
A Multi-Batch L-BFGS Method for Machine Learning

Albert S. Berahas

Northwestern University
Evanston, IL

albertberahas@u.northwestern.edu

Jorge Nocedal

Northwestern University
Evanston, IL

j-nocedal@northwestern.edu

Martin Takáč

Lehigh University
Bethlehem, PA

takac.mt@gmail.com

Abstract

The question of how to parallelize the stochastic gradient descent (SGD) method has received much attention in the literature. In this paper, we focus instead on batch methods that use a sizeable fraction of the training set at each iteration to facilitate parallelism, and that employ second-order information. In order to improve the learning process, we follow a *multi-batch* approach in which the batch changes at each iteration. This can cause difficulties because L-BFGS employs gradient differences to update the Hessian approximations, and when these gradients are computed using different data points the process can be unstable. This paper shows how to perform stable quasi-Newton updating in the multi-batch setting, illustrates the behavior of the algorithm in a distributed computing platform, and studies its convergence properties for both the convex and nonconvex cases.

1 Introduction

It is common in machine learning to encounter optimization problems involving millions of parameters and very large datasets. To deal with the computational demands imposed by such applications, high performance implementations of stochastic gradient and batch quasi-Newton methods have been developed [1, 11, 9]. In this paper we study a batch approach based on the L-BFGS method [20] that strives to reach the right balance between efficient learning and productive parallelism.

In supervised learning, one seeks to minimize empirical risk,

$$F(w) := \frac{1}{n} \sum_{i=1}^n f(w; x^i, y^i) \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n f_i(w),$$

where $(x^i, y^i)_{i=1}^n$ denote the training examples and $f(\cdot; x, y) : \mathbb{R}^d \rightarrow \mathbb{R}$ is the composition of a prediction function (parametrized by w) and a loss function. The training problem consists of finding an optimal choice of the parameters $w \in \mathbb{R}^d$ with respect to F , i.e.,

$$\min_{w \in \mathbb{R}^d} F(w) = \frac{1}{n} \sum_{i=1}^n f_i(w). \quad (1.1)$$

At present, the preferred optimization method is the stochastic gradient descent (SGD) method [23, 5], and its variants [14, 24, 12], which are implemented either in an asynchronous manner (e.g. when

using a parameter server in a distributed setting) or following a synchronous mini-batch approach that exploits parallelism in the gradient evaluation [2, 22, 13]. A drawback of the asynchronous approach is that it cannot use large batches, as this would cause updates to become too dense and compromise the stability and scalability of the method [16, 22]. As a result, the algorithm spends more time in communication as compared to computation. On the other hand, using a synchronous mini-batch approach one can achieve a near-linear decrease in the number of SGD iterations as the mini-batch size is increased, up to a certain point after which the increase in computation is not offset by the faster convergence [26].

An alternative to SGD is a batch method, such as L-BFGS, which is able to reach high training accuracy and allows one to perform more computation per node, so as to achieve a better balance with communication costs [27]. Batch methods are, however, not as efficient learning algorithms as SGD in a sequential setting [6]. To benefit from the strength of both methods some high performance systems employ SGD at the start and later switch to a batch method [1].

Multi-Batch Method. In this paper, we follow a different approach consisting of a single method that selects a *sizeable* subset (batch) of the training data to compute a step, and changes this batch at each iteration to improve the learning abilities of the method. We call this a *multi-batch* approach to differentiate it from the mini-batch approach used in conjunction with SGD, which employs a very small subset of the training data. When using large batches it is natural to employ a quasi-Newton method, as incorporating second-order information imposes little computational overhead and improves the stability and speed of the method. We focus here on the L-BFGS method, which employs gradient information to update an estimate of the Hessian and computes a step in $O(d)$ flops, where d is the number of variables. The multi-batch approach can, however, cause difficulties to L-BFGS because this method employs gradient differences to update Hessian approximations. When the gradients used in these differences are based on different data points, the updating procedure can be unstable. Similar difficulties arise in a parallel implementation of the standard L-BFGS method, if some of the computational nodes devoted to the evaluation of the function and gradient are unable to return results on time — as this again amounts to using different data points to evaluate the function and gradient at the beginning and the end of the iteration. The goal of this paper is to show that stable quasi-Newton updating can be achieved in both settings without incurring extra computational cost, or special synchronization. The key is to perform quasi-Newton updating based on the overlap between consecutive batches. The only restriction is that this overlap should not be too small, something that can be achieved in most situations.

Contributions. We describe a novel implementation of the batch L-BFGS method that is robust in the absence of sample consistency; i.e., when different samples are used to evaluate the objective function and its gradient at consecutive iterations. The numerical experiments show that the method proposed in this paper — which we call the *multi-batch* L-BFGS method — achieves a good balance between computation and communication costs. We also analyze the convergence properties of the new method (using a fixed step length strategy) on both convex and nonconvex problems.

2 The Multi-Batch Quasi-Newton Method

In a pure batch approach, one applies a gradient based method, such as L-BFGS [20], to the deterministic optimization problem (1.1). When the number n of training examples is large, it is natural to parallelize the evaluation of F and ∇F by assigning the computation of the component functions f_i to different processors. If this is done on a distributed platform, it is possible for some of the computational nodes to be slower than the rest. In this case, the contribution of the slow (or unresponsive) computational nodes could be ignored given the stochastic nature of the objective function. This leads, however, to an inconsistency in the objective function and gradient at the beginning and at the end of the iteration, which can be detrimental to quasi-Newton methods. Thus, we seek to find a *fault-tolerant* variant of the batch L-BFGS method that is capable of dealing with slow or unresponsive computational nodes.

A similar challenge arises in a *multi-batch* implementation of the L-BFGS method in which the entire training set $T = \{(x^i, y^i)_{i=1}^n\}$ is not employed at every iteration, but rather, a subset of the data is used to compute the gradient. Specifically, we consider a method in which the dataset is randomly divided into a number of batches — say 10, 50, or 100 — and the minimization is performed with respect to a different batch at every iteration. At the k -th iteration, the algorithm chooses a batch

$S_k \subset \{1, \dots, n\}$, computes

$$F^{S_k}(w_k) = \frac{1}{|S_k|} \sum_{i \in S_k} f_i(w_k), \quad \nabla F^{S_k}(w_k) = g_k^{S_k} = \frac{1}{|S_k|} \sum_{i \in S_k} \nabla f_i(w_k), \quad (2.2)$$

and takes a step along the direction $-H_k g_k^{S_k}$, where H_k is an approximation to $\nabla^2 F(w_k)^{-1}$. Allowing the sample S_k to change freely at every iteration gives this approach flexibility of implementation and is beneficial to the learning process, as we show in Section 4. (We refer to S_k as the sample of training points, even though S_k only indexes those points.)

The case of unresponsive computational nodes and the multi-batch method are similar. The main difference is that node failures create unpredictable changes to the samples S_k , whereas a multi-batch method has control over sample generation. In either case, the algorithm employs a stochastic approximation to the gradient and can no longer be considered deterministic. We must, however, distinguish our setting from that of the classical SGD method, which employs small mini-batches and noisy gradient approximations. Our algorithm operates with much larger batches so that distributing the function evaluation is beneficial and the compute time of $g_k^{S_k}$ is not overwhelmed by communication costs. This gives rise to gradients with relatively small variance and justifies the use of a second-order method such as L-BFGS.

Robust Quasi-Newton Updating. The difficulties created by the use of a different sample S_k at each iteration can be circumvented if consecutive samples S_k and S_{k+1} overlap, so that $O_k = S_k \cap S_{k+1} \neq \emptyset$. One can then perform stable quasi-Newton updating by computing gradient differences based on this overlap, i.e., by defining

$$y_{k+1} = g_{k+1}^{O_k} - g_k^{O_k}, \quad s_{k+1} = w_{k+1} - w_k, \quad (2.3)$$

in the notation given in (2.2). The correction pair (y_k, s_k) can then be used in the BFGS update. When the overlap set O_k is not too small, y_k is a useful approximation of the curvature of the objective function F along the most recent displacement, and will lead to a productive quasi-Newton step. This observation is based on an important property of Newton-like methods, namely that there is much more freedom in choosing a Hessian approximation than in computing the gradient [7, 3]. Thus, a smaller sample O_k can be employed for updating the inverse Hessian approximation H_k than for computing the batch gradient $g_k^{S_k}$ in the search direction $-H_k g_k^{S_k}$. In summary, by ensuring that unresponsive nodes do not constitute the vast majority of all working nodes in a fault-tolerant parallel implementation, or by exerting a small degree of control over the creation of the samples S_k in the multi-batch method, one can design a robust method that naturally builds upon the fundamental properties of BFGS updating.

We should mention in passing that a commonly used strategy for ensuring stability of quasi-Newton updating in machine learning is to enforce gradient consistency [25], i.e., to use the same sample S_k to compute gradient evaluations at the beginning and the end of the iteration. Another popular remedy is to use the same batch S_k for multiple iterations [19], alleviating the gradient inconsistency problem at the price of slower convergence. In this paper, we assume that achieving such *sample consistency is not possible* (in the fault-tolerant case) or *desirable* (in a multi-batch framework), and wish to design a new variant of L-BFGS that imposes minimal restrictions in the sample changes.

2.1 Specification of the Method

At the k -th iteration, the multi-batch BFGS algorithm chooses a set $S_k \subset \{1, \dots, n\}$ and computes a new iterate

$$w_{k+1} = w_k - \alpha_k H_k g_k^{S_k}, \quad (2.4)$$

where α_k is the step length, $g_k^{S_k}$ is the batch gradient (2.2) and H_k is the inverse BFGS Hessian matrix approximation that is updated at every iteration by means of the formula

$$H_{k+1} = V_k^T H_k V_k + \rho_k s_k s_k^T, \quad \rho_k = \frac{1}{y_k^T s_k}, \quad V_k = I - \rho_k y_k s_k^T.$$

To compute the correction vectors (s_k, y_k) , we determine the overlap set $O_k = S_k \cap S_{k+1}$ consisting of the samples that are common at the k -th and $k+1$ -st iterations. We define

$$F^{O_k}(w_k) = \frac{1}{|O_k|} \sum_{i \in O_k} f_i(w_k), \quad \nabla F^{O_k}(w_k) = g_k^{O_k} = \frac{1}{|O_k|} \sum_{i \in O_k} \nabla f_i(w_k),$$

and compute the correction vectors as in (2.3). In this paper we assume that α_k is constant.

In the limited memory version, the matrix H_k is defined at each iteration as the result of applying m BFGS updates to a multiple of the identity matrix, using a set of m correction pairs $\{s_i, y_i\}$ kept in storage. The memory parameter m is typically in the range 2 to 20. When computing the matrix-vector product in (2.4) it is not necessary to form that matrix H_k since one can obtain this product via the two-loop recursion [20], using the m most recent correction pairs $\{s_i, y_i\}$. After the step has been computed, the oldest pair (s_j, y_j) is discarded and the new curvature pair is stored.

A pseudo-code of the proposed method is given below, and depends on several parameters. The parameter r denotes the fraction of samples in the dataset used to define the gradient, i.e., $r = \frac{|S|}{n}$. The parameter o denotes the length of overlap between consecutive samples, and is defined as a fraction of the number of samples in a given batch S , i.e., $o = \frac{|O|}{|S|}$.

Algorithm 1 Multi-Batch L-BFGS

Input: w_0 (initial iterate), $T = \{(x^i, y^i), \text{ for } i = 1, \dots, n\}$ (training set), m (memory parameter), r (batch, fraction of n), o (overlap, fraction of batch), $k \leftarrow 0$ (iteration counter).

- 1: Create initial batch S_0 ▷ As shown in Firgure 1
 - 2: **for** $k = 0, 1, 2, \dots$ **do**
 - 3: Calculate the search direction $p_k = -H_k g_k^{S_k}$ ▷ Using L-BFGS formula
 - 4: Choose the step length $\alpha_k > 0$
 - 5: Compute $w_{k+1} = w_k + \alpha_k p_k$
 - 6: Create the next batch S_{k+1}
 - 7: Compute the curvature pairs $s_{k+1} = w_{k+1} - w_k$ and $y_{k+1} = g_{k+1}^{O_k} - g_k^{O_k}$
 - 8: Replace the oldest pair (s_i, y_i) by s_{k+1}, y_{k+1}
 - 9: **end for**
-

2.2 Sample Generation

We now discuss how the sample S_{k+1} is created at each iteration (Line 8 in Algorithm 1).

Distributed Computing with Faults. Consider a distributed implementation in which slave nodes read the current iterate w_k from the master node, compute a local gradient on a subset of the dataset, and send it back to the master node for aggregation in the calculation (2.2). Given a time (computational) budget, it is possible for some nodes to fail to return a result. The schematic in Figure 1a illustrates the gradient calculation across two iterations, k and $k+1$, in the presence of faults. Here $\mathcal{B}_i, i = 1, \dots, B$ denote the batches of data that each slave node i receives (where $T = \cup_i \mathcal{B}_i$), and $\tilde{\nabla} f(w)$ is the gradient calculation using all nodes that responded within the preallocated time.




Figure 1: Sample and Overlap formation.

Let $\mathcal{J}_k \subset \{1, 2, \dots, B\}$ and $\mathcal{J}_{k+1} \subset \{1, 2, \dots, B\}$ be the set of indices of all nodes that returned a gradient at the k -th and $k+1$ -st iterations, respectively. Using this notation $S_k = \cup_{j \in \mathcal{J}_k} \mathcal{B}_j$ and $S_{k+1} = \cup_{j \in \mathcal{J}_{k+1}} \mathcal{B}_j$, and we define $O_k = \cup_{j \in \mathcal{J}_k \cap \mathcal{J}_{k+1}} \mathcal{B}_j$. The simplest implementation in this setting preallocates the data on each compute node, requiring minimal data communication, i.e., only

one data transfer. In this case the samples S_k will be independent if node failures occur randomly. On the other hand, if the same set of nodes fail, then sample creation will be biased, which is harmful both in theory and practice. One way to ensure independent sampling is to shuffle and redistribute the data to all nodes after a certain number of iterations.

Multi-batch Sampling. We propose two strategies for the *multi-batch* setting.

