hat{F}(\bar{w}^{(\tau)})\right)^{2\alpha}\right] + \frac{L^4 \eta^2}{N} \mathbb{E}\left[\left(\sum_{\tau=1}^t \|W^{(\tau)} - \bar{W}^{(\tau)}\|^2\right)\right]\right).$$

Thus, by Lemma 2,

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E} \left[ F(\bar{w}^{(t)}) \right] = O \left( \frac{\|w\|^2}{\eta T} + \hat{F}(w) + \frac{L^2 c^2 \eta^2}{n^{3-2\alpha}} \frac{1}{T} \sum_{t=1}^T t^2 \left( \frac{\|w\|^2}{\eta t} + \hat{F}(w) \right)^{2\alpha} + \frac{L^4 \eta^4}{N} \frac{1}{T} \sum_{t=1}^T t^2 \left( \frac{\|w\|^2}{\eta t} + \hat{F}(w) \right) \right).$$

By Assumption 5 and assuming  $\varepsilon \leq \rho(\varepsilon)^2 / \eta T$ , the statement of the theorem follows.  $\square$

## C Proofs for Section 2.2

**Lemma 15** (Iterates of consensus error). *Consider DGD with the loss functions and mixing matrix satisfying Assumptions 1, 2, 6 and Assumption 4 with  $\alpha = 1$  and  $c = h$ . By choosing  $\eta_{t-1} \leq \frac{1-\lambda}{4hNM_{(t-1)}}$  the consensus error at iteration  $t > 1$  satisfies*

$$\left\| W^{(t)} - \bar{W}^{(t)} \right\|_F^2 < \beta_1 \left\| W^{(t-1)} - \bar{W}^{(t-1)} \right\|_F^2 + \beta_2 \eta_{t-1}^2 h^2 N^2 \hat{F}^2(\bar{w}^{(t-1)}) \quad (20)$$

where we define

$$\beta_1 := (3 + \lambda)/4, \beta_2 := (4/(1 - \lambda) - 2), \lambda := \max\{|\lambda_2(A)|^2, |\lambda_N(A)|^2\}, M_{(t-1)} := \max\{\hat{F}(W^{(t-1)}), \hat{F}(\bar{w}^{(t-1)})\}.$$

*Proof.* By Lemma 9 and the inequality (15), the consensus error satisfies for any  $\beta > 0$ ,

$$\begin{aligned} \left\| W^{(t)} - \bar{W}^{(t)} \right\|_F^2 &< (1 + \beta)\lambda \left\| W^{(t-1)} - \bar{W}^{(t-1)} \right\|_F^2 \\ &+ 2(1 + \beta^{-1})\eta_{t-1}^2 \left\| \nabla \hat{F}(W^{(t-1)}) - \nabla \hat{F}(\bar{W}^{(t-1)}) \right\|_F^2 + 2(1 + \beta^{-1})\eta_{t-1}^2 \left\| \nabla \hat{F}(\bar{W}^{(t-1)}) \right\|_F^2. \end{aligned} \quad (21)$$

For the second term in (21), we have the following chain of inequalities,

$$\begin{aligned} &\left\| \nabla \hat{F}(W^{(t-1)}) - \nabla \hat{F}(\bar{W}^{(t-1)}) \right\|_F^2 \\ &= \sum_{\ell=1}^N \left\| \nabla \hat{F}_\ell(w_i^{(t-1)}) - \nabla \hat{F}_\ell(\bar{w}^{(t-1)}) \right\|^2 \\ &\leq \sum_{\ell=1}^N \max_{v_\ell \in [w_\ell^{(t-1)}, \bar{w}^{(t-1)}]} \left\| \nabla^2 \hat{F}_\ell(v_\ell) \right\|^2 \|w_\ell^{(t-1)} - \bar{w}^{(t-1)}\|^2 \end{aligned} \quad (22)$$

$$\leq h^2 \sum_{\ell=1}^N \max_{v_\ell \in [w_\ell^{(t-1)}, \bar{w}^{(t-1)}]} (\hat{F}_\ell(v_\ell))^2 \|w_\ell^{(t-1)} - \bar{w}^{(t-1)}\|^2 \quad (23)$$

$$\begin{aligned} &= h^2 \sum_{\ell=1}^N \left( \max_{v_\ell \in [w_\ell^{(t-1)}, \bar{w}^{(t-1)}]} \hat{F}_\ell(v_\ell) \right)^2 \|w_\ell^{(t-1)} - \bar{w}^{(t-1)}\|^2 \\ &\leq h^2 \sum_{\ell=1}^N \max\{\hat{F}_\ell^2(w_\ell^{(t-1)}), \hat{F}_\ell^2(\bar{w}^{(t-1)})\} \|w_\ell^{(t-1)} - \bar{w}^{(t-1)}\|^2 \end{aligned} \quad (24)$$

$$\begin{aligned} &\leq h^2 \max\left\{ \max_{k \leq N} \hat{F}_k^2(w_k^{(t-1)}), \max_{k \leq N} \hat{F}_k^2(\bar{w}^{(t-1)}) \right\} \sum_{\ell=1}^N \|w_\ell^{(t-1)} - \bar{w}^{(t-1)}\|^2 \\ &\leq h^2 N^2 M_{(t-1)}^2 \left\| W^{(t-1)} - \bar{W}^{(t-1)} \right\|_F^2. \end{aligned} \quad (25)$$

The Taylor's remainder theorem gives (22) and  $v_i \in [w_i^{(t-1)}, \bar{w}^{(t-1)}]$  denotes a point that lies on the line connecting  $w_i^{(t-1)}$  and  $\bar{w}^{(t-1)}$ . Also, (23) is valid due to the self-boundedness of the Hessian stated in Assumption 6. The inequality (24) follows by the assumption of convexity of  $\hat{F}_i$ , due to the fact that for a convex function  $f : \mathbb{R}^d \rightarrow \mathbb{R}$  and any two points  $w_1, w_2 \in \mathbb{R}^d$ , it holds that  $\max_{v \in [w_1, w_2]} f(v) \leq \max\{f(w_1), f(w_2)\}$ . To derive (25), we used  $\max_{i \leq N} \hat{F}_i(w_i) \leq N \hat{F}(W)$  and  $\max_{i \leq N} \hat{F}_i(\bar{w}) \leq N \cdot \hat{F}(\bar{w})$ , which hold since the loss functions are non-negative.

