
Batch-Expansion Training: An Efficient Optimization Framework*

Michał Dereziński
Dept. of Computer Science
UC Santa Cruz
mderezin@ucsc.edu

Dhruv Mahajan
Facebook Research
Menlo Park
dhruvm@fb.com

S. Sathiya Keerthi
Microsoft Corporation
Mountain View
keerthi@microsoft.com

S. V. N. Vishwanathan
Dept. of Computer Science
UC Santa Cruz
vishy@ucsc.edu

Markus Weimer
Microsoft Corporation
Mountain View
mweimer@microsoft.com

Abstract

We propose Batch-Expansion Training (BET), a framework for running a batch optimizer on a gradually expanding dataset. As opposed to stochastic approaches, batches do not need to be resampled i.i.d. at every iteration, thus making BET more resource efficient in a distributed setting, and when disk-access is constrained. Moreover, BET can be easily paired with most batch optimizers, does not require any parameter-tuning, and compares favorably to existing stochastic and batch methods. We show that when the batch size grows exponentially with the number of outer iterations, BET achieves optimal $\tilde{O}(1/\epsilon)$ data-access convergence rate for strongly convex objectives. Experiments in parallel and distributed settings show that BET performs better than standard batch and stochastic approaches.

parameter updates at a much greater frequency. Both approaches present different trade-offs. Stochastic updates often provide very good early performance, since they can update the parameters a considerable number of times before even a single batch update finishes. On the other hand, by accessing the full dataset at each step, batch algorithms can better utilize second-order information about the loss function, while taking advantage of parallel and distributed architectures. Finding approaches that provide the best of both worlds is an important area of research. In this paper, we propose a new framework which – while closer to batch in spirit – enjoys the benefits of stochastic methods. In addition, our framework addresses a practical issue observed in real-world industrial settings, which we now describe.

In industrial server farms, compute resources for large jobs become available only in a phased manner. When running a batch optimizer, which requires loading the entire dataset, one has to wait until all the machines become available before beginning the computation. Moreover, the training data, which typically consists of user logs, is distributed across multiple locations. This data needs to be normalized, often by communicating summary statistics or subsets of the data across the network. Since stochastic optimization algorithms only deal with a subset of the data at a time, one can largely avoid the bottleneck of data normalization by preparing mini-batches at a time. However, now each data point needs to be visited multiple times randomly, and this requires performing many random accesses from a hard-disk or network attached storage (NAS), which is inherently slow. Moreover, extending stochastic optimization to the distributed setting is still an active area of research. This raises the question of whether the compute and data-availability delays could be avoided in batch methods, and how that would affect the training time.

1 Introduction

State-of-the-art optimization algorithms used in machine learning broadly tend to fall into two main categories: batch methods, which visit the entire dataset once before performing an expensive parameter update, and stochastic methods, which rely on a small subset of training data, to apply quick

*This research was supported by NSF grant IIS-1546452.

We propose Batch-Expansion Training (BET), an adaptive, parameter-free meta-algorithm, which can be paired with most batch optimization methods to accelerate their performance in sequential, parallel, as well as distributed settings, addressing the issues discussed above. Our approach hinges on the following observation: initially, when only a subset of the data is available, the statistical error (the error that arises because one is observing a sample from the true underlying data distribution) is large. Therefore, we can tolerate a large optimization error (the error that arises because of the iterative nature of the batch optimizer). However, as the sample size increases, the statistical error decreases and the corresponding optimization error that we can tolerate also decreases. BET exploits this by initially training models with large optimization error on smaller subsets of the data, and iteratively loading more training data, and driving down the optimization error.

Relying on classical results in statistical learning theory, we show that any optimizer exhibiting linear convergence rate can be effectively accelerated by periodically doubling the data size after a certain number of iterations. We propose a simple parameter-free algorithm, which dynamically decides the optimal number of iterations between each data expansion, adapting to the performance of a batch optimizer provided as a black box method. Experiments using Non-linear Conjugate Gradient (CG), as well as the L-BFGS method, demonstrate the versatility of our framework. We show that for strongly convex losses BET achieves the same asymptotic $\tilde{O}(1/\epsilon)$ convergence rate as SGD in terms of data accesses (and is strictly faster than regular batch updates). However, unlike stochastic methods, BET reuses all of the data that has already been loaded, which helps it scale very well in parallel and distributed settings, as confirmed by the experimental results.

2 Related Work

There has been a recent explosion of interest in both batch and stochastic optimization methods for machine learning. Algorithms like Stochastic Variance Reduced Gradient method [15] and related approaches [22, 10] mix SGD-like steps with some batch computations to control the stochastic noise. Others have proposed to parallelize stochastic training through large mini-batches [16, 11]. However, these methods do not address the issues of compute and data-availability delays that we discussed above. Additionally, two-stage approaches have been proposed [1, 24], which employ SGD at the beginning followed by a batch optimizer (e.g., L-BFGS). These methods are much more limited than BET, which allows for multiple stages of optimization, with clear practical benefits. Interleaving computation with data loading was shown to have significant practical benefits by [18]. However, that work is confined to training on a single machine and did not provide any theoretical convergence guarantees. In contrast, we largely focus on the distributed

setting and provide convergence guarantees.

