
Adaptive Balancing of Gradient and Update Computation Times using Global Geometry and Approximate Subproblems

Sai Praneeth Karimireddy
EPFL

Sebastian U. Stich
EPFL

Martin Jaggi
EPFL

Abstract

First-order optimization methods comprise two important primitives: i) the computation of gradient information and ii) the computation of the update that leads to the next iterate. In practice there is often a wide mismatch between the time required for the two steps, leading to underutilization of resources. In this work, we propose a new framework, Approx Composite Minimization (ACM) that uses *approximate* update steps to ensure balance between the two operations. The accuracy is *adaptively* chosen in an online fashion to take advantage of changing conditions. Our unified analysis for approximate composite minimization generalizes and extends previous work to new settings. Numerical experiments on Lasso regression and SVMs demonstrate the effectiveness of the novel scheme.

1 Introduction

In the last decade, first-order methods and especially stochastic first-order methods have proven to be the superior choice to solve problems of very large size that arise in modern machine learning applications (cf. [27]). The state of the art comprises i) stochastic sub-gradient methods [28], ii) variance reduced stochastic gradient methods [2, 8, 12, 15, 20], iii) coordinate methods (e.g. dual ascent) such as [9, 11, 18, 19, 24, 29, 30].

In a different line of research, distributed algorithm variants have been developed for larger datasets which exceed the capacity of a single machine or device. Here the state-of-the-art algorithms are [16, 31] based on a

primal-dual structure of the problem and [10, 13, 22] which work in the primal alone.

All of these methods are iterative algorithms that can be described in the following framework: in each iteration t , the algorithm i) constructs a *model* $m_t(\mathbf{x}): \mathbb{R}^d \rightarrow \mathbb{R}$ of the objective function (m_t is an approximation of the objective function) and then ii) computes an update step by minimizing this model. The model is typically constructed from gradient information (full gradient, a subset of its coordinates, or a stochastic approximation) and accounting for smoothness assumptions and the structure of the regularizers.

In this work, we are studying the cases when there is a *mismatch* between the time required to compute the gradient information and the time required to compute the update step (i.e. the optimization of the model). We will assume that the cost of the steps involved in computing the gradient information is fixed (i.e. following from the specification of the optimization problem) and outside of our control. However, a parameter that is in our control is the time we spend optimizing the model in each iteration. That is, we will show that it is not required to solve the model *exactly*, but it is enough to compute an *approximate* solution in each iteration. By making the accuracy a *tunable parameter*, we ensure that we do not spend too much time computing the update. Therefore, we optimally balance the two steps.

Whilst the idea of using approximations to speed up the computation is not new (cf. e.g. [7, 26, 33]), the idea of exploiting this for balancing the computational cost is novel. We also extend the framework of [33] to several new settings (Table 1).

Paper outline and contributions. In Section 2 we specify the problem setting and present two motivating examples that exemplify the need of sufficiently complex models—we propose to use models that take *global geometry* (curvature) into account. Section 3 outlines the proposed unified ACM framework and its convergence analysis is given in Section 4. Section 5 extends the framework to the empirical risk minimiza-

Table 1: Summary of the settings where our framework is applicable.

Method	Approximate subproblems	Composite functions	Strongly convex	General convex	Dual version	Parallel block-coordinate updates
PCDM [24]	✗	✓	✓	✓	✗	✓
Inexact [33]	✓	✓	✓	✓	✗	✗
SDNA [21]	✗	✓	✓	✗	✓	✗
PSNM [17]	✗	✗	✓	✗	✗	✓
ACM	§3, (Def 3)	✓	§4 (Thm 1)	§4 (Thm 2)	§5 (Thms 3,4)	§6 (Lem 4 with Thms 1,2,3, and 4)

tion setting where we prove primal-dual guarantees. In Section 6 we demonstrate the generality of the ACM framework by exemplarily considering specific solvers for the subproblems and we show how the convergence guarantees of these algorithms easily follow from our analysis. ACM specifically supports *changing accuracies* for different iterations. In Section 7 we capitalize this property by designing mechanisms that *adaptively control* the subproblem accuracy, always balancing the computation times. In Section 8 we present experimental results that show the numerical advantage our schemes and conclude in Section 9. All missing proofs and figures can be found in the appendix.

2 Setup and Motivation

We address optimization problems of the form

$$\min_{\mathbf{x} \in \mathbb{R}^d} \left[F(\mathbf{x}) \stackrel{\text{def}}{=} f(\mathbf{x}) + g(\mathbf{x}) \right], \quad (1)$$

where $f: \mathbb{R}^d \rightarrow \mathbb{R}$ is a smooth convex function and $g: \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$ is an arbitrary extended-valued convex function. $\mathbf{x}^* \in \mathbb{R}^d$ denotes an optimal solution, and for $\epsilon > 0$, a point $\mathbf{y} \in \mathbb{R}^d$ with $F(\mathbf{y}) - F(\mathbf{x}^*) \leq \epsilon$ is an ϵ -approximate solution.

2.1 Imbalance in the computations

In this section we illustrate the main problem we tackle in this paper—inefficiencies caused by mismatches in the computation times of different steps in the optimization algorithm. First-order methods typically optimize an over-approximation $U: \mathbb{R}^n \rightarrow \mathbb{R}$ of the objective (1), with

$$U(\Delta\mathbf{x}) \stackrel{\text{def}}{=} u(\mathbf{x} + \Delta\mathbf{x}) + g(\mathbf{x} + \Delta\mathbf{x}) \geq F(\mathbf{x} + \Delta\mathbf{x}) \quad (2)$$

where $u(\mathbf{x} + \Delta\mathbf{x}) \stackrel{\text{def}}{=} f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \Delta\mathbf{x} \rangle + \frac{L}{2} \|\Delta\mathbf{x}\|_2^2$ and L is the smoothness parameter of f .

We will now discuss two key examples of a significant imbalance in the gradient computation vs. the optimization of the model (2).

