

Dominikus Noll · Pierre Apkarian

Spectral bundle methods for non-convex maximum eigenvalue functions: first-order methods

Dedicated to R. T. Rockafellar on the occasion of his 70th anniversary

Received: November 7, 2003 / Accepted: March 17, 2005

Published online: July 14, 2005 – © Springer-Verlag 2005

Abstract. Many challenging problems in automatic control may be cast as optimization programs subject to matrix inequality constraints. Here we investigate an approach which converts such problems into non-convex eigenvalue optimization programs and makes them amenable to non-smooth analysis techniques like bundle or cutting plane methods. We prove global convergence of a first-order bundle method for programs with non-convex maximum eigenvalue functions.

Key words. Bilinear matrix inequality (BMI) – Linear matrix inequality (LMI) – Eigenvalue optimization – First-order spectral bundle method – ϵ -subgradient

1. Introduction

The importance of linear matrix inequalities (LMIs) and bilinear matrix inequalities (BMIs) for applications in automatic control has been recognized during the past decade. Semidefinite programming (SDP) used to solve LMI problems has found a widespread interest due to its large spectrum of applications [9]. But many challenging engineering design problems lead to BMI feasibility or optimization programs no longer amenable to convex methods. Some prominent BMI problems are parametric robust feedback control [60, 61], static and reduced-order controller design [1, 14], design of structured controllers [13], decentralized synthesis, or synthesis with finite precision controllers [72]. These problems are in fact known to be NP-hard (see e.g. [45]).

Due to their significance for industrial applications, many solution strategies for BMI problems have been proposed. The most ambitious ones use ideas from global optimization such as branch-and-bound methods [20, 8, 22] or concave programming [6, 7] in order to address the presence of multiple local minima. On the other hand, in many situations, semidefinite programming relaxations, or heuristic approaches like coordinate descent schemes or alternating techniques (analysis versus synthesis) have been used with considerable success; see [10, 71, 28, 21, 31].

The general nature of BMI problems, which include for instance quadratic constraint quadratic (QCQP) programming, all polynomial problems and mixed binary programming, makes it evident that a general strategy may hardly be expected. We have observed

D. Noll: Université Paul Sabatier, Institut de Mathématiques, 118, route de Narbonne, 31062 Toulouse, France. e-mail: noll@mip.ups-tlse.fr

P. Apkarian: CERT-ONERA, Control System Department, 2, avenue Edouard Bélin, 31055 Toulouse, France

that BMIs in control applications may usually be solved by local methods. This is significant, since global methods have a prohibitive computational load and are therefore of very limited applicability. We have contributed several local nonlinear programming approaches, which are capable to deal with matrix inequality constraints [1, 2, 14, 15, 47, 3, 67–69]; see [30] for another approach. Here we investigate a strategy which converts BMI problems into non-convex eigenvalue optimization programs, which are then solved by non-smooth analysis tools.

1.1. BMIs and eigenvalue programs

We consider affine $\mathcal{A} : \mathbb{R}^n \rightarrow \mathbb{S}^m$ and bilinear operators $\mathcal{B} : \mathbb{R}^n \rightarrow \mathbb{S}^m$ into the space \mathbb{S}^m of symmetric $m \times m$ matrices,

$$\mathcal{A}(x) = A_0 + \sum_{i=1}^n A_i x_i, \quad \mathcal{B}(x) = \mathcal{A}(x) + \sum_{1 \leq i < j \leq n} B_{ij} x_i x_j. \quad (1)$$

Then a BMI-optimization program is of the form

$$\begin{aligned} & \text{minimize } c^\top x, \quad x \in \mathbb{R}^n \\ & \text{subject to } \mathcal{B}(x) \preceq 0 \end{aligned} \quad (2)$$

where $\preceq 0$ means negative semidefinite. Semidefinite programming is a special case where an affine operator \mathcal{A} is used. The simpler BMI-feasibility problem seeks $x \in \mathbb{R}^n$ satisfying $\mathcal{B}(x) \preceq 0$. This is a special case of (2) if solved as $\min\{t : \mathcal{B}(x) \preceq t I_m\}$.

These problems are clearly related to eigenvalue optimization. We consider the unconstrained

$$\text{minimize } \lambda_1(\mathcal{B}(x)), \quad x \in \mathbb{R}^n, \quad (3)$$

and the constrained eigenvalue optimization program

$$\begin{aligned} & \text{minimize } c^\top x, \quad x \in \mathbb{R}^n \\ & \text{subject to } \lambda_1(\mathcal{B}(x)) \leq 0. \end{aligned} \quad (4)$$

Here $\lambda_1 : \mathbb{S}^m \rightarrow \mathbb{R}$ is the maximum eigenvalue function, which is convex but non-smooth in general. The non-convexity of (3) and (4) is induced by the operator \mathcal{B} . Clearly (2) is equivalent to (4), while the BMI-feasibility problem may be solved using (3). Programs of the form (4) may be transformed into (3) via exact penalization, even though it may be preferable to use the structure of (4) explicitly.

Eigenvalue optimization has an interest of its own even outside control applications. Much pioneering work has been contributed by M. Overton in a series of papers [48–52] beginning in the 1980s, where Newton type methods for (3) are considered. Further to be mentioned among the earliest contributions are J. Cullum et al. [12], R. Fletcher [16], A. Shapiro [63, 64] and A. Shapiro and M.K.H. Fan [65].

Bundle methods have been invented by C. Lemaréchal [38] and Wolfe [73] and developed mostly in the 1980s by numerous contributions, in particular from Lemaréchal [39, 41] and K. Kiwiel (see e.g. [32, 35, 34]). A survey is [40], see also [29, 32, 62]. E. Polak

and co-workers have contributed several methods for eigenvalue optimization and for minimizing the H_∞ -norm, see [56–58] and the references given there. Bundle methods have recently been revived in the context of semidefinite programming; see [55, 42, 25, 43, 44, 37, 59]. Early contributions to non-convex programs are Kiwiel [35, 33, 36], a recent non-convex bundle algorithm is presented in Fuduli et al. [18, 19], a convex filter bundle method is given in [17].

1.2. Purpose

As a rule, bundle methods maintain a finite set of affine estimates of the objective function f at the current iterate. This set is used to predict a descent step. This trial step is either accepted (serious step) if the actual descent is sufficient, or it is used to improve the local model (null step) by adding another affine approximation of f . It has been observed in [12] (and used in [54, 55, 25, 42] and also [23, 26, 27, 24]) that maximum eigenvalue functions f allow for specific bundling strategies, where infinite sets of affine estimates of f are manageable. Here we develop a similar strategy for non-convex eigenvalue functions. Compared to the convex case, the main difficulty is that approximate subgradients of f do no longer provide global information. This complicates the analysis of the bundling procedure. In Section 3 we develop a first-order algorithm for (3), which extends the approach of [12] and [55] to the non-convex case. In Section 4 we solve (4) using an improvement function as proposed in [32].

2. Preparation and preliminary results

We recall some general notions and definitions and prepare the setting for the analysis of eigenvalue functions.

2.1. General definitions

Along with the operators \mathcal{A}, \mathcal{B} of (1) we consider more general matrix-valued operators $\mathcal{F} : \mathbb{R}^n \rightarrow \mathbb{S}^m$ of class \mathcal{C}^2 . For such \mathcal{F} let $\mathcal{F}'(x) : \mathbb{R}^n \rightarrow \mathbb{S}^m$ denote its derivative, whose adjoint $\mathcal{F}'(x)^*$ maps \mathbb{S}^m into \mathbb{R}^n . For affine \mathcal{A} we write $A := \mathcal{A}'$, which is the linear part of \mathcal{A} , defined as $Ax = \sum_{i=1}^n A_i x_i$. Its adjoint is $A^* : \mathbb{S}^m \rightarrow \mathbb{R}^n$, defined as $A^*Z = (A_1 \bullet Z, \dots, A_n \bullet Z)$. The scalar product in \mathbb{S}^m is $X \bullet Z = \text{tr}(XZ)$. The second derivative $\mathcal{F}''(x)$ of \mathcal{F} is a linear operator $\mathbb{R}^n \rightarrow L(\mathbb{R}^n, \mathbb{S}^m)$. For bilinear \mathcal{B} it is independent of x and defined via the tensor $[d, \mathcal{B}''d] = 2 \sum_{1 \leq i < j \leq n} B_{ij} d_i d_j \in \mathbb{S}^m$. With these preparations we are ready to consider the class of (generally non-convex) maximum eigenvalue functions f of the form $f = \lambda_1 \circ \mathcal{F}$. For an affine operator \mathcal{A} , $f = \lambda_1 \circ \mathcal{A}$ is convex and has been analyzed by many authors; see in particular [12, 53–55].

We use notions from convexity and non-smooth analysis as in [29] or [11]. Given $\epsilon \geq 0$ and a convex function ϕ on some Euclidean space E , following [29, ch. XI] the ϵ -subdifferential of ϕ at x is

$$\partial_\epsilon \phi(x) = \{y \in E^* : \langle y, h \rangle \leq \phi(x+h) - \phi(x) - \epsilon \text{ for all } h \in E\}.$$

For $\epsilon = 0$ this is the usual subdifferential in convex analysis, which we denote $\partial\phi(x)$ as usual. The ϵ -subdifferential gives rise to the ϵ -directional derivative, (cf. [29, ch. XI])

$$\phi'_\epsilon(x; d) = \max\{\langle y, d \rangle : y \in \partial_\epsilon \phi(x)\} = \inf_{t>0} \frac{\phi(x + td) - \phi(x) - \epsilon}{t},$$

where $\epsilon = 0$ reproduces the usual directional derivative $\phi'(x; d)$ of convex analysis. Notice that the ϵ -subdifferential of the maximum eigenvalue function λ_1 has the equivalent representation (cf. [29, Example XI.1.2.5]):

$$\partial_\epsilon \lambda_1(X) = \{G \in \mathbb{S}^m : G \geq 0, \text{tr}(G) = 1, G \bullet X \geq \lambda_1(X) - \epsilon\}. \tag{5}$$

This is a consequence of the fact that λ_1 is the support function of the set $\mathcal{C}_m = \{X \in \mathbb{S}^m : X \geq 0, \text{tr}(X) = 1\}$.

2.2. Concepts relating to $f = \lambda_1 \circ \mathcal{F}$

Let us introduce a notation which we will use systematically. Given $x, d \in \mathbb{R}^n$, let $X = \mathcal{F}(x) \in \mathbb{S}^m$ and $D = \mathcal{F}'(x)d \in \mathbb{S}^m$. If x_k, d_k arise in an algorithm, we use X_k and D_k accordingly.

Let us now extend ϵ -subgradients and ϵ -directional derivatives to functions $f = \lambda_1 \circ \mathcal{F}$. Let

$$\partial_\epsilon f(x) = \mathcal{F}'(x)^*[\partial_\epsilon \lambda_1(\mathcal{F}(x))] = \mathcal{F}'(x)^*[\partial_\epsilon \lambda_1(X)].$$

For $\epsilon = 0$ we obtain $\partial f(x) = \mathcal{F}'(x)^*[\partial \lambda_1(X)]$, which is the Clarke subdifferential of the composite function $f = \lambda_1 \circ \mathcal{F}$. Notice that $\partial f(x)$ so defined coincides with the convex subdifferential when \mathcal{F} happens to be an affine operator \mathcal{A} , so there is no ambiguity in the notations $\partial f(x)$ and $\partial_\epsilon f(x)$.

Given $\partial_\epsilon f(x)$, the ϵ -directional derivative is defined as

$$\begin{aligned} f'_\epsilon(x; d) &= \max\{g^\top d : g \in \partial_\epsilon f(x)\} \\ &= \max\{G \bullet D : G \in \partial_\epsilon \lambda_1(X)\} \\ &= (\lambda_1)'_\epsilon(X; D). \end{aligned}$$

For $\epsilon = 0$, we obtain the Clarke directional derivative, which we denote $f'(x; d)$ as usual. It is well-known that for $\epsilon > 0$ the ϵ -directional derivative is difficult to compute. Following Cullum et al. [12] and Oustry [55], we consider the so-called ϵ -enlargement $\delta_\epsilon \lambda_1(X)$ of the subdifferential of the maximum eigenvalue function, which is somewhere in between the purely local subdifferential $\partial \lambda_1(X)$ and the global $\partial_\epsilon \lambda_1(X)$ and turns out a good estimate of the latter.

Following [12] and in particular [55, Def. 1], fix $\epsilon > 0$ and let $r(\epsilon)$ be the largest index i such that $\lambda_i(X) > \lambda_1(X) - \epsilon$, called the ϵ -multiplicity of $\lambda_1(X)$. Notice that $r(\epsilon)$ is always at the end of a block of equal eigenvalues, that is, $\lambda_1(X) \geq \dots \geq \lambda_{t-1}(X) > \lambda_t(X) = \dots = \lambda_{r(\epsilon)}(X) > \lambda_{r(\epsilon)+1}(X) \geq \dots \geq \lambda_m(X)$, where of course $t = 1$ and $t = r(\epsilon)$ are admitted. We therefore define the spectral separation of ϵ as

$$\Delta_\epsilon(X) = \lambda_{r(\epsilon)}(X) - \lambda_{r(\epsilon)+1}(X) > 0. \tag{6}$$

Let Q_ϵ be a $m \times r(\epsilon)$ matrix whose columns form an orthonormal basis of the invariant subspace of X spanned by the eigenvectors of the first $r(\epsilon)$ eigenvalues of X . Now define

$$\delta_\epsilon \lambda_1(X) = \{G : G = Q_\epsilon Y Q_\epsilon^\top, Y \succeq 0, \text{tr}(Y) = 1, Y \in \mathbb{S}^{r(\epsilon)}\},$$

then $\delta_\epsilon \lambda_1(X) \subset \partial_\epsilon \lambda_1(X)$. We extend this concept to the class of functions $f = \lambda_1 \circ \mathcal{F}$ by setting

$$\delta_\epsilon f(x) = \mathcal{F}'(x)^*[\delta_\epsilon \lambda_1(X)],$$

so $\partial f(x) \subset \delta_\epsilon f(x) \subset \partial_\epsilon f(x)$, and the ϵ -enlarged subdifferential $\delta_\epsilon f(x)$ may be considered as an inner approximation of the ϵ -subdifferential $\partial_\epsilon f(x)$.

