

ON REPRESENTER THEOREMS AND CONVEX REGULARIZATION*

CLAIRE BOYER[†], ANTONIN CHAMBOLLE[‡], YOHANN DE CASTRO[§],
VINCENT DUVAL[¶], FRÉDÉRIC DE GOURNAY^{||}, AND PIERRE WEISS[#]

Abstract. We establish a general principle which states that regularizing an inverse problem with a convex function yields solutions that are convex combinations of a small number of *atoms*. These atoms are identified with the extreme points and elements of the extreme rays of the regularizer level sets. An extension to a broader class of quasi-convex regularizers is also discussed. As a side result, we characterize the minimizers of the total gradient variation, which was previously an unresolved problem.

Key words. inverse problems, convex regularization, representer theorem, vector space, total variation

AMS subject classifications. 52A05, 49N45, 46E27

DOI. 10.1137/18M1200750

1. Introduction. Let E denote a real vector space. Let $\Phi : E \rightarrow \mathbb{R}^m$ be a linear mapping called a *sensing operator* and $u \in E$ denote a signal. The main results in this paper describe the structural properties of certain solutions of the following problem:

$$(1.1) \quad \inf_{u \in E} f(\Phi u) + R(u),$$

where $R : E \rightarrow \mathbb{R} \cup \{+\infty\}$ is a convex function called a *regularizer* and f is an arbitrary convex or nonconvex function called the *data fitting term*. In many applications, one looks for “sparse solutions” that are linear sums of a few atoms. This article investigates the theoretical legitimacy of this usage.

Representer theorems and Tikhonov regularization. The name “representer theorem” comes from the field of machine learning [43]. To provide a first concrete example,¹ assume that $\Phi \in \mathbb{R}^{m \times n}$ is a finite-dimensional measurement operator and $L \in \mathbb{R}^{p \times n}$ is a linear transform. Solving an inverse problem using Tikhonov regularization amounts to finding the minimizers of

$$(1.2) \quad \min_{u \in \mathbb{R}^m} \frac{1}{2} \|\Phi u - y\|_2^2 + \frac{1}{2} \|Lu\|_2^2.$$

*Received by the editors July 16, 2018; accepted for publication (in revised form) January 15, 2019; published electronically May 9, 2019. A two-page paper containing preliminary results about this work (without proofs) was submitted to the iTWIST workshop.

<http://www.siam.org/journals/siopt/29-2/M120075.html>

Funding: Supported by EPSRC grant N. EP/K032208/1. The work of the second author was partially supported by a grant from the Simons Foundation.

[†]LPSM, Sorbonne Université, and DMA, ENS Paris, France (claire.boyer@upmc.fr).

[‡]CMAF, CNRS, École Polytechnique, Palaiseau, France (antonin.chambolle@cmap.polytechnique.fr).

[§]CERMICS, École des Ponts ParisTech, Marne la Vallée, and INRIA, France (yohann.decastro@enpc.fr).

[¶]INRIA, and Université Paris-Dauphine, Université PSL, CNRS, CEREMADE, France (vincent.duval03@m4x.org).

^{||}IMT & ITAV, INSA, Université de Toulouse, France (frederic.de-gournay@insa-toulouse.fr).

[#]ITAV & IMT, CNRS, Université de Toulouse, France (pierre.armand.weiss@gmail.com).

¹Here, we follow the presentation of [27].

Provided that $\ker \Phi \cap \ker L = \{0\}$, it is possible to show that, whatever the data y , solutions are always of the form

$$(1.3) \quad u^* = \sum_{i=1}^m \alpha_i \psi_i + u_K,$$

where $u_K \in \ker(L)$ and $\psi_i = (\Phi^T \Phi + L^T L)^{-1}(\phi_i)$, where $\phi_i^T \in \mathbb{R}^n$ is the i th row of Φ . This result characterizes structural properties of the minimizers without actually needing to solve the problem. In addition, when E is an infinite-dimensional Hilbert space, (1.3) sometimes allows us to compute exact solutions by simply solving a finite-dimensional linear system. This is a critical observation that explains the practical success of kernel methods and radial basis functions [49].

Representer theorems and convex regularization. The Tikhonov regularization (1.2) is a powerful tool when the number m of observations is large and the operator Φ is not too ill-conditioned. However, recent results in the fields of compressed sensing [17], matrix completion [11], and super-resolution [46, 10]—to name a few—suggest that much better results may be obtained in general by using convex regularizers, with level sets containing singularities. Popular examples of regularizers in the finite-dimensional setting include the indicator of the nonnegative orthant [18], the ℓ^1 -norm [17] or its composition with a linear operator [41], and the nuclear norm [11]. Those results were nicely unified in [13] and one of the critical arguments behind all these techniques is a representer theorem of type (1.3). In most situations, however, this argument is only implicit. The main objective of this paper is to generalize (1.3) to arbitrary convex functions R . This paper covers all the aforementioned problems, as well as new ones for problems formulated over the space of measures.

To the best of our knowledge, the name “representer theorem” is new to the field of convex regularization, and its first mention is due to Unser, Fageot, and Ward [48]. Describing the solutions of (1.1) is, however, an old problem which has been studied since at least the 1940s in the case of Radon measure recovery.

Total variation regularization of Radon measures. A typical example of the inverse problem in the space of measures is

$$(1.4) \quad \min_{\mu \in \mathcal{M}(\Omega)} |\mu|(\Omega) \quad \text{s.t.} \quad \Phi \mu = y,$$

where $\Omega \subseteq \mathbb{R}^N$, $\mathcal{M}(\Omega)$ denotes the space of Radon measures, $|\mu|(\Omega)$ is the total variation of the measure μ (see section 4) and $\Phi \mu$ is a vector of *generalized moments*, i.e., $\Phi \mu = (\int_{\Omega} \varphi_i(x) d\mu(x))_{1 \leq i \leq m}$, where $\{\varphi_i\}_{1 \leq i \leq m}$ is a family of continuous functions (which “vanish at infinity” if Ω is not compact).

Problems of the form (1.4) have received considerable attention since the pioneering works of Beurling [6] and Krein [31], sometimes under the name *L-moment problem* (see the monograph [32]). To the best of our knowledge, the first “representer theorem” for problems of the form (1.1) is given for (1.4) by Zuhovickii [50] (see [51, Theorem 3] for an English version). It essentially states that

$$(1.5) \quad \text{there exists a solution to (1.4) of the form } \sum_{i=1}^r a_i \delta_{x_i} \text{ with } r \leq m.$$

A more precise result was given by Fisher and Jerome in [23]. When considering the problem (1.4), and for a bounded domain Ω , the result reads as follows:

the extreme points of the solution set to (1.4) are of the form

$$(1.6) \quad \sum_{i=1}^r a_i \delta_{x_i} \text{ with } r \leq m.$$

Incidentally, the Fisher–Jerome theorem considers more general problems of the form

$$(1.7) \quad \min_{u \in E} |Lu|(\Omega) \quad \text{s.t.} \quad Lu \in \mathcal{M}(\Omega) \quad \text{and} \quad \Phi u = y,$$

where $E \subseteq \mathcal{D}'(\Omega)$ is a suitably defined Banach space of distributions, $L : \mathcal{D}'(\Omega) \rightarrow \mathcal{D}'(\Omega)$ maps E onto $\mathcal{M}(\Omega)$, and $\Phi : E \rightarrow \mathbb{R}^m$ is a continuous linear operator. We refer the reader to section 4 for precise assumptions. Let us mention that the initial results by Fisher and Jerome were extended to a significantly more general setting in [48].

It is important to note that the Fisher–Jerome theorem [23] provides a much finer description of the solution set than Zuhovickii’s result [50]. Indeed, the well-known Krein–Milman theorem states that if E is endowed with the topology of a locally convex Hausdorff vector space and $C \subset E$ is compact convex, then C is the closed convex hull of its extreme points,

$$(1.8) \quad \text{cl conv}(\text{ext}(C)) = C.$$

In other words, the solutions described by the Fisher–Jerome theorem are sufficient to recover *the whole set of solutions*. Let us mention that the Krein–Milman theorem was extended by Klee [30] to unbounded sets: if C is locally compact, closed, convex, and contains no line, then

$$(1.9) \quad \text{cl conv}(\text{ext}(C) \cup \text{rext}(C)) = C,$$

where $\text{rext}(C)$ denotes the union of the extreme rays of C (see section 2).

“Representer theorems” for convex sets. As the Dirac masses are the extreme points of the total variation unit ball, each of the abovementioned “representer theorems” for inverse problems actually reflects some phenomenon in the geometry of convex sets. In that regard, the celebrated Minkowski–Carathéodory theorem [28, Theorem III.2.3.4] is fundamental: any point of a compact convex set in an m -dimensional space is a convex combination of (at most) $m + 1$ of its extreme points. In [29, Theorem (3)], Klee removed the boundedness assumption and obtained the following extension: any point of a closed convex set in an m -dimensional space is a convex combination of (at most) $m + 1$ extreme points, or m points, each an extreme point or a point in an extreme ray.

One aim of the present paper is to point out the connection between the Fisher–Jerome theorem and a lesser known theorem by Dubins [20] (see also [8, Exercise II.7.3.f]): *the extreme points of the intersection of C with an affine space of codimension m are a convex combination of (at most) $m + 1$ extreme points of C ,*² provided C is linearly bounded and linearly closed (see section 2). That theorem was extended by Klee [29] to deal with the unbounded case.

