



HAL
open science

Model identification and local linear convergence of coordinate descent

Quentin Klopfenstein, Quentin Bertrand, Alexandre Gramfort, Joseph Salmon, Samuel Vaiter

► **To cite this version:**

Quentin Klopfenstein, Quentin Bertrand, Alexandre Gramfort, Joseph Salmon, Samuel Vaiter. Model identification and local linear convergence of coordinate descent. 2020. hal-03019711

HAL Id: hal-03019711

<https://hal.science/hal-03019711v1>

Preprint submitted on 23 Nov 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Model identification and local linear convergence of coordinate descent

Quentin Klopfenstein^{*1}, Quentin Bertrand^{*2}, Alexandre Gramfort², Joseph Salmon³, and Samuel Vaiter⁴

¹*Université de Bourgogne, Institut de Mathématiques de Bourgogne, Dijon, France*

²*Université Paris-Saclay, Inria, CEA, Palaiseau, France*

³*IMAG, Université de Montpellier, CNRS, Montpellier, France*

⁴*CNRS and IMB, Université de Bourgogne, Dijon, France*

^{*}*Equal contribution*

October 16, 2020

Abstract

For composite nonsmooth optimization problems, Forward-Backward algorithm achieves model identification (*e.g.*, support identification for the Lasso) after a finite number of iterations, provided the objective function is regular enough. Results concerning coordinate descent are scarcer and model identification has only been shown for specific estimators, the support-vector machine for instance. In this work, we show that cyclic coordinate descent achieves model identification in finite time for a wide class of functions. In addition, we prove explicit local linear convergence rates for coordinate descent. Extensive experiments on various estimators and on real datasets demonstrate that these rates match well empirical results.

1 Introduction

1.1 Coordinate descent

Over the last two decades, coordinate descent (CD) algorithms have become a powerful tool to solve large scale optimization problems (Friedman et al., 2007, 2010). Many applications coming from machine learning or compressed sensing have lead to optimization problems that can be solved efficiently via CD algorithms: the Lasso (Tibshirani, 1996; Chen et al., 1998), the elastic net (Zou and Hastie, 2005) or support-vector machine (Boser et al., 1992). All the previously cited estimators are based on an optimization problem which can be

written:

$$x^* \in \arg \min_{x \in \mathbb{R}^p} \{ \Phi(x) \triangleq f(x) + \underbrace{\sum_{j=1}^p g_j(x_j)}_{\triangleq g(x)} \} , \quad (1)$$

with f a convex smooth (*i.e.*, with a Lipschitz gradient) function and g_j proper closed and convex functions. In the past twenty years, the popularity of CD algorithms has greatly increased due to the well suited structure of the new optimization problems mentioned above (*i.e.*, separability of the nonsmooth term), as well as the possible parallelization of the algorithms (Fercoq and Richtárik, 2015).

The key idea behind CD (Algorithm 1) is to solve small and simple subproblems iteratively until convergence. More formally, for a function $\Phi : \mathbb{R}^p \mapsto \mathbb{R}$, the idea is to minimize successively one dimensional functions $\Phi_{|x_j} : \mathbb{R} \mapsto \mathbb{R}$, updating only one coordinate at a time, while the others remain unchanged. There exists many variants of CD algorithms, the main branching being:

- **The index selection.** There are different ways to choose the index of the updated coordinate at each iteration. The main variants can be divided in three categories, **cyclic** CD (Tseng and Yun, 2009) when the indices are chosen in the set $[p] \triangleq \{1, \dots, p\}$ cyclically. **Random** CD (Nesterov, 2012), where the indices are chosen following a given random distribution. Finally, **greedy** CD (Nutini et al., 2015) *picks* an index, optimizing a given criterion: largest decrease of the objective function, or largest gradient norm (Gauss-Southwell rule), for instance.
- **The update rule.** There also exists several possible schemes for the coordinate update: exact minimization, coordinate gradient descent or prox-linear update (see Shi et al. 2016, Sec. 2.2 for details).

In this work, we will focus on cyclic CD with prox-linear update rule (Algorithm 1): a popular instance, *e.g.*, the one coded in popular packages such as `glmnet` (Friedman et al., 2007) or `sklearn` (Pedregosa et al., 2011).

Among the methods of coordinate selection, **random** CD has been extensively studied, especially by Nesterov (2012) for the minimization of a smooth function f . It was the first paper proving global non-asymptotic $1/k$ convergence rate in the case of a smooth and convex f . This work was later extended to composite optimization $f + \sum_j g_j$ for nonsmooth separable functions (Richtárik and Takáč, 2014; Fercoq and Richtárik, 2015). Refined convergence rates were also shown by Shalev-Shwartz and Tewari (2011); Shalev-Shwartz and Zhang (2013). These convergence results have then been extended to coordinate descent with equality constraints (Necoara and Patrascu, 2014) that induce non-separability as found in the SVM dual problem in the presence of the bias term. Different distributions have been considered for the index selection such as uniform distribution (Fercoq and Richtárik, 2015; Nesterov, 2012; Shalev-Shwartz and Tewari,

2011; Shalev-Shwartz and Zhang, 2013), importance sampling (Leventhal and Lewis, 2010; Zhang, 2004) and arbitrary sampling (Necoara and Patrascu, 2014; Qu and Richtárik, 2016a,b).

On the opposite, theory on **cyclic** coordinate descent is more fuzzy, the analysis in the cyclic case being more difficult. First, Luo and Tseng (1992); Tseng (2001); Tseng and Yun (2009); Razaviyayn et al. (2013) have shown convergence results for (block) CD algorithms for nonsmooth optimization problems (without rates¹). Then, Beck and Tetrushvili (2013) showed $1/k$ convergence rates for Lipschitz convex functions and linear convergence rates in the strongly convex case. Saha and Tewari (2013) proved $1/k$ convergence rates for composite optimization $f + \|\cdot\|_1$ under "isotonicity" condition. Sun and Hong (2015); Hong et al. (2017) have extended the latter results and showed $1/k$ convergence rates with improved constants for composite optimization $f + \sum_j g_j$. Li et al. (2017) have extended the work of Beck and Tetrushvili (2013) to the nonsmooth case and refined their convergence rates in the smooth case. Finally, as far as we know, the work by Xu and Yin (2017) is the first one tackling the problem of local linear convergence. They have proved local linear convergence under the very general Kurdyka-Lojasiewicz hypothesis, relaxing convexity assumptions. Following the line of work by Liang et al. (2014), we use a more restrictive framework (see Section 1.4) that allows to achieve finer results: model identification as well as improved local convergence results.

1.2 Model identification

Nonsmooth optimization problems coming from machine learning such as the Lasso or the support-vector machine (SVM) generally generate solutions lying onto a low-complexity model (see Definition 1 for details). For the Lasso, for example, a solution x^* has typically only a few non-zeros coefficients: it lies on the model set $T_{x^*} = \{u \in \mathbb{R}^p : \text{supp}(u) \subseteq \text{supp}(x^*)\}$, where $\text{supp}(x)$ is the support of x , *i.e.*, the set of indices corresponding to the non-zero coefficients. A question of interest in the literature is: does the algorithm achieve model identification after a finite number of iterations? Formally, does it exist $K > 0$ such that for all $k > K$, $x^{(k)} \in T_{x^*}$? For the Lasso the question boils down to "does it exist $K > 0$ such that for all $k > K$, $\text{supp}(x^{(k)}) \subseteq \text{supp}(x^*)$?" This finite time identification property is paramount for features selection (Tibshirani, 1996), but also for potential acceleration methods (Massias et al., 2018) of the CD algorithm, as well as model calibration (Bertrand et al., 2020).

Finite model identification was first proved in Bertsekas (1976) for the projected gradient method with non-negative constraints. In this case, after a finite number of steps the sparsity pattern of the iterates is the same as the sparsity pattern of the solution. It means that for k large enough, $x_i^{(k)} = 0$ for all i such that $x_i^* = 0$. Then, many other results of finite model identification have been shown in different settings and for various algorithms. For the projected

¹Note that some local rates are shown in Tseng and Yun (2009) but under some strong hypothesis.

gradient descent algorithm, identification was proved for polyhedral constraints (Burke and Moré, 1988), for general convex constraints (Wright, 1993), and even non-convex constraints (Hare and Lewis, 2004). More recently, identification was proved for proximal gradient algorithm (Mercier and Vijayasundaram, 1979; Combettes and Wajs, 2005), for the ℓ_1 regularized problem (Hare, 2011). Liang et al. (2014, 2017); Vaiter et al. (2018) have shown model identification and local linear convergence for proximal gradient descent. These results have then been extended to other popular machine learning algorithms such as SAGA, SVRG (Poon et al., 2018) and ADMM (Poon and Liang, 2019). To our knowledge, CD has not been extensively studied with a similar generality. Some identification results have been shown for CD, but only on specific models (She and Schmidt, 2017; Massias et al., 2019) or variants of CD (Wright, 2012), in general, under restrictive hypothesis. The authors are not aware of generic model identification results for CD Algorithm 1.