There is a natural analogue of the ϵ -directional derivative based on the new set $\delta_\epsilon f(x)$. Indeed, following [12, 55], define the ϵ -enlarged directional derivative of λ_1 by

$$(\tilde{\lambda}_1)'_\epsilon(X; D) = \max\{G \bullet D : D \in \delta_\epsilon \lambda_1(X)\} = \sigma_{\delta_\epsilon \lambda_1(X)}(D),$$

where σ_K denotes the support function of a convex set K . Extend this to functions $f = \lambda_1 \circ \mathcal{F}$ by setting

$$\begin{aligned} \tilde{f}'_\epsilon(x; d) &= \max\{g^\top d : g \in \delta_\epsilon f(x)\} \\ &= \max\{G \bullet D : G \in \delta_\epsilon \lambda_1(X)\} \\ &= (\tilde{\lambda}_1)'_\epsilon(X; D). \end{aligned}$$

The advantage of $\delta_\epsilon f(x)$ over the larger ϵ -subdifferential $\partial_\epsilon f(x)$ is that an explicit formula is available (cf. [12], [55, Prop. 3]):

$$\tilde{f}'_\epsilon(x; d) = \sigma_{\delta_\epsilon f(x)}(d) = \lambda_1 \left(Q_\epsilon^\top (\mathcal{F}'(x)d) Q_\epsilon \right) = \lambda_1 \left(Q_\epsilon^\top D Q_\epsilon \right).$$

Notice that $\epsilon = 0$ leads back to the standard Clarke directional derivative, $\tilde{f}'_0 = f'$.

Remark. The new ϵ -directional derivatives $f'_\epsilon(x; d)$ introduced for composite functions $f = \lambda_1 \circ \mathcal{F}$ coincides with the $f'_\epsilon(x; d)$ known from convex analysis [29] in those cases where \mathcal{F} is affine. Similarly, the new notion $\tilde{f}'_\epsilon(x; d)$ for nonconvex $f = \lambda_1 \circ \mathcal{F}$ coincides with the concept introduced in [55] for convex $f = \lambda_1 \circ \mathcal{A}$ if \mathcal{F} happens to be affine. So there is no ambiguity in these notations.

One of the main contributions from the work [55] is the following

Theorem 1. [55, Thm. 4] *Let $\epsilon \geq 0$, $\eta \geq 0$ and $X \in \mathbb{S}^m$ and define*

$$\rho(\epsilon, \eta) = \left(\frac{2\eta}{\Delta_\epsilon(X)} \right)^{1/2} + \frac{2\eta}{\Delta_\epsilon(X)}. \quad (7)$$

Then

$$\partial_\eta \lambda_1(X) \subset \delta_\epsilon \lambda_1(X) + \rho(\epsilon, \eta) B, \quad (8)$$

where B is the unit ball in \mathbb{S}^m .

An immediate consequence of this theorem is the estimate:

$$f'_\eta(x; d) \leq \tilde{f}'_\epsilon(x; d) + \rho(\epsilon, \eta) \|D\|. \quad (9)$$

2.3. Steepest descent

The direction of steepest descent plays an important role in the analysis of smooth functions. In the case of a non-smooth function it may be obtained by solving the program

$$\min_{\|d\| \leq 1} f'(x; d).$$

Fenchel duality shows that

$$\begin{aligned} \min_{\|d\| \leq 1} f'(x; d) &= \min_{\|d\| \leq 1} \max_{g \in \partial f(x)} g^\top d \\ &= \max_{g \in \partial f(x)} \min_{\|d\| \leq 1} g^\top d \\ &= \max_{g \in \partial f(x)} -g^\top \frac{g}{\|g\|}, \end{aligned}$$

so the direction d of steepest descent is obtained as the solution of a convex program:

$$d = -\frac{g}{\|g\|}, \quad g = \operatorname{argmin}\{\|g\| : g \in \partial f(x)\}.$$

As in the classical case, d will be a direction of descent if there is any, and in that case the relation

$$f'(x; d) = -\|g\| = -\operatorname{dist}(0, \partial f(x)) < 0$$

is satisfied. What is important is that the very same conclusions hold for the ϵ -subdifferential and the ϵ -enlarged subdifferential.

Definition 1. *If $0 \notin \partial_\epsilon f(x)$, the direction of steepest ϵ -descent d is*

$$d = -\frac{g}{\|g\|}, \quad g = \operatorname{argmin}\{\|g\| : g \in \partial_\epsilon f(x)\}$$

and satisfies $f'_\epsilon(x; d) = -\operatorname{dist}(0, \partial_\epsilon f(x)) < 0$. Similarly, if $0 \notin \delta_\epsilon f(x)$, then the direction of steepest ϵ -enlarged descent is

$$d = -\frac{g}{\|g\|}, \quad g = \operatorname{argmin}\{\|g\| : g \in \delta_\epsilon f(x)\} \tag{10}$$

and satisfies $\tilde{f}'_\epsilon(x; d) = -\operatorname{dist}(0, \delta_\epsilon f(x)) < 0$.

These notions are crucial for our present work. In particular, $\delta_\epsilon f$ will be used to compute qualified descent steps via (10). For a practical implementation it will be convenient to accept approximate solutions of program (10). We have the following

Lemma 1. *Let $0 < \omega \leq 1$ and $0 \notin \delta_\epsilon f(x)$, and consider a direction d which solves (10) approximately in the sense that*

$$d = -\frac{g}{\|g\|}, \quad \tilde{f}'_\epsilon(x; d) \leq -\omega \|g\|. \tag{11}$$

Then

$$-\operatorname{dist}(0, \delta_\epsilon f(x)) \leq \tilde{f}'_\epsilon(x; d) \leq -\omega \operatorname{dist}(0, \delta_\epsilon f(x)). \tag{12}$$

Proof. Let $\tilde{d} = -\tilde{g}/\|\tilde{g}\|$ be the solution of (10), then $\|\tilde{g}\| = \text{dist}(0, \delta_\epsilon f(x))$ and $\tilde{f}'_\epsilon(x; \tilde{d}) = -\|\tilde{g}\|$. As $g \in \delta_\epsilon f(x)$, we have $\|\tilde{g}\| \leq \|g\|$. Therefore (11) gives $\tilde{f}'_\epsilon(x; d) \leq -\omega\|g\| \leq -\omega\|\tilde{g}\| = -\omega \text{dist}(0, \delta_\epsilon f(x)) < 0$. The left hand estimate follows from $\tilde{f}'_\epsilon(x; d) = \max\{g^\top d : g \in \delta_\epsilon f(x)\} \geq \tilde{g}^\top d \geq -\|\tilde{g}\| = -\text{dist}(0, \delta_\epsilon f(x))$. \square

3. First-order analysis

In this chapter we derive and analyze a first-order bundle algorithm for minimizing non-convex maximum eigenvalue functions $f = \lambda_1 \circ \mathcal{F}$ for \mathcal{C}^2 -operators \mathcal{F} . Occasionally we will specify \mathcal{F} to a bilinear \mathcal{B} or even to an affine \mathcal{A} .

3.1. Optimality conditions

It is well-known that the ϵ -subdifferential may be used to obtain approximate optimality conditions, which lead to finite termination criteria in a convex minimization algorithm. See [29, 40, 41, 32] on how this is done. How about the meaning of an approximate optimality condition like $0 \in \partial_\epsilon f(x)$ for non-convex maximum eigenvalue functions f ? It is not surprising that without convexity, the consequences of $0 \in \partial_\epsilon f(x)$ are weaker:

Lemma 2. *Let $f = \lambda_1 \circ \mathcal{B}$ with \mathcal{B} bilinear. Let*

$$c := \max_{\|d\|=1} \max_{Z \in \mathcal{C}_m} \left| \sum_{i < j} Z \bullet B_{ij} d_i d_j \right|,$$

where $\mathcal{C}_m = \{Z \in \mathbb{S}^m : Z \succeq 0, \text{tr}(Z) = 1\}$. Let $\theta > 0, \epsilon \geq 0, \sigma \geq 0$ and define $r = r(\epsilon, \sigma, \theta, c)$ as

$$r = \frac{-\sigma + \sqrt{\sigma^2 + 4\theta c \epsilon}}{2c}.$$

Then every x such that $\text{dist}(0, \partial_\epsilon f(x)) \leq \sigma$ is $(1+\theta)\epsilon$ -optimal within the ball $B(x, r) := \{x' \in \mathbb{R}^n : \|x' - x\| \leq r\}$. That means, $f(x) \leq \min_{x' \in B(x, r)} f(x') + (1 + \theta)\epsilon$.

Proof. Let $X = \mathcal{B}(x)$. By assumption there exists $G \in \partial_\epsilon \lambda_1(X)$ such that $g = \mathcal{B}'(x)^* G$ satisfies $\|g\| \leq \sigma$. Now let $x' \in B(x, r)$ be written as $x' = x + td$ with $\|d\| = 1, t = \|x - x'\|$. Put $X' = \mathcal{B}(x')$. By definition of the ϵ -subdifferential, $G \bullet X' \leq \lambda_1(X')$ and $G \bullet X \geq \lambda_1(X) - \epsilon = f(x) - \epsilon$. Therefore

$$\begin{aligned} f(x) &\leq G \bullet X + \epsilon \\ &= G \bullet X' + G \bullet (X - X') + \epsilon \\ &\leq \lambda_1(X') + G \bullet (\mathcal{B}(x) - \mathcal{B}(x')) + \epsilon \\ &= f(x') + G \bullet \left(t\mathcal{B}'(x)d + t^2[d, \mathcal{B}''d] \right) + \epsilon \\ &\leq f(x') + \sigma\|x - x'\| + c\|x - x'\|^2 + \epsilon \\ &\leq f(x') + \sigma r + cr^2 + \epsilon = f(x') + (1 + \theta)\epsilon \end{aligned}$$

by the definition of r . This proves the claim. \square

The result includes several limiting cases. Clearly $\sigma = 0$ is interesting, where we get $r(\epsilon, 0, \theta, c) = (\theta\epsilon c^{-1})^{1/2}$. The case $c = 0$ is also possible. It corresponds to an affine operator \mathcal{A} , where the second derivative \mathcal{B}'' vanishes. We obtain $r(\epsilon, 0, \theta, 0) = +\infty$, meaning that the estimate is a global one. A third limiting case is when $\sigma > 0$ and $c = 0$. Here $r(\epsilon, \sigma, \theta, 0) = \theta\epsilon\sigma^{-1}$. Notice however that in the case $c < \infty$ the look-ahead character of this result is limited. Even while it is true that for $\sigma = 0$ the radius $r \sim \epsilon^{1/2}$ decreases slower than the discrepancy in the values, which behaves like $\sim \epsilon$, we have no way to know whether a local minimum is within the horizon r of our current iterate x .

Lemma 3. *Let $f = \lambda_1 \circ \mathcal{F}$. Suppose $\text{dist}(0, \partial_{\epsilon_k} f(x_k)) \rightarrow 0, x_k \rightarrow \bar{x}$ and $\epsilon_k \rightarrow 0$. Then $0 \in \partial f(\bar{x})$.*

Proof. The proof is straightforward. By the definitions we have $G_k \in \partial_{\epsilon_k} \lambda_1(X_k)$ for some G_k such that $\mathcal{F}'(x_k)^* G_k = g_k \rightarrow 0$. In particular, $G_k \bullet X \leq \lambda_1(X)$ for every $X \in \mathbb{S}^m$ and $G_k \bullet X_k \geq \lambda_1(X_k) - \epsilon_k$. Passing to the limit in every convergent subsequence of G_k , say $G_k \rightarrow \bar{G}$, gives $\bar{G} \bullet X \leq \lambda_1(X)$ for every X and $\bar{G} \bullet \bar{X} = \lambda_1(\bar{X})$, so $\bar{G} \in \partial \lambda_1(\bar{X})$. By the continuity of \mathcal{F}' we have $\mathcal{F}'(\bar{x})^* \bar{G} = 0$. That means $0 \in \partial f(\bar{x})$. \square

Taken together these two Lemmas justify a stopping test based on the smallness of $\text{dist}(0, \partial_\epsilon f(x))$ in tandem with the smallness of ϵ . Such a test is used e.g. in the non-convex bundle method [18, 19], where the Goldstein ϵ -subdifferential is used.

