# A Flexible Coordinate Descent Method 

Kimon Fountoulakis • Rachael Tappenden

the date of receipt and acceptance should be inserted later


#### Abstract

We present a novel randomized block coordinate descent method for the minimization of a convex composite objective function. The method uses (approximate) partial second-order (curvature) information, so that the algorithm performance is more robust when applied to highly nonseparable or ill conditioned problems. We call the method Flexible Coordinate Descent (FCD). At each iteration of FCD, a block of coordinates is sampled randomly, a quadratic model is formed about that block and the model is minimized approximately/inexactly to determine the search direction. An inexpensive line search is then employed to ensure a monotonic decrease in the objective function and acceptance of large step sizes. We present several high probability iteration complexity results to show that convergence of FCD is guaranteed theoretically. Finally, we present numerical results on large-scale problems to demonstrate the practical performance of the method.


Keywords. large scale optimization; second-order methods; curvature information; block coordinate descent; nonsmooth problems; iteration complexity; randomized.

AMS Classification. 49M15; 49M37; 65K05; 90C06; 90C25; 90C53.

## 1 Introduction

In this work we are interested in solving the following convex composite optimization problem

$$
\begin{equation*}
\min _{x \in \mathbf{R}^{N}} F(x):=f(x)+\Psi(x) \tag{1}
\end{equation*}
$$

where $f(x)$ is a smooth convex function and $\Psi(x)$ is a (possibly) nonsmooth, (coordinate) separable, real valued convex function. Problems of the form (1) arise in many important scientific fields, and applications include machine learning [43], regression [38] and compressed sensing [5, 4, 11]. Often the term $f(x)$ is a data fidelity term, and the term $\Psi(x)$ represents some kind of regularization.

Frequently, problems of the form (1) are large-scale, i.e., the size of $N$ is of the order of a million or a billion. Large-scale problems impose restrictions on the types of methods that can be employed for the solution of (11). In particular, the methods should have low per iteration computational cost, otherwise completing even a single iteration of the method might require unreasonable time. The methods must also rely only on simple operations such as matrix vector products, and ideally, they should offer fast progress towards optimality.

First order methods, and in particular randomized coordinate descent methods, have found great success in this area because they take advantage of the underlying problem structure (separability and

[^0]Address(es) of author(s) should be given
block structure), and satisfy the requirements of low computational cost and low storage requirements. For example, in [30] the authors show that their randomized coordinate descent method was able to solve sparse problems with millions of variables in a reasonable amount of time.

Unfortunately, randomized coordinate descent methods have two significant drawbacks. Firstly, due to their coordinate nature, coordinate descent methods are usually efficient only on problems with a high degree of partial separability ${ }^{\dagger}$ and performance suffers when there is a high dependency between variables. Secondly, as first-order methods, coordinate descent methods do not usually capture essential curvature information and have been shown to struggle on complicated sparse problems [15, 16 .

The purpose of this work is to overcome these drawbacks by equipping a randomized block coordinate descent method with approximate partial second-order information. In particular, at every iteration of FCD a search direction is obtained by solving a local quadratic model approximately. The quadratic model is defined by a matrix representing approximate second order information, so that essential curvature information is incorporated. The model need only be minimized approximately to give an inexact search direction, so that the process is efficient in practice. (The termination condition for the quadratic subproblem is inspired by [3], and we discuss this in more detail later.) A line search is then employed along this inexact direction, in order to guarantee a monotonic decrease in the objective function and large step-sizes.

FCD is computationally efficient. At each iteration of FCD, a block of coordinates is randomly selected, which is inexpensive. The method allows the use of inexact search directions, i.e., the quadratic model is minimized approximately. Intuitively, it is expected that this will lead to a reduction in algorithm running time, compared with minimizing the model exactly at every iteration. Also, the line search step depends on a subset of the coordinates only (corresponding to the randomly selected block of coordinates), so it is much cheaper compared with working in the full dimensional space. We note that the per iteration computational cost of the FCD method may be slightly higher than other randomized coordinate descent methods. This is due to the incorporation of the matrix representing partial second order information - the matrix gives rise to a quadratic model that is (in general) nonseparable, so the associated subproblem may not have a closed form solution. However, in the numerical results presented later in this work, we show that the method is more robust (and efficient) in practice, and often FCD will require fewer iterations to reach optimality, compared with other state-of-the-art methods. i.e., we show that FCD is able to solve difficult problems, on which other coordinate descent methods may struggle.

The FCD method is supported by theoretical convergence guarantees in the form of high probability iteration complexity results. These iteration complexity results provide an explicit expression for the number of iterations $k$, which guarantees that, for any given error tolerance $\epsilon>0$, and confidence level $\rho \in(0,1)$, the probability that FCD achieves an $\epsilon$-optimal solution, exceeds $1-\rho$, i.e., $\mathbf{P}\left(F\left(x_{k}\right)-F^{*} \leq\right.$ $\epsilon) \geq 1-\rho$.

### 1.1 Literature review

Coordinate descent methods are some of the oldest and simplest iterative methods, and they are often better known in the literature under various names such as Jacobi methods, or Gauss-Seidel methods, among others. It has been observed that these methods may suffer from poor practical performance, particularly on ill-conditioned problems, and until recently, often higher order methods have been favoured by the optimization community. However, as we enter the era of big data, coordinate descent methods are coming back into favour, because of their ability to provide approximate solutions to real world instances of very large/huge scale problems in a reasonable amount of time.

Currently, state-of-the-art randomized coordinate descent methods include that of Richtàrik and Takàč [30, where the method can be applied to unconstrained convex composite optimization problems of the form (1). The algorithm is supported by theoretical convergence guarantees in the form of high probability iteration complexity results, and [30 also reports very impressive practical performance on highly separable large scale problems. Their work has also been extended to the parallel case [31, to include acceleration techniques [14, and to include the use of inexact updates [36].

Other important works on coordinate descent methods include that of Tseng and Yun in [39, 40, 41, Nesterov [24] for huge-scale problems, work in [21] and [37] that improves the complexity analysis of [30]

[^1]and 31 respectively, coordinate descent methods for group lasso problems 27, 35] or for problems with general regularizers [34,42], a coordinate descent method that uses a sequential active set strategy [32], and coordinate descent for constrained optimization problems [22].

Unfortunately, on ill-conditioned problems, or problems that are highly nonseparable, first order methods can display very poor practical performance, and this has prompted the study of methods that employ second order information. To this end, recently there has been a flurry of research on Newton-type methods for problems of the general form (1), or in the special case where $\Psi(x)=\|x\|_{1}$. For example, Karimi and Vavasis [18] have developed a proximal quasi-Newton method for $l_{1}$-regularized least squares problems, Lee, Sun and Saunders [19,20 have proposed a family of Newton-type methods for solving problems of the form (11) and Scheinberg and Tang [33] present iteration complexity results for a proximal Newton-type method. Moreover, the authors in [3] extended standard inexact Newton-type methods to the case of minimization of a composite objective involving a smooth convex term plus an $l_{1}$-regularizer term.

Facchinei and coauthors in [8, 12, 13] have also made a significant contribution to the literature on randomized coordinate descent methods. Two new algorithms are introduced in those works, namely FLEXA [12,13] and HyFLEXA [8. Both methods can be applied to problems of the form (1] where $f$ is allowed to be nonconvex, and for HyFLEXA, $\Psi$ is allowed to be nonseparable. We stress that nonconvexity of $f$, and nonseparability of $\Psi$ are outside the scope of this work. Nevertheless, these methods are pertinent to the current work, so we discuss them now.

For FLEXA and HyFLEXA, (as is standard in the current literature), the block structure of the problem is fixed throughout the algorithm: the data vector $x \in \mathbf{R}^{N}$ is partitioned into $n(\leq N)$ blocks, $x=\left(x^{1}, \ldots, x^{n}\right) \in \mathbf{R}^{N}$. At each iteration, one block, all blocks, or some number in between, are selected in such a way that some specified error bound is satisfied. Next, a local convex model is formed and that model is approximately minimized to find a search direction. Finally, a convex combination of the current point, and the search direction is taken, to give a new point, and the process is repeated. The model they propose is block separable, so the subproblems for each block are independent. (See (3) in 13 for the local/block quadratic model, or F1-F3 in [8] for a description of the strongly convex local/block model). This has the advantage that the subproblem for each block can be solved/minimized in parallel, but has the disadvantage that no interaction between different blocks is captured by either algorithm. We also note that, while these algorithms are equipped with global convergence guarantees, iteration complexity results have not been established for either method, which is a drawback.

We make the following two comments about FLEXA and HyFLEXA. The first is that, although, as presented in [13] and [8], coordinate selection must respect the block structure of the problem, if the regularizer is assumed to be coordinate separable (which is the assumption we make in this work) then the convergence analysis in [13] and [8] can be adapted to hold in the 'non-fixed block' setting. The second comment is that, while the works in [13] and [8] only present global convergence guarantees, a variation of the method, presented in [29], is equipped with iteration complexity results. The algorithm in this paper incorporates partial (possibly inexact) second order information, allows inexact updates with a verifiable stopping condition, and is supported by iteration complexity guarantees.

Another important method that has recently been proposed is a randomized coordinate decent method by Qu et al. 28. One of the most notable features of their algorithm is that the block structure is not fixed throughout the algorithm. That is, at each iteration of their algorithm, a random subset of coordinates $S_{k}$ is selected from the set of coordinates $\{1, \ldots, N\}$. Then, a quadratic model is minimized exactly to obtain an update, and a new point is generated by applying the update to the current point. Unfortunately, this algorithm is only applicable to strongly convex smooth functions, or strongly convex empirical risk minimization problems, but is not suitable for general convex problems of the form (1). Moreover, the matrices used to define their quadratic model must obey several (strong) assumptions, and the quadratic model must be minimized exactly. These restrictions may make the algorithm inconvenient from a practical perspective.

The purpose of this work is to build upon the positive features of some existing coordinate descent methods and create a general flexible framework that can be used to solve any problem of the form (11). We will adopt some of the ideas in [3, 13,28] (including variable block structure; the incorporation of partial approximate second order information; the practicality of approximate solves and inexact updates; cheap line search) and create a new flexible randomized coordinate descent framework that is not only computationally practical, but is also supported by strong theoretical guarantees.

### 1.2 Contributions

In this section we list several of the contributions of this paper, which correspond to the central ideas of our FCD algorithm.

1. Blocks can vary throughout iterations. Existing randomized coordinate descent methods initially partition the data $x \in \mathbf{R}^{N}$ into $n(\leq N)$ blocks $\left(x=\left(x^{1}, \ldots, x^{n}\right) \in \mathbf{R}^{N}\right)$ and at each iteration, one/all/several of the blocks are selected according to some probability distribution, and those blocks are then updated. For those methods, the block partition is fixed throughout the entire algorithm. So, for example, coordinates in block $x^{1}$ will always be updated all together as a group, independently of any other block $x^{i} i \neq 1$. One of the main contributions of this work is that we allow the blocks to vary throughout the algorithm. i.e., the method is not restricted to a fixed block structure. No partition of $x$ need be initialized. Rather, when FCD is initialized, a parameter $1 \leq \tau \leq N$ is chosen and then, at each iteration of FCD, a subset $S \subseteq\{1, \ldots, N\}$ of size $|S|=\tau$ is randomly selected/generated, and the coordinates of $x$ corresponding to the indices in $S$ are updated. To the best of our knowledge, the only other paper that allows for this random mixing of coordinates/varying blocks strategy, is the work by Qu et al. [28] ${ }^{2}$ This variable block structure is crucial with regards to our next major contribution.

Remark 1 We note that the algorithms of Tseng and Yun 39, 40, 41 also allow a certain amount of coordinate mixing. However, their algorithms enforce deterministic rules for coordinate subset selection (either a 'generalized Gauss-Seidel', or Gauss-Southwell rule), which is different from our randomized selection rule. This difference is important because their deterministic strategy can be expensive/inconvenient to implement in practice, whereas our random scheme is simple and cheap to implement.
2. Model: Incorporation of partial second order information. FCD uses a quadratic model to determine the search direction. The model is defined by any symmetric positive definite matrix $H^{S}\left(x_{k}\right)$, which depends on both the subset of coordinates $S$, and also on the current point $x_{k}$. i.e., $H^{S}\left(x_{k}\right)$ is not fixed; rather, it can change/vary throughout the algorithm. We stress that this is a key advantage over most existing methods, which enforce the symmetric positive definite matrix defining their quadratic model to be fixed with respect to the fixed block structure, and/or iteration counter. To the best of our knowledge, the only works that allow the matrix $H^{S}\left(x_{k}\right)$ to vary with respect to the subset $S$ is this one (FCD), and the work by Qu et al. [28]. A crucial observation is that, if $H^{S}\left(x_{k}\right)$ approximates the Hessian, our approach allows every element of the Hessian to be accessed. On the other hand, other existing methods can only access blocks along the diagonal of (some perturbation of) the Hessian, (including [8, 12, 13, 31, 30, 37.)
Furthermore, the only restriction we make on the choice of the matrix $H^{S}\left(x_{k}\right)$ is that it be symmetric and positive definite. This is a much weaker assumption than made by Qu et al. [28], who insist that the matrix defining their quadratic model be a principle submatrix of some fixed/predefined $N \times N$ matrix $M$ say, and the large matrix $M$ must be symmetric and sufficiently positive definite. (See Section 2.1 in 28.)
3. Inexact search directions. To ensure that FCD is computationally practical, it is imperative that the iterates are inexpensive, and this is achieved through the use of inexact updates. That is, the model can be approximately minimized, giving inexact search directions. This strategy makes intuitive sense because, if the current point is far from the true solution, it may be wasteful/unnecessary to exactly minimize the model. Any algorithm can be used to approximately minimize the model, and if $H^{S}\left(x_{k}\right)$ approximates the Hessian, then the search direction obtained by minimizing the model is an approximate Newton-type direction. Importantly, the stopping conditions for the inner solve are easy to verify; they depend upon quantities that are easy/inexpensive to obtain, or may be available as a byproduct of the inner search direction solver. Block coordinate descent methods that use inexact search directions are uncommon in the literature because it is often notoriously difficult to check the conditions that determine whether an inexact search direction is 'acceptable'. A positive feature of our algorithm is that it is computationally practical to identify and verify the inexact directions used in FCD.

[^2]Remark 2 The precise form of our inexactness termination criterion is important because, not only is it computationally practical (i.e., easy to verify), it also allowed us to derive iteration complexity results for FCD (see point 5). We considered several other termination criterion formulations, but we were unable to obtain corresponding iteration complexity results (although global convergence results were possible). Currently, the majority of randomized CD methods require exact solves for their inner subproblems and we believe that a major reason for this is because iteration complexity results are significantly easier to obtain in the exact case. Coming up with practical inexact termination criteria for randomized CD methods that also allow the derivation of iteration complexity results seems to be a major hurdle in this area, although some progress is being made, e.g., [6, 9,36 .
4. Line search. FCD includes a line search to ensure a monotonic decrease of the objective function as iterates progress. (The line search is needed because we only make weak assumptions on the matrix $H^{S}\left(x_{k}\right)$.) The line search is inexpensive to perform because, at each iteration, it depends on a subset of coordinates $S$ only. One of the major advantages of incorporating second-order information combined with a line search is to allow, in practice, the selection of large step sizes (close to one). This is because unit step sizes can substantially improve the practical efficiency of a method. In fact, for all experiments that we performed, unit step sizes were accepted by the line search for the majority of the iterations. A commonly held view in this field is that function evaluations costs are unacceptably high, and so a line-search should not be used. Instead, block coordinate descent methods almost always use a fixed step size. However, FCD does indeed use a line search, and the numerical experiments presented in Section 9 show that our algorithm is competitive with, and often outperforms, other state-of-the-art block coordinate descent methods that use a fixed step size. Thus, we would argue that the previously mentioned reservations regarding the use of a line search are not necessarily well founded.
5. Convergence theory. We provide a complete convergence theory for FCD. In particular, we provide high probability iteration complexity guarantees for the algorithm in both the convex and strongly convex cases, and in the cases of both inexact or exact search directions. Our results show that FCD converges at a sublinear rate for convex functions, and at a linear rate for strongly convex functions. Complexity theory for block coordinate descent methods that use inexact updates is also uncommon, which is another strength of FCD.

### 1.3 Paper Outline

The paper is organized as follows. In Section 2 we introduce the notation and definitions that are used throughout this paper, including the definition of the quadratic model, and stationarity conditions. A thorough description of the FCD algorithm is presented in Section 3, including how the blocks are selected/sampled at each iteration (Section 3.1), a description of the search direction, several suggestions for selecting the matrices $H^{S}\left(x_{k}\right)$, and analysis of the subproblem/model termination conditions (Section 3.2), a definition of the line search (Section 3.3) and we also present several concrete problem examples (Section 3.4). In Section 4 we show that the line search in FCD will always be satisfied for some positive step $\alpha$, and that the objective function value is guaranteed to decrease at every iteration. Section 5 introduces several technical results, and in Section 6 we present our main iteration complexity results. In Section 7 we give a simplified analysis of FCD in the smooth case. Finally, several numerical experiments are presented in Section 9, which show that the algorithm performs very well in practice, even on ill-conditioned problems.

## 2 Preliminaries

Here we introduce the notation and definitions that are used in this paper, and we also present some important technical results. Throughout the paper we denote the standard Euclidean norm by $\|\cdot\| \equiv$ $\sqrt{\langle\cdot, \cdot\rangle}$ and the ellipsoidal norm by $\|\cdot\|_{A} \equiv \sqrt{\langle\cdot, A \cdot\rangle}$, where $A$ is a symmetric positive definite matrix. Furthermore, $\mathbf{R}_{++}^{N}$ denotes a vector in $\mathbf{R}^{N}$ with (strictly) positive components.

Subgradient and subdifferential. For a function $\Phi: \mathbf{R}^{N} \rightarrow \mathbf{R}$ the elements $s \in \mathbf{R}^{N}$ that satisfy

$$
\Phi(y) \geq \Phi(x)+\langle s, y-x\rangle, \quad \forall y \in \mathbf{R}^{N}
$$

are called the subgradients of $\Phi$ at point $x$. In words, all elements defining a linear function that supports $\Phi$ at a point $x$ are subgradients. The set of all $s$ at a point $x$ is called the subdifferential of $\Phi$ and it is denoted by $\partial \Phi(x)$.

Convexity. A function $\Phi: \mathbf{R}^{N} \rightarrow \mathbf{R}$ is strongly convex with convexity parameter $\mu_{\Phi}>0$ if

$$
\begin{equation*}
\Phi(y) \geq \Phi(x)+\langle s, y-x\rangle+\frac{\mu_{\Phi}}{2}\|y-x\|^{2}, \quad \forall x, y \in \mathbf{R}^{N}, \tag{2}
\end{equation*}
$$

where $s \in \partial \Phi(x)$. If $\mu_{\Phi}=0$. then function $\Phi$ is said to be convex.
Convex conjugate and proximal mapping. The proximal mapping of a convex function $\Psi$ at the point $x \in \mathbf{R}^{N}$ is defined as follows:

$$
\begin{equation*}
\operatorname{prox}_{\Psi}(x):=\arg \min _{y \in \mathbf{R}^{N}} \Psi(y)+\frac{1}{2}\|y-x\|^{2} . \tag{3}
\end{equation*}
$$

Furthermore, for a convex function $\Phi: \mathbf{R}^{N} \rightarrow \mathbf{R}$, it's convex conjugate $\Phi^{*}$ is defined as $\Phi^{*}(y):=$ $\sup _{u \in \mathbb{R}^{N}}\langle u, y\rangle-\Phi(u)$. From Chapter 1 of [1] we see that $\operatorname{prox}_{\Psi}(\cdot)$, and it's convex conjugate $\operatorname{prox}_{\Psi^{*}}(\cdot)$, are nonexpansive:

$$
\begin{equation*}
\left\|\operatorname{prox}_{\Psi}(y)-\operatorname{prox}_{\Psi}(x)\right\| \leq\|y-x\|, \quad \text { and } \quad\left\|\operatorname{prox}_{\Psi^{*}}(y)-\operatorname{prox}_{\Psi^{*}}(x)\right\| \leq\|y-x\| \tag{4}
\end{equation*}
$$

Coordinate decomposition of $\mathbf{R}^{N}$. Let $U \in \mathbb{R}^{N \times N}$ be a column permutation of the $N \times N$ identity matrix and further let $U=\left[U_{1}, U_{2}, \ldots, U_{N}\right]$ be a decomposition of $U$ into $N$ submatrices (column vectors), where $U_{i} \in \mathbf{R}^{N \times 1}$. It is clear that any vector $x \in \mathbb{R}^{N}$ can be written uniquely as $x=\sum_{i=1}^{N} U_{i} x^{i}$, where $x^{i} \in \mathbf{R}$ denotes the $i$ th coordinate of $x$.

