
Efficient Distributed Hessian Free Algorithm for Large-scale Empirical Risk Minimization via Accumulating Sample Strategy

Majid Jahani
Lehigh University

Xi He
Lehigh University

Chenxin Ma
Lehigh University

Aryan Mokhtari
UT Austin

Dheevatsa Mudigere
Facebook

Alejandro Ribeiro
University of Pennsylvania

Martin Takáč
Lehigh University

Abstract

In this paper, we propose a Distributed Accumulated Newton Conjugate gradiEnt (DANCE) method in which sample size is gradually increasing to quickly obtain a solution whose empirical loss is under satisfactory statistical accuracy. Our proposed method is multistage in which the solution of a stage serves as a warm start for the next stage which contains more samples (including the samples in the previous stage). The proposed multi-stage algorithm reduces the number of passes over data to achieve the statistical accuracy of the full training set. Moreover, our algorithm in nature is easy to be distributed and shares the strong scaling property indicating that acceleration is always expected by using more computing nodes. Various iteration complexity results regarding descent direction computation, communication efficiency and stopping criteria are analyzed under convex setting. Our numerical results illustrate that the proposed method outperforms other comparable methods for solving learning problems including neural networks.

1 Introduction

In the field of machine learning, solving the expected risk minimization problem has received lots of attentions over the last decades, which is in the form of

$$\min_{w \in \mathbb{R}^d} L(w) = \min_{w \in \mathbb{R}^d} \mathbb{E}_z[f(w, z)], \quad (1)$$

where z is a $d + 1$ dimensional random variable containing both feature variables and a response variable. $f(w, z)$ is a loss function with respect to w and any fixed value of z .

In most practical problems, the distribution of z is either unknown or leading great difficulties evaluating the expected loss. One general idea is to estimate the expectation with a statistical average over a large number of independent and identically distributed data samples of z , denoted by $\{z_1, z_2, \dots, z_N\}$ where N is the total number of samples. Thus, the problem in (1) can be rewritten as the Empirical Risk Minimization (ERM) problem

$$\min_{w \in \mathbb{R}^d} L_N(w) = \min_{w \in \mathbb{R}^d} \frac{1}{N} \sum_{i=1}^N f_i(w), \quad (2)$$

where $f_i(w) = f(w, z_i)$.

Many studies have been done on developing optimization algorithms to find an optimal solution of the above problem under different setting. For example, the studies by [Beck and Teboulle, 2009, Nesterov, 2013, Drusvyatskiy et al., 2018, Ma et al., 2017] are some of the gradient-based methods which require at least one pass over all data samples to evaluate the gradient $\nabla L_N(w)$. As the sample size N becomes larger, these methods would be less efficient compared to stochastic gradient methods where the gradient is approximated based on a small number of samples [Johnson and Zhang, 2013, Roux et al., 2012, Defazio et al., 2014, Shalev-Shwartz and Zhang, 2013, Konečný and Richtárik, 2017, Nguyen et al., 2017].

Second order methods are well known to share faster convergence rate by utilizing the Hessian information. Recently, several papers by [Byrd et al., 2016, Berahas et al., , Jahani et al., 2019, Berahas et al., 2016, Schraudolph et al., 2007, Berahas et al., 2019, Mokhtari and Ribeiro, 2015, Roosta-Khorasani and Mahoney, 2018] have studied how to apply second orders methods to solve ERM

Table 1: Comparison of computational complexity between different algorithms for convex functions

Method	Complexity
AdaNewton	$\mathcal{O}(2Nd^2 + d^3 \log_2(N))$
k -TAN	$\mathcal{O}(2Nd^2 + d^2 \log_2(N) \log k)$
DANCE	$\mathcal{O}((\log_2(N))^3 N^{1/4} d^2)$

problem. However, evaluating the Hessian inverse or its approximation is always computationally costly, leading to a significant difficulty on applying these methods on large-scale problems.

The above difficulty can be addressed by applying the idea of adaptive sample size methods by recent works of [Mokhtari and Ribeiro, 2017, Eisen et al., 2018, Mokhtari et al., 2016], which is based on the following two facts. First, the empirical risk and the statistical loss have different minimizers, and it is not necessary to go further than the difference between the mentioned two objectives, which is called *statistical accuracy*. More importantly, if we increase the size of the samples in the ERM problem the solutions should not significantly change as samples are drawn from a fixed but unknown probability distribution. The key idea of adaptive samples size methods is to solve an ERM problem with a small number of samples upto its statistical accuracy and use the obtained solution as a warm start for the next ERM problem which contains more samples. In particular, [Mokhtari et al., 2016] reduced the complexity of Newton’s method by incorporating the adaptive sample size idea; however, their approach still requires computing $\log N$ Hessian inversions which is costly when the problem dimension d is large. In order to decrease the cost of computing the Hessian inverse, [Eisen et al., 2018] proposed the k -Truncated Adaptive Newton (k -TAN) approach in which the inverse of Hessian is approximated by truncating the k largest eigenvalues of the Hessian. The cost per iteration of this approach is $\mathcal{O}((\log k + n)d^2)$ which may not be satisfactory either when d is large or k is close to d .

