
On Matching Pursuit and Coordinate Descent

Francesco Locatello^{*12} Anant Raj^{*1} Sai Praneeth Karimireddy³
Gunnar Rätsch² Bernhard Schölkopf¹ Sebastian U. Stich³ Martin Jaggi³

Abstract

Two popular examples of first-order optimization methods over linear spaces are coordinate descent and matching pursuit algorithms, with their randomized variants. While the former targets the optimization by moving along coordinates, the latter considers a generalized notion of directions. Exploiting the connection between the two algorithms, we present a unified analysis of both, providing affine invariant sublinear $\mathcal{O}(1/t)$ rates on smooth objectives and linear convergence on strongly convex objectives. As a byproduct of our affine invariant analysis of matching pursuit, our rates for steepest coordinate descent are the tightest known. Furthermore, we show the first accelerated convergence rate $\mathcal{O}(1/t^2)$ for matching pursuit and steepest coordinate descent on convex objectives.

1. Introduction

In this paper we address the following convex optimization problem:

$$\min_{\mathbf{x} \in \text{lin}(\mathcal{A})} f(\mathbf{x}), \quad (1)$$

where f is a convex function. The minimization is over a linear space, which is parametrized as the set of linear combinations of elements from a given set \mathcal{A} . These elements of \mathcal{A} are called *atoms*. In the most general setting, \mathcal{A} is assumed to be a compact but not necessarily finite subset of a Hilbert space, i.e., a linear space equipped with an inner product, complete in the corresponding norm. Problems of the form (1) are tackled by a multitude of first-order optimization methods and are of paramount interest in the

machine learning community (Seber & Lee, 2012; Meir & Rätsch, 2003; Schölkopf & Smola, 2001; Menard, 2018; Tibshirani, 2015).

Traditionally, matching pursuit (MP) algorithms were introduced to solve the inverse problem of representing a measured signal by a sparse combination of atoms from an over-complete basis (Mallat & Zhang, 1993). In other words, the solution of the optimization problem (1) is formed as a linear combination of few of the elements of the atom set \mathcal{A} – i.e. a sparse approximation. At each iteration, the MP algorithm picks a direction from \mathcal{A} according to the gradient information, and takes a step. This procedure is not limited to atoms of fixed dimension. Indeed, $\text{lin}(\mathcal{A})$ can be an arbitrary linear subspace of the ambient space and we are interested in finding the minimizer of f only on this domain, see e.g. (Gillis & Luce, 2018). Conceptually, MP stands in the middle between coordinate descent (CD) and gradient descent, as the algorithm is allowed to descend the function along a prescribed set of directions which does not necessarily correspond to coordinates. This is particularly important for machine learning applications as it translates to a sparse representation of the iterates in terms of the elements of \mathcal{A} while maintaining the convergence guarantees (Lacoste-Julien et al., 2013; Locatello et al., 2017b).

The first analysis of the MP algorithm in the optimization sense to solve the template (1) without incoherence assumptions was done by (Locatello et al., 2017a). To prove convergence, they exploit the connection between MP and the Frank-Wolfe (FW) algorithm (Frank & Wolfe, 1956), a popular projection-free algorithm for the constrained optimization case. On the other hand, steepest coordinate descent is a special case of MP (when the atom set is the L1 ball). This is particularly important as the CD rates can be deduced from the MP rates. Furthermore, the literature on coordinate descent is currently much richer than the one on MP. Therefore, understanding the connection of the two classes of CD and MP-type algorithms is a main goal of this paper, and results in benefits for both sides of the spectrum. In particular, the contributions of this paper are:

- We present an affine invariant convergence analysis for Matching Pursuit algorithms solving (1). Our approach is tightly related to the analysis of coordinate descent

^{*}Equal contribution ¹Max Planck Institute for Intelligent Systems, Tübingen, Germany ²Dept. of Computer Science, ETH Zurich, Zurich, Switzerland ³EPFL, Lausanne, Switzerland. Correspondence to: Francesco Locatello <francesco.locatello@tuebingen.mpg.de>, Anant Raj <anant.raj@tuebingen.mpg.de>.

and relies on the properties of the atomic norm in order to generalize from coordinates to atoms.

- Using our analysis, we present the tightest known linear and sublinear convergence rates for steepest coordinate descent, improving the constants in the rates of (Stich et al., 2017; Nutini et al., 2015).
- We discuss the convergence guarantees of Random Pursuit (RP) methods which we analyze through the lens of MP. In particular, we present a unified analysis of both MP and RP which allows us to carefully trade off the use of (approximate) steepest directions over random ones.
- We prove the first known accelerated rate for MP, as well as for steepest coordinate descent. As a consequence, we also demonstrate an improvement upon the accelerated random CD rate by performing a steepest coordinate update instead.

Related Work: Matching Pursuit was introduced in the context of sparse recovery (Mallat & Zhang, 1993), and later, fully corrective variants similar to the one used in Frank-Wolfe (Holloway, 1974; Lacoste-Julien & Jaggi, 2015; Kerdreux et al., 2018) were introduced under the name of orthogonal matching pursuit (Chen et al., 1989; Tropp, 2004). The classical literature for MP-type methods is typically focused on recovery guarantees for sparse signals and the convergence depends on very strong assumptions (from an optimization perspective), such as incoherence or restricted isometry properties of the atom set (Tropp, 2004; Davenport & Wakin, 2010). Convergence rates with incoherent atom sets are predated in (Gribonval & Vandergheynst, 2006; Temlyakov, 2013; 2014; Nguyen & Petrova, 2014). Also boosting can be seen as a generalized coordinate descent method over a hypothesis class (Rätsch et al., 2001; Meir & Rätsch, 2003).

The idea of following a prescribed set of directions also appears in the field of derivative free methods. For instance, the early method of Pattern-Search (Hooke & Jeeves, 1961; Dennis & Torczon, 1991; Torczon, 1997) explores the search space by probing function values along predefined directions (“patterns” or atoms). This method is in some sense orthogonal to the approach here: by probing the function values along all atoms, one aims to find a direction along which the function decreases (and the absolute value of the scalar product with the gradient is potentially small). MP does not access the function value, but computes the gradient and then picks the atom with the smallest scalar product with the gradient, and then moves to a point where the function value decreases.

The description of random pursuit appears already in the work of Mutsenyeks & Rastrigin (1964) and was first analyzed by Karmanov (1974b;a); Zieliński & Neumann (1983).

More recently random pursuit was revisited in (Stich et al., 2013; Stich, 2014).

Acceleration of first-order methods was first developed in (Nesterov, 1983). An accelerated CD method was described in (Nesterov, 2012). The method was extended in (Lee & Sidford, 2013) for non-uniform sampling, and later in (Stich, 2014) for optimization along arbitrary random directions. Recently, optimal rates have been obtained for accelerated CD (Nesterov & Stich, 2017; Allen-Zhu et al., 2016). A close setup is the accelerated algorithm presented in (El Halabi et al., 2017), which minimizes a composite problem of a convex function on \mathbb{R}^n with a non-smooth regularizer which acts as prior for the structure of the space. Contrary to our setting, the approach is restricted to the atoms being linearly independent. Simultaneously at ICML 2018, Lu et al. (2018) propose an accelerated rate for the semi-greedy coordinate descent which is a special case of our accelerated MP algorithm.

Notation: Given a non-empty subset \mathcal{A} of some Hilbert space, let $\text{conv}(\mathcal{A})$ be the convex hull of \mathcal{A} , and let $\text{lin}(\mathcal{A})$ denote its linear span. Given a closed set \mathcal{A} , we call its diameter $\text{diam}(\mathcal{A}) = \max_{\mathbf{z}_1, \mathbf{z}_2 \in \mathcal{A}} \|\mathbf{z}_1 - \mathbf{z}_2\|$ and its radius $\text{radius}(\mathcal{A}) = \max_{\mathbf{z} \in \mathcal{A}} \|\mathbf{z}\|$. $\|\mathbf{x}\|_{\mathcal{A}} := \inf\{c > 0: \mathbf{x} \in c \cdot \text{conv}(\mathcal{A})\}$ is the atomic norm of \mathbf{x} over a set \mathcal{A} (also known as the gauge function of $\text{conv}(\mathcal{A})$). We call a subset \mathcal{A} of a Hilbert space symmetric if it is closed under negation.

2. Revisiting Matching Pursuit

Let \mathcal{H} be a Hilbert space with associated inner product $\langle \mathbf{x}, \mathbf{y} \rangle$, $\forall \mathbf{x}, \mathbf{y} \in \mathcal{H}$. The inner product induces the norm $\|\mathbf{x}\|^2 := \langle \mathbf{x}, \mathbf{x} \rangle$, $\forall \mathbf{x} \in \mathcal{H}$. Let $\mathcal{A} \subset \mathcal{H}$ be a compact and symmetric set (the “set of atoms” or dictionary) and let $f: \mathcal{H} \rightarrow \mathbb{R}$ be convex and L -smooth (L -Lipschitz gradient in the finite dimensional case). If \mathcal{H} is an infinite-dimensional Hilbert space, then f is assumed to be Fréchet differentiable.

Algorithm 1 Generalized Matching Pursuit

- 1: **init** $\mathbf{x}_0 \in \text{lin}(\mathcal{A})$
 - 2: **for** $t = 0 \dots T$
 - 3: Find $\mathbf{z}_t := (\text{Approx-})\text{LMO}_{\mathcal{A}}(\nabla f(\mathbf{x}_t))$
 - 4: $\mathbf{x}_{t+1} := \mathbf{x}_t - \frac{\langle \nabla f(\mathbf{x}_t), \mathbf{z}_t \rangle}{L \|\mathbf{z}_t\|^2} \mathbf{z}_t$
 - 5: **end for**
-

In each iteration, MP queries a linear minimization oracle (LMO) to find the steepest descent direction among the set \mathcal{A} :

$$\text{LMO}_{\mathcal{A}}(\mathbf{y}) := \arg \min_{\mathbf{z} \in \mathcal{A}} \langle \mathbf{y}, \mathbf{z} \rangle, \quad (2)$$

for a given query vector $\mathbf{y} \in \mathcal{H}$. This key subroutine is shared with the Frank-Wolfe method (Frank & Wolfe, 1956;

Jaggi, 2013) as well as steepest coordinate descent. Indeed, finding the steepest coordinate is equivalent to minimizing Equation 2. The MP update step minimizes a quadratic upper bound $g_{\mathbf{x}_t}(\mathbf{x}) = f(\mathbf{x}_t) + \langle \nabla f(\mathbf{x}_t), \mathbf{x} - \mathbf{x}_t \rangle + \frac{L}{2} \|\mathbf{x} - \mathbf{x}_t\|^2$ of f at \mathbf{x}_t on the direction \mathbf{z} returned by the LMO, where L is an upper bound on the smoothness constant of f with respect to the Hilbert norm $\|\cdot\|$. For $f(\mathbf{x}) = \frac{1}{2} \|\mathbf{y} - \mathbf{x}\|^2$, $\mathbf{y} \in \mathcal{H}$, Algorithm 1 recovers the classical MP algorithm (Mallat & Zhang, 1993).

The LMO. Greedy and projection-free optimization algorithms such as Frank-Wolfe and Matching Pursuit rely on the property that the result of the LMO is a descent direction, which is translated to an *alignment assumption* of the search direction returned by the LMO (i.e., \mathbf{z}_t in Algorithm 1) and the gradient of the objective at the current iteration (see (Locatello et al., 2017b), (Pena & Rodriguez, 2015, third premise) and (Torczon, 1997, Lemma 12 and proof of Proposition 6.4)). Specifically, for Algorithm 1, a symmetric atom set \mathcal{A} ensures that $\langle \nabla f(\mathbf{x}_t), \mathbf{z}_t \rangle < 0$, as long as \mathbf{x}_t is not optimal yet. Indeed, we then have that $\min_{\mathbf{z} \in \mathcal{A}} \langle \nabla f(\mathbf{x}_t), \mathbf{z} \rangle = \min_{\mathbf{z} \in \text{conv}(\mathcal{A})} \langle \nabla f(\mathbf{x}_t), \mathbf{z} \rangle < 0$ where the inequality comes from symmetry as $\mathbf{z} = \mathbf{0} \in \text{conv}(\mathcal{A})$. Note that an alternative sufficient condition instead of symmetry is that \mathcal{A} is the atomic ball of a norm (the so called atomic norm (Chandrasekaran et al., 2012)).

Steepest Coordinate Descent. In the case when \mathcal{A} is the L1-ball, the MP algorithm becomes identical to steepest coordinate descent (Nesterov, 2012). Indeed, due to symmetry of \mathcal{A} , one can rewrite the LMO problem as $i_t = \arg \max_i |\nabla_i f(x)|$, where ∇_i is the i -th component of the gradient, i.e. $\langle \nabla f(x), \mathbf{e}_i \rangle$ with \mathbf{e}_i being one of the natural vectors. Then the update step can be written as:

$$\mathbf{x}_{t+1} := \mathbf{x}_t - \frac{1}{L} \nabla_{i_t} f(\mathbf{x}_t) \mathbf{e}_{i_t}.$$

Note that by assuming a symmetric atom set and solving the LMO problem as defined in (2) the steepest atom is aligned with the negative gradient, therefore the positive stepsize $-\frac{\langle \nabla f(\mathbf{x}_t), \mathbf{z}_t \rangle}{L}$ decreases the objective.

