

RAPID: RAPIDLY ACCELERATED PROXIMAL GRADIENT ALGORITHMS FOR CONVEX MINIMIZATION

Ziming Zhang and Venkatesh Saligrama

ECE, Boston University, MA 02215

ABSTRACT

In this paper, we propose a new algorithm to speed-up the convergence of accelerated proximal gradient (APG) methods. In order to minimize a convex function $f(\mathbf{x})$, our algorithm introduces a simple line search step after each proximal gradient step in APG so that a biconvex function $f(\theta\mathbf{x})$ is minimized over scalar variable $\theta > 0$ while fixing variable \mathbf{x} . We propose two new ways of constructing the auxiliary variables in APG based on the intermediate solutions of the proximal gradient and the line search steps. We prove that at arbitrary iteration step $t(t \geq 1)$, our algorithm can achieve a smaller upper-bound for the gap between the current and optimal objective values than those in the traditional APG methods such as FISTA [1], making it converge faster in practice. We apply our algorithm to many important convex optimization problems such as sparse linear regression. Our experimental results demonstrate that our algorithm converges faster than APG, even comparable to some sophisticated solvers.

1. INTRODUCTION

As a general convex minimization algorithm, accelerated proximal gradient (APG) has been increasingly attracting attention, and it has been widely used in many different research areas such as signal processing [2], computer vision [3], and data mining [4]. In general, APG solves the following problem:

$$\min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) = f_1(\mathbf{x}) + f_2(\mathbf{x}) \quad (1)$$

where $\mathcal{X} \subseteq \mathbb{R}^n$ denotes the closed and convex feasible set for variable $\mathbf{x} \in \mathbb{R}^n$, and $f(\mathbf{x}) : \mathbb{R}^n \rightarrow \mathbb{R}$ is a convex function, which consists of a convex and differentiable function f_1 with Lipschitz constant $L \geq 0$ and a convex but non-differentiable function f_2 .

In APG, proximal gradient is used to update variables based on the *proximity operator* [2], denoted as **prox**. The basic idea of proximity operator is to approximate a convex function using a strongly convex function whose minimizer in the feasible set is returned as an approximate solution for the original minimization problem. At the optimal solution, the solution returned by proximity operator is identical to the optimal. Taking FISTA [1] as an example, we can see that APG generates an *auxiliary variable* for proximal gradient so that

the convergence rate of APG for general convex optimization is $O(\frac{1}{T^2})$, which was proved to be optimal for first-order gradient descent methods [5].

Generally speaking, the computational bottleneck in APG comes from the following two aspects:

(i) **Computation of proximal gradients.** Evaluating the gradients of f_1 could be time-consuming, because the evaluation is over the entire dataset. This situation is more prominent for high dimensional and large-scale datasets. Also, projecting a point into the feasible set may be difficult. Many recent approaches have attempted to reduce this computational complexity, such as inexact proximal gradient methods [6], stochastic proximal gradient methods [7, 8], and distributed proximal gradient methods [9].

(ii) **Number of iterations.** In order to minimize the number of iterations, intuitively in each iteration the resulting function value should be as close to the global minimum as possible. One way to achieve this is to optimize the step size in the proximal gradient, which unfortunately may be very difficult for many convex functions. Instead, in practice line search [1, 10, 11, 12] is widely used to estimate the step size so that the function value is decreasing. In general, the line search in the proximal gradient step has to evaluate the function repeatedly by changing the step size so that numerically the learned step size is close to the optimal. Alternatively many types of restarting schemes [10, 13, 14] have been utilized to reduce the number of iterations empirically. Here additional restarting conditions are established and evaluated periodically. If such conditions are satisfied, the algorithm is re-initialized using current solutions.

1.1. Our Contribution

In this paper, we focus on reducing the number of iterations, and simply assume that the non-differentiable function f_2 is *simple* [13] for performing proximal gradient efficiently. One instance of a simple function is ℓ_1 norm penalization that is often used in sparse regression.

