

2015/13



Smooth Strongly Convex Interpolation and Exact Worst-case Performance of First-order Methods

A.B. Taylor, J.M. Hendrickx, F. Glineur

The logo for CORE, featuring the word 'CORE' in a bold, black, sans-serif font. A light blue curved line starts above the 'O' and arcs over the 'R' and 'E'.

DISCUSSION PAPER

Center for Operations Research
and Econometrics

Voie du Roman Pays, 34
B-1348 Louvain-la-Neuve
Belgium

<http://www.uclouvain.be/core>

CORE

Voie du Roman Pays 34, L1.03.01

B-1348 Louvain-la-Neuve, Belgium.

Tel (32 10) 47 43 04

Fax (32 10) 47 43 01

E-mail: immaq-library@uclouvain.be

<http://www.uclouvain.be/en-44508.html>

Smooth Strongly Convex Interpolation and Exact Worst-case Performance of First-order Methods

A.B. Taylor · J.M. Hendrickx · F. Glineur

Date of current version: March 11, 2015

Abstract We show that the exact worst-case performance of fixed-step first-order methods for smooth (possibly strongly) convex functions can be obtained by solving convex programs.

Finding the worst-case performance of a black-box first-order method is formulated as an optimization problem over a set of smooth (strongly) convex functions and initial conditions. We develop closed-form necessary and sufficient conditions for smooth (strongly) convex interpolation, which provide a finite representation for those functions. This allows us to reformulate the worst-case performance estimation problem as an equivalent finite dimension-independent semidefinite optimization problem, whose exact solution can be recovered up to numerical precision. Optimal solutions to this performance estimation problem provide both worst-case performance bounds and explicit functions matching them, as our smooth (strongly) convex interpolation procedure is constructive.

Our works build on those of Drori and Teboulle in [8] who introduced and solved relaxations of the performance estimation problem for smooth convex functions.

We apply our approach to different fixed-step first-order methods with several performance criteria, including objective function accuracy and residual gradient norm. We conjecture several numerically supported worst-case bounds on the performance of the gradient, fast gradient and optimized fixed-step methods, both in the smooth convex and the smooth strongly convex cases, and deduce tight estimates of the optimal step size for the gradient method.

1 Introduction to performance estimation

Consider the standard unconstrained minimization problem

$$\min_{x \in \mathbb{R}^d} f(x),$$

where f is a smooth convex function, possibly strongly convex. First-order black-box methods, which only rely on the computation of f and its gradient at a sequence of iterates, can be designed to solve this type of problem iteratively. A central question is then to estimate the accuracy of solutions computed by such a method. More precisely, given a class of problems and a first-order method, one wishes to establish the worst-case accuracy of solutions that can be obtained after applying a given number of iterations, i.e. the performance of the method on the given class of problems.

Many first-order algorithms have been proposed in the literature for smooth convex or smooth strongly convex functions, for which one usually provides a theoretical *upper bound* on the worst-case accuracy after a number of iterations (see e.g. [18] or [5, Chap.6] for recent overviews). We note however that many analyses focus on the asymptotic rate of convergence of these bounds, rather than trying

This paper presents research results of the Belgian Network DYSCO (Dynamical Systems, Control, and Optimization), funded by the Interuniversity Attraction Poles Programme, initiated by the Belgian State, Science Policy Office, and of the Concerted Research Action (ARC) programme supported by the Federation Wallonia-Brussels (contract ARC 14/19-060). The scientific responsibility rests with its authors. A.B.T. is a F.R.I.A. fellow.

A.B. Taylor, J.M. Hendrickx, F. Glineur
 Université catholique de Louvain, ICTEAM Institute/CORE, B-1348 Louvain-la-Neuve, Belgium
 E-mail: Adrien.Taylor@uclouvain.be; Julien.Hendrickx@uclouvain.be; Francois.Glineur@uclouvain.be

to compute exact numerical values. Similarly, *lower bounds* on the performance of first-order black-box methods on given classes of problems can be found in the literature (see e.g. the seminal [16]), again often with a focus on asymptotic rates of convergence. In many situations, the asymptotic rate of the performance of the best available methods match those lower bounds.

Nevertheless, the *exact* numerical value of the worst-case performance of a given method is usually unknown. This can be due to the fact that upper bounds are not assessed precisely, i.e. are known only up to a (possibly unspecified) constant. Another reason is that lower bounds for specific methods are not very frequently developed, and that general lower bounds (valid for all methods) can be quite weak for specific methods, especially if those methods do not feature the best possible asymptotic rate of convergence. Finally, even if exact numerical values are known for both lower and upper bounds, and share the same (optimal) asymptotic rate of convergence, a significant gap between the numerical values of those lower and upper bounds can subsist. If one cares about the worst-case efficiency of a first-order method in practice, this gap can translate into a very large uncertainty on the concrete behavior of a method.

This work is not concerned with asymptotic rates of convergence. It will focus on the computation of the exact worst-case performance of a given first-order black-box method, on a given class of functions, after a given number of iterations. We prove that this question can be formulated and solved exactly as a (finite-dimensional) convex optimization problem, with the following attractive features:

- Our formulation is a semidefinite optimization problem, whose dimension is proportional to the number of iterations of the method to be analyzed.
- Any dual feasible solution of our formulation provides an upper bound on the worst-case performance. This solution can be converted easily into a standard proof establishing a bound on the performance (i.e. a series of valid inequalities).
- Any primal feasible solution of our formulation provides a lower bound on the worst case performance. This solution can be easily converted into a concrete function on which the method exhibits the corresponding performance.
- Hence our formulation is exact, i.e. its optimal value provides the exact worst-case performance.

Our formulation covers both smooth convex functions and smooth strongly convex functions in a unified fashion. It covers a very large class of first-order methods which includes the majority of standard methods. It can be applied to a variety of performance measures, such as objective function accuracy, gradient norm, or distance to an optimal solution.

1.1 Formal definition

Our goal is to express the worst-case performance of an optimization algorithm as the solution of an optimization problem. This approach was pioneered by Drori and Teboulle in [8], who called it a Performance Estimation Problem (PEP). We now provide a formal definition for this problem.

We consider unconstrained minimization problems involving a given class of objective functions, and only treat first-order black-box method. This means that the method can only gather information about the objective function using an oracle \mathcal{O}_f , which returns first-order information about specific points, i.e. $\mathcal{O}_f(x) = \{f(x), f'(x)\}$. Formally, the first N iterates generated by a first-order black-box method \mathcal{M} (which correspond to N calls of the oracle), starting from an initial point x_0 , can be described with

$$\begin{aligned} x_1 &= \mathcal{M}_1(x_0, \mathcal{O}_f(x_0)), \\ x_2 &= \mathcal{M}_2(x_0, \mathcal{O}_f(x_0), \mathcal{O}_f(x_1)), \\ &\vdots \\ x_N &= \mathcal{M}_N(x_0, \mathcal{O}_f(x_0), \dots, \mathcal{O}_f(x_{N-1})). \end{aligned}$$

In order to measure the performance of a given method \mathcal{M} on a specific function f with a specific starting point, we introduce a performance criterion \mathcal{P} , to be minimized, that will only depend on the function f and the sequence of the iterates $\{x_0, x_1, \dots, x_N\}$ generated by the method. Since we are in a black-box setting, we will require that the criterion can be computed from the output of the oracle \mathcal{O}_f , which has only access to the iterates as well as to an additional point x_* , defined to be any minimizer of function f (the latter being necessary if the criterion has to compare iterates to an optimal solution).

Examples of this performance criterion $\mathcal{P}(\mathcal{O}_f, x_0, \dots, x_N, x_*)$ include the objective function accuracy $f(x_N) - f(x_*)$, the norm of the gradient $\|\nabla f(x_N)\|$, or the distance to an optimal solution $\|x_N - x_*\|$ (see also Section 4.3 for an example of criterion that does not only depend on the last iterate).

Finally, we consider a given class \mathcal{F} of smooth convex or smooth strongly convex functions, over which we wish to estimate the worst-case performance of a method after N iterations. As methods try to minimize the performance criterion, their worst-case performance is obtained by maximizing \mathcal{P} over functions in \mathcal{F} , which can be written as

$$\begin{aligned} w(\mathcal{F}, R, \mathcal{M}, N, \mathcal{P}) = & \sup_{f, x_0, \dots, x_N, x_*} \mathcal{P}(\mathcal{O}_f, x_0, \dots, x_N, x_*) & (\text{PEP}) \\ & \text{s.t. } f \in \mathcal{F}, \\ & x_* \text{ is optimal for } f, \\ & x_1, \dots, x_N \text{ is generated by method } \mathcal{M} \text{ from } x_0, \\ & \|x_0 - x_*\|_2 \leq R. \end{aligned}$$

Note that a last parameter R was introduced to bound the distance between the initial point x_0 and the optimal solution x_* . Indeed, it is well-known that, in most situations, performance of a first-order method cannot be sensibly assessed without such a constraint (see also the discussion of Section 3.2).

1.2 Finite-dimensional reformulation using interpolation

Because it involves an unknown function f as a variable, problem (PEP) is infinite-dimensional. Nevertheless, using the black-box property of the method (and of the performance criterion), we will show that a completely equivalent finite-dimensional problem can readily be formulated, by restricting variable f to the knowledge of the output of its oracle \mathcal{O}_f on the iterates $\{x_0, x_1, \dots, x_N\}$ and x_* .

More concretely, defining a set $I = \{0, 1, 2, \dots, N, *\}$ for the indices of iterates, we can reformulate (PEP) into a problem involving only the iterates $\{x_i\}_{i \in I}$, their function values $\{f_i\}_{i \in I}$ and their gradients $\{g_i\}_{i \in I}$ as

$$\begin{aligned} w_f(\mathcal{F}, R, \mathcal{M}, N, \mathcal{P}) = & \sup_{\{x_i, g_i, f_i\}_{i \in I}} \mathcal{P}(\{x_i, g_i, f_i\}_{i \in I}), & (\text{f-PEP}) \\ & \text{s.t. there exists } f \in \mathcal{F} \text{ such that } \mathcal{O}_f(x_i) = \{f_i, g_i\} \forall i \in I, \\ & g_* = 0, \\ & x_1, \dots, x_N \text{ is generated by method } \mathcal{M} \text{ from } x_0, \\ & \|x_0 - x_*\|_2 \leq R, \end{aligned}$$

Observe that assuming optimality of x_* is equivalent to requiring $g_* = 0$. The crucial part of this reformulation is the first constraint, which can be understood as requiring that the set variables $\{x_i, g_i, f_i\}_{i \in I}$ can be *interpolated* by a function belonging to the class \mathcal{F} .

1.3 Paper organization and main contributions

We focus in this paper on the class of smooth (strongly) convex functions. Therefore the main technical tool required for an exact formulation of problem (PEP) as (f-PEP) will be a set of necessary and sufficient conditions for the existence of a smooth strongly convex interpolating function, which is the main result obtained in Section 2. This set of conditions, which is of independent interest, was previously only known for general nonsmooth convex functions. Our approach is fully constructive, as we also exhibit a procedure to interpolate a smooth (strongly) convex function from a set of points with their associated gradients and function values, when such an interpolating function exists.

In Section 3, we show how the resulting finite-dimensional (f-PEP) problem can be reformulated exactly into a (convex) semidefinite optimization problem, which provides the first tractable and provably exact formulation of the performance estimation problem. We allow consideration of both smooth convex and smooth strongly convex functions, as well as a large class of performance criteria.

Section 4 then tests our approach numerically on several standard first-order methods, including the constant-step gradient method, the fast gradient method and the optimized method from [11]. We are

able to confirm several bounds appearing previously in [8], and to conjecture several new worst-case performance bounds, including bounds for strongly convex functions, and bounds on the gradient norm (either for the final iterate, or the smallest norm among all iterates). Another byproduct of our results is a tight estimate of the optimal step size for the gradient method on smooth convex and smooth strongly convex functions.

1.4 Prior work

Drori and Teboulle [8] were first to consider the notion of a performance estimation problem. They focus exclusively on the case of smooth convex functions equipped with the performance criterion $f(x_N) - f_*$, and introduce the idea of reducing (PEP) to a finite-dimensional problem involving only the iterates x_i , their gradients g_i and function values f_i , along with the optimum point x_* and optimal value f_* . They treat several standard first-order algorithms, namely, the standard fixed-step gradient algorithm, the heavy ball method and the accelerated gradient method [17]. In their approach, (PEP) is expressed as a non-convex Quadratic Matrix Program [2], which is then relaxed and dualized. The resulting convex problem is then used to provide bounds on the worst-case performance (and, in some cases, is solved analytically). As will be shown later in this paper (see Section 4), because of the use of a relaxation and the dualization of a non-convex problem, these bounds are in general not tight, although they are in many special cases.

A Section in [8] is also devoted to the optimization of the coefficients of a general first-order black-box method. More precisely, a numerical optimization solver is used to identify a method performing best according to their relaxation of the performance estimation problem. This approach is taken further in [11], which provides an analytical description of this optimized method. Again we stress that, due to the non-tightness of the relaxation in general, these optimized methods are however not guaranteed to have the best possible performances.

Another computational approach for the analysis and design of first-order algorithms is proposed in [13], in which optimization procedures are regarded as dynamical systems. Integral Quadratic Constraints (IQC), which are usually used to obtain stability guarantees on complicated dynamical systems, are adapted in order to obtain sufficient conditions for the convergence of optimization algorithms. This methodology is more conservative than that of [8] (and than ours), i.e. its performance bounds are worse, but it presents the advantage of deriving a bound valid for all iterations after solving a single semidefinite optimization problem.