Approximate linear oracles. Exactly solving the LMO defined in (2) can be costly in practice, both in the MP and the CD setting, as \mathcal{A} can contain (infinitely) many atoms. On the other hand, approximate versions can be much more efficient. Algorithm 1 allows for an *approximate* LMO. Different notions of such a LMO were explored for MP and OMP in (Mallat & Zhang, 1993) and (Tropp, 2004), respectively, for the Frank-Wolfe framework in (Jaggi, 2013; Lacoste-Julien et al., 2013) and for coordinate descent (Stich et al., 2017). For given quality parameter $\delta \in (0, 1]$ and given direction $\mathbf{d} \in \mathcal{H}$, the approximate LMO for Algorithm 1 returns a vector $\tilde{\mathbf{z}} \in \mathcal{A}$ such that:

$$\langle \mathbf{d}, \tilde{\mathbf{z}} \rangle \leq \delta \langle \mathbf{d}, \mathbf{z} \rangle, \quad (3)$$

relative to $\mathbf{z} = \text{LMO}_{\mathcal{A}}(\mathbf{d})$ being an exact solution.

2.1. Affine Invariant Algorithm

In this section, we will present our new affine invariant algorithm for the optimization problem (1). Hence, we first explain in Definition 1 that what does it mean for an optimization algorithm to be affine invariant:

Definition 1. *An optimization method is called affine invariant if it is invariant under affine transformations of the input problem: If one chooses any re-parameterization of the domain \mathcal{Q} by a surjective linear or affine map $\mathbf{M} : \hat{\mathcal{Q}} \rightarrow \mathcal{Q}$, then the “old” and “new” optimization problems $\min_{\mathbf{x} \in \mathcal{Q}} f(\mathbf{x})$ and $\min_{\hat{\mathbf{x}} \in \hat{\mathcal{Q}}} \hat{f}(\hat{\mathbf{x}})$ for $\hat{f}(\hat{\mathbf{x}}) := f(\mathbf{M}\hat{\mathbf{x}})$ look the same to the algorithm.*

In other words, a step of the algorithm in the original optimization problem is the same as a step in the transformed problem. We will further demonstrate in the appendix that the proposed Algorithm 2 which we discuss later in detail is indeed an affine invariant algorithm. In order to obtain an affine invariant algorithm, we define an affine invariant notion of smoothness using the atomic norm. This notion is inspired by the curvature constant employed in FW and MP, see (Jaggi, 2013; Locatello et al., 2017a). We define:

$$L_{\mathcal{A}} := \sup_{\substack{\mathbf{x}, \mathbf{y} \in \text{lin}(\mathcal{A}) \\ \mathbf{y} = \mathbf{x} + \gamma \mathbf{z} \\ \|\mathbf{z}\|_{\mathcal{A}} = 1, \gamma \in \mathbb{R}_{>0}}} \frac{2}{\gamma^2} [f(\mathbf{y}) - f(\mathbf{x}) - \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle]. \quad (4)$$

This definition combines the complexity of the function f as well as the set \mathcal{A} into a single number, and is affine invariant under transformations of our input problem (1). It yields the same upper bound to the function as the one given by the traditional smoothness definition, that is $L_{\mathcal{A}}$ -smoothness with respect to the atomic norm $\|\cdot\|_{\mathcal{A}}$, when \mathbf{x}, \mathbf{y} are constrained to the set $\text{lin}(\mathcal{A})$:

$$f(\mathbf{y}) \leq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{L_{\mathcal{A}}}{2} \|\mathbf{y} - \mathbf{x}\|_{\mathcal{A}},$$

For example, if \mathcal{A} is the L1-ball we obtain $f(\mathbf{x} + \gamma \mathbf{z}) \leq f(\mathbf{x}) + \gamma \langle \nabla f(\mathbf{x}), \mathbf{z} \rangle + \gamma^2 \frac{L_1}{2}$ where $\|\mathbf{z}\|_1 = 1$. Based on the affine-invariant notion of smoothness defined above, we now present pseudocode of our affine-invariant method in Algorithm 2.

Algorithm 2 Affine Invariant Generalized Matching Pursuit

- 1: **init** $\mathbf{x}_0 \in \text{lin}(\mathcal{A})$
 - 2: **for** $t = 0 \dots T$
 - 3: Find $\mathbf{z}_t := (\text{Approx-})\text{LMO}_{\mathcal{A}}(\nabla f(\mathbf{x}_t))$
 - 4: $\mathbf{x}_{t+1} = \mathbf{x}_t - \frac{\langle \nabla f(\mathbf{x}_t), \mathbf{z}_t \rangle}{L_{\mathcal{A}}} \mathbf{z}_t$
 - 5: **end for**
-

The above algorithm looks very similar to the generalized MP (Algorithm 1), however, the main difference is that

while the original algorithm is not affine invariant over the domain $\mathcal{Q} = \text{lin}(\mathcal{A})$ (Def 1), the new Algorithm 2 is so, due to using the generalized smoothness constant $L_{\mathcal{A}}$.

Note. For the purpose of the analysis, we call \mathbf{x}^* the minimizer of problem (1). If the optimum is not unique, we pick the one with largest atomic norm as it represent the worst case for the analysis. All the proofs are deferred to the appendix.

2.1.1. NEW AFFINE INVARIANT SUBLINEAR RATE

In this section, we will provide the theoretical justification of our proposed approach for smooth functions (sublinear rate) and its theoretical comparison with existing previous analysis for special cases. We define the level set radius measured with the atomic norm as:

$$R_{\mathcal{A}}^2 := \max_{\substack{\mathbf{x} \in \text{lin}(\mathcal{A}) \\ f(\mathbf{x}) \leq f(\mathbf{x}_0)}} \|\mathbf{x} - \mathbf{x}^*\|_{\mathcal{A}}^2. \quad (5)$$

When we measure this radius with the $\|\cdot\|_2$ we call it R_2^2 , and when we measure it with $\|\cdot\|_1$ we call it R_1^2 . Note that measuring smoothness using the atomic norm guarantees that for the Lipschitz constant $L_{\mathcal{A}}$ the following holds:

Lemma 2. *Assume f is L -smooth w.r.t. a given norm $\|\cdot\|$, over $\text{lin}(\mathcal{A})$ where \mathcal{A} is symmetric. Then,*

$$L_{\mathcal{A}} \leq L \text{ radius}_{\|\cdot\|}(\mathcal{A})^2. \quad (6)$$

For example, in the coordinate descent setting we measure smoothness with the atomic norm being the L1-norm. Lemma 2 implies that $L_{\mathcal{A}} \leq L_1 \leq L_2$ where L_2 is the smoothness constant measured with the L2-norm. Note that the radius of the L1-ball measured with $\|\cdot\|_1$ is 1. Therefore, we put ourselves in a more general setting than Algorithm 1, showing convergence of the affine invariant Algorithm 2

We are now ready to prove the convergence rate of Algorithm 2 for smooth functions.

Theorem 3. *Let $\mathcal{A} \subset \mathcal{H}$ be a closed and bounded set. We assume that $\|\cdot\|_{\mathcal{A}}$ is a norm over $\text{lin}(\mathcal{A})$. Let f be convex and $L_{\mathcal{A}}$ -smooth w.r.t. the norm $\|\cdot\|_{\mathcal{A}}$ over $\text{lin}(\mathcal{A})$, and let $R_{\mathcal{A}}$ be the radius of the level set of \mathbf{x}_0 measured with the atomic norm. Then, Algorithm 2 converges for $t \geq 0$ as*

$$f(\mathbf{x}_{t+1}) - f(\mathbf{x}^*) \leq \frac{2L_{\mathcal{A}}R_{\mathcal{A}}^2}{\delta^2(t+2)},$$

where $\delta \in (0, 1]$ is the relative accuracy parameter of the employed approximate LMO (3).

Discussion. The proof of Theorem 3 extends the convergence analysis of steepest coordinate descent. As opposed to the classical proof in (Nesterov, 2012), the atoms are here

not orthogonal to each other, do not have the same norm and do not correspond to the coordinates of the ambient space. Indeed, $\text{lin}(\mathcal{A})$ could be a subset of the ambient space and the only assumptions on \mathcal{A} are that is closed, bounded and $\|\cdot\|_{\mathcal{A}}$ is a norm over $\text{lin}(\mathcal{A})$. We do not make any incoherence assumption. The key element of our proof is the definition of smoothness using the atomic norm. Furthermore, we use the properties of the atomic norm to obtain a proof which shares the spirit of the Nesterov's one without having to rely on strong assumptions on \mathcal{A} .

Relation to Previous MP Sublinear Rate. The sublinear convergence rate presented in Theorem 3 is fundamentally different in spirit from the one proved in (Locatello et al., 2017a). Indeed, their convergence analysis builds on top of the proof technique used for Frank-Wolfe in (Jaggi, 2013). They introduce a dependency from the atomic norm of the iterates as a way to constrain the part of the space in which the optimization is taking place which artificially induce a notion of duality gap. They do so by defining $\rho := \max \{\|\mathbf{x}^*\|_{\mathcal{A}}, \|\mathbf{x}_0\|_{\mathcal{A}}, \dots, \|\mathbf{x}_T\|_{\mathcal{A}}\} < \infty$. (Locatello et al., 2017a) also used an affine invariant notion of smoothness, thus obtaining an affine invariant rate. On the other hand, their notion of smoothness depends explicitly on ρ . While this constant can be further upper bounded with the level set radius, it is not known a priori, which makes the estimation of the smoothness constant problematic as it is needed in the algorithm and the proof technique more involved. We propose a much more elegant solution, which uses a different affine invariant definition of smoothness which explicitly depend on the atomic norm. Furthermore, we managed to get rid of the dependency on the sequence of the iterates by using only properties of the atomic norm without any additional assumption (finiteness of ρ).

Relation to Steepest Coordinate Descent. From our analysis, we can readily recover existing rates for coordinate descent. Indeed, if \mathcal{A} is the L1-ball in an n dimensional space, the rate of Theorem 3 with exact oracle can be written as:

$$f(\mathbf{x}_{t+1}) - f(\mathbf{x}^*) \leq \frac{2L_1R_1^2}{t+2} \leq \frac{2L_2R_1^2}{t+2} \leq \frac{2L_2nR_2^2}{t+2},$$

where the first inequality is our rate, the second inequality is the rate of (Stich et al., 2017) and the last inequality is the rate given in (Nesterov, 2012). Therefore, by measuring smoothness with the atomic norm, we have shown a tighter dependency on the dimensionality of the space. Indeed, the atomic norm gives the tightest norm to measure the product between the smoothness of the function and the level set radius among the known rates. Therefore, our rate for steepest coordinate descent is the tightest known.

Coordinate Descent and Affine Transformations. But what does it mean to have an affine invariant rate for coordinate descent? By definition, it means that if one applies an affine transformation to the L1-ball, the coordinate descent algorithm in the natural basis and on the transformed domain \hat{Q} are equivalent. Note that in the transformed problem, the coordinates do not correspond to the natural coordinates anymore. Indeed, in the transformed domain the coordinates are $\hat{e}_i = \mathbf{M}^{-1}\mathbf{e}_i$ where \mathbf{M}^{-1} is the inverse of the affine map $\mathbf{M} : \hat{Q} \rightarrow Q$. If one would instead perform coordinate descent in the transformed space using the natural coordinates, one would obtain not only different atoms but also a different iterate sequence. In other words, while Matching Pursuit is fully affine invariant, the definition of CD is not, as the choice of the coordinates is not part of the definition of the optimization problem. The two algorithms do coincide for one particular choice of basis, the canonical coordinate basis for \mathcal{A} .

2.1.2. SUBLINEAR RATE OF RANDOM PURSUIT

There is a significant literature on optimization methods which do not require full gradient information. A notable example is random coordinate descent, where only a random component of the gradient is known. As long as the direction that is selected by the LMO is not orthogonal to the gradient we have convergence guarantees due to the inexact oracle definition. We now abstract from the random coordinate descent setting and analyze a randomized variant of matching pursuit, the *random pursuit* algorithm, in which the atom \mathbf{z} is randomly sampled from a distribution over \mathcal{A} , rather than picked by a linear minimization oracle. This approach is particularly interesting, as it is deeply connected to the random pursuit algorithm analyzed in (Stich et al., 2013). For now we assume that we can compute the projection of the gradient onto a single atom $\langle \nabla f, \mathbf{z} \rangle$ efficiently. In order to present a general recipe for any atom set, we exploit the notion of inexact oracle and define the inexactness of the expectation of the sampled direction for a given sampling distribution:

$$\hat{\delta}^2 := \min_{\mathbf{d} \in \text{lin}(\mathcal{A})} \frac{\mathbb{E}_{\mathbf{z} \in \mathcal{A}} \langle \mathbf{d}, \mathbf{z} \rangle^2}{\|\mathbf{d}\|_{\mathcal{A}^*}^2}. \quad (7)$$

This constant was already used in (Stich, 2014) to measure the convergence of random pursuit (β^2 in his notation). Note that for uniform sampling from the corners of the L1-ball, we have $\hat{\delta}^2 = \frac{1}{n}$. Indeed, $\mathbb{E}_{\mathbf{z} \in \mathcal{A}} \langle \mathbf{d}, \mathbf{z} \rangle^2 = \frac{1}{n}$ for any \mathbf{d} . This definition holds for any sampling scheme as long as $\hat{\delta}^2 \neq 0$. Note that by using this quantity we do not get the tightest possible rate, as at each iteration, we consider how much worse a random update could be compared to the optimal (steepest) update.

We are now ready to present the sublinear convergence rate of random matching pursuit.

