

---

# Semi-Cyclic Stochastic Gradient Descent

---

Hubert Eichner<sup>1</sup> Tomer Koren<sup>1</sup> H. Brendan McMahan<sup>1</sup> Nathan Srebro<sup>2</sup> Kunal Talwar<sup>1</sup>

## Abstract

We consider convex SGD updates with a block-cyclic structure, i.e., where each cycle consists of a small number of blocks, each with many samples from a possibly different, block-specific, distribution. This situation arises, e.g., in Federated Learning where the mobile devices available for updates at different times during the day have different characteristics. We show that such block-cyclic structure can significantly deteriorate the performance of SGD, but propose a simple approach that allows prediction with the same guarantees as for i.i.d., non-cyclic, sampling.

## 1. Introduction

When using Stochastic Gradient Descent (SGD), it is important that samples are drawn at random. In contrast, cycling over a specific permutation of data points (e.g., sorting data by label, or even alphabetically or by time of collection) can be detrimental to the performance of SGD. But in some applications, such as in Federated Learning (Konečný et al., 2016; McMahan et al., 2017) or when running SGD in “real time”, some cyclic patterns in the data samples are hard to avoid. In this paper we investigate how such cyclic structure hurts the effectiveness of SGD, and how this can be corrected. We model this semi-cyclicity through a “block-cyclic” structure, where during the course of training, we cycle over blocks in a fixed order and samples for each block are drawn from a block-specific distribution.

Our primary motivation is Federated Learning. In this setting, mobile devices collaborate in the training of a shared model while keeping the training data decentralized. Devices communicate updates (e.g., gradients) to a coordinating server, which aggregates the updates and applies them to the global model. In each iteration (or round) of Federated

<sup>1</sup>Google. <sup>2</sup>Toyota Technological Institute at Chicago. Part of this work was done while NS was visiting Google. Authors are listed alphabetically. Correspondence to: <{huberte,tkoren,mcmahan,kunal}@google.com, nati@ttic.edu>.

SGD, typically a few hundred devices are chosen randomly by the server to participate; critically, however, only devices that are idle, charging, and on a free wireless connection are selected (McMahan & Ramage, 2017; Bonawitz et al., 2019). This ensures Federated Learning does not impact the user’s experience when using their phone, but also can produce significant diurnal variation in the devices available, since devices are more likely to meet the training eligibility requirements at night local time. For example, for a language model (Hard et al., 2018), devices from English speakers in India and America are likely available at different times of day.

When the number of blocks in each cycle is high and the number of samples per block is low, the setting approaches fully-cyclic SGD, which is known to be problematic. But what happens when the number of blocks is fairly small and each blocks consists of a large number of samples? E.g., if there are only two blocks corresponding to “day” and “night”, with possibly millions of samples per block. An optimist would hope that this “slightly cyclic” case is much easier than the fully cyclic case, and the performance of SGD degrades gracefully as the number of blocks increases.

Unfortunately, in Section 3 we show that even with only two blocks, the block-cyclic sampling can cause an arbitrarily bad slowdown. One might ask whether alternative forms of updates, instead of standard stochastic gradient updates, can alleviate this degradation in performance. We show that such a slowdown is unavoidable for any iterative method based on semi-cyclic samples.

Instead, the solution we suggest is to embrace the heterogeneity in the data and resort to a *pluralistic* solution, allowing a potentially different model for each block in the cycle (e.g., a “daytime” model and a “nighttime” model). A naive pluralistic approach would still suffer a significant slowdown as it would not integrate information between blocks (Section 4). In Section 5 we show a remarkably simple and practical pluralistic approach that allows us to obtain exactly the same guarantees as with i.i.d. non-cyclic sampling, thus entirely alleviating the problem introduced by such data heterogeneity—as we also demonstrate empirically in Section 8. In Section 6 we go even further and show how we can maintain the same guarantee without any deterioration while also being competitive with separately

learned predictors, hedging our bets in case the differences between components are high. Section 7 discusses the relevance of the pluralistic approach in the context of Federated Learning.

## 2. Setup

We consider a stochastic convex optimization problem

$$F(w) = \mathbb{E}_{z \sim \mathcal{D}}[f(w, z)] \quad \text{where} \quad \mathcal{D} = \frac{1}{m} \sum_{i=1}^m \mathcal{D}_i, \quad (1)$$

where each component  $\mathcal{D}_i$  represents the data distribution associated with one of  $m$  blocks  $i = 1, \dots, m$ . For simplicity, we assume a uniform mixture; our results can be easily extended to non-uniform mixing weights and corresponding non-uniform block lengths. In a learning setting,  $z = (x, y)$  represents a labeled example and the instantaneous objective  $f(w, (x, y)) = \text{loss}(h_w(x), y)$  is the loss incurred if using the model  $w$ .

We assume  $f(w, z)$  is convex and 1-Lipschitz with respect to  $w$ , which lives in some high-, possibly infinite-dimensional, Euclidean or Hilbert space. Our goal is to compete with the best possible predictor of some bounded norm  $B$ , that is to learn a predictor  $\hat{w}$  such that  $F(\hat{w}) \leq F(w_*) + \epsilon$  where  $F(w_*) = \inf_{\|w\| \leq B} F(w)$ . We consider Stochastic Gradient Descent (SGD) iterates on the above objective:

$$w_{t+1} \leftarrow w_t - \eta_t \nabla f(w_t, z_t). \quad (2)$$

If the samples  $z_t$  are chosen independently from the data distribution  $\mathcal{D}$ , then with an appropriately chosen stepsizes  $\eta_t$ , SGD attains the mini-max optimal error guarantee (e.g., Shalev-Shwartz et al., 2009):

$$\mathbb{E}[F(\bar{w})] \leq F(w_*) + O\left(\sqrt{\frac{B^2}{T}}\right) \quad (3)$$

where  $\bar{w} = \frac{1}{T} \sum_{i=1}^T w_i$  and  $T$  is the total number of iterations, and thus also the total number of samples used.