Several works explore the idea of batch expansion in various forms. [7, 14] propose using stochastic mini-batches of gradually increasing size paired with optimizers like gradient descent or Newton-CG. The algorithms proposed there strongly rely on stochastic sampling, which makes them less resource-efficient than BET, and not applicable to certain distributed settings. The convergence guarantees offered in [7] are limited to using gradient descent as the inner optimizer, and heavily rely on the independence conditions present in stochastic sampling. The idea of gradually increasing batch size without resampling the batches was first discussed by [27], however their convergence analysis is also limited to gradient descent. Recently, [19, 9] proposed variants of the Newton’s method and of the SAGA [10] algorithm which run on increasing batch sizes. Both of those methods pose challenges in large-scale distributed settings (expensive hessian computation required for Newton’s method and sequential stochastic updates in SAGA). The key advantage of our meta-algorithm is that it can be used with any inner batch optimizer as a black box, including quasi-Newton methods like L-BFGS, making this approach more broadly applicable and scalable. Similarly, our complexity analysis applies to a wide range of inner optimizers.

3 Batch-Expansion Training

We consider a standard composite convex optimization problem arising in regularized linear prediction. Given a dataset $\mathbf{Z} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$, we aim to approximately minimize the average regularized loss

$$\hat{f} \triangleq \frac{1}{N} \sum_{i=1}^N \ell_{z_i}(\mathbf{w}) + \frac{\lambda}{2} \|\mathbf{w}\|^2, \quad (1)$$

where $z_i = (\mathbf{x}_i, y_i)$ and $\ell_{z_i}(\mathbf{w}) \triangleq \ell(\langle \mathbf{w}, \phi(\mathbf{x}_i) \rangle, y_i)$ is the loss of predicting with a linear model $\langle \mathbf{w}, \phi(\mathbf{x}_i) \rangle$ against a target label y_i . Any iterative optimization algorithm in this setting will produce a sequence of models $\{\mathbf{w}_t\}_{t=1}^T$ with the goal that \mathbf{w}_T has small optimization error $\hat{g}(\mathbf{w}_T)$ with respect to the exact optimum $\hat{\mathbf{w}}^*$, where

$$\hat{g}(\mathbf{w}) \triangleq \hat{f}(\mathbf{w}) - \hat{f}(\hat{\mathbf{w}}^*) \quad \text{and} \quad \hat{\mathbf{w}}^* \triangleq \arg \min_{\mathbf{w}} \hat{f}(\mathbf{w}).$$

Note that our true goal is to predict well on an unseen example $z = (\mathbf{x}, y)$ coming from an underlying distribution. The right regularization λ for this task can be determined experimentally or from the statistical guarantees of loss function ℓ , as discussed in [25, 23]. In this paper, however, we will assume that an acceptable λ was chosen and concentrate on the problem of minimizing \hat{f} .

3.1 Linear convergence of the batch optimizer

Adding ℓ_2 -norm regularization in function \hat{f} makes it λ -strongly convex. In this setting, many popular batch op-

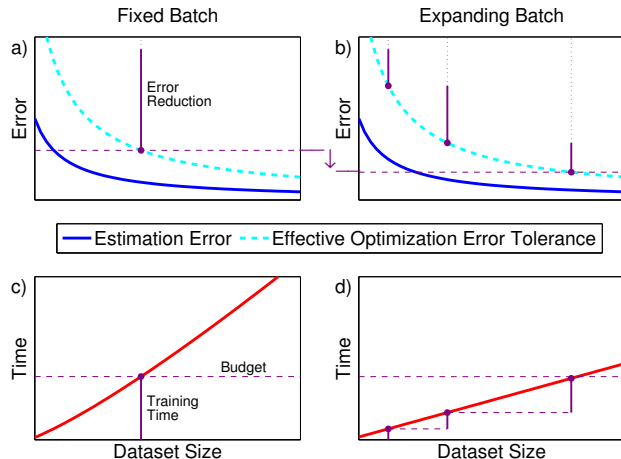


Figure 1: Estimation error and effective optimization error tolerance are inversely proportional to the dataset size (plots a and b). Fixed Batch preselects a data size, then reduces the optimization error (vertical line in a), whereas Expanding Batch divides the work into stages with different data sizes (vertical lines in b). Optimization is faster for smaller data sizes (see plots c and d), so Expanding Batch reaches smaller effective optimization error tolerance within the same time budget compared to Fixed Batch (plots a and b).

imization algorithms enjoy linear convergence rate [21]. Namely, given arbitrary model \mathbf{w} and any $c > 1$, after at most $\mathcal{O}(\log(c))$ iterations - where a single iteration can look at the entire dataset - we can reduce its optimization error by a multiplicative factor of c , obtaining \mathbf{w}' such that

$$\hat{g}(\mathbf{w}') \leq c^{-1} \cdot \hat{g}(\mathbf{w}).$$

Note that the runtime of a single iteration will depend on data size N , but the number of needed iterations does not. Setting $c = 2$, we observe that only a constant number of batch iterations is necessary for a linearly converging optimizer to halve the optimization error of \mathbf{w} . This constant depends on the convergence rate enjoyed by the method. For example, in the case of batch gradient descent we need $\mathcal{O}(1/\lambda)$ iterations to reduce the optimization error by a factor of 2 [6] (note the dependence on the strong convexity coefficient), while other methods (like L-BFGS) can achieve better rates of convergence [17]. Furthermore, the time complexity of performing a single iteration for many of those algorithms (including GD and L-BFGS) is linearly proportional to the data size. We refer to methods exhibiting both linear convergence (with respect to a given loss) and linear time complexity of a single iteration as *linear optimizers*. From now on, we only consider this class of methods.