Define $[N]:=\{1, \ldots, N\}$. Then we let $S \subseteq[N]$ and $U_{S} \in \mathbb{R}^{N \times|S|}$ be the collection of columns from matrix $U$ that have column indices in the set $S$. We denote the vector

$$
\begin{equation*}
x^{S}:=\sum_{i \in S} U_{i}^{T} x=U_{S}^{T} x \tag{5}
\end{equation*}
$$

Coordinate decomposition of $\Psi$. The function $\Psi: \mathbf{R}^{N} \rightarrow \mathbf{R}$ is assumed to be coordinate separable. That is, we assume that $\Psi(x)$ can be decomposed as:

$$
\begin{equation*}
\Psi(x)=\sum_{i=1}^{N} \Psi_{i}\left(x^{i}\right) \tag{6}
\end{equation*}
$$

where the functions $\Psi_{i}: \mathbf{R} \rightarrow \mathbf{R}$ are convex. We denote by $\Psi_{S}: \mathbf{R}^{|S|} \rightarrow \mathbf{R}$, where $S \subseteq[N]$, the function

$$
\begin{equation*}
\Psi_{S}\left(x^{S}\right)=\sum_{i \in S} \Psi_{i}\left(x^{i}\right) \tag{7}
\end{equation*}
$$

where $x^{S} \in \mathbf{R}^{|S|}$ is the collection of components from $x$ that have indices in set $S$. The following relationship will be used repeatedly in this work. For $x \in \mathbf{R}^{N}$, index set $S \subseteq[N]$, and $t, \hat{t} \in \mathbf{R}^{|S|}$, it holds that:

$$
\begin{equation*}
\Psi\left(x+U_{S} t^{S}\right)-\Psi\left(x+U_{S} \hat{t}^{S}\right) \stackrel{\sqrt{6}}{\stackrel{+7}{7}} \Psi_{S}\left(x^{S}+t^{S}\right)-\Psi_{S}\left(x^{S}+\hat{t}^{S}\right) \tag{8}
\end{equation*}
$$

Clearly, when $\hat{t}^{S} \equiv 0$, we have the special case $\Psi\left(x+U_{S} t^{S}\right)-\Psi(x)=\Psi_{S}\left(x^{S}+t^{S}\right)-\Psi_{S}\left(x^{S}\right)$.
Subset Lipschitz continuity of $f$. Throughout the paper we assume that the gradient of $f$ is Lipschitz, uniformly in $x$, for all subsets $S \subseteq[N]$. This means that, for all $x \in \mathbf{R}^{N}, S \subseteq[N]$ and $t^{S} \in \mathbf{R}^{|S|}$ we have

$$
\begin{equation*}
\left\|\nabla_{S} f\left(x+U_{S} t^{S}\right)-\nabla_{S} f(x)\right\| \leq L_{S}\left\|t^{S}\right\| \tag{9}
\end{equation*}
$$

where $\nabla_{S} f(x) \stackrel{5}{=} U_{S}^{T} \nabla f(x)$. An important consequence of (9) is the following standard inequality [23, p.57]:

$$
\begin{equation*}
f\left(x+U_{S} t^{S}\right) \leq f(x)+\left\langle\nabla_{S} f(x), t^{S}\right\rangle+\frac{L_{S}}{2}\left\|t^{S}\right\|^{2} . \tag{10}
\end{equation*}
$$

Radius of the Levelset. Let $X_{*}$ denote the set of optimal solutions of (1), and let $x_{*}$ be any member of that set. We define

$$
\begin{equation*}
\mathcal{R}(x)=\max _{y} \max _{x_{*} \in X_{*}}\left\{\left\|y-x_{*}\right\|: F(y) \leq F(x)\right\} \tag{11}
\end{equation*}
$$

which is a measure of the size of the level set of $F$ given by $x$. In this work we make the standard assumption that $\mathcal{R}\left(x_{0}\right)$ is bounded at the initial iterate $x_{0}$. We also define a scaled version of (11)

$$
\begin{equation*}
\mathcal{R}_{w}(x)=\max _{y} \max _{x_{*} \in X_{*}}\left\{\left\|y-x_{*}\right\|_{w}: F(y) \leq F(x)\right\} \tag{12}
\end{equation*}
$$

where $w \in \mathbf{R}_{++}^{N}$ and $\|u\|_{w}=\left(\sum_{i=1}^{N} w_{i} u_{i}^{2}\right)^{1 / 2}$.

### 2.1 Piecewise Quadratic Model

For fixed $x \in \mathbf{R}^{N}$, we define a piecewise quadratic approximation of $F$ around the point $(x+t) \in \mathbf{R}^{N}$ as follows: $F(x+t) \approx f(x)+Q(x ; t)$, where

$$
\begin{equation*}
Q(x ; t):=\langle\nabla f(x), t\rangle+\frac{1}{2}\|t\|_{H(x)}^{2}+\Psi(x+t) \tag{13}
\end{equation*}
$$

and $H(x)$ is any symmetric positive definite matrix. We also define

$$
\begin{equation*}
Q_{S}\left(x, t^{S}\right):=\left\langle\nabla_{S} f(x), t^{S}\right\rangle+\frac{1}{2}\left\|t^{S}\right\|_{H^{S}(x)}^{2}+\Psi_{S}\left(x^{S}+t^{S}\right), \tag{14}
\end{equation*}
$$

and $H^{S}(x) \in \mathbf{R}^{|S| \times|S|}$ is a square submatrix (principal minor) of $H(x)$ that corresponds to the selection of columns and rows from $H(x)$ with column and row indices in set $S$. Notice that $Q_{S}\left(x, t^{S}\right)$ is the quadratic model for the collection of coordinates in $S$.

Similarly to (8), we have the following important relationship. For $x \in \mathbf{R}^{N}$, index set $S \subseteq[N]$, and $t, \hat{t} \in \mathbf{R}^{|S|}$, it holds that:

$$
\begin{align*}
Q\left(x ; U_{S} t^{S}\right)-Q\left(x ; U_{S} \hat{t}^{S}\right) \stackrel{\boxed{13}}{=} & \left\langle\nabla f(x), U_{S} t^{S}\right\rangle+\frac{1}{2}\left\|U_{S} t^{S}\right\|_{H(x)}^{2}+\Psi\left(x+U_{S} t^{S}\right) \\
& -\left\langle\nabla f(x), U_{S} \hat{t}^{S}\right\rangle-\frac{1}{2}\left\|U_{S} \hat{t}^{S}\right\|_{H(x)}^{2}-\Psi\left(x+U_{S} \hat{t}^{S}\right) \\
\stackrel{8}{=} & \left\langle\nabla_{S} f(x), t^{S}\right\rangle-\left\langle\nabla_{S} f(x), \hat{t}^{S}\right\rangle+\frac{1}{2}\left\|t^{S}\right\|_{H^{S}(x)}^{2}-\frac{1}{2}\left\|\hat{t}^{S}\right\|_{H^{S}(x)}^{2} \\
& +\Psi_{S}\left(x^{S}+t^{S}\right)-\Psi_{S}\left(x^{S}+\hat{t}^{S}\right) \\
\stackrel{14}{=} & Q_{S}\left(x, t^{S}\right)-Q_{S}\left(x, \hat{t}^{S}\right)
\end{align*}
$$

### 2.2 Stationarity conditions.

The following theorem gives the equivalence of some stationarity conditions of problem (1).
Theorem 1 The following are equivalent first order optimality conditions for problem (1), Section 2 in [26].
(i) $\nabla f(x)+s=0$ and $s \in \partial \Psi(x)$,
(ii) $-\nabla f(x) \in \partial \Psi(x)$,
(iii) $\nabla f(x)+\operatorname{prox}_{\Psi^{*}}(x-\nabla f(x))=0$,
(iv) $x=\operatorname{prox}_{\Psi}(x-\nabla f(x))$.

Let us define the continuous function

$$
\begin{equation*}
g(x ; t):=\nabla f(x)+H(x) t+\operatorname{prox}_{\Psi^{*}}(x+t-\nabla f(x)-H(x) t) . \tag{16}
\end{equation*}
$$

By Theorem 1. the points that satisfy $g(x ; 0)=\nabla f(x)+\operatorname{prox}_{\Psi^{*}}(x-\nabla f(x))=0$ are stationary points for problem (1). Hence, $g(x ; 0)$ is a continuous measure of the distance from the set of stationary points of problem (1). Furthermore, we define

$$
\begin{equation*}
g_{S}\left(x ; t^{S}\right):=\nabla_{S} f(x)+H^{S}(x) t^{S}+\operatorname{prox}_{\Psi_{S}^{*}}\left(x^{S}+t^{S}-\nabla_{S} f(x)-H^{S}(x) t^{S}\right) \tag{17}
\end{equation*}
$$

## 3 The Algorithm

In this section we present the Flexible Coordinate Descent (FCD) algorithm for solving problems of the form (1). There are three key steps in the algorithm: (Step 4) the coordinates are sampled randomly; (Step 5) the model (14) is solved approximately until rigorously defined stopping conditions are satisfied; (Step 6) a line search is performed to find a step size that ensures a sufficient reduction in the objective value. Once these key steps have been performed, the current point $x_{k}$ is updated to give a new point $x_{k+1}$, and the process is repeated.

In the algorithm description, and later in the paper, we will use the following definition for the set $\Omega$ :

$$
\begin{equation*}
\Omega:=\partial Q_{S}\left(x_{k} ; t_{k}^{S}\right) \tag{18}
\end{equation*}
$$

We are now ready to present pseudocode for the FCD algorithm, while a thorough description of each of the key steps in the algorithm will follow in the rest of this section.

```
Algorithm 1 Flexible Coordinate Descent (FCD)
    Input Choose \(x_{0} \in \mathbf{R}^{N}, \theta \in(0,1)\) and \(1 \leq \tau \leq N\).
    for \(k=1,2, \ldots\) do
        Sample a subset of coordinates \(S \subseteq[N]\), of size \(|S|=\tau\), with probability \(\mathbf{P}(S)>0\)
        If \(g_{S}\left(x_{k} ; 0\right)=0\) go to Step 3 ; else select \(\eta_{k}^{S} \in[0,1)\) and approximately solve
                    \(t_{k}^{S} \approx \underset{t^{S}}{\arg \min } Q_{S}\left(x_{k} ; t^{S}\right)\),
``` until the following stopping conditions hold:
\[
\begin{equation*}
Q\left(x_{k} ; U_{S} t_{k}^{S}\right)<Q\left(x_{k} ; 0\right) \tag{20}
\end{equation*}
\]
and
\[
\begin{equation*}
\inf _{u \in \Omega}\left\|u-g_{S}\left(x_{k} ; t_{k}^{S}\right)\right\|^{2} \leq\left(\eta_{k}^{S}\left\|g_{S}\left(x_{k} ; 0\right)\right\|\right)^{2}-\left\|g_{S}\left(x_{k} ; t_{k}^{S}\right)\right\|^{2} \tag{21}
\end{equation*}
\]