1.3 Notation

General notation. We write $\|\cdot\|$ the Euclidean norm on vectors. For, $x, \gamma \in \mathbb{R}^p$, the weighted norm is denoted $\|x\|_\gamma \triangleq \sqrt{\sum_{j=1}^p \gamma_j x_j^2}$. For a differentiable function $\psi : \mathbb{R}^p \mapsto \mathbb{R}^p$, at $x \in \mathbb{R}^p$, we write $\mathcal{J}\psi(x) \in \mathbb{R}^{p \times p}$ the Jacobian of ψ at x . For a set S , we denote by S^c its complement. We denote $[p] = \{1, \dots, p\}$. Let $(e_j)_{j=1}^p$ be the vectors of the canonical base of \mathbb{R}^p . We denote the coordinatewise multiplication of two vectors u and v by $u \odot v$ and by $u \odot M$ the row wise multiplication between a vector and a matrix. We denote by $\mathcal{B}(x, \epsilon)$ the ball of center x and radius ϵ . The spectral radius of a matrix M is denoted $\rho(M)$.

Convex analysis. We recall the definition of the proximity operator of a convex function g , for any $\gamma > 0$:

$$\text{prox}_{\gamma g}(x) = \arg \min_{y \in \mathbb{R}^p} \frac{1}{2\gamma} \|x - y\|^2 + g(y) .$$

Let $\mathcal{C} \subset \mathbb{R}^p$ be a convex set, $\text{aff}(\mathcal{C})$ denotes its affine hull, the smallest affine set containing \mathcal{C} , and $\text{ri}(\mathcal{C})$ denotes its relative interior (the interior of its affine hull). The indicator function of \mathcal{C} is the function defined for any $x \in \mathbb{R}^p$ by

$$\delta_{\mathcal{C}}(x) = \begin{cases} 0 & \text{if } x \in \mathcal{C} \\ +\infty & \text{otherwise} \end{cases} . \quad (2)$$

The domain of a function f is defined as $\text{dom}(f) = \{x \in \mathbb{R}^p : f(x) < +\infty\}$. For a convex function f , $\partial f(x)$ denotes its subdifferential at x and is given by $\partial f(x) = \{s \in \mathbb{R}^p : f(y) \geq f(x) + \langle s, y - x \rangle, \forall y \in \text{dom}(f)\}$. We denote by L_j the coordinatewise Lipschitz constants of $\nabla_j f$, *i.e.*, for all $x \in \mathbb{R}^p$, $h_j \in \mathbb{R}$:

$$\|\nabla_j f(x + e_j h_j) - \nabla_j f(x)\| \leq L_j |h_j| . \quad (3)$$

Coordinate descent. We denote $0 < \gamma_j \leq 1/L_j$ the local step size and $\gamma = (\gamma_1, \dots, \gamma_p)^\top$. To prove model identification we need to “keep track” of the iterates: following the notation from [Beck and Tetrushvili \(2013\)](#) coordinate descent can be written:

Algorithm 1 PROXIMAL COORDINATE DESCENT

```

input :  $\gamma_1, \dots, \gamma_p \in \mathbb{R}_+, n_{iter} \in \mathbb{N}, x^{(0)} \in \mathbb{R}^p$ 
for  $k = 0, \dots, n_{iter}$  do // index selection
     $x^{(0,k)} \leftarrow x^{(k)}$ 
    for  $j = 1, \dots, p$  do
         $x^{(j,k)} \leftarrow x^{(j-1,k)}$ 
         $x_j^{(j,k)} \leftarrow \text{prox}_{\gamma_j g_j} \left( x_j^{(j-1,k)} - \gamma_j \nabla_j f(x^{(j-1,k)}) \right)$ 
     $x^{(k+1)} \leftarrow x^{(p,k)}$ 
return  $x^{n_{iter}+1}$ 

```

1.4 Assumptions on composite problem

We consider the optimization problem defined in [Equation \(1\)](#) with the following assumptions:

Assumption 1 (Smoothness). f is a convex and differentiable function, with a Lipschitz gradient.

Assumption 2 (Proper, closed, convex). For any $j \in [p]$, g_j is proper, closed and convex.

Assumption 3 (Existence). The problem admits at least one solution:

$$\arg \min_{x \in \mathbb{R}^p} \Phi(x) \neq \emptyset . \quad (4)$$

Assumption 4 (Non degeneracy). The problem is non-degenerate: for any $x^* \in \arg \min_{x \in \mathbb{R}^p} \Phi(x)$

$$-\nabla f(x^*) \in \text{ri}(\partial g(x^*)) . \quad (5)$$

[Assumption 4](#) can be seen as a generalization of qualification constraints ([Hare and Lewis, 2007](#), Sec. 1).

1.5 Contributions

With mild assumptions on the g_j functions, for the **cyclic** proximal coordinate descent algorithm:

- We prove finite time model identification ([Th. 1](#)).
- We provide local linear convergence rates ([Th. 2](#)).
- We illustrate our results on multiple real datasets and estimators ([Section 4](#)) showing that our theoretical rates match the empirical ones.

2 Model identification for CD

As stated before, the solutions of the Lasso are structured. Using an iterative algorithm like [Algorithm 1](#) to find an approximate solution (since we stop after a finite number of iterations) brings the question of structure recovery. For the Lasso, the underlying structure, also called model ([Candès and Recht, 2012](#)), is identified by the Forward-Backward algorithm. It means that after a finite number of iterations, the iterative algorithm leads to an approximate solution that shares a similar structure than the true solution of the optimization problem ([Liang et al., 2014](#); [Vaïter et al., 2018](#); [Fadili et al., 2018](#)). For the Lasso, the underlying model is related to the notion of support: *i.e.*, the non-zero coefficients for the Lasso, and it can be generalized for the case of completely separable functions as follows:

Definition 1 (Generalized support, [Sun et al. 2019](#)). We call *generalized support* $\mathcal{S}_x \subseteq [p]$ the set of indices $j \in [p]$ where g_j is differentiable at x_j :

$$\mathcal{S}_x \triangleq \{j \in [p] : \partial g_j(x_j) \text{ is a singleton}\} . \quad (6)$$

This notion can be unified with the definition of model subspace from [Vaïter et al. \(2015, Sec. 3.1\)](#):

Definition 2 (Model subspace, [Vaïter et al. 2015](#)). We denote the model subspace at x :

$$T_x = \{u \in \mathbb{R}^p : \forall j \in \mathcal{S}_x, u_j = 0\} . \quad (7)$$

See [Lemma 1](#) in [Appendix A](#) for details.

Examples in machine learning.

The ℓ_1 norm. The function $g(x) = \sum_{i=1}^p |x_i|$ is certainly the most popular nonsmooth convex regularizer promoting sparsity. Indeed, the ℓ_1 norm generates structured solution with model subspace ([Vaïter et al., 2018](#)). We have that $\mathcal{S}_x = \{j \in [p] : x_j \neq 0\}$ since $|\cdot|$ is differentiable everywhere but not at 0, and the model subspace reads:

$$T_x = \{u \in \mathbb{R}^p : \text{supp}(u) \subseteq \text{supp}(x)\} . \quad (8)$$

The box constraints indicator function $\delta_{[0,C]}$. This indicator function appears for instance in box constrained optimization problems such as the dual problem of the SVM. Let $\mathcal{I}_x^0 = \{j \in [p] : x_j = 0\}$ and $\mathcal{I}_x^C = \{j \in [p] : x_j = C\}$, then

$$T_x = \{u \in \mathbb{R}^p : \mathcal{I}_x^0 \subseteq \mathcal{I}_u^0 \text{ and } \mathcal{I}_x^C \subseteq \mathcal{I}_u^0\}.$$

For the SVM, model identification boils down to finding the active set of the box constrained quadratic optimization problem after a finite number of iterations.

We now turn to our identification result. To ensure model identification, we need the following (mild) assumption:

Assumption 5 (Locally \mathcal{C}^2). For all $j \in \mathcal{S}_{x^*}$, g_j is locally \mathcal{C}^2 around x_j^* , and f is locally \mathcal{C}^2 around x^* .

It is satisfied for the Lasso and the dual SVM problem mentioned above, but also for sparse logistic regression or elastic net. The following theorem shows that the CD ([Algorithm 1](#)) has the model identification property with local constant step size $0 < \gamma_j \leq 1/L_j$:

Theorem 1 (Model identification of CD). *Consider a solution $x^* \in \arg \min_{x \in \mathbb{R}^p} \Phi(x)$ and $\mathcal{S} = \mathcal{S}_{x^*}$. Suppose*

1. *Assumptions 1 to 5 hold.*
2. *The sequence $(x^{(k)})_{k \geq 0}$ generated by [Algorithm 1](#) converges to x^* .*

Then, [Algorithm 1](#) identifies the model after a finite number of iterations, which means that there exists $K > 0$ such that for all $k \geq K$, $x_{\mathcal{S}^c}^{(k)} = x_{\mathcal{S}^c}^$.*

This result implies that for k large enough, $x^{(k)}$ shares the support of x^* (potentially smaller).