Our first contribution is to propose a new general algorithm, *Rapidly Accelerated Proximal Gradient (RAPID)*, to speed up the empirical convergence of APG, where an additional simple line search step is introduced after the proximal gradient step. Fig. 1(a) illustrates the basic idea of our al-

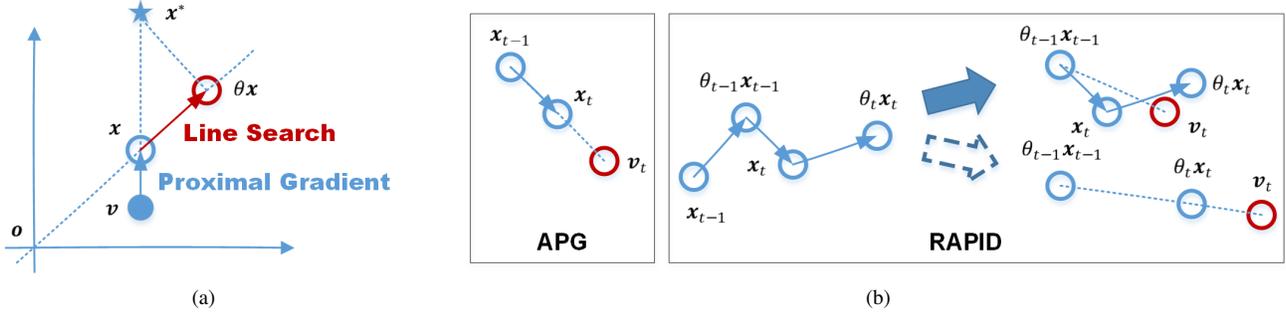


Fig. 1. (a) Illustration of the basic idea of our algorithm for improving convergence of APG in 2D, where \mathbf{o} denotes the origin of the 2D coordinate system, \mathbf{v} , \mathbf{x} , $\theta\mathbf{x}$, and \mathbf{x}^* denote the solutions at the initial point, after the proximal gradient step, after the line search step, and at the minimum, respectively, and the directed lines denote the updating directions. (b) Illustration of the differences between APG and our RAPID in terms of constructing the auxiliary variable \mathbf{v}_t in iteration t , where directed lines denote the generating orders of intermediate solutions, and the dotted lines denote the directions of the auxiliary variable starting at the solution in iteration $t - 1$. This figure is best viewed in color.

gorithm in 2D. After the proximal gradient step, another line search is applied along the direction of the current solution \mathbf{x} . Ideally, we would like to find a scalar $\theta > 0$ so that $\theta = \arg \min_{\{\hat{\theta} | \hat{\theta}\mathbf{x} \in \mathcal{X}\}} f(\hat{\theta}\mathbf{x})$. Therefore, we can guarantee $f(\theta\mathbf{x}) \leq f(\mathbf{x})$. The positiveness of θ guarantees that both \mathbf{x} and $\theta\mathbf{x}$ point to the same direction so that the information from gradient descent can be preserved. Geometrically, this additional line search tries to push \mathbf{x} towards the optimal solution, making the distance between the current and optimal solutions smaller. Also, we propose two ways of constructing the auxiliary variables based on the intermediate solutions in the previous and current iterations, as illustrated in Fig. 1(b).

Our second contribution is that theoretically we prove that at an arbitrary iteration t , the upper bound of the objective error $f(\theta_t\mathbf{x}_t) - \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$ in our algorithm is consistently smaller than that of $f(\mathbf{x}_t) - \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$ in traditional APG methods such as FISTA. This result implies that in order to achieve the same precision, the number of iterations in our algorithm is most likely no more than that in APG. In other words, empirically our algorithm tends to converge faster than APG.

We apply our general algorithm to several interesting convex optimization problems, *i.e.* LASSO [15], group LASSO [16], and least square loss with trace norm [17, 18], and compare our performance with APG and SLEP [19]. Our experimental results demonstrate present empirical evidence of our theoretical results on faster convergence.