Figure 1b illustrates the sample creation process in the first strategy. The dataset is shuffled and batches are generated by collecting subsets of the training set, in order. Every set (except S_0) is of the form $S_k = \{O_{k-1}, N_k, O_k\}$, where O_{k-1} and O_k are the overlapping samples with batches S_{k-1} and S_{k+1} respectively, and N_k are the samples that are unique to batch S_k . After each pass through the dataset, the samples are reshuffled, and the procedure described above is repeated. In our implementation samples are drawn without replacement, guaranteeing that after every pass (epoch) all samples are used. This strategy has the advantage that it requires no extra computation in the evaluation of $g_k^{O_k}$ and $g_{k+1}^{O_{k+1}}$, but the samples $\{S_k\}$ are not independent.

The second sampling strategy is simpler and requires less control. At every iteration k , a batch S_k is created by randomly selecting $|S_k|$ elements from $\{1, \dots, n\}$. The overlapping set O_k is then formed by randomly selecting $|O_k|$ elements from S_k (subsampling). This strategy is slightly more expensive since $g_{k+1}^{O_{k+1}}$ requires extra computation, but if the overlap is small this cost is not significant.

3 Convergence Analysis

In this section, we analyze the convergence properties of the multi-batch L-BFGS method (Algorithm 1) when applied to the minimization of *strongly convex* and *nonconvex* objective functions, using a fixed step length strategy. We assume that the goal is to minimize the empirical risk F given in (1.1), but note that a similar analysis could be used to study the minimization of the expected risk.

3.1 Strongly Convex case

Due to the stochastic nature of the multi-batch approach, every iteration of Algorithm 1 employs a gradient that contains errors that do not converge to zero. Therefore, by using a fixed step length strategy one cannot establish convergence to the optimal solution w^* , but only convergence to a neighborhood of w^* [18]. Nevertheless, this result is of interest as it reflects the common practice of using a fixed step length and decreasing it only if the desired testing error has not been achieved. It also illustrates the tradeoffs that arise between the size of the batch and the step length.

In our analysis, we make the following assumptions about the objective function and the algorithm.

Assumptions A.

1. F is twice continuously differentiable.
2. There exist positive constants $\hat{\lambda}$ and $\hat{\Lambda}$ such that $\hat{\lambda}I \preceq \nabla^2 F^O(w) \preceq \hat{\Lambda}I$ for all $w \in \mathbb{R}^d$ and all sets $O \subset \{1, 2, \dots, n\}$.
3. There is a constant γ such that $\mathbb{E}_S [\|\nabla F^S(w)\|]^2 \leq \gamma^2$ for all $w \in \mathbb{R}^d$ and all sets $S \subset \{1, 2, \dots, n\}$.
4. The samples S are drawn independently and $\nabla F^S(w)$ is an unbiased estimator of the true gradient $\nabla F(w)$ for all $w \in \mathbb{R}^d$, i.e., $\mathbb{E}_S[\nabla F^S(w)] = \nabla F(w)$.

Note that Assumption A.2 implies that the entire Hessian $\nabla^2 F(w)$ also satisfies

$$\lambda I \preceq \nabla^2 F(w) \preceq \Lambda I, \quad \forall w \in \mathbb{R}^d,$$

for some constants $\lambda, \Lambda > 0$. Assuming that every sub-sampled function $F^O(w)$ is strongly convex is not unreasonable as a regularization term is commonly added in practice when that is not the case.

We begin by showing that the inverse Hessian approximations H_k generated by the multi-batch L-BFGS method have eigenvalues that are uniformly bounded above and away from zero. The proof technique used is an adaptation of that in [8].

Lemma 3.1. *If Assumptions A.1-A.2 above hold, there exist constants $0 < \mu_1 \leq \mu_2$ such that the Hessian approximations $\{H_k\}$ generated by Algorithm 1 satisfy*

$$\mu_1 I \preceq H_k \preceq \mu_2 I, \quad \text{for } k = 0, 1, 2, \dots$$

Utilizing Lemma 3.1, we show that the multi-batch L-BFGS method with a constant step length converges to a neighborhood of the optimal solution.

Theorem 3.2. *Suppose that Assumptions A.1-A.4 hold and let $F^* = F(w^*)$, where w^* is the minimizer of F . Let $\{w_k\}$ be the iterates generated by Algorithm 1 with $\alpha_k = \alpha \in (0, \frac{1}{2\mu_1\lambda})$, starting from w_0 . Then for all $k \geq 0$,*

$$\mathbb{E}[F(w_k) - F^*] \leq (1 - 2\alpha\mu_1\lambda)^k [F(w_0) - F^*] + [1 - (1 - \alpha\mu_1\lambda)^k] \frac{\alpha\mu_2^2\gamma^2\Lambda}{4\mu_1\lambda}$$

$$\xrightarrow{k \rightarrow \infty} \frac{\alpha\mu_2^2\gamma^2\Lambda}{4\mu_1\lambda}.$$

The bound provided by this theorem has two components: (i) a term decaying linearly to zero, and (ii) a term identifying the neighborhood of convergence. Note that a larger step length yields a more favorable constant in the linearly decaying term, at the cost of an increase in the size of the neighborhood of convergence. We will consider again these tradeoffs in Section 4, where we also note that larger batches increase the opportunities for parallelism and improve the limiting accuracy in the solution, but slow down the learning abilities of the algorithm.

One can establish convergence of the multi-batch L-BFGS method to the optimal solution w^* by employing a sequence of step lengths $\{\alpha_k\}$ that converge to zero according to the schedule proposed by Robbins and Monro [23]. However, that provides only a sublinear rate of convergence, which is of little interest in our context where large batches are employed and some type of linear convergence is expected. In this light, Theorem 3.2 is more relevant to practice.

3.2 Nonconvex case

The BFGS method is known to fail on nonconvex problems [17, 10]. Even for L-BFGS, which makes only a finite number of updates at each iteration, one cannot guarantee that the Hessian approximations have eigenvalues that are uniformly bounded above and away from zero. To establish convergence of the BFGS method in the nonconvex case *cautious* updating procedures have been proposed [15]. Here we employ a cautious strategy that is well suited to our particular algorithm; we skip the update, i.e., set $H_{k+1} = H_k$, if the curvature condition

$$y_k^T s_k \geq \epsilon \|s_k\|^2 \quad (3.5)$$

is not satisfied, where $\epsilon > 0$ is a predetermined constant. Using said mechanism we show that the eigenvalues of the Hessian matrix approximations generated by the multi-batch L-BFGS method are bounded above and away from zero (Lemma 3.3). The analysis presented in this section is based on the following assumptions.

Assumptions B.

1. F is twice continuously differentiable.
2. The gradients of F are Λ -Lipschitz continuous, and the gradients of F^O are Λ_O -Lipschitz continuous for all $w \in \mathbb{R}^d$ and all sets $O \subset \{1, 2, \dots, n\}$.
3. The function $F(w)$ is bounded below by a scalar \hat{F} .
4. There exist constants $\gamma \geq 0$ and $\eta > 0$ such that $\mathbb{E}_S [\|\nabla F^S(w)\|]^2 \leq \gamma^2 + \eta \|\nabla F(w)\|^2$ for all $w \in \mathbb{R}^d$ and all sets $S \subset \{1, 2, \dots, n\}$.
5. The samples S are drawn independently and $\nabla F^S(w)$ is an unbiased estimator of the true gradient $\nabla F(w)$ for all $w \in \mathbb{R}^d$, i.e., $\mathbb{E}[\nabla F^S(w)] = \nabla F(w)$.

Lemma 3.3. *Suppose that Assumptions B.1-B.2 hold and let $\epsilon > 0$ be given. Let $\{H_k\}$ be the Hessian approximations generated by Algorithm 1, with the modification that $H_{k+1} = H_k$ whenever (3.5) is not satisfied. Then, there exist constants $0 < \mu_1 \leq \mu_2$ such that*

$$\mu_1 I \preceq H_k \preceq \mu_2 I, \quad \text{for } k = 0, 1, 2, \dots$$

We can now follow the analysis in [4, Chapter 4] to establish the following result about the behavior of the gradient norm for the multi-batch L-BFGS method with a cautious update strategy.

Theorem 3.4. *Suppose that Assumptions B.1-B.5 above hold, and let $\epsilon > 0$ be given. Let $\{w_k\}$ be the iterates generated by Algorithm 1, with $\alpha_k = \alpha \in (0, \frac{\mu_1}{\mu_2\eta\Lambda})$, starting from w_0 , and with the*

modification that $H_{k+1} = H_k$ whenever (3.5) is not satisfied. Then,

$$\mathbb{E} \left[\frac{1}{L} \sum_{k=0}^{L-1} \|\nabla F(w_k)\|^2 \right] \leq \frac{\alpha \mu_2^2 \gamma^2 \Lambda}{\mu_1} + \frac{2[F(w_0) - \hat{F}]}{\alpha \mu_1 L}$$

$$\xrightarrow{L \rightarrow \infty} \frac{\alpha \mu_2^2 \gamma^2 \Lambda}{\mu_1}.$$

This result bounds the average norm of the gradient of F after the first $L - 1$ iterations, and shows that the iterates spend increasingly more time in regions where the objective function has a small gradient.

4 Numerical Results

In this Section, we present numerical results that evaluate the proposed robust multi-batch L-BFGS scheme (Algorithm 1) on logistic regression problems. Figure 2 shows the performance on the webspam dataset¹, where we compare it against three methods: (i) multi-batch L-BFGS without enforcing sample consistency (L-BFGS), where gradient differences are computed using different samples, i.e., $y_k = g_{k+1}^{S_{k+1}} - g_k^{S_k}$; (ii) multi-batch gradient descent (Gradient Descent), which is obtained by setting $H_k = I$ in Algorithm 1; and, (iii) serial SGD, where at every iteration one sample is used to compute the gradient. We run each method with 10 different random seeds, and, where applicable, report results for different batch (r) and overlap (o) sizes. The proposed method is more stable than the standard L-BFGS method; this is especially noticeable when r is small. On the other hand, serial SGD achieves similar accuracy as the robust L-BFGS method and at a similar rate (e.g., $r = 1\%$), at the cost of n communications per epochs versus $\frac{1}{r(1-o)}$ communications per epoch. Figure 2 also indicates that the robust L-BFGS method is not too sensitive to the size of overlap. Similar behavior was observed on other datasets, in regimes where $r \cdot o$ was not too small; see Appendix B.1. We mention in passing that the L-BFGS step was computed using the a vector-free implementation proposed in [9].

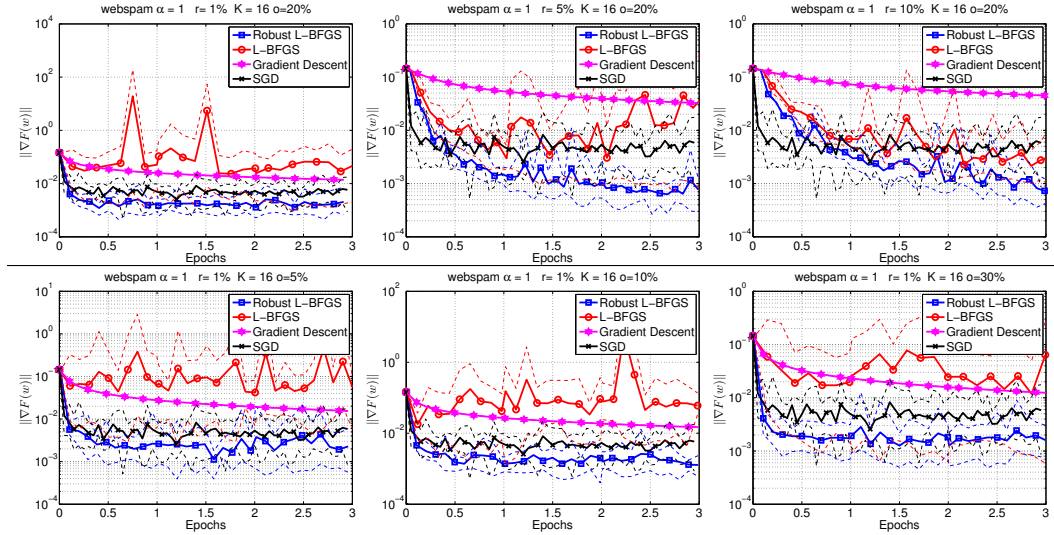


Figure 2: **webspam dataset.** Comparison of Robust L-BFGS, L-BFGS (multi-batch L-BFGS without enforcing sample consistency), Gradient Descent (multi-batch Gradient method) and SGD for various batch (r) and overlap (o) sizes. Solid lines show average performance, and dashed lines show worst and best performance, over 10 runs (per algorithm). $K = 16$ MPI processes.

We also explore the performance of the robust multi-batch L-BFGS method in the presence of node failures (faults), and compare it to the multi-batch variant that does not enforce sample consistency (L-BFGS). Figure 3 illustrates the performance of the methods on the webspam dataset, for various

¹LIBSVM: <https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/binary.html>.

probabilities of node failures $p \in \{0.1, 0.3, 0.5\}$, and suggests that the robust L-BFGS variant is more stable; see Appendix B.2 for further results.




Figure 3: **webspam dataset**. Comparison of Robust L-BFGS and L-BFGS (multi-batch L-BFGS without enforcing sample consistency), for various node failure probabilities p . Solid lines show average performance, and dashed lines show worst and best performance, over 10 runs (per algorithm). $K = 16$ MPI processes.

Lastly, we study the strong and weak scaling properties of the robust L-BFGS method on artificial data (Figure 4). We measure the time needed to compute a gradient (Gradient) and the associated communication (Gradient+C), as well as, the time needed to compute the L-BFGS direction (L-BFGS) and the associated communication (L-BFGS+C), for various batch sizes (r). The figure on the left shows strong scaling of multi-batch L-BFGS on a $d = 10^4$ dimensional problem with $n = 10^7$ samples. The size of input data is 24GB, and we vary the number of MPI processes, $K \in \{1, 2, \dots, 128\}$. The time it takes to compute the gradient decreases with K , however, for small values of r , the communication time exceeds the compute time. The figure on the right shows weak scaling on a problem of similar size, but with varying sparsity. Each sample has $10 \cdot K$ non-zero elements, thus for any K the size of local problem is roughly 1.5GB (for $K = 128$ size of data 192GB). We observe almost constant time for the gradient computation while the cost of computing the L-BFGS direction decreases with K ; however, if communication is considered, the overall time needed to compute the L-BFGS direction increases slightly. For more details see Appendix C.