Let us fix x and ϵ and see what happens when $0 \notin \delta_\epsilon f(x)$. Choose $0 < \omega \leq 1$ and suppose d is an approximate direction of steepest ϵ -enlarged descent satisfying (11). The function $\eta \rightarrow \rho(\epsilon, \eta)$ in (7) is unbounded and monotonically increasing on $[0, \infty)$. Therefore if $\mathcal{F}'(x)d \neq 0$, there exists a unique $\eta = \eta(\epsilon)$ such that $\rho(\epsilon, \eta(\epsilon)) \|\mathcal{F}'(x)d\| = -\frac{1}{2} \tilde{f}'_\epsilon(x; d) > 0$. Using (9) gives the following consequence of Theorem 1:

$$f'_{\eta(\epsilon)}(x; d) \leq \frac{1}{2} \tilde{f}'_\epsilon(x; d) < 0. \tag{13}$$

The same estimate also holds in the case $\mathcal{F}'(x)d = 0$. Straightforward calculus with (7) now shows that in the case $\mathcal{F}'(x)d \neq 0$

$$\eta(\epsilon) = \frac{\Delta_\epsilon(X)}{8} \left(-1 + \sqrt{1 - \frac{2 \tilde{f}'_\epsilon(x; d)}{\|\mathcal{F}'(x)d\|}} \right)^2. \tag{14}$$

We summarize these observations in the following

Lemma 4. *Suppose $0 \notin \delta_\epsilon f(x)$ for some $\epsilon > 0$. Let $0 < \omega \leq 1$ and let d be a direction of approximate steepest ϵ -enlarged descent satisfying (11), (12). Then choosing $\eta = \eta(\epsilon)$ as in (14) gives $f'_\eta(x; d) \leq \frac{1}{2} \tilde{f}'_\epsilon(x; d) < 0$, i.e., d is a direction of η -descent. Moreover, $\eta(\epsilon)$ satisfies the estimate*

$$\eta(\epsilon) \geq \begin{cases} \frac{\Delta_\epsilon(X) \omega^2}{18 \|\mathcal{F}'(x)d\|^2} \text{dist}(0, \delta_\epsilon f(x))^2, & \text{for } 0 > \tilde{f}'_\epsilon(x; d) \geq -(3/2) \|\mathcal{F}'(x)d\| \\ \frac{\Delta_\epsilon(X) \omega}{16 \|\mathcal{F}'(x)d\|} \text{dist}(0, \delta_\epsilon f(x)), & \text{for } \tilde{f}'_\epsilon(x; d) \leq -(3/2) \|\mathcal{F}'(x)d\| \end{cases} \tag{15}$$

Proof. The proof starts with (14) and uses the elementary estimate

$$\left(-1 + \sqrt{1+t}\right)^2 \geq \begin{cases} \frac{1}{9}t^2, & 0 \leq t \leq 3, \\ \frac{1}{4}t, & t \geq 3 \end{cases},$$

which is applied to $t = -2\tilde{f}'_\epsilon(x; d)/\|\mathcal{F}'(x)d\| > 0$. Bringing in (12) yields (15). \square

3.2. Actual and predicted decrease

Our approach to (3) is to estimate the decrease of f in direction d with the help of the convex model $t \mapsto \lambda_1(X + tD)$. The following definition will be helpful.

Definition 2. Let $f = \lambda_1 \circ \mathcal{F}$, and fix $x, d \in \mathbb{R}^n$. Then $\alpha_t = f(x + td) - f(x)$ is called the actual decrease of f at x in direction d with step t , while $\pi_t = \lambda_1(X + tD) - f(x)$ is called the predicted decrease at x in direction d with step t .

Naturally, a true decrease of f in direction d with step t only occurs when $\alpha_t < 0$, and similarly for π_t . During the following, we will use the interplay between π_t and α_t in order to find suitable steps t which allow us to quantify α_t . We have the formula

$$\begin{aligned} \alpha_t &= f(x + td) - f(x) \\ &= f(x + td) - \lambda_1(X + tD) + \lambda_1(X + tD) - f(x) \\ &= f(x + td) - \lambda_1(X + tD) + \pi_t \end{aligned}$$

so for non-convex f we will have to estimate the mismatch $\alpha_t - \pi_t = f(x + td) - \lambda_1(X + tD)$.

Expanding the \mathcal{C}^2 operator \mathcal{F} in a neighborhood of x gives $\mathcal{F}(x + td) = X + tD + t^2H + t^2K_t$, where $K_t \rightarrow 0$ as $t \rightarrow 0$. [In the case of a bilinear \mathcal{B} , $K_t = 0$ and $H = [d, \mathcal{B}''d]$ is independent of x .] Then $f(x + td) = \lambda_1(X + tD + t^2(H + K_t))$ and by Weyl's theorem we obtain

$$t^2\lambda_m(H + K_t) \leq \lambda_1(X + tD + t^2(H + K_t)) - \lambda_1(X + tD) \leq t^2\lambda_1(H + K_t).$$

Altogether, $|f(x + td) - \lambda_1(X + tD)| \leq t^2\|H + K_t\|$. This motivates the following

Definition 3. Let $x, d \in \mathbb{R}^n$, $\|d\| = 1$ and $0 < t < +\infty$ and expand \mathcal{F} as above. Define $L_{x,d,t} := \sup\{\|H + K_\tau\| : 0 \leq \tau \leq t\} < +\infty$. When \mathcal{F} is of class \mathcal{C}^2_b , i.e., has bounded second derivative, then $t = +\infty$ is allowed and the $L_{x,d,\infty}$ are uniformly bounded on bounded sets of x . For bilinear \mathcal{B} , $L_{x,d,t} = L_d = \|H\| = \|[d, \mathcal{B}''d]\|$ is independent of x and $t > 0$ and therefore uniformly bounded.

We summarize our findings by the following

Lemma 5. Let $0 < T \leq \infty$ such that $L_{x,d,T} < \infty$. Then for every $t \leq T$ the actual decrease α_t satisfies

$$\alpha_t = \ell t^2 + \pi_t \quad \text{for some } |\ell| \leq L_{x,d,T}. \tag{16}$$

If \mathcal{F} is of class \mathcal{C}^2_b , we may choose $T = \infty$. In the bilinear case $L_{x,d,T} = L_d = \|[d, \mathcal{B}''d]\|$, and in the affine case, $L_{x,d,T} = L_d = 0$.

This result allows us to proceed as follows. We quantify the predicted decrease $\pi_t < 0$ of the convex model $t \mapsto \lambda_1(X + tD)$. Then we find suitable steps t for which the mismatch ℓt^2 is smaller in absolute value than the predicted decrease π_t , so that $\alpha_t = \ell t^2 + \pi_t < 0$ in a controlled fashion. This program is carried out in the next section.

3.3. Directional analysis

Let x be the current iterate in a tentative algorithm, and let d with $\|d\| = 1$ be a direction such that for some $\eta > 0$, $f'_\eta(x; d) < 0$, i.e., d is a direction of η -descent at x . This could happen with $\eta = \eta(\epsilon)$ and d an approximate direction of steepest ϵ -enlarged descent as in (11). Then $0 \notin \partial_\eta f(x)$. As usual let $X = \mathcal{F}(x)$ and $D = \mathcal{F}'(x)d$, then by definition, $f'_\eta(x; d) = (\lambda_1)'_\eta(X; D)$, so $(\lambda_1)'_\eta(X; D) < 0$. Following [29, XI.1], there exists a hyperplane supporting the epigraph of λ_1 , which passes through the point $(X, \lambda_1(X) - \eta)$ and touches the epigraph of λ_1 at a point $(X + t_\eta D, \lambda_1(X + t_\eta D))$, except when $t \mapsto \lambda_1(X + tD)$ is affine on $[0, \infty)$, or has an asymptote with slope $f'_\eta(x; d)$. In those cases let $t_\eta = \infty$ for consistency. If there are several steps with this property, then for definiteness let t_η be the smallest one. We shall say that the step t_η realizes the η -directional derivative. Notice that $\lambda'_1(X + t_\eta D; D) \leq (\lambda_1)'_\eta(X; D) \leq -\lambda'_1(X + t_\eta D; -D)$, so for almost all t_η we have equality $(\lambda_1)'_\epsilon(X; D) = \lambda'_1(X + t_\eta D; D)$. In general there exists at least a subgradient $G \in \partial \lambda_1(X + t_\eta D)$ such that $(\lambda_1)'_\eta(X; D) = G \bullet D$.

All this being a purely directional situation, we could also describe the case by introducing the function $\phi(t) = \lambda_1(X + tD)$. Then the line passing through $(0, \phi(0) - \eta)$ touches the epigraph of ϕ at $(t_\eta, \phi(t_\eta))$. The reader may want to inspect Figure 2.1.2 on page 105 of [29] for some illustration of these on-goings.

For the following assume that ϕ is not affine on $[0, \infty)$. The case $t_\eta = \infty$ with an asymptote is allowed. Suppose we backtrack and consider all lines passing through $(0, \phi(0) - \eta')$ for some $0 \leq \eta' \leq \eta$, touching the epigraph of ϕ at the corresponding points $(t_{\eta'}, \phi(t_{\eta'}))$, where again $t_{\eta'}$ is the smallest step if there are several. Then we introduce a function $\eta' \rightarrow t_{\eta'}$ with the following properties: $t_0 = 0$, and $t_\eta = t_{\eta'}$ at $\eta' = \eta$. It is monotonically increasing by convexity of ϕ . (Notice that its inverse, $t \rightarrow \eta(t)$, is not a function in the strict sense unless ϕ is differentiable. But we can introduce the notation to indicate any choice such that $\eta(t(\eta)) = \eta$.)

Recall that we assume $f'_\eta(x; d) < 0$. Then the decrease of ϕ on $[0, t_\eta]$ is

$$\phi(t_\eta) - \phi(0) = -\eta + t_\eta f'_\eta(x; d) < -\eta < 0. \tag{17}$$

Therefore the secant joining the points $(0, \phi(0))$ and $(t_\eta, \phi(t_\eta))$ has slope $-\sigma$, where

$$\sigma = \frac{\eta - t_\eta f'_\eta(x; d)}{t_\eta} = \frac{\eta}{t_\eta} - f'_\eta(x; d) > 0. \tag{18}$$

We can then get a pessimistic estimate of π_t by using the weaker decrease of the secant, which for a step $t \leq t_\eta$ decreases by $-\sigma t$. Therefore $\alpha_t = \ell t^2 + \pi_t \leq Lt^2 - \sigma t$ for every $t \leq t_\eta$, where $L := L_{x,d,t_\eta}$. (If $t_\eta = \infty$ and $L_{x,d,\infty} = \infty$, then this estimate will only be true for some finite T , $L := L_{x,d,T}$ and all $t \leq T$.) The minimum of the term $Lt^2 - \sigma t$ on $[0, t_\eta]$ is attained either at $t = \sigma/2L$ if $\sigma/2L \leq t_\eta$, or at $t = t_\eta$. Substituting these two possibilities into α_t gives the two cases (a) and (b) of the following

Lemma 6. *Let d be an approximate direction of steepest ϵ -enlarged descent satisfying (11) and let $\eta = \eta(\epsilon)$ be as in (14), so $f'_\eta(x; d) < 0$. Suppose $t_\eta < \infty$.*

(a) *If σ defined by (18) satisfies $\sigma \leq 2L_{x,d,t_\eta}$, then the step $t_\sigma := \sigma/2L_{x,d,t_\eta}$ gives a guaranteed decrease*

$$\begin{aligned} \alpha_{t_\sigma} &\leq -\frac{\sigma^2}{4L_{x,d,t_\eta}} = -\frac{1}{4L_{x,d,t_\eta}} \left(\frac{\eta}{t_\eta} - f'_\eta(x; d) \right)^2 \\ &\leq -\frac{1}{4L_{x,d,t_\eta}} \left(\frac{\eta}{t_\eta} - \frac{1}{2} \tilde{f}'_\epsilon(x; d) \right)^2 < 0. \end{aligned} \quad (19)$$

(b) *If on the other hand $\sigma > 2L_{x,d,t_\eta}$, then the step t_η guarantees the decrease*

$$\alpha_{t_\eta} \leq -\frac{\sigma t_\eta}{2} = \frac{1}{2} \pi_{t_\eta} = -\frac{1}{2} \left(\eta - t_\eta f'_\eta(x; d) \right) < 0. \quad (20)$$

The remaining case $t_\eta = \infty$ will be settled in Lemma 7 below. Before we do this, let us comment on Lemma 6, as it is crucial for the convergence analysis obtained in Section 3.7 and the central estimate in (22).

At first sight we would probably prefer case (b) over case (a) in Lemma 6, as it seems to give a better rate of decrease, $\mathcal{O}(\eta)$ versus $\mathcal{O}(\eta^2)$. Moreover, the constant L is likely to be large, so α_{t_σ} is expected to be small. In the same vein, t_η could be large, so the main contribution in α_{t_σ} probably comes from the term $f'_\eta(x; d)$ respectively $\tilde{f}'_{\epsilon_k}(x_k; d_k)$, which also occurs with order 2, as opposed to the second branch in (15). But (a) has a surprising advantage over (b). Namely, when later on in Section 3.7 our algorithm will take steps t_k which force the terms α_{t_k} to tend to zero, $t_k = t_{\sigma_k}$ in case (a) will imply $\tilde{f}'_{\epsilon_k}(x_k; d_k) \rightarrow 0$. On the other hand, steps $t_k = t_{\eta_k}$ in case (b) only lead to $\eta_k \rightarrow 0$, which implies $\Delta_{\epsilon_k}(X_k) \tilde{f}'_{\epsilon_k}(x_k; d_k) \rightarrow 0$ by estimate (15). Dealing with the extra term $\Delta_{\epsilon_k}(X_k)$ will cause some trouble.

The reason why we have to accept this extra term is that in case (b) of Lemma 6, we have no information as to the size of t_η , so the term $t_\eta f'_\eta(x; d) < 0$ may not be used to quantify the actual descent α_{t_η} . What remains in case (b) is $\alpha_{t_\eta} \leq -\eta/2$, which leads to a quantifiable descent via formula (15). Here, for small η , we have to take the pessimistic lower branch in that formula, which gives $\alpha_{t_\eta} \leq -\eta/2 \leq -K|\Delta_\epsilon(X) \tilde{f}'_\epsilon(x; d)|^2$ for some $K > 0$. This goes into the central estimate (22) in Section 3.7.

Likewise, if case (a) of Lemma 6 occurs, the term which allows to quantify descent, as step t_σ is taken, is $\alpha_{t_\sigma} \leq -K|\tilde{f}'_\epsilon(x; d)|^2 < 0$. Here it is the term $-\eta/t_\eta$, which does not seem to allow an a priori estimate of the actual descent. [Despite these theoretical facts, there are heuristic arguments, to be presented in Section 3.9, which indicate that the difference in the order of η in cases (a) and (b) of Lemma 6 may still be relevant numerically.] The corresponding estimate will be the left hand side of the maximum in (22).

