

# *Open Journal of Mathematical Optimization*

Lénaïc Chizat

**Convergence Rates of Gradient Methods for Convex Optimization in the Space of Measures**

Volume 3 (2022), article no. 8 (19 pages)

<https://doi.org/10.5802/ojmo.20>

Article submitted on June 7, 2021, accepted on November 13, 2022.

© The author(s), 2022.



This article is licensed under the  
CREATIVE COMMONS ATTRIBUTION 4.0 INTERNATIONAL LICENSE.  
<http://creativecommons.org/licenses/by/4.0/>



# Convergence Rates of Gradient Methods for Convex Optimization in the Space of Measures

Lénaïc Chizat

Institute of Mathematics

École polytechnique fédérale de Lausanne (EPFL), Switzerland

lenaic.chizat@epfl.ch

---

## Abstract

We study the convergence rate of Bregman gradient methods for convex optimization in the space of measures on a  $d$ -dimensional manifold. Under basic regularity assumptions, we show that the suboptimality gap at iteration  $k$  is in  $O(\log(k)k^{-1})$  for multiplicative updates, while it is in  $O(k^{-q/(d+q)})$  for additive updates for some  $q \in \{1, 2, 4\}$  determined by the structure of the objective function. Our flexible proof strategy, based on approximation arguments, allows us to painlessly cover all Bregman Proximal Gradient Methods (PGM) and their acceleration (APGM) under various geometries such as the hyperbolic entropy and  $L^p$  divergences. We also prove the tightness of our analysis with matching lower bounds and confirm the theoretical results with numerical experiments on low dimensional problems. Note that all these optimization methods must additionally pay the computational cost of discretization, which can be exponential in  $d$ .

Digital Object Identifier 10.5802/ojmo.20

Keywords Gradient Descent, Convergence rate, Space of measures, Convex Optimization, Banach space.

## 1 Introduction

Convex optimization in the space of measures is a theoretical framework that leads to fruitful point of views on a large variety of problems, ranging from sparse deconvolution [12] and two-layer neural networks [9] to global optimization [31] and many more [10]. Various algorithms have been proposed to solve such problems including moments methods [31], conditional gradient [12, 21], (non-convex) particle gradient flows [16] and noisy versions [33, 37].

In this paper, we consider perhaps the simplest methods: gradient descent and its extensions that handle non-smooth regularizers and non-Euclidean geometries, the Bregman Proximal Gradient Method (PGM) (an extension of mirror descent [34] that handles composite objectives) and its acceleration (APGM) [40]. Our aim is to establish well-posedness and convergence rates for these methods when minimizing composite functions over the space of measures  $\mathcal{M}(\Theta)$  over a  $d$ -dimensional manifold  $\Theta$ , of the form

$$F(\mu) := R\left(\int \Phi d\mu\right) + H(\mu) \tag{1}$$

where  $\Phi$  is continuous and Hilbert space-valued,  $R$  convex and smooth and  $H$  is convex and “simple” (see precise assumptions in Section 3.1). For such problems, minimizers are typically at an infinite (Bregman) distance from the initialization, and thus all the standard convergence bounds are inapplicable.

Our contributions are the following:

- We recall and adapt (A)PGM in Section 3, taking care of the subtleties that appear in our context (definition of the iterates and lack of strong convexity of the divergence);
- We prove in Section 4 upper-bounds on the convergence rate for (A)PGM under various structural assumptions, summarized in Table 2. These rates depend on the choice of the Bregman divergence and on the precise structure of the objective function;
- Tight lower bounds of two kinds are proved in Section 5: *proof technique*-dependent lower bounds, and *algorithm*-dependent lower bounds (the latter are stronger but do not cover all cases);



© Lénaïc Chizat;  
licensed under Creative Commons License Attribution 4.0 International

- Numerical experiments on synthetic toy problems in Section 6 often show an excellent agreement between the theoretical rates and the ones observed in practice. Even for cases with an apparent mismatch, a closer look at the structure of the problem shows that the theory still shades light on the observed rates.

Our motivation for studying this problem is threefold. First, our results make a case for APGM with the hyperbolic geometry instead of FISTA to solve convex problems in the space of measures, as they show that the former enjoys a faster convergence rate. Second, we believe that a precise understanding of (A)PGM in this context is useful to develop and analyze more complex methods, such as the particle-based (a.k.a. moving grid) approaches mentioned above<sup>1</sup>. Third, this setting offers a rich test case to deepen our understanding of Bregman gradient methods in Banach spaces, and the behavior of optimization algorithms when all minimizers are at an infinite distance from the initialization, beyond the well-explored Hilbert space setting.

## Related work

The comparison between additive updates ( $L^2$  geometry) and multiplicative updates (entropy geometry) is well-known in finite dimensional spaces [30]. For instance, for convex optimization in the  $n$ -dimensional simplex, the two methods typically converge at the same rate but the “constant” factor is polynomial in  $n$  for additive updates while it is logarithmic in  $n$  for multiplicative updates, see [13, §4]. We obtain in this paper an infinite dimensional ( $n = \infty$ ) version of this separation; but where the distinction is directly in the rates rather than in the constants.

Analysis of convex optimization in infinite dimensional (Banach) spaces is a classical subject [6, 7]. Here, we study a concrete class of problems defined on the space of measures which exhibit specific features. This problem-specific approach for infinite dimensional problems has proved fruitful for the analysis of gradient methods for least-squares (e.g. [22, 43] and references therein), for partly smooth problems [32] and for the Iterative Soft Thresholding Algorithm (ISTA) in Hilbert spaces [11, 24].

The latter is close to our subject since ISTA is in fact an instance of PGM with the  $L^2$ -divergence – and FISTA [8] is analogous to APGM with the  $L^2$ -divergence. These prior works perform the analysis in a Hilbert space, while we work in the space of measures or in  $L^1$ , which are non-reflexive Banach space. This is also the context of [15] who, for a modified version of FISTA, obtained in particular the convergence rate of Table 2 when  $p = 2$  and  $q = 1$ , and also discuss discretization. Our analysis allows to compare various algorithms and shows that FISTA is always slower than APGM with the hyperbolic entropy geometry [25] when the solution is truly sparse, see the rates in Table 2. This is clearly observed in numerical experiments and suggests that the latter forms a stronger baseline for our class of problems.

To prove our upper bounds, we use the abstract proof strategy proposed by [28], recalled in Section 2. In that paper, the authors study different classes of problems (total variation denoising of image and earth mover’s distance) under Hilbertian geometry.

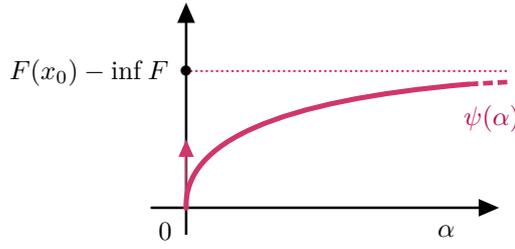
## Notation

The domain of a function  $F : V \rightarrow \mathbb{R} \cup \{+\infty\}$  is  $\text{dom } F = \{x \in V ; F(x) < +\infty\}$ . Throughout,  $\Theta$  is a compact  $d$ -dimensional manifold,  $\mathcal{M}(\Theta)$  (resp.  $\mathcal{M}_+(\Theta)$ ) is the set of finite signed (resp. nonnegative) Borel measures on  $\Theta$  and  $\mathcal{P}(\Theta)$  is the set of Borel probability measures. For  $\mu \in \mathcal{M}(\Theta)$ ,  $\|\mu\|$  is its total variation norm. For a Hilbert space  $\mathcal{F}$ ,  $\mathcal{C}^p(\Theta; \mathcal{F})$  is the set of  $p$ -times continuously differentiable functions from  $\Theta$  to  $\mathcal{F}$ .  $\text{Lip}(f)$  is the Lipschitz constant of a function  $f$ . For  $\tau \in \mathcal{P}(\Theta)$  and  $p \geq 1$ ,  $L^p(\tau)$  is the space of (equivalence classes of) measurable functions  $f : \Theta \rightarrow \mathbb{R}$  such that  $\int_{\Theta} |f(\theta)|^p d\tau(\theta) < +\infty$  or, for  $p = +\infty$ , such that  $|f|$  is  $\tau$ -almost everywhere bounded by some  $K > 0$ . The asymptotic notation  $a(k) \lesssim b(k)$  means that there exists  $c > 0$  independent of  $k$  such that  $a(k) \leq c \cdot b(k)$ , and  $a(k) \asymp b(k)$  means  $[a \lesssim b \text{ and } b \lesssim a]$ .

## 2 Strategy to derive upper bounds on convergence rates

This section introduces the strategy, adapted from [28], that we adopt to derive upper bounds on the convergence rates.

<sup>1</sup> In fact, the idea of writing this paper came from a technical step in a proof of [16], which studies particle-based methods.



■ **Figure 1** Shape of  $\psi$  defined in Eq. (3) in our situation of interest where  $D(x_k, x_0)$  explodes for any minimizing sequence  $(x_k)_{k \in \mathbb{N}}$  (Proposition 1). When an optimization method satisfies Eq. (2) for some sequence  $(\alpha_k)_{k \in \mathbb{N}}$  then  $\psi(\alpha_k)$  bounds its convergence rate in objective values.

Let  $F$  be a lower bounded convex function defined on a real vector space. Suppose that an iterative method designed to minimize  $F$  initialized at  $x_0 \in \text{dom } F$  generates a sequence  $x_1, x_2, \dots \in \text{dom } F$  that satisfies

$$F(x_k) - F(x) \leq \alpha_k \cdot D(x, x_0), \quad \forall x \in \text{dom } F, \quad \forall k \geq 1, \quad (2)$$

where  $(\alpha_k)_{k \in \mathbb{N}^*}$  is a positive sequence converging to 0 and  $D$  is a *divergence*, i.e.  $D(x, x_0) \in [0, +\infty]$  and  $D(x_0, x_0) = 0$ . Most first order methods enjoy guarantees of this form. For instance, PGM and APGM enjoy such guarantees with respectively  $\alpha_k \lesssim k^{-1}$  and  $\alpha_k \lesssim k^{-2}$  under suitable assumptions, see Section 3.3.

While Eq. (2) is sometimes the endpoint of the analysis in the optimization literature, this is our starting point: we are interested in cases where for any minimizer  $x^*$  of  $F$ , the quantity  $D(x^*, x_0)$  is infinite, which makes the bound of Eq. (2) inapplicable. Even if there exists a quasi-minimizer  $x$  with a small suboptimality gap and satisfying  $D(x, x_0) < +\infty$ , choosing a fixed  $x$  independent of  $k$  in Eq. (2) leads to a poor upper bound which often does not match the observed practical behavior. Instead, we should exploit the flexibility offered by the guarantee of Eq. (2) and choose a different reference point at each time step<sup>2</sup>. This means that we reformulate the guarantee in the equivalent form:

$$F(x_k) - \inf F \leq \psi(\alpha_k) \quad \text{where} \quad \psi(\alpha) := \inf_x \left\{ F(x) - \inf F + \alpha D(x, x_0) \right\}. \quad (3)$$

Studying  $\psi$  is particularly fruitful to understand optimization algorithms satisfying Eq. (2). In particular, its behavior at 0 determines the asymptotic convergence rate. This function can be interpreted as the value at  $x_0$  of the (Bregman) Moreau envelope [29] of  $(F - \inf F)$  with regularization parameter  $\alpha$ , and it intervenes in many area of applied mathematics. For instance, when  $D(x, x_0)$  is a squared Hilbertian norm,  $\psi$  has a variety of behaviors which characterize the performance of kernel ridge regression in machine learning (see e.g. [5, Chap. 7.5]). Before we head in a more concrete setting, let us gather a few relevant properties of the function  $\psi$  that hold in full generality.

► **Proposition 1.** *Assume that  $F(x_0) < +\infty$  and  $D(\cdot, x_0) \geq 0$  with equality at  $x_0$ . Then the function  $\psi$  is concave on  $[0, +\infty[$  and satisfies  $0 = \psi(0) \leq \psi(\alpha) \leq F(x_0) - \inf F$ . Moreover,*

- i.  $\psi$  is right-continuous at 0 if and only if there exists a minimizing sequence  $(x_k)_{k \in \mathbb{N}}$  such that  $F(x_k) \rightarrow \inf F(x)$  and  $D(x_k, x_0) < +\infty, \forall k \in \mathbb{N}$  ;
- ii.  $\psi'(0) := \lim_{\alpha \rightarrow 0^+} \psi(\alpha)/\alpha$  is finite if and only if there exists a minimizing sequence  $(x_k)_{k \in \mathbb{N}}$  such that  $F(x_k) \rightarrow \inf F(x)$  and  $D(x_k, x_0)$  is bounded.