5: Perform a backtracking line search along the direction \(t_{k}^{S}\) starting from \(\alpha=1\). i.e., find \(\alpha \in(0,1]\) such that
\[
\begin{equation*}
F\left(x_{k}\right)-F\left(x_{k}+\alpha U_{S} t_{k}^{S}\right) \geq \theta\left(\ell\left(x_{k} ; 0\right)-\ell\left(x_{k} ; \alpha U_{S} t_{k}^{S}\right)\right) \tag{22}
\end{equation*}
\]
where
\[
\begin{equation*}
\ell\left(x_{k} ; t\right):=f\left(x_{k}\right)+\left\langle\nabla f\left(x_{k}\right), t\right\rangle+\Psi\left(x_{k}+t\right) \tag{23}
\end{equation*}
\]
```

        Update }\mp@subsup{x}{k+1}{}=\mp@subsup{x}{k}{}+\alpha\mp@subsup{U}{S}{}\mp@subsup{t}{k}{S
    ```
    end for

\subsection*{3.1 Selection of coordinates (Step 3)}

One of the central ideas of this algorithm is that the selection of coordinates to be updated at each iteration is chosen randomly. This allows the coordinates to be selected very quickly and cheaply. For all the results in this work, we use a ' \(\tau\)-nice sampling' 31. This sampling is described briefly below, and a full description can be found in 31.

Let \(2^{[N]}\) denote the set of all subsets of \([N]\), which includes the empty set \(\emptyset\) and \([N]\) itself. At every iteration of FCD, a subset is sampled from \(2^{[N]}\). The probability of a set \(S \in 2^{[N]}\) being selected is denoted by \(\mathbf{P}(S)\), and we are interested in probability distributions with \(\mathbf{P}(S)>0 \forall S \in 2^{[N]} \backslash \emptyset\) and \(\mathbf{P}(\emptyset)=0\). i.e., all non-empty sets have positive probability of being selected.

In what follows we describe a (specific instance of a) uniform probability distribution (called a \(\tau\) nice sampling). In particular, the uniform probability distribution is constructed in such a way that subsets in \(2^{[N]}\) with the same cardinality have equal probability of being selected, i.e., \(\mathbf{P}(S)=\mathbf{P}\left(S^{\prime}\right)>0\) \(\forall S, S^{\prime} \in 2^{[N]}\) with \(|S|=\left|S^{\prime}\right|\).

Definition 1 ( \(\tau\)-nice sampling) Given an integer \(1 \leq \tau \leq N\), a set \(S\) with \(|S|=\tau\) is selected uniformly with probability one, and the sampling is uniquely characterised by its density function:
\[
\begin{equation*}
\mathbf{P}(S)=\frac{1}{\binom{N}{\tau}} \forall S \in 2^{[N]} \text { with }|S|=\tau, \quad \mathbf{P}(S)=0 \forall S \in 2^{[N]} \text { with }|S| \neq \tau . \tag{24}
\end{equation*}
\]

Assumption 2 In Step 3 of FCD (Algorithm1), we assume that the sampling procedure used to generate the subset of coordinates \(S\) in that given in Definition 1.

Below is a technical result (see 31) that uses the probability distribution (24) and is used in the worstcase iteration complexity results of this paper. Let \(\theta_{i}\) with \(i=1,2, \ldots, N\) be some constants. Then
\[
\begin{equation*}
\mathbf{E}\left[\sum_{i \in S} \theta_{i}\right]=\frac{\tau}{N} \sum_{i=1}^{N} \theta_{i} . \tag{25}
\end{equation*}
\]

\subsection*{3.2 The search direction and Hessian approximation (Step 4)}

The search direction (Step 4 of FCD) is determined as follows. Given a set of coordinates \(S\), FCD forms a model for \(S(14\), and minimizes the model approximately until the stopping conditions (20) and 21 are satisfied, giving an 'inexact' search direction. We also describe the importance of the choice of the matrix \(H\), which defines the model and is an approximate second order information term. Henceforth, we will use the shorthand \(H_{k}^{S} \equiv H^{S}\left(x_{k}\right)\).

\subsection*{3.2.1 The search direction}

The subproblem (19), (where \(Q_{S}\left(x_{k} ; t^{S}\right)\) is defined in 14 ) is approximately solved, and the search direction \(t_{k}^{S}\) is accepted when the stopping conditions 20 and 21 are satisfied, for some \(\eta_{k}^{S} \in[0,1)\). Notice that
\[
\begin{align*}
& Q\left(x ; U_{S} t^{S}\right)-Q(x ; 0) \stackrel{\stackrel{15}{=}}{=} Q_{S}\left(x ; t^{S}\right)-Q_{S}(x ; 0) \\
& \stackrel{144}{=}\left\langle\nabla_{S} f(x), t^{S}\right\rangle+\frac{1}{2}\left\|t^{S}\right\|_{H^{S}}^{2}+\Psi_{S}\left(x^{S}+t^{S}\right)-\Psi_{S}\left(x^{S}\right) \tag{26}
\end{align*}
\]

Hence, from (26), the stopping conditions (20) and (21) depend on subset \(S\) only, and are therefore inexpensive to verify, meaning that they are implementable.

\section*{Remark 3}
(i) At some iteration \(k\), it is possible that \(g_{S}\left(x_{k} ; 0\right)=0\). In this case, it is easy to verify that the optimal solution of subproblem 19 is \(t_{k}^{S}=0\). Therefore, before calculating \(t_{k}^{S}\) we check first if condition \(g_{S}\left(x_{k} ; 0\right)=0\) is satisfied.
(ii) Notice that, unless at optimality (i.e., \(g\left(x_{k} ; 0\right)=0\) ), there will always exist a subset \(S\) such that \(g_{S}\left(x_{k} ; 0\right) \neq 0\), which implies that \(t_{k}^{S} \neq 0\). Hence, FCD will not stagnate.

\subsection*{3.2.2 The Hessian approximation}

One of the most important features of this method is that the quadratic model (14) incorporates second order information in the form of a symmetric positive definite matrix \(H_{k}^{S}\). This is key because, depending upon the choice of \(H_{k}^{S}\), it makes the method robust. Moreover, at each iteration, the user has complete freedom over the choice of \(H_{k}^{S} \succ 0\).

We now provide a few suggestions for the choice of \(H_{k}^{S}\). (This list is not intended to be exhaustive.) Notice that in each case there is a trade off between a matrix that is inexpensive to work with, and one that is a more accurate representation of the true block Hessian.
1. Identity. Here, \(H_{k}^{S}=I\), so that no second order information is employed by the method.
2. Multiple of the identity. One could set \(H_{k}^{S}=\nu I\), for some \(\nu>0\). In particular, a popular choice is \(\nu=L_{S}\), where \(L_{S}\) is the Lipschitz constant defined in (9).
3. Diagonal matrix. When \(H_{k}^{S}\) is a diagonal matrix, \(H_{k}^{S}\) and it's inverse are inexpensive to work with. In particular, \(H_{k}^{S}=\operatorname{diag}\left(\nabla_{S}^{2} f\left(x_{k}\right)\right)\) captures partial curvature information, and is an effective choice in practice, particularly when \(\operatorname{diag}\left(\nabla^{2} f(x)\right)\) is a good approximation to \(\nabla^{2} f(x)\). Moreover, if \(f\) is quadratic, then \(\nabla^{2} f\left(x_{k}\right)\) is constant for all \(k\), so \(\operatorname{diag}\left(\nabla^{2} f(x)\right)\) can be computed and stored at the start of the algorithm and elements can be accessed throughout the algorithm as necessary.
4. Principal minor of the Hessian. When \(H_{k}^{S}=\nabla_{S}^{2} f\left(x_{k}\right)\), second order information is incorporated into FCD, which can be very beneficial for ill-conditioned problems. However, this choice of \(H_{k}^{S}\) is usually more computationally expensive to work with. In practice, \(\nabla_{S}^{2} f\left(x_{k}\right)\) may be used in a matrix-free way and not explicitly stored. For example, there may be an analytic formula for performing matrix-vector products with \(\nabla_{S}^{2} f\left(x_{k}\right)\), or techniques from automatic differentiation could be employed, see [25] Section 7].
5. Quasi-Newton approximation. One may choose \(H_{k}^{S}\) to be an approximation to \(\nabla_{S}^{2} f\left(x_{k}\right)\) based on a limited-memory BFGS update scheme, see [25, Section 8]. This approach might be suitable in cases where the problem is ill-conditioned, but performing matrix-vector products with \(\nabla_{S}^{2} f\left(x_{k}\right)\) is expensive.

Remark 4 We make the following remarks regarding the choice of \(H_{k}^{S}\).
1. If any of the matrices mentioned above are not positive definite, then they can be adjusted to make them so. For example, if \(H_{k}^{S}\) is diagonal, any zero that appears on the diagonal can be replaced with a positive number. Moreover, if \(\nabla_{S}^{2} f\left(x_{k}\right)\) is not positive definite, a multiple of the identity can be added to it.
2. If \(H_{k}^{S}\) is chosen to be a diagonal matrix, then the quadratic model 14 is separable, so the subproblems for each coordinate in \(S\) are independent. Moreover, in some cases this gives rise to a closed form solution for the search direction \(t_{k}^{S}\). (For example, if \(\Psi(x) \equiv\|x\|_{1}\), then soft thresholding may be used.)

An advantage of the FCD algorithm (if Option 4 is used for \(H_{k}^{S}\) ) is that all elements of the Hessian can be accessed. This is because the set of coordinates can change at every iteration, and so too can matrix \(H_{k}^{S}\). This makes FCD extremely flexible and is particularly advantageous when there are large off diagonal elements in the Hessian. The importance of incorporation of information from off-diagonal elements is demonstrated in numerical experiments in Section 9 .

\subsection*{3.2.3 Termination criteria for model minimization}

In Step 4 of FCD, the termination criteria 20 and 21 are used to determine whether an acceptable search direction has been found. For composite functions of the form (1), both conditions are required to ensure that \(t_{S}^{k}\) is a descent direction. Moreover, (20) is important for intuitively obvious reasons: the model should be decreased at each iteration. We will now attempt to explain/justify the use of the second termination condition (21).

The parameter \(\eta_{k}^{S}\) plays a key role; it determines how 'inexactly' the quadratic model \(\sqrt{19}\) may be solved (or equivalently, how 'inexact' the search direction \(t_{k}^{S}\) is allowed to be). If one sets \(\eta_{k}^{S}=0\), then the model \(\sqrt{19}\) is minimized exactly, leading to exact search directions. On the other hand, if \(\eta_{k}^{S}>0\) then the model is approximately minimized, leading to inexact search directions.

To obtain iteration complexity guarantees for FCD, the termination conditions must be chosen carefully. In particular, it becomes obvious in Lemma 3, that 20) and 21) are the appropriate conditions. We now proceed to show that the conditions are mathematically sound. By rearranging (21) we obtain
\[
\begin{equation*}
\inf _{u \in \Omega}\left\|u-g_{S}\left(x_{k} ; t_{k}^{S}\right)\right\|^{2}+\left\|g_{S}\left(x_{k} ; t_{k}^{S}\right)\right\|^{2} \leq\left(\eta_{k}^{S}\left\|g_{S}\left(x_{k} ; 0\right)\right\|\right)^{2} \tag{27}
\end{equation*}
\]

The left hand side in (27) is a continuous function because the 'inf' operator is an orthonormal projection onto a closed subspace \(\Omega:=\nabla_{S} f\left(x_{k}\right)+\partial \Psi_{S}\left(x_{k_{S}}^{S}+t_{k}^{S}\right)\), see 18 , and \(\left\|g_{S}\left(x_{k} ; t_{k}^{S}\right)\right\|^{2}\) is continuous as a function of \(t_{k}^{S}\). Furthermore, there exists a point \(t_{k}^{S}\) such that 27 is satisfied, and this is the minimizer of (19). In particular, if \(t_{k}^{S}\) is the minimizer of (19), then from Theorem 1 we have that \(g_{S}\left(x_{k} ; t_{k}^{S}\right)=0 \in \Omega\). Hence, the left hand side in (27) is zero and the condition is satisfied for any \(\eta^{k} \in[0,1)\) and \(g_{S}\left(x_{k} ; 0\right)\). Since the right hand side in (27) is a non-negative constant, the left hand side is continuous and the minimizer of (19) satisfies 27, then there exists a closed ball centered at the minimizer such that within
this ball (27) is satisfied. Therefore, any convergent method which can be used to minimize (19) will eventually satisfy the termination condition (27) without solving (19) exactly, unless the right hand side in (27) is zero.

We stress that this projection operation is inexpensive for the applications that we consider, e.g., when \(\Psi\) is the \(\ell_{1}\)-norm or the \(\ell_{2}\)-norm squared, or elastic net regularization, which is a combinations of these two. In particular, the subdifferential \(\Omega\) has an 'analytic' form, which also allows for a closed form solution of the left hand side of the termination criterion. Below we provide an example for the case where \(\Psi=\tau\|x\|_{1}\). Using the definition of \(\Omega\) we have that the left hand side problem in 27) \(\inf _{u \in \Omega}\left\|u-g_{S}\left(x_{k} ; t_{k}^{S}\right)\right\|^{2}\) is equivalent to
\[
\inf _{u-\nabla_{S} f\left(x_{k}\right)-H^{S}\left(x_{k}\right) t_{k}^{S} \in \partial \Psi_{S}^{*}\left(x_{k}^{S}+t_{k}^{S}\right)}\left\|u-g_{S}\left(x_{k} ; t_{k}^{S}\right)\right\|^{2}
\]

By making a change of coordinates from \(u\) to \(v\) as \(u:=\nabla_{S} f\left(x_{k}\right)+H^{S}\left(x_{k}\right) t_{k}^{S}+v\), we rewrite the previous problem as
\[
\inf _{v \in \partial \Psi_{S}^{*}\left(x_{k}^{S}+t_{k}^{S}\right)}\left\|v+\nabla_{S} f\left(x_{k}\right)+H^{S}\left(x_{k}\right) t_{k}^{S}-g_{S}\left(x_{k} ; t_{k}^{S}\right)\right\|^{2}
\]

Using the definition of \(g_{S}\left(x_{k} ; t_{k}^{S}\right)\) from (17) we have that the previous problem is equivalent to
\[
\begin{equation*}
\inf _{v \in \partial \Psi_{S}^{*}\left(x_{k}^{S}+t_{k}^{S}\right)}\left\|v-\operatorname{prox}_{\Psi_{S}^{*}}\left(x_{k}^{S}+t_{k}^{S}-\nabla_{S} f\left(x_{k}\right)-H^{S}\left(x_{k}\right) t_{k}^{S}\right)\right\|^{2} \tag{28}
\end{equation*}
\]

The subdifferential \(\partial \Psi_{S}^{*}\left(x_{k}^{S}+t_{k}^{S}\right)\) has the following coordinate-wise form
\[
\partial \Psi_{i}^{*}\left(x_{k}^{i}+t_{k}^{i}\right)=\left\{\begin{array}{ll}
\tau & \text { if } x_{k}^{i}+t_{k}^{i}>0 \\
{[-\tau, \tau]} & \text { if } x_{k}^{i}+t_{k}^{i}=0, \\
-\tau & \text { if } x_{k}^{i}+t_{k}^{i}<0
\end{array} \quad \forall i \in S\right.
\]

Using the above analytic form of the subdifferential it is easy to see that problem (28) has a closed form solution that is inexpensive to compute.

Remark 5 We stress that if the regularizer is the \(\ell_{1}\)-norm, or the squared \(\ell_{2}\)-norm, which are two very popular regularizers, our inexactness criterion 27 is inexpensive to implement in practice. However, for other general regularizers, this may not be the case, and it may not be possible to compute (27) analytically. In such a case, it may be necessary to compute the minimizer of the quadratic model exactly, or to use some other iterative method, both of which come with an associated computational cost.

\subsection*{3.3 The line search (Step 5)}

The stopping conditions (20) and (21) ensure that \(t_{k}^{S}\) is a descent direction, but if the full step \(x_{k}+U_{S} t_{k}^{S}\) is taken, a reduction in the function value (1) is not guaranteed. To this end, we include a line search step in our algorithm in order to guarantee a monotonic decrease in the function \(F\). Essentially, the line search guarantees the sufficient decrease of \(F\) at every iteration, where sufficient decrease is measured by the loss function (23).

In particular, for fixed \(\theta \in(0,1)\), we require that for some \(\alpha \in(0,1], 22)\) is satisfied. (In Lemma 1 we prove the existence of such an \(\alpha\).) Notice that
\[
\begin{align*}
\ell\left(x ; U_{S} t^{S}\right)-\ell(x ; 0) & \stackrel{\sqrt[23]{=}}{=}\left\langle\nabla_{S} f(x), t^{S}\right\rangle+\Psi\left(x+U_{S} t^{S}\right)-\Psi(x)  \tag{29}\\
& \stackrel{8}{=}\left\langle\nabla_{S} f(x), t^{S}\right\rangle+\Psi_{S}\left(x^{S}+t^{S}\right)-\Psi_{S}\left(x^{S}\right)
\end{align*}
\]
which shows that the calculation of the right hand side of 22 only depends upon block \(S\), so it is inexpensive. Moreover, the line search condition (Step 5) involves the difference between function values \(F\left(x_{k}\right)-F\left(x_{k}+\alpha U_{S} t_{k}^{S}\right)\). Fortunately, while function values can be expensive to compute, computing the difference between function values at successive iterates need not be. (See Section 3.4 and the numerical results in Section 9 and Table 6.)

\subsection*{3.4 Concrete examples}

Now we give two examples to show that the difference between function values at successive iterates can be computed efficiently, which demonstrates that the line-search is implementable. The first example considers a convex composite function, where the smooth term is a quadratic loss term and the second example considers a logistic regression problem.

Quadratic loss plus regularization example. Suppose that \(f(x)=\frac{1}{2}\|A x-b\|^{2}\) and \(\Psi(x)\) is not equivalent to zero, where \(A \in \mathbb{R}^{m \times N}, b \in \mathbb{R}^{m}\) and \(x \in \mathbf{R}^{N}\). Then
\[
\begin{align*}
F\left(x_{k}\right)-F\left(x_{k}+\alpha U_{S} t_{k}^{S}\right) & \stackrel{8}{=} f\left(x_{k}\right)+\Psi_{S}\left(x^{S}\right)-f\left(x_{k}+\alpha U_{S} t_{k}^{S}\right)-\Psi_{S}\left(x_{k}^{S}+\alpha t_{k}^{S}\right)  \tag{30}\\
& =\Psi_{S}\left(x^{S}\right)-\alpha\left\langle\nabla_{S} f(x), t_{k}^{S}\right\rangle-\frac{\alpha^{2}}{2}\left\|A_{i} t_{k}^{S}\right\|_{2}^{2}-\Psi_{S}\left(x_{k}^{S}+\alpha t_{k}^{S}\right)
\end{align*}
\]

The calculation \(F\left(x_{k}\right)-F\left(x_{k}+\alpha U_{S} t_{k}^{S}\right)\), as a function of \(\alpha\), only depends on subset \(S\). Hence, it is inexpensive. Moreover, often the quantities in (30) are already needed in the computation of the search direction \(t\), so regarding the line search step, they essentially come "for free".

