Nonsmooth Minimization Using Smooth Envelope Functions

Pontus Giselsson and Mattias Fält

Abstract

Recently, the forward-backward and Douglas-Rachford envelope functions were proposed in the literature. The stationary points of these envelope functions have a close relationship with the solutions of the possibly nonsmooth optimization problem to be solved. The envelopes were shown to be smooth and convex under some additional assumptions. Therefore, these envelope functions create powerful bridges between nonsmooth and smooth optimization.

In this paper, we present a general envelope function that has these envelope functions as special cases. Under additional assumptions, we provide properties of the general envelope function that improve corresponding known results for the special cases. As a new special case, we present an envelope function for the generalized alternating projections method (GAP), named the GAP envelope. It enables for convex feasibility problems with two sets of which one is affine, to be solved by finding any stationary point of the smooth and sometimes convex GAP envelope. We finally note that primal-dual embedding of cone programs imply that most convex optimization problems can be solved this way.

1 Introduction

Many convex optimization problems can be solved by finding a fixed-point to a nonexpansive operator $S : \mathbb{R}^n \to \mathbb{R}^n$. That is, finding a point $x \in \mathrm{fix}S$, where

$$\mathrm{fix}S = \{x \in \mathbb{R}^n \mid Sx = x\}. \tag{1}$$

This is the basis for many first-order methods such as forward-backward splitting [9], Douglas-Rachford splitting [12, 27], the alternating direction method of multipliers (ADMM) [16, 22, 5] and its linearized versions [8], the three operator splitting method [10] and generalized alternating projections [25, 1, 29, 14, 7] that generalizes [39].

All these methods seek a fixed-point by an averaged iteration of the nonexpansive mapping S. So, they can be written on the following general form

$$x^{k+1} = (1-\alpha)x^k + \alpha S x^k. \tag{2}$$

This method is known to converge to a fixed-point of S, see [9]. The rate of convergence can, however, be very slow in practice.

One way to improve convergence of such first-order methods is to precondition the problem data. This approach has been extensively studied in the literature and has proven very successful in practice; see, e.g., [4, 6, 26, 17, 19, 20, 18] for a limited selection of such approaches. The underlying idea is to incorporate static second-order information in the respective algorithms.

In this paper, we build on the recently proposed forward-backward envelope in [34, 37] and Douglas-Rachford envelope in [33] that have pioneered the possibility for smooth optimization of nonsmooth problems in first-order methods. The assumption on the underlying composite convex optimization problem is that one of the functions is twice continuously differentiable with a Lipschitz continuous gradient. These envelope functions enable for second-order methods such as truncated Newton methods or quasi-Newton methods to be incorporated in the respective basic methods, see [34, 37]. This can lead to significantly improved local convergence.

In this paper, we show that a unifying property of forward-backward splitting and Douglas-Rachford splitting is that they are on the form (2), where $S = S_2S_1$, S_1 and S_2 are nonexpansive and gradients of some functions f_1 and f_2 respectively, and f_1 is twice continuously differentiable. We propose a general differentiable envelope function for such fixed-point iterations that has the forward-backward and Douglas-Rachford envelopes as special cases. Other special cases include the Moreau envelope and the ADMM envelope (which is a special case of the Douglas-Rachford envelope since ADMM is Douglas-Rachford splitting applied to the Fenchel dual problem, see [15]).

We analyze this general envelope function in the more restrictive setting of f_1 being quadratic, or equivalently $S_1 = \nabla f_1$ being affine, i.e., of the form $S_1 = P(\cdot)+q$, with P linear. We show that if P is nonsingular, the stationary points of the envelope coincide with the fixed-points of $S = S_2S_1$. We provide quadratic upper and lower bounds to the envelope function that improve corresponding results for the known special cases in the literature. The bounds imply, e.g., that the gradient of the envelope function is always 2-Lipschitz continuous. If in addition the linear operator P that defines S_1 is positive semidefinite, the envelope function is convex. Since the fixed-points of S and stationary points of the envelope coincide, a fixed-point to S can, when P is positive semidefinite, be found by minimizing a smooth and convex envelope function.

In [34, 37, 33] it was shown that forward-backward splitting and Douglas-Rachford splitting can be seen as variable metric gradient methods applied to the respective envelope functions. In the setting we consider (with S_1 being affine), they show that it instead is a scaled gradient method with fixed metric. We show that this holds also in this general setting. To do this, we interpret the averaged iteration in (2) as a scaled gradient method applied to the envelope function. Since the envelope function has nice smoothness properties and is in some cases convex, more efficient methods to find a fixed-point to S, or equivalently a stationary point of the envelope, probably exist. For instance, quasi-Newton, nonlinear conjugate gradient, or truncated Newton methods, some of which has been proposed to be used with the forward-backward envelope in [34, 37] can be used to improve local convergence (see [31] for details on the methods). Devising new algorithm or suggesting which existing ones that are most efficient is, however, outside the scope of this paper.

We also provide a new envelope function that is a special case of the general envelope, namely the generalized alternating projections (GAP) envelope. Generalized alternating projections [25, 1, 29, 14, 7] (which is also referred to as the method of alternating relaxed projections, e.g., in [3]) solves feasibility problems involving a finite number of nonempty closed and convex sets. This is done by alternating relaxed projections on the sets. It can use either under-relaxation, in which the step does not go all the way to the projection point, or over-relaxation when the step goes past the projection point, up towards the reflection point. Our envelope function applies to problems with two sets, with one nonempty closed and convex and one affine. Since the general envelope function always has a Lipschitz continuous gradient, so has the GAP envelope. If in addition, the first relaxed projection (onto the affine set) is an under-relaxation, the GAP envelope is convex. Therefore all feasibility problems with an affine subspace and a convex set can be solved by minimizing a smooth convex function.

This class of feasibility problems is not as restrictive as it first may sound. Actually, all convex convex cone programs can be posed as feasibility problems of this form using, e.g., the self-dual homogeneous embedding [40]. Many different types of convex optimization problems, such as LP, QP, SOCP, SDP, and lasso-type problems, can be cast as cone programs. (This is one reason why the CVX modeling languages [24, 11, 38] transform the stated problems to cone programs [23], before invoking a solver.) Therefore, the GAP envelope is another envelope function (besides the known forward-backward and Douglas-Rachford envelopes) and that enables for solving many nonsmooth convex optimization problems (at least those that can be solved by the CVX modeling languages) using smooth optimization techniques.

2 Preliminaries

2.1 Notation

We denote by $\mathbb{R}$ the set of real numbers, $\mathbb{R}^n$ the set of real column-vectors of length n, and $\mathbb{R}^{m\times n}$ the set of real matrices with m rows and n columns. Further $\overline{\mathbb{R}} := \mathbb{R} \cup \{\infty\}$ denotes the extended real line. We denote inner-products on $\mathbb{R}^n$ by $\langle\cdot,\cdot\rangle$ and their induced norms by $\|\cdot\|$. We will also use scaled norms $\|x\|_P := \langle Px, x\rangle$ where P is a positive definite operator (defined in Definition 2). We will use the same notation for scaled semi-norms, i.e., $\|x\|_P := \langle Px, x\rangle$ where P is a positive semidefinite operator (defined in Definition 1). The identity operator is denoted by I. The conjugate function is denoted and defined by $f^*(y) \triangleq \sup_x \{\langle y, x\rangle - f(x)\}$. The adjoint operator to a linear operator $L : \mathbb{R}^n \to \mathbb{R}^m$ is defined as the unique operator $L^* : \mathbb{R}^m \to \mathbb{R}^n$ that satisfies $\langle Lx, y\rangle = \langle x, L^*y\rangle$. The linear operator $L : \mathbb{R}^n \to \mathbb{R}^n$ is self-adjoint if $L = L^*$. The notation $\operatorname{argmin}_x f(x)$ refers to any element that minimizes f while the notation $\operatorname{Argmin}_x f(x)$ refers to the set of minimizers. Finally, ι_C denotes the indicator function for the set C that satisfies $\iota_C(x) = 0$ if $x \in C$ and $\iota_C(x) = \infty$ if $x \notin C$.

2.2 Background

In this section, we introduce some standard definitions that can be found, e.g. in [2, 35].

2.2.1 Operator properties

Definition 1 (Positive semidefiniteness) *A linear operator $L : \mathbb{R}^n \to \mathbb{R}^n$ is positive semidefinite if it is self-adjoint and all eigenvalues $\lambda_i(L) \geq 0$.*

Remark 1 *An equivalent characterization of a positive semidefinite operator is that $\langle Lx, x\rangle \geq 0$ for all $x \in \mathbb{R}^n$.*

Definition 2 (Positive definiteness) *A linear operator $L : \mathbb{R}^n \to \mathbb{R}^n$ is positive definite it is self-adjoint and if all eigenvalues $\lambda_i(L) \geq m$ with $m > 0$.*

Remark 2 *An equivalent characterization of a positive definite operator L is that $\langle Lx, x\rangle \geq m\|x\|^2$ for some $m > 0$ and all $x \in \mathbb{R}^n$.*

Definition 3 (Lipschitz mappings) *A mapping $T : \mathbb{R}^n \to \mathbb{R}^n$ is δ-*Lipschitz continuous *with $\delta \geq 0$ if*

$$\|Tx - Ty\| \leq \delta\|x - y\|$$

holds for all $x, y \in \mathbb{R}^n$. If $\delta = 1$ then T is nonexpansive *and if $\delta \in [0, 1)$ then T is δ-*contractive.

Definition 4 (Averaged mappings) *A mapping $T : \mathbb{R}^n \to \mathbb{R}^n$ is α-*averaged *if there exists a nonexpansive mapping $S : \mathbb{R}^n \to \mathbb{R}^n$ and $\alpha \in (0, 1]$ such that $T = (1 - \alpha)I + \alpha S$.*

Definition 5 (Negatively averaged mappings) *A mapping $T : \mathbb{R}^n \to \mathbb{R}^n$ is β-*negatively averaged *with $\beta \in (0, 1]$ if $-T$ is β-averaged.*

Remark 3 *For notational convenience later, we have included $\alpha = 1$ and $\beta = 1$ in the definitions of (negative) averagedness. But 1-averagedness and 1-negative averagedness is precisely nonexpansiveness. For values of $\alpha \in (0, 1)$ and $\beta \in (0, 1)$ averagedness is a stronger property than nonexpansiveness. For more on negatively averaged operators, see [18] where they were introduced.*

Note that if a gradient operator ∇f is α-averaged and β-negatively averaged. Then it must hold that $\alpha + \beta \geq 1$. This follows immediately from Lemma 4 and Lemma 5 in Appendix C.

Definition 6 (Cocoercivity) *A mapping $T : \mathbb{R}^n \to \mathbb{R}^n$ is δ-cocoercive with $\delta > 0$ if δT is $\frac{1}{2}$-averaged.*

Remark 4 *This cocoercivity definition implies that cocoercive mappings T can be expressed as*

$$T = \tfrac{1}{2\delta}(I + S) \tag{3}$$

for some nonexpansive operator S. We also note that 1-cocoercivity is equivalent to $\frac{1}{2}$-averagedness (which is also called firm nonexpansiveness).

We conclude this subsection with a result relating Lipschitz continuity and cocoercivity to averagedness and negative averagedness.

Proposition 1 *Suppose that $\nabla f : \mathbb{R}^n \to \mathbb{R}^n$ is the gradient of some function $f : \mathbb{R}^n \to \mathbb{R}$. Then the following hold:*

(i) ∇f is δ-Lipschitz continuous with $\delta \in [0,1]$ if and only if it is $\frac{\delta+1}{2}$-averaged and $\frac{\delta+1}{2}$-negatively averaged.

(ii) ∇f is $\frac{1}{\delta}$-cocoercive with $\delta \in [0,1]$ if and only if it is $\frac{1}{2}$-averaged and $\frac{\delta+1}{2}$-negatively averaged.