**Proof.** The function  $\psi$  is concave as the pointwise infimum of affine functions. The lower bound is immediate and the upper bound is obtained by taking  $x_0$  as a candidate in the infimum. Let us prove ii (the proof of i follows a similar scheme and is simpler). By concavity, the limit defining  $\psi'(0)$  always exists and belongs to  $]0, +\infty]$ . If a sequence  $(x_k)$  exists as in the statement, then for any  $\alpha > 0$ , pick  $x_k$  such that  $F(x_k) - \inf F \leq \alpha^2$  and then  $\psi(\alpha) \leq \alpha D(x_k, x_0) + \alpha^2$  so  $\psi(\alpha)/\alpha \leq D(x_k, x_0) + \alpha$ . Since the upper bound is uniformly bounded as  $\alpha \rightarrow 0$  it follows that  $\psi'(0)$  is finite. Conversely, if  $\psi'(0)$  is finite, take a decreasing sequence  $(\alpha_k)$  that converges to 0 and let  $(x_k)$  be a sequence of quasi-minimizers for Eq. (3) satisfying  $F(x_k) - \inf F + \alpha_k D(x_k, x_0) \leq \psi(\alpha_k) + \alpha_k^2$ . By dividing by  $\alpha_k$ , we see that  $(F(x_k) - \inf F)/\alpha_k + D(x_k, x_0)$  is bounded as  $k \rightarrow \infty$  which implies that  $F(x_k) \rightarrow \inf F$  and  $D(x_k, x_0)$  is bounded. ◀

<sup>2</sup> In Section 4, these reference points will be constructed as mollifications of the optimal measure  $\mu^*$ .

Figure 1 illustrates the general shape of the function  $\psi$ . Observe that if  $\psi'(0) < +\infty$  then the bound of Eq. (3) is  $F(x_k) - \inf F \leq \psi'(0)\alpha_k + o(\alpha_k)$  and thus the convergence rate given by Eq. (2) is not modified (only the constant changes). However, Proposition 1 shows that when any minimizing sequence  $(x_k)$  satisfies  $D(x_k, x_0) \rightarrow +\infty$ , then  $\psi'(0) = +\infty$  and thus the *convergence rate* is modified. This is the situation we are interested in this paper, in the context of optimization in the space of measures.

### 3 Gradient methods for optimization in the space of measures

In the rest of this paper, we apply the general method of Section 2 to a class of optimization problems in the space of measures where it leads to a zoo of – often tight – convergence rates.

#### 3.1 Objective function

Let  $\Theta$  be a compact Riemannian manifold without boundary, with distance  $\text{dist}$  and with a reference probability measure  $\tau \in \mathcal{P}(\Theta)$  that is proportional to the volume measure. We consider an objective function on the space of measures  $\bar{F} : \mathcal{M}(\Theta) \rightarrow \mathbb{R} \cup \{+\infty\}$  of the form

$$\bar{F}(\mu) := \bar{G}(\mu) + \bar{H}(\mu) \quad \text{where} \quad \bar{G}(\mu) := R\left(\int \Phi d\mu\right).$$

Typically,  $\bar{G}$  is a data-fitting term and  $\bar{H}$  a regularizer. We make the following assumptions, where  $\iota_C$  is the convex indicator of a convex set  $C$  and  $\lambda \geq 0$  a regularization parameter:

**(A1).**  $\Phi \in \mathcal{C}^0(\Theta; \mathcal{F})$  where  $\mathcal{F}$  is a Hilbert space,  $R : \mathcal{F} \rightarrow \mathbb{R}$  is convex and differentiable with a Lipschitz gradient  $\nabla R$ , and  $\bar{H}$  is a sum of functions from the following list:  $\iota_{\mathcal{P}(\Theta)}$ ,  $\iota_{\mathcal{M}_+(\Theta)}$ ,  $\lambda\|\mu\|$  and  $\iota_{\{\mu; \lambda\|\mu\| \leq 1\}}$ .

One specific property of  $\bar{H}$  that we use in our proof is that it should be non-decreasing under convolutions by a probability kernel, but we prefer to work with these specific instances rather than giving abstract conditions. We finally denote by  $F : L^1(\tau) \rightarrow \mathbb{R} \cup \{+\infty\}$  the function defined, for  $f \in L^1(\tau)$ , by

$$F(f) := \bar{F}(f\tau),$$

and similarly  $H(f) := \bar{H}(f\tau)$  and  $G(f) := \bar{G}(f\tau)$  so that  $F = G + H$ . These “bar” notations convey the idea that  $\bar{F}, \bar{G}, \bar{H}$  are the lower-semicontinuous (l.s.c.) extensions of  $F, G, H$  for the weak\* topology induced by  $\mathcal{C}^0(\Theta)$  on  $\mathcal{M}(\Theta)$ .

Here are examples of problems that fall under this setting:

- **(Sparse deconvolution).** The goal is, given a signal  $y^* \in L^2(\tau)$ , to find a sparse measure  $\mu$  such that the convolution of  $\mu$  with a filter  $\phi \in \mathcal{C}(\Theta)$  approximately recovers  $y^*$ . Here the domain is typically the  $d$ -dimensional torus  $\Theta = \mathbb{T}^d$  endowed with the Lebesgue measure  $\tau$  and the objective is [14, 20]

$$\bar{F}(\mu) := \int_{\Theta} \left| \int_{\Theta} \phi(\theta_1 - \theta_2) d\mu(\theta_2) - y^*(\theta_1) \right|^2 d\tau(\theta_1) + \lambda\|\mu\|.$$

Adding the nonnegativity constraint  $\iota_{\mathcal{M}_+(\Theta)}$  is also relevant in certain applications.

- **(Two-layer relu neural networks).** The goal is, given  $n$  observations  $(x_i, y_i) \in \mathbb{R}^d \times \mathbb{R}$ , to find a regressor written as a linear combination of simple “ridge” functions. Consider a loss  $\ell : \mathbb{R}^2 \rightarrow \mathbb{R}$  convex and smooth in its second argument, let  $\phi(s) = (s)_+$  and let

$$\bar{F}(\mu) := \frac{1}{n} \sum_{i=1}^n \ell\left(y_i, \int_{\Theta} \phi([x_i; 1]^\top \theta) d\mu(\theta)\right) + \bar{H}(\mu).$$

Key differences with the previous setting are that  $\Theta = \mathbb{S}^d$  is the sphere, with  $d$  potentially large, and that the object that is truly sought after is the regressor  $x \mapsto \int \phi([x; 1]^\top \theta) d\mu(\theta)$  rather than the measure  $\mu$ . Typical choices for  $\ell$  are the logistic loss  $\ell(y, z) = \log(1 + \exp(-yz))$  when  $y_i \in \{-1, 1\}$  or the square loss  $\ell(y, z) = \frac{1}{2}|y - z|^2$ . The signed setting with regularization  $\bar{H}(\mu) = \lambda\|\mu\|$  is the most common one [4, 9] but the regularization  $\iota_{\mathcal{P}(\Theta)}$  also appears in the context of max-margin problems [17].

The following smoothness lemma will be useful to analyze optimization algorithms and is analogous to the usual “Lipschitz gradient” property in convex optimization. Since the dual of  $(\mathcal{M}(\Theta), \|\cdot\|)$  is a bit exotic, we avoid using the notion of gradient at all.

► **Lemma 2 (Smoothness).** *Under Assumption (A1), if  $\Phi \in \mathcal{C}^p(\Theta, \mathcal{F})$  for  $p \in \mathbb{N}$ , then the differential of  $\bar{G}$  at  $\mu \in \mathcal{M}(\Theta)$  can be represented by the function  $\bar{G}'[\mu] \in \mathcal{C}^p(\Theta, \mathbb{R})$  defined by*

$$\bar{G}'[\mu](\theta) = \left\langle \nabla R \left( \int \Phi d\mu \right), \Phi(\theta) \right\rangle_{\mathcal{F}}$$

in the sense that it holds  $\bar{G}(\mu + \nu) - \bar{G}(\mu) = \int \bar{G}'[\mu](\theta) d\nu(\theta) + o(\|\mu - \nu\|)$ . Moreover,  $\mu \mapsto \bar{G}'[\mu]$  is Lipschitz continuous as a function from  $\mathcal{M}(\Theta)$  to  $\mathcal{C}^p(\Theta, \mathbb{R})$ . The following smoothness inequality holds with  $\text{Lip}(\bar{G}') \leq \|\Phi\|_{\infty}^2 \cdot \text{Lip}(\nabla R)$  and for all  $\mu, \nu \in \mathcal{M}(\Theta)$ ,

$$0 \leq \bar{G}(\nu) - \bar{G}(\mu) - \int \bar{G}'[\mu] d[\nu - \mu] \leq \frac{1}{2} \text{Lip}(\bar{G}') \|\nu - \mu\|^2.$$

Those results hold true when replacing  $(\mathcal{M}(\Theta), \bar{G}(\mu), \bar{G}'[\mu])$  by  $(L^1(\tau), G(f), G'[f] := \bar{G}'[f\tau])$ .

**Proof.** For the first part, the differentiability of  $R$  implies that

$$\bar{G}(\mu + \nu) - \bar{G}(\mu) = \left\langle \nabla R \left( \int \Phi d\mu \right), \int \Phi d\nu \right\rangle_{\mathcal{F}} + o\left(\left\| \int \Phi d\nu \right\|_{\mathcal{F}}\right) = \int \bar{G}'[\mu] d\nu + o(\|\Phi\|_{\infty} \|\nu\|).$$

For the regularity of  $\mu \mapsto \bar{G}'[\mu]$ , we have for  $\mu, \nu \in \mathcal{M}(\Theta)$ ,

$$\|\bar{G}'[\mu] - \bar{G}'[\nu]\|_{\mathcal{C}^p} \leq \|\Phi\|_{\mathcal{C}^p(\Theta, \mathcal{F})} \cdot \text{Lip}(\nabla R) \cdot \|\Phi\|_{\infty} \cdot \|\mu - \nu\|.$$

The smoothness inequality can be shown by bounding a 1-dimensional integral as in the Euclidean case [36, Thm. 2.1.5]. Finally,  $L^1(\tau) \ni f \mapsto f\tau \in \mathcal{M}(\Theta)$  is an isometry, so those results hold mutandis mutatis in  $L^1(\tau)$ . ◀

### 3.2 Bregman divergences

Let us consider  $\eta : \mathbb{R} \rightarrow [0, \infty]$  a differentiable function that we will refer to as the *distance-generating* function. For  $f \in L^1(\tau)$  we write  $\eta(f) := \eta \circ f$  and we define

$$\bar{\eta}(f) := \int_{\Theta} \eta(f(\theta)) d\tau(\theta) = \int \eta(f) d\tau.$$

Let  $D_{\eta}$  (resp.  $D_{\bar{\eta}}$ ) be the Bregman divergence associated to  $\eta$  (resp.  $\bar{\eta}$ ), given for  $f, g \in L^1(\tau)$  by

$$D_{\eta}(a, b) := \eta(a) - \eta(b) - \eta'(b)(a - b) \quad \text{and} \quad D_{\bar{\eta}}(f, g) := \int_{\Theta} D_{\eta}(f(\theta), g(\theta)) d\tau(\theta).$$

We consider the following assumptions on the distance-generating function  $\eta$ :

**(A2).**  $\eta : \mathbb{R} \rightarrow [0, \infty]$  is strictly convex, l.s.c., continuously differentiable in  $\text{int}(\text{dom } \eta)$ , such that  $\eta'(\text{int}(\text{dom } \eta)) = \mathbb{R}$  and for any  $c > 0$  it holds  $\eta(cx) \asymp \eta(x)$  as  $x \rightarrow \infty$ . Moreover, either:

**(A2)<sub>+</sub>.**  $\text{dom } \eta = [0, +\infty[$  and  $\eta(1) = \eta'(1) = 0$ , or

**(A2)<sub>±.</sub>**  $\text{dom } \eta = \mathbb{R}$ ,  $\eta$  is even and  $\eta(0) = \eta'(0) = 0$ .

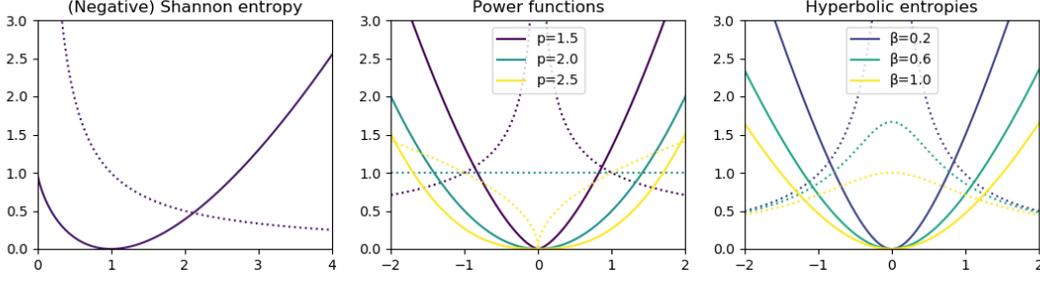
Specifying the values of  $\eta$  and  $\eta'$  at a point in  $\text{int}(\text{dom } \eta)$  is just for convenience and is not restrictive since  $D_{\eta}$  is not affected by affine perturbations of  $\eta$ . Also, the assumption  $\eta(cx) \asymp \eta(x)$  is only needed to simplify the statement of the results. Under assumption **(A2)<sub>+</sub>**, we have that  $\eta'(0) := \lim_{t \rightarrow 0^+} \eta(t)/t = -\infty$  which automatically enforces a nonnegativity constraint in the methods in the next section.

Here are examples of distance-generating functions that fall under these assumptions:

- (Power functions  $\eta_p$ ). Defined on  $\mathbb{R}$  for  $p > 1$  by  $\eta_p(s) := \frac{|s|^p}{p(p-1)}$ , satisfy **(A2)<sub>±</sub>**;
- (Shannon entropy  $\eta_{\text{ent}}$ ). Defined on  $\mathbb{R}_+$  by  $\eta_{\text{ent}}(s) := s \log(s) - s + 1$ , satisfies **(A2)<sub>+</sub>**;
- (Hyperbolic entropy  $\eta_{\text{hyp}}$ ). Defined on  $\mathbb{R}$  by  $\eta_{\text{hyp}} := s \text{arcsinh}(s/\beta) - \sqrt{s^2 + \beta^2} + \beta$  with  $\beta > 0$ , satisfies **(A2)<sub>±</sub>** (introduced by [25]).

