

A QUASI-NEWTON METHOD FOR LARGE SCALE SUPPORT VECTOR MACHINES

Aryan Mokhtari and Alejandro Ribeiro

Department of Electrical and Systems Engineering, University of Pennsylvania

ABSTRACT

This paper adapts a recently developed regularized stochastic version of the Broyden, Fletcher, Goldfarb, and Shanno (BFGS) quasi-Newton method for the solution of support vector machine classification problems. The proposed method is shown to converge almost surely to the optimal classifier at a rate that is linear in expectation. Numerical results show that the proposed method exhibits a convergence rate that degrades smoothly with the dimensionality of the feature vectors.

1. INTRODUCTION

Given a training set with points whose class is known the goal of a support vector machine (SVM) is to find a hyperplane that best separates the training set. If future samples are statistically identical to the training set this hyperplane provides the best classification accuracy. Computation of the separating hyperplane entails solution of a convex optimization problem that can be implemented without much difficulty in problems of moderate size [1]. Large scale problems in which the dimension of the points to be classified is large require a commensurably large training set. In these situations, computing the gradients that are required for numerical determination of the separating hyperplanes becomes infeasible and motivates the use of stochastic gradient descent methods which build unbiased gradient estimates based on small data subsamples [1–4].

However practical, stochastic gradient descent methods need a large number of iterations to converge. This translates into the need of very large training sets, or, since the size of the training set is in general limited by data collection, in the computation of hyperplanes that are not as good classifiers as they could be given the available data. In this paper we resort to quasi-Newton methods [5–12] to make better use of the provided training set. In particular, we adapt a recently developed regularized stochastic version of the Broyden, Fletcher, Goldfarb, and Shanno (BFGS) method [9] for the solution of SVM classification problems (Section 2). The proposed method is shown to converge almost surely over realizations of the training set to the optimal classifier (Theorem 1) at a rate that is linear in expectation (Theorem 2). Numerical results show that the method exhibits a convergence time that degrades smoothly with the dimensionality of the feature vectors. (Section 4).

2. STOCHASTIC QUASI-NEWTON METHOD

Consider a training set $\mathcal{S} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$ containing N pairs of the form (\mathbf{x}_i, y_i) , where $\mathbf{x}_i \in \mathbb{R}^n$ is a feature vector and $y_i \in \{-1, 1\}$ the corresponding vector's class. We want to find a hyperplane supported by a vector $\mathbf{w} \in \mathbb{R}^n$ which separates the training set so that $\mathbf{w}^T \mathbf{x}_i > 0$ for all points with $y_i = 1$ and $\mathbf{w}^T \mathbf{x}_i < 0$ for all points with $y_i = -1$. Since this vector may not exist if the data is not perfectly separable we introduce the loss function $l((\mathbf{x}, y); \mathbf{w})$ measuring the distance between the point \mathbf{x}_i and the hyperplane supported by \mathbf{w} and proceed to select the hyperplane supporting vector as the one with minimum aggregate loss

$$\mathbf{w}^* := \operatorname{argmin}_{\mathbf{w}} \frac{\lambda}{2} \|\mathbf{w}\|^2 + \frac{1}{N} \sum_{i=1}^N l((\mathbf{x}_i, y_i); \mathbf{w}), \quad (1)$$

where we also added the regularization term $\lambda \|\mathbf{w}\|^2/2$ for some constant $\lambda > 0$. The vector \mathbf{w}^* in (1) balances the minimization of the sum of

distances to the separating hyperplane, as measured by the loss function $l((\mathbf{x}, y); \mathbf{w})$, with the minimization of the L_2 norm $\|\mathbf{w}\|_2$ to enforce desirable properties in \mathbf{w}^* [13]. Common selections for the loss function are the squared hinge loss $l((\mathbf{x}, y); \mathbf{w}) = \max(0, 1 - y(\mathbf{w}^T \mathbf{x}))^2$ and the log loss $l((\mathbf{x}, y); \mathbf{w}) = \log(1 + \exp(-y(\mathbf{w}^T \mathbf{x})))$, e.g. [1].

To model (1) as a stochastic optimization problem let $\theta_i := (\mathbf{x}_i, y_i)$ be a given training point and consider a uniform probability distribution on the training set $\mathcal{S} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N = \{\theta_i\}_{i=1}^N$. Upon defining the function $f(\mathbf{w}, \theta) := \lambda \|\mathbf{w}\|^2/2 + l((\mathbf{x}_i, y_i); \mathbf{w})$ we can rewrite (1) as

$$\mathbf{w}^* := \operatorname{argmin}_{\mathbf{w}} \mathbb{E}_{\theta} [f(\mathbf{w}, \theta)] := \operatorname{argmin}_{\mathbf{w}} F(\mathbf{w}). \quad (2)$$

In (2), we (re-)interpret the sum in (1) as an expectation over the uniform discrete distribution on the set \mathcal{S} . We refer to $f(\mathbf{w}, \theta)$ as the instantaneous functions and to $F(\mathbf{w}) := \mathbb{E}_{\theta} [f(\mathbf{w}, \theta)]$ as the average function.

Since the loss functions $l((\mathbf{x}_i, y_i); \mathbf{w})$ are convex, the functions $f(\mathbf{w}, \theta) := \lambda \|\mathbf{w}\|^2/2 + l((\mathbf{x}_i, y_i); \mathbf{w})$ are strongly convex. Thus, the average objective $F(\mathbf{w})$ in (2) is also strongly convex and the optimal separating hyperplane \mathbf{w}^* can be found by stochastic gradient descent algorithms. However, the number of iterations required to run these algorithms, which translates to the number of training features (\mathbf{x}_i, y_i) that need to be acquired, becomes prohibitive for large dimensional problems. To reduce the number of iterations required for convergence we develop a regularized stochastic version of the BFGS method.