Notice that as a rule we have to expect case (b) in Lemma 6. In particular, for convex $f = \lambda_1 \circ \mathcal{A}$, we have $L = 0$, so here case (b) is always on (compare the discussion in [55]).

Let us catch up with the case where $t_\eta = \infty$. This may happen if $\phi : t \mapsto \lambda_1(X + tD)$ is affine on $[0, \infty)$ or has an asymptote with slope $f'_\eta(x; d) < 0$. Then ϕ is below the

line with slope $-\sigma = f'_\eta(x; d)$ passing through the point $(0, \phi(0))$, and we may use this line to estimate π_t . The result is the same as in Lemma 6, with $t_\eta = \infty$ and $\eta/t_\eta = 0$.

Lemma 7. *Suppose $f'_\eta(x; d) < 0$ and $t_\eta = \infty$. Fix $T > 0$ such that $L := L_{x,d,T} < \infty$. Then the largest possible decrease of f in direction d amongst steps of length $t \leq T$ is obtained either (a) at $t_\sigma = \sigma/2L$ with $\alpha_{t_\sigma} \leq -\sigma^2/4L$ if $t_\sigma \leq T$, or (b) at $t = T$ with $\alpha_t \leq LT^2 - \sigma T \leq -\frac{1}{2}T\sigma$.*

For \mathcal{F} of class \mathcal{C}_b^2 we let $T = \infty$. In particular, this works for bilinear \mathcal{B} , so there is no restriction on the steplength t_σ . For general \mathcal{C}^2 operators the result seems to pose a little problem, as we need to know T to compute $L = L_{x,d,T}$, and L in order to check whether $t_\sigma \leq T$. Notice, however, that $L_{x,d,T}T$ increases as $T \rightarrow \infty$, while $t_\sigma > T$ for all T meant $L_{x,d,T}T$ remained bounded by $\sigma/2$. This is only possible when $L_{x,d,T} = 0$. So there is always the possibility to escape this dilemma. But section 3.8 will show an even simpler way to deal with general \mathcal{C}^2 operators.

3.4. Decrease of the order $\mathcal{O}(\eta)$

Let us look more systematically at those cases where a quantifiable decrease of the order $\mathcal{O}(\eta)$ is possible. This requires a sufficiently small trial step t_η . We have the following

Lemma 8. *Let $0 < \rho_0 < 1$. Then the trial step t_η gives a decrease $\alpha_{t_\eta} \leq -(1 - \rho_0)\eta$ provided $t_\eta \leq \vartheta$, where ϑ is the critical stepsize*

$$\vartheta := \frac{-f'_\eta(x; d) + \sqrt{f'_\eta(x; d)^2 + 4L\rho_0\eta}}{2L}. \tag{21}$$

Here $L = L_{x,d,t_\eta}$, and the limiting case $L_{x,d,t_\eta} = 0$ is allowed and gives $\vartheta = \infty$.

Proof. Since $|\ell| \leq L$, the inequality $\alpha_{t_\eta} = \ell t_\eta^2 - \eta + t_\eta f'_\eta(x; d) \leq -(1 - \rho_0)\eta$ is satisfied as soon as the stronger quadratic inequality $Lt_\eta^2 + t_\eta f'_\eta(x; d) - \rho_0\eta \leq 0$ holds. The corresponding quadratic equation has two real solutions, and the positive solution is ϑ in (21). Therefore the quadratic inequality holds as soon as $t_\eta \leq \vartheta$. \square

Remark. Notice that Lemma 8 is not required for the convergence proof in Section 3.7. In particular, it does not enter into estimate (22). However, checking whether the step t_η satisfies $t_\eta \leq \vartheta$ gives the algorithm a chance to take steps with a better rate of decrease than the pessimistic estimates in Lemmas 6, 7 and in formula (22). This is done in steps 6 and 8 of the algorithm. For instance, in the convex case $f = \lambda_1 \circ \mathcal{A}$, we have $\vartheta = \infty$, so here Lemma 8 guarantees that the descent steps have the rate $\mathcal{O}(\eta)$, a fact already established in [55].

3.5. The ϵ -management

The dependence of the estimate (14) on the gap $\Delta_\epsilon(X)$ suggests that we choose $\Delta_\epsilon(X)$ as large as possible. This strategy is indeed best when we aim at a finite termination theorem (see [12], [55]). If we want to prove convergence, the situation is, as we shall see, more subtle. Here we want $\Delta_\epsilon(X)$ large, but with the proviso that $\epsilon \rightarrow 0$.

Lemma 9. *Let $\bar{\epsilon} > 0$. Then there exists $\epsilon \leq \bar{\epsilon}$ such that $\Delta_\epsilon(X) \geq \bar{\epsilon}/m$.*

Proof. By definition of $r(\bar{\epsilon})$ we have $\lambda_1(X) \geq \dots \geq \lambda_{r(\bar{\epsilon})}(X) \geq \lambda_1(X) - \bar{\epsilon} > \lambda_{r(\bar{\epsilon})+1}(X)$. That means the $r(\bar{\epsilon})$ gaps $\lambda_i(X) - \lambda_{i+1}(X)$, $i = 1, \dots, r(\bar{\epsilon})$, add up to $\lambda_1(X) - \lambda_{r(\bar{\epsilon})+1}(X)$, which exceeds $\bar{\epsilon}$. So at least one of these gaps is larger than $\bar{\epsilon}/r(\bar{\epsilon})$, hence larger than $\bar{\epsilon}/m$. Suppose a gap exceeding $\bar{\epsilon}/m$ is $\lambda_i(X) - \lambda_{i+1}(X)$. Then we put $\epsilon = \lambda_1(X) - \lambda_i(X)$. \square

In practice we could either choose i smallest possible with $\lambda_i(X) - \lambda_{i+1}(X) \geq \bar{\epsilon}/m$, or we could pick i so that $\lambda_i(X) - \lambda_{i+1}(X)$ is largest possible. In the first case we get a small ϵ , which reduces the numerical burden to compute Q_ϵ . In the second case we render (14) the most efficient, possibly with a larger ϵ .

Definition 4. *The maximum eigenvalue function $f = \lambda_1 \circ \mathcal{F}$ is called linearly bounded below if for every $x \in \mathbb{R}^n$, the mapping $e \mapsto \lambda_1(\mathcal{F}(x) + \mathcal{F}'(x)e)$ is bounded below.*

The significance of this definition is in the following

Lemma 10. *Suppose $f = \lambda_1 \circ \mathcal{F}$ is linearly bounded below. Let $x \in \mathbb{R}^n$ be such that $0 \notin \partial f(x)$. Then there exists $\bar{\epsilon} > 0$ such that $0 \in \delta_{\bar{\epsilon}} f(x)$ but $0 \notin \delta_\epsilon f(x)$ for all $0 \leq \epsilon < \bar{\epsilon}$.*

Proof. Let $\bar{\epsilon} = \lambda_1(X) - \lambda_m(X)$, then $\delta_{\bar{\epsilon}} f(x) = \mathcal{F}'(x)^*(\mathcal{C}_m)$, where as before $\mathcal{C}_m = \{X \in \mathbb{S}^m : X \geq 0, \text{tr}(X) = 1\}$. By [55, Lemma 6] we have $0 \in \mathcal{F}'(x)^*(\mathcal{C}_m)$ because $e \mapsto \lambda_1(\mathcal{F}(x) + \mathcal{F}'(x)e)$ is bounded below. So we have proved $0 \in \delta_{\bar{\epsilon}} f(x)$. On the other hand, $0 \notin \delta_\epsilon f(x)$ if ϵ is sufficiently small. Take for instance ϵ so small that $r(\epsilon)$ equals the multiplicity of $\lambda_1(X)$. Then $\delta_\epsilon f(x) = \partial f(x)$, hence $0 \notin \delta_\epsilon f(x)$ by our hypothesis. By reducing $\bar{\epsilon}$ we can make it smallest possible with $0 \in \delta_{\bar{\epsilon}} f(x)$. \square

It is easy to force a maximum eigenvalue function to be linearly bounded below on a bounded set C . More generally, if f is bounded below by $\gamma \in \mathbb{R}$ on a set C , then define $\tilde{\mathcal{F}}(x) = \text{diag}(\gamma, \mathcal{F}(x)) \in \mathbb{S}^{m+1}$. Clearly $\tilde{f} = \lambda_1 \circ \tilde{\mathcal{F}}$ agrees with f on C and is linearly bounded below. During the following we will assume that f is linearly bounded below on the level set $\{x \in \mathbb{R}^n : f(x) \leq f(x_0)\}$.

3.6. Line search

We need one more element before our first-order bundle algorithm may be presented in detail. The analysis so far suggests several steps t which give a quantifiable decrease α_t . But we have to make sure that such a step is found by a finite procedure.

Suppose the parameter ϵ has been chosen, $\tilde{f}'_\epsilon(x; d) < 0$ has been computed and $\eta = \eta(\epsilon)$ has been found as in (14). Then we need to find t_η realizing $f'_\eta(x; d)$. If done in a precise way, this search is likely to be very costly. We therefore choose a relaxation.

Fix a tolerance parameter $0 < \theta_0 < 1$ and seek $t > 0$ such that $\lambda'_1(X + tD; D) < \frac{1}{4} \tilde{f}'_\epsilon(x; d) < 0$ and $\eta(t) > \theta_0 \eta$. The set of these t is nonempty and contains t_η as an interior point. This means a line search procedure like bisection can find $t > 0$ with these properties in a finite number of steps. So altogether replacing the original η by

$\theta_0\eta$ still gives the same order of decrease, but has the benefit to locate an approximation of the realizing abscissa in a finite procedure. This relaxation means that some of the estimates in previous Lemmas will get an extra factor θ_0 or θ_0^2 .

Naturally, while trying to locate t_η , we should not forget our original purpose of reducing f . After all, t_η relates only to $t \mapsto \lambda_1(X + tD)$, and not directly to f . We should therefore evaluate α_t at each intermediate step t visited during the search for t_η . Using the parameter $\rho_0 \in (0, 1)$ from Section 3.4, we may accept an intermediate t immediately if $\alpha_t \leq -(1 - \rho_0)\eta$, because this is the best order of decrease we can hope to achieve in general.

Suppose t_η respectively its substitute has been found. Keep among the intermediate steps the one with maximum decrease α_t and call it ζ . If necessary continue and compute the decrease α_{t_σ} at t_σ , where σ is as in (18) and compare α_{t_σ} to α_ζ . This covers all possible cases. For instance, if $t_\eta < \vartheta$ with ϑ as in (21), we will get a decrease of the order $\mathcal{O}(\eta)$.

Remark. Except in the bilinear case, we need to estimate $L_{x,d,t}$ for various stepsizes t . Strictly speaking this could not be done in a finite procedure. However, t and therefore $\|K_t\|$ are expected to be small, so $L_{x,d,t} \sim \|H\|$ for sufficiently small t , meaning that good estimates are available. In any case, the constants $L_{x,d,t}$ need not be known exactly. An upper bound L for all the $L_{x,d,t}$ with x ranging over the level set $\{x : f(x) \leq f(x_0)\}$ is all that is needed to prove convergence.

3.7. First-order algorithm

In this section we present our first-order bundle algorithm for the unconstrained minimization of $f = \lambda_1 \circ \mathcal{F}$ (see Figure 1) and prove convergence.

Let us consider sequences x_k, d_k and $\epsilon_k^\sharp, \bar{\epsilon}_k, \epsilon_k$ generated by the bundle algorithm. Suppose the sequence x_k is bounded, the $f(x_k)$ are bounded below, and that \mathcal{F} is of class \mathcal{C}_b^2 . Then $f(x_j) - f(x_0) = \sum_{k=0}^{j-1} f(x_{k+1}) - f(x_k) = \sum_{k=0}^{j-1} \alpha_{t_k}$ is bounded, and since each coefficient α_{t_k} is negative, the series $\sum_{k=0}^\infty \alpha_{t_k}$ converges. By boundedness of \mathcal{F}'' and the x_k , the constants $L_{x_k, d_k, t_{\eta_k}}$ are uniformly bounded. Then Lemmas 6 and 7 and the choice of t_k in step 8 give

$$\alpha_{t_k} \leq -K \max \left\{ \left(\eta_k / t_{\eta_k} - \tilde{f}'_{\epsilon_k}(x_k; d_k) \right)^2, \eta_k \right\} < 0 \tag{22}$$

for some $K > 0$ independent of k . This implies $\tilde{f}'_{\epsilon_k}(x_k; d_k) \rightarrow 0$ for the subsequence where the maximum is attained by the left hand expression, corresponding to case (a) in Lemmas 6 and 7. On the other hand, for the subsequence where the maximum in (22) is attained at the second term we obtain $\eta_k \rightarrow 0$, which by formula (15) implies $\Delta_{\epsilon_k}(X_k) \tilde{f}'_{\epsilon_k}(x_k; d_k) \rightarrow 0$. This corresponds to case (b) in Lemmas 6,7. Therefore, by (12), and since $\delta_{\epsilon_k} f(x_k) \subset \partial_{\epsilon_k} f(x_k)$, this implies $\text{dist}(0, \partial_{\epsilon_k} f(x_k)) \rightarrow 0$ in the first case, and $\Delta_{\epsilon_k}(X_k) \text{dist}(0, \partial_{\epsilon_k} f(x_k)) \rightarrow 0$ in the second.