2 Smooth strongly convex interpolation

This section develops a necessary and sufficient condition for the existence of a smooth strongly convex function interpolating through a given set of data triples $\{x_i, g_i, f_i\}_{i \in I}$, i.e. deciding whether there exists a smooth strongly convex function f such that $f(x_i) = f_i$ and $g_i \in \partial f(x_i)$ for all $i \in I$.

This result generalizes the well-known set of conditions guaranteeing the existence of a convex, possibly nonsmooth interpolating function (see Theorem 1 in Subsection 2.3). It is the main technical ingredient of our exact convex reformulation of performance estimation problems.

2.1 Definitions and problem statement

We start by defining the functional class of interest, using the standard point of view from convex analysis — we refer to classic books [1, 10, 20, 21] for details. Given two parameters μ and L satisfying $0 \leq \mu < L \leq +\infty$, we consider proper closed convex functions (i.e. whose epigraph are non-empty closed convex) satisfying both a smoothness condition (depending on the parameter L , which is the Lipschitz constant of the gradient) and a strong convexity condition (depending on the parameter μ). We explicitly allow the case $L = +\infty$, while μ on the other hand is always assumed to be finite. In the rest of this paper, to deal with the case $L = +\infty$, we use the conventions $1/+\infty = 0$ and $+\infty - \mu = +\infty$.

Definition 1 (L -smooth μ -strongly convex functions) Consider a proper and closed convex function $f : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$, and constants $\mu \in \mathbb{R}^+$, $L \in \mathbb{R}^+ \cup \{+\infty\}$ with $\mu < L$. We say that function f is L -smooth μ -strongly convex (which we denote by $f \in \mathcal{F}_{\mu, L}$) if and only if the following two conditions are satisfied:

- (a) inequality $\frac{1}{L}\|g_1 - g_2\|_2 \leq \|x_1 - x_2\|_2$ holds for all pairs $x_1, x_2 \in \mathbb{R}^d$ and corresponding subgradients $g_1, g_2 \in \mathbb{R}^d$ (i.e. such that $g_1 \in \partial f(x_1)$ and $g_2 \in \partial f(x_2)$).
- (b) function $f(x) - \frac{\mu}{2}\|x\|_2^2$ is convex.

This definition is not entirely standard, as it involves subgradients and allows constant L to be equal to $+\infty$. In the case of a finite L , condition (a) immediately implies uniqueness of the subgradient at each point, hence the function is differentiable, and we recover the well-known Lipschitz condition on the gradient of a smooth function. On the other hand, when $L = \infty$, condition (a) becomes vacuous, and the function can be non-differentiable. Condition (b) can easily be seen to be algebraically equivalent to the standard definition of strong convexity (and the case $\mu = 0$ corresponds to a convex but not strongly convex function). The class of proper closed convex functions simply corresponds to $\mathcal{F}_{0,\infty}$. The case $L = \mu$ can be safely discarded, as it only involves simple quadratic functions whose minimization is trivial. The reason for this slightly non-standard definition of $\mathcal{F}_{\mu,L}$ and the possibility of choosing $L = +\infty$ will become clear later, when dealing with operation of Fenchel conjugation.

As explained above and in the introduction, our approach to express the original infinite dimensional (PEP) in a finite-dimensional fashion relies on an interpolating condition for smooth strongly convex functions. This directly motivates the following definition.

Definition 2 ($\mathcal{F}_{\mu,L}$ -interpolation) Let I be an index set, and consider the set of triples $S = \{(x_i, g_i, f_i)\}_{i \in I}$ where $x_i, g_i \in \mathbb{R}^d$ and $f_i \in \mathbb{R}$ for all $i \in I$. Set S is $\mathcal{F}_{\mu,L}$ -interpolable if and only if there exists a function $f \in \mathcal{F}_{\mu,L}$ such that we have both $g_i \in \partial f(x_i)$ and $f(x_i) = f_i$ for all $i \in I$.

2.2 Necessity and sufficiency of conditions for smooth convex interpolation

Our goal is therefore to identify a set of necessary of sufficient conditions involving the set of data triples and characterizing the existence of an interpolating function. Finding necessary conditions is relatively easy: starting from any set of necessary conditions that holds on the whole domain of a smooth strongly convex function, one can simply restrict this set to those conditions involving only points x_i with $i \in I$ (i.e. to discretize it). For example, it is well-known that the class of L -smooth convex functions $\mathcal{F}_{0,L}$ is characterized by the pair of inequalities

$$\begin{aligned} f(y) &\geq f(z) + \nabla f(z)^T(y - z), & \forall y, z \in \mathbb{R}^d, \\ \|\nabla f(y) - \nabla f(z)\|_2 &\leq L\|y - z\|_2, & \forall y, z \in \mathbb{R}^d. \end{aligned} \quad (\text{C1})$$

Therefore, specializing those conditions for $y = x_i$ and $z = x_j$ with $i, j \in I$ leads to the following set of inequalities, which is *necessary* for the existence of an interpolating function in $\mathcal{F}_{0,L}$

$$\begin{aligned} f_i &\geq f_j + g_j^T(x_i - x_j), & \forall i, j \in I, \\ \|g_i - g_j\|_2 &\leq L\|x_i - x_j\|_2, & \forall i, j \in I. \end{aligned} \quad (\text{C1f})$$

Now, perhaps surprisingly, it turns out that this latter set of conditions is *not sufficient* to guarantee interpolability by a function in $\mathcal{F}_{0,L}$, despite the fact that the originating conditions (C1) *are sufficient* to guarantee that $f \in \mathcal{F}_{0,L}$. In order to see that, consider the following:

$$\begin{aligned} (x_1, g_1, f_1) &= (-1, -2, 1), \\ (x_2, g_2, f_2) &= (0, -1, 0). \end{aligned}$$

This pair can clearly not be interpolated by a function of the appropriate class, as there is an unavoidable non-differentiability at x_1 . However, it satisfies Conditions (C1f) with $L = 1$, which is therefore not sufficient to guarantee smooth convex interpolation.

Similarly, we can carry out the same exercise for the following conditions, also well-known to be equivalent to inclusion on $\mathcal{F}_{0,L}$ when imposed on the whole space:

$$\begin{aligned} f_i &\geq f_j + g_j^T(x_i - x_j), & \forall i, j \in I, \\ f_i &\leq f_j + g_j^T(x_i - x_j) + \frac{L}{2}\|x_i - x_j\|_2^2, & \forall i, j \in I. \end{aligned} \quad (\text{C2f})$$

With an appropriate use of an additional dimension, one can readily observe that some information may be hidden to this pair of inequalities: consider

$$\begin{aligned}(x_1, g_1, f_1) &= \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 \\ 0 \end{pmatrix}, 0 \right), \\ (x_2, g_2, f_2) &= \left(\begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 \\ 1 \end{pmatrix}, 1 \right),\end{aligned}$$

from which no smooth convex interpolation can be made (again, unavoidable non-differentiability at both x_1 and x_2). However, both Conditions C1f and C2f are satisfied with $L = 1$.

Those examples illustrate the weakness of a naive approach that consists in discretizing standard necessary and sufficient conditions defined on the whole space. If those discretized conditions were used in a performance estimation problem over a given class of functions $\mathcal{F}_{\mu,L}$, they would implicitly allow the performance of functions that do not belong to the class $\mathcal{F}_{\mu,L}$ to be taken into account. This would correspond to a relaxation of the original performance estimation problem, and would only lead to upper bounds on the worst-case performance.

In the next subsections, we follow a more principled approach in order to tackle the $\mathcal{F}_{\mu,L}$ -interpolation problem. We start with a special case of convex interpolation, that of proper convex functions without smoothness or strong convexity requirement (i.e. the class $\mathcal{F}_{0,\infty}$), for which a solution is well-known.

2.3 Convex interpolation

In order to build the previous interpolation conditions for the class of smooth strongly convex functions, we begin by constructing interpolation conditions for the simpler class of convex functions $\mathcal{F}_{0,\infty}$. As this result will be one of building blocks of the smooth strongly convex interpolation procedure, a simple constructive proof of this theorem is provided.

Theorem 1 (Convex interpolation) *Set $\{(x_i, g_i, f_i)\}_{i \in I}$ is $\mathcal{F}_{0,\infty}$ -interpolable if and only if*

$$f_i \geq f_j + g_j^\top (x_i - x_j) \quad \forall i, j \in I. \quad (1)$$

Proof (Necessity.) Assume there exists a convex function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ such that $f_i = f(x_i)$ and $g_i \in \partial f(x_i) \forall i \in I$. Definition of subgradient then immediately implies that

$$f_i \geq f_j + g_j^\top (x_i - x_j) \quad \forall i, j \in I.$$

(Sufficiency.) Define the following piecewise-linear convex function

$$f(x) = \max_{j \in I} \left\{ f_j + g_j^\top (x - x_j) \right\}.$$

Since f is the pointwise maximum of a finite number of affine functions, its epigraph is a non-empty polyhedron, hence f is convex, closed and proper. In addition, $f(x_i) = f_i$ holds by construction. Indeed, we first see that

$$\begin{aligned}f_i &= f_i + g_i^\top (x_i - x_i) \\ &\leq \max_{j \in I} \left\{ f_j + g_j^\top (x_i - x_j) \right\} = f(x_i).\end{aligned}$$

Therefore, we have $f_i \leq f(x_i)$. In addition to this, we have

$$\begin{aligned}f(x_i) &= \max_{j \in I} \left\{ f_j + g_j^\top (x_i - x_j) \right\}, \\ &\leq f_i \quad \text{using Condition (1) for each } j,\end{aligned}$$

which allows to conclude that $f(x_i) = f_i$. The construction also implies that $g_i \in \partial f(x_i)$:

$$\begin{aligned}f(x) &= \max_{j \in I} \left\{ f_j + g_j^\top (x - x_j) \right\} \quad \forall x \in \mathbb{R}^d, \\ &\geq f_i + g_i^\top (x - x_i) \quad \forall k \in I, x \in \mathbb{R}^d, \\ &\geq f(x_i) + g_i^\top (x - x_i) \quad \forall k \in I, x \in \mathbb{R}^d.\end{aligned}$$

□

One should note that the effective domain is $\text{dom } f = \mathbb{R}^d$ — that is, the function takes finite values for all $x \in \mathbb{R}^d$. This is of course not the only way of reconstructing a valid f . For example, we could choose

$$f(x) = \begin{cases} \max_j \{f_j + g_j^\top(x - x_j)\} & \text{if } x \in \text{conv}(\{x_i\}_{i \in I}), \\ +\infty & \text{otherwise.} \end{cases}$$

Remark 1 Our interpolation problem is an extension of the classical finite convex integration problem, which is concerned with the recovery of a convex function from a set of points x_i associated with a subgradient g_i (i.e. function values are not specified). Finite convex integration is treated in details in [12] (only in the convex case $\mu = 0$ and $L = +\infty$). In this context, the so-called cyclic monotonicity conditions are necessary and sufficient in order to be able to recover a convex function. Finite convex integration problem is the finite version of the continuous convex integrability problem, which is treated in [20].

2.4 Conjugation and minimal curvature subtraction

In this section, we review some concepts and results needed for our generalization of convex interpolation to smooth convex interpolation. We begin with the concept of *conjugation* operation. This operation is a key element in our approach, since it provides a way to reduce the general smooth strongly convex interpolation problem to a simpler convex interpolation problem.

Definition 3 Given a function $f : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$, the (Legendre-Fenchel) conjugate $f^* : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{\pm\infty\}$ of f is defined as:

$$f^*(y) = \sup_{x \in \mathbb{R}^d} y^\top x - f(x).$$

Conjugate functions enjoy numerous useful properties. Among other, they are always closed and convex. In fact, conjugation realizes a one-to-one correspondence in the set of proper closed convex functions (i.e. an involution), as recalled by the next theorem.

Theorem 2 ([20], Theorem 12.2) Consider a function $f \in \mathcal{F}_{0,\infty}$. Then, $f^* \in \mathcal{F}_{0,\infty}$ and $f^{**} = f$.

The next two theorems are standard and link properties and values of proper closed convex functions in terms of the values and properties of their conjugates.

Theorem 3 ([20], Theorem 23.5) Given a function $f \in \mathcal{F}_{0,\infty}$ and its conjugate f^* , the following propositions are equivalent conditions on x and g :

- (a) $f(x) + f^*(g) = g^\top x$,
- (b) $g \in \partial f(x)$,
- (c) $x \in \partial f^*(g)$.

This theorem provides a useful interpretation of conjugation as an operation reversing the roles of the coordinates and the subgradients: any subgradient (resp. coordinate) in one space becomes a coordinate (resp. subgradient) in the second space. This link between a function and its conjugate is exploited in the next section in order to express the interpolation conditions on the conjugate function. The next theorem emphasizes the effect of this link for the class of smooth convex functions.

Theorem 4 Consider a function $f \in \mathcal{F}_{0,\infty}$. We have $f \in \mathcal{F}_{0,L}$ if and only if $f^* \in \mathcal{F}_{1/L,\infty}$.

Theorem 4 is basically Proposition 12.60 of [21] in the case $L < +\infty$ and reduces to Theorem 2 in the case $L = +\infty$. This also explains why we need to include the case $L = +\infty$ in our interpolation problem: this is so that we can include the conjugates of smooth but non strongly convex functions in $\mathcal{F}_{0,L}$.