To be precise let $t \geq 0$ be an iteration index and assume that at time t we are given a sample of L realizations of the random variables θ . Group these samples in the vector $\tilde{\theta}_t := [\theta_{t1}; \dots; \theta_{tL}]$ and let \mathbf{w}_t denote the current hyperplane normal vector iterate. We then define the stochastic gradient of $F(\mathbf{w})$ associated with samples $\tilde{\theta}_t$ at point \mathbf{w}_t as

$$\hat{\mathbf{s}}(\mathbf{w}_t, \tilde{\theta}_t) = \frac{1}{L} \sum_{l=1}^L \nabla f(\mathbf{w}_t, \theta_{tl}). \quad (3)$$

Further introduce a step size sequence ϵ_t , a positive definite curvature approximation matrix $\hat{\mathbf{B}}_t$, and a regularization constant $\Gamma > 0$. The regularized stochastic BFGS algorithm is then defined by the iteration

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \epsilon_t \left(\hat{\mathbf{B}}_t^{-1} + \Gamma \mathbf{I} \right) \hat{\mathbf{s}}(\mathbf{w}_t, \tilde{\theta}_t). \quad (4)$$

The update in (4) proceeds along the negative stochastic gradient direction $-\hat{\mathbf{s}}(\mathbf{w}_t, \tilde{\theta}_t)$ pre-multiplied by the positive definite matrix $\hat{\mathbf{B}}_t^{-1} + \Gamma \mathbf{I}$ and modulated by the step size ϵ_t .

For the algorithm in (4) to have better convergence properties than gradient descent we need the matrix $\hat{\mathbf{B}}_t$ to approximate the Hessian of the objective function $\mathbf{H}(\mathbf{w}_t) := \nabla^2 F(\mathbf{w}_t)$ so that (4) approximates a stochastic version of Newton's method – the role of $\Gamma \mathbf{I}$ is to provide a guarantee of minimum progress as we discuss in the convergence analysis in Section 3. To define such approximation we use a stochastic version of the secant condition used in deterministic BFGS. Start by defining the variable and stochastic gradient variations at time t as

$$\mathbf{v}_t := \mathbf{w}_{t+1} - \mathbf{w}_t, \quad \hat{\mathbf{r}}_t := \hat{\mathbf{s}}(\mathbf{w}_{t+1}, \tilde{\theta}_t) - \hat{\mathbf{s}}(\mathbf{w}_t, \tilde{\theta}_t), \quad (5)$$

respectively, and select the matrix $\hat{\mathbf{B}}_{t+1}$ to be used in the next time step so that it satisfies the secant condition $\hat{\mathbf{B}}_{t+1} \mathbf{v}_t = \hat{\mathbf{r}}_t$. The rationale for this selection is that the Hessian $\mathbf{H}(\mathbf{w}_t)$ satisfies this condition for \mathbf{w}_{t+1} tending to \mathbf{w}_t . Notice however that the secant condition $\hat{\mathbf{B}}_{t+1} \mathbf{v}_t = \hat{\mathbf{r}}_t$

is not enough to completely specify $\hat{\mathbf{B}}_{t+1}$. To resolve this indeterminacy, matrices $\hat{\mathbf{B}}_{t+1}$ in BFGS are also required to be as close as possible to $\hat{\mathbf{B}}_t$ in terms of minimizing the Gaussian differential entropy,

$$\begin{aligned} \hat{\mathbf{B}}_{t+1} = \underset{\mathbf{Z}}{\operatorname{argmin}} \operatorname{tr} \left[\hat{\mathbf{B}}_t^{-1} \mathbf{Z} \right] - \log \det \left[\hat{\mathbf{B}}_t^{-1} \mathbf{Z} \right] - n, \\ \text{s. t.} \quad \mathbf{Z} \mathbf{v}_t = \hat{\mathbf{r}}_t, \quad \mathbf{Z} \succeq \mathbf{0}. \end{aligned} \quad (6)$$

The constraint $\mathbf{Z} \succeq \mathbf{0}$ restricts the feasible space to positive semidefinite matrices whereas the constraint $\mathbf{Z} \mathbf{v}_t = \hat{\mathbf{r}}_t$ requires \mathbf{Z} to satisfy the secant condition. The objective $\operatorname{tr}(\hat{\mathbf{B}}_t^{-1} \mathbf{Z}) - \log \det(\hat{\mathbf{B}}_t^{-1} \mathbf{Z}) - n$ is the differential entropy between Gaussian variables with covariances $\hat{\mathbf{B}}_t$ and \mathbf{Z} .

Observe that $\hat{\mathbf{B}}_{t+1}$ stays positive definite as long as the matrix $\hat{\mathbf{B}}_t \succ \mathbf{0}$ is positive definite, e.g. [10]. However, it is possible for the smallest eigenvalue of $\hat{\mathbf{B}}_t$ to become arbitrarily close to zero which means that the largest eigenvalue of $\hat{\mathbf{B}}_t^{-1}$ becomes very large. To avoid this problem we introduce a regularization of (6) that requires the smallest eigenvalue of $\hat{\mathbf{B}}_{t+1}$ to be larger than a positive constant δ ,

$$\begin{aligned} \hat{\mathbf{B}}_{t+1} = \underset{\mathbf{Z}}{\operatorname{argmin}} \operatorname{tr} \left[\hat{\mathbf{B}}_t^{-1} (\mathbf{Z} - \delta \mathbf{I}) \right] - \log \det \left[\hat{\mathbf{B}}_t^{-1} (\mathbf{Z} - \delta \mathbf{I}) \right] - n, \\ \text{s. t.} \quad \mathbf{Z} \mathbf{v}_t = \hat{\mathbf{r}}_t, \quad \mathbf{Z} \succeq \mathbf{0}. \end{aligned} \quad (7)$$

Since the logarithm determinant $\log \det[\hat{\mathbf{B}}_t^{-1} (\mathbf{Z} - \delta \mathbf{I})]$ diverges as the smallest eigenvalue of \mathbf{Z} approaches δ , the smallest eigenvalue of the Hessian approximation matrices $\hat{\mathbf{B}}_{t+1}$ computed as solutions of (7) exceeds the lower bound δ . Thus, the largest eigenvalue of $\hat{\mathbf{B}}_{t+1}^{-1}$ is bounded above by $1/\delta$. The following lemma shows that solutions of (7) can be computed by a simple algebraic formula (see [14] for proofs of results in this paper).