But here we study *Block-Cyclic SGD*, which consists of  $K$  cycles (e.g., days) of  $mn$  iterations each, for a total of  $T = Kmn$  iterations of the update (2). In each cycle, the first  $n$  samples are drawn from  $\mathcal{D}_1$ , the second  $n$  from  $\mathcal{D}_2$  and so forth. That is, samples are drawn according to

$$z_{t(k, i, j)} \sim \mathcal{D}_i \quad \text{where } t(k, i, j) = (k-1)nm + (i-1)n + j \quad (4)$$

where  $k = 1..K$  indexes the cycles,  $i = 1..m$  indexes blocks and  $j = 1..n$  indexes iterations within a block and thus  $t(k, i, j) \in \{1, \dots, T\}$  indexes the overall sequence of samples and corresponding updates. The samples are therefore no longer identically distributed and the standard SGD analysis is no longer valid. We study the effect of such block-cyclic sampling on the SGD updates (2).

One can think of  $\mathcal{D}$  as the population distribution, with each step of SGD being based on a fresh sample; or of  $\mathcal{D}$  as an empirical distribution over a finite training set that is partitioned to  $m$  groups. It is important however not to confuse a *cycle* over the different components with an *epoch* over all samples—the notion of an *epoch* or the size of the support of  $F$  (i.e., the number of training points if  $F(w)$  is viewed as an empirical objective) do not enter our analysis. Either view is consistent with our development, though our discussion will mostly refer to the population view of SGD based on fresh samples.

## 3. Lower Bounds for Block-Cyclic Sampling

How badly can the block-cyclic sampling (4) hurt the SGD updates (2)? Unfortunately, we will see that the effect can be quite detrimental, especially when the number of samples  $n$  in each block is very large compared to the overall number of cycles  $K$ , and even if the number of components  $m$  is small. This is a typical regime in Federated Learning if we consider daily cycles, as we would hope training would take a small number of days, but the number of SGD iterations per block can be very high. For example, recent models for Gboard were trained for  $1.2 \times 10^6$  sequential steps over 4-5 days ( $K = 4$  or  $5$ ,  $nm = 3 \times 10^5$ , and e.g.  $m = 6$  and  $n = 5 \times 10^4$  if the day is divided into six 4-hour blocks) (Hard et al., 2018).<sup>1</sup>

To see how even  $K = 2$  can be very bad, consider an extreme setting where the objective  $f(w, z)$  depends only on the component of  $z$ , but is the same for all samples from each  $\mathcal{D}_i$ . Formally, let  $z = \{1, 2\}$  with  $P(z = i) = 1$ . That is there are only two possible examples which we call “1”, which is the only example in the support of the first component  $\mathcal{D}_1$  and “2”, which is the only example in the support of  $\mathcal{D}_2$ . The problem then reduces to a deterministic optimization problem of minimizing  $F(w) = \frac{1}{2}f(w, 1) + \frac{1}{2}f(w, 2)$ , but in each block we see gradients only of one of the two terms. With a large number of iterations per block  $n$ , we might be able to optimize each term separately very well. But we also know that in order to optimize such a composite objective to within  $\epsilon$ , even if we could optimize each component separately, requires alternating between the two components at least  $\Omega(B/\epsilon)$  times (Woodworth & Srebro, 2016), establishing a lower bound on the number of required cycles, independent of the number of iterations per cycle. This extreme deterministic situation establishes the limits of what can be done with a limited number of cycles, and captures the essence of the difficulty with block-cyclic sampling.

In fact, the lower bound we shall see applies not only to the

<sup>1</sup>The cited work uses the Federated Averaging algorithm with communication every 400 sequential steps, as opposed to averaging gradients every step, so while not exactly matching our setting, it provides a rough estimate for realistic scenarios.

precise SGD updates (2) but to *any* method that uses block-cyclic samples as in (4), and at each iteration access  $f(\cdot, z_t)$ , i.e. the objective on example  $z_t$ . More precisely, consider any method that at iteration  $t$  chooses a “query point”  $w_t$  and evaluates  $f(w_t, z_t)$  and the gradient (or a sub-gradient)  $\nabla f(w_t, z_t)$ , where  $w_t$  may be chosen in *any* way based on all previous function and gradient evaluations (the SGD iterates (2) is just one examples of such a method). The output  $\hat{w}$  of the algorithm can be any function of the query points  $w_t$ , the evaluated function values and the gradients.

**Theorem 1.** *Consider any (possibly randomized) optimization method of the form described in the previous paragraph, i.e. where access to the objective is by evaluating  $f(w_t, z_t)$  and  $\nabla f(w_t, z_t)$  on semi-cyclic samples (4) and where  $w_t$  is chosen based on  $\{(f(w_s, z_s), \nabla f(w_s, z_s)), s < t\}$  and the output  $\hat{w}$  based on all iterates<sup>2</sup>. For any  $B, n, K$  and  $m > 1$  there exists a 1-Lipschitz convex problem over high enough dimension such that  $\mathbb{E}[F(\hat{w})] \geq F(w_*) + \Omega(B/K)$ , where the expectation is over  $z_t$  and any randomization in the method.*

*Proof (sketch).* We construct a problem where all examples in the same block are the same (i.e.,  $\mathcal{D}_i$  is concentrated on a single example  $i$ ), thus reducing the problem to that of minimizing a composite objective  $\frac{1}{m}F_i(w)$ , and choose the components such that making progress requires alternating between them. See Appendix A for a complete proof of the Theorem, as well as of Theorem 2 below.  $\square$

Theorem 1 establishes that once we have a block-cyclic structure with multiple blocks per cycle, we cannot ensure excess error better then:

$$F(\hat{w}) - F(w_*) = \Omega\left(\frac{B}{K}\right) = \Omega\left(\sqrt{\frac{B^2}{T}} \sqrt{\frac{mn}{K}}\right). \quad (5)$$

Compared to using i.i.d. sampling as in Eq. (3), this is worse by a factor of  $\sqrt{mn/K}$ , which is very large when the number of samples per block  $n$  is much larger then the number of cycles  $K$ , as we would expect in many applications.