When  $\eta$  is smooth, it holds  $D_{\eta}(a, b) = \eta''(b) \cdot \|a - b\|^2/2 + o(\|a - b\|^2)$  so locally,  $D_{\eta}$  is equivalent to a squared Riemannian metric on the real axis given by  $\eta''$ . For the examples listed above, it holds  $\eta_p''(s) = |s|^{p-2}$ ,  $\eta_{\text{ent}}''(s) = s^{-1}$  and  $\eta_{\text{hyp}}''(s) = (s^2 + \beta^2)^{-1/2}$ , see Figure 2 for an illustration. The hyperbolic entropy  $\eta_{\text{hyp}}$  can be interpreted as a “signed” version of  $\eta_{\text{ent}}$  (see Proposition 7 for a precise version of this remark).

The next lemma states the strong convexity of these divergences with respect to the  $L^1(\tau)$  norm, which is needed in the next section. It is a generalization of Pinsker inequality, recovered when  $K = 1$  and with  $\eta_{\text{ent}}$ . Notice that when  $p < 2$ , the bound worsens as the norm increases.



■ **Figure 2** Divergence-generating functions (plain) and their second-order derivatives (dashed).

► **Lemma 3** (Strong convexity of Bregman divergences). *Assume that  $f, g$  have  $L^1(\tau)$ -norm bounded by  $K$ . Then for  $p \in ]1, 2]$ ,*

$$D_{\bar{\eta}_p}(f, g) \geq \frac{K^{p-2}}{2} \|f - g\|_{L^1(\tau)}^2.$$

The inequality also holds for  $\eta_{\text{ent}}$  with  $p = 1$  (assuming  $f, g \geq 0$ ). Finally for  $\eta_{\text{hyp}}$ , it holds

$$D_{\bar{\eta}_{\text{hyp}}}(f, g) \geq \frac{(K + \beta)^{-1}}{2} \|f - g\|_{L^1(\tau)}^2.$$

**Proof.** For  $\beta > 0$  and  $p \in [1, 2]$ , consider the function  $\eta_{p,\beta} : \mathbb{R} \rightarrow \mathbb{R}_+$  satisfying  $\eta_{p,\beta}(0) = \eta'_{p,\beta}(0) = 0$  and  $\eta''_{p,\beta}(s) = |s^2 + \beta^2|^{\frac{p-2}{2}}$ . This function is smooth and, for  $p > 1$ , converges monotonously from below to  $\eta_p$  as  $\beta \rightarrow 0$  (remember that  $\eta_p$  satisfies  $\eta_p(0) = \eta'_p(0) = 0$  and  $\eta''_p(s) = |s|^{p-2} \geq \eta''_{p,\beta}(s)$ ). Our first step is to prove a Pinsker-like inequality for the Bregman divergence  $D_{\bar{\eta}_{p,\beta}}$ . For  $f, g \in L^1(\tau)$  such that  $\|g\|_{L^1(\tau)} = 1$ , it holds by the Cauchy–Schwarz inequality

$$\begin{aligned} 1 &= \left( \int |g| d\tau \right)^2 = \left( \int |g| |f^2 + \beta^2|^{\frac{p-2}{4}} |f^2 + \beta^2|^{\frac{2-p}{4}} d\tau \right)^2 \\ &\leq \left( \int |g|^2 |f^2 + \beta^2|^{\frac{p-2}{2}} d\tau \right) \left( \int |f^2 + \beta^2|^{\frac{2-p}{2}} d\tau \right). \end{aligned}$$

But since  $2 - p \geq 0$ , we have  $|f^2 + \beta^2|^{\frac{2-p}{2}} \leq (|f| + \beta)^{2-p}$ . Thanks to Jensen's inequality for concave functions (we use  $2 - p \leq 1$ ), this leads to

$$\int |f^2 + \beta^2|^{\frac{2-p}{2}} d\tau \leq \int (|f| + \beta)^{2-p} d\tau \leq (\|f\|_{L^1} + \beta)^{2-p}.$$

Combining these inequalities and by homogeneity, it follows that if  $\|f\|_{L^1(\tau)} \leq K$ , then  $\forall g \in L^1(\tau)$  it holds

$$\int |g|^2 |f^2 + \beta^2|^{\frac{p-2}{2}} d\tau \geq (K + \beta)^{p-2} \|g\|_{L^1(\tau)}^2.$$

This equation shows that  $f \mapsto \bar{\eta}_{p,\beta}(f) := \int \eta_{p,\beta}(f) d\tau$  is  $(K + \beta)^{p-2}$ -strongly convex for the  $L^1(\tau)$ -norm over the  $L^1(\tau)$ -ball of radius  $K$ , see e.g. [44, Thm. 3]. This means that for all  $f, g \in L^1(\tau)$  of norm smaller than  $K$ , it holds

$$\lambda \bar{\eta}_{p,\beta}(f) + (1 - \lambda) \bar{\eta}_{p,\beta}(g) \geq \bar{\eta}_{p,\beta}(\lambda f + (1 - \lambda)g) + \frac{(K + \beta)^{p-2}}{2} \lambda(1 - \lambda) \|f - g\|_{L^1(\tau)}^2, \quad \forall \lambda \in [0, 1].$$

By the monotone convergence theorem, we have when  $p > 1$  that  $\lim_{\beta \rightarrow 0} \bar{\eta}_{p,\beta}(f) \rightarrow \bar{\eta}_p(f)$  and so strong convexity also holds for  $\bar{\eta}_p$ , taking the pointwise limit of the strong convexity inequality as  $\beta \rightarrow 0$ .

It follows [44, Thm. 1] that, over this ball of functions, the Bregman divergence  $D_{\bar{\eta}_{p,\beta}}$  (for  $p \in ]1, 2]$  and  $\beta > 0$  or for  $p = 1$  and  $\beta > 0$ ) satisfies the Pinsker-like inequality

$$D_{\bar{\eta}_{p,\beta}}(f, g) \geq \frac{(K + \beta)^{p-2}}{2} \|f - g\|_{L^1(\tau)}^2. \quad (4)$$

Specializing to  $p = 1$  and  $\beta > 0$  proves the Lemma for  $\eta_{\text{hyp}}$  and specializing to  $p > 1$  and  $\beta = 0$  proves the Lemma for  $\eta_p$  with  $p \in ]1, 2]$ . It only remains to prove the case  $\eta_{\text{ent}}$ , i.e. the classical Pinsker inequality. Note

that here we cannot take the limit  $\beta \rightarrow 0$  because  $\eta_{1,\beta}$  does not converge to  $\eta_{\text{ent}}$  (in fact it diverges, except at  $s = 0$ ). One way to recover this case in an analogous way, is to define instead  $\eta_{1,\beta} : \mathbb{R}_+ \rightarrow \mathbb{R}_+$  as the function satisfying  $\eta_{1,\beta}(1) = \eta'_{1,\beta}(1) = 0$  and  $\eta''_{1,\beta}(s) = |s^2 + \beta^2|^{-1/2}$ . This function is smooth and converges monotonously to  $\eta_{\text{ent}}$  as  $\beta \rightarrow 0$  and  $\bar{\eta}_{1,\beta}$  is  $(K + \beta)^{-1}$  strongly convex. Thus we recover Pinsker's inequality with an analogous argument in the limit  $\beta \rightarrow 0$ . ◀

### 3.3 Gradient methods and their classical guarantees

We now detail two classical algorithms that enjoy guarantees of the form Eq. (2) for a large class of composite optimization problems. Algorithm 1 (PGM) is closely related to mirror descent [34] and is discussed in [2, 7]. Algorithm 2 (APGM) is taken from [40] who presents it as a generalization of [2] itself an extension of Nesterov's second method [35]. For the sake of concreteness, we instantiate these algorithms in the context of optimization in the space of measures, where small adaptations have to be made.

---

#### Algorithm 1: (Bregman) Proximal Gradient Method (PGM)

---

**Initialization:**  $f_0 \in \text{dom } H$ , *step-size*  $s > 0$   
**for**  $k=0,1,\dots$  **do**  
  |  $f_{k+1} = \arg \min_f \{G(f_k) + \int G'[f_k](f - f_k)d\tau + H(f) + \frac{1}{s}D_{\bar{\eta}}(f, f_k)\}$   
**end**  
**Output:**  $f_{k+1}$

---



---

#### Algorithm 2: Accelerated (Bregman) Proximal Gradient Method (APGM)

---

**Initialization:**  $f_0 = h_0 \in \text{dom } H$ ,  $\gamma_0 = 1$ , *step-size*  $s > 0$   
**for**  $k=0,1,\dots$  **do**  
  |  $g_k = (1 - \gamma_k)f_k + \gamma_k h_k$   
  |  $h_{k+1} = \arg \min_f \{G(g_k) + \langle \nabla G(g_k), f - g_k \rangle + H(f) + \frac{\gamma_k}{s}D_{\bar{\eta}}(f, h_k)\}$   
  |  $f_{k+1} = (1 - \gamma_k)f_k + \gamma_k h_{k+1}$   
  |  $\gamma_{k+1} = \frac{1}{2}(\sqrt{\gamma_k^4 + 4\gamma_k^2 - \gamma_k^2})$   
**end**  
**Output:**  $f_{k+1}$

---

In the next proposition, we verify that the updates are well-defined under suitable assumptions. Table 1 lists some update formulas which are directly implementable, after discretization.

► **Proposition 4** (Well-defined updates). *Assume (A1) and (A2). If  $\eta'(h_k) \in L^\infty(\tau)$  and  $g_k \in L^1(\tau)$ , then there exists a unique solution  $h_{k+1} \in L^1(\tau)$  to the optimization problem*

$$\min_{f \in L^1(\tau)} \int G'[g_k]f d\tau + H(f) + \frac{1}{s} \left( \int \eta(f) d\tau - \int \eta'(h_k) f d\tau \right)$$

which moreover satisfies  $\eta'(h_{k+1}) \in L^\infty(\tau)$ . It is characterized by the fact that there exists  $\phi \in \partial H(h_{k+1}) \subset L^\infty(\tau)$  such that

$$G'[g_k] + \frac{1}{s}(\eta'(h_{k+1}) - \eta'(h_k)) + \phi = 0. \quad (5)$$

**Proof.** Let  $J_k$  be the function to minimize. Thanks to our assumptions that  $\eta'(\text{int}(\text{dom } \eta)) = \mathbb{R}$  and since  $H$  is lower-bounded, by [39, Cor. 2B], the sublevels of  $J_k$  are compact with respect to the weak topology (induced on  $L^1(\tau)$  by  $L^\infty(\tau)$ ). Moreover,  $J_k$  is convex and l.s.c. for the same topology; in particular because  $G'[g_k], \eta'(h_k) \in L^\infty(\tau)$  and for the term  $\int \eta(f) d\tau$ , this follows from [39, Cor. 2A]. Thus by the direct method of the calculus of variations, there exists a minimizer  $f_{k+1} \in L^1(\tau)$ . Since  $\eta$  is strictly convex, so is  $J_k$  and this minimizer is unique. The condition of Eq. (5) is always a sufficient optimality condition since, by the subdifferential inclusion rule, it implies that  $0 \in \partial J_k(f)$ . It thus remain to show that it is also necessary, in which case the property  $\eta'(h_{k+1}) \in L^\infty(\tau)$  immediately follows.

■ **Table 1** Update step  $h_{k+1}$  for APGM, and PGM (when  $f_k = g_k = h_k$ ) where  $\text{sft}_{\kappa}(a) = \text{sign}(a)(|a| - \kappa a)$  is a soft-thresholding. In (ii),  $\kappa \in \mathbb{R}$  is the *unique* number such that the update satisfies the constraint. In (iv),  $\kappa \geq 0$  is the *smallest* number such that the update satisfies the constraint (see [18] for efficient algorithms to compute  $\kappa$  in practice).

$\bar{H} \backslash \eta$	Assumption (A2) $_{\pm}$ (dom $\eta = \mathbb{R}$ )	Assumption (A2) $_{+}$ (dom $\eta = \mathbb{R}_{+}$ )
(i) $\iota_{\mathcal{M}_{+}(\Theta)} + \lambda \ \cdot\ $	$[\eta']^{-1}(\eta'(h_k) - sG'[g_k] - s\lambda)_{+}$	$[\eta']^{-1}(\eta'(h_k) - sG'[g_k] - s\lambda)$
(ii) $\iota_{\mathcal{P}(\Theta)}$	$[\eta']^{-1}(\eta'(h_k) - sG'[g_k] - \kappa)_{+}$	$[\eta']^{-1}(\eta'(h_k) - sG'[g_k] - \kappa)$
(iii) $\lambda \ \cdot\ $	$[\eta']^{-1}\left(\text{sft}_{\lambda s}(\eta'(h_k) - sG'[g_k])\right)$	$[\eta']^{-1}(\eta'(h_k) - sG'[g_k] - s\lambda)$
(iv) $\iota_{\{\mu; \ \mu\  \leq K\}}$	$[\eta']^{-1}\left(\text{sft}_{\kappa}(\eta'(h_k) - sG'[g_k])\right)$	$[\eta']^{-1}(\eta'(h_k) - sG'[g_k] - \kappa)$

This is done on a case by case basis for the functions  $\bar{H}$  admissible under Assumption (A1). Consider for instance the nonnegativity constraint  $\bar{H} = \iota_{\mathcal{M}_{+}(\Theta)}$  and  $\eta$  that satisfies Assumption (A2) $_{\pm}$ . Then with the update  $h_{k+1}$  given in Table 1 (take  $\lambda = 0$ ), it holds

$$\phi := \frac{1}{s}\eta'(h_k) - G'[g_k] - \frac{1}{s}\eta'(h_{k+1}) = \min\{0, \frac{1}{s}\eta'(h_k) - G'[g_k]\}.$$

Clearly  $\phi \in L^{\infty}(\tau)$ ,  $\phi \leq 0$  and  $\int \phi h_{k+1} d\tau = 0$  and thus  $\phi \in \partial H(h_{k+1})$ , which shows that  $h_{k+1}$  is a minimizer and satisfies Eq. (5). The other cases for  $\bar{H}$  and  $\eta$  that are admissible under (A1) and (A2) (such as those listed in Table 1) can be treated similarly and follow computations which are standard in the finite dimensional setting. ◀

Let us now recall the guarantees for these methods. We stress that, as discussed in Section 2, these guarantees do not necessarily lead to convergence rates.

