
Stochastic Optimization with Laggard Data Pipelines

Naman Agarwal

Google AI Princeton
Princeton, NJ 08540
namanagarwal@google.com

Rohan Anil

Google Research
Mountain View, CA 94043
rohananil@google.com

Tomer Koren

Tel Aviv University & Google
Tel Aviv, Israel
tkoren@tauex.tau.ac.il

Kunal Talwar

Apple*
Cupertino, CA 95014
ktalwar@apple.com

Cyril Zhang

Microsoft Research†
New York, NY 10012
cyrilzhang@microsoft.com

Abstract

State-of-the-art optimization is steadily shifting towards massively parallel pipelines with extremely large batch sizes. As a consequence, CPU-bound preprocessing and disk/memory/network operations have emerged as new performance bottlenecks, as opposed to hardware-accelerated gradient computations. In this regime, a recently proposed approach is data echoing (Choi et al., 2019), which takes repeated gradient steps on the same batch while waiting for fresh data to arrive from upstream. We provide the first convergence analyses of “data-echoed” extensions of common optimization methods, showing that they exhibit provable improvements over their synchronous counterparts. Specifically, we show that in convex optimization with stochastic minibatches, data echoing affords speedups on the curvature-dominated part of the convergence rate, while maintaining the optimal statistical rate.

1 Introduction

Recent empirical successes in large-scale machine learning have been powered by massive data parallelism and hardware acceleration, with batch sizes trending beyond 10K+ images [46] or 1M+ tokens [9]. Numerous interdisciplinary sources [5, 12, 24, 33] indicate that the performance bottlenecks of contemporary deep learning pipelines can lie in many places other than gradient computation. In other words, since the initial breakthroughs in hardware-accelerated deep learning [14, 28, 37], GPUs (and TPUs, FPGAs, etc.) have become too fast for upstream data loaders and preprocessors to keep up with.

Choi et al. [13] propose *data echoing*, a simple and versatile way to improve training performance in this regime. Each stage of the data pipeline runs asynchronously, oblivious to whether its input has been refreshed upstream. In particular, the optimization algorithm may choose to take additional gradient steps before a minibatch is refreshed, rather than spend idle time waiting for more data. The authors present a large-scale proof-of-concept empirical study, and find that data echoing affords a 3.25× speedup in a network-bound ImageNet setting.

Some natural curiosities arise from this practice: *When might this overfit? How carefully should one adjust the step size of an echoed gradient? Does acceleration work? A theoretical understanding of convergence guarantees for these data-echoed optimization algorithms is missing.*

*Work performed while at Google Brain.

†Work performed while at Google AI Princeton and Princeton University.




Figure 1: Schematic of data echoing, inspired by Choi et al. [13]. If the upstream data pipeline is $K = 4$ times slower than SGD, then SGD can potentially take that many steps on the same batch before the next one arrives.

	$T = 1$	T general
$K = 1$	SGD	
K small	Compute-bound ERM	Data echoing (Thm. 7)
K large	Data-bound ERM	Approx-Prox [43]
$K \rightarrow \infty$	Statistical ERM	Minibatch-Prox [43]

Table 1: Regimes of echoing factor K and number of batches T which our analyses interpolates.

In this paper, we settle the issues of convergence and generalization for echoed gradient methods in convex optimization. We show that these methods can match the optimization performance of their non-stochastic counterparts, while achieving optimal statistical rates. As state-of-the-art batch sizes continue to grow, along with the distributed systems that enable them, we hope that this will provide a first theoretical grounding towards understanding the algorithmic and statistical challenges in these hardware-motivated optimization settings.

1.1 Technical contributions

Our model of data echoing is parameterized by the batch size B , the number of fresh i.i.d. batches T , and the *echoing factor* K , which is the number of gradient steps an algorithm can take on the (convex) loss on each batch. This reflects the hardware-determined setting where the data loader is at least K times slower than the optimizer.

Convergence in all data echoing regimes. We first show that echoed SGD, with the correctly tuned step size, achieves a factor- K speedup on the curvature term of the standard convergence rate, while keeping the optimal statistical term. Next, we develop an echoed method that is oblivious to the echoing factor K , getting the same rates for echoed SGD with an appropriately chosen proximal regularizer. Finally, we show that Nesterov’s accelerated gradient descent, when echoed, achieves the optimal rates on quadratic losses. As a side contribution, we fix a small error in a technical lemma in [11], used in establishing the stability of AGD on quadratics. For general convex losses, we arrive at the same open question as these authors.

Full interpolation between known regimes. To set up notation, suppose that we go over T batches of data, and perform K echoed gradient steps for each batch. In the special case of $T = 1$ fresh batches, the problem becomes empirical risk minimization with a limited computational budget of K gradient steps. When K is small, the error is dominated by a *curvature* term, while for large enough K this falls below the *statistical* term.

Motivated by the communication-limited setting, Wang et al. [43] focus on the case where T is general and $K \rightarrow \infty$, analyzing the convergence of *exact* optimization of the prox-regularized minibatch loss. They develop a mild “approx-prox” guarantee when K is large enough to enable an *exact+perturbation* analysis. Our analysis generalizes and strengthens these results, handling all values of K ; Table 1 summarizes this discussion. When $B \rightarrow \infty$, the statistical problem disappears, and we recover the classical setting of full gradient descent with KT oracle calls [8, 35].