In order to derive an upper-bound on the last term in (21), we use Assumption 4 (with  $\alpha = 1, c = h$ ):

$$\left\| \nabla \hat{F}(\bar{W}^{(t-1)}) \right\|_F^2 = \sum_{\ell=1}^N \left\| \nabla \hat{F}_\ell(\bar{w}^{(t-1)}) \right\|^2 \leq h^2 \sum_{\ell=1}^N (\hat{F}_\ell(\bar{w}^{(t-1)}))^2 \leq h^2 N^2 \hat{F}^2(\bar{w}^{(t-1)}).$$

Replacing the upper-bounds back in (21), we conclude

$$\begin{aligned} \left\| W^{(t)} - \bar{W}^{(t)} \right\|_F^2 &< ((1 + \beta)\lambda + 2(1 + \beta^{-1})\eta_{t-1}^2 h^2 N^2 M_{(t-1)}^2) \left\| W^{(t-1)} - \bar{W}^{(t-1)} \right\|_F^2 \\ &\quad + 2(1 + \beta^{-1})\eta_{t-1}^2 h^2 N^2 \hat{F}^2(\bar{w}^{(t-1)}). \end{aligned}$$

Choose  $\beta = \frac{1-\lambda}{2\lambda}$ . Then by lemma's assumption  $\eta_{t-1} \leq \frac{1-\lambda}{4hNM_{(t-1)}}$ , we can verify the following two inequalities:

$$\begin{aligned} (1 + \beta)\lambda + 2(1 + \beta^{-1})\eta_{t-1}^2 h^2 N^2 M_{(t-1)}^2 &\leq \frac{3 + \lambda}{4}, \\ 2(1 + 1/\beta) &\leq \frac{4}{1 - \lambda} - 2. \end{aligned}$$

This concludes the proof.  $\square$

By recursively evaluating (20), we obtain a bound on the consensus error at iteration  $T$ , which we present next.

**Lemma 16** (Last iterate consensus error). *Under the assumptions and notations of Lemma 15, the consensus error at iteration  $T$  satisfies*

$$\left\| W^{(T)} - \bar{W}^{(T)} \right\|_F^2 < \beta_1^{T-1} \left\| W^{(1)} - \bar{W}^{(1)} \right\|_F^2 + \beta_2 h^2 N^2 \sum_{t=1}^{T-1} \beta_1^{t-1} \eta_{T-t}^2 \hat{F}^2(\bar{w}_{T-t}).$$

The next lemma obtains a sandwich relation between  $F(\bar{w}^{(T)})$  and  $\hat{F}(W^{(T)})$ . This is convenient as it allows replacing  $M_{(t)} := \max(\hat{F}(W^{(t)}), \hat{F}(\bar{w}^{(t)}))$  by either of the two terms with only paying a constant factor of two. See also the remark after the statement of the theorem.

**Lemma 17.** *Under the assumptions and notations of Lemma 15, with zero initialization  $W^{(1)} = \bar{W}^{(1)} = 0$  and by choosing  $\eta_t \leq \frac{(1-\lambda)\sqrt{1-\beta_1}}{8h^2NM_{(t)}\sqrt{\beta_2}}$  for  $t \in [T-1]$ , it holds at iteration  $T$  that*

$$\frac{1}{2} \hat{F}(\bar{w}^{(T)}) \leq \hat{F}(W^{(T)}) \leq 2\hat{F}(\bar{w}^{(T)}). \quad (26)$$

*Proof.* First, we prove  $\hat{F}(W^{(T)}) \leq 2\hat{F}(\bar{w}^{(T)})$ . If  $\hat{F}(W^{(T)}) \leq \hat{F}(\bar{w}^{(T)})$ , there is nothing to prove. Thus, assume  $\hat{F}(W^{(T)}) \geq \hat{F}(\bar{w}^{(T)})$ . Then by applying Taylor's remainder theorem, the self-boundedness Assumption 4 with  $c = h, \alpha = 1$ , convexity of  $\hat{F}$ , Lemma 16 and the restriction on the step-size, in respective order, we have the following inequalities,

$$\begin{aligned} \hat{F}(W^{(T)}) &\leq |\hat{F}(W^{(T)}) - \hat{F}(\bar{w}^{(T)})| + \hat{F}(\bar{w}^{(T)}) \\ &\leq \max_{v \in [\bar{W}^{(T)}, W^{(T)}]} \|\nabla \hat{F}(v)\| \cdot \|W^{(T)} - \bar{W}^{(T)}\| + \hat{F}(\bar{w}^{(T)}) \\ &\leq h \cdot \max_{v \in [\bar{W}^{(T)}, W^{(T)}]} \hat{F}(v) \cdot \|W^{(T)} - \bar{W}^{(T)}\| + \hat{F}(\bar{w}^{(T)}) \\ &\leq h \cdot \max\{\hat{F}(W^{(T)}), \hat{F}(\bar{w}^{(T)})\} \cdot \|W^{(T)} - \bar{W}^{(T)}\| + \hat{F}(\bar{w}^{(T)}) \\ &\leq \hat{F}(W^{(T)}) \left( \beta_2 h^4 N^2 \sum_{t=1}^{T-1} \beta_1^{t-1} \eta_{T-t}^2 \hat{F}^2(\bar{w}_{T-t}) \right)^{1/2} + \hat{F}(\bar{w}^{(T)}) \\ &\leq \frac{1}{2} \hat{F}(W^{(T)}) + \hat{F}(\bar{w}^{(T)}). \end{aligned}$$

Thus  $\hat{F}(W^{(T)}) \leq 2\hat{F}(\bar{w}^{(T)})$ . By exchanging  $W^{(T)}$  and  $\bar{w}^{(T)}$  and in a similar style we derive  $\hat{F}(W^{(T)}) \geq \frac{1}{2}\hat{F}(\bar{w}^{(T)})$ . This completes the proof of the lemma.  $\square$



*Remark 8.* Lemma 16 above requires tuning  $\eta_t \propto 1/M_{(t)} := 1/\max(\hat{F}(W^{(t)}), \hat{F}(\bar{w}^{(t)}))$ . Lemma 17 shows that abiding by this choice for  $t = 1, \dots, T-1$  and any  $T > 1$  guarantees  $\hat{F}(W^{(T)}) \leq 2\hat{F}(\bar{w}^{(T)})$ . Hence,  $M_{(T)} \geq 2\hat{F}(\bar{w}^{(T)})$ . Since this holds for all  $T$  and at  $t = 1$ ,  $\hat{F}(W^{(1)}) = \hat{F}(\bar{w}^{(1)})$ , it follows by recursion that Lemma 16 holds provided  $\eta_t \propto 1/\hat{F}(\bar{w}^{(t)})$ . We use observation in the proofs below.

We are ready to prove Theorem 4. First, we prove that DGD is a descent algorithm in the next lemma.