In this paper, we propose an increasing sample size second-order method which solves the Newton step in ERM problems more efficiently. Our proposed algorithm, called Distributed Accumulated Newton Conjugate gradiEnt (DANCE), starts with a small number of samples and minimizes their corresponding ERM problem. This subproblem is solved up to a specific accuracy, and the solution of this stage is used as a warm start for the next stage. We increase the number of samples in the next stage which contains all the previous samples, and use the previous solution as a warm start for solving the new ERM. Such procedure is run iteratively until either all the samples have been included, or we find that it is unnecessary to further

increase the sample size. Our DANCE method combines the idea of increasing sample size and the inexact damped Newton method discussed in the works of [Zhang and Lin, 2015] and [Ma and Takáč, 2016]. Instead of solving the Newton system directly, we apply preconditioned conjugate gradient (PCG) method as the solver for each Newton step. Also, it is always a challenging problem to run first order algorithms such as SGD and Adam by [Kingma and Ba, 2014] in a distributed fashion. The DANCE method is designed to be easily parallelized and shares the strong scaling property, i.e., linear speed-up property. Since it is possible to split gradient and Hessian-vector product computations across different machines, it is always expected to get extra acceleration via increasing the number of computational nodes. We formally characterize the required number of communication rounds to reach the statistical accuracy of the full dataset. For a distributed setting, we show that DANCE is communication efficient in both theory and practice. In particular, Table 1 highlights the advantage of DANCE with respect to other adaptive sample size methods which will be discussed in more details in Section 4.

We organize this paper as following. In Section 2, we introduce the necessary assumptions and the definition of statistical accuracy. Section 3 describes the proposed algorithm and its distributed version. Section 4 explores the theoretical guarantees on complexity of DANCE. In Section 5, we demonstrate the outstanding performance of our algorithm in practice. In Section 6, we close the paper by concluding remarks.

2 Problem Formulation

In this paper, we focus on finding the optimal solution w^* of the problem in (1). As described earlier, due to difficulties in the expected risk minimization, as an alternative, we aim to find a solution for the empirical loss function $L_N(w)$, which is the empirical mean over N samples. Now, consider the empirical loss $L_n(w)$ associated with $n \leq N$ samples. In [Bousquet and Bottou, 2008] and [Bottou, 2010], it has been shown that the difference between the expected loss and the empirical loss L_n with high probability (w.h.p.) is upper bounded by the statistical accuracy V_n , i.e., w.h.p.

$$\sup_{w \in \mathbb{R}^d} |L(w) - L_n(w)| \leq V_n. \quad (3)$$

In other words, there exists a constant ϑ such that the inequality (3) holds with probability of at least $1 - \vartheta$. Generally speaking, statistical accuracy V_n depends on n (although it depends on ϑ too, but for simplicity in notation we just consider the size of the samples), and is of order $V_n = \mathcal{O}(1/n^\gamma)$ where $\gamma \in [0.5, 1]$ [Vapnik, 2013,

Bousquet, 2002, Bartlett et al., 2006].

For problem (2), if we find an approximate solution w_n which satisfies the inequality $L_n(w_n) - L_n(\hat{w}_n) \leq V_n$, where \hat{w}_n is the true minimizer of L_n , it is not necessary to go further and find a better solution (a solution with less optimization error). The reason comes from the fact that for a more accurate solution the summation of estimation and optimization errors does not become smaller than V_n . Therefore, when we say that w_n is a V_n -suboptimal solution for the risk L_n , it means that $L_n(w_n) - L_n(\hat{w}_n) \leq V_n$. In other words, w_n solves problem (2) within its statistical accuracy.

It is crucial to note that if we add an additional term in the magnitude of V_n to the empirical loss L_n , the new solution is also in the similar magnitude as V_n to the expected loss L . Therefore, we can regularize the non-strongly convex loss function L_n by $cV_n\|w\|^2/2$ and consider it as the following problem:

$$\min_{w \in \mathbb{R}^d} R_n(w) := \frac{1}{n} \sum_{i=1}^n f_i(w) + \frac{cV_n}{2} \|w\|^2. \quad (4)$$

The noticeable feature of the new empirical risk R_n is that R_n is cV_n -strongly convex¹, where c is a positive constant depending on the VC dimension of the problem. Thus, we can utilize any practitioner-favorite algorithm. Specifically, we are willing to apply the inexact damped Newton method, which will be discussed in the next section. Due to the fact that a larger strong-convexity parameter leads to a faster convergence, we could expect that the first few steps would converge fast since the values of cV_n in these steps are large (larger statistical accuracy), as will be discussed in Theorem 1. From now on, when we say w_n is a V_n -suboptimal solution of the risk R_n , it means that $R_n(w_n) - R_n(w_n^*) \leq V_n$, where w_n^* is the true optimal solution of the risk R_n . Our final aim is to find w_N which is V_N -suboptimal solution for the risk R_N which is the risk over the whole dataset.

In the rest of this section, first we define the self-concordant functions which have the property that its third derivative can be controlled by its second derivative. By assuming that function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ has continuous third derivative, we define self-concordant function as follows.

Definition 1. A convex function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is ρ_f -self-concordant if for any $w \in \text{dom}(f)$ and $u \in \mathbb{R}^d$

$$|u^T (f'''(w)[u])u| \leq \rho_f (u^T \nabla^2 f(w) u)^{\frac{3}{2}}, \quad (5)$$

where $f'''(w)[u] := \lim_{t \rightarrow 0} \frac{1}{t} (\nabla^2 f(w + tu) - \nabla^2 f(w))$. As it is discussed in [Nesterov, 2013], any self-concordant function f with parameter ρ_f can be

¹ cV_n depends on number of samples, probability, and VC dimension of the problem. For simplicity in notation, we just consider the number of samples.

rescaled to become standard self-concordant (with parameter 2). Some of the well-known empirical loss functions which are self-concordant are linear regression, Logistic regression and squared hinge loss. In order to prove our results the following conditions are considered in our analysis.

Assumption 1. The loss functions $f(w, z)$ are convex w.r.t w for all values of z . In addition, their gradients $\nabla f(w, z)$ are M -Lipschitz continuous

$$\|\nabla f(w, z) - \nabla f(w', z)\| \leq M\|w - w'\|, \quad \forall z. \quad (6)$$

Assumption 2. The loss functions $f(w, z)$ are self-concordant w.r.t w for all values of z .

