



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Inexact Coordinate Descent: Complexity and Preconditioning

Citation for published version:

Tappenden, R, Richtárik, P & Gondzio, J 2013 'Inexact Coordinate Descent: Complexity and Preconditioning' ArXiv.

Link:

[Link to publication record in Edinburgh Research Explorer](#)

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Inexact Coordinate Descent: Complexity and Preconditioning*

Rachael Tappenden Peter Richtárik Jacek Gondzio

*School of Mathematics
University of Edinburgh
United Kingdom*

December 11, 2014

Abstract

In this paper we consider the problem of minimizing a convex function using a randomized block coordinate descent method. One of the key steps at each iteration of the algorithm is determining the update to a block of variables. Existing algorithms assume that in order to compute the update, a particular subproblem is solved *exactly*. In this work we relax this requirement, and allow for the subproblem to be solved *inexactly*, leading to an inexact block coordinate descent method. Our approach incorporates the best known results for exact updates as a special case. Moreover, these theoretical guarantees are complemented by practical considerations: the use of iterative techniques to determine the update as well as the use of preconditioning for further acceleration.

Keywords: inexact methods, block coordinate descent, convex optimization, iteration complexity, preconditioning, conjugate gradients.

AMS: 65F08; 65F10; 65F15; 65Y20; 68Q25; 90C25

1 Introduction

Due to a dramatic increase in the size of optimization problems being encountered, first order methods are becoming increasingly popular. These large-scale problems are often highly structured and it is important for any optimization method to take advantage of the underlying structure. Applications where such problems arise and where first order methods have proved successful include machine learning [17, 33], compressive sensing [8, 43], group lasso [26, 36], matrix completion [5, 27], and truss topology design [28].

Block coordinate descent methods seem a natural choice for these very large-scale problems due to their low memory requirements and low per-iteration computational cost. Furthermore, they are often designed to take advantage of the underlying structure of the optimization problem [41, 42] and many of these algorithms are supported by high probability iteration complexity results [24, 25, 28, 29, 38].

*This work was supported by the EPSRC grant EP/I017127/1 “Mathematics for vast digital resources”. Peter Richtárik was also supported by the Centre for Numerical Algorithms and Intelligent Software (funded by EPSRC grant EP/G036136/1 and the Scottish Funding Council).

1.1 The Problem

If the block size is larger than one, determining the update to use at a particular iteration in a block coordinate descent method can be computationally expensive. The purpose of this work is to reduce the cost of this step. To achieve this, we extend the work in [29] to include the case of an *inexact* update.

In this work we study randomized block coordinate descent methods applied to the problem of minimizing a composite objective function. That is, a function formed as the sum of a smooth convex and a simple nonsmooth convex term:

$$\min_{x \in \mathbb{R}^N} \{F(x) := f(x) + \Psi(x)\}. \quad (1)$$

We assume that the problem has a minimum ($F^* > -\infty$), f has (block) coordinate Lipschitz gradient, and Ψ is a (block) separable proper closed convex extended real valued function (all these concepts will be defined precisely in Section 2).

Our algorithm (namely, the Inexact Coordinate Descent (ICD) method) is supported by high probability iteration complexity results. That is, for confidence level $\rho \in (0, 1)$ and error tolerance $\epsilon > 0$, we give an explicit expression for the number of iterations k that guarantee that the ICD method produces a random iterate x_k for which

$$\mathbb{P}(F(x_k) - F^* \leq \epsilon) \geq 1 - \rho.$$

We will show that in the inexact case it is not always possible to achieve a solution with small error and/or high confidence.

Our theoretical guarantees are complemented by practical considerations. In Section 3.3 we explain our inexactness condition in detail and in Section 3.4 we give examples to show when the inexactness condition is implementable. Further, in Section 6 we give several examples, derive the update subproblems, and suggest algorithms that could be used to solve the subproblems inexactly. Finally, we present some encouraging computational results.

1.2 Literature Review

As problem sizes increase, first order methods are benefiting from revived interest. On very large problems however, the computation of a single gradient step is expensive, and methods are needed that are able to make progress before a standard gradient algorithm takes a single step. For instance, a randomized variant of the Kaczmarz method for solving linear systems has recently been studied, equipped with iteration complexity bounds [21, 22, 16, 37], and found to be surprisingly efficient. This method can be seen as a special case of a more general class of decomposition algorithms, block coordinate descent methods, which have recently gained much popularity [20, 24, 25, 29, 30, 31, 40]. One of the main differences between various (serial) coordinate descent schemes is the way in which the coordinate is chosen at each iteration. Traditionally cyclic schemes [32] and greedy schemes [28] were studied. More recently, a popular alternative is to select coordinates randomly, because the coordinate can be selected cheaply, and useful iteration complexity results can be obtained [19, 29, 30, 31, 39, 35].

Another current trend in this area is to consider methods that incorporate some kind of ‘inexactness’, perhaps using approximate gradients, or using inexact updates. For example, [18] considers methods based on inexact dual gradient information, while [33] considers the minimization of an

unconstrained convex composite function where error is present in the gradient of the smooth term, or in the proximity operator for the non-smooth term. Other works study methods that use inexact updates when the objective function is convex, smooth and unconstrained [1], smooth and constrained [3] or for ℓ_1 -regularized quadratic least squares problem [15].

1.3 Contribution

In this paper we extend the work of Richtárik and Takáč [29] and present a block coordinate descent method that employs inexact updates having the potential to reduce the overall algorithm running time. Furthermore, we focus in detail on the quadratic case, which benefits greatly from inexact updates, and show how preconditioning can be used to complement the inexact update strategy.

F	Exact Method [29]	Inexact Method [this paper]	Theorem
C-N	$\frac{c_1}{\epsilon} \left(1 + \log \frac{1}{\rho}\right) + 2$	$\frac{c_1}{\epsilon - u} + \frac{c_1}{\epsilon - \alpha c_1} \log \left(\frac{\epsilon - \frac{\beta c_1}{\epsilon - \alpha c_1}}{\epsilon \rho - \frac{\beta c_1}{\epsilon - \alpha c_1}} \right) + 2$	9(i)
C-N	$c_2 \log \left(\frac{F(x_0) - F^*}{\epsilon \rho} \right)$	$\frac{c_2}{1 - \alpha c_2} \log \left(\frac{F(x_0) - F^* - \frac{\beta c_2}{1 - \alpha c_2}}{\epsilon \rho - \frac{\beta c_2}{1 - \alpha c_2}} \right)$	9(ii)
SC-N	$\frac{n}{\mu} \log \left(\frac{F(x_0) - F^*}{\epsilon \rho} \right)$	$\frac{n}{\mu - \alpha n} \log \left(\frac{F(x_0) - F^* - \frac{\beta n}{\mu - \alpha n}}{\epsilon \rho - \frac{\beta n}{\mu - \alpha n}} \right)$	11
C-S	$\frac{\hat{c}_1}{\epsilon} \left(1 + \log \frac{1}{\rho}\right) + 2$	$\frac{\hat{c}_1}{\epsilon - \hat{u}} + \frac{\hat{c}_1}{\epsilon - \alpha \hat{c}_1} \log \left(\frac{\epsilon - \frac{\beta \hat{c}_1}{\epsilon - \alpha \hat{c}_1}}{\epsilon \rho - \frac{\beta \hat{c}_1}{\epsilon - \alpha \hat{c}_1}} \right) + 2$	12
SC-S	$\frac{1}{\mu_f} \log \left(\frac{f(x_0) - f^*}{\epsilon \rho} \right)$	$\frac{1}{\mu_f - \alpha} \log \left(\frac{f(x_0) - f^* - \frac{\beta}{\mu_f - \alpha}}{\epsilon \rho - \frac{\beta}{\mu_f - \alpha}} \right)$	13

Table 1: Comparison of the iteration complexity results for coordinate descent methods using an inexact update and using an exact update (C=Convex, SC=Strongly Convex, N=Nonsmooth, S = Smooth).

Table 1 compares *some* of the new complexity results obtained in this paper for an inexact update with the complexity results for an exact update presented in [29]. The following notation is used in the table: by μ_ϕ we denote the strong convexity parameter of function ϕ (with respect to a certain norm specified later), $\mu = (\mu_f + \mu_\Psi)/(1 + \mu_\Psi)$ and $\mathcal{R}_w(x_0)$ can be roughly considered to be distance from x_0 to a solution of (1) measured in a specific weighted norm parameterized by the vector w (to be defined precisely in (14)). The constants are $c_1 = 2n \max\{\mathcal{R}_w^2(x_0), F(x_0) - F^*\}$, $\hat{c}_1 = 2\mathcal{R}_w^2(x_0)$ and $c_2 = 2n\mathcal{R}_w^2(x_0)/\epsilon$, and n is the number of blocks. Parameters $\alpha, \beta \geq 0$ control the level of inexactness (to be defined precisely in Section 3.2) and u and \hat{u} are constants depending on α, β and c_1 , and α, β and \hat{c}_1 , respectively.

Table 1 shows that for fixed ϵ and ρ , an inexact method will require more iterations than an exact one. However, it is expected that in certain situations an inexact update will be significantly cheaper to compute than an exact update, leading to better overall running time. Moreover, the new complexity results for the inexact method generalize those for the exact method. Specifically,

for inexactness parameters $\alpha = \beta = 0$ we recover the complexity results in [29].

1.4 Outline

The first part of this paper focuses on the theoretical aspects of a block coordinate descent method when an inexact update is employed. In Section 2 the assumptions and notation are laid out and in Section 3 the ICD method is presented. In Section 4 iteration complexity results for ICD applied to (1) are presented in both the convex and strongly convex cases. Iteration complexity results for ICD applied to a convex smooth minimization problem ($\Psi = 0$ in (1)) are presented in Section 5, in both the convex and strongly convex cases.

The second part of the paper considers the practicality of an inexact update. Section 6 provides several examples of how to derive the formulation for the update step subproblem, as well as giving suggestions for algorithms that can be used to solve the subproblem inexactly. Numerical experiments are presented in Section 7 and Appendix A provides a detailed analysis of the spectrum of the preconditioned matrix used in the numerical experiments.

2 Assumptions and Notation

In this section we introduce the notation and definitions that are used throughout the paper.

2.1 Block structure of \mathbb{R}^N

The problem under consideration is assumed to have block structure and this is modelled by decomposing the space \mathbb{R}^N into n subspaces as follows. Let $U \in \mathbb{R}^{N \times N}$ be a column permutation of the $N \times N$ identity matrix and further let $U = [U_1, U_2, \dots, U_n]$ be a decomposition of U into n submatrices, where U_i is $N \times N_i$ and $\sum_{i=1}^n N_i = N$. It is clear (e.g., see [30] for a brief proof) that any vector $x \in \mathbb{R}^N$ can be written uniquely as

$$x = \sum_{i=1}^n U_i x^{(i)}, \quad (2)$$

where $x^{(i)} \in \mathbb{R}^{N_i}$. Moreover, these vectors are given by

$$x^{(i)} := U_i^T x. \quad (3)$$

For simplicity we will sometimes write $x = (x^{(1)}, x^{(2)}, \dots, x^{(n)})$ instead of (2). We equip \mathbb{R}^{N_i} with a pair of conjugate norms, induced by a quadratic form involving a symmetric positive definite matrix B_i :

$$\|t\|_{(i)} := \langle B_i t, t \rangle^{\frac{1}{2}}, \quad \|t\|_{(i)}^* = \langle B_i^{-1} t, t \rangle^{\frac{1}{2}}, \quad t \in \mathbb{R}^{N_i}, \quad (4)$$

where $\langle \cdot, \cdot \rangle$ is the standard Euclidean dot product.

2.2 Smoothness of f

Throughout this paper we assume that the gradient of f is block Lipschitz, uniformly in x , with positive constants l_1, \dots, l_n . This means that, for all $x \in \mathbb{R}^N$, $i \in \{1, 2, \dots, n\}$ and $t \in \mathbb{R}^{N_i}$ we have

$$\|\nabla_i f(x + U_i t) - \nabla_i f(x)\|_{(i)}^* \leq l_i \|t\|_{(i)}, \quad (5)$$

where

$$\nabla_i f(x) := (\nabla f(x))^{(i)} \stackrel{(3)}{=} U_i^T \nabla f(x) \in \mathbb{R}^{N_i}. \quad (6)$$

An important consequence of (5) is the following standard inequality [23, p.57]:

$$f(x + U_i t) \leq f(x) + \langle \nabla_i f(x), t \rangle + \frac{l_i}{2} \|t\|_{(i)}^2. \quad (7)$$

2.3 Block separability of Ψ

The function $\Psi : \mathbb{R}^N \rightarrow \mathbb{R} \cup \{+\infty\}$ is assumed to be block separable. That is, we assume that it can be decomposed as:

$$\Psi(x) = \sum_{i=1}^n \Psi_i(x^{(i)}), \quad (8)$$

where the functions $\Psi_i : \mathbb{R}^{N_i} \rightarrow \mathbb{R} \cup \{+\infty\}$ are convex and closed.

2.4 Norms on \mathbb{R}^N

For fixed positive scalars w_1, w_2, \dots, w_n , let $w = (w_1, \dots, w_n)$ and define a pair of conjugate norms in \mathbb{R}^N by

$$\|x\|_w^2 := \sum_{i=1}^n w_i \|x^{(i)}\|_{(i)}^2, \quad (\|y\|_w^*)^2 := \max_{\|x\|_w \leq 1} \langle y, x \rangle^2 = \sum_{i=1}^n w_i^{-1} (\|y^{(i)}\|_{(i)}^*)^2. \quad (9)$$

In the subsequent analysis we will often use $w = l$ (for $\Psi \neq 0$) and/or $w = lp^{-1}$ (for $\Psi = 0$), where $l = (l_1, \dots, l_n)$ is a vector of Lipschitz constants, $p = (p_1, \dots, p_n)$ is a vector of positive probabilities and lp^{-1} denotes the vector $(l_1/p_1, \dots, l_n/p_n)$.