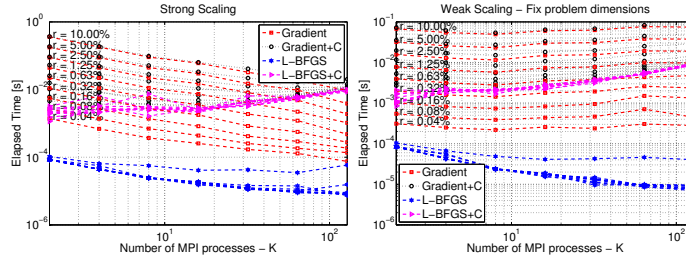


Figure 4: Strong and weak scaling of multi-batch L-BFGS method.

5 Conclusion

This paper describes a novel variant of the L-BFGS method that is robust and efficient in two settings. The first occurs in the presence of node failures in a distributed computing implementation; the second arises when one wishes to employ a different batch at each iteration in order to accelerate learning. The proposed method avoids the pitfalls of using inconsistent gradient differences by performing quasi-Newton updating based on the overlap between consecutive samples. Numerical results show that the method is efficient in practice, and a convergence analysis illustrates its theoretical properties.

Acknowledgements

The first two authors were supported by the Office of Naval Research award N000141410313, the Department of Energy grant DE-FG02-87ER25047 and the National Science Foundation grant DMS-1620022. Martin Takáč was supported by National Science Foundation grant CCF-1618717.

References

- [1] A. Agarwal, O. Chapelle, M. Dudík, and J. Langford. A reliable effective terascale linear learning system. *The Journal of Machine Learning Research*, 15(1):1111–1133, 2014.
- [2] D. P. Bertsekas and J. N. Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*, volume 23. Prentice hall Englewood Cliffs, NJ, 1989.
- [3] R. Bollapragada, R. Byrd, and J. Nocedal. Exact and inexact subsampled newton methods for optimization. *arXiv preprint arXiv:1609.08502*, 2016.
- [4] L. Bottou, F. E. Curtis, and J. Nocedal. Optimization methods for large-scale machine learning. *arXiv preprint arXiv:1606.04838*, 2016.
- [5] L. Bottou and Y. LeCun. Large scale online learning. In *NIPS*, pages 217–224, 2004.
- [6] O. Bousquet and L. Bottou. The tradeoffs of large scale learning. In *NIPS*, pages 161–168, 2008.
- [7] R. H. Byrd, G. M. Chin, W. Neveitt, and J. Nocedal. On the use of stochastic Hessian information in optimization methods for machine learning. *SIAM Journal on Optimization*, 21(3):977–995, 2011.
- [8] R. H. Byrd, S. L. Hansen, J. Nocedal, and Y. Singer. A stochastic quasi-newton method for large-scale optimization. *SIAM Journal on Optimization*, 26(2):1008–1031, 2016.
- [9] W. Chen, Z. Wang, and J. Zhou. Large-scale L-BFGS using MapReduce. In *NIPS*, pages 1332–1340, 2014.
- [10] Y.-H. Dai. Convergence properties of the BFGS algorithm. *SIAM Journal on Optimization*, 13(3):693–701, 2002.
- [11] J. Dean, G. Corrado, R. Monga, K. Chen, M. Devin, M. Mao, A. Senior, P. Tucker, K. Yang, Q. V. Le, et al. Large scale distributed deep networks. In *NIPS*, pages 1223–1231, 2012.
- [12] A. Defazio, F. Bach, and S. Lacoste-Julien. SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives. In *NIPS*, pages 1646–1654, 2014.
- [13] I. Goodfellow, Y. Bengio, and A. Courville. Deep learning. Book in preparation for MIT Press, 2016.
- [14] R. Johnson and T. Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *NIPS*, pages 315–323, 2013.
- [15] D.-H. Li and M. Fukushima. On the global convergence of the BFGS method for nonconvex unconstrained optimization problems. *SIAM Journal on Optimization*, 11(4):1054–1064, 2001.
- [16] H. Mania, X. Pan, D. Papailiopoulos, B. Recht, K. Ramchandran, and M. I. Jordan. Perturbed iterate analysis for asynchronous stochastic optimization. *arXiv preprint arXiv:1507.06970*, 2015.
- [17] W. F. Mascarenhas. The BFGS method with exact line searches fails for non-convex objective functions. *Mathematical Programming*, 99(1):49–61, 2004.
- [18] A. Nedić and D. Bertsekas. Convergence rate of incremental subgradient algorithms. In *Stochastic Optimization: Algorithms and Applications*, pages 223–264. Springer, 2001.
- [19] J. Ngiam, A. Coates, A. Lahiri, B. Prochnow, Q. V. Le, and A. Y. Ng. On optimization methods for deep learning. In *ICML*, pages 265–272, 2011.
- [20] J. Nocedal and S. Wright. *Numerical Optimization*. Springer New York, 2 edition, 1999.
- [21] M. J. Powell. Some global convergence properties of a variable metric algorithm for minimization without exact line searches. *Nonlinear programming*, 9(1):53–72, 1976.
- [22] B. Recht, C. Re, S. Wright, and F. Niu. HOGWILD!: A lock-free approach to parallelizing stochastic gradient descent. In *NIPS*, pages 693–701, 2011.
- [23] H. Robbins and S. Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, pages 400–407, 1951.
- [24] M. Schmidt, N. Le Roux, and F. Bach. Minimizing finite sums with the stochastic average gradient. *Mathematical Programming*, page 1–30, 2016.
- [25] N. N. Schraudolph, J. Yu, and S. Günter. A stochastic quasi-Newton method for online convex optimization. In *AISTATS*, pages 436–443, 2007.
- [26] M. Takáč, A. Bijral, P. Richtárik, and N. Srebro. Mini-batch primal and dual methods for SVMs. In *ICML*, pages 1022–1030, 2013.
- [27] Y. Zhang and X. Lin. DiSCO: Distributed optimization for self-concordant empirical loss. In *ICML*, pages 362–370, 2015.

A Proofs and Technical Results

A.1 Assumptions

We first restate the Assumptions that we use in the Convergence Analysis section (Section 3). Assumption A and B are used in the strongly convex and nonconvex cases, respectively.

Assumptions A

1. F is twice continuously differentiable.
2. There exist positive constants $\hat{\lambda}$ and $\hat{\Lambda}$ such that

$$\hat{\lambda}I \preceq \nabla^2 F^O(w) \preceq \hat{\Lambda}I, \quad (\text{A.6})$$

for all $w \in \mathbb{R}^d$ and all sets $O \subset \{1, 2, \dots, n\}$.

3. There is a constant γ such that

$$\mathbb{E}_S [\|\nabla F^S(w)\|]^2 \leq \gamma^2, \quad (\text{A.7})$$

for all $w \in \mathbb{R}^d$ and all batches $S \subset \{1, 2, \dots, n\}$.

4. The samples S are drawn independently and $\nabla F^S(w)$ is an unbiased estimator of the true gradient $\nabla F(w)$ for all $w \in \mathbb{R}^d$, i.e.,

$$\mathbb{E} [\nabla F^S(w)] = \nabla F(w). \quad (\text{A.8})$$

Note that Assumption A.2 implies that the entire Hessian $\nabla^2 F(w)$ also satisfies

$$\lambda I \preceq \nabla^2 F(w) \preceq \Lambda I, \forall w \in \mathbb{R}^d, \quad (\text{A.9})$$

for some constants $\lambda, \Lambda > 0$.

Assumptions B

1. F is twice continuously differentiable.
2. The gradients of F are Λ -Lipschitz continuous and the gradients of F^O are Λ_O -Lipschitz continuous for all $w \in \mathbb{R}^d$ and all sets $O \subset \{1, 2, \dots, n\}$.
3. The function $F(w)$ is bounded below by a scalar \hat{F} .
4. There exist constants $\gamma \geq 0$ and $\eta > 0$ such that

$$\mathbb{E}_S [\|\nabla F^S(w)\|]^2 \leq \gamma^2 + \eta \|\nabla F(w)\|^2, \quad (\text{A.10})$$

for all $w \in \mathbb{R}^d$ and all batches $S \subset \{1, 2, \dots, n\}$.

5. The samples S are drawn independently and $\nabla F^S(w)$ is an unbiased estimator of the true gradient $\nabla F(w)$ for all $w \in \mathbb{R}^d$, i.e.,

$$\mathbb{E} [\nabla F^S(w)] = \nabla F(w). \quad (\text{A.11})$$

A.2 Proof of Lemma 3.1 (Strongly Convex Case)

The following Lemma shows that the eigenvalues of the matrices generated by the multi-batch L-BFGS method are bounded above and away from zero if F is strongly convex.

Lemma 3.1. *If Assumptions A.1-A.2 above hold, there exist constants $0 < \mu_1 \leq \mu_2$ such that the Hessian approximations $\{H_k\}$ generated by the multi-batch L-BFGS method (Algorithm 1) satisfy*

$$\mu_1 I \preceq H_k \preceq \mu_2 I, \quad \text{for } k = 0, 1, 2, \dots$$

Proof. Instead of analyzing the inverse Hessian approximation H_k , we study the direct Hessian approximation $B_k = H_k^{-1}$. In this case, the limited memory quasi-Newton updating formula is given as follows

1. Set $B_k^{(0)} = \frac{y_k^T y_k}{s_k^T y_k} I$ and $\tilde{m} = \min\{k, m\}$; where m is the memory in L-BFGS.

2. For $i = 0, \dots, \tilde{m} - 1$ set $j = k - \tilde{m} + 1 + i$ and compute

$$B_k^{(i+1)} = B_k^{(i)} - \frac{B_k^{(i)} s_j s_j^T B_k^{(i)}}{s_j^T B_k^{(i)} s_j} + \frac{y_j y_j^T}{y_j^T s_j}.$$

3. Set $B_{k+1} = B_k^{(\tilde{m})}$.

The curvature pairs s_k and y_k are updated via the following formulae

$$y_{k+1} = g_{k+1}^{O_k} - g_k^{O_k}, \quad s_k = w_{k+1} - w_k. \quad (\text{A.12})$$

A consequence of Assumption A.2 is that the eigenvalues of any sub-sampled Hessian ($|O|$ samples) are bounded above and away from zero. Utilizing this fact, the convexity of component functions and the definitions (A.12), we have

$$y_k^T s_k \geq \frac{1}{\hat{\Lambda}} \|y_k\|^2 \Rightarrow \frac{\|y_k\|^2}{y_k^T s_k} \leq \hat{\Lambda}. \quad (\text{A.13})$$

On the other hand, strong convexity of the sub-sampled functions, the consequence of Assumption A.2 and definitions (A.12), provide a lower bound,

$$y_k^T s_k \leq \frac{1}{\hat{\lambda}} \|y_k\|^2 \Rightarrow \frac{\|y_k\|^2}{y_k^T s_k} \geq \hat{\lambda}. \quad (\text{A.14})$$

Combining the upper and lower bounds (A.13) and (A.14)

$$\hat{\lambda} \leq \frac{\|y_k\|^2}{y_k^T s_k} \leq \hat{\Lambda}. \quad (\text{A.15})$$

The above proves that the eigenvalues of the matrices $B_k^{(0)} = \frac{y_k^T y_k}{s_k^T y_k} I$ at the start of the L-BFGS update cycles are bounded above and away from zero, for all k . We now use a Trace-Determinant argument to show that the eigenvalues of B_k are bounded above and away from zero.

Let $Tr(B)$ and $\det(B)$ denote the trace and determinant of matrix B , respectively, and set $j_i = k - \tilde{m} + i$. The trace of the matrix B_{k+1} can be expressed as,

$$\begin{aligned} Tr(B_{k+1}) &= Tr(B_k^{(0)}) - Tr \sum_{i=1}^{\tilde{m}} \left(\frac{B_k^{(i)} s_{j_i} s_{j_i}^T B_k^{(i)}}{s_{j_i}^T B_k^{(i)} s_{j_i}} \right) + Tr \sum_{i=1}^{\tilde{m}} \frac{y_{j_i} y_{j_i}^T}{y_{j_i}^T s_{j_i}} \\ &\leq Tr(B_k^{(0)}) + \sum_{i=1}^{\tilde{m}} \frac{\|y_{j_i}\|^2}{y_{j_i}^T s_{j_i}} \\ &\leq Tr(B_k^{(0)}) + \tilde{m} \hat{\Lambda} \\ &\leq C_1, \end{aligned} \quad (\text{A.16})$$

for some positive constant C_1 , where the inequalities above are due to (A.15), and the fact that the eigenvalues of the initial L-BFGS matrix $B_k^{(0)}$ are bounded above and away from zero.

Using a result due to Powell [21], the determinant of the matrix B_{k+1} generated by the multi-batch L-BFGS method can be expressed as,

$$\begin{aligned} \det(B_{k+1}) &= \det(B_k^{(0)}) \prod_{i=1}^{\tilde{m}} \frac{y_{j_i}^T s_{j_i}}{s_{j_i}^T B_k^{(i-1)} s_{j_i}} \\ &= \det(B_k^{(0)}) \prod_{i=1}^{\tilde{m}} \frac{y_{j_i}^T s_{j_i}}{s_{j_i}^T s_{j_i}} \frac{s_{j_i}^T s_{j_i}}{s_{j_i}^T B_k^{(i-1)} s_{j_i}} \\ &\geq \det(B_k^{(0)}) \left(\frac{\hat{\lambda}}{C_1} \right)^{\tilde{m}} \\ &\geq C_2, \end{aligned} \quad (\text{A.17})$$

for some positive constant C_2 , where the above inequalities are due to the fact that the largest eigenvalue of $B_k^{(i)}$ is less than C_1 and Assumption A.2.

The trace (A.16) and determinant (A.17) inequalities derived above imply that largest eigenvalues of all matrices B_k are bounded above, uniformly, and that the smallest eigenvalues of all matrices B_k are bounded away from zero, uniformly. \square

A.3 Proof of Theorem 3.2 (Strongly Convex Case)

Utilizing the result from Lemma 3.1, we now prove a linear convergence result to a neighborhood of the optimal solution, for the case where Assumptions A hold.