Let $F = \bigcup_k F_k$ and $S = \bigcup_k S_k$ be the flat respectively steep iterates. There are two cases. Suppose first that there is an infinite number of flat iterates $k \in F$. Then $\epsilon_k^\sharp \rightarrow 0$ by step 9, because $\epsilon_k^\sharp = \gamma_{\ell_k}$, with ℓ_k the number of flat steps that occurred among the

Spectral bundle Algorithm for (3)

1. Choose an initial iterate x_0 and fix $0 < \theta_0, \rho_0 < 1$ and $0 < \omega \leq 1$.
Let $\gamma_k > 0$ be a sequence converging slowly to 0. Fix $\epsilon_0^\sharp = \rho_0$. Initialize $S_0 = \emptyset, F_0 = \emptyset$ and let $\text{slope} \in \{\text{steep}, \text{flat}\}$ be a binary variable.
2. Given the current iterate x_k , stop if $0 \in \partial f(x_k)$. Otherwise let $\bar{\epsilon}_k > 0$ such that $0 \in \delta_{\bar{\epsilon}_k} f(x_k)$ but such that $0 \notin \delta_\epsilon f(x_k)$ for $\epsilon < \bar{\epsilon}_k$. Choose $\epsilon_k \leq \min\{\bar{\epsilon}_k, \epsilon_k^\sharp\}$ such that $\Delta_{\epsilon_k}(X_k) \geq \min\{\bar{\epsilon}_k/m, \epsilon_k^\sharp/m\}$.
3. Compute a direction d_k of approximate steepest ϵ_k -enlarged descent

$$d_k = -\frac{g_k}{\|g_k\|}, \quad \tilde{f}'_{\epsilon_k}(x_k; d_k) \leq -\omega \|g_k\|.$$
4. If $|\tilde{f}'_{\epsilon_k}(x_k; d_k)| \leq \epsilon_k^\sharp$ put $\text{slope} = \text{flat}$, otherwise put $\text{slope} = \text{steep}$.
5. Compute $\eta_k = \eta(\epsilon_k)$ according to (14).
6. Search for t_{η_k} using a backtracking line search. If during the search t satisfying $\alpha_t \leq -(1 - \rho_0)\eta_k$ is found, put $t_k = t$ and goto 9. Otherwise stop as soon as t satisfying $\lambda'_1(X_k + tD_k; D_k) < \frac{1}{4}\tilde{f}'_{\epsilon_k}(x_k; d_k)$ and $\eta(t) \geq \theta_0\eta_k$ is found. Replace η_k by $\eta(t)$ and t_{η_k} by t . Let ζ_k be the step which gave the best α_{ζ_k} during the search.
7. Compute $L = L_{x_k, d_k, t_{\eta_k}}$, the trial step t_{σ_k} in (6) and $\alpha_{t_{\sigma_k}}$.
8. Compute ϑ_k by (21) and α_{ϑ_k} . Let $t_k \in \{\zeta_k, \vartheta_k, t_{\eta_k}\}$ the step which gives the best decrease.
9. Put $x_{k+1} = x_k + t_k d_k$. If $\text{slope} == \text{steep}$ let $F_{k+1} = F_k, \epsilon_{k+1}^\sharp = \epsilon_k^\sharp$.
 $S_{k+1} = S_k \cup \{k\}$. If $\text{slope} == \text{flat}$, let $F_{k+1} = F_k \cup \{k\}, S_{k+1} = S_k$
and put $\epsilon_{k+1}^\sharp = \gamma_\ell$, where $\ell = \text{card}(F_{k+1})$. Replace k by $k + 1$ and go back to step 2.

Fig. 1.

first $k + 1$ steps, so $\ell_k \rightarrow \infty$. Therefore also $\epsilon_k \rightarrow 0$. On the other hand, the flat steps $k \in F$ satisfy $|\tilde{f}'_{\epsilon_k}(x_k; d_k)| \leq \epsilon_k^\sharp \rightarrow 0$. Then $\text{dist}(0, \partial_{\epsilon_k} f(x_k)) \rightarrow 0$ because of (12). Therefore if \bar{x} is an accumulation point of the flat subsequence $x_k, k \in F$, we conclude with Lemma 3 that $0 \in \partial f(\bar{x})$.

Let us next assume that all but a finite number of steps are steep, i.e., $k \in S$ for all $k \geq k_0$. By step 4, $|\tilde{f}'_{\epsilon_k}(x_k; d_k)| > \epsilon_k^\sharp, k \geq k_0$, and by step 9 the algorithm stops driving ϵ_k^\sharp to 0. That means $\tilde{f}'_{\epsilon_k}(x_k; d_k)$ stays away from 0. Then we must have $\Delta_{\epsilon_k}(X_k) \rightarrow 0$. (In particular, this rules out case (a) in Lemma 6.) Now $\min\{\bar{\epsilon}_k, \epsilon_k^\sharp\} = \bar{\epsilon}_k$ eventually, because by step 2 this minimum tends to 0, while ϵ_k^\sharp stays away from 0. Hence $\bar{\epsilon}_k \rightarrow 0$, and since $0 \in \delta_{\bar{\epsilon}_k} f(x_k) \subset \delta_{\epsilon_k} f(x_k)$, Lemma 3 implies $0 \in \partial f(\bar{x})$ for every accumulation point \bar{x} of the entire sequence of iterates x_k . Altogether we have proved the following

Theorem 2. Let $f = \lambda_1 \circ \mathcal{F}$ be a (non-convex) maximum eigenvalue function with \mathcal{F} of class \mathcal{C}_b^2 . Let x_0 be fixed so that the level set $\{x \in \mathbb{R}^n : f(x) \leq f(x_0)\}$ is compact. Suppose the sequence x_k , starting with x_0 , is generated by the first-order bundle algorithm.

Then x_k is bounded and the $f(x_k)$ decrease monotonically. If all but finitely many iterates are steep, ($k \in S, k \geq k_0$), then every accumulation point of x_k is a critical point. If F is infinite, then every accumulation point of the flat subsequence $x_k, k \in F$, is a critical point of f .

In [55] the author considers the convex case $f = \lambda_1 \circ \mathcal{A}$ and aims at finite termination. The following is therefore a complement to [12] and [55, Thm. 7]:

Corollary 1. Suppose $f = \lambda_1 \circ \mathcal{A}$ is convex. Let x_0 be such that the level set $\{x \in \mathbb{R}^n : f(x) \leq f(x_0)\}$ is compact. Then every accumulation point of the sequence of iterates x_k generated by the first-order bundle algorithm starting with x_0 is a minimum of f .

Proof. By convexity every critical point \bar{x} of f is a (global) minimum. Therefore, if all but finitely many k are steep, the Theorem tells us that we are done. Suppose then that the subsequence $k \in F$ is infinite. Then by the above $x_k, k \in F$, has an accumulation point, \bar{x} , which is a minimum of f . Since the algorithm is of decent type, potential other accumulation points \tilde{x} of the steep subsequence $k \in S$ satisfy $f(\tilde{x}) = f(\bar{x})$, hence by convexity are also minima. That proves the claim. \square

Clearly this argument also applies when f is non-convex and the accumulation point \bar{x} above is a global minimum of f in $\{x \in \mathbb{R}^n : f(x) \leq f(x_0)\}$.

Remark. In the convex case the bundle algorithm essentially agrees with that of Cullum et al. [12] and Oustry [55], except that we force $\epsilon_k \rightarrow 0$ in certain cases in order to assure convergence. In [26] Helmberg and Rendl propose an alternative way to maintain an approximation $\hat{\mathcal{W}}$ of $\partial_\epsilon \lambda_1(X)$, using an orthogonal matrix P other than Q_ϵ . Moreover, they include the possibility to remember previous steps via an aggregate subgradient. Both approaches are compared in [25], and some merits of $\hat{\mathcal{W}}$ are observed. In [26] the authors prove convergence of their method and in 3.1 claim that convergence based on $\delta_\epsilon f(x)$ requires partial knowledge of the multiplicity \bar{r} of $\lambda_1(\bar{X})$ at the limit \bar{X} . More precisely, they claim that $r(\epsilon_k) \geq \bar{r}$ at iterates X_k near \bar{X} is needed. While this is true if the ϵ -management from [12, 55] is used, Corollary 1 shows that our way of choosing ϵ_k gives at least subsequence convergence without guessing \bar{r} correctly. (Naturally, if $f = \lambda_1 \circ \mathcal{A}$ has a strict minimum, our method gives convergence. Helmberg and Rendl's method converges without this hypothesis.)

Remark. The reader will have understood already that we intend the covering sequence γ_k to converge so slowly that the flat case almost never occurs. Nonetheless, it may seem a little puzzling that when F and S are both infinite, it is the flat subsequence which gets the merit of (subsequence) convergence, while the steep subsequence, seemingly doing all the work on the way, does not get rewarded by subsequence convergence. Indeed, the estimates suggest that steps where $|\tilde{f}'_{\epsilon_k}(x_k; d_k)|$ is large give the best decrease in the cost, and those are the ones in S .

Let us therefore examine the case where both F and S are infinite more closely. As we shall see, in most cases the subsequence S will still have many convergent subsequences.

Notice first that it is possible that $\tilde{f}'_{\epsilon_k}(x_k; d_k) \rightarrow 0$, for a subsequence $k \in S' \subset S$, even though the speed is necessarily slower than that of $\epsilon_k^{\sharp} \rightarrow 0$. Since $\epsilon_k \rightarrow 0$, a consequence of the fact that F is infinite, we conclude in that case (via Lemma 3) that every accumulation point \tilde{x} of S' is critical. Altogether, subsequences like S' are welcome.

Let us next examine a subsequence $S'' \subset S$ where $\tilde{f}'_{\epsilon_k}(x_k; d_k) \leq \tau < 0$, $k \in S''$. In that case we know that $\Delta_{\epsilon_k}(X_k) \rightarrow 0$, $k \in S''$. In particular, this may not happen if case (a) in Lemma 6 is on. Assuming that we are in case (b) of that Lemma, suppose we have $\bar{\epsilon}_k \rightarrow 0$, $k \in S''$. Then we are again done, ending up with another good subsequence S'' exhibiting subsequence convergence.

So finally the bad case is when $\bar{\epsilon}_k$, $k \in S''$, stay away from 0. Then by step 2 of the algorithm, we will eventually have $\epsilon_k^{\sharp} < \bar{\epsilon}_k$, $k \in S''$. Now we have to remember that $\sum_k \alpha_{t_k}$ is even summable, a fact we have never exploited so far. From estimate (14) we deduce that $\sum_{k \in S''} \Delta_{\epsilon_k}(X_k)^2 < \infty$, hence $\sum_{k \in S''} (\epsilon_k^{\sharp})^2 < \infty$ by the above. Put differently, a subsequence S'' of this last type must be extremely sparse, because as we agreed, γ_k tends to 0 very slowly. We may for instance decide that it converges so slowly that $\sum_k \gamma_k^2 = \infty$. Then also $\sum_k (\epsilon_k^{\sharp})^2 = \infty$.

This observation at least partially resolves the following dilemma caused by our algorithm. Suppose our method proposes iterates x_k with $k \in F$ and $k \in S$ fairly mixed. Then in order to be on the safe side, we would probably stop the process when $k \in F$. But what to do when all the iterates are in S ? In practice this will be satisfactory, as we are probably in the case where F is finite. But of course we can never be sure, having to stop after a finite number of steps. It will then be reassuring to know that the probability to be in a subsequence S'' of the last type, where subsequence convergence may fail, is in some sense very low. For instance, the probability to meet an element of S'' in an interval of fixed length ℓ , say in $I_n = [n, n + \ell]$, will tend to 0 (as $n \rightarrow \infty$), which makes it very likely that stopping the algorithm in I_n gives an iterate belonging to some of the "good" subsequences of S .

Remark. The above observations rise the question of how ϵ_k should be chosen in practice, and whether the failure of subsequence convergence may be avoided by a different ϵ -management. Our tests in [5] show that if the problem size is not too large, it may be worthwhile to compute *all* the eigenvalue gaps $\Delta_{\epsilon_k}(X_k)$ with $\epsilon_k \leq \bar{\epsilon}_k$ and their corresponding descent α_t , and retain the best one. In that reference we also show that with this extra information the convergence properties of the algorithm are somewhat better and pathological subsequences S'' as above may be avoided. Needless to say, for large size X , this option is impractical, and on picking $r(\epsilon_k) \ll m$, we have to be content to work with low dimensional approximations of the large matrix inequality. This situation is at the basis of the methods in [23, 24, 26] and also [55, 25].

3.8. Bounding t_η

For general \mathcal{C}^2 operators \mathcal{F} our algorithm has to be mildly modified. Here we encounter the problem that the constants L_{x,d,t_η} may become arbitrarily large due to the fact that the trial steps t_η may become arbitrarily large. In particular, we then cannot allow cases where $t_\eta = \infty$. As we have seen, this does not occur for bilinear \mathcal{B} , so we might regard

this case as of minor importance for the applications we have in mind. But eigenvalue programs with general \mathcal{C}^2 operators \mathcal{F} have frequently been treated in the literature; for applications in automatic control see for instance [4], [5]. We therefore include a discussion of this case here. The way it is handled is by brute force. We oblige the steps t_η to be uniformly bounded by a constant $T > 0$.

In order to understand the difficulty, consider the known case of an affine operator \mathcal{A} . If the minimization of $f = \lambda_1 \circ \mathcal{A}$ is to be well-defined, i.e., if f is to be bounded below, then the linear part A of \mathcal{A} needs to satisfy $\lambda_1(Ad) \geq 0$ for every d . That is, Ad is not negative definite for any d . Moreover, if the set of minimizers of $f = \lambda_1 \circ \mathcal{A}$ is to be bounded, we even require coercivity of f , which means $\lambda_1(Ad) > 0$ for every $d \neq 0$.

What happens for general non-convex \mathcal{C}^2 operators? Here we notice a difference to the affine case. Even when $f = \lambda_1 \circ \mathcal{F}$ is nicely bounded below and coercive, that is $f(x) \rightarrow +\infty$ as $\|x\| \rightarrow \infty$, there is no reason why its linearizations about a point x , that is, $e \mapsto \lambda_1(\mathcal{F}(x) + \mathcal{F}'(x)e)$, should share this property. In fact some of the linearizations may fail to be coercive, and this may lead to arbitrarily large values t_η as above. This effect motivates the following

Definition 5. A representation $f = \lambda_1 \circ \mathcal{F}$ of f is called linearly coercive if $e \mapsto \lambda_1(\mathcal{F}(x) + \mathcal{F}'(x)e)$ is coercive for every x .