The next lemma gives a simple way of expressing smooth strongly convex functions in terms of smooth functions.

Theorem 5 Consider a function $f \in \mathcal{F}_{\mu,L}$, the following propositions are equivalent:

- (a) $f \in \mathcal{F}_{\mu,L}$,
- (b) $f(x) - \frac{\mu}{2}\|x\|_2^2 \in \mathcal{F}_{0,L-\mu}$.

This theorem holds true by Definition 1 when $L = +\infty$. The case $L < +\infty$ can be found in the proof of Theorem 2.1.11 in [18].

2.5 Necessary and sufficient conditions for smooth strongly convex interpolation

We now focus on transforming the smooth strongly convex interpolation problem into a convex interpolation problem. In order to do so, we mainly use the two previously defined operations: conjugation (using Theorem 4) and minimal curvature subtraction (using Theorem 5). The reasoning is the following:

- (i) Reformulate the $\mathcal{F}_{\mu,L}$ interpolation problem into a $\mathcal{F}_{0,L-\mu}$ interpolation problem using minimal curvature subtraction.
- (ii) Write the $\mathcal{F}_{0,L-\mu}$ interpolation problem into a $\mathcal{F}_{1/(L-\mu),\infty}$ interpolation problem using Legendre-Fenchel conjugation.
- (iii) Transform the $\mathcal{F}_{1/(L-\mu),\infty}$ interpolation problem into a $\mathcal{F}_{0,\infty}$ interpolation problem using again minimal curvature subtraction.

The effect of minimal curve subtraction on our interpolation problem, used in steps (i) and (iii), is described by the following Lemma.

Lemma 1 *Consider a finite set $\{(x_i, g_i, f_i)\}_{i \in I}$ with $x_i, g_i \in \mathbb{R}^d$ and $f_i \in \mathbb{R}$. The following propositions are equivalent for any constants $0 \leq \mu < L \leq +\infty$:*

- (a) $\{(x_i, g_i, f_i)\}_{i \in I}$ is $\mathcal{F}_{\mu,L}$ -interpolable,
- (b) $\{(x_i, g_i - \mu x_i, f_i - \frac{\mu}{2}\|x_i\|_2^2)\}_{i \in I}$ is $\mathcal{F}_{0,L-\mu}$ -interpolable.

Proof (a) \Rightarrow (b) It follows from Theorem 5 that if there exists $f \in \mathcal{F}_{\mu,L}$ interpolating the set, then $h(x) = f(x) - \frac{\mu}{2}\|x\|_2^2$ exists and satisfies $h \in \mathcal{F}_{0,L-\mu}$ and $\forall i \in I$:

$$h(x_i) = f_i - \frac{\mu}{2}\|x_i\|_2^2, \quad g_i - \mu x_i \in \partial h(x_i).$$

Hence, the function $h \in \mathcal{F}_{0,L-\mu}$ interpolates the set $\{(x_i, g_i - \mu x_i, f_i - \frac{\mu}{2}\|x_i\|_2^2)\}_{i \in I}$.

(a) \Leftarrow (b) If such a $h \in \mathcal{F}_{0,L-\mu}$ exists and satisfies the interpolation conditions (b), then one can reconstruct a function $f(x) = h(x) + \frac{\mu}{2}\|x\|_2^2$, $f \in \mathcal{F}_{\mu,L}$ which interpolates the set $\{(x_i, g_i, f_i)\}_{i \in I}$. \square

The effect of conjugation in step (ii) of the reduction procedure is precisely given in the following lemma.

Lemma 2 *Consider a finite set $\{(x_i, g_i, f_i)\}_{i \in I}$ with $x_i, g_i \in \mathbb{R}^d$ and $f_i \in \mathbb{R}$. The following propositions are equivalent for any constant $0 < L \leq +\infty$:*

- (a) $\{(x_i, g_i, f_i)\}_{i \in I}$ is $\mathcal{F}_{0,L}$ -interpolable,
- (b) $\{(g_i, x_i, x_i^\top g_i - f_i)\}_{i \in I}$ is $\mathcal{F}_{1/L,\infty}$ -interpolable.

Proof (a) \Rightarrow (b) It follows from Theorem 4 that if there exists $f \in \mathcal{F}_{0,L}$ then f^* exists and satisfies $f^* \in \mathcal{F}_{1/L,\infty}$. In addition to that, Theorem 3 states that if both $f \in \mathcal{F}_{0,L}$ and f^* exists, then they satisfy $\forall i \in I$ the three conditions:

$$f(x_i) + f^*(g_i) = g_i^\top x_i, \quad g_i \in \partial f(x_i), \quad x_i \in \partial f^*(g_i).$$

(b) \Rightarrow (a) If a function $f^* \in \mathcal{F}_{1/L,\infty}$ exists and satisfies the interpolation conditions (b), then the conjugate f^{**} (which is convex, proper and closed by construction) satisfies $f^{**} \in \mathcal{F}_{0,L}$ by Theorem 4, and the interpolation conditions by Theorem 3: $\forall i \in I$:

$$f^{**}(x_i) + f^*(g_i) = g_i^\top x_i, \quad g_i \in \partial f^{**}(x_i), \quad x_i \in \partial f^*(g_i).$$

We obtain the desired result by choosing $f = f^{**}$. \square

We are now properly armed in order to define all interpolations equivalences. Let us now use steps (i), (ii) and (iii) in order to prove the main theorem of this section.

Theorem 6 ($\mathcal{F}_{\mu,L}$ -interpolability) *Set $\{(x_i, g_i, f_i)\}_{i \in I}$ is $\mathcal{F}_{\mu,L}$ -interpolable if and only if the following set of conditions holds for every pair of indices $i \in I$ and $j \in I$*

$$f_i - f_j - g_j^\top (x_i - x_j) \geq \frac{1}{2(1-\mu/L)} \left(\frac{1}{L} \|g_i - g_j\|_2^2 + \mu \|x_i - x_j\|_2^2 - 2\frac{\mu}{L} (g_j - g_i)^\top (x_j - x_i) \right).$$

Proof We begin by showing the equivalence between the following propositions:

- (a) $\{(x_i, g_i, f_i)\}_{i \in I}$ is $\mathcal{F}_{\mu, L}$ -interpolable,
- (b) $\{(x_i, g_i - \mu x_i, f_i - \frac{\mu}{2} \|x_i\|_2^2)\}_{i \in I}$ is $\mathcal{F}_{0, L-\mu}$ -interpolable,
- (c) $\{(g_i - \mu x_i, x_i, x_i^\top g_i - f_i - \frac{\mu}{2} \|x_i\|_2^2)\}_{i \in I}$ is $\mathcal{F}_{1/(L-\mu), \infty}$ -interpolable,
- (d) $\left\{ \left(g_i - \mu x_i, \frac{Lx_i}{L-\mu} - \frac{g_i}{L-\mu}, \frac{Lx_i^\top g_i}{L-\mu} - f_i - \frac{L\mu \|x_i\|_2^2}{2(L-\mu)} - \frac{\|g_i\|_2^2}{2(L-\mu)} \right) \right\}_{i \in I}$ is $\mathcal{F}_{0, \infty}$ -interpolable,
- (e) $\left\{ \left(\frac{Lx_i}{L-\mu} - \frac{g_i}{L-\mu}, g_i - \mu x_i, \frac{\mu x_i^\top g_i}{L-\mu} + f_i - \frac{L\mu \|x_i\|_2^2}{2(L-\mu)} - \frac{\|g_i\|_2^2}{2(L-\mu)} \right) \right\}_{i \in I}$ is $\mathcal{F}_{0, \infty}$ -interpolable.

Both (a) \Leftrightarrow (b) and (c) \Leftrightarrow (d) are direct application of Lemma 1, whereas both (b) \Leftrightarrow (c) and (d) \Leftrightarrow (e) are direct application of Lemma 2. Theorem 6 follows from equivalence between propositions (a) and (e) applied to the necessary and sufficient conditions for convex interpolation of Theorem 1. Finally, it is straightforward to check that condition (e) reduces to the statement of the Theorem. \square

Remark 2 Note that one can also easily construct an interpolating function $f(x)$ for the original set of points from Theorem 6(a). It follows from Theorem 1 that a possible interpolating function for the set $\left\{ \left(\tilde{x}_i, \tilde{g}_i, \tilde{f}_i \right) \right\}_{i \in I}$ of Theorem 6(c) is given by

$$h(\tilde{x}) = \max_i \left\{ \tilde{f}_i + \tilde{g}_i^\top (\tilde{x} - \tilde{x}_i) + \frac{1}{2(L-\mu)} \|\tilde{x} - \tilde{x}_i\|_2^2 \right\} = \max_i h_i(\tilde{x}).$$

This can be conjugated into an interpolating function $h^*(x)$ of the set given by Theorem 6(b). Using Theorem 16.5 from [20], this can equivalently be written in the form

$$h^*(x) = \text{conv} (h_i^*(x)),$$

where the conv operation takes the convex hull of the epigraphs of h_i^* 's. Hence an interpolating function for the original set $\{(x_i, g_i, f_i)\}_{i \in I}$ is given by

$$f(x) = \text{conv} (h_i^*(x)) + \frac{\mu}{2} \|x\|_2^2.$$

It is straightforward to establish the equivalent interpolation conditions for both the nonsmooth strongly convex case ($L = +\infty$) and the smooth but non-strongly convex case ($\mu = 0$). In the first case — given by Corollary 1 — we find the discrete version of the well-known inequality characterizing L -smooth convex functions, which turns out to be necessary and sufficient

$$f(x) \geq f(y) + \nabla f^\top(y)(x - y) + \frac{1}{2L} \|\nabla f(x) - \nabla f(y)\|_2^2.$$

Corollary 1 *The finite set $\{(x_i, g_i, f_i)\}_{i \in I}$ is $\mathcal{F}_{0, L}$ -interpolable if and only if*

$$f_i \geq f_j + g_j^\top (x_i - x_j) + \frac{1}{2L} \|g_i - g_j\|_2^2, \quad \forall i, j \in I.$$

Strongly convex interpolation conditions are given in Corollary 2, which corresponds to the well-known inequality characterizing the subgradients of strongly convex functions.

Corollary 2 *The finite set $\{(x_i, g_i, f_i)\}_{i \in I}$ is $\mathcal{F}_{\mu, \infty}$ -interpolable if and only if*

$$f_i \geq f_j + g_j^\top (x_i - x_j) + \frac{\mu}{2} \|x_i - x_j\|_2^2, \quad \forall i, j \in I.$$

Finally, we note that our Theorem 6 can readily be extended to handle the L -smooth μ -strongly convex finite and continuous integration problems (i.e. interpolation without specified function values), with the use of a cyclic monotonicity condition, which directly extends the results from [12, 20]. In particular, the following inequality that is standard in the analysis of gradient methods on smooth strongly convex functions (see e.g. Theorem 2.1.12 of [18])

$$(g_i - g_j)^\top (x_i - x_j) \geq \frac{1}{1 + \mu/L} \left(\frac{1}{L} \|g_i - g_j\|_2^2 + \mu \|x_i - x_j\|_2^2 \right),$$

which holds for all $\mathcal{F}_{\mu, L}$ -interpolable set $\{(x_i, g_i, f_i)\}_{i \in I}$, can be simply obtained from Theorem 6 (by summing original inequality and the one with indices i and j exchanged).

3 A convex formulation for performance estimation

As explained in the introduction, our performance estimation problem can now be expressed in terms of the iterates and optimal point $\{x_i, g_i, f_i\}_{i \in \{0, \dots, N, *\}}$ only, using the interpolation conditions given by Theorem 6.

As our class of functions $\mathcal{F}_{\mu, L}$ and the first-order methods we study are invariant with respect to both an additive shift in the function values and a translation in their domain, we can assume without loss of generality $x_* = 0$ and $f_* = 0$, which will simplify our derivations. We also assume $g_* = 0$, from optimality conditions. The problem can now be stated in its finite dimensional formulation:

$$w_d(\mathcal{F}_{\mu, L}, R, \mathcal{M}, N, \mathcal{P}) = \sup_{\{x_i, g_i, f_i\}_{i \in I \setminus \{*\}}} \mathcal{P}(\{x_i, g_i, f_i\}_{i \in I}), \quad (\text{f-PEP2})$$

$$\text{s.t. } \{x_i, g_i, f_i\}_{i \in I} \text{ is } \mathcal{F}_{\mu, L}\text{-interpolable,}$$

$$x_1, \dots, x_N \text{ is generated by method } \mathcal{M},$$

$$\|x_0 - x_*\|_2 \leq R.$$

In the non strongly convex case ($\mu = 0$) and for the gradient method, we see that interpolation conditions from Corollary 1 induce the formulation proposed in [8](G). Our results now establish that it is exact, i.e. not a relaxation.

This optimization problem is strictly equivalent to the original (PEP) in terms of optimal value, since every solution to (f-PEP2) can be interpolated by a solution of (PEP) and reciprocally every solution of (PEP) can be discretized to provide a solution to (f-PEP2). It is however not convex, as it involves several non-convex quadratic constraints (e.g. $g_j^\top x_i$ terms in the interpolation conditions). In the next section, we show how (f-PEP2) can be cast as a convex semidefinite program [23] when dealing with a certain class of first-order black-box methods, those with fixed steps.