Lemma 1 Consider the semidefinite program in (7) where the matrix $\hat{\mathbf{B}}_t \succ \mathbf{0}$ is positive definite and define the corrected gradient variation

$$\tilde{\mathbf{r}}_t := \hat{\mathbf{r}}_t - \delta \mathbf{v}_t, \quad (8)$$

If $\tilde{\mathbf{r}}_t^T \mathbf{v}_t = (\hat{\mathbf{r}}_t - \delta \mathbf{v}_t)^T \mathbf{v}_t > 0$, the solution $\hat{\mathbf{B}}_{t+1}$ of (7) can be written as

$$\hat{\mathbf{B}}_{t+1} = \hat{\mathbf{B}}_t + \frac{\tilde{\mathbf{r}}_t \tilde{\mathbf{r}}_t^T}{\mathbf{v}_t^T \tilde{\mathbf{r}}_t} - \frac{\hat{\mathbf{B}}_t \mathbf{v}_t \mathbf{v}_t^T \hat{\mathbf{B}}_t}{\mathbf{v}_t^T \hat{\mathbf{B}}_t \mathbf{v}_t} + \delta \mathbf{I}. \quad (9)$$

When $\delta = 0$ the update in (9) coincides with standard non-regularized BFGS [7, 10, 11, 15]. Therefore, the differences between BFGS and regularized BFGS are the replacement of the gradient variation $\hat{\mathbf{r}}_t$ by the corrected variation $\tilde{\mathbf{r}}_t := \hat{\mathbf{r}}_t - \delta \mathbf{v}_t$ and the addition of the regularization term $\delta \mathbf{I}$. Notice that the expression in (9) is the solution to (7) only when the inner product $\tilde{\mathbf{r}}_t^T \mathbf{v}_t = (\hat{\mathbf{r}}_t - \delta \mathbf{v}_t)^T \mathbf{v}_t > 0$.

2.1. Regularized stochastic BFGS support vector machines

To solve the SVM problem in (1) using regularized stochastic BFGS we need the stochastic gradient in (3). For that, select a sample of L feature vectors $\tilde{\mathbf{x}} = [\mathbf{x}_1; \dots; \mathbf{x}_L]$ and corresponding classes $\tilde{\mathbf{y}} = [y_1; \dots; y_L]$ from the training set and compute the stochastic gradient as [cf. (3)]

$$\hat{\mathbf{s}}(\mathbf{w}, (\tilde{\mathbf{x}}, \tilde{\mathbf{y}})) = \lambda \mathbf{w} + \frac{1}{L} \sum_{i=1}^L \nabla_{\mathbf{w}} l((\mathbf{x}_i, y_i); \mathbf{w}). \quad (10)$$

Start at time t with current iterate \mathbf{w}_t and recall that $\hat{\mathbf{B}}_t$ stands for the Hessian approximation computed by stochastic BFGS in the previous iteration. Proceed to collect feature vectors $\tilde{\mathbf{x}}_t = [\mathbf{x}_{t1}; \dots; \mathbf{x}_{tL}]$ and their corresponding class vectors $\tilde{\mathbf{y}}_t = [y_{t1}; \dots; y_{tL}]$ and for each pair $(\tilde{\mathbf{x}}_t, \tilde{\mathbf{y}}_t)$ determine the stochastic gradients $\hat{\mathbf{s}}(\mathbf{w}_t, (\tilde{\mathbf{x}}_t, \tilde{\mathbf{y}}_t))$ as per (10). Descend then along the direction $(\hat{\mathbf{B}}_t^{-1} + \Gamma \mathbf{I}) \hat{\mathbf{s}}(\mathbf{w}_t, (\tilde{\mathbf{x}}_t, \tilde{\mathbf{y}}_t))$ as per (4). This leads to the next iterate \mathbf{w}_{t+1} , but to complete the iteration we still need to compute the updated Hessian approximation $\hat{\mathbf{B}}_{t+1}$. To do so compute the stochastic gradient $\hat{\mathbf{s}}(\mathbf{w}_{t+1}, (\tilde{\mathbf{x}}_t, \tilde{\mathbf{y}}_t))$ associated with the *same* set of

Algorithm 1 Regularized stochastic BFGS support vector machines

Require: Variable \mathbf{w}_0 . Hessian approximation $\hat{\mathbf{B}}_0 \succ \delta \mathbf{I}$.