3.2 Dataset size selection

The general task of an optimization algorithm is to return a model with small optimization error $\hat{g}(\mathbf{w}) \leq \epsilon$. For select-

ing an effective optimization error tolerance ϵ in a machine learning problem, we often look at the estimation error exhibited by the objective, i.e. how much \hat{f} deviates from the expected regularized loss measured on a random unseen example. As discussed in [23, 5], the effective optimization error tolerance should be proportional to the estimation error of \hat{f} , because optimizing beyond that point does not yield any improvement on unseen data. Note that under standard statistical assumptions, the more data we use, the smaller the estimation error becomes [25], and thus, the effective optimization error tolerance should also be decreased (see Figure 1a).

Consider the scenario where data is abundant, but we have a limited time budget for training the model¹. In this case, a practitioner would select the data size so that we can reach the smallest effective optimization error tolerance (dashed curve in Figure 1a) within the allotted training time. Vertical line in Figure 1a shows the reduction in optimization error that the algorithm has to achieve to obtain the desired tolerance. If we use a batch optimizer for this task, the training time is determined by two factors. First, as the dataset grows, each iteration takes longer (e.g. for linear optimizers the iteration time is proportional to the data size, as discussed in Section 3.1). Second, the algorithm has to perform larger number of iterations, the smaller the effective optimization error tolerance. Combined, those two effects result in training time, which, for linear optimizers, grows faster than linearly with the data size (see Figure 1c).

In this paper, we propose that instead of finding the optimal dataset size at the beginning, we first load a small subset of data, train the model until we reach the corresponding effective optimization error tolerance, then we load more data, and optimize further, etc. Vertical lines in Figure 1b illustrate how optimization error is reduced in multiple stages working with increasing data sizes. This procedure benefits from faster iterations in the early stages, as shown in Figure 1d, with vertical lines corresponding to the training time for each data size, and the horizontal line showing the time budget (note that the budget shown for this method is the same as the one used in Figure 1c). Thus, our approach is able to reach smaller effective optimization error tolerance within the same time budget compared to fixing the dataset size at the beginning (see Figures 1a and 1b).

3.3 Exponentially increasing batches

We now precisely formulate the idea of training with gradually increasing data size. Suppose our goal is to return a model $\hat{\mathbf{w}}$ with optimization error $\hat{g}(\hat{\mathbf{w}}) \leq \epsilon$. In the procedure described above, we seek to obtain a sequence of gradually improving models $\mathbf{w}_1, \dots, \mathbf{w}_{T-1}$ before we reach $\mathbf{w}_T = \hat{\mathbf{w}}$. Any of the intermediate models (say, model \mathbf{w}_t for $t < T$) has a large optimization error relative to $\hat{\mathbf{w}}$, so

¹This is a reasonable practice in many web applications.

our model. However, that constant may depend on the type of loss function, the dataset, as well as the optimizer used, and moreover, in practice the right number of iterations may in fact vary to some extent between the stages. Therefore, we need a practical method of deciding the right time to double the data. Consider the following experiment: we run two optimization tracks in parallel, first one for the batch of size n_t , the other for half of that batch. One update on the bigger batch takes longer than an update on the smaller one. Which track will make better progress towards the optimum of the bigger batch in the same amount of time? If the starting model is far enough from the optimum $\hat{\mathbf{w}}_t^*$, then the faster updates will initially have an advantage. However, as the convergence proceeds, only the slower track can get arbitrarily close, so at some point it will move ahead of the fast one (in terms of the loss $\hat{f}_t(\mathbf{w})$). Denote the starting model as $\mathbf{w}_{t,0}$. The secondary track (running on half of the batch) also starts at the same point, denoted as $\mathbf{w}'_{t-1,0} = \mathbf{w}_{t,0}$, and they are both updated:

$$\begin{aligned}\mathbf{w}_{t,s+1} &\leftarrow \text{Update}(\mathbf{w}_{t,s}, n_t), \\ \mathbf{w}'_{t-1,s+1} &\leftarrow \text{Update}(\mathbf{w}'_{t-1,s}, n_{t-1}),\end{aligned}$$

where $\text{Update}(\mathbf{w}, n)$ is one step of the optimizer, with respect to model \mathbf{w} , on the batch $\{z_i\}_{i=1}^n$. Let c_s and c'_s be the total runtimes of the first and second track after s iterations, respectively, and let $s_1(s) = \max\{s_0 : c_{s_0} \leq c'_s\}$ be the number of iterations that the slower track completes in the time it took the faster track to do s . To compare the performance of the two tracks we will use the condition:

$$\hat{f}_t(\mathbf{w}_{t,s_1(s)}) < \hat{f}_t(\mathbf{w}'_{t-1,s}). \quad (4)$$

Algorithm 2 Two-Track Optimizer

```

Initialize  $\mathbf{w}_{1,0} = \mathbf{w}'_{0,0}$  arbitrarily,  $s \leftarrow 0$ ,  $t \leftarrow 1$ 
Pick any  $2 \leq n_1 = 2n_0 < N$ 
while  $n_t < N$  do
     $\mathbf{w}_{t,s+1} \leftarrow \text{Update}(\mathbf{w}_{t,s}, n_t)$ 
     $\mathbf{w}'_{t-1,s+1} \leftarrow \text{Update}(\mathbf{w}'_{t-1,s}, n_{t-1})$ 
     $s \leftarrow s + 1$ 
     $s_1 \leftarrow \max\{s_0 : c_{s_0} \leq c'_s\}$ 
    if  $\hat{f}_t(\mathbf{w}_{t,s_1}) < \hat{f}_t(\mathbf{w}'_{t-1,s})$  then
         $n_{t+1} \leftarrow 2n_t$ 
         $\mathbf{w}'_{t,0}, \mathbf{w}_{t+1,0} \leftarrow \mathbf{w}_{t,s}$ 
         $t \leftarrow t + 1$ 
         $s \leftarrow 0$ 
    end if
end while
while stopping condition not met do
     $\mathbf{w}_{t,s+1} \leftarrow \text{Update}(\mathbf{w}_{t,s}, N)$ 
     $s \leftarrow s + 1$ 
end while
return  $\mathbf{w}_{t,s}$ 
    
```

Algorithm 2 describes Batch-Expansion Training implemented using the Two-Track strategy. Note that in this

algorithm we run the two tracks sequentially, in an alternating order, however running them in parallel would further improve the overall performance.