Case A: Gradient computation is slower. Consider the coordinate descent algorithm on the L1-regularized logistic regression problem. Given an $n \times d$

data matrix A , let A_i refer to its i -th column and denote the residual as $\mathbf{v} \stackrel{\text{def}}{=} A\mathbf{x}$. Then the logistic loss can be written as $F(\mathbf{x}) = \sum_{i=1}^n \log(1 + e^{v_i}) + \|\mathbf{x}\|_1$ and the i -th coordinate of the gradient $\nabla f(\mathbf{x})$ is $\nabla_i f(\mathbf{x}) = \langle A_i, \nabla l_{\log}(\mathbf{v}) \rangle$, where

$$\nabla l_{\log}(\mathbf{v}) = \left(\frac{1}{1+e^{-v_1}}, \dots, \frac{1}{1+e^{-v_n}} \right).$$

Assuming that \mathbf{v} is already stored in active memory, computing the i -th gradient requires i) fetching A_i from memory, ii) performing n exponentiations and divisions, iii) one dot product, and iv) possibly communicating back the update steps to facilitate the next gradient computation. The coordinate update step on the other hand is $[\mathbf{x}_i - \frac{1}{L} \nabla_i f(\mathbf{x})]_{\lambda/L}$ where $[c]_{\gamma}$ is the *shrinkage* operator. If $\Delta\mathbf{x}_i$ is the update made to coordinate i , \mathbf{v} can be updated as $\mathbf{v} - \Delta x_i A_i$. In total for minimizing the model we only make i) one coordinate update and ii) one vector addition, in contrast with the more involved gradient coordinate computation. Thus, cache misses, complex operations, as well as communication overhead all lead to the gradient computation being much slower than the update step. Increasing the number of coordinates being updated i.e. block-coordinate updates can alleviate overhead to some extent due to cache misses, though it is ineffective against the other sources of delay.

Case B: Model minimization is slower. When the data is not extremely high dimensional and the loss function f is simple, it is quite fast to compute the gradient. For example, a very common task in image processing is to reconstruct the original image given a corrupted linear measurement y . For such applications, the commonly used loss function is $F(\mathbf{x}) = \frac{1}{2} \|\mathbf{Ax} - \mathbf{b}\|_2^2 + \lambda \|\mathbf{x}\|_{TV}$ [5]. Even though f is just a quadratic, there is no closed form solution for minimizing $U(\Delta\mathbf{x})$ because of $g(\mathbf{x}) = \|\mathbf{x}\|_{TV}$. In this case a lot of time is spent in minimizing the model $U(\Delta\mathbf{x})$ whereas the gradient computation is fast.

2.2 A solution: curvature models

When the model is too “simple”, i.e. computationally cheap to minimize compared to the expensive gradient information, then traditional methods will waste

significant resources for just the gradient information. Hence, the model should be sufficiently sophisticated in order to be able to *extract the maximum progress* using a single gradient computation. In other words, we would prefer that exact minimization of the model takes more time than computing the gradient information. In this case, we can minimize the model approximately to a *tunable degree of accuracy* to ensure that the two essential parts are balanced in terms of computational cost. The model, of course, cannot be arbitrarily hard to optimize either. In particular, it should be at least possible to make non-zero progress on minimizing the model in time less than it takes to compute the gradient information.

One way to create good models is to capture some second-order information of the objective, i.e. the global geometry (curvature) using a $d \times d$ matrix M as for instance in [21, 33]. This means that we replace u in (2) with

$$u_M(\mathbf{x} + \Delta\mathbf{x}) \stackrel{\text{def}}{=} f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \Delta\mathbf{x} \rangle + \frac{1}{2} \Delta\mathbf{x}^\top M \Delta\mathbf{x}.$$

For example, in the L1-regularized logistic regression discussed before, we can use $M \stackrel{\text{def}}{=} \frac{1}{4} A^\top A$.

Related literature. There are a number of algorithms which incorporate curvature information. These include methods based on using the diagonal of the Hessian information to compute the sample probabilities as well as the step size [3, 18, 19, 24]. Recently, a more direct approach to incorporating the Hessian information through *preconditioning* has become more popular [17, 32, 33]. However, such preconditioning means that the computational effort involved in solving the subproblem is significant. To overcome this, one resorts to approximate solutions [32, 33]. The idea of preconditioning has also been extended to dual ascent algorithms as in [21]. However, there exact solutions to the subproblems were required.

2.3 Notation and Definitions

For a fixed positive semi-definite matrix M ($M \succcurlyeq 0$), we use $\|\mathbf{x}\|_M^2 \stackrel{\text{def}}{=} \mathbf{x}^\top M \mathbf{x}$. With this semi-norm, the regularity assumptions on f can be written as follows.

Definition 1 (M -smoothness). *A differentiable function $h: \mathbb{R}^d \rightarrow \mathbb{R}$ is M -smooth with respect to a fixed matrix $M \succcurlyeq 0$ if for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$*

$$h(\mathbf{y}) \leq h(\mathbf{x}) + \langle \nabla h(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{1}{2} \|\mathbf{y} - \mathbf{x}\|_M^2.$$

Definition 2 (λ -strong convexity). *A differentiable function $h: \mathbb{R}^d \rightarrow \mathbb{R}$ is λ -strongly convex w.r.t. $M \succcurlyeq 0$ and $\lambda > 0$ if for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$*

$$h(\mathbf{y}) \geq h(\mathbf{x}) + \langle \nabla h(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{\lambda}{2} \|\mathbf{y} - \mathbf{x}\|_M^2.$$

Note that an M -smooth function can only be λ -strongly convex w.r.t. M for $\lambda \in (0, 1]$.

3 Approx Composite Minimization

In this section we describe the general method as outlined in Algorithm 1. At each time step t , an approximation $m_t(\Delta\mathbf{x}; M) \geq F(\mathbf{x}_t + \Delta\mathbf{x}) - F(\mathbf{x}_t)$ is constructed as follows:

$$m_t(\Delta\mathbf{x}; \mathbf{x}_t, M) \stackrel{\text{def}}{=} \langle \nabla f(\mathbf{x}_t), \Delta\mathbf{x} \rangle + \frac{1}{2} \|\Delta\mathbf{x}\|_M^2 + g(\mathbf{x}_t + \Delta\mathbf{x}) - g(\mathbf{x}_t). \quad (3)$$

When obvious from context, we drop M and \mathbf{x}_t and simply refer it as $m_t(\Delta\mathbf{x})$. Let $m_t(\Delta\mathbf{x}_t^*) \stackrel{\text{def}}{=} m_t^*$ be the minimum of the subproblem obtained at $\Delta\mathbf{x}_t^*$.