2.5 Strong convexity of F

A function $\phi : \mathbb{R}^N \rightarrow \mathbb{R} \cup \{+\infty\}$ is strongly convex w.r.t. $\|\cdot\|_w$ with convexity parameter $\mu_\phi(w) > 0$ if for all $x, y \in \text{dom } \phi$,

$$\phi(y) \geq \phi(x) + \langle \phi'(x), y - x \rangle + \frac{\mu_\phi(w)}{2} \|y - x\|_w^2, \quad (10)$$

where ϕ' is any subgradient of ϕ at x . The case with $\mu_\phi(w) = 0$ reduces to convexity.

In some of the results presented in this work we assume that F is strongly convex. Strong convexity of F may come from f or Ψ or both and we will write $\mu_f(w)$ (resp. $\mu_\Psi(w)$) for the strong convexity parameter of f (resp. Ψ), with respect to $\|\cdot\|_w$. Following from (10)

$$\mu_F(w) \geq \mu_f(w) + \mu_\Psi(w). \quad (11)$$

Using (7) and (10) it can be shown that

$$\mu_f(l) \leq 1, \quad \text{and} \quad \mu_f(lp^{-1}) < 1. \quad (12)$$

We will also make use of the following characterisation of strong convexity. For all $x, y \in \text{dom } \phi$ and $\lambda \in [0, 1]$,

$$\phi(\lambda x + (1 - \lambda)y) \leq \lambda \phi(x) + (1 - \lambda)\phi(y) - \frac{\mu_\phi(w)\lambda(1-\lambda)}{2} \|x - y\|_w^2. \quad (13)$$

2.6 Level set radius

The set of optimal solutions of (1) is denoted by X^* and x^* is any element of that set. We define

$$\mathcal{R}_w(x) := \max_y \max_{x^* \in X^*} \{\|y - x^*\|_w : F(y) \leq F(x)\}, \quad (14)$$

which is a measure of the size of the level set of F given by x . We assume that $\mathcal{R}_w(x_0)$ is finite for the initial iterate x_0 .

3 The Algorithm

Let us start by presenting the algorithm; a more detailed description will follow.

Algorithm 1 ICD: Inexact Coordinate Descent

- 1: Input: Inexactness parameters $\alpha, \beta \geq 0$, and probabilities $p_1, \dots, p_n > 0$.
 - 2: **for** $k = 0, 1, 2, \dots$ **do**
 - 3: Choose $\delta_k = (\delta_k^{(1)}, \dots, \delta_k^{(n)}) \in \mathbb{R}^n$ according to (20)
 - 4: Choose block $i \in \{1, 2, \dots, n\}$ with probability p_i
 - 5: Compute the inexact update $T_{\delta_k}^{(i)}(x_k)$ to block i of x_k
 - 6: Update block i of x_k : $x_{k+1} = x_k + U_i T_{\delta_k}^{(i)}(x_k)$
 - 7: **end for**
-

3.1 Generic description

Given iterate $x_k \in \mathbb{R}^N$, Algorithm 1 picks block $i \in \{1, 2, \dots, n\}$ with probability p_i , computes the update vector $T_{\delta_k}^{(i)}(x_k) \in \mathbb{R}^{N_i}$ (we comment on how this is computed later in this section) and then adds it to the i th block of x_k , producing the new iterate x_{k+1} . The iterates $\{x_k\}$ are random vectors and the values $\{F(x_k)\}$ are random variables. The update vector depends on x_k , the current iterate, and on δ_k , a vector of parameters controlling the “level of inexactness” with which the update is computed. The rest of this section is devoted to giving a precise definition of $T_{\delta_k}^{(i)}(x_k)$. Note that from (1) and (7) we have, for all $x \in \mathbb{R}^N$, $i \in \{1, 2, \dots, n\}$ and $t \in \mathbb{R}^{N_i}$:

$$F(x + U_i t) = f(x + U_i t) + \Psi(x + U_i t) \leq f(x) + V_i(x, t) + \Psi_{-i}(x), \quad (15)$$

where

$$V_i(x, t) := \langle \nabla_i f(x), t \rangle + \frac{l_i}{2} \|t\|_{(i)}^2 + \Psi_i(x^{(i)} + t), \quad (16)$$

$$\Psi_{-i}(x) := \sum_{j \neq i} \Psi_j(x^{(j)}). \quad (17)$$

That is, (15) gives an upper bound on $F(x + U_i t)$, viewed as a function of $t \in \mathbb{R}^{N_i}$.

The inexact update computed in Step 5 of Algorithm 1 is the *inexact* minimizer of the upper bound (15) on $F(x_k + U_i t)$ (to be defined precisely below). However, since only the second term of this bound depends on t , the update is computed by minimizing, *inexactly*, $V_i(x, t)$ in t .

3.2 Inexact update

The approach of this paper best applies to situations in which it is much easier to approximately minimize $t \mapsto V_i(x, t)$ than to either (i) approximately minimize $t \mapsto F(x + U_i t)$ and/or (ii) exactly minimize $t \mapsto V_i(x, t)$. For $x \in \mathbb{R}^N$ and $\delta = (\delta^{(1)}, \dots, \delta^{(n)}) \geq 0$ we define $T_\delta(x) := (T_\delta^{(1)}(x), \dots, T_\delta^{(n)}(x)) \in \mathbb{R}^N$ to be any vector satisfying

$$V_i(x, T_\delta^{(i)}(x)) \leq \min \left\{ V_i(x, 0), \delta^{(i)} + \min_{t \in \mathbb{R}^{N_i}} V_i(x, t) \right\}, \quad i = 1, \dots, n. \quad (18)$$

(We allow here for an abuse of notation — $\delta^{(i)}$ is a scalar, rather than a vector in \mathbb{R}^{N_i} as $x^{(i)}$ for $x \in \mathbb{R}^N$ — because we wish to emphasize that the scalar $\delta^{(i)}$ is associated with the i th block.) That is, we require that the inexact update $T_\delta^{(i)}(x)$ of the i th block of x is (i) no worse than a vacuous update, and that it is (ii) close to the optimal update $T_0^{(i)}(x) = \arg \min_t V_i(x, t)$, where the degree of suboptimality/inexactness is bounded by $\delta^{(i)}$.

As the following lemma shows, the update (18) leads to a monotonic algorithm.

Lemma 1. For all $x \in \mathbb{R}^N$, $\delta \in \mathbb{R}_+^n$ and $i \in \{1, 2, \dots, n\}$,

$$F(x + U_i T_\delta^{(i)}(x)) \leq F(x). \quad (19)$$

Proof:

$$\begin{aligned} F(x + U_i T_\delta^{(i)}(x)) &\stackrel{(15)}{\leq} f(x) + V_i(x, T_\delta^{(i)}(x)) + \Psi_{-i}(x) \\ &\stackrel{(18)}{\leq} f(x) + V_i(x, 0) + \Psi_{-i}(x) \stackrel{(16)+(17)}{=} F(x). \quad \square \end{aligned}$$

Furthermore, in this work we provide iteration complexity results for ICD, where $\delta_k = (\delta_k^{(1)}, \dots, \delta_k^{(n)})$ is chosen in such a way that the expected suboptimality is bounded above by a linear function of the residual $F(x_k) - F^*$. That is, we have the following assumption.

Assumption 2. For constants $\alpha, \beta \geq 0$, the vector $\delta_k = (\delta_k^{(1)}, \dots, \delta_k^{(n)})$ is chosen to satisfy

$$\bar{\delta}_k := \sum_{i=1}^n p_i \delta_k^{(i)} \leq \alpha(F(x_k) - F^*) + \beta, \quad (20)$$

Notice that, for instance, Assumption 2 holds if we require $\delta_k^{(i)} \leq \alpha(F(x_k) - F^*) + \beta$ for all blocks i and iterations k .

The motivation for allowing inexact updates of the form (18) is that calculating exact updates is impossible in some cases (for example, not all problems have a closed form solution), and computationally intractable in others. The purpose of allowing inexactness in the update step of ICD is that an *iterative method* can be used to solve for the update $T_\delta^{(i)}(x)$, thus significantly expanding the range of problems that can be successfully tackled by coordinate descent. In this case, there is an outer coordinate descent loop, and an inner iterative loop to determine the update. Assumption 2 shows that the stopping tolerance on the inner loop, must be bounded above via (20).

CD methods provide a mechanism to break up very large/huge scale problem, into smaller pieces that are a fraction of the total dimension. Moreover, often the subproblems that arise to

solve for the update have a similar/the same form as the original huge scale problem. (For example, see the numerical experiments in Section 7, and the examples given in Section 6.) There are many iterative methods that cannot scale up to the original huge dimensional problem, but are excellent at solving the medium scale update subproblems. ICD allows these algorithms to solve for the update at each iteration, and if the updates are solved efficiently, then the overall ICD algorithm running time is kept low.

3.3 The role of α and β in ICD

The condition (18) shows that the updates in ICD are *inexact*, while Assumption 2 gives the *level of inexactness* that is allowed in the computed update. Moreover, Assumption 2 allows us to provide a unified analysis; formulating the error/inexactness expression in this general way (20) gives insight into the role of both *multiplicative and additive error*, and *how this error propagates* through the algorithm as iterates progress.

Formulation (20) is interesting from a theoretical perspective because it allows us to present a *sensitivity analysis* for ICD, which is interesting in its own right. However, we stress that (18), coupled with Assumption 2, is much more than just a technical tool; α and β are actually parameters of the ICD algorithm (Algorithm 1) that can be assigned *explicit numerical values* in many cases.

We now explain (18), Assumption 2 and the role of parameters α and β in slightly more detail. (Note that α and β must be chosen sufficiently small to guarantee converge of the ICD algorithm. However, we postpone discussion of the magnitude of α and β until Section 3.5.) There are four cases.

1. **Case I:** $\alpha = \beta = 0$. This corresponds to the *exact case* where *no error* is allowed in the computed update.
2. **Case II:** $\alpha = 0, \beta > 0$. This case corresponds to additive error only, where the error level $\beta > 0$ is fixed at the start of Algorithm 1. In this case, (18) and (20) show that the error allowed in the inexact update $T_\delta^{(i)}(x_k)$ is *on average* β . For example, one can set $\delta_k^{(i)} = \beta$, for all blocks i and all iterations k so that (20) becomes $\bar{\delta}_k = \sum_i p_i \beta = \beta$. Notice that the tolerance allowable on each block need not be the same; if one sets $\delta_k^{(i)} \leq \beta$, for all blocks i and iterates k then $\bar{\delta}_k \leq \beta$, so (20) holds true. Moreover, one need not set $\delta_k^{(i)} > 0$ for all i , so that the update vector $T_\delta^{(i)}(x_k)$ could be exact for some blocks ($\delta_k^{(i)} = 0$), and inexact for others ($\delta_k^{(j)} > 0$). (This may be sensible, for example, when $\Psi_i(x^{(i)}) \neq \Psi_j(x^{(j)})$ for some $i \neq j$ and that (16) has a closed form solution for $T^{(i)}(x_k)$ but not for $T^{(j)}(x_k)$.) Furthermore, consider the extreme case where only one block update is inexact $T_\delta^{(i)}(x_k)$, ($T_0^{(j)}(x_k)$ for all $j \neq i$). If the coordinates are selected with uniform probability, then the inexactness level on block i can be as large as $\delta_k^{(i)} = n\beta$ and Assumption 2 holds.
3. **Case III:** $\alpha > 0, \beta = 0$. In this case only multiplicative error is allowed in the computed update $T_\delta^{(i)}(x_k)$, where the error allowed in the update at iteration k is related to the error in the function value $(F(x_k) - F^*)$. The multiplicative error level α is fixed at the start of Algorithm 1, and $\alpha(F(x_k) - F^*)$ is an upper bound on the average error in the update $T_\delta^{(i)}(x_k)$ over all blocks i at iteration k . In particular, notice that setting $\delta_k^{(i)} \leq \alpha(F(x_k) - F^*)$, for all i and k satisfies Assumption 2. As for Case II, one is allowed to set $\delta_k^{(i)} = 0$ for some block(s)

i , or to set $\delta_k^{(i)} > \alpha(F(x_k) - F^*)$ for some blocks i and iterations k as long as Assumption 2 is satisfied.

4. **Case IV:** $\alpha > 0, \beta > 0$. This is the most general case, corresponding to the inclusion of both multiplicative and additive error, where the error level parameters α and β are fixed at the start of Algorithm 1. Notice that Assumption 2 is satisfied when the error $\delta_k^{(i)}$ in the computed update $T_{\delta}^{(i)}(x_k)$ obeys $\delta_k^{(i)} \leq \alpha(F(x_k) - F^*) + \beta$. Moreover, as for Cases II and III, one is allowed to set $\delta_k^{(i)} = 0$ for some block(s) i , or to set $\delta_k^{(i)} > \alpha(F(x_k) - F^*)$ for some blocks i and iterations k as long as Assumption 2 is satisfied. Notice that, as iterations progress, the multiplicative error α may become dominated by the additive error β , in the sense that, $\alpha(F(x_k) - F^*) \rightarrow 0$ as $k \rightarrow \infty$ so the upper bound on $\bar{\delta}_k$ tends to β .

Cases I–IV above show that the parameters α and β directly relate to the stopping criterion used in the algorithm employed to solve for the update $T_{\delta}^{(i)}(x_k)$ at each iteration of ICD. The following section gives examples of algorithms that can be used within ICD, where α and β can be given explicit numerical values and the stopping tolerances are verifiable.

3.4 Computing the inexact update

In this section we focus on the computation of the inexact update (Step 5 of Algorithm 1). We discuss several cases where it is possible to verify Assumption 2, and thus provide specific instances to show that ICD is indeed implementable.

In order to compute an inexact update, the quantity $V_i(x, T_0^{(i)}(x_k))$ in (18) is needed. Moreover, to incorporate multiplicative error, the optimal objective value F^* must also be known.