Proof. Claim *(i)*: Follows immediately from Lemma 3, Lemma 4, and Lemma 5. Claim *(ii)*: Lemma 4, and Lemma 5 imply that $\frac{1}{2}$-averagedness and $\frac{\delta+1}{2}$-negative averagedness is equivalent to that

$$0 \leq f(x) - f(y) - \langle \nabla f(y), x - y \rangle \leq \tfrac{\delta}{2}\|x - y\|^2$$

holds for all $x, y \in \mathbb{R}^n$. This is equivalent to that ∇f is $\frac{1}{\delta}$-cocoercive, see [30, Theorem 2.1.5] and [2, Definition 4.4]. □

2.2.2 Function properties

Definition 7 (Strong convexity) *Let $P : \mathbb{R}^n \to \mathbb{R}^n$ be positive definite. A proper and closed function $f : \mathbb{R}^n \to \overline{\mathbb{R}}$ is σ-*strongly convex *w.r.t. $\|\cdot\|_P$ with $\sigma > 0$ if $f - \frac{\sigma}{2}\|\cdot\|_P^2$ is convex.*

Remark 5 *If f is differentiable, σ-strong convexity w.r.t. $\|\cdot\|_P$ can equivalently be defined as that*

$$\tfrac{\sigma}{2}\|x - y\|_P^2 \leq f(x) - f(y) - \langle \nabla f(y), x - y \rangle \tag{4}$$

holds for all $x, y \in \mathbb{R}^n$. If $P = I$, i.e., if the norm is the induced norm, we merely say that f is σ-strongly convex. If $\sigma = 0$, the function is convex.

There are many smoothness definitions for functions in the literature. We will use the following that implies that the function is in every point majorized and minimized by a norm-squared function.

Definition 8 (Smoothness) *Let $P : \mathbb{R}^n \to \mathbb{R}^n$ be positive semidefinite. A function $f : \mathbb{R}^n \to \mathbb{R}$ is β-smooth w.r.t. $\|\cdot\|_P$ with $\beta \geq 0$, if it is differentiable and*

$$-\tfrac{\beta}{2}\|x - y\|_P^2 \leq f(x) - f(y) - \langle \nabla f(y), x - y \rangle \leq \tfrac{\beta}{2}\|x - y\|_P^2 \tag{5}$$

holds for all $x, y \in \mathbb{R}^n$.

2.2.3 Connections

We will later show that our envelope function satisfies upper and lower bounds of the form

$$\tfrac{1}{2}\langle M(x-y), x-y\rangle \leq f(x) - f(y) - \langle \nabla f(y), x-y\rangle \leq \tfrac{1}{2}\langle L(x-y), x-y\rangle \quad (6)$$

for all $x, y \in \mathbb{R}^n$ and for different linear operators $M : \mathbb{R}^n \to \mathbb{R}^n$ and $L : \mathbb{R}^n \to \mathbb{R}^n$. Depending on M and L, we get different properties of f and its gradient ∇f. Some of these are stated below. The results follow immediately from Lemma 3 in Appendix C and the definitions of smoothness and strong convexity in Definition 7 and Definition 8 respectively.

Proposition 2 *Assume that* $L = -M = \beta I$ *with* $\beta \geq 0$ *in* (6). *Then* (6) *is equivalent to that* ∇f *is* β*-Lipschitz continuous.*

Proposition 3 *Assume that* $M = \sigma I$ *and* $L = \beta I$ *with* $0 \leq \sigma \leq \beta$ *in* (6). *Then* (6) *is equivalent to that* ∇f *is* β*-Lipschitz continuous and* f *is* σ*-strongly convex.*

Proposition 4 *Assume that* $L = -M$ *and that* L *is positive definite. Then* (6) *is equivalent to that* f *is* 1*-smooth w.r.t.* $\|\cdot\|_L$.

Proposition 5 *Assume that* M *and* L *are positive definite. Then* (6) *is equivalent to that* f *is* 1*-smooth w.r.t.* $\|\cdot\|_L$ *and* 1*-strongly convex w.r.t.* $\|\cdot\|_M$.

3 Envelope functions

To find a fixed-point of a nonexpansive mapping S using an averaged iteration of that mapping, is the basis for many first-order optimization methods. Based on ideas from [34, 33], we present another method to find such a fixed-point. We create an envelope function whose stationary points coincide with the fixed-points of the operator S. For forward-backward splitting and Douglas-Rachford splitting, such envelopes have been proposed in [34] and [33] respectively. These envelope functions turn out to be special cases of the envelopes we propose, see Section 4. The envelope functions often possess favorable properties such as convexity and Lipschitz continuity of the gradient. Then, any method to find a stationary point (in the convex case, a minimizer) of the envelope function can be used to find a fixed-point to the nonexpansive mapping S.

To formulate our envelope function, we assume that the nonexpansive operator S is a composition of S_2 and S_1, i.e., $S = S_2S_1$. We make the following basic assumptions on S_1 and S_2, that sometimes will be sharpened or relaxed:

Assumption 1 *Suppose that:*

(i) $S_1 : \mathbb{R}^n \to \mathbb{R}^n$ *and* $S_2 : \mathbb{R}^n \to \mathbb{R}^n$ *are nonexpansive*

(ii) $S_1 = \nabla f_1$ *and* $S_2 = \nabla f_2$ *for some differentiable functions* $f_1 : \mathbb{R}^n \to \mathbb{R}$ *and* $f_2 : \mathbb{R}^n \to \mathbb{R}$

(iii) $S_1 : \mathbb{R}^n \to \mathbb{R}^n$ *is affine, i.e.,* $S_1 x = Px + q$ *and* $f_1(x) = \frac{1}{2}\langle Px, x\rangle + \langle q, x\rangle$*, where* $P \in \mathbb{R}^{n\times n}$ *is a a self-adjoint nonexpansive linear operator and* $q \in \mathbb{R}^n$

Remark 6 *Part (iii) of the assumption means that* P *is symmetric with eigenvalues in the interval* $[-1, 1]$.

Now, we are ready to define the general envelope function whose properties we will investigate in this paper:

$$F(x) := \tfrac{1}{2}\langle Px, x\rangle - f_2(\nabla f_1(x)). \tag{7}$$

The gradient of this function is given by

$$\nabla F(x) = Px - \nabla^2 f_1(x)\nabla f_2(\nabla f_1(x)) = Px - PS_2(S_1 x) = P(x - S_2 S_1 x). \tag{8}$$

The set of stationary points to the envelope function F is the set of points for which the gradient is zero. This set is denoted as follows:

$$X^\star := \{x \mid \nabla F(x) = 0\}. \tag{9}$$

3.1 Basic properties of the envelope function

Here, we list some basic properties of the envelope function (7). The first two results are special cases and direct corollaries of a more general result in Theorem 1, and therefore not proven here.

Proposition 6 *Suppose that Assumption 1 holds. Then the gradient of* F *is 2-Lipschitz continuous. That is,* ∇F *satisfies*

$$\|\nabla F(x) - \nabla F(y)\| \le 2\|x - y\|$$

for all $x, y \in \mathbb{R}^n$.

Proposition 7 *Suppose that Assumption 1 holds and that* P*, the operator defining the linear part of* S_1*, is positive semidefinite. Then* F *is convex.*

So, if P is positive semidefinite, then the envelope function F is convex and differentiable with a Lipschitz continuous gradient. The set of stationary points of F also has a close relationship with the fixed-points of $S = S_2 S_1$. This is shown next.

Proposition 8 *Suppose that Assumption 1 holds and that* P *is nonsingular. Then* $X^\star = \mathrm{fix}(S_2 S_1)$ *where* $X^\star$ *is defined in* (9) *and the fixed-point set is defined in* (1)*. If in addition* P *is positive definite, then* $\mathrm{Argmin}_x F(x) = X^\star = \mathrm{fix}(S_2 S_1)$.

Proof. The first claim follows directly from (8). The second claim follows from (8) and that F is convex when P is positive (semi)definite, see Proposition 7. □

These three results show that if P is positive definite, a fixed-point to $S_2 S_1$ can be found by minimizing the differentiable convex function F, which has a 2-Lipschitz continuous gradient.

3.2 Finer properties of the envelope function

Here, we establish some finer properties of the envelope function. We start with a general result on upper and lower bounds for the envelope function. This result uses stronger assumptions on S_2 than nonexpansiveness, namely that it is α-averaged and β-negatively averaged with $\alpha, \beta \in (0, 1]$, see Definition 4 and Definition 5. We state this as an assumption.

Assumption 2 *The operator S_2 is α-averaged and β-negatively averaged with $\alpha \in (0, 1]$ and $\beta \in (0, 1]$.*

Theorem 1 *Suppose that Assumption 1 and Assumption 2 hold. Further, let $\delta_\alpha = 2\alpha - 1$ and $\delta_\beta = 2\beta - 1$. Then the envelope function F in (7) satisfies*

$$F(x) - F(y) - \langle \nabla F(y), x - y\rangle \geq \tfrac{1}{2}\langle (P - \delta_\beta P^2)(x - y), x - y\rangle$$

and

$$F(x) - F(y) - \langle \nabla F(y), x - y\rangle \leq \tfrac{1}{2}\langle (P + \delta_\alpha P^2)(x - y), x - y\rangle$$

for all $x, y \in \mathbb{R}^n$.

A proof to this result is found in Appendix A.

As seen in Section 2.2.3, such bounds have many implications on the properties of the function. Next, we provide some in the form of corollaries.

Corollary 1 *Suppose that Assumption 1 and Assumption 2 hold and that P is positive semidefinite. Let $\delta_\alpha = 2\alpha - 1$ and $\delta_\beta = 2\beta - 1$. Then*

$$\tfrac{1}{2}\|x - y\|^2_{P-\delta_\beta P^2} \leq F(x) - F(y) - \langle \nabla F(y), x - y\rangle \leq \tfrac{1}{2}\|x - y\|^2_{P+\delta_\alpha P^2}$$

where $P - \delta_\beta P^2$ is positive semidefinite.

Proof. It follows directly from Theorem 1 and Lemma 6 in Appendix C. □

Corollary 2 *Suppose that Assumption 1 and Assumption 2 hold and that either of the following holds:*

(i) P is positive definite and contractive

(ii) P is positive definite and $\beta \in (0, 1)$ in the negative averagedness

Let $\delta_\alpha = 2\alpha - 1$ and $\delta_\beta = 2\beta - 1$. Then F is 1-strongly convex w.r.t. $\|\cdot\|_{P-\delta_\beta P^2}$ and 1-smooth w.r.t. $\|\cdot\|_{P+\delta_\alpha P^2}$.

Proof. To show the strong convexity claim, it is sufficient to apply Theorem 1 and show that $P - \delta_\beta P^2$ is positive definite, i.e., that $\lambda_{\min}(P - \delta_\beta P^2)$ is positive. In *(i)*, $\lambda_i(P) \in (0, 1)$ and $\delta_\beta \in (-1, 1]$ and in *(ii)*, $\lambda_i(P) \in (0, 1]$ and $\delta_\beta \in (-1, 1)$. From Lemma 6 it follows that in both cases, $\lambda_{\min}(P - \delta_\beta P^2)$ is positive. The smoothness claim follows immediately from Theorem 1 and Definition 8. □

Next, we show a less tight characterization of the envelope function that does not take the shape of the upper and lower bounds into account.

Corollary 3 *Suppose that Assumption 1 and Assumption 2 hold. Let* $m = \lambda_{\min}(P)$*,* $L = \lambda_{\max}(P)$*,* $\delta_\alpha = 2\alpha - 1 \in [-0.5, 1]$*, and* $\delta_\beta = 2\beta - 1 \in [-0.5, 1]$*. Then*

$$\tfrac{\beta_l}{2}\|x - y\|^2 \leq F(x) - F(y) - \langle \nabla F(y), x - y\rangle \leq \tfrac{\beta_u}{2}\|x - y\|^2$$

where $\beta_l = \min(m(1 - \delta_\beta m), L(1 - \delta_\beta L))$ *and* $\beta_u = L(1 + \delta_\alpha L)$*.*

Proof. This follows from Theorem 1, Lemma 6, and Lemma 7. □

We restricted δ_α and δ_β to $[-0.5, 1]$ (i.e, α and β to $[0.25, 1]$) in this result for convenience of the statement. Similar results for other δ_β and δ_α (and a sharpening of the result when $\delta_\beta \in [-0.5, 0]$) can be concluded from Lemma 6 and Lemma 7.

From Corollary 3, the following two results are immediate.

