

GRADIENT METHODS WITH MEMORY

Yurii Nesterov, Mihai Florea

REPRINT | 3230

CORE

Voie du Roman Pays 34, L1.03.01



B-1348 Louvain-la-Neuve

Tel (32 10) 47 43 04

Email: lidam-library@uclouvain.be

<https://uclouvain.be/en/research-institutes/lidam/core/core-reprints.html>

Gradient methods with memory

Yurii Nesterov ^a and Mihai I. Florea ^b

^aCenter for Operations Research and Econometrics (CORE), Catholic University of Louvain (UCL), Louvain-la-Neuve Belgium; ^bDepartment of Mathematical Engineering (INMA), Catholic University of Louvain (UCL), Louvain-la-Neuve, Belgium

ABSTRACT

In this paper, we consider gradient methods for minimizing smooth convex functions, which employ the information obtained at the previous iterations in order to accelerate the convergence towards the optimal solution. This information is used in the form of a piecewise linear model of the objective function, which provides us with much better prediction abilities as compared with the standard linear model. To the best of our knowledge, this approach was never really applied in Convex Minimization to differentiable functions in view of the high complexity of the corresponding auxiliary problems. However, we show that all necessary computations can be done very efficiently. Consequently, we get new optimization methods, which are better than the usual Gradient Methods both in the number of oracle calls and in the computational time. Our theoretical conclusions are confirmed by preliminary computational experiments.

ARTICLE HISTORY

Received 22 April 2020
Accepted 30 November 2020

KEYWORDS

Convex optimization;
gradient methods; relative
smoothness; rate of
convergence; piece-wise
linear model

1. Introduction

1.1. Motivation

First-order gradient methods for minimizing smooth convex functions generate a sequence of test points based on the information obtained from the oracle: the function values and the gradients. Most methods either use the information from the last test point or accumulate it in the form of an aggregated linear function (see, for example, Chapter 2 in [7]). This approach is very different from the technique used in Nonsmooth Optimization, where the piece-wise linear model of the objective is a standard and powerful tool. It is enough to mention the Bundle Method, the Level Method, cutting plane schemes, etc. The reason for this situation is quite clear. The presence of piece-wise linear models in the auxiliary problems, which we need to solve at each iteration of the method, usually significantly increases the complexity of the corresponding computations. This is acceptable in Nonsmooth Optimization, which has the reputation of a difficult field. By contrast, Smooth Optimization admits very simple and elegant schemes, with a very small computational cost of each iteration, preventing us from introducing there such a heavy machinery.

CONTACT Mihai I. Florea  mihai.florea@uclouvain.be

© 2021 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group
This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way.

After the preparation of this manuscript, we became aware of a highly specialized attempt in [2], using quadratic lower bounds instead of linear ones. Although the results presented there seem promising, the study limits itself to studying smooth unconstrained problems with strongly convex objectives. The study also states that an extension of the results to a wider context is a difficult open problem.

The main goal of this paper is the demonstration that the above situation is not as clear as it looks like. We will show that the Gradient Method,¹ equipped with a piece-wise linear model of the objective function, has much more chances to accelerate on particular optimization problems. At the same time, it appears that the corresponding auxiliary optimization problems can be easily solved by an appropriate version of the Frank–Wolfe algorithm. All our claims are supported by a complexity analysis. In the end, we present preliminary computational results, which show that very often the new schemes have much better computational time.

1.2. Contents

In Section 2, we analyse the Gradient Method with Memory as applied to the composite form of smooth convex optimization problems [5]. In order to measure the level of smoothness of our objective function, we introduce the *relative smoothness* condition [1,4], based on an arbitrary strictly convex distance function. The main novelty here is the piece-wise linear model of the objective function, formed around the current test point. We analyse the corresponding auxiliary optimization problem and propose a condition on its approximate solution which does not destroy the rate of convergence of the algorithm. In Section 3, we analyse the complexity of solving the auxiliary optimization problem using the Frank–Wolfe algorithm. More precisely, we consider the anti-dual² of the auxiliary problem.

In this section, we restrict ourselves to strongly convex distance functions. We show that our auxiliary problem can be easily solved by the Frank–Wolfe method. Its complexity is proportional to the maximal squared norm of the gradient in the current model of the objective divided by the desired accuracy.

In Section 4, we specify our complexity results for the Euclidean setup, when all distances are measured by a Euclidean norm. We show that, for some strategies of updating the piece-wise linear model, the complexity of the auxiliary computations is very low.

Finally, in Section 5 we present preliminary computational results. We compare the usual Gradient Method with two gradient methods with memory, which use different strategies for updating the piece-wise model of the objective function. Our conclusion is that the new schemes are always better, both in terms of the calls of oracle and in the total computational time.

1.3. Notation and generalities

In what follows, denote by \mathbb{E} a finite-dimensional real vector space and by \mathbb{E}^* its dual space, the space of linear functions on \mathbb{E} . The value of function $s \in \mathbb{E}^*$ at point $x \in \mathbb{E}$ is denoted by $\langle s, x \rangle$. Let us fix some arbitrary (possibly non-Euclidean) norm $\| \cdot \|$ on the space \mathbb{E} and

define the dual norm $\|\cdot\|_*$ on \mathbb{E}^* in the standard way:

$$\|s\|_* \stackrel{\text{def}}{=} \sup_{h \in \mathbb{E}} \{\langle s, h \rangle : \|h\| \leq 1\}.$$

Let us choose a simple closed convex prox-function $d(\cdot)$, which is differentiable at the interior of its domain.³ This function must be strictly convex:

$$d(y) > d(x) + \langle \nabla d(x), y - x \rangle, \quad x \in \text{int}(\text{dom } d), y \in \text{dom } d, x \neq y. \quad (1)$$

Using this function, we can define the *Bregman distance* between two points x and y :

$$\beta_d(x, y) = d(y) - d(x) - \langle \nabla d(x), y - x \rangle, \quad x \in \text{int}(\text{dom } d), y \in \text{dom } d. \quad (2)$$

Clearly, $\beta_d(x, y) \stackrel{(1)}{>} 0$ for $x \neq y$ and $\beta_d(x, x) = 0$.

We will use Bregman distances for measuring the level of *relative smoothness* of convex functions (see [4]). Namely, for a differentiable closed convex function f with open $\text{dom } f \subseteq \text{dom } d$ we define two constants, $L_d(f) \geq \mu_d(f) \geq 0$, such that

$$\begin{aligned} f(y) - f(x) - \langle \nabla f(x), y - x \rangle &\geq \mu_d(f) \beta_d(x, y), \\ f(y) - f(x) - \langle \nabla f(x), y - x \rangle &\leq L_d(f) \beta_d(x, y), \end{aligned} \quad x, y \in \text{dom } f, \quad (3)$$

See [1] and [4] for definitions, motivations, and examples.

