Performance Estimation of First-Order Methods on Quadratic Functions ⋆

Nizar Bousselmi * **Julien M. Hendrickx** * **François Glineur** *

* *ICTEAM, UCLouvain, 1348 Louvain-la-Neuve, Belgium (e-mail: {nizar.bousselmi,julien.hendrickx,francois.glineur}@uclouvain.be).*

Abstract: We are interested in determining the worst performance exhibited by a given first-order optimization method on the class of quadratic functions. Since its introduction, the *Performance Estimation Problem* (PEP) methodology has allowed the computation of the exact worst-case performance of first-order optimization methods on several functions classes, including smooth convex, strongly convex or nonconvex functions.
In this work, we extend the PEP framework to the class of quadratic functions, and apply it to analyze the difference of performance of the gradient method between convex quadratic and general smooth convex functions.

Keywords: Performance estimation, First-order methods, Quadratic functions, Linear matrix inequalities

1. INTRODUCTION

The *Performance Estimation Problem* (PEP) methodology (introduced by Drori and Teboulle (2014)) allows to compute the exact worst-case performance of a first-order optimization method on a given class of functions. More precisely, given a method and a performance criterion (lower is better), a PEP is an optimization problem that maximizes this criterion among all possible functions belonging to some class. Thus, it provides the worst possible behavior of the method on the class of functions.

It has been shown in Taylor et al. (2017) that a PEP can be reformulated as a convex semidefinite program for a wide range of function classes $\mathcal{C}$. This provided several tight results on the performance of first-order methods. In particular, the worst-case behavior of the *Gradient Method* (GM) on the class $\mathcal{F}_{\mu,L}$ of L-smooth μ-strongly convex functions was exhaustively covered.

In this work, we extend the PEP framework to function classes defined by matrices. This typically allows to study the worst-case performance of first-order methods on the class $\mathcal{Q}_{\mu,L}$ of homogeneous quadratic functions of the form $f(x) = \frac{1}{2}x^TQx$ with $\mu I \preceq Q \preceq LI$ for given parameters μ and L ($0 \leq \mu \leq L$). Another type of classes newly analyzable through our extension of the PEP are function classes $\mathcal{C}_1$ and $\mathcal{C}_2$ of the form $g(Ax)$ and $h(x) + g(Ax)$. These three classes turn out to be included in $\mathcal{F}_{\mu,L}$ if we define the smoothness and strong convexity parameters of A, g and h in a proper way. Since the worst-case functions of $\mathcal{F}_{\mu,L}$ for (GM), found in Taylor et al. (2017), are sometimes but not always quadratic or of the form $g(Ax)$, we will quantify the performance gap between the general class $\mathcal{F}_{\mu,L}$ and the classes $\mathcal{Q}_{\mu,L}$, $\mathcal{C}_1$, $\mathcal{C}_2$ or other function classes involving matrices that we can now analyze through the PEP framework.

⋆ N. Bousselmi is supported by the French Community of Belgium through a FRIA fellowship (F.R.S-FNRS).

Theorems are stated without proofs in this extended abstract, they will be appear in a forthcoming paper.

2. PEP FORMULATION

Typically, a PEP can be formulated as follows. Given the class of functions $\mathcal{C}$, the optimization method $\mathcal{M}$ performing N iterations, the initial distance R and the classical performance criterion $f(x_N) - f^*$ (objective function accuracy after N iterations), the PEP is

$$\begin{aligned} \max_{x_0,\dots,x_N,f} \quad & f(x_N) - f^* \\ \text{s.t.} \quad & f \in \mathcal{C}, \\ & x_k \text{ generated by applying } \mathcal{M} \text{ to } f \text{ from } x_0, \\ & \|x_0 - x^*\| \leq R. \end{aligned} \tag{PEP}$$

We can study any method $\mathcal{M}$ that computes each iterate as a linear combination of the initial point x_0 and the gradients of the previous iterations, i.e.