Corollary 4 *Suppose that Assumption 1 and Assumption 2 hold. Let* $\delta_\alpha = 2\alpha - 1 \in [-0.5, 1]$*,* $\delta_\beta = 2\beta - 1 \in [-0.5, 1]$*,* $m = \lambda_{\min}(P)$*, and* $L = \lambda_{\max}(P)$ *and suppose that either of the following two conditions holds:*

(i) P *is positive definite with* $\lambda_{\min}(P) \in (0, 1)$ *and* $\lambda_{\max}(P) \in [m, 1)$

(ii) P *is positive definite with* $\lambda_{\min}(P) \in (0, 1]$ *and* $\delta_\beta = 2\beta - 1 \in [-0.5, 1)$

Then F *is* $\min(m(1 - \delta_\beta m), L(1 - \delta_\beta L))$*-strongly convex (w.r.t.* $\|\cdot\|$*) and* $L(1 + \delta_\alpha L)$*-smooth (w.r.t.* $\|\cdot\|$*).*

Corollary 5 *Suppose that Assumption 1 and Assumption 2 hold and that* P *is positive semidefinite, i.e., that* $\lambda_{\min}(P) \geq 0$*. Let* $L = \lambda_{\max}(P)$*,* $\delta_\beta = 2\beta - 1 \in [-0.5, 1]$*, and* $\delta_\alpha = 2\alpha - 1 \in [-0.5, 1]$*. Then* F *is convex and it is* $L(1 + \delta_\alpha L)$*-smooth (or equivalently* ∇F *is* $L(1 + \delta_\alpha L)$*-Lipschitz continuous).*

The results in Theorem 1 and its corollaries hold for α-averaged and β-negatively averaged operators S_2. In Proposition 1, some properties that are equivalent to averagedness and negative averagedness are stated. Therefore, we can use these equivalent properties instead when stating the above results. This is done in the following to propositions.

Proposition 9 *Suppose that Assumption 1 holds and that* S_2 *is* δ*-Lipschitz continuous with* $\delta \in [0, 1]$*. Then all results in this section hold with* $\delta_\beta = \delta_\alpha = \delta$*.*

Proposition 10 *Suppose that Assumption 1 holds and that* S_2 *is* $\frac{1}{\delta}$*-cocoercive with* $\delta \in [0, 1]$*. Then all results in this section hold with* $\delta_\beta = \delta$ *and* $\delta_\alpha = 0$*.*

3.3 Relation to averaged operator iteration

As noted in [34, 33], the forward-backward and Douglas-Rachford splitting methods are variable metric gradient methods applied to their respective envelope functions. In our setting with S_1 being affine, it reduces to a fixed-metric scaled gradient method. Here, we show that this observation holds also in our setting.

We apply the following scaled gradient method to the envelop function F:

$$x^{k+1} = x^k - \alpha P^{-1} \nabla F(x^k).$$

This scaled gradient method is equivalent to the underlying averaged iteration:

$$\begin{aligned} x^{k+1} &= x^k - \alpha P^{-1} \nabla F(x^k) \\ &= x^k - \alpha P^{-1} P(S_2 S_1 x^k - x^k) \\ &= x^k - \alpha(S_2 S_1 x^k - x^k) \\ &= (1-\alpha)x^k + \alpha S_2 S_1 x^k. \end{aligned}$$

Therefore, the basic method can be interpreted as a scaled gradient method applied to the envelope function.

This is most probably not the most efficient way to find a stationary point of the envelope function (or equivalently a fixed-point to S_2S_1). At least in the convex setting (for the envelope), there are numerous methods that can minimize smooth functions such as truncated Newton methods, quasi-Newton methods, and nonlinear conjugate gradient descent. See [31] for an overview of such methods and [34, 37] for some of these methods applied to the forward-backward envelope. Evaluating which ones that are most efficient and devising new methods to improve performance is outside the scope of this paper.

4 Special cases

In this section, we present four special cases of our envelope function. Namely the Moreau envelope [28], the forward-backward envelope [34, 37], the Douglas-Rachford envelope [33], and the ADMM envelope (which is a special case of the Douglas-Rachford envelope).

The forward-backward and Douglas-Rachford envelopes in [34, 37, 33] are stated in a more general setting than our envelope in (7). Translated to our setting, they only require that f_1 that defines S_1 through $S_1 = \nabla f_1$ is twice continuously differentiable (as opposed to quadratic in our case). To get the forward-backward [34, 37] and Douglas-Rachford [33] envelopes in their full generality as special cases, we state the following more general envelope function

$$F(x) = \langle \nabla f_1(x), x\rangle - f_1(x) - f_2(\nabla f_1(x)). \tag{10}$$

When $f_1(x) = \frac{1}{2}\langle Px, x\rangle + \langle q, x\rangle$ it reduces to (7) since then $\langle \nabla f_1(x), x\rangle - f_1(x) = \langle Px + q, x\rangle - (\frac{1}{2}\langle Px, x\rangle + \langle q, x\rangle) = \frac{1}{2}\langle Px, x\rangle$. The gradient of the envelope function in (10) is

$$\begin{aligned}\nabla F(x) &= \nabla^2 f_1(x)x + \nabla f_1(x) - \nabla f_1(x) - \nabla^2 f_1(x)\nabla f_2(\nabla f_1(x))\\ &= \nabla^2 f_1(x)(x - \nabla f_2(\nabla f_1(x)))\\ &= \nabla^2 f_1(x)(x - S_2S_1x).\end{aligned}$$

If $\nabla^2 f_1(x)$ is nonsingular for all x, the set of stationary points of the envelope coincides with the fixed-point set of $S = S_2S_1$. We do not provide any properties of the envelope functions in this setting (it is left as future work), but merely show that that it generalizes the previously known envelope functions.

In the more restricted setting with $S_1 = \nabla f_1$ being affine, we provide envelope function properties that coincide with or improve corresponding results in the literature for the special cases.

4.1 Preliminaries

Before we present the special cases, we introduce some functions whose gradients are operators that are used in the respective underlying methods. Most importantly, we will introduce a function whose gradient is the proximal operator, which is defined as follows:

$$\mathrm{prox}_{\gamma f}(z) := \underset{x}{\mathrm{argmin}}\{f(x) + \tfrac{1}{2\gamma}\|x - z\|^2\},$$

where $\gamma > 0$ is a parameter. To do this, we introduce the following function which is a scaling and regularization of f:

$$r_{\gamma f}(x) := \gamma f(x) + \tfrac{1}{2}\|x\|^2 \tag{11}$$

This is related to the proximal operator of f as follows:

Proposition 11 *Suppose that* $f : \mathbb{R}^n \to \mathbb{R} \cup \{\infty\}$ *is proper closed and convex. The proximal operator* $\mathrm{prox}_{\gamma f}$ *then satisfies*

$$\mathrm{prox}_{\gamma f} = \nabla r_{\gamma f}^*$$

where r *is defined in* (11)*.*

This result is from [20, Proposition 1] and implies that the proximal operator is the gradient of a convex function.

A special case is when $f = \iota_C$, where ι_C is the indicator function for the nonempty closed and convex set C. The proximal operator then reduces to the projection operator. The projection operator onto C is denoted by Π_C and the corresponding regularized function is denoted and defined by

$$r_C(x) := \iota_C(x) + \tfrac{1}{2}\|x\|^2. \tag{12}$$

With this notation, $\Pi_C(x) = \nabla r_C^*(x)$. Next, we introduce a linear combination between r^* and $\frac{1}{2}\|\cdot\|^2$, namely

$$p_{\gamma f}^{\alpha}(x) := \alpha r_{\gamma f}^*(x) + \tfrac{1-\alpha}{2}\|x\|^2, \tag{13}$$

where we typically require that $\alpha \in (0, 2]$. The gradient of $p^\alpha_{\gamma f}$ is denoted by $P^\alpha_{\gamma f}$ and is given by

$$P^\alpha_{\gamma f}(x) := \nabla p^\alpha_{\gamma f}(x) = \alpha \text{prox}_{\gamma f}(x) + (1 - \alpha)x. \tag{14}$$

This is called a relaxed proximal mapping. Some special cases of this will have their own notation. Letting $\alpha = 2$, we get the reflected proximal operator

$$R_{\gamma f}(x) := P^2_{\gamma f}(x) = 2\text{prox}_{\gamma f}(x) - x. \tag{15}$$

When $f = \iota_C$, we will use notation p^α_C, P^α_C, and R_C for (13), (14), and (15) respectively. That is

$$\begin{aligned} p^\alpha_C(x) &:= \alpha r^*_C(x) + \tfrac{1-\alpha}{2}\|x\|^2, &(16)\\ P^\alpha_C(x) &:= \nabla p^\alpha_C(x) = \alpha \Pi_C(x) + (1 - \alpha)x &(17)\\ R_C(x) &:= 2\Pi_C(x) - x. &(18) \end{aligned}$$

We refer to (17) as a relaxed projection, and (18) as a reflection. So, the proximal and projected operators and their relaxed and reflected variants are gradients of functions.

We conclude with the straightforward observation that

$$(x - \gamma \nabla f(x)) = \nabla \left(\tfrac{1}{2}\|x\|^2 - \gamma f(x)\right).$$

That is, the gradient step operator is the gradient of the function $\frac{1}{2}\|x\|^2 - \gamma f(x)$.

4.2 The proximal point algorithm

The proximal point algorithm solves problems of the form

$$\text{minimize } f(x)$$

where $f : \mathbb{R}^n \to \mathbb{R} \cup \{\infty\}$ is proper closed and convex.

The algorithm repeatedly applies the proximal operator of f and is given by

$$x^{k+1} = \text{prox}_{\gamma f}(x^k), \tag{19}$$

where $\gamma > 0$ is a parameter. This algorithm is mostly of conceptual interest since it is often as computationally demanding to evaluate the prox as to minimize the function f itself.

Its envelope function, which is called the Moreau envelope [28], is a scaled version of our envelope F in (7). The scaling factor is γ^{-1} and F in (7) is obtained by letting $S_1 x = \nabla f_1(x) = x$, i.e., $P = I$ and $q = 0$, and $f_2 = r^*_{\gamma f}$, where $r_{\gamma f}$ is defined in (11). The resulting envelope function f^γ is given by

$$f^\gamma(x) = \gamma^{-1} F(x) = \gamma^{-1}\left(\tfrac{1}{2}\|x\|^2 - r^*_{\gamma f}(x)\right), \tag{20}$$

and its gradient satisfies

$$\nabla f^\gamma(x) = \gamma^{-1}\left(x - \text{prox}_{\gamma f}(x)\right).$$

The following properties of the Moreau envelope follow directly from Corollary 5 and Proposition 10 since the proximal operator is 1-cocoercive (see Remark 4 and [2, Proposition 12.27]).

Proposition 12 *The Moreau envelope f^γ in* (20) *is differentiable and convex and ∇f^γ is γ^{-1}-Lipschitz continuous.*

This coincides with previously known properties of the Moreau envelope, see [2, Chapter 12].

4.3 Forward-backward splitting

Forward-backward splitting solves problems of the form

$$\text{minimize } f(x) + g(x) \tag{21}$$

where $f : \mathbb{R}^n \to \mathbb{R}$ is convex with an L-Lipschitz (or equivalently $\frac{1}{L}$-cocoercive) gradient, and $g : \mathbb{R}^n \to \mathbb{R} \cup \{\infty\}$ is proper closed and convex.

The algorithm performs a forward step then a backward step and is given by

$$x^{k+1} = \text{prox}_{\gamma g}(I - \gamma \nabla f)x^k, \tag{22}$$

where $\gamma \in (0, \frac{2}{L})$ is a parameter.