► **Proposition 5.** *Assume (A1) and (A2) and that  $\eta$  satisfies the conclusion of Lemma 3 for some  $p \in [1, 2]$ ,  $\beta \geq 0$ . Consider an initialization  $f_0 \in \text{dom } H$  such that  $\eta'(f_0) \in L^{\infty}(\tau)$  and, for some  $K \geq \|f_0\|_{L^1(\tau)}$ , a step-size  $s \leq (K + \beta)^{p-2} (\|\Phi\|_{\infty}^2 \text{Lip}(\nabla R))^{-1}$ .*

i. *Let  $(f_k)_{k \in \mathbb{N}}$  be generated by Algorithm 1 (PGM). If  $\sup_k \|f_k\|_{L^1(\tau)} \leq K$ , then Eq. (2) holds with  $\alpha_k = 1/(sk)$ , i.e.*

$$F(f_k) - F(f) \leq \frac{1}{sk} D_{\bar{\eta}}(f, f_0), \quad \forall f \in L^1(\tau), \forall k \geq 1.$$

ii. *Let  $(f_k, g_k, h_k)_{k \in \mathbb{N}}$  be generated by Algorithm 2 (APGM). If  $\sup_k \|h_k\|_{L^1(\tau)} \leq K$ , then Eq. (2) holds for  $f_k$  with  $\alpha_k = 4/((k+1)^2 s)$ , i.e.*

$$F(f_k) - F(f) \leq \frac{4}{s(k+1)^2} D_{\bar{\eta}}(f, f_0), \quad \forall f \in L^1(\tau), \forall k \geq 1.$$

Moreover, it holds  $0 < \gamma_k \leq 1$  and  $\gamma_k \leq 2/(k+2) \forall k \geq 0$ .

**Proof.** By Proposition 4, the updates are well-defined. The proof of [40, Thm. 1] goes through, in particular thanks to Lemma 2 (smoothness) and since  $D_{\eta}/(K + \beta)^{p-2}$  is 1-strongly convex with respect to  $\|\cdot\|_{L^1(\tau)}$  whenever this property is needed in the proof. A particularly simple exposition of the proof for APGM can be found in [19, Thm. 4.24]. ◀

► **Remark 6.** A difficulty in Proposition 5 is that when  $p < 2$ , one needs to assume *a priori* bounds on the  $L^1$ -norm of certain iterates to obtain convergence guarantees, because the metric induced by the divergence  $D_{\bar{\eta}}$  becomes weaker as the  $L^1$ -norm increases. Since Algorithm 1 (PGM) is a descent method,  $\|f_k\|_{L^1}$  is bounded, uniformly in  $k$ , as soon as the objective is coercive for the  $L^1$ -norm. But for Algorithm 2 (APGM), even if variants exist where  $(F(f_k))_k$  is monotonous [19], this property does not seem to be sufficient to control  $\|h_k\|_{L^1}$ , even for coercive objectives. Of course, uniform bounds are always trivially satisfied when  $\bar{H}$  includes the constraint  $\iota_{\|\mu\| \leq K}$  or  $\iota_{\mathcal{P}(\Theta)}$ .

### 3.4 Reparameterized gradient descent as a Bregman descent

In this paragraph, we recall a link between Bregman gradient descent, a.k.a. mirror descent (an instance of Algorithm 1) and  $L^2$  gradient descent dynamics on certain *reparameterized* objectives. The purpose is to show that the convergence rates proved in Section 4 with  $\eta_{\text{ent}}$  and  $\eta_{\text{hyp}}$  are also relevant to understand  $L^2$  gradient descent in certain contexts. While these remarks are well-known [1, 3, 41], we find it instructive to state them clearly in our context. In order to reduce the discussion to its simplest setting, we consider the continuous time dynamics in the unregularized setting, and we assume that they are well-defined.

The optimality conditions of the update of Algorithm 1 (see Proposition 4) can be written as  $\eta'(f_{k+1}) - \eta'(f_k) = -sF'[f_k]$ . As the step-size  $s$  vanishes, this leads to a continuous trajectory  $(f_t)_{t \geq 0}$ , which we refer to as the  $\eta$ -mirror flow, that solves

$$\frac{d}{dt}\eta'(f_t) = -F'[f_t] \iff \frac{d}{dt}f_t = -[\eta''(f_t)]^{-1}F'[f_t]. \quad (6)$$

► **Proposition 7** (Reparameterized mirror flows as gradient flows).

- i. (Square parameterization). Let  $(f_t)_{t \geq 0}$  be the  $L^2(\tau)$  gradient flow of  $\widehat{F} : f \mapsto F(f^2)$  initialized such that  $\log(f_0) \in L^\infty(\tau)$ . Then  $h_t := f_t^2$  is the  $\eta_{\text{ent}}$ -mirror flow of  $4F$ .
- ii. (Difference of squares parameterization). Let  $(f_t, g_t)_{t \geq 0}$  be the  $(L^2(\tau))^2$  gradient flow of  $\widehat{F} : (f, g) \mapsto F(f^2 - g^2)$  initialized such that  $\log(f_0 g_0) \in L^\infty(\tau)$ . Then  $h_t := f_t^2 - g_t^2$  is the  $\eta_{\text{hyp}}$ -mirror flow of  $4F$  with parameter  $\beta = 2f_0 g_0$  (here  $\beta$  is function instead of a scalar).

We can make the following remarks:

- Combining (i) and (ii), we find that if  $(h_t^+, h_t^-)_{t \geq 0}$  follows a  $\eta_{\text{ent}}$ -mirror flow for  $(h^+, h^-) \mapsto F(h^+ - h^-)$  then  $h_t^+ - h_t^-$  is a  $\eta_{\text{hyp}}$ -mirror flow for  $F$ . This confirms the interpretation of  $\eta_{\text{hyp}}$  as a “signed” version of the entropy (see also [25, Thm. 23]).
- These exact equivalences are lost in discrete time, with an error term that scales as the squared step-size. It is thus difficult to convert the most efficient guarantees for (Bregman) PGM into guarantees with the same convergence rate for gradient descent.

**Proof.**

- i. The  $L^2(\tau)$  gradient flow of  $\widehat{F}$  satisfies  $\frac{d}{dt}f_t = -\widehat{F}'[f_t^2] = -2f_t F'[f_t^2]$ . Thus the function  $h_t := f_t^2$  evolves according to

$$\frac{d}{dt}h_t = 2f_t \frac{d}{dt}f_t = -4f_t^2 F'[f_t^2] = -4h_t F'[h_t]$$

which is precisely the  $\eta_{\text{ent}}$ -mirror flow of  $4F$  since  $\eta_{\text{ent}}''(s) = s^{-1}$  for  $s > 0$ .

- ii. The  $(L^2(\tau))^2$  gradient flow of  $\widehat{F}$  satisfies  $\frac{d}{dt}f_t = -2f_t F'[f_t^2 - g_t^2]$  and  $\frac{d}{dt}g_t = 2g_t F'[f_t^2 - g_t^2]$ . As a consequence, we have for  $h_t := f_t^2 - g_t^2$  and  $\tilde{h}_t := f_t^2 + g_t^2$  that  $\frac{d}{dt}h_t = -4\tilde{h}_t F'[h_t]$ . To conclude, it remains to show that  $\tilde{h}_t = [\eta_{\text{hyp}}''(h_t)]^{-1} = (h_t^2 + \beta^2)^{1/2}$  for some  $\beta > 0$ . To prove this, observe that

$$\frac{d}{dt}(\tilde{h}_t^2 - h_t^2) = \frac{d}{dt}(4f_t^2 g_t^2) = 8f_t' f_t g_t^2 + 8g_t' g_t f_t^2 = 0.$$

Hence  $\tilde{h}_t^2 - h_t^2 = \tilde{h}_0^2 - h_0^2 = 4f_0^2 g_0^2$ , which proves that  $\tilde{h}_t = (h_t^2 + \beta^2)^{1/2}$  with  $\beta = 2f_0 g_0$ . ◀

## 4 Upper bounds on the convergence rates

This section contains the main result of this paper which is Theorem 8 and summarized in Table 2. As discussed in Section 2 and thanks to Proposition 5, in order to derive convergence rates for Algorithms 1 and 2 it is sufficient to control the function

$$\psi(\alpha) := \inf_{f \in L^1(\tau)} \left\{ F(f) - \inf F + \alpha D_\eta(f, f_0) \right\}. \quad (7)$$

For the class of problems we consider, the behavior of  $\psi$  highly depends on the context. The simplest situation is when  $F$  admits a minimizer  $f^* \in L^q(\tau)$  with  $q > 1$  (since  $\tau$  is finite, it holds  $L^q(\tau) \subset L^1(\tau)$ ). Then for the distance-generating functions  $\eta_p$  for  $1 < p \leq q$  or  $\eta_{\text{hyp}}$ , it is easy to see that  $D_{\tilde{\eta}}(f^*, f_0) < +\infty$  (this further

■ **Table 2** (a) Upper bounds on the convergence rates of  $F(f_k) - \inf F$  for Algorithm 1 (PGM) and 2 (APGM). (b) The value of  $q$  that appears in the rate depends on the regularity of  $\Phi$  and on whether  $\bar{G}'$  vanishes at optimality or not. This defines 4 settings referred to as (I), (I\*), (II) and (II\*). Upper bounds are derived in Theorem 8, lower bounds are proved in Section 5.

(a) Convergence rates			(b) Value of $q$		
	PGM	APGM		$\Phi$ Lipschitz	$\nabla_{\theta}\Phi$ Lipschitz
$\eta_{\text{ent}}, \eta_{\text{hyp}}$	$\log(k)k^{-1}$	$\log(k)k^{-2}$	$\bar{G}'[\mu^*]$ arbitrary	$q = 1$ (I)	$q = 2$ (II)
$\eta_p, p > 1$	$k^{-\frac{q}{(p-1)d+q}}$	$k^{-\frac{2q}{(p-1)d+q}}$	$\bar{G}'[\mu^*] = 0$	$q = 2$ (I*)	$q = 4$ (II*)

requires  $f^* \geq 0$  for  $\eta_{\text{ent}}$ ). Thus by Proposition 1, it holds  $\psi(\alpha) \lesssim \alpha$  and the convergence rates  $(\alpha_k)$  given in Proposition 5 are preserved.

In the more subtle case where the minimizer of  $\bar{F}$  is only assumed to be in  $\mathcal{M}(\Theta)$ , the variety of behaviors is captured by the following result.

► **Theorem 8.** *Under Assumptions (A1) and (A2), let  $f_0$  be such that  $\eta'(f_0) \in L^\infty(\tau)$ . Assume that there exists  $\mu^* \in \mathcal{M}(\Theta)$  such that  $\bar{F}(\mu^*) = \inf F$ . Under setting (A2)<sub>+</sub> (i.e.  $\text{dom } \eta = \mathbb{R}_+$ ) assume moreover  $\mu^* \in \mathcal{M}_+(\Theta)$ . Then it holds*

$$\psi(\alpha) \lesssim \inf_{\epsilon > 0} \epsilon^q + \alpha \epsilon^d \eta(\epsilon^{-d}). \quad (8)$$

where  $q \in \{1, 2, 4\}$  is determined by Table 2(b) (the largest  $q$ , the strongest the bound). Namely, the bound holds:

- with  $q = 1$  if  $\Phi$  is Lipschitz continuous;
- with  $q = 2$  if  $\Phi$  is Lipschitz continuous and  $\bar{G}'[\mu^*] = 0$ , or if  $\nabla_{\theta}\Phi$  is Lipschitz continuous;
- with  $q = 4$  if  $\nabla_{\theta}\Phi$  is Lipschitz continuous and  $\bar{G}'[\mu^*] = 0$ .

Given the bound of Eq. (8), it is straightforward to compute the rates given in Table 2. The exponent  $q \in \{1, 2, 4\}$  can be interpreted as follows: it is such that  $F(\mu_\epsilon) - F(\mu^*) \lesssim \epsilon^q$  where  $\mu_\epsilon$  is the convolution of  $\mu^*$  with a box kernel of radius  $\epsilon > 0$ . An asymptotic analysis leads to lower bounds for this exponent under several assumptions (Table 2(b)), but in practice, non asymptotic effects may play an important role (see experiments in Section 6.2).