1. **F^* is known:** There are many instances when F^* is known a priori, which means that the bound $F(x_k) - F^*$ is computable at every iteration, and subsequently multiplicative error can be incorporated into ICD. In most cases, the update subproblem (18) has the same form as the original problem (see Section 6), so that $V_i(x_k, T_0^{(i)})$ is also known. This is the case, for example, when solving a consistent system of equations (minimizing a quadratic function), where $F^* = 0$, so that $F(x_k) - F^* = F(x_k)$. The update subproblem has the same form (a consistent system of equations must be solved to calculate $T_{\delta_k}^{(i)}$), so we also have $V_i(x_k, T_0^{(i)}) = 0$. Hence, for any $\alpha, \beta \geq 0$, one can compute $\delta_k^{(i)}$ (20), that satisfies Assumption 2 and (18). That is, at each iteration of ICD, accept any inexact update $T_{\delta_k}^{(i)}$ that satisfies:

$$V_i(x_k, T_{\delta_k}^{(i)}) - V_i(x_k, T_0^{(i)}) = V_i(x_k, T_{\delta_k}^{(i)}) \leq \delta_k^{(i)} \leq \alpha F(x_k) + \beta. \quad (21)$$

2. **Primal-dual algorithm:** If the update $T_{\delta_k}^{(i)}$ is found using an algorithm that terminates on the duality gap, then (18) and Assumption 2 are easy to verify. In particular, suppose that $\alpha = 0$ and $\beta > 0$ is fixed when initializing ICD. Then, for block i and iteration k , we accept $T_{\delta_k}^{(i)}$ such that

$$V_i(x_k, T_{\delta_k}^{(i)}) - V_i(x_k, T_0^{(i)}) \leq V_i(x_k, T_{\delta_k}^{(i)}) - V_i^{\text{DUAL}}(x_k, T_{\delta_k}^{(i)}) \leq \delta_k^{(i)} \leq \beta, \quad (22)$$

where $V_i^{\text{DUAL}}(x_k, T_{\delta_k}^{(i)})$ is the value of the dual at the point $T_{\delta_k}^{(i)}$.

Remark 3. 1. Both (21) and (22) are termination criteria for the iterative method used for the inner loop to determine the inexact update $T_\delta^{(i)}(x_k)$. They show that the error bound is indeed *implementable*, and that the inexactness parameters α and β relate to the stopping tolerance used in the inner loop of ICD.

2. Selecting an appropriate stopping criterion and tolerance (i.e., deciding upon numerical values for α and β) is a problem that frequently arises when solving optimization problems, and the decision is left to the discretion of the user.
3. It may be possible to find other stopping conditions such that (18) is verifiable. Moreover, it may be possible for the algorithm to converge *in practice* if a stopping condition is used for which (18) cannot be checked.

3.5 Technical result

The following result plays a key role in the complexity analysis of ICD.

Theorem 4. Fix $x_0 \in \mathbb{R}^N$ and let $\{x_k\}_{k \geq 0}$ be a sequence of random vectors in \mathbb{R}^N with x_{k+1} depending on x_k only. Let $\varphi : \mathbb{R}^N \rightarrow \mathbb{R}$ be a nonnegative function, define $\xi_k := \varphi(x_k)$ and assume that $\{\xi_k\}_{k \geq 0}$ is nonincreasing. Further, let $\rho \in (0, 1)$, $\epsilon > 0$ and $\alpha, \beta \geq 0$ be such that one of the following two conditions holds:

- (i) $\mathbb{E}[\xi_{k+1} | x_k] \leq (1 + \alpha)\xi_k - \frac{\xi_k^2}{c_1} + \beta$, for all $k \geq 0$, where $c_1 > 0$,
 $\frac{c_1}{2} \left(\alpha + \sqrt{\alpha^2 + \frac{4\beta}{c_1\rho}} \right) < \epsilon < \min\{(1 + \alpha)c_1, \xi_0\}$ and $\sigma := \sqrt{\alpha^2 + \frac{4\beta}{c_1}} < 1$;
- (ii) $\mathbb{E}[\xi_{k+1} | x_k] \leq \left(1 + \alpha - \frac{1}{c_2}\right)\xi_k + \beta$, for all $k \geq 0$ for which $\xi_k \geq \epsilon$,
 where $\alpha c_2 < 1 \leq (1 + \alpha)c_2$, and $\frac{\beta c_2}{\rho(1 - \alpha c_2)} < \epsilon < \xi_0$.

If (i) holds and we define $u := \frac{c_1}{2}(\alpha + \sigma)$ and choose

$$K \geq \frac{c_1}{\epsilon - \alpha c_1} \log \left(\frac{\epsilon - \frac{\beta c_1}{\epsilon - \alpha c_1}}{\epsilon \rho - \frac{\beta c_1}{\epsilon - \alpha c_1}} \right) + \min \left\{ \frac{1}{\sigma} \log \left(\frac{\xi_0 - u}{\epsilon - u} \right), \frac{c_1}{\epsilon - u} - \frac{c_1}{\xi_0 - u} \right\} + 2, \quad (23)$$

(where the second term in the minimum is chosen if $\sigma = 0$), or if (ii) holds and we choose

$$K \geq \frac{c_2}{1 - \alpha c_2} \log \left(\frac{\xi_0 - \frac{\beta c_2}{1 - \alpha c_2}}{\epsilon \rho - \frac{\beta c_2}{1 - \alpha c_2}} \right), \quad (24)$$

then $\mathbb{P}(\xi_K \leq \epsilon) \geq 1 - \rho$.

Proof. First notice that the thresholded sequence $\{\xi_k^\epsilon\}_{k \geq 0}$ defined by

$$\xi_k^\epsilon = \begin{cases} 0, & \text{if } \xi_k < \epsilon, \\ \xi_k, & \text{otherwise,} \end{cases} \quad (25)$$

satisfies $\xi_k^\epsilon > \epsilon \Leftrightarrow \xi_k > \epsilon$. Therefore, by Markov's inequality, $\mathbb{P}(\xi_k > \epsilon) = \mathbb{P}(\xi_k^\epsilon > \epsilon) \leq \frac{\mathbb{E}[\xi_k^\epsilon]}{\epsilon}$. Letting $\theta_k := \mathbb{E}[\xi_k^\epsilon]$, it thus suffices to show that

$$\theta_K \leq \epsilon \rho. \quad (26)$$

(The rationale behind this “thresholding trick” is that the sequence $\mathbb{E}[\xi_k^\epsilon]$ decreases faster than $\mathbb{E}[\xi_k]$ and hence will reach $\epsilon\rho$ sooner.) Assume now that (i) holds. It can be shown (for example, see Theorem 1 of [29] for the case $\alpha = \beta = 0$) that

$$\mathbb{E}[\xi_{k+1}^\epsilon | x_k] \leq (1 + \alpha)\xi_k^\epsilon - \frac{(\xi_k^\epsilon)^2}{c_1} + \beta, \quad \mathbb{E}[\xi_{k+1}^\epsilon | x_k] \leq \left(1 + \alpha - \frac{\epsilon}{c_1}\right) \xi_k^\epsilon + \beta. \quad (27)$$

By taking expectations in (27) (in x_k) and using Jensen’s inequality, we obtain

$$\theta_{k+1} \leq (1 + \alpha)\theta_k - \frac{\theta_k^2}{c_1} + \beta, \quad k \geq 0, \quad (28)$$

$$\theta_{k+1} \leq \left(1 + \alpha - \frac{\epsilon}{c_1}\right)\theta_k + \beta, \quad k \geq 0. \quad (29)$$

Notice that (28) is better than (29) precisely when $\theta_k > \epsilon$. It is easy to see that the inequality $(1 + \alpha)\theta_k - \frac{\theta_k^2}{c_1} + \beta \leq \theta_k$ holds if and only if $\theta_k \geq u$. In other words, (28) leads to θ_{k+1} that is better than θ_k only for $\theta_k \geq u$. We will now compute $k = k_1$ for which $u < \theta_k \leq \epsilon$. Inequality (28) can be equivalently written as

$$\theta_{k+1} - u \leq (1 - \sigma)(\theta_k - u) - \frac{(\theta_k - u)^2}{c_1}, \quad k \geq 0. \quad (30)$$

where $\sigma < 1$. Writing (28) in the form (30) eliminates the constant term β , which allows us to provide a simple analysis. (Moreover, this “shifted” form leads to a better result; see the remarks after the Theorem for details.) Letting $\hat{\theta}_k := \theta_k - u$, by monotonicity we have $\hat{\theta}_{k+1}\hat{\theta}_k \leq \hat{\theta}_k^2$, whence

$$\frac{1 - \sigma}{\hat{\theta}_{k+1}} - \frac{1}{\hat{\theta}_k} = \frac{(1 - \sigma)\hat{\theta}_k - \hat{\theta}_{k+1}}{\hat{\theta}_{k+1}\hat{\theta}_k} \geq \frac{(1 - \sigma)\hat{\theta}_k - \hat{\theta}_{k+1}}{\hat{\theta}_k^2} \stackrel{(30)}{\geq} \frac{1}{c_1}. \quad (31)$$

If we choose $r \in \{1, \frac{1}{1-\sigma}\}$, then

$$\frac{1}{\hat{\theta}_k} \stackrel{(31)}{\geq} r \left(\frac{1}{\hat{\theta}_{k-1}} + \frac{1}{c_1} \right) \geq r^k \frac{1}{\hat{\theta}_0} + \frac{1}{c_1} \sum_{j=1}^k r^j = \begin{cases} r^k \left(\frac{1}{\xi_0 - u} + \frac{1}{c_1 \sigma} \right) - \frac{1}{c_1 \sigma}, & r = \frac{1}{1-\sigma}, \\ \frac{1}{\xi_0 - u} + \frac{k}{c_1}, & r = 1. \end{cases}$$

In particular, using the above estimate with $r = 1$ and $r = \frac{1}{1-\sigma}$ gives

$$\hat{\theta}_{k_1} \leq \epsilon - u \quad (\text{and hence } \theta_{k_1} \leq \epsilon) \quad (32)$$

for

$$k_1 := \min \left\{ \left\lceil \log \left(\frac{\frac{1}{\epsilon - u} + \frac{1}{c_1 \sigma}}{\frac{1}{\xi_0 - u} + \frac{1}{c_1 \sigma}} \right) / \log \left(\frac{1}{1 - \sigma} \right) \right\rceil, \left\lceil \frac{c_1}{\epsilon - u} - \frac{c_1}{\xi_0 - u} \right\rceil \right\}, \quad (33)$$

where the left term in (33) applies when $\sigma > 0$ only.

Applying the inequalities (i) $\lceil t \rceil \leq 1 + t$; (ii) $\log(\frac{1}{1-t}) \geq t$ (holds for $0 < t < 1$; we use the inverse version, which is surprisingly tight for small t); and (iii) the fact that $t \mapsto \frac{C+t}{D+t}$ is decreasing on $[0, \infty)$ if $C \geq D > 0$, we arrive at the following bound

$$k_1 \geq 1 + \min \left\{ \frac{1}{\sigma} \log \left(\frac{\xi_0 - u}{\epsilon - u} \right), \frac{c_1}{\epsilon - u} - \frac{c_1}{\xi_0 - u} \right\}. \quad (34)$$

Letting $\gamma := 1 - \frac{\epsilon - \alpha c_1}{c_1}$ (notice that $\gamma \in (0, 1)$), for any $k_2 \geq 0$ we have

$$\begin{aligned} \theta_{k_1+k_2} &\stackrel{(29)}{\leq} \gamma \theta_{k_1+k_2-1} + \beta \leq \gamma^{k_2} \theta_{k_1} + \beta(\gamma^{k_2-1} + \gamma^{k_2-2} + \dots + 1) \\ &\stackrel{(32)}{\leq} \gamma^{k_2} \epsilon + \beta \frac{1 - \gamma^{k_2}}{1 - \gamma} = \gamma^{k_2} \left(\epsilon - \frac{\beta}{1 - \gamma} \right) + \frac{\beta}{1 - \gamma}. \end{aligned} \quad (35)$$

In (35), notice that the second to last term can be made as small as we like (by taking k_2 large), but we can never force $\theta_{k_1+k_2} \leq \frac{\beta}{1-\gamma}$. Therefore, in order to establish (26), we need to ensure that $\frac{\beta c_1}{\epsilon - \alpha c_1} < \epsilon \rho$. Rearranging this gives the condition $\frac{c_1}{2}(\alpha + \sqrt{\alpha^2 + \frac{4\beta}{c_1 \rho}}) < \epsilon$, which holds by assumption. Now we can find k_2 for which the right hand side in (35) is at most $\epsilon \rho$:

$$k_2 := \left\lceil \log \left(\frac{\epsilon - \frac{\beta}{1-\gamma}}{\epsilon \rho - \frac{\beta}{1-\gamma}} \right) / \log \left(\frac{1}{\gamma} \right) \right\rceil \leq 1 + \frac{c_1}{\epsilon - \alpha c_1} \log \left(\frac{\epsilon - \frac{\beta c_1}{\epsilon - \alpha c_1}}{\epsilon \rho - \frac{\beta c_1}{\epsilon - \alpha c_1}} \right). \quad (36)$$

In view of (26), it is enough to take $K = k_1 + k_2$ iterations. The expression in (23) is obtained by adding the upper bounds on k_1 and k_2 in (34) and (36).