3.5 Discussion

We found that the choice to increase data size by a factor of 2 at each stage (rather than by a different factor) is not crucial for the optimization performance (both theoretically and in practice), therefore this parameter does not require tuning. The initial subset size n_0 also does not affect performance significantly - generally, the larger n_0 we select, the more updates will be performed before first data expansion, but as long as the initial subset is small enough, total optimization time will remain close to optimal. Thus, Algorithm 2 does not require any tuning to achieve good performance. Moreover, our method can be paired with many popular batch optimizers, and it will automatically adapt its behavior to the selected inner optimizer, as shown in Section 5.1.

It is important to note that the fraction of data accessed by the algorithm is only gradually expanded as optimization proceeds. Moreover, BET iterates multiple times over the data points that have already been loaded. Thus, it is very resource efficient in a way that can be beneficial with:

Slow disk-access. Loading data from disk to memory can be a significant bottleneck [18]. Performing multiple iterations over the data points while extra data is being loaded in parallel provides speed-up.

Resource ramp-up. In distributed computing, often not all resources are made available immediately at the beginning of the optimization [20], which similarly leads to gradual data availability.

4 Complexity Analysis

In this section, we provide theoretical guarantees for the time complexity of Batch-Expansion Training and compare them to other approaches. For the remainder of this section, we assume that the inner optimizer for some $\kappa > 1$ and for every t , \mathbf{w} exhibits linear convergence:

$$\hat{g}_t(\text{Update}(\mathbf{w}, n_t)) \leq (1 - (1/\kappa)) \hat{g}_t(\mathbf{w}).$$

For the sake of complexity analysis, we discuss a parameterized variant of our approach, described in Algorithm 3, and establish complexity results for it. Here, the number of updates needed at each stage is a fixed parameter $\hat{\kappa}$.

Assumptions. We assume that the feature mapping $\phi(\cdot)$ (see (1) and the discussion below it) is B -bounded, i.e. it satisfies $\|\phi(\mathbf{x}_i)\| \leq B$, and that the loss function $\ell_{z_i}(\mathbf{w}) = \ell(z_i, \mathbf{w})$ is L -Lipschitz in z_i for all \mathbf{w} . Moreover, we will use the fact that \hat{f} is λ -strongly convex, due to the use of ℓ_2 -norm regularization. The result below holds with high probability with respect to a random permutation of data.

Algorithm 3 Optimal BET

Input: Target tolerance ϵ
 Pick ϵ_0, n_0
 Initialize $\mathbf{w}_0 \leftarrow 0$, $t \leftarrow 0$ and $\hat{\kappa} = \lceil \kappa \log(6) \rceil$
while $3 \cdot \epsilon_t > \epsilon$ **do**
 $n_{t+1} \leftarrow 2n_t$
 $\mathbf{w}_{t,0} \leftarrow \mathbf{w}_t$
 for $s = 1.. \hat{\kappa}$ **do**
 $\mathbf{w}_{t,s} \leftarrow \text{Update}(\mathbf{w}_{t,s-1}, n_{t+1})$
 end for
 $\mathbf{w}_{t+1} \leftarrow \mathbf{w}_{t,\hat{\kappa}}$
 $\epsilon_{t+1} \leftarrow \epsilon_t/2$
 $t \leftarrow t + 1$
end while
 $T \leftarrow t$
return \mathbf{w}_T

Theorem 4.1 For any n_0, δ, ϵ there exists ϵ_0 s.t. data-access complexity of Algorithm 3 is

$$\mathcal{O}\left(\frac{\kappa}{\lambda \epsilon} \cdot (\log \log(1/\epsilon) + \log(1/\delta))\right)$$

and $\hat{f}(\mathbf{w}_T) - \hat{f}(\hat{\mathbf{w}}^*) \leq \epsilon$ w.p. at least $1 - \delta$.

Proof Recall that we denote the approximation error estimate at stage t by $\hat{g}_t(\mathbf{w}) \triangleq \hat{f}_t(\mathbf{w}) - \hat{f}_t(\hat{\mathbf{w}}_t^*)$ and the full approximation error as $\hat{g}(\mathbf{w}) \triangleq \hat{f}(\mathbf{w}) - \hat{f}(\hat{\mathbf{w}}^*)$. Moreover, note that Algorithm 3 defines the optimization error tolerances recursively as $\epsilon_{t+1} = \epsilon_t/2$. First, we give the following uniform convergence result:

Lemma 1 For any n_0, δ, T , there exists ϵ_0 such that

$$\epsilon_0 = \mathcal{O}(L^2 B^2 \log(T/\delta) \cdot \lambda^{-1}),$$

and with probability $1 - \delta$, for all \mathbf{w} and all $0 \leq t < T$:

$$\hat{g}_0(\mathbf{w}_0) \leq \epsilon_0, \quad (5)$$

$$\hat{g}_{t+1}(\mathbf{w}) \leq 2 \cdot \hat{g}_t(\mathbf{w}) + \epsilon_t, \quad (6)$$

$$\hat{g}(\mathbf{w}) \leq 2 \cdot \hat{g}_T(\mathbf{w}) + \epsilon_T. \quad (7)$$