Algorithm	Standard	Data-echoed
SGD	$O\left(\frac{\beta D^2}{T} + \frac{\rho D}{\sqrt{BT}}\right)$ (classical; see Lan [29])	$O\left(\frac{\beta D^2}{KT} + \frac{\rho D}{\sqrt{BT}}\right)$ (Theorem 7)
Minibatch-Prox	$O\left(e^{-K/\kappa} + \frac{\rho D}{\sqrt{BT}}\right)$ (Wang et al. [43]; K large)	$O\left(\frac{\beta D^2}{KT} + \frac{\rho D}{\sqrt{BT}}\right)$ (Theorem 10)
Stochastic AGD	$O\left(\frac{\beta D^2}{T^2} + \frac{\rho D}{\sqrt{BT}}\right)$ (Lan [29])	$O\left(\frac{\beta D^2}{K^2 T^2} + \frac{\rho D}{\sqrt{BT}}\right)$ (Theorem 13; quadratics)

Table 2: Single-step and *data-echoed* convergence rates of stochastic optimization algorithms studied in this paper. Notice that the optimization terms depend analogously on the total number of steps KT , and the statistical terms have optimal dependence on the total number of i.i.d. samples BT .

Stability-based analysis. We provide a modular proof framework for data echoing convergence bounds, based on uniform stability [7] and a potential-based notion of regret, which isolates the “bias” (curvature) and “variance” (generalization) components of the problem. This recipe (Theorem 4) can be used to sharpen bounds in more restricted settings, or analyze future data-echoed algorithms.

1.2 Motivation and context

It is well-known in the practice of GPU training that model parameter updates are not necessarily the performance bottleneck; this is why SSD storage is critical for pipelines on the scale of ImageNet [17]. For quantitative studies of I/O performance in deep learning, see [12, 45]. Many empirical advances have stemmed from innovations in data augmentation [15, 22, 41]. Unlike neural network training and inference, these data transformations can be highly sequential and/or heterogeneous, and must be done on CPU. Unlocking GPU parallelism for CPU-bound computations is often a significant engineering effort [16, 20, 27, 32].

Extremely large batch sizes have become the norm in training state-of-the-art models [4, 9, 18, 40, 46]. An overwhelming theme has been that *constant factors matter*; for example, memory-bound optimizers [2, 10, 39] care about factors of 2-3. When selecting hyperparameters in large-batch training setups, it is common to balance the curvature- and noise-dominated terms [26, 34, 38, 42]. This underscores the need to better understand the fine-grained dependences on B , T , and K , especially the resources at stake are on the scale of GPU-years.

The idea of repeated steps on a batch/workers has also been investigated in the context of federated learning [25], where a related concept is referred to as *local SGD* or *federated averaging*. There are two key distinctions: federated learning considers multiple copies of local SGD running on different workers, which synchronize intermittently through averaging; under the most simplified assumptions, each individual gradient step within a worker is taken on a fresh batch. While the improvements obtained in this recent and concurrent line of work (see [44] and references therein) bear resemblance to our bounds, we do not see a direct reduction in either direction. Indeed, due to the distinctions mentioned, getting similar improvements to the curvature term in federated learning is not possible beyond quadratics, as shown by [44]. Obtaining optimal rates for convex functions in the federated learning setting remains an interesting open problem.

1.3 The bias-variance problem in data echoing

As mentioned earlier, data echoing presents a natural tradeoff between the optimization gains from repeating gradient steps vs. the potential loss of generalization due to overfitting to stale batches. To understand this in detail, let us revisit the standard convergence guarantee for SGD on smooth functions:

$$\mathbb{E}[F(w_{\text{out}})] - F(w^*) \leq O\left(\frac{1}{T} + \frac{1}{\sqrt{BT}}\right).$$

We interpret the first term as a *bias (curvature)* term, which diminishes at a faster rate due to smoothness. The second term is the *variance (statistical)* term, which arises due to the stochasticity

in the data, and thus naturally scales as the inverse square root of the batch size. Viewing B as fixed, the variance term is intrinsic to the data; therefore, we cannot expect data-echoing (or any algorithm) to give us improvements on that term for free. In fact, it is possible to make this term degrade, by overfitting on a batch. On the other hand, we can expect the bias term, which is governed by progress on the curvature of the underlying population loss, to decrease as we are given more echoing steps K . In light of this, the best analogous convergence rate one should hope to achieve in the data-echoing setting is

$$O\left(\frac{1}{KT} + \frac{1}{\sqrt{BT}}\right).$$

Our results establish exactly this rate for the data-echoed version of gradient descent. The data-echoed version of accelerated gradient descent is also shown to possess similar gains but with a faster rate of K^2T^2 . The challenge is to prevent overfitting; obtaining such rates requires careful control (depending on K) of step sizes. Later, we alleviate this need via data-echoed proximal GD, whose parameters are independent of K .

1.4 Overview of techniques

All of our theorems follow the same analysis structure. In particular, we formalize a notion of *potential-bounded regret* (Definition 3), which connects an algorithm’s function-value progress on a minibatch to a decrement on a certain potential function with respect to an arbitrary point. This potential function depends on the algorithm in question, but the key property is that it telescopes when summed over batches; this provides a fast rate on the bias term with respect to T .