2. Gradient method with memory

In this paper, we are solving the following composite minimization problem:

$$\min_{x \in \text{dom } \psi} \left\{ F(x) \equiv f(x) + \psi(x) \right\}, \quad (4)$$

where function f satisfies the relative smoothness condition (3), possibly with $\mu_d(f) = 0$. The function $\psi : \mathbb{E} \rightarrow \mathbb{R} \cup \{+\infty\}$ is a proper closed convex function with $\text{dom } \psi \subseteq \text{dom } f$ and $\text{int}(\text{dom } \psi)$ non-empty. The function ψ is simple (in the sense of satisfying Assumptions 2.1 and 2.2, stated in the sequel) but it does not have to be differentiable or even continuous. For instance, ψ can incorporate the indicator function of the feasible set. We assume that a solution $x^* \in \text{dom } \psi$ of problem (4) does exist, denoting $F^* = F(x^*)$.

The simplest method for solving the problem (4) is the usual Gradient Method:

Choose $x_0 \in \text{int}(\text{dom } \psi)$. For $k \geq 0$, iterate:

$$x_{k+1} = \arg \min_{y \in \text{dom } \psi} \left\{ f(x_k) + \langle \nabla f(x_k), y - x_k \rangle + \psi(y) + L \beta_d(x_k, y) \right\}. \quad (5)$$

The constant L in this method has to be big enough in order to ensure

$$f(x_{k+1}) \leq f(x_k) + \langle \nabla f(x_k), x_{k+1} - x_k \rangle + L \beta_d(x_k, x_{k+1}).$$

In view of (3), this is definitely true for $L \geq L_d(f)$. However, we are interested in choosing L as small as possible since this would significantly increase the rate of convergence of the scheme.

Method (5) is based on the simplest *linear model* of function $f(\cdot)$ around the point x_k . In our paper, we suggest to replace it by a piece-wise linear model, defined by the information collected at other test points.

Namely, for each $k \geq 0$ define a *discrete set* \mathcal{Z}_k of m_k feasible points ($m_k \geq 1$):

$$\mathcal{Z}_k = \{z_i \in \text{dom } \psi, i = 1, \dots, m_k\}.$$

Then we can use a more sophisticated model of the smooth part of the objective function,

$$f(y) \geq \ell_k(y) \stackrel{\text{def}}{=} \max_{z_i \in \mathcal{Z}_k} \{f(z_i) + \langle \nabla f(z_i), y - z_i \rangle\}, \quad y \in \text{dom } \psi. \quad (6)$$

This model is always better than the initial linear model provided that

$$x_k \in \mathcal{Z}_k. \quad (7)$$

In what follows, we always assume that this condition is satisfied.

Thus, we come to the following natural generalization of the method (5), which we call the *Gradient Method with Memory* (GMM):

<p>Choose $x_0 \in \text{int}(\text{dom } \psi)$.</p> <p>For $k \geq 0$, iterate:</p> $x_{k+1} = \arg \min_{y \in \text{dom } \psi} \{ \ell_k(y) + \psi(y) + L\beta_d(x_k, y) \}.$	(8)
---	-----

Remark 2.1: Note that for any $x \in \text{dom } f$ we have

$$\begin{aligned} f(x_k) + \langle \nabla f(x_k), x - x_k \rangle + L\beta_d(x_k, x) &\stackrel{(7)}{\leq} \ell_k(x) + L\beta_d(x_k, x) \\ &\stackrel{(6)}{\leq} f(x) + L\beta_d(x_k, x) \stackrel{(3)}{\leq} f(x_k) + \langle \nabla f(x_k), x - x_k \rangle + (L + L_d(f))\beta_d(x_k, x). \end{aligned}$$

Therefore, we can count on a better convergence of method (8) only if we will be able to choose the parameter L significantly smaller than $L_d(f)$.

At each iteration of method (8), we need to solve a non-trivial auxiliary minimization problem. Therefore, the practical efficiency of this method crucially depends on the complexity of this computation. In what follows, we suggest to solve this problem approximately using a special method for its dual problem.

Let us start by presenting the corresponding technique. For the sake of notation, we omit the index of the iteration. Thus, our auxiliary problem is as follows:

$$\min_{y \in \text{dom } \psi} \max_{\lambda \in \Delta_m} \left\{ \sum_{i=1}^m \lambda^{(i)} [f_i + \langle g_i, y - z_i \rangle] + \psi(y) + L\beta_d(\bar{x}, y) \right\},$$

where $f_i = f(z_i)$, $g_i = \nabla f(z_i)$, $i = 1, \dots, m$, and Δ_m is the standard simplex in \mathbb{R}^m . Introducing now the vector $f_* \in \mathbb{R}^m$ with coordinates

$$f_*^{(i)} = \langle g_i, z_i \rangle - f_i, \quad i = 1, \dots, m, \quad (9)$$

we get the following representation of our problem:

$$\min_{y \in \text{dom } \psi} \max_{\lambda \in \Delta_m} \left\{ \langle \lambda, G^T y - f_* \rangle + \psi(y) + L\beta_d(\bar{x}, y) \right\}, \quad (10)$$

where $G = (g_1, \dots, g_m) \in \mathbb{E}^* \times \mathbb{R}^m$. Note that the pay-off function in this saddle point problem can be written as follows:

$$\begin{aligned} & \langle \lambda, G^T y - f_* \rangle + \psi(y) + L\beta_d(\bar{x}, y) \\ &= \langle \lambda, G^T y - f_* \rangle + \psi(y) + L[d(y) - d(\bar{x}) - \langle \nabla d(\bar{x}), y - \bar{x} \rangle] \\ &= Ld(y) - \langle L\nabla d(\bar{x}) - G\lambda, y \rangle + \psi(y) - \langle \lambda, f_* \rangle + L[\langle \nabla d(\bar{x}), \bar{x} \rangle - d(\bar{x})]. \end{aligned}$$

Hence, we need to introduce the following dual function:

$$\Phi_L(s) = \max_{y \in \text{dom } \psi} \{ \langle s, y \rangle - Ld(y) - \psi(y) \}, \quad s \in \mathbb{E}^*. \quad (11)$$

Our main joint assumption on functions $d(\cdot)$ and $\psi(\cdot)$ is as follows.

Assumption 2.1: For any $L > 0$, function $\Phi_L(s)$ is defined at any $s \in \mathbb{E}^*$.