3.1 Convex formulation of PEPs for fixed-step first-order algorithms

We hereby restrict ourselves to the class of fixed-step first-order methods, which correspond to the following iterates

$$x_i = x_0 - \frac{1}{L} \sum_{k=0}^{i-1} h_{i,k} g_k, \quad (2)$$

where the step sizes are taken relative to the starting point x_0 . Note that many classical methods such as fixed-step gradient method (GM) and fast-gradient methods (FGM) are included in this class of algorithms (see the details in Section 4).

In order to obtain a convex formulation for (f-PEP2), we use a Gram matrix. We denote

$$P = [g_0 \ g_1 \ \dots \ g_N \ x_0],$$

and formulate the problem in terms of the entries of the positive semidefinite Gram matrix $G = P^\top P$, along with the function values f_i . Note that only iterate x_0 appears in the Gram matrix, as other iterates depend directly from x_0 and the gradient values g_i , as expressed by the fixed-step method.

For notational convenience, we define vectors $h_i \in \mathbb{R}^{N+2}$ for any i between 0 and N and $h_* \in \mathbb{R}^{N+2}$ as follows

$$Lh_i^\top = [-h_{i,0} \ -h_{i,1} \ \dots \ -h_{i,i-1} \ 0 \ \dots \ 0 \ L], \quad h_*^\top = [0 \ \dots \ 0],$$

such that $x_i = Ph_i$. In order to lighten the notations we also define $u_i = e_{i+1} \in \mathbb{R}^{N+2}$, the canonical basis vectors and u_* the vector of zeros. Using those notations, we rewrite the interpolation constraints coming from Theorem 6 in the following form:

$$\begin{aligned} f_i &\geq f_j + \frac{L}{L-\mu} (u_j^\top G h_i - u_j^\top G h_j) + \frac{1}{2(L-\mu)} (u_i - u_j)^\top G (u_i - u_j) \\ &\quad + \frac{\mu}{L-\mu} (u_i^\top G h_j - u_i^\top G h_i) + \frac{L\mu}{2(L-\mu)} (h_i - h_j)^\top G (h_i - h_j), \quad i, j \in I. \end{aligned}$$

We can equivalently formulate all constraints using the trace operator, which allows writing them more compactly:

$$\begin{aligned} f_j - f_i + \text{Tr}(GA_{ij}) &\leq 0, & i, j \in I, \\ \text{Tr}(GA_R) - R^2 &\leq 0, \\ G &\succeq 0, \end{aligned}$$

where the matrices A_{ij} and A_R are defined in the following way:

$$\begin{aligned} 2A_{ij} &= \frac{L}{L-\mu} \left(u_j(h_i - h_j)^\top + (h_i - h_j)u_j^\top \right) + \frac{1}{L-\mu} (u_i - u_j)(u_i - u_j)^\top \\ &\quad + \frac{\mu}{L-\mu} \left(u_i(h_j - h_i)^\top + (h_j - h_i)u_i^\top \right) + \frac{L\mu}{L-\mu} (h_i - h_j)(h_i - h_j)^\top, \\ A_R &= u_{N+1}u_{N+1}^\top. \end{aligned}$$

The interpolation and starting point constraints are now formulated linearly in the entries of the Gram matrix G and in the function values f . Note that this formulation allows recovering a dimension-independent worst-case, as in the original (PEP) — that is, we compute the worst-case functions without any restriction on the dimension d of their domains. This independence justified in different situations, such as for example when considering large-scale optimization problems — that is when $d \gg N$, which is the setting we study in this paper. An upper bound on the dimension of the domain of function f could be formulated by imposing a (non-convex) rank constraint $\text{rank}(G) \leq d$. Of course, the reverse statement also applies: any feasible solution of rank d can be interpolated into a d -dimensional function $f \in \mathcal{F}_{\mu,L}$.

Looking now at the performance criterion \mathcal{P} , we observe that any concave semidefinite-representable function G and f leads to a worst-case estimation problem that can be cast as a convex semidefinite optimization problem (see e.g. [4]). In particular, linear functions of the entries of f and G are suitable. Classical performance criteria such as $f(x_N) - f_*$, $\|\nabla f(x_N)\|_2^2$ and $\|x_N - x_*\|_2^2$ are indeed covered by this formulation. We focus below on the case of a linear performance criterion, but note that other criteria can be useful (see for example a concave piecewise linear criteria used in Section 4.3).

One can also note that, when L is finite, any continuous performance criterion will force the optimal value of the (f-PEP2) to be attained, because the constraints on variables f and G force the domain to be closed and bounded. On the other hand when $L = +\infty$, the problem is unbounded and it is possible to design feasible functions which drive standard performance criteria arbitrarily away from 0. Nevertheless, performance estimation on such nonsmooth functions could still be tackled after introduction of another appropriate Lipschitz condition on the functions, such as $\|g_i\|_2 \leq L$, a topic which we leave for further research. In the rest of this text, we will restrict ourselves to the smooth case $L < +\infty$.

Proposition 1 *Consider a first-order fixed-step algorithm given by $H = \{h_i\}_{i \in \{1, \dots, N\}}$ and define the linear criterion $\mathcal{P} = b^T f + \text{Tr}(CG)$ with $b \in \mathbb{R}^{N+1}$ and $C \in \mathbb{S}^{N+2}$. The worst-case performance for criterion \mathcal{P} of the method H applied on the class of function $\mathcal{F}_{\mu,L}$ is the optimal value of the convex PEP:*

$$\begin{aligned} w_{sdp}(L, \mu, R, H, N, C, b) &= \max_{G \in \mathbb{S}^{N+2}, f \in \mathbb{R}^{N+1}} b^T f + \text{Tr}(CG) && \text{(sdp-PEP)} \\ \text{s.t. } f_j - f_i + \text{Tr}(GA_{ij}) &\leq 0, && i, j \in I, \\ \text{Tr}(GA_R) - R^2 &\leq 0, \\ G &\succeq 0. \end{aligned}$$

Proof On the one hand, given H , \mathcal{P} , N and R , an optimal solution of (PEP) on the class $\mathcal{F}_{\mu,L}$ can be expressed as a feasible point of (sdp-PEP) by construction of the matrices G , A_{ij} and A_R . Therefore, the optimal solution of (PEP) is feasible for (sdp-PEP). Hence,

$$w(\mathcal{F}_{\mu,L}, R, H, N, \mathcal{P}) \leq w_{sdp}(L, \mu, R, H, N, C, b).$$

On the other hand, from a solution (G, f) of (sdp-PEP), one can recover a matrix P such that $G = P^\top P$, using positive semidefiniteness of G (e.g. using Cholesky decomposition). Choosing $x_i = Ph_i$, $g_i = Pu_i$

($i = 0, \dots, N, *$) and $f_* = 0$ we obtain a $\mathcal{F}_{\mu,L}$ -interpolable set. Hence it is possible to construct a function $f \in \mathcal{F}_{\mu,L}$ feasible for (PEP). Therefore, we have

$$w_{sdp}(L, \mu, R, H, N, C, b) \leq w(\mathcal{F}_{\mu,L}, R, H, N, \mathcal{P}),$$

and we conclude that $w_{sdp}(L, \mu, R, H, N, C, b) = w(\mathcal{F}_{\mu,L}, R, H, N, \mathcal{P})$. \square

In general, it is not easy to find an analytical optimal solution to (sdp-PEP). Hence, we are also interested in a generic and easier way of obtaining analytical upper bounds on the performance of a given algorithm. A classical way of doing so is to work with the Lagrangian dual of (sdp-PEP):

$$\begin{aligned} \min_{\lambda_{ij}, \tau} \quad & \tau R^2 & (\text{d-sdp-PEP}) \\ \text{s.t.} \quad & \tau A_R - C + \sum_{i,j \in I} \lambda_{ij} A_{ij} \succeq 0, \\ & b - \sum_{i,j \in I} \lambda_{ij} (u_j - u_i) = 0, \\ & \lambda_{ij} \geq 0, \quad i, j \in I, \\ & \tau \geq 0, \end{aligned}$$

whose feasible solutions will provide theoretical upper bounds on the worst-case behaviour of every fixed-step first-order method (using weak duality). Note that the final dual formulation used in [8], which deals with the case $\mu = 0$, can be recovered by taking $\lambda_{ij} = 0$ for $i + 1 \neq j$ or $i \neq *$ in our dual, i.e. it is a restriction of (d-sdp-PEP) with a potentially larger optimal value.

The next lemma guarantees that no duality gap occurs between (sdp-PEP) and (d-sdp-PEP) under the technical assumption that $h_{i,i-1} \neq 0$ ($i \in \{1, \dots, N\}$). This assumption is reasonable as it only implies that the last gradient obtained from the oracle has to be used at every iteration. The lemma will also also guarantees the existence of a dual feasible point attaining the optimal value of the primal-dual pair of estimation problems (sdp-PEP) and (d-sdp-PEP), i.e. a tight upper bound on the worst-case performance of the considered method.

Proposition 2 *The optimal value of the dual problem (d-sdp-PEP) with $0 \leq \mu < L$ is attained and equal to $w_{sdp}(H, C, b, N, R, L, \mu)$ under the assumption that $h_{i,i-1} \neq 0$ ($i \in \{1, \dots, N\}$).*

Proof We use the classical Slater condition [6] on the primal problem in order to guarantee a zero duality gap — that is, we show that (sdp-PEP) has a feasible point with $G \succ 0$. The reasoning is divided in two parts; we consider first the case $\mu = 0$ and $L = 2 + 2 \cos(\pi/(N+2))$ and we generalize it to general $\mu < L$ afterwards. Consider the quadratic function

$$f(x) = \frac{1}{2} x^\top Q x,$$

with the tridiagonal positive definite matrix

$$Q = \begin{pmatrix} 2 & 1 & 0 & 0 & \dots & 0 \\ 1 & 2 & 1 & 0 & \dots & 0 \\ 0 & 1 & 2 & 1 & \ddots & \vdots \\ \vdots & & \ddots & \ddots & \ddots & \ddots \end{pmatrix} \succ 0.$$

We show how to construct a full-rank G feasible for (sdp-PEP) using the values of the quadratic function f . In order to do so, we exhibit a full-rank matrix

$$P = [x_0 \ g_0 \ g_1 \ \dots \ g_N]$$

corresponding to the application of a given method (with $h_{i,i-1} \neq 0$) to the quadratic function f . Indeed, choosing $x_0 = R e_1$, we can show that P is upper triangular with non-zero diagonal entries. We have:

$$\begin{aligned} g_0 &= Q x_0 = 2e_1 + e_2, \\ x_1 &= x_0 - \frac{h_{1,0}}{L} g_0, \\ g_1 &= Q x_1 = g_0 - \frac{h_{1,0}}{L} Q g_0 = 2e_1 + e_2 - \frac{h_{1,0}}{L} (4e_1 + 4e_2 + e_3). \end{aligned}$$

Hence g_1 has a non-zero element associated with e_3 whereas the only non-zero elements of g_0 are associated with e_1 and e_2 . Now, we assume that g_{i-1} has a non-zero element corresponding to e_{i+1} and zero elements corresponding to $e_k \forall k > i + 1$, while all previous gradients have zero components corresponding to $e_k \forall k > i$. We have:

$$g_i^\top e_{i+2} = x_i^\top Q e_{i+2} = x_i^\top (e_{i+1} + 2e_{i+2} + e_{i+3}),$$

with $x_i^\top e_{i+2} = x_i^\top e_{i+3} = 0$ and $x_i^\top e_{i+1} \neq 0$ because of the recurrence assumption and the iterative form of the algorithm:

$$\begin{aligned} x_i^\top e_{i+1} &= \underbrace{x_0^\top e_{i+1}}_{=0} - \sum_{k=0}^{i-2} \frac{h_{i,k}}{L} \underbrace{g_k^\top e_{i+1}}_{=0} - \frac{h_{i,i-1}}{L} \underbrace{g_{i-1}^\top e_{i+1}}_{\neq 0}, \\ x_i^\top e_{i+2} &= \underbrace{x_0^\top e_{i+2}}_{=0} - \sum_{k=0}^{i-2} \frac{h_{i,k}}{L} \underbrace{g_k^\top e_{i+2}}_{=0} - \frac{h_{i,i-1}}{L} \underbrace{g_{i-1}^\top e_{i+2}}_{=0}, \\ x_i^\top e_{i+3} &= \underbrace{x_0^\top e_{i+3}}_{=0} - \sum_{k=0}^{i-2} \frac{h_{i,k}}{L} \underbrace{g_k^\top e_{i+3}}_{=0} - \frac{h_{i,i-1}}{L} \underbrace{g_{i-1}^\top e_{i+3}}_{=0}. \end{aligned}$$

Hence, g_i has a non-zero element associated with e_{i+2} . We obtain that the following components are zero by computing $g_i^\top e_{i+2+k}$ for $k > 0$:

$$g_i^\top e_{i+2+k} = x_i^\top Q e_{i+2+k} = x_i^\top (e_{i+1+k} + 2e_{i+2+k} + e_{i+3+k}),$$

which is zero because of the algorithmic structure of x_i :

$$x_i^\top e_{i+1+k} = \underbrace{x_0^\top e_{i+1+k}}_{=0} - \sum_{k=0}^{i-2} \frac{h_{i,k}}{L} \underbrace{g_k^\top e_{i+1+k}}_{=0} - \frac{h_{i,i-1}}{L} \underbrace{g_{i-1}^\top e_{i+1+k}}_{=0}.$$

Hence matrix P is an upper triangular matrix with positive entries on the diagonal, and is therefore full-rank. In order to make this statement hold for general $\mu < L$, observe that the structure of the matrix is preserved using the operation:

$$Q' = (Q - \lambda_{\min}(Q)I_{N+2}) \frac{(L - \mu)}{\lambda_{\max}(Q) - \lambda_{\min}(Q)} + \mu I_{N+2}.$$

The corresponding quadratic function is μ -strongly convex with a L -Lipschitz gradient. Therefore, the interior of the domain of (sdp-PEP) is non-empty and Slater's condition applies for $\mu < L$, ensuring that no duality gap occurs and that the dual optimal value is attained. \square

3.2 Homogeneity of the optimal values with respect to L and R

To conclude this section, we observe that, for most performance criteria, one can predict the dependence of the worst-case performance on parameters L and R , considerably reducing the number of numerical experiments one has to perform (i.e. one only has to solve the performance estimation problem in the case $R=1$ and $L = 1$). With a standard reasoning involving appropriate scaling operations, one can obtain for the standard criteria $f(x_N) - f_*$, $\|\nabla f(x_N)\|_2$ and $\|x_N - x_*\|_2$ the homogeneity relations

$$\begin{aligned} w_{sdp}(L, \mu, R, H, N, f(x_N) - f_*) &= LR^2 w_{sdp}(1, \kappa, 1, H, N, f(x_N) - f_*), \\ w_{sdp}(L, \mu, R, H, N, \|\nabla f(x_N)\|_2) &= LR w_{sdp}(1, \kappa, 1, H, N, \|\nabla f(x_N)\|_2), \\ w_{sdp}(L, \mu, R, H, N, \|x_N - x_*\|_2) &= R w_{sdp}(1, \kappa, 1, H, N, \|x_N - x_*\|_2), \end{aligned}$$

where $\kappa = \mu/L$ is the inverse condition number.