The envelope function, which is called the forward-backward envelope [34, 37], is a scaled version of our envelope F in (10) and applies when f is twice continuously differentiable and ∇F is Lipschitz continuous. The scaling factor is γ^{-1} and F in (10) is obtained by letting $f_1 = \frac{1}{2}\|\cdot\|^2 - \gamma f$ and $f_2 = r_{\gamma g}^*$, where $r_{\gamma g}$ is defined in (11). The resulting forward-backward envelope function is

$$F_\gamma^{\text{FB}}(x) = \gamma^{-1}\left(\langle x - \gamma\nabla f(x), x\rangle - (\tfrac{1}{2}\|x\|^2 - \gamma f(x)) - r_{\gamma g}^*(x - \gamma \nabla f(x))\right).$$

The gradient of this function is

$$\begin{aligned}
\nabla F_\gamma^{\text{FB}}(x) &= \gamma^{-1}\big((I - \gamma\nabla^2 f(x))x + (x - \gamma\nabla f(x)) - (x - \gamma\nabla f(x)) \\
&\quad - (I - \gamma\nabla^2 f(x))\text{prox}_{\gamma g}(x - \gamma\nabla f(x))\big) \\
&= \gamma^{-1}(I - \gamma\nabla^2 f(x))\left(x - \text{prox}_{\gamma g}(x - \gamma\nabla f(x))\right)
\end{aligned}$$

which coincides with the gradient in [34, 37]. As described in [34, 37], the stationary points of the envelope coincides with the fixed-points of $x - \text{prox}_{\gamma g}(x - \gamma\nabla f(x))$ if $(I - \gamma\nabla^2 f(x))$ is nonsingular.

4.3.1 S_1 affine

We provide properties of the forward-backward envelope in the more restrictive setting where $S_1 = \nabla f_1 = (I - \gamma\nabla f)$ is affine. This happens if f is convex quadratic, i.e., $f(x) = \frac{1}{2}\langle Hx, x\rangle + \langle h, x\rangle$ with $H \in \mathbb{R}^{n\times n}$ positive semidefinite and $h \in \mathbb{R}^n$. Then $S_1 x = Px + q$ with $P = (I - \gamma H)$ and $q = -\gamma h$.

In this setting, the following result follows immediately from Corollary 1 and Proposition 10 (where Proposition 10 is invoked since $S_2 = \mathrm{prox}_{\gamma g}$ is 1-cocoercive, see Remark 4 and [2, Proposition 12.27]).

Proposition 13 *Assume that $f(x) = \frac{1}{2}\langle Hx, x\rangle + \langle h, x\rangle$ and $\gamma \in (1, \frac{1}{L})$ where $L = \lambda_{\max}(H)$. Then the forward-backward envelope F_γ^{FB} satisfies*

$$\tfrac{1}{2\gamma}\|x-y\|_{P-P^2}^2 \leq F_\gamma^{\mathrm{FB}}(x) - F_\gamma^{\mathrm{FB}}(y) - \langle \nabla F_\gamma^{\mathrm{FB}}(y), x-y\rangle \leq \tfrac{1}{2\gamma}\|x-y\|_P^2$$

for all $x, y \in \mathbb{R}^n$, where $P = (I - \gamma H)$ is positive definite. If in addition $\lambda_{\min}(H) = m > 0$, then $P - P^2$ is positive definite and F_γ^{FB} is γ^{-1}-strongly convex w.r.t. $\|\cdot\|_{P-P^2}$.

Less tight bounds for the forward-backward envelope are provided next. These follow immediately from Corollary 4, Corollary 5, and Proposition 10.

Proposition 14 *Assume that $f(x) = \frac{1}{2}\langle Hx, x\rangle + \langle h, x\rangle$, that $\gamma \in (0, \frac{1}{L})$ where $L = \lambda_{\max}(H)$, and that $m = \lambda_{\min}(H) \geq 0$. Then the forward-backward envelope F_γ^{FB} is $\gamma^{-1}(1-\gamma m)$-smooth and $\min\left((1-\gamma m)m, (1-\gamma L)L\right)$-strongly convex (both w.r.t. to the induced norm $\|\cdot\|$).*

This result is a slight improvement of the corresponding result in [34, Theorem 2.3]. The strong convexity moduli are the same, but this smoothness constant is a factor two smaller.

4.4 Douglas-Rachford splitting

Douglas-Rachford splitting solves problems of the form

$$\text{minimize } f(x) + g(x) \tag{23}$$

where $f : \mathbb{R}^n \to \mathbb{R} \cup \{\infty\}$ and $g : \mathbb{R}^n \to \mathbb{R} \cup \{\infty\}$ are proper closed and convex functions.

The algorithm performs two reflection steps (15), then an averaging according to

$$z^{k+1} = (1-\alpha)z^k + \alpha R_{\gamma g} R_{\gamma f} z^k \tag{24}$$

where $\gamma > 0$ and $\alpha \in (0, 1)$ are parameters. The objective is to find a fixed-point $\bar{z}$ to $R_{\gamma g} R_{\gamma f}$, from which a solution to (23) can be computed as $\mathrm{prox}_{\gamma f}\bar{z}$, see [2, Proposition 25.1].

The envelope function from [33], which is called the Douglas-Rachford envelope, is a scaled version of the basic envelope function F in (10) and applies when f is twice continuously differentiable and ∇F is Lipschitz continuous. The scaling factor is $(2\gamma)^{-1}$ and F is obtained by letting $f_1 = p_{\gamma f}^2$ with gradient $\nabla f_1 = S_1 = R_{\gamma f}$ and $f_2 = p_{\gamma g}^2$, where $p_{\gamma g}^2$ is defined in (13). The Douglas-Rachford envelope function becomes

$$F_\gamma^{\mathrm{DR}}(z) = (2\gamma)^{-1}\left(\langle R_{\gamma f}(z), z\rangle - p_{\gamma f}^2(z) - p_{\gamma g}^2(R_{\gamma f} z)\right). \tag{25}$$

The gradient of this function is

$$\nabla F_\gamma^{\mathrm{DR}}(z) = (2\gamma)^{-1}\big(\nabla R_{\gamma f}(z)z + R_{\gamma f} - R_{\gamma f} - \nabla R_{\gamma f}(z)R_{\gamma g}(R_{\gamma f}(z))\big)$$
$$= (2\gamma)^{-1}\nabla R_{\gamma f}(z)(z - R_{\gamma g}R_{\gamma f}(z)).$$

which coincides with the gradient in [33] since $\nabla R_{\gamma f} = 2\nabla\mathrm{prox}_{\gamma f} - I$ and

$$z - R_{\gamma g}R_{\gamma f}z = z - 2\mathrm{prox}_{\gamma g}(2\mathrm{prox}_{\gamma f}(z) - z) + 2\mathrm{prox}_{\gamma f}(z) - z$$
$$= 2(\mathrm{prox}_{\gamma f}(z) - \mathrm{prox}_{\gamma g}(2\mathrm{prox}_{\gamma f}(z) - z)).$$

As described in [33], the stationary points of the envelope coincides with the fixed-points of $x - R_{\gamma g}R_{\gamma f}$ if $\nabla R_{\gamma f}$ is nonsingular.

4.4.1 S_1 affine

We state properties of the Douglas-Rachford envelope in the more restrictive setting where $S_1 = R_{\gamma f}$ is affine. This holds if f is convex quadratic, i.e., of the form

$$f(x) = \tfrac{1}{2}\langle Hx, x\rangle + \langle h, x\rangle.$$

The operator S_1 becomes

$$S_1(z) = R_{\gamma f}(z) = 2(I + \gamma H)^{-1}(z - \gamma h) - z,$$

which confirms that it is affine. We implicitly define P and q through $S_1 = R_{\gamma f} = P(\cdot) + q$, and note that they are given by $P = 2(I + \gamma H)^{-1} - I$ and $q = -2\gamma(I + \gamma H)^{-1}h$.

In this setting, the following result follows immediately from Corollary 1 since $S_2 = R_{\gamma g}$ is nonexpansive (1-averaged and 1-negatively averaged).

Proposition 15 *Assume that $f(x) = \frac{1}{2}\langle Hx, x\rangle + \langle h, x\rangle$ and $\gamma \in (0, \frac{1}{L})$ where $L = \lambda_{\max}(H)$. Then the Douglas-Rachford envelope F_γ^{DR} satisfies*

$$\tfrac{1}{4\gamma}\|z - y\|_{P-P^2}^2 \leq F_\gamma^{\mathrm{DR}}(z) - F_\gamma^{\mathrm{DR}}(z) - \langle \nabla F_\gamma^{\mathrm{DR}}(y), z - y\rangle \leq \tfrac{1}{4\gamma}\|z - y\|_{P+P^2}^2$$

for all $y, z \in \mathbb{R}^n$, where $P = 2(I + \gamma H)^{-1} - I$ is positive definite. If in addition $\lambda_{\min}(H) = m > 0$, then $P - P^2$ is positive definite and F_γ^{DR} is $(2\gamma)^{-1}$-strongly convex w.r.t. $\|\cdot\|_{P-P^2}$.

The following less tight characterization of the Douglas-Rachford envelope follows from Corollary 4 and Corollary 5.

Proposition 16 *Assume that $f(x) = \frac{1}{2}\langle Hx, x\rangle + \langle h, x\rangle$, that $\gamma \in (0, \frac{1}{L})$ where $L = \lambda_{\max}(H)$, and that $m = \lambda_{\min}(H) \geq 0$. Then the Douglas-Rachford envelope F_γ^{DR} is $\frac{1-\gamma m}{(1+\gamma m)^2}\gamma^{-1}$-smooth and $\min\left(\frac{(1-\gamma m)m}{(1+\gamma m)^2}, \frac{(1-\gamma L)L}{(1+\gamma L)^2}\right)$-strongly convex.*

The strong convexity modulus of this result coincides with the corresponding one in [33, Theorem 2]. The smoothness constant in this result is $\frac{1}{1+\gamma m}$ times that in [33, Theorem 2], i.e., it is slightly smaller.

4.5 ADMM

The alternating direction method of multipliers (ADMM) solves problems of the form (23). It is well known [15] that ADMM can be interpreted as Douglas-Rachford applied to the dual of (23), namely to

$$\text{minimize } f^*(\mu) + g^*(-\mu). \tag{26}$$

So the algorithm is given by

$$v^{k+1} = (1-\alpha)v^k + \alpha R_{\rho(g^*\circ -I)} R_{\rho f} v^k \tag{27}$$

where $\rho > 0$ is a parameter, and $R_{\rho f}$ the reflected proximal operator (15) and $(g^* \circ -I)$ is the composition that satisfies $(g^* \circ -I)(\mu) = g^*(-\mu)$.

In accordance with the Douglas-Rachford envelope (25), the ADMM envelope is defined as

$$F_\rho^{\mathrm{ADMM}}(v) = (2\rho)^{-1}\left(\langle R_{\rho f^*}(v), v\rangle - p^2_{\rho f^*}(v) - p^2_{\rho(g^*\circ -I)}(R_{\rho f^*} v)\right). \tag{28}$$

and its gradient becomes

$$\nabla F_\rho^{\mathrm{ADMM}}(v) = (2\rho)^{-1} \nabla R_{\rho f^*}(v)(v - R_{\rho(g^*\circ -I)} R_{\rho f^*}(v)).$$

In this section, we relate the ADMM algorithm and its envelope function to the Douglas-Rachford counterparts. To do so, we need the following lemma which is proven in Appendix B.

Lemma 1 *Let $g : \mathbb{R}^n \to \mathbb{R}\cup\{\infty\}$ and be proper closed and convex and $\rho > 0$. Then*

$$\begin{aligned} R_{\rho g^*}(x) &= -\rho R_{\rho^{-1}g}(\rho^{-1}x) \\ R_{\rho(g^*\circ -I)}(x) &= \rho R_{\rho^{-1}g}(-\rho^{-1}x) \\ p^2_{\rho(g^*\circ -I)}(y) &= -\rho^2 p^2_{\rho^{-1}g}(-\rho^{-1}y) \end{aligned}$$

where $R_{\rho g}$ is defined in (15) *and $p^2_{\rho g}$ is defined in* (13).

First, we show that the z^k sequence in (primal) Douglas-Rachford (24) and the v^k sequence in ADMM (i.e., dual Douglas-Rachford) in (27) differ by a factor only. This is well known [13], but the relation is stated next with a simple proof.

Proposition 17 *Assume that $\rho > 0$ and $\gamma > 0$ satisfy $\rho^{-1} = \gamma$. Further assume that $z^0 = \rho^{-1}v^0$. Then $z^k = \rho^{-1}v^k$ for all $k \geq 1$, where $\{z^k\}$ is the primal Douglas-Rachford sequence defined in* (24) *and the $\{v^k\}$ is the ADMM sequence is defined in* (27).