- 1: **for** $t = 0, 1, 2, \dots$ **do**
- 2: Collect L training points $\tilde{\mathbf{x}}_t = [\mathbf{x}_{t1}, \dots, \mathbf{x}_{tL}]$ and $\tilde{\mathbf{y}}_t = [y_{t1}, \dots, y_{tL}]$
- 3: Compute stochastic gradient $\hat{\mathbf{s}}(\mathbf{w}_t, (\tilde{\mathbf{x}}_t, \tilde{\mathbf{y}}_t))$ [cf. (10)].

$$\hat{\mathbf{s}}(\mathbf{w}_t, (\tilde{\mathbf{x}}_t, \tilde{\mathbf{y}}_t)) = \lambda \mathbf{w}_t + \frac{1}{L} \sum_{i=1}^L \nabla_{\mathbf{w}} l((\mathbf{x}_{ti}, y_{ti}); \mathbf{w}_t).$$

- 4: Descend along direction $(\hat{\mathbf{B}}_t^{-1} + \Gamma \mathbf{I}) \hat{\mathbf{s}}(\mathbf{w}_t, (\tilde{\mathbf{x}}_t, \tilde{\mathbf{y}}_t))$ [cf. (4)]

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \epsilon_t (\hat{\mathbf{B}}_t^{-1} + \Gamma \mathbf{I}) \hat{\mathbf{s}}(\mathbf{w}_t, (\tilde{\mathbf{x}}_t, \tilde{\mathbf{y}}_t)).$$

- 5: Compute $\hat{\mathbf{s}}(\mathbf{w}_{t+1}, (\tilde{\mathbf{x}}_t, \tilde{\mathbf{y}}_t))$ [cf. (10)]

$$\hat{\mathbf{s}}(\mathbf{w}_{t+1}, (\tilde{\mathbf{x}}_t, \tilde{\mathbf{y}}_t)) = \lambda \mathbf{w}_{t+1} + \frac{1}{L} \sum_{i=1}^L \nabla_{\mathbf{w}} l((\mathbf{x}_{ti}, y_{ti}); \mathbf{w}_{t+1}).$$

- 6: Variable and modified stochastic gradient variations [cf. (5) and (8)]

$$\mathbf{v}_t = \mathbf{w}_{t+1} - \mathbf{w}_t,$$

$$\tilde{\mathbf{r}}_t = \hat{\mathbf{s}}(\mathbf{w}_{t+1}, (\tilde{\mathbf{x}}_t, \tilde{\mathbf{y}}_t)) - \hat{\mathbf{s}}(\mathbf{w}_t, (\tilde{\mathbf{x}}_t, \tilde{\mathbf{y}}_t)) - \delta \mathbf{v}_t$$

- 7: Update Hessian approximation matrix [cf. (9)]

$$\hat{\mathbf{B}}_{t+1} = \hat{\mathbf{B}}_t + \frac{\tilde{\mathbf{r}}_t \tilde{\mathbf{r}}_t^T}{\mathbf{v}_t^T \tilde{\mathbf{r}}_t} - \frac{\hat{\mathbf{B}}_t \mathbf{v}_t \mathbf{v}_t^T \hat{\mathbf{B}}_t}{\mathbf{v}_t^T \hat{\mathbf{B}}_t \mathbf{v}_t} + \delta \mathbf{I}.$$

- 8: **end for**
-

random data points samples $(\tilde{\mathbf{x}}_t, \tilde{\mathbf{y}}_t)$ used to compute the stochastic gradient $\hat{\mathbf{s}}(\mathbf{w}_t, (\tilde{\mathbf{x}}_t, \tilde{\mathbf{y}}_t))$. The stochastic gradient variation $\hat{\mathbf{r}}_t$, the variable variation \mathbf{v}_t , and the modified stochastic gradient variation $\tilde{\mathbf{r}}_t$ at time t are now computed using (5) and (8). The Hessian approximation $\hat{\mathbf{B}}_{t+1}$ for the next iteration is defined as the matrix that satisfies the stochastic secant condition $\hat{\mathbf{B}}_{t+1} \mathbf{v}_t = \tilde{\mathbf{r}}_t$ and is closest to $\hat{\mathbf{B}}_t$ in the sense of (7). As per Lemma 1 we can compute $\hat{\mathbf{B}}_{t+1}$ using (9).

The solution of (1) using regularized stochastic BFGS is summarized in Algorithm 1. The two core steps in each iteration are the descent in Step 4 and the update of the Hessian approximation $\hat{\mathbf{B}}_t$ in Step 8. Step 2 comprises the observation of L pairs of data points and feature vectors that are required to compute the stochastic gradients in steps 3 and 5. The stochastic gradient $\hat{\mathbf{s}}(\mathbf{w}_t, (\tilde{\mathbf{x}}_t, \tilde{\mathbf{y}}_t))$ in Step 3 is used in the descent iteration in Step 4. The stochastic gradient of Step 3 along with the stochastic gradient $\hat{\mathbf{s}}(\mathbf{w}_{t+1}, (\tilde{\mathbf{x}}_t, \tilde{\mathbf{y}}_t))$ of Step 5 are used to compute the variations in steps 6 and 7 that permit carrying out the update of the Hessian approximation $\hat{\mathbf{B}}_t$ in Step 8. Iterations are initialized with arbitrary vector \mathbf{w}_0 and matrix $\hat{\mathbf{B}}_0$ having all eigenvalues larger than δ .

3. CONVERGENCE ANALYSIS

Our goal here is to show that as time progresses the sequence of classifiers \mathbf{w}_t approaches the optimal classifier \mathbf{w}^* . In proving this result we make the following assumptions.

Assumption 1 For any set of samples $\tilde{\boldsymbol{\theta}} = [\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_L]$ the instantaneous functions $\hat{f}(\mathbf{w}, \tilde{\boldsymbol{\theta}}) := (1/L) \sum_{i=1}^L f(\mathbf{w}, \boldsymbol{\theta}_i)$ are twice differentiable and their Hessians $\hat{\mathbf{H}}(\mathbf{w}, \tilde{\boldsymbol{\theta}}) = \nabla_{\mathbf{w}}^2 \hat{f}(\mathbf{w}, \tilde{\boldsymbol{\theta}})$ have lower and upper bounded eigenvalues,

$$\tilde{m} \mathbf{I} \preceq \hat{\mathbf{H}}(\mathbf{w}, \tilde{\boldsymbol{\theta}}) \preceq \tilde{M} \mathbf{I}. \quad (11)$$

Assumption 2 There exists a constant S^2 such that for all variables \mathbf{w} the second moment of the norm of the stochastic gradient satisfies

$$\mathbb{E}_{\boldsymbol{\theta}} [\|\hat{\mathbf{s}}(\mathbf{w}_t, \tilde{\boldsymbol{\theta}}_t)\|^2] \leq S^2, \quad (12)$$

Assumption 3 The regularization constant δ is smaller than the smallest Hessian eigenvalue \tilde{m} , i.e., $\delta < \tilde{m}$.