Sketch of proof [[Theorem 1](#)]

- First we show that [Assumptions 1 to 5](#) implies that g is *partly smooth* ([Lewis, 2002](#)) at x^* relative to the affine space $x^* + T_{x^*}$.
- Then we show that for the CD [Algorithm 1](#): $\text{dist}(\partial\Phi(x^{(k)}), 0) \rightarrow 0$, when $k \rightarrow \infty$, enabling us to apply [Hare and Lewis \(2004\)](#)[Thm. 5.3].

A full proof of [Theorem 1](#) can be found in [Appendix A](#). The first point is shown in appendix [Lemma 1](#). We show the second point below:

Proof As written in [Algorithm 1](#), one update of coordinate descent reads:

$$\begin{aligned} \frac{1}{\gamma_j} x_j^{(j-1,k)} - \nabla_j f(x^{(j-1,k)}) - \frac{1}{\gamma_j} x_j^{(j,k)} &\in \partial g_j(x_j^{(j,k)}) \\ \frac{1}{\gamma_j} x_j^{(k)} - \nabla_j f(x^{(j-1,k)}) - \frac{1}{\gamma_j} x_j^{(k+1)} &\in \partial g_j(x_j^{(k+1)}) \end{aligned} .$$

Since g is separable with non-empty subdifferential, the coordinate wise subdifferential of g is equal to the subdifferential of g , we then have

$$\begin{aligned} \frac{1}{\gamma} \odot x^{(k)} - \left(\nabla_j f(x^{(j-1,k)}) \right)_{j \in [p]} - \frac{1}{\gamma} \odot x^{(k+1)} \\ \in \partial g(x^{(k+1)}) \end{aligned} , \tag{9}$$

which leads to

$$\begin{aligned} \frac{1}{\gamma} \odot x^{(k)} - \left(\nabla_j f(x^{(j-1,k)}) \right)_{j \in [p]} - \frac{1}{\gamma} \odot x^{(k+1)} \\ + \nabla f(x^{(k+1)}) \in \partial \Phi(x^{(k+1)}) \end{aligned} . \tag{10}$$

To prove support identification using [Hare and Lewis \(2004, Theorem 5.3\)](#), we need to bound the distance between $\partial\Phi(x^{(k+1)})$ and 0, using [Equation \(10\)](#):

$$\begin{aligned}
& \text{dist}\left(\partial\Phi(x^{(k+1)}), 0\right)^2 \\
& \leq \sum_{j=1}^p \left| \frac{x_j^{(k)}}{\gamma_j} - \nabla_j f(x^{(j-1,k)}) - \frac{x_j^{(k+1)}}{\gamma_j} + \nabla_j f(x^{(k+1)}) \right|^2 \\
& \leq \|x^{(k)} - x^{(k+1)}\|_{\gamma^{-1}}^2 \\
& + \sum_{j=1}^p \left| \nabla_j f(x^{(j-1,k)}) - \nabla_j f(x^{(k+1)}) \right|^2 \\
& \leq \|x^{(k)} - x^{(k+1)}\|_{\gamma^{-1}}^2 + L^2 \sum_{j=1}^p \|x^{(j-1,k)} - x^{(k+1)}\|^2 \\
& \leq \underbrace{\|x^{(k)} - x^{(k+1)}\|_{\gamma^{-1}}^2 + L^2 \sum_{j=1}^p \sum_{j' \geq j}^p \left| x_{j'}^{(k)} - x_{j'}^{(k+1)} \right|^2}_{\rightarrow 0 \text{ when } k \rightarrow \infty}.
\end{aligned}$$

□

We thus have:

- $\text{dist}(\partial\Phi(x^{(k+1)}), 0) \rightarrow 0$
- $\Phi(x^{(k)}) \rightarrow \Phi(x^*)$ because Φ is prox-regular (since it is convex, see [Poliquin and Rockafellar 1996](#)) and subdifferentially continuous.

Then the conditions to apply [Hare and Lewis \(2004, Th. 5.3\)](#) are met and hence we have model identification after a finite number of iterations. □

Comments on [Theorem 1](#). It unifies several results found in the literature: [Massias et al. \(2019\)](#) showed model identification for the Lasso, solved with coordinate descent, but requiring uniqueness assumption. [Nutini et al. \(2017\)](#) showed some identification results under strong convexity assumption on f . [Theorem 1](#) do not rely on any uniqueness, strong convexity, or local strong convexity hypothesis. Even if the solution of the optimization problem defined in [Equation \(1\)](#) is not unique, CD achieves model identification.

3 Local convergence rates

In this section, we prove local linear convergence of the CD [Algorithm 1](#). After model identification, there exists a regime where the convergence towards a solution of [Equation \(1\)](#) is linear. Local linear convergence was already proved in various settings such as for ISTA and FISTA algorithms (*i.e.*, with an ℓ_1 penalty, [Tao et al. 2016](#)) and then for the general Forward-Backward algorithm ([Liang et al., 2014](#)).

Local linear convergence requires an additional assumption: *restricted injectivity*. It is classical for this type of analysis as it can be found in [Liang et al. \(2017\)](#) and [Poon and Liang \(2019\)](#).

Assumption 6. (Restricted injectivity) For a solution $x^* \in \arg \min_{x \in \mathbb{R}^p} \Phi(x)$, the restricted Hessian to its generalized support $\mathcal{S} = \mathcal{S}_{x^*}$ is definite positive, *i.e.*,

$$\nabla_{\mathcal{S}, \mathcal{S}}^2 f(x^*) \succ 0 . \quad (11)$$

For the Lasso, [Assumption 6](#) is a classical necessary condition to ensure uniqueness of the minimizer ([Fuchs, 2004](#)).

In order to study local linear convergence, we consider the fixed point iteration of a complete epoch (an epoch is a complete pass over all the coordinates). A full epoch of CD can be written:

$$x^{(k+1)} = \psi(x^{(k)}) \triangleq \mathcal{P}_p \circ \dots \circ \mathcal{P}_1(x^{(k)}) , \quad (12)$$

where \mathcal{P}_j are coordinatewise sequential applications of the proximity operator $\mathcal{P}_j : \mathbb{R}^p \rightarrow \mathbb{R}^p$:

$$x \mapsto \begin{pmatrix} x_1 \\ \vdots \\ x_{j-1} \\ \text{prox}_{\gamma_j g_j} (x_j - \gamma_j \nabla_j f(x)) \\ x_{j+1} \\ \vdots \\ x_p \end{pmatrix} .$$

Thanks to model identification ([Theorem 1](#)), we are able to prove that once the generalized support is correctly identified, there exists a regime where CD algorithm converges linearly towards the solution:

Theorem 2 (Local linear convergence). *Consider a solution $x^* \in \arg \min_{x \in \mathbb{R}^p} \Phi(x)$ and $\mathcal{S} = \mathcal{S}_{x^*}$. Suppose*

1. *Assumptions 1 to 6 hold.*
2. *The sequence $(x^{(k)})_{k \geq 0}$ generated by [Algorithm 1](#) converges to x^* .*
3. *The model has been identified *i.e.*, there exists $K \geq 0$ such as for all $k \geq K$*

$$x_{\mathcal{S}^c}^{(k)} = x_{\mathcal{S}^c}^* .$$

Then $(x^{(k)})_{k \geq K}$ converges linearly towards x^ . More precisely, for any $\nu \in [\rho(\mathcal{J}\psi_{\mathcal{S}, \mathcal{S}}(x^*)), 1[$, there exists $K > 0$ and a constant C such that for all $k \geq K$,*

$$\|x_{\mathcal{S}}^{(k)} - x_{\mathcal{S}}^*\| \leq C \nu^{(k-K)} \|x_{\mathcal{S}}^{(K)} - x_{\mathcal{S}}^*\| .$$




Figure 1: **Lasso, linear convergence.** Distance to optimum, $\|x^{(k)} - x^*\|$, as a function of the number of iterations k , on 4 different datasets: *leukemia*, *gisette*, *rcv1*, and *real-sim*.

The complete proof of [Theorem 2](#) can be found in [Appendix B](#).

Sketch of proof [[Theorem 2](#)]

- A key element of the proof is to consider a full epoch of CD: it can be written as a fixed point iteration: $x^{(k+1)} = \psi(x^{(k)})$ (see [Equation \(12\)](#)).
- We then show that the proximal operators, $\text{prox}_{\gamma_j g_j}$, evaluated at $x_j^* - \gamma_j \nabla_j f(x^*)$ are differentiable (for $j \in [p]$). Once stated, the differentiability of the proximal operator allows us to write the Taylor expansion of ψ :

$$\begin{aligned} x^{(k+1)} - x^* &= \psi(x^{(k)}) - \psi(x^*) \\ &= \langle \mathcal{J}\psi(x^*), x^{(k)} - x^* \rangle + o(\|x^{(k)} - x^*\|). \end{aligned}$$

- Capitalizing on model identification ([Theorem 1](#)) we start from $x_{\mathcal{S}^c}^{(k)} = x_{\mathcal{S}^c}^*$ and show the bound $\rho(\mathcal{J}\psi_{\mathcal{S},\mathcal{S}}(x^*)) < 1$ on the spectral radius of the restricted Jacobian of ψ at x^* : $\mathcal{J}\psi_{\mathcal{S},\mathcal{S}}(x^*)$.
- Finally, all the conditions are met to apply [Polyak \(1987, Th. 1, Sec. 2.1.2\)](#). The latter reference provides sufficient conditions for local linear convergence of sequences based non linear fixed point iterations.