For smooth objectives (i.e., with Lipschitz gradients) the situation is quantitatively a bit better, but we again can be arbitrarily worse then i.i.d. sampling as the number  $n$  of iterations per block increases:

**Theorem 2.** *Under the same setup as in Theorem 1, for any  $B, n, K$  and  $m > 1$  there exists a 1-Lipschitz convex problem where the gradient  $\nabla_w f(w, z)$  is also 1-Lipschitz, such that  $\mathbb{E}[F(\hat{w})] \geq F(w_*) + \Omega(B^2/K^2)$ .*

That is, the best we can ensure with block-cyclic sampling,

<sup>2</sup>This theorem, as well as Theorem 2, holds even if the method is allowed “prox queries” of the form  $\arg \min_w f(w, z_t) + \lambda_t \|w - w_t\|^2$ .

even if the objective is smooth, is an excess error of:

$$\epsilon = \Omega\left(\frac{B^2}{K^2}\right) = \Omega\left(\sqrt{\frac{B^2}{T}} \sqrt{\frac{B^2 mn}{K^3}}\right), \quad (6)$$

i.e., worse by a  $\sqrt{\frac{B^2 mn}{K^3}}$  factor compared to i.i.d. sampling.

## 4. A Pluralistic Approach

The slowdown discussed in the previous section is due to the difficulty of finding a *consensus* solution  $w$  that is good for all components  $\mathcal{D}_i$ . But thinking of the problem as a learning problem, this should be avoidable. Why should we insist on a single consensus predictor? In a learning setting, our goal is to be able to perform well on future samples  $z \sim \mathcal{D}$ . If the data arrives cyclically in blocks, we should be able to associate each sample  $z$  with the mixture component  $\mathcal{D}_i$  it was generated from. Formally, one can consider a joint distribution  $(i, z) \sim \tilde{\mathcal{D}}$  over samples and components, where  $i$  is uniform over  $1..m$  and  $z| i \sim \mathcal{D}_i$ . In settings where the data follows a block-cyclic structure, it is reasonable to consider the component index  $i$  as observed and leverage this at prediction (test) time.

For example, in Federated Learning, this could be done by associating each client with the hours it is available for training. If it is available during multiple blocks, we can associate with it a distribution over  $i$  reflecting its availability.

Instead of learning a single consensus model for all components, we can therefore take a pluralistic approach and learn a separate model  $w^i$  for each one of the components  $i = 1..m$  with the goal of minimizing:

$$\mathbf{F}(\mathbf{w}) = \mathbb{E}_{(i, z) \sim \tilde{\mathcal{D}}} [f(w^i, z)] = \frac{1}{m} \sum_{i=1}^m F_i(w^i), \quad (7)$$

where  $\mathbf{w} = (w^1, \dots, w^m)$  and  $F_i(w^i) = \mathbb{E}_{z \sim \mathcal{D}_i} [f(w^i, z)]$ .

How can this be done? A simplistic approach would be to learn a predictor  $w^i$  for each component  $D_i$  separately, based only on samples from block  $i$ . This could be done by maintaining  $m$  separate SGD chains, and updating chain  $i$  only using samples from block  $i$ :

$$w_{t(k, i, j+1)}^i \leftarrow w_{t(k, i, j)}^i - \eta_{t(k, i, j)} \nabla f(w_{t(k, i, j)}^i, z_{t(k, i, j)}) \\ \text{where } w_{t(k+1, i, 1)}^i = w_{t(k, i, n+1)}^i. \quad (8)$$

We will therefore only have  $Kn = T/m$  samples per model, but for each model, the samples are now i.i.d., and we can learn  $w^i$  such that:

$$F_i(\hat{w}^i) \leq F_i(w_*^i) + O\left(\sqrt{\frac{B^2}{T/m}}\right) \quad (9)$$

where  $w_*^i = \arg \min_{\|w\| \leq B} F_i(w)$ , and so:

$$\mathbf{F}(\hat{\mathbf{w}}) = \frac{1}{m} \sum_{i=1}^m F_i(\hat{w}^i)$$

$$\leq \frac{1}{m} \sum_{i=1}^m F_i(w_*^i) + O\left(\sqrt{\frac{B^2}{T/m}}\right) \quad (10)$$

$$\leq F(w_*) + O\left(\sqrt{\frac{B^2}{T/m}}\right). \quad (11)$$

That is, ensuring excess error  $\epsilon$  requires  $T = O(mB^2/\epsilon^2)$  iterations, which represents a slow-down by a factor of  $m$  relative to overall i.i.d. sampling.

The pluralistic approach can have other gains, owing to its multi-task nature, if the  $m$  sub-problems  $\min F_i(w)$  are very different from each other. An extreme case is when the different subproblems are in direct conflict, in a sense “competing” with each other, and the optimal  $w_*^i$  for the different sub-populations are essentially incompatible. For example, this could happen if an informative feature has different effects in different subgroups. In this case we might have  $F_i(w_*^i) \ll F_i(w_*)$ , and there is a strong benefit to the pluralistic approach regardless of the semi-cyclic sampling.

The lower bounds of Section 3 involve a subtly different conflicting situation, where there is a decent consensus predictor, but finding it requires, in a sense, “negotiations” between the different blocks, and this requires going back-and-forth between them many times, as is possible with i.i.d. non-cyclic data, but not in a block-cyclic situation. Learning separate predictors would bypass this required conflict resolution.

An intermediate situation is when the sub-problems are “orthogonal” to each other, e.g. different features are used in different problems. In this case we have that  $w_*^i$  are orthogonal to each other, and we might have that  $w_* = \sum_i w_*^i$  and  $F_i(w_*^i) = F_i(w_*)$ . However, in this case, we would have that on average (over  $i$ ),  $\|w_*^i\| = \frac{\|w_*\|}{\sqrt{m}}$ , and so with the pluralistic approach we can learn relative to a norm-bound  $B' = B/\sqrt{m}$  that is  $\sqrt{m}$  times smaller than would be required when using a single consensus model. This precisely negates the slow-down in terms of  $m$  of the pluralistic learn-separate-predictors approach (8), and we recover the same performance as when learning a consensus predictor based on i.i.d. non-cyclic data.