See Appendix A.1 for proof. Next, using this lemma, we show the main result. Let κ be the convergence rate of the inner optimizer. Recall, that we set the number of inner iterations of Algorithm 3 to be

$$\hat{\kappa} \triangleq \lceil \kappa \log(6) \rceil.$$

This gives us the following bound for the progress that the inner optimizer makes at each stage (for any $0 \leq t \leq T$):

$$\begin{aligned} \hat{g}_{t+1}(\mathbf{w}_{t+1}) &\leq \left(1 - \frac{1}{\kappa}\right)^{\hat{\kappa}} \hat{g}_{t+1}(\mathbf{w}_t) \\ &\leq \exp\left(-\frac{\hat{\kappa}}{\kappa}\right) \hat{g}_{t+1}(\mathbf{w}_t) \leq \frac{\hat{g}_{t+1}(\mathbf{w}_t)}{6}. \end{aligned}$$

Suppose that ϵ_0 satisfies Lemma 1. We can show by induction that (with probability $1 - \delta$) for all $t \leq T$, model \mathbf{w}_t

is an ϵ_t -approximate solution for \hat{f}_t , i.e. that $\hat{g}_t(\mathbf{w}_t) \leq \epsilon_t$. Base case is given by (5). The inductive step follows from:

$$\hat{g}_{t+1}(\mathbf{w}_{t+1}) \leq \frac{2 \cdot \hat{g}_t(\mathbf{w}_t) + \epsilon_t}{6} \leq \epsilon_{t+1}.$$

Next, we verify that \mathbf{w}_T is an ϵ -approximate solution for \hat{f} :

$$\hat{g}(\mathbf{w}_T) \leq 2 \cdot \hat{g}_T(\mathbf{w}_T) + \epsilon_T \leq 3 \cdot \epsilon_T \leq \epsilon.$$

Finally, we move on to complexity analysis. The number of iterations in the algorithm is $T = \mathcal{O}(\log(\epsilon_0/\epsilon))$, since:

$$2^T = \frac{2 \cdot \epsilon_0}{\epsilon_{T-1}} < 6 \cdot \frac{\epsilon_0}{\epsilon}.$$

Assuming that one update of the inner optimizer requires C passes over the data, we obtain the data-access complexity:

$$\begin{aligned} \sum_{t=1}^T \hat{\kappa} C n_t &= C \hat{\kappa} n_0 \sum_{t=1}^T 2^t \\ &\leq 2C \hat{\kappa} n_0 \cdot 2^T = \mathcal{O}\left(n_0 \epsilon_0 \cdot \frac{\kappa}{\epsilon}\right). \end{aligned}$$

See Appendix A.2 for details regarding the log terms. \square

Using gradient descent as the inner optimizer, we have $\kappa = \mathcal{O}(1/\lambda)$, so Algorithm 3 reaches data-access complexity $\tilde{\mathcal{O}}(1/(\lambda^2 \epsilon))$. However, the general nature of this approach allows us to choose a different linear optimizer with better guarantees, like $\kappa = \mathcal{O}(1/\sqrt{\lambda})$ for accelerated gradient descent. Methods like L-BFGS and other approximate Newton algorithms have been shown to exhibit linear convergence [17, 12] with a rate that does not suffer from such strict dependence on the strong convexity coefficient λ . Hence, when using those optimizers, for most problems we should expect κ to be a small constant factor, in which case data-access complexity becomes $\tilde{\mathcal{O}}(1/(\lambda \epsilon))$.

5 Experiments

In this section we present experimental results showing the benefits of applying BET acceleration to batch optimizers and demonstrating scalability of the method in parallel and distributed settings. As the optimization problem we use logistic loss with ℓ_2 -norm regularization trained on several standard LIBSVM datasets (see Table 1). All algorithms start with the initial model vector \mathbf{w} set to all zeros. BET was implemented as shown in Algorithm 2, without any parameters that required tuning. The results are presented using test set accuracy and training objective, the latter being shown in terms of the log Relative Functional Value Difference

$$\log \text{RFVD: } \log\left(\frac{\hat{f}(\mathbf{w}) - \hat{f}(\hat{\mathbf{w}}^*)}{\hat{f}(\hat{\mathbf{w}}^*)}\right). \quad (8)$$

All algorithms were implemented in the PETSc [2, 3, 4] framework, with the data split between multiple computing cores to achieve parallelization speed-up (we used 16 cores

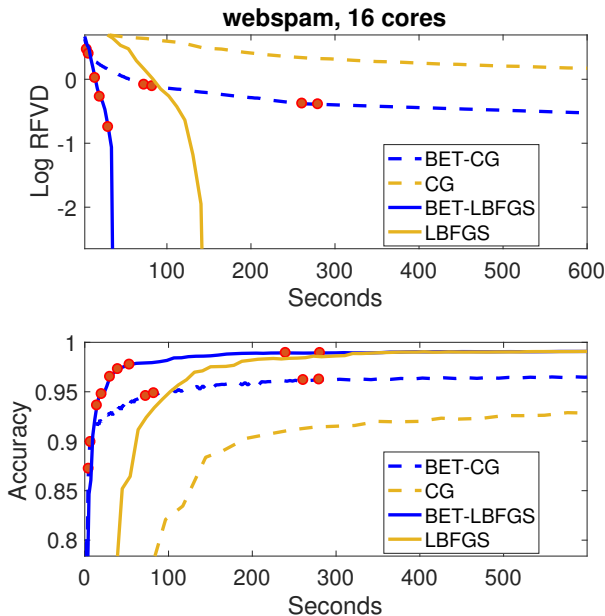


Figure 2: Comparing BET and Batch, for two inner optimizers: Nonlinear-CG (CG) and Limited-memory BFGS (L-BFGS) (webspam dataset, 16 cores). Circular dots in BET denote the batch expansion points.

in most experiments). For splice-site dataset³, data was additionally divided between multiple machines.