Logistic regression example. Suppose that \(f(x) \equiv \sum_{j=1}^{m} \log \left(1+e^{-b_{j} a_{j}^{T} x}\right)\) and \(\Psi(x)\) is not equivalent to zero, where \(a_{j}^{T}\) is the \(j\) th row of a matrix \(A \in \mathbb{R}^{m \times N}\) and \(b_{j}\) is the \(j\) th component of vector \(b \in \mathbb{R}^{m}\). As before, we need to evaluate (30). Let us split calculation of \(F\left(x_{k}\right)-F\left(x_{k}+\alpha U_{S} t_{k}^{S}\right)\) in parts. The first part \(\Psi_{S}\left(x^{S}\right)-\Psi_{S}\left(x_{k}^{S}+\alpha t_{k}^{S}\right)\) is inexpensive, since it depends only on subset \(S\). The second part \(f\left(x_{k}\right)-f\left(x_{k}+\alpha U_{S} t_{k}^{S}\right)\) is more expensive because is depends upon the logarithm. In this case, one can calculate \(f\left(x_{0}\right)\) once at the beginning of the algorithm and then update \(f\left(x_{k}+\alpha U_{S} t_{k}^{S}\right) \forall k \geq 1\) less expensively. In particular, let us assume that the following terms:
\[
\begin{equation*}
e^{-b_{j} a_{j}^{T} x_{0}} \quad \forall j \quad \text { and } \quad f\left(x_{0}\right)=\sum_{j=1}^{m} \log \left(1+e^{-b_{j} a_{j}^{T} x_{0}}\right), \tag{31}
\end{equation*}
\]
are calculated once and stored in memory. Then, at iteration 1 , the calculation of \(f\left(x_{0}+\alpha U_{S} t_{0}^{S}\right)=\) \(\sum_{j=1}^{m} \log \left(1+e^{-b_{j} a_{j}^{T} x_{0}} e^{-\alpha b_{j} a_{j}^{T}\left(U_{S} t_{0}^{S}\right)}\right)\) is required for different values of \(\alpha\) by the backtracking line search algorithm. The most demanding task in calculating \(f\left(x_{0}+\alpha U_{S} t_{0}^{S}\right)\) is the calculation of the products \(b_{j} a_{j}^{T}\left(U_{S} t_{0}^{S}\right) \forall j\) once, which is inexpensive since \(\forall j\) this operation depends only on subset \(S\). Having \(b_{j} a_{j}^{T}\left(U_{S} t_{0}^{S}\right) \forall j\) and (31) calculation of \(f\left(x_{0}\right)-f\left(x_{0}+\alpha U_{S} t_{0}^{S}\right)\) for different values of \(\alpha\) is inexpensive. At the end of the process, \(f\left(x_{1}\right)\) and \(e^{-b_{j} a_{j}^{T} x_{1}} \forall j\) are given for free, and the process can be repeated for the calculation of \(f\left(x_{1}\right)-f\left(x_{1}+\alpha U_{S} t_{1}^{S}\right)\) etc.

Remark 6 The examples above show that, for quadratic loss functions, and logistic regression problems, computing function values is relatively inexpensive. Thus, for problems with this, or similiar structure, FCD is competitive, even though it requires function evaluations for the line search. (The competitiveness of FCD is confirmed in our numerical experiments in Section 9.) However, we stress that, for other applications with more general objective functions, it may not be possible to perform function evaluations efficiently, in which case FCD may not be a suitable algorithm, and a different algorithm could be used.

\section*{4 Bounded step-size and monotonic decrease of \(F\)}

Throughout this section we denote \(H_{k}^{S} \equiv H^{S}\left(x_{k}\right)\). The following assumptions are made about \(H_{k}^{S}\) and \(f\). Assumption 3 explains that the Hessian approximation \(H_{k}^{S}\) must be positive definite for all subsets of coordinates \(S\) at all iterations \(k\). Assumption 4 explains that the gradient of \(f\) must be Lipschitz for all subsets \(S\).

Assumption 3 There exist \(0<\lambda_{S} \leq \Lambda_{S}\), such that the sequence of symmetric \(\left\{H_{k}^{S}\right\}_{k \geq 0}\) satisfies:
\[
\begin{equation*}
0<\lambda_{S} \leq \lambda_{\min }\left(H_{k}^{S}\right) \quad \text { and } \quad \lambda_{\max }\left(H_{k}^{S}\right) \leq \Lambda_{S}, \quad \text { for all } S \subseteq[N] \tag{32}
\end{equation*}
\]

Assumption 4 The function \(f\) is smooth and satisfies (9) for all \(S \subseteq[N]\).

The next assumption regards the relationship between the parameters introduced in Assumptions 3 and 4.

Assumption 5 The relation \(\lambda_{S} \leq L_{S}\) holds for all \(S \subseteq[N]\).
The following lemma shows that if \(t_{k}^{S}\) is nonzero, then \(F\) is decreased. The proof closely follows that of [3, Theorem 3.1].

Lemma 1 Let Assumptions 2, 3, 4 and 5 hold, and let \(\theta \in(0,1)\). Given \(x_{k}\), let \(S\) and \(t_{k}^{S} \neq 0\) be generated by FCD (Algorithm 1) with \(\eta_{k}^{S} \in[0,1)\). Then Step 6 of \(F C D\) will accept a step-size \(\alpha\) that satisfies
\[
\begin{equation*}
\alpha \geq \alpha_{S}, \quad \text { where } \quad \alpha_{S}:=(1-\theta) \frac{\lambda_{S}}{2 L_{S}}>0 \tag{33}
\end{equation*}
\]

Furthermore,
\[
\begin{equation*}
F\left(x_{k}\right)-F\left(x_{k}+\alpha U_{S} t_{k}^{S}\right)>\theta(1-\theta) \frac{\lambda_{S}^{2}}{4 L_{S}}\left\|t_{k}^{S}\right\|^{2} \tag{34}
\end{equation*}
\]

Proof From (20), we have \(0>Q\left(x_{k} ; U_{S} t_{k}^{S}\right)-Q\left(x_{k} ; 0\right) \stackrel{13)}{=} \ell\left(x_{k} ; U_{S} t_{k}^{S}\right)-\ell\left(x_{k} ; 0\right)+\frac{1}{2}\left\|t_{k}^{S}\right\|_{H_{k}^{S}}^{2}\). Rearranging gives
\[
\begin{equation*}
\ell\left(x_{k} ; 0\right)-\ell\left(x_{k} ; U_{S} t_{k}^{S}\right)>\frac{1}{2}\left\|t_{k}^{S}\right\|_{H_{k}^{S}}^{2} \stackrel{32}{\geq} \frac{1}{2} \lambda_{S}\left\|t_{k}^{S}\right\|^{2} . \tag{35}
\end{equation*}
\]

Denote \(x_{k+1}=x_{k}+\alpha U_{S} t_{k}^{S}\). By Assumption 4, for all \(\alpha \in(0,1]\), we have
\[
F\left(x_{k+1}\right) \leq f\left(x_{k}\right)+\alpha\left\langle\nabla_{S} f\left(x_{k}\right), t_{k}^{S}\right\rangle+\frac{L_{S}}{2} \alpha^{2}\left\|t_{k}^{S}\right\|^{2}+\Psi\left(x_{k}+\alpha U_{S} t_{k}^{S}\right)
\]

Adding \(\Psi\left(x_{k}\right)\) to both sides of the above and rearranging gives
\[
\begin{align*}
F\left(x_{k}\right)-F\left(x_{k+1}\right) & \geq-\alpha\left\langle\nabla_{S} f\left(x_{k}\right), t_{k}^{S}\right\rangle-\frac{L_{S}}{2} \alpha^{2}\left\|t_{k}^{S}\right\|^{2}-\Psi\left(x_{k}+\alpha U_{S} t_{k}^{S}\right)+\Psi\left(x_{k}\right) \\
& \stackrel{29}{=} \ell\left(x_{k} ; 0\right)-\ell\left(x_{k} ; \alpha U_{S} t_{k}^{S}\right)-\frac{L_{S}}{2} \alpha^{2}\left\|t_{k}^{S}\right\|^{2} . \tag{36}
\end{align*}
\]

By convexity of \(\Psi(x)\) we have that
\[
\begin{equation*}
\ell\left(x_{k} ; 0\right)-\ell\left(x_{k} ; \alpha U_{S} t_{k}^{S}\right) \geq \alpha\left(\ell\left(x_{k} ; 0\right)-\ell\left(x_{k} ; U_{S} t_{k}^{S}\right)\right) . \tag{37}
\end{equation*}
\]

Then
\[
\begin{align*}
F\left(x_{k}\right)-F\left(x_{k+1}\right) & -\theta\left(\ell\left(x_{k} ; 0\right)-\ell\left(x_{k} ; \alpha U_{S} t_{k}^{S}\right)\right) \\
& \stackrel{\sqrt{36}}{\geq}(1-\theta)\left(\ell\left(x_{k} ; 0\right)-\ell\left(x_{k} ; \alpha U_{S} t^{S}\right)\right)-\frac{L_{S}}{2} \alpha^{2}\left\|t_{k}^{S}\right\|^{2} \\
& \stackrel{37 \pi}{\geq} \alpha(1-\theta)\left(\ell\left(x_{k} ; 0\right)-\ell\left(x_{k} ; U_{S} t_{k}^{S}\right)\right)-\frac{L_{S}}{2} \alpha^{2}\left\|t_{k}^{S}\right\|^{2} \\
& \stackrel{\sqrt[35]{35}}{>} \frac{1}{2}\left(\alpha(1-\theta) \lambda_{S}\left\|t_{k}^{S}\right\|^{2}-L_{S} \alpha^{2}\left\|t_{k}^{S}\right\|^{2}\right) \\
& =\frac{\alpha}{2}\left((1-\theta) \lambda_{S}-L_{S} \alpha\right)\left\|t_{k}^{S}\right\|^{2} \geq 0, \tag{38}
\end{align*}
\]
if \((1-\theta) \lambda_{S}-L_{S} \alpha \geq 0\). Therefore, we observe that if \(\alpha\) satisfies \(0<\alpha \leq(1-\theta) \frac{\lambda_{S}}{L_{S}}\), then \(\alpha\) also satisfies the backtracking line search step of FCD (i.e., a function value reduction is achieved). Suppose that any \(\alpha\) that is rejected by the line search is halved for the next line search trial. Then, it is guaranteed that the \(\alpha\) that is accepted satisfies (33).

We now show that (34) holds. By the previous arguments, the line search condition 22 is guaranteed to be satisfied for some step size \(\alpha\). Then,
\[
\begin{align*}
F\left(x_{k}\right)-F\left(x_{k+1}\right) & \stackrel{\sqrt{22})+\sqrt{37}}{\geq} \theta \alpha\left(\ell\left(x_{k} ; 0\right)-\ell\left(x_{k} ; U_{S} t_{k}^{S}\right)\right) \\
\stackrel{\stackrel{35}{3}}{>} & \frac{\alpha \theta \lambda_{S}}{2}\left\|t_{k}^{S}\right\|^{2} \geq \frac{\alpha_{S} \theta \lambda_{S}}{2}\left\|t_{k}^{S}\right\|^{2} \tag{39}
\end{align*}
\]

Using the definition of \(\alpha_{S}\) in (33) gives the result.
The following lemma bounds the norm of the direction \(t_{k}^{S}\) in terms of \(g_{S}\left(x_{k}, 0\right)\). This proof closely follows that of [3, Theorem 3.1].

Lemma 2 Let Assumptions 2, 3. 4 and 5 hold, and let \(\theta \in(0,1)\). Given \(x_{k}\), let \(S\), \(t_{k}^{S} \neq 0\) and \(\alpha\) be generated by \(F C D\) (Algorithm 1) with \(\eta_{k}^{S} \in[0,1\) ). Then
\[
\begin{equation*}
\left\|t_{k}^{S}\right\| \geq \gamma_{S}\left\|g_{S}\left(x_{k} ; 0\right)\right\|, \quad \text { where } \quad \gamma_{s}:=\frac{1-\eta_{k}^{S}}{1+2 \Lambda_{S}} \tag{40}
\end{equation*}
\]

Moreover,
\[
\begin{equation*}
F\left(x_{k}\right)-F\left(x_{k}+\alpha U_{S} t_{k}^{S}\right)>\theta(1-\theta) \frac{\lambda_{S}^{2}}{4 L_{S}} \gamma_{S}^{2}\left\|g_{S}\left(x_{k} ; 0\right)\right\|^{2} \tag{41}
\end{equation*}
\]

Proof From the termination condition (21) we have that \(\left\|g_{S}\left(x_{k} ; t_{k}^{S}\right)\right\| \leq \eta_{k}^{S}\left\|g_{S}\left(x_{k} ; 0\right)\right\|\). Using the reverse triangular inequality and the fact that \(\operatorname{prox}_{\Psi_{S}^{*}}(\cdot)\) is nonexpansive, we have that
\[
\begin{aligned}
\left(1-\eta_{k}^{S}\right)\left\|g_{S}\left(x_{k} ; 0\right)\right\| \leq & \left\|g_{S}\left(x_{k} ; 0\right)\right\|-\left\|g_{S}\left(x_{k} ; t_{k}^{S}\right)\right\| \\
\leq & \left\|g_{S}\left(x_{k} ; t_{k}^{S}\right)-g_{S}\left(x_{k} ; 0\right)\right\| \\
\stackrel{17}{=} & \| H_{k}^{S} t_{k}^{S}+\operatorname{prox}_{\Psi_{S}^{*}}\left(x_{k}^{S}+t_{k}^{S}-\left(\nabla_{S} f\left(x_{k}\right)+H_{k}^{S} t_{k}^{S}\right)\right) \\
& -\operatorname{prox}_{\Psi_{S}^{*}}\left(x_{k}^{S}-\nabla_{S} f\left(x_{k}\right)\right) \| \\
& \\
\leq & \left\|H_{k}^{S} t_{k}^{S}\right\|+\left\|\left(I-H_{k}^{S}\right) t_{k}^{S}\right\| \\
\leq & \left(1+2\left\|H_{k}^{S}\right\|\right)\left\|t_{k}^{S}\right\| \\
& \\
\leq & \left(1+2 \Lambda_{S}\right)\left\|t_{k}^{S}\right\| .
\end{aligned}
\]

Rearranging gives (40), and combining (40) with (34) gives (41).

\section*{5 Technical results}

In this section we will present several technical results that will be necessary when establishing our main iteration complexity results for FCD. The first result (Theorem 6) is taken from [30] and will be used to finish all our iteration complexity results for FCD. The second result (Lemma 3) provides an upper bound on the decrease in the model that is obtained when an inexact search direction is used in FCD. The final two results provide upper bounds on the difference of function values \(F\left(x_{k+1}\right)-F\left(x_{k}\right)\) (Lemma (4) and the expected distance to the solution \(\mathbf{E}\left[F\left(x_{k+1}\right)-F^{*}\right]\) (Lemma 5).

The following Theorem will be used at the end of all our proofs to obtain iteration complexity results for FCD. (It will be used with \(\xi_{k}=F\left(x_{k}\right)-F^{*}\).)
Theorem 6 (Theorem 1 in [30]) Fix \(\xi_{0} \in \mathbf{R}^{N}\) and let \(\left\{\xi_{k}\right\}_{k \geq 0}\) be a sequence of random vectors in \(\mathbf{R}^{N}\) with \(\xi_{k+1}\) depending on \(\xi_{k}\) only. Let \(\left\{\xi_{k}\right\}_{k \geq 0}\) be a nonnegative non increasing sequence of random variables and assume that it has one of the following properties:
(i) \(\mathbf{E}\left[\xi_{k+1} \mid \xi_{k}\right] \leq \xi_{k}-\frac{\xi_{k}^{2}}{c_{1}}\) for all \(k\), where \(c_{1}>0\) is a constant,
(ii) \(\mathbf{E}\left[\xi_{k+1} \mid \xi_{k}\right] \leq\left(1-\frac{1}{c_{2}}\right) \xi_{k}\) for all \(k\) such that \(\xi_{k} \geq \epsilon\), where \(c_{2}>1\) is a constant.

Choose accuracy level \(0<\epsilon<\xi_{0}\), confidence level \(\rho \in(0,1)\). If (i) holds and we choose \(\epsilon<c_{1}\) and \(K \geq \frac{c_{1}}{\epsilon}\left(1+\ln \frac{1}{\rho}\right)+2-\frac{c_{1}}{\xi_{0}}\), or if property (ii) holds and we choose \(K \geq c_{2} \ln \frac{\xi_{0}}{\epsilon \rho}\), then \(\mathbf{P}\left(\xi_{K} \leq \epsilon\right) \geq 1-\rho\).

The following result gives an upper bound on the decrease in the quadratic model when an inexact update is used in FCD.

Lemma 3 Let Assumptions 2, 3, and 4 hold, let \(\theta \in(0,1)\), and let \(t_{*}^{S}\) denote the exact minimizer of subproblem (19). Given \(x_{k}\), let \(S, t_{k}^{S} \neq 0\) and \(\alpha\) be generated by FCD (Algorithm 1) with \(\eta_{k}^{S} \in[0,1\) ), and let \(x_{k+1}=x_{k}+\alpha U_{S} t_{k}^{S}\). Then
\[
Q\left(x_{k} ; U_{S} t_{k}^{S}\right)-Q\left(x_{k} ; 0\right) \leq Q\left(x_{k} ; U_{S} t_{*}^{S}\right)-Q\left(x_{k} ; 0\right)+\frac{2\left(\eta_{k}^{S}\right)^{2} L_{S}}{\theta(1-\theta) \lambda_{S}^{3} \gamma_{S}^{2}}\left(F\left(x_{k}\right)-F\left(x_{k+1}\right)\right)
\]
where \(\eta_{k}^{S}, \theta\) and \(\gamma_{S}\) are defined in (21), 22) and 40, respectively.

Proof We have that
\[
\begin{equation*}
Q\left(x_{k} ; U_{S} t_{k}^{S}\right)-Q\left(x_{k} ; 0\right)=\left(Q\left(x_{k} ; U_{S} t_{k}^{S}\right)-Q\left(x_{k} ; U_{S} t_{*}^{S}\right)\right)+\left(Q\left(x_{k} ; U_{S} t_{*}^{S}\right)-Q\left(x_{k} ; 0\right)\right) . \tag{42}
\end{equation*}
\]

The term \(Q\left(x_{k} ; U_{S} t_{k}^{S}\right)-Q\left(x_{k} ; U_{S} t_{*}^{S}\right)\) can be bounded using strong convexity of \(Q_{S}\) (see [2, p.460]). That is, for every \(u \in \partial Q_{S}\left(x_{k} ; t_{k}^{S}\right)\) :
\[
\begin{align*}
Q\left(x_{k} ; U_{S} t_{k}^{S}\right)-Q\left(x_{k} ; U_{S} t_{*}^{S}\right) & \stackrel{\boxed{15}}{=} Q_{S}\left(x_{k} ; t_{k}^{S}\right)-Q_{S}\left(x_{k} ; t_{*}^{S}\right) \leq \frac{1}{2 \lambda_{S}}\|u\|^{2}  \tag{43}\\
& \leq \frac{1}{2 \lambda_{S}}\left\|u+g_{S}\left(x_{k} ; t_{k}^{S}\right)-g_{S}\left(x_{k} ; t_{k}^{S}\right)\right\|^{2} \\
& \leq \frac{1}{2 \lambda_{S}}\left(\left\|u-g_{S}\left(x_{k} ; t_{k}^{S}\right)\right\|^{2}+\left\|g_{S}\left(x_{k} ; t_{k}^{S}\right)\right\|^{2}\right) .
\end{align*}
\]

Therefore, using termination condition (21), where \(\Omega\) is defined in (18), we have that
\[
\begin{aligned}
Q\left(x_{k} ; U_{S} t_{k}^{S}\right)-Q\left(x_{k} ; U_{S} t_{*}^{S}\right) & \leq \frac{1}{2 \lambda_{S}}\left(\inf _{u \in \Omega}\left\|u-g_{S}\left(x_{k} ; t_{k}^{S}\right)\right\|^{2}+\left\|g_{S}\left(x_{k} ; t_{k}^{S}\right)\right\|^{2}\right) \\
& \leq \frac{1}{2 \lambda_{S}}\left(\eta_{k}^{S}\left\|g_{S}\left(x_{k} ; 0\right)\right\|\right)^{2}
\end{aligned}
\]

Using Lemma 2 in the previous inequality, followed by (42), gives the result.
In the next Lemma we establish an upper bound on the difference between consecutive function values obtained using FCD.

Lemma 4 Let Assumptions 2, 3, 4 and 5 hold. Let \(\theta \in(0,1)\), fix \(\tilde{\eta} \in[0,1)\), and let \(t_{*}^{S}\) denote the exact minimizer of subproblem 19). Given \(x_{k}\), let \(S, t_{k}^{S} \neq 0\) and \(\alpha\) be generated by FCD (Algorithm 1) with \(\eta_{k}^{S} \in[0, \tilde{\eta}]\), and let \(x_{k+1}=x_{k}+\alpha U_{S} t_{k}^{S}\). Then
\[
\begin{equation*}
F\left(x_{k+1}\right)-F\left(x_{k}\right) \leq \chi(\tilde{\eta})\left(Q\left(x_{k} ; U_{S} t_{*}^{S}\right)-Q\left(x_{k} ; 0\right)\right), \tag{44}
\end{equation*}
\]
where
\[
\begin{equation*}
\chi(\tilde{\eta})=\min _{S \in 2^{\left[{ }^{N}\right]}}:|S|=\tau \quad \frac{\theta(1-\theta) \lambda_{S}^{3} \gamma_{S}^{2}}{2 L_{S}\left(\tilde{\eta}^{2}+\lambda_{S} \gamma_{S}^{2}\left(L_{S}-(1-\theta) \lambda_{S}\right)\right)} . \tag{45}
\end{equation*}
\]

Proof Using block Lipschitz continuity of \(f(10)\), and the definition of \(Q(13)\) we get
\[
\begin{aligned}
& F\left(x_{k+1}\right) \leq F\left(x_{k}\right)+\left(Q\left(x_{k} ; \alpha U_{S} t_{k}^{S}\right)-Q\left(x_{k} ; 0\right)\right)-\frac{\alpha^{2}}{2}\left\|t_{k}^{S}\right\|_{H_{k}^{S}}^{2}+\frac{\alpha^{2} L_{S}}{2}\left\|t_{k}^{S}\right\|^{2} \\
& \quad \stackrel{322}{\leq} F\left(x_{k}\right)+\left(Q\left(x_{k} ; \alpha U_{S} t_{k}^{S}\right)-Q\left(x_{k} ; 0\right)\right)+\frac{\alpha^{2}}{2}\left(\frac{L_{S}}{\lambda_{S}}-1\right)\left\|t_{k}^{S}\right\|_{H_{k}^{S}}^{2} .
\end{aligned}
\]

Using convexity of \(Q\) with respect to \(\alpha\) gives
\[
\begin{align*}
F\left(x_{k+1}\right) & \leq F\left(x_{k}\right)+\alpha\left(Q\left(x_{k} ; U_{S} t_{k}^{S}\right)-Q\left(x_{k} ; 0\right)\right)+\frac{\alpha^{2}}{2}\left(\frac{L_{S}}{\lambda_{S}}-1\right)\left\|t_{k}^{S}\right\|_{H_{k}^{S}}^{2} \\
& \leq F\left(x_{k}\right)+\alpha_{S}\left(Q\left(x_{k} ; U_{S} t_{k}^{S}\right)-Q\left(x_{k} ; 0\right)\right)+\frac{\alpha^{2}}{2}\left(\frac{L_{S}}{\lambda_{S}}-1\right)\left\|t_{k}^{S}\right\|_{H_{k}^{S}}^{2}, \tag{46}
\end{align*}
\]
where the second inequality holds because \(Q\left(x_{k} ; U_{S} t_{k}^{S}\right)-Q\left(x_{k} ; 0\right)<0\) (20), and \(\alpha \geq \alpha_{S}\) (33). By rearranging 39) we get \(\frac{\alpha}{2}\left\|t_{k}^{S}\right\|_{H_{k}^{S}}^{2} \leq \frac{1}{\theta}\left(F\left(x_{k}\right)-F\left(x_{k+1}\right)\right)\). Using this in 46) gives
\[
F\left(x_{k+1}\right) \leq F\left(x_{k}\right)+\alpha_{S}\left(Q\left(x_{k} ; U_{S} t_{k}^{S}\right)-Q\left(x_{k} ; 0\right)\right)+\frac{\alpha}{\theta}\left(\frac{L_{S}}{\lambda_{S}}-1\right)\left(F\left(x_{k}\right)-F\left(x_{k+1}\right)\right)
\]

By Assumption 5, \(L_{S} \geq \lambda_{S}\), so we can replace \(\alpha \in(0,1]\) with 1 to get
\[
\begin{equation*}