This can be ensured, for example, by the strong convexity of function $d(\cdot)$, or by the boundedness of $\text{dom } \psi$, or in many other ways.

Since the objective function in the definition (11) is strictly concave, its solution

$$y_L^*(s) = \arg \max_{y \in \text{dom } \psi} \{ \langle s, y \rangle - Ld(y) - \psi(y) \}$$

is uniquely defined for any $s \in \mathbb{E}^*$. Moreover, function $\Phi_L(\cdot)$ is differentiable and

$$\nabla \Phi_L(s) = y_L^*(s), \quad s \in \mathbb{E}^*. \quad (12)$$

Now we can write down the problem anti-dual to (10)

$$\xi_L^* \stackrel{\text{def}}{=} \min_{\lambda \in \Delta_m} \left\{ \xi_L(\lambda) \stackrel{\text{def}}{=} \Phi_L(L\nabla d(\bar{x}) - G\lambda) + \langle \lambda, f_* \rangle + \alpha \right\}, \quad (13)$$

where $\alpha = L[d(\bar{x}) - \langle \nabla d(\bar{x}), \bar{x} \rangle]$. This is a convex optimization problem with a differentiable objective function. Our second main assumption is as follows.

Assumption 2.2: Function $\Phi_L(\cdot)$ in the problem (13) is easily computable.

We will discuss the reasonable strategies for finding an approximate solution to problem (13) in Section 3. At this moment, it is enough to assume that we are able to compute a point $\bar{\lambda} = \bar{\lambda}(\bar{x}, \mathcal{Z}, L)$ such that

$$\langle \bar{\lambda} - \lambda, \nabla \xi_L(\bar{\lambda}) \rangle \leq \delta, \quad \lambda \in \Delta_m, \quad (14)$$

where $\delta \geq 0$ is some tolerance parameter. Clearly, if $\delta = 0$, then $\bar{\lambda}$ is the optimal solution to the problem (13). Note that condition (14) ensures also a small functional gap:

$$\xi_L(\bar{\lambda}) - \xi_L^* = \max_{\lambda \in \Delta_m} [\xi_L(\bar{\lambda}) - \xi_L(\lambda)] \leq \max_{\lambda \in \Delta_m} \langle \bar{\lambda} - \lambda, \nabla \xi_L(\bar{\lambda}) \rangle \stackrel{(14)}{\leq} \delta. \quad (15)$$

Condition (14) immediately leads to the following result.

Lemma 2.1: Let $\bar{\lambda} \in \Delta_m$ satisfy condition (14). Then for $\bar{s} = L\nabla d(\bar{x}) - G\bar{\lambda}$ we have

$$\sum_{i=1}^m \bar{\lambda}^{(i)} [f_i + \langle g_i, y_L^*(\bar{s}) - z_i \rangle] \geq \max_{1 \leq i \leq m} [f_i + \langle g_i, y_L^*(\bar{s}) - z_i \rangle] - \delta. \quad (16)$$

Proof: Indeed, $\nabla \xi_L(\bar{\lambda}) = f_* - G^T y_L^*(\bar{s})$. Thus, inequality (14) can be rewritten as follows:

$$\langle \bar{\lambda}, G^T y_L^*(\bar{s}) - f_* \rangle \geq \langle \lambda, G^T y_L^*(\bar{s}) - f_* \rangle - \delta, \quad \lambda \in \Delta_m.$$

It remains to note that

$$(G^T y_L^*(\bar{s}) - f_*)^{(i)} = f_i + \langle g_i, y_L^*(\bar{s}) - z_i \rangle, \quad i = 1, \dots, m.$$

■

Now we are able to analyse one iteration of the inexact version of method (8).

Input: Point $\bar{x} \in \text{int}(\text{dom } \psi)$, set of test points \mathcal{Z} , containing \bar{x} , constant $L > 0$, and tolerance $\delta \geq 0$.	(17)
Iteration: Using the input data, form the optimization problem (13) and compute its approximate solution $\bar{\lambda}$ satisfying condition (14).	
Output: Points $\bar{s} = L\nabla d(\bar{x}) - G\bar{\lambda}$ and $x_+ = y_L^*(\bar{s})$.	

Theorem 2.1: (1) Let point x_+ be generated by one iteration (17) of the Inexact Gradient Method with Memory (IGMM), and let $L \geq L_d(f)$. Then for any $y \in \text{dom } \psi$

$$\beta(x_+, y) \leq \beta(\bar{x}, y) + \frac{1}{L} [F(y) - F(x_+) + \delta]; \quad (18)$$

(2) For every $y \in \text{dom } \psi$ that satisfies $\beta_d(z_i, y) \geq \beta_d(\bar{x}, y)$, $i = 1, \dots, m$ we have

$$\beta(x_+, y) \leq \left(1 - \frac{1}{L} \mu_d(f)\right) \beta(\bar{x}, y) + \frac{1}{L} [F(y) - F(x_+) + \delta]. \quad (19)$$

Proof: Note that

$$\begin{aligned} & \beta(x_+, y) - \beta(\bar{x}, y) \\ & \stackrel{(2)}{=} d(y) - d(x_+) - \langle \nabla d(x_+), y - x_+ \rangle - d(y) + d(\bar{x}) + \langle \nabla d(\bar{x}), y - \bar{x} \rangle \\ & \stackrel{(2)}{=} \langle \nabla d(\bar{x}) - \nabla d(x_+), y - x_+ \rangle - \beta(\bar{x}, x_+). \end{aligned}$$

The point x_+ is defined as

$$x_+ = \arg \max_{x \in \text{dom } \psi} \{ \langle L \nabla d(\bar{x}) - G\bar{\lambda}, x \rangle - Ld(x) - \psi(x) \}.$$

The first-order optimality condition for x_+ at point y can be written in the following form:

$$\psi(x_+) \leq \psi(y) + \langle G\bar{\lambda} + L(\nabla d(x_+) - \nabla d(\bar{x})), y - x_+ \rangle.$$

Hence,

$$\beta(x_+, y) - \beta(\bar{x}, y) \leq \frac{1}{L} [\psi(y) - \psi(x_+) + \langle G\bar{\lambda}, y - x_+ \rangle] - \beta(\bar{x}, x_+).$$