This paper is organized as follows. In Section 2, we explain the details of our RAPID algorithm. In Section 3 we prove our theorems on the convergence rate of our algorithm. We show our experimental results and comparison in Section 4, and conclude the paper in Section 5.

2. ALGORITHMS

In general, there are two basic steps in each iteration in APG algorithms: (1) performing *proximal gradients*, and (2) constructing *auxiliary variables*. Proximal gradient is defined as

Algorithm 1: RAPID: Rapidly Accelerated Proximal Gradient Algorithms

Input : $f(\mathbf{x})$, $\lambda_{\mathbf{x}}$, λ_{θ}

Output: \mathbf{x}

$\mathbf{x}_0 \leftarrow \mathbf{0}$; $\mathbf{v}_0 \leftarrow \mathbf{x}_0$; $\theta_0 \leftarrow 1$; $\eta_0 \leftarrow 1$;

for $t = 1, \dots, T$ **do**

$\mathbf{x}_t \leftarrow \mathbf{prox}_{\lambda_{\mathbf{x}}f_2}(\mathbf{v}_{t-1} - \gamma_t \nabla f_1(\mathbf{v}_{t-1}))$;

$\theta_t \leftarrow \arg \min_{\{\theta | \theta\mathbf{x}_t \in \mathcal{X}\}} \left\{ f(\theta\mathbf{x}_t) + \frac{1}{2\lambda_{\theta}} \|\theta\mathbf{x}_t - \mathbf{x}_t\|_2^2 \right\}$;

$\eta_t = \frac{\sqrt{\eta_{t-1}^4 + 4\eta_{t-1}^2 - \eta_{t-1}^2}}{2}$;

Update \mathbf{v}_t using either Eq. 3 or Eq. 4;

end

return $\theta_T\mathbf{x}_T$;

applying the proximity operator to a gradient descent step.

Alg. 1 shows our RAPID algorithm, where in each iteration t ($t \geq 1$) four steps are involved:

Step 1: A proximal gradient step using the auxiliary variable \mathbf{v}_{t-1} , same as APG.

Step 2: A simple line search step along the direction of the current solution \mathbf{x}_t . Actually the definition of θ_t in Alg. 1 is equivalent to the following equation:

$$\mathbf{x}_t^* = \mathbf{prox}_{\lambda_{\theta}f}(\mathbf{x}_t) \in \{\tilde{\mathbf{x}}_t | \exists \theta, \tilde{\mathbf{x}}_t = \theta\mathbf{x}_t \in \mathcal{X}\}. \quad (2)$$

In other words, this line search step essentially *adapts* the current solution \mathbf{x}_t to a better one in a very efficient way (*e.g.* with close-form solutions).

Step 3: Updating parameter η_t used for constructing the new auxiliary variable \mathbf{v}_t , same as APG. Note that any number sequence $\{\eta_t\}$ can be used here as long as the sequence satisfies $\forall t \leq 1, \frac{1-\eta_{t+1}}{\eta_{t+1}^2} \leq \frac{1}{\eta_t^2}$.

Step 4: Updating the new auxiliary variable \mathbf{v}_t using one of the following two equations:

$$\mathbf{v}_t = \eta_t(1 - \eta_{t-1}^{-1})\theta_{t-1}\mathbf{x}_{t-1} + \eta_t\eta_{t-1}^{-1}\mathbf{x}_t + (1 - \eta_t)\theta_t\mathbf{x}_t, \quad (3)$$

$$\mathbf{v}_t = \eta_t(1 - \eta_{t-1}^{-1})\theta_{t-1}\mathbf{x}_{t-1} + (1 - \eta_t + \eta_t\eta_{t-1}^{-1})\theta_t\mathbf{x}_t. \quad (4)$$

In this way, our algorithm guarantees its convergence, but with different convergence rate. See our analysis in Section 3 and comparison results in Section 4.