The meaning of this definition becomes clear with the following

Lemma 11. Suppose $f = \lambda_1 \circ \mathcal{F}$ is linearly coercive. Then for every $R > 0$ there exists $T = T(R)$ such that for every $\|x\| \leq R$, every direction $\|d\| = 1$ and every $\eta > 0$ having $f'_\eta(x; d) < 0$, the abscissa t_η realizing $f'_\eta(x; d)$ satisfies $t_\eta \leq T$.

Proof. Suppose on the contrary that there exist $\|x_k\| \leq R$, $\|d_k\| = 1$ and $\eta_k > 0$ such that $-\lambda'_1(X_k + t_{\eta_k}D_k; -D_k) \leq f'_{\eta_k}(x_k; d_k) < 0$ is satisfied, where t_{η_k} realize $f'_{\eta_k}(x_k; d_k)$ and $t_{\eta_k} \rightarrow \infty$. Then each of the functions $\phi_k : t \mapsto \lambda_1(X_k + tD_k)$ decreases on the interval $[0, t_{\eta_k}]$. In particular, for every intermediate $0 < t < t_{\eta_k}$ we have $\lambda'_1(X_k + tD_k; D_k) < 0$ and $\lambda_1(X_k) > \lambda_1(X_k + tD_k)$.

Passing to subsequences, we may assume $x_k \rightarrow x$, $d_k \rightarrow d$, hence $X_k \rightarrow X$, $D_k \rightarrow D$. Now consider an arbitrary $t > 0$. Then $t_{\eta_k} > t$ for k large enough, hence $\lambda_1(X_k) > \lambda_1(X_k + tD_k)$ for k large enough. Then in the limit, $\lambda_1(X) \geq \lambda_1(X + tD)$. Since $t > 0$ was arbitrary, $t \mapsto \lambda_1(X + tD)$ is bounded above by $\lambda_1(X)$ on $[0, \infty)$. That contradicts linear coercivity. \square

The proof is not constructive, so it is not clear at the moment how T should be computed. However, as we shall see, T will not be required explicitly in the algorithm and rather serves as a theoretical parameter to obtain convergence. Let us now see how a given maximum eigenvalue function $f = \lambda_1 \circ \mathcal{F}$ could be forced to linear coercivity.

Lemma 12. Let $f = \lambda_1 \circ \mathcal{F}$ be given. For $R > 0$ and $\epsilon_0 > 0$, there exists a linearly coercive maximum eigenvalue function $\tilde{f} = \lambda_1 \circ \tilde{\mathcal{F}}$ such that

- (i) $f(x) = \tilde{f}(x)$ for every $\|x\| \leq R$.
- (ii) For every $0 \leq \epsilon \leq \epsilon_0$ and $\|x\| \leq R$, the ϵ -enlarged subdifferentials $\delta_\epsilon f(x)$ and $\delta_\epsilon \tilde{f}(x)$ coincide.
- (iii) If f is coercive, so is \tilde{f} .

Proof. Recall that $\mathcal{F} : \mathbb{R}^n \rightarrow \mathbb{S}^m$ by our standing notation. Now let $\mathcal{A} : \mathbb{R}^n \rightarrow \mathbb{S}^p$ be any affine matrix function such that $\lambda_1 \circ \mathcal{A}$ is coercive, and consider the following augmented functions

$$\tilde{\mathcal{F}}_\nu(x) = \begin{bmatrix} \mathcal{A}(x) - \nu I_p & 0 \\ 0 & \mathcal{F}(x) \end{bmatrix} \in \mathbb{S}^{m+p} \quad (23)$$

with $\nu \in \mathbb{N}$ a parameter to be chosen. Define open sets $\Omega_\nu := \{x \in \mathbb{R}^n : \lambda_1(\mathcal{A}(x)) - \nu < \lambda_1(\mathcal{F}(x))\}$, then the entire sequence satisfies $\cup_{\nu=1}^\infty \Omega_\nu = \mathbb{R}^n$. For ν large enough, let's say for $\nu \geq \nu_0$, the ball $\|x\| \leq R$ is contained in Ω_ν . By construction

$$\lambda_1\left(\tilde{\mathcal{F}}_\nu(x)\right) = \lambda_1(\mathcal{F}(x)) \text{ for every } x \in \Omega_\nu,$$

hence every $f_\nu = \lambda_1 \circ \tilde{\mathcal{F}}_\nu$ with $\nu \geq \nu_0$ is now running as a candidate to becoming the function \tilde{f} we are looking for.

Indeed, consider the linearization of $f_\nu = \lambda_1 \circ \tilde{\mathcal{F}}_\nu$ about some fixed x . This is

$$e \mapsto \lambda_1\left(\tilde{\mathcal{F}}_\nu(x) + \tilde{\mathcal{F}}'_\nu(x)e\right) = \max\{\lambda_1(\mathcal{F}(x) + \mathcal{F}'(x)e); \lambda_1(\mathcal{A}(x+e)) - \nu\},$$

because \mathcal{A} is its own linearization. As \mathcal{A} is coercive, so is $e \mapsto \mathcal{A}(x+e) - \nu I_p$, hence the linearization of each $\tilde{\mathcal{F}}_\nu$ is now coercive, i.e., f_ν is linearly coercive.

Let $\nu \geq \max\{\nu_0, \epsilon_0\}$, then for every $0 \leq \epsilon \leq \epsilon_0$ and $\|x\| \leq R$ we have

$$\lambda_1(\mathcal{A}(x)) - 2\nu \leq \lambda_1(\mathcal{A}(x)) - \nu - \epsilon < \lambda_1(\mathcal{F}(x)) - \epsilon < \lambda_{r(\epsilon)}(\mathcal{F}(x))$$

by the definition of $r(\epsilon)$ and since $x \in \Omega_\nu$. Therefore the first $r(\epsilon)$ eigenvalues of $\tilde{\mathcal{F}}_{2\nu}(x)$ and of $\mathcal{F}(x)$ are the same. The corresponding $r(\epsilon) \times (m+p)$ matrix \tilde{Q}_ϵ is simply $\tilde{Q}_\epsilon = [0_{p \times r(\epsilon)}; Q_\epsilon^\top]^\top$.

The conclusion is that for $0 \leq \epsilon \leq \epsilon_0$, the ϵ -enlarged subdifferential of $f_{2\nu} = \lambda_1 \circ \tilde{\mathcal{F}}_{2\nu}$ at $x \in \Omega_\nu$ is the same as that of $f = \lambda_1 \circ \mathcal{F}$. In consequence, also steepest ϵ -enlarged descent directions at $x \in \Omega_\nu$ are the same for both representations of f . \square

This is in contrast with the η -subdifferentials $\partial_\eta f$ and $\partial_\eta f_{2\nu}$, which are *global* concepts in the sense that they change even when the function is modified far away from the current position x . The conclusion is that changing $\partial_\eta f(x)$ into $\partial(f_{2\nu})_\eta(x)$ without affecting $\delta_\epsilon f(x)$ allows to bound the abscissae t_η . The situations is explained by the following

Example. Consider $\mathcal{B}(x) = x^2 \in \mathbb{S}^1$, then $f(x) = \lambda_1(\mathcal{B}(x)) = x^2$ is coercive, but the linearizations $e \rightarrow \lambda_1(\mathcal{B}(x) + \mathcal{B}'(x)e) = x^2 + 2ex$ are not. Now consider the following construction. Let $\tilde{\mathcal{B}}(x) = \text{diag}(-1+x, -1-x, x^2) \in \mathbb{S}^3$, then $\lambda_1(\tilde{\mathcal{B}}(x)) = x^2$, so f is represented as $\lambda_1 \circ \tilde{\mathcal{B}}$, now on \mathbb{S}^3 . In the terminology above we have chosen $\mathcal{A}(x) = \text{diag}(x, -x)$ and $\nu = 1$ with $\Omega_\nu = \mathbb{R}$. The linearization of $\tilde{\mathcal{B}}$ about a point x is now $\tilde{\mathcal{B}}(x) + \tilde{\mathcal{B}}'(x)e = \text{diag}(-1+x+e, -1-x-e, x^2+2ex)$, so $\lambda_1(\tilde{\mathcal{B}}(x) + \tilde{\mathcal{B}}'(x)e) = \max\{-1+|x+e|, x^2+2xe\} \rightarrow \infty$ as $|e| \rightarrow \infty$. In other words, the representation $f = \lambda_1 \circ \tilde{\mathcal{B}}$ is now linearly coercive. \square

The first-order bundle algorithm for (3) with general \mathcal{C}^2 operators is now obtained as follows. We follow the same steps, but use a linearly coercive representation \tilde{f} of f on the compact set $\{x : f(x) \leq f(x_0)\}$ in step 2 and during the search for t_η in step 6. The rest of the procedure remains unchanged, and the convergence or finite termination properties are the same.

3.9. Comments and extensions

From a practical point of view it may be attractive to modify the bundle method by allowing larger sets of subgradients $\delta_{\epsilon_k} f(x_k) \subset \mathcal{G}_k \subset \partial_{\epsilon_k} f(x_k)$, as proposed in [25]. For instance, some of the elements $g_{k-j} = \mathcal{F}'(x_{k-j})^* G_{k-j}$ from previous steps may be recycled at the step x_k . As opposed to the convex case [25], there are two ways how this could be arranged. We could either keep the old g_{k-j} , or we could keep only the old G_{k-j} and create new $g_{k-j}^\# = \mathcal{F}'(x_k)^* G_{k-j}$ at the actual point x_k . We would then accept as trial subgradients g any convex combination $g = \sum_j \alpha_j g_{k-j} + \alpha g'$, $\alpha_j, \alpha \geq 0$, $\sum_j \alpha_j + \alpha = 1$, with $g' \in \delta_{\epsilon_k} f(x_k)$, such that $g \in \partial_{\epsilon_k} f(x_k)$. At any rate, the sets \mathcal{G}_k so obtained are finite extensions of $\delta_{\epsilon_k} f(x_k)$. Notice that our convergence theory covers this case as well, even though the estimates based on $\delta_{\epsilon_k} f(x_k)$ get the more conservative, the larger the gap between $\delta_{\epsilon_k} f(x_k)$ and \mathcal{G}_k . Moreover, in each case we have to specify in which way the minimum norm element analogous to (11) is computed.

A question of practical importance is whether we can expect a decrease of the order $\mathcal{O}(\eta)$ as in the convex case (see [55]), or whether we will frequently have to be content with the order $\mathcal{O}(\eta^2)$. The following result gives some indication. Concerning terminology, recall that if the maximum eigenvalue $\lambda_1(X)$ has multiplicity r , we call $\text{sep}(X) = \lambda_r(X) - \lambda_{r+1}(X) > 0$ the separation of X .

Lemma 13. *Let $R > 0$ be fixed. There exist constants $K > 0$ and $\alpha > 0$ such that for every x with $\|x\| \leq R$, every d with $\|d\| = 1$ and every $\eta > 0$ with $f'_\eta(x; d) = \lambda'_1(X + t_\eta D; D)$, the following expansion is valid*

$$\frac{\eta}{t_\eta^2} = \frac{1}{2} \lambda''_1(X; D) + \kappa t_\eta$$

for some $|\kappa| \leq K$ and all $0 < t_\eta \leq \alpha \text{sep}(X)$.

Proof. According to Torki [70, Thm. 1.5], the maximum eigenvalue function is twice directionally differentiable and admits an expansion

$$\lambda_1(X + tD) = \lambda_1(X) + t\lambda'_1(X; D) + \frac{t^2}{2} \lambda''_1(X; D) + \mathcal{O}(t^3).$$

Inspecting [66, Thm. 2.8] on which the result is built shows that the $\mathcal{O}(t^3)$ term may more accurately be written as $\mathcal{O}(t^3) = \kappa_1 t^3$, where $|\kappa_1| \leq K_1$ for a constant K_1 which is uniform for a bounded set of X and D , and for $0 \leq t \leq \alpha \text{sep}(X)$, where α is independent of t .

A similar directional expansion holds at the second order level. We have

$$\lambda'_1(X + tD; D) - \lambda'_1(X; D) = t \lambda''_1(X; D) + \kappa_2 t^2$$

where $|\kappa_2| \leq K_2$ with K_2 uniform on a bounded set of X and D , but for t only in $0 \leq t \leq \alpha \operatorname{sep}(X)$.

Substituting these two estimates with $t = t_\eta$ into (17) gives the relationship

$$\eta = \frac{1}{2} t_\eta^2 \lambda_1''(X; D) + \kappa t_\eta^3$$

for $|\kappa| \leq K$ and some constant K which is the same for a bounded set of X and D , and for all $t_\eta \leq \alpha \operatorname{sep}(X)$. \square

Let us substitute this estimate into (21) and check whether $t_\eta < \vartheta$, i.e., whether a decrease of the order $\mathcal{O}(\eta)$ may be expected. Taking squares this is equivalent to

$$t_\eta^2 \leq \frac{|f'_\eta(x; d)| \left(|f'_\eta(x; d)| + \sqrt{f'_\eta(x; d)^2 + 4L\rho_0\eta} \right)}{2L^2} + \frac{\rho_0 \lambda_1''(X; D)}{2L} t_\eta^2 + K t_\eta^3$$

where the first term on the right hand side is positive, and the third term is negligible. Concerning the second term, Torki shows,

$$\lambda_1''(X; D) = \lambda_1(2U^\top Q^\top D(\lambda_1(X)I_r - X)^\dagger DQU),$$

where Q is a $r \times m$ matrix whose columns span the eigenspace of $\lambda_1(X)$, and U is a similar matrix for the eigenspace of $\lambda_1(Q^\top DQ)$. This means that $\lambda_1''(X; D)$ behaves roughly like $\operatorname{sep}(X)^{-1} = (\lambda_r(X) - \lambda_{r+1}(X))^{-1}$, the order of magnitude of the pseudo-inverse. This term is expected to explode as the iterates X_k approach a limit \bar{X} with $\lambda_1(\bar{X})$ of multiplicity greater than 1. At least this will happen when the $\lambda_1(X_k)$ have smaller multiplicity than $\lambda_1(\bar{X})$, as is usually the case. In other terms, $\rho_0 \lambda_1''(X; D)/2L$ is expected to be (way) larger than 1, so $t_\eta \leq \vartheta$ is expected to be satisfied asymptotically.