The regime where we *would* see a significant slow-down is when the distributions  $\mathcal{D}_i$  are fairly similar. By separating the problem into  $m$  distinct problems, and not sharing information between components, we are effectively cutting the amount of data, and updates, by a factor of  $m$ , causing the slow-down. But if the distributions are indeed similar, then at least intuitively, the semi-cyclicity shouldn’t be too much

of a problem, and a simple single-SGD approach might be OK. At an extreme, if all components are identical ( $\mathcal{D}_i = \mathcal{D}$ ), the block-cyclic sampling is actually i.i.d. sampling and we do not have a problem in the first place.

We see, then, that in extreme situations, semi-cyclicity is not a real problem: if the components are extremely “competing”, we would be better off with separate predictors, while if they are identical we can just use a single SGD chain and lose nothing. But how do we know which situation we are in? Furthermore, what we actually expect is a combination of the above scenarios, with some aspects being in direct competition between the components, some being orthogonal while others being aligned and similar. Is there a simple approach that would always allow us to compete with training a single model based on i.i.d. data? That is, a procedure for learning  $\mathbf{w} = (w^1, \dots, w^m)$  such that:

$$\mathbf{F}(\mathbf{w}) \leq F(w_*) + O\left(\sqrt{\frac{B^2}{T}}\right). \quad (12)$$

We would like to make a distinction between our objective here and that of multi-task learning. In multi-task learning (Caruana, 1997, and many others), one considers several different but possibly related tasks (e.g. the tasks specified by our component distributions  $\mathcal{D}_i$ ) and the goal is to learn different “specialized” predictors for each task so as to *improve* over the consensus error  $F(w_*)$  while leveraging the relatedness between them so as to reduce sample complexity. In order to do so, it is necessary to target a particular way in which tasks are related (Baxter, 2000; Ben-David & Schuller, 2003). For example, one can consider shared sparsity (e.g. Turlach et al., 2005), shared linear features or low-rank structure (Ando & Zhang, 2005), shared deep features (e.g. Caruana, 1997), shared kernel or low nuclear-norm structure (e.g. Argyriou et al., 2008; Amit et al., 2007), low-norm additive structure (e.g. Evgeniou & Pontil, 2004), or graph-based relatedness (e.g. Evgeniou et al., 2005; Maurer, 2006). The success of multi-task learning then rests on whether the chosen relatedness assumptions hold, and the sample complexity depends on the specificity of this inductive bias. But we do *not* want to make any assumptions about relatedness. We would like a pluralistic method that *always* achieve the guarantee (12), without any additional or even low-order terms that depend on the relatedness. We hope we can achieve this since in (12), we are trying to compete with the *fixed* solution  $F(w_*)$ , and are resorting to pluralism only in order to overcome data heterogeneity, not in order to leverage it.

## 5. Pluralistic Averaging

We give a surprisingly simple solution to the above problem. It *is* possible to compete with  $F(w_*)$  in a semi-cyclic setting,

without any additional performance deterioration (at least on average) and with no assumptions about the specific relatedness structure between different components. In fact, this can be done by running a single semi-cyclic SGD chain (2), which previously we discussed was problematic. The only modification is that instead of averaging all iterates to obtain a single final predictor (or using a single iterate), we create  $m$  different pluralistic models by averaging, for each component  $i$ , only the iterates corresponding to that block:

$$\tilde{w}^i = \frac{1}{Kn} \sum_{k=1}^K \sum_{j=1}^n w_{t(k,i,j)}. \quad (13)$$

**Theorem 3.** Consider semi-cyclic samples as in (4). The pluralistic averaged solution  $\tilde{\mathbf{w}}$  given in (13) in terms of the iterates of (2) with step size  $\eta_t = B/\sqrt{2T}$  and starting at  $w_1 = 0$ , satisfies

$$\mathbb{E}[\mathbf{F}(\tilde{\mathbf{w}})] \leq F(w_*) + O\left(\sqrt{\frac{B^2}{T}}\right), \quad (14)$$

where the expectation is w.r.t. the samples.

The main insight is that the SGD guarantee can be obtained through an online-to-batch conversion, and that the online guarantee itself does *not* require any distributional assumption and so is valid also for semi-cyclic data. Going from Online Gradient Descent to Stochastic Gradient Descent, the i.i.d. sampling is vital for the online-to-batch conversion. But by averaging only over iterates corresponding to samples from the same component, we are in a sense doing an online-to-batch conversion for each component separately, and thus over i.i.d. samples.

*Proof.* Viewing the updates (2) as implementing online gradient descent (Zinkevich, 2003),<sup>3</sup> we have the following online regret guarantee for any sequence  $z_t$  (and in particular any sequence obtained from any sort of sampling) and any  $w$  with  $\|w\| \leq B$ :

$$\sum_{t=1}^T f(w_t, z_t) \leq \sum_{t=1}^T f(w, z_t) + \sqrt{2B^2 T}. \quad (15)$$

choosing  $w = w_*$  on the right hand side, dividing by  $T$  and

<sup>3</sup>The development in Zinkevich (2003) is for updates that also involve a projection:  $w_{t+1} \leftarrow \Pi_B(w_t - \eta_t \nabla f(w_t, z_t))$  where  $\Pi_B(w) = w/\max(\|w\|/B, 1)$ . See, e.g., Shalev-Shwartz (2012), for a development of SGD without the projection as in (2), and proof of the regret guarantee (15) for these updates. Although we present the analysis without a projection, our entire development is also valid with a projection as in Zinkevich (2003), which results in all the same guarantees.

rearranging the summation we have:

$$\begin{aligned} \frac{1}{m} \sum_{i=1}^m \frac{1}{Kn} \sum_{k,j} f(w_{t(k,i,j)}, z_{t(k,i,j)}) \\ \leq \frac{1}{m} \sum_{i=1}^m \frac{1}{Kn} \sum_{k,j} f(w_*, z_{t(k,i,j)}) + \sqrt{\frac{2B^2}{T}}. \end{aligned} \quad (16)$$