Note that

$$\begin{aligned} \langle G\bar{\lambda}, y - x_+ \rangle &= \langle \bar{\lambda}, G^T(y - x_+) \rangle = \sum_{i=1}^m \bar{\lambda}^{(i)} \langle g_i, y - x_+ \rangle \\ &= \sum_{i=1}^m \bar{\lambda}^{(i)} [\langle g_i, z_i - x_+ \rangle + \langle g_i, y - z_i \rangle] \\ &\stackrel{(3)}{\leq} \sum_{i=1}^m \bar{\lambda}^{(i)} [\langle g_i, z_i - x_+ \rangle + f(y) - f_i - \mu_d(f) \beta_d(z_i, y)] \\ &= f(y) - \sum_{i=1}^m \bar{\lambda}^{(i)} [f_i + \langle g_i, x_+ - z_i \rangle] - \mu_d(f) \sum_{i=1}^m \bar{\lambda}^{(i)} \beta_d(z_i, y). \end{aligned}$$

Under the conditions of Item 1, we drop the last term in the above inequality and by Lemma 2.1 obtain the following:

$$\begin{aligned} \beta(x_+, y) - \beta(\bar{x}, y) &\leq \frac{1}{L} \left[F(y) - \sum_{i=1}^m \bar{\lambda}^{(i)} [f_i + \langle g_i, x_+ - z_i \rangle] - \psi(x_+) - L\beta(\bar{x}, x_+) \right] \\ &\leq \frac{1}{L} \left[F(y) - \max_{1 \leq i \leq m} [f_i + \langle g_i, x_+ - z_i \rangle] - \psi(x_+) - L\beta(\bar{x}, x_+) + \delta \right]. \end{aligned}$$

Under the conditions of Item 2, by the same reasoning we get

$$\begin{aligned} \beta(x_+, y) - \beta(\bar{x}, y) &\leq \frac{1}{L} \left[F(y) - \mu_d(f) \beta_d(\bar{x}, y) - \sum_{i=1}^m \bar{\lambda}^{(i)} [f_i + \langle g_i, x_+ - z_i \rangle] - \psi(x_+) - L\beta(\bar{x}, x_+) \right] \\ &\leq \frac{1}{L} \left[F(y) - \mu_d(f) \beta_d(\bar{x}, y) - \max_{1 \leq i \leq m} [f_i + \langle g_i, x_+ - z_i \rangle] - \psi(x_+) - L\beta(\bar{x}, x_+) + \delta \right]. \end{aligned}$$

In both cases, since $\bar{x} \in \mathcal{Z}$, we have

$$\max_{1 \leq i \leq m} [f_i + \langle g_i, x_+ - z_i \rangle] + L\beta(\bar{x}, x_+) \geq f(\bar{x}) + \langle \nabla f(\bar{x}), x_+ - \bar{x} \rangle + L\beta(\bar{x}, x_+)$$

$$\stackrel{(3)}{\geq} f(x_+).$$

Thus, we obtain inequalities (18) and (19). ■

Remark 2.2: In the end of the proof, we have seen that the statement of Theorem 2.1 remains valid if the condition $L \geq L_d(f)$ is replaced by the following:

$$\max_{1 \leq i \leq m} [f_i + \langle g_i, x_+ - z_i \rangle] + L\beta(\bar{x}, x_+) \geq f(x_+). \quad (20)$$

Denote the output of the iteration (17) by $x_{\delta,L}(\bar{x}, \mathcal{Z})$. Then we can define the following Inexact Gradient Method with Memory.

<p>Choose $x_0 \in \text{int}(\text{dom } \psi)$, $\delta \geq 0$ and $L > 0$. For $k \geq 0$, iterate: (1) Choose the set \mathcal{Z}_k containing x_k. (2) Compute $x_{k+1} = x_{\delta,L}(x_k, \mathcal{Z}_k)$.</p>	(21)
--	------

Let us describe the rate of convergence of this process.

Theorem 2.2: Let sequence $\{x_k\}_{k \geq 1}$ be generated by IGMM (21) with $L \geq L_d(f)$. Then, for any $T \geq 1$ and $y \in \text{dom } \psi$ we have

$$\frac{1}{T} \sum_{k=1}^T F(x_k) \leq F(y) + \frac{L}{T} \beta_d(x_0, y) + \delta. \quad (22)$$

Proof: Indeed, in view of inequality (18), we have

$$\beta_d(x_{k+1}, y) \leq \beta_d(x_k, y) + \frac{1}{L} [F(y) - F(x_{k+1}) + \delta], \quad k \geq 0.$$

Summing up these inequalities for $k = 0, \dots, T-1$, we get inequality (22). ■

In the above result, the only restriction for the sets \mathcal{Z}_k is the inclusion (7). If we apply a more accurate strategy of choosing \mathcal{Z}_k , we can get for this scheme a finer estimate of its rate of convergence.

Theorem 2.3: Let sequence $\{x_k\}_{k \geq 1}$ be generated by IGMM (21) with $L \geq L_d(f)$. Assume that besides the condition (7), the sets \mathcal{Z}_k satisfy also the following condition:

$$\mathcal{Z}_k \subseteq \{x_0, \dots, x_k\}, \quad k \geq 0. \quad (23)$$

Suppose that $\mu_d(f) > 0$ and for all k , $1 \leq k \leq T$, we have

$$F(x_k) - F^* \geq \delta. \quad (24)$$

Then for $\Delta_T^* = \min_{1 \leq k \leq T} [F(x_k) - F^*]$ we get the following rate of convergence

$$\Delta_T^* \leq \delta + \frac{(1-\gamma)^T \mu_d(f)}{1 - (1-\gamma)^T} \beta_d(x_0, x^*)$$

$$\leq \delta + \frac{\mu_d(f)}{e^{\gamma T} - 1} \beta_d(x_0, x^*) \leq \delta + \frac{L}{T} \beta_d(x_0, x^*), \tag{25}$$

where $\gamma = \frac{1}{L} \mu_d(f)$.

Proof: In view of assumption (24) and inequality (18), we have

$$\beta_d(x_{k+1}, x^*) \leq \beta_d(x_k, x^*), \quad 0 \leq k \leq T - 1.$$

Therefore, in view of inequality (19), for $r_k \stackrel{\text{def}}{=} \beta_d(x_k, x^*)$ and all $k = 0, \dots, T - 1$ we have

$$r_{k+1} \leq (1 - \gamma)r_k - \frac{1}{L}[F(x_{k+1}) - F^* - \delta] \leq (1 - \gamma)r_k - \frac{1}{L}[\Delta_T^* - \delta].$$

Applying this inequality recursively, we get

$$\frac{1}{L}[\Delta_T^* - \delta] \frac{1 - (1 - \gamma)^T}{1 - (1 - \gamma)} \leq (1 - \gamma)^T r_0.$$

This can be rewritten as

$$\Delta_T^* \leq \delta + \frac{L\gamma(1 - \gamma)^T}{1 - (1 - \gamma)^T} \beta_d(x_0, x^*) = \delta + \frac{(1 - \gamma)^T \mu_d(f)}{1 - (1 - \gamma)^T} \beta_d(x_0, x^*).$$

■

Note that the rate of convergence given by inequality (25) is continuous as $\mu_d(f) \rightarrow 0$.