Remark. Unfortunately this argument remains heuristic, since Torki's expansion at each step k is only valid on a neighborhood $t \leq \alpha \operatorname{sep}(X_k)$ (for the same fixed $\alpha > 0$ which depends neither on k , nor on t). But it is to be expected that $\operatorname{sep}(X_k) \rightarrow 0$, so we cannot a priori render these estimates uniform over k . In fact, we could do so provided $\operatorname{sep}(X_k) \rightarrow 0$ converged slower than $\|X_k - \bar{X}\| \rightarrow 0$. This meant that the X_k approached \bar{X} transversally to the smooth manifold $\mathcal{M}_r = \{X \in \mathbb{S}^m : \lambda_1(X) = \dots = \lambda_r(X) > \lambda_{r+1}(X)\}$. Only iterates X_k approaching \mathcal{M}_r tangentially cause problems. Ironically, the situation is also o.k. for iterates which are *exactly on* the manifold \mathcal{M}_r . Then $\operatorname{sep}(X_k)$ stays away from 0 and Torki's local estimates again hold uniformly over k . This strange behavior highlights the non-smooth character of the maximum eigenvalue function. In the past, similar phenomena have motivated approaches, where in order to avoid the tangential zone, iterates have been *forced* to lie on the manifold \mathcal{M}_r . We will re-examine this idea in part 2 [46] of this paper.

4. Constrained program

Let us now address the constrained eigenvalue program (4). A first idea, often used in non-smooth optimization, is exact penalization. While this has been reported to induce irregular numerical behavior of bundle methods due to inconveniently large penalty constants (cf. [34]), we believe that the situation is less dramatic for (4) with its single scalar constraint. We have the following

Proposition 1. *Let \bar{x} be a local minimum of (4) such that \bar{x} is not a critical point of $f = \lambda_1 \circ \mathcal{F}$ alone. Then \bar{x} is a KKT-point of (4). If at least one of the associated Lagrange multipliers $\bar{\rho} \geq 0$ is known to satisfy $\bar{\rho} \leq \beta$, then \bar{x} is a critical point of the following unconstrained program of the form (3):*

$$\min\{c^\top x + \beta \lambda_1(\text{diag}(0_{1 \times 1}, \mathcal{F}(x))) : x \in \mathbb{R}^n\}.$$

Proof. Since $c^\top x$ and $f = \lambda_1 \circ \mathcal{F}$ are locally Lipschitz functions, the F. John necessary optimality conditions are satisfied at \bar{x} (see [11, Thm. 6.1.1]). That is, there exist $\bar{\sigma} \geq 0$, $\bar{\rho} \geq 0$, not both zero, such that

$$\bar{\sigma} c + \bar{\rho} \mathcal{F}'(\bar{x})^* \bar{G} = 0, \quad \bar{G} \in \partial \lambda_1(\mathcal{F}(\bar{x})), \quad \bar{\rho} \lambda_1(\mathcal{F}(\bar{x})) = 0, \quad \lambda_1(\mathcal{F}(\bar{x})) \leq 0.$$

Clearly $\bar{\rho} = 0$ is impossible, while $\bar{\sigma} = 0$ would imply that \bar{x} was a KKT point for $\lambda_1 \circ \mathcal{F}$ alone, which was excluded by hypothesis. Therefore \bar{x} is a KKT-point for (4). In other terms, we may assume $\bar{\sigma} = 1$, $\bar{\rho} > 0$ above.

Now we show that the set of $\bar{\rho}$ above is bounded. Indeed, suppose on the contrary that we have KKT-conditions with $\bar{\sigma} = 1$, $\rho_n \rightarrow \infty$ and $G_n \in \partial \lambda_1(\mathcal{F}(\bar{x}))$. By compactness of the subdifferential, we may assume $G_n \rightarrow G_\infty \in \partial \lambda_1(\mathcal{F}(\bar{x}))$ for a subsequence. Dividing by ρ_n and passing to the limit then implies $\mathcal{F}'(\bar{x})^* G_\infty = 0$, which means that \bar{x} is a KKT-point for $\lambda_1 \circ \mathcal{F}$ alone, contradicting our hypothesis. Hence the set of $\bar{\rho}$ is bounded.

What we have shown is that program (4) is calm at \bar{x} in the sense of Clarke [11, 6.4]. Hence it may be solved by exact penalization. That is, we find $\beta > 0$ such that

$$\min c^\top x + \beta \max\{0, \lambda_1(\mathcal{F}(\bar{x}))\}$$

is equivalent to (4). But observe that $\max\{0, \lambda_1(\mathcal{F}(\bar{x}))\} = \lambda_1(\text{diag}(0_{1 \times 1}, \mathcal{F}(x)))$, so the exact penalty program is of the form (3). The fact that every $\beta \geq \bar{\rho}$ will do is standard. \square

Let us now look at a second way to address (4), which builds on Kiwiel’s improvement function [32]. In the convex case $f = \lambda_1 \circ \mathcal{A}$, this has more recently been used by Miller et al. [43, 44], where ideas from [26] have been amalgamated with those of [32]. The emerging numerical method is reported to perform nicely.

Given the current iterate x_k in a minimization algorithm for (4), consider the improvement function

$$\phi(x_k; x) = \max \left\{ c^\top(x - x_k), \lambda_1(\mathcal{F}(x)) \right\} = \lambda_1 \left(\begin{bmatrix} c^\top(x - x_k) & 0 \\ 0 & \mathcal{F}(x) \end{bmatrix} \right). \quad (24)$$

The algorithm given in Figure 2 essentially follows the line of the unconstrained case, where at each instance k the new search direction is computed with respect to the improvement function $\phi(x_k; \cdot)$ instead of f .

Then we have the following

Theorem 3. *Consider program (4). Let x_0 be fixed and suppose \mathcal{F} is of class \mathcal{C}_b^2 . Suppose $c^\top x$ is bounded below on the feasible set $\{x \in \mathbb{R}^n : f(x) \leq 0\}$. Let the sequence x_k starting with x_0 be generated by the spectral bundle algorithm for (4). Then the following cases may occur:*

Spectral bundle algorithm for program (4)

1. Let $\omega, \rho_0, \theta_0, \epsilon_0^\sharp, x_0, \gamma_k, S, F \subset \mathbb{N}$ and $\text{slope} \in \{\text{steep}, \text{flat}\}$ be as in the unconstrained algorithm.
2. Given the current x_k , stop if $0 \in \partial\phi(x_k; x_k)$. Otherwise let $\bar{\epsilon}_k$ such that $0 \in \delta_{\bar{\epsilon}_k}\phi(x_k; x_k)$, but $0 \notin \delta_\epsilon\phi(x_k; x_k)$ for $\epsilon < \bar{\epsilon}_k$. Choose $\epsilon_k \leq \epsilon_k^\sharp$ such that $\Delta_{\epsilon_k}(X_k) \geq \min\{\bar{\epsilon}_k/m, \epsilon_k^\sharp/m\}$.
3. Compute a direction d_k of approximate steepest ϵ_k -enlarged descent at x_k with respect to the function $\phi(x_k; \cdot)$.
4. If $|\tilde{\phi}'(x_k; \cdot)_{\epsilon_k}(x_k; d_k)| \leq \epsilon_k^\sharp$ put $\text{slope} = \text{flat}$, otherwise put $\text{slope} = \text{steep}$.
5. Compute η_k as in (14) using $\phi(x_k; \cdot)$.
6. Search for t_{η_k} as in the unconstrained algorithm using $\phi(x_k; \cdot)$ and the corresponding $\mathcal{F}_k = \text{diag}\{c^\top(x - x_k), \mathcal{F}\}$ in lieu of f and \mathcal{F} .
- 7-9. In analogy with the unconstrained case.

Fig. 2.

1. All iterates are infeasible, i.e., $f(x_k) > 0$ for all k . If \bar{x} is an accumulation point of the entire sequence x_k (when F is finite) and of the flat subsequence $x_k, k \in F$ (when F is infinite), then \bar{x} is a critical point of f .
2. The iterates x_k are strictly feasible, i.e., $f(x_k) < 0$, for some k_0 and all $k \geq k_0$. If \bar{x} is an accumulation point of the entire sequence x_k (if F is finite) and of the flat subsequence $x_k, k \in F$ otherwise, then \bar{x} satisfies the F -John necessary optimality conditions for program (4).

Proof. 1) Let us first examine the situation where the initial point x_0 is infeasible, $f(x_0) > 0$. Then there are two possibilities. Either feasibility is reached in finite time, i.e., $f(x_k) \leq 0$ at some stage k . Or $f(x_k) > 0$ for all k , so that feasibility is never reached. In the second case $\phi(x_k, x) = f(x)$ around x_k , and the algorithm essentially behaves like the unconstrained bundle algorithm, that is, it reduces f . Since the $f(x_k)$ are bounded below by 0, the same conclusions are obtained. More precisely, every accumulation point \bar{x} of x_k is a critical point of f if the set F is finite, and so is every accumulation point of the flat subsequence if F is infinite. Such an accumulation point may or may not be feasible.

2) Let us now suppose that some iterate k is feasible, $f(x_k) \leq 0$. Then $\phi(x_k; x_k) = 0$. Unless the algorithm halts with $0 \in \partial\phi(x_k; x_k)$, the new d_k is a direction of descent of $\phi(x_k; \cdot)$ at x_k , so the line search will give a descent step with $\phi(x_k; x_{k+1}) < \phi(x_k; x_k) = 0$. Then $f(x_{k+1}) \leq \phi(x_k; x_{k+1}) < 0$. Therefore iterate x_{k+1} is even strictly feasible.

3) Suppose next that $f(x_k) < 0$ for some k . Then by repeating the argument in 2), all iterates $x_j, j \geq k$ stay strictly feasible. What is more, $c^\top(x_{k+1} - x_k) \leq \phi(x_k; x_{k+1}) < \phi(x_k; x_k) = 0$, so the algorithm now reduces the objective function at each step and continues to do so during the following steps.

Altogether iterates now decrease in value and remain strictly feasible. Since by hypothesis the problem is bounded below, the series $\sum_k \alpha_{t_k}$ converges. By construction, this yields then the same cases as discussed in the unconstrained algorithm.

4) Suppose for instance that some subsequence $k \in \mathcal{N}$ has $\bar{\epsilon}_k \rightarrow 0$, where $0 \in \delta_{\bar{\epsilon}_k} \phi(x_k; x_k)$. Let \bar{x} be one of its accumulation points, then the argument of Lemma 3 shows $0 \in \partial \phi(\bar{x}; \bar{x})$. In that case \bar{x} satisfies the F. John necessary optimality conditions, so it must be either a KKT-point, or a critical point of f alone.

Feasibility $f(\bar{x}) \leq 0$ is clear. But $f(\bar{x}) < 0$ is not possible, for in that case we would have $\partial \phi(\bar{x}; \bar{x}) = \{c\}$, a contradiction. So $f(\bar{x}) = 0$. Then both $c^\top(x - x_k)$ and $\mathcal{F}(x)$ are active at x_k , hence $0 = \alpha c + (1 - \alpha)g$ for some $0 \leq \alpha \leq 1$ and $g \in \partial f(\bar{x})$. If $\alpha > 0$ we have a KKT point. If $\alpha = 0$, then \bar{x} is a critical point of f alone, whose value is 0.

5) Suppose next that $\tilde{\phi}'_{\epsilon_k}(x_k; \cdot)(x_k; d_k) \rightarrow 0$ in tandem with $\epsilon_k \rightarrow 0$ for a subsequence $k \in \mathcal{N}$. Here the argument of Lemma 3 shows that every accumulation point \bar{x} of $x_k, k \in \mathcal{N}$, has $0 \in \partial \phi(\bar{x}; \bar{x})$, so the conclusion is the same.

6) Finally suppose $0 \in \partial \phi(x_k; x_k)$ at some k , in which case the algorithm halts. Here if $f(x_k) > 0$, we must have a critical point of f alone. If $f(x_k) \leq 0$, the discussion is the same as in 4) above. This completes the proof. \square

Remark. Clearly when $0 \in \partial f(\bar{x})$ with $f(\bar{x}) > 0$, the algorithm fails. While this always happens when the problem is infeasible, it is clear that even in the feasible case we may create situations, where a first-order method like the proposed one must fail. For instance, we could arrange that feasible points can only be reached from x_0 by *increasing* the improvement function, which the method never does. Nonetheless, Theorem 3 seems practically useful, as local minima or critical points of f alone are not expected to occur frequently in practice.

Notice that for a convex constraint, $f(x) = \lambda_1(\mathcal{A}(x)) \leq 0$, the algorithm never fails. We have the following

Corollary 2. *Suppose $f = \lambda_1 \circ \mathcal{A}$ is convex and program (4) is bounded below and has a strictly feasible point. Then every accumulation point of the sequence x_k generated by the constraint bundle algorithm is a minimum of (4).*

Proof. Suppose we had $f(x_k) > 0$ for all k . Then some accumulation point \bar{x} of the x_k is critical for f alone, which by convexity means \bar{x} is a minimum of f . Suppose $f(\bar{x}) > 0$, then the program has no feasible points, a contradiction. But $f(\bar{x}) = 0$ is also impossible, because no point is strictly feasible. This means $f(x_k)$ become strictly feasible after a finite number of iterations. Now Theorem 3 shows that every accumulation point \bar{x} of the x_k satisfies the F. John optimality condition. If \bar{x} is a KKT-point, then it is a minimum by convexity, so we are done. The other possibility is that \bar{x} is a minimum of f . But the proof of Theorem 3 shows that $f(\bar{x}) = 0$, while a minimum of f (if any) must have strictly negative value. So this is impossible. This completes the proof. \square

5. Conclusion

We have proved a global convergence result for constrained and unconstrained optimization problem using non-convex maximum eigenvalue functions, which assures subsequence convergence of the sequence of iterates towards critical points under mild

assumptions. The proposed method computes qualified descent steps using an inner approximation of ϵ -subdifferentials proposed in [12] and [55]. Recent numerical tests, to be published in [4, 5], seem to indicate that the method performs fairly well in practice. An extension to second order methods will be presented in part 2 [46] of this work.