²In the rest of the paper, we omit the “at most,” with the convention that some points may be chosen to be identical.

Although the connection with the Fisher–Jerome theorem is striking, Dubins’s theorem actually provides one extreme point too many. In the case of (1.4), it would yield two Dirac masses for one linear measurement. We provide in this paper a refined analysis of the case of variational problems, which ensures at most m extreme points.

Contributions. The main results of this paper yield a description of some solutions to (1.1) of the following form:

$$u^* = \sum_{i=1}^r \alpha_i \psi_i + u_K,$$

where $r \leq m$, the atoms ψ_i are identified with some extreme points (or points in extreme rays) of the regularizer level sets, and u_K is an element of the so-called constancy space of R , i.e., the set of directions along which R is invariant. The results take the form (1.5) when f is an arbitrary function and the form (1.6) when it is convex. We provide tight bounds on the number of atoms r , which depends on the geometry of the level sets and on the link between the constancy space of R and the measurement operator Φ .

Our general theorems then allow us to revisit many results in the literature (linear programming, semidefinite programming, nonnegative constraints, nuclear norm, analysis priors) to give simple and accurate descriptions of the minimizers. Our analysis also allows us to characterize the solutions of a resisting problem: we provide a representation theorem for the minimizers of the total gradient variation [41] as sums of indicators of simple sets. This provides a simple explanation to the staircasing effect when only a few measurements are used.

Let us mention that, shortly after this work was posted on arXiv, similar results appeared, with somewhat different proofs, in a paper by Bredies and Carioni [9].

2. Notation and preliminaries. Throughout the paper, unless otherwise specified, E denotes a finite- or infinite-dimensional real vector space and $C \subseteq E$ is a convex set. Given two distinct points x and y in E , we let

$$]x, y[= \{tx + (1 - t)y : 0 < t < 1\} \quad \text{and} \quad [x, y] = \{tx + (1 - t)y : 0 \leq t \leq 1\}$$

denote the open and closed segments joining x to y . We recall the following definitions, and we refer the reader to [20, 30] for more details.

Lines, rays, and linearly closed sets. A *line* is an affine subspace of E with dimension 1. An open half line, i.e., a set of the form $\rho = \{p + tv : t > 0\}$, where $p, v \in E$, $v \neq 0$, is called a *ray* (through p). We say that the set C is *linearly closed* (respectively, linearly bounded) if the intersection of C and a line of E is closed (respectively, bounded) for the natural topology of the line. If E is a topological vector space and C is closed for the corresponding topology, then C is linearly closed.

If C is linearly closed and contains some ray $\rho = p + \mathbb{R}_+^* v$, it also contains the endpoint p as well as the rays $q + \mathbb{R}_+ v$ for all $q \in C$. Therefore, if C contains a ray (respectively, line), it recesses in the corresponding direction.

Recession cone and lineality space. The set of all $v \in E$ such that $C + \mathbb{R}_+^* v \subseteq C$ is a convex cone called the *recession cone of C* , which we denote by $\text{rec}(C)$. If C is linearly closed, then so is $\text{rec}(C)$, and $\text{rec}(C)$ is the union of 0 and all the vectors v that direct the rays of C . In particular, C contains a line if and only the vector space

$$(2.1) \quad \text{lin}(C) \stackrel{\text{def.}}{=} \text{rec}(C) \cap (-\text{rec}(C))$$

is nontrivial. The vector space $\text{lin}(C)$ is called the *lineality space of C* . It corresponds to the largest vector space of invariant directions for C .

If E is finite dimensional and C is closed, the recession cone coincides with the asymptotic cone.

Extreme points, extremal rays, faces. An *extreme point* of C is a point $p \in C$ such that $C \setminus \{p\}$ is convex. An *extremal ray* of C is a ray $\rho \in C$ such that if $x, y \in C$ and $]x, y[$ intersects ρ , then $]x, y[\subset \rho$. If C contains the endpoint p of ρ (e.g., if C is linearly closed), this is equivalent to p being an extreme point of C and $C \setminus \rho$ being convex.

Following [20, 29], if $p \in C$, the smallest face of C that contains p is the union of $\{p\}$ and all the open segments in C that have p as an inner point. We denote it by $\mathcal{F}_C(p)$. The (co)dimension of $\mathcal{F}_C(p)$ is defined as the (co)dimension of its affine hull. The collection of all elementary faces, $\{\mathcal{F}_C(p)\}_{p \in C}$, is a partition of C . Extreme points correspond to the zero-dimensional faces of C , while extreme rays are (generally a strict subcollection of the) one-dimensional faces.

Quotient by lines. As noted above, if C is linearly closed, it contains a line if and only if the vector space $\text{lin}(C)$ defined in (2.1) is nontrivial. In that case, letting W denote some linear supplement to $\text{lin}(C)$, we may write

$$(2.2) \quad C = \tilde{C} + \text{lin}(C) \text{ with } \tilde{C} \stackrel{\text{def.}}{=} C \cap W$$

and the corresponding decomposition is unique (i.e., any element of C can be decomposed in a unique way as the sum of an element of \tilde{C} and $\text{lin}(C)$). The convex set \tilde{C} (isomorphic to the projection of C onto the quotient space $E/\text{lin}(C)$) is then linearly closed, and the decomposition of C in elementary faces is exactly given by the partition $\{\mathcal{F}_{\tilde{C}}(p) + \text{lin}(C)\}_{p \in \tilde{C}}$, where $\mathcal{F}_{\tilde{C}}(p)$ is the smallest face of p in \tilde{C} .

One may check that \tilde{C} contains no line, as its recession cone $\text{rec}(\tilde{C})$, the projection of $\text{rec}(C)$ onto W parallel to $\text{lin}(C)$, is a *salient* convex cone.

3. Abstract representer theorems.

3.1. Main result. Our main result describes the facial structure of the solution set to

$$(P) \quad \min_{u \in E} R(u) \quad \text{s.t.} \quad \Phi u = y,$$

where $y \in \mathbb{R}^m$, $\Phi : E \rightarrow \mathbb{R}^m$ is a linear operator, and $m \leq \dim E$, $m < +\infty$. In the following, let t^* denote the *optimal value* of (P), S^* denote its *solution set*, and C^* denote the *corresponding level set* of R ,

$$(3.1) \quad C^* \stackrel{\text{def.}}{=} \{u \in E : R(u) \leq t^*\}.$$

THEOREM 3.1. *Let $R : E \rightarrow \mathbb{R} \cup \{+\infty\}$ be a convex function. Assume that $\inf_E R < t^* < +\infty$, that S^* is nonempty, and that the convex set C^* is linearly closed and contains no line. Let $p \in S^*$ and let j be the dimension of the face $\mathcal{F}_{S^*}(p)$. Then p belongs to a face of C^* with dimension at most $m + j - 1$.*

In particular, if $j < +\infty$, p can be written as a convex combination of

- $m + j$ extreme points of C^* , or
- $m + j - 1$ points of C^* , each an extreme point of C^* or in an extreme ray of C^* .

Moreover, $\text{rec}(S^) = \text{rec}(C^*) \cap \ker(\Phi)$ and therefore $\text{lin}(S^*) = \text{lin}(C^*) \cap \ker(\Phi)$.*

The proof of Theorem 3.1 is given in Appendix A.1. Before extending this theorem to a wider setting, let us formulate some remarks.




FIG. 1. An illustration of Theorem 3.1 for $m = 2$. The solution set $\mathcal{S}^* = C^* \cap \Phi^{-1}(\{y\})$ is made of an extreme point and an extreme ray. The extreme point is a convex combination of $\{e_0, e_1\}$. Depending on their position, the points in the ray are a convex combination of $\{e_0, e_1, e_2\}$ or a pair of points, one in ρ_1 and the other in ρ_2 .

Remark 3.2 (extreme points and extreme rays of \mathcal{S}^*). In particular ($j = 0$), each extreme point of \mathcal{S}^* is a convex combination of m extreme points of C^* or a convex combination of $m - 1$ points of C^* , each an extreme point of C^* or in an extreme ray. Similarly ($j = 1$), each point on an extreme ray of \mathcal{S}^* is a convex combination of $m + 1$ extreme points of C^* or a convex combination of m points of C^* , each an extreme point of C^* or in an extreme ray. Hence, provided the assumptions of Klee’s theorem (see (1.9)) hold, Theorem 3.1 completely characterizes the solution set. An illustration is provided in Figure 1.³

Remark 3.3 (the hypothesis $\inf_E R < t^*$). We have focused on the case $t^* > \inf_E R$ in the theorem, since the case $t^* = \min R$ is easier. In that case, $M = \Phi^{-1}(\{y\})$ can be in an arbitrary position (i.e., not necessarily tangent) w.r.t. $C^* = \operatorname{argmin} R$, and one can only use the general Dubins–Klee theorem [20, 29] to describe their intersection. As a result, the conclusions of Theorem 3.1 are slightly weakened, p belongs to a face of C^* with dimension $m + j$, and one must add one more point in the convex combination (e.g., for $j = 0$, each extreme point of \mathcal{S}^* is a convex combination of $m + 1$ extreme points of C^* , or m points of C^* , each an extreme point of C^* or in an extreme ray).