**Lemma 18** (Descent lemma). *Consider DGD under the assumptions and notations of Lemma 15. Moreover, let Assumption 7 hold, then by choosing  $\eta_t \leq \frac{\delta}{\hat{F}(\bar{w}^{(t)})}$  for  $t \leq T$ , where*

$$\delta := 1 / \max \left\{ \frac{4h^3 N}{\tau^2}, h^2, \frac{6h^2 \beta_2}{1 - \beta_1}, \frac{4h^2 \sqrt{\beta_2}}{\tau(1 - \beta_1)} \right\}, \quad (27)$$

DGD is a descent algorithm, i.e., for all  $T \geq 1$ .

$$\hat{F}(\bar{w}^{(T+1)}) \leq \hat{F}(\bar{w}^{(T)}).$$

*Proof.* With the self-boundedness assumption on the Hessian (Assumption 6) and applying the Taylor's remainder theorem for step  $T+1$  of DGD, we obtain the following,

$$\begin{aligned} & \hat{F}(\bar{w}^{(T+1)}) \\ & \leq \hat{F}(\bar{w}^{(T)}) + \langle \nabla \hat{F}(\bar{w}^{(T)}), \bar{w}^{(T+1)} - \bar{w}^{(T)} \rangle + \frac{1}{2} \max_{v \in [\bar{w}^{(T)}, \bar{w}^{(T+1)}]} \|\nabla^2 \hat{F}(v)\| \|\bar{w}^{(T+1)} - \bar{w}^{(T)}\|^2 \\ & \leq \hat{F}(\bar{w}^{(T)}) - \eta_T \langle \nabla \hat{F}(\bar{w}^{(T)}), \bar{\nabla} \hat{F}(W^{(T)}) \rangle + \frac{\eta_T^2}{2} \max_{v \in [\bar{w}^{(T)}, \bar{w}^{(T+1)}]} \|\nabla^2 \hat{F}(v)\| \|\bar{\nabla} \hat{F}(W^{(T)})\|^2 \\ & \leq \hat{F}(\bar{w}^{(T)}) - \eta_T \langle \nabla \hat{F}(\bar{w}^{(T)}), \bar{\nabla} \hat{F}(W^{(T)}) \rangle + \frac{h\eta_T^2}{2} \max_{v \in [\bar{w}^{(T)}, \bar{w}^{(T+1)}]} \hat{F}(v) \|\bar{\nabla} \hat{F}(W^{(T)})\|^2 \\ & \leq \hat{F}(\bar{w}^{(T)}) - \eta_T \langle \nabla \hat{F}(\bar{w}^{(T)}), \bar{\nabla} \hat{F}(W^{(T)}) \rangle + \frac{h\eta_T^2}{2} \max\{\hat{F}(\bar{w}^{(T)}), \hat{F}(\bar{w}^{(T+1)})\} \|\bar{\nabla} \hat{F}(W^{(T)})\|^2, \end{aligned} \quad (28)$$

where for the third step we used

$$\|\nabla^2 \hat{F}(w)\| \leq \frac{1}{N} \sum_{\ell=1}^N \|\nabla^2 \hat{F}_\ell(w)\| \leq \frac{h}{N} \sum_{\ell=1}^N \hat{F}_\ell(w) = h \hat{F}(w).$$

In the next step of the proof, we upper-bound the second and third terms in (28). For the second term, by noting that  $2\langle a, b \rangle = \|a\|^2 + \|b\|^2 - \|a - b\|^2$ , we can write

$$\langle \nabla \hat{F}(\bar{w}^{(T)}), \bar{\nabla} \hat{F}(W^{(T)}) \rangle = \frac{1}{2} \|\nabla \hat{F}(\bar{w}^{(T)})\|^2 + \frac{1}{2} \|\bar{\nabla} \hat{F}(W^{(T)})\|^2 - \frac{1}{2} \|\nabla \hat{F}(\bar{w}^{(T)}) - \bar{\nabla} \hat{F}(W^{(T)})\|^2. \quad (29)$$

By recalling (25) and Lemma 16 (which we can apply because of Remark 8), we find an upper-bound the last term in (29) as follows,

$$\begin{aligned} \|\nabla \hat{F}(\bar{w}^{(T)}) - \bar{\nabla} \hat{F}(W^{(T)})\|^2 & = \frac{1}{N^2} \|\nabla \hat{F}(W^{(T)}) - \nabla \hat{F}(\bar{W}^{(T)})\|_F^2 \\ & \leq h^2 M_{(T)}^2 \|W^{(T)} - \bar{W}^{(T)}\|_F^2 \\ & \leq h^4 M_{(T)}^2 N^2 \beta_2 \sum_{t=1}^{T-1} \beta_1^{t-1} \eta_{T-t}^2 \hat{F}^2(\bar{w}_{T-t}). \end{aligned} \quad (30)$$

Returning back to (28), thus far we have derived the following,

$$\begin{aligned} \hat{F}(\bar{w}^{(T+1)}) & \leq \hat{F}(\bar{w}^{(T)}) - \frac{\eta_T}{2} \|\nabla \hat{F}(\bar{w}^{(T)})\|^2 - \frac{\eta_T}{2} \|\bar{\nabla} \hat{F}(W^{(T)})\|^2 \\ & \quad + \frac{1}{2} \eta_T h^4 N^2 M_{(T)}^2 \beta_2 \sum_{t=1}^{T-1} \beta_1^{t-1} \eta_{T-t}^2 \hat{F}^2(\bar{w}_{T-t}) + \frac{h\eta_T^2}{2} \|\bar{\nabla} \hat{F}(W^{(T)})\|^2 \cdot \max\{\hat{F}(\bar{w}^{(T)}), \hat{F}(\bar{w}^{(T+1)})\}. \end{aligned} \quad (31)$$

We aim to prove that  $\hat{F}(\bar{w}^{(T+1)}) \leq \hat{F}(\bar{w}^{(T)})$  for all  $T \geq 1$ . If  $\hat{F}(\bar{w}^{(T+1)}) > \hat{F}(\bar{w}^{(T)})$ , applying (31) with the assumption  $\eta_t < \frac{\delta}{\hat{F}(\bar{w}^{(t)})}$  yields,

$$\begin{aligned} \hat{F}(\bar{w}^{(T+1)}) &\leq \hat{F}(\bar{w}^{(T)}) - \frac{\eta_T}{2} \|\nabla \hat{F}(\bar{w}^{(T)})\|^2 + \frac{1}{2(1-\beta_1)} \eta_T \delta^2 h^4 N^2 M_{(T)}^2 \beta_2 \\ &\quad + \frac{h\eta_T^2}{2} \|\bar{\nabla} \hat{F}(W^{(T)})\|^2 \cdot \hat{F}(\bar{w}^{(T+1)}). \end{aligned} \quad (32)$$

Note that it holds due to (30) that,

$$\begin{aligned} \|\bar{\nabla} \hat{F}(W^{(T)})\|^2 &\leq 2\|\nabla \hat{F}(\bar{w}^{(T)})\|^2 + 2\|\bar{\nabla} \hat{F}(W^{(T)}) - \nabla \hat{F}(\bar{w}^{(T)})\|^2 \\ &\leq 2\|\nabla \hat{F}(\bar{w}^{(T)})\|^2 + \frac{2}{1-\beta_1} \delta^2 h^4 N^2 M_{(T)}^2 \beta_2. \end{aligned}$$

