

000
001
002
003
004
005
006
007
008
009
010
011
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031
032
033
034
035
036
037
038
039
040
041
042
043
044
045
046
047
048
049
050
051
052
053

A Parallel SGD method with Strong Convergence

Anonymous Author(s)

Affiliation

Address

email

Abstract

This paper proposes a novel parallel stochastic gradient descent (SGD) method that is obtained by applying parallel sets of SGD iterations (each set operating on one node using the data residing in it) for finding the direction in each iteration of a batch descent method. The method has strong convergence properties. Experiments on datasets with high dimensional feature spaces show the value of this method.

Introduction. We are interested in the large scale learning of linear classifiers. Let $\{x_i, y_i\}$ be the training set associated with a binary classification problem ($y_i \in \{1, -1\}$). Consider a linear classification model, $y = \text{sgn}(w^T x)$. Let $l(w \cdot x_i, y_i)$ be a continuously differentiable, non-negative, convex loss function that has Lipschitz continuous gradient. This allows us to consider loss functions such as least squares, logistic loss and squared hinge loss. Hinge loss is not covered by our theory since it is non-differentiable. Our aim is to minimize the regularized risk functional $f(w) = \frac{\lambda}{2} \|w\|^2 + L(w)$ where $\lambda > 0$ is the regularization constant and $L(w) = \sum_i l(w \cdot x_i, y_i)$ is the total loss. The gradient function, $g = \nabla f$ is Lipschitz continuous.

For large scale learning on a single machine, it is now well established that example-wise methods¹ such as stochastic gradient descent (SGD) and its variations [1, 2, 3] and dual coordinate ascent [4] are much faster than batch gradient-based methods for reaching weights with sufficient training optimality needed for attaining steady state generalization performance. However, example-wise methods are inherently sequential.

For tackling problems involving huge sized data, distributed solution becomes necessary. One approach to parallel SGD solution [5] is via (iterative) parameter mixing [6, 7]. Consider a distributed setting with a master-slave architecture² in which the examples are partitioned over P slave computing nodes. Let: I_p be the set of indices i such that (x_i, y_i) sits in the p -th node; and $L_p(w) = \sum_{i \in I_p} l(w \cdot x_i, y_i)$ be the total loss associated with node p . Thus, $f(w) = \frac{\lambda}{2} \|w\|^2 + \sum_p L_p(w)$. Suppose the master node has the current weight vector w^r and it communicates it to all the nodes. Each node p can form the approximation,

$$\tilde{f}_p = \frac{\lambda}{2} \|w\|^2 + L_p(w) \tag{1}$$

of f using only its examples, and do several SGD epochs (local passes over its examples) on \tilde{f}_p and reach a point w_p . The $w_p \forall p$ can then be communicated back to the master node and averaged to form the next iterate w^{r+1} . One can stop after just one major iteration (go from $r = 0$ to $r = 1$) or repeat many such major iterations. Convergence theory for such methods is limited, and, even that requires a complicated analysis [5]. There are two main issues related to variance and bias: (a) When the number of nodes is large, the \tilde{f}_p are very different from each other, and so, the variability in the w_p is large and the averaged weight vector is also far away from $w^* = \arg \min f(w)$. (b)

¹These methods update w after scanning each example.

²An *AllReduce* arrangement of nodes [8] may also be used.

If we use too many SGD epochs within each node p , then, within each major iteration, SGD will converge to the minimizer of \tilde{f}_p irrespective of the starting point w^r , making the major iterations useless.

A New Parallel SGD method. Our main idea is to use a descent method for batch training and, in each of its iterations, compute the direction by doing SGD iterations in parallel on function approximations that are “better” than the \tilde{f}_p . We begin each iteration r by computing the gradient g^r at the current point w^r .³ One can communicate w^r and g^r to all P (slave) nodes. The direction d^r is formed as follows. Each node p constructs an approximation of $f(w)$ using only information that is available in that node⁴ (call this function as $\hat{f}_p(w)$) and (approximately) optimizes it (starting from w^r) to get the point w_p . Let $d_p = w_p - w^r$. Then d^r is chosen to be any convex combination of $d_p \forall p$. Note that, if each d_p is a descent direction then d^r is also a descent direction.

The key is to choose each approximating functional \hat{f}_p to have *gradient consistency* at w^r :

$$\hat{f}_p(w) = \tilde{f}_p(w) + (g^r - \lambda w^r - \nabla L_p(w^r)) \cdot (w - w^r) \quad (2)$$

Note that $\nabla \hat{f}_p(w^r) = g^r$. This condition gives the necessary tilt to the approximating functions for maintaining consistency with the minimization of the global objective function f .