Now assume that property (ii) holds. By a similar argument as that leading to (27), we obtain

$$\begin{aligned} \theta_K &\leq \left(1 - \frac{1 - \alpha c_2}{c_2} \right) \theta_{K-1} + \beta \leq \left(1 - \frac{1 - \alpha c_2}{c_2} \right)^K \theta_0 + \beta \sum_{j=0}^{K-1} \left(1 - \frac{1 - \alpha c_2}{c_2} \right)^j \\ &\leq \left(1 - \frac{1 - \alpha c_2}{c_2} \right)^K \left(\theta_0 - \frac{\beta c_2}{1 - \alpha c_2} \right) + \frac{\beta c_2}{1 - \alpha c_2} \stackrel{(24)}{\leq} \epsilon \rho. \end{aligned}$$

The proof follows by taking K given by (24). \square

Let us now comment on several aspects of the above result:

1. *Usage.* We will use Theorem 4 to finish the proofs of the complexity results in Section 4; with $\xi_k = \varphi(x_k) := F(x_k) - F^*$, where $\{x_k\}$ is the random process generated by ICD.
2. *Monotonicity and Nonnegativity.* Note that the monotonicity assumption in Theorem 4 is for the choice of x_k and φ described in 1) satisfied due to (19). Nonnegativity is satisfied automatically since $F(x_k) \geq F^*$ for all x_k .
3. *Best of two.* In (33), we notice that the first term applies when $\sigma > 0$ only. If $\sigma = 0$, then $u = 0$, and subsequently the second term in (33) applies, which corresponds to the exact case. Notice that if $\sigma > 0$ is very small (so $u \neq 0$), the iteration complexity result still may be better if the second term is used.
4. *Generalization.* Note that for $\alpha = \beta = 0$, (23) recovers $\frac{c_1}{\epsilon}(1 + \log \frac{1}{\rho}) + 2 - \frac{c_1}{\xi_0}$, which is the result proved in Theorem 1(i) in [29], while (24) recovers $c_2 \log((F(x_0) - F^*)/\epsilon \rho)$, which is the result proved in Theorem 1(ii) in [29]. Since the last term in (23) is negative, the theorem holds also if we ignore it. This is what we have done, for simplicity, in Table 1.
5. *High accuracy with high probability.* In the exact case, the iteration complexity results hold for any error tolerance $\epsilon > 0$ and confidence $\rho \in (0, 1)$. However, in the inexact case, there

are restrictions on the choice of ρ and ϵ for which we can guarantee the result $\mathbb{P}(F(x_k) - F^* \leq \epsilon) \geq 1 - \rho$. Table 5 gives conditions on α and β under which arbitrary confidence level (i.e., small ρ) and accuracy (i.e., small ϵ) is achievable. For instance, if Theorem 4(ii) is used, then one can achieve arbitrary accuracy only if $\beta = 0$, but arbitrary confidence under no assumptions on α and β . The situation for part (i) is worse: ϵ is lower bounded by a positive expression that involves ρ , unless $\alpha = \beta = 0$.

	Theorem 4(i)	Theorem 4(ii)
ϵ can be arbitrarily small if	$\alpha = \beta = 0$	$\beta = 0$
ρ can be arbitrarily small if	$\beta = 0$	any α, β

Table 2: The conditions under which arbitrary confidence ρ and accuracy ϵ are attainable.

6. *Two lower bounds on ϵ .* The inequality $\epsilon > \frac{c_1}{2} \left(\alpha + \sqrt{\alpha^2 + \frac{4\beta}{\rho c_1}} \right)$ (see part (i) of Theorem 4) is equivalent to $\epsilon > \frac{\beta c_1}{\rho(\epsilon - \alpha c_1)}$. Note the similarity of the last expression and the lower bound on ϵ in part (ii) of the theorem. We can see that the lower bound on ϵ is smaller (and hence, is less restrictive) in (ii) than in (i), provided that $c_1 = c_2$.
7. *Two analyses.* It can be seen that analyzing the “shifted” form (30) leads to a better result than analyzing (28) directly, even when $\beta = 0$. Consider the case $\beta = 0$, so that $\sigma = \alpha$ and $u = \alpha c_1$. From equation (31) $\theta_{k+1} \leq A := \alpha c_1 + (1 - \alpha) / (\frac{1}{\theta_k - \alpha c_1} + \frac{1}{c_1})$, whereas analyzing equation (28) directly yields $\theta_{k+1} \leq B := (1 + \alpha) / (\frac{1}{\theta_k} + \frac{1}{c_1})$. It can be shown that $A \leq B$, with equality if $\alpha = 0$.

4 Complexity Analysis: Convex Composite Objective

The following function plays a central role in our analysis:

$$H(x, T) := f(x) + \langle \nabla f(x), T \rangle + \frac{1}{2} \|T\|_l^2 + \Psi(x + T). \quad (37)$$

Comparing (37) with (16) using (2), (3), (6), (8) and (9) we get

$$H(x, T) = f(x) + \sum_{i=1}^n V_i(x, T^{(i)}). \quad (38)$$

It will be useful to establish inequalities relating H evaluated at the vector of exact updates $T_0(x)$ and H evaluated at the vector of inexact updates $T_\delta(x)$.

Lemma 5. *For all $x \in \mathbb{R}^N$ and $\delta \in \mathbb{R}_+^n$,*

$$H(x, T_0(x)) \leq H(x, T_\delta(x)) \leq H(x, T_0(x)) + \sum_{i=1}^n \delta^{(i)}. \quad (39)$$

Proof:

$$\begin{aligned}
H(x, T_0(x)) &\stackrel{(38)}{=} f(x) + \sum_{i=1}^n V_i(x, T_0^{(i)}(x)) \stackrel{(18)}{=} f(x) + \sum_{i=1}^n \min_{t \in \mathbb{R}^{N_i}} V_i(x, t) \\
&\leq f(x) + \sum_{i=1}^n V_i(x, T_\delta^{(i)}(x)) \stackrel{(38)}{=} H(x, T_\delta(x)) \\
&\stackrel{(18)}{\leq} f(x) + \sum_{i=1}^n \left(\delta^{(i)} + \min_{t \in \mathbb{R}^{N_i}} V_i(x, t) \right) \stackrel{(38)}{=} H(x, T_0(x)) + \sum_{i=1}^n \delta^{(i)}. \quad \square
\end{aligned}$$

The following Lemma provides an upper bound on the expected distance between the current and optimal objective value in terms of the function H .

Lemma 6. For $x, T \in \mathbb{R}^N$, let $x_+(x, T)$ be the random vector equal to $x + U_i T^{(i)}$ with probability $\frac{1}{n}$ for each $i \in \{1, 2, \dots, n\}$. Then

$$\mathbb{E}[F(x_+(x, T)) - F^* \mid x] \leq \frac{1}{n}(H(x, T) - F^*) + \frac{n-1}{n}(F(x) - F^*).$$

Proof:

$$\begin{aligned}
\mathbb{E}[F(x_+(x, T)) \mid x] &= \sum_{i=1}^n \frac{1}{n} F(x + U_i T^{(i)}) \\
&\stackrel{(15)}{\leq} \frac{1}{n} \sum_{i=1}^n [f(x) + V_i(x, T^{(i)}) + \psi_i(x)] \\
&\stackrel{(38)+(17)}{=} \frac{1}{n} H(x, T) + \frac{n-1}{n} f(x) + \frac{1}{n} \sum_{i=1}^n \sum_{j \neq i} \Psi_j(x^{(j)}) \\
&= \frac{1}{n} H(x, T) + \frac{n-1}{n} F(x). \quad \square
\end{aligned}$$

Note that if $x = x_k$ and $T = T_\delta(x_k)$, then $x_+(x, T) = x_{k+1}$, as produced by Algorithm 1. The following Lemma, which provides an upper bound on H , will be used repeatedly throughout the remainder of this paper.

Lemma 7. For all $x \in \text{dom } F$ and $\delta \in \mathbb{R}_+^n$ (letting $\Delta = \sum_i \delta^{(i)}$), we have

$$H(x, T_\delta(x)) \leq \Delta + \min_{y \in \mathbb{R}^N} \left\{ F(y) + \frac{1-\mu_f(l)}{2} \|y - x\|_l^2 \right\}. \quad (40)$$

Proof:

$$\begin{aligned}
H(x, T_\delta(x)) &\stackrel{(39)}{\leq} \Delta + \min_{T \in \mathbb{R}^N} H(x, T) \\
&= \Delta + \min_{y \in \mathbb{R}^N} H(x, y - x) \quad (\text{where } y = x + T) \\
&\stackrel{(37)}{=} \Delta + \min_{y \in \mathbb{R}^N} \left\{ f(x) + \langle \nabla f(x), y - x \rangle + \Psi(y) + \frac{1}{2} \|y - x\|_l^2 \right\} \\
&\stackrel{(10)}{\leq} \Delta + \min_{y \in \mathbb{R}^N} \left\{ f(y) - \frac{\mu_f(l)}{2} \|y - x\|_l^2 + \Psi(y) + \frac{1}{2} \|y - x\|_l^2 \right\}. \quad \square
\end{aligned}$$

4.1 Convex case

Now we need to estimate $H(x, T_\delta(x)) - F^*$ from above in terms of $F(x) - F^*$.

Lemma 8. Fix $x^* \in X^*$, $x \in \text{dom } F$, $\delta \in \mathbb{R}_+^n$ and let $R = \|x - x^*\|_l$ and $\Delta = \sum_i \delta^{(i)}$. Then

$$H(x, T_\delta(x)) - F^* \leq \Delta + \begin{cases} (1 - \frac{F(x) - F^*}{2R^2})(F(x) - F^*), & \text{if } F(x) - F^* \leq R^2, \\ \frac{1}{2}R^2 < \frac{1}{2}(F(x) - F^*), & \text{otherwise.} \end{cases} \quad (41)$$

Proof: Because strong convexity is not assumed, $\mu_f(l) = 0$, so

$$\begin{aligned} H(x, T_\delta(x)) &\stackrel{(40)}{\leq} \Delta + \min_{y \in \mathbb{R}^N} \{F(y) + \frac{1}{2}\|y - x\|_l^2\} \\ &\leq \Delta + \min_{\lambda \in [0,1]} \{F(\lambda x^* + (1 - \lambda)x) + \frac{\lambda^2}{2}\|x - x^*\|_l^2\} \\ &\leq \Delta + \min_{\lambda \in [0,1]} \{F(x) - \lambda(F(x) - F^*) + \frac{\lambda^2}{2}R^2\}. \end{aligned}$$

Minimizing in λ gives $\lambda^* = \min\{1, (F(x) - F^*)/R^2\}$ and the result follows. \square

We now state the main complexity result of this section, which bounds the number of iterations sufficient for ICD used with uniform probabilities to decrease the value of the objective to within ϵ of the optimal value with probability at least $1 - \rho$.

Theorem 9. Choose an initial point $x_0 \in \mathbb{R}^N$ and let $\{x_k\}_{k \geq 0}$ be the random iterates generated by ICD applied to problem (1), using uniform probabilities $p_i = \frac{1}{n}$ and inexactness parameters $\delta_k^{(1)}, \dots, \delta_k^{(n)} \geq 0$ that satisfy (20) for $\alpha, \beta \geq 0$. Choose target confidence $\rho \in (0, 1)$ and error tolerance $\epsilon > 0$ so that one of the following two conditions hold:

$$(i) \frac{c_1}{2}(\alpha + \sqrt{\alpha^2 + \frac{4\beta}{c_1\rho}}) < \epsilon < F(x_0) - F^* \text{ and } \alpha^2 + \frac{4\beta}{c_1} < 1, \text{ where } c_1 = 2n \max\{\mathcal{R}_l^2(x_0), F(x_0) - F^*\},$$

$$(ii) \frac{\beta c_2}{\rho(1 - \alpha c_2)} < \epsilon < \min\{\mathcal{R}_l^2(x_0), F(x_0) - F^*\}, \text{ where } c_2 = \frac{2n\mathcal{R}_l^2(x_0)}{\epsilon} \text{ and } \alpha c_2 < 1.$$

If (i) holds and we choose K as in (23), or if (ii) holds and we choose K as in (24), then $\mathbb{P}(F(x_K) - F^* \leq \epsilon) \geq 1 - \rho$.

Proof. Since $F(x_k) \leq F(x_0)$ for all k by (19), we have $\|x_k - x^*\|_l \leq \mathcal{R}_l(x_0)$ for all k and $x^* \in X^*$. Using Lemma 6 and Lemma 8, and letting $\xi_k := F(x_k) - F^*$, we have

$$\mathbb{E}[\xi_{k+1} | x_k] \leq \bar{\delta}_k + \frac{1}{n} \max\left\{1 - \frac{\xi_k}{2\|x_k - x^*\|_l^2}, \frac{1}{2}\right\} \xi_k + \frac{n-1}{n} \xi_k \quad (42)$$

$$\begin{aligned} &= \bar{\delta}_k + \max\left\{1 - \frac{\xi_k}{2n\|x_k - x^*\|_l^2}, 1 - \frac{1}{2n}\right\} \xi_k \\ &\leq \bar{\delta}_k + \max\left\{1 - \frac{\xi_k}{2n\mathcal{R}_l^2(x_0)}, 1 - \frac{1}{2n}\right\} \xi_k. \end{aligned} \quad (43)$$

Consider case (i). From (43) and (20) we obtain

$$\mathbb{E}[\xi_{k+1} | x_k] \leq \bar{\delta}_k + \left(1 - \frac{\xi_k}{c_1}\right) \xi_k \leq (1 + \alpha)\xi_k - \frac{\xi_k^2}{c_1} + \beta, \quad (44)$$

and the result follows by applying Theorem 4(i). Now consider case (ii). Notice that if $\xi_k \geq \epsilon$, then (43) together with (20), imply that

$$\mathbb{E}[\xi_{k+1} | x_k] \leq \bar{\delta}_k + \max\left\{1 - \frac{\epsilon}{2n\mathcal{R}_l^2(x_0)}, 1 - \frac{1}{2n}\right\} \xi_k \leq \left(1 + \alpha - \frac{1}{c_2}\right) \xi_k + \beta.$$

The result follows by applying Theorem 4(ii). \square

4.2 Strongly convex case

Let us start with an auxiliary result.

Lemma 10. *Let F be strongly convex with respect to $\|\cdot\|_l$ with $\mu_f(l) + \mu_\Psi(l) > 0$. Then for all $x \in \text{dom } F$ and $\delta \in \mathbb{R}_+^n$, with $\Delta = \sum_i \delta^{(i)}$, we have*

$$H(x, T_\delta(x)) - F^* \leq \Delta + \left(\frac{1 - \mu_f(l)}{1 + \mu_\Psi(l)} \right) (F(x) - F^*).$$

Proof: Let $\mu_f = \mu_f(l)$, $\mu_\Psi = \mu_\Psi(l)$ and $\lambda^* = (\mu_f + \mu_\Psi)/(1 + \mu_\Psi) \leq 1$. Then,

$$\begin{aligned} H(x, T_\delta(x)) &\stackrel{(40)}{\leq} \Delta + \min_{y \in \mathbb{R}^N} \left\{ F(y) + \frac{1 - \mu_f}{2} \|y - x\|_l^2 \right\} \\ &\leq \Delta + \min_{\lambda \in [0, 1]} \left\{ F(\lambda x^* + (1 - \lambda)x) + \frac{(1 - \mu_f)\lambda^2}{2} \|x - x^*\|_l^2 \right\} \\ &\stackrel{(11)+(13)}{\leq} \Delta + \min_{\lambda \in [0, 1]} \left\{ \lambda F^* + (1 - \lambda)F(x) + \frac{(1 - \mu_f)\lambda^2 - (\mu_f + \mu_\Psi)\lambda(1 - \lambda)}{2} \|x - x^*\|_l^2 \right\} \\ &\leq \Delta + F(x) - \lambda^*(F(x) - F^*). \end{aligned}$$

The last inequality follows from the fact that $(\mu_f + \mu_\Psi)(1 - \lambda^*) - (1 - \mu_f)\lambda^* = 0$. It remains to subtract F^* from both sides of the final inequality. \square

We can now estimate the number of iterations needed to decrease a strongly convex objective F within ϵ of the optimal value with high probability.