Definition 3. *We denote by $\Theta_t \in (0, 1]$ the relative accuracy to which the subproblem (3) is solved at step t , i.e., we compute $\Delta\mathbf{x}_t$ such that*

$$m_t(\Delta\mathbf{x}_t) - m_t^* \leq (1 - \Theta_t)(m_t(\mathbf{0}) - m_t^*).$$

Here $\Theta_t = 1$ means that we solve the problem exactly. We then update the iterate as $\mathbf{x}_{t+1} := \mathbf{x}_t + \Delta\mathbf{x}_t$. Note that Θ_t can adaptively change with each step.

Algorithm 1: Approx Composite Minimization (ACM)

Input: M, \mathbf{x}_0
for $t = 0, \dots$ **do**
 Let $m_t(\Delta\mathbf{x}) \stackrel{\text{def}}{=} \langle \nabla f(\mathbf{x}_t), \Delta\mathbf{x} \rangle + \frac{1}{2} \|\Delta\mathbf{x}\|_M^2 + g(\mathbf{x}_t + \Delta\mathbf{x}) - g(\mathbf{x}_t)$
 Minimize subproblem: Find $\Delta\mathbf{x}_t$ such that
 $m_t(\Delta\mathbf{x}_t) - m_t^* \leq (1 - \Theta_t)(m_t(\mathbf{0}) - m_t^*)$
 $\mathbf{x}_{t+1} \leftarrow \mathbf{x}_t + \Delta\mathbf{x}_t$
end

4 Coverage Analysis

We will examine the cases when F is strongly convex (Def. 2 is satisfied for $\lambda > 0$) and the general convex case separately (when $\lambda = 0$) and generalize the results in [33] to account for varying approximation factors.

Theorem 1. *Given that f is M -smooth (1), and that f and g are λ_f and λ_g strongly convex respectively (2) for $\lambda_f + \lambda_g > 0$, then running Algorithm 1 gives linear convergence, i.e. $F(\mathbf{x}_T) - F(\mathbf{x}^*) \leq \epsilon$ for*

$$T \geq \frac{1 + \lambda_g}{(\lambda_f + \lambda_g) \tilde{\Theta}_T} \log \left(\frac{F(\mathbf{x}_0) - F(\mathbf{x}^*)}{\epsilon} \right),$$

where $\tilde{\Theta}_T = \frac{1}{T} \sum_{t=1}^T \Theta_t$ is the average accuracy.

Remark 1. *The convergence rate from Theorem 1 shows that the rate critically depends on not just λ_f but also λ_g . Thus M must be chosen to not just approximate f but also g , i.e. M must be chosen to approximate the curvature of F as closely as possible. This is an important observation since f and g typically have very different curvatures.*

Our proof hinges on this very useful lemma proved in the appendix.

Lemma 2. *Assuming g is λ_g -strongly convex, for any vector \mathbf{v} and $\lambda \in (0, 1]$, then*

$$\min_{\Delta \mathbf{x}} m(\Delta \mathbf{x}; \mathbf{v}, M) \leq \frac{\lambda + \lambda_g}{1 + \lambda_g} \cdot \min_{\Delta \mathbf{x}} m(\Delta \mathbf{x}; \mathbf{v}, \lambda M).$$

In the general convex case, we obtain a convergence rate of $O(1/T)$ which matches the rate of standard proximal-gradient algorithms. We again leave the proof of the theorem for the appendix.

Theorem 2. *Given that f is M -smooth (1), running Algorithm 1 such that at every step $m_t(\Delta \mathbf{x}_t) \leq 0$ and $\Theta_t > 0$ ensures*

$$F(\mathbf{x}_T) - F(\mathbf{x}^*) \leq \epsilon \quad \text{for } T > \frac{2D}{\bar{\Theta}_T \epsilon},$$

where D is at most the diameter of the initial level set of F and $\bar{\Theta}_T$ is the average accuracy:

$$D \stackrel{\text{def}}{=} \max_{\mathbf{y} | F(\mathbf{y}) \leq F(\mathbf{x}_0)} \|\mathbf{y} - \mathbf{x}^*\|_M^2 \quad \text{and} \quad \bar{\Theta}_T \stackrel{\text{def}}{=} \frac{1}{T} \sum_{t=0}^T \Theta_t.$$

5 Primal-Dual Guarantees

Suppose our objective function comes with the following additional structure, ubiquitous in machine learning and signal processing models:

$$\min_{\mathbf{w}} \sum_{i=1}^n l_i(\mathbf{w}^\top A_i) + \psi(\mathbf{w}).$$

Let A be a $d \times n$ data matrix. Here l_i is the loss defined for each data-point $A_i \in \mathbb{R}^d$ and ψ is a regularizer. We can define the dual objective in terms of the *Fenchel conjugate* of $\{l_i\}_i$ and ψ , $\{l_i^*\}_i$ and ψ^* respectively [9]:

$$\max_{\boldsymbol{\alpha}} \sum_{i=1}^n -l_i^*(-\boldsymbol{\alpha}_i) - \psi^*(A\boldsymbol{\alpha}).$$

These can be more compactly written as

$$\begin{aligned} \mathcal{O}_A(\boldsymbol{\alpha}) &\stackrel{\text{def}}{=} f(A\boldsymbol{\alpha}) + g(\boldsymbol{\alpha}), \\ \mathcal{O}_B(\mathbf{w}) &\stackrel{\text{def}}{=} f^*(\mathbf{w}) + g^*(-A^\top \mathbf{w}). \end{aligned} \quad (4)$$

Our objective to minimize the *duality gap* defined as $\mathcal{O}_B(\mathbf{w}) + \mathcal{O}_A(\boldsymbol{\alpha})$. We assume that $f(A\boldsymbol{\alpha})$ is M -smooth. Depending on the problem, it may be more convenient and efficient either to cast the problem as $\mathcal{O}_A(\boldsymbol{\alpha})$, or map it to $\mathcal{O}_B(\mathbf{w})$ and solve the dual. We can do this since primal-dual algorithms minimize both $\mathcal{O}_B(\mathbf{w})$ and $\mathcal{O}_A(\boldsymbol{\alpha})$ simultaneously. For more details about the setting and applications, we refer to the discussion in [9].