The second piece of the analysis connects function-value decrease on a batch to the population objective via the notion of *uniform stability* (Definition 1). Note that the potential decrease scales inversely with K , whereas the stability constant increases with K (unless a proximal regularizer is added). The key to maintaining the optimal statistical rate is to balance these terms via the choice of an appropriate step size. This type of algorithmic stability analysis has appeared various times in the literature [7, 11, 21]; we show here that it affords a way to analyze echoed gradient methods.

2 Preliminaries

2.1 Problem definition

Given a convex set $\mathcal{W} \subseteq \mathbb{R}^n$ and a domain Ξ with a distribution \mathcal{D} , we consider the following stochastic convex optimization problem:

$$\underset{w \in \mathcal{W}}{\text{minimize}} \quad F(w) \stackrel{\text{def}}{=} \mathbb{E}_{\xi \sim \mathcal{D}} [f(w, \xi)]. \quad (1)$$

Here $f : \mathbb{R}^n \times \Xi \rightarrow \mathbb{R}$ is such that for any ξ , $f(\cdot, \xi)$ is convex, differentiable, ρ -Lipschitz, and β -smooth; i.e., for all $w, w' \in \mathcal{W}$,

$$f(w) - f(w') \leq \langle \nabla f(w'), w - w' \rangle + \frac{\beta}{2} \|w - w'\|^2.$$

When the minimizer exists, we define $w^* = \arg \min_{w \in \mathcal{W}} F(w)$. However, our results pertaining to optimality gaps $F(w) - F(w^*)$ hold for arbitrary w^* , encompassing the case when this minimizer does not exist. We further assume that we have access to an initial point w_0 with a bounded distance D from the comparator; i.e., $\|w_0 - w^*\| \leq D$.

Minibatch optimization. We will work in the stochastic minibatch oracle model: at each time step t , we receive a new batch (of size B) examples $\xi^{(t)} = \{\xi^{(t,i)}\}_{i=1}^B$ sampled i.i.d. from the distribution \mathcal{D} . For any batch of examples $\xi = \{\xi^{(i)}\}$, we define the empirical objective on the batch as

$$\bar{f}_\xi(w) \stackrel{\text{def}}{=} \frac{1}{|\xi|} \sum_{i=1}^B f(w, \xi^{(i)}).$$

Throughout this paper, we will use **boldface** ξ to denote a batch of B examples, and unbolded ξ to represent a single example in Ξ .

Optimization algorithms. We formalize a generic notion of optimization algorithms. Since these algorithms are called repeatedly by the data-echoing procedure, we will augment the output space of optimization algorithms with a notion of *state*, which it internally maintains and passes to the next

run of the same algorithm. Formally, an optimization algorithm is an iterative procedure which takes four arguments: an initial point $w_{\text{init}} \in \mathcal{W}$, an initial state s_{init} , the current batch ξ which determines the current objective \bar{f}_ξ , and the number of steps k . The algorithm outputs a point $w_{\text{out}} \in \mathcal{W}$ and an output state s_{out} . In short, an algorithm \mathcal{A} implements

$$(w_{\text{out}}, s_{\text{out}}) \leftarrow \mathcal{A}(w_{\text{init}}, s_{\text{init}}, \xi, k).$$

We will suppress the notation of one or more of the arguments to \mathcal{A} when they will be clear from the context, and write $f(\mathcal{A}(\cdot))$ as a shorthand for $f(w_{\text{out}})$, ignoring the auxiliary state s_{out} . Note that w_{out} and s_{out} are random variables, determined by the stochastic minibatch ξ .

2.2 Algorithmic stability

Definition 1 (Uniform stability). A deterministic³ algorithm \mathcal{A} is considered to be ϵ -uniformly stable with respect to loss function $f : \mathcal{W} \times \Xi \rightarrow \mathbb{R}$ if, for two batches of data ξ, ξ' differing in exactly one example, we have that

$$\sup_{\xi \in \Xi} |f(\mathcal{A}(\xi), \xi) - f(\mathcal{A}(\xi'), \xi)| \leq \epsilon.$$

The following is a well-known result connecting stability to generalization [7]. Here, we state a version taken from [21]:

Theorem 2. *If an algorithm \mathcal{A} is ϵ -uniformly stable, then it holds that*

$$\left| \mathbb{E}_{\xi \sim \mathcal{D}^B} [\bar{f}_\xi(\mathcal{A}(\xi)) - F(\mathcal{A}(\xi))] \right| \leq \epsilon.$$

3 The data echoing meta-algorithm

Given an minibatch optimization algorithm \mathcal{A} , its data-echoed extension is defined by Algorithm 1.

Algorithm 1 Data echoing meta-algorithm

- 1: **Input:** Optimizer \mathcal{A} ; initializer $w_{\text{init}} := w_0$; initial state $s_{\text{init}} := s_0$; number of inner steps K
 - 2: **for** $t = 0, \dots, T - 1$ **do**
 - 3: Receive a batch of examples $\xi^{(t)} = \{\xi^{(t,i)}\}_{i=1}^B$.
 - 4: Execute \mathcal{A} on $\xi^{(t)}$ starting at w_t for K steps: $(w_{t+1}, s_{t+1}) \leftarrow \mathcal{A}(w_t, s_t, \xi^{(t)}, K)$.
 - 5: **Output:** Average iterate $w_{\text{out}} := \frac{1}{T} \sum_{t=0}^{T-1} w_t$
-

3.1 Data-echoed algorithms

Using the framework of Algorithm 1, we introduce the data-echoed versions of three ubiquitous optimization algorithms. In [13], several types of data echoing are defined; we focus on what the authors call *batch echoing*.