Remark 3.4 (gauge functions or seminorms). A common practice in inverse problems is to consider positively homogeneous regularizers R , such as (semi)norms or gauge functions of convex sets. In that case the extreme points of C^* correspond, up to a rescaling, to the extreme points of $\{u \in E : R(u) \leq 1\}$. In several cases of interest, the extreme points of such convex sets are well understood; see section 4 for examples in Banach spaces or, for instance, [13, section 2.2] for examples in finite-dimensional spaces.

Remark 3.5 (extension to semistrictly quasi-convex functions). Theorem 3.1 can be extended to the case where R is a semistrictly quasi-convex function. A function R is said to be *semistrictly quasi-convex* [14] if it is quasi-convex and if

$$R(x) < R(y) \implies R(\lambda x + (1 - \lambda)y) < R(y) \quad \forall \lambda \in]0, 1[.$$

³Color figures are available in the online version of this paper.

In other words, semistrictly quasi-convex functions are functions that, when restricted to a line, are successively decreasing, constant, and increasing on their domain. In comparison, strictly quasi-convex functions are successively decreasing and increasing, while quasi-convex functions are successively nonincreasing and nondecreasing.

The set of semistrictly quasi-convex functions is a subset of quasi-convex functions and it contains all convex and strictly quasi-convex functions. In the proof of Theorem 3.1, only the semistrictly quasi-convex property is required to ensure that (A.2) holds.

Remark 3.6 (topological properties). The assumption that C^* is linearly closed is fulfilled in most practical cases, since E is usually endowed with the topology of a Banach (or locally convex) vector space and R is assumed to be lower semicontinuous (so as to guarantee the existence of a solution to (\mathcal{P})). Note also that if R is lower semicontinuous on any line (for the natural topology of the line), the set C^* is linearly closed.

3.2. The case of level sets containing lines. The reader might be intrigued by the assumption of Theorem 3.1 that C^* contains no line, since in several applications the regularizer R is invariant by the addition of, e.g., constant functions or low-degree polynomials (see section 4). In that case, one is generally interested in the nonconstant or nonpolynomial part, and it is natural to consider a quotient problem for which the theorem applies. We describe below (see Corollary 3.7) how our result extends to the case where C^* contains lines.




FIG. 2. Taking the quotient by $K = \text{lin}(C^*)$ yields a level set \tilde{C}^* with no line. In this figure, to simplify the notation, we have omitted the isomorphism ψ_K (i.e., in this figure C^* should be replaced with $\psi_K(C^*)$, and similarly for S^* and $\Phi^{-1}(\{y\})$).

If C^* is linearly closed and contains some line, it is translation invariant in the corresponding direction. The collection of all such directions is the lineality space of C^* (see section 2); we denote it by $K \stackrel{\text{def.}}{=} \text{lin}(C^*)$ (typically, if R is the composition of a linear operator and a norm, K is the *kernel* of that linear operator). Let $\pi_K : E \rightarrow E/K$ be the canonical projection map. We recall that there exists a linear isomorphism $\psi_K : E \rightarrow (E/K) \times K$ such that the first component of $\psi_K(p)$ is $\pi_K(p)$ for all $p \in E$. We may now describe the equivalence classes (modulo K) of the solutions.

COROLLARY 3.7. *Let $R : E \rightarrow [-\infty, +\infty]$ be a convex function. Assume that $\inf_E R < t^* < +\infty$, that \mathcal{S}^* is nonempty, and that the convex set C^* is linearly closed. Let $K \stackrel{\text{def.}}{=} \text{lin}(C^*)$ be the lineality space of C^* and $d \stackrel{\text{def.}}{=} \dim \Phi(K)$. Let $p \in \mathcal{S}^*$, let $\pi_K(p)$ denote its equivalence class, and let j be the dimension of the face $\mathcal{F}_{\pi_K(\mathcal{S}^*)}(\pi_K(p))$.*

Then, $\pi_K(p)$ belongs to a face of $\pi_K(C^)$ with dimension at most $m + j - d - 1$. In particular, if $j < +\infty$,*

- $\pi_K(p)$ is a convex combination of $m + j - d$ extreme points of $\pi_K(C^*)$, or
- $\pi_K(p)$ is a convex combination of $m + j - d - 1$ points of $\pi_K(C^*)$, each an extreme point of $\pi_K(C^*)$ or in an extreme ray of $\pi_K(C^*)$.

As a result, letting $\tilde{q}_1, \dots, \tilde{q}_r$ denote those extreme points (or points in extreme rays),

$$(3.2) \quad p = \sum_{i=1}^r \theta_i \psi_K^{-1}(\tilde{q}_i, 0) + u_K, \quad \text{where } \theta_i \geq 0, \quad \sum_{i=1}^r \theta_i = 1, \quad \text{and } u_K \in K.$$

The proof of Corollary 3.7 is given in Appendix A.2.

One can have an explicit representation with elements of E of a solution $p \in \mathcal{S}^*$. Indeed, let W be some linear complement to $K = \text{lin}(C^*)$. One may decompose $C^* = \tilde{C}^* + K$, where $\tilde{C}^* = C \cap W$, and observe that $\pi_K(C^*)$ and \tilde{C}^* are isomorphic. In this case, Corollary 3.7 implies that p can be written as the sum of one point in $\text{lin}(C^*)$ and of a convex combination of

- $m + j - d$ extreme points of \tilde{C}^* , or
- $m + j - 1 - d$ points of \tilde{C}^* , each an extreme point of \tilde{C}^* or in an extreme ray of \tilde{C}^* .

3.3. Extensions to data fitting functions. In this section, we discuss the extension of the above results to more general problems of the form

$$(\mathcal{P}_f) \quad \inf_{u \in E} f(\Phi u) + R(u),$$

where $f : \mathbb{R}^m \rightarrow \mathbb{R} \cup \{+\infty\}$ is an arbitrary fidelity term.

3.3.1. Convex data fitting term. When f is a convex data fitting function f , we get the following result.

COROLLARY 3.8. *Assume that f is convex and that the solution set \mathcal{S}_f^* of (\mathcal{P}_f) is nonempty. Let $p \in \mathcal{S}_f^*$ such that $C^* \stackrel{\text{def.}}{=} \{u \in E, R(u) \leq R(p)\}$ is linearly closed, $K \stackrel{\text{def.}}{=} \text{lin } C^*$ and let j be the dimension of the face $\mathcal{F}_{\pi_K(\mathcal{S}_f^*)}(\pi_K(p))$.*

If $\inf_E R < R(p)$, then the conclusions of Corollary 3.7 (or Theorem 3.1 if $K = \{0\}$) hold. If $\inf_E R = R(p)$, they hold with one more dimension (see Remark 3.3).

Let us recall that, in view of Remark 3.6, if E is a topological vector space and R is lower semicontinuous, C^* is closed regardless of the choice p .

Proof. Let $y = \Phi p$ and consider the following problem:

$$(\mathcal{P}_{\{y\}}) \quad \min_{u \in E} R(u) \quad \text{s.t.} \quad \Phi u = y.$$

Let $\mathcal{S}_{\{y\}}^*$ denote its solution set. It is a convex subset of \mathcal{S}_f^* , with $p \in \mathcal{S}_{\{y\}}^*$. Additionally, if j is the dimension of $\mathcal{F}_{\pi_K(\mathcal{S}_f^*)}(p)$ (respectively, $\mathcal{F}_{\mathcal{S}_f^*}(p)$ if C^* contains no line), then the face $\mathcal{F}_{\pi_K(\mathcal{S}_{\{y\}}^*)}(p)$ (respectively, $\mathcal{F}_{\mathcal{S}_{\{y\}}^*}(p)$) has dimension at most j , since $\mathcal{S}_{\{y\}}^* \subseteq \mathcal{S}_f^*$. It suffices to apply Corollary 3.7 (respectively, Theorem 3.1) to obtain the result. □

Remark 3.9 (the case of a strictly convex function). In the case when f is strictly convex, it is known that $\Phi\mathcal{S}_f^*$ is a singleton, which means that $\mathcal{S}_{\{y\}}^* = \mathcal{S}_f^*$.

Remark 3.10 (the case of quasi-convex functions). The result actually holds whenever the solution set \mathcal{S}_f^* is convex. In particular, this property holds when f is quasi-convex and R is convex.

3.3.2. Nonconvex function. In the general case, i.e., when $f : \mathbb{R}^m \rightarrow \mathbb{R} \cup \{+\infty\}$ is an arbitrary function, it is difficult to describe the structure of the solution set. However, one may choose a solution p_0 (as before, provided it exists) and observe that it is also a solution to (\mathcal{P}) for $y \stackrel{\text{def.}}{=} \Phi p_0$. Then, one may apply Corollary 3.7, but the difficult part is that the dimensions j to consider are with respect to the solution set \mathcal{S}^* of the *convex* problem (\mathcal{P}) . Nevertheless, if one is able to assert that the solution set \mathcal{S}^* has at least one extreme point p , then Corollary 3.7 ensures that p can be written in the form (3.2), where $r \leq m$ and the \tilde{q}_i 's are extreme points (or points in extreme rays) of C^* . Since p must also be a solution to (\mathcal{P}_f) , one obtains that there exists a solution to (\mathcal{P}_f) of the form (3.2).