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

Coefficients $h_{k,i}$ entirely describe the method $\mathcal{M}$. For example, the gradient method with constant step size $\frac{1}{L}$ started from x_0:

For $i = 0 : N-1$

$$\begin{aligned} x_{i+1} &= x_i - \frac{1}{L}\nabla f(x_i) \qquad \text{(GM)} \\ &= x_0 - \frac{1}{L}\sum_{i=0}^{k-1} \nabla f(x_i). \end{aligned}$$

is described with

$$\begin{cases} h_{k,i} = \dfrac{1}{L} & \text{if } i < k, \\ h_{k,i} = 0 & \text{otherwise.} \end{cases}$$

The constraint $f \in \mathcal{C}$ must be expressed in an explicit way with *interpolation conditions* in order to have a tractable problem.

Definition 1. Given a set of triplet $\{(x_i, g_i, f_i)\}_{i \in I}$ with I some set of indices, *interpolation conditions* for the class of functions $\mathcal{C}$ are such that there exists a function $f \in \mathcal{C}$ with

$$\begin{aligned} f(x_i) &= f_i \quad \forall i \in I, \\ \nabla f(x_i) &= g_i \quad \forall i \in I, \end{aligned}$$

if and only if the *interpolation conditions* are satisfied.

When those conditions are available, the PEP can be rewritten as the following finite-dimensional problem

$$\begin{aligned} \max_{x_0,\dots,x_N,x_*,g_0,\dots,g_N,g_*,f_0,\dots,f_N,f_*} \quad & f_N - f_* \\ \text{s.t.} \quad & x_k = x_0 - \sum_{i=0}^{k-1} h_{k,i} \cancel{\nabla} g_i, \\ & ||x_0 - x_*||^2 \le R^2, \\ & \|g_*\|^2 = 0, \\ & \{(x_i, g_i, f_i)\}_{i \in I = \{0,1,\dots,N,*\}} \text{ are interpolable} \\ & \text{by some function } f \in \mathcal{C}. \end{aligned} \tag{PEP}$$

Finally, it was shown in Taylor et al. (2017) that this problem becomes a convex semidefinite problem provided that the iterates x_i and their gradients g_i are represented as elements of the Gram matrix $G = P^T P$, with

$$P = \begin{pmatrix} x_1 & \cdots & x_N & g_1 & \cdots & g_N \end{pmatrix} \in \mathbb{R}^{d \times 2N}.$$

3. PROBLEM STATEMENT

The key step and our main contribution is to obtain interpolation conditions for the class $\mathcal{Q}_{\mu,L}$ of quadratic functions. Indeed, we want to solve the following PEP on the class $\mathcal{Q}_{\mu,L}$,

$$\begin{aligned} \max_{x_0,\dots,x_N,f} \quad & f(x_N) - f^* \\ \text{s.t.} \quad & f \in \mathcal{Q}_{\mu,L}, \\ & x_k \text{ generated by applying } \mathcal{M} \text{ to } f \text{ from } x_0, \\ & ||x_0 - x^*|| \le R. \end{aligned} \tag{PEP-Q}$$

where we need an explicit equivalent reformulation of the condition $f \in \mathcal{Q}_{\mu,L}$ in order to solve (PEP-Q).

As mentioned above, (PEP) can be formulated under the form of a semidefinite program (see Taylor et al. (2017)) involving only the Gram matrix G of the iterates x_i and their gradients g_i and the values f_i of the function at these iterates.

In order to work in the class $\mathcal{Q}_{\mu,L}$, we must consider the set of Gram matrices associated to a quadratic function. Note that in that case we have

$$\nabla f(x) = Qx \quad \forall x \tag{1}$$

Definition 2. A symmetric matrix $G \in \mathbb{S}^{2N}$ is a (μ, L, N)-*quadratic-Gram matrix* if and only if there exist a dimension $d \in \mathbb{N}$, a symmetric matrix $Q \in \mathbb{S}^d$ with $\mu I \preceq Q \preceq LI$ and a sequence $x_i \in \mathbb{R}^d$ for $i = 1, \dots, N$ such that $G = P^T P$ with

$$P = \begin{pmatrix} x_1 & \cdots & x_N & \overbrace{Qx_1}^{g_1} & \cdots & \overbrace{Qx_N}^{g_N} \end{pmatrix} \in \mathbb{R}^{d \times 2N}.$$