As we have mentioned in Remark 2.1, it is important to adjust the value of the constant L during the minimization process. Therefore we present an adaptive version of the method (21).

Choose $x_0 \in \text{int}(\text{dom } \psi)$, $\delta \geq 0$, and some $L_0 \in (0, L_d(f)]$.

For $k \geq 0$, **iterate:**

(1) Choose the set \mathcal{Z}_k containing x_k .

(2) Find the smallest integer $i_k \geq 0$ such that for the point $x_k^+ = x_{\delta, 2^{i_k} L_k}(x_k, \mathcal{Z}_k)$ we have (26)

$$f(x_k^+) \stackrel{(6)}{\leq} \ell_k(x_k^+) + L_k \beta_d(x_k, x_k^+).$$

(3) Set $x_{k+1} = x_k^+$ and $L_{k+1} = 2^{i_k - 1} L_k$.

The rate of convergence of this algorithm can be established exactly in the same way as the one of method (21). The main fact is that during the minimization process we always have

$$L_k \leq 2L_d(f), \quad k \geq 0.$$

Therefore, for any $y \in \text{dom } \psi$ we will have

$$\beta_d(x_{k+1}, y) \leq \beta_d(x_k, y) + \frac{1}{2L_d(f)} [F(y) - F(x_{k+1}) + \delta]$$

with corresponding consequences for the rate of convergence. At the same time, the average number of oracle calls at each iteration of this method is bounded by two (see [7] for justification details).

3. Getting an approximate solution of the anti-dual problem

The complexity of solving the auxiliary problem (13) crucially depends on the properties of the prox-function $d(\cdot)$. In the previous section, we assumed its strict convexity and the solvability of problem (11) (see Assumption 2.1). It is time now to make a stronger assumption, which ensures these two properties.

Assumption 3.1: Function $d(\cdot)$ is differentiable in the interior of its domain and strongly convex with convexity parameter one:

$$d(y) \geq d(x) + \langle \nabla d(x), y - x \rangle + \frac{1}{2} \|y - x\|^2, \quad x \in \text{int}(\text{dom } d), \quad y \in \text{dom } d. \quad (27)$$

Clearly, for all $x, y \in \text{int}(\text{dom } d)$ we have

$$\beta_d(x, y) \geq \frac{1}{2} \|y - x\|^2. \quad (28)$$

The main consequence of Assumption 3.1 is the Lipschitz continuity of the gradient of function $\Phi(\cdot)$. Since usually this fact is proved for a function $\psi(\cdot)$ being an indicator function of a closed convex set, we provide it with a simple proof.

Lemma 3.1: Let function $d(\cdot)$ satisfy Assumption 3.1. Then the gradient $\nabla \Phi_L(s) = y_L^*(s)$, $s \in \mathbb{E}^*$, is Lipschitz continuous:

$$\|\nabla \Phi_L(s_1) - \nabla \Phi_L(s_2)\| \leq \frac{1}{L} \|s_1 - s_2\|_*, \quad s_1, s_2 \in \mathbb{E}^*. \quad (29)$$

Proof: Let us write down the first-order optimality conditions for the optimization problems defining the points $y_1 \stackrel{\text{def}}{=} y_L^*(s_1)$ and $y_2 \stackrel{\text{def}}{=} y_L^*(s_2)$:

$$\begin{aligned} \langle s_1 - L[\nabla d(y_1) - \nabla d(\bar{x})], y - y_1 \rangle - \psi(y) &\leq -\psi(y_1), \quad y \in \text{dom } \psi, \\ \langle s_2 - L[\nabla d(y_2) - \nabla d(\bar{x})], y - y_2 \rangle - \psi(y) &\leq -\psi(y_2), \quad y \in \text{dom } \psi. \end{aligned}$$

Taking in the first inequality $y = y_2$ and $y = y_1$ in the second one, and adding the results, we obtain

$$\langle s_1 - s_2, y_2 - y_1 \rangle \leq L \langle \nabla d(y_1) - \nabla d(y_2), y_2 - y_1 \rangle.$$

Thus,

$$\langle s_1 - s_2, y_1 - y_2 \rangle \geq L \langle \nabla d(y_2) - \nabla d(y_1), y_2 - y_1 \rangle \stackrel{(27)}{\geq} L \|y_1 - y_2\|^2.$$

Therefore, by the Cauchy-Schwartz inequality, we get

$$\|y_L^*(s_1) - y_L^*(s_2)\| \leq \frac{1}{L} \|s_1 - s_2\|_*.$$

■

Thus, in this section, our main problem of interest is as follows:

$$\xi_L^* = \min_{\lambda \in \Delta_m} \left\{ \xi_L(\lambda) \stackrel{\text{def}}{=} \Phi_L(L\nabla d(\bar{x}) - G\lambda) + \langle \lambda, f_* \rangle + \alpha \right\}, \tag{30}$$

where $\alpha = L[d(\bar{x}) - \langle \nabla d(\bar{x}), \bar{x} \rangle]$. This is a convex optimization problem over a simplex, where the objective function has a Lipschitz-continuous gradient.

The most natural algorithm for solving the problem (13) is the *Frank–Wolfe algorithm* [3] (or *Conditional Gradients Method*). For our problem, it looks as follows.

Set $\lambda_0 = \frac{1}{m} \bar{e}_m$.

For $k \geq 0$ **iterate:**

1. Compute the gradient $\nabla \xi_L(\lambda_k)$.
2. Compute $i_k = \arg \min_{1 \leq i \leq m} \nabla_i \xi_L(\lambda_k)$.
3. Set $\lambda_{k+1} = \frac{k}{k+2} \lambda_k + \frac{2}{k+2} e_{i_k}$.

(31)

In this scheme, $\bar{e}_m \in \mathbb{R}^m$ is the vector of all ones, and e_i is i th coordinate vector in \mathbb{R}^m .