Data-echoed gradient descent. We first formalize gradient descent in our optimization framework. The gradient descent procedure only contains the *fixed* learning rate as the state:

$$s_{\text{init}} = s_{\text{out}} := \{\eta\}.$$

The iterations defining the inner algorithm \mathcal{A} are straightforward:

$$w_0 = w_{\text{init}}, \quad \{w_{j+1} = w_j - \eta \nabla \bar{f}_\xi(w_j)\}_{j=0}^{K-1}, \quad w_{\text{out}} = w_K.$$

When Algorithm 1 is instantiated with this choice of \mathcal{A} , we call the overall procedure *data-echoed gradient descent*.

Data-echoed proximal gradient descent. The state of the proximally-regularized gradient descent procedure contains three variables: the fixed learning rate η , the prox parameter γ , and w_{pivot} , the center of the prox term:

$$s_{\text{init}} := \{\eta, \gamma, w_{\text{pivot}}\}.$$

We now define the proximal function

$$\bar{f}_{\text{prox}}(w) = \bar{f}_\xi(w) + \frac{\gamma}{2} \|w - w_{\text{pivot}}\|^2.$$

³A similar definition exists for randomized algorithms [7]. In this work, we focus on deterministic algorithms.

The iterations proceed in same way as gradient descent, but on \bar{f}_{prox} :

$$w_0 = w_{\text{init}}, \quad \{w_{j+1} = w_j - \eta \nabla \bar{f}_{\text{prox}}(w_j)\}_{j=0}^{K-1}, \quad w_{\text{out}} = w_K.$$

The output returned is $s_{\text{out}} = \{\eta, \gamma, \frac{1}{K} \sum_{j=0}^{K-1} w_j\}$. This particular choice of returning the average iterate as the next w_{pivot} simplifies our analysis. With this choice of \mathcal{A} , this overall procedure will be called *data-echoed proximal gradient descent*.

Data-echoed accelerated gradient descent. The state space for accelerated gradient consists of a step size η , an initial momentum vector d , and a momentum scale factor λ ; thus $s_{\text{init}} = \{\eta, d, \lambda\}$. Define the following scalar sequences with $\lambda_0 = \lambda$:

$$\lambda_{j+1}^2 - \lambda_{j+1} = \lambda_j^2, \quad \gamma_{j+1} = \frac{\lambda_j - 1}{\lambda_{j+1}}.$$

The updates now follow the progression as in Nesterov's acceleration [36]:

$$w_0 = w_{\text{init}}, \quad d_0 = d, \quad w_{j+1} = (w_j + d_j) - \eta \nabla \bar{f}_{\xi}(w_j + d_j), \quad d_{j+1} = \gamma_{j+1}(w_{j+1} - w_j).$$

Finally, the outputs are given by $s_{\text{out}} = \{\eta, d_K, \lambda_K\}$, $w_{\text{out}} = w_K$.

With this choice of \mathcal{A} , we refer to the overall procedure as *data-echoed accelerated gradient descent*.

4 Convergence analyses of echoed methods

We will analyze the data-echoing algorithms by separating their optimization properties from their stability properties. For the latter, we use the standard notion of uniform stability, as defined earlier. For the optimization part, we use a notion of potential-bounded regret, which we define next.

Definition 3 (Potential-bounded regret). We say that an algorithm \mathcal{A} has *potential-bounded regret* with potential function $V_{\mathcal{A}}$ if given a β smooth convex function f on a domain \mathcal{W} and a starting point w_{init} , \mathcal{A} produces a point w_{out} such that for all $w^* \in \mathcal{W}$, it holds that

$$f(w_{\text{out}}) - f(w^*) \leq V_{\mathcal{A}}(w_{\text{init}}, s_{\text{init}}, w^*) - V_{\mathcal{A}}(w_{\text{out}}, s_{\text{out}}, w^*).$$

This inequality is a fundamental lemma in the standard analysis of mirror descent (see [6], or Section B.2 from [1]), but we extend it to *nested stateful algorithms* instead of a single step. For the echoed algorithms we analyze in this work, squared Euclidean norms will be suitable potentials.

We state and prove our main generic theorem below:

Theorem 4. *Let \mathcal{A} be an ϵ -uniformly stable algorithm. Furthermore, suppose \mathcal{A} has the potential-bounded regret property with respect to $V_{\mathcal{A}}$. Then, for any $w^* \in \mathcal{W}$, Algorithm 1 with inner algorithm \mathcal{A} satisfies*

$$\mathbb{E}[F(w_{\text{out}})] - F(w^*) \leq \frac{V_{\mathcal{A}}(w_0, s_0, w^*) - \mathbb{E}[V_{\mathcal{A}}(w_T, s_T, w^*)]}{T} + \epsilon.$$

Proof. From the potential-bounded regret property of the algorithm \mathcal{A} , we get that

$$\bar{f}_{\xi^{(t)}}(w_{t+1}) - \bar{f}_{\xi^{(t)}}(w^*) \leq V_{\mathcal{A}}(w_t, s_t, w^*) - V_{\mathcal{A}}(w_{t+1}, s_{t+1}, w^*).$$