► **Remark 9 (Additive vs. Multiplicative updates).** A consequence of Theorem 8 is that algorithms with “additive updates” (obtained with  $\eta_2$  as a distance-generating function, e.g. ISTA, FISTA) suffer from the “curse of dimensionality in the convergence rates, see Table 2(a). In comparison, algorithms with “multiplicative updates” (obtained with  $\eta_{\text{ent}}$  or  $\eta_{\text{hyp}}$  as a distance-generating function) always converge at a faster rate which is independent of the dimension  $d$ . Note that Theorem 8 only proves upper bounds on the rates, but we will see that they are tight in Section 5.

**Proof.** The upper-bound in Eq. (8) corresponds to an upper bound on  $F(f_\epsilon) - \inf F + \alpha D_{\bar{\eta}}(f_\epsilon, f_0)$  for a specific family of candidates  $f_\epsilon \in L^1(\tau)$ . A special case of this argument for sparse  $\mu^*$ ,  $\eta_{\text{ent}}$  and  $q = 1$  appeared in [16] and extended to  $\mu^* \in \mathcal{M}_+(\Theta)$  in [23]. In the following we write  $F, G, H$  for  $\bar{F}, \bar{G}, \bar{H}$  to lighten notations.

**Step 1. Smoothing with a box kernel.** For  $\epsilon > 0$  (smaller than the injectivity radius of the exponential map in  $\Theta$ ) consider the transition kernel  $(\gamma_{\theta, \epsilon})_{\theta \in \Theta}$  where  $\gamma_{\theta, \epsilon} \in \mathcal{P}(\Theta)$  is proportional to the restriction of  $\tau$  to the closed geodesic ball  $B_\epsilon(\theta)$  of radius  $\epsilon$  centered at  $\theta$ . We define  $\gamma_\epsilon \in \mathcal{M}(\Theta^2)$  as  $\gamma_\epsilon(d\theta, d\theta') := \gamma_{\theta, \epsilon}(d\theta')\mu^*(d\theta)$ . By construction, the first marginal of  $\gamma_\epsilon$  is  $\mu^*$  and we call  $\mu_\epsilon$  its second marginal, which is absolutely continuous with respect to  $\tau$  with density  $f_\epsilon = \frac{d\mu_\epsilon}{d\tau}$ . Since  $\frac{d\gamma_{\theta, \epsilon}}{d\tau}(\theta') = \tau(B_\epsilon(\theta))^{-1} \mathbb{1}_{B_\epsilon(\theta)}(\theta')$ , it holds

$$\mu_\epsilon(d\theta') = \int_{\Theta} \gamma_{\theta, \epsilon}(d\theta')\mu^*(d\theta), \quad \text{and} \quad f_\epsilon(\theta') = \int_{\Theta} \frac{\mathbb{1}_{B_{\theta, \epsilon}}(\theta')}{\tau(B_\epsilon(\theta))} d\mu^*(\theta).$$

Note that  $\tau(B_\epsilon(\theta)) \asymp \epsilon^d$ , see e.g. [26, Thm.3.3].

**Step 2. Bounding  $F(\mu_\epsilon) - F(\mu^*)$ .** For our admissible regularizers, it is easy to verify that  $H(\mu_\epsilon) \leq H(\mu^*)$ . By convexity of  $G$ , we have

$$F(\mu_\epsilon) - F(\mu^*) \leq G(\mu_\epsilon) - G(\mu^*) \leq \int G'[\mu_\epsilon] d[\mu_\epsilon - \mu^*] = \int_{\Theta} (G'[\mu_\epsilon](\theta) - G'[\mu_\epsilon](\theta')) d\gamma(\theta, \theta').$$

It is clear that the magnitude and regularity of  $G'[\mu_\epsilon]$  plays a role in the magnitude of this quantity. To go further, let us consider the various cases in Table 2 (b) successively.

(I). If  $\Phi$  is Lipschitz then for  $\theta, \theta' \in \Theta$  it holds

$$|G'[\mu_\epsilon](\theta) - G'[\mu_\epsilon](\theta')| \leq \left\| \nabla R \left( \int \Phi d\mu_\epsilon \right) \right\|_{\mathcal{F}} \cdot \text{Lip}(\Phi) \cdot \text{dist}(\theta, \theta').$$

Since  $\nabla R$  is Lipschitz continuous, we deduce that there exists  $K > 0$  such that  $|G'[\mu_\epsilon](\theta) - G'[\mu_\epsilon](\theta')| \leq K \text{dist}(\theta, \theta')$ . It follows

$$\begin{aligned} F(\mu_\epsilon) - F(\mu^*) &\leq K \int_{\Theta \times \Theta} \text{dist}(\theta, \theta') d\gamma(\theta, \theta') \\ &= K \int_{\Theta} \frac{1}{\tau(B_\epsilon(\theta))} \int_{B_\epsilon(\theta)} \text{dist}(\theta, \theta') d\tau(\theta') d\mu^*(\theta) \leq K \cdot \epsilon \cdot \|\mu^*\|. \end{aligned}$$

(I\*). If  $\Phi$  is Lipschitz and moreover  $G'[\mu^*] = 0$ , it holds

$$G'[\mu_\epsilon](\theta) = \left\langle \nabla R \left( \int \Phi d\mu_\epsilon \right) - \nabla R \left( \int \Phi d\mu^* \right), \Phi(\theta) \right\rangle_{\mathcal{F}}.$$

Since  $\nabla R$  is Lipschitz continuous, the first factor is bounded by  $\text{Lip}(\nabla R) \cdot \|\int \Phi d[\mu_\epsilon - \mu^*]\| = \text{Lip}(\nabla R) \cdot \sup_{\|\tilde{\Phi}\| \leq 1} \int \langle \tilde{\Phi}, \Phi(\theta) \rangle d[\mu_\epsilon - \mu^*](\theta)$ . Under assumption (I\*), the functions  $\{\theta \mapsto \langle \tilde{\Phi}, \Phi(\theta) \rangle ; \|\tilde{\Phi}\| \leq 1\}$  are uniformly Lipschitz, so by the reasoning above, it follows that  $\|\int \Phi d[\mu_\epsilon - \mu^*]\| \lesssim \epsilon$ , so  $\text{Lip}(G'[\mu_\epsilon]) \lesssim \epsilon$ . It follows that the previous bound improves to  $F(\mu_\epsilon) - F(\mu^*) \lesssim \epsilon^2$ .

(II). Here we have that  $G'[\mu_\epsilon] \in \mathcal{C}^1(\Theta; \mathbb{R})$  and  $\nabla_\theta G'[\mu_\epsilon]$  is Lipschitz. By the Mean Value Theorem on Riemannian manifolds (see e.g. [27, Thm. 4.6]), there exists a constant  $K' \geq 0$  such that for all  $\theta \in \Theta$

$$\left| \frac{1}{\tau(B_\epsilon(\theta))} \int_{B_\epsilon(\theta)} G'[\mu_\epsilon](\theta') d\tau(\theta') - G'[\mu_\epsilon](\theta) \right| \leq K' \epsilon^2.$$

It follows

$$F(\mu_\epsilon) - F(\mu^*) \leq \int d\mu^*(\theta) \int (G'[\mu_\epsilon](\theta) - G'[\mu_\epsilon](\theta')) d\gamma_{\theta, \epsilon}(\theta') \leq K' \|\mu^*\| \epsilon^2.$$

(II\*). If  $\nabla \Phi$  is Lipschitz and moreover  $G'[\mu^*] = 0$  then an improvement as in (I\*) applies. The functions  $\{\theta \mapsto \langle \tilde{\Phi}, \Phi(\theta) \rangle ; \|\tilde{\Phi}\| \leq 1\}$  are differentiable with a uniformly Lipschitz derivative so arguments as in the previous paragraph show that  $\|\int \Phi d[\mu_\epsilon - \mu^*]\| \lesssim \epsilon^2$ , so  $\text{Lip}(\nabla_\theta G'[\mu_\epsilon]) \lesssim \epsilon^2$ . Thus, going through the argument for (II), with all the constants multiplied by  $\epsilon^2$ , we obtain that the bound improves to  $F(\mu_\epsilon) - F(\mu^*) \lesssim \epsilon^4$ .

**Step 3. Bounding  $D_{\tilde{\eta}}(f_\epsilon, f_0)$ .** It holds  $D_{\tilde{\eta}}(f_\epsilon, f_0) = \int \eta(f_\epsilon) d\tau - \int \eta(f_0) d\tau - \int \eta'(f_0)(f_\epsilon - f_0) d\tau$ . Since  $\eta'(f_0) \in L^\infty$ , all these terms are bounded by a constant independent of  $\epsilon$  except  $\int \eta(f_\epsilon) d\tau$ , so it remains to bound the latter. If  $\mu^* = 0$  then this quantity is bounded by a constant independent of  $\epsilon$  and we are done. Otherwise let us first assume that  $\mu^* \in \mathcal{M}_+(\Theta)$ . By Jensen's inequality, one has  $\forall \theta' \in \Theta$

$$\eta(f_\epsilon(\theta')) = \eta \left( \int_{\Theta} \frac{\mu^*(\Theta) \mathbb{1}_{B_\epsilon(\theta)}(\theta')}{\tau(B_\epsilon(\theta))} \frac{d\mu^*(\theta)}{\mu^*(\Theta)} \right) \leq \int_{\Theta} \eta \left( \frac{\mu^*(\Theta) \mathbb{1}_{B_\epsilon(\theta)}(\theta')}{\tau(B_\epsilon(\theta))} \right) \frac{d\mu^*(\theta)}{\mu^*(\Theta)}.$$

It follows, by Fubini's theorem,

$$\begin{aligned} \int \eta(f_\epsilon(\theta')) d\tau(\theta') &\leq \int_{\Theta} \frac{d\mu^*(\theta)}{\mu^*(\Theta)} \int_{\Theta} d\tau(\theta') \eta \left( \frac{\mu^*(\Theta) \mathbb{1}_{B_\epsilon(\theta)}(\theta')}{\tau(B_\epsilon(\theta))} \right) \\ &\leq \left( \sup_{\theta} \tau(B_\epsilon(\theta)) \left| \eta \left( \frac{\mu^*(\Theta)}{\tau(B_\epsilon(\theta))} \right) \right| \right) + |\eta(0)| \end{aligned}$$

Now we use the fact that  $\sup_{\theta} \tau(B_\epsilon(\theta)) \asymp \epsilon^d$  and our assumption that  $\eta(cx) \asymp \eta(x)$  for any fixed  $c > 0$  as  $x \rightarrow \infty$  to bound this quantity by  $O(\epsilon^d \eta(\epsilon^{-d}))$ .

For the general case where  $\mu^* \in \mathcal{M}(\Theta)$  let  $\mu^*_+, \mu^*_- \in \mathcal{M}_+(\Theta)$  be the Jordan decomposition of  $\mu^* = \mu^*_+ - \mu^*_-$ . It holds  $|f_\epsilon(\theta)| \leq \max\{f_{+, \epsilon}(\theta), f_{-, \epsilon}(\theta)\}$  where  $f_{+, \epsilon}$  and  $f_{-, \epsilon}$  are obtained by applying the smoothing procedure of Step. 1 to  $\mu^*_+$  and  $\mu^*_-$  respectively. Using the fact that under Assumption (A2) $_{\pm}$ ,  $\eta$  is even and increasing on  $\mathbb{R}_+$ , we obtain the bound

$$\int \eta(f_\epsilon) d\tau \leq \int \eta(f_{+, \epsilon}) d\tau + \int \eta(f_{-, \epsilon}) d\tau.$$

We finally bound each of these terms as when  $\mu^* \in \mathcal{M}_+(\Theta)$  to conclude.  $\blacktriangleleft$

## 5 Lower bounds

We will consider two types of lower bounds: (i) lower bounds on  $\psi(\alpha)$  in order to confirm that the analysis in Theorem 8 is tight, and (ii) direct lower bounds on the convergence rates of Algorithms 1 and 2. Of course, the latter imply the former, but studying  $\psi$  directly has its own interest and makes it simpler to cover all the cases.

### 5.1 Tight lower bounds on $\psi$

Let us show that the bounds on  $\psi$  in Theorem 8 cannot be improved without additional assumptions.

► **Proposition 10** (Lower-bounds). *For each of the 4 settings of Table 2(b) (determining the value of  $q \in \{1, 2, 4\}$ ), there exists an objective function  $\bar{F}$  satisfying Assumption (A1) and  $\inf \bar{F} = \bar{F}(\mu^*)$  with  $\mu^* \in \mathcal{M}_+(\Theta)$ , such that for any distance-generating function  $\eta$  satisfying Assumption (A2) and  $f_0$  such that  $\eta'(f_0) \in L^\infty(\tau)$ , it holds*

$$\psi(\alpha) \gtrsim \inf_{\epsilon > 0} \epsilon^q + \alpha \epsilon^d \eta(\epsilon^{-d}). \quad (9)$$

**Proof.** Let us build explicit objective functions with  $\Theta = \mathbb{T}^d$  the  $d$  dimensional torus,  $\mathcal{F} = \mathbb{R}$  and  $\bar{H} = \nu_{\mathcal{P}(\Theta)}$ . Let  $\theta_0 \in \mathbb{T}^d$  and for  $f \in L^1(\Theta)$  such that  $f\tau \in \mathcal{P}(\Theta)$ , let  $\epsilon_f > 0$  be such that  $\int_{B_{\epsilon_f}(\theta_0)} f d\tau = \frac{1}{2}$ . Such an  $\epsilon_f$  exists because  $\epsilon \mapsto \int_{B_\epsilon(\theta_0)} f d\tau$  continuously interpolates between 0 when  $\epsilon = 0$  and 1 when  $\epsilon$  is large. For any  $\eta$  satisfying Assumption (A2), using the fact that  $\eta \geq 0$  and Jensen's inequality, it holds

$$\int \eta(f) d\tau \geq \int_{B_{\epsilon_f}(\theta_0)} \eta(f) d\tau \geq \tau(B_{\epsilon_f}(\theta_0)) \eta\left(\frac{1}{\tau(B_{\epsilon_f}(\theta_0))} \int_{B_{\epsilon_f}(\theta_0)} f d\tau\right) \gtrsim \epsilon_f^d \eta(\epsilon_f^{-d}).$$

Thus for any  $f_0$  such that  $\eta'(f_0) \in L^\infty(\tau)$  it holds  $D_{\bar{\eta}}(f, f_0) \gtrsim \epsilon_f^d \eta(\epsilon_f^{-d})$ .