Theorem 4. *Let $\mathcal{A} \subset \mathcal{H}$ be a closed and bounded set. We assume that $\|\cdot\|_{\mathcal{A}}$ is a norm. Let f be convex and $L_{\mathcal{A}}$ -smooth w.r.t. the norm $\|\cdot\|_{\mathcal{A}}$ over $\text{lin}(\mathcal{A})$ and let $R_{\mathcal{A}}$ be the radius of the level set of \mathbf{x}_0 measured with the atomic norm. Then, Algorithm 2 converges for $t \geq 0$ as*

$$\mathbb{E}_{\mathbf{z}} [f(\mathbf{x}_{t+1})] - f(\mathbf{x}^*) \leq \frac{2L_{\mathcal{A}}R_{\mathcal{A}}^2}{\hat{\delta}^2(t+2)},$$

when the LMO is replaced with random sampling of \mathbf{z} from a distribution over \mathcal{A} .

Gradient-Free Variant. If it is possible to obtain a fully gradient-free optimization scheme. In addition to having replaced the LMO in Algorithm 1 by the random sampling as above, as can additionally also replace the line search step on the quadratic upper bound given by smoothness, with instead an approximate line search on f . As long as the update scheme guarantees as much decrease as the above algorithm, the convergence rate of Theorem 4 holds.

Discussion. This approach is very general, as it allows to guarantee convergence for *any* sampling scheme and *any* set \mathcal{A} provided that $\hat{\delta}^2 \neq 0$. In the coordinate descent case we have that for the worst possible gradient for random has $\hat{\delta}^2 = \frac{1}{n}$. Therefore, the speed-up of steepest can be up to a factor equal to the number of dimensions in the best case. Similarly, if \mathbf{z} is sampled from a spherical distribution, $\hat{\delta}^2 = \frac{1}{n}$ (Stich et al., 2013). More examples of computation of $\hat{\delta}^2$ can be found in (Stich, 2014, Section 4.2). Last but not least, note that $\hat{\delta}^2$ is affine invariant as long as the sampling distribution over the atoms is preserved.

2.1.3. STRONG CONVEXITY AND AFFINE INVARIANT LINEAR RATES

Similar to the affine invariant notion of smoothness, we here define the affine invariant notion of strong convexity.

$$\mu_{\mathcal{A}} := \inf_{\substack{\mathbf{x}, \mathbf{y} \in \text{lin}(\mathcal{A}) \\ \mathbf{x} \neq \mathbf{y}}} \frac{2}{\|\mathbf{y} - \mathbf{x}\|_{\mathcal{A}}^2} D(\mathbf{y}, \mathbf{x}).$$

where $D(\mathbf{y}, \mathbf{x}) := f(\mathbf{y}) - f(\mathbf{x}) - \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle$. We can now show the linear convergence rate of both the matching pursuit algorithm and its random pursuit variant.

Theorem 5. *Let $\mathcal{A} \subset \mathcal{H}$ be a closed and bounded set. We assume that $\|\cdot\|_{\mathcal{A}}$ is a norm. Let f be $\mu_{\mathcal{A}}$ -strongly convex and $L_{\mathcal{A}}$ -smooth w.r.t. the norm $\|\cdot\|_{\mathcal{A}}$, both over $\text{lin}(\mathcal{A})$. Then, Algorithm 2 converges for $t \geq 0$ as*

$$\varepsilon_{t+1} \leq \left(1 - \delta^2 \frac{\mu_{\mathcal{A}}}{L_{\mathcal{A}}}\right) \varepsilon_t.$$

where $\varepsilon_t := f(\mathbf{x}_t) - f(\mathbf{x}^*)$. If the LMO direction is sampled randomly from \mathcal{A} , Algorithm 2 converges for $t \geq 0$ as

$$\mathbb{E}_{\mathbf{z}} [\varepsilon_{t+1} | \mathbf{x}_t] \leq \left(1 - \hat{\delta}^2 \frac{\mu_{\mathcal{A}}}{L_{\mathcal{A}}}\right) \varepsilon_t.$$

Relation to Previous MP Linear Rate. Again, the proof of Theorem 5 extends the convergence analysis of steepest coordinate descent using solely the affine invariant definition of strong convexity and the properties of the atomic norm. Note that again we define the strong convexity constant without relying on $\rho = \max \{ \|\mathbf{x}^*\|_{\mathcal{A}}, \|\mathbf{x}_0\|_{\mathcal{A}} \dots, \|\mathbf{x}_T\|_{\mathcal{A}} \} < \infty$ as in (Locatello et al., 2017a). We now show that our choice of the strong convexity parameter is the tightest w.r.t. any choice of the norm and that we can precisely recover the non affine invariant rate of (Locatello et al., 2017a). Let us recall their notion of *minimal directional width*, which is the crucial constant to measure the geometry of the atom set for a fixed norm:

$$\text{mDW}(\mathcal{A}) := \min_{\mathbf{d} \in \text{lin}(\mathcal{A})} \max_{\substack{\mathbf{z} \in \mathcal{A} \\ \mathbf{d} \neq 0}} \left\langle \frac{\mathbf{d}}{\|\mathbf{d}\|}, \mathbf{z} \right\rangle.$$

Note that for CD we have that $\text{mDW}(\mathcal{A}) = \frac{1}{\sqrt{n}}$. Now, we relate the affine invariant notion of strong convexity with the minimal directional width and the strong convexity w.r.t. any chosen norm. This is important, as we want to make sure to perfectly recover the convergence rate given in (Locatello et al., 2017a).

Lemma 6. *Assume f is μ -strongly convex w.r.t. a given norm $\|\cdot\|$ over $\text{lin}(\mathcal{A})$ and \mathcal{A} is symmetric. Then:*

$$\mu_{\mathcal{A}} \geq \text{mDW}(\mathcal{A})^2 \mu.$$

We then recover their non-affine-invariant rate as:

$$\varepsilon_{t+1} \leq \left(1 - \delta^2 \frac{\mu \text{mDW}(\mathcal{A})^2}{L \text{radius}_{\|\cdot\|}(\mathcal{A})^2} \right) \varepsilon_t.$$

Relation to Coordinate Descent. When we fix \mathcal{A} as the L1-ball and use an exact oracle our rate becomes:

$$\varepsilon_{t+1} \leq \left(1 - \frac{\mu_1}{L_1} \right) \varepsilon_t \leq \left(1 - \frac{\mu_1}{L} \right) \varepsilon_t \leq \left(1 - \frac{\mu}{nL} \right) \varepsilon_t,$$

where the first is our rate, the second is the rate of steepest CD (Nutini et al., 2015) and the last is the one for randomized CD (Nesterov, 2012) (n is the dimension of the ambient space). Therefore, our linear rate for coordinate descent is the tightest known.

3. Accelerating Generalized Matching Pursuit

As we established in the previous sections, matching pursuit can be considered a generalized greedy coordinate descent where the allowed directions do not need to form an orthogonal basis. This insight allows us to generalize the analysis of accelerated coordinate descent methods and to accelerate matching pursuit (Lee & Sidford, 2013; Nesterov & Stich, 2017). However it is not clear at the outset how to even accelerate greedy coordinate descent, let alone the matching

pursuit method. Recently Song et al. (2017) proposed an accelerated greedy coordinate descent method by using the linear coupling framework of (Allen-Zhu & Orecchia, 2014). However the updates they perform at each iteration are not guaranteed to be sparse which is critical for our application. We instead extend the acceleration technique in (Stich et al., 2013) which in turn is based on (Lee & Sidford, 2013). They allow the updates to the two sequences of iterates \mathbf{x} and \mathbf{v} to be chosen from any distribution. If this distribution is chosen to be over coordinate directions, we get the familiar accelerated coordinate descent, and if we instead chose the distribution to be over the set of atoms, we would get an accelerated random pursuit algorithm. To obtain an accelerated *matching* pursuit algorithm, we need to additionally *decouple* the updates for \mathbf{x} and \mathbf{v} and allow them to be chosen from different distributions. We will update \mathbf{x} using the greedy coordinate update (or the matching pursuit update), and use a random coordinate (or atom) direction to update \mathbf{v} .

The possibility of decoupling the updates was noted in (Stich, 2014, Corollary 6.4) though its implications for accelerating greedy coordinate descent or matching pursuit were not explored. From here on out, we shall assume that the linear space spanned by the atoms \mathcal{A} is finite dimensional. This was not necessary for the non-accelerated matching pursuit and it remains open if it is necessary for accelerated MP. When sampling, we consider only a non-symmetric version of the set \mathcal{A} with all the atoms in the same half space. Line search ensures that sampling either \mathbf{z} or $-\mathbf{z}$ yields the same update. For simplicity, we focus on an exact LMO.

3.1. From Coordinates to Atoms

For the acceleration of MP we make some stronger assumption w.r.t. the rates in the previous section. In particular, we will not obtain an affine invariant rate which remains an open problem. The key challenges for an affine invariant accelerated rate are strong convexity of the model, which can be solved using arguments similar to (d’Aspremont et al., 2018) and the fact that our proof relies on defining a new norm which deform the space in order to obtain favorable sampling properties as we will explain in this section. The main difference between working with atoms and working with coordinates is that projection along coordinate basis vectors is ‘unbiased’. Let \mathbf{e}_i represent the i th coordinate basis vector. Then for some vector \mathbf{d} , if we project along a random basis vector \mathbf{e}_i ,

$$\mathbb{E}_{i \in [n]} [\langle \mathbf{e}_i, \mathbf{d} \rangle \mathbf{e}_i] = \frac{1}{n} \mathbf{d}.$$

However if instead of coordinate basis, we choose from a set of atoms \mathcal{A} , then this is no longer true. We can correct for this by morphing the geometry of the space. Suppose we sample the atoms from a distribution \mathcal{Z} defined over \mathcal{A} .

Let us define

$$\tilde{\mathbf{P}} := \mathbb{E}_{\mathbf{z} \sim \mathcal{Z}}[\mathbf{z}\mathbf{z}^\top].$$

We assume that the distribution \mathcal{Z} is such that $\text{lin}(\mathcal{A}) \subseteq \text{range}(\tilde{\mathbf{P}})$. This intuitively corresponds to assuming that there is a non-zero probability that the sampled $\mathbf{z} \sim \mathcal{Z}$ is along the direction of every atom $\mathbf{z}_t \in \mathcal{A}$ i.e.

$$\mathbb{P}_{\mathbf{z} \sim \mathcal{Z}}[\langle \mathbf{z}, \mathbf{z}_t \rangle > 0] > 0, \forall \mathbf{z}_t \in \mathcal{A}.$$

Further let $\mathbf{P} = \tilde{\mathbf{P}}^\dagger$ be the pseudo-inverse of $\tilde{\mathbf{P}}$. Note that both \mathbf{P} and $\tilde{\mathbf{P}}$ are positive semi-definite matrices. We can equip our space with a new inner product $\langle \cdot, \mathbf{P} \cdot \rangle$ and the resulting norm $\|\cdot\|_{\mathbf{P}}$. With this new dot product,

$$\mathbb{E}_{\mathbf{z} \sim \mathcal{Z}}[\langle \mathbf{z}, \mathbf{P}\mathbf{d} \rangle \mathbf{z}] = \mathbb{E}_{\mathbf{z} \sim \mathcal{Z}}[\mathbf{z}\mathbf{z}^\top] \mathbf{P}\mathbf{d} = \tilde{\mathbf{P}} \mathbf{P}\mathbf{d} = \mathbf{d}.$$

The last equality follows from our assumption that $\text{lin}(\mathcal{A}) \subseteq \text{range}(\tilde{\mathbf{P}})$.

Algorithm 3 Accelerated Random Pursuit

- 1: **init** $\mathbf{x}_0 = \mathbf{v}_0 = \mathbf{y}_0$, $\beta_0 = 0$, and ν'
 - 2: **for** $t = 0, 1 \dots T$
 - 3: Solve $\alpha_{t+1}^2 L\nu' = \beta_t + \alpha_{t+1}$
 - 4: $\beta_{t+1} := \beta_t + \alpha_{t+1}$
 - 5: $\tau_t := \frac{\alpha_{t+1}}{\beta_{t+1}}$
 - 6: Compute $\mathbf{y}_t := (1 - \tau_t)\mathbf{x}_t + \tau_t \mathbf{v}_t$
 - 7: Sample $\mathbf{z}_t \sim \mathcal{Z}$
 - 8: $\mathbf{x}_{t+1} := \mathbf{y}_t - \frac{\langle \nabla f(\mathbf{y}_t), \mathbf{z}_t \rangle}{L\|\mathbf{z}_t\|_2^2} \mathbf{z}_t$
 - 9: $\mathbf{v}_{t+1} := \mathbf{v}_t - \alpha_{t+1} \langle \nabla f(\mathbf{y}_t), \mathbf{z}_t \rangle \mathbf{z}_t$
 - 10: **end for**
-

Algorithm 4 Accelerated Matching Pursuit

- 1: **init** $\mathbf{x}_0 = \mathbf{v}_0 = \mathbf{y}_0$, $\beta_0 = 0$, and ν
 - 2: **for** $t = 0, 1 \dots T$
 - 3: Solve $\alpha_{t+1}^2 L\nu = \beta_t + \alpha_{t+1}$
 - 4: $\beta_{t+1} := \beta_t + \alpha_{t+1}$
 - 5: $\tau_t := \frac{\alpha_{t+1}}{\beta_{t+1}}$
 - 6: Compute $\mathbf{y}_t := (1 - \tau_t)\mathbf{x}_t + \tau_t \mathbf{v}_t$
 - 7: Find $\mathbf{z}_t := \text{LMO}_{\mathcal{A}}(\nabla f(\mathbf{y}_t))$
 - 8: $\mathbf{x}_{t+1} := \mathbf{y}_t - \frac{\langle \nabla f(\mathbf{y}_t), \mathbf{z}_t \rangle}{L\|\mathbf{z}_t\|_2^2} \mathbf{z}_t$
 - 9: Sample $\tilde{\mathbf{z}}_t \sim \mathcal{Z}$
 - 10: $\mathbf{v}_{t+1} := \mathbf{v}_t - \alpha_{t+1} \langle \nabla f(\mathbf{y}_t), \tilde{\mathbf{z}}_t \rangle \tilde{\mathbf{z}}_t$
 - 11: **end for**
-

3.2. Analysis

Modeling explicitly the dependency on the structure of the set is crucial to accelerate MP. Indeed, acceleration works by defining two different quadratic subproblems, one upper bound given by smoothness, and one lower bound given by a model of the function. The constraints on the set of possible descent direction implicitly used in MP influence both these

subproblems. While the smoothness quadratic upper bound contains information about \mathcal{A} in its definition ($\mathbf{y} = \mathbf{x} + \gamma \mathbf{z}$ and $\|\mathbf{z}\|_{\mathcal{A}} = 1$), the model of the function needs explicit modeling of \mathcal{A} . This is particularly crucial when sampling a direction in the model update, which can be thought as a sort of exploration part of the algorithm. In both the algorithms, the update of the parameter \mathbf{v} corresponds to optimizing the modeling function ψ which can be given as :

$$\psi_{t+1}(\mathbf{x}) = \psi_t(\mathbf{x}) + \alpha_{t+1} \left(f(\mathbf{y}_t) + \langle \tilde{\mathbf{z}}_t^\top \nabla f(\mathbf{y}_t), \tilde{\mathbf{z}}_t^\top \mathbf{P}(\mathbf{x} - \mathbf{y}_t) \rangle \right), \quad (8)$$

where $\psi_0(\mathbf{x}) = \frac{1}{2} \|\mathbf{x} - \mathbf{x}_0\|_{\mathbf{P}}^2$.