3.4. Ensuring the existence of extreme points. It is important to note that, in Theorem 3.1, the existence of a face of \mathcal{S}^* with dimension j (or of an extreme point for $j = 0$) is not guaranteed. The convex set \mathcal{S}^* might not even have any finite-dimensional face! For instance, let E be the space of Lebesgue-integrable functions on $[0, 1]$. If $R(u) = \int_0^1 |u(x)|dx$, $\Phi : E \rightarrow \mathbb{R}$ is defined by $u \mapsto \int_0^1 u(x)dx$, and $y = 1$, then

$$(3.3) \quad \mathcal{S}^* = \left\{ u \in E : \int_0^1 |u(x)|dx \leq 1 \text{ and } \int_0^1 u(x)dx = 1 \right\}.$$

It is possible to prove that such a set \mathcal{S}^* does not have any extreme point. As a consequence \mathcal{S}^* does not have any finite-dimensional face (otherwise an extreme point of the closure of a face would be an extreme point of \mathcal{S}^*).

However, Theorem 3.1 (in fact the Dubins–Klee theorem [20, 29]) asserts that *if there is* a finite-dimensional face in \mathcal{S}^* , then C^* indeed has extreme points (and possibly extreme rays), and the convex combinations of such points generate the abovementioned face.

As a result, it is crucial to be able to assert a priori the existence of some finite-dimensional face for \mathcal{S}^* , and this is where topological arguments come into play. If E is endowed with the topology of a locally convex (Hausdorff) vector space, Theorems 3.3 and 3.4 of [30], which generalize the celebrated Krein–Milman theorem, state that \mathcal{S}^* *has an extreme point provided*

- \mathcal{S}^* is nonempty and convex,
- \mathcal{S}^* contains no line, and
- \mathcal{S}^* is closed and locally compact.

The last two conditions hold in particular if \mathcal{S}^* is compact. Moreover, as in Corollary 3.7, the second condition can be ensured by considering a suitable quotient map, provided it preserves the other topological properties (e.g., if $\text{lin}(C)$ has a topological complement).

Remark 3.11. Whereas local compactness is a very strong property for topological vector spaces (implying their finite-dimensionality; see [8, Chapter 1, Theorem 3]), it is not so difficult to ensure the local compactness of \mathcal{S}^* in practice. Indeed, very often, even the existence of solutions is usually ensured using compactness arguments

for a suitable weak or weak- $*$ topology. The unbounded cases require more specific arguments, but let us mention that there are examples of cones that are locally compact without being contained in any finite-dimensional vector space. In subsection 4.2.2, we discuss the example of the cone $\mathcal{M}^+(\Omega)$ of nonnegative measures over a compact set for the weak- $*$ topology. Another example of a locally compact convex cone is

$$\mathcal{C} = \left\{ x \in \mathbb{R}^{\mathbb{N}} \text{ such that } x_n \geq 0 \text{ and } \sum_{n \in \mathbb{N}} x_n \omega_n \leq \sum_{n \in \mathbb{N}} x_n < +\infty \right\} \subseteq \ell^1(\mathbb{N}),$$

for some nondecreasing positive sequence $(\omega_n)_n$ converging to $+\infty$ (note that $\omega_0 \leq 1$ for the cone to be nonempty). The cone \mathcal{C} is locally compact for the strong topology. Indeed, consider the intersection K of the cone \mathcal{C} and the strong unit ball, namely $K = \{x = (x_n)_n : x_n \geq 0, \sum_n \omega_n x_n \leq \sum_n x_n \leq 1\}$ and consider a sequence of elements of K denoted $(x^k)_k \subset K$. Using a diagonal argument, each $(x_n^k)_k$ converges to some $\bar{x}_n \geq 0$. Furthermore, using that $\{n : \omega_n < 1\}$ is finite and Fatou's lemma, it holds that

$$\begin{aligned} \sum_n \bar{x}_n &\leq \liminf_k \sum_n x_n^k \leq 1 \quad \text{and} \\ \sum_{n:\omega_n \geq 1} (\omega_n - 1)\bar{x}_n &\leq \liminf_k \sum_{n:\omega_n \geq 1} (\omega_n - 1)x_n^k \\ &\leq \liminf_k \sum_{n:\omega_n < 1} (1 - \omega_n)x_n^k = \sum_{n:\omega_n < 1} (1 - \omega_n)\bar{x}_n, \end{aligned}$$

and we deduce that $\bar{x} = (\bar{x}_n)_n \in K$. Furthermore, one has

$$\|x^k - \bar{x}\|_1 = \sum_{n=0}^M |x_n^k - \bar{x}_n| + \sum_{n>M} |x_n^k - \bar{x}_n|$$

and M will be chosen later. Finally,

$$\sum_{n>M} |x_n^k - \bar{x}_n| \leq \sum_{n>M} (x_n^k + \bar{x}_n) \leq (1/\omega_M) \sum_{n>M} \omega_n (x_n^k + \bar{x}_n) \leq 2/\omega_M$$

since $\omega_n/\omega_M \geq 1$ for $n > M$. Therefore, choosing M large enough ensures that the second term, $\sum_{n>M} |x_n^k - \bar{x}_n|$, is less than some $\varepsilon > 0$. Choosing k large enough leads to $\|x^k - \bar{x}\|_1 \leq 2\varepsilon$.

4. Application to some popular regularizers. We now show that the extreme points and extreme rays of numerous convex regularizers can be described analytically, allowing us to describe important analytical properties of the solutions of some popular problems. The list given below is far from being exhaustive, but it gives a taste of the diversity of applications targeted by our main results.

4.1. Finite-dimensional examples. We first consider examples where one has $\dim E < +\infty$. In that case, Φ is continuous, and since the considered regularizations R are lower semicontinuous, we deduce that

- the level set C^* is closed,
- the solution set $\mathcal{S}^* = C^* \cap \Phi^{-1}(\{y\})$ is closed and locally compact (even compact in most cases), and hence it admits extreme points provided it contains no line.

4.1.1. Nonnegativity constraints. In a large number of applications, the signals to recover are known to be nonnegative. In that case, one may be interested in solving nonnegatively constrained problems of the form

$$(4.1) \quad \inf_{u \in \mathbb{R}_+^n} f(\Phi u - y).$$

An important instance of this class of problems is the nonnegative least squares [33], which finds its motivation in a large number of applications. Applying the results of section 3 to problem (4.1) yields the following result.

PROPOSITION 4.1. *If the solution set of (4.1) is nonempty, then it contains a solution that is m -sparse. In addition if f is convex and the solution set is compact, then its extreme points are m -sparse.*

Choosing R as the characteristic function of \mathbb{R}_+^n , the result simply stems from the fact that the extreme rays of the positive orthant are the half lines $\{\alpha e_i, \alpha \geq 0\}$, where $(e_i)_{1 \leq i \leq n}$ denote the elements of the canonical basis. We have to consider m atoms and not $m - 1$ since $t^* = \inf R = 0$; see Remark 3.3.

It may come as a surprise to some readers, since the usual way to promote sparsity consists in using ℓ^1 -norms. This type of result is one of the main ingredients of [18] that shows that the ℓ^1 -norm can sometimes be replaced by the indicator of the positive orthant when sparse *positive* signals are looked after.

4.1.2. Linear programming. Let $\psi \in \mathbb{R}^n$ be a vector and $\Phi \in \mathbb{R}^{m \times n}$ be a matrix and consider the following linear program in standard (or equational) form:

$$(4.2) \quad \inf_{\substack{u \in \mathbb{R}_+^n \\ \Phi u = y}} \langle \psi, u \rangle.$$

Applying Theorem 3.1 to problem (4.2), we get the following well-known result (see, e.g., [34, Theorem 4.2.3]).

PROPOSITION 4.2. *Assume that the solution set of (4.2) is nonempty and compact. Then, its extreme points are m -sparse, i.e., of the form*

$$(4.3) \quad u = \sum_{i=1}^m \alpha_i e_i, \quad \alpha_i \geq 0,$$

where e_i denotes the i th element of the canonical basis.

In the linear programming literature, solutions of this kind are called *basic solutions*. To prove the result, we can reformulate (4.2) as follows:

$$(4.4) \quad \inf_{\substack{(u,t) \in \mathbb{R}_+^n \times \mathbb{R} \\ \Phi u = y \\ \langle \psi, u \rangle = t}} t.$$

Letting $R(u, t) = t + \iota_{\mathbb{R}_+^n}(u)$, we get $\inf R = -\infty$. Hence, if a solution exists, we only need to analyze the extreme points and extreme rays of

$$C^* = \{(x, t) \in \mathbb{R}^n \times \mathbb{R}, R(x, t) \leq t^*\} = \mathbb{R}_+^n \times]-\infty, t^*],$$

where t^* denotes the optimal value. The extreme rays of this set (a shifted nonnegative orthant) are of the form $\{\alpha e_i, \alpha > 0\} \times \{t^*\}$ or $\{0\} \times]-\infty, t^*[$. In addition, C^* possesses only one extreme point $(0, t^*)$. Applying Theorem 3.1, we get the desired result.