In order to estimate the rate of convergence of this method, we introduce the following accuracy measure:

$$\delta_L(\bar{\lambda}) = \max_{\lambda \in \Delta_m} \langle \nabla \xi_L(\bar{\lambda}), \bar{\lambda} - \lambda \rangle.$$

For the sequence $\{\lambda_k\}_{k \geq 0}$ generated by the method (31), denote

$$\delta_L^*(T) = \min_{0 \leq k \leq T} \delta_L(\lambda_k), \quad T \geq 0.$$

For estimating the rate of convergence of method (31), we need to choose an appropriate norm in \mathbb{R}^m . Since the feasible set of the problem (13) is the standard simplex, it is reasonable to use the ℓ_1 -norm:

$$\|\lambda\|_1 = \sum_{i=1}^m |\lambda^{(i)}|, \quad \lambda \in \mathbb{R}^m.$$

Then, for measuring the gradients of function $\xi_L(\cdot)$, we can use the ℓ_∞ -norm:

$$\|\lambda\|_\infty = \max_{1 \leq i \leq m} |\lambda^{(i)}|, \quad \lambda \in \mathbb{R}^m.$$

In this case, the Lipschitz constant for the gradients of function $\xi_L(\cdot)$ can be estimated as follows:

$$\begin{aligned} & \|\nabla \xi_L(\lambda_1) - \nabla \xi_L(\lambda_2)\|_\infty \\ &= \max_{1 \leq i \leq m} |\langle g_i, \nabla \Phi(L\nabla d(\bar{x}) - G\lambda_2) - \nabla \Phi(L\nabla d(\bar{x}) - G\lambda_1) \rangle| \\ &\leq \max_{1 \leq i \leq m} \|g_i\|_* \cdot \|\nabla \Phi(L\nabla d(\bar{x}) - G\lambda_2) - \nabla \Phi(L\nabla d(\bar{x}) - G\lambda_1)\| \\ &\stackrel{(29)}{\leq} \max_{1 \leq i \leq m} \|g_i\|_* \cdot \frac{1}{L} \|G(\lambda_1 - \lambda_2)\|_* \leq \frac{1}{L} \max_{1 \leq i \leq m} \|g_i\|_*^2 \cdot \|\lambda_1 - \lambda_2\|_1. \end{aligned}$$

Thus, the gradients of function $\xi_L(\cdot)$ are Lipschitz continuous with the constant

$$L(\xi_L) = \frac{1}{L} \max_{1 \leq i \leq m} \|g_i\|_*^2. \quad (32)$$

Since the diameter of the standard simplex in \mathbb{R}^m in ℓ_1 -norm is two, in accordance to the estimate (3.13) in [6], we can guarantee the following rate of convergence:

$$\delta_L^*(T) \leq \frac{18}{L \cdot T} \max_{1 \leq i \leq m} \|g_i\|_*^2, \quad t \geq 1. \quad (33)$$

(Here we replace the constant $\frac{136}{11 \ln 2}$ from [6] by a bigger value 18.) In accordance to the condition (14), this means that we need

$$N_L(\delta) = \frac{18}{L \cdot \delta} \max_{1 \leq i \leq m} \|g_i\|_*^2 \quad (34)$$

iterations of the method (31) in order to generate an appropriate dual solution $\bar{\lambda}$.

4. Unconstrained minimization in Euclidean setup

In this section we consider the simplest unconstrained minimization problem

$$f^* = \min_{x \in \mathbb{E}} f(x), \quad (35)$$

where $f(\cdot)$ is a smooth convex function. For measuring distances in \mathbb{E} , we introduce a Euclidean norm

$$\|x\| = \langle Bx, x \rangle^{1/2}, \quad x \in \mathbb{E},$$

where $B = B^* > 0$ is a linear operator from \mathbb{E} to \mathbb{E}^* . Then the dual norm is defined as follows:

$$\|g\|_* = \langle g, B^{-1}g \rangle^{1/2}, \quad g \in \mathbb{E}^*.$$

Let us choose now the distance function $d(x) = \frac{1}{2}\|x\|^2$. Then the Bregman distance is given by

$$\beta_d(x, y) = \frac{1}{2}\|x - y\|^2, \quad x, y \in \mathbb{E}.$$

In this case, the relative smoothness condition (3) is equivalent to strong convexity and Lipschitz continuity of the gradient:

$$\begin{aligned} f(y) - f(x) - \langle \nabla f(x), y - x \rangle &\geq \frac{1}{2}\mu_d(f)\|x - y\|^2, \\ f(y) - f(x) - \langle \nabla f(x), y - x \rangle &\leq \frac{1}{2}L_d(f)\|x - y\|^2. \end{aligned} \quad x, y \in \text{dom } f, \quad (36)$$

Let us write down now the specific form of the objective function $\xi_L(\cdot)$ in problem (13). Note that in our case

$$\Phi_L(s) = \max_{y \in \mathbb{E}} \left\{ \langle s, y \rangle - \frac{L}{2}\|y\|^2 \right\} = \frac{1}{2L}\|s\|_*^2, \quad s \in \mathbb{E}^*.$$

Therefore,

$$\xi_L(\lambda) = \frac{1}{2L} \|LB\bar{x} - G\lambda\|_*^2 + \langle \lambda, f_* \rangle + \alpha,$$

where $\alpha = -\frac{1}{2}L\|\bar{x}\|^2$. The gradient of function $\xi_L(\cdot)$ can be computed as follows:

$$\begin{aligned} \nabla \xi_L(\lambda) &= \frac{1}{L} G^* B^{-1} (G\lambda - LB\bar{x}) + f_*, \quad \lambda \in \mathbb{R}^m \\ &= \frac{1}{L} Q\lambda - \bar{f}, \end{aligned} \tag{37}$$

where $Q = G^* B^{-1} G$ and $\bar{f} = G^* \bar{x} - f_*$. Note that

$$\bar{f}^{(i)} \stackrel{(9)}{=} f_i + \langle g_i, \bar{x} - z_i \rangle, \quad i = 1, \dots, m.$$

Thus, in the Euclidean setup, our auxiliary problem (13) can be written as follows:

$$\min_{\lambda \in \Delta_m} \left[\xi_L(\lambda) = \frac{1}{2L} \langle \lambda, Q\lambda \rangle - \langle \lambda, \bar{f} \rangle \right]. \tag{38}$$

The stopping criterion (14) for this problem is as follows:

$$\begin{aligned} \langle \bar{\lambda}, \nabla \xi_L(\bar{\lambda}) \rangle &\stackrel{(37)}{=} \left\langle \bar{\lambda}, \frac{1}{L} Q\bar{\lambda} - \bar{f} \right\rangle \stackrel{(14)}{\leq} \delta + \min_{1 \leq i \leq m} \nabla_i \xi_L(\bar{\lambda}) \\ &= \delta + \min_{1 \leq i \leq m} \left(\frac{1}{L} Q\bar{\lambda} - \bar{f} \right)^{(i)}. \end{aligned}$$

Note that the main output of the minimization process for problem (38) is

$$x_+ = y_*(LB\bar{x} - G\bar{\lambda}) \stackrel{(12)}{=} \frac{1}{L} B^{-1} (LB\bar{x} - G\bar{\lambda}) = \bar{x} - \frac{1}{L} B^{-1} G\bar{\lambda}.$$

Then

$$\bar{f} - \frac{1}{L} Q\bar{\lambda} = \bar{f} - \frac{1}{L} G^* B^{-1} G\bar{\lambda} = \bar{f} + G^*(x_+ - \bar{x}).$$

Hence, in the Euclidean case, the stopping criterion (14) can be written as follows:

$$\sum_{i=1}^m \bar{\lambda}^{(i)} [f_i + \langle g_i, x_+ - z_i \rangle] \leq \delta + \max_{1 \leq i \leq m} [f_i + \langle g_i, x_+ - z_i \rangle]. \tag{39}$$

Now we can estimate the computational expenses of the method (31) as applied to the auxiliary problem (38).