Theorem 11. *Let F be strongly convex with respect to the norm $\|\cdot\|_l$ with $\mu_f(l) + \mu_\Psi(l) > 0$ and let $\mu := \frac{\mu_f(l) + \mu_\Psi(l)}{1 + \mu_\Psi(l)}$. Choose an initial point $x_0 \in \mathbb{R}^N$ and let $\{x_k\}_{k \geq 0}$, be the random iterates generated by ICD applied to problem (1), used with uniform probabilities $p_i = \frac{1}{n}$ for $i = 1, 2, \dots, n$ and inexactness parameters $\delta_k^{(1)}, \dots, \delta_k^{(n)} \geq 0$ satisfying (20), for $0 \leq \alpha < \frac{\mu}{n}$ and $\beta \geq 0$. Choose confidence level $\rho \in (0, 1)$ and error tolerance ϵ satisfying $\frac{\beta n}{\rho(\mu - \alpha n)} < \epsilon < F(x_0) - F^*$. Then for K given by (24), we have $\mathbb{P}(F(x_K) - F^* \leq \epsilon) \geq 1 - \rho$.*

Proof. Letting $\xi_k = F(x_k) - F^*$, we have

$$\begin{aligned} \mathbb{E}[\xi_{k+1} \mid x_k] &\stackrel{(\text{Lemma 6})}{\leq} \frac{1}{n} (H(x_k, T_{\delta_k}(x_k)) - F^*) + \frac{n-1}{n} \xi_k \\ &\stackrel{(\text{Lemma 10})}{\leq} \bar{\delta}_k + \frac{1}{n} \left(\frac{1 - \mu_f(l)}{1 + \mu_\Psi(l)} \xi_k \right) + \frac{n-1}{n} \xi_k \\ &\stackrel{(20)}{\leq} \left(1 + \alpha - \frac{\mu}{n} \right) \xi_k + \beta. \end{aligned}$$

By (12), $\mu \leq 1$, and the result follows from Theorem 4(ii) with $c_2 = \frac{n}{\mu}$. \square

5 Complexity Analysis: Smooth Objective

In this section we provide simplified iteration complexity results when the objective function is smooth ($\Psi \equiv 0$ so $F \equiv f$). Furthermore, we provide complexity results for arbitrary (rather than uniform) probabilities $p_i > 0$.

5.1 Convex case

In the smooth exact case, we can write down a closed-form expression for the update:

$$T_0^{(i)}(x) \stackrel{(18)}{=} \arg \min_{t \in \mathbb{R}^{N_i}} V_i(x, t) \stackrel{(16)}{=} \arg \min_{t \in \mathbb{R}^{N_i}} \{ \langle \nabla_i f(x), t \rangle + \frac{l_i}{2} \|t\|_{(i)}^2 \} = -\frac{1}{l_i} B_i^{-1} \nabla_i f(x).$$

Substituting this into $V_i(x, \cdot)$ yields

$$V_i(x, T_0^{(i)}(x)) = \langle \nabla_i f(x), T_0^{(i)}(x) \rangle + \frac{l_i}{2} \|T_0^{(i)}(x)\|_{(i)}^2 = -\frac{1}{2l_i} (\|\nabla_i f(x)\|_{(i)}^*)^2. \quad (45)$$

We can now estimate the decrease in f during one iteration of ICD:

$$\begin{aligned} f(x + U_i T_\delta^{(i)}(x)) - f(x) &\stackrel{(7)}{\leq} \langle \nabla_i f(x), T_\delta^{(i)}(x) \rangle + \frac{l_i}{2} \|T_\delta^{(i)}(x)\|_{(i)}^2 \\ &\stackrel{(16)}{=} V_i(x, T_\delta^{(i)}(x)) \\ &\stackrel{(18)}{\leq} \min\{0, \delta^{(i)} + V_i(x, T_0^{(i)}(x))\} \\ &\stackrel{(45)}{=} \min\{0, \delta^{(i)} - \frac{1}{2l_i} (\|\nabla_i f(x)\|_{(i)}^*)^2\}. \end{aligned} \quad (46)$$

The main iteration complexity result of this section can be now established.

Theorem 12. *Choose an initial point $x_0 \in \mathbb{R}^N$ and let $\{x_k\}_{k \geq 0}$ be the random iterates generated by ICD applied to the problem of minimizing f , used with probabilities $p_1, \dots, p_n > 0$ and inexactness parameters $\delta_k^{(1)}, \dots, \delta_k^{(n)} \geq 0$ satisfying (20) for $\alpha, \beta \geq 0$, where $\alpha^2 + \frac{4\beta}{c_1} < 1$ and $c_1 = 2\mathcal{R}_{l_{p-1}}^2(x_0)$. Choose target confidence $\rho \in (0, 1)$, error tolerance ϵ satisfying $\frac{c_1}{2}(\alpha + \sqrt{\alpha^2 + \frac{4\beta}{c_1\rho}}) < \epsilon < f(x_0) - f^*$, and let the iteration counter K be given by (23). Then $\mathbb{P}(f(x_K) - f^* \leq \epsilon) \geq 1 - \rho$.*

Proof. We first estimate the expected decrease of the objective function during one iteration of the method:

$$\begin{aligned} \mathbb{E}[f(x_{k+1}) \mid x_k] &= f(x_k) + \sum_{i=1}^n p_i [f(x_k + U_i T_{\delta_k^{(i)}}^{(i)}(x_k)) - f(x_k)] \\ &\stackrel{(46)}{\leq} f(x_k) + \sum_{i=1}^n p_i \left(\delta_k^{(i)} - \frac{1}{2l_i} (\|\nabla_i f(x_k)\|_{(i)}^*)^2 \right) \\ &\stackrel{(9)}{=} f(x_k) - \frac{1}{2} (\|\nabla f(x_k)\|_{l_{p-1}}^*)^2 + \sum_{i=1}^n p_i \delta_k^{(i)} \\ &\leq f(x_k) - \frac{1}{2} (\|\nabla f(x_k)\|_{l_{p-1}}^*)^2 + \alpha(f(x_k) - f^*) + \beta. \end{aligned} \quad (47)$$

Since $f(x_k) \leq f(x_0)$ for all k ,

$$f(x_k) - f^* \leq \max_{x^* \in X^*} \langle \nabla f(x_k), x_k - x^* \rangle \leq \|\nabla f(x_k)\|_{l_{p-1}}^* \mathcal{R}_{l_{p-1}}(x_0). \quad (48)$$

Substituting (48) into (47) we obtain

$$\mathbb{E}[f(x_{k+1}) - f^* \mid x_k] \leq f(x_k) - f^* - \frac{1}{2} \left(\frac{f(x_k) - f^*}{\mathcal{R}_{l_{p-1}}(x_0)} \right)^2 + \alpha(f(x_k) - f^*) + \beta. \quad (49)$$

It remains to apply Theorem 4(i). □

5.2 Strongly convex case

In this section we assume that f is strongly convex with respect to $\|\cdot\|_{lp^{-1}}$ with convexity parameter $\mu_f(lp^{-1})$. Using (10) with $x = x_k$ and $y = x^*$, and letting $h = x^* - x_k$, we obtain

$$\begin{aligned} f^* - f(x_k) &\geq \langle \nabla f(x_k), h \rangle + \frac{\mu_f(lp^{-1})}{2} \|h\|_{lp^{-1}}^2 \\ &= \mu_f(lp^{-1}) \left(\left\langle \frac{1}{\mu_f(lp^{-1})} \nabla f(x_k), h \right\rangle + \frac{1}{2} \|h\|_{lp^{-1}}^2 \right). \end{aligned} \quad (50)$$

By minimizing the right hand side of (50), and rearranging, we obtain

$$f(x_k) - f^* \leq \frac{1}{2\mu_f(lp^{-1})} (\|\nabla f(x_k)\|_{lp^{-1}}^*)^2. \quad (51)$$

We can now give an efficiency estimate for the case of a strongly convex objective.

Theorem 13. *Let f be strongly convex with respect to the norm $\|\cdot\|_{lp^{-1}}$ with convexity parameter $\mu_f(lp^{-1}) > 0$. Choose an initial point $x_0 \in \mathbb{R}^N$ and let $\{x_k\}_{k \geq 0}$ be the random iterates generated by ICD applied to the problem of minimizing f , used with probabilities $p_1, \dots, p_n > 0$ and inexactness parameters $\delta_k^{(1)}, \dots, \delta_k^{(n)} \geq 0$ that satisfy (20) for $0 \leq \alpha < \mu_f(lp^{-1})$ and $\beta \geq 0$. Choose the target confidence $\rho \in (0, 1)$, let the target accuracy ϵ satisfy $\frac{\beta}{\rho(\mu_f(lp^{-1}) - \alpha)} < \epsilon < f(x_0) - f^*$, let $c_2 = 1/\mu_f(lp^{-1})$ and let iteration counter K be as in (24). Then $\mathbb{P}(f(x_K) - f^* \leq \epsilon) \geq 1 - \rho$.*

Proof. The expected decrease of the objective function during one iteration of the method can be estimated as follows:

$$\begin{aligned} \mathbb{E}[f(x_{k+1}) - f^* | x_k] &\stackrel{(47)}{\leq} (1 + \alpha)(f(x_k) - f^*) - \frac{1}{2} (\|\nabla f(x_k)\|_{lp^{-1}}^*)^2 + \beta \\ &\stackrel{(51)}{\leq} (1 + \alpha - \mu_f(lp^{-1}))(f(x_k) - f^*) + \beta \end{aligned}$$

It remains to apply Theorem 4(ii) with $\varphi(x_k) = f(x_k) - f^*$ (and notice that $c_2 > 1$ by (12)). \square

6 Practical aspects of an inexact update

The goal of the second part of this paper is to demonstrate the practical importance of employing an inexact update in the (block) ICD method.

6.1 Solving smooth problems via ICD

In the first part of this section we assume that $\Psi = 0$, so the function $F(x) = f(x)$ is smooth and convex. In this case the overapproximation is

$$F(x_k + U_i t) = f(x_k + U_i t) \stackrel{(7)+(4)}{\leq} f(x_k) + \langle \nabla_i f(x_k), t \rangle + \frac{l_i}{2} \langle B_i t, t \rangle \equiv f(x_k) + V_i(x_k, t). \quad (52)$$

Differentiating (52) with respect to t and setting the result to 0, shows that determining the update to block i at iteration k is equivalent to solving the system of equations

$$B_i t = -\frac{1}{l_i} \nabla_i f(x_k). \quad (53)$$

Recall that B_i is positive definite so the exact update is

$$T_0^{(i)}(x_k) = -\frac{1}{l_i} B_i^{-1} \nabla_i f(x_k), \quad (54)$$

and, as mentioned in Section 3.4, $V_i(x_k, T_0^{(i)}(x_k)) = 0$. Clearly, solving systems of equations is central to the block coordinate descent method in the smooth case.

Exact CD [28] requires the exact update (54), which depends on the inverse of an $N_i \times N_i$ matrix. A standard approach to solving for $T_0^{(i)}(x_k)$ in (54) is to form the Cholesky factors of B_i followed by two triangular solves. This can be extremely expensive for medium N_i , or dense B_i .

The results in this work allow (53) to be solved using an iterative method to find an *inexact* update $T_{\delta_k}^{(i)}(x_k)$. If we compute t for which $V_i(x_k, t) - V_i(x_k, T_0^{(i)}(x_k)) = V_i(x_k, t) = \|B_i t - \frac{1}{l_i} \nabla_i f(x_k)\|_2^2 \leq \beta$, then we terminate the iterative method and accept the inexact update $T_{\delta_k}^{(i)} \equiv t$.¹

Because B_i is positive definite, a natural choice is to solve (52) using conjugate gradients [13]. (This is the method we adopt in the numerical experiments presented in Section 7). It is widely accepted that using an iterative technique has many advantages over a direct method for solving systems of equations, so we expect that an inexact update can be determined quickly, and subsequently the overall ICD algorithm running time reduces. Moreover, applying a preconditioner to (52) can enable even faster convergence of conjugate gradients. Finding good preconditioners is an active area of research; see for example [2, 10, 12].

6.1.1 A special case: a quadratic function

A special case of the above is when we have the unconstrained quadratic minimization problem

$$\min_{x \in \mathbb{R}^N} f(x) = \frac{1}{2} \|Ax - b\|_2^2, \quad (55)$$

where $A \in \mathbb{R}^{M \times N}$, and $b \in \mathbb{R}^M$. In this case, the overapproximation (7) becomes

$$f(x + U_i t) = \frac{1}{2} \|A(x + U_i t) - b\|_2^2 = f(x) + \langle \nabla_i f(x), t \rangle + \frac{1}{2} \langle A_i^T A_i t, t \rangle, \quad (56)$$

where $A_i = U_i A$. Comparing (56) with (52), we see that in the quadratic case, (56) is an exact upper bound on $f(x + U_i t)$ if we choose $l_i = 1$ and $B_i = A_i^T A_i$ for all blocks $i = 1, \dots, n$. The matrix B_i is required to be (strictly) positive definite so A_i is assumed to have full (column) rank.² Substituting $l_i = 1$ and $B_i = A_i^T A_i$ into (53) gives

$$A_i^T A_i t = -A_i^T (Ax - b). \quad (57)$$

Therefore, when ICD is applied to a problem of the form (55), the update is found by solving (57).