The above is valid for any sequence  $z_t$ . Taking  $z_t$  to be random, regardless of their distribution, we can take expectations on both sides. For the semi-cyclic sampling (4), and since  $w_t$  is independent of  $z_t$  we have:

$$\begin{aligned} \mathbb{E} \left[ \frac{1}{m} \sum_{i=1}^m \frac{1}{Kn} \sum_{k,j} \mathbb{E}_{z \sim \mathcal{D}_i} [f(w_{t(k,i,j)}, z)] \right] \\ \leq \frac{1}{m} \sum_{i=1}^m \frac{1}{Kn} \sum_{k,j} \mathbb{E}_{z \sim \mathcal{D}_i} [f(w_*, z)] + \sqrt{\frac{2B^2}{T}}, \end{aligned} \quad (17)$$

where the outer expectation on the left-hand-side is w.r.t.  $w_t$ . On the right hand side we have that  $\mathbb{E}_{z \sim \mathcal{D}_i} [f(w_*, z)] = F_i(w_*)$ . On the left hand side, we can use the convexity of  $f$  and apply Jensen's inequality:

$$\begin{aligned} \frac{1}{Kn} \sum_{k,j} \mathbb{E}_{z \sim \mathcal{D}_i} [f(w_{t(k,i,j)}, z)] \\ \geq \mathbb{E}_{z \sim \mathcal{D}_i} \left[ f\left(\frac{1}{Kn} \sum_{k,j} w_{t(k,i,j)}, z\right) \right] = F_i(\tilde{w}_i). \end{aligned} \quad (18)$$

Substituting (18) back into (17) we have:

$$\mathbb{E} \left[ \frac{1}{m} \sum_{i=1}^m F_i(\tilde{w}_i) \right] \leq \frac{1}{m} \sum_{i=1}^m F_i(w_*) + \sqrt{\frac{2B^2}{T}}. \quad (19)$$

recalling the definition (7) of  $\mathbf{F}(\tilde{\mathbf{w}})$  and that  $F(w) = \frac{1}{m} \sum_i F_i(w)$ , we obtain the expectation bound (14).  $\square$

An important, and perhaps surprising, feature of the guarantee of Theorem 3 is that it does *not* depend on the number of blocks  $m$ , and does not deteriorate as  $m$  increases. That is, in terms of learning, we could partition our cycle to as many blocks as we would like, ensuring homogeneity inside each block and without any accuracy cost. E.g. we could learn a separate model for each hour, or minute, or second of the day. The cost here is only a memory and engineering cost of storing, distributing and handling the plenitude of models, not an accuracy nor direct computational cost. This is in sharp contrast to separate training for each component, as well as to most multi-task approaches.

## 6. Pluralistic Hedging

The pluralistic averaging approach shows that we can always compete with the best single possible solution  $w_*$ , even when faced with semi-cyclic data. But as discussed in Section 4, perhaps in some extreme cases learning separate predictors  $w^i$ , each based only on data from the  $i$ th component as in (8), might be better. That is, depending on the similarity and conflict between components, the guarantees (9) might be better than that of Theorem 3, at least for some of the predictors. But if we do not know in advance which regime we are at, nor for which components we are better off with a separately-learned model (because they are very different from the others) and which are better off leveraging also the other components, can we still ensure the  $m$  individual guarantees (9) and the guarantee of Theorem 3 simultaneously?

Here we give a method for doing so, based on running both the single SGD chain (2) and the separate SGD chains (8) and carefully combining them using learned weights for each chain. Let  $q > 0$  be the weight for the full SGD chain  $w_t$  (2), which will be kept fixed throughout, and let  $q_t^i > 0$  be the weights assigned to the block-specific SGD chains  $w_t^i$  (8) on step  $t$ . We learn the weights using a multiplicative update rule. At step  $t = t(k, i, j)$ , in which block  $i$  is active, this update takes the form:

$$q_{t+1}^i \leftarrow q_t^i \cdot (1 + \nu(f(w_t, z_t) - f(w_t^i, z_t)));$$

$$\forall i' \neq i, \quad q_{t+1}^{i'} \leftarrow q_t^{i'},$$

where  $\nu > 0$  is a learning rate (separate from those of the SGD chains). Then, we let  $p_t = q_t^i / (q_t^i + q) \in [0, 1]$  and choose between the full and block-specific SGD chains by:

$$u_t \leftarrow \begin{cases} w_t^i & \text{with probability } p_t; \\ w_t & \text{otherwise.} \end{cases}$$

Finally, we obtain the final predictors via pluralistic averaging; namely, for each component  $i$ , we average only the iterate within the corresponding blocks:

$$\tilde{u}^i = \frac{1}{Kn} \sum_{k=1}^K \sum_{j=1}^n u_{t(k, i, j)}. \quad (20)$$

For the averaged solution  $\tilde{\mathbf{u}} = (\tilde{u}^1, \dots, \tilde{u}^m)$ , we have the following.

**Theorem 4.** Set  $\nu = \frac{1}{2B} \sqrt{(m/T) \log(BT/m)}$ ,  $q = 1 - \eta$ , and initialize  $q_1^i = \eta$  for each  $i$ . Then, for all  $i$  we have

$$\mathbb{E}[F_i(\tilde{u}^i)] \leq F_i(w_*^i) + 4\sqrt{\frac{B^2 \log(BT/m)}{T/m}},$$

and further, provided that  $m \leq B^2 Kn$ ,

$$\mathbb{E}[\mathbf{F}(\tilde{\mathbf{u}})] \leq F(w_*) + 4\sqrt{\frac{B^2}{T}}.$$

The requirement that  $m \leq B^2 Kn$  is extremely mild, as we would generally expect a large number of iterations per block  $n$ , and only a mild number of blocks  $m$ , i.e.  $n \ll m$ .

The proofs in this section use the following notation. For all  $i \in [m]$ , let  $S_i \subseteq [T]$  be the set of time steps where we got a sample from distribution  $\mathcal{D}_i$  (so that  $|S_i| = T/m$ ). We first prove the following.