The immediate conclusion of Assumption 1 is that both $L(w)$ and $L_n(w)$ are convex and M -smooth. Also, we can note that $R_n(w)$ is cV_n -strongly convex and $(cV_n + M)$ -smooth. Moreover, by Assumption 2, $R_n(w)$ is also self-concordant.

3 Distributed Accumulated Newton Conjugate Gradient Method

The goal in inexact damped Newton method, as discussed in [Zhang and Lin, 2015], is to find the next iterate based on an approximated Newton-type update. It has two important differences comparing to Newton's method. First, as it is clear from the word "damped", the learning rate of the inexact damped Newton type update is not 1, since it depends on the approximation of Newton decrement. The second distinction is that there is no need to compute exact Newton direction (which is very expensive to calculate in one step). Alternatively, an approximated inexact Newton direction is calculated by applying an iterative process to obtain a direction with desirable accuracy under some measurements.

In order to utilize the important features of ERM, we combine the idea of increasing sample size and the inexact damped Newton method. In our proposed method, we start with handling a small number of samples, assume m_0 samples. We then solve its corresponding ERM to its statistical accuracy, i.e. V_{m_0} , using the inexact damped Newton algorithm. In the next step, we increase the number of samples geometrically with rate of $\alpha > 1$, i.e., αm_0 samples. The approximated solution of the previous ERM can be used as a warm start point to find the solution of the new ERM. The sample size increases until it equals the number of full samples.

Consider the iterate w_m within the statistical accuracy of the set with m samples, i.e. \mathcal{S}_m for the risk R_m . In DANCE, we increase the size of the training set to $n = \alpha m$ and use the inexact damped Newton to find

the iterate w_n which is V_n -suboptimal solution for the sample set \mathcal{S}_n , i.e. $R_n(w_n) - R_n(w_n^*) \leq V_n$ after K_n iterations. To do so, we initialize $\tilde{w}_0 = w_m$ and update the iterates according to the following

$$\tilde{w}_{k+1} = \tilde{w}_k - \frac{1}{1+\delta_n(\tilde{w}_k)} v_k, \quad (7)$$

where v_k is an ϵ_k -Newton direction. The outcome of applying (7) for $k = K_n$ iterations is the approximate solution w_n for the risk R_n , i.e., $w_n := \tilde{w}_{K_n}$.

To properly define the approximate Newton direction v_k , first consider that the gradient and Hessian of the risk R_n can be evaluated as

$$\nabla R_n(w) = \frac{1}{n} \sum_{i=1}^n \nabla f_i(w) + cV_n w \quad (8)$$

and

$$\nabla^2 R_n(w) = \frac{1}{n} \sum_{i=1}^n \nabla^2 f_i(w) + cV_n I, \quad (9)$$

respectively. The favorable descent direction would be the Newton direction $-\nabla^2 R_n(\tilde{w}_k)^{-1} \nabla R_n(\tilde{w}_k)$; however, the cost of computing this direction is prohibitive. Therefore, we use v_k which is an ϵ_k -Newton direction satisfying the condition

$$\|\nabla^2 R_n(\tilde{w}_k) v_k - \nabla R_n(\tilde{w}_k)\| \leq \epsilon_k. \quad (10)$$

As we use the descent direction v_k which is an approximation for the Newton step, we also redefine the Newton decrement $\delta_n(\tilde{w}_k)$ based on this modification. To be more specific, we define $\delta_n(\tilde{w}_k) := (v_k^T \nabla^2 R_n(\tilde{w}_k) v_k)^{1/2}$ as the approximation of (exact) Newton decrement $(\nabla R_n(\tilde{w}_k)^T \nabla^2 R_n(\tilde{w}_k)^{-1} \nabla R_n(\tilde{w}_k))^{1/2}$, and use it in the update in (7).

In order to find v_k which is an ϵ_k -Newton direction, we use Preconditioned CG (PCG). As it is discussed in [Zhang and Lin, 2015, Nocedal and Wright, 2006], PCG is an efficient iterative process to solve Newton system with the required accuracy. The preconditioned matrix that we considered is in the form of $P = \tilde{H}_n + \mu_n I$, where $\tilde{H}_n = \frac{1}{|\mathcal{A}_n|} \sum_{i \in \mathcal{A}_n} \nabla^2 R_n^i(w)$, $\mathcal{A}_n \subset \mathcal{S}_n$, and μ_n is a small regularization parameter. In this case, v_k is an approximate solution of the system $P^{-1} \nabla^2 R_n(\tilde{w}_k) v_k = P^{-1} \nabla R_n(\tilde{w}_k)$. The reason for using preconditioning is that the condition number of $P^{-1} \nabla^2 R_n(\tilde{w}_k)$ may be close to 1 in the case when \tilde{H}_n is close to $\nabla^2 R_n(\tilde{w}_k)$; consequently, PCG can be faster than CG. The PCG steps are summarized in Algorithm 2. In every iteration of Algorithm 2, a system needs to be solved in step 10. Due to the structure of matrix P , and as it is discussed in [Ma and Takáč, 2016], this matrix can be considered as $|\mathcal{A}_n|$ rank 1 updates on a diagonal matrix, and

Algorithm 1 DANCE

- 1: Initialization: Sample size increase constant α , initial sample size $n = m_0$ and $w_n = w_{m_0}$ with $\|\nabla R_n(w_n)\| < (\sqrt{2c})V_n$
 - 2: **while** $n \leq N$ **do**
 - 3: Update $w_m = w_n$ and $m = n$
 - 4: Increase sample size: $n = \min\{\alpha m, N\}$
 - 5: Set $\tilde{w}_0 = w_m$ and set $k = 0$
 - 6: **repeat**
 - 7: Calculate v_k and $\delta_n(\tilde{w}_k)$ by **Algorithm 2**
 - 8: **PCG**
 - 9: Set $\tilde{w}_{k+1} = \tilde{w}_k - \frac{1}{1+\delta_n(\tilde{w}_k)} v_k$
 - 10: $k = k + 1$
 - 11: **until** satisfy stop criteria leading to $R_n(\tilde{w}_k) - R_n(w_n^*) \leq V_n$
 - 12: Set $w_n = \tilde{w}_k$
 - 12: **end while**
-

now, using Woodbury Formula [Press et al., 2007] is a very efficient way to solve the mentioned system. The following lemma states the required number of iterations for PCG to find an ϵ_k -Newton direction v_k which is used in every stage of DANCE algorithm.