6.2 Solving nonsmooth problems via ICD

The nonsmooth case is not as simple as the smooth case, because the update subproblem will have a different form for each nonsmooth term Ψ . However, we will see that in many cases, the subproblem will have the same, or similar, form to the original objective function. We demonstrate this through the use of the following concrete examples.

¹Note that if f is a quadratic function corresponding to a consistent system of equations, then we could have used the more general stopping condition $\alpha F(x_k) + \beta$, because then we also know that $F^* = 0$.

²If a block A_i does not have full column rank then we simply adjust our choice of l_i and B_i accordingly, although this means that we have an overapproximation to $f(x + U_i t)$, rather than equality as in (56).

6.2.1 Group Lasso

A widely studied optimization problem arising in statistics and machine learning is the so-called group lasso problem, which has the form

$$\min_{x \in \mathbb{R}^N} \frac{1}{2} \|Ax - b\|_2^2 + \lambda \sum_{i=1}^n \sqrt{d_i} \|x^{(i)}\|_2, \quad (58)$$

where $\lambda > 0$ is a regularization parameter and d_i for all i is a weighting parameter that depends on the size of the i th block. Formulation (58) fits the structure (1) with $f(x) = \frac{1}{2} \|Ax - b\|_2^2$ and $\Psi(x) = \sum_{i=1}^n \lambda \sqrt{d_i} \|x^{(i)}\|_2$. It can be shown that choosing $B_i = A_i^T A_i$ ³ and $l_i = 1$ for all i , satisfies the overapproximation (15) giving $F(x_k + U_i t) \leq f(x_k) + \langle A_i^T r_k, t \rangle + \frac{1}{2} \langle A_i^T A_i t, t \rangle + \lambda \sqrt{d_i} \|x_k^{(i)} + t\|_2$, where $r_k = Ax_k - b$, so

$$V_i(x_k, t) = \frac{1}{2} \|A_i t - r_k\|_2^2 + \lambda \sqrt{d_i} \|x_k^{(i)} + t\|_2. \quad (59)$$

We see that (after a simple change of variables) (59) has the same form as the original problem (58). We can apply *any* algorithm to approximately minimize (59) that uses one of the stopping conditions described in Section 3.4.

7 Numerical Experiments

In this section we present preliminary numerical results to demonstrate the practical performance of Inexact Coordinate Descent and compare the results with Exact Coordinate Descent. We note that a thorough practical investigation of Exact CD is given in [29] where its usefulness on huge-scale problems is evidenced. We do not intend to reproduce such results for ICD, rather, we investigate the affect of inexact updates compared with exact updates, which should be apparent on medium scale problems. We do this the full knowledge that if exact CD scales well to very large sizes (shown in [29]) then so too will ICD.

Each experiment presented in this section was implemented in MATLAB and run (under linux) on a desktop computer with a quad core i5-3470CPU, 3.20GHz processor with 24Gb of RAM.

7.1 Problem description for a smooth objective

In this numerical experiment, we assume that the function $F = f$ is quadratic (55) and $\Psi = 0$. Further, as ICD can work with blocks of data, we impose block structure on the system matrix. In particular, we assume that the matrix A has block angular structure. Matrices with this structure frequently arise in optimization, from optimal control, scheduling and planning problems to stochastic optimization problems, and exploiting this structure is an active area of research [6, 11, 34]. To this end, we define

$$A = \begin{bmatrix} C \\ D \end{bmatrix} \in \mathbb{R}^{M \times N}, \quad (60)$$

³Here we assume that $A_i^T A_i \succ 0$

with the partitioning

$$C = \begin{bmatrix} C_1 & & \\ & \ddots & \\ & & C_n \end{bmatrix} \in \mathbb{R}^{m \times N}, \quad D = [D_1 \quad \dots \quad D_n] \in \mathbb{R}^{\ell \times N} \quad \text{and} \quad A_i = \begin{bmatrix} C_i \\ D_i \end{bmatrix} \in \mathbb{R}^{M \times N_i}. \quad (61)$$

Moreover, we assume that each block $C_i \in \mathbb{R}^{M_i \times N_i}$, and the linking blocks $D_i \in \mathbb{R}^{\ell \times N_i}$. We assume that $\ell \ll N$, and that there are n blocks with $m = \sum_{i=1}^n M_i$ so $M = m + \ell$, and $N = \sum_{i=1}^n N_i$.

Notice that if $D = \mathbf{0}$, where $\mathbf{0}$ is the $\ell \times N$ matrix of all zeros, then problem (55) is completely (block) separable so it can be solved easily. The linking constraints D make problem (55) nonseparable, which makes it non-trivial to solve.

The system of equations (57) must be solved at each iteration of ICD (where $B_i = A_i^T A_i = C_i^T C_i + D_i^T D_i$) because it determines the update to apply to the i th block. We solve this system *inexactly* using an *iterative method*. In particular we use the conjugate gradient method (CG) in the numerical experiments presented in this section.

It is well known that the performance of CG is improved by the use of an appropriate preconditioner. To this end, we compare ICD using CG with ICD using preconditioned CG (PCG). If $M_i \geq N_i$ and $\text{rank}(C_i) = N_i$, then the block $C_i^T C_i$ is positive definite so we propose the preconditioner (for the i th system)

$$\mathcal{P}_i := C_i^T C_i. \quad (62)$$

If $M_i < N_i$ then \mathcal{P}_i is rank deficient and is therefore singular. In such a case, we perturb (62) by adding a multiple of the identity matrix, and propose the nonsingular preconditioner

$$\hat{\mathcal{P}}_i = \mathcal{P}_i + \rho I = C_i^T C_i + \rho I, \quad (63)$$

where $\rho > 0$.

Applying the preconditioners (defined in (62) for $M_i \geq N_i$, and (63) for $M_i < N_i$) to (57), should result in the system having better spectral properties than the original, and this will lead to faster convergence of the conjugate gradient algorithm. A full theoretical justification (eigenvalue analysis) for the preconditioners is presented in Appendix A.

Remark: Notice that the preconditioners (62) and (63) are likely to be significantly more sparse than B_i , and consequently we expect that these preconditioners will be cost effective to apply in practice. To see this, notice that the blocks C_i are generally much sparser than the linking blocks D_i so that $\mathcal{P}_i = C_i^T C_i$ is much sparser than $C_i^T C_i + D_i^T D_i$.

7.1.1 Experiment parameters and results

The purpose of this experiment is to study the use of an iterative technique (CG or PCG) to determine the update used at each iteration of the inexact block coordinate descent method, and compare this approach with Exact CD. For Exact CD, the system (57) was solved by forming the Cholesky Decomposition of B_i for each i and then performing two triangular solves to find the exact update.

In the first two experiments, simulated data was used to generate A and the solution vector x_* . For each matrix A , each block C_i has approximately 20 nonzeros per column, and the density of the linking constraints D_i is approximately $0.1 \ell N_i$. The data vector b was generated from $b = Ax_*$, so

the optimal value is known in advance: $F^* = 0$. The stopping condition and tolerance ϵ for ICD are: $F(x_K) - F^* = \frac{1}{2}\|Ax_K - b\|_2^2 < \epsilon = 0.1$.

The inexactness parameters are set to $\alpha = 0$ and $\beta = 0.1$. Therefore, the update for each block is accepted when $\frac{1}{2}\|A_i T_{\delta_k}^{(i)} - r\|_2^2 \leq \beta = \delta_k^{(i)} = 0.1$ for all i, k . Moreover, each block was chosen with uniform probability $\frac{1}{n}$ in all experiments in this section.

In the first experiment the blocks C_i are tall. The incomplete Cholesky decomposition of the preconditioner \mathcal{P}_i was found using MATLAB’s ‘ichol’ function with a drop tolerance set to 0.1. The results of this experiment are shown in the Table 3 and all results are averages over 20 runs.

In the second experiment the blocks C_i are wide.⁴ The incomplete Cholesky decomposition of the perturbed preconditioner $\hat{\mathcal{P}}_i = \mathcal{P}_i + \rho I$ (with $\rho = 0.5$) was found using MATLAB’s ‘ichol’ function with a drop tolerance set to 0.1. The results are shown in the Table 4 and all results are averages over 20 runs.

We briefly explain the terminology used in the tables presented in this section. ‘Time’ represents the cpu time in seconds. Further, the term ‘block updates’ refers to the total number of block updates computed throughout the algorithm; dividing this number by n gives the number of ‘epochs’, which is (approximately) equivalent to the total number of full dimensional matrix-vector products required by the algorithm. The abbreviation ‘o.o.m.’ is the out of memory token.

Table 3: Results of Exact CD, ICD with CG and ICD with PCG on a quadratic objective with block angular structure using simulated data. For all of these problems, the blocks C_i are tall, and the preconditioner (62) is used for ICD with PCG. The size of A ranges from $10^6 \times 10^5$ to $10^7 \times 10^6$. All results are averages over 20 runs.

				Exact CD		ICD with CG			ICD with PCG		
n	M_i	N_i	ℓ	Block Updates	Time	Block Updates	CG Iterations	Time	Block Updates	PCG Iterations	Time
100	10^4	10^3	1	4,820.1	37.42	4,726.3	15,126	13.95	5,230.6	11,379	12.59
100	10^4	10^3	10	7,056.7	53.94	7,181.1	14,480	17.88	6,864.0	13,516	15.95
100	10^4	10^3	100	19,129	151.97	19,411	37,841	46.32	19,446	41,344	51.12
10	10^5	10^4	1	3129.4	2488.2	3,307.5	5,316.4	64.39	3,246.8	4,201.4	62.71
10	10^5	10^4	10	4588	3738.6	4,753.6	9,907.6	109.79	4,655.4	7,646.8	104.65
10	10^5	10^4	100	12,431	15,302	15,938	35,943	446.81	15,417	29,272	391.12
100	10^5	10^4	1	o.o.m.	o.o.m.	44,799	59,340	821.64	43,427	49,801	783.11
100	10^5	10^4	10	o.o.m.	o.o.m.	63,654	101,163	1,302.0	59,351	82,097	1,267.3
100	10^5	10^4	100	o.o.m.	o.o.m.	207,314	329276	4982.8	204070	302,308	4806.1

The results presented in Table 3 show that ICD with either CG or PCG significantly outperforms Exact CD in terms of cpu time. When the blocks are of size $M_i \times N_i = 10^4 \times 10^3$, ICD is approximately 3 times faster than Exact CD. The results are even more striking as the block size increases. Notice that ICD was able to solve problems of all sizes, whereas Exact CD ran out of memory on the problems of size $10^7 \times 10^6$. Further, we notice that PCG is faster than CG in terms of cpu time, demonstrating the benefits of preconditioning. These results strongly support the ICD method.

⁴To ensure that C_i has full rank, a multiple of the identity I_{m_i} is added to the first m_i columns of C_i .

Table 4: Results of Exact CD, ICD with CG and ICD with PCG on a quadratic objective with block angular structure using simulated data. For all of these problems, the blocks C_i are wide, and the preconditioner (63) with $\rho = 0.5$ is used for ICD with PCG. The size of A ranges from $10^5 \times 10^5$ to $10^6 \times 10^6$. All results are averages over 20 runs.

n	M_i	N_i	ℓ	Exact CD		ICD with CG			ICD with PCG		
				Block Updates	Time	Block Updates	CG Iterations	Time	Block Updates	PCG Iterations	Time
10	9,999	10^4	1	34.2	190.62	821.2	1957	12.55	471.4	1597	9.29
10	9,990	10^4	10	31.3	191.96	1,500.8	4,793.3	45.81	867.7	3612	24.55
10	9,000	10^4	10^3	25.5	287.79	703.6	4,052.8	58.31	439.0	4,309.8	46.74
10	7,500	10^4	2,500	39.7	336.69	532.0	3183	76.77	386.5	4592	70.09
100	9,999	10^4	1	o.o.m.	o.o.m.	13077	27321	185.31	8280	25715	143.63
100	9,900	10^4	10^2	o.o.m.	o.o.m.	12,979	50,685	397.47	6,159	47,034	245.89
100	9,000	10^4	10^3	o.o.m.	o.o.m.	6974	39535	453.18	4797	52,665	496.35
100	7,500	10^4	2,500	o.o.m.	o.o.m.	4936	28986	542.69	4246	57001	740.75

The results presented in Table 4 show that ICD outperforms Exact CD. ICD is able to solve all problem instances, whereas Exact CD gives the out of memory token on the large problems. We see that when ℓ is small, ICD with PCG has an advantage over ICD with CG. However, when ℓ is large, the the preconditioner $\hat{\mathcal{P}}_i$ is not as good an approximation to $A_i^T A_i$ and so ICD with CG is preferable.

Remark: Notice that in several of the numerical experiments, Exact CD returned the out of memory token. Exact CD requires the matrices $B_i = C_i^T C_i + D_i^T D_i$ for all i to be formed explicitly, and the Cholesky factors to be found and stored. Even if A_i is sparse, B_i need not be, and the Cholesky factor could be dense, making it very expensive to work with. Moreover, this problem does not arise for ICD with CG (and arises to a much lesser extent for PCG) because B_i is never explicitly formed. Instead, only sparse matrix vector products: $B_i x \equiv C_i^T (C_i x) + D_i^T (D_i x)$ are required. This is why ICD performs extremely well, even when the blocks are very large.

7.1.2 Real-world data

In the third experiment we test ICD on a quadratic objective with block angular structure, where the matrices arise from real-world applications. In particular, we have taken several matrices from the Florida Sparse Matrix Collection [7] that have block angular structure. The matrices used are given in Table 5. Note that in each case we have taken the transpose of the original matrix to ensure that the matrix is tall. Further, in each case the upper block (recall (60)) is diagonal, so we have scaled each of the matrices so that $C = I$. Note that in this case $\mathcal{P}_i = I$ so there is no need for preconditioning. We compare Exact CD with ICD using CG. All the stopping conditions and algorithm parameters are the same as those given in Section 7.1.1.