4 Numerical performance estimation of standard first-order algorithms

In this section, we apply the convex PEP formulation to study convergence of the fixed-step gradient method (GM), the standard fast gradient method (FGM) and the optimized gradient method (OGM) proposed by [11]. We begin the section by the study of the GM for smooth convex optimization, whose worst-case is conjectured [8] to be attained on a simple one-dimensional function. Numerical experiments with our exact formulation confirm this conjecture. Further experiments on the worst-case complexity for different methods, problem classes and performance criteria lead to a series of conjectures based on worst-case functions possessing a similar shape. We conclude this section with the study of a nonlinear performance criteria corresponding to the smallest gradient norm among all iterates computed by the method.

All numerical results in this section were obtained on an Intel 3.5Ghz desktop computer using a combination of the YALMIP modeling environment in MATLAB [14], the MOSEK [15] and SeDuMi [22] semidefinite solvers and the VSDP (verified semidefinite programming) toolbox [9].

4.1 Gradient method

4.1.1 A conjecture on smooth convex functions by Drori and Teboulle [8]

Consider the classical fixed-step gradient method (GM) with normalized step sizes, applied to a smooth convex function in $\mathcal{F}_{0,L}$.

Gradient Method (GM)

Input: $f \in \mathcal{F}_{0,L}$, $x_0 \in \mathbb{R}^d$, $y_0 = x_0$.

For $i = 0 : N - 1$

$$x_{i+1} = x_i - \frac{h}{L} \nabla f(x_i)$$

The following conjecture on the convergence of the worst-case objective function values was made in [8].

Conjecture 1 ([8], Conjecture 3.1.) Any sequence $\{x_i\}$ generated by the GM with $0 \leq h \leq 2$ on a function $f \in \mathcal{F}_{0,L}$ satisfies :

$$f(x_N) - f_* \leq \frac{LR^2}{2} \max\left(\frac{1}{2Nh+1}, (1-h)^{2N}\right).$$

A proof of the conjecture is provided in [8] for (relative) step size values $0 \leq h \leq 1$, leaving the case $1 < h < 2$ open. We also recall that the upper bound in this conjecture cannot be improved, as it matches the performance of the GM on two specific one-dimensional functions. Indeed, define

$$f_1(x) = \begin{cases} \frac{LR}{2Nh+1}|x| - \frac{LR^2}{2(2Nh+1)^2} & \text{if } |x| \geq \frac{R}{2Nh+1}, \\ \frac{L}{2}x^2 & \text{else.} \end{cases}$$

$$f_2(x) = \frac{L}{2}x^2.$$

It is straightforward to check that the final objective value accuracy of GM on f_1 is equal to $\frac{1}{2} \frac{LR^2}{2Nh+1}$, and that on f_2 it is equal to $\frac{1}{2} LR^2 (1-h)^{2N}$. This means that the conjecture can be reformulated as saying that the worst-case behavior of the GM according to objective function accuracy is achieved by function f_1 or f_2 , depending on which of the two is worst (which will depend only on the normalized step size h and number of iterations N).

Intuitively, the behavior of GM on piecewise affine-quadratic f_1 corresponds to a situation in which iterates slowly approach the optimal value without oscillating around it (i.e. no overshooting), whereas GM applied on purely quadratic f_2 generates a sequence oscillating around the optimal point. Those behaviors are illustrated on Figure 1. We also note that iterates for f_1 never leave the affine piece of

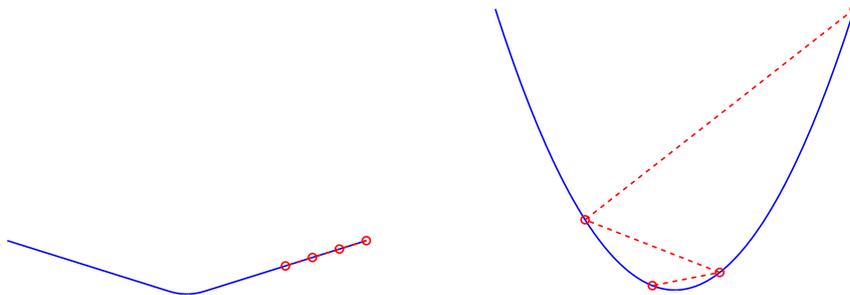


Fig. 1 Behaviour of the gradient method on f_1 (left) and f_2 (right), for $L = R = 1$. We observe that GM does not overshoot the optimal solution on f_1 , while it does so at each iteration on f_2 .

the function. Interestingly, the existence of a one-dimensional worst-case function with a simple affine-quadratic shape will also be observed for the other algorithms studied in this section, both in the smooth convex and in the smooth strongly convex cases.

Empirical results from the numerical resolution of (sdp-PEP) strongly support Conjecture 1. Indeed, when comparing its predictions with numerically computed worst-case bounds, we obtained a maximal relative error of magnitude 10^{-7} (all pairs of values $N \in \{1, 2, \dots, 30\}$ and $h \in \{0.05, 0.010, \dots, 1.95\}$ were tested).

Before going into the details of other methods, we underline another observation coming from [8]: Conjecture 1 also suggests the existence of an optimal step size $h_{\text{opt}}(N)$ for the GM — optimal in the sense of achieving the lowest worst-case. That is, if you know in advance how many iterations of the GM you will perform, it suggests using a step size $h_{\text{opt}}(N)$ that is the unique minimizer of the right hand side of the Conjecture 1 for a fixed value of N . It is obtained by solving¹ the following non-linear equation in h_{opt} (for which no closed form solution seems to be available):

$$\frac{1}{2Nh_{\text{opt}} + 1} = (1 - h_{\text{opt}})^{2N}.$$

This optimal step size can be interpreted in terms of the trade-off between what we obtain on f_1 and f_2 . On the one hand, we ensure that we are not going too slowly to the optimal point on f_1 , and on the other hand we do not want to overshoot too much on f_2 .

Assuming Conjecture 1 holds true, one can show that the optimal step size is an increasing function of N with $3/2 \leq h_{\text{opt}}(N) < 2$ and $h_{\text{opt}}(N) \rightarrow 2$ as $N \rightarrow \infty$. More precisely, working out the expression defining h_{opt} gives the following tight lower and upper estimates:

$$2 - \frac{\log 4N}{2N} \sim 1 + (1 + 4N)^{-1/(2N)} \leq h_{\text{opt}}(N) \leq 1 + (1 + 2N)^{-1/(2N)} \sim 2 - \frac{\log 2N}{2N}.$$

It is interesting to compare the results of [8] with ours for values of the normalized step size h that are close to h_{opt} . Indeed, while the results of the two formulations are similar most of the time, it turns out that those from [8] are significantly more conservative in the zone around h_{opt} , as presented in Table 1 for different values of N . This also formally establishes the fact that the formulation from [8] is a strict relaxation of the performance estimation problem.

These numerical results have been obtained with MOSEK, a standard semidefinite optimization solver. Despite convexity of the formulation, it might happen that the solution returned by such as solver is inaccurate, and in particular (slightly) infeasible. In that case, the objective value of the approximate primal (resp. dual) solution is no longer guaranteed to be a lower (resp. upper) bound on the exact optimal value, hence potentially negating the advantage of an exact convex formulation. For this reason, all numerical results reported in this section have been double checked with an interval arithmetic-based semidefinite optimization solver [9] that returns an interval that is guaranteed to contain the optimal value. These guaranteed bounds are reported in Table 2 for the case $h = 1.5$, which compares them with Conjecture 1.

We can observe that the use of a verified solver does not impact our conclusions about the validity of the conjecture. Moreover, this table is typical of what we observed for all conjectures in this section: all

¹ This equation possesses several solutions, but the optimum is the unique point where the two terms feature derivatives of opposite signs (a necessary and sufficient condition for the maximum of two smooth convex functions). This point can easily be computed numerically with an appropriate bisection method.

N	h_{opt}	Conjecture 1	Value computed in [8]	Rel. error	Value from (sdp-PEP)	Rel. error
1	1.5000	$LR^2/8.00$	$LR^2/8.00$	0.00	$LR^2/8.00$	7e-09
2	1.6058	$LR^2/14.85$	$LR^2/14.54$	2e-02	$LR^2/14.85$	5e-09
5	1.7471	$LR^2/36.94$	$LR^2/32.57$	1e-01	$LR^2/36.94$	1e-08
10	1.8341	$LR^2/75.36$	$LR^2/59.80$	3e-01	$LR^2/75.36$	3e-08
20	1.8971	$LR^2/153.77$	$LR^2/109.58$	4e-01	$LR^2/153.77$	6e-08
30	1.9238	$LR^2/232.85$	$LR^2/156.23$	5e-01	$LR^2/232.85$	7e-08
40	1.9388	$LR^2/312.21$	$LR^2/201.10$	6e-01	$LR^2/312.21$	3e-08
50	1.9486	$LR^2/391.72$	$LR^2/244.70$	6e-01	$LR^2/391.72$	1e-07
100	1.9705	$LR^2/790.22$	$LR^2/451.72$	7e-01	$LR^2/790.22$	1e-07

Table 1 Gradient Method with $\mu = 0$, worst-case computed with relaxation from [8] and worst-case obtained by exact formulation (sdp-PEP) for the criterion $f(x_N) - f^*$. Error is measured relatively to the conjectured result. Results obtained with MOSEK [15].

N	1	2	5	10	15	20	30
Relative upper	2e-09	7e-10	2e-09	1e-09	9e-10	1e-09	9e-10
Conjecture	1.2e-01	7.1e-02	3.1e-02	1.6e-02	1.1e-02	8.2e-03	5.5e-03
Relative lower	2e-09	3e-09	9e-09	9e-08	2e-07	3e-07	9e-07

Table 2 Gradient method with relative step size $h = 1.5$: numerical values from Conjecture 1 and guaranteed relative error interval obtained numerically with VSDP [9] and SeDuMi [22].

numerical results reported were validated², and in what follows we will no longer mention this verification explicitly.

Finally, we compare on Figure 2 results obtained with Conjecture 1 with a classical analytical bound from the literature [18] for the GM with unit normalized step size $h = 1$ (which is usually recommended, and sometimes called optimal):

$$f(x_N) - f_* \leq \frac{2LR^2}{N+4}. \quad (3)$$

We observe an improvement by a factor approximately 8 between Conjecture 1 with $h = 1$ and bound (3), and a additional factor 2 between the conjecture with $h = 1$ and the conjecture with the optimal normalized step size h_{opt} (one can check these constants become exact asymptotically).

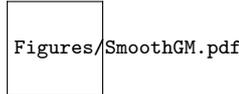


Fig. 2 Comparison between the classical analytical bound (3) on GM with standard $h = 1$ (red), Conjecture 1 with standard $h = 1$ (dashed, black) and Conjecture 1 with optimal step sizes (blue). The criterion used is $f(x_N) - f_*$.

4.1.2 A generalized conjecture for strongly convex functions

In view of the encouraging results obtained for the GM in the smooth case, we now study the behavior of the GM on the class of strongly convex functions $\mathcal{F}_{\mu,L}$ using our formulation (sdp-PEP) with the same performance criterion, objective function accuracy. It turns out that the solution for every problem consisted again in a one-dimensional worst-case function (rank $G = 1$) of the same piecewise quadratic type. We therefore introduce the following general definitions for functions $f_{1,\tau}$ and f_2 :

$$f_{1,\tau}(x) = \begin{cases} \frac{\kappa}{2}x^2 + a|x| + b & \text{if } |x| \geq \tau, \\ \frac{1}{2}x^2 & \text{else,} \end{cases}$$

$$f_2(x) = \frac{1}{2}x^2,$$

² Except for tests encountering numerical difficulties, i.e for which VSDP returned no valid interval, which occurred more and more frequently as the value of the worst-case bound became close to zero.

where parameters $a(\tau)$ and $b(\tau)$ are chosen to ensure continuity of $f_{1,\tau}$ and its first derivative, $\kappa = \mu/L$ is the inverse condition number and τ is a parameter controlling the radius of the central quadratic piece, with the largest curvature (note that previously introduced function f_1 corresponds to a special case with $\kappa = 0$).