Lemma 7. *The update of \mathbf{v} in Algorithm 3 and 4 minimizes the model*

$$\mathbf{v}_t \in \arg \min_{\mathbf{x}} \psi_t(\mathbf{x}).$$

We will be first discussing the theory for the *greedy* accelerated method in detail. As evident from the algorithm 4, another important constant which is required for both the analysis and to actually run the algorithm is ν for which:

$$\nu \leq \max_{\mathbf{d} \in \text{lin}(\mathcal{A})} \frac{\mathbb{E} \left[(\tilde{\mathbf{z}}_t^\top \mathbf{d})^2 \|\tilde{\mathbf{z}}_t\|_{\mathbf{P}}^2 \right] \|\mathbf{z}(\mathbf{d})\|_2^2}{(\mathbf{z}(\mathbf{d})^\top \mathbf{d})^2},$$

where $\mathbf{z}(\mathbf{d})$ is defined to be

$$\mathbf{z}(\mathbf{d}) = \text{LMO}_{\mathcal{A}}(-\mathbf{d}).$$

The quantity ν relates the geometry of the atom set with the sampling procedure in a similar way as $\hat{\delta}^2$ in Equation (7) but instead of measuring how much worse a random update is when compared to a steepest update.

Theorem 8. *Let f be a convex function and \mathcal{A} be a symmetric compact set. Then the output of algorithm 4 for any $t \geq 1$ converges with the following rate:*

$$\mathbb{E}[f(\mathbf{x}_t)] - f(\mathbf{x}^*) \leq \frac{2L\nu}{t(t+1)} \|\mathbf{x}^* - \mathbf{x}_0\|_{\mathbf{P}}^2.$$

Proof. We extend the proof technique of (Lee & Sidford, 2013; Stich et al., 2013) to allow for general atomic updates. The analysis can be found in Appendix C.1 \square

Once we understand the convergence of the greedy approach, the analysis of accelerated random pursuit can be derived easily. Here, we state the rate of convergence for accelerated random pursuit:

Theorem 9. *Let f be a convex function and \mathcal{A} be a symmetric set. Then the output of the algorithm 3 for any $t \geq 1$ converges with the following rate:*

$$\mathbb{E}[f(\mathbf{x}_t)] - f(\mathbf{x}^*) \leq \frac{2L\nu'}{t(t+1)} \|\mathbf{x}^* - \mathbf{x}_0\|_{\mathbf{P}}^2,$$

where

$$\nu' \leq \max_{\mathbf{d} \in \text{lin}(\mathcal{A})} \frac{\mathbb{E} \left[(\mathbf{z}_t^\top \mathbf{d})^2 \|\mathbf{z}_t\|_{\mathbf{P}}^2 \right]}{\mathbb{E} \left[(\mathbf{z}_t^\top \mathbf{d})^2 / \|\mathbf{z}_t\|_2^2 \right]}.$$

Discussion on Greedy Accelerated Coordinate Descent.

The convergence rate for greedy accelerated coordinate descent can directly be obtained from the rate from accelerated matching pursuit. Let the atom set \mathcal{A} consist of the standard basis vectors $\{\mathbf{e}_i, i \in [n]\}$ and \mathcal{Z} be a uniform distribution over this set. Then algorithm 3 reduces to the accelerated randomized coordinate method (ACDM) of (Lee & Sidford, 2013; Nesterov & Stich, 2017) and we recover their rates. Instead if we use algorithm 4, we obtain a novel accelerated greedy coordinate method with a (potentially) better convergence rate.¹

Lemma 10. *When $\mathcal{A} = \{\mathbf{e}_i, i \in [n]\}$ and \mathcal{Z} is a uniform distribution over \mathcal{A} , then $\mathbf{P} = n\mathbf{I}$, $\nu' = n$ and $\nu \in [1, n]$.*

4. Empirical Evaluation

In this section we aim at empirically validate our theoretical findings. In both experiments we use 1 and the intrinsic dimensionality of $\text{lin}(\mathcal{A})$ as ν and ν' respectively. Note that a value of ν smaller than ν' represents the best case for the steepest update. We implicitly assume that the worst case in which a random update is as good as the steepest one never happens.

Toy Data: First, we report the function value while minimizing the squared distance between the a random 100 dimensional signal with both positive and negative entries and its sparse representation in terms of atoms. We sample a random dictionary containing 200 atoms which we then make symmetric. The result is depicted in Figure 1. As anticipated from our analysis, the accelerated schemes converge much faster than the non-accelerated variants. Furthermore, in both cases the steepest update converge faster than the random one, due to a better dependency on the dimensionality of the space.

Real Data: We use the under-sampled Urban HDI Dataset from which we extract the dictionary of atoms using the hierarchical clustering approached of (Gillis et al., 2015). This dataset contains 5'929 pixels, each associated with 162 hyperspectral features. The number of dictionary elements is 6, motivated by the fact that 6 different physical materials are depicted in this HSI data (Gillis & Luce, 2018). We approximate each pixel with a linear combination of the dictionary elements by minimizing the square distance between the observed pixel and our approximation. We report in Figure 2 the loss as an average across all the pixels:

$$\min_{\mathbf{x}_i \in \text{lin}(\mathcal{A})} \frac{1}{N} \sum_{i=1}^N \|\mathbf{x}_i - \mathbf{b}_i\|^2$$

We notice that as expected, the steepest matching pursuit

¹Simultaneously (and independently) (Lu et al., 2018) derived the same accelerated greedy coordinate algorithm.




Figure 1. loss for synthetic data




Figure 2. loss for hyperspectral data

converges faster than the random pursuit, but as expected both of them converge at the same regime. On the other hand, the accelerated scheme converge much faster than the non-accelerated variants. Note that the acceleration kicks in only after a few iterations as the accelerated rate has a worse dependency on the intrinsic dimensionality of the linear span than the non accelerated algorithms. We notice that the speedup of steepest MP is much more evident in the synthetic data. The reason is that this experiment is much more high dimensional than the hyperspectral data. Indeed, the span of the dictionary is a 6 dimensional manifold in the latter and the full ambient space in the former and the steepest update yields a better dependency on the dimensionality.

5. Conclusions

In this paper we presented a unified analysis of matching pursuit and coordinate descent algorithms. As a consequence, we exploit the similarity between the two to obtain the best of both worlds: tight sublinear and linear rates for steepest coordinate descent and the first accelerated rate for matching pursuit and steepest coordinate descent. Furthermore, we discussed the relation between the steepest and the random directions by viewing the latter as an approximate version of the former. An affine invariant accelerated proof remains an open problem.

Acknowledgements

FL is supported by Max Planck/ETH Center for Learning Systems and partially supported by ETH core funding (to GR). SPK and MJ acknowledge support by SNSF project 175796, and SUS by the Microsoft Research Swiss JRC.

References

- Allen-Zhu, Z. and Orecchia, L. Linear Coupling of Gradient and Mirror Descent: A Novel, Simple Interpretation of Nesterov’s Accelerated Method. *arXiv.org*, July 2014.
- Allen-Zhu, Z., Qu, Z., Richtárik, P., and Yuan, Y. Even faster accelerated coordinate descent using non-uniform sampling. In *ICML 2016 - Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *PMLR*, pp. 1110–1119. PMLR, 20–22 Jun 2016.
- Chandrasekaran, V., Recht, B., Parrilo, P. A., and Willsky, A. S. The Convex Geometry of Linear Inverse Problems. *Foundations of Computational Mathematics*, 12(6):805–849, October 2012.
- Chen, S., Billings, S. A., and Luo, W. Orthogonal least squares methods and their application to non-linear system identification. *International Journal of control*, 50(5):1873–1896, 1989.
- d’Aspremont, A., Guzmán, C., and Jaggi, M. Optimal affine invariant smooth minimization algorithms. *SIAM Journal on Optimization (and arXiv:1301.0465)*, 2018.
- Davenport, M. A. and Wakin, M. B. Analysis of orthogonal matching pursuit using the restricted isometry property. *IEEE Transactions on Information Theory*, 56(9):4395–4401, 2010.
- Dennis, Jr., J. and Torczon, V. Direct Search methods on parallel machines. *SIAM Journal on Optimization*, 1(4): 448–474, 1991. doi: 10.1137/0801027.
- El Halabi, M., Hsieh, Y.-P., Vu, B., Nguyen, Q., and Cevher, V. General proximal gradient method: A case for non-euclidean norms. Technical report, 2017.
- Frank, M. and Wolfe, P. An Algorithm for Quadratic Programming. *Naval Research Logistics Quarterly*, 3:95–110, 1956.
- Gillis, N. and Luce, R. A fast gradient method for non-negative sparse regression with self-dictionary. *IEEE Transactions on Image Processing*, 27(1):24–37, 2018.
- Gillis, N., Kuang, D., and Park, H. Hierarchical clustering of hyperspectral images using rank-two nonnegative matrix factorization. *IEEE Transactions on Geoscience and Remote Sensing*, 53(4):2066–2078, 2015.
- Gribonval, R. and Vandergheynst, P. On the exponential convergence of matching pursuits in quasi-incoherent dictionaries. *IEEE Transactions on Information Theory*, 52(1):255–261, 2006.
- Holloway, C. A. An extension of the frank and Wolfe method of feasible directions. *Mathematical Programming*, 6(1):14–27, 1974.
- Hooke, R. and Jeeves, T. A. “Direct Search” solution of numerical and statistical problems. *Journal of the ACM*, 8(2):212–229, 1961. ISSN 0004-5411. doi: 10.1145/321062.321069.
- Jaggi, M. Revisiting Frank-Wolfe: Projection-Free Sparse Convex Optimization. In *ICML*, pp. 427–435, 2013.
- Karmanov, V. G. On convergence of a random search method in convex minimization problems. *Theory of Probability and its applications*, 19(4):788–794, 1974a. (in Russian).
- Karmanov, V. G. Convergence estimates for iterative minimization methods. *USSR Computational Mathematics and Mathematical Physics*, 14(1):1–13, 1974b. ISSN 0041-5553.
- Kerdreux, T., Pedregosa, F., and d’Aspremont, A. Frank-wolfe with subsampling oracle. In *ICML 2018 - Proceedings of the 35th International Conference on Machine Learning*, 2018.
- Lacoste-Julien, S. and Jaggi, M. On the Global Linear Convergence of Frank-Wolfe Optimization Variants. In *NIPS 2015*, pp. 496–504, 2015.
- Lacoste-Julien, S., Jaggi, M., Schmidt, M., and Pletscher, P. Block-Coordinate Frank-Wolfe Optimization for Structural SVMs. In *ICML 2013 - Proceedings of the 30th International Conference on Machine Learning*, 2013.
- Lee, Y. T. and Sidford, A. Efficient accelerated coordinate descent methods and faster algorithms for solving linear systems. In *FOCS ’13 - Proceedings of the 2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, FOCS ’13, pp. 147–156, 2013.
- Locatello, F., Khanna, R., Tschannen, M., and Jaggi, M. A unified optimization view on generalized matching pursuit and frank-wolfe. In *AISTATS - Proc. International Conference on Artificial Intelligence and Statistics*, 2017a.
- Locatello, F., Tschannen, M., Rätsch, G., and Jaggi, M. Greedy algorithms for cone constrained optimization with convergence guarantees. In *NIPS - Advances in Neural Information Processing Systems 30*, 2017b.