Algorithm 2: Approximate Dual Ascent (ADA)

Input: M, \mathbf{x}_0
Initialize: $\boldsymbol{\alpha}_0 \leftarrow \mathbf{0} \in \mathbb{R}^n, \mathbf{v}_0 \leftarrow \mathbf{0} \in \mathbb{R}^d$
for $t = 1, \dots$ **do**
 Let $m_t(\Delta \boldsymbol{\alpha}) \stackrel{\text{def}}{=} \langle \nabla f(\mathbf{v}_t), A\Delta \boldsymbol{\alpha} \rangle + \frac{1}{2} \|\Delta \boldsymbol{\alpha}\|_M^2 + g(\boldsymbol{\alpha}_t + \Delta \boldsymbol{\alpha}) - g(\boldsymbol{\alpha}_t)$
 Minimize subproblem: Find $\Delta \boldsymbol{\alpha}_t$ such that $m_t(\Delta \boldsymbol{\alpha}_t) - m_t^* \leq (1 - \Theta_t)(m_t(\mathbf{0}) - m_t^*)$
 Update steps:
 $\boldsymbol{\alpha}_{t+1} \leftarrow \boldsymbol{\alpha}_t + \Delta \boldsymbol{\alpha}_t$
 $\mathbf{v}_{t+1} \leftarrow \mathbf{v}_t + A\Delta \boldsymbol{\alpha}_t$
end

Just as in the Section 4, we can state two theorems—one for the case when the objective is strongly convex and one for the general convex case. The proofs of these theorems follow along the lines of [16, 31] with some simplifications.

Theorem 3. *For an objective of the form (4), let us assume that $f(A\boldsymbol{\alpha})$ is M -smooth (1), and that $f(A\boldsymbol{\alpha})$ and $g(\boldsymbol{\alpha})$ are λ_f and λ_g strongly convex respectively (2). Then running Algorithm 2 gives linear convergence i.e. for $\lambda = \frac{\lambda_g + \lambda_f}{1 + \lambda_g}$ and $\bar{\Theta}_T = \frac{1}{T} \sum_{t=0}^T \Theta_t$,*

$$\mathcal{O}_B(\nabla f(\mathbf{v}_t)) + \mathcal{O}_A(\boldsymbol{\alpha}_t) \leq \epsilon \quad \text{for}$$

$$T \geq \frac{1}{\lambda \bar{\Theta}_T} \log \left(\frac{(1 + \lambda_g)(\mathcal{O}_A(\mathbf{0}) - \mathcal{O}_A(\boldsymbol{\alpha}^*))}{\lambda_g \Theta_T \epsilon} \right).$$

Theorem 4. *Given that f is M -smooth (1), running Algorithm 2 ensures convergence at a rate of $O(\frac{1}{T})$ i.e. $\mathcal{O}_B(\mathbf{w}(\bar{\mathbf{v}}_t)) + \mathcal{O}_A(\bar{\boldsymbol{\alpha}}_t) \leq \epsilon$ for*

$$T \geq t_0 + \max \left[\frac{4D}{\epsilon}, \frac{1}{\hat{\Theta}_t} \right], \quad t_0 \geq \frac{4D}{\bar{\Theta}_{t_0} \epsilon},$$

where $\bar{\boldsymbol{\alpha}}_t \stackrel{\text{def}}{=} \frac{1}{t-t_0} \sum_{i=t_0+1}^t \boldsymbol{\alpha}_i$, $\bar{\mathbf{v}}_t \stackrel{\text{def}}{=} \frac{1}{t-t_0} \sum_{i=t_0+1}^t \mathbf{v}_i$, and $\mathbf{w}(\bar{\mathbf{v}}_t) \stackrel{\text{def}}{=} \nabla f(\bar{\mathbf{v}}_t)$. D is the diameter of the level set of $\mathcal{O}_A(\boldsymbol{\alpha})$, $Q \stackrel{\text{def}}{=} \{\boldsymbol{\alpha} \mid \mathcal{O}_A(\boldsymbol{\alpha}) \leq \mathcal{O}_A(\mathbf{0})\}$. $D \stackrel{\text{def}}{=} \max_{\mathbf{a}, \mathbf{b} \in Q} \|\mathbf{a} - \mathbf{b}\|_M^2$. Further $\tilde{\Theta}_t \stackrel{\text{def}}{=} \frac{1}{t} \sum_{t'=0}^t \Theta_{t'}$ and $\hat{\Theta}_t \stackrel{\text{def}}{=} \min_{t' \in [t]} \Theta_{t'}$.

Remark 3. *We obtain sharper bounds than in [16, 31] in both settings. In particular, we show that strong convexity constant λ_f of f is also useful for faster convergence. The rate in [31] is equivalent to setting $\lambda_f = 0$. Similarly, in the general convex setting, our rate is again simpler and has better constants.*

6 Extension to Coordinate Updates

Thus far in our discussion we have largely restricted ourselves to the case where the model was constructed

using the full gradient. However, when the function $g(\mathbf{x}) = \sum_{j=1}^d g_j(x_j)$ is coordinate-wise separable, a popular strategy is to update just a single coordinate in each iteration. This approach leads to state of the art algorithms for several problems [11, 18, 19, 24, 29, 30]. Coordinate methods have been widely successful not just due to their faster convergence, but also because they are less resource intensive and provide faster update times [35]. However, the imbalance between the time for gradient information computation and the time for update computation is often exacerbated in coordinate descent methods. This is because the time to compute one directional derivative of the gradient scales in general linearly with the dimension, whereas computing the one-dimensional update only requires constant time. Updating multiple coordinates simultaneously might reduce some of the overhead—such as the delay due to memory accesses or communication costs—but cannot completely settle this imbalance.

In this section we see how we can extend our ideas from Sections 3 and 5 seamlessly to this scenario and achieve a better balance between the gradient and update computation times. It is clear, that one single step of coordinate descent can be seen as an approximate solution to the full dimensional problem (3). Thus, the convergence results can be directly recovered from the theorems derived in the previous sections. However, there is a small caveat: the model (3) depends on the full gradient, whereas in coordinate descent methods it is sufficient to compute just the directional derivatives of gradient along the descent directions. To capture this more precisely, we will now refine our notation.

Typically, the number of coordinates being updated at each iteration is fixed by external factors—for e.g. when the data is distributed by columns amongst multiple machines, the coordinates being updated are constrained by data available on each machine, or in the single machine case it is constrained by the number of columns which fit in a cache block. For this reason we assume that the number of coordinates being updated is fixed. Further, we will not only support sequential updates, but also updates to multiple blocks of coordinates in *parallel*. This setting is useful when i) the data is distributed over multiple machines which can update a different block of coordinates in parallel, or ii) when there are multiple threads (in a multiprocessor) each with a separate cache. Examples of this setting can be found in [17, 31].