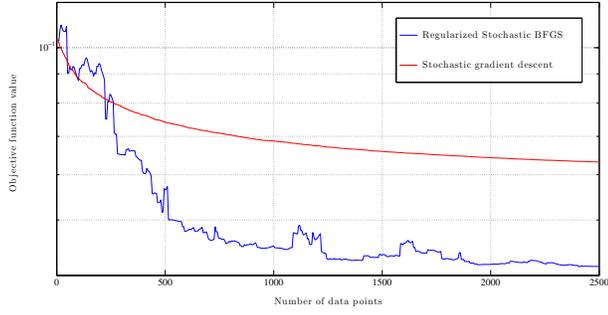


Fig. 1. Convergence of stochastic gradient descent and regularized stochastic BFGS for feature vectors of dimension $n = 4$. Convergence of stochastic BFGS is faster than convergence of stochastic gradient descent (sample size $L = 5$; stepsizes $\epsilon_t = \epsilon_0 \tau / (\tau + t)$ with $\epsilon_0 = 3 \times 10^{-2}$ and $\tau = 10^2$; stochastic BFGS parameters $\delta = 10^{-3}$ and $\Gamma = 10^{-4}$).

Recall that according to Lemma 1 the update in (9) is a solution to (7) as long as the inner product $(\hat{\mathbf{r}}_t - \delta \mathbf{v}_t)^T \mathbf{v}_t = \hat{\mathbf{r}}_t^T \mathbf{v}_t > 0$ is positive. Our first result is to show that selecting $\delta < \tilde{m}$ as required by Assumption 3 guarantees that this inequality is satisfied for all times t .

Lemma 2 Consider the modified stochastic gradient variation $\tilde{\mathbf{r}}_t$ defined in (8) and the variable variation \mathbf{v}_t defined in (5). If assumptions 1 and 3 are true, then, for all times t it holds

$$\tilde{\mathbf{r}}_t^T \mathbf{v}_t = (\hat{\mathbf{r}}_t - \delta \mathbf{v}_t)^T \mathbf{v}_t \geq (\tilde{m} - \delta) \|\mathbf{v}_t\|^2 > 0. \quad (13)$$

The result in Lemma 2 guarantees that the regularized stochastic BFGS algorithm as defined by recursive application of (4), (5), (8), and (9) results in matrices $\hat{\mathbf{B}}_t$ that solve (7). In particular, this implies that $\hat{\mathbf{B}}_t$ is positive definite with smallest eigenvalue not smaller than δ , i.e., $\hat{\mathbf{B}}_t \preceq \delta \mathbf{I}$. This implies that all the eigenvalues of $\hat{\mathbf{B}}_t^{-1}$ are between 0 and $1/\delta$ and that, as a consequence, the matrix $\hat{\mathbf{B}}_t^{-1} + \Gamma \mathbf{I}$ is such that

$$\Gamma \mathbf{I} \preceq \hat{\mathbf{B}}_t^{-1} + \Gamma \mathbf{I} \preceq \left(\Gamma + \frac{1}{\delta}\right) \mathbf{I}. \quad (14)$$

Having matrices $\hat{\mathbf{B}}_t^{-1} + \Gamma \mathbf{I}$ that are strictly positive definite with eigenvalues uniformly upper bounded by $\Gamma + (1/\delta)$ leads to the conclusion that if $\hat{\mathbf{s}}(\mathbf{w}_t, \hat{\boldsymbol{\theta}}_t)$ is a descent direction, the same holds true of $(\hat{\mathbf{B}}_t^{-1} + \Gamma \mathbf{I}) \hat{\mathbf{s}}(\mathbf{w}_t, \hat{\boldsymbol{\theta}}_t)$. The stochastic gradient $\hat{\mathbf{s}}(\mathbf{w}_t, \hat{\boldsymbol{\theta}}_t)$ is not a descent direction in general, but we know that this is true for its conditional expectation $\mathbb{E}[\hat{\mathbf{s}}(\mathbf{w}_t, \hat{\boldsymbol{\theta}}_t) | \mathbf{w}_t] = \nabla_{\mathbf{w}} F(\mathbf{w}_t)$. Therefore, we conclude that $(\hat{\mathbf{B}}_t^{-1} + \Gamma \mathbf{I}) \hat{\mathbf{s}}(\mathbf{w}_t, \hat{\boldsymbol{\theta}}_t)$ is an average descent direction because $\mathbb{E}[(\hat{\mathbf{B}}_t^{-1} + \Gamma \mathbf{I}) \hat{\mathbf{s}}(\mathbf{w}_t, \hat{\boldsymbol{\theta}}_t) | \mathbf{w}_t] = (\hat{\mathbf{B}}_t^{-1} + \Gamma \mathbf{I}) \nabla_{\mathbf{w}} F(\mathbf{w}_t)$. Having a displacement $\mathbf{w}_{t+1} - \mathbf{w}_t$ that is a descent direction on average implies convergence towards optimal arguments as we claim in the following theorem.