Theorem 3.2. *Suppose that Assumptions A.1-A.4 above hold, and let $F^* = F(w^*)$, where w^* is the minimizer of F . Let $\{w_k\}$ be the iterates generated by the multi-batch L-BFGS method (Algorithm 1) with*

$$\alpha_k = \alpha \in (0, \frac{1}{2\mu_1\lambda}),$$

starting from w_0 . Then for all $k \geq 0$,

$$\begin{aligned} \mathbb{E}[F(w_k) - F^*] &\leq (1 - 2\alpha\mu_1\lambda)^k [F(w_0) - F^*] + [1 - (1 - \alpha\mu_1\lambda)^k] \frac{\alpha\mu_2^2\gamma^2\Lambda}{4\mu_1\lambda} \\ &\xrightarrow{k \rightarrow \infty} \frac{\alpha\mu_2^2\gamma^2\Lambda}{4\mu_1\lambda}. \end{aligned}$$

Proof. We have that

$$\begin{aligned} F(w_{k+1}) &= F(w_k - \alpha H_k \nabla F^{S_k}(w_k)) \\ &\leq F(w_k) + \nabla F(w_k)^T (-\alpha H_k \nabla F^{S_k}(w_k)) + \frac{\Lambda}{2} \|\alpha H_k \nabla F^{S_k}(w_k)\|^2 \\ &\leq F(w_k) - \alpha \nabla F(w_k)^T H_k \nabla F^{S_k}(w_k) + \frac{\alpha^2 \mu_2^2 \Lambda}{2} \|\nabla F^{S_k}(w_k)\|^2, \end{aligned} \quad (\text{A.18})$$

where the first inequality arises due to (A.9), and the second inequality arises as a consequence of Lemma 3.1.

Taking the expectation (over S_k) of equation (A.18)

$$\begin{aligned} \mathbb{E}_{S_k}[F(w_{k+1})] &\leq F(w_k) - \alpha \nabla F(w_k)^T H_k \nabla F(w_k) + \frac{\alpha^2 \mu_2^2 \Lambda}{2} \mathbb{E}_{S_k} [\|\nabla F^{S_k}(w_k)\|^2] \\ &\leq F(w_k) - \alpha\mu_1 \|\nabla F(w_k)\|^2 + \frac{\alpha^2 \mu_2^2 \gamma^2 \Lambda}{2}, \end{aligned} \quad (\text{A.19})$$

where in the first inequality we make use of Assumption A.5, and the second inequality arises due to Lemma 3.1 and Assumption A.4.

Since F is λ -strongly convex, we can use the following relationship between the norm of the gradient squared, and the distance of the k -th iterate from the optimal solution.

$$2\lambda[F(w_k) - F^*] \leq \|\nabla F(w_k)\|^2.$$

Using the above with (A.19),

$$\begin{aligned} \mathbb{E}_{S_k}[F(w_{k+1})] &\leq F(w_k) - \alpha\mu_1 \|\nabla F(w_k)\|^2 + \frac{\alpha^2 \mu_2^2 \gamma^2 \Lambda}{2} \\ &\leq F(w_k) - 2\alpha\mu_1\lambda[F(w_k) - F^*] + \frac{\alpha^2 \mu_2^2 \gamma^2 \Lambda}{2}. \end{aligned} \quad (\text{A.20})$$

Let

$$\phi_k = \mathbb{E}[F(w_k) - F^*], \quad (\text{A.21})$$

where the expectation is over all batches S_0, S_1, \dots, S_{k-1} and all history starting with w_0 . Thus (A.20) can be expressed as,

$$\phi_{k+1} \leq (1 - 2\alpha\mu_1\lambda)\phi_k + \frac{\alpha^2\mu_2^2\gamma^2\Lambda}{2}, \quad (\text{A.22})$$

from which we deduce that in order to reduce the value with respect to the previous function value, the step length needs to be in the range

$$\alpha \in \left(0, \frac{1}{2\mu_1\lambda}\right).$$

Subtracting $\frac{\alpha\mu_2^2\gamma^2\Lambda}{4\mu_1\lambda}$ from either side of (A.22) yields

$$\begin{aligned} \phi_{k+1} - \frac{\alpha\mu_2^2\gamma^2\Lambda}{4\mu_1\lambda} &\leq (1 - 2\alpha\mu_1\lambda)\phi_k + \frac{\alpha^2\mu_2^2\gamma^2\Lambda}{2} - \frac{\alpha\mu_2^2\gamma^2\Lambda}{4\mu_1\lambda} \\ &= (1 - 2\alpha\mu_1\lambda)\left[\phi_k - \frac{\alpha\mu_2^2\gamma^2\Lambda}{4\mu_1\lambda}\right]. \end{aligned} \quad (\text{A.23})$$

Recursive application of (A.23) yields

$$\phi_k - \frac{\alpha\mu_2^2\gamma^2\Lambda}{4\mu_1\lambda} \leq (1 - 2\alpha\mu_1\lambda)^k \left[\phi_0 - \frac{\alpha\mu_2^2\gamma^2\Lambda}{4\mu_1\lambda}\right],$$

and thus,

$$\phi_k \leq (1 - 2\alpha\mu_1\lambda)^k \phi_0 + \left[1 - (1 - \alpha\mu_1\lambda)^k\right] \frac{\alpha\mu_2^2\gamma^2\Lambda}{4\mu_1\lambda}. \quad (\text{A.24})$$

Finally using the definition of ϕ_k (A.21) with the above expression yields the desired result,

$$\mathbb{E}\left[F(w_k) - F^*\right] \leq \left(1 - 2\alpha\mu_1\lambda\right)^k \left[F(w_0) - F^*\right] + \left[1 - (1 - \alpha\mu_1\lambda)^k\right] \frac{\alpha\mu_2^2\gamma^2\Lambda}{4\mu_1\lambda}. \quad \square$$

A.4 Proof of Lemma 3.3 (Nonconvex Case)

The following Lemma shows that the eigenvalues of the matrices generated by the multi-batch L-BFGS method are bounded above and away from zero (nonconvex case).

Lemma 3.3. *Suppose that Assumptions B.1-B.2 hold and let $\epsilon > 0$ be given. Let $\{H_k\}$ be the Hessian approximations generated by the multi-batch L-BFGS method (Algorithm 1), with the modification that the Hessian approximation H_k update is performed only when*

$$y_k^T s_k \geq \epsilon \|s_k\|^2, \quad (\text{A.25})$$

else $H_{k+1} = H_k$. Then, there exist constants $0 < \mu_1 \leq \mu_2$ such that

$$\mu_1 I \preceq H_k \preceq \mu_2 I, \quad \text{for } k = 0, 1, 2, \dots$$

Proof. Similar to the proof of Lemma 3.1, we study the direct Hessian approximation $B_k = H_k^{-1}$.

The curvature pairs s_k and y_k are updated via the following formulae

$$y_{k+1} = g_{k+1}^{O_k} - g_k^{O_k}, \quad s_k = w_{k+1} - w_k. \quad (\text{A.26})$$

The skipping mechanism (A.25) provides both an upper and lower bound on the quantity $\frac{\|y_k\|^2}{y_k^T s_k}$, which in turn ensures that the initial L-BFGS Hessian approximation is bounded above and away from zero. The lower bound is attained by repeated application of Cauchy's inequality to condition (A.25). We have from (A.25) that

$$\epsilon \|s_k\|^2 \leq y_k^T s_k \leq \|y_k\| \|s_k\|,$$

and therefore

$$\|s_k\| \leq \frac{1}{\epsilon} \|y_k\|.$$

It follows that

$$s_k^T y_k \leq \|s_k\| \|y_k\| \leq \frac{1}{\epsilon} \|y_k\|^2$$

and hence

$$\frac{\|y_k\|^2}{s_k^T y_k} \geq \epsilon. \quad (\text{A.27})$$

The upper bound is attained by the Lipschitz continuity of sample gradients,

$$\begin{aligned} y_k^T s_k &\geq \epsilon \|s_k\|^2 \\ &\geq \epsilon \frac{\|y_k\|^2}{\Lambda_{O_k}^2}, \end{aligned}$$

Re-arranging the above expression yields the desired upper bound,

$$\frac{\|y_k\|^2}{s_k^T y_k} \leq \frac{\Lambda_{O_k}^2}{\epsilon}. \quad (\text{A.28})$$

Combining (A.27) and (A.28),

$$\epsilon \leq \frac{\|y_k\|^2}{y_k^T s_k} \leq \frac{\Lambda_{O_k}^2}{\epsilon}.$$

The above proves that the eigenvalues of the matrices $B_k^{(0)} = \frac{y_k^T y_k}{s_k^T y_k} I$ at the start of the L-BFGS update cycles are bounded above and away from zero, for all k . The rest of the proof follows the same trace-determinant argument as in the proof of Lemma 3.1, the only difference being that the last inequality in A.17 comes as a result of the cautious update strategy. \square

A.5 Proof of Theorem 3.4 (Nonconvex Case)

Utilizing the result from Lemma 3.3, we can now establish the following result about the behavior of the gradient norm for the multi-batch L-BFGS method with a cautious update strategy.

Theorem 3.4. *Suppose that Assumptions B.1-B.5 above hold. Let $\{w_k\}$ be the iterates generated by the multi-batch L-BFGS method (Algorithm 1) with*

$$\alpha_k = \alpha \in (0, \frac{\mu_1}{\mu_2^2 \eta \Lambda}),$$

where w_0 is the starting point. Also, suppose that if

$$y_k^T s_k < \epsilon \|s_k\|^2,$$

for some $\epsilon > 0$, the inverse L-BFGS Hessian approximation is skipped, $H_{k+1} = H_k$. Then, for all $k \geq 0$,

$$\begin{aligned} \mathbb{E} \left[\frac{1}{L} \sum_{k=0}^{L-1} \|\nabla F(w_k)\|^2 \right] &\leq \frac{\alpha \mu_2^2 \gamma^2 \Lambda}{\mu_1} + \frac{2[F(w_0) - \hat{F}]}{\alpha \mu_1 L} \\ &\xrightarrow{L \rightarrow \infty} \frac{\alpha \mu_2^2 \gamma^2 \Lambda}{\mu_1}. \end{aligned}$$

Proof. Starting with (A.19),

$$\begin{aligned}
\mathbb{E}_{S_k}[F(w_{k+1})] &\leq F(w_k) - \alpha\mu_1\|\nabla F(w_k)\|^2 + \frac{\alpha^2\mu_2^2\Lambda}{2}\mathbb{E}_{S_k}\left[\|\nabla F^{S_k}(w_k)\|\right]^2 \\
&\leq F(w_k) - \alpha\mu_1\|\nabla F(w_k)\|^2 + \frac{\alpha^2\mu_2^2\Lambda}{2}(\gamma^2 + \eta\|\nabla F(w_k)\|^2) \\
&= F(w_k) - \alpha\left(\mu_1 - \frac{\alpha\mu_2^2\eta\Lambda}{2}\right)\|\nabla F(w_k)\|^2 + \frac{\alpha^2\mu_2^2\gamma^2\Lambda}{2} \\
&\leq F(w_k) - \frac{\alpha\mu_1}{2}\|\nabla F(w_k)\|^2 + \frac{\alpha^2\mu_2^2\gamma^2\Lambda}{2},
\end{aligned}$$

where the second inequality holds due to Assumption B.4, and the fourth inequality is obtained by using the upper bound on the step length. Taking an expectation over all batches S_0, S_1, \dots, S_{k-1} and all history starting with w_0 yields

$$\phi_{k+1} - \phi_k \leq -\frac{\alpha\mu_1}{2}\mathbb{E}\|\nabla F(w_k)\|^2 + \frac{\alpha^2\mu_2^2\gamma^2\Lambda}{2}, \quad (\text{A.29})$$

where $\phi_k = \mathbb{E}[F(w_k)]$. Summing (A.29) over the first $L - 1$ iterations

$$\begin{aligned}
\sum_{k=0}^{L-1} [\phi_{k+1} - \phi_k] &\leq -\frac{\alpha\mu_1}{2} \sum_{k=0}^{L-1} \mathbb{E}\|\nabla F(w_k)\|^2 + \sum_{k=0}^{L-1} \frac{\alpha^2\mu_2^2\gamma^2\Lambda}{2} \\
&= -\frac{\alpha\mu_1}{2} \mathbb{E}\left[\sum_{k=0}^{L-1} \|\nabla F(w_k)\|^2\right] + \frac{\alpha^2\mu_2^2\gamma^2\Lambda L}{2}. \quad (\text{A.30})
\end{aligned}$$

The left-hand-side of the above inequality is a telescoping sum

$$\begin{aligned}
\sum_{k=0}^{L-1} [\phi_{k+1} - \phi_k] &= \phi_L - \phi_0 \\
&= \mathbb{E}[F(w_L)] - F(w_0) \\
&\geq \hat{F} - F(w_0).
\end{aligned}$$

Substituting the above expression into (A.30) and re-arranging terms

$$\mathbb{E}\left[\sum_{k=0}^{L-1} \|\nabla F(w_k)\|^2\right] \leq \frac{\alpha\mu_2^2\gamma^2\Lambda L}{\mu_1} + \frac{2[F(w_0) - \hat{F}]}{\alpha\mu_1}.$$

Dividing the above equation by L completes the proof. \square

B Extended Numerical Experiments - Real Datasets

In this Section, we present further numerical results, on the datasets listed in Table 1, in both the multi-batch and fault-tolerant settings. Note, that some of the datasets are too small, and thus, there is no reason to run them on a distributed platform; however, we include them as they are part of the standard benchmarking datasets.

Notation. Let n denote the number of training samples in a given dataset, d the dimension of the parameter vector w , and K the number of MPI processes used. The parameter r denotes the fraction of samples in the dataset used to define the gradient, i.e., $r = \frac{|S|}{n}$. The parameter o denotes the length of overlap between consecutive samples, and is defined as a fraction of the number of samples in a given batch S , i.e., $o = \frac{|O|}{|S|}$.

Table 1: Datasets together with basic statistics. All datasets are available at <https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/binary.html>.

| Dataset | n | d | Size (MB) | K |
|--------------|------------|------------|-----------|----|
| ijcnn (test) | 91,701 | 22 | 14 | 4 |
| cov | 581,012 | 54 | 68 | 4 |
| news20 | 19,996 | 1,355,191 | 134 | 4 |
| rcvtest | 677,399 | 47,236 | 1,152 | 16 |
| url | 2,396,130 | 3,231,961 | 2,108 | 16 |
| kdda | 8,407,752 | 20,216,830 | 2,546 | 16 |
| kddb | 19,264,097 | 29,890,095 | 4,894 | 16 |
| webspam | 350,000 | 16,609,143 | 23,866 | 16 |
| splice-site | 50,000,000 | 11,725,480 | 260,705 | 16 |

We focus on logistic regression classification; the objective function is given by

$$\min_{w \in \mathbb{R}^d} F(w) = \frac{1}{n} \sum_{i=1}^n \log(1 + e^{-y^i (w^T x^i)}) + \frac{\sigma}{2} \|w\|^2,$$

where $(x^i, y^i)_{i=1}^n$ denote the training examples and $\sigma = \frac{1}{n}$ is the regularization parameter.