Let $\mathbb{E}_t[\cdot]$ denote the expectation conditioned on all randomness in the minibatches up to (and including) time t . We now get from the uniform stability of \mathcal{A} that

$$\mathbb{E}[F(w_{t+1})] = \mathbb{E}_{t-1} \mathbb{E}_{\xi^{(t)}} [F(w_{t+1})] \leq \mathbb{E}_{t-1} \left[\mathbb{E}_{\xi^{(t)}} [\bar{f}_{\xi^{(t)}}(w_{t+1})] + \epsilon \right].$$

Thus we have

$$\begin{aligned} \mathbb{E}[F(w_{t+1})] - F(w^*) &\leq \mathbb{E}_{t-1} \mathbb{E}_{\xi^{(t)}} [\bar{f}_{\xi^{(t)}}(w_{t+1}) - \bar{f}_{\xi^{(t)}}(w^*)] + \epsilon \\ &\leq \mathbb{E}_{t-1} \mathbb{E}_{\xi^{(t)}} [V_{\mathcal{A}}(w_t, s_t, w^*) - V_{\mathcal{A}}(w_{t+1}, s_{t+1}, w^*)] + \epsilon \\ &\leq \mathbb{E}[V_{\mathcal{A}}(w_t, s_t, w^*)] - \mathbb{E}[V_{\mathcal{A}}(w_{t+1}, s_{t+1}, w^*)] + \epsilon. \end{aligned}$$

Summing the above over time and using the convexity of F gives us that

$$\begin{aligned} \mathbb{E}[F(w_{\text{out}})] - F(w^*) &\leq \sum_{t=0}^{T-1} \frac{\mathbb{E}[F(w_{t+1})] - F(w^*) + \epsilon}{T} \\ &\leq \frac{V_{\mathcal{A}}(w_0, s_0, w^*) - \mathbb{E}[V_{\mathcal{A}}(w_T, s_T, w^*)]}{T} + \epsilon. \quad \blacksquare \end{aligned}$$

In the rest of the section, we present various applications of our main data echoing theorem. In each case, we will consider a standard algorithm, derive its stability and potential bounded regret properties, then use Theorem 4 to derive the convergence rate for its echoed version. All regret proofs can be found in Appendix A, and stability proofs in Appendix B; the corresponding convergence rates for the echoed algorithms are proven in Appendix C.

4.1 Echoed gradient descent

We begin by establishing the following properties of gradient descent. In the rest of the theorem and lemma statements in this section w^* is an arbitrary point in \mathcal{W} .

Lemma 5 (Potential-bounded regret for GD). *Let f be a β -smooth convex function. Then K steps of gradient descent on f , with a step size $\eta \leq 1/\beta$, satisfies the potential-bounded regret property with $V(w, s, w^*) := \frac{1}{2}\|w - w^*\|^2$:*

$$f(w_{\text{out}}) - f(w^*) \leq \frac{1}{\eta K} \left(\frac{\|w_{\text{init}} - w^*\|^2}{2} - \frac{\|w_{\text{out}} - w^*\|^2}{2} \right).$$

Lemma 6 (Stability of GD). *For a β -smooth function f , and any $0 \leq \eta \leq 1/\beta$, gradient descent on f , run with step size η for K steps, is ϵ -uniformly stable with $\epsilon = 2\eta K \rho^2 / B$.*

Combining Lemmas 5 and 6, we conclude the following convergence bound for data-echoed GD:

Theorem 7 (Data-echoed GD). *T outer steps of data-echoed gradient descent, with a step size of $\eta = \min \left\{ \frac{1}{\beta}, \frac{\rho}{KD} \sqrt{\frac{B}{T}} \right\}$ and K internal steps, produces a point w_{out} satisfying*

$$\mathbb{E}[F(w_{\text{out}})] - F(w^*) \leq \frac{\beta D^2}{2KT} + \frac{2\rho D}{\sqrt{BT}}.$$

4.2 Echoed proximal gradient descent

For proximal GD, we derive the following bounds on potential-bounded regret and stability:

Lemma 8. *Let f be a β -smooth convex function. Consider the potential function*

$$V(w, \{\eta, \gamma, w_{\text{pivot}}\}, w^*) = \frac{\|w - w^*\|^2}{2\eta K} + \frac{\gamma \|w - w_{\text{pivot}}\|^2}{2}.$$

Then K -step proximal gradient descent, with step-size $\eta \leq 1/(\beta + \gamma)$ has regret bounded by

$$f(w_{\text{out}}) - f(w^*) \leq V(w_{\text{out}}, s_{\text{out}}, w^*) - V(w_{\text{init}}, s_{\text{init}}, w^*).$$

Lemma 9 (Stability of prox-GD). *For a β -smooth function f , any $\lambda \geq 0$ and any $0 \leq \eta \leq 1/(\beta + \lambda)$, K steps of proximal gradient descent are ϵ -uniformly stable with $\epsilon = \frac{2\rho^2}{B\gamma} (1 - (1 - \eta\gamma)^K)$.*

The proofs of both lemmas are included in the the supplementary material. Combining Lemmas 8 and 9, we get the following guarantee on the performance of data-echoed prox-GD (proof included in the supplementary material):