Replacing this in (32) and noting that  $M_{(T)} \leq 2\hat{F}(\bar{w}^{(T)})$  by Lemma 17, we can simplify the inequality (32) as follows,

$$\begin{aligned} &\hat{F}(\bar{w}^{(T+1)}) \\ &\leq \hat{F}(\bar{w}^{(T)}) + \eta_T \|\nabla \hat{F}(\bar{w}^{(T)})\|^2 (h\eta_T \hat{F}(\bar{w}^{(T+1)}) - \frac{1}{2}) + C' h^2 N^2 \delta^2 \eta_T \hat{F}^2(\bar{w}^{(T)}) (1 + \eta_T \hat{F}(\bar{w}^{(T+1)})) \\ &\leq \hat{F}(\bar{w}^{(T)}) + \eta_T h^2 \hat{F}^2(\bar{w}^{(T)}) ((h\eta_T + \delta^2 \eta_T N^2 C') \hat{F}(\bar{w}^{(T+1)}) - \frac{\tau^2}{2h^2} + \delta^2 N^2 C'), \end{aligned}$$

where for the ease of notation we define  $C' := 4h^2\beta_2(1-\beta_1)^{-1}$ . Recalling  $\eta_T < \frac{\delta}{\hat{F}(\bar{w}^{(T)})}$  and noting that by the assumption of the lemma  $\delta \leq \frac{\tau^2}{4h^3}$ ,  $\delta < \frac{1}{h^2}$  and  $\delta < \frac{\tau}{2hN\sqrt{C'}}$  we conclude that,

$$\hat{F}(\bar{w}^{(T+1)}) \leq \hat{F}(\bar{w}^{(T)}) + \frac{\tau}{2h} \hat{F}(\bar{w}^{(T)}) \left( \frac{\hat{F}(\bar{w}^{(T+1)})}{\hat{F}(\bar{w}^{(T)})} - 1 \right).$$

Dividing both sides by  $\hat{F}(\bar{w}^{(T)})$  leads to the contradiction due to the fact that  $\tau \leq h$  and thus  $\tau/2h < 1$ . Thus  $\hat{F}(\bar{w}^{(T+1)}) \leq \hat{F}(\bar{w}^{(T)})$ . This completes the proof.  $\square$

### C.1 Proof of Theorem 4

**Theorem 19** (Restatement of Theorem 4). *Consider DGD with the loss functions and mixing matrix satisfying Assumptions 1, 2, 6, 7 and Assumption 4 with  $\alpha = 1$  and  $c = h$ . Assume that the step-size satisfies  $\eta < \frac{\delta}{\hat{F}(1)}$  for  $\delta$  defined in (27). Also, recall positive constants  $\beta_1, \beta_2$  depending only on the mixing matrix as defined in Lemma 15. Then DGD is a descent algorithm i.e. for all  $T \geq 1$  it holds that*

$$\hat{F}(\bar{w}^{(T+1)}) \leq \hat{F}(\bar{w}^{(T)}).$$

Moreover, the train loss and the consensus error of DGD at iteration  $T$  satisfy the following for all  $w \in \mathbb{R}^d$ ,

$$\hat{F}(\bar{w}^{(T)}) \leq 4\hat{F}(w) + \frac{2\|w\|^2}{\eta T}, \quad (33)$$

$$\frac{1}{N^2} \|W^{(T)} - \bar{W}^{(T)}\|_F^2 \leq \frac{8\beta_2 h^2}{1-\beta_1} (4\eta^2 \hat{F}^2(w) + \frac{\|w\|^4}{T^2}). \quad (34)$$

*Proof.* First, we note that by Lemma 18, under the assumption  $\eta_t \leq \delta/\hat{F}(\bar{w}^{(t)})$  for  $t \leq T$ , we have  $\hat{F}(\bar{w}^{(T+1)}) \leq \hat{F}(\bar{w}^{(T)})$ . Thus fixing  $\eta \leq \delta/\hat{F}(\bar{w}^{(1)})$ , ensures that  $\hat{F}(\bar{w}^{(T+1)}) \leq \hat{F}(\bar{w}^{(T)})$ , for all  $T$ .

Next, we derive the train loss and consensus error under the assumptions of the theorem. Start with,

$$\|\bar{w}^{(t+1)} - w\|^2 = \|\bar{w}^{(t)} - w\|^2 + \eta^2 \|\bar{\nabla} \hat{F}(W^{(t)})\|^2 - 2\eta \langle \bar{w}^{(t)} - w, \bar{\nabla} \hat{F}(W^{(t)}) \rangle. \quad (35)$$

For the second term, by self-boundedness of gradient, we can write,

$$\|\bar{\nabla} \hat{F}(W^{(t)})\| = \frac{1}{n} \left\| \sum_{\ell=1}^n \nabla \hat{F}_\ell(w_\ell^{(t)}) \right\| \leq \frac{h}{n} \sum_{\ell=1}^n \hat{F}_\ell(w_\ell^{(t)}) = h\hat{F}(W^{(t)}).$$

For the third term in (35), we have,

$$\begin{aligned}
 -\langle \bar{w}^{(t)} - w, \bar{\nabla} \hat{F}(W^{(t)}) \rangle &= -\frac{1}{N} \sum_{\ell=1}^N \langle \bar{w}^{(t)} - w, \nabla \hat{F}_\ell(w_\ell^{(t)}) \rangle \\
 &= -\frac{1}{N} \sum_{\ell=1}^N \langle \bar{w}^{(t)} - w_\ell^{(t)}, \nabla \hat{F}_\ell(w_\ell^{(t)}) \rangle - \frac{1}{N} \sum_{\ell=1}^N \langle w_\ell^{(t)} - w, \nabla \hat{F}_\ell(w_\ell^{(t)}) \rangle \\
 &\leq \frac{1}{N} \sum_{\ell=1}^N \|w_\ell^{(t)} - \bar{w}^{(t)}\| \|\nabla \hat{F}_\ell(w_\ell^{(t)})\| - \frac{1}{N} \sum_{\ell=1}^N \langle w_\ell^{(t)} - w, \nabla \hat{F}_\ell(w_\ell^{(t)}) \rangle \\
 &\leq \frac{1}{N} \sum_{\ell=1}^N \|w_\ell^{(t)} - \bar{w}^{(t)}\| \|\nabla \hat{F}_\ell(w_\ell^{(t)})\| + \frac{1}{N} \sum_{\ell=1}^N (\hat{F}_\ell(w) - \hat{F}_\ell(w_\ell^{(t)})) \tag{36}
 \end{aligned}$$

$$\begin{aligned}
 &\leq \frac{h}{N} \sum_{\ell=1}^N \hat{F}_\ell(w_\ell^{(t)}) \|w_\ell^{(t)} - \bar{w}^{(t)}\| + \frac{1}{N} \sum_{\ell=1}^N (\hat{F}_\ell(w) - \hat{F}_\ell(w_\ell^{(t)})) \tag{37} \\
 &\leq h \hat{F}(W^{(t)}) \|W^{(t)} - \bar{W}^{(t)}\|_F + \hat{F}(w) - \hat{F}(W^{(t)}).
 \end{aligned}$$

Here (36) follows by convexity of  $\hat{F}_i$ , and (37) follows by the assumption on self-boundedness of the gradient.

