Automatica 62 (2015) 1–10

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Metric selection in fast dual forward–backward splitting☆

Pontus Giselsson[a], Stephen Boyd[b]

[a] Department of Automatic Control, Lund University, Sweden

[b] Electrical Engineering Department, Stanford University, CA, United States

ARTICLE INFO

Article history:
Received 11 December 2013
Received in revised form
30 April 2015
Accepted 28 August 2015

Keywords:
First order optimization algorithms
Metric selection
Preconditioning
Model predictive control
Distributed optimization

ABSTRACT

The performance of fast forward–backward splitting, or equivalently fast proximal gradient methods, depends on the conditioning of the optimization problem data. This conditioning is related to a metric that is defined by the space on which the optimization problem is stated; selecting a space on which the optimization data is better conditioned improves the performance of the algorithm. In this paper, we propose several methods, with different computational complexity, to find a space on which the algorithm performs well. We evaluate the proposed metric selection procedures by comparing the performance to the case when the Euclidean space is used. For the most ill-conditioned problem we consider, the computational complexity is improved by two to three orders of magnitude. We also report comparable to superior performance compared to state-of-the-art optimization software.

1. Introduction

Fast gradient methods have been around since the early 1980s when the seminal paper Nesterov (1983) was published. The algorithm in Nesterov (1983) is applicable to unconstrained smooth optimization problems and has since been extended and generalized in various directions. In Nesterov (2003), new acceleration schemes were presented as well as fast gradient methods for constrained optimization. In Nesterov (2005), smoothing techniques for nonsmooth problems are presented. Fast proximal gradient methods, or equivalently fast forward–backward splitting methods, that solve composite convex optimization problems of the form

$$\text{minimize} \quad f(x) + g(x) \tag{1}$$

where f is required to be smooth, are proposed in Beck and Teboulle (2009) and Nesterov (2013). In Tseng (2008), generalizations and unifications of many fast forward–backward splitting methods are presented.

The smooth part of the composite objective function, f in (1), is in fast forward–backward splitting approximated by the r.h.s. of

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

where the norm and inner-product are given by the space on which the problem is defined. The condition that (2) holds for all x and y is referred to as f being β-smooth. Since the r.h.s. of the smoothness condition (2) is the only information the algorithm has about the smooth function, the smaller the gap in (2) (i.e. the better the r.h.s. of (2) approximates f), the better the performance of the algorithm is likely to be. In this paper, we show how to select a space (or metric, we will use these notions interchangeably since the metric defines the space) on which the fast forward–backward splitting method performs well, when solving the dual of strongly convex composite optimization problems. The spaces we consider are Euclidean spaces with inner product $\langle x, y\rangle = x^Ty$ and scaled norm $\|x\|_K = \sqrt{x^TKx}$, where K is a positive definite metric matrix. We show how to select the metric K such that the gap in (2) for the smooth part of the dual problem is minimized. Using this metric in the algorithm often leads to improved performance compared to using the Euclidean metric with $K = I$.

Recently, Patrinos and Bemporad (2014); Richter, Jones, and Morari (2013) proposed to use fast dual forward–backward splitting for embedded model predictive control. They apply fast forward–backward splitting with the standard Euclidean metric on two different dual problems. We show how these algorithms can be improved by choosing a metric that reduces the gap in (2). The performance improvement is confirmed by applying the methods to a pitch control problem in an AFTI-16 aircraft. This benchmark

☆ The material in this paper was partially presented at the 19th World Congress of the International Federation of Automatic Control, August 24–29, 2014, Cape Town, South Africa and the 53rd Conference on Decision and Control, December 15–17, 2014, Los Angeles, CA, USA. This paper was recommended for publication in revised form by Associate Editor Lalo Magni under the direction of Editor Ian R. Petersen.

E-mail addresses: pontusg@control.lth.se (P. Giselsson), boyd@stanford.edu (S. Boyd).