4.1.3. ℓ^1 analysis priors. An important class of regularizers in the finite-dimensional setting $E = \mathbb{R}^n$ contains the functions of the form $R(u) = \|Lu\|_1$, where L is a linear operator from \mathbb{R}^n to \mathbb{R}^p . They are sometimes called *analysis priors*, since the signal u is “analyzed” through the operator L . Remarkable practical and theoretical results have been obtained using this prior in the fields of inverse problems and compressed sensing, even though—to the best of the authors’ knowledge—many of its properties are still quite obscure.

Since R is 1-homogeneous, it suffices to describe the extremality properties of the unit ball $C = \{u \in \mathbb{R}^n, \|Lu\|_1 \leq 1\}$ to use our theorems. The lineality space is simply equal to $\text{lin}(C) = \ker(L)$. Let $K = \ker(L)$, K^\perp denote the orthogonal complement of K in \mathbb{R}^n , and $L^+ : \mathbb{R}^n \rightarrow K^\perp$ denote the pseudoinverse of L . We can decompose C as $C = K + C_{K^\perp}$ with $C_{K^\perp} = C \cap K^\perp$. Our ability to characterize the extreme points of C_{K^\perp} depends on whether L is surjective or not. Indeed, we have

$$(4.5) \quad \text{ext}(C_{K^\perp}) = L^+(\text{ext}(\text{ran}(L) \cap B_1^p)),$$

where B_1^p is the unit ℓ^1 -ball defined as

$$B_1^p = \{z \in \mathbb{R}^p, \|z\|_1 \leq 1\}.$$

Property (4.5) simply stems from the fact that C_{K^\perp} and $D = \text{ran}(L) \cap B_1^p$ are in bijection through the operators L and L^+ .

The case of a surjective operator L . When L is surjective (hence $p \leq n$), the problem becomes rather elementary.

PROPOSITION 4.3. *If L is surjective, the extreme points u of C_{K^\perp} are $\text{ext}(C_{K^\perp}) = (\pm L^+ e_i)_{1 \leq i \leq p}$, where e_i denotes the i th element of the canonical basis. Consider problem (\mathcal{P}_f) and assume that at least one solution exists. Then problem (\mathcal{P}_f) has solutions of the form*

$$(4.6) \quad u^* = \sum_{i \in I} \alpha_i L^+ e_i + u_K,$$

where $u_K \in \ker(L)$ and $I \subset \{1, \dots, p\}$ is a set of cardinality $|I| \leq m - \dim(\Phi \ker(L))$.

The proof of Proposition 4.3 follows from Corollary 3.7 and subsection 3.3.2, with $j = 0$ and observing that π_K is the orthogonal projection on K^\perp .

The case of an arbitrary operator L . When L is not surjective, the description of the extreme points $\text{ext}(D)$ becomes untractable in general. A rough upper bound on the number of extreme points can be obtained as follows. We assume that L has full rank n and that $\text{ran}(L)$ is in general position. The extreme points of $\text{ran}(L) \cap B_1^p$ correspond to the intersections of some faces of the ℓ^1 -ball with a subspace of dimension n . In order for some k -face to intersect the subspace $\text{ran}(L)$ on a singleton, k should satisfy $n + k - p = 0$, i.e., $k = p - n$. The k -faces of the ℓ^1 -ball contain $(k+1)$ -sparse elements. The number of $(k+1)$ -sparse supports in dimension p is $\binom{p}{k+1}$. For a fixed support, the number of sign patterns is upper-bounded by 2^{k+1} . Hence, the maximal number of extreme points satisfies $|\text{ext}(\text{ran}(L) \cap B_1^p)| \leq 2^{k+1} \binom{p}{k+1}$. This upper bound is pessimistic since the subspace may not cross all extreme points, but it provides an idea of the combinatorial explosions that may happen in general.

Notice that enumerating the number of faces of a polytope is usually a hard problem. For instance, the Motzkin conjecture [36], which upper bounds the number of k faces of a d polytope with z vertices, was formulated in 1957 and only solved by McMullen [35] in 1970.

4.1.4. Matrix examples. In several applications, one deals with optimization problems in matrix spaces. The following regularizations/convex sets are commonly used.

Semidefinite matrix constraint. Similarly to subsection 4.1.1, one may consider in $\mathbb{R}^{n \times n}$ the following constrained problem:

$$(4.7) \quad \inf_{M \succeq 0} f(\Phi M - y),$$

where $M \succeq 0$ means that M must be symmetric positive semidefinite (p.s.d.). The extreme rays of the positive semidefinite cone C^* are the p.s.d. matrices of rank 1 (see, for instance, [15, section 2.9.2.7]). Hence, arguing as in subsection 4.1.1, we may deduce that if there exists a solution to (4.7), there is also a solution that has rank (at most) m .

However, that conclusion is not optimal, as in that case a theorem by Barvinok [5, Theorem 2.2] ensures that there exists a solution M with

$$(4.8) \quad \text{rank}(M) \leq \frac{1}{2} (\sqrt{8m+1} - 1).$$

To understand the gap with Barvinok's result, let us note that the p.s.d. cone has a very special structure that makes the Minkowski–Carathéodory theorem (or its extension by Klee) too pessimistic. By [15, section 2.9.2.3], given $M \succeq 0$, the smallest face of the p.s.d. cone that contains M (i.e., the set of p.s.d. matrices that have the same kernel) has dimension

$$(4.9) \quad d = \frac{1}{2} \text{rank}(M)(\text{rank}(M) + 1).$$

Equivalently, if the smallest face that contains M has dimension d , then $\text{rank}(M) = \frac{1}{2}(\sqrt{8d+1} - 1)$; hence, M is a convex combination of $\frac{1}{2}(\sqrt{8d+1} - 1)$ points in extreme rays, a value which is less than the value d predicted by Klee's extension of Carathéodory's theorem.

As a result, we recover Barvinok's result by noting that, as ensured by the first claim of Theorem 3.1,⁴ any extreme point M of the solution set belongs to a face of dimension m . Then, taking into account (4.9) improves upon the second claim of Theorem 3.1, and we immediately obtain (4.8).

Semidefinite programming. Semidefinite programs are problems of the form

$$(4.10) \quad \inf_{\substack{M \succeq 0 \\ \Phi(M)=y}} \langle A, M \rangle,$$

where $A \in \mathbb{R}^{n \times n}$ is a matrix and $\langle A, M \rangle \stackrel{\text{def.}}{=} \text{Tr}(AM)$. Arguing as in Proposition 4.2, if the solution set of (4.10) is nonempty, our main result allows us to state that its extreme points are matrices of rank m . In view of the above discussion, it is possible to refine this statement and show that (4.8) holds.

The nuclear norm. The nuclear norm of a matrix $M \in \mathbb{R}^{p \times n}$ is often denoted by $\|M\|_*$ and defined as the sum of the singular values of M . Lately, it has gained considerable attention as a regularizer thanks to its applications in matrix completion [11] or blind inverse problems [2]. The geometry of the unit ball $\{M \in \mathbb{R}^{p \times n}, \|M\|_* \leq 1\}$

⁴Or, more precisely, by its variant when $t^* = \inf R$; see Remark 3.3.

is well studied due to its central role in the field of semidefinite programming [38]. Its extreme points are the rank-1 matrices $M = uv^T$, with $\|u\|_2 = \|v\|_2 = 1$.

Combining Theorem 3.1 with this result explains why regularizing problems over the space of matrices with the nuclear norm allows *rank- m solutions* to be recovered.

The rank-sparsity ball. The *rank-sparsity ball* is the set

$$\{M \in \mathbb{R}^{m \times n}, \|M\|_* + \|M\|_1 \leq 1\},$$

where $\|M\|_1$ is the ℓ^1 -norm of the entries of M . The corresponding regularization is sometimes used in order to favor sparse and low-rank matrices. Drusvyatskiy, Vavasis, and Wolkowicz [19] have described the extreme points of this unit ball and proved that the *extreme points M of the rank sparsity ball satisfy $\frac{r(r+1)}{2} - |I| \leq 1$* , where $|I|$ denotes the number of nonzero entries in M and r denotes its rank. This result partly explains why using the rank-sparsity gauge promotes sparse and low-rank solutions. Let us stress that this effect might be better obtained by using different strategies [39, 40].

Bistochastic matrices. A doubly stochastic matrix is a matrix with nonnegative rows and columns summing to one. The set of such matrices is called the Birkhoff polytope. The Birkhoff–von Neumann theorem states that its extreme points are the permutation matrices. We refer the interested reader to [26] for the use of such matrices in DNA sequencing.

4.2. Examples in infinite dimension. In this section, we provide results in infinite-dimensional spaces, which mirror those described in finite dimensions.

4.2.1. Problems formulated in Hilbert or Banach sequence spaces. The case of Hilbert spaces (or countable sequences) can be treated within our formalism, and all the examples presented previously have their natural counterpart in this setting. In the same vein, one can also treat Banach sequence spaces ℓ^p for $1 \leq p \leq \infty$. We do not reproduce the results here due to space limitations. Let us mention, however, that the two works [47, 1] treat this specific case with ℓ^1 regularizers.