The set of all (μ, L, N)-quadratic-Gram matrices is denoted $\mathcal{G}_{\mu,L,N}$. It can be shown that any conic combination of (μ, L, N)-quadratic-Gram matrices is also a (μ, L, N)-quadratic-Gram matrix, hence the set $\mathcal{G}_{\mu,L,N}$ is a convex cone.

In the following, we provide an explicit convex description of this set in order to be able to include those constraints to (PEP-Q). In other words, we show a convex formulation of the condition $f \in \mathcal{Q}_{\mu,L}$.

4. INTERPOLATION CONDITIONS

Several observations can be made about the form of the (μ, L, N)-quadratic-Gram matrices. Indeed, if a matrix G belongs to $\mathcal{G}_{\mu,L,N}$, then, by diagonalization of Q, it can be written under the form

$$\begin{aligned} G &= \begin{pmatrix} X^T X & X^T Q X \\ X^T Q X & X^T Q^2 X \end{pmatrix} \\ &= \begin{pmatrix} Y^T Y & Y^T D Y \\ Y^T D Y & Y^T D^2 Y \end{pmatrix} \\ &= \sum_{k=1}^{d} \begin{pmatrix} u_k u_k^T & \lambda_k u_k u_k^T \\ \lambda_k u_k u_k^T & \lambda_k^2 u_k u_k^T \end{pmatrix} \end{aligned} \tag{2}$$

where $X = (x_1 \cdots x_N) \in \mathbb{R}^{d \times N}$, $Q = VDV^T$ is the eigenvalue decomposition of Q, $Y = V^T X \in \mathbb{R}^{d \times N}$, $D = \mathrm{diag}(\lambda_1, \dots, \lambda_d)$, $\lambda_k \in [\mu, L]$ and $u_k = \begin{pmatrix} y_{1,k} \\ \vdots \\ y_{N,k} \end{pmatrix}$.

Vector u_k contains the k-th component of all vectors y_i. Expression (2) informs us that each block $X^T X$, $X^T Q X$, $X^T Q^2 X$ can be expressed as the sum of d positive definite rank-1 matrices $u_k u_k^T$, $\lambda_k u_k u_k^T$ and $\lambda_k^2 u_k u_k^T$.

By characterizing the Gram matrices exhibiting this structure, we are able to obtain the following explicit description of (μ, L, N)-quadratic-Gram matrices.

Theorem 2. Given a symmetric matrix

$$G = \begin{pmatrix} A & B \\ B^T & C \end{pmatrix} \in \mathbb{S}^{2N}$$

with $A, C \in \mathbb{S}^N$ and $B \in \mathbb{R}^{N \times N}$, the conditions

$$G \succeq 0 \tag{C1}$$

$$B = B^T \tag{C2}$$

$$B \succeq \frac{\mu L}{\mu + L} A + \frac{1}{\mu + L} C \tag{C3}$$

are necessary and sufficient for

$$G \in \mathcal{G}_{\mu,L,N}.$$

Observe that the quadratic interpolation conditions (C1), (C2) and (C3) of Theorem 2 do not involve the function values f_i. Actually, the variables f_i are directly encoded in the diagonal of the block matrix $B = X^T Q X$. Indeed, thanks to (1), we have

$$f(x) = \frac{1}{2} x^T Q x = \frac{1}{2} x^T \nabla f(x)$$

and the iterates x_i, g_i and f_i are linked through

$$f_i = \frac{1}{2} x_i^T g_i.$$

Since B contains the scalar products $x_i^T g_j$, whenever we need the value of f_i, we just use

$$f_i = \frac{1}{2} B_{ii}.$$

It is now possible to replace the condition $f \in \mathcal{Q}_{\mu,L}$ in (PEP-Q) by the interpolation conditions obtained in Theorem 2, which allows to reformulate the whole problem as a tractable optimization problem. This problem is a convex semidefinite program involving linear matrix constraint and can be comfortably written and solved with the *Python* library `PEPit` (see Goujaud et al. (2022)). Note that as we only consider homogeneous quadratic functions in the class $\mathcal{Q}_{\mu,L}$, we can assume implicitly that $x^* = 0$ and $f^* = f(x^*) = 0$, which simplifies the formulation.