F\left(x_{k+1}\right) \leq F\left(x_{k}\right)+\alpha_{S}\left(Q\left(x_{k} ; U_{S} t_{k}^{S}\right)-Q\left(x_{k} ; 0\right)\right)+\frac{1}{\theta}\left(\frac{L_{S}}{\lambda_{S}}-1\right)\left(F\left(x_{k}\right)-F\left(x_{k+1}\right)\right) . \tag{47}
\end{equation*}
\]

Using Lemma 3 and the definition of \(\tilde{\eta}\) we get
\[
\begin{aligned}
F\left(x_{k+1}\right) & \leq F\left(x_{k}\right)+\alpha_{S}\left(Q\left(x_{k} ; U_{S} t_{*}^{S}\right)-Q\left(x_{k} ; 0\right)\right)+\frac{1}{\theta}\left(\frac{2 \alpha_{S} \tilde{\eta}^{2} L_{S}}{(1-\theta) \lambda_{S}^{3} \gamma_{S}^{2}}+\frac{L_{S}}{\lambda_{S}}-1\right)\left(F\left(x_{k}\right)-F\left(x_{k+1}\right)\right) \\
& \stackrel{331}{\leq} F\left(x_{k}\right)+\alpha_{S}\left(Q\left(x_{k} ; U_{S} t_{*}^{S}\right)-Q\left(x_{k} ; 0\right)\right)+\frac{1}{\theta}\left(\frac{\tilde{\eta}^{2}}{\lambda_{S}^{2} \gamma_{S}^{2}}+\frac{L_{S}}{\lambda_{S}}-1\right)\left(F\left(x_{k}\right)-F\left(x_{k+1}\right)\right) .
\end{aligned}
\]

Rearranging the previous, using the definitions of \(\chi(\tilde{\eta})\) and \(\alpha_{S}\), gives the result.

Remark 7 Notice that when exact directions are used in FCD (i.e., \(\tilde{\eta}=0\) ), we have
\[
\begin{equation*}
\chi(0)=\min _{S \in 2^{[N]}:|S|=\tau} \frac{\theta(1-\theta) \lambda_{S}^{2}}{2 L_{S}^{2}-2(1-\theta) \lambda_{S} L_{S}} \tag{48}
\end{equation*}
\]
i.e., in general \(\chi(\tilde{\eta})\) depends cubically on \(\lambda_{S}\) (45), but in the case of exact directions, \(\chi(0)\) only depends upon \(\lambda_{S}^{2}\). However, if \(L_{S}=\lambda_{S}\), but inexact directions are used ( \(\left.\tilde{\eta} \neq 0\right)\), then
\[
\begin{equation*}
\chi(\tilde{\eta})=\min _{S \in 2^{[N]}:|S|=\tau} \frac{1-\theta}{2} \frac{\theta \lambda_{S}^{2} \gamma_{S}^{2}}{\tilde{\eta}^{2}+\theta \lambda_{S}^{2} \gamma_{S}^{2}} \tag{49}
\end{equation*}
\]
which again shows a dependence upon \(\lambda_{S}^{2}\). Moreover, if \(L_{S}=\lambda_{S}\), then \(\chi(0)=(1-\theta) / 2\) (i.e., constant). This is summarized in the following table.
\begin{tabular}{|c|c|c|}
\hline & Inexact Directions \((\tilde{\eta} \neq 0)\) & Exact Directions \((\tilde{\eta}=0)\) \\
\hline\(\lambda_{S} \neq L_{S}\) & cubic & squared \\
\hline\(\lambda_{S}=L_{S}\) & squared & constant \\
\hline
\end{tabular}

Table 1: Dependence of \(\chi(\tilde{\eta})\) 45) upon \(\lambda_{S}\) (Assumption 3).

The final result in this section gives an upper bound on the expected distance between the current function value and the optimal value.
Lemma 5 Let Assumptions 2, 3. 4 and 5 hold. Let \(\theta \in(0,1)\), fix \(\tilde{\eta} \in[0,1)\), and let \(\mu_{f} \geq 0\) be the strong convexity constant of \(f\). Given \(x_{k}\), let \(S, t_{k}^{S} \neq 0\) and \(\alpha\) be generated by FCD (Algorithm 1) with \(\eta_{k}^{S} \in[0, \tilde{\eta}]\), and let \(x_{k+1}=x_{k}+\alpha U_{S} t_{k}^{S}\). Then
\[
\begin{equation*}
\mathbf{E}\left[F\left(x_{k+1}\right)-F^{*} \mid x_{k}\right] \leq\left(1-\frac{\tau \chi(\tilde{\eta}) \zeta}{N}\right)\left(F\left(x_{k}\right)-F^{*}\right)-\left(\mu_{f}-\Lambda_{\max } \zeta\right) \frac{\tau \chi(\tilde{\eta}) \zeta}{2 N}\left\|x_{k}-x_{*}\right\|^{2} \tag{50}
\end{equation*}
\]
where \(\zeta \in[0,1], \chi(\tilde{\eta})\) is defined in 45 and
\[
\begin{equation*}
\Lambda_{\max }:=\max _{S \in 2^{[N]}:|S|=\tau} \Lambda_{S} \tag{51}
\end{equation*}
\]

Proof The assumptions of Lemma 4 are satisfied, so (44) holds. Notice that \(t_{*}^{S}\) is the vector that makes the difference \(Q\left(x_{k} ; U_{S} t_{*}^{S}\right)-Q\left(x_{k} ; 0\right)\), as small/negative as possible over the subspace defined by the columns of \(U_{S}\). Therefore, for any vector \(y_{k}^{S} \neq t_{*}^{S}, Q\left(x_{k} ; U_{S} t_{*}^{S}\right)-Q\left(x_{k} ; 0\right) \leq Q\left(x_{k} ; U_{S} y_{k}^{S}\right)-Q\left(x_{k} ; 0\right)\). Choosing \(y_{k}^{S}=\zeta\left(x_{*}^{S}-x_{k}^{S}\right)\) for any \(\zeta \in[0,1]\) gives
\[
\begin{aligned}
F\left(x_{k+1}\right)-F\left(x_{k}\right) & \leq \chi(\tilde{\eta})\left(Q\left(x_{k} ; \zeta U_{S}\left(x_{*}^{S}-x_{k}^{S}\right)\right)-Q\left(x_{k} ; 0\right)\right) \\
& \stackrel{113}{\leq} \chi(\tilde{\eta})\left(\zeta\left(\nabla f\left(x_{k}\right), U_{S}\left(x_{*}^{S}-x_{k}^{S}\right)\right\rangle+\frac{\zeta^{2}}{2}\left\|x_{*}^{S}-x_{k}^{S}\right\|_{H_{k}^{S}}^{2}\right. \\
& \left.+\Psi_{S}\left(x_{k}^{S}+\zeta\left(x_{*}^{S}-x_{k}^{S}\right)\right)-\Psi_{S}\left(x_{k}^{S}\right)\right)
\end{aligned}
\]

Using convexity of \(\Psi,\left(32\right.\) and the definition \(\Lambda_{\max }\) we get
\[
\begin{align*}
F\left(x_{k+1}\right)-F\left(x_{k}\right) & \leq \chi(\tilde{\eta}) \zeta\left(\left\langle\nabla f\left(x_{k}\right), U_{S}\left(x_{*}^{S}-x_{k}^{S}\right)\right\rangle+\Psi_{S}\left(x_{*}^{S}\right)-\Psi_{S}\left(x_{k}^{S}\right)\right)  \tag{52}\\
& +\frac{\chi(\tilde{\eta}) \Lambda_{\max } \zeta^{2}}{2}\left\|x_{*}^{S}-x_{k}^{S}\right\|^{2}
\end{align*}
\]

Taking expectation of 52 and using convexity of \(f\) gives
\[
\begin{aligned}
& \mathbf{E}\left[F\left(x_{k+1}\right)-F\left(x_{k}\right) \mid x_{k}\right] \stackrel{\sqrt{25}}{\leq} \frac{\tau \chi(\tilde{\eta}) \zeta}{N}\left(\left\langle\nabla f\left(x_{k}\right), x_{*}-x_{k}\right\rangle+\Psi\left(x_{*}\right)-\Psi\left(x_{k}\right)\right) \\
&+\frac{\tau \chi \Lambda_{\max }}{N} \frac{\zeta^{2}}{2}\left\|x_{k}-x_{*}\right\|^{2} \\
& \leq-\frac{\tau \chi(\tilde{\eta}) \zeta}{N}\left(F\left(x_{k}\right)-F^{*}\right)-\frac{\tau \chi(\tilde{\eta}) \zeta}{N} \frac{\mu_{f}}{2}\left\|x_{k}-x_{*}\right\|_{2}^{2} \\
&+\frac{\tau \chi(\tilde{\eta}) \Lambda_{\max }}{N} \frac{\zeta^{2}}{2}\left\|x_{k}-x_{*}\right\|^{2}
\end{aligned}
\]

Rearranging and subtracting \(F^{*}\) from both sides gives the result.

\section*{6 Iteration complexity results for inexact FCD}

In this section we present iteration complexity results for FCD applied to problem (1) when an inexact update is used.

\subsection*{6.1 Iteration complexity in the convex case}

Here we present iteration complexity results for FCD in the convex case.
Theorem 7 Let \(F\) be the convex function defined in (1) and let Assumptions 2. 3. 4 and 5 hold. Choose an initial point \(x_{0} \in \mathbf{R}^{N}\), a target confidence \(\rho \in(0,1)\), fix \(\tilde{\eta} \in[0,1)\), and recall that \(\chi(\tilde{\eta})\) is defined in (45). Let the target accuracy \(\epsilon\) and iteration counter \(K\) be chosen in either of the following two ways:
(i) Let \(m_{1}:=\max \left\{\mathcal{R}^{2}\left(x_{0}\right), F\left(x_{0}\right)-F^{*}\right\}, 0<\epsilon<F\left(x_{0}\right)-F^{*}\) and
\[
\begin{equation*}
K \geq \frac{2 N}{\tau \chi(\tilde{\eta})} \frac{m_{1}}{\epsilon}\left(1+\ln \frac{1}{\rho}\right)+2-\frac{2 N}{\tau \chi(\tilde{\eta})} \frac{m_{1}}{\left(F\left(x_{0}\right)-F^{*}\right)} \tag{53}
\end{equation*}
\]
(ii) Let \(0<\epsilon<\min \left\{\mathcal{R}^{2}\left(x_{0}\right), F\left(x_{0}\right)-F^{*}\right\}\) and
\[
\begin{equation*}
K \geq \frac{2 N}{\tau \chi(\tilde{\eta})} \frac{\mathcal{R}^{2}\left(x_{0}\right)}{\epsilon} \ln \frac{F\left(x_{0}\right)-F^{*}}{\epsilon \rho} . \tag{54}
\end{equation*}
\]

If \(\left\{x_{k}\right\}_{k \geq 0}\) are the random points generated by FCD (Algorithm 1 ), with \(\eta_{k}^{S} \in[0, \tilde{\eta}]\) for all \(k\) and all \(S \in 2^{[N]}\), then \(\mathbf{P}\left(F\left(x_{K}\right)-F^{*} \leq \epsilon\right) \geq 1-\rho\).

Proof By Lemma 1 we have that \(F\left(x_{k}\right) \leq F\left(x_{0}\right)\) for all \(k\), so \(\left\|x_{k}-x_{*}\right\| \leq \mathcal{R}\left(x_{0}\right)\) for all \(x_{*} \in X^{*}\), where \(\mathcal{R}\left(x_{0}\right)\) is defined in 11). Then, using Lemma 5 with \(\mu_{f}=0\) gives
\[
\begin{equation*}
\mathbf{E}\left[F\left(x_{k+1}\right)-F^{*} \mid x_{k}\right] \leq\left(1-\frac{\tau \chi(\tilde{\eta}) \zeta}{N}\right)\left(F\left(x_{k}\right)-F^{*}\right)+\frac{\tau \chi(\tilde{\eta}) \Lambda_{\max } \zeta^{2}}{2 N} \mathcal{R}^{2}\left(x_{0}\right) \tag{55}
\end{equation*}
\]

Minimizing the right hand side of the previous with respect to \(\zeta\) gives: \(\zeta=\min \left\{\frac{F\left(x_{k}\right)-F^{*}}{\Lambda_{\max } \mathcal{R}^{2}\left(x_{0}\right)}, 1\right\}\). Then, (55) becomes
\[
\mathbf{E}\left[F\left(x_{k+1}\right)-F^{*} \mid x_{k}\right] \leq \begin{cases}\left(1-\frac{\tau \chi(\tilde{\eta})}{2 N} \frac{F\left(x_{k}\right)-F^{*}}{\Lambda_{\max } \mathcal{R}^{2}\left(x_{0}\right)}\right)\left(F\left(x_{k}\right)-F^{*}\right) & \text { if } \frac{F\left(x_{k}\right)-F^{*}}{\Lambda_{\max } \mathcal{R}^{2}\left(x_{0}\right)} \leq 1 \\ \left(1-\frac{\tau \chi(\tilde{\eta})}{2 N}\right)\left(F\left(x_{k}\right)-F^{*}\right) & \text { otherwise }\end{cases}
\]
or alternatively,
\[
\begin{equation*}
\mathbf{E}\left[F\left(x_{k+1}\right)-F^{*} \mid x_{k}\right] \leq\left(1-\frac{\tau \chi(\tilde{\eta})}{2 N} \min \left\{\frac{F\left(x_{k}\right)-F^{*}}{\Lambda_{\max } \mathcal{R}^{2}\left(x_{0}\right)}, 1\right\}\right)\left(F\left(x_{k}\right)-F^{*}\right) \tag{56}
\end{equation*}
\]

It is enough to study 56 under the condition \(F\left(x_{k}\right)-F^{*} \leq \Lambda_{\max } \mathcal{R}^{2}\left(x_{0}\right){ }^{3}\) Hence, we have the following two cases.
(i) Letting \(c_{1}=\frac{2 N}{\tau \chi(\tilde{\eta})} \max \left\{\Lambda_{\max } \mathcal{R}^{2}\left(x_{0}\right), F\left(x_{0}\right)-F^{*}\right\}\), 56 becomes
\[
\begin{aligned}
\mathbf{E}\left[F\left(x_{k+1}\right)-F^{*} \mid x_{k}\right] & \leq\left(1-\frac{\tau \chi(\tilde{\eta})}{2 N} \frac{F\left(x_{k}\right)-F^{*}}{\Lambda_{\max } \mathcal{R}^{2}\left(x_{0}\right)}\right)\left(F\left(x_{k}\right)-F^{*}\right) \\
& \leq\left(1-\frac{F\left(x_{k}\right)-F^{*}}{c_{1}}\right)\left(F\left(x_{k}\right)-F^{*}\right) .
\end{aligned}
\]

Notice that for this choice of \(c_{1}, \epsilon<F\left(x_{0}\right)-F^{*}<c_{1}\), so it suffices to apply Theorem \(\sqrt[6]{ }\) (i) and the result follows.
(ii) Now we choose \(c_{2}=\frac{2 N}{\tau \chi(\tilde{\eta})} \frac{\mathcal{R}^{2}\left(x_{0}\right)}{\epsilon}>1\). Observe that, whenever \(\epsilon \leq F\left(x_{k}\right)-F^{*}\), by 56 we have \(\mathbf{E}\left[F\left(x_{k+1}\right)-F^{*} \mid x_{k}\right] \leq\left(1-\frac{1}{c_{2}}\right)\left(F\left(x_{k}\right)-F^{*}\right)\). It remains to apply Theorem 6 (ii), and the result follows.

\footnotetext{
\({ }^{3}\) Notice that, at any particular iteration of FCD, either of the two options inside the 'min' in 56) could be the smallest. However, \(\Lambda_{\max } \mathcal{R}^{2}\left(x_{0}\right)\) is fixed throughout the iterations, and \(F\left(x_{k}\right)-F^{*}\) is becoming smaller as the iterations progress by Lemma 1 Therefore, eventually it will be the case that the first term inside the ' min ' in 56 will be the smallest, i.e., \(F\left(x_{k}\right)-F^{*} \leq \Lambda_{\max } \mathcal{R}^{2}\left(x_{0}\right)\).
}

\subsection*{6.2 Iteration complexity in the strongly convex case}

In this section we establish iteration complexity results for FCD when the objective function \(F\), and the smooth function \(f\), are strongly convex, with strong convexity parameters \(\mu_{f}>0\) and \(\mu_{F}>0\), respectively. Moreover, \(\mu_{f} \leq \mu_{F}\).

Before presenting our iteration complexity results, we make the following assumption. It explains that at least one of the chosen matrices \(\left\{H_{k}^{S}\right\}_{k \geq 0}\) has an eigenvalue greater than \(\mu_{f}\).
Assumption 8 We assume that \(\Lambda_{\max } \geq \mu_{f}>0\), where \(\Lambda_{\max }\) is defined in 51 .
We also define the following variable, which appears in the complexity results:
\[
\delta:= \begin{cases}\frac{\mu_{f}+\mu_{F}}{4 \Lambda_{\max }}\left(1+\frac{\mu_{f}}{\mu_{F}}\right) & \text { if } \mu_{F}+\mu_{f}<2 \Lambda_{\max }  \tag{57}\\ 1-\frac{\Lambda_{\max }-\mu_{f}}{\mu_{F}} & \text { otherwise }\end{cases}
\]

It is straightforward to show that \(\delta \in(0,1]\). In particular, if \(\mu_{F}+\mu_{f}<2 \Lambda_{\max }\), the conclusion follows because \(\mu_{F} \geq \mu_{f} \Rightarrow 1+\frac{\mu_{f}}{\mu_{F}} \leq 2\). On the other hand, if \(2 \Lambda_{\max } \leq \mu_{F}+\mu_{f}\), then by inspection, we have \(\left(\mu_{F}+\mu_{f}-\Lambda_{\max }\right) / \mu_{F} \in(0,1]\).

Theorem 9 Let \(f\) and \(F\) be the strongly convex functions defined in (1) with \(\mu_{f}>0\) and \(\mu_{F}>0\) respectively, and let Assumptions 2, 3, 4, 5 and 8 hold. Choose an initial point \(x_{0} \in \mathbf{R}^{N}\), a target confidence \(\rho \in(0,1)\), a target accuracy \(\epsilon>0\), and fix \(\tilde{\eta} \in[0,1)\). Let \(\chi(\tilde{\eta})\) and \(\delta\) be defined in 45) and (57) respectively, and let
\[
\begin{equation*}
K \geq \frac{N}{\tau \chi(\tilde{\eta}) \delta} \ln \left(\frac{F\left(x_{0}\right)-F^{*}}{\epsilon \rho}\right) \tag{58}
\end{equation*}
\]

If \(\left\{x_{k}\right\}_{k \geq 0}\) are the random points generated by FCD (Algorithm 1), with \(\eta_{k}^{S} \in[0, \tilde{\eta}]\) for all \(k\) and all \(S \in 2^{[N]}\), then \(\mathbf{P}\left(F\left(x_{K}\right)-F^{*} \leq \epsilon\right) \geq 1-\rho\).

Proof Define \(\xi_{k}:=F\left(x_{k}\right)-F^{*}\). From Lemma 5we have that
\[
\begin{align*}
\mathbf{E}\left[\xi_{k+1} \mid x_{k}\right] & \leq\left(1-\frac{\tau \chi(\tilde{\eta}) \zeta}{N}\right) \xi_{k}+\left(\Lambda_{\max } \zeta-\mu_{f}\right) \frac{\tau \chi(\tilde{\eta}) \zeta}{2 N}\left\|x_{k}-x_{*}\right\|^{2} \\
& \leq\left(1-\frac{\tau \chi(\tilde{\eta}) \zeta}{N}\right) \xi_{k}+\left(\Lambda_{\max } \zeta-\mu_{f}\right) \frac{\tau \chi(\tilde{\eta}) \zeta}{N \mu_{F}} \xi_{k} . \tag{59}
\end{align*}
\]

Notice that (2) can only be applied in (59) when \(\mu_{f} \leq \Lambda_{\max } \zeta\). But, differentiating the right hand side of (59) with respect to \(\zeta\) gives \(\zeta=\min \left\{\left(\mu_{F}+\mu_{f}\right) / 2 \Lambda_{\text {max }}, 1\right\}{ }_{\square}^{4}\) Then
\[
\mathbf{E}\left[\xi_{k+1} \mid x_{k}\right] \leq \begin{cases}\left(1-\frac{\tau \chi(\tilde{\eta})}{N} \frac{\left(\mu_{F}+\mu_{f}\right)}{4 \Lambda_{\max }}\left(1+\frac{\mu_{f}}{\mu_{F}}\right)\right) \xi_{k} & \text { if } \frac{\mu_{F}+\mu_{f}}{2 \Lambda_{\max }}<1 \\ \left(1-\frac{\tau \chi(\tilde{\eta})}{N}\left(1-\frac{\Lambda_{\max }-\mu_{f}}{\mu_{F}}\right)\right) \xi_{k} & \text { otherwise }\end{cases}
\]

Now, using (57), we can write \(\mathbf{E}\left[\xi_{k+1} \mid x_{k}\right] \leq\left(1-\frac{\tau \chi(\tilde{\eta}) \delta}{N}\right) \xi_{k}\). The result follows by applying Theorem 6 (ii) with \(c_{2}=\frac{N}{\tau \chi(\tilde{\eta}) \delta}>1\).

\section*{7 Iteration complexity of FCD applied to smooth functions}

The iteration complexity results presented in Section 6 simplify in the smooth case. Therefore, in this section we assume that \(\Psi=0\) so that \(F \equiv f\). In the smooth case, the quadratic model becomes
\[
\begin{equation*}
Q\left(x_{k} ; U_{S} t_{k}^{S}\right) \equiv Q_{S}\left(x_{k} ; t_{k}^{S}\right)=\left\langle\nabla_{S} f\left(x_{k}\right), t_{k}^{S}\right\rangle+\frac{1}{2}\left\langle H_{k}^{S} t_{k}^{S}, t_{k}^{S}\right\rangle \tag{60}
\end{equation*}
\]

Notice that the stationarity conditions also simplify in the following way
\[
\begin{equation*}
g_{S}\left(x_{k} ; t_{k}^{S}\right)=\nabla_{S} f\left(x_{k}\right)+H_{k}^{S} t_{k}^{S} \tag{61}
\end{equation*}
\]

\footnotetext{
\({ }^{4}\) We could simply have set \(\zeta=1\) initially, so that \(\mu_{f} \leq \Lambda_{\max } \zeta\) is satisfied by Assumption 8 However, we obtain a better complexity result by taking a smaller \(\zeta\).
}
so that minimizing the (smooth) quadratic model 60), is equivalent to solving the linear system
\[
\begin{equation*}
H_{k}^{S} t_{k}^{S}=-\nabla_{S} f\left(x_{k}\right) . \tag{62}
\end{equation*}
\]

The matrix \(H_{k}^{S}\) is always symmetric and positive definite (by definition), so an obvious approach is to solve (62) using the Conjugate Gradient method (CG). It is possible to use other iterative methods to approximately minimize 60 /solve \(\sqrt{622}\). However, we will see that using CG gives better iteration complexity guarantees, so, for the results presented in this section, we assume that CG is always used to solve (62).

We have the following theorem, which allows us to determine the value of the function \(Q_{S}\left(x_{k} ; U_{S} t_{k}^{S}\right)\) when an inexact direction \(t_{k}^{S}\) is used.
Lemma 6 (Lemma 7 in [16]) Let \(A z=b\), where \(A\) is a symmetric and positive definite matrix. Furthermore, let us assume that this system is solved using CG approximately; CG is terminated prematurely at the \(j\) th iteration. Then if \(C G\) is initialized with the zero solution, the approximate solution \(z_{j}\) satisfies \(z_{j}^{T} A z_{j}=z_{j}^{T} b\).

Further, note that for smooth functions, the FCD stopping condition (21) becomes
\[
\begin{equation*}
\frac{\left\|g_{S}\left(x ; t_{k}^{S}\right)\right\|_{2}}{\left\|g_{S}(x ; 0)\right\|_{2}} \equiv \frac{\left\|\nabla_{S} f\left(x_{k}\right)+H_{k}^{S} t_{k}^{S}\right\|_{2}}{\left\|\nabla_{S} f\left(x_{k}\right)\right\|_{2}} \leq \eta_{k}^{S} . \tag{63}
\end{equation*}
\]

Therefore, in this section we will make the following assumption.
Assumption 10 We assume that if FCD (Algorithm 1) is applied to problem (1) with \(\Psi=0\), then, for all \(k\) and all \(S \in 2^{[N]}\), the inexact search direction \(t_{k}^{S}\) is generated by applying \(j\) iterations of \(C G\) to (62), initialized with the zero vector, until the stopping condition (63) is satisfied, where \(\eta_{k}^{S} \in[0,1)\).