Lemma 1. (Lemma 4 in [Zhang and Lin, 2015]) *Suppose Assumption 1 holds and $\|\tilde{H}_n - \nabla^2 R_n(\tilde{w}_k)\| \leq \mu_n$. Then, Algorithm 2, after $C_n(\epsilon_k)$ iterations calculates v_k such that $\|\nabla^2 R_n(\tilde{w}_k) v_k - \nabla R_n(\tilde{w}_k)\| \leq \epsilon_k$, where*

$$C_n(\epsilon_k) = \left\lceil \sqrt{\left(1 + \frac{2\mu_n}{cV_n}\right) \log \left(\frac{2\sqrt{\frac{cV_n + M}{cV_n}} \|\nabla R_n(\tilde{w}_k)\|}{\epsilon_k} \right)} \right\rceil. \quad (11)$$

Note that ϵ_k has a crucial effect on the speed of the algorithm. When $\epsilon_k = 0$, then v_k is the exact Newton direction, and the update in (7) is the exact damped Newton step (which recovers the update in Ada Newton algorithm in [Mokhtari et al., 2016] when the step-length is 1). Furthermore, the number of total iterations to reach V_N -suboptimal solution for the risk R_N is \mathbf{K} , i.e. $\mathbf{K} = K_{m_0} + K_{\alpha m_0} + \dots + K_N$. Hence, if we start with the iterate w_{m_0} with corresponding m_0 samples, after \mathbf{K} iterations, we reach w_N with statistical accuracy of V_N for the whole dataset. In Theorem 1, the required rounds of communication to reach the mentioned statistical accuracy will be discussed.

Our proposed method is summarized in Algorithm 1. We start with m_0 samples, and an initial point w_{m_0} which is an V_{m_0} -suboptimal solution for the risk R_{m_0} . In every iteration of outer loop of Algorithm 1, we increase the sample size geometrically with rate of α in step 4. In the inner loop of Algorithm 1, i.e. steps 6-10, in order to calculate the approximate Newton direction and approximate Newton decrement, we use

Algorithm 2 PCG

1: Master Node: 2: Input: $\tilde{w}_k \in \mathbb{R}^d$, ϵ_k , and \mathcal{A}_n 3: Let $H = \nabla^2 R_n(\tilde{w}_k)$, $P = \frac{1}{ \mathcal{A}_n } \sum_{i \in \mathcal{A}_n} \nabla^2 R_n^i(\tilde{w}_k) + \mu_n I$ 4: Broadcast: \tilde{w}_k → 5: Reduce: $\nabla R_n^i(\tilde{w}_k)$ to $\nabla R_n(\tilde{w}_k)$ ← 6: Set $r^{(0)} = \nabla R_n(\tilde{w}_k)$, $u^{(0)} = s^{(0)} = P^{-1}r^{(0)}$ 7: Set $v^{(0)} = 0$, $t = 0$ 8: repeat 9: Broadcast: $u^{(t)}$ and $v^{(t)}$ → 10: Reduce: $\nabla^2 R_n^i(\tilde{w}_k)u^{(t)}$ and $\nabla^2 R_n^i(\tilde{w}_k)v^{(t)}$ to $Hu^{(t)}$ and $Hv^{(t)}$ ← 11: Compute $\gamma_t = \frac{\langle r^{(t)}, s^{(t)} \rangle}{\langle u^{(t)}, Hu^{(t)} \rangle}$ 12: Set $v^{(t+1)} = v^{(t)} + \gamma_t u^{(t)}$, $r^{(t+1)} = r^{(t)} - \gamma_t H u^{(t)}$ 13: Compute $\zeta_t = \frac{\langle r^{(t+1)}, s^{(t+1)} \rangle}{\langle r^{(t)}, s^{(t)} \rangle}$ 14: Set $P s^{(t+1)} = r^{(t+1)}$, $u^{(t+1)} = s^{(t+1)} + \zeta_t u^{(t)}$ 15: Set $t = t + 1$ 16: until $\ r^{(t+1)}\ \leq \epsilon_k$ 17: Output: $v_k = v^{(t+1)}$, $\delta_n(\tilde{w}_k) = \sqrt{v_k^T H v^{(t)} + \gamma_t v_k^T H u^{(t)}}$	Worker Nodes ($i = 1, 2, \dots, \mathcal{K}$): Compute $\nabla R_n^i(\tilde{w}_k)$ Compute $\nabla^2 R_n^i(\tilde{w}_k)u^{(t)}$ and $\nabla^2 R_n^i(\tilde{w}_k)v^{(t)}$
---	--

PCG algorithm which is shown in Algorithm 2. This process repeats till we get the point w_N with statistical accuracy of V_N . The practical stopping criteria for Algorithm 1 is discussed in Section A.1.