Finally, as mentioned in the introduction, we actually obtained more general interpolation conditions that the ones of the class of quadratic functions. Indeed, Definition 2 and Theorem 2 provide interpolation conditions for any two sequences x_i and y_i linked by a matrix, i.e. $y_i = Qx_i \quad \forall i$. For example, if we apply a first-order method to the class of functions of the form $f(x) = g(Qx)$, then we will need to compute the gradient of f, i.e.

$$\nabla f(x) = Q \nabla g(Qx). \tag{3}$$

In order to describe this class with interpolation conditions, we need to force x_i and $y_i = Qx_i$ to be linked by a matrix as well as the $u_i = \nabla g(Qx_i)$ and $v_i = Q\nabla g(Qx_i)$. Thanks to Theorem 2, we are able to do it and, thus, to analyze the worst-case performance of this class through PEP.

5. RELATION BETWEEN INTERPOLATION CONDITIONS OF $\mathcal{F}_{\mu,L}$ AND $\mathcal{Q}_{\mu,L}$

The class of quadratic functions $\mathcal{Q}_{\mu,L}$ is included in $\mathcal{F}_{\mu,L}$, therefore, from the interpolation conditions of $\mathcal{Q}_{\mu,L}$, it must be possible to obtain the interpolation conditions of $\mathcal{F}_{\mu,L}$.

In Taylor et al. (2017), the following interpolation conditions for the class $\mathcal{F}_{\mu,L}$ have been obtained $\forall i,j = 0,1,\dots,N$

$$\begin{aligned} f_i - f_j - g_j^T(x_i - x_j) \geq \; & \frac{1}{2(1-\mu/L)} \\ & \bigg(\frac{1}{L}(g_i^T g_i + g_j^T g_j - 2 g_i^T g_j) \\ & + \mu (x_i^T x_i + x_j^T x_j - 2 x_i^T x_j) \\ & - 2\frac{\mu}{L}(g_j^T x_j - g_j^T x_i - g_i^T x_j + g_i^T x_i) \bigg). \end{aligned} \tag{4}$$

In the quadratic case, if we define a matrix $M = -\frac{\mu L}{\mu+L} A + B - \frac{1}{\mu+L} C$, the condition (C3) can be written as the positive semidefiniteness of matrix M which is equivalent to

$$\begin{aligned} M \succeq 0 \Leftrightarrow & \; z^T M z \geq 0 \quad \forall z \in \mathbb{R}^N \\ \Leftrightarrow & \sum_{k=1}^{N} \sum_{l=1}^{N} z_k z_l M_{kl} \geq 0 \quad \forall z \in \mathbb{R}^N. \end{aligned} \tag{5}$$

Choosing $z_i = 1$, $z_j = -1$ and all the other components of z equal to zero in (5) and then using $f_i = \frac{1}{2} x_i^T g_i$ yields the interpolation conditions (4) of the class $\mathcal{F}_{\mu,L}$.

Therefore, the finite set of interpolation conditions of $\mathcal{F}_{\mu,L}$ is explicitly seen as a consequence of the set of interpolation conditions of $\mathcal{Q}_{\mu,L}$.

6. ANALYSIS OF THE GRADIENT METHOD

In Taylor et al. (2017), the worst-case performance and the functions reaching it for the class $\mathcal{F}_{\mu,L}$ have been completely analyzed thanks to the PEP methodology. We would like to compare these results with the behavior of (GM) on the class $\mathcal{Q}_{\mu,L}$ and other classes involving matrices.