- Lu, H., Freund, R. M., and Mirrokni, V. Accelerating greedy coordinate descent methods. In *ICML 2018 - Proceedings of the 35th International Conference on Machine Learning*, 2018.
- Mallat, S. and Zhang, Z. Matching pursuits with time-frequency dictionaries. *IEEE Transactions on Signal Processing*, 41(12):3397–3415, 1993.
- Meir, R. and Rätsch, G. An introduction to boosting and leveraging. In *Advanced lectures on machine learning*, pp. 118–183. Springer, 2003.
- Menard, S. *Applied logistic regression analysis*, volume 106. SAGE publications, 2018.
- Mutseniyeks, V. A. and Rastrigin, L. A. Extremal control of continuous multi-parameter systems by the method of random search. *Eng. Cybernetics*, 1:82–90, 1964.
- Nesterov, Y. A method of solving a convex programming problem with convergence rate $O(1/K^2)$. *Soviet Mathematics Doklady*, 27:372–376, 1983.
- Nesterov, Y. *Introductory Lectures on Convex Optimization*, volume 87 of *Applied Optimization*. Springer US, Boston, MA, 2004.
- Nesterov, Y. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optimization*, 22(2):341–362, 2012.
- Nesterov, Y. and Stich, S. U. Efficiency of the accelerated coordinate descent method on structured optimization problems. *SIAM Journal on Optimization*, 27(1):110–123, 2017.
- Nguyen, H. and Petrova, G. Greedy strategies for convex optimization. *Calcolo*, pp. 1–18, 2014.
- Nutini, J., Schmidt, M., Laradji, I., Friedlander, M., and Koepke, H. Coordinate Descent Converges Faster with the Gauss-Southwell Rule Than Random Selection. In *ICML 2015 - Proceedings of the 32th International Conference on Machine Learning*, pp. 1632–1641, 2015.
- Pena, J. and Rodriguez, D. Polytope conditioning and linear convergence of the frank-wolfe algorithm. *arXiv preprint arXiv:1512.06142*, 2015.
- Rätsch, G., Mika, S., Warmuth, M. K., et al. On the convergence of leveraging. In *NIPS*, pp. 487–494, 2001.
- Schölkopf, B. and Smola, A. J. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2001.
- Seber, G. A. and Lee, A. J. *Linear regression analysis*, volume 329. John Wiley & Sons, 2012.
- Song, C., Cui, S., Jiang, Y., and Xia, S.-T. Accelerated stochastic greedy coordinate descent by soft thresholding projection onto simplex. In *NIPS - Advances in Neural Information Processing Systems*, pp. 4841–4850, 2017.
- Stich, S. U. *Convex Optimization with Random Pursuit*. PhD thesis, ETH Zurich, 2014. Nr. 22111.
- Stich, S. U., Müller, C. L., and Gärtner, B. Optimization of convex functions with random pursuit. *SIAM Journal on Optimization*, 23(2):1284–1309, 2013.
- Stich, S. U., Raj, A., and Jaggi, M. Approximate steepest coordinate descent. In *ICML 2017 - Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *PMLR*, pp. 3251–3259, 2017.
- Temlyakov, V. Chebushev Greedy Algorithm in convex optimization. *arXiv.org*, December 2013.
- Temlyakov, V. Greedy algorithms in convex optimization on Banach spaces. In *48th Asilomar Conference on Signals, Systems and Computers*, pp. 1331–1335. IEEE, 2014.
- Tibshirani, R. J. A general framework for fast stagewise algorithms. *Journal of Machine Learning Research*, 16: 2543–2588, 2015.
- Torczon, V. On the convergence of pattern search algorithms. *SIAM Journal on optimization*, 7(1):1–25, 1997.
- Tropp, J. A. Greed is good: algorithmic results for sparse approximation. *IEEE Transactions on Information Theory*, 50(10):2231–2242, 2004.
- Zieliński, R. and Neumann, P. *Stochastische Verfahren zur Suche nach dem Minimum einer Funktion*. Akademie-Verlag, Berlin, Germany, 1983.

Appendix

A. Sublinear Rates

Theorem' 2. Assume f is L -smooth w.r.t. a given norm $\|\cdot\|$, over $\text{lin}(\mathcal{A})$ where \mathcal{A} is symmetric. Then,

$$L_{\mathcal{A}} \leq L \text{radius}_{\|\cdot\|}(\mathcal{A})^2. \quad (9)$$

Proof. Let $D(\mathbf{y}, \mathbf{x}) := f(\mathbf{y}) - f(\mathbf{x}) + \gamma \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle$ By the definition of smoothness of f w.r.t. $\|\cdot\|$,

$$D(\mathbf{y}, \mathbf{x}) \leq \frac{L}{2} \|\mathbf{y} - \mathbf{x}\|^2.$$

Hence, from the definition of $L_{\mathcal{A}}$,

$$\begin{aligned} L_{\mathcal{A}} &\leq \sup_{\substack{\mathbf{x}, \mathbf{y} \in \text{lin}(\mathcal{A}) \\ \mathbf{y} = \mathbf{x} + \gamma \mathbf{z} \\ \|\mathbf{z}\|_{\mathcal{A}} = 1, \gamma \in \mathbb{R}_{>0}}} \frac{2}{\gamma^2} \frac{L}{2} \|\mathbf{y} - \mathbf{x}\|^2 \\ &= L \sup_{\mathbf{z} \text{ s.t. } \|\mathbf{z}\|_{\mathcal{A}} = 1} \|\mathbf{z}\|^2 \\ &= L \text{radius}_{\|\cdot\|}(\mathcal{A})^2. \quad \square \end{aligned}$$

The definition of the smoothness constant w.r.t. the atomic norm yields the following quadratic upper bound:

$$L_{\mathcal{A}} = \sup_{\substack{\mathbf{x}, \mathbf{y} \in \text{lin}(\mathcal{A}) \\ \mathbf{y} = \mathbf{x} + \gamma \mathbf{z} \\ \|\mathbf{z}\|_{\mathcal{A}} = 1, \gamma \in \mathbb{R}_{>0}}} \frac{2}{\gamma^2} [f(\mathbf{y}) - f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle]. \quad (10)$$

Furthermore, let:

$$R_{\mathcal{A}}^2 = \max_{\substack{\mathbf{x} \in \text{lin}(\mathcal{A}) \\ f(\mathbf{x}) \leq f(\mathbf{x}_0)}} \|\mathbf{x} - \mathbf{x}^*\|_{\mathcal{A}}^2. \quad (11)$$

Now, we show that the algorithm we presented is affine invariant. An optimization method is called *affine invariant* if it is invariant under affine transformations of the input problem: If one chooses any re-parameterization of the domain \mathcal{Q} by a *surjective* linear or affine map $\mathbf{M} : \hat{\mathcal{Q}} \rightarrow \mathcal{Q}$, then the ‘‘old’’ and ‘‘new’’ optimization problems $\min_{\mathbf{x} \in \mathcal{Q}} f(\mathbf{x})$ and $\min_{\hat{\mathbf{x}} \in \hat{\mathcal{Q}}} \hat{f}(\hat{\mathbf{x}})$ for $\hat{f}(\hat{\mathbf{x}}) := f(\mathbf{M}\hat{\mathbf{x}})$ look the same to the algorithm. Note that $\nabla \hat{f} = \mathbf{M}^T \nabla f$.

First of all, let us note that $L_{\mathcal{A}}$ is affine invariant as it does not depend on any norm. Now:

$$\begin{aligned} \mathbf{M}\hat{\mathbf{x}}_{t+1} &= \mathbf{M} \left(\hat{\mathbf{x}}_t + \frac{\langle \nabla \hat{f}(\hat{\mathbf{x}}_t), \hat{\mathbf{z}}_t \rangle}{L_{\mathcal{A}}} \hat{\mathbf{z}}_t \right) \\ &= \mathbf{M}\hat{\mathbf{x}}_t + \frac{\langle \nabla \hat{f}(\hat{\mathbf{x}}_t), \hat{\mathbf{z}}_t \rangle}{L_{\mathcal{A}}} \mathbf{M}\hat{\mathbf{z}}_t \\ &= \mathbf{x}_t + \frac{\langle \nabla \hat{f}(\hat{\mathbf{x}}_t), \hat{\mathbf{z}}_t \rangle}{L_{\mathcal{A}}} \mathbf{z}_t \\ &= \mathbf{x}_t + \frac{\langle \mathbf{M}^T \nabla f(\mathbf{x}_t), \hat{\mathbf{z}}_t \rangle}{L_{\mathcal{A}}} \mathbf{z}_t \\ &= \mathbf{x}_t + \frac{\langle \nabla f(\mathbf{x}_t), \mathbf{M}\hat{\mathbf{z}}_t \rangle}{L_{\mathcal{A}}} \mathbf{z}_t \\ &= \mathbf{x}_t + \frac{\langle \nabla f(\mathbf{x}_t), \mathbf{z}_t \rangle}{L_{\mathcal{A}}} \mathbf{z}_t \\ &= \mathbf{x}_{t+1}. \end{aligned}$$

Therefore the algorithm is affine invariant.

A.1. Affine Invariant Sublinear Rate

Theorem' 3. *Let $\mathcal{A} \subset \mathcal{H}$ be a closed and bounded set. We assume that $\|\cdot\|_{\mathcal{A}}$ is a norm over $\text{lin}(\mathcal{A})$. Let f be convex and $L_{\mathcal{A}}$ -smooth w.r.t. the norm $\|\cdot\|_{\mathcal{A}}$ over $\text{lin}(\mathcal{A})$, and let $R_{\mathcal{A}}$ be the radius of the level set of \mathbf{x}_0 measured with the atomic norm. Then, Algorithm 2 converges for $t \geq 0$ as*

$$f(\mathbf{x}_{t+1}) - f(\mathbf{x}^*) \leq \frac{2L_{\mathcal{A}}R_{\mathcal{A}}^2}{\delta^2(t+2)},$$

where $\delta \in (0, 1]$ is the relative accuracy parameter of the employed approximate LMO (3).

Proof. Recall that $\tilde{\mathbf{z}}_t$ is the atom selected in iteration t by the approximate LMO defined in (3). We start by upper-bounding f using the definition of $L_{\mathcal{A}}$ as follows:

$$\begin{aligned} f(\mathbf{x}_{t+1}) &\leq \min_{\gamma \in \mathbb{R}} f(\mathbf{x}_t) + \gamma \langle \nabla f(\mathbf{x}_t), \tilde{\mathbf{z}}_t \rangle + \frac{\gamma^2}{2} L_{\mathcal{A}} \|\mathbf{z}\|_{\mathcal{A}}^2 \\ &= \min_{\gamma \in \mathbb{R}} f(\mathbf{x}_t) + \gamma \langle \nabla f(\mathbf{x}_t), \tilde{\mathbf{z}}_t \rangle + \frac{\gamma^2}{2} L_{\mathcal{A}} \\ &\leq f(\mathbf{x}_t) - \frac{\langle \nabla f(\mathbf{x}_t), \tilde{\mathbf{z}}_t \rangle^2}{2L_{\mathcal{A}}} \\ &= f(\mathbf{x}_t) - \frac{\langle \nabla_{\parallel} f(\mathbf{x}_t), \tilde{\mathbf{z}}_t \rangle^2}{2L_{\mathcal{A}}} \\ &\leq f(\mathbf{x}_t) - \delta^2 \frac{\langle \nabla_{\parallel} f(\mathbf{x}_t), \mathbf{z}_t \rangle^2}{2L_{\mathcal{A}}}. \end{aligned}$$

Where $\nabla_{\parallel} f$ is the parallel component of the gradient wrt the linear span of \mathcal{A} . Note that $\|\mathbf{d}\|_{\mathcal{A}^*} := \sup \{ \langle \mathbf{z}, \mathbf{d} \rangle, \mathbf{z} \in \mathcal{A} \}$ is the dual of the atomic norm. Therefore, by definition:

$$\langle \nabla_{\parallel} f(\mathbf{x}_t), \mathbf{z}_t \rangle^2 = \| -\nabla_{\parallel} f(\mathbf{x}_t) \|_{\mathcal{A}^*}^2,$$

which gives:

$$\begin{aligned} f(\mathbf{x}_{t+1}) &\leq f(\mathbf{x}_t) - \delta^2 \frac{1}{2L_{\mathcal{A}}} \|\nabla_{\parallel} f(\mathbf{x}_t)\|_{\mathcal{A}^*}^2 \\ &\leq f(\mathbf{x}_t) - \delta^2 \frac{1}{2L_{\mathcal{A}}} \frac{(-\langle \nabla_{\parallel} f(\mathbf{x}_t), \mathbf{x}_t - \mathbf{x}^* \rangle)^2}{R_{\mathcal{A}}^2} \\ &= f(\mathbf{x}_t) - \delta^2 \frac{1}{2L_{\mathcal{A}}} \frac{(\langle \nabla_{\parallel} f(\mathbf{x}_t), \mathbf{x}_t - \mathbf{x}^* \rangle)^2}{R_{\mathcal{A}}^2} \\ &\leq f(\mathbf{x}_t) - \delta^2 \frac{1}{2L_{\mathcal{A}}} \frac{(f(\mathbf{x}_t) - f(\mathbf{x}^*))^2}{R_{\mathcal{A}}^2}, \end{aligned}$$

where the second inequality is Cauchy-Schwarz and the third one is convexity. Which gives:

$$\epsilon_{t+1} \leq \frac{2L_{\mathcal{A}}R_{\mathcal{A}}^2}{\delta^2(t+2)}. \quad \square$$

A.2. Randomized Affine Invariant Sublinear Rate

For random sampling of \mathbf{z} from a distribution over \mathcal{A} , let

$$\hat{\delta}^2 := \min_{\mathbf{d} \in \text{lin} \mathcal{A}} \frac{\mathbb{E}_{\mathbf{z} \in \mathcal{A}} \langle \mathbf{d}, \mathbf{z} \rangle^2}{\|\mathbf{d}\|_{\mathcal{A}^*}^2}. \quad (12)$$

Theorem' 4. Let $\mathcal{A} \subset \mathcal{H}$ be a closed and bounded set. We assume that $\|\cdot\|_{\mathcal{A}}$ is a norm. Let f be convex and $L_{\mathcal{A}}$ smooth w.r.t. the norm $\|\cdot\|_{\mathcal{A}}$ over $\text{lin}(\mathcal{A})$ and let $R_{\mathcal{A}}$ be the radius of the level set of \mathbf{x}_0 measured with the atomic norm. Then, Algorithm 2 converges for $t \geq 0$ as

$$\mathbb{E}_{\mathbf{z}} [f(\mathbf{x}_{t+1})] - f(\mathbf{x}^*) \leq \frac{2L_{\mathcal{A}}R_{\mathcal{A}}^2}{\hat{\delta}^2(t+2)},$$

when the LMO is replaced with random sampling of \mathbf{z} from a distribution over \mathcal{A} .