Let \(t_{k}^{S}\) be the inexact solution obtained by applying CG to (62) starting from the zero vector, and terminating according to 63 ), (i.e., let Assumption 10 hold). Then, by Lemma 6
\[
\begin{equation*}
\left\langle H_{k}^{S} t_{k}^{S}, t_{k}^{S}\right\rangle=-\left\langle\nabla_{S} f\left(x_{k}\right), t_{k}^{S}\right\rangle . \tag{64}
\end{equation*}
\]
so that
\[
\begin{equation*}
Q_{S}\left(x_{k} ; t_{k}^{S}\right)=\frac{1}{2}\left\langle\nabla_{S} f\left(x_{k}\right), t_{k}^{S}\right\rangle \equiv-\frac{1}{2}\left\langle H_{k}^{S} t_{k}^{S}, t_{k}^{S}\right\rangle<0 . \tag{65}
\end{equation*}
\]

Therefore, combining (65) with the fact that \(Q\left(x_{k} ; 0\right)=0\) in the smooth case, the stopping condition \(Q\left(x_{k} ; U_{S} t_{k}^{S}\right)\left(=Q_{S}\left(x_{k} ; t_{k}^{S}\right)\right)<Q\left(x_{k} ; 0\right)\) (i.e., 20 \()\) is always satisfied.

\subsection*{7.1 Technical Results when \(\Psi=0\)}

We begin by presenting several technical results that will be needed when establishing iteration complexity results for FCD in the smooth ( \(\Psi=0\) ) case. The first result is analogous to Lemma 2 for smooth functions, while the second is similar to Lemma 1

Lemma 7 Let Assumptions 国 3, 因 5 and 10 hold. Given \(x_{k}\), let \(S\), and \(t_{k}^{S} \neq 0\) be generated by FCD (Algorithm 1), applied to problem (1) with \(\Psi=0\), and with \(\eta_{k}^{S} \in[0,1)\). Then
\[
\begin{equation*}
\frac{1-\eta_{k}^{S}}{\Lambda_{S}}\left\|\nabla_{S} f\left(x_{k}\right)\right\|_{2} \leq\left\|t_{k}^{S}\right\|_{2} . \tag{66}
\end{equation*}
\]

Proof This proof closely follows that of [3, Theorem 3.1], and of Lemma 2. The termination condition (63) is \(\left\|g_{S}\left(x_{k} ; t_{k}^{S}\right)\right\|_{2} \leq \eta_{k}^{S}\left\|g_{S}\left(x_{k} ; 0\right)\right\|_{2}\). Then
\[
\begin{aligned}
\left(1-\eta_{k}^{S}\right)\left\|g_{S}\left(x_{k} ; 0\right)\right\|_{2} & \leq\left\|g_{S}\left(x_{k} ; 0\right)\right\|_{2}-\left\|g_{S}\left(x_{k} ; t_{k}^{S}\right)\right\|_{2} \\
& \leq\left\|g_{S}\left(x_{k} ; t_{k}^{S}\right)-g_{S}\left(x_{k} ; 0\right)\right\|_{2} \\
& \stackrel{\sqrt{61}}{=}\left\|H_{k}^{S} t_{k}^{S}\right\|_{2} \leq \Lambda_{S}\left\|t_{k}^{S}\right\|_{2} .
\end{aligned}
\]

It remains to recall that \(\nabla_{S} f\left(x_{k}\right) \equiv g_{S}\left(x_{k} ; 0\right)\).

Lemma 8 Let Assumptions 2, 3, 4. 5 and 10 hold, let \(\theta \in(0,1 / 2)\), and fix \(\tilde{\eta} \in[0,1)\). Given \(x_{k}\), let \(S\), \(t_{k}^{S} \neq 0\) be generated by FCD (Algorithm 1) applied to problem (1) with \(\Psi=0\), and with \(\eta_{k}^{S} \in[0, \tilde{\eta}]\). Then Step 6 of FCD will accept a step-size \(\alpha\) that satisfies
\[
\begin{equation*}
\alpha \geq \frac{\theta}{2} \frac{\lambda_{S}}{L_{S}}>0 \tag{67}
\end{equation*}
\]

Moreover,
\[
\begin{equation*}
f\left(x_{k}\right)-f\left(x_{k}+\alpha U_{S} t_{k}^{S}\right) \geq \frac{\vartheta(\tilde{\eta})}{2}\left\|\nabla_{S} f\left(x_{k}\right)\right\|_{2}^{2} \tag{68}
\end{equation*}
\]
where
\[
\begin{equation*}
\vartheta(\tilde{\eta}):=\min _{S:|S|=\tau} \frac{\theta \lambda_{S}^{2}}{L_{S}} \frac{(1-\tilde{\eta})^{2}}{\Lambda_{S}^{2}} \tag{69}
\end{equation*}
\]

Proof The proof follows that of Lemma 9 in 16. By Lipschitz continuity of the gradient of \(f\) :
\[
\begin{aligned}
f\left(x_{k}+\alpha U_{S} t_{k}^{S}\right) & \leq f\left(x_{k}\right)+\alpha\left\langle\nabla_{S} f\left(x_{k}\right), t_{k}^{S}\right\rangle+\frac{L_{S} \alpha^{2}}{2}\left\|t_{k}^{S}\right\|_{2}^{2} \\
& \leq f\left(x_{k}\right)-\alpha\left\|t_{k}^{S}\right\|_{H_{k}^{S}}^{2}+\frac{L_{S} \alpha^{2}}{2 \lambda_{S}}\left\|t_{k}^{S}\right\|_{H_{k}^{S}}^{2} .
\end{aligned}
\]

The right hand side of the above inequality is minimized for \(\alpha^{*}=\lambda_{S} / L_{S}\), which gives \(f\left(x_{k}+\alpha^{*} U_{S} t_{k}^{S}\right) \leq\) \(f\left(x_{k}\right)-\frac{\lambda_{S}}{2 L_{S}}\left\|t_{k}^{S}\right\|_{H_{k}^{S}}^{2}\). Observe that this step-size satisfies the backtracking line-search condition 22) because \(f\left(x_{k}+\alpha^{*} U_{S} t_{k}^{S}\right) \leq f\left(x_{k}\right)-\frac{1}{2} \frac{\lambda_{S}}{L_{S}}\left\|t_{k}^{S}\right\|_{H_{k}^{S}}^{2}<f\left(x_{k}\right)-\theta \frac{\lambda_{S}}{L_{S}}\left\|t_{k}^{S}\right\|_{H_{k}^{S}}^{2}\). Suppose that any \(\alpha\) that is rejected by the line search is halved for the next line search trial. Then, it is guaranteed that the \(\alpha\) that is accepted in Step 6 of FCD, satisfies (67). This, combined with Lemma 7 , results in the following decrease in the objective function:
\[
\begin{equation*}
f\left(x_{k}\right)-f\left(x_{k}+\alpha U_{S} t_{k}^{S}\right) \geq \frac{\theta \lambda_{S}}{2 L_{S}}\left\|t_{k}^{S}\right\|_{H_{k}^{S}}^{2} \geq \frac{\theta \lambda_{S}^{2}}{2 L_{S}}\left\|t_{k}^{S}\right\|_{2}^{2} \geq \frac{\theta \lambda_{S}^{2}}{2 L_{S}} \frac{\left(1-\eta_{k}^{S}\right)^{2}}{\Lambda_{S}^{2}}\left\|\nabla_{S} f\left(x_{k}\right)\right\|_{2}^{2} \tag{70}
\end{equation*}
\]

Using the definition (69) in (70) gives 68).
Remark 8 The quantity \(\vartheta(\tilde{\eta})\) in 69 will appear in the complexity result for FCD in the smooth case. Notice that, while \(\chi(\tilde{\eta})\) depends upon \(\lambda_{S}^{3}, \vartheta(\tilde{\eta})\) only depends upon \(\lambda_{S}^{2}\). Interestingly, the dependence of \(\vartheta(\tilde{\eta})\) on \(\lambda_{S}\) is unaffected by an inexact search direction.
\begin{tabular}{|c|c|c|}
\hline & \(\lambda_{S}=L_{S}\) & \(\lambda_{S}=\Lambda_{S}\) \\
\hline\(\vartheta(\tilde{\eta})\) & \(\min _{S:|S|=\tau} \frac{\theta(1-\tilde{\eta})^{2} L_{S}}{\Lambda_{S}^{2}}\) & \(\min _{S:|S|=\tau} \frac{\theta(1-\tilde{\eta})^{2}}{L_{S}}\) \\
\hline
\end{tabular}

\subsection*{7.2 Iteration Complexity in the Convex Smooth Case}

We are now ready to present iteration complexity results for FCD applied to smooth convex functions of the form (1) when \(\Psi=0\).
Theorem 11 Let Assumptions 2, 3. 4, 5 and 10 hold. Choose an initial point \(x_{0} \in \mathbf{R}^{N}\), a target confidence \(\rho \in(0,1)\), accuracy \(0<\epsilon<\max \left\{\mathcal{R}^{2}\left(x_{0}\right), f\left(x_{0}\right)-f^{*}\right\}\), fix \(\tilde{\eta} \in[0,1)\). Let \(\vartheta(\tilde{\eta})\) be as defined in (69), with \(\theta \in(0,1 / 2)\) and let
\[
\begin{equation*}
K \geq \frac{2 N \mathcal{R}^{2}\left(x_{0}\right)}{\tau \vartheta(\eta) \epsilon}\left(1+\ln \frac{1}{\rho}\right)+2-\frac{2 N \mathcal{R}^{2}\left(x_{0}\right)}{\tau \vartheta(\eta)\left(f\left(x_{0}\right)-f^{*}\right)} \tag{71}
\end{equation*}
\]

If \(\left\{x_{k}\right\}_{k \geq 0}\) are the random points generated by FCD (Algorithm 1) applied to problem (1) with \(\Psi=0\), where \(\eta_{k}^{S} \in[0,1)\) satisfies \(\eta_{k}^{S} /\left(1-\eta_{k}^{S}\right)^{2} \leq \eta<1\) for all \(k\) and all \(S \in 2^{[N]}\), then \(\mathbf{P}\left(F\left(x_{K}\right)-F^{*} \leq \epsilon\right) \geq 1-\rho\).
Proof By Lemma \(8 f\left(x_{k}\right) \leq f\left(x_{0}\right)\) for all \(k\), and by convexity of \(f\), we have \(f\left(x_{k}\right)-f^{*} \leq \max _{x^{*} \in X^{*}}\left\langle\nabla f\left(x_{k}\right), x_{k}-\right.\) \(\left.x^{*}\right\rangle \leq\left\|\nabla f\left(x_{k}\right)\right\|_{2} \mathcal{\mathcal { R }}\left(x_{0}\right)\). Taking expectation of (70), and applying the previous, gives:
\[
f\left(x_{k}\right)-\mathbf{E}\left[f\left(x_{k}+\alpha U_{S} t_{k}^{S}\right) \mid x_{k}\right] \geq \frac{\vartheta(\tilde{\eta}) \tau}{2 N}\left\|f\left(x_{k}\right)\right\|_{2}^{2} \geq \frac{\vartheta(\tilde{\eta}) \tau}{2 N}\left(\frac{f\left(x_{k}\right)-f^{*}}{\mathcal{R}\left(x_{0}\right)}\right)^{2}
\]

Rearranging and applying Theorem6(i) with \(c_{1}=2 N \mathcal{R}^{2}\left(x_{0}\right) / \tau \vartheta(\tilde{\eta})\), gives the result.

\subsection*{7.3 Iteration Complexity in the Strongly Convex Smooth Case}

We now present iteration complexity results for FCD applied to smooth, strongly convex functions of the form (11), when \(\Psi=0\).

Theorem 12 Let \(\Psi=0\), let \(f=F\) be the strongly convex function defined in (1) with \(\mu_{f}>0\) and let Assumptions 2, 3, 4. 5 and 10 hold. Choose an initial point \(x_{0} \in \mathbf{R}^{N}\), a target confidence \(\rho \in(0,1)\), a target accuracy \(\epsilon>0\), and fix \(\tilde{\eta} \in[0,1)\). Let \(\vartheta(\tilde{\eta})\) be as defined in 69), with \(\theta \in(0,1 / 2)\) and let
\[
\begin{equation*}
K \geq \frac{N}{\tau \vartheta(\tilde{\eta}) \mu_{f}} \ln \left(\frac{F\left(x_{0}\right)-F^{*}}{\epsilon \rho}\right) \tag{72}
\end{equation*}
\]

If \(\left\{x_{k}\right\}_{k \geq 0}\) are the random points generated by FCD (Algorithm 1) applied to problem (1) with \(\Psi=0\), where \(\eta_{k}^{S} \in[0, \tilde{\eta}]\), then \(\mathbf{P}\left(F\left(x_{K}\right)-F^{*} \leq \epsilon\right) \geq 1-\rho\).

Proof Notice that, by strong convexity, \(f\left(x_{k}\right)-f^{*} \leq \frac{1}{2 \mu_{f}}\left\|\nabla f\left(x_{k}\right)\right\|_{2}^{2}\), (see, e.g., [2, p.460]). Taking the expectation of 68), and using the previous, gives:
\[
\mathbf{E}\left[f\left(x_{k}\right)-f\left(x_{k}+\alpha_{S} U_{S} t_{k}^{S}\right) \mid x_{k}\right] \geq \frac{\tau \vartheta(\tilde{\eta})}{2 N}\left\|\nabla f\left(x_{k}\right)\right\|_{2}^{2} \geq \frac{\tau \vartheta(\tilde{\eta}) \mu_{f}}{N}\left(f\left(x_{k}\right)-f^{*}\right)
\]

Rearranging, and applying Theorem 6 (ii) with \(c_{2}=N /\left(\mu_{f} \tau \vartheta(\tilde{\eta})\right)>1\) gives the result.

\section*{8 Discussion/Comparison}

In this section we discuss the complexity results derived in this paper, and compare them with the current state of the art results.

\subsection*{8.1 General Framework}

The FCD framework has been designed to be extremely flexible. Moreover, this framework generalizes several existing algorithms.

To see this, suppose that FCD is set-up in the following way. Choose \(|S|=\tau=1, H_{k}^{i}=L_{i}\) (where \(L_{i}\) is the Lipschitz constant for the \(i\) th coordinate), suppose coordinate \(i\) is selected with uniform probability and suppose that (19) is solved exactly at every iteration. Then the update generated by FCD is given by
\[
t_{k}^{i}=\arg \min _{t \in \mathbf{R}}\left\{\left\langle\nabla_{i} f(x), t\right\rangle+\frac{L_{i}}{2} t^{2}+\Psi_{i}\left(x^{i}+t\right)\right\}
\]

However, this is equivalent to the update generated by UCDC (see Algorithm 1 and equation (17)) in 30. So FCD generalizes UCDC, (and accepts a step-size \(\alpha=1\) for all iterations in this case, i.e., no line-search is needed by FCD.)

Now suppose that we assume that \(\Psi=0\), and that \(f\) is strongly convex. For simplicity, let us assume that \(f\) is quadratic. Moreover, suppose that FCD is set-up in the following way. Choose \(|S|=\tau=1\), \(H_{k}^{i}=\nabla_{i i}^{2} f\) (for a quadratic function the Hessian \(\nabla^{2} f(\cdot)\) is independent of \(x_{k}\) ), suppose coordinate \(i\) is selected with uniform probability and suppose that 19 is solved exactly at every iteration. Then the update generated by FCD is \(t_{k}^{i}=-\left(\nabla_{i} f\left(x_{k}\right)\right) /\left(\nabla_{i i}^{2} f\right)\). However, this is equivalent to the update generated by SDNA in [28] (Methods \(1-3\) in [28] are equivalent when \(\tau=1\) ). So FCD generalizes SDNA, (and also accepts a step-size \(\alpha=1\) for all iterations in this case, i.e., no line-search is needed by FCD.)

Clearly, a central advantage of FCD is it's flexibility. It recovers several existing state-of-the-art algorithms, depending on the particular choice of algorithm parameters. We stress that, for UCDC, one must always compute and use the Lipschitz constants, while for FCD, \(H_{k}^{S}\) can be any positive semidefinite matrix. Further, SDNA only applies to problems that are smooth and strongly convex, whereas FCD can be applied to general convex problems of the form (1).

\subsection*{8.2 Practicality vs Iteration Complexity}

We remark that it is expected that the complexity results for FCD may be slightly worse than for other existing methods. There are several reasons for this:
1. FCD is a general framework that enables substantial user flexibility, as well as computational practicality. The user may choose the block size (i.e., the number of coordinates to be updated at each iteration) the matrix \(H_{k}^{S}\). Therefore, the user has full flexibility regarding how expensive each iteration of FCD is. For example, choosing \(\tau\) small, and \(H_{k}^{S}\) to be diagonal, will lead to cheap iterations. Choosing a larger block size and/or a full matrix \(H_{k}^{S}\), will lead to iterations that are more expensive. However, it is anticipated that fewer iterations may be needed when \(H_{k}^{S}\) is a good approximation to a principal minor of the Hessian. Because of the user friendly, general and flexible framework, the complexity rates for FCD may be slightly worse than for methods that apply to specific problem instances, with specially tuned parameters.
2. Another difference between FCD, and other current methods, is that many methods do not incorporate any second order information, or they enforce that \(H_{k}^{S}=L_{S} I\). The restrictions made in other methods are usually imposed to ensure that their surrogate function/model, overapproximates \(F\left(x_{k}+U_{S} t_{k}^{S}\right)\), so that minimizing their model will result in a reduction in the original function value. The assumption that FCD makes regarding \(H_{k}^{S}\) is much weaker (any symmetric positive definite \(H_{k}^{S}\) is allowed), which means that FCD requires a line search, to guarantee a reduction in the objective function value at every iteration, and ultimately, convergence. We prefer an algorithm with fewer assumptions, so that it is widely applicable and practical, but we pay the price of a (potentially) slightly weaker complexity result.
3. FCD is one of the (very) few coordinate descent methods that allow inexact updates. It is expected that inexact updates will lead to cheaper iterations (it is intuitive that solving a subproblem inexactly, will usually require less computational effort than solving it exactly), although this may result in slightly more iterations compared with an 'exact' method. We stress that, despite the potential increase in the number of iterations, it is still expected that inexact updates will result in a lower computational run time overall.

We stress that, despite the slightly weaker iteration complexity results for FCD, the algorithm works extremely well in practice, and often outperforms algorithms with 'better' complexity rates, (see the numerical experiments in Section 91. Moreover, the user has some control over the complexity results for FCD through their choice of parameters. For example, \(\lambda_{S}\) and \(\Lambda_{S}\) (Assumption 3) are user defined parameters, which appear in the complexity rates for FCD (see 45) and (69), so the complexity rate for FCD can be larger or smaller, depending on how the user chooses these parameters.

\subsection*{8.3 Comparison of complexity results}

We will now compare the complexity results obtained in this work for FCD, with those of UCDC in 30] and PCDM 31.

\subsection*{8.3.1 Comparison of complexity results with \(U C D C\) [30]}

UCDC is a serial method that allows exact updates only, so in this section we assume that \(|S|=\tau=1\), and that FCD uses only exact updates \((\tilde{\eta} \equiv 0)\). This means that we have \(L_{i}\) for \(i=1, \ldots, N\), and \(H_{k}^{i}\) is a scalar, for all \(k\). The complexity results in [30] use 'scaled' quantities that depend on the vector of Lipschitz constants \(\left(L=\left(L_{1}, \ldots, L_{N}\right)\right)\) while the quantities in this work depend on the 'unscaled' 2-norm (i.e., \(e=(1, \ldots, 1)) \cdot{ }^{5}\) Moreover, in [30], the authors suppose that strong convexity can come from \(f\) or \(\Psi\), or both. Here, to allow easy comparison, we suppose that \(\mu_{\Psi}=0\), which means that, in this section, if strong convexity is assumed, then \(\mu_{f}=\mu_{F}\). To this end, define the following constants, which depend a 'scaling' vector \(w \in \mathbf{R}_{++}^{N}\) :
\[
\begin{equation*}
c_{1}(w)=\frac{2 N}{\epsilon} \max \left\{\mathcal{R}_{w}^{2}\left(x_{0}\right), F\left(x_{0}\right)-F^{*}\right\}, \text { and } c_{2}(w)=\frac{2 N \mathcal{R}_{w}^{2}\left(x_{0}\right)}{\epsilon} . \tag{73}
\end{equation*}
\]

\footnotetext{
\({ }^{5}\) Using the notation of 30, we have \(n=N\). Further, although UCDC allows arbitrary probabilities in the smooth case, we restrict our attention to uniform probabilities only, so \(N\) appears in the C-S and SC-S results in Table 2
}

Table 2 presents the complexity results obtained for UCDC 30, and the complexity results obtained in this work for FCD. The following notation is used in the table. Constants \(\xi_{0}=F\left(x_{0}\right)-F^{*}, \mu_{\phi}(w)\) denotes the strong convexity parameter of function \(\phi\) with respect to a ' \(w\)-weighted' norm, \(\mu=\left(\mu_{f}+\mu_{\Psi}\right) /\left(1+\mu_{\Psi}\right)\) and \(\mathcal{R}_{w}\left(x_{0}\right)\) is defined in 12).
\begin{tabular}{|c|c|c|c|}
\hline\(F\) & UCDC [30] & FCD [this paper] & Theorem \\
\hline C-N & \(c_{1}(L)\left(1+\log \frac{1}{\rho}-\frac{\epsilon}{\xi_{0}}\right)+2\) & \(\frac{c_{1}(e)}{\chi(0)}\left(1+\log \frac{1}{\rho}-\frac{\epsilon}{\xi_{0}}\right)+2\) & \(7(\mathrm{i})\) \\