(I). Consider  $\bar{G}(\mu) = \int_{\Theta} \text{dist}(\theta_0, \theta) d\mu(\theta)$ . This satisfies Assumption (A1)(I) with  $R$  the identity on  $\mathbb{R}$  and  $\Phi(\theta) = \text{dist}(\theta_0, \theta)$  which is clearly Lipschitz. Since  $\bar{H} = \nu_{\mathcal{P}(\Theta)}$ ,  $\bar{F}$  admits a unique minimizer  $\mu^* = \delta_{\theta_0}$  and it holds  $\bar{F}(\mu^*) = \inf \bar{F} = 0$ . For any  $f \in \text{dom } H$ , we have the following lower bound where  $\epsilon_f$  is defined above:

$$F(f) - \inf F = \int_{\Theta} \text{dist}(\theta, \theta_0) f(\theta) d\tau(\theta) \geq \int_{\Theta \setminus B_{\epsilon_f}(\theta_0)} \text{dist}(\theta, \theta_0) f(\theta) d\tau(\theta) \geq \frac{\epsilon_f}{2}.$$

This proves the lower bound of Eq. (9) with  $q = 1$ .

(II). Consider  $\bar{G}(\mu) = \int_{\Theta} \tilde{\Phi}(\theta) d\mu(\theta)$ , where  $\tilde{\Phi}$  is any function which is continuously twice differentiable, coincides with  $\text{dist}(\theta_0, \cdot)^2$  on  $B_{1/2}(\theta_0)$  and is larger than  $1/4$  outside of this ball. We cannot directly take  $\text{dist}(\theta_0, \cdot)^2$  because this function is not smooth everywhere on  $\Theta$  due to the existence of a cut locus, but it is smooth on  $B_{1/2}(\theta_0)$ . Assumptions (A1)(II) are satisfied. Again  $\mu^* = \delta_{\theta_0}$  is the unique minimizer and  $F(\mu^*) = 0$ . Analogous computations show that  $F(f) - \inf F \geq \min\{\frac{1}{4}, \epsilon_f^2\}$ . This proves the lower bound of Eq. (9) with  $q = 2$ .

(I\*). Consider  $\bar{G}(\mu) = \frac{1}{2} (\int \Phi d\mu)^2$  where  $\Phi(\theta) = \text{dist}(\theta, \theta_0)$ . Clearly,  $\mu^* = \delta_{\theta_0}$  is the unique minimizer and  $\bar{G}'[\mu^*] = 0$  so Assumption (A1)(I\*) is satisfied. By direct computations, it holds  $F(f) - \inf F \gtrsim \epsilon_f^2$ . This proves the lower bound of Eq. (9) with  $q = 2$ .

(II\*). Consider  $\bar{G}(\mu) = \frac{1}{2} (\int \tilde{\Phi} d\mu)^2$  where  $\tilde{\Phi}$  is defined as in the analysis of (II). Again,  $\mu^*$  is the unique minimizer and  $\bar{G}'[\mu^*] = 0$  so Assumption (A1)(II\*) is satisfied. By direct computations, it holds  $F(f) - \inf F \gtrsim \epsilon_f^4$ , which proves the lower bound with  $q = 4$ . ◀

► **Remark 11** (Exact decay of  $\psi$  for a natural class of problems.). There exists in fact a broad class of problems satisfying Assumption (A1)(II) for which the bound on  $\psi$  with  $q = 2$  is exact. These are problems with a sparse solution  $\mu^*$  that satisfy an additional non-degeneracy condition at optimality, that appear naturally in certain contexts [38]. For these problems, it is shown in [16, Prop. 3.2] that  $\bar{F}(\mu) - \inf \bar{F} \gtrsim (\sup_g \int g d[\mu - \mu^*])^2$  where the supremum is over 1-Lipschitz functions  $g : \Theta \rightarrow \mathbb{R}$  uniformly bounded by 1. Reasoning as in the proof of Proposition 10, this implies  $F(f) - \inf F \gtrsim \epsilon_f^2$  which leads to

$$\psi(\alpha) \asymp \inf_{\epsilon > 0} \epsilon^2 + \alpha \epsilon^d \eta(\epsilon^{-d}).$$

## 5.2 Direct lower bounds on the convergence rates

In this section, we directly lower-bound the convergence rates of Algorithm 1 and Algorithm 2. We focus on the  $L^2$  geometry ( $\eta_2$ ) for which we prove all the lower bounds (there are 8 cases to consider) and on the relative entropy geometry ( $\eta_{\text{ent}}$ ) for which we omit certain settings for the sake of conciseness. In all the cases considered, the lower bounds match the upper bounds (up to logarithmic terms for  $\eta_{\text{ent}}$ ). Let us start with PGM and  $\eta_2$ .

► **Proposition 12.** *For each of the settings (I), (I\*), (II) and (II\*) under Assumption (A1), there exists a function  $F$  such that the iterates of Algorithm 1 (PGM)  $(f_k)_{k \geq 0}$  with the distance-generating function  $\eta_2$ , and initialized with  $f_0 = 1$ , with any step-size  $s > 0$ , satisfy*

$$F(f_k) - \inf F \gtrsim k^{-\frac{q}{d+q}}$$

where  $q$  is the constant associated to the setting via Table 2(b).

**Proof.**

(I). As in the proof of Proposition 10, we consider  $\Theta = \mathbb{T}^d$ ,  $\theta_0 \in \mathbb{T}$ ,  $\bar{H} = \nu_{\mathcal{P}(\Theta)}$  and  $\bar{G}(\mu) = \int \Phi d\mu$  with  $\Phi(\theta) = \text{dist}(\theta, \theta_0)$ . We set  $s = 1$  as the step-size plays no role in what follows. In this case, the update equation of Algorithm 1 writes  $f_{k+1} = (f_k - \Phi - \kappa_k)_+$  where  $\kappa_k \in \mathbb{R}$  is such that  $f_{k+1} \in \text{dom } H$ . Thanks to the symmetries of the problem, a direct recursion shows that it holds

$$f_k = (m(k) - k\Phi)_+$$

for some  $m(k) \in \mathbb{R}_+$ . Let  $r(k) := m(k)/k$  which is such that  $B_{r(k)}(\theta_0)$  is the support of  $f_k$ . We have

$$1 = \int f_k d\tau \asymp \int_0^{r(k)} u^{d-1} (m(k) - ku) du \asymp \frac{m(k)}{d} r(k)^d - \frac{1}{d+1} k \cdot r(k)^{d+1} \asymp m(k)^{d+1} k^{-d}$$

thus  $m(k) \asymp k^{d/(d+1)}$  and  $r(k) \asymp k^{-1/(d+1)}$ . We can compute the objective

$$F(f_k) \asymp \int_0^{r(k)} u^{d-1} u (m(k) - ku) du \asymp \frac{1}{d+1} m(k) r(k)^{d+1} - k \frac{1}{d+2} r(k)^{d+2} \asymp k^{-1/(d+1)}.$$

Which proves the case (I) (here  $q = 1$ ) since  $\inf F = 0$ .

(II). Consider a smooth function  $\tilde{\Phi}$  which equals  $\text{dist}(\theta, \theta_0)^2$  on the ball  $B_{1/2}(\theta_0)$  as in the proof of Proposition 10. Again the iterates have the form  $f_k = (m(k) - k\tilde{\Phi})_+$  and now let  $r(k)^2 = m(k)/k$ . For  $k$  large enough, so that  $\tilde{\Phi} = \text{dist}(\cdot, \theta_0)^2$  over  $B_{r(k)}(\theta_0)$ , it holds

$$1 = \int f_k d\tau \asymp \int_0^{r(k)} u^{d-1} (m(k) - ku^2) du \asymp \frac{m(k)}{d} r(k)^d - \frac{k}{d+2} \cdot r(k)^{d+2} \asymp m(k)^{(d+2)/2} k^{-d/2}$$

thus  $m(k) \asymp k^{d/(d+2)}$  and  $r(k) \asymp k^{-1/(d+2)}$ . It also holds

$$F(f_k) \asymp \int_0^{r(k)} u^{d-1} u^2 (m(k) - ku^2) du \asymp \frac{m(k)}{d+2} r(k)^{d+2} - \frac{k}{d+4} r(k)^{d+4} \asymp k^{-2/(d+2)}$$

Which proves the case (II) (here  $q = 2$ ).

(I\*). We consider the function  $G(\mu) = \frac{1}{2} (\int \Phi d\mu)^2$ . Now the reasoning is slightly more subtle because of the non-linearity. It holds  $f_k = (m(k) - s(k)\Phi)_+$  where  $s(k+1) = s(k) + \int \Phi f_k d\tau$ . Again,  $m(k)$  is such that the iterate is feasible. Our upper bounds imply that  $\int \Phi f_k d\tau \lesssim k^{-1/(d+2)}$ , thus  $s(k+1) - s(k) \lesssim k^{-1/(d+2)}$  and it follows  $s(k) \lesssim k^{\frac{d+1}{d+2}}$  and  $F(f_k) \gtrsim (s(k)^{-1/(d+1)})^2 \asymp k^{-2/(d+2)}$  as desired.

(II\*). For the case (II\*) we combine the ideas from (II) and (I\*), let us just emphasize on the differences. We take the function  $G(\mu) = \frac{1}{2} (\int \tilde{\Phi})^2$ . For  $k$  large enough, it holds  $f_k = (m(k) - s(k)\tilde{\Phi})_+$  and, thanks to our upper-bound,  $\int \tilde{\Phi} f_k d\tau \lesssim k^{-2/(d+4)}$ . The recursion becomes  $s(k+1) - s(k) \lesssim k^{-2/(d+4)}$  and it follows  $s(k) \lesssim k^{(d+2)/(d+4)}$  and  $F(f_k) \gtrsim (s(k)^{-2/(d+2)})^2 \asymp k^{-4/(d+4)}$  as desired. ◀

Let us now prove similar lower bounds for APGM, again for the specific choice of distance-generating function  $\eta_2$ .

► **Proposition 13.** *For each of the settings (I), (II), (I\*) and (II\*) under Assumption (A1), there exists a function  $F$  such that the iterates of Algorithm 2 (APGM)  $(f_k)_{k \geq 0}$  with the distance-generating function  $\eta_2$ , and initialized with  $f_0 = 1$ , with any step-size  $s > 0$ , satisfy*

$$F(f_k) - \inf F \gtrsim k^{-\frac{2q}{d+q}}$$

where  $q$  is the constant associated to the setting via Table 2(b).

**Proof.** For (I), we consider the same set up as in the proof of Proposition 12(I) (also fixing  $s = 1$  for conciseness). Since  $G$  is linear, the update of  $h_k$  reads  $h_{k+1} = (m(k) - s(k)\Phi)_+$  with  $s_k = \sum_{i=0}^k \frac{1}{\gamma_i}$  where  $(\gamma_k)_k$  is defined in Algorithm 2. Since  $\gamma_k \gtrsim 1/k$ , we have  $s(k) \asymp k^2$  which implies that  $F(h_k) \gtrsim k^{-2/(d+1)}$ . Finally, it is clear that  $f_k$  is in the convex hull of  $\{h_0, \dots, h_k\}$  and since  $G$  is linear, it holds  $F(f_k) \geq \min_{i=1, \dots, k} F(h_k) \gtrsim k^{-2/(d+1)}$ . The proof for the case (II) follows exactly the same scheme but with the function  $\tilde{\Phi}$  considered in the proof of Proposition 12 and we omit the details.

In the case (I\*), we have  $(m(k) - s(k)\Phi)_+$  with  $s(k+1) - s(k) = \frac{1}{\gamma_k} \int \Phi g_k d\tau$ . Thanks to our upper-bounds, we have  $\int \Phi f_k d\tau \lesssim k^{-2/(d+2)}$ , and since  $\gamma_k \rightarrow 0$  and  $\|h_k\|_{L^1} = 1$ , it follows  $\int \Phi g_k d\tau \lesssim k^{-2/(d+2)}$ . Thus  $s(k+1) - s(k) \lesssim k^{d/(d+2)}$  which leads to  $s(k) \lesssim k^{2(d+1)/(d+2)}$ . It follows  $F(h_k) \gtrsim s(k)^{-2/(d+1)} \gtrsim k^{-4/(d+2)}$ . Since  $h_k$  is optimal for  $F$  over the convex hull of  $\{h_0, \dots, h_k\}$ , which contains  $f_k$  (it has the most mass in small balls around  $\theta_0$ ) we obtain the same lower bound on  $F(f_k)$ . The proof for the case (II\*) follows the same scheme but with the function  $\tilde{\Phi}$  and we omit the details. ◀

It is also instructive to look at lower bounds with the entropy  $\eta_{\text{ent}}$ . We observe that here there exist cases where the guarantee given by Proposition 5 is off by a  $\log(k)$  factor (because this factor is present in the lower bound of Proposition 10).