Proof. Recall that $\bar{\mathbf{z}}_t$ is the atom selected in iteration t by the approximate LMO defined in (3). We start by upper-bounding f using the definition of $L_{\mathcal{A}}$ as follows

$$\begin{aligned} \mathbb{E}_{\mathbf{z}} f(\mathbf{x}_{t+1}) &\leq \mathbb{E}_{\mathbf{z}} \left[\min_{\gamma \in \mathbb{R}} f(\mathbf{x}_t) + \gamma \langle \nabla f(\mathbf{x}_t), \mathbf{z} \rangle + \frac{\gamma^2}{2} L_{\mathcal{A}} \|\mathbf{z}\|_{\mathcal{A}}^2 \right] \\ &= \mathbb{E}_{\mathbf{z}} \left[\min_{\gamma \in \mathbb{R}} f(\mathbf{x}_t) + \gamma \langle \nabla f(\mathbf{x}_t), \mathbf{z} \rangle + \frac{\gamma^2}{2} L_{\mathcal{A}} \right] \\ &\leq f(\mathbf{x}_t) - \frac{\mathbb{E}_{\mathbf{z}} [\langle \nabla f(\mathbf{x}_t), \mathbf{z} \rangle^2]}{2L_{\mathcal{A}}} \\ &= f(\mathbf{x}_t) - \frac{\mathbb{E}_{\mathbf{z}} [\langle \nabla_{\parallel} f(\mathbf{x}_t), \mathbf{z} \rangle^2]}{2L_{\mathcal{A}}} \\ &\leq f(\mathbf{x}_t) - \hat{\delta}^2 \frac{\langle \nabla_{\parallel} f(\mathbf{x}_t), \mathbf{z}_t \rangle^2}{2L_{\mathcal{A}}}. \end{aligned}$$

The rest of the proof proceeds as in Theorem 3. □

B. Linear Rates

B.1. Affine Invariant Linear Rate

Let us first the fine the affine invariant notion of strong convexity based on the atomic norm:

$$\mu_{\mathcal{A}} := \inf_{\substack{\mathbf{x}, \mathbf{y} \in \text{lin}(\mathcal{A}) \\ \mathbf{x} \neq \mathbf{y}}} \frac{2}{\|\mathbf{y} - \mathbf{x}\|_{\mathcal{A}}^2} [f(\mathbf{y}) - f(\mathbf{x}) - \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle].$$

Let us recall the definition of *minimal directional width* from (Locatello et al., 2017a):

$$\text{mDW}(\mathcal{A}) := \min_{\substack{\mathbf{d} \in \text{lin}(\mathcal{A}) \\ \mathbf{d} \neq \mathbf{0}}} \max_{\mathbf{z} \in \mathcal{A}} \left(\frac{\mathbf{d}}{\|\mathbf{d}\|}, \mathbf{z} \right).$$

Then, we can relate our new definition of strong convexity with the $\text{mDW}(\mathcal{A})$ as follows.

Theorem' 6. Assume f is μ strongly convex wrt a given norm $\|\cdot\|$ over $\text{lin}(\mathcal{A})$ and \mathcal{A} is symmetric. Then:

$$\mu_{\mathcal{A}} \geq \text{mDW}(\mathcal{A})^2 \mu.$$

Proof. First of all, note that for any $\mathbf{x}, \mathbf{y} \in \text{lin}(\mathcal{A})$ with $\mathbf{x} \neq \mathbf{y}$ we have that:

$$\langle \nabla f(\mathbf{x}), \mathbf{x} - \mathbf{y} \rangle^2 \leq \|\nabla f(\mathbf{x})\|_{\mathcal{A}^*}^2 \|\mathbf{x} - \mathbf{y}\|_{\mathcal{A}}^2.$$

Therefore:

$$\begin{aligned}
 \mu_{\mathcal{A}} &= \inf_{\substack{\mathbf{x}, \mathbf{y} \in \text{lin } \mathcal{A} \\ \mathbf{x} \neq \mathbf{y}}} \frac{2}{\|\mathbf{y} - \mathbf{x}\|_{\mathcal{A}}^2} D(\mathbf{x}, \mathbf{y}) \\
 &\geq \inf_{\substack{\mathbf{x}, \mathbf{y}, \mathbf{d} \in \text{lin } \mathcal{A} \\ \mathbf{x} \neq \mathbf{y}, \mathbf{d} \neq \mathbf{0}}} \frac{\|\mathbf{d}\|_{\mathcal{A}^*}^2}{\langle \mathbf{d}, \mathbf{x} - \mathbf{y} \rangle^2} 2D(\mathbf{x}, \mathbf{y}) \\
 &\geq \inf_{\substack{\mathbf{x}, \mathbf{y}, \mathbf{d} \in \text{lin } \mathcal{A} \\ \mathbf{x} \neq \mathbf{y}, \mathbf{d} \neq \mathbf{0}}} \frac{\|\mathbf{d}\|_{\mathcal{A}^*}^2}{\langle \mathbf{d}, \mathbf{x} - \mathbf{y} \rangle^2} \mu \|\mathbf{x} - \mathbf{y}\|^2 \\
 &\geq \inf_{\substack{\mathbf{x}, \mathbf{y}, \mathbf{d} \in \text{lin } \mathcal{A} \\ \mathbf{x} \neq \mathbf{y}, \mathbf{d} \neq \mathbf{0}}} \frac{\|\mathbf{d}\|_{\mathcal{A}^*}^2}{\langle \mathbf{d}, \frac{\mathbf{x} - \mathbf{y}}{\|\mathbf{x} - \mathbf{y}\|} \rangle^2} \mu \\
 &\geq \inf_{\substack{\mathbf{x}, \mathbf{y}, \mathbf{d} \in \text{lin } \mathcal{A} \\ \mathbf{x} \neq \mathbf{y}, \mathbf{d} \neq \mathbf{0}}} \frac{\|\mathbf{d}\|_{\mathcal{A}^*}^2}{\|\mathbf{d}\|^2} \mu \\
 &\geq \inf_{\substack{\mathbf{d} \in \text{lin } \mathcal{A} \\ \mathbf{d} \neq \mathbf{0}}} \max_z \frac{\langle \mathbf{d}, \mathbf{z} \rangle^2}{\|\mathbf{d}\|^2} \mu \\
 &= \text{mDW}(\mathcal{A})^2 \mu. \quad \square
 \end{aligned}$$

Theorem 5. (Part 1). Let $\mathcal{A} \subset \mathcal{H}$ be a closed and bounded set. We assume that $\|\cdot\|_{\mathcal{A}}$ is a norm. Let f be $\mu_{\mathcal{A}}$ -strongly convex and $L_{\mathcal{A}}$ -smooth w.r.t. the norm $\|\cdot\|_{\mathcal{A}}$, both over $\text{lin}(\mathcal{A})$. Then, Algorithm 2 converges for $t \geq 0$ as

$$\epsilon_{t+1} \leq \left(1 - \delta^2 \frac{\mu_{\mathcal{A}}}{L_{\mathcal{A}}}\right) \epsilon_t.$$

where $\epsilon_t := f(\mathbf{x}_t) - f(\mathbf{x}^*)$.

Proof. Recall that $\tilde{\mathbf{z}}_t$ is the atom selected in iteration t by the approximate LMO defined in (3). We start by upper-bounding f using the definition of $L_{\mathcal{A}}$ as follows

$$\begin{aligned}
 f(\mathbf{x}_{t+1}) &\leq \min_{\gamma \in \mathbb{R}} f(\mathbf{x}_t) + \gamma \langle \nabla f(\mathbf{x}_t), \tilde{\mathbf{z}}_t \rangle + \frac{\gamma^2}{2} L_{\mathcal{A}} \|\mathbf{z}\|_{\mathcal{A}}^2 \\
 &= \min_{\gamma \in \mathbb{R}} f(\mathbf{x}_t) + \gamma \langle \nabla f(\mathbf{x}_t), \tilde{\mathbf{z}}_t \rangle + \frac{\gamma^2}{2} L_{\mathcal{A}} \\
 &\leq f(\mathbf{x}_t) - \frac{\langle \nabla f(\mathbf{x}_t), \tilde{\mathbf{z}}_t \rangle^2}{2L_{\mathcal{A}}} \\
 &= f(\mathbf{x}_t) - \frac{\langle \nabla_{\parallel} f(\mathbf{x}_t), \tilde{\mathbf{z}}_t \rangle^2}{2L_{\mathcal{A}}} \\
 &\leq f(\mathbf{x}_t) - \delta^2 \frac{\langle \nabla_{\parallel} f(\mathbf{x}_t), \mathbf{z}_t \rangle^2}{2L_{\mathcal{A}}} \\
 &= f(\mathbf{x}_t) - \delta^2 \frac{\langle -\nabla_{\parallel} f(\mathbf{x}_t), \mathbf{z}_t \rangle^2}{2L_{\mathcal{A}}}.
 \end{aligned}$$

Where $\|\mathbf{d}\|_{\mathcal{A}^*} := \sup \{ \langle \mathbf{z}, \mathbf{d} \rangle, \mathbf{z} \in \mathcal{A} \}$ is the dual of the atomic norm. Therefore, by definition:

$$\langle -\nabla_{\parallel} f(\mathbf{x}_t), \mathbf{z}_t \rangle^2 = \|\nabla_{\parallel} f(\mathbf{x}_t)\|_{\mathcal{A}^*}^2,$$

which gives:

$$f(\mathbf{x}_{t+1}) \leq f(\mathbf{x}_t) - \delta^2 \frac{1}{2L_{\mathcal{A}}} \|\nabla_{\parallel} f(\mathbf{x}_t)\|_{\mathcal{A}^*}^2.$$

From strong convexity we have that:

$$f(\mathbf{y}) \geq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{\mu_{\mathcal{A}}}{2} \|\mathbf{y} - \mathbf{x}\|_{\mathcal{A}}^2.$$

Fixing $\mathbf{y} = \mathbf{x}_t + \gamma(\mathbf{x}^* - \mathbf{x}_t)$ and $\gamma = 1$ in the LHS and minimizing the RHS we obtain:

$$\begin{aligned} f(\mathbf{x}^*) &\geq f(\mathbf{x}_t) - \frac{1}{2\mu_{\mathcal{A}}} \frac{\langle \nabla f(\mathbf{x}_t), \mathbf{x}^* - \mathbf{x}_t \rangle}{\|\mathbf{x}^* - \mathbf{x}_t\|_{\mathcal{A}}^2} \\ &\geq f(\mathbf{x}_t) - \frac{1}{2\mu_{\mathcal{A}}} \|\nabla_{\parallel} f(\mathbf{x}_t)\|_{\mathcal{A}^*}^2, \end{aligned}$$

where the last inequality is obtained by the fact that $\langle \nabla f(\mathbf{x}_t), \mathbf{x}^* - \mathbf{x}_t \rangle = \langle \nabla_{\parallel} f(\mathbf{x}_t), \mathbf{x}^* - \mathbf{x}_t \rangle$ and Cauchy-Schwartz. Therefore:

$$\|\nabla f(\mathbf{x}_t)\|_{\mathcal{A}^*} \geq 2\epsilon_t \mu_{\mathcal{A}},$$

which yields:

$$\epsilon_{t+1} \leq \epsilon_t - \delta^2 \frac{\mu_{\mathcal{A}}}{L_{\mathcal{A}}} \epsilon_t. \quad \square$$

B.2. Randomized Affine Invariant Linear Rate

Theorem' 5. (Part 2). Let $\mathcal{A} \subset \mathcal{H}$ be a closed and bounded set. We assume that $\|\cdot\|_{\mathcal{A}}$ is a norm. Let f be $\mu_{\mathcal{A}}$ -strongly convex and $L_{\mathcal{A}}$ -smooth w.r.t. the norm $\|\cdot\|_{\mathcal{A}}$, both over $\text{lin}(\mathcal{A})$. Then, Algorithm 2 converges for $t \geq 0$ as

$$\mathbb{E}_{\mathbf{z}} [\epsilon_{t+1} | \mathbf{x}_t] \leq \left(1 - \hat{\delta}^2 \frac{\mu_{\mathcal{A}}}{L_{\mathcal{A}}}\right) \epsilon_t,$$

where $\epsilon_t := f(\mathbf{x}_t) - f(\mathbf{x}^*)$, and the LMO direction \mathbf{z} is sampled randomly from \mathcal{A} , from the same distribution as used in the definition of $\hat{\delta}$.