Suppose at iteration t , we have κ machines such that machine $k \in [\kappa]$ can compute the coordinates $\pi_k \subseteq [d]$ of the gradient i.e. it can compute

$$\nabla_{\pi_k} f(\mathbf{x}) \stackrel{\text{def}}{=} \sum_{j \in \pi_k} \nabla_j f(\mathbf{x}) \mathbf{e}_j.$$

Here the sets π_k need not necessarily need to form a complete partition of $[d]$. Note that this notation also captures the case of sequential updates on one single machine: in this case π_1 just denotes the set of coordinates that are updated in this iteration. We further assume that machine k has access to the principal submatrix of M corresponding to the coordinates π_k ,

$$M_{\pi_k} \stackrel{\text{def}}{=} \sum_{i,j \in \pi_k} M_{i,j} \mathbf{e}_i \mathbf{e}_j^\top.$$

Then we can form a set of κ subproblems such that each machine $k \in [\kappa]$ can solve $m_t^\sigma(\Delta \mathbf{x}; \pi_k)$ defined as

$$\begin{aligned} m_t^\sigma(\Delta \mathbf{x}; \pi_k) &\stackrel{\text{def}}{=} \langle \nabla_{\pi_k} f(\mathbf{x}), \Delta \mathbf{x} \rangle + \frac{\sigma}{2} \|\Delta \mathbf{x}\|_{M_{\pi_k}}^2 \\ &\quad + \sum_{j \in \pi_k} g_j(x_j + \Delta x_j) - g_j(x_j). \end{aligned}$$

Here $\sigma \in [1, \kappa]$ is a parameter which measures the *separability* of the matrix M . Crucially, the problems $m_t^\sigma(\Delta \mathbf{x}; \pi_k)$ for $k \in \{1, \dots, \kappa\}$ can be solved in parallel. The solutions from these subproblems can then be combined as

$$\mathbf{x}_{t+1} = \mathbf{x}_t + \sum_{k=1}^{\kappa} [\Delta \mathbf{x}]_{\pi_k},$$

where $[\Delta \mathbf{x}]_{\pi_k}$ is a Θ_t -approximate solution to $m_t^\sigma(\Delta \mathbf{x}; \pi_k)$ as in Definition 3. In this manner we can use global geometry (curvature) information encoded in M to better utilize the gradient information we computed even in the block coordinate setting.

We can relate the Θ_t progress on the coordinate models $m_t^\sigma(\Delta \mathbf{x}; \pi_k)$ to the progress on the global model $m_t(\Delta \mathbf{x})$. Recall that $m_t(\Delta \mathbf{x})$ was our shorthand for $m_t(\Delta \mathbf{x}; \mathbf{x}_t, M)$ defined in (3). A more formal study with additional technical details and definitions is relegated to Section A.1.

Lemma 4. *Suppose that $[\Delta \mathbf{x}]_{\pi_k}$ is an Θ_t -approximate solution to $m_t^\sigma(\Delta \mathbf{x}; \pi_k)$ for $k \in [\kappa]$ and $\Delta \mathbf{x}_t = \sum_{k=1}^{\kappa} [\Delta \mathbf{x}]_{\pi_k}$. Then under the conditions specified in Lemma 8 in Section A.1,*

$$\mathbb{E}[m_t(\Delta \mathbf{x}_t) - m_t^*] \leq \left(1 - \frac{\kappa \Theta_t}{m \sigma \nu}\right) (m_t(\mathbf{0}) - m_t^*),$$

where the expectation is over the random selection of the sets $\{\pi_1, \dots, \pi_\kappa\}$, and $m \in [1, d]$ and $\nu \geq 1$ are parameters depending on the sampling (see Definitions 4 and 6 in Section A.1).

Thus, using just the block-coordinate models, we can make progress on the global problem $m_t(\Delta \mathbf{x}_t; \mathbf{x}_t, M)$. We can then combine Lemma 4 with Theorems 1, 2, 3, and 4 to derive the corresponding parallel-block coordinate versions of the algorithm.

7 Adaptive Accuracy

Requiring only approximate solutions means that we can use iterative methods for the sub-problems which can be much faster and cheaper than exact solvers [33]. More importantly, using iterative algorithms enables us to adaptively control the number of iterations run on each of the subproblem. This allows us to tune the time spent on computing the update, and thus better balance against the time spent on computing gradient information. Intuitively, the more expensive is the cost of computing the gradient information, the more time we should spend trying to utilize it better. The time for the former can however be orders of magnitude different in different real world systems [16]. Moreover, as we will argue later in the section, the ‘right’ amount of time to be spent computing the update under identical system conditions also varies during the duration of the algorithm. Due to these reasons we propose to use simple strategies to *automatically* and *adaptively* choose the number of iterations, and hence tune the time spent computing the update.

Let us assume that each iteration of the solver takes one unit of time. From the proofs of Theorems 1,2,3, and 4, we know that the progress we make at each step is

$$F(\mathbf{x}_{t+1}) \leq F(\mathbf{x}_t) + \Theta_t m_t^*.$$

Suppose we spent r iterations i.e. r units of times on this subproblem. Since we want to design strategies which control r , let us parameterize the above equation in terms of r . Since $m_t^* \leq 0$, we denote the progress made as

$$\Theta_t m_t^* \stackrel{\text{def}}{=} -p_t(r).$$

Further let c_t be the ratio between the time spent in computing the gradient information and the time required to perform one iteration of the sub-solver. This includes all fixed costs such as communication time, etc. Our objective should to be to pick an r which makes the maximum progress possible per unit of time spent:

$$r_t^* \stackrel{\text{def}}{=} \arg \max_r \frac{p_t(r)}{r + c_t}. \quad (5)$$

7.1 Fixed strategies for picking r_t

It is reasonable to assume that i) $p_t(r)$ is increasing meaning that running the sub-solver for more iterations leads to improved accuracy, and ii) $p_t(r)$ is *sub-linear* and concave function meaning that doubling the number of iterations results in at most double the accuracy (refer Fig 1). Based on this knowledge, two fixed rules can be formulated to pick a fixed r_t .