Fig. 1(b) illustrates the differences in constructing the auxiliary variable between APG and our RAPID. In APG, the auxiliary variable \mathbf{v}_t is constructed along the gradient of $\mathbf{x}_t - \mathbf{x}_{t-1}$ starting from \mathbf{x}_t . Similarly, in RAPID we would like to construct the auxiliary variable \mathbf{v}_t using $\theta_t\mathbf{x}_t$ and the other intermediate solutions in the previous and current iterations (*i.e.* \mathbf{x}_{t-1} , $\theta_{t-1}\mathbf{x}_{t-1}$, and \mathbf{x}_t). It turns out that *all possible combinations of intermediate solutions for constructing \mathbf{v}_t end up with Eq. 3*, with guaranteed better upper bounds over $f(\theta_t\mathbf{x}_t) - \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$ than those over $f(\mathbf{x}_t) - \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$ in APG in arbitrary iteration t (see Theorem 1 in Section 3). Under mild conditions, we can adopt the same way as APG to construct \mathbf{v}_t using the final solutions in the previous and current iterations, *i.e.* $\theta_{t-1}\mathbf{x}_{t-1}$ and $\theta_t\mathbf{x}_t$, which is exactly Eq. 4. However, for this setting we lose the theoretical guarantee of better upper bounds than APG, as shown in Theorem 2 in Section 3. Nevertheless, surprisingly in our experiments RAPID with Eq. 4 outperforms than that with Eq. 3 with significant improvement in terms of empirical convergence (see Section 4 for details).

3. ANALYSIS

In this section, we present our theoretical results on the convergence rate of our RAPID algorithm in Theorem 1 and 2, which is clearly better than those of conventional APG methods such as FISTA, leading to faster convergence in practice. Please refer to [20] for the details of our proofs.

Lemma 1 (Sandwich [21]). *Let \tilde{f} be the linear approximation of f in \mathbf{v} w.r.t. f_1 , *i.e.* $\tilde{f}(\mathbf{w}; \mathbf{v}) = f_1(\mathbf{v}) + \langle \nabla f(\mathbf{v}), \mathbf{w} - \mathbf{v} \rangle + f_2(\mathbf{w})$, where $\langle \cdot, \cdot \rangle$ denotes the inner product between two vectors. Then*

$$f(\mathbf{w}) \leq \tilde{f}(\mathbf{w}; \mathbf{v}) + \frac{L}{2}\|\mathbf{w} - \mathbf{v}\|_2^2 \leq f(\mathbf{w}) + \frac{L}{2}\|\mathbf{w} - \mathbf{v}\|_2^2. \quad (5)$$

Lemma 2 (3-Point Property [21]). *If $\hat{\mathbf{w}} = \arg \min_{\mathbf{w} \in \mathbb{R}^d} \frac{1}{2}\|\mathbf{w} - \mathbf{w}_0\|_2^2 + \phi(\mathbf{w})$, then for any $\mathbf{w} \in \mathbb{R}^d$,*

$$\phi(\hat{\mathbf{w}}) + \frac{1}{2}\|\hat{\mathbf{w}} - \mathbf{w}_0\|_2^2 \leq \phi(\mathbf{w}) + \frac{1}{2}\|\mathbf{w} - \mathbf{w}_0\|_2^2 - \frac{1}{2}\|\mathbf{w} - \hat{\mathbf{w}}\|_2^2. \quad (6)$$

Lemma 3. *In Alg. 1, at an arbitrary iteration t , we have*

$$f(\theta_t\mathbf{x}_t) \leq f(\mathbf{x}_t) - \frac{1}{2\lambda_\theta}\|\mathbf{x}_t - \theta_t\mathbf{x}_t\|_2^2. \quad (7)$$

Theorem 1. *Let $\mathbf{x}^* = \arg \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$ and $\lambda_\theta = \frac{1}{L}$. If Alg. 1 updates \mathbf{v}_t using Eq. 3, in iteration T ($\forall T \geq 1$) in Alg. 1,*

we have

$$f(\theta_T\mathbf{x}_T) - f(\mathbf{x}^*) \leq \frac{2L}{(T+1)^2} \left[\|\mathbf{x}^* - \mathbf{z}_0\|_2^2 - \sum_{t=1}^T \frac{\|\mathbf{x}_t - \theta_t\mathbf{x}_t\|_2^2}{\eta_{t-1}^2} \right], \quad (8)$$

where \mathbf{z}_0 is a constant.