\cline { 2 - 4 } \(\mathrm{C}-\mathrm{N}\) & \(c_{2}(L) \log \left(\frac{\xi_{0}}{\epsilon \rho}\right)\) & \(\frac{c_{2}(e)}{\chi(0)} \log \left(\frac{\xi_{0}}{\epsilon \rho}\right)\) & 7 (ii) \\
\hline \(\mathrm{SC}-\mathrm{N}\) & \(\frac{N}{\mu_{f}(L)} \log \left(\frac{\xi_{0}}{\epsilon \rho}\right)\) & \(\frac{N}{\chi(0) \delta} \log \left(\frac{\xi_{0}}{\epsilon \rho}\right)\) & 9 \\
\hline \(\mathrm{C}-\mathrm{S}\) & \(c_{2}(L) \log \left(\frac{\xi_{0}}{\epsilon \rho}\right)\) & \(\frac{c_{2}(e)}{\vartheta(0)} \log \left(\frac{\xi_{0}}{\epsilon_{\rho}}\right)\) & 11 \\
\hline \(\mathrm{SC}-\mathrm{S}\) & \(\frac{N}{\mu_{f}(L)} \log \left(\frac{\xi_{0}}{\epsilon \rho}\right)\) & \(\frac{N}{\vartheta(0) \mu_{f}} \log \left(\frac{\xi_{0}}{\epsilon \rho}\right)\) & 12 \\
\hline
\end{tabular}

Table 2: Comparison of the iteration complexity results for coordinate descent methods using an inexact update and using an exact update ( \(\mathrm{C}=\) Convex, \(\mathrm{SC}=\) Strongly Convex, \(\mathrm{N}=\) Nonsmooth, \(\mathrm{S}=\) Smooth \()\).

Case \(C\) - \(N(i)\). In this case, the difference between the complexity rates for UCDC and FCD appears in the constants \(c_{1}(L)\) and \(c_{1}(e) / \chi(0)\). Clearly, if \(\mathcal{R}_{L}^{2}\left(x_{0}\right) \leq \xi_{0}^{F}\) then \(c_{1}(L)=c_{1}(e)\), so that the the complexity rate of FCD is simply \(1 / \chi(0)\) times worse than UCDC. (Recall Table 1 for a comparison of \(\chi(\cdot)\) values.) On the other hand, if \(\mathcal{R}_{L}^{2}\left(x_{0}\right)>\xi_{0}^{F}\), then it is difficult to compare the complexity rates directly. Notice that the dependence of FCD on the Lipschitz constants is explicit (the constants \(L_{i}\) appear in \(\chi(0)\) ), whereas the dependence of UCDC on the Lipschitz constants is implicit in the weighted term \(\mathcal{R}_{L}^{2}\left(x_{0}\right)\). However, consider the case when \(L_{i}=L_{j}\) for all \(i, j\). Then \(\|y-x\|_{L}^{2}=\sum_{i=1}^{N} L_{i}\left\|y_{i}-x_{i}\right\|_{2}^{2}=L_{i}\|y-x\|_{2}^{2}\), so that \(\mathcal{R}_{L}^{2}\left(x_{0}\right)=L_{i} \mathcal{R}^{2}\left(x_{0}\right)\) in this case. Then we simply compare \(L_{i}\) vs \(1 / \chi(0)\). Now suppose we choose \(H_{k}^{i}=L_{i}\) for all \(k\), so that \(\chi(0)=(1-\theta) / 2\). If \(\theta \approx 0\), then FCD has a better complexity rate than UCDC whenever \(L_{i}>2\). Of course, in other circumstances, UCDC may have a better complexity rate than FCD - this is simply one particular example.

Case \(C\) - \(N(i i)\). Here we compare \(c_{2}(L)\) and \(c_{2}(\mathbf{1})\). Both rates contain the term \(2 N / \epsilon\), so it remains to compare \(\mathcal{R}_{L}^{2}\left(x_{0}\right)\) and \(\mathcal{R}^{2}\left(x_{0}\right) / \chi(0)\). Thus, see the discussion for \(\mathrm{C}-\mathrm{N}(\mathrm{i})\).

Case SC-N. Substituting \(\mu_{f}=\mu_{F}\) into (57), shows that \(\delta=\mu_{f} / \Lambda_{\max }\), so in this case we compare \(1 / \mu_{f}(L)\) with \(\Lambda_{\max } / \chi(0) \mu_{f}\). Moreover, the Lipschitz constants appear explicitly in \(\chi(0)\) in the analysis in this work, while they are again implicit in the term \(\mu_{f}(L)\) in 30. Again, this makes it difficult to compare the rates directly. However, consider the following example, where \(\mu_{f}=\Lambda_{\max }\). Then \(\lambda_{S}=L_{S}\) (because \(\lambda_{S} \leq \Lambda_{\max }=\mu_{f} \leq L_{S}\) ), so \(\Lambda_{\max } / \chi(0) \mu_{f}=2 /(1-\theta)\). Then, if \(\theta \approx 0\), FCD has a better complexity rate than UCDC, whenever \(\mu_{f}(L) \leq 1 / 2\). (Note that it always holds that \(\mu_{f}(L) \leq 1\) [30, (12)], and this does not necessarily imply that \(\left.\left(\mu_{f}(e) \equiv\right) \mu_{f} \leq 1 / 2\right)\).

Case \(C\)-S. Similar to Case C-N(ii), we compare \(\mathcal{R}_{L}^{2}\left(x_{0}\right)\) and \(\mathcal{R}^{2}\left(x_{0}\right) / \vartheta(0)\). Suppose we have the same set up as Case C-N(i), so that we compare \(L_{i}\) with \(1 / \vartheta(0)\). Choosing \(\lambda_{S}=\Lambda_{S}\) and \(\theta \approx 1 / 2\) shows that \(1 / \vartheta(0) \approx 2 L_{i}\), so that the complexity rates for FCD is approximately twice that for UCDC.

Case \(S C\)-S. Here we compare the quantities \(1 / \mu_{f}(L)\) and \(1 / \mu_{f} \vartheta(0)\). All that can be said for UCDC is that \(1 / \mu_{f}(L) \geq 1\) because \(\mu_{f}(L) \leq 1\). Sometimes it may be the case that \(1 / \mu_{f}(L) \approx 1\), whereas in other cases, we may have \(1 / \mu_{f}(L) \gg 1\). Note that, for FCD, if we choose \(\lambda_{S}=\Lambda_{S}\), then \(1 / \mu_{f} \vartheta(0)=L_{S} / \theta \mu_{f}\). For a well conditioned problem, (i.e., \(\mu_{f} \approx L_{S}\) ), choosing \(\theta \approx 1\) gives \(1 / \mu_{f} \vartheta(0) \approx 1\). On the other hand, if \(L_{S} \gg \theta \mu_{f}\), then \(1 / \mu_{f} \vartheta(0) \gg 1\). So, while the rates for FDC and UCDC cannot be compared directly, the complexity rates are similar, when taking into account the problem conditioning.

\subsection*{8.3.2 Comparison of complexity results with PCDM [31]}

Now we compare FCD with PCDM 31. PCDM is a parallel coordinate descent method, that uses exact updates, so in this case we suppose that \(|S|=\tau \geq 1\), and that FCD also uses exact updates. Furthermore, we will suppose that the coordinates are sampled according to Definition 1 (which is called a \(\tau\)-nice sampling in [31). Define \(\beta=1+\frac{(\omega-1)(\tau-1)}{\max \{1, N-1\}}\), where \(\omega\) is a measure of the 'separability' of the objective function. Then, the iteration complexity for PCDM is presented in Table 3.
\begin{tabular}{|c|c|c|}
\hline & C-N & SC-N \\
\hline PCDM [31] & \(\frac{2 N}{\tau \epsilon} \max \left\{\beta \mathcal{R}_{L}^{2}\left(x_{0}\right), F\left(x_{0}\right)-F^{*}\right\}\left(1+\log \frac{1}{\rho}-\frac{\epsilon}{\xi_{0}}\right)+2\) & \(\frac{\beta N}{\tau \mu_{f}(L)} \log \left(\frac{\xi_{0}}{\epsilon \rho}\right)\) \\
\hline FCD & \(\frac{2 N}{\tau \epsilon \chi(0)} \max \left\{\mathcal{R}^{2}\left(x_{0}\right), F\left(x_{0}\right)-F^{*}\right\}\left(1+\log \frac{1}{\rho}-\frac{\epsilon}{\xi_{0}}\right)+2\) & \(\frac{N}{\tau \delta \chi(0)} \log \left(\frac{\xi_{0}}{\epsilon \rho}\right)\) \\
\hline
\end{tabular}

Table 3: Complexity results for PCDM 31.
\(C\) - \(N\). Note that \(\beta \geq 1\), so if \(\beta \mathcal{R}_{L}^{2}\left(x_{0}\right) \leq F\left(x_{0}\right)-F^{*}\), then the complexity rate for FCD is \(1 / \chi(0)\) times worse than for PCDM. On the other hand, suppose that \(\mathcal{R}_{L}^{2}\left(x_{0}\right)>F\left(x_{0}\right)-F^{*}\). Then we must compare the quantities \(\beta \mathcal{R}_{L}^{2}\left(x_{0}\right)\) and \(\mathcal{R}^{2}\left(x_{0}\right) / \chi(0)\). Similar arguments to the Case C-N(i) in Section 8.3.1 can be made here. However, we note that \(\beta\) depends upon \(\tau, N\) and \(\omega\). While \(\tau\) and \(N\) can be controlled (they are parameters of the algorithm), \(\omega\) depends upon the problem data, and could be very large, implying large \(\beta\).
\(S C\) - \(N\). In this case we compare the quantities \(\beta / \mu_{f}(L)\) and \(1 / \delta \chi(0)\). The case for FCD and \(1 / \delta \chi(0)\) is made in Section 8.3.1. Case SC-N. Now consider PCDM. For a well conditioned problem, and separable problem, we may have \(\mu_{f}(L) \approx 1 \approx \omega\), so that \(\beta / \mu_{f}(L) \approx 1\). However, for a poorly conditioned problem, \(\omega\) may be very large, and \(\mu_{f}(L)\) may be very small, so that \(\beta / \mu_{f}(L) \gg 1\). Again, while the rates for FDC and PCDM cannot be compared directly, the complexity rates are similar, when taking into account the problem conditioning.

\section*{9 Numerical Experiments}

In this section we compare the performance of FCD with UCDC [30] on the \(\ell_{1}\) - and \(\ell_{2}\)-regularized logistic regression problem, which has the form (11) with
\[
\begin{equation*}
f(x)=\sum_{j=1}^{m} \log \left(1+e^{-b_{j} x^{T} a_{j}}\right) \quad \text { and } \quad \Psi(x)=c\|x\|_{1} \text { or } \Psi(x)=c\|x\|_{2} \tag{74}
\end{equation*}
\]

Here \(c>0, a_{j} \in \mathbb{R}^{N} \forall j=1,2, \ldots, m\) are the training samples and \(b_{j} \in\{-1,+1\}\) are the corresponding labels. Problems of the form (74) are important in machine learning and are used for training a linear classifier \(x \in \mathbb{R}^{N}\) that separates input data into two distinct clusters; see, for example, 43. We present the performance of FCD and UCDC on six real-world large-scale data sets from the LIBSVM data collection [7]. Details of the data sets are given in Table 4, where \(A \in \mathbb{R}^{m \times N}\) is a matrix whose rows represent the training samples.

\subsection*{9.1 Implementations of FCD and UCDC}
\(F C D\). At every iteration of \(\mathrm{FCD}, \tau\) coordinates are sampled uniformly at random without replacement, i.e., Assumption 2 holds. FCD is implemented with two different choices for \(H_{k}^{S}\), as shown in Table 5 All other parameters are the same for FCD v. 1 and FCD v.2. Clearly, for FCD v.1, subproblem (19) is separable and for \(\ell_{1}\)-regularization it has the closed form solution
\[
\begin{equation*}
t_{k}^{S}=\mathcal{S}\left(x_{k}^{S}-\left(H_{k}^{S}\right)^{-1} \nabla_{S} f\left(x_{k}^{S}\right), c \operatorname{diag}\left(\left(H_{k}^{S}\right)^{-1}\right)\right) \tag{75}
\end{equation*}
\]

Table 4: Properties of the \(\ell_{1}\)-regularized logistic regression data sets considered here. Note that nnz \((A)\) denotes the number of nonzeros in matrix \(A\); consequently, the fourth column gives the (relative) sparsity of \(A\).
\begin{tabular}{|l|c|c|c|}
\hline Problem & \(m\) & \(N\) & \(\mathrm{nnz}(A) / m N\) \\
\hline kdd2010 (algebra) & \(8,407,752\) & \(20,216,830\) & \(1.79 e-6\) \\
kdd2010 (bridge to algebra) & \(19,264,097\) & \(29,890,095\) & \(9.83 e-7\) \\
news20.binary & 19,996 & \(1,355,191\) & \(3.35 e-4\) \\
rcv1.binary & 20,242 & 47,236 & \(1.16 e-3\) \\
url & \(2,396,130\) & \(3,231,961\) & \(3.57 e-5\) \\
webspam & 350,000 & \(16,609,143\) & \(2.24 e-4\) \\
\hline
\end{tabular}

Table 5: Algorithm parameters for FCD v.1, FCD v.2, UCDC v.1, UCDC v.2. For FCD v. 2 we set \(\rho=10^{-6}, L_{j}\) is the Lipschitz constant for the \(j\)-th coordinate (see 99 ), and \(\bar{L}^{S}=\sum_{j \in S} L_{j}\).
\begin{tabular}{|l|c|c|}
\hline Algorithm & \(\tau\) & \(H_{k}^{S}(\forall k, S)\) \\
\hline FCD v.1 & \(\lceil 0.001 N\rceil\) & \(\operatorname{diag}\left(\nabla_{S}^{2} f\left(x_{k}\right)\right)\), \\
FCD v.2 & \(\lceil 0.001 N\rceil\) & \(\nabla_{S}^{2} f\left(x_{k}\right)+\rho I_{N_{i}}\) \\
UCDC v.1 & 1 & \(L_{j}\) \\
UCDC v.2 & \(\lceil 0.001 N\rceil\) & \(\bar{L}^{S} I\) \\
\hline
\end{tabular}
where
\[
\begin{equation*}
\mathcal{S}(u, v)=\operatorname{sign}(u) \max (|u|-v, 0) \tag{76}
\end{equation*}
\]
is the well-known soft-thresholding operator, which is applied component wise when \(u\) and \(v\) are vectors. Since subproblem (19) is solved exactly via (75), there is no need to verify the stopping conditions (21). For \(\ell_{2}\)-regularization and FCD v. 1 subproblem \(\sqrt{19}\) has a trivial and inexpensive solution.

For FCD v.2, parameter \(\rho\) ensures that \(H_{k}^{S}\) is positive definite, and consequently, subproblem (19) is well defined \({ }^{6}\) Note that the matrix \(H_{k}^{S}\) is never explicitly formed, we only perform matrix-vector products with it in a matrix-free manner. Moreover, for FCD v.2, subproblem (19) is solved iteratively using an Orthant Wise Limited-memory Quasi-Newton (OWL) method in the case of \(\ell_{1}\)-regularization, which can be downloaded from http://www.di.ens.fr/~mschmidt/Software/L1General.html. OWL was chosen because it has been shown in [3] to result in a robust and efficient deterministic version of FCD, i.e. \(\tau=N\) (one block of size \(N\) ). While for \(\ell_{2}\)-regularization and FCD v. 2 an iterative method like Conjugate Gradients is used.
\(U C D C\). We compare FCD with the current state-of-the-art coordinate descent method, namely the Uniform Coordinate Descent for Composite functions (UCDC) method, [30, Algorithm 2]. For UCDC, the block size \(\tau\) and the decomposition of \(\mathbb{R}^{N}\) into \(\lceil N / \tau\rceil\) blocks must be fixed a-priori, and at each iteration of UCDC the blocks are selected with uniform probability. Again, UCDC is also implemented with two different choices of \(H_{k}^{S}\), as seen in Table 5 . The reasons for this set-up is as follows. For UCDC, the block Lipschitz constants are explicitly required in the algorithm, and, while the Lipschitz constants for coordinates can be computed with relative ease, the Lipschitz constants for blocks of coordinates can be far more expensive to compute. To this end, UCDC v. 1 is the case of coordinate updates where the Lipschitz constants are straightforward to compute, and UCDC v. 2 considers blocks of \(\tau\) coordinates, where \(\bar{L}^{S}:=\sum_{j \in S} L_{j}\) is used as an inexpensive overapproximation of the true block Lipschitz constant. Here, the (coordinate) Lipschitz constants for UCDC are computed via [30, Table 10]. Note that, for both UCDC v. 1 and v.2, the choice of \(H_{k}^{S}\) leads to a separable subproblem 19 , which is solved exactly using it's closed for solution.

Remark 9 Notice that Algorithm 2 in [30] is a special case of FCD where the subproblem (19) is solved exactly. A line search is unnecessary in this case because, for \(H_{k}^{S}\) as stated, 19p is an overapproximation of \(F\) along block coordinate direction \(t^{S}\), and minimizing this overapproximation exactly will lead to a decrease in the objective function; see [30].

\footnotetext{
\({ }^{6}\) Clearly, there is a trade-off when selecting \(\rho\). The larger \(\rho\) is, the smaller the condition number of \(H_{k}\), so that 19 p will be solved quickly using an iterative solver. However, if \(\rho\) is too large then \(H_{k}^{S} \approx \rho I\) so that essential second order information from \(\nabla^{2} f\left(x_{k}\right)\) may be lost.
}

\subsection*{9.2 Termination Criteria and Parameter Tuning}

In these experiments FCD and UCDC are terminated when their running time exceeds an allowed maximum. Note that, using subgradients as a measure of the distance from optimality or any other operation of similar cost are considered to be too expensive a task for large scale problems, and are therefore not appropriate here. Here, all methods were terminated after the relative error of the objective function dropped below \(10^{-8}\) or after 28 hours of wall-clock time.

Furthermore, for FCD we set parameter \(\eta_{k}^{S}\) in 21 equal to \(0.9 \forall i, k\). The maximum number of backtracking line search iterations is set to \(200, \theta=10^{-3}\) and each iteration the step-size is halved. For all instances we computed \(c\) after performing fivefold cross validation as proposed in [17. For UCDC the coordinate Lipschitz constants \(L_{j} \forall j\) are calculated once at the beginning of the algorithm and this task is included in the overall running time. Finally, all methods are initialized with the zero solution.

\subsection*{9.3 Results for \(\ell_{1}\)-regularized logistic regression}

The results of this experiment is shown in the following figures. Figure 1 shows the results of FCD and UCDC on several of the smaller data sets, and Figure 2 shows the results of FCD and UCDC on several of the larger data sets for \(\ell_{1}\)-regularized logistic regression. Note that the calculation of \(F(x)\) is included in the running time of \(F C D\) when it is required by line search. The reported running time of UCDC includes the time required to compute the Lipschitz constants at the beginning of the algorithm.


Fig. 1: Performance of FCD and UCDC on 3 large scale \(\ell_{1}\)-regularized logistic regression problems. The first row shows the function value \(\frac{F(x)-F^{*}}{F^{*}}\) vs the number of iterations, while the second row shows the function value \(\frac{F(x)-F^{*}}{F^{*}}\) vs time. Columns 1,2 and 3 correspond to data sets news20.binary, rcv1 and url, respectively. All plots are in log-scale. For some figures, i.e., Figures 1 d , 1 e and 1 f , we measured \(\frac{F(x)-F^{*}}{F^{*}}\) every 100 iterations of the algorithms, which explains the initial rapid decrease in \(\frac{F(x)-F^{*}}{F^{*}}\).

FCD performs very well on all data sets. Notice that both versions of FCD achieve a lower objective function value than UCDC in the allocated time. The subfigures corresponding to function value vs iterations demonstrate that an iteration of FCD is slightly more expensive in terms of computation time


Fig. 2: Performance of FCD and UCDC on 3 large scale \(\ell_{1}\)-regularized logistic regression problems. The first row shows the function value \(\frac{F(x)-F^{*}}{F^{*}}\) vs the number of iterations, while the second row shows the function value \(\frac{F(x)-F^{*}}{F^{*}}\) vs time. Columns 1,2 and 3 correspond to data sets webspam, 'kdda' (kdd2010 (algebra)) and 'kddb' (kdd2010 (bridge to algebra)), respectively. All plots are in log-scale. For some figures, i.e., Figures 2 d 2e and 2 f we measured \(\frac{F(x)-F^{*}}{F^{*}}\) every 100 iterations of the algorithms, which explains the initial rapid decrease in \(\frac{F(x)-F^{*}}{F^{*}}\).