http://dx.doi.org/10.1016/j.automatica.2015.09.010
0005-1098/

has previously been studied in Bemporad, Casavola, and Mosca (1997) and Kapasouris, Athans, and Stein (1990) and is a challenging problem for first order methods since it is fairly ill-conditioned. We report computation time improvements by two to three orders of magnitude. Besides this, we also compare the performance to the ADMM-based (see Boyd, Parikh, Chu, Peleato, & Eckstein, 2011 for more on ADMM—the alternating direction method of multipliers) algorithm in Jerez et al. (2014) and O'Donoghue, Stathopoulos, and Boyd (2013). We also compare our algorithms, that are implemented in the MATLAB toolbox QPgen Giselsson (2014a), to several other toolboxes and software for embedded optimization, namely: DuQuad, see Necoara and Patrascu (2015), FiOrdOs, see Ullmann and Richter (2012), FORCES, see Domahidi, Zgraggen, Zeilinger, Morari, and Jones (2012), CVXGEN, see Mattingley and Boyd (2012), qpOASES, see Ferreau, Bock, and Diehl (2008) and the MPT Toolbox, see Herceg, Kvasnica, Jones, and Morari (2013). Finally, we also compare to the general commercial QP-solver MOSEK, see Mosek (2013). QPgen, with the proposed fast dual forward–backward splitting method, performs little to much better than the other methods on this example.

Fast dual forward–backward splitting can also be used for distributed optimization when the objective to be minimized is separable. In the context of gradient methods, this has been known since Benders (1962), Danzig and Wolfe (1961) and Everett (1963). Recently such approaches have been proposed for distributed model predictive control (DMPC) (Doan, Keviczky, & De Schutter, 2011; Giselsson, 2013; Giselsson, Doan, Keviczky, De Schutter, & Rantzer, 2013; Negenborn, 2007), and resource optimization over networks (Beck, Nedic, Ozdaglar, & Teboulle, 2014; Ghadimi, Shames, & Johansson, 2013; Necoara & Nedelcu, 2015). Often, centralized coordination is needed when selecting the step-size for the gradient-step. This is relaxed in Beck et al. (2014), where the authors noted that the smooth part of the dual problem consists of a sum of local functions. Each of these can compute its own step-size, share with its neighbors and sum, to get a fully distributed step-size selection. This procedure can be augmented by the results of this paper to select local metrics instead of step-sizes. This leads to more efficient algorithms which is confirmed by a numerical example which shows improvements of about one order of magnitude. We also compare the performance to the dual Newton conjugate gradient method in Kozma, Klintberg, Gros, and Diehl (2014), which is outperformed in our numerical example.

This paper unifies and extends the conference publications Giselsson (2014b,c) and Giselsson and Boyd (2014).

2. Notation and preliminaries

We denote by $\mathbb{R}$, $\mathbb{R}^n$, $\mathbb{R}^{m\times n}$, the sets of real numbers, column vectors, and matrices. We use notation $(x,y,z) := [x^T\ y^T\ z^T]^T$ for stacked real column vectors. We also use notation $\bar{\mathbb{R}} = \mathbb{R}\cup\{\infty\}$ for the extended real line. $\mathbb{S}^n \subseteq \mathbb{R}^{n\times n}$ is the set of symmetric matrices, and $\mathbb{S}^n_{++} \subseteq \mathbb{S}^n$, $[\mathbb{S}^n_{+}] \subseteq \mathbb{S}^n$, are the sets of positive [semi] definite matrices. We use Euclidean spaces with the standard inner product $\langle x,y\rangle = x^Ty$ and different norms. When using the induced norm $\|x\| = \sqrt{\langle x,x\rangle}$, we get the standard Euclidean space. We also consider spaces $\mathbb{E}_H$ with Euclidean inner product and scaled norm $\|x\|_H = \sqrt{\langle x,Hx\rangle}$, where $H \in \mathbb{S}^n_{++}$. The dual space to $\mathbb{E}_H$ is denoted by $\mathbb{E}_H^*$. The dual norm to $\|y\|_H$ is $\|y\|_H^* = \max_x\{\langle y,x\rangle_2 : \|x\|_H = 1\} = \|y\|_{H^{-1}}$, i.e., $\mathbb{E}_H^* = \mathbb{E}_{H^{-1}}$. Further, the class of closed, proper, and convex functions $f : \mathbb{E}_H \to \bar{\mathbb{R}}$ is denoted by $\Gamma_0(\mathbb{E}_H)$. The conjugate function $f^* : \mathbb{E}_H^* \to \bar{\mathbb{R}}$ to $f \in \Gamma_0(\mathbb{E}_H)$ is defined as $f^*(y) = \sup_x\{\langle y,x\rangle - f(x)\}$. The adjoint operator to a bounded linear operator $\mathcal{A} : \mathbb{E}_H \to \mathbb{E}_K$ is denoted by $\mathcal{A}^* : \mathbb{E}_K^* \to \mathbb{E}_H^*$ and is defined as the unique operator that satisfies $\langle \mathcal{A}x, y\rangle = \langle \mathcal{A}^*y, x\rangle$ for all $x \in \mathbb{E}_H$ and $y \in \mathbb{E}_K^*$. Since the ambient space for $\mathbb{E}_H$ is the standard Euclidean space, we often denote the matrix that corresponds to the operator $\mathcal{A} : \mathbb{E}_H \to \mathbb{E}_K$ by $A \in \mathbb{R}^{m\times n}$. We use notation $I_{\mathcal{X}}$ for the indicator function for the set $\mathcal{X}$, and $I_{g(x)\le 0}$ for the indicator function for the set $\mathcal{X} = \{x \mid g(x) \le 0\}$.