**Lemma 1.** For each  $i \in [m]$  we have

$$\mathbb{E} \left[ \sum_{t \in S_i} f(u_t, z_t) - \sum_{t \in S_i} f(w_t, z_t) \right] \leq 2,$$

and

$$\mathbb{E} \left[ \sum_{t \in S_i} f(u_t, z_t) - \sum_{t \in S_i} f(w_t^i, z_t) \right] \leq 2B \sqrt{\frac{T}{m} \log \frac{BT}{m}}.$$

Here, the expectation is taken w.r.t. the  $z_t$  as well as the internal randomization of the algorithm.

The proof (in Appendix B) is based on classic analysis of the PROD algorithm (Cesa-Bianchi et al., 2007; see also Even-Dar et al., 2008; Sani et al., 2014). We can now prove the main theorem of this section.

*Proof of Theorem 4.* The proof follows from a combination of Lemma 1 with regret bounds for the SGD chains (2) and (8), viewed as trajectories of the Online Gradient Descent algorithm. Standard regret bounds for the latter (Zinkevich, 2003; see also Shalev-Shwartz, 2012; Hazan, 2016) imply that for any sequence  $z_1, \dots, z_T$  and for any  $\|w_*\| \leq B$ , it holds that

$$\sum_{t=1}^T f(w_t, z_t) - \sum_{t=1}^T f(w_*, z_t) \leq 2B\sqrt{T}, \quad (21)$$

and, for each  $i \in [m]$  and for any  $w_*^i$  such that  $\|w_*^i\| \leq B$ ,

$$\sum_{t \in S_i} f(w_t^i, z_t) - \sum_{t \in S_i} f(w_*^i, z_t) \leq 2B \sqrt{\frac{T}{m}}. \quad (22)$$

Now, fix  $i \in [m]$ ; Lemma 1 together with Eq. (22) imply

$$\frac{1}{T/m} \sum_{t \in S_i} \mathbb{E}[f(u_t, z_t) - f(w_*^i, z_t)] \leq 4\sqrt{\frac{B^2 \log(BT/m)}{T/m}}.$$

Using the facts that  $\mathbb{E}_t[f(u_t, z_t)] = F_i(u_t)$  and  $\mathbb{E}_t[f(w_*^i, z_t)] = F_i(w_*^i)$  for any  $t \in S_i$ , we have

$$\frac{1}{T/m} \sum_{t \in S_i} (\mathbb{E}[F_i(u_t)] - F_i(w_*^i)) \leq 4\sqrt{\frac{B^2 \log(BT/m)}{T/m}}.$$

Appealing to the convexity of  $F_i$  and applying Jensen's inequality on the left-hand side to lower bound  $\frac{1}{T/m} \sum_{t \in S_i} F_i(u_t) \geq F_i(\tilde{u}^i)$ , we obtain

$$\mathbb{E}[F_i(\tilde{u}^i)] - F_i(w_\star) \leq 4\sqrt{\frac{B^2 \log(BT/m)}{T/m}},$$

which implies the first guarantee of the theorem. For the second claim, the first bound of Lemma 1 implies

$$\sum_{i=1}^m \mathbb{E}\left[\sum_{t \in S_i} f(u_t, z_t)\right] - \mathbb{E}\left[\sum_{t=1}^T f(w_t, z_t)\right] \leq 2m.$$

Summing this with Eq. (21) and dividing through by  $T$  gives

$$\frac{1}{m} \sum_{i=1}^m \frac{1}{T/m} \sum_{t \in S_i} \mathbb{E}[f(u_t, z_t) - f(w_\star, z_t)] \leq \frac{2m}{T} + 2\sqrt{\frac{B^2}{T}}.$$

Next, as before, substitute the conditional expectations and use Jensen's to lower bound the left-hand side; this yields

$$\frac{1}{m} \sum_{i=1}^m (\mathbb{E}[F_i(\tilde{u}^i)] - F_i(w_\star)) \leq \frac{2m}{T} + 2\sqrt{\frac{B^2}{T}}.$$

Recalling now the definitions  $\mathbf{F}(\tilde{\mathbf{u}}) = \frac{1}{m} \sum_{i=1}^m F_i(\tilde{u}^i)$  and  $F(w_\star) = \frac{1}{m} \sum_{i=1}^m F_i(w_\star)$ , we have shown that

$$\mathbb{E}[\mathbf{F}(\tilde{\mathbf{u}})] - F(w_\star) \leq \frac{2m}{T} + 2\sqrt{\frac{B^2}{T}}.$$

Finally, note  $m \leq B\sqrt{T} = B\sqrt{Kmn}$  implies  $m \leq B^2Kn$ .  $\square$

## 7. Applications in Federated Learning

Our primary motivation for studying the problem of semi-cyclic SGD is the Federated Learning setting. Here, mobile devices participate when certain conditions — idleness, charging, free wireless connection — are met. These constraints lead to various sources of cyclic bias. For instance, the number of devices eligible for training at a given location is typically much higher at night time (Bonawitz et al., 2019). Hence, at any given point in time, the majority of devices available globally will come from locations where it's night. For example, a model trained on English text may experience cyclic behavior from blocks representing mostly users from India, the UK, and the United States.

Even if we consider only devices from a specific time zone, we may encounter demographic bias — devices charging during the day may represent a different distribution of users than those charging at night. An example of such cyclic behavior was shown in Yang et al. (2018). Here,  $\sim 3x$  fewer examples were available to train on in the day time than at night. The authors speculate that this imbalance indicates differences in the training data distribution at day and night

time — the model sees more night time data, leading to worse accuracy on day time data.

Applying the pluralistic approach from Section 5 to Federated Learning is straightforward. Notably, there are no changes to the training procedure, other than maintaining sufficient incremental model versions. For evaluation and inference, the number and extent of blocks needs to be defined by observing the period of the cycle (e.g. by taking an iterate and computing its evaluation loss over time). Devices are then associated with a specific block based on the phase of the cycle where they typically meet training eligibility requirements; this determines which model is served for either evaluation or inference.