► **Proposition 14.** *For the settings (I) and (II) under Assumption (A1), there exists a function  $F$  such that the iterates of Algorithm 1 (PGM)  $(f_k)_{k \geq 0}$  with the distance-generating function  $\eta_{\text{ent}}$ , and initialized with  $f_0 = 1$ , with any step-size  $s > 0$ , satisfy*

$$F(f_k) - \inf F \asymp k^{-1}.$$

**Proof.** We consider the same setting as in the proof of Proposition 12(I) (and  $s = 1$  for simplicity). In this case, the update reads  $f_{k+1} \propto \exp(f_k - \Phi)$  so by an immediate recursion  $f_k \propto \exp(-k\Phi)$ . This is essentially a (multi-dimensional) Laplace distribution and when  $k$  is large, up to exponentially small terms in  $k$ , we can compute the integrals over  $\mathbb{R}^d$  instead of  $\mathbb{T}^d$ . For the normalizing factor, we have

$$\int_{\mathbb{R}^d} \exp(-k\|x\|_2^2) dx \asymp \int_0^\infty u^{d-1} \exp(-ku) du = k^{2-d} \int_0^\infty s^{d-1} e^{-s} ds = k^{2-d} \Gamma(d)$$

where  $\Gamma$  is the Gamma function. For the (unnormalized) value of  $F(f_k)$ , we have

$$\int_{\mathbb{R}^d} \|x\|_2 \exp(-k\|x\|_2^2) dx \asymp \int_0^\infty u^d \exp(-ku) du = k^{1-d} \int_0^\infty s^d e^{-s} ds = k^{1-d} \Gamma(d+1).$$

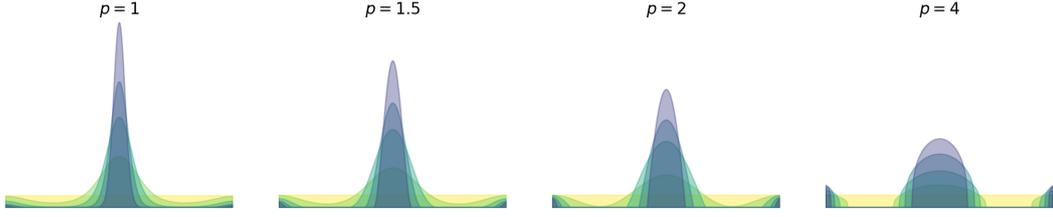
By computing the ratio, it follows that  $F(f_k) - \inf F \asymp k^{-1}$ . In Setting (II), we take the function  $\tilde{\Phi}$  as before which is equal to  $\text{dist}(\cdot, \theta_0)^2$  near  $\theta_0$ . Now  $f_k \propto \exp(-k\tilde{\Phi})$  which is essentially, when  $k$  is large, a Gaussian distribution of variance  $1/k \asymp F(f_k) - \inf F$ . ◀

► **Remark 15.** Although the convergence rates obtained with  $\eta_{\text{ent}}$  and  $\eta_{\text{hyp}}$  are independent of the dimension  $d$  (see Table 2), this favorable behavior crucially relies on the assumption that  $\bar{F}$  admits a minimizer  $\mu^* \in \mathcal{M}(\Theta)$ . When this is not the case, [42] show that there is an example where the continuous time dynamics induced by  $\eta_{\text{ent}}$  also suffer from the curse of dimensionality (our setting is slightly different but their argument would apply here). In addition, the discrete time dynamics are not stable in this case because the norm of the iterates grows unbounded, see Remark 6.

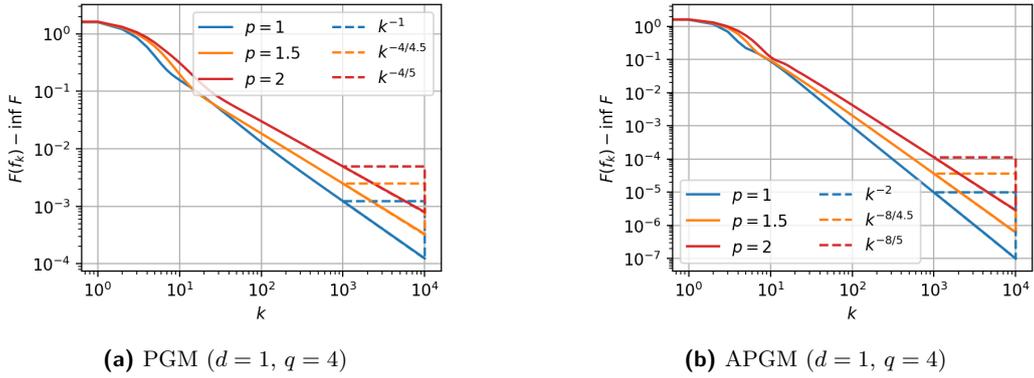
## 6 Numerical experiments

In this section we compare our theoretical rates with the practical behavior of PGM (Algorithms 1) and APGM (Algorithm 2) on simple toy problems. The purpose is to show that, although our analysis is asymptotic (in  $k$  and in the spatial discretization), it describes well the convergence of those algorithms in certain practical scenarios. The code to reproduce these experiments can be found online<sup>3</sup>.

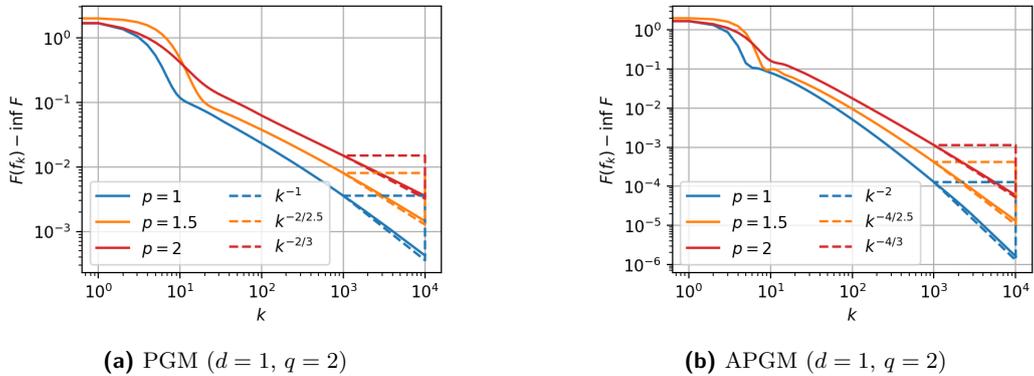
<sup>3</sup> <https://github.com/lchizat/2021-measures-PGM>



**Figure 3** Behavior of Algorithm 1 (PGM) for a nonnegative sparse deconvolution problem with solution  $\mu^* = \delta_0$ , with  $d = 1$  and for various Bregman divergences  $\eta_p$  ( $p = 1$  stands for  $\eta_{\text{ent}}$ ). We plot the function  $f_k$  for  $k \in \{0, 6, 6^2, 6^3, 6^4\}$ , we use the same step-size and the same axes in all cases. The associated convergence plots are in Figure 4.



**Figure 4** Convergence of PGM and APGM vs. theoretical rates (up to log factors) in a sparse deconvolution problem with  $\bar{H} = \nu_{\mathcal{M}_+(\Theta)}$  and  $d = 1$ . Here  $p$  refers to the parameter of  $\eta_p$  and  $p = 1$  refers to  $\eta_{\text{ent}}$ . The objective has structure (II\*) so  $q = 4$  in the rates of Table 2.

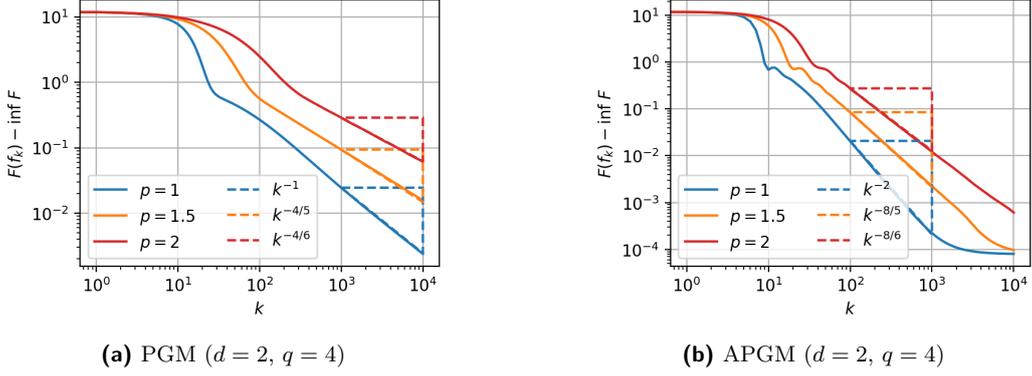


**Figure 5** Convergence of PGM and APGM vs. theoretical rates (up to log factors) in a sparse deconvolution problem with  $\bar{H} = \lambda \|\mu\|$  and  $d = 1$ . Here  $p$  refers to the parameter of  $\eta_p$  and  $p = 1$  refers to  $\eta_{\text{hyp}}$ . The objective has structure (II) so  $q = 2$  in the rates of Table 2.

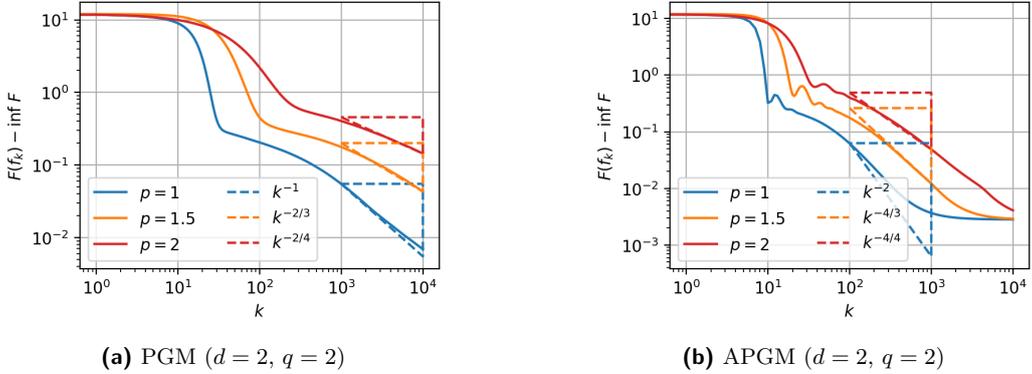
## 6.1 Sparse deconvolution

We consider the sparse deconvolution problem introduced in Section 3.1 where  $\phi$  is a Dirichlet kernel  $\phi(\theta) = \sum_{k \in \{-2, 1, 0, 1, 2\}^d} \exp(2\sqrt{-1}\pi k^\top \theta)$  and  $y^*(\theta) = \phi(\theta) = \int \phi(\theta - \theta') d\mu^*(\theta')$  with  $\mu^* = \delta_0$ . The domain  $\mathbb{T}^1$  is discretized into a regular grid of  $m = 300$  points and  $\mathbb{T}^2$  into a regular grid of  $m = 60 \times 60$  points. Figure 3 illustrates the behavior of the various Bregman divergences for this problem, where it is seen that the iterates  $f_k$  (weakly) converge faster to the Dirac solution as  $p$  is smaller (in the following discussion, we use  $p = 1$  to refer to the entropy or hypentropy distance-generating function).

Figures 4, 5, 6 and 7 report the convergence rates in a variety of settings, which we compare to our theoretical predictions (without the logarithmic factors, since they do not change the asymptotic slope on a log-log plot). In both cases,  $\inf F$  admits a closed form so we can exactly plot  $\bar{F}(f_k \tau_m) - \min \bar{F}$  and additionally observe the



**Figure 6** Convergence of PGM and APGM vs. theoretical rates (up to log factors) in a sparse deconvolution problem with  $\bar{H} = \nu_{\mathcal{M}_+(\Theta)}$  and  $d = 2$ . Here  $p$  refers to the parameter of  $\eta_p$  and  $p = 1$  refers to  $\eta_{\text{ent}}$ . The objective has structure (II\*) so  $q = 4$  in the rates of Table 2.



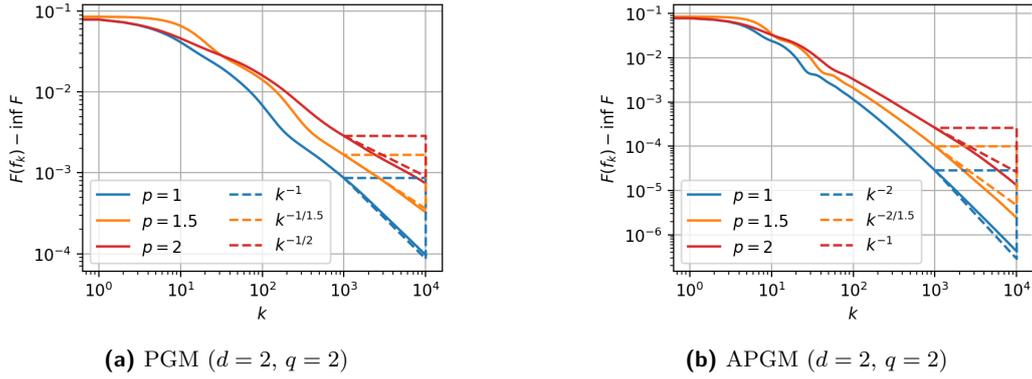
**Figure 7** Convergence of PGM and APGM vs. theoretical rates (up to log factors) in a sparse deconvolution problem with  $\bar{H} = \lambda\|\mu\|$  and  $d = 2$ . Here  $p$  refers to the parameter of  $\eta_p$  and  $p = 1$  refers to  $\eta_{\text{hyp}}$ . The objective has structure (II) so  $q = 2$  in the rates of Table 2.

effect of the discretization (here  $\tau_m$  is the discretized reference measure). Observe that in the 2D experiments, APGM with  $p = 1$  quickly reaches the discretization error, and on Figure 7 (b), it does not have enough “time” to attain the theoretical asymptotic rate before the effect of the discretization comes in. While our analysis is asymptotic, it thus corresponds in practice to a non-asymptotic and transient behavior.