- (1) Computation of matrix Q : $O(m^2 n)$ arithmetic operations. For certain strategies for *updating* the sets Z_k , it can be reduced to $O(mn)$ operations.
- (2) Computation of the vector \bar{f} : $O(mn)$ operations.
- (3) Computation of the initial gradient $u_0 = \frac{1}{L}\lambda_0 - \bar{f}$: $O(m^2)$ operations. For certain updating strategies it can be $O(m)$.
- (4) Expenses at each iteration:

- Computing the index i_k : $O(m)$ operations.
- Updating the point λ_k : $O(m)$ operations.
- Updating the gradient $u_k = \frac{1}{L}Q\lambda_k - \tilde{f}$: $O(m)$ operations.

Thus, taking into account the upper bound (34) for the number of iterations in method (31), we obtain the following bound for the arithmetic complexity of problem (38) with reasonable updating strategies for the sets \mathcal{Z}_k :

$$O\left(mn + \frac{m}{L \cdot \delta} \max_{1 \leq i \leq m} \|g_i\|_*^2\right). \quad (40)$$

Taking into account that we can expect that in the problem (35) we have

$$\frac{1}{2L_d(f)} \|\nabla f(z_i)\|_*^2 \stackrel{(36)}{\leq} f(z_i) - f^* \rightarrow 0$$

as $i \rightarrow \infty$, the bound in (40) suggests that the overhead of solving the inner problem (38) decreases to a small constant in $O(mn)$ as the algorithm approaches the optimum.

5. Numerical experiments

In this section we present preliminary computational results for method (26) as applied to the following unconstrained minimization problem:

$$\min_{x \in \mathbb{R}^n} \left[f(x) = \mu \ln \left(\sum_{j=1}^M e^{\langle a_j, x \rangle - b_j} / \mu \right) \right]. \quad (41)$$

The data defining this function is randomly generated in the following way. First of all, we generate a collection of random vectors

$$\hat{a}_1, \dots, \hat{a}_M$$

with entries uniformly distributed in the interval $[-1, 1]$. Using the same distribution, we generate values $b_j, j = 1, \dots, m$. Using this data, we form the preliminary function

$$\hat{f}(x) = \mu \ln \left(\sum_{j=1}^M e^{\langle a_j, x \rangle - b_j} / \mu \right)$$

and compute $g = \nabla \hat{f}(0)$. Then, we define

$$a_j = \hat{a}_j - g, \quad j = 1, \dots, n.$$

Clearly, in this case we have $\nabla f(0) = 0$, so the unique solution of our test problem (41) is $x^* = 0$. The starting point x_0 is chosen in accordance to the uniform distribution on the Euclidean sphere of radius one.

Thus, the problem (41) has three parameters, the dimension n , the number of linear functions $M \geq n$, and the smoothness coefficient $\mu > 0$. In our experiments, we always choose $M = 6n$. Let us present our computational results for different values of n and μ .

Table 1. Smoothness parameter $\mu = 0.05$.

$n = 100$	GM	Cyclic	Max-Norm
Iter	2683	801	664
NFunc	5371	1606	1332
FW/Iter	2	55	66
Time (s)	1.94	0.88	0.73
IT(ms)	0.72	1.10	1.10
$n = 250$			
Iter	2148	227	227
NFunc	4302	459	459
FW/Iter	2	243	243
Time (s)	10.20	1.36	1.36
IT(ms)	4.75	5.99	5.99
$n = 500$			
Iter	2902	268	268
NFunc	5809	537	537
FW/Iter	2	428	428
Time (s)	52.16	5.5	5.5
IT(ms)	17.97	20.52	20.52

In the definition of method (26) we have some freedom in the choice of the bundle \mathcal{Z}_k . Let us bound its maximal size by a parameter $m \geq 1$. Then $m = 1$ corresponds to the usual Gradient Method. In the first series of our experiments (shown in Tables 1 and 2) we always choose

$$m = n.$$

We also have some freedom in the updating strategy for the sets \mathcal{Z}_k . Clearly, at the first m steps we can simply add all new points in the bundle. However, at the next iterations we need to decide on the strategy of replacement of the old information. In our experiments we implemented two strategies:

- Cyclic replacement (CYCLIC).
- Replacement of the linear function with the maximal norm of the gradient (MAX-NORM).

The second strategy is motivated by the formula (32) for the Lipschitz constant of the gradient of function $\xi_L(\cdot)$. For both strategies, at each iteration we need to update only one column of matrix Q_k (see (37)), which costs $O(mn)$ operations.

Let us present the results of our numerical experiments. All methods were stopped when the residual in the function value was smaller than $\epsilon = 10^{-6}$. The parameter δ for the stopping criterion (14) was chosen as $\delta = \epsilon/2$.

In Tables 1 and 2, the first line indicates the total number of iterations. The second line displays the total number of oracle calls. The third line shows the average number of Frank–Wolfe steps per iterations (for the Gradient Method we just put two). The next line indicates the total computational time (in seconds). Finally, at the last line we can see the average time spent on one iteration of the corresponding method (in milliseconds).

As we can see from these tables, in all our experiments the gradient methods with memory were better than the standard Gradient Method, both in the number of iterations, and, what is rather surprising, in the total computational time. The MAX-NORM version usually

Table 2. Smoothness parameter $\mu = 0.01$.

$n = 100$	GM	Cyclic	Max-Norm
Iter	43893	4171	6710
NFunc	87795	8351	13427
FW/Iter	2	59	17
Time (s)	31.73	4.50	6.95
IT(ms)	0.72	1.08	1.04
$n = 250$			
Iter	116479	45183	25492
NFunc	232967	90377	50990
FW/Iter	2	8	13
Time (s)	540.64	311.95	176.02
IT(ms)	4.64	6.90	6.90
$n = 500$			
Iter	105610	38144	29916
NFunc	211229	76297	59840
FW/Iter	2	10	17
Time (s)	1894.19	1010.30	788.97
IT(ms)	17.94	26.49	26.37

Table 3. Gradient method with cyclic memory replacement.