One step. The simplest strategy is to just perform one inner step per round. If $c_t = 0$, $\frac{p_t(r)}{r}$ is a decreasing function (Fig 1, left). Thus when the gradient computation cost is low, the one step strategy is optimal.




Figure 1: The progress $p(r)$ (in blue) is bounded, increasing, and concave; attaining its maxima at $r = \infty$. Progress per unit time $\left(\frac{p(r)}{r+c}\right)$ shown in red is i) a decreasing function if $c = 0$ (left) or ii) increases first and then decreases if $c > 0$ (right).




Figure 2: The slope of $p_2(r)$ (in blue, right) is always lesser than that of $p_1(r)$ (in blue, left) and the optimum number or rounds is lesser for p_2 than for p_1 .

However, this is no more true when c_t is significant (Fig 1, right).

Comparable steps. ‘Best practice’ dictates that we should spend roughly equal time performing the update as we take for gradient computation [16]. So this strategy tries to pick $r_t = c_t$. While being a good guiding principle, this strategy fails to take advantage of the trade-offs present in more realistic regimes. For example it ignores the fact that if the local solver is bad or the particular subproblem is especially hard, there might not be much to gain running for c_t steps and it might make sense to terminate early.

7.2 Adaptive strategies for picking r_t

Consider two functions p_1 and p_2 with the corresponding optimum number of iterations being r_1^* and r_2^* as per (5). If the slope of $p_2(r)$ is always smaller than that of $p_1(r)$, then the optimum number of rounds typically (but not always) reduces and $r_2^* < r_1^*$ (refer Fig 2). Since $p(r)$ denotes the progress made on the subproblem, this means that the optimal number of rounds depends on the *ease of minimizing the subproblem*. During the course of our algorithm, the nature of $p_t(r)$ may change substantially and hence the optimum number of rounds also changes. Thus it is imperative to use strategies which are *adaptive to the hardness of $p_t(r)$* .

Optimal. If we could access the entire function $p_t(r)$, and c_t *beforehand*, it is possible to pick the op-

timal r_t^* at every step. However, knowing $p_t(r)$ for all r requires expensive computations. Using an upper bound on $p_t(r)$ obtained from convergence rates is also impractical since it requires knowledge of parameters which are often inaccessible.

Gradient based strategies. At each step, we can assume access to the total time spent ($r_t + c_t$), $p_t(r_t)$, and $p'_t(r_t)$ as feedback. This is because $p_t(r_t)$ and $p'_t(r_t)$ are just the value of the subproblem $-m_t(\Delta \mathbf{x}_t)$ and the rate at which it decreases, measured at the end. Using this feedback, we can compute the gradient g_t of $\frac{p_t(r)}{r+c_t}$ at r_t :

$$g_t = \frac{p'_t(r_t)(r_t + c_t) - p_t(r_t)}{(r_t + c_t)^2}.$$

We can use the gradient g_t as an indicator about the direction we need to change r_t .

1. Additive change. If $g_t > 0$, increase $r_{t+1} \leftarrow r_t + 1$ else $r_{t+1} \leftarrow r_t - 1$.
2. Multiplicative-additive change. If $g_t > 0$, increase $r_{t+1} \leftarrow 2r_t$ else $r_{t+1} \leftarrow r_t - 1$. This is inspired by the TCP protocols used to determine the rate at which to send packets to make maximum use of network capacity [6].
3. Gradient change. Increment with gradient $r_{t+1} \leftarrow r_t + g_t$.

We only focus on very simple strategies to compute r here—many more sophisticated methods based on hyper-parameter optimization or bandit optimization are possible [14].

8 Empirical Evaluation

In this section we present empirical results for our framework, demonstrating the insights gained. Recall from Section 7 that c_t is the ratio between the time taken for computing the gradient and the time taken by the local solver to perform one iteration. In practice, the value of c_t can vary by orders of magnitudes [31]. We artificially set c_t to take values from 0.5 to 1024 in powers of 2 and observe its effect on the different algorithms. This way we can simulate a wide range of real world conditions in our experiments.

8.1 Experimental Setting

We focus on two important tasks in our experiments—Lasso and SVM. We vary the value of c_t and measure the total time it takes to reach a predetermined sub-optimality value. The total time is the sum of the measured update time as well as the simulated gradient computation time. The gradient computation time is calculated as $\frac{t_u c_t}{r_t}$ where t_u is the measured time for update and r_t is the number of iterations performed by the solver.

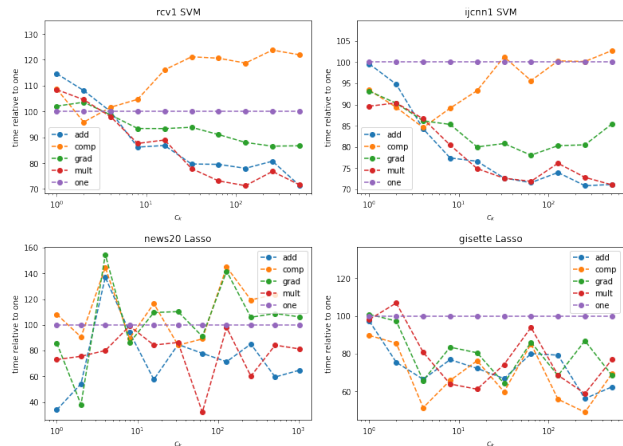


Figure 3: Time taken relative to `one` in percentage, to reach sub optimality (duality gap for SVM) of $1e-4$. Here c_t is the ratio between time for gradient and 1 step of subsolver. Adaptive rules nearly always outperform fixed rules `one` and `comp`.

Lasso. For Lasso, our objective function is $F(\mathbf{x}) = f(\mathbf{x}) + g(\mathbf{x})$ where $g(\mathbf{x}) = \lambda \|\mathbf{x}\|_1$ and $f(\mathbf{x}) = \frac{1}{2n} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2$. Here λ is a regularization parameter chosen to be $\frac{1}{n}$ as is standard [25]. We run this on the `gisette` and `news20` datasets.¹ For each value of c_t , we measure the total time it takes for each algorithm to reach an sub-optimality of $1e-4$. The minimum value is calculated by letting the algorithm run for 1k effective passes over the data.