Dataset, size	Train/Test	Dim.	λ
url, 1GB	1.8M/0.5M	3.2M	1e-8
covtype, 19GB	0.5M/69k	170k	1e-6
webspam, 30GB	250k/100k	16.6M	1e-6
splice-site, 1.5TB	25M/4.6M	11.7M	2e-10

Table 1: A list of datasets and regularization used for the experiments. For covtype, features were expanded to all monomials of degree 3.

5.1 Adapting to inner optimizers

To demonstrate the flexibility of our framework, in this section we use two different inner optimizers with BET:

1. Nonlinear Conjugate Gradient method (CG), using Fletcher-Reeves [13] formula, with exact line-search;
2. Limited-memory BFGS (L-BFGS) [8].

Both of those methods are linear optimizers that employ strategies for enhancing the basic gradient descent direction. Note, that CG uses a memory vector updated at each iteration. When the loss function changes from \hat{f}_t to \hat{f}_{t+1} , one might expect that memory vector to become invalid, rendering CG ineffective. However, we observe that BET is still very effective at accelerating memory-based algorithms.

³Due to resource constraints, we used a 1.5TB portion of the full 3TB splice-site dataset.

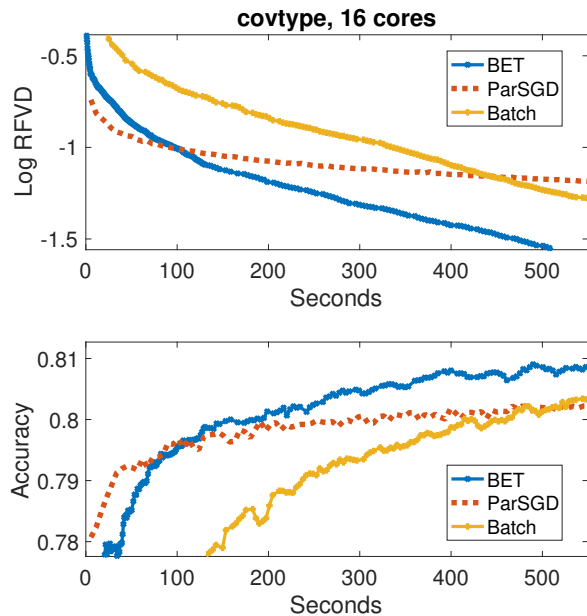


Figure 3: Comparing BET, Batch (L-BFGS) and Parallel SGD for covtype dataset, running on 16 cores. More datasets are presented in Appendix B.

Figure 2 shows the performance comparison of using BET with the two inner optimizers on webspam dataset, contrasted with both of them ran in regular batch mode. First, we can see that L-BFGS is a much more effective optimizer than Nonlinear CG, thus in all of the following experiments we used L-BFGS as the inner optimizer. However, both methods significantly benefit from BET. In fact, in the early phase of the optimization, performance of BET is similar with either of the optimizers. However, once batch expansion reaches close to the full dataset, quality of the underlying optimizer starts to play an important role. Circular dots on the BET plots in Figure 2 mark the points when batch size is doubled during the optimization. We observed that the average number of iterations per stage of BET is larger when the inner optimizer is CG. This matches our expectation, since theory suggests that the number of iterations per stage should be inversely proportional to the convergence rate of the inner optimizer. Moreover, within one optimization run, we saw significant fluctuations in the number of iterations needed before each batch-expansion stage, which means that there is no universal number of iterations per stage which will work well in all settings. Thus, the adaptive capabilities of the two-track algorithm are crucial to achieving good performance of batch-expansion.

5.2 Combining the benefits of batch and stochastic

In this experiment, we examine the trade-offs between batch and stochastic methods, showing how BET fits into this comparison. A batch optimizer typically starts off slower than a stochastic one, but once it gets closer to the optimum, we expect to see a fast convergence behavior. Thus, depend-

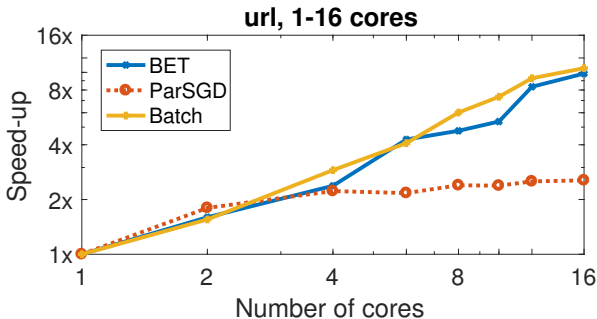


Figure 4: Parallelization speed-up on url dataset for Batch, BET and Parallel SGD.

ing on which tolerance level we select, we would choose the appropriate optimizer. This is shown on a sample plot for the covtype dataset in Figure 3, where we use Parallel SGD [28] as the stochastic algorithm and L-BFGS as the Batch method, with all algorithms running on a 16 core machine (for the sake of clarity, only a portion of the full convergence time is shown). We also tested mini-batch SVRG as an alternative stochastic approach. However, due to the high communication cost, this method exhibited much slower per iteration time, and for this reason we did not include the results in the paper. To obtain the reported performance for Parallel SGD we had to tune the step size on a chunk of data, which adds to the overall optimization time, whereas BET is a parameter-free method. Our algorithm gets the best of both worlds, since it behaves like a stochastic method at the beginning, and then like a batch method towards the end. More plots are presented in Appendix B, showing datasets which favor either batch or stochastic methods. In summary, BET performs as well as the best method for each dataset, with no tuning necessary.