Although the value of parameter τ could in principle be estimated from the numerical solutions of our problems, it turns out it can be computed optimally with the simple assumption that all iterates of the GM starting at $x_0 = R$ must barely remain in the low-curvature zone of f_1 (which is similar to the corresponding behaviour of the GM for smooth convex functions). Deriving the value of τ maximizing the gap $f_{1,\tau}(x_N) - f_*$ then leads to

$$\tau = \frac{\kappa(1 - h\kappa)^{2N}}{(\kappa - 1)(1 - h\kappa)^{2N} + 1}. \quad (4)$$

We are now able to extend Conjecture 1 to the GM applied to strongly convex functions.

Conjecture 2 Any sequence $\{x_i\}$ generated by the GM with $0 \leq h \leq 2$ on a function $f \in \mathcal{F}_{\mu,L}$ satisfies :

$$f(x_N) - f_* \leq \frac{LR^2}{2} \max(f_{1,\tau}(x_{1,N}), f_2(x_{2,N})),$$

with τ chosen according to (4) and $x_{1,N}$ (resp. $x_{2,N}$) the final iterate from GM applied to $f_{1,\tau}$ (resp. f_2), or equivalently

$$f_{1,\tau}(x_{1,N}) = \frac{\kappa}{(\kappa - 1) + (1 - \kappa h)^{-2N}}, \quad f_2(x_{2,N}) = (1 - h)^{2N}.$$

This conjecture can again be reformulated as saying that the worst-case behaviour of the GM according to objective function accuracy is achieved by function $f_{1,\tau}$ or f_2 , depending on which of the two is worst. Proceeding now to its numerical validation, we first point out that our results are intrinsically limited to accuracy that can be reached by numerical SDP solvers. For this reason, we only report on situations for which Conjecture 2 predicts a final accuracy larger than 10^{-6} , ensuring a few significant digits for the numerical results. The resulting estimated relative differences between Conjecture 2 and the numerical results obtained with (sdp-PEP) are given in Table 3, for different values of κ . We observe that the Conjecture is very well supported by our numerical results, with a largest relative error around 10^{-6} , reached for the largest value of κ considered here. This is expected as GM tends to perform better as κ increases (i.e. final accuracy $f(x_N) - f_*$ approaches zero), which renders a precise comparison between numerical results and the conjecture more and more difficult. We now investigate some consequences of

$\kappa =$	0	.001	.005	.010	.015	0.1	0.2	0.5
Rel. error	6e-10	7e-10	4e-10	6e-10	8e-10	2e-07	9e-08	1e-06

Table 3 Maximum relative estimated differences between Conjecture 2 and corresponding numerical results obtained with SeDuMi [22]. The maximum is taken over all $N \in \{1, \dots, 30\}$ and $h \in \{0.05, \dots, 1.95\}$ for which the conjecture predicts a worst-case larger than 10^{-6} .

the conjecture. First of all, we note that the result of Conjecture 2 tends to the result of Conjecture 1 as $\kappa \rightarrow 0$. Hence our formulation (sdp-PEP) allows closing an apparent gap between worst-case analyses in the smooth convex and the smooth strongly convex cases. Indeed, to the best of our knowledge, existing worst-case bounds for the smooth strongly convex case do not converge to the smooth case as $\mu \rightarrow 0$. In addition to that, as for Conjecture 1, our new Conjecture 2 suggests optimal step sizes $h_{\text{opt}}(N, \kappa)$, which can be obtained by solving the system (for $0 < \kappa < 1$):

$$\frac{\kappa}{(\kappa - 1) + (1 - \kappa h_{\text{opt}})^{-2N}} = (1 - h_{\text{opt}})^{2N}$$

(note that one recovers the previous system for h_{opt} when $\kappa \rightarrow 0$). For a given N , as κ increases from 0 to 1, those optimal step sizes decrease from $h_{\text{opt}}(N, 0)$ (optimal step size in the smooth case) to $h_{\text{opt}}(N, 1) = 1$ (the latter being expected since it can only correspond to the case of function f_2 in the original (PEP), for which the GM with $h = 1$ converges in one iteration). For a given κ , we find that

$h_{\text{opt}}(N, \kappa)$ increases as N increases, as in the smooth convex case, according to the following lower and upper estimates

$$1 + \left(\frac{\kappa - 1}{\kappa} + \frac{1}{\kappa} \left(\frac{1 + \kappa}{1 - \kappa} \right)^{2N} \right)^{-1/(2N)} \leq h_{\text{opt}}(N, \kappa) \leq \min \left\{ 1 + \left(\frac{\kappa - 1}{\kappa} + \frac{1}{\kappa} (1 - \kappa)^{-2N} \right)^{-1/(2N)}, \frac{2}{1 + \kappa} \right\},$$

which both tend to $\frac{2}{1 + \kappa}$ as N increases, which corresponds to the standard recommended step size $\frac{2}{L + \mu}$ for the GM.

We now illustrate the improvements provided by Conjecture 2 with respect to the classical worst-case bound from [18] for the GM applied to functions from $\mathcal{F}_{\mu, L}$:

$$f(x_N) - f_* \leq \frac{LR^2}{2} \left(1 - \frac{2h\kappa}{1 + \kappa} \right)^N. \quad (5)$$

corresponding to a normalized step size $h = \frac{2}{1 + \kappa}$. We plot this bound on Figure 3 along with the values of Conjecture with the same step size and with the optimal step size h_{opt} . We observe that the asymptotic

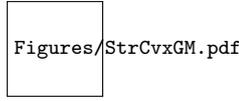


Fig. 3 Comparison between the classical analytical bound (5) on GM with the standard normalized step size $h = \frac{2}{1 + \kappa}$ (red) and Conjecture 2 with standard $h = \frac{2}{1 + \kappa}$ (dashed, black) and optimal normalized step size (blue), for $\kappa = 0.1$. The criterion used is $f(x_N) - f_*$.

convergence rate of the true worst-case appears to match the rate predicted by the standard bound (5). However, the optimal step size allows gaining a constant factor of approximately 10 in comparison with the standard bound.

4.1.3 A conjecture on the residual gradient norm

We now consider a different performance criterion, given by the norm of the gradient computed at the last iterate. Numerical experiments with our formulation suggest that results similar to those presented in the previous sections can be obtained both in the smooth convex and smooth strongly convex cases, based again on one-dimensional piecewise quadratic worst-case functions. Using the same definition for functions f_1 and f_2 and choosing now for parameter τ

$$\tau = \frac{\kappa(1 - h\kappa)}{(\kappa - 1)(1 - h\kappa)^N + 1}, \quad (6)$$

we propose the following conjecture.

Conjecture 3 Any sequence $\{x_i\}$ generated by the GM with $0 \leq h \leq 2$ on a function $f \in \mathcal{F}_{\mu, L}$ satisfies :

$$\|\nabla f(x_N)\|_2 \leq LR \max(|f'_{1, \tau}(x_{1, N})|, |f'_{2, \tau}(x_{2, N})|),$$

with τ chosen according to (6) and $x_{1, N}$ (resp. $x_{2, N}$) the final iterate from GM applied on $f_{1, \tau}$ (resp. f_2), or equivalently

$$|f'_{1, \tau}(x_{1, N})| = \frac{\kappa}{(\kappa - 1) + (1 - \kappa h)^{-N}}, \quad |f'_{2, \tau}(x_{2, N})| = |1 - h|^N.$$

As for Conjecture 2, we limit our numerical validation to the cases where the worst-case values predicted by the Conjecture are larger than 10^{-6} . The resulting estimated relative errors are given in Table 4, where we observe a largest relative error around 10^{-7} .

We note that, as κ tends to zero (i.e. the smooth case), Conjecture 3 tends to

$$\|\nabla f(x_N)\|_2 \leq LR \max\left(\frac{1}{Nh + 1}, |1 - h|^N\right).$$

$\kappa =$	0	.001	.005	.010	.015	0.1	0.2	0.5
Rel. error	4e-10	6e-10	8e-10	5e-10	2e-09	8e-08	9e-08	9e-07

Table 4 Maximum relative estimated differences between Conjecture 3 and corresponding numerical results obtained with SeDuMi [22]. The maximum is taken over all $N \in \{1, \dots, 30\}$ and $h \in \{0.05, \dots, 1.95\}$ for which the conjecture predicts a worst-case larger than 10^{-6} .

From that, we see that the optimal step size $h_{\text{opt}}^{\nabla}(N, 0)$ for the GM is again an increasing function of N with $\sqrt{2} \leq h_{\text{opt}}^{\nabla}(N, 0) < 2$ and $h_{\text{opt}}^{\nabla}(N, 0) \rightarrow 2$ as $N \rightarrow \infty$. The optimal step size for the case $\kappa > 0$ is a decreasing function of κ and satisfies $h_{\text{opt}}^{\nabla}(N, \kappa) \rightarrow 1$ as $\kappa \rightarrow 1$. As in the previous case, $h_{\text{opt}}^{\nabla}(N, \kappa)$ is bounded above by $\frac{2}{1+\kappa}$, which we can confirm with the following lower and upper bounds on h_{opt}^{∇} :

$$1 + \left(\frac{\kappa - 1}{\kappa} + \left(\frac{1 + \kappa}{1 - \kappa} \right)^N \right)^{-1/N} \leq h_{\text{opt}}^{\nabla}(N, \kappa) \leq \min \left\{ 1 + \left(\frac{\kappa - 1}{\kappa} + \frac{1}{\kappa} (1 - \kappa)^{-N} \right)^{-1/N}, \frac{2}{1 + \kappa} \right\}.$$

In the smooth case, those bounds reduce to the simpler

$$2 - \frac{\log 2N}{N} \sim 1 + (1 + 2N)^{-1/N} \leq h_{\text{opt}}^{\nabla} \leq 1 + (1 + N)^{-1/N} \sim 2 - \frac{\log N}{N}.$$

We now compare with a standard analytical worst-case bound. Noting that the iterates of the GM method satisfy [18]:

$$\|x_N - x_*\|_2 \leq R \left(1 - \frac{2h\kappa}{1 + \kappa} \right)^{N/2},$$

and using the L -Lipschitz property of the gradient along with $\nabla f(x_*) = 0$, we obtain:

$$\|\nabla f(x_N)\|_2 \leq L \|x_N - x_*\|_2 \leq LR \left(1 - \frac{2h\kappa}{1 + \kappa} \right)^{N/2}. \quad (7)$$

Figure 4 illustrates the comparison between bound (7) with standard normalized step size $h = \frac{2}{1+\kappa}$ and values of Conjecture with the same step size and with the optimal step size. We observe that the analytic bound derived above matches the asymptotic rate observed in our computations. Using the optimal step size provides an improvement by a factor approximately 8.

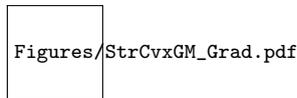


Fig. 4 Comparison between the classical analytical bound (7) on GM with the standard normalized step size $h = \frac{2}{1+\kappa}$ (red) and Conjecture 3 with standard $h = \frac{2}{1+\kappa}$ (dashed, black) optimal normalized step size (blue), for $\kappa = 0.1$. The criterion used is $\|\nabla f(x_N)\|_2$.

4.2 Fast gradient method and optimized gradient method

In this section we assess the performance in the smooth convex case $\mu = 0$ of two accelerated first-order methods: the so-called fast gradient method (FGM) due to Nesterov [17], and an optimized gradient method (OGM) recently proposed by Kim and Fessler [11].