B.1 Multi-batch L-BFGS Implementation

For the experiments in this section (Figures 5-13), we run four methods:

- (Robust L-BFGS) robust multi-batch L-BFGS (Algorithm 1),
- (L-BFGS) multi-batch L-BFGS without enforcing sample consistency; gradient differences are computed using different samples, i.e., $y_k = g_{k+1}^{S_{k+1}} - g_k^{S_k}$,
- (Gradient Descent) multi-batch gradient descent; obtained by setting $H_k = I$ in Algorithm 1,
- (SGD) serial SGD; at every iteration one sample is used to compute the gradient.

In Figures 5-13 we show the evolution of $\|\nabla F(w)\|$ for different step lengths α , and for various batch ($|S| = r \cdot n$) and overlap ($|O| = o \cdot |S|$) sizes. Each Figure (5-13) consists of 10 plots that illustrate the performance of the methods with the following parameters:

- Top 3 plots: $\alpha = 1$, $o = 20\%$ and $r = 1\%, 5\%, 10\%$
- Middle 3 plots: $\alpha = 0.1$, $o = 20\%$ and $r = 1\%, 5\%, 10\%$
- Bottom 4 plots: $\alpha = 1$, $r = 1\%$ and $o = 5\%, 10\%, 20\%, 30\%$

As is expected for quasi-Newton methods, robust L-BFGS performs best with a step-size $\alpha = 1$, for the most part.

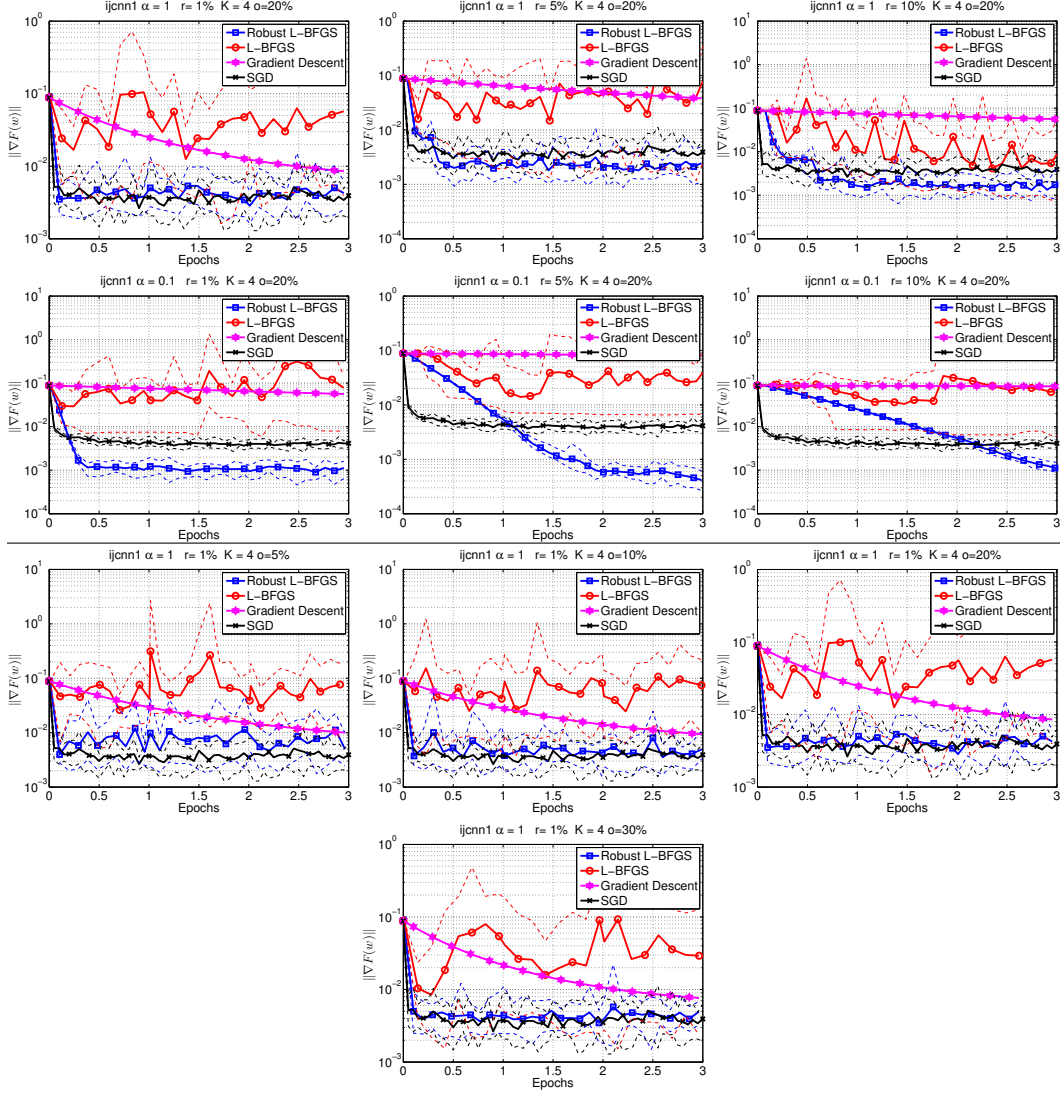


Figure 5: **ijcnn1 dataset**. Comparison of Robust L-BFGS, L-BFGS (multi-batch L-BFGS without enforcing sample consistency), Gradient Descent (multi-batch Gradient method) and SGD. Top part: we used $\alpha \in \{1, 0.1\}$, $r \in \{1\%, 5\%, 10\%\}$ and $o = 20\%$. Bottom part: we used $\alpha = 1$, $r = 1\%$ and $o \in \{5\%, 10\%, 20\%, 30\%\}$. Solid lines show average performance, and dashed lines show worst and best performance, over 10 runs (per algorithm). $K = 4$ MPI processes.

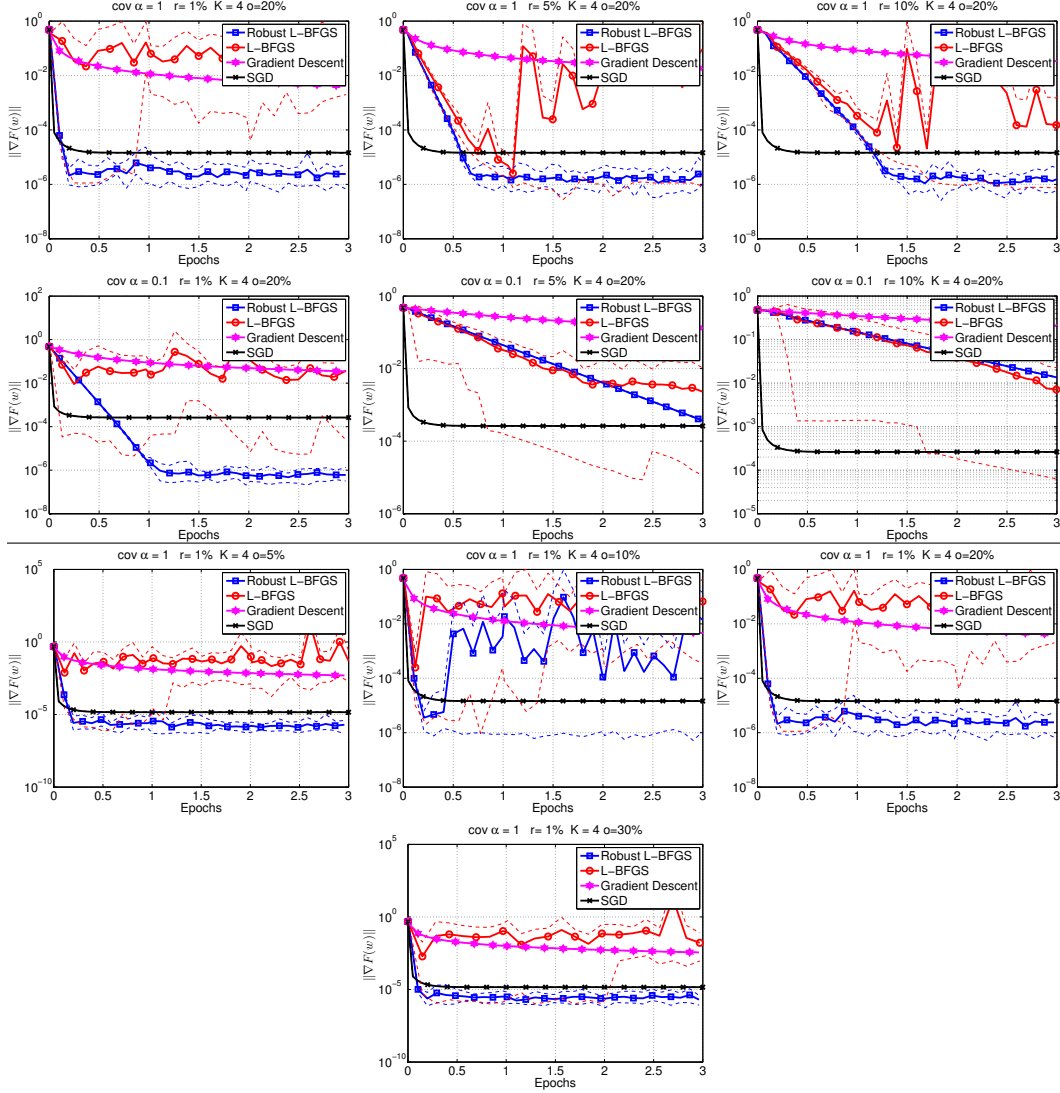


Figure 6: **cov dataset**. Comparison of Robust L-BFGS, L-BFGS (multi-batch L-BFGS without enforcing sample consistency), Gradient Descent (multi-batch Gradient method) and SGD. Top part: we used $\alpha \in \{1, 0.1\}$, $r \in \{1\%, 5\%, 10\%\}$ and $o = 20\%$. Bottom part: we used $\alpha = 1$, $r = 1\%$ and $o \in \{5\%, 10\%, 20\%, 30\%\}$. Solid lines show average performance, and dashed lines show worst and best performance, over 10 runs (per algorithm). $K = 4$ MPI processes.

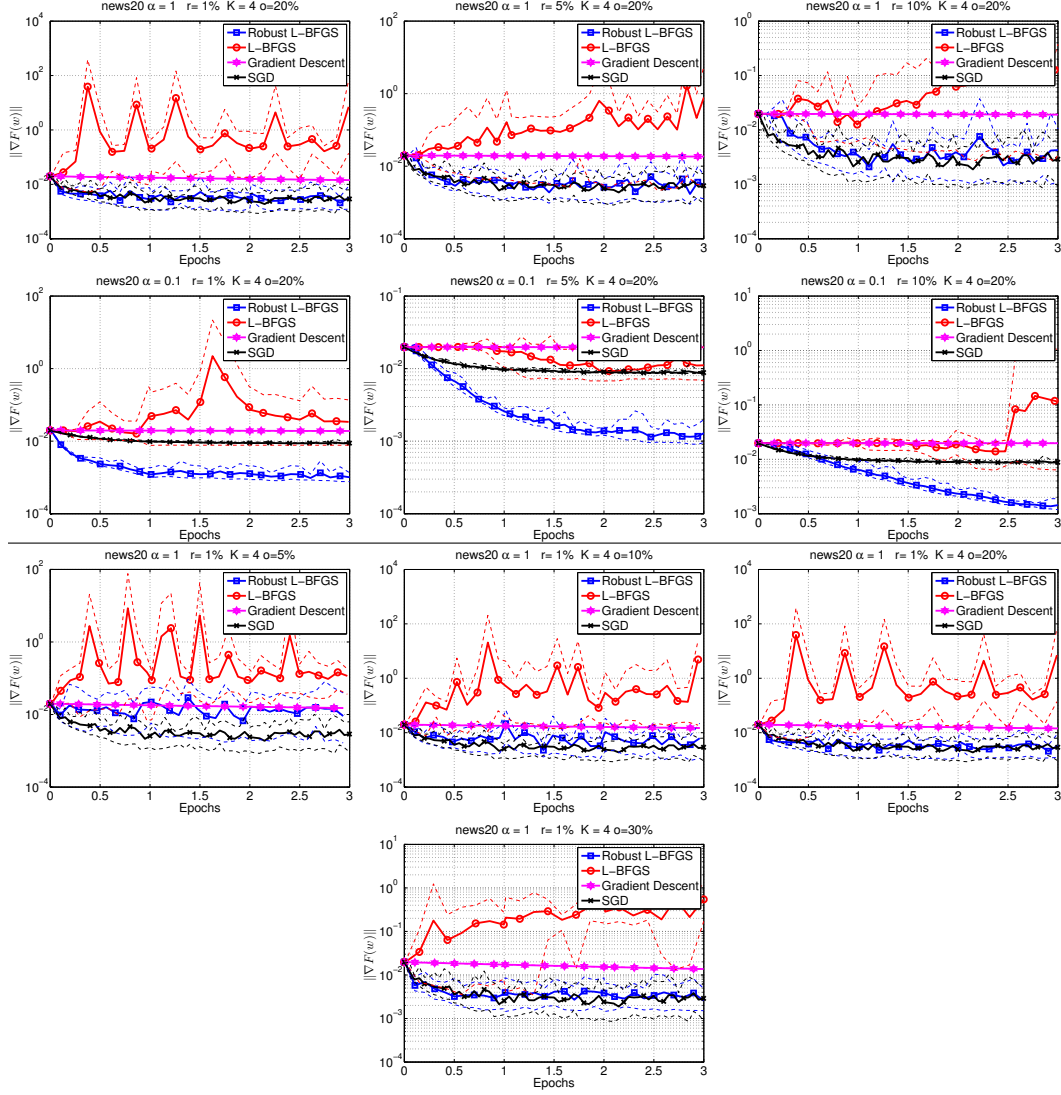


Figure 7: **news20 dataset**. Comparison of Robust L-BFGS, L-BFGS (multi-batch L-BFGS without enforcing sample consistency), Gradient Descent (multi-batch Gradient method) and SGD. Top part: we used $\alpha \in \{1, 0.1\}$, $r \in \{1\%, 5\%, 10\%\}$ and $o = 20\%$. Bottom part: we used $\alpha = 1$, $r = 1\%$ and $o \in \{5\%, 10\%, 20\%, 30\%\}$. Solid lines show average performance, and dashed lines show worst and best performance, over 10 runs (per algorithm). $K = 4$ MPI processes.