Proof. Lemma 1 implies that

$$\begin{aligned} v^{k+1} &= (1-\alpha)v^k + \alpha R_{\rho(g^*\circ -I)} R_{\rho f^*} v^k \\ &= (1-\alpha)v^k + \alpha\rho R_{\rho^{-1}g}(-\rho^{-1}(-\rho R_{\rho^{-1}f}(\rho^{-1}v^k))) \\ &= (1-\alpha)v^k + \alpha\rho R_{\rho^{-1}g}(R_{\rho^{-1}f}(\rho^{-1}v^k))) \end{aligned}$$

Multiply by ρ^{-1}, let $z^k = \rho^{-1} v^k$, and identify $\gamma = \rho^{-1}$ to get

$$z^{k+1} = (1-\alpha) z^k + \alpha R_{\gamma g}(R_{\gamma f}(z^k))).$$

This concludes the proof. $\square$

There is also a tight relationship between the ADMM and Douglas-Rachford envelopes. Essentially, they have opposite signs.

Proposition 18 *Assume that $\rho > 0$ and $\gamma > 0$ satisfy $\rho = \gamma^{-1}$ and that $z = \rho^{-1} v = \gamma v$. Then*

$$F_\rho^{\mathrm{ADMM}}(v) = -F_\gamma^{\mathrm{DR}}(z).$$

Proof. Using Lemma 1 several times, $\gamma = \rho^{-1}$, and $z = \rho^{-1} v$, we conclude that

$$\begin{aligned}
F_\rho^{\mathrm{ADMM}}(v) &= (2\rho)^{-1} \Big(\langle R_{\rho f^*}(v), v\rangle - p_{\rho f^*}^2(v) - p_{\rho(g^* \circ -I)}^2(R_{\rho f^*}(v)) \Big) \\
&= (2\rho)^{-1} \Big(-\rho \langle R_{\rho^{-1} f}(\rho^{-1} v), v\rangle + \rho^2 p_{\rho^{-1}(f \circ -I)}(-\rho^{-1} v) \\
&\qquad + \rho^2 p_{\rho^{-1} g}(-\rho^{-1}(-\rho R_{\rho^{-1} f}(\rho^{-1} v))) \Big) \\
&= -\tfrac{\rho}{2} \Big(\langle R_{\rho^{-1} f}(\rho^{-1} v), \rho^{-1} v\rangle - p_{\rho^{-1} f}^2(\rho^{-1} v) + p_{\rho^{-1} g}^2(R_{\rho^{-1} f}(\rho^{-1} v)) \Big) \\
&= -(2\gamma)^{-1} \big(\langle R_{\gamma f}(z), z\rangle - p_{\gamma f}^2(z) + p_{\gamma g}^2(R_{\gamma f}(z)) \big) \\
&= -F_\gamma^{\mathrm{DR}}(z)
\end{aligned}$$

This concludes the proof. $\square$

This result implies that the ADMM envelope is concave when the DR envelope is convex, and vice versa. We know from Section 4.4 that the operator $S_1 = R_{\rho f^*}$ is affine when f^* is quadratic. This happens when

$$f(x) = \begin{cases} \frac{1}{2}\langle Hx, x\rangle + \langle h, x\rangle & \text{if } Ax = b \\ \infty & \text{else} \end{cases}$$

and H is positive definite on the nullspace of A. From Proposition 15 and Proposition 16, we conclude that an appropriate choice of ρ implies that the ADMM envelope is convex. Therefore, and the Douglas-Rachford envelope is concave in this setting.

Remark 7 *The standard ADMM formulation is applied to solve problems of the form*

$$\begin{array}{ll} \textit{minimize} & \hat{f}(x) + \hat{g}(z) \\ \textit{subject to} & Ax + Bz = c \end{array}$$

Using infimal post-compositions, also called image functions, the dual of this is on the form (26)*, see e.g., [21, Appendix B] for details. So also this setting is implicitly considered.*

5 The GAP envelope

In this section, we provide an envelope function to a generalization of the classic alternating projections method in [39]. The generalization uses relaxed projections and is sometimes referred to as the method of alternating relaxed projections (MARP) [3], but we will refer to it as generalized alternating projections (GAP). The algorithm is analyzed in [25, 1, 29, 14, 7] and a more general formulation is treated in [9].

GAP solves feasibility problems with a finite number of nonempty closed and convex sets that have a nonempty intersection. Here, we consider feasibility problems with two sets:

$$\text{find } x \in C \cap D$$

where $C \subset \mathbb{R}^n$ and $D \subset \mathbb{R}^n$ are nonempty closed and convex.

The generalized alternating projections method is given by

$$x^{k+1} = (1-\alpha)x^k + \alpha P_C^{\alpha_2} P_D^{\alpha_1} x^k. \tag{29}$$

where P_C^{α} is the relaxed projection in (17), and $\alpha \in (0,1]$ and $\alpha_1, \alpha_2 \in (0,2]$. These assumptions imply that $P_C^{\alpha_2}$ is $\frac{\alpha_2}{2}$-averaged if $\alpha_2 \in (0,2)$ and nonexpansive if $\alpha_2 \in (0,2]$ (and similarly for $P_D^{\alpha_1}$). If $\alpha_1 = 2$ or $\alpha_2 = 2$, the composition $P_C^{\alpha_2} P_D^{\alpha_1}$ is nonexpansive and we need $\alpha \in (0,1)$ to arrive at an averaged iteration that guarantees convergence to a fixed-point. If $\alpha_1 = \alpha_2 = 2$, the algorithm is Douglas-Rachford splitting (see Section 4.4) applied to a feasibility problem. In this case, we have $\Pi_D(\text{fix}(P_C^{\alpha_2} P_D^{\alpha_1})) = C \cap D$. For all other feasible choices of α_1 and α_2, the fixed-point set satisfies $\text{fix}(P_C^{\alpha_2} P_D^{\alpha_1}) = C \cap D$. In either case, the algorithm performs an averaged iteration to find a fixed-point to the nonexpansive operator $P_C^{\alpha_2} P_D^{\alpha_1}$.

The algorithm is on the the general form we consider and we identify S_2 in Assumption 1 with $P_C^{\alpha_2}$ and S_1 with $P_D^{\alpha_1}$. We consider in particular the case when $S_1 = P_D^{\alpha_1}$ is affine, i.e., $S_1 = P(\cdot) + q$. This holds if D is an affine set, i.e., if $D = \{x \in \mathbb{R}^n \mid Ax = b\}$ for some linear operator A. Let N denote the linear part of the projection onto the affine set Π_D, i.e.,

$$N = \Pi_{D_0} \tag{30}$$

where $D_0 = \{x \in \mathbb{R}^n \mid Ax = 0\}$, and let d denote the constant part, to get $\Pi_D x = Nx + d$. The operator S_1 then satisfies

$$S_1 x = P_D^{\alpha_1} x = (1-\alpha_1)x + \alpha_1 \Pi_D = (1-\alpha_1)x + \alpha_1(Nx + d).$$

This implies that P and q that define the affine operator $S_1 = P(\cdot) + q$ satisfy

$$P = (1-\alpha_1)I + \alpha_1 N, \qquad\qquad q = \alpha_1 d. \tag{31}$$

The GAP envelope function follows from the general envelope in (7) and is given by

$$F_{\alpha_1,\alpha_2}^{\text{GAP}}(x) = \tfrac{1}{2}\langle Px, x\rangle - p_C^{\alpha_2}(P_D^{\alpha_1} x)$$

where $p_C^{\alpha_2}$ is defined in (16) and P is from (31). Since $P_D^{\alpha_1} = Px + q$ and $\nabla p_C^{\alpha_2} = P_C^{\alpha_2}$, its gradient satisfies

$$\begin{aligned}\nabla F^{\text{GAP}}_{\alpha_1,\alpha_2}(x) &= Px - P\nabla p_C^{\alpha_2}(Px+q)\\ &= P(x - P_C^{\alpha_2}P_D^{\alpha_1}x).\end{aligned}$$

So if P is nonsingular, the stationary points of the GAP envelope coincides with the fixed-points of $P_C^{\alpha_2}P_D^{\alpha_1}$. The following proposition follows immediately from Proposition 8.

Proposition 19 *Suppose that $\alpha_1, \alpha_2 \in (0, 2]$ and that $\alpha_1 \neq 1$. Then the set of stationary points to the gap envelope $F^{\text{GAP}}_{\alpha_1,\alpha_2}$ is the fixed-point set of $P_C^{\alpha_2}P_D^{\alpha_1}$.*

Next, we state some properties of the GAP envelope.

Proposition 20 *Suppose that $\alpha_1 \in (0, 2]$ and $\alpha_2 \in (0, 2]$. Then the GAP envelope $F^{\text{GAP}}_{\alpha_1,\alpha_2}$ satisfies*

$$\begin{aligned}\tfrac{1}{2}\langle M(x-y), x-y\rangle &\leq F^{\text{GAP}}_{\alpha_1,\alpha_2}(x) - F^{\text{GAP}}_{\alpha_1,\alpha_2}(y) - \langle \nabla F^{\text{GAP}}_{\alpha_1,\alpha_2}(y), x-y\rangle\\ &\leq \tfrac{1}{2}\langle L(x-y), x-y\rangle\end{aligned}$$

where

$$M = \alpha_1(1-\alpha_1)(I-N) \tag{32}$$

and

$$L = (1-\alpha_1)(1+(\alpha_2-1)(1-\alpha_1))I + \alpha_1(1+(\alpha_2-1)(2-\alpha_1))N \tag{33}$$

where N is defined in (30).

Proof. The operator $P_C^{\alpha_2}$ is $\frac{\alpha_2}{2}$-averaged and 1-negatively averaged (nonexpansive). So we can apply Theorem 1 with $\delta_\beta = 1$, $\delta_\alpha = \alpha_2 - 1$, and P in (31). Using $N = N^2$ (which holds since N is a projection onto a linear subspace), we conclude that

$$\begin{aligned}M = P - P^2 &= (1-\alpha_1)I + \alpha_1 N - ((1-\alpha_1)I + \alpha_1 N)^2\\ &= (1-\alpha_1)I + \alpha_1 N - ((1-\alpha_1)^2 I + 2\alpha_1(1-\alpha_1)N + \alpha_1^2 N)\\ &= ((1-\alpha_1) - (1-\alpha_1)^2)I + (\alpha_1 - (2\alpha_1 - \alpha^2))N\\ &= ((1-\alpha_1) - (1-2\alpha_1+\alpha_1^2))I + (\alpha_1^2 - \alpha_1))N\\ &= \alpha_1(1-\alpha_1)I + \alpha_1(\alpha_1 - 1))N\\ &= \alpha_1(1-\alpha_1)(I-N)\end{aligned}$$

and that

$$\begin{aligned}L = P + (\alpha_2-1)P^2 &= (1-\alpha_1)I + \alpha_1 N + (\alpha_2-1)((1-\alpha_1)I + \alpha_1 N)^2\\ &= ((1-\alpha_1) + (\alpha_2-1)(1-\alpha_1)^2)I + (\alpha_1 + (\alpha_2-1)(2\alpha_1(1-\alpha_1) + \alpha_1^2))N\\ &= (1-\alpha_1)(1+(\alpha_2-1)(1-\alpha_1))I + \alpha_1(1+(\alpha_2-1)(2-\alpha_1))N.\end{aligned}$$

This concludes the proof. $\square$

Since N is a projection operator onto a linear subspace, it has only two distinct eigenvalues, namely zero and one. Therefore, there are only two distinct eigenvalues of M and L in (32) and (33). Expressions for these eigenvalues are given in the following proposition.

Proposition 21 *The eigenvalues of M in* (32) *are*

$$\lambda_i(M) = \begin{cases} 0 & \text{for } i \text{ such that } \lambda_i(N) = 1 \\ \alpha_1(1-\alpha_1) & \text{for } i \text{ such that } \lambda_i(N) = 0 \end{cases} \tag{34}$$

and the eigenvalues of L in (33) *are*

$$\lambda_i(L) = \begin{cases} \alpha_2 & \text{for } i \text{ such that } \lambda_i(N) = 1 \\ (1-\alpha_1)(1+(\alpha_2-1)(1-\alpha_1)) & \text{for } i \text{ such that } \lambda_i(N) = 0 \end{cases} \tag{35}$$

with N defined in (30).