Distributed Implementation Similar to the algorithm in [Zhang and Lin, 2015], Algorithms 1 and 2 can also be implemented in a distributed environment. Suppose the entire dataset is stored across \mathcal{K} machines, i.e., each machine stores N_i data samples such that $\sum_{i=1}^{\mathcal{K}} N_i = N$. Under this setting, each iteration in Algorithm 1 can be executed on different machines in parallel with $\sum_{i=1}^{\mathcal{K}} n_i = n$, where n_i is the batchsize on i^{th} machine. To implement Algorithm 2 in a distributed manner, a broadcast operation is needed at each iteration to guarantee that each machine will share the same \tilde{w}_k value. Moreover, the gradient and Hessian-vector product can be computed locally and later reduce to the master machine. With the increasing of batch size, computation work on each machine will increase while we still have the same amount of communication need. As a consequence, the computation expense will gradually dominate the communication expense before the algorithm terminates. Therefore the proposed algorithm could take advantage of utilizing more machines to shorten the running time of Algorithm 2.

4 Complexity Analysis

In this section, we study the convergence properties of our algorithm. To do so, we analyze the required number of communication rounds and total computational complexity of DANCE to solve every subproblem up

to its statistical accuracy.

We analyze the case when we have w_m which is a V_m -suboptimal solution of the risk R_m , and we are interested in deriving a bound for the number of required communication rounds to ensure that w_n is a V_n -suboptimal solution for the risk R_n .

Theorem 1. *Suppose that Assumptions 1 and 2 hold. Consider w_m which satisfies $R_m(w_m) - R_m(w_m^*) \leq V_m$ and also the risk R_n corresponding to sample set $\mathcal{S}_n \supset \mathcal{S}_m$ where $n = \alpha m$, $\alpha > 1$. Set the parameter ϵ_k (the error in (10)) as following*

$$\epsilon_k = \beta \left(\frac{cV_n}{M+cV_n} \right)^{1/2} \|\nabla R_n(\tilde{w}_k)\|, \quad (12)$$

where $\beta \leq \frac{1}{20}$. Then, in order to find the variable w_n which is an V_n -suboptimal solution for the risk R_n , i.e. $R_n(w_n) - R_n(w_n^*) \leq V_n$, the number of communication rounds T_n satisfies in the following:

$$T_n \leq K_n (1 + C_n(\epsilon_k)), \quad w.h.p. \quad (13)$$

where $K_n = \left\lceil \frac{R_n(w_m) - R_n(w_n^*)}{\frac{1}{2}\omega(1/6)} \right\rceil + \left\lceil \log_2 \left(\frac{2\omega(1/6)}{V_n} \right) \right\rceil$. Here $\lceil t \rceil$ shows the smallest nonnegative integer larger than or equal to t .

As a result, the update in (7) needs to be done for $K_n = \mathcal{O}(\log_2 n)$ times in order to attain the solution w_n which is V_n -suboptimal solution for the risk R_n . Also, based on the result in (13), by considering the risk R_n , we can note that when the strong-convexity parameter for the mentioned risk (cV_n) is large, less number of iterations (communication rounds) are needed (or

equally faster convergence is achieved) to reach the iterate with V_n -suboptimal solution; and this happens in the first steps. More importantly, DANCE is a multi-stage algorithm which any stage of DANCE converges linearly (the reason is that in every stage of our DANCE method, we use inexact damped-Newton method, and as it is shown in [Zhang and Lin, 2015] that with ϵ_k in (12), the inexact damped-Newton method converges linearly). Furthermore, if we consider DANCE as 1-stage algorithm with full samples in each iteration, it covers DiSCO which converges linearly.

Corollary 1. *Suppose that Assumptions 1 and 2 hold. Further, assume that w_m is a V_m -suboptimal solution for the risk R_m and consider R_n as the risk corresponding to sample set $\mathcal{S}_n \supset \mathcal{S}_m$ where $n = 2m$. If we set parameter ϵ_k (the error in (10)) as (12), then with high probability T_n communication rounds*

$$\tilde{T}_n \leq \left(\left\lceil \frac{\left(3 + \left(1 - \frac{1}{2^\gamma}\right)\left(2 + \frac{\epsilon}{2}\|w^*\|^2\right)\right)V_m}{\frac{1}{2}\omega(1/6)} \right\rceil + \left\lceil \log_2\left(\frac{2\omega(1/6)}{V_n}\right) \right\rceil \right) \left(1 + \left\lceil \sqrt{1 + \frac{2\mu}{cV_n}} \log_2\left(\frac{2(cV_n + M)}{\beta cV_n}\right) \right\rceil \right), \quad (14)$$

are needed to reach the point w_n with statistical accuracy of V_n for the risk R_n .

Corollary 2. *By assuming $\alpha = 2$, the total number of communication rounds to reach a point with the statistical accuracy of V_N of the full training set is w.h.p.*

$$\tilde{T} = \tilde{O}(\gamma(\log_2 N)^2 \sqrt{N^\gamma} \log_2 N^\gamma) \quad (15)$$

where $\gamma \in [0.5, 1]$.

The rounds of communication for DiSCO in [Zhang and Lin, 2015]² is $\tilde{T}_{DiSCO} = \tilde{O}((R_N(w_0) - R_N(w_N^*) + \gamma(\log_2 N))\sqrt{N^\gamma} \log_2 N^\gamma)$ where $\gamma \in [0.5, 1]$. Comparing these bounds shows that the communication complexity of DANCE is independent of the choice of initial variable w_0 and the suboptimality $R_N(w_0) - R_N(w_N^*)$, while the overall communication complexity of DiSCO depends on the initial suboptimality. In addition, implementation of each iteration of DiSCO requires processing all the samples in the dataset, while DANCE only operates on an increasing subset of samples at each phase. Therefore, the computation complexity of DANCE is also lower than DiSCO for achieving the statistical accuracy of the training set.

Theorem 2. *The total complexity of DANCE Algorithm in order to reach a point with the statistical accuracy of V_N of the full training set is w.h.p.*

$$\tilde{O}((\log_2(N))^3 N^{1/4} d^2) \quad (16)$$

²In order to have fair comparison, we put $f = R_N$, $\epsilon = V_N$, and $\lambda = cV_N$ in their analysis, and also the constants are ignored for the communication complexity.