Theorem 10 (Data-echoed prox-GD). *T outer steps of echoed gradient descent, with a prox parameter of $\gamma = \frac{\rho}{D} \sqrt{\frac{T}{B}}$, step size $\eta = \frac{1}{\beta + \gamma}$, and K internal steps, produces a point w_{out} satisfying*

$$\mathbb{E}[F(w_{\text{out}})] - F(w^*) \leq \sqrt{1 + \frac{1}{K}} \cdot \frac{2\rho \|w_{\text{init}} - w^*\|}{\sqrt{BT}} + \frac{\beta \|w_{\text{init}} - w^*\|^2}{2KT}.$$

Note that using this algorithm, the correct choice of step size η no longer depends on the echoing factor K . In fact, even if K varies across the execution of the proximal algorithm, a straightforward extension of our analysis shows that proximal gradient descent can achieve $\sum_t K_t$ instead of the KT factor in the denominator of the bias term. This resilience to indeterminate echoing factors is especially appealing for the motivating setting of asynchronous pipelines.

4.3 Echoed accelerated gradient descent

For the case of Nesterov’s accelerated gradient descent, we consider a slightly modified version of our data-echoing meta-procedure. This arises from the fact that even the stochastic setting of accelerated gradient descent, algorithms output the final iterate and not the average iterate. The resulting slightly modified procedure is outlined in Algorithm 2.

Algorithm 2 Data-echoing meta-algorithm (final iterate)

- 1: **Input:** Optimizer \mathcal{A} ; initializer $w_{\text{init}} := w_0$; initial state $s_{\text{init}} := s_0$; number of inner steps K .
 - 2: **for** $t = 0, \dots, T - 1$ **do**
 - 3: Receive a batch of examples $\xi^{(t)} = \{\xi^{(t,i)}\}_{i=1}^B$.
 - 4: Execute \mathcal{A} on $\xi^{(t)}$ starting at w_t for K steps: $w_{t+1}, s_{t+1} \leftarrow \mathcal{A}(w_t, s_t, \xi^{(t)}, K_t)$.
 - 5: **Output:** Final iterate $w_{\text{out}} := w_T$
-

We also add a slight extension to our potential-based regret abstraction:

Lemma 11 (Potential-bounded regret for AGD). *Let f be a β -smooth convex function. Running accelerated gradient descent for K steps, with a step size $\eta \leq 1/\beta$, gives the regret bound*

$$\begin{aligned} & (\lambda_{\text{out}}^2 - \lambda_{\text{out}})(f(w_{\text{out}}) - f(w)) - (\lambda_{\text{init}}^2 - \lambda_{\text{init}})(f(w_{\text{init}}) - f(w)) \\ & \leq \frac{1}{2\eta} (\|w_{\text{init}} + \lambda_{\text{init}}d_{\text{init}} - w\|^2 - \|w_{\text{out}} + \lambda_{\text{out}}d_{\text{out}} - w\|^2). \end{aligned}$$

Further, to bound the stability, we note the following lemma which was essentially proved in [11]. Since we believe there is a small typo in the main argument in the original presentation of the proof, we provide an alternate derivation in the supplementary material.

Lemma 12 (Stability of AGD). *Suppose that f is a β -smooth convex quadratic function of w for any ξ . Then, for any $0 \leq \eta \leq 1/\beta$ and initial state s_{init} , K steps of accelerated gradient descent are ϵ -uniformly stable with $\epsilon = O(\eta\rho^2K^2/B)$.*

Combining Lemmas 11 and 12 we obtain the following guarantee for data-echoed AGD:

Theorem 13. *Suppose f is a convex quadratic in w , for all ξ . Then, T outer steps of echoed AGD, with echoing factor K and step size $\eta = \Theta(\min\{\frac{1}{\beta}, \frac{\rho}{K^2D\sqrt{B/T^{3/2}}}\})$, produces a point w_{out} satisfying*

$$\mathbb{E}[F(w_{\text{out}})] - F(w^*) = O\left(\frac{\beta\|w_0 - w^*\|^2}{K^2T^2} + \frac{\rho\|w_0 - w^*\|}{\sqrt{BT}}\right).$$

5 Experiments

We demonstrate numerical experiments on convex machine learning benchmarks. This acts as a validation of our theoretical findings, as well as a way to examine “beyond worst-case” phenomena not captured by our minimax convergence guarantees. This can be seen as a combination of the experiments of Figures 4-6 in [13], where we have exchanged the state-of-the-art setting for a more robust one, allowing for a closer dissection of the bias-variance decomposition.

Methodology. We consider two logistic regression problems as a benchmark, the scaled CoverType dataset from the UCI repository [19], and MNIST [30]. We record the number of iterations (including as well as excluding the data-echoing iterations) needed for SGD to reach within 1% of the optimum training loss, as we increase the echoing factor K , and thus decrease the *rate* of fresh independent samples usable by SGD. For each choice of (B, K) , we tune a constant learning rate by grid search, to minimize this time. All details can be found in the supplementary material.

Results and discussion. Figures 2 and 3 show our findings. As batch size B increases, there is a phase transition from a variance-dominated regime (the $O(\rho D/\sqrt{BT})$ term in our analysis is larger) to a bias-dominated regime (the $O(\beta D/KT)$ term is larger). In the former regime, data-echoed SGD saturates on the stale data, and the optimal learning rate scales inversely with K , as predicted by the theory. In the latter regime, echoing attains a nearly embarrassingly-parallel speedup, and the optimal learning rate is close to constant. These experiments provide an end-to-end example of how the bias-variance decomposition can help to understand and diagnose the benefits and limitations of data-echoed algorithms.