The results of the numerical experiments on these matrices are shown in Table 6. (To determine n (the number of blocks) and N_i the size of the blocks, we have simply taken the prime factorization of N .) ICD with CG performs extremely well on these test problems. In most cases ICD with CG needs more iterations than Exact CD to converge, yet ICD requires only a fraction of the cpu time needed by Exact CD.

Table 5: Block angular matrices from the Florida Sparse Matrix Collection [7]. (`cep1` is from the Meszaros Group while all others are from the Mittelman Group.) Note that the dimensions given in the table are for the *transpose* of the original test matrix.

	M	N	ℓ
<code>cep1</code>	4769	1521	3,248
<code>neos</code>	515,905	479,119	36,786
<code>neos1</code>	133,473	131,528	1,945
<code>neos2</code>	134,128	132,568	1,560
<code>neos3</code>	518,832	512,209	6,623

Table 6: Results showing the performance of Exact CD and ICD with CG applied to a quadratic function with the block angular matrices described in Table 5. For the small problem `cep1`, Exact CD is the best algorithm. For all other matrices, ICD with CG is significantly better than Exact CD in terms of the cpu time.

			Exact CD		ICD with CG		
	n	N_i	Block Updates	Time	Block Updates	CG Iterations	Time
<code>cep1</code>	9	169	446	0.18	448	828	0.61
	3	507	376	0.29	342	678	0.52
<code>neos</code>	283	1,693	622,659	3,258.8	869,924	3,919,172	2,734.65
<code>neos1</code>	41	3,208	148,228	8,759.6	143,156	592,070	773.70
	8	16,441	25,503	52,113	25,853	116,468	446.26
<code>neos2</code>	73	1,816	329,749	4,669.1	439,296	1,825,835	997.04
	8	16,571	82,784	11,518	55,414	255,129	972.27
<code>neos3</code>	107	4,787	81,956	9,032.1	82,629	433,354	700.82

7.2 A numerical experiment for a nonsmooth objective

In this numerical experiment we consider the l_1 -regularized least squares problem

$$\min_{x \in \mathbb{R}^N} \frac{1}{2} \|Ax - b\|_2^2 + \lambda \|x\|_1, \quad (64)$$

where $A \in \mathbb{R}^{M \times N}$, $b \in \mathbb{R}^M$ and $\lambda > 0$. Problem (64) fits into the framework (1) with $f = \frac{1}{2} \|Ax - b\|_2^2$ and $\Psi = \lambda \|x\|_1 = \lambda \sum_{i=1}^n \|x^{(i)}\|_1$. For this experiment we set $B_i = A_i^T A_i$ and $l_i = 1$ for $i = 1, \dots, n$. (It can be shown that this choice of B_i and l_i satisfy the overapproximation (7).) Further, for this experiment we use uniform probabilities, $p_i = \frac{1}{n}$ for all i , and we set $\alpha = 0$ and $\beta > 0$. The algorithm stopping condition is $F(x_k) - F^* < \epsilon = 10^{-4}$, (the data was constructed so that F^* is known), and the regularization parameter was set to $\lambda = 0.01$.

The exact update for the i th block is computed via

$$V_i(x_k, t) = \langle A_i^T r_k, t \rangle + \frac{1}{2} t^T A_i^T A_i t + \lambda \|x_k^{(i)} + t\|_1 = \frac{1}{2} \|A_i t + r_k\|_2^2 + \lambda \|x_k^{(i)} + t\|_1, \quad (65)$$

where $r_k := Ax_k - b$ and $\nabla_i f(x) = A_i^T r_k$. Notice that (65) *does not have a closed form solution*, meaning that only an *inexact update* can be used in this case. Recall that the inexact update must

satisfy (18), and for (65), we do not know the optimal value $V_i(x_k, T_0^{(i)})$. In this case, to ensure that (18) is satisfied, we simply find the inexact update $T_{\delta_k}^{(i)}$ using an algorithm that terminates on the duality gap. That is, we accept $T_{\delta_k}^{(i)}$ using a stopping condition of the same form as that given by (22).

In the numerical experiments presented in this section, we use the BCGP algorithm [4] to solve for the update at each iteration of ICD. This is a gradient based method that solves problems of the form (65), and terminates on the duality gap.

We conduct two numerical experiments. In the first experiment A is of size $0.5N \times N$ where $N = 10^5$. In this case (64) is convex (but not strongly convex.) This means that the complexity result of Theorem 9 applies. In the second experiment A is of size $2N \times N$ where $N = 10^5$. In this case (64) is strongly convex, and the complexity result of Theorem 11 apply.

The purpose of these experiments is to investigate the effect of different levels of inexactness (different values of β) on the algorithm runtime. In particular we used three different values: $\beta \in \{10^{-4}, 10^{-6}, 10^{-8}\}$. To make this a fair test, for each problem instance, the block ordering was fixed in advance. (i.e., before the algorithm begins we form and store a vector whose k th element is a index between 1 and n that has been chosen with uniform probability, corresponding to the block to be updated at iteration k of ICD.) Then, ICD was run three times using this block ordering, once for each value of $\beta \in \{10^{-4}, 10^{-6}, 10^{-8}\}$. In all cases we use $\delta_k^{(i)} = \beta$ for all i and k .

Figures 1 and 2 show the results of experiments 1 ($M < N$) and 2 ($M > N$) respectively. The experiments were performed many times on simulated data and the plots shown are a particular instance that is *representative of the typical behaviour* observed using ICD on this problem description. We see that when the same block ordering is used, all algorithms essentially require the same number of iterations until termination regardless of the parameter β . (This is to be expected.) Moreover, it is clear that using a smaller value of β , corresponding to more ‘inexactness’ in the computed update, leads to a reduction in the algorithm running time, without affecting the ultimate convergence of ICD. This shows that using an inexact update (an iterative method) has significant practical advantages.

References

- [1] G. C. Bento, J. X. Da Cruz Neto, P. R. Oliveira, and A. Soubeyran. The self regulation problem as and inexact steepest descent method for multicriteria optimization. Technical report, July 2012. arXiv:1207.0775v1 [math.OC].
- [2] M. Benzi, G. H. Golub, and J. Liesen. Numerical solution of saddle point problems. *Acta Numerica*, pages 1–137, 2005.
- [3] S. Bonettini. Inexact block coordinate descent methods with application to non-negative matrix factorization. *IMA Journal of Numerical Analysis*, 31:1431–1452, 2011.
- [4] R. Broughton, I. Coope, P. Renaud, and R. Tappenden. A box-constrained gradient projection algorithm for compressed sensing. *Signal Processing*, 91(8):1985–1992, 2011.
- [5] E. Candès and B. Recht. Exact matrix completion via convex optimization. *Commun. ACM*, 55(6):111–119, 2012.




Figure 1: Plots of the objective function value vs the number of iterations, and the cputime vs the objective function values on the l_1 -regularised quadratic loss problem (64). For these plots, the matrix A is $0.5N \times N$, where $N = 10^5$. In the plots in the first row, $n = 10$ and $N_i = 10^4$ for all $i = 1, \dots, n$. In the second row, $n = 100$ and $N_i = 10^3$ for all $i = 1, \dots, n$.

- [6] J. Castro and J. Cuesta. Quadratic regularizations in an interior-point method for primal block-angular problems. *Math. Program., Ser. A*, 130:415–445, 2011.
- [7] T. A. Davis and Y. Hu. The University of Florida sparse matrix collection. *ACM Transactions on Mathematical Software*, 38(1):1:1 – 1:25, 2011.
- [8] D. Donoho. Compressed sensing. *IEEE Trans. on Information Theory*, 52(4):1289 – 1306, April 2006.
- [9] G. H. Golub and C. F. van Loan. *Matrix Computations*. Johns Hopkins University Press, third edition, 1986.




Figure 2: Plots of the objective function value vs the number of iterations, and the cputime vs the objective function values on the l_1 -regularised quadratic loss problem (64). For these plots, the matrix A is $2N \times N$, where $N = 10^5$. In the plots in the first row, $n = 10$ and $N_i = 10^4$ for all $i = 1, \dots, n$. In the second row, $n = 100$ and $N_i = 10^3$ for all $i = 1, \dots, n$.

- [10] G. H. Golub and Q. Ye. Inexact preconditioned conjugate gradient method with inner-outer iteration. *SIAM J. Sci. Comput.*, 21(4):1305–1320, 1999.
- [11] J. Gondzio and R. Sarkissian. Parallel interior-point solver for structured linear programs. *Mathematical Programming*, 96(3):561–584, 2003.
- [12] S. Gratton, A. Sartenaer, and J. Tshimanga. On a class of limited memory preconditioners for large scale linear systems with multiple right-hand sides. *SIAM J. Optim.*, 21(3):912–935, 2011.

- [13] M. R. Hestenes and E. Steifel. Methods of conjugate gradients for solving linear systems. *J. Res. Natl. Bur. Stand.*, 49:409–436, 1952.
- [14] R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, 1985.
- [15] X. Hua and N. Yamashita. An inexact coordinate descent method for the weighted l_1 -regularized convex optimization problem. Technical report, School of Mathematics and Physics, Kyoto University, Kyoto 606-8501, Japan, September 2012.
- [16] D. Leventhal and A. S. Lewis. Randomized methods for linear constraints: Convergence rates and conditioning. *Mathematics of Operations Research*, 35(3):641–654, August 2010.
- [17] P. Machart, S. Anthoine, and L. Baldassarre. Optimal computational trade-off of inexact proximal methods. Technical report, INRIA, 00704398, October 2012. Version 3.
- [18] I. Necoara and V. Nedelcu. Inexact dual gradient methods with guaranteed primal feasibility: application to distributed MPC. Technical report, Politehnica University of Bucharest, 060042 Bucharest, Romania, September 2012.
- [19] I. Necoara, Y. Nesterov, and F. Glineur. Efficiency of randomized coordinate descent methods on optimization problems with linearly coupled constraints. Technical report, June 2012. pp.1–21.
- [20] I. Necoara and A. Patrascu. A random coordinate descent algorithm for optimization problems with composite objective function and linear coupled constraints. Technical report, University Politehnica Bucharest, Spl. Independentei 313, Romania, June 2012. pp.1–20.
- [21] D. Needell. Randomized Kaczmarz solver for noisy linear systems. *BIT Numerical Mathematics*, 50:395–403, 2010.
- [22] D. Needell and J. Tropp. Paved with good intentions: Analysis of a randomized Kaczmarz method. Technical report, August 2012. ArXiv:1208.3805v1 [math.NA].
- [23] Y. Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Applied Optimization. Kluwer Academic Publishers, 2004.
- [24] Y. Nesterov. Gradient methods form minimizing composite objective function. Technical report, Core discussion paper #2007/76, Université catholique de Louvain, Center for Operations Research and Econometrics (CORE), September 2007.
- [25] Y. Nesterov. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM J. Optim.*, 22(2):341–362, 2012.
- [26] Z. Qin, K. Scheinberg, and D. Goldfarb. Efficient block-coordinate descent algorithms for the group lasso. Technical report, Department of Industrial Engineering and Operations Research, Columbia University, 2010.
- [27] B. Recht and C. Ré. Parallel stochastic gradient algorithms for large-scale matrix completion. Technical report, Computer Sciences Department, University of Wisconsin-Madison, 1210 W Dayton St, Madison, WI 53706, April 2011.

- [28] P. Richtárik and M. Takáč. Efficient serial and parallel coordinate descent methods for huge-scale truss topology design. In Diethard Klatte, Hans-Jakob Lüthi, and Karl Schmedders, editors, *Operations Research Proceedings 2011*, Operations Research Proceedings, pages 27–32. Springer Berlin Heidelberg, 2012.
- [29] P. Richtárik and M. Takáč. Iteration complexity of randomized block-coordinate descent methods for minimizing a composite function. *Mathematical Programming*, DOI: 10.1007/s10107-012-0614-z, pages 1–38, 2012.
- [30] P. Richtárik and M. Takáč. Parallel coordinate descent methods for big data optimization. Technical report, November 2012. arXiv:1212.0873.
- [31] P. Richtárik and M. Takáč. Efficiency of randomized coordinate descent methods on minimization problems with a composite objective function. In *4th Workshop on Signal Processing with Adaptive Sparse Structured Representations*, June 2011.
- [32] A. Saha and A. Tewari. On the finite time convergence of cyclic coordinate descent methods. Technical report, May 2010. arXiv:1005.2146v1 [cs.LG].
- [33] M. Schmidt, N. Le Roux, and F. Bach. Convergence rates of inexact proximal-gradient methods for convex optimization. Technical report, INRIA, 00618152, December 2011.
- [34] G. L. Schultz and R. R. Meyer. An interior point method for block angular optimization. *SIAM J. Optim.*, 1(4):583–602, 1991.
- [35] S. Shalev-Schwartz and T. Zhang. Stochastic dual coordinate ascent methods for regularized loss minimization. *Journal of Machine Learning Research*, 14:567–599, 2013.
- [36] N. Simon and R. Tibshirani. Standardization and the group lasso penalty. Technical report, Stanford University, March 2011.
- [37] T. Strohmer and R. Vershynin. A randomized Kaczmarz algorithm with exponential convergence. *Journal of Fourier Analysis and Applications*, 15:262–278, 2009.
- [38] M. Takáč, A. Bijral, P. Richtárik, and N. Srebro. Mini-batch primal and dual methods for SVMs. In *30th International Conference on Machine Learning*, 2013.
- [39] Q. Tao, K. Kong, D. Chu, and G. Wu. Stochastic coordinate descent methods for regularized smooth and nonsmooth losses. In P. A. Flach, T. De Bie, and N. Cristianini, editors, *Machine Learning and Knowledge Discovery in Databases*, volume 7523 of *Lecture Notes in Computer Science*, pages 537–552. Springer, 2012.
- [40] P. Tseng. Convergence of block coordinate descent method for nondifferentiable minimization. *Journal of Optimization Theory and Applications*, 109:475–494, June 2001.
- [41] P. Tseng and S. Yun. A coordinate gradient descent method for nonsmooth separable minimization. *Math. Program., Ser. B*, 117:387–423, 2009.
- [42] S. J. Wright. Accelerated block-coordinate relaxation for regularized optimization. *SIAM J. Optim.*, 22(1):159–186, 2012.