Table 1 shows that the total complexity of the k -Truncated method [Eisen et al., 2018] is lower than the one for AdaNewton [Mokhtari et al., 2016]. Further, as $(\log_2(N))^3 \ll N^{3/4}$, the total complexity of DANCE is lower than both AdaNewton and k -TAN methods. All in all, the theoretical results highlight that DANCE is more efficient than DiSCO in terms of communication, and has lower complexity than previous adaptive sample size methods including AdaNewton and k -TAN.

5 Numerical Experiments

In this section, we present numerical experiments on several large real-world datasets to show that our restarting DANCE algorithm can outperform other existed methods on solving both convex and non-convex problems. Also, we compare the results obtained from utilizing different number of machines to demonstrate the strong scaling property for DANCE. All the algorithms are implemented in Python with PyTorch [Paszke et al., 2017] library and we use MPI for Python [Dalcin et al., 2011] distributed environment³. For all plots in this section, vertical pink dashed lines represent restarts in our DANCE method.

Convex problems. First, we compare DANCE with two algorithms SGD (mini-batch)⁴ and DiSCO [Zhang and Lin, 2015], for solving convex problems. The experiments in this section are performed on a cluster with 16 Xeon E5-2620 CPUs (2.40GHz).

We use logistic regression model for two binary classification tasks based on *rcv1* and *giset* [Chang and Lin, 2011] datasets for our convex test problem. We use logistic loss function defined as $f_i(w) := \log(1 + \exp(-y_i w^T x_i))$, where $x_i \in \mathbb{R}^d$ is data sample and $y_i \in \{-1, 1\}$ is binary label corresponding to $x_i, i \in [m]$. Then we minimize the empirical loss function as (4). Note that there is a fixed ℓ_2 -regularization parameter in DiSCO and SGD and we set $c = 0.1$ in (4) to form the ℓ_2 -regularization parameter for our DANCE method. In Section C.1, we numerically show that DANCE is robust w.r.t. different setting of hyper-parameters⁵.

We run our algorithm and compare algorithms with different datasets using 8 nodes. The starting batchsize on each node for our DANCE algorithm is set to 16 while other two algorithms go over the whole dataset at each iteration. For DANCE implementation, number of samples used to form the new ERM loss are doubled from previous iteration after each restarting.

³All codes to reproduce these experimental results are available at <https://github.com/OptMLGroup/DANCE>.

⁴The batch size is 10 in our experiments

⁵Please see Figures 6, 7, 8 and 9




Figure 1: Performance of different algorithms on a logistic regression problem with rcv1 and gisette datasets. In the left two figures, the plot $DANCE^*$ is the training accuracy based on the entire training set, while the plot $DANCE$ represents the training accuracy based on the current sample size.

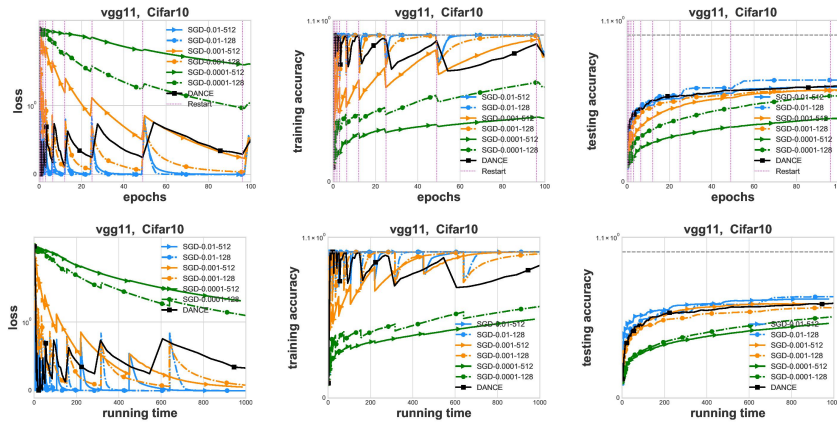


Figure 2: Comparison between DANCE and SGD with various hyper-parameters setting on Cifar10 dataset and vgg11 network. vgg11 represents simonyan2014very a 28 layers convolutional neural network (see details at Appendix B). Figures on the top and bottom show how loss values, training accuracy and test accuracy are changing with respect to epochs and running time. Note that we force both algorithms to restart (double training sample size) after achieving the following number of epochs: 0.2, 0.8, 1.6, 3.2, 6.4, 12, 24, 48, 96. For SGD, we varies learning rate from 0.01, 0.001, 0.0001 and batchsize from 128, 512.

In Figure 1, we observe consistently that DANCE has a better performance over the other two methods from the beginning stages (please see Figure 5 for additional results). Both training and test accuracy for DANCE converges to optimality after processing a small number of samples. This observation suggests that DANCE finds a good initial solution and updates it over time. Compared with DiSCO, our restarting approach helps to reduce computational cost for the first iterations, where the second order methods usually performs less efficiently comparing to first order methods. The key difference comes from utilizing the idea of increasing sample size where DANCE goes over small number of samples and finds a suboptimal solution, and use it as a warm-start for the next stage. In this way, less passes over data is needed in the beginning but with satisfactory accuracy. On the other hand, DiSCO uses total number of samples from the beginning which some passes over data is needed in order to reach the neighborhood of global solution. Therefore, DANCE behaves efficiently and reaches the optimal solution with less passes over data.