Fast Gradient Method (FGM)	Optimized Gradient Method (OGM)
Input: $f \in \mathcal{F}_{0,L}$, $x_0 \in \mathbb{R}^d$, $y_0 = x_0$, $\theta_0 = 1$.	Input: $f \in \mathcal{F}_{0,L}$, $x_0 \in \mathbb{R}^d$, $y_0 = x_0$, $\theta_0 = 1$.
For $i = 0 : N - 1$	For $i = 0 : N - 1$
$y_{i+1} = x_i - \frac{1}{L} \nabla f(x_i)$	$y_{i+1} = x_i - \frac{1}{L} \nabla f(x_i)$
$\theta_{i+1} = \frac{1 + \sqrt{4\theta_i^2 + 1}}{2}$	$\theta_{i+1} = \begin{cases} \frac{1 + \sqrt{4\theta_i^2 + 1}}{2}, & i \leq N - 2 \\ \frac{1 + \sqrt{8\theta_i^2 + 1}}{2}, & i = N - 1 \end{cases}$
$x_{i+1} = y_{i+1} + \frac{\theta_i - 1}{\theta_{i+1}}(y_{i+1} - y_i)$	$x_{i+1} = y_{i+1} + \frac{\theta_i - 1}{\theta_{i+1}}(y_{i+1} - y_i) + \frac{\theta_i}{\theta_{i+1}}(y_{i+1} - x_i)$

Both of these algorithms are defined in terms of two sequences: $\{y_i\}_i$ is a primary sequence, and $\{x_i\}_i$ is a secondary sequence, where the gradient is evaluated. We first show that both of these algorithms can be expressed as fixed-step first-order methods. Recall we define such a method using

$$x_i = x_0 - \frac{1}{L} \sum_{k=0}^{i-1} h_{i,k} \nabla f(x_k).$$

Hence, we will focus on the secondary sequence $\{x_i\}_i$, and substitute the y_i 's in the algorithm formulation. For FGM, we obtain:

$$\begin{aligned} x_{i+1} &= x_i - \frac{g_i}{L} + \frac{\theta_i - 1}{\theta_{i+1}} \left(x_i - x_{i-1} - \frac{g_i}{L} + \frac{g_{i-1}}{L} \right), \\ &= x_i + \frac{\theta_i - 1}{\theta_{i+1}} (x_i - x_{i-1}) - \left(\frac{\theta_i - 1}{\theta_{i+1}} + 1 \right) \frac{g_i}{L} + \frac{\theta_i - 1}{\theta_{i+1}} \frac{g_{i-1}}{L}, \end{aligned}$$

which allows obtaining the step sizes relative to x_0 by recurrence:

$$h_{i,k} = \begin{cases} \frac{\theta_i - 1}{\theta_{i+1}} (h_{i-1,k} - h_{i-2,k}) & \text{if } k \leq i - 3, \\ \frac{\theta_i - 1}{\theta_{i+1}} (h_{i-1,k} - 1) & \text{if } k = i - 2, \\ \frac{\theta_i - 1}{\theta_{i+1}} + 1 & \text{if } k = i - 1, \end{cases}$$

with initial conditions $h_{1,0} = 1$, $h_{1,k} = 0$ if $k < 0$ and $h_{0,k} = 0 \forall k$. Similarly, for OGM:

$$h_{i,k} = \begin{cases} \frac{\theta_i - 1}{\theta_{i+1}} (h_{i-1,k} - h_{i-2,k}) & \text{if } k \leq i - 3, \\ \frac{\theta_i - 1}{\theta_{i+1}} (h_{i-1,k} - 1) & \text{if } k = i - 2, \\ \frac{2\theta_i - 1}{\theta_{i+1}} + 1 & \text{if } k = i - 1, \end{cases}$$

with the same initial conditions. This approach provides estimates for the last secondary iterate x_N . If an estimate for last primary iterate y_N is needed, one has just to replace the expression of x_N by y_N , which is done by using the following alternative coefficients for the last step:

$$h_{N,k} = \begin{cases} h_{N-1,k} & \text{if } k \leq N - 2, \\ 1 & \text{if } k = N - 1, \end{cases}$$

for both FGM and OGM.

Again, our numerical experiments strongly suggest we make the same assumption about the shape of the worst-case functions, i.e. one-dimensional and piecewise quadratic (iterates staying in the affine zone of $f_{1,\tau}$). Using this, we are able compute values of τ achieving the worst-case final accuracy $f_1(x_N) - f_*$:

$$\tau = \frac{1}{2 \sum_{k=0}^{N-2} h_{N-1,k} + 3},$$

for the primary sequences and

$$\tau = \frac{1}{2 \sum_{k=0}^{N-1} h_{N,k} + 1},$$

for the secondary sequences. It is interesting to note that the same formula for τ holds for both the classical FGM and the more recent optimized OGM.

Our numerical results about those worst-case structures lead us to the following conjectures.

Conjecture 4 Any (primary) sequence $\{y_i\}$ generated by the FGM (resp. OGM) on a function $f \in \mathcal{F}_{0,L}$ satisfies:

$$f(y_N) - f_* \leq \frac{LR^2}{2} f_{1,\tau}(y_{1,N}),$$

with $y_{1,N}$ the final (primary) iterate from FGM (resp. OGM) applied on $f_{1,\tau}$, and

$$f_{1,\tau}(y_{1,N}) = \frac{1}{2 \sum_{k=0}^{N-2} h_{N-1,k} + 3},$$

where quantities $h_{N-1,k}$ are the fixed coefficients of the last step of FGM (resp. OGM).

Conjecture 5 Any (secondary) sequence $\{x_i\}$ generated by the FGM (resp. OGM) on a function $f \in \mathcal{F}_{0,L}$ satisfies:

$$f(x_N) - f_* \leq \frac{LR^2}{2} f_{1,\tau}(x_{1,N}),$$

with $x_{1,N}$ the final (secondary) iterate from FGM (resp. OGM) applied on $f_{1,\tau}$, and

$$f_{1,\tau}(x_{1,N}) = \frac{1}{2 \sum_{k=0}^{N-1} h_{N,k} + 1},$$

where quantities $h_{N,k}$ are the fixed coefficients of the last step of FGM (resp. OGM).

The numerical validations for both Conjecture 4 and 5 with a grid $N \in \{1, \dots, 100\}$ allowed to obtain a relative error of magnitude 10^{-4} — the worst value is obtained among the large number of iterations for the OGM, for which the precision of the solvers becomes more critical when compared with the actual worst-case values. A subset of those numerical comparisons are given in Table 5 and Table 6.

N	Primary sequence $\{y_i\}_i$			Secondary sequence $\{x_i\}_i$		
	Conjecture 4	Value from (sdp-PEP)	Rel. error	Conjecture 5	Value from (sdp-PEP)	Rel. error
1	$LR^2/6.00$	$LR^2/6.00$	4e-09	$LR^2/6.00$	$LR^2/6.00$	4e-09
2	$LR^2/10.00$	$LR^2/10.00$	8e-10	$LR^2/11.13$	$LR^2/11.13$	5e-11
5	$LR^2/28.66$	$LR^2/28.66$	2e-08	$LR^2/33.03$	$LR^2/33.03$	6e-09
10	$LR^2/81.07$	$LR^2/81.07$	4e-09	$LR^2/90.69$	$LR^2/90.69$	7e-09
20	$LR^2/263.65$	$LR^2/263.65$	1e-08	$LR^2/283.55$	$LR^2/283.55$	2e-09
30	$LR^2/548.51$	$LR^2/548.51$	3e-09	$LR^2/578.58$	$LR^2/578.58$	5e-10
40	$LR^2/934.89$	$LR^2/934.89$	2e-08	$LR^2/975.10$	$LR^2/975.10$	1e-08
50	$LR^2/1422.45$	$LR^2/1422.45$	2e-09	$LR^2/1472.76$	$LR^2/1472.76$	2e-10
100	$LR^2/5371.90$	$LR^2/5371.90$	1e-08	$LR^2/5472.54$	$LR^2/5472.54$	1e-07

Table 5 Fast Gradient Method, convergence of primary and secondary sequences for the criterion $f_N - f^*$. The error is measured relatively to the conjectured results. Results obtained with MOSEK [15].

N	Primary sequence $\{y_i\}_i$			Secondary sequence $\{x_i\}_i$		
	Conjecture 4	Value from (sdp-PEP)	Rel. error	Conjecture 5	Value from (sdp-PEP)	Rel. error
1	$LR^2/8.00$	$LR^2/8.00$	1e-08	$LR^2/6.00$	$LR^2/6.00$	4e-09
2	$LR^2/16.16$	$LR^2/16.16$	3e-08	$LR^2/12.47$	$LR^2/12.47$	6e-09
5	$LR^2/53.80$	$LR^2/53.80$	2e-07	$LR^2/45.42$	$LR^2/45.42$	2e-08
10	$LR^2/159.07$	$LR^2/159.07$	4e-07	$LR^2/143.23$	$LR^2/143.23$	5e-10
20	$LR^2/525.09$	$LR^2/525.09$	7e-07	$LR^2/494.68$	$LR^2/494.68$	4e-08
30	$LR^2/1095.62$	$LR^2/1095.62$	1e-06	$LR^2/1050.81$	$LR^2/1050.81$	2e-07
40	$LR^2/1869.22$	$LR^2/1869.21$	3e-06	$LR^2/1810.08$	$LR^2/1810.08$	9e-08
50	$LR^2/2845.15$	$LR^2/2845.14$	5e-06	$LR^2/2771.72$	$LR^2/2771.72$	1e-07
100	$LR^2/10748.13$	$LR^2/10747.97$	1e-05	$LR^2/10603.52$	$LR^2/10603.52$	9e-08

Table 6 Optimized Gradient Method, convergence of primary and secondary sequences for the criterion $f_N - f^*$. The error is measured relatively to the conjectured results. Results obtained with MOSEK [15].

In order to develop some insight in those rather non-intuitive worst-case forms, let us compare the numerical values obtained with Conjectures 4 and 5 with analytical bounds for the FGM. The analytical bound we use for the primary sequence of FGM is developed in [3]:

$$f(y_N) - f_* \leq \frac{2LR^2}{(N+1)^2}, \quad (8)$$

whereas the analytical bound we use for the secondary sequence was very recently derived in [11]:

$$f(x_N) - f_* \leq \frac{2LR^2}{(N+2)^2}. \quad (9)$$

The comparison is given by Figure 5. The asymptotic behaviours of the sequences are well captured by the analytical bounds (8) and (9), but we observe that the estimation of the transient worst-cases are indeed improved by our conjectures: a factor approximately 1.15 on the analytical bounds for both sequences after 30 iterations is gained.

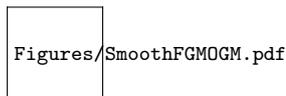


Fig. 5 Comparison (for the FGM) between the analytical bound (8) (magenta) and Conjecture 4 (red) and the analytical bound (9) (black) and Conjecture 5 (blue).

4.3 Estimation of the smallest residual gradient norm among all iterates

First-order methods are often used in dual approaches, where, in addition to objective function accuracy, residual norm of the gradient plays an important role. Indeed, this quantity controls primal feasibility of the iterates (see e.g. [7]).

Considering for example the accelerated FGM in the smooth case, we know from the previous section that the worst-case accuracy for a function in $\mathcal{F}_{0,L}$ is $\frac{2LR^2}{(N+1)^2}$. From that bound, it is easy to obtain a similar bound on the last residual gradient, using Corollary 1:

$$\|\nabla f(y_N)\|_2 \leq \sqrt{2L(f(y_N) - f_*)} \leq \frac{2LR}{N+1}. \quad (10)$$

Observe that this asymptotic rate is significantly worse than that of the objective function accuracy. Moreover, it is often the case that this residual norm is not monotonically decreasing among iterates. Hence, in this section, we will estimate the worst-case performance of FGM according to the best observed residual gradient norm among all iterates:

$$\min_{i \in \{0, \dots, N\}} \|\nabla f(y_i)\|_2.$$

In order to do so, only a slightly modified version of (sdp-PEP) is needed: this min-type objective function is representable using a new variable t for the objective and $N+1$ additional linear inequalities $t \leq \|\nabla f(y_i)\|_2^2 \Leftrightarrow t \leq G_{i,i}$. Note that since this concave piecewise linear objective function remains continuous, the maximum is still attained.

This criterion was suggested in [19], which proposes a variant of FGM that consists in performing $N/2$ steps of the standard FGM followed by $N/2$ steps of the GM with $h = 1$. It is then theoretically established that this variant of FGM satisfies

$$\min_{i \in \{0, \dots, N\}} \|\nabla f(y_i)\|_2 \leq \frac{8LR}{N^{3/2}}. \quad (11)$$

which is an improvement compared to the rate of convergence of the previous bound.

We now compare FGM with this modified variant MFGM using our performance estimation formulation. Figure 6 compares the behaviours of those methods in both their last (for FGM) and best iterates, as well as the above analytic bounds (10) and (11). This experiment confirms that the residual gradient of the last iterate of FGM decreases according to (10), i.e. with a $\mathcal{O}(N^{-1})$ rate. We also observe that FGM without modification achieves the same $\mathcal{O}(N^{-3/2})$ convergence as the MFGM variant. In addition, numerical results in Table 7 suggest that FGM performs slightly better than MFGM.

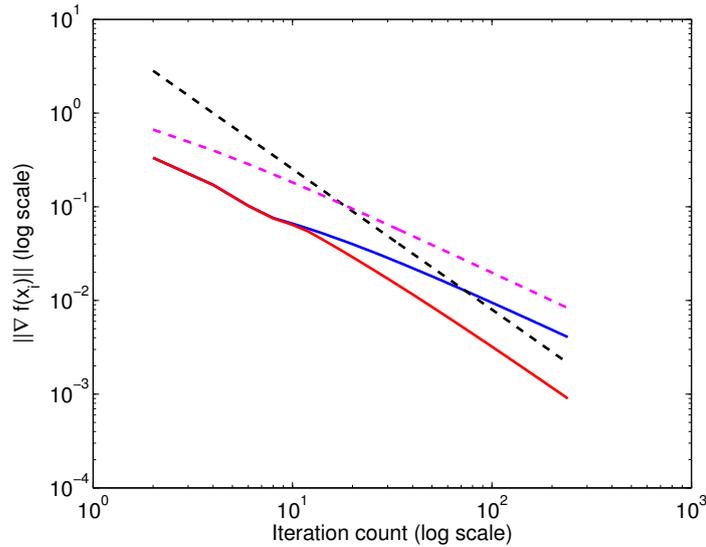


Fig. 6 Comparison of residual gradient convergence rate for the FGM and the MFGM from [19]. Theoretical guarantees are dashed. Analytical bound on FGM (10) in its last iterate (dashed, magenta); numerical worst-case for FGM at its last iterate (blue); numerical worst-case for FGM at its best iterate (red); analytical bound on MFGM (11) for the best iterate (dashed, black).