Bundle	$n = 100$			$n = 200$			$n = 400$		
	Iter	FW	Sec	Iter	FW	Sec	Iter	FW	Sec
1	2683	2	1.94	1753	2	4.78	1676	2	19.86
2	2543	3	1.84	1433	7	3.91	1669	5	19.81
4	1755	13	1.31	829	39	2.31	1084	18	12.92
8	1363	19	1.03	633	60	1.78	758	37	9.08
16	1220	22	0.97	579	68	1.67	720	44	8.69
32	1202	24	1.00	593	71	1.76	719	47	8.78
64	1138	28	1.06	423	114	1.39	636	58	8.02
128	880	60	1.06	247	320	0.94	349	127	4.63
256	207	361	0.36	173	507	0.70	202	260	2.70

outperforms the CYCLIC version. It is interesting that the auxiliary algorithm (31) works very well. The average time spent on one iteration of the methods with memory is never increased more than on 50% of the time of the simple Gradient Method. This is partially explained by the fact that in our test problems the data is fully dense, so each call of oracle is very expensive ($O(Mn)$ operations).

Let us look now at how small bundles can accelerate the Gradient Method. In Tables 3 and 4, the first line with parameter Bundle = 1 corresponds to the Gradient method with line search. The next lines display the results for different sizes of the bundle. We list the number of iterations, the average number of Frank–Wolfe steps per iteration, and the total computational time in seconds. Table 3 displays the results for the IGMM (26) with the cyclic replacement strategy for each bundle size. In Table 4, we show the results for the Max-Norm replacement strategy. The accuracy parameters for the experiments shown in Tables 3 and 4 are $\epsilon = 10^{-4}$ and $\delta = \epsilon/2$. The smoothness parameter for our objective function is chosen as $\mu = 0.05$.

As we can see from these tables, the Max-Norm replacement strategy was always better than the cyclic one. Even for small bundle sizes, the total number of iterations decreases very quickly. What is more important, this decrease is also seen in total computation time.

Table 4. Gradient method with Max-Norm memory replacement.

Bundle	$n = 100$			$n = 200$			$n = 400$		
	Iter	FW	Sec	Iter	FW	Sec	Iter	FW	Sec
1	2683	2	1.94	1753	2	4.76	1676	2	19.86
2	1288	4	0.95	737	11	2.05	796	7	9.47
4	592	27	0.45	316	79	0.91	385	38	4.61
8	400	48	0.33	269	121	0.78	283	86	3.42
16	362	58	0.33	418	91	1.22	335	86	4.05
32	786	33	0.67	354	116	1.08	354	89	4.34
64	838	39	0.80	379	127	1.23	421	85	5.31
128	529	100	0.66	285	277	1.09	352	127	4.64
256	207	361	0.36	173	507	0.70	202	260	2.70

The number of auxiliary Frank–Wolfe steps remains on an acceptable level and cannot increase significantly the computational time of each iteration as compared with the Gradient Method. Recall that our test function has an expensive oracle, requiring $O(Mn)$ operations for computing the function value and the gradient. For the Max-Norm version of IGMM, the optimal size of the bundle is probably between 8 and 16. Another candidate is 256, but it needs many more Frank–Wolfe steps.

Maybe our preliminary conclusions are problem specific. However, we believe that in any case they demonstrate a high potential of our approach in increasing the efficiency of gradient methods, both in accelerated and, hopefully, non-accelerated variants.

Notes

1. By Gradient Method we denote the extended scheme described in [1], which encompasses Gradient Descent and the Proximal Gradient Method.
2. The dual of the problem with the objective multiplied by -1 .
3. Recall that a function is closed if its epigraph is a closed set.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This work was supported by the European Research Council (ERC) under Advanced Grant 788368.

Notes on contributors

Yurii Nesterov Born: 1956, Moscow. Master degree 1977, Moscow State University. Doctor degree 1984. Professor at Center for Operations Research and Econometrics, UCLouvain, Belgium. Author of 5 monographs and more than 100 refereed papers in the leading optimization journals. Dantzig Prize, John von Neumann Theory Prize 2009, Charles Broyden prize 2010, Francqui Chair (Liege University 2011–2012), SIAM Outstanding paper award 2014, EURO Gold Medal 2016.

Main direction is the development of efficient numerical methods for convex and nonconvex optimization problems supported by the global complexity analysis: general interior-point methods (theory of self-concordant functions), fast gradient methods (smoothing technique), global complexity analysis of second-order and tensor schemes (cubic regularization of Newton's method).

Mihai I. Florea Bachelor degree 2009, Tohoku University. Master degree 2013, Uppsala University. Doctor degree 2018, Aalto University. Postdoctoral Fellow with the Department of Mathematical Engineering, UCLouvain, Belgium. Japanese Government (MEXT) Scholarship 2004–2009, Tohoku University Head of School of Engineering Prize 2009, Aalto University Dissertation Award 2018.

Research interests include first-order methods for convex minimization and applications to medical imaging.

ORCID

Yurii Nesterov  <http://orcid.org/0000-0002-0542-8757>

Mihai I. Florea  <http://orcid.org/0000-0003-4275-2449>

References

- [1] H. Bauschke, J. Bolte, and M. Teboulle, *A descent lemma beyond Lipschitz gradient continuity: First-order methods revisited and applications*, *Math. Oper. Res.* 42 (2017), pp. 330–348.
- [2] D. Drusvyatskiy, M. Fazel, and S. Roy, *An optimal first order method based on optimal quadratic averaging*, *SIAM J. Optim.* 28 (2018), pp. 251–271.
- [3] M. Frank and P. Wolfe, *An algorithm for quadratic programming*, *Naval Res. Logist. Q.* 3 (1956), pp. 149–154.
- [4] H. Lu, R. Freund, and Yu. Nesterov, *Relatively smooth convex optimization by first-order methods, and applications*, *SIAM J. Optim.* 28S (2018), pp. 333–354.
- [5] Yu. Nesterov, *Gradient methods for minimizing composite functions*, *Math. Program.* 140 (2013), pp. 125–161.
- [6] Yu. Nesterov, *Complexity bounds for primal–dual methods minimizing the model of objective function*, *Math. Program.* 171 (2018), pp. 311–330.
- [7] Yu. Nesterov, *Lectures on Convex Optimization*, Springer, Berlin, Germany, 2018.