SVM. Here we require primal-dual convergence for the hinge loss $\mathcal{O}_B(\mathbf{w}) = \frac{1}{n} \sum_{i=1}^n l(\mathbf{w}^\top A_i, y_i) + \frac{\lambda}{2} \|\mathbf{w}\|^2$, c.f. [9]. The regularization parameter λ is again chosen to be $\frac{1}{n}$. We run this on `rcv1` and `ijcnn1` datasets till we reach a duality gap of $1e-4$.

Subproblem solver. To simplify the comparison, we use random coordinate descent as a solver for all the subproblems. Further we create smaller subproblems of dimension 100 using the algorithm from Section 6 (see also Sec. A.1). Hence, one iteration of the subproblem solver consists of 100 random coordinate updates. To estimate the gradient $p'_t(r_t)$, we use the average progress made in the last 10 steps of the solver.

Strategies for r_t . To demonstrate our discussion from the previous section, we compare the different strategies which perform $100r_t$ CD iterations at each step: i) `one` where $r_t = 1$, ii) `comp` where $r_t = c_t$, iii) `add` where r_t is computed using the additive change rule, iv) `mult` where r_t is computed using the multiplicative-additive rule, and v) `grad` where r_t is computed using the gradient rule.

¹All datasets are available from <https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/>




Figure 4: Average value of r chosen by different strategies. It remains relatively constant as c_t increases across problems and strategies.

8.2 Discussion

As seen in Fig 3, the adaptive rules based on the gradient consistently scale better for increasing values for the ratio c_t . In almost all the cases the simple **mult** rule performs better than both the standard **one** rule as well the **comp** rule.

For small values of c_t (≤ 2) in the **rcv1** SVM experiment as well as in the **gisette** Lasso example using the **one** rule and setting $r_t = 1$ is effective. However, as c_t increases, the **one** rule becomes much less competitive. This is because for small values of c_t , the function $\frac{p_t(r)}{r+c_t}$ attains its maxima close to 1.

We also plot the average values of r_t over the period of optimization for the **ijcnn1** and the **news20** cases (cf. Fig 4). The plots from the other two datasets were similar. The average values of r_t are pretty identical in both the cases. It also remains close to 1–2, and is mostly independent of c_t for the **mult** and the **add** rule. Thus the improvements we see in Fig 3 for these rules must have been because of the use of *adaptivity*.

To understand how adaptivity affects our algorithms, we look at how r_t varies over time for **ijcnn1** with the **mult** rule with $c_t = 64$ in (Fig 5, left). The strikingly clear downward trend in r_t can be explained via our discussion Section 7.2 about changing hardness of sub-problems. The sub-problems in **ijcnn1** are increasing in their difficulty (Fig 5, right), pushing for lower values of r_t .

For **news20** dataset, the value of r_t remain relatively constant (Appendix Section C). This explains why fixed rules such as **comp** perform comparably to the adaptive ones in **news20**. The variation of r_t in **rcv1** is similar to that of **ijcnn1** and so they have similar results. The growing gap between adaptive rules and the fixed rules in **rcv1** and **ijcnn1** clearly demonstrates the effectiveness of adaptive rules in taking advantage of changing conditions. Refer to the Appendix Section C for additional figures and discussion.

Another aspect to notice in Fig 3 is that the **add** and **mult** rules perform quite similarly. This is because




Figure 5: (Left) Time varying r_t for **mult** with $c = 64$ for **ijcnn1**. ((Right) The number of rounds required to solve the subproblem to an accuracy of 0.1. The hardness of the subproblems increase with t and in response, r_t chosen by **mult** decreases. A moving average with window of size 100 is used to aid visualization.

during the majority of the runtime, r_t is close to 1. When $r_t = 1$, both the rules are identical. Changing the multiplicative constant could alleviate this issue. Similarly in the **grad** algorithm, due to the highly noisy nature of the gradient signal, we noticed large oscillations. Using step-sizes or a running average instead could stabilize the algorithm. We believe that such fine tuning of parameters would lead to better algorithms and even more gains.

9 Concluding Remarks

We present ACM, a single framework that provides a unified analysis for first-order optimization algorithms for composite problems. The framework allows to incorporate curvature information and only requires the computation of weak approximate solutions to the subproblems in a stochastic sense—hence, it specifically also covers randomized methods such as block-coordinate descent and also parallel algorithms. Moreover, the accuracy parameters are allowed to change over time. This combines and improves upon the results of [16, 17, 21, 33].

We leverage our framework to provide speedups when the gradient computation and the update computation times are unbalanced. In particular, we give a simple adaptive procedure to adaptively tune the accuracy that is required in for the optimization of the model in each iteration. This procedure ensures optimal balance between of the two computations.

The effectiveness of the adaptive scheme is exemplary demonstrated on a set of numerical experiments for Lasso (**gisette** and **news20** dataset) and SVM (**rcv1**, **ijcnn1**) with randomized coordinate descent as sub-problem solver. The experiments shows that simply tuning a static accuracy parameter will in general not obtain optimal rate. In contrast, the parameters of our adaptive scheme vary as the optimization progresses, and achieve significant speedups.

Acknowledgements SPK thanks Mojmir Mutny for discussions relating to parallelizing preconditioned block-coordinate methods, as well as Anastasia Koloskova for help with the experiments.