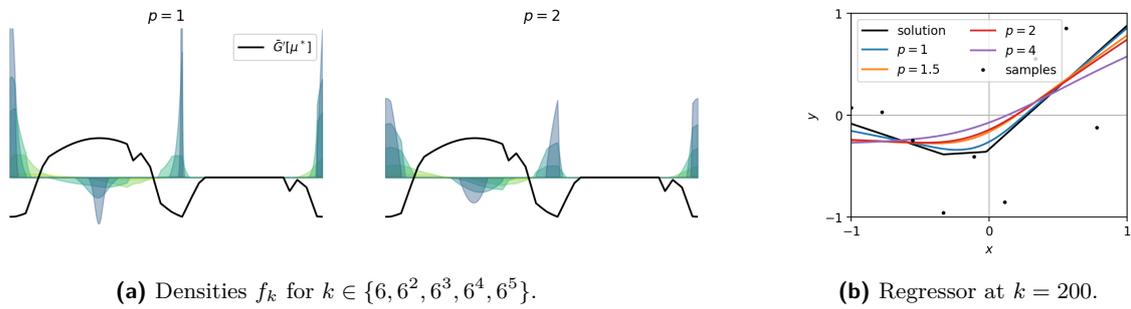
## 6.2 Two-layer neural networks

We consider a two-layer ReLU neural network with the objective function introduced in Section 3.1 where we consider  $n = 10$  input samples  $x_i$  on a regular grid on  $[-1, 1]$  and observed variables  $y_i = |x_i| - \frac{1}{2} + Z_i$  where  $Z_i$  are independent and uniform on  $[-1, 1]$  (see the samples on Figure 9 (b)). The domain is  $\mathbb{S}^1$  discretized into a regular grid of  $m = 2000$  points. This setting gives an example where  $\Phi$  does not have a Lipschitz gradient and is only Lipschitz (observe the irregularity of  $\bar{G}'[\mu^*]$  on Figure 9 (a)). Since we use the regularization  $\bar{H} = \lambda\|\mu\|$ , we are in the setting (I) from Table 2 (b), and the parameter for the rate is  $q = 1$ .

Figure 8 shows the rates of convergence for PGM and APGM. Although the general picture is consistent with the theory, we observe that our guarantees are a bit over-conservative. For PGM, we roughly measure (between iteration  $k = 10^3$  and  $k = 10^5$ ) the rates exponents  $(-1.00, -0.72, -0.58)$  for respectively  $p = (1, 1.5, 2)$  which corresponds to a parameter  $q \approx 1.5$  rather than  $q = 1$ . For APGM, we roughly measure the rates exponents  $(-1.97, -1.71, -1.41)$  instead of the predicted  $(-2, -1.33, -1)$ . Figure 9 (a) helps understanding this discrepancy: as can be seen from the proof of Theorem 8 what truly determines the asymptotic rate is how much the objective function increases when  $\mu^*$  is mollified, and we quantified this using the regularity and magnitude of  $\bar{G}'[\mu]$  near  $\mu^*$ . Here it appears that  $\bar{G}'[\mu^*]$  is smooth near 2 out of the 3 points in the support of  $\mu^*$ , while it is non-smooth at the third point (the one in the middle). The fact that we have a mix of both levels of regularity (i.e. smooth



**Figure 8** Convergence of PGM and APGM vs. theoretical rates (up to log factors) for a 2-layer ReLU neural network with  $\bar{H} = \lambda \|\mu\|$  and  $d = 1$ . Here  $p$  refers to the parameter of  $\eta_p$  and  $p = 1$  refers to  $\eta_{\text{hyp}}$ . The objective has structure (I) so  $q = 1$  in the rates of Table 2.



**Figure 9** Dynamics of PGM on a two-layer neural network, for various values of  $p$  ( $p = 1$  corresponds to  $\eta_{\text{hyp}}$ ). Observe in (b) how the dynamics with  $\eta_{\text{hyp}}$  fits the kinks of the optimal regressor much faster than with  $p > 1$ .

vs. merely Lipschitz) may explain why the convergence is a bit faster than with the parameter  $q = 1$ , which corresponds to only taking into account the Lipschitz regularity.

## 7 Conclusion

We have studied the convergence rates of PGM and APGM for convex optimization in the space of measures. Our analysis exhibits the influence of the regularity of the objective function on the convergence rates. It also confirms that the geometry induced by  $\eta_{\text{ent}}$  and  $\eta_{\text{hyp}}$  is better suited than the  $L^2$  geometry to solve such problems. An important question for future research is to better understand the unregularized case, where the phenomenon of algorithmic regularization is at play.

## Acknowledgments

I am thankful to Adrien Taylor for fruitful discussions during the preparation of this paper. In particular, I learnt about Algorithm 2 (APGM) from him.

## References

- 1 Ehsan Amid and Manfred K. Warmuth. Winnowing with gradient descent. In *Conference on Learning Theory*, pages 163–182. PMLR, 2020.
- 2 Alfred Auslender and Marc Teboulle. Interior gradient and proximal methods for convex and conic optimization. *SIAM J. Optim.*, 16(3):697–725, 2006.
- 3 Shahar Azulay, Edward Moroshko, Mor Shpigel Nacson, Blake Woodworth, Nathan Srebro, Amir Globerson, and Daniel Soudry. On the Implicit Bias of Initialization Shape: Beyond Infinitesimal Mirror Descent. <https://arxiv.org/abs/2102.09769>, 2021.
- 4 Francis Bach. Breaking the curse of dimensionality with convex neural networks. *J. Mach. Learn. Theory*, 18(1):629–681, 2017.

- 5 Francis Bach. Learning Theory from First Principles Draft, 2021.
- 6 Heinz H. Bauschke, Jonathan M. Borwein, and Patrick L. Combettes. Essential smoothness, essential strict convexity, and Legendre functions in Banach spaces. *Commun. Contemp. Math.*, 3(04):615–647, 2001.
- 7 Heinz H. Bauschke, Jonathan M. Borwein, and Patrick L. Combettes. Bregman monotone optimization algorithms. *SIAM J. Control Optim.*, 42(2):596–636, 2003.
- 8 Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imaging Sci.*, 2(1):183–202, 2009.
- 9 Yoshua Bengio, Nicolas Roux, Pascal Vincent, Olivier Delalleau, and Patrice Marcotte. Convex Neural Networks. In Y. Weiss, B. Schölkopf, and J. Platt, editors, *Advances in Neural Information Processing Systems*, volume 18. MIT Press, 2006.
- 10 Nicholas Boyd, Geoffrey Schiebinger, and Benjamin Recht. The alternating descent conditional gradient method for sparse inverse problems. *SIAM J. Optim.*, 27(2):616–639, 2017.
- 11 Kristian Bredies and Dirk A. Lorenz. Linear convergence of iterative soft-thresholding. *J. Fourier Anal. Appl.*, 14(5-6):813–837, 2008.
- 12 Kristian Bredies and Hanna Katriina Pikkarainen. Inverse problems in spaces of measures. *ESAIM, Control Optim. Calc. Var.*, 19(1):190–218, 2013.
- 13 Sébastien Bubeck. Convex Optimization: Algorithms and Complexity. *Found. Trends Mach. Learn.*, 8(3-4):231–357, 2015.
- 14 Emmanuel J. Candès and Carlos Fernandez-Granda. Towards a mathematical theory of super-resolution. *Commun. Pure Appl. Math.*, 67(6):906–956, 2014.
- 15 Antonin Chambolle and Robert Tovey. "FISTA" in Banach spaces with adaptive discretisations. <https://arxiv.org/abs/2101.09175>, 2021.
- 16 Lénaïc Chizat. Sparse optimization on measures with over-parameterized gradient descent. *Math. Program.*, pages 1–46, 2021.
- 17 Lénaïc Chizat and Francis Bach. Implicit bias of gradient descent for wide two-layer neural networks trained with the logistic loss. In *Conference on Learning Theory*, pages 1305–1338. PMLR, 2020.
- 18 Laurent Condat. Fast projection onto the simplex and the  $\ell_1$  ball. *Math. Program.*, 158(1):575–585, 2016.
- 19 Alexandre d’Aspremont, Damien Scieur, and Adrien Taylor. Acceleration methods. <https://arxiv.org/abs/2101.09545>, 2021.
- 20 Yohann De Castro and Fabrice Gamboa. Exact reconstruction using Beurling minimal extrapolation. *J. Math. Anal. Appl.*, 395(1):336–354, 2012.
- 21 Quentin Denoyelle, Vincent Duval, Gabriel Peyré, and Emmanuel Soubies. The sliding Frank–Wolfe algorithm and its application to super-resolution microscopy. *Inverse Probl.*, 36(1), 2019.
- 22 Aymeric Dieuleveut. *Stochastic approximation in Hilbert spaces*. PhD thesis, PSL Research University, 2017.
- 23 Carles Domingo-Enrich, Samy Jelassi, Arthur Mensch, Grant Rotskoff, and Joan Bruna. A mean-field analysis of two-player zero-sum games. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 20215–20226. Curran Associates, Inc., 2020.
- 24 Guillaume Garrigos, Lorenzo Rosasco, and Silvia Villa. Thresholding gradient methods in Hilbert spaces: support identification and linear convergence. *ESAIM, Control Optim. Calc. Var.*, 26, 2020.
- 25 Udaya Ghai, Elad Hazan, and Yoram Singer. Exponentiated gradient meets gradient descent. In *Algorithmic Learning Theory*, pages 386–407. PMLR, 2020.
- 26 Alfred Gray and Lieven Vanhecke. Riemannian geometry as determined by the volumes of small geodesic balls. *Acta Math.*, 142(1):157–198, 1979.
- 27 Alfred Gray and Tom J. Willmore. Mean-value theorems for Riemannian manifolds. *Proc. R. Soc. Edinb., Sect. A, Math.*, 92(3-4):343–364, 1982.
- 28 Matt Jacobs, Flavien Léger, Wuchen Li, and Stanley Osher. Solving large-scale optimization problems with a convergence rate independent of grid size. *SIAM J. Numer. Anal.*, 57(3):1100–1123, 2019.
- 29 Chao Kan and Wen Song. The Moreau envelope function and proximal mapping in the sense of the Bregman distance. *Nonlinear Anal., Theory Methods Appl.*, 75(3):1385–1399, 2012.
- 30 Jyrki Kivinen and Manfred K. Warmuth. Exponentiated gradient versus gradient descent for linear predictors. *Inf. Comput.*, 132(1):1–63, 1997.
- 31 Jean B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM J. Optim.*, 11(3):796–817, 2001.
- 32 Jingwei Liang, Jalal M. Fadili, and Gabriel Peyré. Local linear convergence of forward–backward under partial smoothness. In *Proceedings of the 27th International Conference on Neural Information Processing Systems-Volume 2*, pages 1970–1978, 2014.
- 33 Song Mei, Andrea Montanari, and Phan-Minh Nguyen. A mean field view of the landscape of two-layer neural networks. *Proc. Natl. Acad. Sci. USA*, 115(33):E7665–E7671, 2018.

- 34 Arkadij Semenovič Nemirovsky and David Borisovich Yudin. *Problem complexity and method efficiency in optimization*. Wiley-Interscience series in discrete mathematics. John Wiley & Sons, 1983.
- 35 Yurii Nesterov. On an approach to the construction of optimal methods of minimization of smooth convex functions. *Ekonom. i. Mat. Metody*, 24(3):509–517, 1988.
- 36 Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer, 2003.
- 37 Atsushi Nitanda, Denny Wu, and Taiji Suzuki. Particle Dual Averaging: Optimization of Mean Field Neural Networks with Global Convergence Rate Analysis. <https://arxiv.org/abs/2012.15477>, 2020.
- 38 Clarice Poon, Nicolas Keriven, and Gabriel Peyré. The geometry of off-the-grid compressed sensing. <https://arxiv.org/abs/1802.08464>, 2018.
- 39 Ralph Rockafellar. Integrals which are convex functionals. II. *Pac. J. Math.*, 39(2):439–469, 1971.
- 40 Paul Tseng. Approximation accuracy, gradient methods, and error bound for structured convex optimization. *Math. Program.*, 125(2):263–295, 2010.
- 41 Tomas Vaskevicius, Varun Kanade, and Patrick Rebeschini. Implicit Regularization for Optimal Sparse Recovery. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- 42 Stephan Wojtowytsch and Weinan E. Can shallow neural networks beat the curse of dimensionality? A mean field training perspective. *IEEE Trans. Artif. Intell.*, 1(2):121–129, 2020.
- 43 Yuan Yao, Lorenzo Rosasco, and Andrea Caponnetto. On early stopping in gradient descent learning. *Computing*, 26(2):289–315, 2007.
- 44 Yao-Liang Yu. The strong convexity of Von Neumann’s entropy, 2013. Unpublished note.