Algorithm 1 gives all the steps of our method. There, $sgd(v^0; \hat{f}_p, s, pars)$ denotes the output point obtained by applying s epochs of a SGD method to \hat{f}_p , starting from v^0 , where $pars$ denotes parameters associated with the SGD method, e.g., learning rate. Note that the output point is stochastic since there is randomness present in the iterations, e.g., the order in which examples are presented.

Algorithm 1: Distributed method for minimizing f (*com*: communication; *cmp*: = computation; *agg*: aggregation)

Choose w^0 ; $s, pars, 0 \leq \theta < \frac{\pi}{2}$;

for $r = 0, 1 \dots$ **do**

1. Compute g^r (*com*: w^r ; *cmp*: Two passes over data; *agg*: g^r); By-product: $\{z_i = w^r \cdot x_i\}$;
2. Exit if $g^r = 0$;
3. **for** $p = 1, \dots, P$ (*in parallel*) **do**
 4. Set $v^0 = w^r$;
 5. Set $w_p = sgd(v^0; \hat{f}_p, s, pars)$;
 6. If $\angle -g^r, d_p \geq \theta$, set $d_p = -g^r$; ($\angle a, b$ denotes the angle between vectors a and b);

end

7. Set d^r as any convex combination of $\{d_p\}$ (*agg*: d_p);
8. Do line search to find t (for each t : *comm*: t ; *cmp*: l and $\partial l / \partial t$; *agg*: $f(w^r + td^r)$ and its derivative wrt t);
9. Set $w^{r+1} = w^r + td^r$;

end

For step 8 we require that the following standard line search conditions are satisfied:

$$\textbf{Armijo:} \quad f^{r+1} \leq f^r + \alpha g^r \cdot (w^{r+1} - w^r) \quad (3)$$

$$\textbf{Wolfe:} \quad g^{r+1} \cdot d^r \geq \beta g^r \cdot d^r \quad (4)$$

where $0 < \alpha < \beta < 1$ and $f^r = f(w^r)$.

Theorem 1. Algorithm 1 has global linear rate of convergence (*glrc*), i.e., $\exists 0 < \delta < 1$ such that $(f(w^{r+1}) - f(w^*)) \leq \delta (f(w^r) - f(w^*)) \forall r$, where $w^* = \arg \min_w f(w)$. It follows that algorithm 1 finds a point w^r satisfying $f(w^r) - f(w^*) \leq \epsilon$ in $O(s \log(1/\epsilon))$ time.

Theorem 1 can actually be stated stronger than what is given above. If steps 4-6 of algorithm 1 are replaced by any sub-algorithm that finds a d_p satisfying $\angle -g^r, d_p < \theta$ then the *glrc* result

³It is worth noting that recently proposed powerful SGD methods [3] also compute full batch gradients once every few SGD epochs.

⁴This information includes the examples sitting in the node as well as w^r, g^r etc.

108 still holds. While convergence follows from standard optimization theory, proving *glrc* under such
 109 general conditions seems to be a new result. Previously, *glrc* seems to have been established only
 110 for special cases such as the gradient descent method [9].

111 If the condition $\langle -g^r, d_p \rangle \geq \theta$ never gets triggered in step 6, then algorithm 1 can be viewed as a
 112 clean parallel SGD method. Step 6 is a “safe” artifact step that is added to account for the stochas-
 113 ticity of (*sgd*). One may ask: *how frequently does the condition $\langle -g^r, d_p \rangle \geq \theta$ happen?*. This can be
 114 answered both, from a theoretical as well as a practical angle. Theoretically, we can show that, by
 115 making s , the number of *sgd* epochs in step 5 large, the probability of the condition happening can
 116 be made arbitrarily small.

117 **Theorem 2.** Let $\hat{w}_p^* = \arg \min_w \hat{f}_p(w)$. Suppose *sgd* has strong stochastic convergence in the sense
 118 that $\mathbb{E} \|w_p - \hat{w}_p^*\|^2 \leq K\alpha^s \|w^r - \hat{w}_p^*\|^2 \forall s$, where $0 \leq \alpha < 1$ and $K \geq 0$. Let $\frac{\pi}{2} > \theta > \cos^{-1} \frac{\lambda}{L}$
 119 where L is a Lipschitz constant for g . Then, for any $0 < \gamma < 1$, $\exists s = O(\log(1/\gamma))$ such that
 120 $\text{Prob}(\langle -g^r, d^r \rangle \geq \theta) < \gamma$.

121 Theorems 1 and 2 can be combined to give algorithmic time complexity as $O(\log(1/(\epsilon\gamma)))$. The
 122 proof is based on formalizing the following observations. (1) Since \hat{f}_p has a ‘curvature’ of at least
 123 λ and $\nabla \hat{f}_p(w^r) = g^r$, $\hat{w}_p^* - w^r$ makes an angle less than $\cos^{-1} \frac{\lambda}{L}$ with $-g^r$. (2) Strong stochastic
 124 convergence of *sgd* implies that, as s is made large, w_p comes close to \hat{w}_p^* with high probability.
 125 Recent SGD methods [3, 2] possess the strong convergence property needed in Theorem 2.