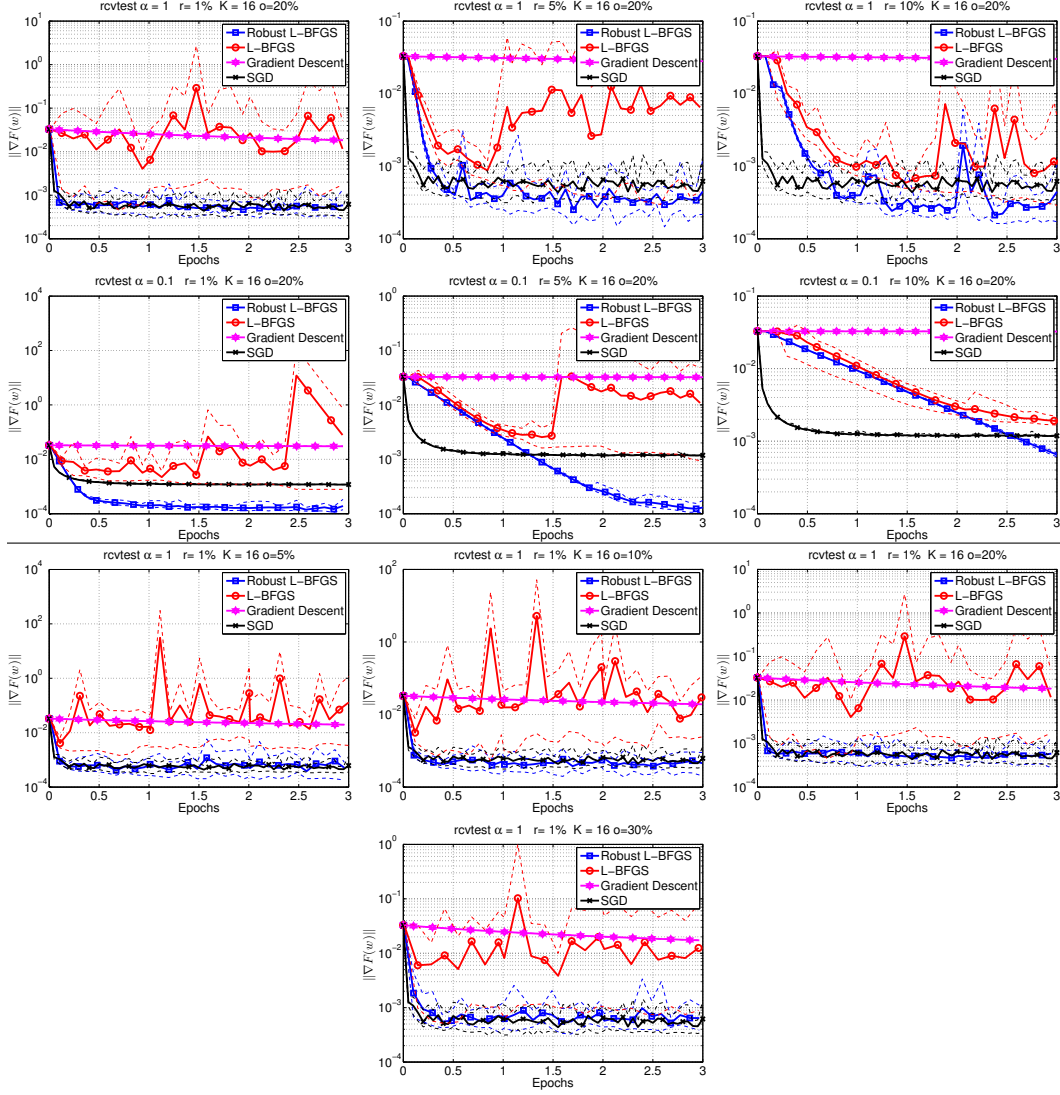


Figure 8: **rcvtest dataset**. Comparison of Robust L-BFGS, L-BFGS (multi-batch L-BFGS without enforcing sample consistency), Gradient Descent (multi-batch Gradient method) and SGD. Top part: we used $\alpha \in \{1, 0.1\}$, $r \in \{1\%, 5\%, 10\%\}$ and $o = 20\%$. Bottom part: we used $\alpha = 1$, $r = 1\%$ and $o \in \{5\%, 10\%, 20\%, 30\%\}$. Solid lines show average performance, and dashed lines show worst and best performance, over 10 runs (per algorithm). $K = 16$ MPI processes.

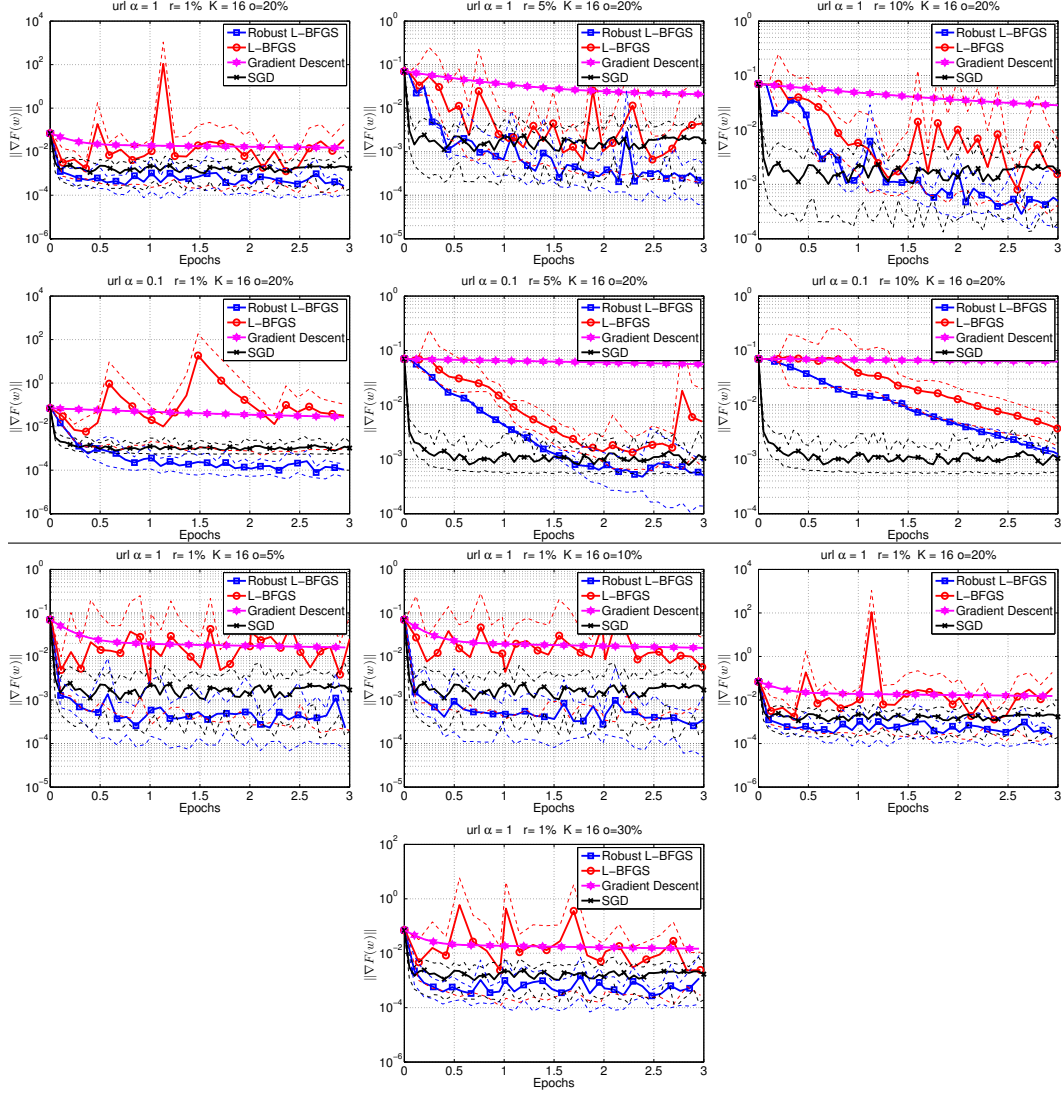


Figure 9: **url dataset**. Comparison of Robust L-BFGS, L-BFGS (multi-batch L-BFGS without enforcing sample consistency), Gradient Descent (multi-batch Gradient method) and SGD. Top part: we used $\alpha \in \{1, 0.1\}$, $r \in \{1\%, 5\%, 10\%\}$ and $o = 20\%$. Bottom part: we used $\alpha = 1$, $r = 1\%$ and $o \in \{5\%, 10\%, 20\%, 30\%\}$. Solid lines show average performance, and dashed lines show worst and best performance, over 10 runs (per algorithm). $K = 16$ MPI processes.

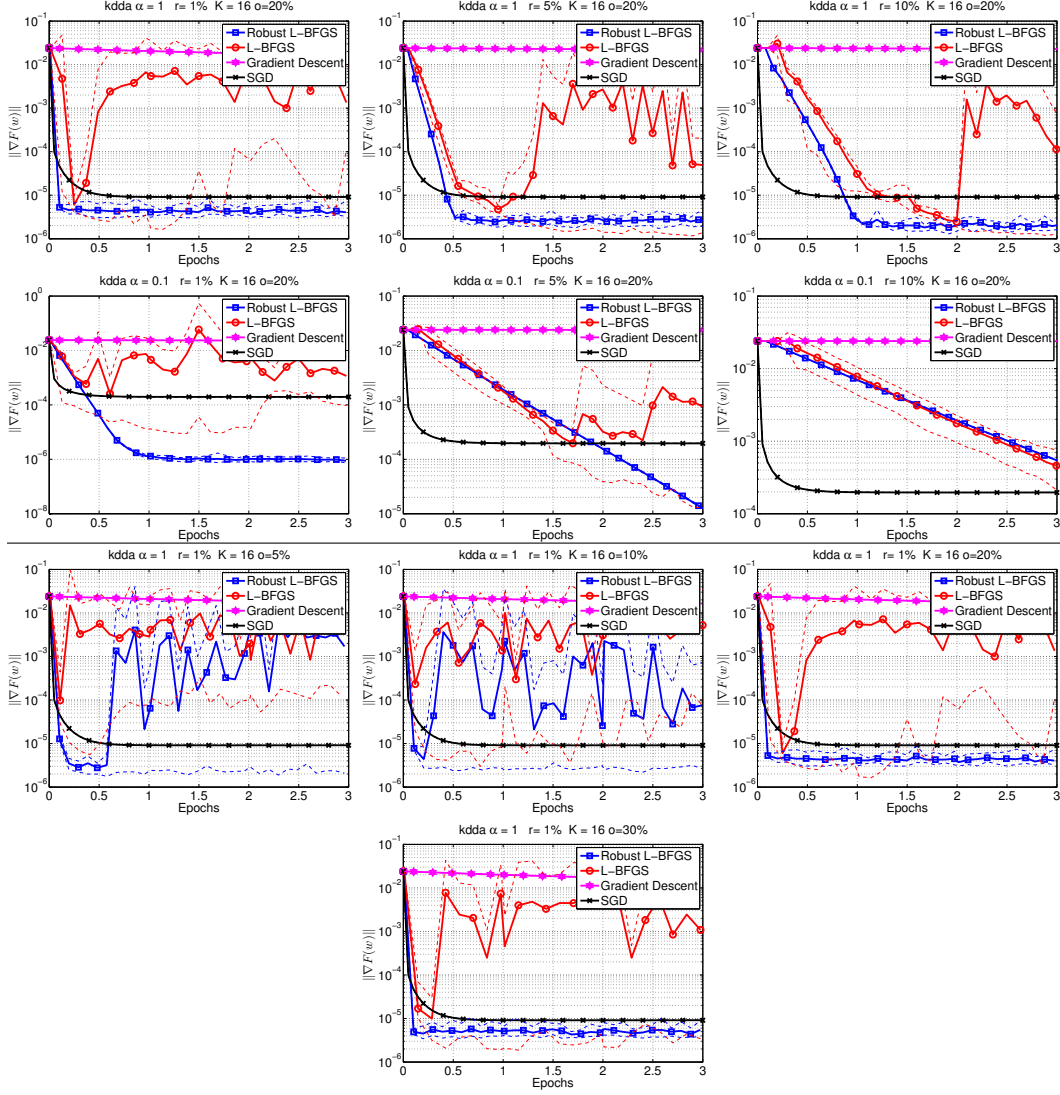


Figure 10: **kdda dataset**. Comparison of Robust L-BFGS, L-BFGS (multi-batch L-BFGS without enforcing sample consistency), Gradient Descent (multi-batch Gradient method) and SGD. Top part: we used $\alpha \in \{1, 0.1\}$, $r \in \{1\%, 5\%, 10\%\}$ and $o = 20\%$. Bottom part: we used $\alpha = 1$, $r = 1\%$ and $o \in \{5\%, 10\%, 20\%, 30\%\}$. Solid lines show average performance, and dashed lines show worst and best performance, over 10 runs (per algorithm). $K = 16$ MPI processes.