Figure 2: Convergence times as a function of echoing factor K , for logistic regression on the CoverType dataset. Learning rates (yellow) are tuned for each (B, K) to minimize convergence times. Convergence times are presented in number of SGD steps KT (blue), as well as number of independent samples consumed BT (red). Note that the red curves reflect wall-clock time for data-echoing when the data loader is K times slower than the optimizer. As batch size B increases, we move from the noise-dominated regime (red curve plateaus) to the curvature-dominated regime (blue curve plateaus).




Figure 3: Convergence times, as in Figure 2, for logistic regression on the MNIST dataset. Note that the phase transition from noise-dominated to curvature-dominated regimes happens in a batch size range commonly used in deep learning benchmarks with this dataset.

A note on deep neural nets. Our theoretical setting was originally motivated by hardware constraints most frequently encountered in the massively parallel training of deep neural networks. Beyond the convex setting, we note that the experimental design problem become significantly more challenging. Some potential confounds include the learning rate choice affecting the generalization gap [23], and counterintuitive interactions between learning rate and batch normalization [3, 31]. In [13], the authors study the *end-to-end* performance gains of data echoing. Indeed, those experiments need many tweaks (like *example-wise* echoing, data re-augmentation, and individually tuned momentum and learning rate schedules) to obtain their most impressive speedups.

6 Conclusion

We have established first theoretical analysis in the nascent field of optimization algorithms for asynchronous data pipelines, where we have found that gradient descent and well-known variants can be adapted to resist overfitting to stale data. An immediate open problem is to develop a corresponding theory for local convergence and saddle point avoidance in the non-convex setting. This work provides further motivation to show the $O(\eta\rho^2K^2/B)$ -uniform stability of AGD for smooth convex functions, which was conjectured in [11] with different motives. More broadly, we hope that the design and analysis of algorithms in optimization for machine learning can derive fruitful inspiration from nascent hardware considerations, like those that motivated this work.

Broader Impact

This work is theoretical in nature, and is concerned with the very general framework of stochastic optimization. As such, there are no foreseen ethical or societal consequences for the research presented herein. We hope that by providing theoretical groundwork and algorithmic techniques for efficient large-scale optimization in settings informed by modern developments in optimization, works like this one will contribute to alleviating the steep resource and energy costs of large-scale machine learning.

Acknowledgments and Disclosure of Funding

We are grateful to George Dahl and Yoram Singer for helpful discussions. CZ was supported by Elad Hazan’s NSF grants IIS-1523815 and CCF1704860, and a Google internship.

References

- [1] Z. Allen-Zhu and L. Orecchia. Linear coupling: An ultimate unification of gradient and mirror descent. *arXiv preprint arXiv:1407.1537*, 2014.
- [2] R. Anil, V. Gupta, T. Koren, and Y. Singer. Memory-efficient adaptive optimization for large-scale learning. *arXiv preprint arXiv:1901.11150*, 2019.
- [3] S. Arora, Z. Li, and K. Lyu. Theoretical analysis of auto rate-tuning by batch normalization. *arXiv preprint arXiv:1812.03981*, 2018.
- [4] A. Bapna, C. A. Cherry, D. D. Lepikhin, G. Foster, M. Krikun, M. Johnson, M. Chen, N. Ari, O. Firat, W. Macherey, Y. Wu, Y. Cao, and Z. Chen. Massively multilingual neural machine translation in the wild: Findings and challenges, 2019.
- [5] T. Ben-Nun and T. Hoefler. Demystifying parallel and distributed deep learning: An in-depth concurrency analysis. *ACM Computing Surveys (CSUR)*, 52(4):1–43, 2019.
- [6] A. Ben-Tal and A. Nemirovski. *Lectures on modern convex optimization: analysis, algorithms, and engineering applications*. SIAM, 2001.
- [7] O. Bousquet and A. Elisseeff. Stability and generalization. *Journal of machine learning research*, 2(Mar):499–526, 2002.
- [8] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [9] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
- [10] X. Chen, N. Agarwal, E. Hazan, C. Zhang, and Y. Zhang. Extreme tensoring for low-memory preconditioning. *arXiv preprint arXiv:1902.04620*, 2019.
- [11] Y. Chen, C. Jin, and B. Yu. Stability and convergence trade-off of iterative optimization algorithms. *arXiv preprint arXiv:1804.01619*, 2018.
- [12] S. W. Chien, S. Markidis, C. P. Sishtla, L. Santos, P. Herman, S. Narasimhamurthy, and E. Laure. Characterizing deep-learning i/o workloads in tensorflow. In *2018 IEEE/ACM 3rd International Workshop on Parallel Data Storage & Data Intensive Scalable Computing Systems (PDSW-DISCS)*, pages 54–63. IEEE, 2018.
- [13] D. Choi, A. Passos, C. J. Shallue, and G. E. Dahl. Faster neural network training with data echoing. *arXiv preprint arXiv:1907.05550*, 2019.
- [14] D. Ciregan, U. Meier, and J. Schmidhuber. Multi-column deep neural networks for image classification. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3642–3649. IEEE, 2012.
- [15] E. D. Cubuk, B. Zoph, D. Mane, V. Vasudevan, and Q. V. Le. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 113–123, 2019.