than UCDC (i.e., fewer iterations of FCD are completed in the allocated time compared with UCDC). However, the expense of each iteration of FCD is offset by a greater reduction in function value. Clearly, including partial curvature information (as in FCD) is useful.

The results shown in Figure 2 are even more striking. Note that data sets kdda and kddb are very large. Despite this, FCD manages to decrease the objective function quickly. On the other hand, UCDC v. 1 barely manages to decrease the objective function at all in the allocated time, while UCDC v. 2 only does a little better. Again, this highlights the importance of partial second order information. The cost of including the curvature information is again offset by the rapid reduction in the objective function value. It is interesting to note that, on these data sets, FCD v. 1 and v. 2 perform similarly, which shows that even a little curvature information is beneficial. (Recall that the two variants only differ in the choice of \(H_{k}^{S}\); see Table 5.) Clearly, FCD outperforms UCDC on these problems.

Overall, FCD v. 1 makes rapid progress decreasing the objective value initially, whereas FCD v. 2 is better at decreasing the objective function closer to the solution. We believe that FCD v. 2 is faster for more accurate solutions due to the second-order information that is capture by using the Hessian to construct the subproblem of the algorithm.

We now comment on an important observation about the running time required by the line search for FCD. It is often believed that a line search might be a task that is too expensive for large-scale problems. Therefore, constant or diminishing step-sizes are usually preferred in order to avoid extra function evaluations, which are often required by line search techniques. However, Table 6 shows that, for the experiments corresponding to Figures 1 and 2 , the additional time required to perform the function evaluations required by line search for FCD are often negligible. In particular, notice that columns three and five report the percentage of iterations in which a step-size of \(\alpha=1\) was accepted by the line-search step in FCD. Perhaps surprisingly, for the kdd2010(algebra) dataset, FCD v. 1 and v. 2 accepted a step-size of \(\alpha=1\) for \(99.99 \%\) of iterations. At worst, FCD v. 2 applied to rcv1.binary, accepted a step size of 1 for \(81.26 \%\) of the iterations. Moreover, the second and fourth columns of Table 6 state the percentage of total
time that was spent on line search computations (including all necessary function evaluations. At worst and only for one experiment, FCD v. 1 spent \(55.38 \%\) of the total time on line search computations for the rcv1.binary problem, whereas FCD v. 2 spent only \(3.62 \%\) of the total time on line search computations for the same problem. On average, over both algorithm variants, and each of the 6 problems, \(9.91 \%\) of the total running time was spent on line search computations. We find it striking the fact that FCD v.1, which uses only diagonal Hessian information, accepts unit step-sizes for at least \(95.95 \%\) of its iterations. We believe that this reveals the power of line-search and it shows that it shouldn't be discarded easily in practice.

Remark 10 Furthermore, the cost of the line-search in terms of worst-case function evaluations is not so high. Because we use a standard line search, it is straightforward to establish that, in the worst-case, one requires \(\log \left(\alpha_{S}\right)\) function evaluations for the line search at each iteration, where \(\alpha_{S}\) is defined in (33). In particular, because \(\alpha\) decreases geometrically (for example, it is halved at every iteration) and \(\alpha_{S}\) is the minimum \(\alpha\) such that the termination criteria of line-search are satisfied, we simply have \((1 / 2)^{s}<\alpha_{S}\), which means that, in the worst case, we require \(s>\log _{2}\left(\alpha_{S}\right)\) function evaluations for each line search.

Table 6: Comparison of the line search costs for FCD for the experiments in Figures 1 and 2. Columns 2 and 4 report the percentage of the overall running time that was spent on line search calculations (including all necessary function evaluations). Columns 3 and 5 report the percentage of iterations where a step-size equal to 1 satisfied the line search termination conditions.
\begin{tabular}{|l|c|c|c|c|}
\hline \multirow{2}{*}{\multicolumn{1}{|c|}{ Problem }} & \multicolumn{2}{c|}{ FCD v.1 } & \multicolumn{2}{c|}{ FCD v.2 } \\
\cline { 2 - 5 } & \% of time & \% of iterations & \% of time & \% of iterations \\
\hline kdd2010 (algebra) & \(7.31 \%\) & \(99.99 \%\) & \(7.07 \%\) & \(99.99 \%\) \\
kdd2010 (bridge to algebra) & \(6.67 \%\) & \(99.99 \%\) & \(4.37 \%\) & \(99.99 \%\) \\
news20.binary & \(11.09 \%\) & \(99.97 \%\) & \(1.44 \%\) & \(95.95 \%\) \\
rcv1.binary & \(55.38 \%\) & \(81.26 \%\) & \(3.62 \%\) & \(99.46 \%\) \\
url & \(7.62 \%\) & \(99.92 \%\) & \(11.63 \%\) & \(99.81 \%\) \\
webspam & \(2.23 \%\) & \(99.99 \%\) & \(0.56 \%\) & \(99.89 \%\) \\
\hline
\end{tabular}

\subsection*{9.4 Results for \(\ell_{2}\)-regularized logistic regression}

Figure 3 shows the results of FCD and UCDC on several of the smaller data sets, and Figure 4 shows the results of FCD and UCDC on several of the larger data sets for \(\ell_{2}\)-regularized logistic regression. Similarly to the \(\ell_{1}\)-regularized experiments in Subsection 9.3 FCD v. 1 or v. 2 or both are faster in terms of running time compared to UCDC. However, due to strong convexity imposed by \(\ell_{2}\)-regularization we observe that FCD v. 2 is clearly better than FCD v.1, even for highly accurate solutions, despite that FCD v. 2 uses only diagonal Hessian information. Strong convexity is also the reason that the performance gap between UCDC and FCD decreased for \(\ell_{2}\)-regularized logistic regression.

Finally, Table 7 shows performance statistics for line-search in FCD for \(\ell_{2}\)-regularized logistic regression. Note that the performance of line-search is even better for \(\ell_{2}\)-regularized logistic regression compared to \(\ell_{1}\)-regularized logistic regression. We believe that the reason is strong convexity that is guaranteed by \(\ell_{2}\)-regularization.

\section*{10 Conclusion}

In this work we have presented a randomized block coordinate descent method (FCD) that can be applied to convex composite functions of the form (1). The proposed method allows the coordinates to be selected randomly via a non-fixed block structure, incorporates partial curvature information via a user defined matrix \(H_{k}^{S}\), and allows inexact updates to be used. These features make the algorithm extremely flexible. The algorithm is supported by high probability iteration complexity results. Moreover, we presented several synthetic and real world, large scale problems where FCD is shown to outperform the current state-of-the-art coordinate descent method.


Fig. 3: Performance of FCD and UCDC on 3 large scale \(\ell_{2}\)-regularized logistic regression problems. The first row shows the function value \(\frac{F(x)-F^{*}}{F^{*}}\) vs the number of iterations, while the second row shows the function value \(\frac{F(x)-F^{*}}{F^{*}}\) vs time. Columns 1,2 and 3 correspond to data sets news20.binary, rcv1 and url, respectively. All plots are in log-scale. For some figures, i.e., Figures 3 d , 3e and 3 f , we measured \(\frac{F(x)-F^{*}}{F^{*}}\) every 100 iterations of the algorithms, which explains the initial rapid decrease in \(\frac{F(x)-F^{*}}{F^{*}}\).

Table 7: Comparison of the line search costs for FCD for the experiments in Figures 3 and 4 . Columns 2 and 4 report the percentage of the overall running time that was spent on line search calculations (including all necessary function evaluations). Columns 3 and 5 report the percentage of iterations where a step-size equal to 1 satisfied the line search termination conditions.
\begin{tabular}{|l|c|c|c|c|}
\hline \multirow{2}{*}{\multicolumn{1}{|c|}{ Problem }} & \multicolumn{2}{|c|}{ FCD v.1 } & \multicolumn{2}{c|}{ FCD v.2 } \\
\cline { 2 - 5 } & \% of time & \% of iterations & \% of time & \% of iterations \\
\hline kdd2010 (algebra) & \(7.32 \%\) & \(99.99 \%\) & \(0.96 \%\) & \(99.99 \%\) \\
kdd2010 (bridge to algebra) & \(6.77 \%\) & \(99.99 \%\) & \(1.73 \%\) & \(99.99 \%\) \\
news20.binary & \(11.92 \%\) & \(99.99 \%\) & \(0.74 \%\) & \(99.98 \%\) \\
rcv1.binary & \(16.93 \%\) & \(99.99 \%\) & \(2.58 \%\) & \(99.90 \%\) \\
url & \(8.55 \%\) & \(99.99 \%\) & \(3.31 \%\) & \(99.99 \%\) \\
webspam & \(3.05 \%\) & \(99.99 \%\) & \(0.36 \%\) & \(99.98 \%\) \\
\hline
\end{tabular}

\section*{11 Acknowledgments}

We would like to thank the anonymous reviewers for their helpful comments and suggestions, which led to improvements of an earlier version of this paper.

\section*{References}
1. P. Alart, O. Maisonneuve, and R. T. Rockafellar. Nonsmooth Mechanics and Analysis: Theoretical and Numerical Advances. Springer US, 2006.
2. S. Boyd and L. Vandenberghe. Convex Optimization. Cambridge University Press New York, NY, USA, 2004.
3. R. H. Byrd, J. Nocedal, and F. Oztoprak. An inexact successive quadratic approximation method for convex l-1 regularized optimization. Mathematical Programming, 157(2):375-396, 2016.
4. E. Candès. Compressive sampling. In International Congress of Mathematics, volume 3, pages 1433-1452, Madrid, Spain, 2006.


Fig. 4: Performance of FCD and UCDC on 3 large scale \(\ell_{1}\)-regularized logistic regression problems. The first row shows the function value \(\frac{F(x)-F^{*}}{F^{*}}\) vs the number of iterations, while the second row shows the function value \(\frac{F(x)-F^{*}}{F^{*}}\) vs time. Columns 1,2 and 3 correspond to data sets webspam, 'kdda' (kdd2010 (algebra)) and 'kddb' (kdd2010 (bridge to algebra)), respectively. All plots are in log-scale. For some figures, i.e., Figures 4 d , 4 e and 4 f we measured \(\frac{F(x)-F^{*}}{F^{*}}\) every 100 iterations of the algorithms, which explains the initial rapid decrease in \(\frac{F(x)-F^{*}}{F^{*}}\).
5. E. J. Candés, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. IEEE Trans. Inf. Theory, 52(2):489-509, 2006.
6. A. Cassioli, D. Di Lorenzo, and M. Sciandrone. On the convergence of inexact block coordinate descent methods for constrained optimization. European Journal of Operational Research, 231(2):274-281, 2013.
7. Chih-Chung Chang and Chih-Jen Lin. LIBSVM: A library for support vector machines. ACM Transactions on Intelligent Systems and Technology, 2:27:1-27:27, 2011. Software available at http://www.csie.ntu.edu.tw/~cjlin/ libsvm.
8. A. Daneshmand, F. Facchinei, V. Kungurtsev, and G. Scutari. Hybrid random/deterministic parallel algorithms for convex and nonconvex big data optimization. IEEE Transactions on Signal Processing, 63(15):3914-3929, 2015.
9. O. Devolder, F. Glineur, and Yu. Nesterov. Intermediate gradient methods for smooth convex problems with inexact oracle. Technical report, CORE-2013017, 2013.
10. O. Devolder, F. Glineur, and Yu. Nesterov. First-order methods of smooth convex optimization with inexact oracle. Mathematical Programming, 146(1-2):37-75, 2014.
11. D. Donoho. Compressed sensing. IEEE Trans. on Information Theory, 52(4):1289-1306, April 2006.
12. F. Facchinei, S. Sagratella, and G. Scutari. Flexible parallel algorithms for big data optimization. Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on, pages 7208-7212, May 2014.
13. F. Facchinei, G. Scutari, and S. Sagratella. Parallel selective algorithms for nonconvex big data optimization. IEEE Transactions on Signal Processing, 63(7):1874-1889, April 2015.
14. O. Fercoq and P. Richtárik. Accelerated, parallel and proximal coordinate descent. SIAM Journal of Optimization, 25(4):19972023, 2015.
15. K. Fountoulakis and J. Gondzio. Performance of first- and second-order methods for \(\ell_{1}\)-regularized least squares problems. Computational Optimization and Applications, 65(3):605-635, 2016.
16. K. Fountoulakis and J. Gondzio. A second-order method for strongly convex \(\ell_{1}\)-regularization problems. Mathematical Programming, 156(1):189-219, 2016.
17. C.-W. Hsu, C.-C. Chang, and C.-J. Lin. A practical guide to support vector classification. Technical report, Department of Computer Science, National Taiwan University, 2003.
18. S. Karimi and S. Vavasis. IMRO: a proximal quasi-Newton method for solving \(l_{1}\)-regularized least square problem. SIAM Journal on Optimization, 27(2):583615, 2017.
19. J. D. Lee, Y. Sun, and M. A. Saunders. Proximal Newton-type methods for convex optimization. Advances in Neural Information Processing Systems, pages 836-844, 2012.
20. J. D. Lee, Y. Sun, and M. A. Saunders. Proximal Newton-type methods for minimizing composite functions. SIAM Journal of Optimization, 24(3):14201443, 2014.
21. Z. Lu and L. Xiao. On the complexity analysis of randomized block-coordinate descent methods. Mathematical Programming, 152(1):615-642, 2014.
22. I. Necoara and A. Patrascu. A random coordinate descent algorithm for optimization problems with composite objective function and linear coupled constraints. Computational Optimization and Applications, 57(2):307-337, 2014.
23. Yu. Nesterov. Introductory Lectures on Convex Optimization: A Basic Course. Applied Optimization. Kluwer Academic Publishers, 2004.
24. Yu. Nesterov. Efficiency of coordinate descent methods on huge-scale optimization problems. SIAM Journal on Optimization, 22(2):341-362, 2012.
25. J. Nocedal and S. J. Wright. Numerical Optimization. Springer Series in Operations Research. Springer-Verlag, New York, 1999.
26. N. Parikh and S. Boyd. Proximal algorithms. Foundations and Trends in Optimization, 1(3):123-231, 2013.
27. Z. Qin, K. Scheinberg, and D. Goldfarb. Efficient block-coordinate descent algorithms for the group lasso. Mathematical Programming Computation, 5(2):143-169, 2013.
28. Z. Qu, P. Richtárik, M. Takáč, and O. Fercoq. SDNA: stochastic dual newton ascent for empirical risk minimization. ICML'16 Proceedings of the 33rd International Conference on International Conference on Machine Learning, 48:1823-1832, 2016.
29. M. Razaviyayn, M. Hong, and J. S. Pang Z. Q. Luo. Parallel successive convex approximation for nonsmooth nonconvex optimization. Advances in Neural Information Processing Systems, 2014.
30. P. Richtárik and M. Takáč. Iteration complexity of randomized block-coordinate descent methods for minimizing a composite function. Mathematical Programming, 144(1-2):1-38, 2014.
31. P. Richtárik and M. Takáč. Parallel coordinate descent methods for big data optimization. Mathematical Programming, pages 1-52, April 2015.
32. M. De Santis, S. Lucidi, and F. Rinaldi. A fast active set block coordinate descent algorithm for \(\ell_{1}\)-regularized least squares. SIAM Journal on Optimization, 26(1):781809, 2016.
33. K. Scheinberg and X. Tang. Practical inexact proximal quasi-Newton method with global complexity analysis. Mathematical Programming, 160(1-2):495529, 2016.
34. S. Shalev-Schwartz and T. Zhang. Stochastic dual coordinate ascent methods for regularized loss minimization. Journal of Machine Learning Research, 14:567-599, 2013.
35. N. Simon and R. Tibshirani. Standardization and the group lasso penalty. Statistica Sinica, 22(3):983, 2012.
36. R. Tappenden, P. Richtárik, and J. Gondzio. Inexact coordinate descent: Complexity and preconditioning. Journal of Optimization Theory and Applications, 170(1):144-176, 2016.
37. R. Tappenden, M. Takáč, and P. Richtárik. On the complexity of parallel coordinate descent. Optimization Methods and Software, 33(2):372-395, 2018.
38. R. Tibshirani. Regression shrinkage and selection via the lasso. Journal of the Roy. Statist. Soc., 58(1):267-288, 1996.
39. P. Tseng and S. Yun. A coordinate gradient descent method for nonsmooth separable minimization. Math. Program., Ser. B, 117:387-423, 2009.
40. Paul Tseng and Sangwoon Yun. Block-coordinate gradient descent method for linearly constrained nonsmooth separable optimization. Journal of Optimization Theory and Applications, 140(3):513-535, 2009.
41. Paul Tseng and Sangwoon Yun. A coordinate gradient descent method for linearly constrained smooth optimization and support vector machines training. Computational Optimization and Applications, 47(2):179-206, 2010.
42. S. J. Wright. Accelerated block-coordinate relaxation for regularized optimization. SIAM Journal of Optimization, 22(1):159-186, 2012.
43. G. X. Yuan, C. H. Ho, and C. J. Lin. Recent advances of large-scale linear classification. Proceedings of the IEEE, 100(9):2584-2603, 2012.```


[^0]:    K. Fountoulakis is with the International Computer Science Institute, Department of Statistics, University of California Berkeley, 1947 Center St, Ste. 600, Berkeley, CA 94704, USA. e-mail: kfount@berkeley.edu.
    R. Tappenden (corresponding author) is with the School of Mathematics and Statistics, University of Canterbury, Private Bag 8400, Christchurch 8041, NZ. e-mail: rachael.tappenden@canterbury.ac.nz.

[^1]:    ${ }^{1}$ See [31] Equation (2)] or 37, Definition 13] for a precise definition of 'partial separability'.

[^2]:    2 An earlier version of this work, which, to the best of our knowledge, was the first to propose varying random block selection, is cited by 28.