Thus, the inequality (35) can be written as follows,

$$\|\bar{w}^{(t+1)} - w\|^2 \leq \|\bar{w}^{(t)} - w\|^2 + \eta^2 h^2 \hat{F}^2(W^{(t)}) + 2\eta h \hat{F}(W^{(t)}) \|W^{(t)} - \bar{W}^{(t)}\|_F + 2\eta \hat{F}(w) - 2\eta \hat{F}(W^{(t)}). \tag{38}$$

Moreover, by Lemma 16 and the assumption on  $\eta$ ,

$$\|W^{(t)} - \bar{W}^{(t)}\|_F \leq (\beta_2 h^2 \sum_{t=1}^{T-1} \beta_1^{t-1} \eta_{T-t}^2 \hat{F}^2(\bar{w}_{T-t}))^{1/2} \leq \frac{1}{4h}$$

and

$$\eta \hat{F}(W^{(t)}) \leq \frac{1}{2h^2}.$$

Thus (38) changes into,

$$\|\bar{w}^{(t+1)} - w\|^2 \leq \|\bar{w}^{(t)} - w\|^2 - \eta \hat{F}(W^{(t)}) + 2\eta \hat{F}(w).$$

Telescoping sum leads to

$$\frac{1}{T} \sum_{t=1}^T \hat{F}(W^{(t)}) \leq 2\hat{F}(w) + \frac{\|\bar{w}^{(1)} - w\|^2}{\eta T}. \tag{39}$$

By Lemma 17, we have  $\hat{F}(\bar{w}^{(t)}) \leq 2\hat{F}(W^{(t)})$ . Finally, as we proved in the beginning, DGD is a descent algorithm, implying

$$\hat{F}(\bar{w}^{(T)}) \leq \frac{1}{T} \sum_{t=1}^T \hat{F}(\bar{w}^{(t)})$$

In view of (39), this yields the claim of the theorem for the train loss (33). Finally, appealing to Lemma 16, gives (34). This completes the proof of the theorem.  $\square$

## D Proofs for Section 2.3

**Lemma 20** (Train loss under PL condition). *Let Assumptions 1, 3 and 8 hold, and let*

$\eta \leq \min\{\frac{1}{\mu}, \sqrt{\frac{(1-\alpha_1)\mu}{4L^4\alpha_2}}, \frac{1}{L}\}$  and  $\bar{\zeta} := \max\{\frac{1+\alpha_1}{2}, 1 - \frac{\eta\mu}{2}\}$ , where  $\alpha_1 := \frac{3+\lambda}{4}$ ,  $\alpha_2 := 4(\frac{2}{1-\lambda} - 1)$  same as in Lemma 9, Then for  $t \geq 1$

$$\hat{F}(\bar{w}^{(t)}) \leq \bar{\zeta}^{t-1} \hat{F}(\bar{w}^{(1)}). \tag{40}$$

*Proof.* By  $L$ -smoothness we have

$$\begin{aligned}
 \hat{F}(\bar{w}^{(t+1)}) &\leq \hat{F}(\bar{w}^{(t)}) - \eta \langle \nabla \hat{F}(\bar{w}^{(t)}), \bar{\nabla} \hat{F}(W^{(t)}) \rangle + \frac{\eta^2 L}{2} \|\bar{\nabla} \hat{F}(W^{(t)})\|^2 \\
 &= \hat{F}(\bar{w}^{(t)}) - \frac{\eta - \eta^2 L}{2} \|\bar{\nabla} \hat{F}(W^{(t)})\|^2 - \frac{\eta}{2} \|\nabla \hat{F}(\bar{w}^{(t)})\|^2 + \frac{\eta}{2} \|\bar{\nabla} \hat{F}(W^{(t)}) - \nabla \hat{F}(\bar{w}^{(t)})\|^2 \\
 &\leq \hat{F}(\bar{w}^{(t)}) - \frac{\eta}{2} \|\nabla \hat{F}(\bar{w}^{(t)})\|^2 + \frac{\eta}{2} \|\bar{\nabla} \hat{F}(W^{(t)}) - \nabla \hat{F}(\bar{w}^{(t)})\|^2 \\
 &\leq \hat{F}(\bar{w}^{(t)}) - \frac{\eta}{2} \|\nabla \hat{F}(\bar{w}^{(t)})\|^2 + \frac{\eta L^2}{N} \|W^{(t)} - \bar{W}^{(t)}\|_F^2.
 \end{aligned}$$

By  $\mu$ -PL condition we have,

$$\hat{F}(\bar{w}^{(t+1)}) \leq (1 - \eta\mu) \hat{F}(\bar{w}^{(t)}) + \frac{\eta L^2}{N} \|W^{(t)} - \bar{W}^{(t)}\|_F^2.$$

By Lemma 9,

$$\frac{1}{N} \|W^{(t)} - \bar{W}^{(t)}\|_F^2 < \alpha_2 \eta^2 L \sum_{i=1}^{t-1} \alpha_1^{i-1} \hat{F}(\bar{w}^{(t-i)}). \quad (41)$$

which results in,

$$\hat{F}(\bar{w}^{(t+1)}) \leq (1 - \eta\mu) \hat{F}(\bar{w}^{(t)}) + \alpha_2 L^3 \eta^3 \sum_{i=1}^{t-1} \alpha_1^{i-1} \hat{F}(\bar{w}^{(t-i)}).$$

By induction assume  $\hat{F}(\bar{w}^{(t)}) \leq \bar{\zeta}^{t-1} \hat{F}(\bar{w}^{(1)})$  then using the assumptions on  $\bar{\zeta}$  and  $\eta$  yield the following inequalities,

$$\begin{aligned}
 \hat{F}(\bar{w}^{(t+1)}) &\leq (1 - \eta\mu) \bar{\zeta}^{t-1} \hat{F}(\bar{w}^{(1)}) + \alpha_2 \eta^3 L^3 \hat{F}(\bar{w}^{(1)}) \sum_{i=1}^{t-1} \alpha_1^{i-1} \bar{\zeta}^{t-i-1} \\
 &\leq (1 - \eta\mu) \bar{\zeta}^{t-1} \hat{F}(\bar{w}^{(1)}) + \alpha_2 \eta^3 L^3 \hat{F}(\bar{w}^{(1)}) \frac{\bar{\zeta}^{t-2}}{1 - \alpha_1/\bar{\zeta}} \\
 &= (1 - \eta\mu + \alpha_2 \eta^3 L^3 / (\bar{\zeta} - \alpha_1)) \bar{\zeta}^{t-1} \hat{F}(\bar{w}^{(1)}) \\
 &\leq (1 - \eta\mu + 2\alpha_2 \eta^3 L^3 / (1 - \alpha_1)) \bar{\zeta}^{t-1} \hat{F}(\bar{w}^{(1)}) \\
 &\leq (1 - \eta\mu + \eta\mu/2) \bar{\zeta}^{t-1} \hat{F}(\bar{w}^{(1)}) \\
 &\leq \bar{\zeta}^t \hat{F}(\bar{w}^{(1)}).
 \end{aligned}$$