- [16] S. Dalton, I. Frosio, and M. Garland. Gpu-accelerated atari emulation for reinforcement learning. *arXiv preprint arXiv:1907.08467*, 2019.
- [17] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009.
- [18] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [19] D. Dua and C. Graff. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- [20] J. A. Guirao, K. Łęcki, J. Lisiecki, S. Panev, M. Szolucha, A. Wolant, and M. Zientkiewicz. Fast AI Data Preprocessing with NVIDIA DALI, 2019. URL <https://developer.nvidia.com/blog/fast-ai-data-preprocessing-with-nvidia-dali/>.
- [21] M. Hardt, B. Recht, and Y. Singer. Train faster, generalize better: Stability of stochastic gradient descent. *arXiv preprint arXiv:1509.01240*, 2015.
- [22] E. Hoffer, T. Ben-Nun, I. Hubara, N. Giladi, T. Hoefler, and D. Soudry. Augment your batch: better training with larger batches. *arXiv preprint arXiv:1901.09335*, 2019.
- [23] Z. Jiang, C. Zhang, K. Talwar, and M. C. Mozer. Exploring the memorization-generalization continuum in deep learning. *arXiv preprint arXiv:2002.03206*, 2020.
- [24] N. P. Jouppi, C. Young, N. Patil, D. Patterson, G. Agrawal, R. Bajwa, S. Bates, S. Bhatia, N. Boden, A. Borchers, et al. In-datacenter performance analysis of a tensor processing unit. In *Proceedings of the 44th Annual International Symposium on Computer Architecture*, pages 1–12, 2017.
- [25] P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings, et al. Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977*, 2019.
- [26] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, and D. Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.
- [27] M. Khadatare, Z. Khan, and H. Bayraktar. Leveraging the Hardware JPEG Decoder and NVIDIA nvJPEG Library on NVIDIA A100 GPUs, 2020. URL <https://developer.nvidia.com/blog/leveraging-hardware-jpeg-decoder-and-nvjpeg-on-a100/>.
- [28] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [29] G. Lan. An optimal method for stochastic composite optimization. *Mathematical Programming*, 133(1-2):365–397, 2012.
- [30] Y. LeCun, C. Cortes, and C. Burges. Mnist handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2, 2010.
- [31] Z. Li and S. Arora. An exponential learning rate schedule for deep learning. *arXiv preprint arXiv:1910.07454*, 2019.
- [32] J. Liang, V. Makoviychuk, A. Handa, N. Chentanez, M. Macklin, and D. Fox. Gpu-accelerated robotic simulation for distributed reinforcement learning. In *Conference on Robot Learning*, pages 270–282, 2018.
- [33] R. Mayer and H.-A. Jacobsen. Scalable deep learning on distributed infrastructures: Challenges, techniques, and tools. *ACM Computing Surveys (CSUR)*, 53(1):1–37, 2020.
- [34] S. McCandlish, J. Kaplan, D. Amodei, and O. D. Team. An empirical model of large-batch training. *arXiv preprint arXiv:1812.06162*, 2018.

- [35] Y. Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2013.
- [36] Y. E. Nesterov. A method of solving a convex programming problem with convergence rate $o(k^2)$. In *Doklady Akademii Nauk*, volume 269, pages 543–547. Russian Academy of Sciences, 1983.
- [37] R. Raina, A. Madhavan, and A. Y. Ng. Large-scale deep unsupervised learning using graphics processors. In *Proceedings of the 26th annual international conference on machine learning*, pages 873–880, 2009.
- [38] C. J. Shallue, J. Lee, J. Antognini, J. Sohl-Dickstein, R. Frostig, and G. E. Dahl. Measuring the effects of data parallelism on neural network training. *arXiv preprint arXiv:1811.03600*, 2018.
- [39] N. Shazeer and M. Stern. Adafactor: Adaptive learning rates with sublinear memory cost. *arXiv preprint arXiv:1804.04235*, 2018.
- [40] N. Shazeer, A. Mirhoseini, K. Maziarz, A. Davis, Q. Le, G. Hinton, and J. Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*, 2017.
- [41] C. Shorten and T. M. Khoshgoftaar. A survey on image data augmentation for deep learning. *Journal of Big Data*, 6(1):60, 2019.
- [42] S. L. Smith, P.-J. Kindermans, C. Ying, and Q. V. Le. Don’t decay the learning rate, increase the batch size. *arXiv preprint arXiv:1711.00489*, 2017.
- [43] J. Wang, W. Wang, and N. Srebro. Memory and communication efficient distributed stochastic optimization with minibatch-prox. *arXiv preprint arXiv:1702.06269*, 2017.
- [44] B. Woodworth, K. K. Patel, S. U. Stich, Z. Dai, B. Bullins, H. B. McMahan, O. Shamir, and N. Srebro. Is local sgd better than minibatch sgd? *arXiv preprint arXiv:2002.07839*, 2020.
- [45] C. Ying, S. Kumar, D. Chen, T. Wang, and Y. Cheng. Image classification at supercomputer scale. *arXiv preprint arXiv:1811.06992*, 2018.
- [46] Y. You, I. Gitman, and B. Ginsburg. Large batch training of convolutional networks. *arXiv preprint arXiv:1708.03888*, 2017.