Proof. We start by upper-bounding f using the definition of $L_{\mathcal{A}}$ as follows

$$\begin{aligned} \mathbb{E}_{\mathbf{z}} [f(\mathbf{x}_{t+1})] &\leq \mathbb{E}_{\mathbf{z}} \left[\min_{\gamma \in \mathbb{R}} f(\mathbf{x}_t) + \gamma \langle \nabla f(\mathbf{x}_t), \tilde{\mathbf{z}}_t \rangle + \frac{\gamma^2}{2} L_{\mathcal{A}} \|\mathbf{z}\|_{\mathcal{A}}^2 \right] \\ &= \mathbb{E}_{\mathbf{z}} \left[\min_{\gamma \in \mathbb{R}} f(\mathbf{x}_t) + \gamma \langle \nabla f(\mathbf{x}_t), \tilde{\mathbf{z}}_t \rangle + \frac{\gamma^2}{2} L_{\mathcal{A}} \right] \\ &\leq f(\mathbf{x}_t) - \mathbb{E}_{\mathbf{z}} \left[\frac{\langle \nabla f(\mathbf{x}_t), \tilde{\mathbf{z}}_t \rangle^2}{2L_{\mathcal{A}}} \right] \\ &\leq f(\mathbf{x}_t) - \hat{\delta}^2 \frac{\langle \nabla f(\mathbf{x}_t), \mathbf{z}_t \rangle^2}{2L_{\mathcal{A}}} \\ &= f(\mathbf{x}_t) - \hat{\delta}^2 \frac{\langle \nabla_{\parallel} f(\mathbf{x}_t), \tilde{\mathbf{z}}_t \rangle^2}{2L_{\mathcal{A}}} \\ &= f(\mathbf{x}_t) - \hat{\delta}^2 \frac{\langle -\nabla_{\parallel} f(\mathbf{x}_t), \tilde{\mathbf{z}}_t \rangle^2}{2L_{\mathcal{A}}}. \end{aligned}$$

The rest of the proof proceeds as in Part 1 of the proof of Theorem 5. □

C. Accelerated Matching Pursuit

Our proof follows the technique for acceleration given in (Lee & Sidford, 2013; Nesterov & Stich, 2017; Nesterov, 2004; Stich et al., 2013)

C.1. Proof of Convergence

We define $\|\mathbf{x}\|_{\mathbf{P}}^2 = \mathbf{x}^{\top} \mathbf{P} \mathbf{x}$. We start our proof by first defining the model function ψ_t . For $t = 0$, we define :

$$\psi_0(\mathbf{x}) = \frac{1}{2} \|\mathbf{x} - \mathbf{v}_0\|_{\mathbf{P}}^2.$$

Then for $t > 1$, ψ_t is inductively defined as

$$\psi_{t+1}(\mathbf{x}) = \psi_t(\mathbf{x}) + \alpha_{t+1} \left(f(\mathbf{y}_t) + \langle \tilde{\mathbf{z}}_t^{\top} \nabla f(\mathbf{y}_t), \tilde{\mathbf{z}}_t^{\top} \mathbf{P}(\mathbf{x} - \mathbf{y}_t) \rangle \right). \quad (13)$$

Algorithm 5 Accelerated Matching Pursuit

```

1: init  $\mathbf{x}_0 = \mathbf{v}_0 = \mathbf{y}_0$ ,  $\beta_0 = 0$ , and  $\nu$ 
2: for  $t = 0, 1 \dots T$ 
3:   Solve  $\alpha_{t+1}^2 L\nu = \beta_t + \alpha_{t+1}$ 
4:    $\beta_{t+1} = \beta_t + \alpha_{t+1}$ 
5:    $\tau_t = \frac{\alpha_{t+1}}{\beta_{t+1}}$ 
6:   Compute  $\mathbf{y}_t = (1 - \tau_t)\mathbf{x}_t + \tau_t\mathbf{v}_t$ 
7:   Find  $\mathbf{z}_t := \text{LMO}_{\mathcal{A}}(\nabla f(\mathbf{y}_t))$ 
8:    $\mathbf{x}_{t+1} = \mathbf{y}_t - \frac{\langle \nabla f(\mathbf{y}_t), \mathbf{z}_t \rangle}{L\|\mathbf{z}_t\|_2^2} \mathbf{z}_t$ 
9:   sample  $\tilde{\mathbf{z}}_t \sim \mathcal{Z}$ 
10:   $\mathbf{v}_{t+1} = \mathbf{v}_t - \alpha_{t+1} \langle \nabla f(\mathbf{y}_t), \tilde{\mathbf{z}}_t \rangle \tilde{\mathbf{z}}_t$ 
11: end for
    
```

Proof of Lemma 7. We will prove the statement inductively. For $t = 0$, $\psi_0(\mathbf{x}) = \frac{1}{2}\|\mathbf{x} - \mathbf{v}_0\|_{\mathbf{P}}^2$ and so the statement holds. Suppose it holds for some $t \geq 0$. Observe that the function $\psi_t(\mathbf{x})$ is a quadratic with Hessian \mathbf{P} . This means that we can reformulate $\psi_t(\mathbf{x})$ with minima at \mathbf{v}_t as

$$\psi_t(\mathbf{x}) = \psi_t(\mathbf{v}_t) + \frac{1}{2}\|\mathbf{x} - \mathbf{v}_t\|_{\mathbf{P}}^2.$$

Using this reformulation,

$$\begin{aligned}
 \arg \min_{\mathbf{x}} \psi_{t+1}(\mathbf{x}) &= \arg \min_{\mathbf{x}} \left\{ \psi_t(\mathbf{x}) + \alpha_{t+1} \left(f(\mathbf{y}_t) + \langle \tilde{\mathbf{z}}_t^\top \nabla f(\mathbf{y}_t), \tilde{\mathbf{z}}_t^\top \mathbf{P}(\mathbf{x} - \mathbf{y}_t) \rangle \right) \right\} \\
 &= \arg \min_{\mathbf{x}} \left\{ \psi_t(\mathbf{v}_t) + \frac{1}{2}\|\mathbf{x} - \mathbf{v}_t\|_{\mathbf{P}}^2 + \alpha_{t+1} \left(f(\mathbf{y}_t) + \langle \tilde{\mathbf{z}}_t^\top \nabla f(\mathbf{y}_t), \tilde{\mathbf{z}}_t^\top \mathbf{P}(\mathbf{x} - \mathbf{y}_t) \rangle \right) \right\} \\
 &= \arg \min_{\mathbf{x}} \left\{ \frac{1}{2}\|\mathbf{x} - \mathbf{v}_t\|_{\mathbf{P}}^2 + \alpha_{t+1} \langle \tilde{\mathbf{z}}_t^\top \nabla f(\mathbf{y}_t), \tilde{\mathbf{z}}_t^\top \mathbf{P}(\mathbf{x} - \mathbf{v}_t) \rangle \right\} \\
 &= \mathbf{v}_t - \alpha_{t+1} \langle \nabla f(\mathbf{y}_t), \tilde{\mathbf{z}}_t \rangle \tilde{\mathbf{z}}_t \\
 &= \mathbf{v}_{t+1}. \quad \square
 \end{aligned}$$

Lemma 11 (Upper bound on $\psi_t(\mathbf{x})$).

$$\mathbb{E}[\psi_t(\mathbf{x})] \leq \beta_t f(\mathbf{x}) + \psi_0(\mathbf{x}).$$

Proof. We will also show this through induction. The statement is trivially true for $t = 0$ since $\beta_0 = 0$. Assuming the statement holds for some $t \geq 0$,

$$\begin{aligned}
 \mathbb{E}[\psi_{t+1}(\mathbf{x})] &= \mathbb{E} \left[\psi_t(\mathbf{x}) + \alpha_{t+1} \left(f(\mathbf{y}_t) + \langle \tilde{\mathbf{z}}_t^\top \nabla f(\mathbf{y}_t), \tilde{\mathbf{z}}_t^\top \mathbf{P}(\mathbf{x} - \mathbf{y}_t) \rangle \right) \right] \\
 &= \mathbb{E} \left[\psi_t(\mathbf{x}) \right] + \alpha_{t+1} \mathbb{E} \left[\left(f(\mathbf{y}_t) + \langle \tilde{\mathbf{z}}_t^\top \nabla f(\mathbf{y}_t), \tilde{\mathbf{z}}_t^\top \mathbf{P}(\mathbf{x} - \mathbf{y}_t) \rangle \right) \right] \\
 &\leq \beta_t f(\mathbf{x}) + \psi_0(\mathbf{x}) + \alpha_{t+1} \left(f(\mathbf{y}_t) + \nabla f(\mathbf{y}_t)^\top \mathbb{E} \left[\tilde{\mathbf{z}}_t \tilde{\mathbf{z}}_t^\top \right] \mathbf{P}(\mathbf{x} - \mathbf{y}_t) \right) \\
 &= \beta_t f(\mathbf{x}) + \psi_0(\mathbf{x}) + \alpha_{t+1} \left(f(\mathbf{y}_t) + \nabla f(\mathbf{y}_t)^\top \mathbf{P}^{-1} \mathbf{P}(\mathbf{x} - \mathbf{y}_t) \right) \\
 &= \beta_t f(\mathbf{x}) + \psi_0(\mathbf{x}) + \alpha_{t+1} \left(f(\mathbf{y}_t) + \nabla f(\mathbf{y}_t)^\top (\mathbf{x} - \mathbf{y}_t) \right) \\
 &\leq \beta_t f(\mathbf{x}) + \psi_0(\mathbf{x}) + \alpha_{t+1} f(\mathbf{x}).
 \end{aligned}$$

In the above, we used the convexity of the function $f(\mathbf{x})$ and the definition of \mathbf{P} . □

Lemma 12 (Bound on progress). *For any $t \geq 0$ of algorithm 5,*

$$f(\mathbf{x}_{t+1}) - f(\mathbf{y}_t) \leq -\frac{1}{2L\|\mathbf{z}_t\|_2^2} \nabla f(\mathbf{y}_t)^\top [\mathbf{z}_t \mathbf{z}_t^\top] \nabla f(\mathbf{y}_t).$$

Proof. The update \mathbf{x}_{t+1} along with the smoothness of $f(\mathbf{x})$ guarantees that for $\gamma_{t+1} = \frac{\langle \nabla f(\mathbf{y}_t), \mathbf{z}_t \rangle}{L \|\mathbf{z}_t\|_2^2}$,

$$\begin{aligned} f(\mathbf{x}_{t+1}) &= f(\mathbf{y}_t + \gamma_{t+1} \mathbf{z}_t) \\ &\leq f(\mathbf{y}_t) + \gamma_{t+1} \langle \nabla f(\mathbf{y}_t), \mathbf{z}_t \rangle + \frac{L \gamma_{t+1}^2}{2} \|\mathbf{z}_t\|^2 \\ &= f(\mathbf{y}_t) - \frac{1}{2L \|\mathbf{z}_t\|_2^2} \nabla f(\mathbf{y}_t)^\top [\mathbf{z}_t \mathbf{z}_t^\top] \nabla f(\mathbf{y}_t). \end{aligned}$$

□

Lemma 13 (Lower bound on $\psi_t(\mathbf{x})$). *Given a filtration \mathcal{F}_t upto time step t ,*

$$\mathbb{E}[\min_{\mathbf{x}} \psi_t(\mathbf{x}) | \mathcal{F}_t] \geq \beta_t f(\mathbf{x}_t).$$

Proof. This too we will show inductively. For $t = 0$, $\psi_t(\mathbf{x}) = \frac{1}{2} \|\mathbf{x} - \mathbf{v}_0\|_{\mathbf{P}}^2 \geq 0$ with $\beta_0 = 0$. Assume the statement holds for some $t \geq 0$. Recall that $\psi_t(\mathbf{x})$ has a minima at \mathbf{v}_t and can be alternatively formulated as $\psi_t(\mathbf{v}_t) + \frac{1}{2} \|\mathbf{x} - \mathbf{v}_t\|_{\mathbf{P}}^2$. Using this,

$$\begin{aligned} \psi_{t+1}^* &= \min_{\mathbf{x}} [\psi_t(\mathbf{x}) + \alpha_{t+1} (\langle \tilde{\mathbf{z}}_t^\top \nabla f(\mathbf{y}_t), \tilde{\mathbf{z}}_t^\top \mathbf{P}(\mathbf{x} - \mathbf{y}_t) \rangle + f(\mathbf{y}_t))] \\ &= \min_{\mathbf{x}} \left[\psi_t(\mathbf{v}_t) + \alpha_{t+1} \left(\langle \tilde{\mathbf{z}}_t^\top \nabla f(\mathbf{y}_t), \tilde{\mathbf{z}}_t^\top \mathbf{P}(\mathbf{x} - \mathbf{y}_t) \rangle + \frac{1}{2\alpha_{t+1}} \|\mathbf{x} - \mathbf{v}_t\|_{\mathbf{P}}^2 + f(\mathbf{y}_t) \right) \right] \\ &= \psi_t^* + \alpha_{t+1} f(\mathbf{y}_t) + \alpha_{t+1} \min_{\mathbf{x}} \left[\langle \mathbf{P} \tilde{\mathbf{z}}_t \tilde{\mathbf{z}}_t^\top \nabla f(\mathbf{y}_t), \mathbf{x} - \mathbf{y}_t \rangle + \frac{1}{2\alpha_{t+1}} \|\mathbf{x} - \mathbf{v}_t\|_{\mathbf{P}}^2 \right]. \end{aligned}$$