□




Figure 2: **Sparse logistic regression, linear convergence.** Distance to optimum, $\|x^{(k)} - x^*\|$, as a function of the number of iterations k , on 4 different datasets: *leukemia*, *gisette*, *rcv1*, and *real-sim*.

4 Experiments

We now illustrate [Theorems 1](#) and [2](#) on multiple datasets and estimators: the Lasso, the logistic regression and the SVM. In this section, we consider a design matrix $A \in \mathbb{R}^{n \times p}$ and a target $y \in \mathbb{R}^n$ for regression (Lasso) and $y \in \{-1, 1\}^n$ for classification (logistic regression and support-vector machine). We used classical datasets from `libsvm` ([Chang and Lin, 2011](#)) summarized in [Table 1](#).

In [Figures 1](#) to [3](#) the distance of the iterates to the optimum, $\|x^{(k)} - x^*\|$ as a function of the number of iterations k is plotted as a solid blue line. The vertical red dashed line represents the iteration k^* where the model has been identified by CD ([Algorithm 1](#)) illustrating [Theorem 1](#). The yellow dashed line represents the theoretical linear rate from [Theorem 2](#). [Theorem 2](#) gives the slope of the dashed yellow line, the (arbitrary) origin point of the theoretical rate line is chosen such that blue and yellow lines coincide at identification time, *i.e.*, all lines intersect at this point. More precisely, if k^* denotes the iteration where model identification happens, the equation of the dashed yellow line is:

$$h(k) = \|x^{(k^*)} - x^*\| \times \rho(\mathcal{J}\psi_{S,S}(x^*))^{(k-k^*)} . \quad (13)$$

Once a solution x^* has been computed, one can calculate $\mathcal{J}\psi_{S,S}(x^*)$ and its spectral radius for each estimator.

For the experiments we used three different estimators that we detail here.

Lasso. ([Tibshirani, 1996](#)) The most famous estimator based on a nonsmooth optimization problem may be the Lasso. For a design matrix $A \in \mathbb{R}^{n \times p}$ and a

Table 1: Characteristics of the datasets.

Datasets	#samples n	#features p	density
leukemia	38	7129	1
gisette	6000	4955	1
rcv1	20,242	19,959	3.6×10^{-3}
real-sim	72,309	20,958	2.4×10^{-3}
20news	5184	155,148	1.9×10^{-3}




Figure 3: **Support vector machine, linear convergence.** Distance to optimum, $\|x^{(k)} - x^*\|$, as a function of the number of iterations k , on 4 different datasets: *leukemia*, *gisette*, *rcv1* and *20news*.

target $y \in \mathbb{R}^n$ it writes:

$$\arg \min_{x \in \mathbb{R}^p} \frac{1}{2n} \|Ax - y\|^2 + \lambda \|x\|_1 . \quad (14)$$

The CD update for the Lasso is given by

$$x_j \leftarrow \text{ST}_{\gamma_j \lambda} (x_j - \gamma_j A_{:,j}^\top (y - Ax)) , \quad (15)$$

where $\text{ST}_\lambda(x) = \text{sign}(x) \cdot \max(|x| - \lambda, 0)$. The solution of Equation (14) is obtained using Algorithm 1 with constant stepsizes $1/\gamma_j = \frac{\|A_{:,j}\|^2}{n}$.

Sparse logistic regression. The sparse logistic regression is an estimator for classification tasks. It is the solution of the following optimization problem, for a design matrix $A \in \mathbb{R}^{n \times p}$ and a target variable $y \in \{-1, 1\}^n$, with $\sigma(z) \triangleq \frac{1}{1+e^{-z}}$:

$$\arg \min_{x \in \mathbb{R}^p} -\frac{1}{n} \sum_{i=1}^n \log \sigma(y_i x^\top A_{i,:}) + \lambda \|x\|_1 . \quad (16)$$

The CD update for the sparse logistic regression is

$$x_j \leftarrow \text{ST}_{\gamma_j \lambda} (x_j - \gamma_j A_{:,j}^\top (y \odot (\sigma(y \odot Ax) - 1))) . \quad (17)$$

The constant stepsizes for the CD algorithm to solve Equation (16) are given by $1/\gamma_j = \frac{\|A_{:,j}\|^2}{4n}$.

Support-vector machine. (Boser et al., 1992) The support-vector machine (SVM) primal optimization problem is, for a design matrix $A \in \mathbb{R}^{n \times p}$ and

Table 2: C values for SVM.

dataset	leukemia	gisetete	rcv1	20news
C value	10	$1.5 \cdot 10^{-2}$	$1.5 \cdot 10^{-2}$	$5 \cdot 10^{-1}$

a target variable $y \in \{-1, 1\}^n$:

$$\arg \min_{x \in \mathbb{R}^p} \frac{1}{2} \|x\|^2 + C \sum_{i=1}^n \max(1 - y_i x^\top A_{i,:}, 0) . \quad (18)$$

The SVM can be solved using the following dual optimization problem:

$$\begin{aligned} \arg \min_{w \in \mathbb{R}^n} \frac{1}{2} w^\top (y \odot A) (y \odot A)^\top w - \sum_{i=1}^n w_i \\ \text{subject to } 0 \leq w_i \leq C . \end{aligned} \quad (19)$$

The CD update for the SVM reads:

$$w_i \leftarrow \mathcal{P}_{[0, C]} \left(w_i - \gamma_i ((y \odot A)_{i,:}^\top (y \odot Aw) - 1) \right) , \quad (20)$$

where $\mathcal{P}_{[0, C]}(x) = \min(\max(0, x), C)$. The stepsizes of the CD algorithm to solve Equation (19) are given by $1/\gamma_i = \|(y \odot A)_{i,:}\|^2$. The values of the regularization parameter C for each dataset from Figure 3 are given in Table 2.

Comments on Figures 1 to 3. Finite time model identification and local linear convergence are illustrated on the Lasso, the sparse logistic regression and the SVM in Figures 1 to 3. As predicted by Theorem 1, the relative model is identified after a finite number of iterations. For the Lasso (Figure 1) and the sparse logistic regression (Figure 2), we observe that as the regularization parameter gets smaller, the number of iterations needed by the CD algorithm to identify the model increases. To our knowledge, this is a classical empirical observation, that is not backed up by theoretical results. After identification, the convergence towards a solution is linear as predicted by Theorem 2. The theoretical local speed of convergence provided by Theorem 2 seems like a sharp estimation of the true speed of convergence as illustrated by the three figures.

Note that on Figures 1 to 3 high values of λ (or small values of C) were required for the restricted injectivity Assumption 6 to hold. Indeed, despite its lack of theoretical foundation, it is empirically observed that, in general, the larger the value of λ , the smaller the cardinal of the generalized support: $|\mathcal{S}|$. It makes the restricted injectivity Assumption 6: $\nabla_{\mathcal{S}, \mathcal{S}}^2 f(x^*) \succ 0$ easier to be satisfied. For instance, for $\lambda = \lambda_{\max}/20$, the restricted injectivity Assumption 6 was not verified for a lot of datasets for the Lasso and the sparse logistic regression (Figures 1 and 2). In the same vein, values of C for the SVM had to be chosen small enough, in order to make $|\mathcal{S}|$ not too large (Figure 3).

Note that finite time model identification is crucial to ensure local linear convergence, see for instance *20news* dataset on Figure 3. However there exists very few quantitative theoretical results for the convergence speed of the model

identification. Nutini et al. (2019); Sun et al. (2019) tried to obtain some rates on the identification, quantifying “how much the problem is qualified”, *i.e.*, how much Assumption 4 is satisfied. But these theoretical results do not seem to explain fully the experimental results of the CD: in particular the identification speed of the model compared to other algorithms.

Limits. We would like to point out the limit of our analysis illustrated for the case of $\lambda = \lambda_{\max}/15$ for the sparse logistic regression and the *rcv1* dataset in Figure 2. In this case, the solution may no longer be unique. The support gets larger and Assumption 6 is no longer met. In this case, the largest eigenvalue of $\mathcal{J}\psi_{\mathcal{S},\mathcal{S}}(x^*)$ is exactly one, which leads to the constant rate observed in Figure 2. Despite the largest eigenvalue being exactly 1, a regime of locally linear convergence toward a (potentially non unique) minimizer is still observed. Linear convergence of non-strongly convex functions starts to be more and more understood (Necoara et al., 2019). Figure 2 with $\lambda = \lambda_{\max}/15$ for *rcv1* suggests extensions of Necoara et al. (2019) could be possible in the nonsmooth case.