This completes the proof of the lemma.  $\square$

## D.1 Proof of Lemma 5

**Lemma 21** (Restatement of Lemma 5). *Let Assumptions 1, 3 and 8 hold and let the step-size  $\eta \leq \min\{\frac{1-\alpha_1}{\mu}, \frac{1}{2L^2} \sqrt{\frac{(1-\alpha_1)\mu}{\alpha_2}}, \frac{1}{L}\}$ , where the constants  $\alpha_1 \in (0, 1)$  and  $\alpha_2 > 0$  are defined same as in Lemma 20. Define  $\zeta := 1 - \frac{\eta\mu}{2}$ , then under the data separability assumption, the iterates of DGD satisfy for all  $t \geq 1$ ,*

$$\begin{aligned}
 \hat{F}(\bar{w}^{(t)}) &\leq \zeta^{t-1} \hat{F}(\bar{w}^{(1)}), \\
 \frac{1}{N} \|W^{(t)} - \bar{W}^{(t)}\|_F^2 &\leq \frac{2\alpha_2 \eta^2 L^2 \hat{F}(\bar{w}^{(1)})}{1 - \alpha_1} \zeta^{t-1}.
 \end{aligned}$$

*Proof.* The bound on the train loss follows directly by Lemma 20, after noting that  $\eta \leq \frac{1-\alpha_1}{\mu}$  implies  $\frac{1+\alpha_1}{2} \leq 1 - \eta\mu/2$ . The consensus error is derived by (41) and using the bound on  $\hat{F}(\bar{w}^{(t)})$ .  $\square$

## D.2 Proof of Theorem 6

**Theorem 22** (Restatement of Theorem 6). *Let Assumptions 1-4 and 8 hold, and let  $\eta$  and  $\zeta$  be as in Lemma 20. Then the iterates of DGD under the data separability assumption satisfy for all  $T \geq 1$ ,*

$$\mathbb{E}\left[F(\bar{w}^{(T)})\right] = O\left(\zeta^T + \frac{L^2 c^2}{n^{3-2\alpha} \mu^{2\alpha}} (\eta T)^{2-2\alpha} + \frac{\eta^2 L^4}{\mu^2 N}\right).$$

*Proof.* By simplifying Lemma 1 using the convergence bounds in Lemma 21, we end up with the following,

$$\mathbb{E}\left[F(\bar{w}^{(T)})\right] = O\left(\zeta^T + \frac{L^2 c^2 \eta^2}{n^{3-2\alpha} (1-\zeta)^{2\alpha}} T^{2-2\alpha} + \frac{\eta^2 L^4}{(1-\sqrt{\zeta})^2 N}\right).$$

Based on the definition of  $\zeta$ , we have  $(\frac{\eta}{1-\zeta})^{2\alpha} = (\frac{2}{\mu})^{2\alpha}$  and  $\frac{\eta^2}{(1-\sqrt{\zeta})^2} \leq \frac{4}{\mu^2}$ . This proves the statement of the theorem.  $\square$

## E Proof of Theorem 7

**Theorem 23** (Restatement of Theorem 7). *Consider FDRL(Algorithm 1) on separable dataset, and choose  $\eta = O(1/\sqrt{t})$ . Then for all  $\ell \in [N]$*

$$\lim_{t \rightarrow \infty} \frac{w_\ell^{(t)}}{\|w_\ell^{(t)}\|} = \frac{w_{\text{MM}}}{\|w_{\text{MM}}\|},$$

where recall that  $w_{\text{MM}}$  denotes the solution to hard-margin SVM problem.

*Proof.* Replace  $\frac{v_\ell^{(t)}}{\|v_\ell^{(t)}\|}$  in step 2 of Algorithm 1 by arbitrary perturbations  $\varepsilon_\ell^{(t)}$  of unit norm. Then note that the sequence  $\{w_\ell^{(t)}\}$  generated by step 2 is identical to decentralized GD with  $\eta \|\varepsilon_\ell^{(t)}\| \rightarrow 0$ . Thus by [Nedić and Olshevsky, 2014, Lemma 1], consensus is asymptotically achieved for all  $\ell \in [N]$ , i.e.,

$$\lim_{t \rightarrow \infty} \|w_\ell^{(t)} - \bar{w}^{(t)}\| = 0.$$

Thus

$$\lim_{t \rightarrow \infty} \|w_\ell^{(t+1)} - w_\ell^{(t)}\| = \lim_{t \rightarrow \infty} \|\bar{w}^{(t+1)} - \bar{w}^{(t)}\| = \lim_{t \rightarrow \infty} \eta \|\bar{\varepsilon}^{(t)}\| = 0.$$

This implies that for all  $i \in [N]$  we have  $\lim_{t \rightarrow \infty} \|\nabla \hat{F}_\ell(w_\ell^{(t+1)}) - \nabla \hat{F}_\ell(w_\ell^{(t)})\| = 0$ , thus by appealing again to [Nedić and Olshevsky, 2014, Lemma 1] and applying it to step (5) of Algorithm 1, we find that,

$$\lim_{t \rightarrow \infty} \|v_\ell^{(t)} - \bar{v}^{(t)}\| = 0.$$

Aggregations of gradients in step (5) imply that  $\bar{v}^{(t)} = \frac{\mathbf{1}^\top \nabla \mathcal{L}(W^{(t)})}{N} \rightarrow \nabla \hat{F}(\bar{w}^{(t)})$ . Thus step (2) of FDLR for every agent  $i$  converges to  $\bar{w}^{(t)} - \eta \frac{\nabla \hat{F}(\bar{w}^{(t)})}{\|\nabla \hat{F}(\bar{w}^{(t)})\|}$ , i.e.,

$$\left\| \left( w_\ell^{(t)} - \eta \frac{v_\ell^{(t)}}{\|v_\ell^{(t)}\|} \right) - \left( \bar{w}^{(t)} - \eta \frac{\nabla \hat{F}(\bar{w}^{(t)})}{\|\nabla \hat{F}(\bar{w}^{(t)})\|} \right) \right\| \xrightarrow{t \rightarrow \infty} 0.$$

Thus for all  $\ell$ , the sequence  $\{w_\ell^{(t)}\}$  converges to the solution of normalized GD, i.e., the max-margin separator  $w_{\text{MM}}$ , for linearly separable datasets ([Nacson et al., 2019, Theorem 5]). This leads to the statement of the theorem.  $\square$

## F Auxiliary Results

**Proposition 24** (Bounds on the exponential loss). *Consider linear classification with the exponential loss  $f(w, (a, y)) = \exp(-y \cdot w^\top a)$  over linearly separable dataset  $(a_i, y_i)_{i=1}^n$  with binary labels  $y_i$  and with  $\max_i \|a_i\| \leq r$  for a constant  $r$ . The training loss in this case satisfies for all  $w \in \mathbb{R}^d$ ,*

$$\|\nabla \hat{F}(w)\| \in [c' F(w), c F(w)], \quad \|\nabla^2 \hat{F}(w)\| \leq h F(w), \quad (42)$$

for constants  $c, c'$  and  $h$  independent of  $w$ .