4.2.2. Linear programming and the moment problem. Let Ω be a compact metric space, let $\mathcal{M}(\Omega)$ be the set of Radon measures on Ω , and let $\mathcal{M}_+(\Omega) \subseteq \mathcal{M}(\Omega)$ be the cone of nonnegative measures on Ω . Let ψ and $(\phi_i)_{1 \leq i \leq m}$ denote a collection of continuous functions on Ω . Now, let $\Phi : \mathcal{M}(\Omega) \rightarrow \mathbb{R}^m$ be defined by $(\Phi\mu)_i = \langle \phi_i, \mu \rangle$, where $\langle \phi_i, \mu \rangle \stackrel{\text{def.}}{=} \int_{\Omega} \phi_i d\mu$, and consider the following linear program in the standard form:

$$(4.11) \quad \inf_{\substack{\mu \in \mathcal{M}_+(\Omega) \\ \Phi\mu = y}} \langle \psi, \mu \rangle.$$

Applying Theorem 3.1 to problem (4.11), we get Proposition 4.4. We do not provide a proof here, since it mimics very closely the one given for linear programming in finite dimension. The extreme rays of $\mathcal{M}_+(\Omega)$ can be described, arguing as in [3, Theorem 15.9], as the rays directed by the Dirac masses.

PROPOSITION 4.4. *Assume that the solution set (4.11) is nonempty. Then, its extreme points are m -sparse, i.e., of the form*

$$(4.12) \quad u = \sum_{i=1}^m \alpha_i \delta_{x_i}, \quad x_i \in \Omega, \quad \alpha_i \geq 0.$$

To make sure that the above proposition is nontrivial, one may wish to ensure that the solution set \mathcal{S}^* does indeed have extreme points, using arguments from subsection 3.4. It is straightforward that \mathcal{S}^* is convex and does not contain any line. Now, let us endow $\mathcal{M}(\Omega)$ with the weak-* topology (i.e., the coarsest topology for which $\mu \mapsto \int_{\Omega} \eta d\mu$ is continuous for every $\eta \in \mathcal{C}(\Omega)$). By lower semicontinuity, \mathcal{S}^* is closed. Moreover, \mathcal{S}^* is locally compact since the closed convex cone \mathcal{M}_+ is itself locally compact (taking any $\mu \in \mathcal{M}_+(\Omega)$, its neighborhood $\{\nu \in \mathcal{M}_+(\Omega) : \nu(\Omega) \leq \mu(\Omega) + 1\}$ is compact in the weak-* topology).

Proposition 4.4 is well known (see, e.g., [44]). Note that if we optimize the linear form $\langle \psi, u \rangle$ over the set of probability measures instead of the set of nonnegative measures, we get the so-called moment problem [45] for which we can obtain a similar result.

4.2.3. The total variation ball. Let Ω denote an open subset of \mathbb{R}^d and let $\mathcal{M}(\Omega)$ denote the set of Radon measures on Ω . The total variation ball $B_{\mathcal{M}} = \{u \in \mathcal{M}(\Omega), \|u\|_{\mathcal{M}(\Omega)} \leq 1\}$ plays a critical role in problems such as super-resolution [10, 46, 21]. It is compact for the weak-* topology and its extreme points are the Dirac masses: $\text{ext}(B_{\mathcal{M}}) = \{\pm \delta_x, x \in \Omega\}$. Hence, total variation regularized problems of the form

$$\inf_{u \in \mathcal{M}} f(\Phi u) + \|u\|_{\mathcal{M}}$$

yield m -sparse solutions (under an existence assumption). A few variations around this central result were provided in [22].

Demixing of sines and spikes. In [22, Page 262], Fernandez-Granda presents a regularization of the type

$$\|\mu\|_{\mathcal{M}} + \eta \|v\|_1,$$

where $\eta > 0$ is a tuning parameter, μ a complex measure, and $v \in \mathbb{C}^n$ a sparse vector. Define E as the set of (μ, v) , where μ is a complex Radon measure on a domain Ω and $v \in \mathbb{C}^n$. Consider the unit ball

$$B \stackrel{\text{def}}{=} \{(\mu, v) \in E : \|\mu\|_{\mathcal{M}} + \eta \|v\|_1 \leq 1\}.$$

Its extreme points are the points

- $(a\delta_t, 0)$ for all $t \in \Omega$ (δ_t denotes the Dirac mass at point t) and all $a \in \mathbb{C}$ such that $|a| = 1$,
- $(0, ae_k)$ for all $k = 1, \dots, n$ and all $a \in \mathbb{C}$ such that $|a| = 1/\eta$ and e_k denotes the vector with 1 at entry k and 0 otherwise.

Group total variation: Point sources with a common support. In [22, Page 266], Fernandez-Granda presents a regularization of the type

$$\|\mu\|_{\mathcal{M}^n} := \sup_{F: \Omega \rightarrow \mathbb{C}^n, \|F(t)\|_2 \leq 1, t \in \Omega} \int_{\Omega} \langle F(t), \nu(t) \rangle d|\mu|(t),$$

where F is continuous and vanishing at infinity, and μ is a vectorial Radon measure on Ω such that $|\mu|$ -a.e. $\mu = \nu \cdot |\mu|$ with ν a measurable function from Ω onto \mathbb{S}^{n-1} the n -sphere and with $|\mu|$ a positive finite measure on Ω . Consider the unit ball

$$B \stackrel{\text{def}}{=} \{\mu, \|\mu\|_{\mathcal{M}^n} \leq 1\}.$$

Its extreme points are $a\delta_t$ for all $t \in \Omega$ (δ_t denotes the Dirac mass at point t) and all $a \in \mathbb{C}^n$ such that $\|a\|_2 = 1$.

4.2.4. Analysis priors in Banach spaces. The analysis of extreme points of analysis priors in an infinite-dimensional setting is more technical. Fisher and Jerome [23] proposed an interesting result, which can be seen as an extension of (4.6). This result was recently revisited in [48] and [25]. Below, we follow the presentation in [25].

Let Ω denote an open set in \mathbb{R}^d . Let $\mathcal{D}'(\Omega)$ denote the set of distributions on Ω and let $L : \mathcal{D}'(\Omega) \rightarrow \mathcal{D}'(\Omega)$ denote a linear operator with kernel $K = \ker(L)$. We let $E = \{u \in \mathcal{D}'(\Omega), Lu \in \mathcal{M}(\Omega)\}$ and let $\|\cdot\|_K$ denote a seminorm on E , which, restricted to K , is a norm. We define a function space $\mathcal{B}(\Omega)$ as

$$\mathcal{B}(\Omega) = \{u \in E, \|Lu\|_{\mathcal{M}(\Omega)} + \|u\|_K < +\infty\},$$

and equip it with the norm $\|u\|_{\mathcal{B}(\Omega)} = \|Lu\|_{\mathcal{M}(\Omega)} + \|u\|_K$. We assume that L is surjective, i.e., $\mathcal{M}(\Omega) = L(\mathcal{B}(\Omega))$, and that K has a topological complement (with respect to $\mathcal{B}(\Omega)$), which we denote by K^\perp . This setting encompasses all surjective Fredholm operators, for instance. Under the stated assumptions, we can define a pseudoinverse L^+ of L relative to K^\perp [7].

The representer theorems in [23, 48, 25] can be obtained using Theorem 3.1 as exemplified below.

PROPOSITION 4.5. *Let $B = \{u \in \mathcal{B}(\Omega), \|Lu\|_{\mathcal{M}(\Omega)} \leq 1\}$. Then the extreme points of the set $C_{K^\perp} = B \cap K^\perp$ are of the form $\pm L^+ \delta_x$ for $x \in \Omega$.*

Let $f : \mathbb{R}^m \rightarrow \mathbb{R} \cup \{+\infty\}$ denote a convex function and define

$$\mathcal{S}^* = \operatorname{argmin}_{u \in \mathcal{B}(\Omega)} f(\Phi u) + \|Lu\|_{\mathcal{M}(\Omega)}.$$

Assume that \mathcal{S}^ is nonempty and does not contain 0. Then the extreme points (if they exist) of $\pi_K(\mathcal{S}^*)$ are of the form $u = \sum_{i=1}^m \alpha_i L^+ \delta_{x_i}$.*

Proof. The proof mimics the finite-dimensional case (4.6). First notice that $B = L^{-1}(B_{\mathcal{M}})$, where $L^{-1}(\{\mu\})$ is the preimage of μ by L and $B_{\mathcal{M}}$ is the unit total variation ball. We have $L^{-1}(B_{\mathcal{M}}) = L^+(B_{\mathcal{M}}) + K$ and we can identify C_{K^\perp} with $L^+(B_{\mathcal{M}})$. Since L^+ is bijective from $\mathcal{M}(\Omega)$ to K^\perp , the extreme points of C_{K^\perp} are the image by L^+ of the Dirac masses.

The end of the proposition follows from Corollary 3.8 and from the fact that the lineality space of $\{u \in \mathcal{B}(\Omega), \|Lu\|_{\mathcal{M}(\Omega)} \leq 1\}$ is equal to K . □

Let us mention that, although the description of the extreme points follows directly from the results of section 3, proving the existence of minimizers and the existence of extreme points is a considerably more difficult problem that needs a careful choice of topologies. Unser, Fageot, and Ward [48] provide a systematic way to construct Banach spaces and pseudoinverse L^+ for “spline admissible operators” L such as the fractional Laplacian. In addition, they prove the existence of solutions by adding weak- $*$ continuity assumptions on the sensing operator Φ .