Theorem 1 Consider the regularized stochastic BFGS algorithm as defined by (4), (5), (8), and (9). If assumptions 1-3 hold true and the sequence of stepsizes satisfies is nonsummable but square summable, i.e., if $\sum_{t=0}^{\infty} \epsilon_t = \infty$, and $\sum_{t=0}^{\infty} \epsilon_t^2 < \infty$, the limit infimum of the squared Euclidean distance to optimality $\|\mathbf{w}_t - \mathbf{w}^*\|^2$ satisfies

$$\liminf_{t \rightarrow \infty} \|\mathbf{w}_t - \mathbf{w}^*\|^2 = 0 \quad a.s. \quad (15)$$

over realizations of the random samples $\{\hat{\boldsymbol{\theta}}_t\}_{t=1}^{\infty}$.

Theorem 1 establishes convergence of the stochastic regularized BFGS algorithm summarized in Algorithm 1. In the proof of this result the lower bound in the eigenvalues of $\hat{\mathbf{B}}_t$ enforced by the regularization in (9) plays a fundamental role. Roughly speaking, the lower bound in

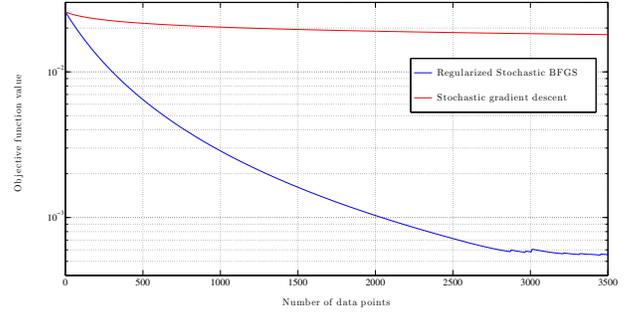


Fig. 2. Convergence of stochastic gradient descent and regularized stochastic BFGS for feature vectors of dimension $n = 40$. Stochastic BFGS is still practicable whereas stochastic gradient descent becomes too slow for practical use (parameters as in Fig. 1).

the eigenvalues of $\hat{\mathbf{B}}_t$ results in an upper bound on the eigenvalues of $\hat{\mathbf{B}}_t^{-1}$ which limits the effect of random variations on the stochastic gradient $\hat{\mathbf{s}}(\mathbf{w}_t, \hat{\boldsymbol{\theta}}_t)$. If this regularization is not implemented, i.e., if we keep $\delta = 0$, we may observe catastrophic amplification of random variations of the stochastic gradient. This effect is indeed observed in the numerical experiments in Section 4. The addition of the identity matrix bias $\Gamma \mathbf{I}$ in (4) is also instrumental in the proof of Theorem 1. This bias limits the effects of randomness in the curvature estimate $\hat{\mathbf{B}}_t$. If random variations in the curvature estimate $\hat{\mathbf{B}}_t$ result in a matrix $\hat{\mathbf{B}}_t^{-1}$ with small eigenvalues the term $\Gamma \mathbf{I}$ dominates and (4) reduces to stochastic gradient descent. This ensures continued progress towards the optimal argument \mathbf{w}^* .

The convergence claim in Theorem 1 is complemented by a expected convergence rate result which we state in the following theorem.

Theorem 2 Consider the regularized stochastic BFGS algorithm as defined by (4)-(9) and let the sequence of stepsizes be given by $\epsilon_t = \epsilon_0 \tau / (\tau + t)$ with the parameter ϵ_0 sufficiently small and the parameter τ sufficiently large so as to satisfy the inequality

$$2 \epsilon_0 \tau \Gamma > 1. \quad (16)$$

If assumptions 1 and 2 hold true the difference between the expected objective value $\mathbb{E}[F(\mathbf{w}_t)]$ at time t and the optimal objective $F(\mathbf{w}^*)$ satisfies

$$\mathbb{E}[F(\mathbf{w}_t)] - F(\mathbf{w}^*) \leq \frac{\xi}{\tau + t}, \quad (17)$$

where the constant ξ satisfies

$$\xi = \max \left\{ \frac{\epsilon_0^2 \tau^2 K}{2\epsilon_0 \tau \Gamma - 1}, (1 + \tau)(F(\mathbf{w}_0) - F(\mathbf{w}^*)) \right\}. \quad (18)$$

Theorem 2 shows the convergence rate of regularized stochastic BFGS is at least linear in terms of the expectation of the objective function. This rate is typical of stochastic optimization algorithms and, in that sense, no better than stochastic gradient descent. While the convergence rate doesn't change, improvements in convergence time are marked as we illustrate with the numerical experiments of the following section.

4. NUMERICAL ANALYSIS

We test Algorithm 1 when using the squared hinge loss $l((\mathbf{x}, y); \mathbf{w}) = \max(0, 1 - y(\mathbf{x}^T \mathbf{w}))^2$ in (1). The training set $\mathcal{S} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$ contains $N = 10^4$ feature vectors half of which belong to the class $y_i = -1$ with the other half belonging to the class $y_i = 1$. For the class $y_i = -1$ each of the n components of each of the feature vectors $\mathbf{x}_i \in \mathbb{R}^n$ is chosen uniformly at random from the interval $[-0.8, 0.2]$. Likewise, each of the n components of each of the feature vectors $\mathbf{x}_i \in \mathbb{R}^n$ is chosen uniformly