Theorem 2. *Let $\mathbf{x}^* = \arg \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$ and $\lambda_\theta = \frac{1}{L}$. If Alg. 1 updates \mathbf{v}_t using Eq. 4, and suppose in any iteration t ($t \geq 1$), $\mathbf{v}_t^* = (1 - \eta_t)\theta_t\mathbf{x}_t + \eta_t\mathbf{x}^*$ and $\|\mathbf{v}_t^* - \mathbf{x}_{t+1}\|_2^2 = \|\mathbf{v}_t^* - \theta_{t+1}\mathbf{x}_{t+1}\|_2^2 + \xi_{t+1}$, then in iteration T ($\forall T \geq 1$) we have*

$$f(\theta_T\mathbf{x}_T) - f(\mathbf{x}^*) \leq \frac{2L}{(T+1)^2} \left[\|\mathbf{x}^* - \mathbf{z}_0\|_2^2 - \sum_{t=1}^T \left(\frac{\|\mathbf{x}_t - \theta_t\mathbf{x}_t\|_2^2}{\eta_{t-1}^2} + \xi_t \right) \right], \quad (9)$$

where \mathbf{z}_0 is a constant.

4. EXPERIMENTS

In this section, we explain how to apply our algorithm to solve sparse linear regression (*i.e.* LASSO, group LASSO, and least square fitting with trace-norm). We also compare our empirical performance with APG and an existing solver SLEP [19].

4.1. Problem Setting

We denote $\mathbf{A} = \{\mathbf{a}_i\}_{i=1, \dots, N} \in \mathbb{R}^{N \times d}$ as a data matrix, $\mathbf{y} = \{y_i\}_{i=1, \dots, N} \in \mathbb{R}^N$ as a regression target vector, $\mathbf{Y} = \{\mathbf{y}_j\}_{j=1, \dots, M} \in \mathbb{R}^{N \times M}$ as a matrix consisting of M regression tasks, $\mathbf{x} \in \mathbb{R}^d$ as a linear regressor, $\mathbf{X} = \{\mathbf{x}_j\}_{j=1, \dots, M} \in \mathbb{R}^{d \times M}$ as a matrix consisting of M linear regressors, and $\lambda \geq 0$ as a regularization parameter. As follows, for each method we list its loss function, regularizer, proximity operation, and optimal line search scalar θ , which are used in Alg. 1.

(i) Loss functions (*i.e.* f_1 , **convex and differentiable**). Least square loss is used in all the methods, *i.e.* $\frac{1}{2}\|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2$ for LASSO and group LASSO, and $\frac{1}{2}\sum_{j=1}^M \|\mathbf{A}\mathbf{x}_j - \mathbf{y}_j\|_2^2$ for trace-norm.

(ii) Regularizers (*i.e.* f_2 , **convex but non-differentiable**). The corresponding regularizers in LASSO, group LASSO, and trace-norm are $\lambda\|\mathbf{x}\|_1$, $\lambda\sum_{g \in \mathcal{G}} \|\mathbf{x}^{(g)}\|_2$, and $\lambda\|\mathbf{X}\|_* = \text{trace}(\sqrt{\mathbf{X}^T\mathbf{X}}) = \sum_{j=1}^{\min(d, M)} \sigma_j$, respectively. Here, $\|\cdot\|_1$ and $\|\cdot\|_2$ denote ℓ_1 and ℓ_2 norms, $(\cdot)^T$ denotes the matrix transpose operator, $g \in \mathcal{G}$ denotes a group index, $\mathbf{x}^{(g)}$ denotes a group of variables without overlaps, and σ_j denotes the j^{th} singular value for matrix $\sqrt{\mathbf{X}^T\mathbf{X}}$.