Proof. First note that $\lambda_i(a_1 I + a_2 N) = a_1 + a_2\lambda_i(N)$. This implies that $\lambda_i(M) = \alpha_1(1-\alpha_1)(1-\lambda_i(N))$, and (34) is proven. It also implies that

$$\lambda_i(L) = (1-\alpha_1)(1+(\alpha_2-1)(1-\alpha_1)) + \alpha_1(1+(\alpha_2-1)(2-\alpha_1))\lambda_i(N).$$

For $\lambda_i(N) = 0$, we see that (35) holds. In the case of $\lambda_i(N) = 1$, we conclude that

$$\begin{aligned} \lambda_i(L) &= (1-\alpha_1)(1+(\alpha_2-1)(1-\alpha_1)) + \alpha_1(1+(\alpha_2-1)(2-\alpha_1)) \\ &= 1-\alpha_1+\alpha_2(1-\alpha_1)^2-(1-\alpha_1)^2+\alpha_1+\alpha_1\alpha_2(2-\alpha_1)-\alpha_1(2-\alpha_1) \\ &= 1+\alpha_2(1-2\alpha_1+\alpha_1^2)-1+2\alpha_1-\alpha_1^2+\alpha_1\alpha_2(2-\alpha_1)-2\alpha_1-\alpha_1^2 \\ &= \alpha_2(1-2\alpha_1+\alpha_1^2)+\alpha_2(2\alpha_1-\alpha_1^2) \\ &= \alpha_2. \end{aligned}$$

This concludes the proof. □

Using this, we can show that for $\alpha_1 \in [1,2]$, the GAP envelope is convex on the nullspace of A and concave on its orthogonal complement, the rangespace of A^*.

Proposition 22 *Let $\mathcal{N}(A)$ denote the nullspace of A and let $\mathcal{R}(A^*)$ denote its orthogonal complement, the rangespace of A^*. Then the GAP envelope is convex and α_2-smooth when restricted to $\mathcal{R}(A^*)$. If $\alpha_1 \in [1,2]$, the GAP envelope is concave and $\alpha_1(\alpha_1-1)$-smooth when restricted to $\mathcal{N}(A)$.*

Proof. The subspace $\mathcal{R}(A^*)$ is spanned by the eigenvectors corresponding to $\lambda_i(N) = 1$. Therefore, Proposition 21 implies that for all $x, y \in \mathcal{R}(A^*)$, the lower bound in Proposition 20 becomes $\langle M(x-y), x-y\rangle = 0$ and the upper bound in Proposition 20 satisfies $\langle L(x-y), x-y\rangle = \alpha_2\|x-y\|^2$. This proves the first claim.

The second claim is proven similarly. The subspace $\mathcal{N}(A)$ is spanned by the eigenvectors corresponding to $\lambda_i(N) = 0$. Therefore, Proposition 21 implies that for all $x, y \in \mathcal{N}(A)$, the lower bound in Proposition 20 becomes $\langle M(x - y), x - y\rangle = \alpha_1(1-\alpha_1)\|x-y\|^2$ and the upper bound in Proposition 20 satisfies $\langle L(x-y), x-y\rangle = (1-\alpha_1)(1+(\alpha_2-1)(1-\alpha_1))\|x-y\|^2$. Noting that $(1-\alpha_1)(1+(\alpha_2-1)(1-\alpha_1)) \leq 0$ when $\alpha_1 \in [1,2]$ and $\alpha_2 \in (0,2]$ proves the second claim. □

The following proposition is a straightforward consequence of Proposition 20 and Proposition 21 and is stated without a proof.

Proposition 23 *Suppose that $\alpha_1 \in (0,2]$ and $\alpha_2 \in (0,2]$. Then the GAP envelope $F^{\mathrm{GAP}}_{\alpha_1,\alpha_2}$ satisfies*

$$\tfrac{\beta_l}{2}\|x-y\|^2 \leq F^{\mathrm{GAP}}_{\alpha_1,\alpha_2}(x) - F^{\mathrm{GAP}}_{\alpha_1,\alpha_2}(y) - \langle \nabla F^{\mathrm{GAP}}_{\alpha_1,\alpha_2}(y), x-y\rangle \leq \tfrac{\beta_u}{2}\|x-y\|^2$$

where $\beta_l = \min((1-\alpha_1)\alpha_1, 0)$ and $\beta_u = \max((1-\alpha_1)(1+(\alpha_2-1)(1-\alpha_1)), \alpha_2)$. If in addition $\alpha_1 \in (0,1]$, then it is convex.

If the first relaxed projection is under-relaxed, i.e., if $\alpha_1 \in (0,1]$, then the GAP envelope is convex. From Proposition 19, we also know that if $\alpha_1 \neq 1$ its set of stationary points is the fixed-point set of $P_C^{\alpha_2}P_D^{\alpha_1}$. For convex functions, all stationary points are minimizers. This therefore implies that all convex feasibility problems where one set is affine, can be solved by minimizing the smooth convex GAP envelope function by setting $\alpha_1 \in (0,1)$. In Section 6, we will see that most convex optimization problems can actually be cast on this feasibility form.

6 Cone programming

In this section we show that many convex optimization problems can be written as a convex feasibility problem involving one product of cones and one affine set. This observation is not new (see, e.g., [40]) but is included here since it is a crucial step in showing that almost all convex optimization problems can be solved by minimizing the smooth convex unconstrained GAP envelope.

The CVX optimization modeling languages [24, 11, 38] are based on transforming the provided (possibly nonsmooth and constrained) convex optimization problem to a convex cone program of the form

$$\begin{array}{ll} \text{minimize} & c^T x \\ \text{subject to} & Ax + s = b \\ & s \in \mathcal{K} \end{array}$$

where $\mathcal{K}$ is a (product of) closed and convex cones.

This transformation can be done for many convex optimization classes including QPs, LPs, SOCPs, SDPs, Lasso-type problems, etc, see [23]. To arrive at a feasibility problem originating from the cone program, we pose its dual:

$$\begin{array}{ll} \text{maximize} & -b^T y \\ \text{subject to} & -A^T y = c \\ & y \in \mathcal{K}^* \end{array}$$

Assuming strong duality, the primal and dual objectives agree, i.e., $c^T x + b^T y = 0$. Using this relation and embedding both the primal and dual problem into one master problem, we get the following primal-dual feasibility problem

$$\begin{array}{ll} \text{find} & (x, s, y) \\ \text{subject to} & \begin{bmatrix} A & I & 0 \\ 0 & 0 & -A^T \\ c^T & 0 & b^T \end{bmatrix} \begin{bmatrix} x \\ s \\ y \end{bmatrix} = \begin{bmatrix} b \\ c \\ 0 \end{bmatrix} \\ & (s, y) \in \mathcal{K} \times \mathcal{K}^* \end{array}$$

This feasibility problem has one affine subspace and one product of convex cones. (There are many other ways to construct a feasibility problem with an affine subspace and a product of convex cones. One example is the homogeneous self-dual embedding which was proposed in [40] and used in the ADMM-based solver SCS [32].)

There are various ways to solve this embedding problem, or reformulations of it, using the forward-backward envelope in [34, 37] or the Douglas-Rachford envelope in [33]. It can also be solved by finding a stationary point of the sometimes convex GAP envelope in Section 5. This provides another bridge (besides the forward-backward and Douglas-Rachford envelopes) that allows almost all convex optimization problems (at least those that can be formulated as convex cone programs and therefore be solved using the CVX modeling languages) to be solved using smooth optimization techniques.

7 Conclusions

We have presented a unified framework for envelope functions. Special cases include the Moreau envelope, the forward-backward envelope, the Douglas-Rachford and ADMM envelopes. We also presented a new envelope function, namely the generalized alternating projections (GAP) envelope. Under additional assumptions, we have provided quadratic upper and lower bounds to the general envelope function. These coincide with or improve corresponding results for the known special cases in the literature. We have also shown that almost all convex optimization problems can be solved by solving a feasibility problem and that this feasibility problem can be solved by minimizing the smooth and sometimes convex GAP envelope function.

8 Acknowledgments

Both authors are financially supported by the Swedish Foundation for Strategic Research and members of the LCCC Linneaus Center at Lund University.

References

[1] S. Agmon. The relaxation method for linear inequalities. *Canadian Journal of Mathematics*, 6(3):382–392, 1954.

[2] H. H. Bauschke and P. L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer, 2011.

[3] H. H. Bauschke, H. M. Phan, and X. Wang. The method of alternating relaxed projections for two nonconvex sets. *Vietnam Journal of Mathematics*, 42:421–450, 2014.

[4] M. Benzi. Preconditioning techniques for large linear systems: A survey. *Journal of Computational Physics*, 182(2):418–477, 2002.

[5] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.

[6] J. H. Bramble, J. E. Pasciak, and A. T. Vassilev. Analysis of the inexact Uzawa algorithm for saddle point problems. *SIAM Journal on Numerical Analysis*, 34(3):1072–1092, 1997.

[7] L. M. Bregman. Finding the common point of convex sets by the method of successive projection. *Dokl Akad. Nauk SSSR*, 162(3):487–490, 1965.

[8] A. Chambolle and T. Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, 2011.

[9] P. L. Combettes. Solving monotone inclusions via compositions of nonexpansive averaged operators. *Optimization*, 53(5–6):475–504, 2004.

[10] D. Davis and W. Yin. A three-operator splitting scheme and its optimization applications. `http://arxiv.org/abs/1504.01032`, 2015.

[11] S. Diamond and S. Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 2016. To appear.

[12] J. Douglas and H. H. Rachford. On the numerical solution of heat conduction problems in two and three space variables. *Trans. Amer. Math. Soc.*, 82:421–439, 1956.

[13] J. Eckstein. *Splitting methods for monotone operators with applications to parallel optimization*. PhD thesis, MIT, 1989.

[14] I. I. Eremin. Generalization of the Motskin-Agmon relaxation method. *Usp. mat. Nauk*, 20(2):183–188, 1965.

[15] D. Gabay. Applications of the method of multipliers to variational inequalities. In M. Fortin and R. Glowinski, editors, *Augmented Lagrangian Methods: Applications to the Solution of Boundary-Value Problems*. North-Holland: Amsterdam, 1983.

[16] D. Gabay and B. Mercier. A dual algorithm for the solution of nonlinear variational problems via finite element approximation. *Computers and Mathematics with Applications*, 2(1):17–40, 1976.

[17] E. Ghadimi, A. Teixeira, I. Shames, and M. Johansson. Optimal parameter selection for the alternating direction method of multipliers (ADMM): Quadratic problems. *IEEE Transactions on Automatic Control*, 60(3):644–658, March 2015.

[18] P. Giselsson. Tight global linear convergence rate bounds for Douglas-Rachford splitting. 2015. Submitted. Available: `http://arxiv.org/abs/1506.01556`.

[19] P. Giselsson and S. Boyd. Metric selection in fast dual forward-backward splitting. *Automatica*, 62:1–10, 2015.

[20] P. Giselsson and S. Boyd. Linear convergence and metric selection in Douglas-Rachford splitting and ADMM. 2016. Accepted for publication in Transactions on Automatic Control. Available: `http://arxiv.org/abs/1410.8479`.

[21] P. Giselsson, M. Fält, and S. Boyd. Line search for averaged operator iteration. Available: `http://arxiv.org/abs/1603.06772`, 2016.

[22] R. Glowinski and A. Marroco. Sur l'approximation, par éléments finis d'ordre un, et la résolution, par pénalisation-dualité d'une classe de problémes de dirichlet non linéaires. *ESAIM: Mathematical Modelling and Numerical Analysis - Modélisation Mathématique et Analyse Numérique*, 9:41–76, 1975.

[23] M. Grant and S. Boyd. Graph implementations for nonsmooth convex programs. In V. Blondel, S. Boyd, and H. Kimura, editors, *Recent Advances in Learning and Control*, Lecture Notes in Control and Information Sciences, pages 95–110. Springer-Verlag Limited, 2008.

[24] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. `http://cvxr.com/cvx`, March 2014.

[25] L. G. Gubin, B. T. Polyak, and E. V. Raik. The method of projections for finding the common point of convex sets. *USSR Computational Mathematics and Mathematical Physics*, 7(6):1–24, 1967.