## 8. Experiments

To illustrate the challenges of optimizing on block-cyclic data, we train and evaluate a small logistic regression model on the Sentiment140 Twitter dataset (Go et al., 2009), a binary classification task over  $1.6 \times 10^6$  examples. We split the data into training (90%) and test (10%) sets, partition it into  $m = 6$  components based on the timestamps (but not dates) of the Tweets: 12am - 4am, 4am - 8am, etc, then divide each component across  $K = 10$  cycles (days). For more details, see Appendix C. For simplicity, we keep the model architecture (linear bag of words classifier over top 1024 words) and minibatch size (128) fixed; we used a learning rate of  $\eta = 0.464$  (determined through log-space grid search) except for the per-component SGD approach (8) where  $\eta = 1.0$  was optimal due to fewer iterations.

To illustrate the differences between the proposed methods, as well as to capture the diurnal variation expected in practical Federated Learning settings, we vary the label balance as a function of the time of day, ensuring the  $\mathcal{D}_i$  are somewhat different. In particular, we randomly drop negative posts from the midnight group so the overall rate is 2/3 positive sentiment, we randomly drop positive posts from the noon group so the overall rate is 1/3 positive sentiment, and we linearly interpolate for the other 4 components. For discussion, and a figure giving positive and negative label counts over components, see Appendix C.2. We write  $\hat{F}_i$  for the empirical accuracy on component  $i$  of the test set, with  $\hat{F}$  measuring the overall test accuracy.

We consider the following approaches:

1. A non-pluralistic **block-cyclic consensus** model, subject to the lower bounds of Section 3. The accuracy of this procedure for day  $k$  is given by the expected test-set accuracy of randomly picking  $i \in \{1, \dots, m = 6\}$ , and evaluating  $F(w_{t(k,i,n)})$ , i.e. evaluating the model after completing a random block  $i$  on each day  $k$ .
2. The **per-component SGD** approach (8), where we run  $m = 6$  SGD chains, with each model training on

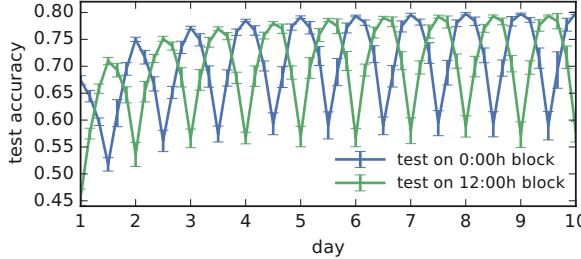


Figure 1. Accuracy of a model trained on block-cyclic data tested on two blocks of data - posts from the midnight component (12am-4am), and posts from the noon component (12pm-4pm). Mean and standard deviations are computed from ten training repetitions.

only one block per day. For each day  $k$ , we evaluate  $\frac{1}{m} \sum_i \hat{F}_i(w_{kn}^i)$ , where  $w_{kn}^i$  is the model for component  $i$  trained on the first  $k$  blocks for component  $i$ .

3. The **pluralistic single SGD chain** approach of Section 5. The SGD chain is the same as for the consensus model, but evaluation is pluralistic as above: for each day  $k$  we evaluate  $\frac{1}{m} \sum_i \hat{F}_i(w_{t(k,i,n)})$ , i.e. for each test component we use the model at the end of the corresponding training block.
4. An (impractical in real settings) **idealized i.i.d. model**, trained on the complete shuffled training set. For purposes of evaluation, we treat the data as if it had cyclic block structure and evaluate identically to the block-cyclic consensus model.

See Appendix C.3 for more details on evaluation. Note that we do not use a model average, but rather take the most recent relevant iterate, matching how such models are typically deployed in practice.

Figure 1 illustrates how the block-cyclic structure of data can hurt accuracy of a consensus model. We plot how the accuracy on two test data components—the midnight, and the noon component—changes as a function of time of day when evaluated on the iterates of the SGD chain (2). Training is quick to leverage the label bias and gravitate towards a model too specific for the block being processed, instead of finding a better predictive solution based on other features that would help predictions across all groups.

Figure 2 compares results from the four approaches. Observe: 1) Pluralistic models with separate SGD chains take longer to converge because they are trained on less data. Depending on how long training proceeds, the size of the data set, and the number of components, these models may or may not surpass the idealized i.i.d. SGD and pluralistic single SGD chain models in accuracy. 2) The pluralistic models from a single SGD chain consistently perform better than a single consensus model from the same chain. 3) The performance of pluralistic models is en par with or better than the idealized i.i.d. SGD model, reflecting the fact that these models better match the target data distribution than a single model (i.i.d. or block cyclic consensus) can.

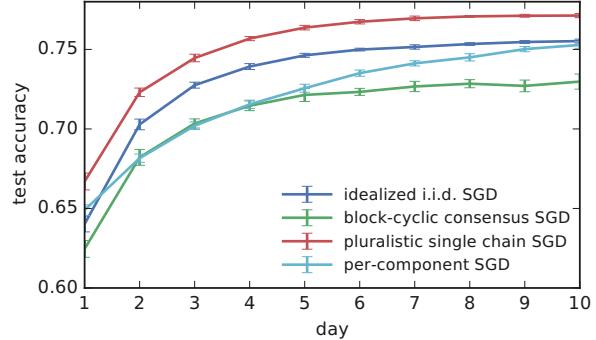


Figure 2. Comparison of approaches for a sentiment classification task, using i.i.d. and block-cyclic data. Means and standard deviations shown are computed from ten training repetitions.

## 9. Summary

We considered training in the presence of block-cyclic data, showed that ignoring this source of data heterogeneity can be detrimental, but that a remarkably simple pluralistic approach can entirely resolve the problem and ensure, even in the worst case, the same guarantee as with homogeneous i.i.d. data, independent of the number of blocks. When the component distributions are different, pluralism can outperform the “ideal” i.i.d. baseline, as our experiments illustrate. An important lesson we want to stress is that pluralism can be a critical tool for dealing with heterogeneous data, by embracing this heterogeneity instead of wastefully insisting on a consensus solution.

Dealing with heterogeneous data can be difficult in many machine learning settings, and especially in Federated Learning where learning happens on heterogeneous devices each handling its own distinct data. This can produce variabilities in processing time and latency, amount of data per client, and client availability and reliability, which can all be correlated in different ways with different data components. We expect pluralistic solutions might be relevant in many of these situations.