5.3 Parallelization speed-up

The goal of this experiment is to analyze the parallelization speed-up of BET and compare it with the speed-up of Batch. We set this up as follows: we run to convergence and find the final accuracy, then for each method, we vary the number of cores and measure the time it takes to reach within 0.25% of optimum accuracy. For this experiment we selected the url dataset to demonstrate a stark contrast with Parallel SGD. The method achieves close to linear speed-up only for up to 4 cores, after which its performance flattens. This behavior can be attributed to the fact that as we increase the number of cores, a single iteration of Parallel SGD becomes less and less effective, which in some cases may negate the parallelization speed-up altogether. Batch methods, on the other hand, behave much more reliably with parallelization, because their iterations produce the same effect regardless of the number of cores. Figure 4 shows the speed-up factors for BET, Batch and Parallel SGD. BET achieves similar speed-up as Batch, due to the fact that parallelization happens in the inner optimizer, which is the same for both methods.

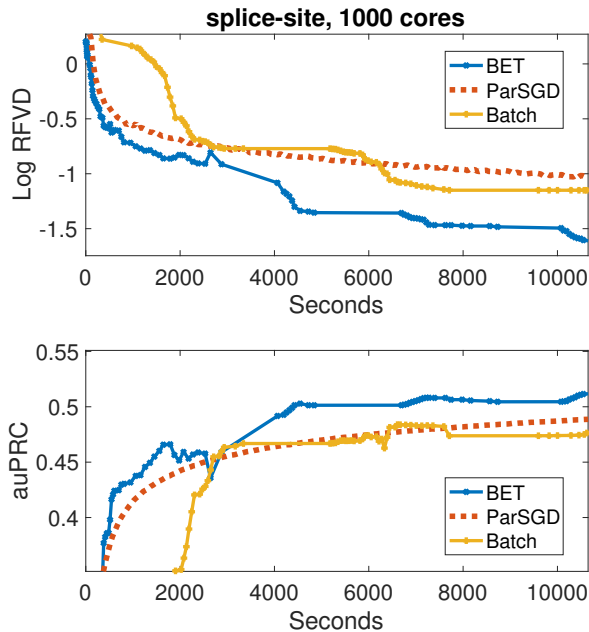


Figure 5: Distributed experiment comparing BET to Batch (L-BFGS) and Parallel SGD on splice-site dataset.

5.4 Distributed optimization

Batch optimization shows the most benefits when dealing with large-scale datasets, which do not fit into the memory of a single machine. BET easily scales up to this setting as seen in Figure 5, running on the splice-site data, with 20 machines, and 50 cores per machine, compared against Batch and Parallel SGD in the same setup. Since the dataset is highly skewed, the results are shown in terms of area under Precision-Recall curve (auPRC).

Note that BET reaches close to optimum test auPRC much more quickly than Batch or Parallel SGD. Moreover, until that point batch-expansion has not reached the full training data, which means that if the data were loaded in parallel with the optimization, the optimization time would be partially absorbed into the loading time. In our experiments, each one of the 1000 processes independently reads a separate chunk of data, however this procedure still takes at least 1500 seconds even on a high-performance cluster. Thus, a portion of BET's optimization time can be absorbed into data loading.

6 Conclusions

We proposed Batch-Expansion Training, a simple parameter-free method of accelerating batch optimizers in a way that is both theoretically and experimentally efficient. BET does not require tuning and can be paired with any optimizer, offering advantages in parallel and distributed settings. Extending our framework to non-convex optimization is an interesting direction for future research.