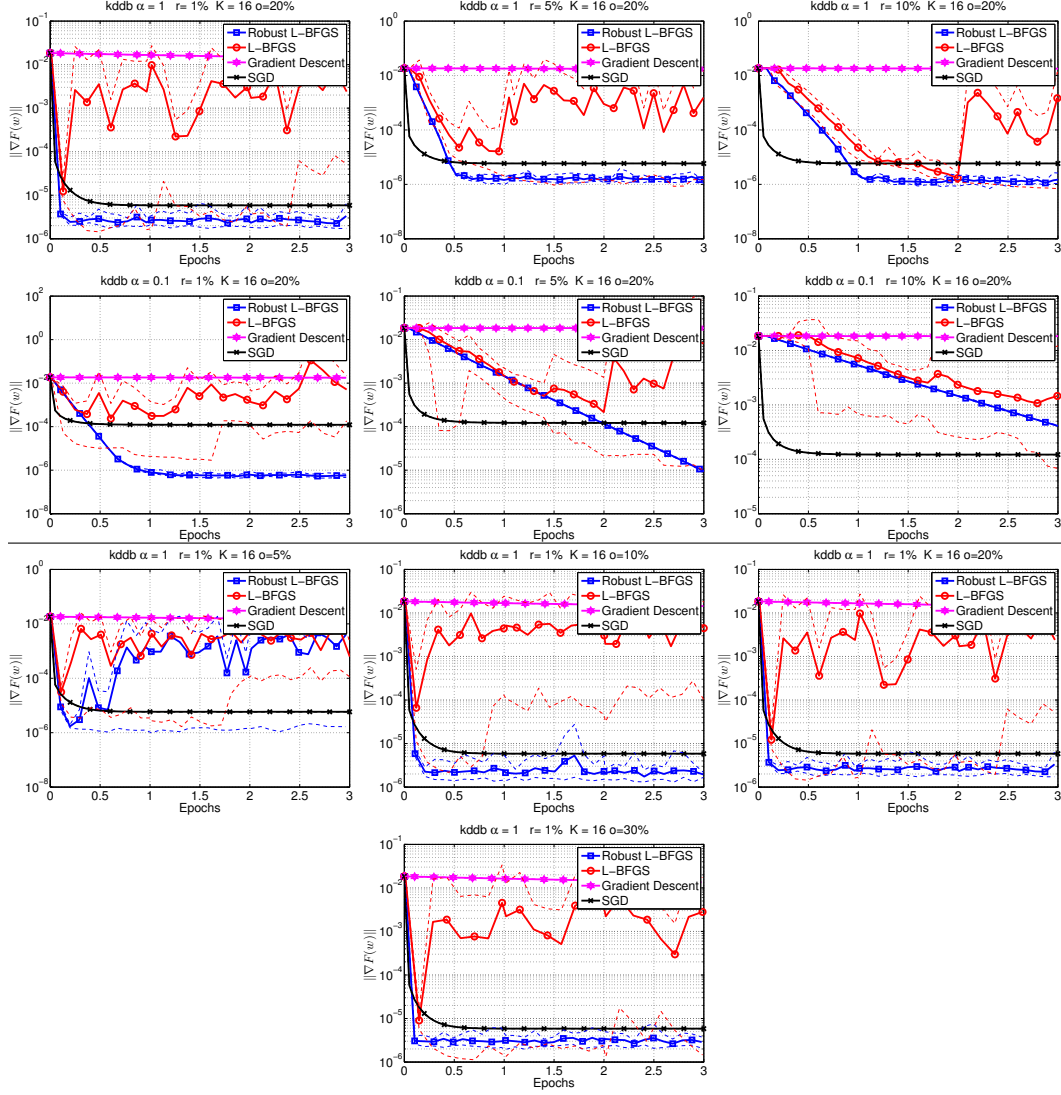


Figure 11: **kddb dataset**. Comparison of Robust L-BFGS, L-BFGS (multi-batch L-BFGS without enforcing sample consistency), Gradient Descent (multi-batch Gradient method) and SGD. Top part: we used $\alpha \in \{1, 0.1\}$, $r \in \{1\%, 5\%, 10\%\}$ and $o = 20\%$. Bottom part: we used $\alpha = 1$, $r = 1\%$ and $o \in \{5\%, 10\%, 20\%, 30\%\}$. Solid lines show average performance, and dashed lines show worst and best performance, over 10 runs (per algorithm). $K = 16$ MPI processes.

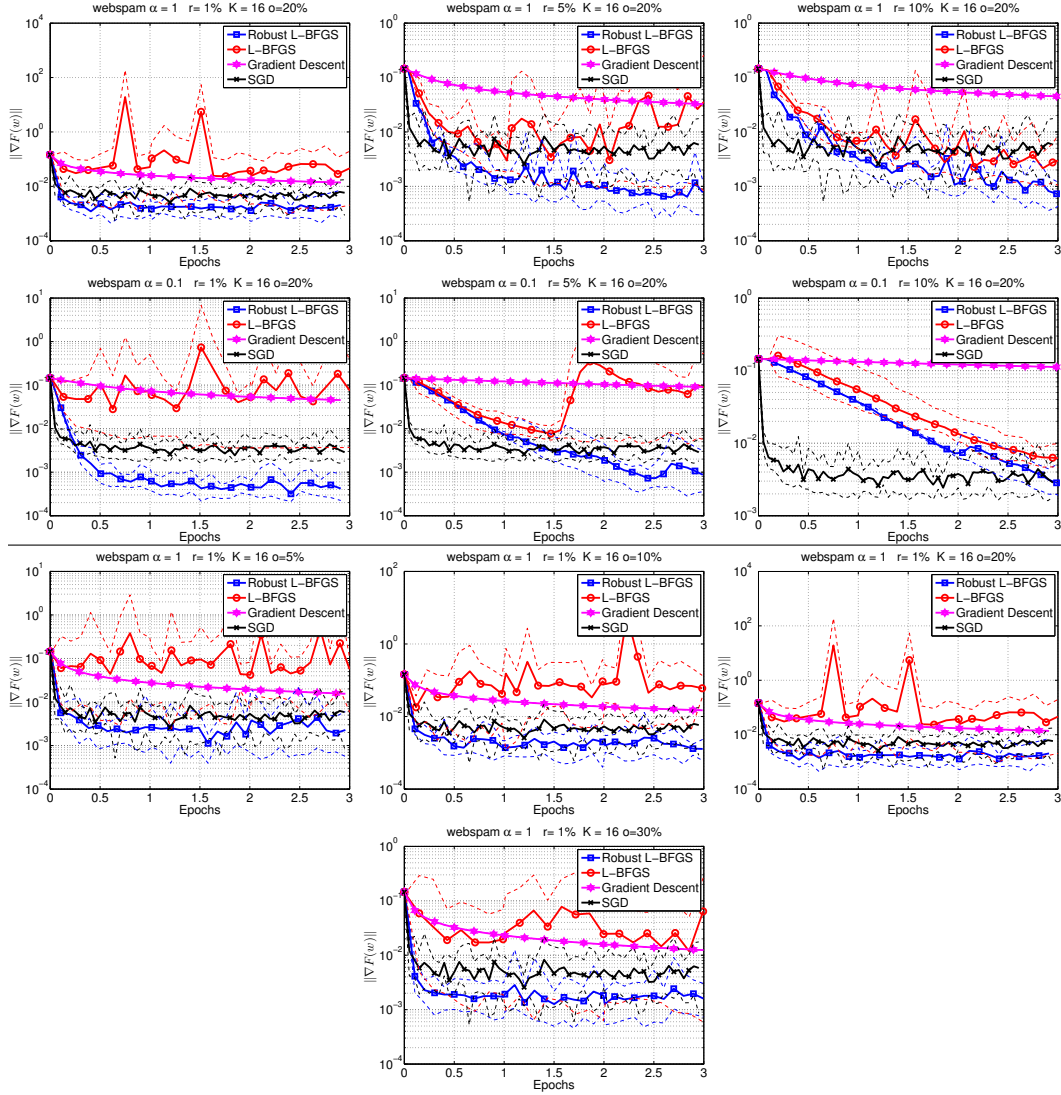


Figure 12: **webspam dataset**. Comparison of Robust L-BFGS, L-BFGS (multi-batch L-BFGS without enforcing sample consistency), Gradient Descent (multi-batch Gradient method) and SGD. Top part: we used $\alpha \in \{1, 0.1\}$, $r \in \{1\%, 5\%, 10\%\}$ and $o = 20\%$. Bottom part: we used $\alpha = 1$, $r = 1\%$ and $o \in \{5\%, 10\%, 20\%, 30\%\}$. Solid lines show average performance, and dashed lines show worst and best performance, over 10 runs (per algorithm). $K = 16$ MPI processes.

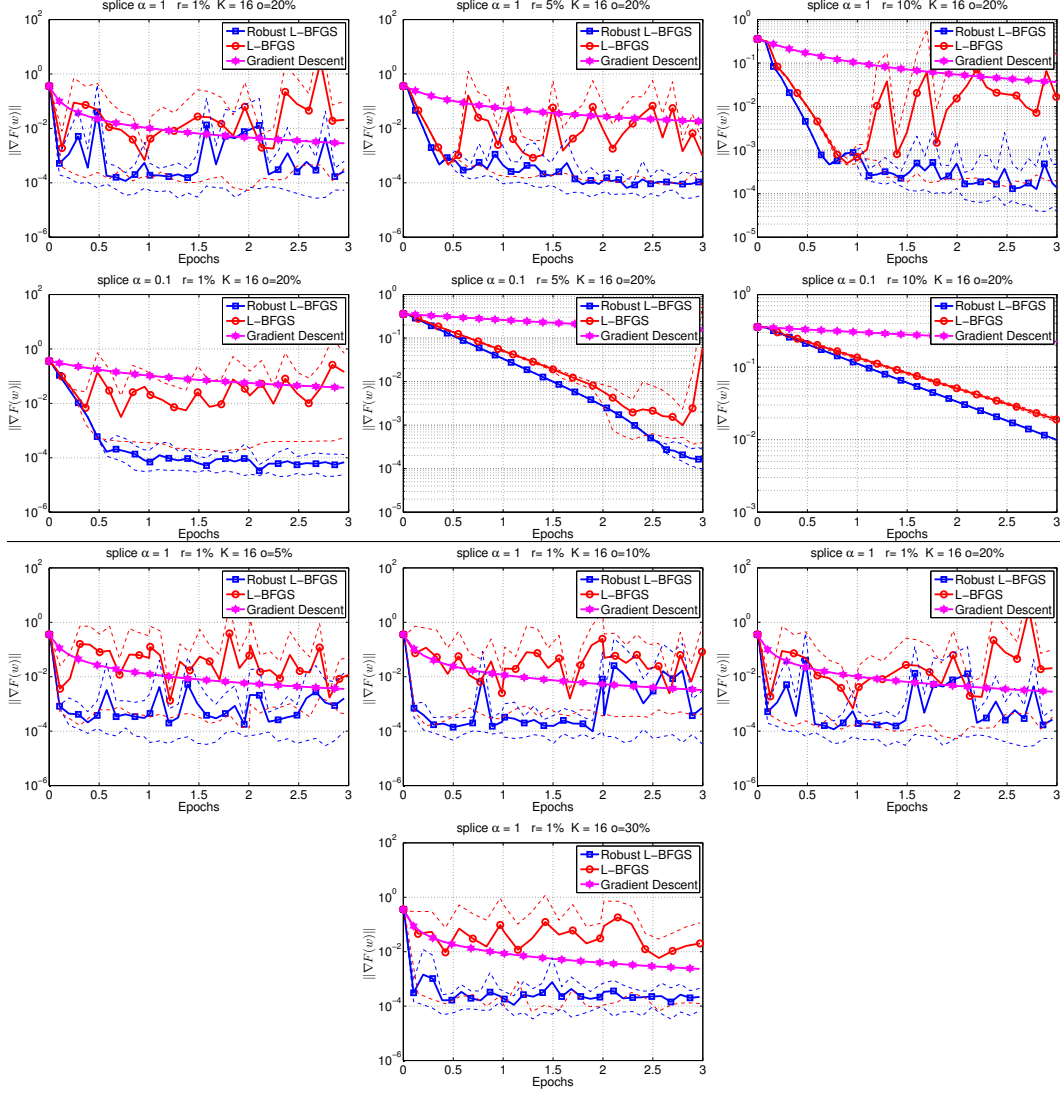


Figure 13: **splice-cite dataset**. Comparison of Robust L-BFGS, L-BFGS (multi-batch L-BFGS without enforcing sample consistency), Gradient Descent (multi-batch Gradient method) and SGD. Top part: we used $\alpha \in \{1, 0.1\}$, $r \in \{1\%, 5\%, 10\%\}$ and $o = 20\%$. Bottom part: we used $\alpha = 1$, $r = 1\%$ and $o \in \{5\%, 10\%, 20\%, 30\%\}$. Solid lines show average performance, and dashed lines show worst and best performance, over 10 runs (per algorithm). $K = 16$ MPI processes. (No Serial SGD experiments due to memory limitations of our cluster.)

B.2 Fault-tolerant L-BFGS Implementation

If we run a distributed algorithm, for example on a shared computer cluster, then we may experience delays. Such delays can be caused by other processes running on the same compute node, node failures and for other reasons. As a result, given a computational (time) budget, these delays may cause nodes to fail to return a value. To illustrate this behavior, and to motivate the robust fault-tolerant L-BFGS method, we run a simple benchmark MPI code on two different environments:

- **Amazon EC2** – Amazon EC2 is a cloud system provided by Amazon. It is expected that if load balancing is done properly, the execution time will have small noise; however, the network and communication can still be an issue. (4 MPI processes)
- **Shared Cluster** – In our shared cluster, multiple jobs run on each node, with some jobs being more demanding than others. Even though each node has 16 cores, the amount of resources each job can utilize changes over time. In terms of communication, we have a GigaBit network. (11 MPI processes, running on 11 nodes)

We run a simple code on the cloud/cluster, with MPI communication. We generate two matrices $A, B \in \mathbb{R}^{n \times n}$, then synchronize all MPI processes and compute $C = A \cdot B$ using the GSL C-BLAS library. The time is measured and recorded as computational time. After the matrix product is computed, the result is sent to the master/root node using asynchronous communication, and the time required is recorded. The process is repeated 3000 times.

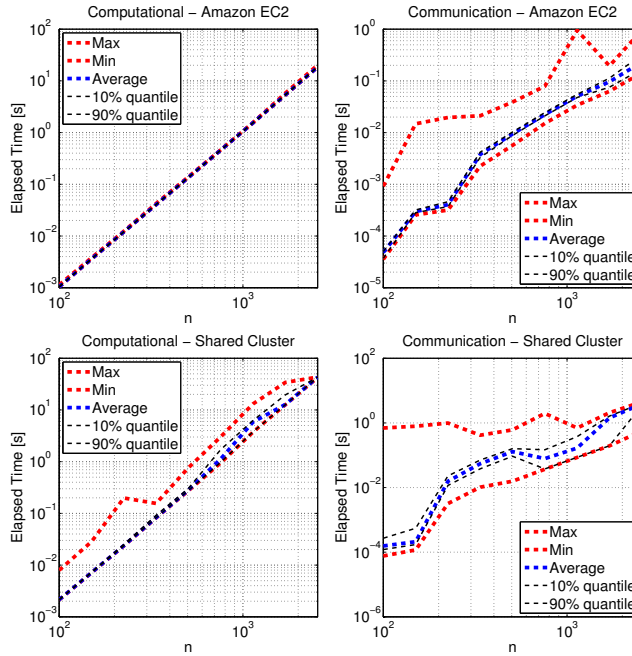


Figure 14: Distribution of Computation and Communication Time for Amazon EC2 and Shared Cluster. Figures show worst and best time, average time and 10% and 90% quantiles. Amazon Cloud EC: In the experiment: 4 MPI processes; Shared Cluster: 11 MPI processes.