(iii) Proximity operators. Based on their corresponding regularizers, the proximity operators can be calculated efficiently as follows, where $\mathbf{u} = \{u_j\}$ and $\mathbf{U} = \{\mathbf{u}_j\}$ are denoted as the variable vector and matrix after the gradient descent:

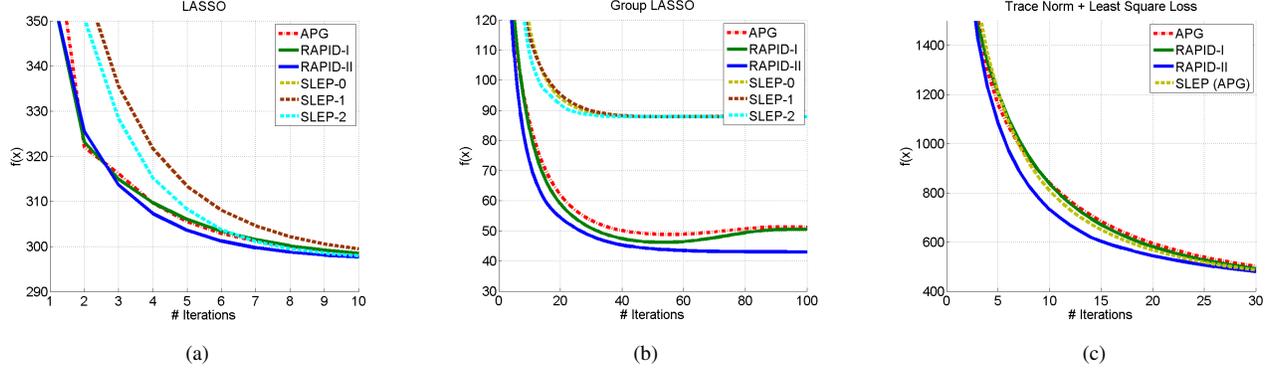


Fig. 2. Empirical convergence rate comparison using synthetic data on (a) LASSO, (b) group LASSO, and (c) trace-norm. This figure is best viewed in color.

- LASSO:

$$\text{prox}_{\lambda f_{\|\cdot\|_1}}(u_j) = \text{sign}(u_j) \cdot \max\{0, |u_j| - \lambda\}, \quad (10)$$

where $\text{sign}(u_j) = 1$ if $u_j \geq 0$, otherwise, $\text{sign}(u_j) = -1$, and $|u_j|$ denotes its absolute value.

- Group LASSO:

$$\text{prox}_{\lambda f_{\|\cdot\|_2}}(u_j^{(g)}) = \frac{u_j^{(g)}}{\|u^{(g)}\|_2} \cdot \max\{0, \|u^{(g)}\|_2 - \lambda\}. \quad (11)$$

- Trace norm: Letting $\mathbf{U} = \mathbf{P} \cdot \text{diag}(\boldsymbol{\sigma}) \cdot \mathbf{Q}^T$, where $\mathbf{P} \in \mathbb{R}^{d \times n}$ and $\mathbf{Q} \in \mathbb{R}^{M \times n}$ are two matrices, $\boldsymbol{\sigma}$ is a vector with n singular values of $\mathbf{U} \in \mathbb{R}^{d \times M}$, and $\text{diag}(\cdot)$ denotes a diagonal matrix with the values along the diagonal. Then we have

$$\text{prox}_{\lambda f_{\|\cdot\|_*}}(\mathbf{U}) = \mathbf{P} \cdot \text{prox}_{\lambda f_{\|\cdot\|_1}}(\boldsymbol{\sigma}) \cdot \mathbf{Q}^T. \quad (12)$$

Here $\text{prox}_{\lambda f_{\|\cdot\|_1}}(\cdot)$ is an entry-wise operator.