References

- [1] Alekh Agarwal, Olivier Chapelle, Miroslav Dudík, and John Langford. A reliable effective terascale linear learning system. *J. Mach. Learn. Res.*, 15(1):1111–1133, January 2014.
- [2] Satish Balay, Shrirang Abhyankar, Mark F. Adams, Jed Brown, Peter Brune, Kris Buschelman, Lisandro Dalcin, Victor Eijkhout, William D. Gropp, Dinesh Kaushik, Matthew G. Knepley, Lois Curfman McInnes, Karl Rupp, Barry F. Smith, Stefano Zampini, Hong Zhang, and Hong Zhang. PETSc Web page. <http://www.mcs.anl.gov/petsc>, 2016.
- [3] Satish Balay, Shrirang Abhyankar, Mark F. Adams, Jed Brown, Peter Brune, Kris Buschelman, Lisandro Dalcin, Victor Eijkhout, William D. Gropp, Dinesh Kaushik, Matthew G. Knepley, Lois Curfman McInnes, Karl Rupp, Barry F. Smith, Stefano Zampini, Hong Zhang, and Hong Zhang. PETSc users manual. Technical Report ANL-95/11 - Revision 3.7, Argonne National Laboratory, 2016.
- [4] Satish Balay, William D. Gropp, Lois Curfman McInnes, and Barry F. Smith. Efficient management of parallelism in object oriented numerical software libraries. In E. Arge, A. M. Bruaset, and H. P. Langtangen, editors, *Modern Software Tools in Scientific Computing*, pages 163–202. Birkhäuser Press, 1997.
- [5] Leon Bottou and Olivier Bousquet. The tradeoffs of large scale learning. In *Proceedings of the 20th International Conference on Neural Information Processing Systems*, NIPS’07, pages 161–168, USA, 2007. Curran Associates Inc.
- [6] Sébastien Bubeck. Convex optimization: Algorithms and complexity. *Found. Trends Mach. Learn.*, 8(3-4):231–357, November 2015.
- [7] Richard H. Byrd, Gillian M. Chin, Jorge Nocedal, and Yuchen Wu. Sample size selection in optimization methods for machine learning. *Math. Program.*, 134(1):127–155, August 2012.
- [8] Richard H. Byrd, Peihuang Lu, Jorge Nocedal, and Ciyou Zhu. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, 16(5):1190–1208, 1995.
- [9] Hadi Daneshmand, Aurelien Lucchi, and Thomas Hofmann. Starting small - learning with adaptive sample sizes. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1463–1471, New York, New York, USA, 20–22 Jun 2016. PMLR.
- [10] Aaron Defazio, Francis R. Bach, and Simon Lacoste-Julien. SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 1646–1654, 2014.
- [11] Ofer Dekel, Ran Gilad-Bachrach, Ohad Shamir, and Lin Xiao. Optimal distributed online prediction using mini-batches. *J. Mach. Learn. Res.*, 13(1):165–202, January 2012.
- [12] Murat A. Erdogdu and Andrea Montanari. Convergence rates of sub-sampled newton methods. In *Proceedings of the 28th International Conference on Neural Information Processing Systems*, NIPS’15, pages 3052–3060, Cambridge, MA, USA, 2015. MIT Press.
- [13] R. Fletcher and C. M. Reeves. Function minimization by conjugate gradients. *The Computer Journal*, 7(2):149–154, February 1964.
- [14] Michael P. Friedlander and Mark Schmidt. Hybrid deterministic-stochastic methods for data fitting. *SIAM Journal on Scientific Computing* 34(3), 2012.
- [15] Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In Christopher J. C. Burges, Léon Bottou, Zoubin Ghahramani, and Kilian Q. Weinberger, editors, *NIPS*, pages 315–323, 2013.
- [16] Mu Li, Tong Zhang, Yuqiang Chen, and Alexander J. Smola. Efficient mini-batch training for stochastic optimization. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD ’14, pages 661–670, New York, NY, USA, 2014. ACM.
- [17] Chih-Jen Lin, Ruby C. Weng, and S. Sathiya Keerthi. Trust region newton methods for large-scale logistic regression. In *Proceedings of the 24th International Conference on Machine Learning*, ICML ’07, pages 561–568, New York, NY, USA, 2007. ACM.
- [18] Shin Matsushima, S.V.N. Vishwanathan, and Alexander J. Smola. Linear support vector machines via dual cached loops. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD ’12, pages 177–185, New York, NY, USA, 2012. ACM.
- [19] Aryan Mokhtari, Hadi Daneshmand, Aurélien Lucchi, Thomas Hofmann, and Alejandro Ribeiro. Adaptive newton method for empirical risk minimization to statistical accuracy. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 4062–4070, 2016.

- [20] Shravan Narayanamurthy, Markus Weimer, Dhruv Mahajan, Tyson Condie, Sundararajan Sellamanickam, and Keerthi Selvaraj. Towards resource-elastic machine learning. NIPS 2013 BigLearn Workshop, 2013.
- [21] Jorge Nocedal and Steve J. Wright. *Numerical optimization*. Springer Series in Operations Research and Financial Engineering. Springer, Berlin, 2006. NEOS guide <http://www-fp.mcs.anl.gov/otc/Guide/>.
- [22] Mark Schmidt, Nicolas Le Roux, and Francis Bach. Minimizing finite sums with the stochastic average gradient. *Math. Program.*, 162(1-2):83–112, March 2017.
- [23] Shai Shalev-Shwartz and Nathan Srebro. Svm optimization: Inverse dependence on training set size. In *Proceedings of the 25th International Conference on Machine Learning, ICML '08*, pages 928–935, New York, NY, USA, 2008. ACM.
- [24] Shai Shalev-Shwartz and Tong Zhang. Stochastic dual coordinate ascent methods for regularized loss. *J. Mach. Learn. Res.*, 14(1):567–599, February 2013.
- [25] Karthik Sridharan, Shai Shalev-shwartz, and Nathan Srebro. Fast rates for regularized objectives. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, editors, *Advances in Neural Information Processing Systems 21*, pages 1545–1552. Curran Associates, Inc., 2009.
- [26] Ilya Tolstikhin, Nikita Zhitovskiy, and Gilles Blanchard. Permutational rademacher complexity. In *Proceedings of the 26th International Conference on Algorithmic Learning Theory - Volume 9355*, pages 209–223, New York, NY, USA, 2015. Springer-Verlag New York, Inc.
- [27] Martin Zinkevich. Theoretical analysis of a warm start technique. NIPS 2011 BigLearn Workshop, 2011.
- [28] Martin Zinkevich, Markus Weimer, Alexander J. Smola, and Lihong Li. Parallelized Stochastic Gradient Descent. In John D. Lafferty, Christopher K. I. Williams, John S. Taylor, Richard S. Zemel, and Aron Culotta, editors, *NIPS*, pages 2595–2603. Curran Associates, Inc., 2010.