Since we defined $\mathbf{y}_t = (1 - \tau_t) \mathbf{x}_t + \tau_t \mathbf{v}_t$, rearranging the terms gives us that

$$\mathbf{y}_t - \mathbf{v}_t = \frac{1 - \tau_t}{\tau_t} (\mathbf{x}_t - \mathbf{y}_t).$$

Let us take now compute $\mathbb{E}[\psi_{t+1}^* | \mathcal{F}_t]$ by combining the above two equations:

$$\begin{aligned} \mathbb{E}[\psi_{t+1}^* | \mathcal{F}_t] &= \psi_t^* + \alpha_{t+1} f(\mathbf{y}_t) + \frac{\alpha_{t+1}(1 - \tau_t)}{\tau_t} \langle \mathbf{P} \mathbb{E}_t[\tilde{\mathbf{z}}_t \tilde{\mathbf{z}}_t^\top] \nabla f(\mathbf{y}_t), \mathbf{y}_t - \mathbf{x}_t \rangle \\ &\quad + \alpha_{t+1} \mathbb{E}_t \min_{\mathbf{x}} \left[\langle \mathbf{P} \tilde{\mathbf{z}}_t \tilde{\mathbf{z}}_t^\top \nabla f(\mathbf{y}_t), \mathbf{x} - \mathbf{v}_t \rangle + \frac{1}{2\alpha_{t+1}} \|\mathbf{x} - \mathbf{v}_t\|_{\mathbf{P}}^2 \right] \\ &= \psi_t^* + \alpha_{t+1} f(\mathbf{y}_t) + \frac{\alpha_{t+1}(1 - \tau_t)}{\tau_t} \langle \nabla f(\mathbf{y}_t), \mathbf{y}_t - \mathbf{x}_t \rangle \\ &\quad + \alpha_{t+1} \mathbb{E}_t \min_{\mathbf{x}} \left[\langle \mathbf{P} \tilde{\mathbf{z}}_t \tilde{\mathbf{z}}_t^\top \nabla f(\mathbf{y}_t), \mathbf{x} - \mathbf{v}_t \rangle + \frac{1}{2\alpha_{t+1}} \|\mathbf{x} - \mathbf{v}_t\|_{\mathbf{P}}^2 \right] \\ &= \psi_t^* + \alpha_{t+1} f(\mathbf{y}_t) + \frac{\alpha_{t+1}(1 - \tau_t)}{\tau_t} \langle \nabla f(\mathbf{y}_t), \mathbf{y}_t - \mathbf{x}_t \rangle \\ &\quad - \frac{\alpha_{t+1}^2}{2} \nabla f(\mathbf{y}_t)^\top \mathbb{E}_t [\tilde{\mathbf{z}}_t \tilde{\mathbf{z}}_t^\top \mathbf{P} \mathbf{P}^{-1} \mathbf{P} \tilde{\mathbf{z}}_t \tilde{\mathbf{z}}_t^\top] \nabla f(\mathbf{y}_t) \\ &= \psi_t^* + \alpha_{t+1} f(\mathbf{y}_t) + \frac{\alpha_{t+1}(1 - \tau_t)}{\tau_t} \langle \nabla f(\mathbf{y}_t), \mathbf{y}_t - \mathbf{x}_t \rangle \\ &\quad - \frac{\alpha_{t+1}^2}{2} \nabla f(\mathbf{y}_t)^\top \mathbb{E}_t [\tilde{\mathbf{z}}_t \tilde{\mathbf{z}}_t^\top \mathbf{P} \tilde{\mathbf{z}}_t \tilde{\mathbf{z}}_t^\top] \nabla f(\mathbf{y}_t). \end{aligned}$$

Let us define a constant $\nu \geq 0$ such that it is the smallest number for which the below inequality holds for all t ,

$$\nu \nabla f(\mathbf{y}_t)^\top \frac{[\mathbf{z}_t \mathbf{z}_t^\top]}{2L \|\mathbf{z}_t\|_2^2} \nabla f(\mathbf{y}_t) \geq \nabla f(\mathbf{y}_t)^\top \mathbb{E} [\tilde{\mathbf{z}}_t \tilde{\mathbf{z}}_t^\top \mathbf{P} \tilde{\mathbf{z}}_t \tilde{\mathbf{z}}_t^\top] \nabla f(\mathbf{y}_t).$$

Also recall from Lemma 12 that

$$f(\mathbf{x}_{t+1}) - f(\mathbf{y}_t) \leq -\frac{1}{2L\|\mathbf{z}_t\|_2^2} \nabla f(\mathbf{y}_t)^\top [\mathbf{z}_t \mathbf{z}_t^\top] \nabla f(\mathbf{y}_t).$$

Using the above two statements in our computation of ψ_{t+1}^* , we get

$$\begin{aligned} \mathbb{E}[\psi_{t+1}^* | \mathcal{F}_t] &= \psi_t^* + \alpha_{t+1} f(\mathbf{y}_t) + \frac{\alpha_{t+1}(1-\tau_t)}{\tau_t} \langle \nabla f(\mathbf{y}_t), \mathbf{y}_t - \mathbf{x}_t \rangle \\ &\quad - \frac{\alpha_{t+1}^2}{2} \nabla f(\mathbf{y}_t)^\top \mathbb{E}_t[\tilde{\mathbf{z}}_t \tilde{\mathbf{z}}_t^\top \mathbf{P} \tilde{\mathbf{z}}_t \tilde{\mathbf{z}}_t^\top] \nabla f(\mathbf{y}_t) \\ &\geq \psi_t^* + \alpha_{t+1} f(\mathbf{y}_t) + \frac{\alpha_{t+1}(1-\tau_t)}{\tau_t} \langle \nabla f(\mathbf{y}_t), \mathbf{y}_t - \mathbf{x}_t \rangle \\ &\quad - \frac{\alpha_{t+1}^2 \nu}{2} \nabla f(\mathbf{y}_t)^\top [\mathbf{z}_t \mathbf{z}_t^\top] \nabla f(\mathbf{y}_t) \\ &\geq \psi_t^* + \alpha_{t+1} f(\mathbf{y}_t) + \frac{\alpha_{t+1}(1-\tau_t)}{\tau_t} \langle \nabla f(\mathbf{y}_t), \mathbf{y}_t - \mathbf{x}_t \rangle \\ &\quad + \alpha_{t+1} L \nu (f(\mathbf{x}_{t+1}) - f(\mathbf{y}_t)) \\ &\geq \psi_t^* + \alpha_{t+1} f(\mathbf{y}_t) + \frac{\alpha_{t+1}(1-\tau_t)}{\tau_t} (f(\mathbf{y}_t) - f(\mathbf{x}_t)) \\ &\quad + \alpha_{t+1}^2 L \nu (f(\mathbf{x}_{t+1}) - f(\mathbf{y}_t)). \end{aligned}$$

Let us pick α_{t+1} such that it satisfies $\alpha_{t+1}^2 \nu L = \beta_{t+1}$. Then the above equation simplifies to

$$\begin{aligned} \mathbb{E}[\psi_{t+1}^* | \mathcal{F}_t] &\geq \psi_t^* + \frac{\alpha_{t+1}}{\tau_t} f(\mathbf{y}_t) - \frac{\alpha_{t+1}(1-\tau_t)}{\tau_t} f(\mathbf{x}_t) + \beta_{t+1} (f(\mathbf{x}_{t+1}) - f(\mathbf{y}_t)) \\ &= \psi_t^* - \beta_t f(\mathbf{x}_t) + \beta_{t+1} f(\mathbf{y}_t) - \beta_{t+1} f(\mathbf{y}_t) + \beta_{t+1} f(\mathbf{x}_{t+1}) \\ &= \psi_t^* - \beta_t f(\mathbf{x}_t) + \beta_{t+1} f(\mathbf{x}_{t+1}). \end{aligned}$$

We used that $\tau_t = \alpha_{t+1}/\beta_{t+1}$. Finally we use the inductive hypothesis to conclude that

$$\mathbb{E}[\psi_{t+1}^* | \mathcal{F}_t] \geq \psi_t^* - \beta_t f(\mathbf{x}_t) + \beta_{t+1} f(\mathbf{x}_{t+1}) \geq \beta_{t+1} f(\mathbf{x}_{t+1}). \quad \square$$

Lemma 14 (Final convergence rate). *For any $t \geq 1$ the output of algorithm 5 satisfies:*

$$\mathbb{E}[f(\mathbf{x}_t)] - f(\mathbf{x}^*) \leq \frac{2L\nu}{t(t+1)} \|\mathbf{x}^* - \mathbf{x}_0\|_{\mathbf{P}}^2.$$

Proof. Putting together Lemmas 11 and 13, we have that

$$\beta_t \mathbb{E}[f(\mathbf{x}_t)] \leq \mathbb{E}[\psi_t^*] \leq \mathbb{E}[\psi_t(\mathbf{x}^*)] \leq \beta_t f(\mathbf{x}^*) + \psi_0(\mathbf{x}^*).$$

Rearranging the terms we get

$$\mathbb{E}[f(\mathbf{x}_t)] - f(\mathbf{x}^*) \leq \frac{1}{2\beta_t} \|\mathbf{x}^* - \mathbf{x}_0\|_{\mathbf{P}}^2.$$

To finish the proof of the theorem, we only have to compute the value of β_t . Recall that

$$\alpha_{t+1}^2 L \nu = \beta_t + \alpha_{t+1}.$$

We will inductively show that $\alpha_t \geq \frac{t}{2L\nu}$. For $t = 0$, $\beta_0 = 0$ and $\alpha_1 = \frac{1}{2L\nu}$ which satisfies the condition. Suppose that for some $t \geq 0$, the inequality holds for all iterations $i \leq t$. Recall that $\beta_t = \sum_{i=1}^t \alpha_i$ i.e. $\beta_t \geq \frac{t(t+1)}{4L\nu}$. Then

$$(\alpha_{t+1} L \nu)^2 - \alpha_{t+1} L \nu = \beta_t L \nu \geq \frac{t(t+1)}{4}.$$

The positive root of the quadratic $x^2 - x - c = 0$ for $c \geq 0$ is $x = \frac{1}{2}(1 + \sqrt{4c + 1})$. Thus

$$\alpha_{t+1}L\nu \geq \frac{1}{2}\left(1 + \sqrt{t(t+1) + 1}\right) \geq \frac{t+1}{2}.$$

This finishes our induction and proves the final rate of convergence. □

Lemma 15 (Understanding ν).

$$\nu \leq \max_{\mathbf{d} \in \text{lin}(\mathcal{A})} \frac{\mathbb{E}\left[(\tilde{\mathbf{z}}_t^\top \mathbf{d})^2 \|\tilde{\mathbf{z}}_t\|_{\mathbf{P}}^2\right] \|\mathbf{z}(\mathbf{d})\|_2^2}{(\mathbf{z}(\mathbf{d})^\top \mathbf{d})^2},$$

Proof. Recall the definition of ν as a constant which satisfies the following inequality for all iterations t

$$\nu \nabla f(\mathbf{y}_t)^\top \frac{[\mathbf{z}_t \mathbf{z}_t^\top]}{2L \|\mathbf{z}_t\|_2^2} \nabla f(\mathbf{y}_t) \geq \nabla f(\mathbf{y}_t)^\top \mathbb{E}[\tilde{\mathbf{z}}_t \tilde{\mathbf{z}}_t^\top \mathbf{P} \tilde{\mathbf{z}}_t \tilde{\mathbf{z}}_t^\top] \nabla f(\mathbf{y}_t).$$

which then yields the following sufficient condition for ν :

$$\nu \leq \max_{\mathbf{d} \in \text{lin}(\mathcal{A})} \frac{\mathbb{E}\left[(\tilde{\mathbf{z}}_t^\top \mathbf{d})^2 \|\tilde{\mathbf{z}}_t\|_{\mathbf{P}}^2\right] \|\mathbf{z}(\mathbf{d})\|_2^2}{(\mathbf{z}(\mathbf{d})^\top \mathbf{d})^2},$$

where $\mathbf{z}(\mathbf{d})$ is defined to be

$$\mathbf{z}(\mathbf{d}) = \text{LMO}_{\mathcal{A}}(-\mathbf{d}).$$

□

Proof of Theorem 9. The proof of Theorem 9 is exactly the same as that of the previous except that now the update to \mathbf{v}_t is also a random variable. The only change needed is the definition of ν' where we need the following to hold:

$$\nu' \nabla f(\mathbf{y}_t)^\top \frac{1}{2L} \mathbb{E}_t \left[\frac{[\mathbf{z}_t \mathbf{z}_t^\top]}{\|\mathbf{z}_t\|_2^2} \right] \nabla f(\mathbf{y}_t) \geq \nabla f(\mathbf{y}_t)^\top \mathbb{E}[\mathbf{z}_t \mathbf{z}_t^\top \mathbf{P} \mathbf{z}_t \mathbf{z}_t^\top] \nabla f(\mathbf{y}_t).$$

Proof of Lemma 10. When $\mathcal{A} = \{\mathbf{e}_i, i \in [n]\}$ and \mathcal{Z} is a uniform distribution over \mathcal{A} , then $\tilde{\mathbf{P}} = 1/n\mathbf{I}$ and $\mathbf{P} = n\mathbf{I}$. A simple computation shows that $\nu' = n$ and $\nu \in [1, n]$. Note that here ν could be upto n times smaller than ν' meaning that our accelerated greedy coordinate descent algorithm could be \sqrt{n} times faster than the accelerated random coordinate descent. In the worst case $\nu = \nu'$, but in practice one can pick a smaller ν compared to ν' as the worst case gradient rarely happen. It is possible to tune ν and ν' empirically but we do not explore this direction.