The results of the experiment described above are captured in Figure 14. As expected, on the Amazon EC2 cloud, the matrix-matrix multiplication takes roughly the same time for all replications and the noise in communication is relatively small. In this example the cost of communication is negligible when compared to the cost of computation. On our shared cluster, one cannot guarantee that all resources are exclusively used for a specific process, and thus, the computation and communication time is considerably more stochastic and unbalanced. For some cases the difference between the minimum and maximum computation (communication) time varies by an order of magnitude or more. Hence, on such a platform a fault-tolerant algorithm that only uses information from nodes that return an update within a preallocated budget is a natural choice.

In Figures 15-19 we show a comparison of the proposed robust multi-batch L-BFGS method and the multi-batch L-BFGS method that does not enforce sample consistency (L-BFGS). In these experiments, p denotes the probability that a single node (MPI process) will not return a gradient evaluated on local data within a given time budget. We illustrate the performance of the methods for $\alpha = 0.1$ and $p \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$. We observe that the robust implementation is not affected much by the failure probability p .

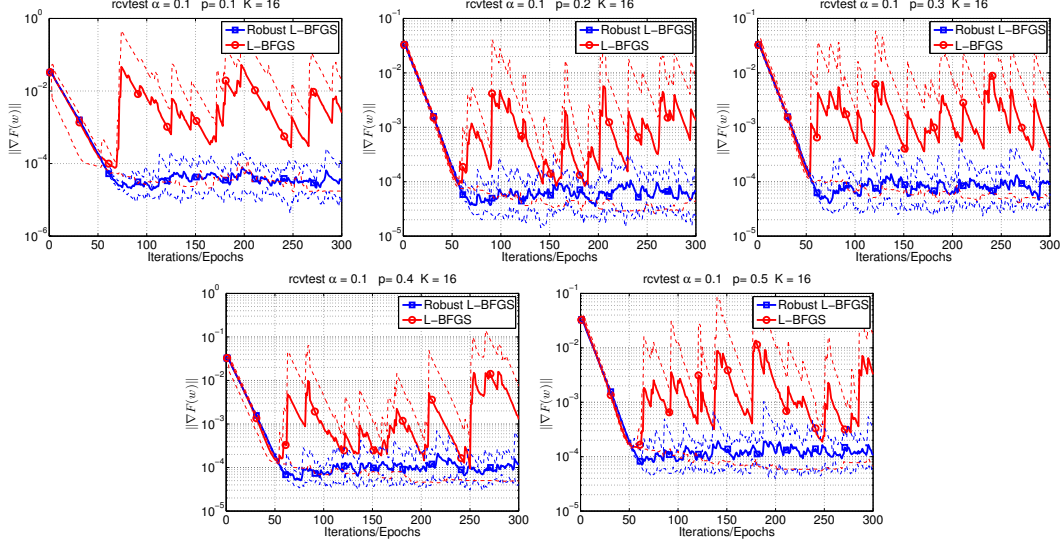


Figure 15: **rcvtest dataset**. Comparison of Robust L-BFGS and L-BFGS in the presence of faults. We used $\alpha = 0.1$ and $p \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$. Solid lines show average performance, and dashed lines show worst and best performance, over 10 runs (per algorithm). $K = 16$ MPI processes.

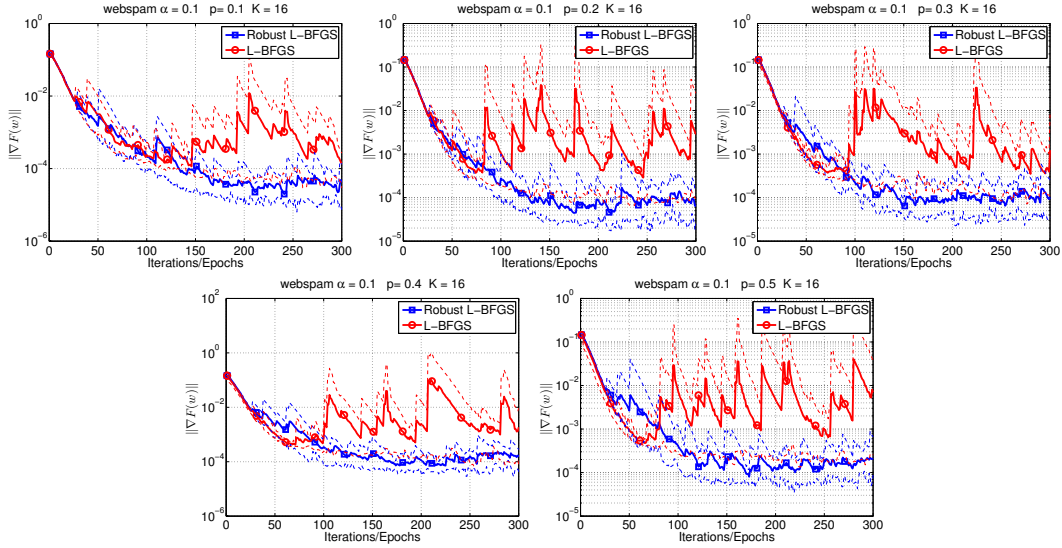


Figure 16: **webspam dataset**. Comparison of Robust L-BFGS and L-BFGS in the presence of faults. We used $\alpha = 0.1$ and $p \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$. Solid lines show average performance, and dashed lines show worst and best performance, over 10 runs (per algorithm). $K = 16$ MPI processes.




Figure 17: **kdda dataset**. Comparison of Robust L-BFGS and L-BFGS in the presence of faults. We used $\alpha = 0.1$ and $p \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$. Solid lines show average performance, and dashed lines show worst and best performance, over 10 runs (per algorithm). $K = 16$ MPI processes.

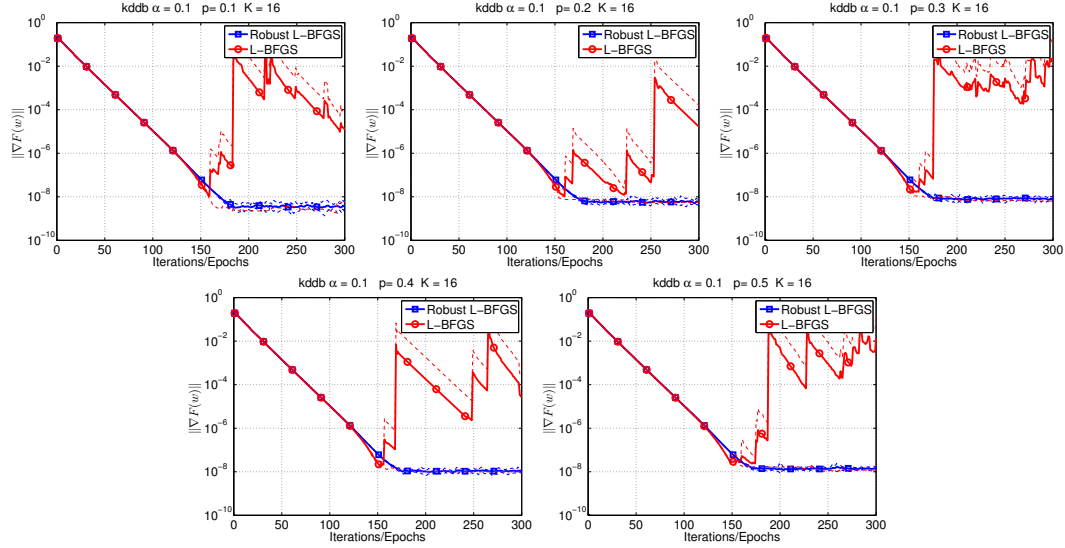


Figure 18: **kddb dataset**. Comparison of Robust L-BFGS and L-BFGS in the presence of faults. We used $\alpha = 0.1$ and $p \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$. Solid lines show average performance, and dashed lines show worst and best performance, over 10 runs (per algorithm). $K = 16$ MPI processes.

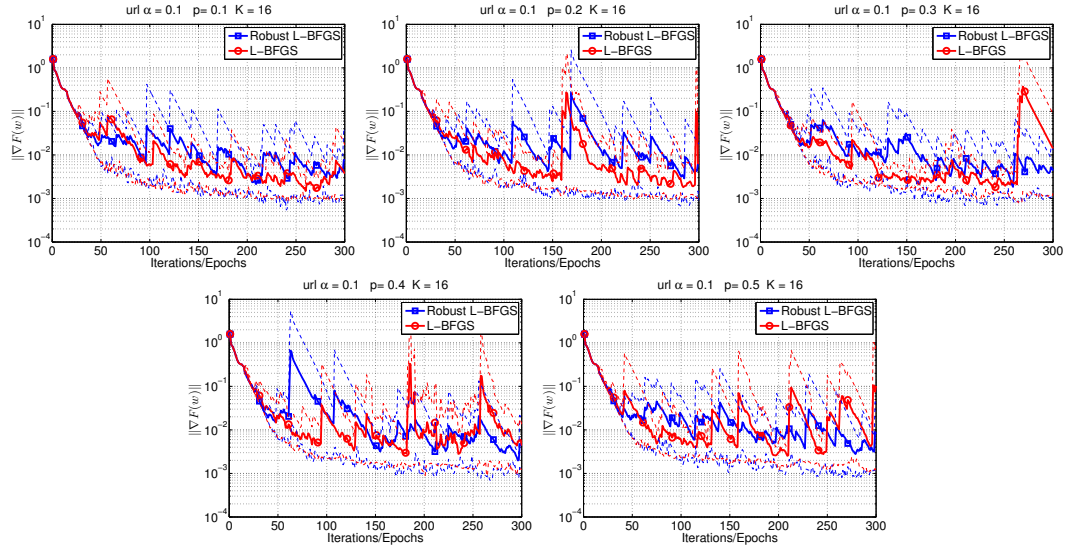


Figure 19: **url dataset**. Comparison of Robust L-BFGS and L-BFGS in the presence of faults. We used $\alpha = 0.1$ and $p \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$. Solid lines show average performance, and dashed lines show worst and best performance, over 10 runs (per algorithm). $K = 16$ MPI processes.

C Scaling of Robust Multi-Batch L-BFGS Implementation

In this Section, we study the strong and weak scaling properties of the robust multi-batch L-BFGS method on an artificial dataset. For various values of r and K , we measure the time needed to compute a gradient (Gradient) and the time needed to compute and communicate the gradient (Gradient+C), as well as, the time needed to compute the L-BFGS direction (L-BFGS) and the associated communication overhead (L-BFGS+C).

C.1 Strong Scaling

Figure 20 depicts the strong scaling properties of our proposed algorithm. We generate a dataset with $n = 10^7$ samples and $d = 10^4$ dimensions, where each sample has 160 randomly chosen non-zero elements (dataset size 24GB). We run our code for different values of r (different batch sizes S_k), with $K = 1, 2, \dots, 128$ number of MPI processes.

One can observe that the compute time for the gradient and the L-BFGS direction decreases as K is increased. However, when communication time is considered, the combined cost increases slightly as K is increased. Notice that for large K , even when $r = 10\%$ (i.e., 10% of all samples processed in one iteration, $\sim 18\text{MB}$ of data), the amount of local work is not sufficient to overcome the communication cost.




Figure 20: Strong scaling of robust multi-batch L-BFGS on a problem with artificial data; $n = 10^7$ and $d = 10^4$. Each sample has 160 non-zero elements. +C indicates that we include communication time to the gradient computation and L-BFGS update computation.

C.2 Weak Scaling - Fixed Problem Dimension, Increasing Data Size

In order to illustrate the weak scaling properties of the algorithm, we generate a data-matrix $X \in R^{10^7 \times 10^4}$, and run it on a shared cluster with $K = 1, 2, 4, 8, \dots, 128$ MPI processes. For a given number of MPI processes (K), each sample contains $10 \cdot K$ non-zero elements. Effectively, the dimension of the problem is fixed, but sparsity of the data is decreased as more MPI processes are used. The size of the input data is $1.5 \cdot K$ GB (i.e., 1.5GB per MPI process).

The compute time for the gradient is almost constant, this is because the amount of work per MPI process (rank) is almost identical; see Figure 21. On the other hand, because we are using a Vector-Free L-BFGS implementation [9] for computing the L-BFGS direction, the amount of time needed for each node to compute the L-BFGS direction is decreasing as K is increased. However, increasing K does lead to larger communication overhead, which can be observed in Figure 21. For $K = 128$ (192GB of data) and $r = 10\%$, almost 20GB of data are processed per iteration in less than 0.1 seconds, which implies that one epoch would take around 1 second.




Figure 21: Weak scaling of robust multi-batch L-BFGS on a problem with artificial data; $n = 10^7$ and $d = 10^4$. Each sample has $10 \cdot K$ non-zero elements. $+C$ indicates that we also include communication time to the gradient computation and L-BFGS update computation.

C.3 Increasing Problem Dimension, Fixed Data Size and K

In this experiment, we investigate the effect of a change in the dimension d of the problem on the performance of the algorithm. We fix the size of data ($29GB$) and the number of MPI processes ($K = 8$). We generate data with $n = 10^7$ samples, where each sample has 200 non-zero elements. Figure 22 shows that increasing the dimension d has a mild effect on the computation time of the gradient, while the effect on the time needed to compute the L-BFGS direction is more apparent. However, if communication time is taken into consideration, the time required for the gradient computation and the L-BFGS direction computation increase as d is increased.




Figure 22: Scaling of robust multi-batch L-BFGS on a problem with artificial data; $n = 10^7$ samples, with increasing d and $K = 8$ MPI processes. Each sample had 200 non-zero elements. $+C$ indicates that we also include communication time to the gradient computation and L-BFGS update computation.