Non-convex problems. Even though the complexity analysis in Section 4 only covers the convex case, the DANCE algorithm is also able to handle nonconvex problems efficiently. In this section, we compare our method with several stochastic first order algorithms, stochastic gradient descent (SGD), SGD with momentum (SGDMom), and Adam [Kingma and Ba, 2014], on training convolution neural networks (CNNs) on two image classification datasets *Mnist* and *Cifar10*. The details of the datasets and the CNNs architecture applied on each dataset are presented in Appendix B. To perform a fair comparison with respect to first order variants, we assume data comes in an online streaming manner, e.g., only a few data samples can be accessed at the beginning, and new data samples arrives at a fixed rate. Such setting is common in industrial production, where business data is collected in a streaming fashion. We feed new data points to all algorithms only if the amount of new samples is equal to the number of existing samples. The experiments in this section are run on an *AWS p2.xlarge instance with an NVIDIA K80 GPU*.

In Figure 2, we compare DANCE with the build-in

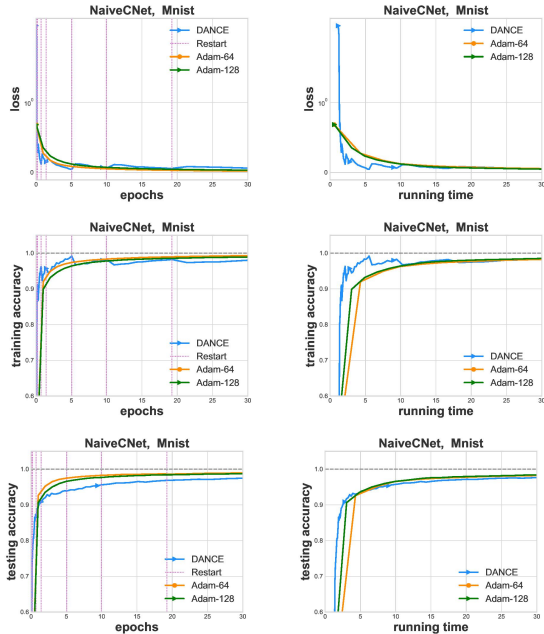


Figure 3: Comparison between DANCE and Adam on Mnist dataset and NaiveCNet. For DANCE, the initial batchsize is 1024. For Adam, the learning rate is 10^{-4} and the batchsize is either 64 or 128.


SGD optimizer in pyTorch on Cifar dataset to train a 28 layers CNN (Vgg11) architecture. Note that there are several hyper-parameters we need to tune for SGD to reach the best performance, such as batch size and learning rate, which are not required for DANCE. Since we have the online streaming data setting, we don't need to determine a restarting criterion. The results show that SGD is sensitive to hyper-parameters tuning, i.e., different combination of hyper-parameters affect the performance of SGD a lot and tune them well to achieve the best performance could be painful. However, our DANCE algorithm does not have such weakness and its performance is comparable to SGD with the best parameters setting. We also show that the DANCE algorithm leads to a faster decreasing on the loss value, which is similar to our convex experiments. Again, this is due to fast convergence rate of the second order methods. One could also found the additional experiments regarding the comparison with SGD with momentum and Adam in terms of Mnist with NaiveCNet at Appendix C.

Regarding Figure 3, the performance of build-in Adam optimizer and our DANCE algorithm are compared regarding Mnist dataset and a 4 layer NaiveCNet (see the details in Appendix B). In this experiment, we do not assume that the data samples follow an online streaming manner for Adam, i.e., the Adam algorithm does not have a restarting setting and therefore it runs on whole dataset directly. Also, this experiment is performed only on CPUs. We set the learning-rate for

Adam as 10^{-4} and varies the running batch-size from 64 and 128. The evolution of loss, training accuracy, testing accuracy with respect to epochs and running time regarding the whole dataset are reported in Figure 3 for different algorithms. One could observe that under the same epochs, Adam eventually achieves the better testing accuracy, while if we look at running time, our DANCE algorithm would be faster due to the distributed implementation.

Strong scaling Moreover, we demonstrate that our DANCE method shares a strong scaling property. As shown in Figure 4, whenever we increase the number of nodes, we obtain acceleration towards optimality.

We use the starting batchsize from 256 upto 4096, and the speed-up compared to the serial run (1 node) is reported. It indicates that as we increase the batchsize, the speed-up becomes



closer to ideal linear speed-up. The advantage of the setting is to utilize the large batch over multiple nodes efficiently but not sacrifice the convergence performance. Regarding first order methods like SGD, it is hard to achieve nice linear scaling since the small batch is often required, which makes the computation time to be comparable with communication cost.

Figure 4: Performance of DANCE w.r.t. different number of nodes.

6 Conclusion

We proposed DANCE an efficient distributed Hessian free algorithm with an increasing sample size strategy to solve the empirical risk minimization problem. Our algorithm obtains a solution within the statistical accuracy of the ERM problem in very few epochs and also can be implemented in a distributed environment. We analyzed the communication-efficiency of DANCE and highlighted its efficiency with respect to DiSCO [Zhang and Lin, 2015] in term of communication and relative to AdaNewton and k -TAN methods in terms of total computational complexity. The presented numerical experiments demonstrated the fast convergence of DANCE for both convex and non-convex problems.

7 Acknowledgement

This work was supported by U.S. National Science Foundation, under award number NSF:CCF:1618717, NSF:CMMI:1663256 and NSF:CCF:1740796.