126 Let us now discuss the practical view of step 6. For *sgd*, suppose we use an SGD method such as
 127 the ones in [3, 2] which move in directions that are “close” to the negative batch gradient and also
 128 have low variance. Then, the points generated by *sgd* lie close to the batch negative gradient flow
 129 and hence nicely lead to $w_p - w^r$ being a descent direction. This actually happens even when s
 130 is small. With large s , of course, w_p moves close to \hat{w}_p^* as explained earlier. Thus, the chance of
 131 $\hat{f}_p(w_p) \geq \hat{f}_p(w^r)$ happening is very low, irrespective of s . Note from the definition of \hat{f}_p that,
 132 $w_p - w^r$ is a descent direction of f at w^r if and only if $\hat{f}_p(w_p) < \hat{f}_p(w^r)$.

133 For a practical implementation, we can simply set $\theta = 0$, going by standard practice in numerical
 134 optimization; thus, directions that lead to descent are accepted. Parameters (*pars*) of *sgd* can be set
 135 as recommended by the individual SGD method used. The number of epochs s can be set based on
 136 a communication-computation trade-off; it is wise to choose s so that the cost of communication
 137 (of w^r , g^r etc.) between nodes is reasonably commensurate with the cost of computation (*sgd*
 138 epochs) in each node. For step 7, one can use simple averaging. For the line search step we can
 139 first calculate $w \cdot x_i$ (note the by-product in step 1) and $d^r \cdot x_i$ for all i in a distributed fashion.
 140 Then the calculation of $f(w^r + td^r)$ and its derivative with respect to t is cheap; so one can use any
 141 good one dimensional search algorithm to find a point satisfying the Armijo-Wolfe conditions. The
 142 parameters in these conditions can be set to: $\alpha = 10^{-4}$ and $\beta = 0.9$.

143 **Experiments.** We use an AllReduce tree running on a Hadoop cluster [8]. We
 144 use the Area under Precision-Recall Curve (AUPRC) and $(f - f^*)/f^*$ (in log scale)
 145 as the evaluation criteria. (We obtained f^* by optimizing with very small tolerances
 146 to get a very accurate solution.) Experiments are conducted on the *kdd2010* dataset
 147 in <http://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/binary.html>. This
 148 dataset has 8.41 million examples, 20.21 million examples and 0.3 billion non-zero elements in
 149 the data matrix. We use squared hinge loss with L_2 regularization. For *sgd* we use the SVRG
 150 method [3]. Let $FS - s$ denote our method with s being the number of *sgd* epochs.

151 We compare our method against *SQM* (Statistical Query Model) [10, 8], which is currently one of the
 152 most effective distributed methods. *SQM* is a batch, gradient-based descent method. The gradient is
 153 computed in a distributed way with each node computing the gradient component corresponding to
 154 its set of examples, followed by an aggregation of the components via an AllReduce tree. *Hybrid* is
 155 same as *SQM*, but uses parameter mixing for initializing. Each node p does one epoch of SGD [1]
 156 on its examples; then the weights from various nodes are averaged to form a weight vector that is
 157 used to initialize *SQM*. Our implementations of *SQM* and *Hybrid* are close to that in [8]; the main
 158 difference is that, instead of L-BFGS we use the better-performing TRON [11] as the core optimizer.
 159

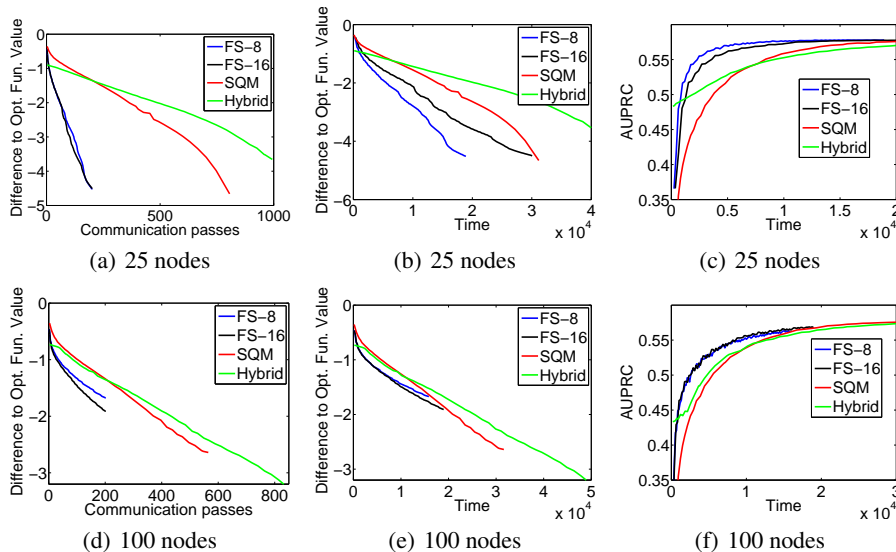


Figure 1: Comparison of methods. Left and middle: $(f - f^*)/f^*$ (log scale) versus Number of communication passes and Time (in seconds); Right: AUPRC as a function of Time, for *kdd2010* with 25 nodes (top) and 100 nodes (bottom).