[26] Q. Hu and J. Zou. Nonlinear inexact Uzawa algorithms for linear and nonlinear saddle-point problems. *SIAM Journal on Optimization*, 16(3):798–825, 2006.

[27] P. L. Lions and B. Mercier. Splitting algorithms for the sum of two nonlinear operators. *SIAM Journal on Numerical Analysis*, 16(6):964–979, 1979.

[28] J. J. Moreau. Proximit et dualit dans un espace hilbertien. *Bulletin de la Socit Mathmatique de France*, 93:273–299, 1965.

[29] T. S. Motzkin and I. Shoenberg. The relaxation method for linear inequalities. *Canadian Journal of Mathematics*, 6(3):383–404, 1954.

[30] Y. Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Springer Netherlands, 1st edition, 2003.

[31] J. Nocedal and S. Wright. *Numerical optimization*. Springer series in operations research and financial engineering. Springer, New York, NY, 2nd edition, 2006.

[32] B. O'Donoghue, E. Chu, N. Parikh, and S. Boyd. Conic optimization via operator splitting and homogeneous self-dual embedding. *Journal of Optimization Theory and Applications*, 2016.

[33] P. Patrinos, L. Stella, and A. Bemporad. Douglas-Rachford splitting: Complexity estimates and accelerated variants. In *Proceedings of the 53rd IEEE Conference on Decision and Control*, pages 4234–4239, Los Angeles, CA, December 2014.

[34] P. Patrinos, L. Stella, and A. Bemporad. Forward-backward truncated Newton methods for convex composite optimization. Available: `http://arxiv.org/abs/1402.6655`, 2014.

[35] R. T. Rockafellar and R. J-B. Wets. *Variational Analysis*. Springer, Berlin, 1998.

[36] M. Sion. On general minimax theorems. *Pacific Journal of Mathematics*, 8(1):171–176, 1958.

[37] L. Stella, A. Themelis, and P. Patrinos. Forward-backward quasi-Newton methods for nonsmooth optimization problems. Available: `http://arxiv.org/abs/1604.08096`, 2016.

[38] M. Udell, K. Mohan, D. Zeng, J. Hong, S. Diamond, and S. Boyd. Convex optimization in Julia. *SC14 Workshop on High Performance Technical Computing in Dynamic Languages*, 2014.

[39] J. von Neumann. *Functional Operators. Volume II. The Geometry of Orthogonal Spaces*. Princeton University Press: Annals of Mathematics Studies, 1950. Reprint of 1933 lecture notes.

[40] Y. Ye, M. J. Todd, and S. Mizuno. An o($\sqrt{n}L$)-iteration homogeneous and self-dual linear programming algorithm. *Mathematics of Operations Research*, 19(1):53–67, February 1994.

A Proof to Theorem 1

First, we establish that

$$-\delta_\alpha\|x-y\|_{P^2}^2 \leq \langle P\nabla f_2(Px+q) - P\nabla f_2(Py+q), x-y\rangle \leq \delta_\beta\|x-y\|_{P^2}^2. \tag{36}$$

We have

$$\begin{aligned}
&\langle P\nabla f_2(Px+q) - P\nabla f_2(Py+q), x-y\rangle \\
&\qquad\qquad = \langle \nabla f_2(Px+q) - \nabla f_2(Py+q), P(x-y)\rangle \\
&\qquad\qquad = \langle \nabla f_2(Px+q) - \nabla f_2(Py+q), (Px+q)-(Py+q))\rangle
\end{aligned}$$

This implies that

$$\begin{aligned}
-(2\alpha-1)\|x-y\|_{P^2}^2 &= -(2\alpha-1)\|(Px+q)-(Py-q)\|^2 \\
&\leq \langle P\nabla f_2(Px+q) - P\nabla f_2(Py+q), x-y\rangle \\
&\leq (2\beta-1)\|(Px+q)-(Py-q)\|^2 \\
&= (2\beta-1)\|x-y\|_{P^2}^2
\end{aligned}$$

where Lemma 4 and Lemma 5 are used in the inequalities. Recalling that $\delta\alpha = 2\alpha-1$ and $\delta_\beta = 2\beta-1$, this shows that (36) holds. Further, for any $\delta \in \mathbb{R}$ we have

$$\begin{aligned}
\langle \nabla F(x) - \nabla F(y), x-y\rangle &= \langle P(x-\nabla f_2\nabla f_1(x)) - P(x-\nabla f_2\nabla f_1(y)), x-y\rangle \\
&= \langle P(x-y), x-y\rangle \\
&\quad - \langle P\nabla f_2(Px+q) - P\nabla f_2(Py+q), x-y\rangle \\
&= \langle (P-\delta P^2)(x-y), x-y\rangle + \delta\|x-y\|_{P^2}^2 \\
&\quad - \langle P\nabla f_2(Px+q) - P\nabla f_2(Py+q), x-y\rangle.
\end{aligned} \tag{37}$$

Let $\delta = -\delta_\alpha$, then (37) and (36) imply

$$\langle \nabla F(x) - \nabla F(y), x-y\rangle \leq \langle (P+\delta_\alpha P^2)(x-y), x-y\rangle.$$

Let $\delta = \delta_\beta$, then (37) and (36) imply

$$\langle \nabla F(x) - \nabla F(y), x-y\rangle \geq \langle (P-\delta_\beta P^2)(x-y), x-y\rangle.$$

Applying Lemma 2 in Appendix C gives the result.

B Proof to Lemma 1

Using the Moreau decomposition [2, Theorem 14.3]

$$\mathrm{prox}_{\rho g^*}(x) = x - \rho\mathrm{prox}_{\rho^{-1}g}(\rho^{-1}x),$$

we conclude that

$$\begin{aligned}
R_{\rho g^*}(x) &= 2\mathrm{prox}_{\rho g^*}(x) - x \\
&= 2(x - \rho\mathrm{prox}_{\rho^{-1}g}(\rho^{-1}x)) - x \\
&= -\rho\left(2(\mathrm{prox}_{\rho^{-1}g}(\rho^{-1}x)) - (\rho^{-1}x)\right) \\
&= -\rho R_{\rho^{-1}g}(\rho^{-1}x)
\end{aligned}$$

and

$$\begin{aligned}
R_{\rho(g^*\circ -I)}(x) &= 2\mathrm{prox}_{\rho(g^*\circ -I)}(x) - x \\
&= -2\mathrm{prox}_{\rho g^*}(-x) - x \\
&= -2(-x - \rho\mathrm{prox}_{\rho^{-1}g}(-\rho^{-1}x)) - x \\
&= 2\rho\mathrm{prox}_{\rho^{-1}g}(-\rho^{-1}x)) + x \\
&= \rho(2\mathrm{prox}_{\rho^{-1}g}(-\rho^{-1}x) - (-\rho^{-1}x)) \\
&= \rho R_{\rho^{-1}g}(-\rho^{-1}x).
\end{aligned}$$

To show the third claim, we first derive an expression for $r^*_{\rho(g^*\circ -I)}$. We have

$$\begin{aligned}
r^*_{\rho(g^*\circ -I)}(y) &= (\rho(g^*\circ -I) + \tfrac{1}{2}\|\cdot\|^2)^*(y) \\
&= \sup_z\{\langle y, z\rangle - \rho\sup_x\{\langle z, x\rangle - g(-x)\} - \tfrac{1}{2}\|z\|^2\} \\
&= \sup_z\{\langle y, z\rangle + \rho\inf_x\{\langle z, -x\rangle + g(-x)\} - \tfrac{1}{2}\|z\|^2\} \\
&= \sup_z\{\langle y, z\rangle + \rho\inf_v\{\langle z, v\rangle + g(v)\} - \tfrac{1}{2}\|z\|^2\} \\
&= \sup_z\inf_v\{\langle y, z\rangle + \rho\langle z, v\rangle + \rho g(v) - \tfrac{1}{2}\|z\|^2\} \\
&= \inf_v\sup_z\{\langle y + \rho v, z\rangle + \rho g(v) - \tfrac{1}{2}\|z\|^2\} \\
&= \inf_v\{\tfrac{1}{2}\|y + \rho v\|^2 + \rho g(v)\} \\
&= \inf_v\{\langle y, \rho v\rangle + \tfrac{1}{2}\|\rho v\|^2 + \rho g(v)\} + \tfrac{1}{2}\|y\|^2 \\
&= -\sup_v\{\langle -y, \rho v\rangle - \tfrac{1}{2}\|\rho v\|^2 - \rho g(v)\} + \tfrac{1}{2}\|y\|^2 \\
&= -\rho^2\sup_v\{\langle -\rho^{-1}y, v\rangle - \tfrac{1}{2}\|v\|^2 - \rho^{-1}g(v)\} + \tfrac{1}{2}\|y\|^2 \\
&= -\rho^2 r^*_{\rho^{-1}g}(-\rho^{-1}y) + \tfrac{1}{2}\|y\|^2,
\end{aligned}$$

where the sup-inf swap is valid by the minimax theorem in [36] since we can construct a compact set for the z variable due to strong convexity of $\|\cdot\|^2$. This implies that

$$\begin{aligned}
p^2_{\rho(g^*\circ -I)}(y) &= 2r^*_{\rho(g^*\circ -I)}(y) - \tfrac{1}{2}\|y\|^2 \\
&= -2\rho^2 r^*_{\rho^{-1}g}(-\rho^{-1}y) + \tfrac{1}{2}\|y\|^2 \\
&= -\rho^2(2r^*_{\rho^{-1}g}(-\rho^{-1}y) - \tfrac{1}{2}\|-\rho^{-1}y\|^2) \\
&= -\rho^2 p^2_{\rho^{-1}g}(-\rho^{-1}y).
\end{aligned}$$

This concludes the proof.

C Technical lemmas

Lemma 2 *Assume that $f : \mathbb{R}^n \to \mathbb{R}$ is differentiable and that $M : \mathbb{R}^n \to \mathbb{R}^n$ and $L : \mathbb{R}^n \to \mathbb{R}^n$ are linear operators. Then*

$$-\tfrac{1}{2}\langle M(x-y), x-y\rangle \le f(x) - f(y) - \langle \nabla f(y), x-y\rangle \le \tfrac{1}{2}\langle L(x-y), x-y\rangle \tag{38}$$

if and only if

$$-\langle M(x-y), x-y\rangle \le \langle \nabla f(x) - \nabla f(y), x-y\rangle \le \langle L(x-y), x-y\rangle \tag{39}$$

Proof. Adding two copies of (38) with x and y interchanged gives

$$-\langle M(x-y), x-y\rangle \le \langle \nabla f(x) - f(y), x-y\rangle \le \langle L(x-y), x-y\rangle. \tag{40}$$

This shows that (38) implies (39). To show the other direction, we use integration. Let $h(\tau) = f(x + \tau(y-x))$, then

$$\nabla h(\tau) = \langle y - x, \nabla f(x + \tau(y-x))\rangle$$

since $f(y) = h(1)$ and $f(x) = h(0)$, we get

$$f(y) - f(x) = h(1) - h(0) = \int_0^1 \nabla h(\tau) d\tau = \int_0^1 \langle y - x, \nabla f(x + \tau(y-x))\rangle d\tau$$

Therefore

$$\begin{aligned}
f(y) - f(x) - \langle \nabla f(x), y - x\rangle &= \int_0^1 \langle \nabla f(x + \tau(y-x)), y - x\rangle d\tau - \langle \nabla f(x), y - x\rangle \\
&= \int_0^1 \langle \nabla f(x + \tau(y-x)) - \nabla f(x), y - x\rangle d\tau \\
&= \int_0^1 \tau^{-1}\langle \nabla f(x + \tau(y-x)) - \nabla f(x), \tau(y - x)\rangle d\tau \\
&= \int_0^1 \tau^{-1}\langle \nabla f(x + \tau(y-x)) - \nabla f(x), (x + \tau(y - x)) - x\rangle d\tau.
\end{aligned}$$

Using the upper bound in (39), we get

$$\begin{aligned}
&\int_0^1 \tau^{-1}\langle \nabla f(x + \tau(y-x)) - \nabla f(x), (x + \tau(y - x)) - x\rangle d\tau \\
&\qquad \le \int_0^1 \tau^{-1}\langle L\tau(x-y), \tau(x-y)\rangle d\tau \\
&\qquad = \langle L(x-y), x-y\rangle \int_0^1 \tau d\tau \\
&\qquad = \tfrac{1}{2}\langle L(x-y), x-y\rangle.
\end{aligned}$$

Similarly, using the lower bound in (39), we get

$$\begin{aligned}\int_0^1 &\tau^{-1}\langle \nabla f(x+\tau(y-x)) - \nabla f(x), (x+\tau(y-x)) - x\rangle d\tau\\ &\geq -\int_0^1 \tau^{-1}\langle M\tau(x-y), \tau(x-y)\rangle d\tau\\ &= -\langle M(x-y), x-y\rangle \int_0^1 \tau d\tau\\ &= -\tfrac{1}{2}\langle M(x-y), x-y\rangle.\end{aligned}$$

This concludes the proof. □

Lemma 3 *Assume that $f : \mathbb{R}^n \to \mathbb{R}$ is differentiable and that L is positive definite. Then that f is L-smooth, i.e., that f satisfies*

$$|f(x) - f(y) - \langle \nabla f(y), x-y\rangle| \leq \tfrac{\beta}{2}\|x-y\|_L^2 \tag{41}$$

holds for all $x, y \in \mathbb{R}^n$ is equivalent to that ∇f is β-Lipschitz continuous w.r.t. $\|\cdot\|_L$, i.e., that

$$\|\nabla f(x) - \nabla f(y)\|_{L^{-1}} \leq \beta\|x-y\|_L \tag{42}$$

holds for all $x, y \in \mathbb{R}^n$.