- [43] S. J. Wright, R. D. Nowak, and M. A. T. Figueiredo. Sparse reconstruction by separable approximation. *Trans. Sig. Proc.*, 57:2479–2493, July 2009.
- [44] F. Zhang. *Matrix Theory: Basic Results and Techniques*. Springer, 1999.

A Eigenvalues of the preconditioned matrix

The purpose of this section is to provide a full theoretical justification for the choice of preconditioner presented in Section 7.1.

The convergence speed of many iterative methods, such as CG, depends on the spectrum of the system matrix. The purpose of a preconditioner is to shift the spectrum of the preconditioned matrix so that the eigenvalues of the resulting system are clustered around one, with few outliers. In this section we study the eigenvalues of the preconditioned matrix under the problem setup described in Sections 6.1.1 and 7.1.

Applying (62) and (63) to $B_i (= C_i^T C_i + D_i^T D_i)$ gives

$$\mathcal{P}_i^{-1} B_i = I + \mathcal{P}_i^{-1} D_i^T D_i, \text{ and } \hat{\mathcal{P}}_i^{-1} B_i = \hat{\mathcal{P}}_i^{-1} \mathcal{P}_i + \hat{\mathcal{P}}_i^{-1} D_i^T D_i \quad (66)$$

To investigate the quality of a preconditioner, we study the eigenvalues of the preconditioned matrices $\mathcal{P}_i^{-1} B_i$ and $\hat{\mathcal{P}}_i^{-1} B_i$ defined in (66).

The nonzero eigenvalues of the $N_i \times N_i$ matrix $\mathcal{P}_i^{-1} D_i^T D_i$ are the same as the nonzero eigenvalues of $D_i \mathcal{P}_i^{-1} D_i^T$ (See Theorem 2.8 in [44]). We prefer to work with $D_i \mathcal{P}_i^{-1} D_i^T$ because it is symmetric and positive semidefinite, so it has real, nonnegative eigenvalues. (Furthermore, if D_i has full (row) rank, then $D_i \mathcal{P}_i^{-1} D_i^T$ is positive definite.)

We can say more about the eigenvalues of the preconditioned matrix by considering the blocks of A and investigating the relationship between the matrices C and D , defined in (61). Recall that C contains blocks $C_i \in \mathbb{R}^{M_i \times N_i}$. The remainder of this section is broken into two parts. The first part considers the case when $M_i \geq N_i$ while the second part considers the case when $M_i < N_i$. In each case C_i is assumed to have full rank.⁵

A.1 Tall blocks

In this section it is assumed that $C_i \in \mathbb{R}^{M_i \times N_i}$ where $M_i \geq N_i$ and that $D_i \in \mathbb{R}^{\ell \times N_i}$ with $1 \leq \ell < N_i$. Furthermore, we assume that C_i has full column rank (the rows of C_i contain a basis for \mathbb{R}^{N_i}) so each row of D_i is a linear combination of the rows of C_i . i.e., for $Z_i \in \mathbb{R}^{\ell \times M_i}$ we can write

$$D_i = Z_i C_i. \quad (67)$$

We have the following result.

Theorem 14. *Let $r_i = \text{rank}(D_i)$ and $r_i \leq N_i$. Then $\mathcal{P}_i^{-1} B_i = I + \mathcal{P}_i^{-1} D_i^T D_i$ (66) has*

(i) r_i eigenvalues that are strictly greater than one.

(ii) $N_i - r_i$ eigenvalues equal to one.

⁵Note that the eigenvalues of $\mathcal{P}_i^{-1} D_i^T D_i$ can be determined exactly by solving the generalized eigenvalue problem $D_i^T D_i v = \lambda \mathcal{P}_i v$.

We have the following bound on the diagonal elements of $\mathcal{P}_i^{-1}D_i^T D_i$.

Lemma 15. *Let $\mathcal{P}_i \in \mathbb{R}^{N_i \times N_i}$ and $D_i \in \mathbb{R}^{\ell \times N_i}$ be the matrices defined in (62) and (67) respectively and let z_j^T denote the j th row of Z_i . Let $C_i = Y_i R_i$ denote the thin QR factorization ([9]) of C_i , so $Y_i \in \mathbb{R}^{M_i \times N_i}$ has orthonormal columns and $R_i \in \mathbb{R}^{N_i \times N_i}$ is upper triangular. Then*

$$\text{trace}(D_i \mathcal{P}_i^{-1} D_i^T) = \sum_{j=1}^{\ell} \|z_j^T Y_i\|_2^2 \leq \|Z_i\|_F^2. \quad (68)$$

Proof. The trace is simply the sum of the diagonal entries of a (square) matrix, so:

$$D_i \mathcal{P}_i^{-1} D_i = Z_i C_i (C_i^T C_i)^{-1} C_i^T Z_i^T = Z_i Y_i R_i (R_i^T Y_i^T Y_i R_i)^{-1} R_i^T Y_i^T Z_i^T = (Z_i Y_i) (Z_i Y_i)^T.$$

Now $(D_i \mathcal{P}_i^{-1} D_i^T)_{jj} = \|Y_i^T z_j\|_2^2$ and so $\|Z_i\|_F^2 = \sum_{j=1}^{\ell} \|z_j\|_2^2$. Because $Y_i Y_i^T$ is a projection matrix, $\|Y_i^T z_j\|_2^2 = \|Y_i Y_i^T z_j\|_2^2 \leq \|z_j\|_2^2$, and the result follows. \square

Remark: When C_i is square and has full rank, Y_i is an orthogonal matrix, and subsequently $\text{trace}(D_i \mathcal{P}_i^{-1} D_i^T) = \sum_{j=1}^{\ell} \|z_j\|_2^2 = \|Z_i\|_F^2$.

We now present the main result of this section.

Theorem 16. *Suppose that $A \in \mathbb{R}^{M \times N}$ has primal block angular structure, with rectangular blocks $C_i \in \mathbb{R}^{M_i \times N_i}$ ($M_i \geq N_i$) of full rank ($N_i = \text{rank}(C_i)$) along the diagonal. Suppose that $B_i \equiv A_i^T A_i$, D_i and \mathcal{P}_i are defined in (61), (67) and (62) respectively, and let $r_i = \text{rank}(D_i)$. Then $\mathcal{P}_i^{-1} B_i$ has*

(i) $N_i - r_i$ eigenvalues equal to one,

(ii) r_i eigenvalues that are strictly greater than 1, and sum to $r_i + \sum_{j=1}^{\ell} \|Y_i^T z_j\|_2^2$.

Proof. The proof follows from Theorem 14 and Lemma 68. \square

A.2 Wide blocks

In this section we assume that $C_i \in \mathbb{R}^{M_i \times N_i}$ where $M_i < N_i$ with full row rank $M_i = \text{rank}(C_i)$, and that $D_i \in \mathbb{R}^{\ell \times N_i}$ where $\ell \geq N_i - M_i$. Then the rows of C_i form a basis for a subspace $\mathcal{W} := \text{span}\{c_1^{(i)}, \dots, c_{M_i}^{(i)}\} \subset \mathbb{R}^{N_i}$, where $(c_j^{(i)})^T$ is the j th row of C_i .

Furthermore, let W be an $\ell \times N_i$ matrix whose rows $w_j^T \in \mathcal{W}$, for $j = 1, \dots, \ell$, and let W^\perp be an $\ell \times N_i$ matrix whose rows $(w_j^\perp)^T \in \mathcal{W}^\perp$, for $j = 1, \dots, \ell$, where \mathcal{W}^\perp denotes the orthogonal complement of \mathcal{W} . Then one can write

$$D_i = W + W^\perp. \quad (69)$$

For ICD, $B_i = A_i^T A_i$ must have full rank because it defines a norm (see Section 2.1). However, when C_i is wide, \mathcal{P}_i defined in (62) is rank deficient so we use the preconditioner $\hat{\mathcal{P}}_i$ defined in (63). The preconditioned matrix is defined in (66). We study the eigenvalues of $\hat{\mathcal{P}}_i^{-1} \mathcal{P}_i$ and $\hat{\mathcal{P}}_i^{-1} D_i^T D_i$ separately, before stating the main result of this section, which describes the eigenvalues of $\hat{\mathcal{P}}_i^{-1} B_i$.

Theorem 17. *Let C_i be a real $M_i \times N_i$ matrix with $M_i < N_i$ and full row rank $M_i = \text{rank}(C_i)$. Let \mathcal{P}_i and $\hat{\mathcal{P}}_i$ be defined in (62) and (63) respectively, and let $M_i = \text{rank}(\mathcal{P}_i)$. Then $\hat{\mathcal{P}}_i^{-1} \mathcal{P}_i$ has $N_i - M_i$ zero eigenvalues and M_i positive eigenvalues that tend to 1 as $\rho \rightarrow 0$.*

Proof. The preconditioner $\hat{\mathcal{P}}_i$ satisfies $\text{rank}(\hat{\mathcal{P}}_i^{-1}\mathcal{P}_i) = \text{rank}(\mathcal{P}_i) = M_i$, so $\hat{\mathcal{P}}_i^{-1}\mathcal{P}_i$ has M_i nonzero eigenvalues and $N_i - M_i$ zero eigenvalues. Furthermore, the M_i nonzero eigenvalues are positive. (Notice that $\hat{\mathcal{P}}_i^{-\frac{1}{2}}\mathcal{P}_i\hat{\mathcal{P}}_i^{-\frac{1}{2}}$ is positive semidefinite.)

Let $\lambda_1, \dots, \lambda_{M_i}$ denote the M_i nonzero eigenvalues of \mathcal{P}_i . The eigenvalue decomposition of \mathcal{P}_i is $\mathcal{P}_i = V\Lambda V^T$, where $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_{M_i}, 0, \dots, 0)$. Moreover, the eigenvalue decomposition for $\hat{\mathcal{P}}_i$ is $\hat{\mathcal{P}}_i = \mathcal{P}_i + \rho I = V\hat{\Lambda}V^T$ where $\hat{\Lambda} = \Lambda + \rho I$.

Finally, $\hat{\mathcal{P}}_i^{-1}\mathcal{P}_i = V\hat{\Lambda}^{-1}V^T V\Lambda V^T = V\hat{\Lambda}^{-1}\Lambda V^T$, where $\hat{\Lambda}^{-1}\Lambda$ is a diagonal matrix with diagonal entries $(\lambda_1/(\lambda_1 + \rho), \dots, \lambda_{M_i}/(\lambda_{M_i} + \rho), 0, \dots, 0)$ and as $\rho \rightarrow 0$, $\frac{\lambda_j}{\lambda_j + \rho} \rightarrow 1$ for $j = 1, \dots, M_i$. \square

Theorem 18. *Let D_i and $\hat{\mathcal{P}}_i$ be as defined in (69) and (63) respectively, and let $\rho > 0$. Then*

$$\text{tr}(D_i\hat{\mathcal{P}}_i^{-1}D_i^T) = \sum_{j=1}^{\ell} \left(\|\hat{\Lambda}_1^{-\frac{1}{2}}V_1^T(w_j + w_j^\perp)\|_2^2 + \frac{1}{\rho}\|V_2^T(w_j^\perp)\|_2^2 \right). \quad (70)$$

Proof. Recall that $\hat{\mathcal{P}}_i = V\hat{\Lambda}V^T$ where $\hat{\Lambda} = \Lambda + \rho I$ and let $V = [V_1 \ V_2]$ be a partitioning of V , where $V_1 \in R^{N_i \times M_i}$, and $V_2 \in R^{N_i \times (N_i - M_i)}$. The columns of V_2 form a basis for the null space of C_i , so $V_2^T w = 0$. The results follows by expanding $(D_i\hat{\mathcal{P}}_i^{-1}D_i^T)_{jj} = (w_j + w_j^\perp)^T \hat{\mathcal{P}}_i^{-1}(w_j + w_j^\perp)$. \square

Theorems 17 and 18 demonstrate the importance of the parameter ρ . A small value of ρ will lead to a good clustering of the eigenvalues around one, but if ρ is too small then $(w_j^\perp)^T \hat{\mathcal{P}}_i^{-1}(w_j^\perp)$ will become arbitrarily large. Hence, there is a trade-off here.

Now we state the main result of this section, which gives bounds on the eigenvalues of $\hat{\mathcal{P}}_i^{-1}D_i^T D_i$.

Theorem 19. *Let C_i be an $M_i \times N_i$ matrix with $M_i < N_i$ and $M_i = \text{rank}(C_i)$ and let D_i be an $\ell \times N_i$ matrix with $r_i = \text{rank}(D_i)$. Let $\hat{\mathcal{P}}_i$ be the preconditioner defined in (63) and let A_i be defined in (61) with $N_i \geq s_i = \text{rank}(A_i)$ and $B_i = A_i^T A_i$. Then $\mathcal{M}_i = \hat{\mathcal{P}}_i^{-1}B_i$ has*

(i) $N_i - s_i$ eigenvalues equal to zero.

(ii) $s_i - r_i$ eigenvalues in the interval $(0, 1)$

(iii) r_i eigenvalues in the interval $\left(1, 1 + \sum_{j=1}^{\ell} \left(\|\hat{\Lambda}_1^{-\frac{1}{2}}V_1^T(w_j + w_j^\perp)\|_2^2 + \frac{1}{\rho}\|V_2^T(w_j^\perp)\|_2^2 \right)\right)$

Proof. Part (i) holds because B_i is $N_i \times N_i$ with $\text{rank}(B_i) = \text{rank}(A_i) = s_i$. Part (ii) follows from Theorem 17 and Theorem 17. For part (iii), notice that $\lambda_{\max}(D_i\hat{\mathcal{P}}_i^{-1}D_i^T) \leq \text{trace}(D_i\hat{\mathcal{P}}_i^{-1}D_i^T)$. Using (18) and [14, Theorem 4.3.1], gives the result. \square

Remark: For ICD we require $\text{rank}(A_i) = N_i$, because this ensures that B_i is a positive definite matrix. Notice that in this case, Theorem 19 shows that all eigenvalues of $\mathcal{M}_i = \hat{\mathcal{P}}_i^{-1}B_i$ are strictly greater than zero (i.e., $N_i = s_i$).