Figure 1 compares our method against *SQM* and *Hybrid*. First let us look at the variation of the objective function accuracy with respect to the number of communication passes.⁵ Looking at the plots on the left side of Figure 1 we can see that our method requires far less number of communication passes than *SQM* and *Hybrid* to achieve the same accuracy in objective function. This difference between the methods also extends to the plots of objective function accuracy versus computing time (middle sub-plots in Figure 1); however, the difference in performance is less pronounced. This is due to the increased computational work (*sgd* epochs) done by our method at the nodes; note that *SQM* and *Hybrid* use the nodes only to compute gradient components.

SQM and *Hybrid* also have the advantage of better convergence when coming close to the optimum since they are directly based on second order modeling of f . Our method is good at forming approximate global views of f right from the beginning, thus making good progress in the early iterations. This is also directly reflected in the AUPRC plots (the right side of Figure 1). Our method reaches stable generalization performance much quicker than *SQM* and *Hybrid*.

When the number of nodes is increased, *SQM* and *Hybrid* come closer to our method; see this by comparing the sub-plots in the top and bottom of Figure 1. This is due to the fact that, when the number of nodes becomes large, \hat{f}_p does not approximate f very well, leading to an increased number of major iterations. The value of s , the number of SGD epochs plays a key role in determining the rate of linear convergence.

Conclusion and discussion. We have given a parallelization of stochastic gradient descent with strong convergence properties and demonstrated its effectiveness.

We can extend our method in several ways. (a) Suppose f is non-convex, e.g., f from neural networks and deep learning. In the definition of \hat{f}_p (see (2) and (1)), if $L_p(w)$ is replaced by a convex approximation then convergence of algorithm 1 can be shown; however, complexity results (which require *glrc*) are not easy to prove. From a practical point of view, it also makes sense to try non-convex \hat{f}_p , but care is needed to stop the optimization of \hat{f}_p via *sgd* early to make sure that the d_p are descent directions. (b) For convex f , our method can also use other algorithms (e.g., L-BFGS, TRON etc.) as a replacement for *sgd* to optimize \hat{f}_p , leading to interesting possibilities. (c) It is also useful to explore automatic ways of switching from our method to *SQM* when nearing the optimum.

⁵One communication pass corresponds to a vector of size equal to the feature dimension being passed between nodes.

216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269

References

- [1] L. Bottou, “Large-scale machine learning with stochastic gradient descent,” in *COMP-STAT’2010*, pp. 177–187, 2010.
- [2] N. Le Roux, M. Schmidt, and F. Bach, “A stochastic gradient method with an exponential convergence rate for strongly convex optimization with finite training sets,” in *arXiv*, 2012.
- [3] R. Johnson and T. Zhang, “Accelerating stochastic gradient descent using predictive variance reduction,” *NIPS*, 2013.
- [4] C. Hsieh, K. Chang, C. Lin, S. Keerthi, and S. Sundararajan, “A dual coordinate descent method for large-scale linear svm,” in *ICML*, pp. 408–415, 2008.
- [5] M. Zinkevich, M. Weimer, A. Smola, and L. Li, “Parallelized stochastic gradient descent,” in *NIPS*, pp. 2595–2603, 2010.
- [6] G. Mann, R. McDonald, M. Mohri, N. Silberman, and D. Walker, “Efficient large-scale distributed training of conditional maximum entropy models,” in *NIPS*, pp. 1231–1239, 2009.
- [7] K. Hall, S. Gilpin, and G. Mann, “Mapreduce/bigtable for distributed optimization,” in *NIPS Workshop on Learning on Cores, Clusters, and Clouds*, 2010.
- [8] A. Agarwal, O. Chapelle, M. Dudik, and J. Langford, “A reliable effective terascale linear learning system,” in *arXiv*, 2011.
- [9] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge, UK: Cambridge University Press, 2004.
- [10] C. Chu, S. Kim, Y. Lin, Y. Yu, G. Bradski, A. Ng, and K. Olukotun, “Map-reduce for machine learning on multicore,” *NIPS*, pp. 281–288, 2006.
- [11] C. Lin, R. Weng, and S. Keerthi, “Trust region newton method for large-scale logistic regression,” *JMLR*, pp. 627–650, 2008.