A function $f \in \Gamma_0(\mathbb{E}_H)$ is *β-strongly convex* (w.r.t. $\mathbb{E}_H$) if $f - \frac{\beta}{2}\|\cdot\|_H^2$ is convex. A function $f \in \Gamma_0(\mathbb{E}_H)$ is *β-smooth* (w.r.t. $\mathbb{E}_H$) if it is differentiable and $\frac{\beta}{2}\|\cdot\|_H^2 - f$ is convex. An equivalent characterization of β-smoothness w.r.t. $\mathbb{E}_H$ is that

$$f(x) \le f(y) + \langle\nabla f(y), x-y\rangle + \frac{\beta}{2}\|x-y\|_H^2 \tag{3}$$

holds for all $x, y \in \mathbb{E}_H$. As seen in the following proposition, these notions are related through the conjugate function.

Proposition 1. *Suppose that $f \in \Gamma_0(\mathbb{E}_H)$. Then the following are equivalent:*

(i) *f is β-strongly convex (w.r.t. $\mathbb{E}_H$).*
(ii) *f^* is $\frac{1}{\beta}$-smooth (w.r.t. $\mathbb{E}_H^* = \mathbb{E}_{H^{-1}}$).*

A proof to this can be found, e.g., in Zalinescu (2002, Proposition 3.5.3).

3. Problem formulation

We consider optimization problems of the form

$$\begin{array}{ll} \text{minimize} & f(x) + g(y) \\ \text{subject to} & \mathcal{A}x = y \end{array} \tag{4}$$

and assume that the following assumption holds throughout the paper:

Assumption 2. (a) The extended valued function $f \in \Gamma_0(\mathbb{E}_H)$ is 1-strongly convex (w.r.t. $\mathbb{E}_H$).
(b) The extended valued function $g \in \Gamma_0(\mathbb{E}_K)$.
(c) $\mathcal{A} : \mathbb{E}_H \to \mathbb{E}_K$ is a bounded linear operator.

Remark 3. A function that satisfies Assumption 2(a) is $f(x) = \frac{1}{2}x^THx + \hat{f}$ where $H \in \mathbb{S}^n_{++}$ and $\hat{f} \in \Gamma_0(\mathbb{E}_H)$. Since $\hat{f}$ (and g) are allowed to be extended valued, they can, e.g., be indicator functions for nonempty, closed, and convex constraint sets. Further, the operator $\mathcal{A} : \mathbb{E}_H \to \mathbb{E}_K$ has an associated matrix $A : \mathbb{R}^n \to \mathbb{R}^m$ that satisfies $\mathcal{A}x = Ax$ for all $x \in \mathbb{R}^n$.

To arrive at the dual problem, we introduce Lagrange multipliers $\mu \in \mathbb{E}_{K^{-1}}$, to get Lagrangian

$$L(x,y,\mu) = f(x) + g(y) + \langle \mathcal{A}x - y, \mu\rangle.$$

By minimizing the Lagrangian over x, and y, we get

$$\begin{aligned} \inf_{x,y} L(x,y,\mu) &= \inf_x\left\{\langle \mathcal{A}^*\mu, x\rangle + f(x)\right\} + \inf_y\{\langle -y,\mu\rangle + g(y)\} \\ &= -\sup_x\left\{\langle -\mathcal{A}^*\mu, x\rangle - f(x)\right\} - \sup_y\{\langle \mu, y\rangle - g(y)\} \\ &= -f^*(-\mathcal{A}^*\mu) - g^*(\mu). \end{aligned}$$

Negating this, we get the negated dual problem to (4) (see, e.g., Rockafellar, 1970, §31 for more details):

$$\text{minimize } d(\mu) + g^*(\mu) \tag{5}$$

where

$$d(\mu) := f^*(-\mathcal{A}^*\mu). \tag{6}$$

Note that $d, g^* \in \Gamma_0(\mathbb{E}_{K^{-1}})$. The efficiency of solving this dual problem using fast forward–backward splitting is highly dependent on which metric that is used. This paper is about choosing metrics to make the algorithm perform well.