*Proof.* Using  $\hat{F}(w) = \frac{1}{n} \sum_{i=1}^n \exp(-y_i \cdot w^\top a_i)$ , one can deduce that,

$$\begin{aligned}\nabla \hat{F}(w) &= -\frac{1}{n} \sum_{i=1}^n y_i a_i \exp(-y_i w^\top a_i), \\ \nabla^2 \hat{F}(w) &= \frac{1}{n} \sum_{i=1}^n a_i a_i^\top \exp(-y_i w^\top a_i).\end{aligned}$$

Therefore it holds that,

$$\|\nabla \hat{F}(w)\| = \frac{1}{n} \left\| \sum_{i=1}^n y_i a_i \exp(-y_i w^\top a_i) \right\| \leq \frac{1}{n} \sum_{i=1}^n \|y_i a_i \exp(-y_i w^\top a_i)\| = \frac{1}{n} \sum_{i=1}^n \|y_i a_i\| \exp(-y_i w^\top a_i) \leq r \hat{F}(w).$$

A similar approach for the Hessian of  $\hat{F}$  results in the following inequality,

$$\|\nabla^2 \hat{F}(w)\| \leq r^2 \hat{F}(w).$$

Moreover, due to linear separability there exists a  $w^* \in \mathbb{R}^d$  such that,

$$\frac{y_i w^{*\top} a_i}{\|w^*\|} \geq \gamma, \quad \forall i \in [N],$$

where  $\gamma > 0$  denotes the margin. Therefore, using the supremum definition of norm we can write,

$$\begin{aligned}\|\nabla \hat{F}(w)\| &= \frac{1}{n} \left\| \sum_{i=1}^n y_i a_i \exp(-y_i w^\top a_i) \right\| \\ &= \sup_{\substack{v \in \mathbb{R}^d \\ \text{s.t. } \|v\|=1}} \left\langle \frac{1}{n} \sum_{i=1}^n y_i a_i \exp(-y_i w^\top a_i), v \right\rangle \\ &\geq \left\langle \frac{1}{n} \sum_{i=1}^n y_i a_i \exp(-y_i w^\top a_i), \frac{w^*}{\|w^*\|} \right\rangle \\ &\geq \frac{1}{n} \sum_{i=1}^n \gamma \cdot \exp(-y_i w^\top a_i) \\ &= \gamma \hat{F}(w).\end{aligned}$$

This completes the proof. □

**Proposition 25** (Bounds on the logistic loss). *Consider linear classification with the logistic loss  $f(w, (a, y)) = \log(1 + \exp(-y \cdot w^\top a))$  over linearly separable dataset  $(a_i, y_i)_{i=1}^n$  with binary labels  $y_i$  and with  $\max_i \|a_i\| \leq r$  for a constant  $r$ . The training loss in this case satisfies for all  $w \in \mathbb{R}^d$ ,*

$$\|\nabla \hat{F}(w)\| \in [c' \Phi(w), cF(w)], \quad \|\nabla^2 \hat{F}(w)\| \leq hF(w),$$

for  $\Phi(w) := \frac{1}{n} \sum_{i=1}^n \frac{\exp(-y_i w^\top a_i)}{1 + \exp(-y_i w^\top a_i)}$  and constants  $c, c'$  and  $h$  independent of  $w$ .

*Proof.* The training loss is now  $\hat{F}(w) = \frac{1}{n} \sum_{i=1}^n \log(1 + \exp(-y_i \cdot w^\top a_i))$ . Thus,

$$\begin{aligned}\nabla \hat{F}(w) &= \frac{1}{n} \sum_{i=1}^n (-y_i a_i) \frac{\exp(-y_i w^\top a_i)}{1 + \exp(-y_i w^\top a_i)}, \\ \nabla^2 \hat{F}(w) &= \frac{1}{n} \sum_{i=1}^n a_i a_i^\top \frac{\exp(-y_i w^\top a_i)}{(1 + \exp(-y_i w^\top a_i))^2}.\end{aligned}$$



By considering the norm and noting that  $\exp(t)/(1 + \exp(t)) \leq \log(1 + \exp(t))$ ,

$$\begin{aligned} \|\nabla \hat{F}(w)\| &= \frac{1}{n} \left\| \sum_{i=1}^n (-y_i a_i) \frac{\exp(-y_i w^\top a_i)}{1 + \exp(-y_i \cdot w^\top a_i)} \right\| \\ &\leq \frac{1}{n} \sum_{i=1}^n \|y_i a_i\| \frac{\exp(-y_i w^\top a_i)}{1 + \exp(-y_i \cdot w^\top a_i)} \\ &\leq \frac{r}{n} \sum_{i=1}^n \log(1 + \exp(-y_i w^\top a_i)) = r \hat{F}(w). \end{aligned}$$

Likewise, since  $\exp(t)/(1 + \exp(t))^2 \leq 2 \log(1 + \exp(t))$ , we can conclude that the operator norm of the Hessian satisfies,

$$\nabla^2 \hat{F}(w) \leq 2r^2 \hat{F}(w).$$

This completes the proof of upper-bounds for the gradient and Hessian. For the lower-bound on gradient note that by using the supremum definition of norm and recalling the max-margin separator satisfies  $\frac{y_i w^{\star \top} a_i}{\|w^*\|} \geq \gamma$  for the margin  $\gamma > 0$  and all  $i \in [n]$ , we can write,

$$\begin{aligned} \|\nabla \hat{F}(w)\| &= \frac{1}{n} \left\| \sum_{i=1}^n y_i a_i \frac{\exp(-y_i w^\top a_i)}{1 + \exp(-y_i \cdot w^\top a_i)} \right\| = \sup_{\substack{v \in \mathbb{R}^d \\ \text{s.t. } \|v\|=1}} \left\langle \frac{1}{n} \sum_{i=1}^n y_i a_i \frac{\exp(-y_i w^\top a_i)}{1 + \exp(-y_i \cdot w^\top a_i)}, v \right\rangle \\ &\geq \left\langle \frac{1}{n} \sum_{i=1}^n y_i a_i \frac{\exp(-y_i w^\top a_i)}{1 + \exp(-y_i \cdot w^\top a_i)}, \frac{w^*}{\|w^*\|} \right\rangle \\ &\geq \frac{1}{n} \sum_{i=1}^n \gamma \frac{\exp(-y_i w^\top a_i)}{1 + \exp(-y_i \cdot w^\top a_i)}. \end{aligned}$$