Conclusion and future work. In conclusion, we show finite time model identification for coordinate descent Algorithm 1 (Theorem 1). Thanks to this identification property we were able to show local linear rates of convergence (Theorem 2). These two theoretical results were illustrated on popular estimators (Lasso, sparse logistic regression and SVM dual) and popular machine learning datasets (Section 4).

A first natural extension of this paper would be to investigate block coordinate minimization: Theorem 1 could be extended for blocks under general partial smoothness assumption (Hare and Lewis, 2004). However, it seems that Theorem 2 would require a more careful analysis. A second extension could be to show linear convergence without the restricted injectivity (Assumption 6), paving the way for a generalization of Necoara et al. (2019) as suggested by Figure 2.

Acknowledgements

This work was partly supported by ANR GraVa ANR-18-CE40-0005 and ERC Starting Grant SLAB ERC-StG-676943.

References

- A. Beck and L. Tetruashvili. On the convergence of block coordinate type methods. *SIAM J. Imaging Sci.*, 23(4):651–694, 2013.
- Q. Bertrand, Q. Kloppenstein, M. Blondel, S. Vaiter, A. Gramfort, and J. Salmon. Implicit differentiation of lasso-type models for hyperparameter optimization. *ICML*, 2020.
- D. P. Bertsekas. On the Goldstein-Levitin-Polyak gradient projection method. *IEEE Trans. Autom. Control*, 21(2):174–184, 1976.
- B. E. Boser, I. M. Guyon, and V. N. Vapnik. A training algorithm for optimal margin classifiers. In *Proceedings of the fifth annual workshop on Computational learning theory*, pages 144–152. ACM, 1992.
- J. V. Burke and J. J. Moré. On the identification of active constraints. *SIAM J. Numer. Anal.*, 25(5):1197–1211, 1988.
- E. J. Candès and B. Recht. Simple bounds for recovering low-complexity models. *Math. Program.*, pages 1–13, 2012.
- C.-C. Chang and C.-J. Lin. Libsvm: a library for support vector machines. *ACM transactions on intelligent systems and technology (TIST)*, 2(3):27, 2011.
- S. S. Chen, D. L. Donoho, and M. A. Saunders. Atomic decomposition by basis pursuit. *SIAM J. Sci. Comput.*, 20(1):33–61, 1998.
- P. L. Combettes and V. R. Wajs. Signal recovery by proximal forward-backward splitting. *Multiscale Modeling & Simulation*, 4(4):1168–1200, 2005.
- J. Fadili, J. Malick, and G. Peyré. Sensitivity analysis for mirror-stratifiable convex functions. *SIAM J. Optim.*, 28(4):2975–3000, 2018.
- O. Fercoq and P. Richtárik. Accelerated, parallel and proximal coordinate descent. *SIAM J. Optim.*, 25(3):1997 – 2013, 2015.
- J. Friedman, T. J. Hastie, H. Höfling, and R. Tibshirani. Pathwise coordinate optimization. *Ann. Appl. Stat.*, 1(2):302–332, 2007.
- J. Friedman, T. J. Hastie, and R. Tibshirani. Regularization paths for generalized linear models via coordinate descent. *J. Stat. Softw.*, 33(1):1–22, 2010.
- J.-J. Fuchs. On sparse representations in arbitrary redundant bases. *IEEE Trans. Inf. Theory*, 50(6):1341–1344, 2004.
- W. L. Hare. Identifying active manifolds in regularization problems. In *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, pages 261–271. Springer, 2011.
- W. L. Hare and A. S. Lewis. Identifying active constraints via partial smoothness and prox-regularity. *Journal of Convex Analysis*, 11(2):251–266, 2004.

- W. L. Hare and A. S. Lewis. Identifying active manifolds. *Algorithmic Operations Research*, 2(2):75–75, 2007.
- M. Hong, X. Wang, M. Razaviyayn, and Z-Q. Luo. Iteration complexity analysis of block coordinate descent methods. *Mathematical Programming*, 163(1-2): 85–114, 2017.
- D. Leventhal and A. S. Lewis. Randomized methods for linear constraints: convergence rates and conditioning. *Mathematics of Operations Research*, 35(3):641–654, 2010.
- A. S. Lewis. Active sets, nonsmoothness, and sensitivity. *SIAM Journal on Optimization*, 13(3):702–725, 2002.
- X. Li, T. Zhao, R. Arora, H. Liu, and M. Hong. On faster convergence of cyclic block coordinate descent-type methods for strongly convex minimization. *Journal of Machine Learning Research*, 18(1):6741–6764, 2017.
- J. Liang, J. Fadili, and G. Peyré. Local linear convergence of forward–backward under partial smoothness. In *Advances in neural information processing systems*, pages 1970–1978, 2014.
- J. Liang, J. Fadili, and G. Peyré. Activity Identification and Local Linear Convergence of Forward–Backward-type Methods. *SIAM J. Optim.*, 27(1): 408–437, 2017.
- Z-Q. Luo and P. Tseng. On the convergence of the coordinate descent method for convex differentiable minimization. *Journal of Optimization Theory and Applications*, 72(1):7–35, 1992.
- M. Massias, A. Gramfort, and J. Salmon. Celer: a fast solver for the lasso with dual extrapolation. In *ICML*, volume 80, pages 3315–3324, 2018.
- M. Massias, S. Vaiter, A. Gramfort, and J. Salmon. Dual extrapolation for sparse generalized linear models. *arXiv preprint arXiv:1907.05830*, 2019.
- B. Mercier and G. Vijayasundaram. *Lectures on topics in finite element solution of elliptic problems*. Springer Berlin, 1979.
- 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.
- I. Necoara, Y. Nesterov, and F. Glineur. Linear convergence of first order methods for non-strongly convex optimization. *Mathematical Programming*, 175(1-2):69–107, 2019.
- Y. Nesterov. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM J. Optim.*, 22(2):341–362, 2012.

- J. Nutini, M. W. Schmidt, I. H. Laradji, M. P. Friedlander, and H. A. Koepke. Coordinate descent converges faster with the Gauss-Southwell rule than random selection. In *ICML*, pages 1632–1641, 2015.
- J. Nutini, I. Laradji, and M. Schmidt. Let’s make block coordinate descent go fast: Faster greedy rules, message-passing, active-set complexity, and super-linear convergence. *arXiv preprint arXiv:1712.08859*, 2017.
- J. Nutini, M. Schmidt, and W. Hare. “active-set complexity” of proximal gradient: How long does it take to find the sparsity pattern? *Optimization Letters*, 13(4):645–655, 2019.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.*, 12:2825–2830, 2011.
- R. Poliquin and R. Rockafellar. Prox-regular functions in variational analysis. *Transactions of the American Mathematical Society*, 348(5):1805–1838, 1996.
- B. T. Polyak. Introduction to optimization. optimization software. *Inc., Publications Division, New York*, 1, 1987.
- C. Poon and J. Liang. Trajectory of alternating direction method of multipliers and adaptive acceleration. In *Advances In Neural Information Processing Systems*, pages 7357–7365, 2019.
- C. Poon, J. Liang, and C.-B. Schönlieb. Local convergence properties of SAGA/Prox-SVRG and acceleration. In *ICML*, volume 90, pages 4121–4129, 2018.
- Z. Qu and P. Richtárik. Coordinate descent with arbitrary sampling i: Algorithms and complexity. *Optimization Methods and Software*, 31(5):829–857, 2016a.
- Z. Qu and P. Richtárik. Coordinate descent with arbitrary sampling ii: Expected separable overapproximation. *Optimization Methods and Software*, 31(5):858–884, 2016b.
- M. Razaviyayn, M. Hong, and Z.-Q. Luo. A unified convergence analysis of block successive minimization methods for nonsmooth optimization. *SIAM J. Optim.*, 23(2):1126–1153, 2013.
- 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.
- A. Saha and A. Tewari. On the nonasymptotic convergence of cyclic coordinate descent methods. *SIAM J. Optim.*, 23(1):576–601, 2013.