This work considered only *convex* objectives and sequential SGD. As with many other training techniques, we expect this study of convex semi-cyclic training to be indicative also of non-convex scenarios like training deep networks (except that a random predictor would need to be used instead of an averaged predictor in the online-to-batch conversion) and also serve as a basis for further analysis of the non-convex case. We are also interested in analyzing the effect of block-cyclic data on methods beyond sequential SGD, and most prominently parallel SGD (aka Federated Averaging). Unfortunately, there is currently no satisfying and tight analysis of parallel SGD even for i.i.d. data, making a detailed analysis of semi-cyclicity for this method beyond our current reach. Nevertheless, we again hope our analysis will be both indicative and serve as a basis for future exploration of parallel SGD and other distributed optimization approaches.

## Acknowledgements

NS was supported by NSF-BSF award 1718970 and a Google Faculty Research Award.

## References

Amit, Y., Fink, M., Srebro, N., and Ullman, S. Uncovering shared structures in multiclass classification. In *Proceedings of the 24th international conference on Machine learning*, pp. 17–24. ACM, 2007.

Ando, R. K. and Zhang, T. A framework for learning predictive structures from multiple tasks and unlabeled data. *Journal of Machine Learning Research*, 6(Nov):1817–1853, 2005.

Argyriou, A., Evgeniou, T., and Pontil, M. Convex multi-task feature learning. *Machine Learning*, 73(3):243–272, 2008.

Arjevani, Y. and Shamir, O. Communication complexity of distributed convex learning and optimization. In *Advances in neural information processing systems 28*, pp. 1756–1764, 2015.

Baxter, J. A model of inductive bias learning. *Journal of Artificial Intelligence Research*, 12:149–198, 2000.

Ben-David, S. and Schuller, R. Exploiting task relatedness for multiple task learning. In *Learning Theory and Kernel Machines*, pp. 567–580. Springer, 2003.

Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingberman, A., Ivanov, V., Kiddon, C., Konecný, J., Mazzocchi, S., McMahan, H. B., Overveldt, T. V., Petrou, D., Ramage, D., and Roselander, J. Towards federated learning at scale: System design. In *Conference on Systems and Machine Learning (SysML)*, 2019. URL <https://arxiv.org/abs/1902.01046>.

Carmon, Y., Duchi, J. C., Hinder, O., and Sidford, A. Lower bounds for finding stationary points i. *arXiv preprint arXiv:1710.11606*, 2017.

Caruana, R. Multitask learning. *Machine learning*, 28(1): 41–75, 1997.

Cesa-Bianchi, N., Mansour, Y., and Stoltz, G. Improved second-order bounds for prediction with expert advice. *Machine Learning*, 66(2): 321–352, Mar 2007. ISSN 1573-0565. doi: 10.1007/s10994-006-5001-7. URL <https://doi.org/10.1007/s10994-006-5001-7>.

Even-Dar, E., Kearns, M., Mansour, Y., and Wortman, J. Regret to the best vs. regret to the average. *Machine Learning*, 72(1):21–37, Aug 2008. ISSN 1573-0565. doi: 10.1007/s10994-008-5060-z. URL <https://doi.org/10.1007/s10994-008-5060-z>.

Evgeniou, T. and Pontil, M. Regularized multi-task learning. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 109–117. ACM, 2004.

Evgeniou, T., Micchelli, C. A., and Pontil, M. Learning multiple tasks with kernel methods. *Journal of Machine Learning Research*, 6(Apr):615–637, 2005.

Go, A., Bhayani, R., and Huang, L. Twitter sentiment classification using distant supervision. *CS224N Project Report, Stanford*, 150, 01 2009. Data from <http://help.sentiment140.com/for-students>.

Hard, A., Rao, K., Mathews, R., Beaufays, F., Augenstein, S., Eichner, H., Kiddon, C., and Ramage, D. Federated learning for mobile keyboard prediction. *arXiv preprint 1811.03604*, 2018.

Hazan, E. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.

Konečný, J., McMahan, H. B., Ramage, D., and Richtárik, P. Federated optimization: Distributed machine learning for on-device intelligence. *arXiv preprint arXiv:1610.02527*, 2016.

Maurer, A. Bounds for linear multi-task learning. *Journal of Machine Learning Research*, 7(Jan):117–139, 2006.

McMahan, H. B. and Ramage, D. Federated learning: Collaborative machine learning without centralized training data, April 2017. URL <https://ai.googleblog.com/2017/04/federated-learning-collaborative.html>. Google AI Blog.

McMahan, H. B., Moore, E., Ramage, D., Hampson, S., and y Arcas, B. A. Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, pp. 1273–1282, 2017.

Sani, A., Neu, G., and Lazaric, A. Exploiting easy data in online optimization. In *Advances in Neural Information Processing Systems*, pp. 810–818, 2014.

Shalev-Shwartz, S. Online learning and online convex optimization. *Foundations and Trends® in Machine Learning*, 4(2):107–194, 2012.

Shalev-Shwartz, S., Shamir, O., Srebro, N., and Sridharan, K. Stochastic convex optimization. In *Proceedings of the 22nd Annual Conference on Learning Theory (COLT)*, 2009.

Turlach, B. A., Venables, W. N., and Wright, S. J. Simultaneous variable selection. *Technometrics*, 47(3):349–363, 2005.

Woodworth, B. and Srebro, N. Tight complexity bounds for optimizing composite objectives. In *Advances in neural information processing systems*, pp. 3639–3647, 2016.

Woodworth, B. and Srebro, N. Lower bound for randomized first order convex optimization. *arXiv preprint arXiv:1709.03594*, 2017.

Woodworth, B., Wang, J., Smith, A., McMahan, B., and Srebro, N. Graph oracle models, lower bounds, and gaps for parallel stochastic optimization. In *Advances in Neural Information Processing Systems 31*, pp. 8505–8515. 2018.

Yang, T., Andrew, G., Eichner, H., Sun, H., Li, W., Kong, N., Ramage, D., and Beaufays, F. Applied federated learning: Improving google keyboard query suggestions. *CoRR*, abs/1812.02903, 2018. URL <http://arxiv.org/abs/1812.02903>.

Zinkevich, M. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning (ICML)*, pp. 928–936, 2003.