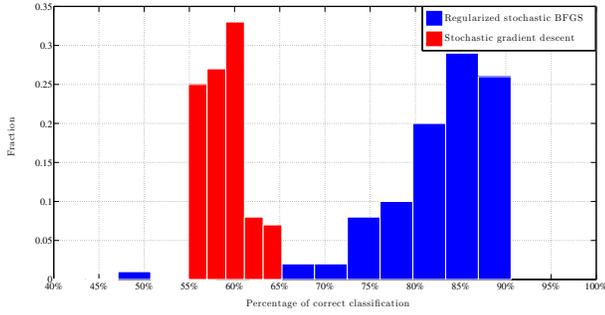


Fig. 3. Histogram of correct classification percentages for dimension $n = 4$ and training set of size $N = 2.5 \times 10^3$. Vectors computed by stochastic BFGS outperform those computed via stochastic gradients and are not far from the accuracy of clairvoyant classifiers (test sets contain 10^4 samples; histogram is across 10^3 realizations; parameters as in Fig. 1).

at random from the interval $[-0.2, 0.8]$ for the class $y_i = 1$. The overlap in the range of the feature vectors is such that the classification accuracy expected from a clairvoyant classifier that knows the statistic model of the data set is less than 100%. Exact values can be computed from the Irwin-Hall distribution [16]. For $n = 4$ this amounts to 98%.

We set the parameter λ in (1) to $\lambda = 10^{-3}$. Since the Hessian eigenvalues of $f(\mathbf{w}, \boldsymbol{\theta}) := \lambda \|\mathbf{w}\|^2/2 + l(\mathbf{x}_i, y_i; \mathbf{w})$ are, at least, equal to λ this implies that the eigenvalue lower bound \tilde{m} is such that $\tilde{m} \geq \lambda = 10^{-3}$. Thus, we set the BFGS regularization parameter to $\delta = \lambda = 10^{-3}$. Further set the minimum progress parameter in (3) to $\Gamma = 10^{-4}$ and the sample size for computation of stochastic gradients to $L = 5$. Stepsizes are of the form $\epsilon_t = \epsilon_0 \tau / (\tau + t)$ with $\epsilon_0 = 3 \times 10^{-2}$ and $\tau = 10^2$. We compare the behavior of stochastic gradient descent and stochastic BFGS for a small dimensional problem with $n = 4$ and a large problem with $n = 40$. For stochastic gradient descent the sample size in (3) is $L = 1$ and we use the same stepsize sequence used for stochastic BFGS.

An illustration of the relative performances of stochastic gradient descent and BFGS for $n = 4$ is presented in Fig. 1. The value of the objective function $F(\mathbf{w}_t)$ is represented with respect to the number of feature vectors processed, which is given by the product Lt between the iteration index and the sample size used to compute stochastic gradients. This is done because the sample sizes in stochastic BFGS ($L = 5$) and stochastic gradient descent ($L = 1$) are different. The curvature correction of stochastic BFGS results in significant reductions in convergence time. E.g., Stochastic BFGS achieves an objective value of $F(\mathbf{w}_t) = 6.5 \times 10^{-2}$ upon processing of $Lt = 315$ feature vectors. To achieve the same objective value $F(\mathbf{w}_t) = 6.5 \times 10^{-2}$ stochastic gradient descent processes 1.74×10^3 feature vectors. Conversely, after processing $Lt = 2.5 \times 10^3$ feature vectors the objective values achieved by stochastic BFGS and gradient descent are $F(\mathbf{w}_t) = 4.14 \times 10^{-2}$ and $F(\mathbf{w}_t) = 6.31 \times 10^{-2}$, respectively.

The performance difference between the two methods is larger for feature vectors of larger dimension n . The plot of the value of the objective function $F(\mathbf{w}_t)$ with respect to the number of feature vectors processed Lt is shown in Fig. 2 for $n = 40$. The convergence time of stochastic BFGS increases but is still acceptable. For stochastic gradient descent the algorithm becomes unworkable. After processing 3.5×10^3 stochastic BFGS reduces the objective value to $F(\mathbf{w}_t) = 5.55 \times 10^{-4}$ while stochastic gradient descent has barely made progress at $F(\mathbf{w}_t) = 1.80 \times 10^{-2}$.

Differences in convergence times translate into differences in classification accuracy when we process all N vectors in the training set. This is shown for dimension $n = 4$ and training set size $N = 2.5 \times 10^3$ in Fig. 3. To build Fig. 3 we process $N = 2.5 \times 10^3$ feature vectors with stochastic BFGS and stochastic gradient descent with the same parameters used in Fig. 1. We then use these vectors to classify 10^4 observations in the test set and record the percentage of samples that are correctly classified. The

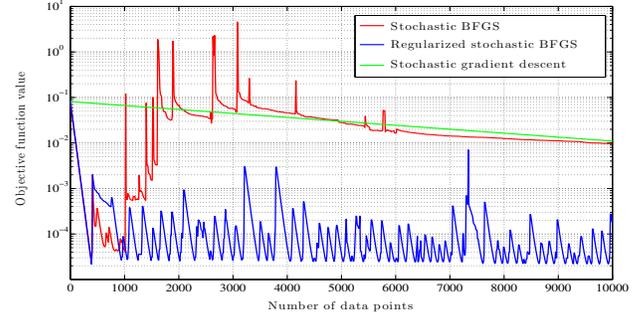


Fig. 4. Comparison of gradient descent, regularized stochastic BFGS, and (non regularized) stochastic BFGS. The regularization is fundamental to control the erratic behavior of stochastic BFGS (sample size $L = 5$; constant stepsize $\epsilon_t = 10^{-1}$; stochastic BFGS parameters $\delta = 10^{-3}$ and $\Gamma = 10^{-4}$, feature vector dimension $n = 10$).

process is repeated 10^3 times to estimate the probability distribution of the correct classification percentage represented by the histograms shown. The dominance of stochastic BFGS with respect to stochastic gradient descent is almost uniform. The vector \mathbf{w}_t computed by stochastic gradient descent classifies correctly at most 65% of the of the feature vectors in the test set. The vector \mathbf{w}_t computed by stochastic BFGS exceeds this accuracy with probability 0.98. Perhaps more relevant, the classifier computed by stochastic BFGS achieves a mean classification accuracy of 82.2% which is not far from the clairvoyant classification accuracy of 98%. Although performance is markedly better in general, stochastic BFGS fails to compute a working classifier with probability 0.02.