References

- [1] N. Agarwal, B. Bullins, and E. Hazan. Second Order Stochastic Optimization in Linear Time. *arXiv:1602.03943 [cs, stat]*, Feb. 2016.
- [2] Z. Allen-Zhu. Katyusha: The First Direct Acceleration of Stochastic Gradient Methods. *arXiv:1603.05953 [cs, math, stat]*, Mar. 2016.
- [3] Y. Bian, X. Li, Y. Liu, and M.-H. Yang. Parallel Coordinate Descent Newton Method for Efficient ℓ_1 -Regularized Minimization. *arXiv:1306.4080 [cs]*, June 2013.
- [4] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, New York, NY, USA, 2004.
- [5] A. Chambolle, M. Novaga, D. Cremers, and T. Pock. An introduction to total variation for image analysis. In *In Theoretical Foundations and Numerical Methods for Sparse Recovery*, De Gruyter, 2010.
- [6] D.-M. Chiu and R. Jain. Analysis of the Increase and Decrease Algorithms for Congestion Avoidance in Computer Networks. *Comput. Netw. ISDN Syst.*, 17(1):1–14, June 1989.
- [7] A. d’Aspremont. Smooth Optimization with Approximate Gradient. *SIAM Journal on Optimization*, 19(3):1171–1183, Jan. 2008.
- [8] A. Defazio, F. Bach, and S. Lacoste-Julien. SAGA: A Fast Incremental Gradient Method With Support for Non-Strongly Convex Composite Objectives. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 1646–1654. Curran Associates, Inc., 2014.
- [9] C. Dünnner, S. Forte, M. Takac, and M. Jaggi. Primal-Dual Rates and Certificates. In *ICML’16 - Proceedings of the 33rd International Conference on International Conference on Machine Learning*, pages 783–792, June 2016.
- [10] O. Fercoq, Z. Qu, P. Richtarik, and M. Takac. Fast distributed coordinate descent for non-strongly convex losses. In *2014 IEEE International Workshop on Machine Learning for Signal Processing (MLSP)*, pages 1–6, Sept. 2014.
- [11] O. Fercoq and P. Richtarik. Accelerated, Parallel, and Proximal Coordinate Descent. *SIAM Journal on Optimization*, 25(4):1997–2023, Jan. 2015.
- [12] R. Johnson and T. Zhang. Accelerating Stochastic Gradient Descent using Predictive Variance Reduction. In *NIPS - Advances in Neural Information Processing Systems 26*, pages 315–323. Curran Associates, Inc., 2013.
- [13] J. D. Lee, Q. Lin, T. Ma, and T. Yang. Distributed Stochastic Variance Reduced Gradient Methods and A Lower Bound for Communication Complexity. *arXiv:1507.07595 [cs, math, stat]*, July 2015.
- [14] L. Li, K. Jamieson, G. DeSalvo, A. Rostamizadeh, and A. Talwalkar. Hyperband: A Novel Bandit-Based Approach to Hyperparameter Optimization. *arXiv:1603.06560 [cs, stat]*, Mar. 2016.
- [15] Q. Lin, Z. Lu, and L. Xiao. An Accelerated Proximal Coordinate Gradient Method. In *NIPS - Advances in Neural Information Processing Systems 27*, pages 3059–3067. Curran Associates, Inc., 2014.
- [16] C. Ma, J. Konecny, M. Jaggi, V. Smith, M. I. Jordan, P. Richtarik, and M. Takac. Distributed optimization with arbitrary local solvers. *Optimization Methods and Software*, 32(4):813–848, July 2017.
- [17] M. Mutny and P. Richtarik. Parallel Stochastic Newton Method. *arXiv:1705.02005 [math]*, May 2017.
- [18] Y. Nesterov. Efficiency of Coordinate Descent Methods on Huge-Scale Optimization Problems. *SIAM Journal on Optimization*, 22(2):341–362, Jan. 2012.
- [19] Y. Nesterov and S. Stich. Efficiency of the Accelerated Coordinate Descent Method on Structured Optimization Problems. *SIAM Journal on Optimization*, 27(1):110–123, Jan. 2017.
- [20] L. M. Nguyen, J. Liu, K. Scheinberg, and M. Takac. SARAH: A Novel Method for Machine Learning Problems Using Stochastic Recursive Gradient. *arXiv:1703.00102 [cs, math, stat]*, Feb. 2017.
- [21] Z. Qu, P. Richtarik, M. Takac, and O. Fercoq. SDNA: Stochastic Dual Newton Ascent for Empirical Risk Minimization. *arXiv:1502.02268 [cs]*, Feb. 2015.
- [22] S. J. Reddi, J. Konecny, P. Richtarik, B. Póczos, and A. Smola. AIDE: Fast and Communication Efficient Distributed Optimization. *arXiv:1608.06879 [cs, math, stat]*, Aug. 2016.
- [23] P. Richtarik and M. Takac. Distributed Coordinate Descent Method for Learning with Big Data. *J. Mach. Learn. Res.*, 17(1):2657–2681, Jan. 2016.
- [24] P. Richtarik and M. Takac. Parallel coordinate descent methods for big data optimization. *Mathematical Programming*, 156(1-2):433–484, Mar. 2016.
- [25] N. L. Roux, M. Schmidt, and F. Bach. A Stochastic Gradient Method with an Exponential Convergence Rate for Finite Training Sets. In *NIPS - Neural Information Processing Systems*, NIPS’12, pages 2663–2671, USA, 2012. Curran Associates Inc.
- [26] M. Schmidt, N. L. Roux, and F. R. Bach. Convergence Rates of Inexact Proximal-Gradient Methods for Convex Optimization. In J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24*, pages 1458–1466. Curran Associates, Inc., 2011.

- [27] S. Shalev-Shwartz and S. Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, New York, NY, USA, 2014.
- [28] S. Shalev-Shwartz, Y. Singer, N. Srebro, and A. Cotter. Pegasos: Primal estimated sub-gradient solver for SVM. *Mathematical Programming*, 127(1):3–30, Mar. 2011.
- [29] S. Shalev-Shwartz and T. Zhang. Stochastic Dual Coordinate Ascent Methods for Regularized Loss. *J. Mach. Learn. Res.*, 14(1):567–599, Feb. 2013.
- [30] S. Shalev-Shwartz and T. Zhang. Accelerated proximal stochastic dual coordinate ascent for regularized loss minimization. *Mathematical Programming*, 155(1-2):105–145, Jan. 2016.
- [31] V. Smith, S. Forte, C. Ma, M. Takac, M. I. Jordan, and M. Jaggi. CoCoA: A General Framework for Communication-Efficient Distributed Optimization. *arXiv:1611.02189 [cs]*, Nov. 2016.
- [32] S. U. Stich, C. L. Müller, and B. Gärtner. Variable metric random pursuit. *Mathematical Programming*, 156(1-2):549–579, Mar. 2016.
- [33] R. Tappenden, P. Richtarik, and J. Gondzio. Inexact Coordinate Descent: Complexity and Preconditioning. *Journal of Optimization Theory and Applications*, 170(1):144–176, July 2016.
- [34] S. Tu, S. Venkataraman, A. C. Wilson, A. Gittens, M. I. Jordan, and B. Recht. Breaking Locality Accelerates Block Gauss-Seidel. *arXiv:1701.03863 [math]*, Jan. 2017.
- [35] S. J. Wright. Coordinate descent algorithms. *Mathematical Programming*, 151(1):3–34, June 2015.