4.2.5. The total gradient variation. Since its introduction into the field of image processing [41], the total gradient variation has proved to be an extremely valuable regularizer in diverse fields of data science and engineering. It is defined, for any locally integrable function u , as

$$TV(u) \stackrel{\text{def}}{=} \sup \left(\int u \operatorname{div}(\phi) dx, \phi \in C_c^1(\mathbb{R}^d)^d, \sup_{x \in \mathbb{R}^d} \|\phi(x)\|_2 \leq 1 \right).$$




FIG. 3. Illustration for the total gradient variation problem $\min\{TV(u) : \Phi(u) = y\}$. Here, Φ is a linear mapping giving access to three measurements, $y \in \mathbb{R}^3$, by taking the mean of an image u of size 200×200 on three different disks represented in (a). The TV problem is solved using a primal-dual algorithm, also known as the Chambolle–Pock algorithm [12]. The recovered image is displayed in (b): it can be represented as the sum of three indicator functions of simple sets.

If the above quantity is finite, we say that u has bounded variation and its gradient Du is a Radon measure with

$$TV(u) = \int_{\mathbb{R}^d} |Du| = \|Du\|_{(\mathcal{M}(\mathbb{R}^d))^d}.$$

Working in $E = L^{d/(d-1)}(\mathbb{R}^d)$, one is led to consider the convex set

$$C = \{u \in E, TV(u) \leq 1\},$$

referred to as the TV unit ball.

The generalized gradient operator is not a surjective operator. Hence, the analysis of subsection 4.2.4 cannot help us to find the extreme points of the TV ball. Still, those were described in the 1950s by Fleming [24] and refined analyses have been proposed more recently by Ambrosio, Caselles, Masnou, and Morel in [4].

THEOREM 4.6 (extreme points of the TV ball [24, 4]). *The extreme points of the unit TV unit ball are the indicators of simple sets normalized by their perimeter, i.e., functions of the form $u = \pm \frac{\mathbf{1}_F}{TV(\mathbf{1}_F)}$, where F is an indecomposable and saturated subset of \mathbb{R}^d .*

Informally, the simple sets of \mathbb{R}^d are the simply connected sets with no hole. We refer the reader to [4] for more details. Using Theorem 4.6 in conjunction with our results tells us that functions minimizing the total variation subject to a finite number of linear constraints can be expressed as the sum of a small number of indicators of simple sets; see, for instance, Figure 3, which is yet another theoretical result explaining the common observation that total variation tends to produce staircasing [37].

5. Conclusion. In this paper we have developed representer theorems for convex regularized inverse problems (1.1), based on fundamental properties of the geometry of convex sets: the solution set can be entirely described using convex combinations of a small number of extreme points and extreme rays of the regularizer level set.

Obviously, the conclusion of Theorem 3.1 is only nontrivial when C^* has a “sufficiently flat boundary,” in the sense that two or more faces of C^* have dimension larger than m . For instance, if C^* is strictly convex (i.e., has only zero-dimensional

faces, except its interior⁵), then the solution set \mathcal{S}^* is always reduced to a single extreme point of C^* ! Nevertheless, several regularizers that are commonly used in the literature (notably sparsity-promoting ones) have that flatness property, and Theorem 3.1 then provides interesting information on the structure of the solution set, as illustrated in section 4.

To conclude, the structure theorem presented in this paper highlights the importance of describing the extreme points and extreme rays of the regularizer: this yields a fine description of the set of solutions of variational problems of the form (1.1). Our theorem also suggests a principled way to design a regularizer. If a particular family of solutions is expected, then one may construct a suitable regularizer by taking the convex hull of this family. Finally, representer theorems have had a lot of success in the fields of approximation theory and machine learning [42], in the frame of reproducible kernel Hilbert spaces. One reason for this success is that they allow us to design efficient numerical procedures that yield *infinite-dimensional* solutions by solving *finite-dimensional* linear systems. Such numerical procedures have recently been extended to the case of Banach spaces for some simple instances of the problems described in this paper [22, 16, 25]. The price to pay when going from a Hilbert space to a Banach space is that semi-infinite convex programs have to be solved instead of simpler linear systems. We foresee that the results in this paper may help to design new efficient numerical procedures, since they allow us to parameterize the solutions using only extreme points and extreme rays.

Appendix A. Proofs of section 3.

A.1. Proof of Theorem 3.1. The set of solutions \mathcal{S}^* is precisely $C^* \cap \Phi^{-1}(\{y\})$, and the statement of the theorem amounts to describing its elementary faces. Since $\Phi^{-1}(\{y\})$ is an affine space with codimension at most m , the main theorem of [29] almost provides the desired conclusion, but, for our particular case, it yields one extreme point/ray too many. We now show how to obtain the correct number.

Let p be a point of \mathcal{S}^* such that $\mathcal{F}_{\mathcal{S}^*}(p)$ has dimension j . Up to a translation, it is not restrictive to assume that $p = 0$, so that $\mathcal{S}^* = C^* \cap \ker \Phi$. Let T be the union of $\{0\}$ and all the lines ℓ such that $C^* \cap \ell$ contains an open interval containing 0. Note that T is a linear space, the linear hull of $\mathcal{F}_{C^*}(0)$.

We claim that $\text{codim}_T(T \cap \ker \Phi) \leq m - 1$. By contradiction, assume that there is a complement Z to $T \cap \ker \Phi$ in T with dimension m . Then $\Phi|_Z$ has rank m and is a bijection; hence, we may define

$$z \stackrel{\text{def.}}{=} -\frac{\theta}{(1-\theta)}(\Phi|_Z)^{-1}(\Phi u_0) \in Z \subset T,$$

where $\theta \in]0, 1[$ and $u_0 \in C^*$ is such that $\inf R \leq R(u_0) < t^*$. For θ small enough, $z \in C^*$, and hence $R(z) \leq t^*$. Moreover,

$$(A.1) \quad \Phi((1-\theta)z + \theta u_0) = (1-\theta)\Phi z + \theta\Phi u_0 = 0,$$

so that $(1-\theta)z + \theta u_0$ lies in $C^* \cap \ker \Phi$. Since $R(u_0) < R(z)$ and R is convex,

$$(A.2) \quad R((1-\theta)z + \theta u_0) < R(z) \leq t^*,$$

and we obtain a contradiction with the fact that t^* is the minimal value of (\mathcal{P}) . As a result, $\text{codim}_T(T \cap \ker \Phi) \leq m - 1$. Observing that $T \cap \ker \Phi$ is the linear hull of

⁵In this example, to simplify the discussion, we assume that E has finite dimension.

$\mathcal{F}_{S^*}(0)$, and hence $j = \dim(T \cap \ker \Phi)$, we deduce that

$$(A.3) \quad \dim \mathcal{F}_{C^*}(0) \stackrel{\text{def.}}{=} \dim T = \text{codim}_T(T \cap \ker \Phi) + \dim(T \cap \ker \Phi) \leq m - 1 + j,$$

and the first claim of the theorem is proved.

Now, applying statement (3) of the Carathéodory–Klee theorem in [29, Page 425], p is a convex combination of at most $m + j$ (respectively, $m + j - 1$) extreme points (respectively, extreme points or in an extreme ray) of $\mathcal{F}_{C^*}(0)$. The conclusion stems from the fact that the extreme points (respectively, rays) of $\mathcal{F}_{C^*}(0)$ are extreme points (respectively, rays) of C^* (see the proof of the main theorem in [29]).

A.2. Proof of Corollary 3.7. Assume that the vector space

$$K \stackrel{\text{def.}}{=} \text{rec}(C^*) \cap (-\text{rec}(C^*))$$

is nontrivial (otherwise the conclusion follows from Theorem 3.1). We note that, for any $u \in C^*$, the convex function $v \mapsto R(u + v)$ is upper bounded by t^* on K , and hence is constant. As a result, possibly replacing R with $R + \chi_{C^*}$, it is not restrictive to assume that R is invariant by translation along K .

Now, let $\pi_K, \pi_{\Phi(K)}$ be the canonical quotient maps and define \tilde{R} and $\tilde{\Phi}$ by the commutative diagrams

$$\begin{array}{ccc} E & \xrightarrow{R} & [-\infty, \infty] \\ \pi_K \downarrow & \nearrow \tilde{R} & \\ E/K & & \end{array} \quad \begin{array}{ccc} E & \xrightarrow{\Phi} & \mathbb{R}^m \\ \pi_K \downarrow & & \downarrow \pi_{\Phi(K)} \\ E/K & \xrightarrow{\tilde{\Phi}} & \mathbb{R}^m / \Phi(K) \end{array}$$

Note that \tilde{R} is a convex function and that $\tilde{\Phi}$ is a linear map with rank $m - d$, where $d \stackrel{\text{def.}}{=} \dim(\Phi(K))$. It is then natural to consider the problem

$$(\tilde{\mathcal{P}}) \quad \min_{\tilde{u} \in E/K} \tilde{R}(\tilde{u}) \quad \text{s.t.} \quad \tilde{\Phi}\tilde{u} = \tilde{y},$$

where $\tilde{y} \stackrel{\text{def.}}{=} \pi_{\Phi(K)}(y)$. In other words, one still wishes to minimize $R(u)$, but one is satisfied if the constraint $\Phi u = y$ merely holds up to an additional term Φv , where $v \in K$. We observe that (\mathcal{P}) and $(\tilde{\mathcal{P}})$ have the same value t^* , and the level set

$$(A.4) \quad \tilde{C}^* \stackrel{\text{def.}}{=} \left\{ \tilde{u} \in E/K : \tilde{R}(\tilde{u}) \leq t^* \right\} = \pi_K(C^*)$$

is convex linearly closed and contains no line. Let $\tilde{\mathcal{S}}^*$ be the solution set to $(\tilde{\mathcal{P}})$. Theorem 3.1 now describes the elements of the j -dimensional faces of $\tilde{\mathcal{S}}^*$ as convex combinations of $m - d + j$ (respectively, $m - d + j - 1$) extreme points (respectively, extreme points or points in an extreme ray) of \tilde{C}^* , which we denote by $\tilde{q}_1, \tilde{q}_2, \dots, \tilde{q}_r$.