(iv) **Optimal line search scalar θ .** For each problem, we re-define θ as $\theta_t = \arg \min_{\theta} \{f(\theta \mathbf{x}_t)\}$ in arbitrary iteration t by setting $\lambda_{\theta} = +\infty$ in Alg. 1, because there exists a close-form solution in our cases. Letting $\forall t, \frac{\partial f(\theta \mathbf{x}_t)}{\partial \theta} = 0$, we have $\theta_t = \frac{\mathbf{y}^T \mathbf{A} \mathbf{x}_t - f_2(\mathbf{x}_t)}{\|\mathbf{A} \mathbf{x}_t\|_2^2}$ for LASSO or group LASSO, and $\theta_t = \frac{\sum_j \mathbf{y}_j^T \mathbf{A}(\mathbf{x}_j)_t - \lambda \|\mathbf{X}_t\|_*}{\sum_j \|\mathbf{A}(\mathbf{x}_j)_t\|_2^2}$ for trace-norm.

4.2. Results

We test and compare our RAPID algorithm on some synthetic data. For each sparse linear regression method, we generate a 10^3 sample data matrix with 10^3 dimensions per sample as variable \mathbf{A} , and its associated regression target vector (matrix) \mathbf{y} (\mathbf{Y}) randomly by normal distributions. The APG and RAPID methods are modified based on the code in [14]¹, and

¹The code can be downloaded from <https://github.com/bodono/apg>. We do not use the re-starting scheme in the code.

SLEP [19] is a widely used sparse learning toolbox for our comparison. Here, RAPID-I and RAPID-II denote our algorithm using Eq. 3 and Eq. 4, respectively. SLEP-0, SLEP-1, and SLEP-2 are the three settings used in SLEP with different parameters $mFlag$ and $lFlag$ (i.e. $(mFlag, lFlag) = \text{SLEP-}\#$: $(0,0)=0$, $(0,1)=1$, $(1,1)=2$). Please refer to the toolbox manual for more details. For trace-norm, SLEP actually implements the APG algorithm as its solver.

Fig. 2 shows our comparison results. In Fig. 2(a), the performances of SLEP-0 and SLEP-1 are identical, and thus there is only one curve (i.e. the brown one) for both methods. Clearly, our RAPID-II algorithm works best in these three cases in terms of empirical convergence rate. Since we can only guarantee that the upper bound of the difference between the current and optimal objective values in each iteration in RAPID is no bigger than that in APG (see Section 3), sometimes the actual objective value using RAPID may be larger than that using APG, as shown in Fig. 2(a).

5. CONCLUSION

In this paper, we propose an improved APG algorithm, namely, Rapidly Accelerated Proximal Gradient (RAPID), to speed up the convergence of conventional APG algorithms. Our first idea is to introduce a new line search step after the proximal gradient step in APG to push the current solution \mathbf{x}_t towards a new one $\theta \mathbf{x}_t \in \mathcal{X}(\theta > 0)$ so that $f(\theta \mathbf{x}_t)$ is minimized over scalar θ . Our second idea is to propose two different ways of constructing the auxiliary variable in APG using the intermediate solutions in the previous and current iterations. In this way, we can prove that our algorithm is guaranteed to converge with a smaller upper bound of the gap between the current and optimal objective values than those in APG algorithms. We demonstrate our algorithm with the application of sparse linear regression. In summary, our RAPID converges faster than APG, in general, and for some problems RAPID based algorithms can be comparable with the sophisticated existing solvers.