- S. Shalev-Shwartz and A. Tewari. Stochastic methods for l_1 -regularized loss minimization. *Journal of Machine Learning Research*, 12:1865–1892, 2011.
- S. Shalev-Shwartz and T. Zhang. Stochastic dual coordinate ascent methods for regularized loss minimization. *Journal of Machine Learning Research*, 14 (Feb):567–599, 2013.
- J. She and M. Schmidt. Linear convergence and support vector identification of sequential minimal optimization. In *10th NIPS Workshop on Optimization for Machine Learning*, volume 5, 2017.
- H.-J. M. Shi, S. Tu, Y. Xu, and W. Yin. A primer on coordinate descent algorithms. *ArXiv e-prints*, 2016.
- R. Sun and M. Hong. Improved iteration complexity bounds of cyclic block coordinate descent for convex problems. In *Advances in Neural Information Processing Systems*, pages 1306–1314, 2015.
- Y. Sun, H. Jeong, J. Nutini, and M. Schmidt. Are we there yet? manifold identification of gradient-related proximal methods. In *AISTATS*, volume 89, pages 1110–1119, 2019.
- S. Tao, D. Boley, and S. Zhang. Local linear convergence of ISTA and FISTA on the LASSO problem. *SIAM J. Optim.*, 26(1):313–336, 2016.
- R. Tibshirani. Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 58(1):267–288, 1996.
- P. Tseng. Convergence of a block coordinate descent method for nondifferentiable minimization. *J. Optim. Theory Appl.*, 109(3):475–494, 2001.
- P. Tseng and S. Yun. Block-coordinate gradient descent method for linearly constrained nonsmooth separable optimization. *J. Optim. Theory Appl.*, 140 (3):513, 2009.
- S. Vaïter, M. Golbabaee, J. Fadili, and G. Peyré. Model selection with low complexity priors. *Information and Inference: A Journal of the IMA*, 4(3): 230–287, 2015.
- S. Vaïter, G. Peyré, and J. Fadili. Model consistency of partly smooth regularizers. *IEEE Transactions on Information Theory*, 64(3):1725–1737, 2018.
- S. J. Wright. Identifiable surfaces in constrained optimization. *SIAM Journal on Control and Optimization*, 31(4):1063–1079, 1993.
- S. J. Wright. Accelerated block-coordinate relaxation for regularized optimization. *SIAM J. Optim.*, 22(1):159–186, 2012.
- Y. Xu and W. Yin. A globally convergent algorithm for nonconvex optimization based on block coordinate update. *Journal of Scientific Computing*, 72(2): 700–734, 2017.

- T. Zhang. Solving large scale linear prediction problems using stochastic gradient descent algorithms. In *Proceedings of the twenty-first international conference on Machine learning*, page 116, 2004.
- H. Zou and T. J. Hastie. Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 67(2):301–320, 2005.

A Proofs of model identification (Theorem 1)

Model identification often relies on the assumption that the nonsmooth function g is regular enough, or more precisely *partly smooth*. Loosely speaking, a partial smooth function behaves smoothly as it lies on the related model and sharply if we move normal to that model. Formally, we recall the definition of partly smooth functions restricted to the case of proper, lower semicontinuous and convex functions.

Definition 3 (Partial smoothness). Let $g : \mathbb{R}^p \mapsto \mathbb{R}$ be a proper closed convex function. g is said to be partly smooth at x relative to a set $\mathcal{M} \subseteq \mathbb{R}^n$ if there exists a neighbourhood \mathcal{U} of x such that

- **(Smoothness)** $\mathcal{M} \cap \mathcal{U}$ is a C^2 -manifold and g restricted to $\mathcal{M} \cap \mathcal{U}$ is C^2 ,
- **(Sharpness)** The tangent space of \mathcal{M} at x is the model tangent space T_x where $T_x = \text{Lin}(\partial g(x))^\perp$,
- **(Continuity)** The set valued mapping ∂g is continuous at x relative to \mathcal{M} .

The class of partly smooth functions was first defined in Lewis (2002). It encompasses a large number of known nonsmooth machine learning optimization penalties, such as the ℓ_1 -norm or box constraints to only name a few, see Vaiter et al. (2018, Section 2.1) for details. Interestingly, this framework enables powerful theoretical tools on model identification such as Hare and Lewis (2004)[Thm. 5.3]. For separable functions, next lemma gives an explicit link between the generalized support Definition 1 (Sun et al., 2019) and the framework of partial smooth functions (Hare and Lewis, 2004).

Lemma 1. *Let $x^* \in \text{dom}(g)$. If for every $j \in \mathcal{S}_{x^*}$, g_j is locally C^2 around x_j^* (Assumption 5), then g is partly smooth at x^* relative to $x^* + T_{x^*}$.*

Proof. We need to prove the three properties of the partial smoothness (Definition 3).

Smoothness. Let us write $\mathcal{M}_{x^*} = x^* + T_{x^*}$ the affine space directed by the model subspace and pointed by x^* . In particular, it is a C^2 -manifold.

For every $j \in \mathcal{S}_{x^*}$, g_j is locally C^2 around x_j^* , hence there exists a neighborhood U_j of x_j^* such that the restriction of f to U is twice continuously differentiable. For $j \in \mathcal{S}_{x^*}^c$, let's write $U_j = \mathbb{R}$. Take $U = \bigotimes_{j \in [p]} U_j$. This a neighborhood of x^* (it is open, and contains x^*). Consider the restriction $g|_{\mathcal{M}_{x^*}}$ of g to \mathcal{M}_{x^*} . It is C^2 at each point of U since each coordinates (for $j \in \mathcal{S}_{x^*}$) are C^2 around U_j .

Sharpness. Since g is completely separable, we have that $\partial g(x^*) = \partial g_1(x_1^*) \times \dots \times \partial g_p(x_p^*)$. Note that $\partial g_j(x_j^*)$ is a set valued mapping which is equal to the singleton $\{\nabla_j g(x_j^*)\}$ if g_j is differentiable at x_j^* or it is equal to an interval. The model tangent space T_{x^*} of g at x^* is given by

$$T_{x^*} = \text{span}(\partial g(x^*))^\perp \quad \text{where} \quad \text{span}(\partial g(x^*)) = \text{aff}(\partial g(x^*)) - e_{x^*} \quad , \quad (21)$$

with

$$e_{x^*} = \arg \min_{e \in \text{aff}(\partial g(x^*))} \|e\|, \quad (22)$$

called the model vector.

In the particular case of separable functions, we have that

$$\text{aff}(\partial g(x^*)) = \text{aff}(\partial g_1(x_1^*) \times \dots \times \partial g_p(x_p^*)) = \text{aff}(\partial g_1(x_1^*)) \times \dots \times \text{aff}(\partial g_p(x_p^*)) .$$

In this case,

$$\text{aff}(\partial g_j(x_j^*)) = \begin{cases} \{\nabla_j g(x_j^*)\} & \text{if } j \in \mathcal{S}_{x^*} \\ \mathbb{R} & \text{otherwise} \end{cases} \quad \text{and} \quad e_{x_j^*} = \begin{cases} \nabla_j g(x_j^*) & \text{if } j \in \mathcal{S}_{x^*} \\ 0 & \text{otherwise} \end{cases} . \quad (23)$$

Thus we have that

$$\text{span}(\partial g(x^*)) = \text{aff}(\partial g(x^*)) - e_{x^*} = \{x \in \mathbb{R}^p : \forall j' \in \mathcal{S}_{x^*}, x_{j'} = 0\} .$$

Then

$$T_{x^*} = \text{span}(\partial g(x^*))^\perp = \{x \in \mathbb{R}^p : \forall j' \in \mathcal{S}_{x^*}, x_{j'} = 0\} . \quad (24)$$

Continuity. We are going to prove that ∂g is inner semicontinuous at x^* relative to \mathcal{M}_{x^*} , *i.e.*, that for any sequence $(x^{(k)})$ of elements of \mathcal{M}_{x^*} converging to x^* and any $\bar{\eta} \in \partial g(x^*)$, there exists a sequence of subgradients $\eta^{(k)} \in \partial g(x^{(k)})$ converging to $\bar{\eta}$.

Let $x^{(k)}$ be a sequence of elements of \mathcal{M}_{x^*} converging to x^* , or equivalently, let $t^{(k)}$ be a sequence of elements of T_{x^*} converging to 0, and let $\bar{\eta} \in \partial g(x^*)$.

For $j \in \mathcal{S}_{x^*}$, we choose $\eta_j^{(k)} \triangleq g'_j(x_j^* + t_j^{(k)})$, using the smoothness property we have $\eta_j^{(k)} \triangleq \bar{\eta}_j$. For all $j \in \mathcal{S}_{x^*}^c$, $x_j^{(k)} = x_j^*$ we choose $\eta_j^{(k)} \triangleq \bar{\eta}_j$, since $x^{(k)} \in \mathcal{M}_{x^*}$, we have $\eta_j^{(k)} \in \partial g(x^{(k)})$.

We have that $\eta^{(k)} \in \partial g(x^{(k)})$ and $\eta^{(k)}$ converges towards $\bar{\eta}$ since g'_j is C^1 around x_j^* for $j \in \mathcal{S}_{x^*}$, hence, $g'_j(x_j^* + t_j^{(k)})$ converges to $g'_j(x_j^*) = \bar{\eta}_j$. Thus, it proves that g is partly smooth at x^* relative to $x^* + T_{x^*}$. \square

The end of the proof of [Theorem 1](#) is contained in [Section 2](#).