References

- [Bartlett et al., 2006] Bartlett, P. L., Jordan, M. I., and McAuliffe, J. D. (2006). Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473):138–156.
- [Beck and Teboulle, 2009] Beck, A. and Teboulle, M. (2009). A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences*, 2(1):183–202.
- [Berahas et al.,] Berahas, A. S., Jahani, M., and Takáč, M. Sampled quasi-newton methods for deep learning.
- [Berahas et al., 2019] Berahas, A. S., Jahani, M., and Takáč, M. (2019). Quasi-newton methods for deep learning: Forget the past, just sample. *arXiv preprint arXiv:1901.09997*.
- [Berahas et al., 2016] Berahas, A. S., Nocedal, J., and Takáč, M. (2016). A multi-batch l-bfgs method for machine learning. In *Advances in Neural Information Processing Systems*, pages 1055–1063.
- [Bottou, 2010] Bottou, L. (2010). Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT’2010*, pages 177–186. Springer.
- [Bousquet, 2002] Bousquet, o. (2002). *Concentration Inequalities and Empirical Processes Theory Applied to the Analysis of Learning Algorithms*. PhD thesis, Biologische Kybernetik.
- [Bousquet and Bottou, 2008] Bousquet, O. and Bottou, L. (2008). The tradeoffs of large scale learning. In *Advances in Neural Information Processing Systems 20*, pages 161–168.
- [Boyd and Vandenberghe, 2004] Boyd, S. and Vandenberghe, L. (2004). *Convex optimization*. Cambridge university press.
- [Byrd et al., 2016] Byrd, R. H., Hansen, S. L., Nocedal, J., and Singer, Y. (2016). A stochastic quasi-Newton method for large-scale optimization. *SIAM Journal on Optimization*, 26(2):1008–1031.
- [Chang and Lin, 2011] Chang, C.-C. and Lin, C.-J. (2011). LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- [Dalcin et al., 2011] Dalcin, L. D., Paz, R. R., Kler, P. A., and Cosimo, A. (2011). Parallel distributed computing using python. *Advances in Water Resources*, 34(9):1124–1139.
- [Defazio et al., 2014] Defazio, A., Bach, F., and Lacoste-Julien, S. (2014). Saga: A fast incremental gradient method with support for non-strongly convex composite objectives. In *Advances in Neural Information Processing Systems 27*, pages 1646–1654.
- [Drusvyatskiy et al., 2018] Drusvyatskiy, D., Fazel, M., and Roy, S. (2018). An optimal first order method based on optimal quadratic averaging. *SIAM Journal on Optimization*, 28(1):251–271.
- [Eisen et al., 2018] Eisen, M., Mokhtari, A., and Ribeiro, A. (2018). Large scale empirical risk minimization via truncated adaptive Newton method. In *International Conference on Artificial Intelligence and Statistics*, pages 1447–1455.
- [Jahani et al., 2019] Jahani, M., Nazari, M., Rusakov, S., Berahas, A. S., and Takáč, M. (2019). Scaling up quasi-newton algorithms: Communication efficient distributed sr1. *arXiv preprint arXiv:1905.13096*.
- [Johnson and Zhang, 2013] Johnson, R. and Zhang, T. (2013). Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in Neural Information Processing Systems 26*, pages 315–323.
- [Kingma and Ba, 2014] Kingma, D. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- [Konečný and Richtárik, 2017] Konečný, J. and Richtárik, P. (2017). Semi-stochastic gradient descent methods. *Frontiers in Applied Mathematics and Statistics*, 3:9.
- [Ma et al., 2017] Ma, C., Gudapati, N. V. C., Jahani, M., Tappenden, R., and Takáč, M. (2017). Underestimate sequences via quadratic averaging. *arXiv preprint arXiv:1710.03695*.
- [Ma and Takáč, 2016] Ma, C. and Takáč, M. (2016). Distributed inexact damped Newton method: Data partitioning and load-balancing. *arXiv preprint arXiv:1603.05191*.
- [Mokhtari et al., 2016] Mokhtari, A., Daneshmand, H., Lucchi, A., Hofmann, T., and Ribeiro, A. (2016). Adaptive Newton method for empirical risk minimization to statistical accuracy. In *Advances in Neural Information Processing Systems 29*, pages 4062–4070.
- [Mokhtari and Ribeiro, 2015] Mokhtari, A. and Ribeiro, A. (2015). Global convergence of online limited memory bfgs. *Journal of Machine Learning Research*, 16(1):3151–3181.
- [Mokhtari and Ribeiro, 2017] Mokhtari, A. and Ribeiro, A. (2017). First-order adaptive sample size methods to reduce complexity of empirical risk minimization. In *Advances in Neural Information Processing Systems 30*, pages 2057–2065.

- [Nesterov, 2013] Nesterov, Y. (2013). *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media.
- [Nguyen et al., 2017] Nguyen, L. M., Liu, J., Scheinberg, K., and Takáč, M. (2017). Sarah: A novel method for machine learning problems using stochastic recursive gradient. In *Proceedings of the 34th International Conference on Machine Learning*, pages 2613–2621.
- [Nocedal and Wright, 2006] Nocedal, J. and Wright, S. J. (2006). *Sequential quadratic programming*. Springer.
- [Paszke et al., 2017] Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., and Lerer, A. (2017). Automatic differentiation in pytorch.
- [Press et al., 2007] Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P. (2007). *Numerical recipes: the art of scientific computing, 3rd Edition*. Cambridge University Press.
- [Roosta-Khorasani and Mahoney, 2018] Roosta-Khorasani, F. and Mahoney, M. W. (2018). Sub-sampled newton methods. *Mathematical Programming*.
- [Roux et al., 2012] Roux, N. L., Schmidt, M., and Bach, F. R. (2012). A stochastic gradient method with an exponential convergence rate for finite training sets. In *Advances in Neural Information Processing Systems 25*, pages 2663–2671.
- [Schraudolph et al., 2007] Schraudolph, N. N., Yu, J., and Günter, S. (2007). A stochastic quasi-Newton method for online convex optimization. In *Artificial Intelligence and Statistics*, pages 436–443.
- [Shalev-Shwartz and Zhang, 2013] Shalev-Shwartz, S. and Zhang, T. (2013). Stochastic dual coordinate ascent methods for regularized loss minimization. *Journal of Machine Learning Research*.
- [Vapnik, 2013] Vapnik, V. (2013). *The nature of statistical learning theory*. Springer science & business media.
- [Zhang and Lin, 2015] Zhang, Y. and Lin, X. (2015). Disco: Distributed optimization for self-concordant empirical loss. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 362–370.