6. REFERENCES

- [1] Amir Beck and Marc Teboulle, “A fast iterative shrinkage-thresholding algorithm for linear inverse problems,” *SIAM J. Img. Sci.*, vol. 2, no. 1, pp. 183–202, Mar. 2009.
- [2] Patrick L. Combettes and Jean-Christophe Pesquet, “Proximal splitting methods in signal processing,” 2011.
- [3] Chenglong Bao, Yi Wu, Haibin Ling, and Hui Ji, “Real time robust l1 tracker using accelerated proximal gradient approach,” in *CVPR*, 2012, pp. 1830–1837.
- [4] Tianyi Zhou, Dacheng Tao, and Xindong Wu, “Nesvm: A fast gradient method for support vector machines,” in *ICDM*, 2010, pp. 679–688.
- [5] A.S. Nemirovsky and D.B. Yudin, “Problem complexity and method efficiency in optimization,” *Mathematics and Computers in Simulation*, vol. 26, no. 1, pp. 75 –, 1984.
- [6] Mark W. Schmidt, Nicolas Le Roux, and Francis Bach, “Convergence rates of inexact proximal-gradient methods for convex optimization,” in *NIPS*, 2011, pp. 1458–1466.
- [7] Yves F Atchade, Gersende Fort, and Eric Moulines, “On stochastic proximal gradient algorithms,” *arXiv preprint arXiv:1402.2365*, 2014.
- [8] Lorenzo Rosasco, Silvia Villa, and Bang Công Vũ, “Convergence of stochastic proximal gradient algorithm,” *arXiv preprint arXiv:1403.5074*, 2014.
- [9] Annie I Chen and Asuman Ozdaglar, “A fast distributed proximal-gradient method,” in *Communication, Control, and Computing (Allerton), 2012 50th Annual Allerton Conference on*. IEEE, 2012, pp. 601–608.
- [10] Qihang Lin and Lin Xiao, “An adaptive accelerated proximal gradient method and its homotopy continuation for sparse optimization,” in *ICML*, 2014, pp. 73–81.
- [11] Katya Scheinberg, Donald Goldfarb, and Xi Bai, “Fast first-order methods for composite convex optimization with backtracking,” *Foundations of Computational Mathematics*, vol. 14, no. 3, pp. 389–417, 2014.
- [12] Lin Xiao and Tong Zhang, “A proximal-gradient homotopy method for the sparse least-squares problem,” *SIAM Journal on Optimization*, vol. 23, no. 2, pp. 1062–1091, 2013.
- [13] Yurii Nesterov, “Gradient methods for minimizing composite functions,” *Math. Program.*, vol. 140, no. 1, pp. 125–161, 2013.
- [14] Brendan ODonoghue and Emmanuel Candès, “Adaptive restart for accelerated gradient schemes,” *Foundations of Computational Mathematics*, pp. 1–18, 2013.
- [15] R. Tibshirani, “Regression shrinkage and selection via the lasso,” *Journal of the Royal Statistical Society (Series B)*, vol. 58, pp. 267–288, 1996.
- [16] Ming Yuan and Yi Lin, “Model selection and estimation in regression with grouped variables,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 68, no. 1, pp. 49–67, 2006.
- [17] Shuiwang Ji and Jieping Ye, “An accelerated gradient method for trace norm minimization,” in *ICML*, 2009, pp. 457–464.
- [18] Paul Tseng, “Approximation accuracy, gradient methods, and error bound for structured convex optimization,” *Math. Program.*, vol. 125, no. 2, pp. 263–295, 2010.
- [19] J. Liu, S. Ji, and J. Ye, *SLEP: Sparse Learning with Efficient Projections*, Arizona State University, 2009.
- [20] Ziming Zhang and Venkatesh Saligrama, “Rapid: Rapidly accelerated proximal gradient algorithms for convex minimization,” *CoRR*, 2014.
- [21] Luca Baldassarre and Massimiliano Pontil, “Advanced topics in machine learning part II 5. proximal methods,” University Lecture, http://www0.cs.ucl.ac.uk/staff/l.baldassarre/lectures/baldassarre_proximal_methods.pdf.