B Proofs of local linear convergence ([Theorem 2](#))

To simplify the notations in this section, $\mathcal{S} \triangleq \mathcal{S}_{x^*}$. Let us also write the element of \mathcal{S} as follows: $\mathcal{S} = \{j_1, \dots, j_{|\mathcal{S}|}\}$. The first point of this proof is to write the CD algorithm as a fixed point iteration. A full epoch of CD can be written as

$$x^{(k+1)} = \psi(x^{(k)}) = \mathcal{P}_p \circ \dots \circ \mathcal{P}_1(x^{(k)}) . \quad (25)$$

We also define $\pi^{x_{\mathcal{S}^c}^*} : \mathbb{R}^{|\mathcal{S}|} \rightarrow \mathbb{R}^p$ for all $x_{\mathcal{S}} \in \mathbb{R}^{|\mathcal{S}|}$ and all $j \in \mathcal{S}$ by

$$\left(\pi^{x_{\mathcal{S}^c}^*}(x_{\mathcal{S}})\right)_j = \begin{cases} x_j & \text{if } j \in \mathcal{S} \\ x_j^* & \text{if } j \in \mathcal{S}^c \end{cases}, \quad (26)$$

and for all $j_s \in \mathcal{S}$, $\tilde{\mathcal{P}}_{j_s}^{x_{\mathcal{S}^c}^*} : \mathbb{R}^{|\mathcal{S}|} \rightarrow \mathbb{R}^{|\mathcal{S}|}$ is the function defined for all $x_{\mathcal{S}} \in \mathbb{R}^{|\mathcal{S}|}$ and all $j \in \mathcal{S}$ by

$$\left(\tilde{\mathcal{P}}_{j_s}^{x_{\mathcal{S}^c}^*}(x_{\mathcal{S}})\right)_j = \begin{cases} x_j & \text{if } j \neq j_s \\ \text{prox}_{\gamma_j g_j}(x_{j_s} - \gamma_{j_s} \nabla_{j_s} f(\pi^{x_{\mathcal{S}^c}^*}(x_{\mathcal{S}}))) & \text{if } j = j_s \end{cases}. \quad (27)$$

Once the model is identified ([Theorem 1](#)), we have that there exists $K \geq 0$ such that for all $k \geq K$, we have that

$$x_{\mathcal{S}^c}^{(k)} = x_{\mathcal{S}^c}^* \quad \text{and} \quad x_{\mathcal{S}}^{(k+1)} = \tilde{\psi}(x_{\mathcal{S}}^{(k)}) \triangleq \mathcal{P}_{j_1}^{x_{\mathcal{S}^c}^*} \circ \dots \circ \mathcal{P}_{j_1}^{x_{\mathcal{S}^c}^*}(x_{\mathcal{S}}^{(k)}). \quad (28)$$

When no confusion is possible, we denote by $\tilde{\mathcal{P}}_j$ the function $\tilde{\mathcal{P}}_{j_s}^{x_{\mathcal{S}^c}^*}$, hence still dependant on $x_{\mathcal{S}^c}^*$. The following lemma shows that $\tilde{\mathcal{P}}_j$ is differentiable at the optimum.

Lemma 2. *For all $j \in \mathcal{S}$, $\tilde{\mathcal{P}}_j$ is differentiable at $x_{\mathcal{S}}^*$.*

Proof. From [Assumption 5](#), we know there exists a neighborhood of x_j^* denoted \mathcal{U} such that, for $j \in \mathcal{S}$, the restriction of g_j to \mathcal{U} is \mathcal{C}^2 on \mathcal{U} . In particular, it means that x_j^* is a differentiable point of g_j and given a pair $(u, v) \in \mathcal{U} \times \mathbb{R}^p$ such that

$$u = \text{prox}_{\gamma_j g_j}(v) \in \mathcal{U}, \quad (29)$$

we have $\frac{1}{\gamma_j}(v - u) \in \partial g_j(u)$ becomes

$$\frac{1}{\gamma_j}(v - u) = g_j'(u) \Leftrightarrow v = u + \gamma_j g_j'(u) \Leftrightarrow v = (\text{Id} + g_j')(u). \quad (30)$$

Let $H(u) = (\text{Id} + g_j')(u)$, since g_j is twice differentiable at u , we have that

$$H'(u) = 1 + \gamma_j g_j''(u). \quad (31)$$

Thus, $H' : \mathcal{U} \mapsto \mathbb{R}$ is continuous and then $H : \mathcal{U} \mapsto \mathbb{R}$ is continuously differentiable. Hence $F(v, u) \triangleq v - H(u)$ is \mathcal{C}^1 and $F(v, u) = 0$. By convexity of g , we have $g_j''(u) \geq 0$ and

$$\frac{\partial F}{\partial u}(v, u) = -H'(u) = -1 - \gamma_j g_j''(u) \neq 0. \quad (32)$$

Using the implicit functions theorem, we have that there exists an open interval $\mathcal{V} \subseteq \mathbb{R}$ with $v \in \mathcal{V}$ and a function $h : \mathcal{V} \mapsto \mathbb{R}$ which is \mathcal{C}^1 such as $u = h(v)$.

Using (29) we thus have with the choice $u = x_j^*$, $v = x_j^* - \gamma_j \nabla_j f(x^*)$ that the map h coincides with $\text{prox}_{\gamma_j g_j}$ on \mathcal{V} and is differentiable at $v = x_j^* - \gamma_j \nabla_j f(x^*) \in \mathcal{V}$. It follows that $\tilde{\mathcal{P}}_j$ is differentiable at $x_{\mathcal{S}}^*$. \square

For the sake of completeness, we show that in fact $\text{prox}_{\gamma_j g_j}$ is also differentiable on the complement of the generalized support at $x_j^* - \nabla_j f(x^*)$.

Lemma 3. *For all $j \in \mathcal{S}^c$, $\text{prox}_{\gamma_j g_j}$ is constant around $x_j^* - \nabla_j f(x^*)$. Moreover, the map $x \mapsto \text{prox}_{\gamma_j g_j}(x_j - \nabla_j f(x))$ is differentiable at x^* with gradient 0.*

Proof. Let $\partial g_j(x_j^*) = [a; b]$ and let $z_j^* = x_j^* - \nabla_j f(x^*)$, then combining the fixed point equation and [Assumption 4](#) leads to:

$$\frac{1}{\gamma_j}(z_j^* - x_j^*) \in \text{ri}(\partial g_j(x_j^*)) =]a; b[. \quad (33)$$

Thus,

$$z_j^* \in]\gamma_j a + x_j^*; \gamma_j b + x_j^*[. \quad (34)$$

For all $v \in]\gamma_j a + x_j^*; \gamma_j b + x_j^*[$, we have $\frac{1}{\gamma_j}(v - x_j^*) \in]a; b[= \text{ri}(\partial g_j(x_j^*))$, i.e., $\text{prox}_{\gamma_j g_j}(v) = x_j^*$. As f is \mathcal{C}^2 in x^* , we have that $x \mapsto \text{prox}_{\gamma_j g_j}(x_j - \nabla_j f(x))$ is differentiable at x^* with gradient being 0. \square

From [Lemma 2](#), we have that $\tilde{\mathcal{P}}_j$ is differentiable at $x_{\mathcal{S}}^*$ for all $j \in \mathcal{S}$. Since x^* is an optimal point, the following fixed points equation holds:

$$x_j^* = \text{prox}_{\gamma_j g_j}(x_j^* - \gamma_j \nabla_j f(x^*)) . \quad (35)$$

The map $\tilde{\psi}$ is then differentiable at $x_{\mathcal{S}}^*$ since it is obtained as the composition of differentiable functions and that each function $\tilde{\mathcal{P}}_j$ is evaluated at a differentiable point (only one coordinate change at each step).

To compute, the Jacobian of $\tilde{\mathcal{P}}_j$ at $x_{\mathcal{S}}^*$, let us first notice that

$$\mathcal{JP}_j(x_{\mathcal{S}}^*)^\top = (e_1 \mid \dots \mid e_{j-1} \mid v_j \mid e_{j+1} \mid \dots \mid e_s) , \quad (36)$$

where $v_j = \partial_x \text{prox}_{\gamma_j g_j}(z_j^*)(e_j - \gamma_j \nabla_j^2 f(x^*))$ and $z_j^* = x_j^* - \gamma_j \nabla_j f(x^*)$. This matrix can be rewritten as

$$\begin{aligned} \mathcal{J}\tilde{\mathcal{P}}_j(x_{\mathcal{S}}^*) &= \text{Id}_{|S|} - e_j e_j^\top + \partial_x \text{prox}_{\gamma_j g_j}(z_j^*)(e_j e_j^\top - \gamma_j e_j e_j^\top \nabla^2 f(x^*)) \\ &= \text{Id}_{|S|} - e_j e_j^\top \gamma_j \partial_x \text{prox}_{\gamma_j g_j}(z_j^*)(\text{diag}(u) + \nabla^2 f(x^*)) \\ &= \text{Id}_{|S|} - e_j e_j^\top \gamma_j \partial_x \text{prox}_{\gamma_j g_j}(z_j^*) M \\ &= M^{-1/2} \left(\text{Id}_{|S|} - M^{1/2} e_j e_j^\top \gamma_j \partial_x \text{prox}_{\gamma_j g_j}(z_j^*) M^{1/2} \right) M^{1/2} \\ &= M^{-1/2} (\text{Id}_{|S|} - B_j) M^{1/2} , \end{aligned} \quad (37)$$

where

$$M \triangleq \nabla_{\mathcal{S}, \mathcal{S}}^2 f(x^*) + \text{diag}(u) , \quad (38)$$

and $u \in \mathbb{R}^{|S|}$ is defined for all $j \in \mathcal{S}$ by

$$u_j = \begin{cases} \frac{1}{\gamma_j \partial_x \text{prox}_{\gamma_j g_j}(z_j^*)} - \frac{1}{\gamma_j} & \text{if } \text{prox}_{\gamma_j g_j}(z_j^*) \neq 0 \\ 0 & \text{otherwise,} \end{cases} \quad (39)$$

and

$$B_j = M_{:,j}^{1/2} \gamma_j \partial_x \text{prox}_{\gamma_j g_j}(z_j^*) M_{:,j}^{1/2\top}. \quad (40)$$

Since only one coordinate change at each step, the chain rule leads to

$$\begin{aligned} \mathcal{J}\tilde{\psi}(x_{\mathcal{S}}^*) &= \mathcal{J}\mathcal{P}_{j_s}(x_{\mathcal{S}}^*) \mathcal{J}\mathcal{P}_{j_{s-1}}(x_{\mathcal{S}}^*) \dots \mathcal{J}\mathcal{P}_{j_1}(x_{\mathcal{S}}^*) \\ &= M^{-1/2} \underbrace{(\text{Id} - B_{j_s}) \dots (\text{Id} - B_{j_1})}_A M^{1/2} \end{aligned}$$

The next series of lemma will be useful to prove that the spectral radius $\rho(\mathcal{J}\tilde{\psi}(x_{\mathcal{S}}^*)) < 1$.