In the convex case $\mu = 0$, the worst-case performance on the class $\mathcal{F}_{0,L}$ is (from Taylor et al. (2017))

$$f(x_N) - f^* \leq \frac{LR^2}{4N+2}. \tag{6}$$

Note that this worst-case performance is reached by a Huber function, which is not quadratic and does not belong to $\mathcal{Q}_{0,L}$.

Thanks to our extension of PEP for the class $\mathcal{Q}_{\mu,L}$, we can solve (PEP) for the class $\mathcal{Q}_{0,L}$. It yields the following numerical results. Fig. 1 is the worst-case performance of (GM) on $\mathcal{F}_{0,L}$ (red) and $\mathcal{Q}_{0,L}$ (blue) for each number of iterations N.

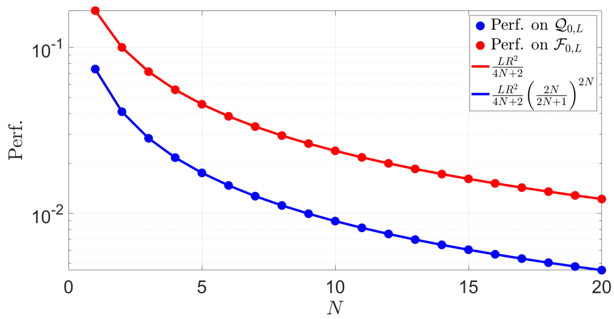

Fig. 1. Worst-case performance of (GM) on $\mathcal{Q}_{\mu,L}$ (blue dots) obtained by PEP and on $\mathcal{F}_{\mu,L}$ (red dots).

It turns out that it is possible to identify the worst-case rate of performance of (GM) on $\mathcal{Q}_{0,L}$, which is equal to the following analytical expression

$$f(x_N) - f^* \leq \frac{LR^2}{4N+2} \left(\frac{2N}{2N+1} \right)^{2N} \tag{7}$$

and this worst performance is achieved by the quadratic function

$$f(x) = \frac{Lx^2}{4N+2}.$$

We observe that the numerical results of PEP (blue dots) in Fig. 1 exactly matches the rate (7) (blue line).

Interestingly, the difference between the worst-case performance of (GM) on $\mathcal{F}_{0,L}$ and $\mathcal{Q}_{0,L}$ is the factor $\left(\frac{2N}{2N+1}\right)^{2N}$. Moreover, this factor exhibits two particular properties

$$\lim_{N\to\infty}\left(\frac{2N}{2N+1}\right)^{2N} = \frac{1}{e},$$
$$\left(\frac{2N}{2N+1}\right)^{2N} \geq \frac{1}{e} \quad \forall N \in \mathbb{N}.$$

Therefore, we can say that, for any number N of iterations, the worst-case performance of (GM) with constant step size $\frac{1}{L}$ on $\mathcal{F}_{0,L}$ is always lower than the performance on $\mathcal{Q}_{0,L}$ multiplied by a factor e.

To be complete, we must now mention that the literature already provides a methodology to analyze the worst-case performance of a first-order method on the class of quadratic functions with eigenvalues between μ and L (see for example Flanders and Shortley (1950); Nemirovsky and Polyak (1984); d'Aspremont et al. (2021)) and, thus, to obtain the rate (7). Indeed, given a quadratic function $\frac{1}{2}x^TQx$, an initial point x_0 and a method $\mathcal{M}$, the maximization of the last iterate x_N can be expressed as the maximization of a polynomial evaluated at the elements of the spectrum of Q, where the coefficients of the polynomial only depend on the method $\mathcal{M}$. Therefore this leads to the maximization of some explicit polynomial whose degree grows with the number of iterations. It can be shown that such a reasoning will provide the same rate (7).