References

1. Apkarian, P., Noll, D., Tuan, H.D.: A prototype primal-dual LMI interior point algorithm for non-convex robust control problems. Rapport interne 01-08 (2002), UPS-MIP, CNRS UMR 5640. <http://mip.ups-tlse.fr/publi/publi.html>
2. Apkarian, P., Noll, D., Tuan, H.D.: Fixed-order \mathcal{H}_∞ control design via an augmented Lagrangian method. *Int. J. Robust and Nonlin. Control* **13** (12), 1137–1148 (2003)
3. Apkarian, P., Noll, D., Thevenet, J.-B., Tuan, H.D.: A spectral quadratic-SDP method with applications to fixed-order H_2 and H_∞ synthesis. *European J. Control.* **10**, 527–538 (2004)
4. Apkarian, P., Noll, D.: Controller design via nonsmooth multi-directional search. *SIAM J. on Control and Optimization*, to appear (2005)
5. Apkarian, P., Noll, D.: Nonsmooth H_∞ -synthesis. Submitted
6. Apkarian, P., Tuan, H.D.: Concave programming in control theory. *J. Global Optim.* **15**, 343–370 (1999)
7. Apkarian, P., Tuan, H.D.: Robust control via concave minimization - local and global algorithms. *IEEE Trans. Autom. Control* **45**, 299–305 (2000)
8. Beran, E.B., Vandenberghe, L., Boyd, S.: A global BMI algorithm based on the generalized Benders decomposition. Proc. European Control Conf., Belgium, 1997
9. Boyd, S., ElGhaoui, L., Feron, E., Balakrishnan, V.: *Linear Matrix Inequalities in System and Control Theory*. vol. **15** of *Studies in Appl. Math.* SIAM, Philadelphia, 1994
10. Boyd, S., Vandenberghe, L.: Semidefinite programming relaxations of non-convex problems in control and combinatorial optimization
11. Clarke, F.H.: *Optimization and Nonsmooth Analysis*. John Wiley, New York, 1983
12. Cullum, J., Donath, W.E., Wolfe, P.: The minimization of certain nondifferential sums of eigenvalues of symmetric matrices. *Math. Progr. Stud.* **3**, 35–55 (1975)
13. ElGhaoui, L., Balakrishnan, V.: Synthesis of fixed-structure controllers via numerical optimization. Proc. IEEE Conf. Dec. Contr., December 1994
14. Fares, B., Noll, D., Apkarian, P.: An augmented Lagrangian method for a class of LMI-constrained problems in robust control theory. *Intern. J. Control* **74**, 348–360 (2001)
15. Fares, B., Noll, D., Apkarian, P.: Robust control via successive semidefinite programming. *SIAM J. Control Optim.* **40**, 1791–1820 (2002)
16. Fletcher, R.: Semidefinite constraints in optimization. *SIAM J. Control Optim.* **23**, 493–513 (1985)
17. Fletcher, R., Leyffer, S.: A bundle filter method for nonsmooth nonlinear optimization. University of Dundee, Report NA/195, 1999
18. Fuduli, A., Gaudioso, M., Giallombardo, G.: Minimizing nonconvex nonsmooth functions via cutting planes and proximity control. *SIAM J. Optim.* **14** (3), 743–756 (2004)
19. Fuduli, A., Gaudioso, M., Giallombardo, G.: A DC piecewise affine model and a bundling technique in nonconvex nonsmooth minimization. *Optimization Methods and Software* **19** (1), 89–102 (2004)
20. Fukuda, M., Kojima, M.: Branch-and-cut algorithms for the bilinear matrix inequality eigenvalue problem. Research Report on Mathematical and Computational Sciences. Series B: Operations Research, 1999
21. Grigoriadis, K.M., Skelton, R.E.: Low-order control design for LMI problems using alternating projection methods. *Automatica* **32**, 1117–1125 (1996)
22. Goh, K.-C., Safonov, M.G., Papavasilopoulos, G.P.: Global optimization for the biaffine matrix inequality problem. *J. Global Optim.* **7**, 365–380 (1995)
23. Helmberg, C.: A cutting plane algorithm for large scale semidefinite relaxations. ZIB-Report ZR-01-26. Padberg Festschrift. The sharpest cut. SIAM, 2001
24. Helmberg, C., Kiwiel, K.C.: A spectral bundle method with bounds. *Math. Programming* **93** (2), 173–194 (2002)
25. Helmberg, C., Oustry, F.: Bundle methods to minimize the maximum eigenvalue function. In: L. Vandenberghe, R. Saigal, H. Wolkowitz, (eds.), *Handbook of Semidefinite Programming. Theory, Algorithms and Applications*, vol. **27**, Int. Series in Oper. Res. and Management Sci. Kluwer Acad. Publ., 2000
26. Helmberg, C., Rendl, F.: A spectral bundle method for semidefinite programming. *SIAM J. Optim.* **10** (3), 673–696 (2000)

27. Helmsberg, C., Rendl, F.: Solving quadratic (0, 1)-problems by semidefinite programs and cutting planes. *Math. Programming* **82**, 291–315 (1998)
28. Helton, J.W., Merino, O.: Coordinate optimization for bi-convex matrix inequalities. *Proc. IEEE Conf. Decision and Control*, San Diego, 1997, pp. 3609–3613
29. Hiriart-Urruty, J.-B., Lemaréchal, C.: *Convex Analysis and Minimization Algorithms*, part II. Springer Verlag, Berlin, 1993
30. Hol, C.W.J., Scherer, C.W., van der Meché, E.G., Bosgra, O.H.: A nonlinear SDP approach to fixed-order controller synthesis and comparison with two other methods applied to an active suspension system. submitted, 2002
31. Iwasaki, T.: The dual iteration for fixed order control. *Proc. Am. Control Conf.*, 1997, pp. 62–66
32. Kiwiel, K.C.: *Methods of descent for nondifferentiable optimization*. Lect. Notes in Math. vol. **1133**, Springer Verlag, Berlin, 1985
33. Kiwiel, K.C.: A linearization algorithm for optimizing control systems subject to singular value inequalities. *IEEE Trans. Autom. Control* **AC-31**, 1986, pp. 595–602
34. Kiwiel, K.C.: A constraint linearization method for nondifferentiable convex minimization. *Numerische Mathematik* **51**, 395–414 (1987)
35. Kiwiel, K.C.: Proximity control in bundle methods for convex nondifferentiable optimization. *Math. Programming* **46**, 105–122 (1990)
36. Kiwiel, K.C.: Restricted-step and Levenberg-Marquardt techniques in proximal bundle methods for non-convex nondifferentiable optimization. *SIAM J. Optim.* **6**, 227–249 (1996)
37. Krishnan, K., Mitchell, J.E.: Cutting plane methods for semidefinite programming. Submitted, 2002
38. Lemaréchal, C.: Extensions diverses des méthodes de gradient et applications. Thèse d'Etat, Paris, 1980
39. Lemaréchal, C.: Bundle methods in nonsmooth optimization. In: *Nonsmooth Optimization*, Proc. IIASA Workshop 1977, C. Lemaréchal, R. Mifflin (eds.), Pergamon Press, 1978
40. Lemaréchal, C.: Nondifferentiable Optimization. chapter VII In: *Handbooks in Operations Research and Management Science*, vol. **1**, Optimization, G.L. Nemhauser, A.H.G. Rinnooy Kan, M.J. Todd (eds.), North Holland, 1989
41. Lemaréchal, C., Nemirovskii, A., Nesterov, Y.: New variants of bundle methods. *Math. Programming* **69**, 111–147 (1995)
42. Lemaréchal, C., Oustry, F.: Nonsmooth algorithms to solve semidefinite programs. In: L. El Ghaoui, L. Niculescu (eds.), *Advances in LMI methods in Control*, SIAM Advances in design and Control Series, 2000
43. Miller, S.A., Smith, R.S.: Solving large structured semidefinite programs using an inexact spectral bundle method. *Proc. IEEE Conf. Dec. Control*, 2000, pp. 5027–5032
44. Miller, S.A., Smith, R.S.: A bundle method for efficiently solving large structured linear matrix inequalities. *Proc. Am. Control Conf.*, Chicago, 2000
45. Nemirovskii, A.: Several NP-Hard problems arising in robust stability analysis. *Mathematics of Control, Signals, and Systems*. **6** (1), 99–105 (1994)
46. Noll, D., Apkarian, P.: Spectral bundle methods for nonconvex maximum eigenvalue functions: second-order methods. *Math. Programming, Series B* **104**, 729–747 (2005)
47. Noll, D., Torki, M., Apkarian, P.: Partially augmented Lagrangian method for matrix inequalities. *SIAM J. on Optimization* **15** (1), 161–184 (2004)
48. Overton, M.L.: On minimizing the maximum eigenvalue of a symmetric matrix. *SIAM J. Matrix Anal. Appl.* **9** (2), 256–268 (1988)
49. Overton, M.L.: Large-scale optimization of eigenvalues. *SIAM J. Optim.* **2** (1), 88–120 (1992)
50. Overton, M.L., Womersley, R.S.: Optimality conditions and duality theory for minimizing sums of the largest eigenvalues of symmetric matrices. *Math. Programming* **62**, 321–357 (1993)
51. Overton, M.L., Womersley, R.S.: Second derivatives for optimizing eigenvalues of symmetric matrices. *SIAM J. Matrix Anal. Appl.* **16** (3), 697–718 (1995)
52. Haerberly, J.-P.A., Overton, M.L.: A hybrid algorithm for optimizing eigenvalues of symmetric definite pencils. *SIAM J. Matrix Anal. Appl.* **15**, 1141–1156 (1994)
53. Oustry, F.: Vertical development of a convex function. *J. Convex Analysis* **5**, 153–170 (1998)
54. Oustry, F.: The \mathcal{U} -Lagrangian of the maximum eigenvalue function. *SIAM J. Optim.* **9** (2), 526–549 (1999)
55. Oustry, F.: A second-order bundle method to minimize the maximum eigenvalue function. *Math. Programming Series A* **89** (1), 1–33 (2000)
56. Polak, E.: On the mathematical foundations of nondifferentiable optimization in engineering design. *SIAM Review* **29**, 21–89 (1987)
57. Polak, E.: *Optimization: Algorithms and consistent approximations*. Springer Verlag, Applied Mathematical Sciences, 1997
58. Polak, E., Wardi, Y.: Nondifferentiable optimization algorithm for designing control systems having singular value inequalities. *Automatica* **18** (3), 267–283 (1982)

59. Rotunno, M., de Callafon, R.A.: A bundle method for solving the fixed order control problem. Proc. IEEE Conf. Dec. Control, Las Vegas, 2002, pp. 3156–3161
60. Scherer, C.: A full block S-procedure with applications. Proc. IEEE Conf. on decision and Control, San Diego, USA, 1997, pp. 2602–2607
61. Scherer, C.: Robust mixed control and linear parameter-varying control with full block scalings. SIAM Advances in Linear Matrix Inequality Methods in Control Series, L. El Ghaoui, S.I. Niculescu (eds.), 2000
62. Schramm, H., Zowe, J.: A version of the bundle method for minimizing a nonsmooth function: conceptual idea, convergence analysis, numerical results. SIAM J. Optim. **2**, 121–152 (1992)
63. Shapiro, A.: Extremal problems on the set of nonnegative matrices. Lin. Algebra Appl. **67**, 7–18 (1985)
64. Shapiro, A.: First and second order analysis of nonlinear semidefinite programs. Math. Programming Ser. B **77**, 301–320 (1997)
65. Shapiro, A., Fan, M.K.H.: On eigenvalue optimization. SIAM J. Optim. **5**(30), 552–569 (1995)
66. Steward, G.W., Ji-guang Sun: Matrix Perturbation Theory. Computer Science and Scientific Computing, Academic Press, 1990
67. Thevenet, J.-B., Apkarian, P., Noll, D.: Description and user's guide of specSDP, a Matlab-interfaced Fortran code to solve linear and nonlinear SDP constrained optimization programs. Technical Report, UPS-MIP, 2005
68. Thevenet, J.-B., Apkarian, P., Noll, D.: Reduced-order output feedback control design with specSDP, a code for linear and nonlinear SDP problems. ICCA05, Budapest, 2005
69. Thevenet, J.-B., Noll, D., Apkarian, P.: Nonlinear spectral SDP method for BMI-constrained problems: Applications to control design. Informatics in Systems and Control, Automation and Robotics, Kluwer Academic Publishers, 2004
70. Torki, M.: Second order directional derivative of all eigenvalues of a symmetric matrix. Nonlin. Analysis, Theory, Methods and Appl. **46**, 1133–1500 (2001)
71. Tuan, H.D., Apkarian, P., Nakashima, Y.: A new Lagrangian dual global optimization algorithm for solving bilinear matrix inequalities. Proc. Amer. Control Conf., 1999
72. Whidborne, J.F., Wu, J., Istepanian, R.H.: Finite word length stability issues in an ℓ_1 Framework. Int. J. Control **73** (2), 166–176 (2000)
73. Wolfe, Ph.: A method of conjugate subgradients for minimizing nondifferentiable convex functions. Math. Programming Studies **3**, 145–173 (1975)