Lemma 4. *The matrix M defined in (38) is symmetric definite positive.*

Proof. Using the non-expansivity of the prox, and the property $\partial_x \text{prox}_{\gamma_j g_j}(z_j^*) > 0$ for $j \in \mathcal{S}$, $\text{diag}(u)$ is a symmetric semidefinite matrix, so M is a sum of a symmetric definite positive matrix and a symmetric semidefinite matrix, hence M is symmetric definite positive. \square

Lemma 5. *For all $j \in \mathcal{S}$, the matrix B_j defined in (40) has spectral norm bounded by 1, i.e., $\|B_j\|_2 \leq 1$.*

Proof. B_j is a rank one matrix which is the product of $\gamma_j \partial_x \text{prox}_{\gamma_j g_j}(z_j^*) M_{:,j}^{1/2}$ and $M_{:,j}^{1/2\top}$, its non-zeros eigenvalue is thus given by

$$\begin{aligned} \|B_j\|_2 &= \left| M_{:,j}^{1/2\top} \gamma_j \partial_x \text{prox}_{\gamma_j g_j}(z_j^*) M_{:,j}^{1/2} \right| \\ &= \left| \gamma_j \partial_x \text{prox}_{\gamma_j g_j}(z_j^*) M_{j,j} \right| \\ &= \left| \gamma_j \partial_x \text{prox}_{\gamma_j g_j}(z_j^*) \left(\underbrace{\nabla_{j,j}^2 f(x^*)}_{0 \leq} + \underbrace{\left(\frac{1}{\gamma_j \partial_x \text{prox}_{\gamma_j g_j}(z_j^*)} - \frac{1}{\gamma_j} \right)}_{0 \leq} \right) \right|. \end{aligned} \quad (41)$$

By positivity of the two terms,

$$\begin{aligned} \|B_j\|_2 &= \gamma_j \partial_x \text{prox}_{\gamma_j g_j}(z_j^*) \underbrace{\nabla_{j,j}^2 f(x^*)}_{\leq L_j \leq \frac{1}{\gamma_j}} + \left(1 - \partial_x \text{prox}_{\gamma_j g_j}(z_j^*) \right) \\ &\leq \partial_x \text{prox}_{\gamma_j g_j}(z_j^*) + \left(1 - \partial_x \text{prox}_{\gamma_j g_j}(z_j^*) \right) \\ &\leq 1. \end{aligned} \quad (42)$$

\square

Lemma 6. *For all $j \in \mathcal{S}$, $B_j / \|B_j\|$ is an orthogonal projector onto $\text{Span}(M_{:,j}^{1/2})$.*

Proof. It is clear that $B_j/\|B_j\|$ is symmetric. We now prove that it is idempotent, i.e., $(B_s/\|B_s\|)^2 = B_s/\|B_s\|$.

$$\begin{aligned} B_j^2/\|B_j\|^2 &= (\gamma_j \partial_x \text{prox}_{\gamma_j g_j}(z_j^*))^2 M_{:,j}^{1/2} M_{:,j}^{1/2\top} M_{:,j}^{1/2} M_{:,j}^{1/2\top} / \|B_j\|^2 \\ &= (\gamma_j \partial_x \text{prox}_{\gamma_j g_j}(z_j^*)) \|B_j\| M_{:,j}^{1/2} M_{:,j}^{1/2\top} / \|B_j\|^2 \\ &= B_j / \|B_j\| . \end{aligned}$$

Hence, $B_j/\|B_j\|$ is an orthogonal projector. \square

Lemma 7. For all $j \in \mathcal{S}$ and for all $x \in \mathbb{R}^S$, if $\|(\text{Id} - B_j)x\| = \|x\|$ then $x \in \text{Span}(M_{:,j}^{1/2})^\perp$.

Proof.

$$\begin{aligned} \text{Id} - B_j &= \text{Id} - \|B_j\| \frac{B_j}{\|B_j\|} \\ &= (1 - \|B_j\|) \text{Id} + \|B_j\|_2 \text{Id} - \|B_j\|_2 \frac{B_j}{\|B_j\|_2} \\ &= (1 - \|B_j\|) \text{Id} + \|B_j\| \underbrace{\left(\text{Id} - \frac{B_j}{\|B_j\|_2} \right)}_{\text{projection onto } M_{:,j}^{1/2\perp}} . \end{aligned} \quad (43)$$

Let $x \notin \text{Span}(M_{:,j}^{1/2})^\perp$, then there exists $\kappa \neq 0$, $x_{M_{:,j}^{1/2\perp}} \in \text{Span}(M_{:,j}^{1/2})^\perp$ such that

$$x = \kappa M_{:,j} + x_{M_{:,j}^{1/2\perp}} . \quad (44)$$

Combining Equations (43) and (44) leads to:

$$\begin{aligned} (\text{Id} - B_j)x &= (1 - \|B_j\|_2)x + \|B_j\|_2 x_{M_{:,j}^{1/2\perp}} \\ \|(\text{Id} - B_j)x\| &\leq \underbrace{[1 - \|B_j\|_2] \|x\|}_{=1 - \|B_j\|_2} + \|B_j\|_2 \underbrace{\|x_{M_{:,j}^{1/2\perp}}\|}_{< \|x\|} \\ &< \|x\| . \end{aligned}$$

\square

Lemma 8. The spectral norm of A is bounded by 1, i.e., $\|(\text{Id} - B_{j_s}) \dots (\text{Id} - B_{j_1})\|_2 = \|A\|_2 < 1$.

Proof. Let $x \in \mathbb{R}^s$ such that $\|(\text{Id} - B_{j_s}) \dots (\text{Id} - B_{j_1})x\| = \|x\|$. Since

$$\|(\text{Id} - B_{j_s}) \dots (\text{Id} - B_{j_1})\|_2 \leq \underbrace{\|(\text{Id} - B_{j_s})\|_2}_{\leq 1} \times \dots \times \underbrace{\|(\text{Id} - B_{j_1})\|_2}_{\leq 1} ,$$

we thus have for all $j \in \mathcal{S}$, $\|(\text{Id} - B_j)x\| = \|x\|$. One can thus successively apply [Lemma 7](#) which leads to:

$$x \in \bigcap_{j \in \mathcal{S}} \text{Span } M_{:,j}^{1/2 \perp} \Leftrightarrow x \in \text{Span} \left(M_{:,j_1}^{1/2}, \dots, M_{:,j_s}^{1/2} \right)^\perp .$$

Moreover $M^{1/2}$ has full rank (see [Lemma 4](#)), thus $x = 0$ and

$$\|(\text{Id} - B_{j_s}) \dots (\text{Id} - B_{j_1})\|_2 < 1 .$$

□

From [Lemma 8](#), $\|A\|_2 < 1$. Moreover A and $\mathcal{J}\tilde{\psi}(x_{\mathcal{S}}^*)$ are similar matrices, then $\rho(\mathcal{J}\tilde{\psi}(x_{\mathcal{S}}^*)) = \rho(A) \leq \|A\|_2 < 1$.

To summarize, $x_{\mathcal{S}}^*$ is the solution of a fixed point equation $\tilde{\psi}(x_{\mathcal{S}}^*, x_{\mathcal{S}^c}^*) = x_{\mathcal{S}}^*$. From [Lemma 2](#), $\tilde{\psi}(\cdot, x_{\mathcal{S}^c}^*)$ is differentiable at $x_{\mathcal{S}}^*$ and the Jacobian at $x_{\mathcal{S}}^*$ satisfies the condition $\rho(\mathcal{J}\tilde{\psi}(x_{\mathcal{S}}^*)) < 1$. Then all conditions are met to apply [Polyak \(1987\)](#)[Theorem 1, Section 2.1.2] which proves local linear convergence.