Proof. We start by proving the result using the induced norm $\|\cdot\|$ only, i.e., in the Hilbert space setting. (This covers, e.g., the setting with inner-product $\langle x,y\rangle_H = \langle Hx, y\rangle$ and scaled norm $\|\cdot\|_H = \sqrt{\langle x,y\rangle_H}$ that will be used later.) To do this, we introduce the functions $h := \frac{1}{\beta}f$ and $r := \frac{1}{2}(h + \frac{1}{2}\|\cdot\|^2)$.

Since $L = I$ in the norm, the condition (42) is β-Lipschitz continuity of ∇f (w.r.t. $\|\cdot\|$). This is equivalent to that $\nabla h = \frac{1}{\beta}\nabla f$ is nonexpansive, which by [2, Proposition 4.2] is equivalent to that $\frac{1}{2}(\nabla h + I) = \nabla\left(\frac{1}{2}(h + \frac{1}{2}\|\cdot\|^2)\right) = \nabla r$ is firmly nonexpansive (or equivalently 1-cocoercive). This, is equivalent to (see [30, Theorem 2.1.5] and [2, Definition 4.4]) that:

$$0 \leq r(x) - r(y) - \langle \nabla r(y), x-y\rangle \leq \tfrac{1}{2}\|x-y\|^2.$$

holds for all $x, y \in \mathbb{R}^n$. Multiplying by 2 and using $2r = h + \frac{1}{2}\|\cdot\|^2$, this is equivalent to that

$$\begin{aligned}0 &\leq h(x) - h(y) - \langle \nabla h(y), x-y\rangle + \tfrac{1}{2}(\|x\|^2 - \|y\|^2 - 2\langle y, x-y\rangle)\\ &= h(x) - h(y) - \langle \nabla h(y), x-y\rangle + \tfrac{1}{2}\|x-y\|^2 \leq \|x-y\|^2.\end{aligned}$$

Multiplying by β and using $f = \beta h$, this is equivalent to

$$-\tfrac{\beta}{2}\|x-y\| \leq f(x) - f(y) - \langle \nabla f(y), x-y\rangle \leq \tfrac{\beta}{2}\|x-y\|^2.$$

This chain of equivalences show that the conditions are equivalent when $L = I$.

Next, we show that the scaled version holds. To do this, introduce the space $\mathbb{H}_H$ with inner-product $\langle x, y\rangle_H = \langle Hx, y\rangle$ and induced norm $\|\cdot\|_H = \sqrt{\langle Hx, x\rangle}$ and the space $\mathbb{E}_L$ inner-product $\langle x, y\rangle$ and induced norm $\|\cdot\|_L = \sqrt{\langle Lx, x\rangle}$. Further let $H = L$ and define $f_h : \mathbb{H}_H \to \mathbb{R}$ and $f_l : \mathbb{E}_L \to \mathbb{R}$ that satisfy $f_h(x) = f_l(x)$ for all $x \in \mathbb{R}^n$. We have already shown that (41) and (42) are equivalent for f_h that is defined on the Hilbert space $\mathbb{H}_H$. To show that it also holds for f_l defined on $\mathbb{E}_L$, we show that the conditions (41) and (42) are equivalent if defined for f_h on $\mathbb{H}_H$ and if defined for f_l on $\mathbb{E}_L$, when $L = H$.

By definition of the gradient, ∇f_l and ∇f_h must satisfy

$$\langle \nabla f_l(y), x - y\rangle = \langle \nabla f_h(y), x - y\rangle_H = \langle H\nabla f_h(y), x - y\rangle$$

for all $x, y \in \mathbb{R}^n$. This implies that $\nabla f_h = H^{-1}\nabla f_l = L^{-1}\nabla f_l$. Therefore that (41) holds for f_l on $\mathbb{E}_L$ is equivalent to that it holds for f_h on $\mathbb{H}_H$.

Further,

$$\begin{aligned}
\|\nabla f_h(x) - \nabla f_h(y)\|_H^2 &= \langle \nabla f_h(x) - \nabla f_h(y), \nabla f_h(x) - \nabla f_h(y)\rangle_H \\
&= \langle L^{-1}(\nabla f(x) - \nabla f(y)), L^{-1}(\nabla f(x) - \nabla f(y))\rangle_L \\
&= \langle \nabla f(x) - \nabla f(y), \nabla f(x) - \nabla f(y)\rangle_{L^{-1}} \\
&= \|\nabla f(x) - \nabla f(y)\|_{L^{-1}}^2.
\end{aligned}$$

So that (42) holds for f_l on $\mathbb{E}_L$ is equivalent to that it holds for f_h on $\mathbb{H}_H$. This concludes the proof. $\square$

Lemma 4 *Assume that f is differentiable. Then ∇f is α-averaged with $\alpha \in (0, 1]$ if and only if*

$$-(2\alpha - 1)\|x - y\|^2 \le \langle \nabla f(x) - \nabla f(y), x - y\rangle \le \|x - y\|^2. \tag{43}$$

Proof. The operator ∇f is α-averaged if and only if $\nabla f = (1 - \alpha)I + \alpha R$ for some nonexpansive operator R. Therefore, ∇f is α-averaged if and only if $\nabla f - (1 - \alpha)I$ is α-Lipschitz continuous, since $\nabla f - (1 - \alpha)I = \alpha R$. Letting $g := f - \frac{1-\alpha}{2}\|\cdot\|^2$, we get $\nabla g = \alpha R$. Therefore ∇g is α-Lipschitz. According to Lemma 3 this is equivalent to that

$$|g(x) - g(y) - \langle \nabla g(y), x - y\rangle| \le \tfrac{\alpha}{2}\|x - y\|^2$$

or equivalently

$$|f(x) - f(y) - \langle \nabla f(y), x - y\rangle - \tfrac{1-\alpha}{2}\|x - y\|^2| \le \tfrac{\alpha}{2}\|x - y\|^2$$

which is equivalent to

$$-\tfrac{2\alpha-1}{2}\|x - y\|^2 \le f(x) - f(y) - \langle \nabla f(y), x - y\rangle \le \tfrac{1}{2}\|x - y\|^2.$$

Applying Lemma 2 gives the result. $\square$

Lemma 5 *Assume that f is differentiable. Then ∇f is β-negatively averaged with $\beta \in (0, 1]$ if and only if*

$$-\|x - y\|^2 \leq \langle \nabla f(x) - \nabla f(y), x - y\rangle \leq (2\beta - 1)\|x - y\|^2. \tag{44}$$

Proof. This follows immediately from 4 since $-\nabla f$ is β-averaged by definition. □

Lemma 6 *Suppose that P is a linear self-adjoint and nonexpansive operator with largest eigenvalue $\lambda_{\max}(P) = L$ and smallest eigenvalue $\lambda_{\min}(P) = m$, satisfying $-1 \leq m \leq L \leq 1$. Further suppose that $\delta \in [-1, 1]$ and let j be the index that minimizes $|\frac{1}{2\delta} - \lambda_i(P)|$, i.e., $j = \operatorname{argmin}_i(|\frac{1}{2\delta} - \lambda_i(P)|)$. The smallest eigenvalue of $P - \delta P^2$ satisfies the following:*

(i) if $\delta \in [0, 1]$, then $\lambda_{\min}(P - \delta P^2) = \min(m - \delta m^2, L - \delta L^2)$

(ii) if $\delta \in [-0.5, 0]$, then $\lambda_{\min}(P - \delta P^2) = m - \delta m^2$

(iii) if $\delta \in [-1, -0.5]$, then $\lambda_{\min}(P - \delta P^2) = \lambda_j(P) - \delta\lambda_j(P)^2$

Proof. From the spectral theorem it follows that the eigenvalues to $\lambda_i(P - \delta P^2) = \lambda_i(P) - \delta\lambda_i(P)^2$. So we need to find the $\lambda_i(P)$ that minimizes the function $\psi(\lambda) = \lambda - \delta\lambda^2$, where $\lambda_i(P) \in [-1, 1]$ for different δ.

For $\delta \in [0, 1]$, the function ψ is concave, and the minimum is found in either of the end points, so $\lambda_{\min}(P - \delta P^2) = \min(m - \delta m^2, L - \delta L^2)$. This shows *(i)*. If instead $\delta \in [-1, 0)$ the function ψ is convex. The unconstrained minimum is at $\frac{1}{2\delta}$. Then, since the level sets of ψ are symmetric around $\frac{1}{2\delta}$, the constrained minimum is the eigenvalue $\lambda_i(P)$ closest to $\frac{1}{2\delta}$. For $\delta \in [-0.5, 0)$ this is $\lambda_{\min}(P) = m$, and for $\delta \in [-1, -0.5]$ this is $\lambda_j(P)$. This concludes the proof. □

Lemma 7 *Suppose that P is a linear self-adjoint and nonexpansive operator with largest eigenvalue $\lambda_{\max}(P) = L$ and smallest eigenvalue $\lambda_{\min}(P) = m$, satisfying $-1 \leq m \leq L \leq 1$. Further suppose that $\delta \in [-1, 1]$ and let j be the index that minimizes $|\frac{1}{2\delta} + \lambda_i(P)|$, i.e., $j = \operatorname{argmin}_i(|\frac{1}{2\delta} + \lambda_i(P)|)$. The largest eigenvalue of $P + \delta P^2$ satisfies the following:*

(li) if $\delta \in [-0.5, 1]$, then $\lambda_{\max}(P + \delta P^2) = L + \delta L^2$

(lii) if $\delta \in [-1, -0.5]$, then $\lambda_{\max}(P + \delta P^2) = \lambda_j(P) + \delta\lambda_j(P)^2$

Proof. From the spectral theorem it follows that the eigenvalues to $\lambda_i(P + \delta P^2) = \lambda_i(P) + \delta\lambda_i(P)^2$. So we need to find the $\lambda_i(P)$ that maximizes the function $\psi(\lambda) = \lambda + \delta\lambda^2$, where $\lambda_i(P) \in [-1, 1]$ for different δ.

For $\delta \in [0, 1]$, the function ψ is convex, and the maximum is found in either of the end points. The function ψ is monotonically increasing on $[-1, 1]$, so the maximum is found at $L + \delta L^2$. For $\delta \in [-1, 0)$, the function ψ is concave. Its unconstrained maximum is at $\frac{1}{-2\delta}$. Since the level sets of ψ are symmetric around $\frac{1}{-2\delta}$, the constrained maximum is the eigenvalue closest to $\frac{1}{-2\delta}$. For $\delta \in [-0.5, 0)$, this is $\lambda_{\max}(P) = L$, and for $\delta \in [-1, -0.5]$ this is $\lambda_j(P)$. This concludes the proof. $\square$