N	FGM, analytic (10)	FGM, last	FGM, best	MFGM, analytic (11)	MFGM, best
2	$LR/1.50$	$LR/3.00$	$LR/3.00$	$LR/0.35$	$LR/3.00$
4	$LR/2.50$	$LR/5.84$	$LR/5.84$	$LR/1.00$	$LR/5.00$
10	$LR/5.50$	$LR/15.14$	$LR/15.62$	$LR/3.95$	$LR/12.66$
20	$LR/10.50$	$LR/25.08$	$LR/34.49$	$LR/11.18$	$LR/30.77$
30	$LR/15.50$	$LR/35.13$	$LR/58.50$	$LR/20.54$	$LR/55.38$
40	$LR/20.50$	$LR/45.19$	$LR/86.17$	$LR/31.62$	$LR/86.41$
50	$LR/25.50$	$LR/55.25$	$LR/117.08$	$LR/44.19$	$LR/119.63$
100	$LR/50.50$	$LR/105.49$	$LR/311.34$	$LR/125.00$	$LR/296.58$
200	$LR/100.50$	$LR/205.77$	$LR/850.59$	$LR/353.55$	$LR/791.87$

Table 7 FGM and MFGM: comparison between theoretical bounds and numerical results for the criteria $\|\nabla f(x_N)\|_2$ (last) and $\min_i \|\nabla f(x_i)\|_2$ (best). Results obtained with [15].

A regularization technique is also describe in [19], with a convergence rate $\mathcal{O}(N^{-2})$ up to a logarithmic factor. A drawback of this approach is that it requires a bound on the distance to the optimal solution, and that the coefficients of the method explicitly depend on this bound. No fixed-step method achieving the same $\mathcal{O}(N^{-2})$ seems to be known.

5 Conclusion

The contribution of this paper is threefold: first, we have presented necessary and sufficient conditions for smooth strongly convex interpolation. Those conditions were derived by showing an explicit way of constructing the interpolating functions. Second, we showed that the exact worst-case performance of any fixed-step first-order algorithm for smooth strongly convex optimization could be formulated as a convex program. In this context, the interpolation conditions allowed to guarantee the existence of a function reaching the bounds obtained with the SDP formulation of the performance estimation problem. Third, we numerically test of our formulation on a variety of functions classes, methods and performance criteria, establishing on a way a series of conjectures on the corresponding worst-case behaviors. In particular, we suggest new tight estimates of the optimal step size for the fixed-step gradient method, which depend on the number of iterations and the condition number.

Our Performance Estimation problem provide a generic tool to analyse fixed-step first-order methods. It allows computing both exact worst-case guarantees and functions reaching them, and provides a unified algorithmic analysis for smooth convex functions and smooth strongly convex functions.

The exact worst-case values provided by our approach require solving a convex problem whose size grows as the square of the number of iterations considered, which may become prohibitive when this number of iterations is large. This can be avoided using iteration-independent bounds, as proposed in [13], but at the cost of obtaining poorer worst-case guarantees.

Further improvements to our approach include an extension of (PEP) to more general methods, such as first-order methods equipped with line search, or first-order methods designed to work on a restricted convex feasible region (projected gradient). Another desirable feature would be the ability to optimize the step sizes of the method considered in (sdp-PEP), as was proposed in [8,11] for the relaxed version of (PEP).

Software. Our semidefinite programming approach to performance estimation has been implemented with MATLAB code, which can be downloaded from <http://perso.uclouvain.be/adrien.taylor>. This routine features an easy-to-use interface, which allows the estimation of the worst-case performance of a given fixed-step first-order algorithm (to be chosen among a pre-defined list or to be specified by its coefficients) on a given class of functions, for a given performance criterion, after any number of steps.

References

1. Bauschke, H., Combettes, P.: *Convex analysis and monotone operator theory in Hilbert spaces*. Springer (2011)
2. Beck, A.: Quadratic matrix programming. *SIAM Journal on Optimization* **17**(4), 1224–1238 (2007)
3. Beck, A., Teboulle, M.: A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences* **2**(1), 183–202 (2009)
4. Ben-Tal, A., Nemirovski, A.: *Lectures on modern convex optimization: analysis, algorithms, and engineering applications*, vol. 2. Siam (2001)
5. Bertsekas, D.P.: *Convex Optimization Theory*. Athena Scientific (2009)
6. Boyd, S., Vandenberghe, L.: *Convex Optimization*. Cambridge University Press (2004)
7. Devolder, O., Glineur, F., Nesterov, Y.: Double smoothing technique for large-scale linearly constrained convex optimization. *SIAM Journal on Optimization* **22**(2), 702–727 (2012)
8. Drori, Y., Teboulle, M.: Performance of first-order methods for smooth convex minimization: a novel approach. *Mathematical Programming* **145**(1-2), 451–482 (2014)
9. Härter, V., Jansson, C., Lange, M.: VSDP: A matlab toolbox for verified semidefinite-quadratic-linear programming. Tech. rep., Technical report, Institute for Reliable Computing, Hamburg University of Technology, 2012. (2012)
10. Hiriart-Urruty, J.B., Lemaréchal, C.: *Convex Analysis and Minimization Algorithms*. Springer Verlag, Heidelberg (1996). Two volumes - 2nd printing
11. Kim, D., Fessler, J.: Optimized first-order methods for smooth convex minimization. arXiv preprint arXiv:1406.5468 (2014)
12. Lambert, D., Crouzeix, J.P., Nguyen, V., Strodiot, J.J.: Finite convex integration. *Journal of Convex Analysis* **11**(1), 131–146 (2004)
13. Lessard, L., Recht, B., Packard, A.: Analysis and design of optimization algorithms via integral quadratic constraints. arXiv preprint arXiv:1408.3595 (2014)
14. Löfberg, J.: YALMIP : A toolbox for modeling and optimization in MATLAB. In: *Proceedings of the CACSD Conference* (2004)
15. Mosek, A.: The MOSEK optimization software. Online at <http://www.mosek.com> **54** (2010)
16. Nemirovsky, A., Yudin, D.: *Problem complexity and method efficiency in optimization*. Willey-Interscience, New York (1983)
17. Nesterov, Y.: A method of solving a convex programming problem with convergence rate $O(1/k^2)$. *Soviet Mathematics Doklady* **27**, 372–376 (1983)
18. Nesterov, Y.: *Introductory lectures on convex optimization : a basic course*. Applied optimization. Kluwer Academic Publ. (2004)
19. Nesterov, Y.: How to make the gradients small. *Optima* **88**, 10–11 (2012)
20. Rockafellar, R.: *Convex Analysis*. Princeton University Press (1996)
21. Rockafellar, R., Wets, R.B.: *Variational Analysis*. Springer (1998)
22. Sturm, J.F.: Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. *Optimization Methods and Software* **11–12**, 625–653 (1999)
23. Vandenberghe, L., Boyd, S.: Semidefinite programming. *SIAM Review* **38**, 49–95 (1994)

Recent titles

CORE Discussion Papers

- 2014/43 Julio DAVILA. The rationality of expectations formation.
- 2014/44 Florian MAYNERIS, Sandra PONCET and Tao ZHANG. The cleaning effect of minimum wages. Minimum wages, firm dynamics and aggregate productivity in China.
- 2014/45 Thierry BRECHET, Natali HRITONENKOVA and Yuri YATSENKO. Domestic environmental policy and international cooperation for global commons.
- 2014/46 Mathieu PARENTI, Philip USHCHEV and Jacques-François THISSE. Toward a theory of monopolistic competition.
- 2014/47 Takatoshi TABUCHI, Jacques-François THISSE and Xiwei ZHU. Does technological progress affect the location of economic activity?
- 2014/48 Paul CASTANEDA DOWER, Victor GINSBURGH and Shlomo WEBER. Colonial legacy, linguistic disenfranchisement and the civil conflict in Sri Lanka.
- 2014/49 Victor GINSBURGH, Jacques MELITZ and Farid TOUBAL. Foreign language learnings: An econometric analysis.
- 2014/50 Koen DECANCO and Dirk NEUMANN. Does the choice of well-being measure matter empirically? An illustration with German data.
- 2014/51 François MANIQUET. Social ordering functions.
- 2014/52 Ivar EKELAND and Maurice QUEYRANNE. Optimal pits and optimal transportation.
- 2014/53 Luc BAUWENS, Manuela BRAIONE and Giuseppe STORTI. Forecasting comparison of long term component dynamic models for realized covariance matrices.
- 2014/54 François MANIQUET and Philippe MONGIN. Judgment aggregation theory can entail new social choice results.
- 2014/55 Pasquale AVELLA, Maurizio BOCCIA and Laurence A. WOLSEY. Single-period cutting planes for inventory routing problems.
- 2014/56 Jean-Pierre FLORENS and Sébastien VAN BELLEGEM. Instrumental variable estimation in functional linear models.
- 2014/57 Abdelrahman ALY and Mathieu VAN VYVE. Securely solving classical networks flow problems.
- 2014/58 Henry TULKENS. Internal vs. core coalitional stability in the environmental externality game: A reconciliation.
- 2014/59 Manuela BRAIONE and Nicolas K. SCHOLTES. Construction of Value-at-Risk forecasts under different distributional assumptions within a BEKK framework.
- 2014/60 Jörg BREITUNG and Christian M. HAFNER. A simple model for now-casting volatility series.
- 2014/61 Timo TERASVIRTA and Yukai YANG. Linearity and misspecification tests for vector smooth transition regression models.
- 2014/62 Timo TERASVIRTA and Yukai YANG. Specification, estimation and evaluation of vector smooth transition autoregressive models with applications.
- 2014/63 Axel GAUTIER and Nicolas PETIT. Optimal enforcement of competition policy: the commitments procedure under uncertainty.
- 2014/64 Sébastien BROOS and Axel GAUTIER. Competing one-way essential complements: the forgotten side of net neutrality.
- 2014/65 Jean HINDRIKS and Yukihiro NISHIMURA. On the timing of tax and investment in fiscal competition models.
- 2014/66 Jean HINDRIKS and Guillaume LAMY. Back to school, back to segregation?
- 2014/67 François MANIQUET et Dirk NEUMANN. Echelles d'équivalence du temps de travail: évaluation de l'impôt sur le revenu en Belgique à la lumière de l'éthique de la responsabilité.
- 2015/01 Yurii NESTEROV and Vladimir SHIKHMAN. Algorithm of Price Adjustment for Market Equilibrium.
- 2015/02 Claude d'ASPREMONT and Rodolphe DOS SANTOS FERREIRA. Oligopolistic vs. monopolistic competition: Do intersectoral effects matter?

Recent titles

CORE Discussion Papers - continued

- 2015/03 Yuuri NESTEROV. Complexity bounds for primal-dual methods minimizing the model of objective function.
- 2015/04 Hassène AISSI, A. Ridha MAHJOUR, S. Thomas MCCORMICK and Maurice QUEYRANNE. Strongly polynomial bounds for multiobjective and parametric global minimum cuts in graphs and hypergraphs.
- 2015/05 Marc FLEURBAEY and François MANIQUET. Optimal taxation theory and principles of fairness.
- 2015/06 Arnaud VANDAELE, Nicolas GILLIS, François GLINEUR and Daniel TUYTTENS. Heuristics for exact nonnegative matrix factorization.
- 2015/07 Luc BAUWENS, Jean-François CARPANTIER and Arnaud DUFAYS. Autoregressive moving average infinite hidden Markov-switching models.
- 2015/08 Koen DECANQ, Marc FLEURBAEY and François MANIQUET. Multidimensional poverty measurement with individual preferences
- 2015/09 Eric BALANDRAUD, Maurice QUEYRANNE, and Fabio TARDELLA. Largest minimally inversion-complete and pair-complete sets of permutations.
- 2015/10 Maurice QUEYRANNE and Fabio TARDELLA. Carathéodory, helly and radon numbers for sublattice convexities.
- 2015/11 Takatoshi TABUSHI, Jacques-François THISSE and Xiwei ZHU. Does technological progress affect the location of economic activity.
- 2015/12 Mathieu PARENTI, Philip USHCHEV, Jacques-François THISSE. Toward a theory of monopolistic competition.
- 2015/13 A.B. TAYLOR, J.M. HENDRICKX and F. GLINEUR. Smooth strongly convex interpolation and exact worst-case performance of first-order methods.

Books

- W. GAERTNER and E. SCHOKKAERT (2012), *Empirical Social Choice*. Cambridge University Press.
- L. BAUWENS, Ch. HAFNER and S. LAURENT (2012), *Handbook of Volatility Models and their Applications*. Wiley.
- J-C. PRAGER and J. THISSE (2012), *Economic Geography and the Unequal Development of Regions*. Routledge.
- M. FLEURBAEY and F. MANIQUET (2012), *Equality of Opportunity: The Economics of Responsibility*. World Scientific.
- J. HINDRIKS (2012), *Gestion publique*. De Boeck.
- M. FUJITA and J.F. THISSE (2013), *Economics of Agglomeration: Cities, Industrial Location, and Globalization*. (2nd edition). Cambridge University Press.
- J. HINDRIKS and G.D. MYLES (2013). *Intermediate Public Economics*. (2nd edition). MIT Press.
- J. HINDRIKS, G.D. MYLES and N. HASHIMZADE (2013). *Solutions Manual to Accompany Intermediate Public Economics*. (2nd edition). MIT Press.
- J. HINDRIKS (2015). *Quel avenir pour nos pensions ? Les grands défis de la réforme des pensions*. De Boeck.

CORE Lecture Series

- R. AMIR (2002), Supermodularity and Complementarity in Economics.
- R. WEISMANTEL (2006), Lectures on Mixed Nonlinear Programming.
- A. SHAPIRO (2010), Stochastic Programming: Modeling and Theory.