This yields the lower bound  $\gamma \Phi(w)$  on the norm of gradient and completes the proof.  $\square$

**Proposition 26** (Realizability of the exponential and logistic loss [Schliserman and Koren, 2022]). *On linearly separable data with margin  $\gamma > 0$ , the exponential loss function satisfies the realizability assumption (Assumption 5) with  $\rho(\varepsilon) = -\frac{1}{\gamma} \log(\varepsilon)$ , where  $\gamma$  denotes the margin. Moreover, the logistic loss function satisfies the realizability assumption with  $\rho(\varepsilon) = -\frac{1}{\gamma} \log(\exp(\varepsilon) - 1)$ .*

## G Additional experiments

### G.1 Experiments on over-parameterized Least-squares

In Fig. 4, we conduct experiments for highly over-parameterized Least-squares ( $f(w, x) = (1 - w^\top x)^2$ ), where  $d$  is typically significantly larger than  $n$  to ensure perfect interpolation of dataset. Note that, the train loss is not strongly-convex in this case, instead it satisfies the PL condition (Assumption 8). Notably, as predicted by Lemma 5, we notice the linear convergence of the train loss and the consensus error in Fig. 4 (Left). On the other hand, for the test loss, we observe its remarkably fast convergence (after approximately 50 iterations) to the optimal value, which is followed by a sharp increase in the subsequent iterations.

### G.2 On the update rule of FDLR

In the final section of the paper, we state a remark regarding the update rule of FDLR. Recall the update rule of DGD,

$$w_\ell^{(t+1)} = \sum_{k \in \mathcal{N}_\ell} A_{\ell k} w_k^{(t)} - \eta_t \nabla \hat{F}_\ell(w_\ell^{(t)}). \quad (43)$$

Notably, we expect FDLR to be perhaps the simplest approach for accommodating normalized gradients in decentralized learning setting since in DGD the agents only have access to their local gradients. In particular consider a Normalized DGD algorithm with the same update as in (43) but with  $\nabla \hat{F}_i(w_i^{(t)})$  replaced by  $\nabla \hat{F}_i(w_i^{(t)}) / \|\nabla \hat{F}_i(w_i^{(t)})\|$ , i.e.,




Figure 4: Consensus error, train loss and test loss for DGD with over-parameterized least-squares(square loss). The test loss achieves its optimum at the very early stages of DGD.




Figure 5: Normalized DGD with the update rule in Eq.(44) for different step-sizes  $\eta$  compared to DGD (Eq.(43)) and to FDLR (Alg 1). The step-sizes for DGD and FDLR are fine-tuned so that best of each algorithm is depicted. Normalized DGD cannot outperform DGD while FDLR is significantly faster than DGD. Here we consider linear classification with the exponential loss function and the dataset is generated according to signed measurements with Gaussian features and  $n = 100, d = 50$ .

$$w_\ell^{(t+1)} = \sum_{k \in \mathcal{N}_\ell} A_{\ell k} w_k^{(t)} - \eta_t \frac{\nabla \hat{F}_\ell(w_\ell^{(t)})}{\|\nabla \hat{F}_\ell(w_\ell^{(t)})\|} \quad (44)$$

The Normalized DGD algorithm above does not lead to faster convergence. This is due to the fact that in DGD the local gradient norm  $\|\nabla \hat{F}_i(w_i^{(t)})\|$  can be different than the global gradient norm  $\|\nabla \hat{F}(w_i^{(t)})\|$ . Thus even if with the update rule (44) the local parameters  $w_i^{(t)}$  converge to the global optimal solution, still the update rule for the averaged parameter  $\bar{w}^{(t)}$  is different than the update rule of centralized normalized GD. Our numerical experiment in Fig. 5 demonstrates the incapability of Normalized DGD in speeding up DGD. In particular, we note that for any choice of step-size Normalized DGD does not lead to acceleration compared to DGD whereas FDLR massively outperforms DGD.

## H A note about convergence rates of DGD

As mentioned in the paper’s introduction, many prior works on investigate convergence of DGD and of its stochastic variant decentralized stochastic gradient descent (DSGD) under various assumptions, e.g. [Jiang et al., 2017, Wang et al.,

2019, Koloskova et al., 2019, Koloskova et al., 2020] and many references therein. Most recently, [Koloskova et al., 2020] has presented a powerful unifying analysis of DSGD under rather weak assumptions. Specialized to convex  $L$ -smooth functions for which there exists  $w^*$  such that  $\|\nabla f_i(w^*)\| = 0$  (i.e. interpolation) [Koloskova et al., 2020, Thm. 2] shows a rate of  $\mathcal{O}(LR_0/T)$  for average DSGD updates. Here,  $R_0 = \|w_1 - w^*\|_2$ . Ignoring logarithmic factors, this rate is the same as what we obtained in (6) (as a consequence of Lemma 2) for DGD specifically applied to logistic loss over separable data. However, our result does *not* directly follow from [Koloskova et al., 2020, Thm. 2]. The reason is that logistic loss on separable data does *not* attain a bounded estimator. In fact, we believe the  $\log^2 T$  dependence of the rate that shows up in our analysis (see Eq. (5)), is a consequence of the infinitely normed-optimizers in our setting and we expect the bound to be tight as suggested by our experiments (see Fig 3) and in agreement with convergence bounds for logistic regression on separable data in the centralized case derived recently in [Ji and Telgarsky, 2018, Thm. 1.1]. On the other hand, the results of [Koloskova et al., 2020] are applicable to finite optimizers, which yields  $\mathcal{O}(1/T)$  convergence rates without log factor. Besides the above, in Theorem 4, we prove novel last-iterate (as opposed to averaged in the literature) convergence bounds for the train loss and faster consensus error rates of  $\tilde{\mathcal{O}}(1/T^2)$ . This is possible by leveraging additional Hessian self-bounded (Ass. 6) and self-lower-bounded (Ass. 7) assumptions, which hold for example for the exponential loss. Finally, we recall that our main focus is on studying finite time *generalization* bounds for DGD (e.g. Thm. 3), which to the best of our knowledge are new in this setting. Having discussed these, it is worth noting that the analysis of [Koloskova et al., 2020] applies under a relaxed assumption on the mixing matrix (see [Koloskova et al., 2020, Ass. 4]) than the corresponding assumptions (e.g. Assumption 1) in the literature. For example, this relaxed assumption covers decentralized local SGD (with multiple local updates per iteration) as a special case and is interesting to extend our results (on logistic regression over separable data) to such settings.