However, as explained earlier, we are now also able to analyze the class of functions of the form $g(Ax)$, which cannot be tackled by the simple polynomial approach described in the previous paragraph. We observe a difference of worst-case performance of (GM) between the general class $\mathcal{F}_{0,L}$ and the class $\mathcal{C}_1$ of functions of the form $f(x) = g(Ax)$ where f is still L-smooth convex. Fig. 2 is the worst-case performance of (GM) on $\mathcal{F}_{0,L}$ (red) and $\mathcal{C}_1$ (blue) for each number of iterations N. Note again that such

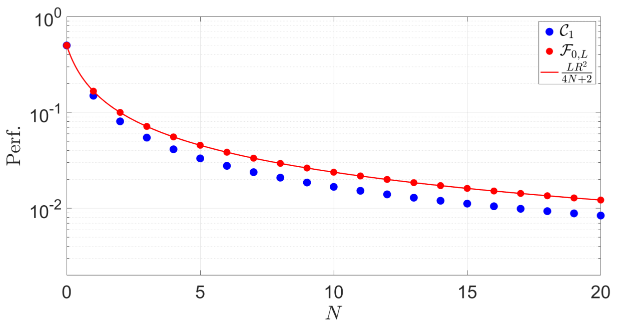

Fig. 2. Worst-case performance of (GM) on C_1 (blue dots) obtained by PEP.

results cannot be obtained by the abovementioned spectral analysis, and that it is possible with PEP to study even more complex functions classes such as $h(x) + g(Ax)$.

7. CONCLUSION

PEP has been shown to be a powerful tool for the analysis of the worst-case behavior of first-order optimization methods on a given class of functions. We showed how to extend PEP to the class of quadratic functions, thanks to Theorem 2, using a list of explicit convex constraints on the Gram matrix G. Moreover, we are able to implement and solve the PEP thanks to the *Python* library `PEPit`.

Our numerical experiments exactly match the analytical expression of the worst-case performance of the gradient method on convex quadratic functions $\mathcal{Q}_{0,L}$ and we compared it to the worst-case performance on smooth strongly convex functions $\mathcal{F}_{0,L}$. An interesting direction for future research would be to obtain a bound on the performance gap of any method between the general class $\mathcal{F}_{\mu,L}$ and the class $\mathcal{Q}_{\mu,L}$.

Moreover, in addition to the class of quadratic functions, we are now able to formulate explicit interpolation conditions for any class of functions involving matrices and to analyze them through PEP. This include for example the simple class of functions $f(x) = g(Ax)$ but also more complicated classes of functions as $f(x) = h(x) + g(Ax)$. Although the worst-case performance on the class of quadratic methods could already be obtained via the spectrum analysis approach, our extension of PEP appears to the best of our knowledge to be the first tool able to analyze classes of function of the forms $f(x) = g(Ax)$ or $f(x) = h(x) + g(Ax)$.

REFERENCES

A.B. Taylor, J.M. Hendrickx, F. Glineur *Smooth strongly convex interpolation and exact worst-case performance of first-order methods.* Mathematical Programming, 2017, vol. 161, no 1, p. 307-345.

B. Goujaud, C. Moucer, F. Glineur, J. Hendrickx, A. Taylor, A. Dieuleveut (2022). *PEPit: computer-assisted worst-case analyses of first-order optimization methods in Python.*

Y. Drori, M. Teboulle *Performance of first-order methods for smooth convex minimization: a novel approach.* Mathematical Programming, 2014, vol. 145, no 1, p. 451-482.

D. Scieur, F. Pedregosa *Universal average-case optimality of Polyak momentum.* International Conference on Machine Learning, 2020, (pp. 8565-8572). PMLR.

A. d'Aspremont, D. Scieur, A. Taylor *Acceleration methods.* Foundations and Trends® in Optimization, 2021, vol. 5, no 1-2, p. 1-245.

D. A. Flanders, G. Shortley *Numerical determination of fundamental modes.* Journal of Applied Physics. 21(12): 1326–1332.

A. S. Nemirovsky, B. T. Polyak *Iterative methods for solving linear ill-posed problems under precise information.* ENG. CYBER. (4): 50–56.