We also investigate the difference between regularized and non-regularized versions of stochastic BFGS for feature vectors of dimension $n = 10$. Observe that non-regularized stochastic BFGS corresponds to making $\delta = 0$ and $\Gamma = 0$ in Algorithm 1. To illustrate the advantage of the regularization induced by the proximity requirement in (7), as opposed to the non regularized proximity requirement in (6), we keep a constant stepsize $\epsilon_t = 10^{-1}$. The corresponding evolutions of the objective function values $F(\mathbf{w}_t)$ with respect to the number of feature vectors processed Lt are shown in Fig. 4 along with the values associated with stochastic gradient descent. As we reach convergence the likelihood of having small eigenvalues appearing in $\hat{\mathbf{B}}_t$ becomes significant. In regularized stochastic BFGS this results in recurrent jumps away from the optimal classifier \mathbf{w}^* . However, the regularization term limits the size of the jumps and further permits the algorithm to consistently recover a reasonable curvature estimate. In Fig. 4 we process 10^4 feature vectors and observe many occurrences of small eigenvalues. However, the algorithm always recovers and heads back to a good approximation of \mathbf{w}^* . In the absence of regularization small eigenvalues in $\hat{\mathbf{B}}_t$ result in larger jumps away from \mathbf{w}^* . This not only sets back the algorithm by a much larger amount than in the regularized case but also results in a catastrophic deterioration of the curvature approximation matrix $\hat{\mathbf{B}}_t$. In Fig. 4 we observe recovery after the first two occurrences of small eigenvalues but eventually there is a catastrophic deviation after which non-regularized stochastic BFGS behaves not better than stochastic gradient descent.

5. CONCLUSIONS

We considered the problem of determining the separating hyperplane of a support vector machine using stochastic optimization. In order to handle large scale problems with reasonable convergence times we adapted a regularized stochastic version of the Broyden, Fletcher, Goldfarb, and Shanno (BFGS) quasi-Newton method [9]. We derived theoretical convergence guarantees that are customary of stochastic optimization and illustrated improvements in convergence time through numerical analysis.

6. REFERENCES

- [1] L. Bottou, "Large-scale machine learning with stochastic gradient descent," *In Proceedings of COMPSTAT'2010*, pp. 177–186, Physica-Verlag HD, 2010.
- [2] S. Shalev-Shwartz, Y. Singer, and N. Srebro, "Pegasos: Primal estimated sub-gradient solver for svm," *In Proceedings of the 24th international conference on Machine learning*, pp. 807–814, ACM, 2007.
- [3] T. Zhang, "Solving large scale linear prediction problems using stochastic gradient descent algorithms," *In Proceedings of the twenty-first international conference on Machine learning*, p. 919926, ACM, 2004.
- [4] N. LeRoux, M. Schmidt, and F. Bach, "A stochastic gradient method with an exponential convergence rate for strongly-convex optimization with finite training sets," *arXiv preprint arXiv*, 1202.6258, 2012.
- [5] C. G. Broyden, J. E. D. Jr., Wang, and J. J. More, "On the local and superlinear convergence of quasi-newton methods," *IMA J. Appl. Math.*, vol. 12, no. 3, pp. 223–245, June 1973.
- [6] A. Bordes, L. Bottou, and P. Gallinari, "Sgd-qn: Careful quasi-newton stochastic gradient descent," *The Journal of Machine Learning Research*, vol. 10, pp. 1737–1754, 2009.
- [7] R. H. Byrd, J. Nocedal, and Y. Yuan, "Global convergence of a class of quasi-newton methods on convex problems," *SIAM J. Numer. Anal.*, vol. 24, no. 5, pp. 1171–1190, October 1987.
- [8] A. Mokhtari and A. Ribeiro, "A dual stochastic dfg algorithm for optimal resource allocation in wireless systems," in *Proc. IEEE 14th Workshop on Signal Process. Advances in Wireless Commun. (SPAWC)*. pp. 21-25, Darmstadt Germany, June 16-19 2013.
- [9] ———, "Regularized stochastic bfgs algorithm," in *Proc. IEEE Global Conf. on Signal and Inform. Process.* pp. 1109-1112, Austin Texas, Dec. 3-5 2013.
- [10] J. Nocedal and S. J. Wright, *Numerical optimization*, 2nd ed. New York, NY: Springer-Verlag, 1999.
- [11] M. J. D. Powell, *Some global convergence properties of a variable metric algorithm for minimization without exact line search*, 2nd ed. London, UK: Academic Press, 1971.
- [12] N. N. Schraudolph, J. Yu, and S. Gnter, "A stochastic quasi-newton method for online convex optimization," *In Proc. 11th Intl. Conf. on Artificial Intelligence and Statistics (AISTATS)*, p. 433 440, Soc. for Artificial Intelligence and Statistics, 2007.
- [13] V. Vapnik, *The nature of statistical learning theory*, 2nd ed. springer, 1999.
- [14] A. Mokhtari and A. Ribeiro, "Res: Regularized stochastic bfgs algorithm," *arXiv preprint arXiv*, 1401.7625, 2014.
- [15] J. J. E. Dennis and J. J. More, "A characterization of super linear convergence and its application to quasi-newton methods," *Mathematics of computation*, vol. 28, no. 126, pp. 549–560, 1974.
- [16] N. L. Johnson, S. Kotz, and N. Balakrishnan, *Continuous Univariate Distributions*, vol. 2, 2nd ed. Wiley-Interscience, 1995.