To conclude, we have obtained $\pi_K(p) = \sum_i \theta_i \tilde{q}_i$ for some $\theta \in \mathbb{R}_+^r$ with $\sum_i \theta_i = 1$. Equivalently, since $\psi_K^{-1}(\cdot, 0)$ provides one element in the corresponding class, this means that $p \in \psi_K^{-1}(\sum_i \theta_i \tilde{q}_i, 0) + K$. We get the claimed result by linearity of ψ_K^{-1} .

Remark A.1. Incidentally, we note that, for $Y = \{y\}$, the face $\mathcal{F}_{S_{\{y\}}^*}(p)$ is isomorphic (through ψ_K) to $\mathcal{F}_{\pi_K(S_{\{y\}}^*)}(\pi_K(p)) \times (K \cap \ker \Phi)$.

Acknowledgments. This work was initially started by two different groups, one composed of A. Chambolle, F. de Gournay and P. Weiss, and the other of C. Boyer, Y. De Castro and V. Duval. The authors realized that they were working on similar topics when a few of them met at the mathematical imaging semester “Variational methods and effective algorithms for imaging and vision” held at the Isaac Newton Institute, Cambridge, UK. They therefore decided to write a joint paper and wish to thank the organizers for giving them this opportunity. This work was started with the help of T. Pock through a few numerical experiments and with discussions with J. Fadili and C. Poon.

REFERENCES

- [1] B. ADCOCK AND A. C. HANSEN, *Generalized sampling and infinite-dimensional compressed sensing*, *Found. Comput. Math.*, 16 (2016), pp. 1263–1323.
- [2] A. AHMED, B. RECHT, AND J. ROMBERG, *Blind deconvolution using convex programming*, *IEEE Trans. Inf. Theory*, 60 (2014), pp. 1711–1732.
- [3] C. ALIPRANTIS AND K. BORDER, *Infinite Dimensional Analysis: A Hitchhiker’s Guide*, Springer, Berlin, 2006, <http://doi.org/10.1007/3-540-29587-9>.
- [4] L. AMBROSIO, V. CASELLES, S. MASNOU, AND J.-M. MOREL, *Connected components of sets of finite perimeter and applications to image processing*, *J. Eur. Math. Soc. (JEMS)*, 3 (2001), pp. 39–92.
- [5] A. I. BARVINOK, *Problems of distance geometry and convex properties of quadratic maps*, *Discrete Comput. Geom.*, 13 (1995), pp. 189–202, <https://doi.org/10.1007/BF02574037>.
- [6] A. BEURLING, *Sur les intégrales de Fourier absolument convergentes et leur application à une transformation fonctionnelle*, in *Neuvième Congrès des Mathématiciens Scandinaves*, Mercator Tryckeri, Helsingfors, 1939, pp. 345–366.
- [7] F. J. BEUTLER, *The operator theory of the pseudo-inverse I. Bounded operators*, *J. Math. Anal. Appl.* 10 (1965), pp. 451–470.
- [8] N. BOURBAKI, *Espaces Vectoriels Topologiques: Chapitres 1 à 5*, Springer, Berlin, 2007, <http://www.springer.com/gp/book/9783540344971>.
- [9] K. BREDIES AND M. CARIONI, *Sparsity of solutions for Variational Inverse Problems with Finite-Dimensional Data*, preprint, <http://arxiv.org/abs/1809.05045>, 2018.
- [10] E. J. CANDÈS AND C. FERNANDEZ-GRANDA, *Towards a mathematical theory of super-resolution*, *Comm. Pure Appl. Math.*, 67 (2014), pp. 906–956.
- [11] E. J. CANDÈS AND B. RECHT, *Exact matrix completion via convex optimization*, *Found. Comput. Math.*, 9 (2009), p. 717.
- [12] A. CHAMBOLLE AND T. POCK, *A first-order primal-dual algorithm for convex problems with applications to imaging*, *J. Math. Imaging Vision*, 40 (2011), pp. 120–145, <https://doi.org/10.1007/s10851-010-0251-1>.
- [13] V. CHANDRASEKARAN, B. RECHT, P. A. PARRILO, AND A. S. WILLSKY, *The convex geometry of linear inverse problems*, *Found. Comput. Math.*, 12 (2012), pp. 805–849.
- [14] A. DANILIDIS AND Y. G. RAMOS, *Some remarks on the class of continuous (semi-)strictly quasiconvex functions*, *J. Optim. Theory Appl.*, 133 (2007), pp. 37–48.
- [15] J. DATTORRO, *Convex Optimization and Euclidean Distance Geometry*, Meboo Publishing, Palo Alto, CA, 2005; available at <https://meboo.convexoptimization.com/Meboo.html>.
- [16] Y. DE CASTRO, F. GAMBOA, D. HENRION, AND J.-B. LASSERRE, *Exact solutions to super resolution on semi-algebraic domains in higher dimensions*, *IEEE Trans. Inf. Theory*, 63 (2017), pp. 621–630.
- [17] D. L. DONOHO, *Compressed sensing*, *IEEE Trans. Inf. Theory*, 52 (2006), pp. 1289–1306.
- [18] D. L. DONOHO AND J. TANNER, *Sparse nonnegative solution of underdetermined linear equations by linear programming*, *Proc. Natl. Acad. Sci. USA*, 102 (2005), pp. 9446–9451.
- [19] D. DRUSVYATSKIY, S. A. VAVASIS, AND H. WOLKOWICZ, *Extreme point inequalities and geometry of the rank sparsity ball*, *Math. Program.*, 152 (2015), pp. 521–544.
- [20] L. E. DUBINS, *On extreme points of convex sets*, *J. Math. Anal. Appl.*, 5 (1962), pp. 237–244.
- [21] V. DUVAL AND G. PEYRÉ, *Exact support recovery for sparse spikes deconvolution*, *Found. Comput. Math.*, 15 (2015), pp. 1315–1355.
- [22] C. FERNANDEZ-GRANDA, *Super-resolution of point sources via convex programming*, *Inf. Inference*, 5 (2016), pp. 251–303.

- [23] S. D. FISHER AND J. W. JEROME, *Spline solutions to L_1 extremal problems in one and several variables*, J. Approx. Theory, 13 (1975), pp. 73–83.
- [24] W. FLEMING, *Functions with generalized gradient and generalized surfaces*, Ann. Mat. Pura Appl., 44 (1957), pp. 93–103.
- [25] A. FLINTH AND P. WEISS, *Exact Solutions of Infinite Dimensional Total-Variation Regularized Problems*, preprint, <https://arxiv.org/abs/1708.02157>, 2017.
- [26] F. FOGEL, R. JENATTON, F. BACH, AND A. D’ASPROMONT, *Convex relaxations for permutation problems*, in Adv. Neural Inf. Process. Syst. 36, Curran Associates, Red Hook, NY, 2015, pp. 1016–1024.
- [27] H. GUPTA, J. FAGEOT, AND M. UNSER, *Continuous-Domain Solutions of Linear Inverse Problems with Tikhonov vs. Generalized TV Regularization*, preprint, <https://arxiv.org/abs/1802.01344>, 2018.
- [28] J.-B. HIRIART-URRUTY AND C. LEMARÉCHAL, *Convex Analysis and Minimization Algorithms I*, Grundlehren Math. Wiss. 305, Springer, Berlin, 1993, <https://doi.org/10.1007/978-3-662-02796-7>.
- [29] V. KLEE, *On a theorem of Dubins*, J. Math. Anal. Appl., 7 (1963), pp. 425–427, [https://doi.org/10.1016/0022-247X\(63\)90063-5](https://doi.org/10.1016/0022-247X(63)90063-5).
- [30] V. L. KLEE, *Extremal structure of convex sets*, Arch. Math. (Basel), 8 (1957), pp. 234–240, <https://doi.org/10.1007/BF01899998>.
- [31] M. KREĬN, *The L -problem in an abstract normed linear space*, in Some Questions in the Theory of Moments, I. Ahiezer and M. Kreĭn, eds., Transl. Math. Monogr. 2, American Mathematical Society, Providence, RI, 1962.
- [32] M. G. KREIN AND A. A. NUDELMAN, *The Markov Moment Problem and Extremal Problems: Ideas and Problems of P. L. Cebysev and A. A. Markov and Their Further Development*, Transl. Math. Monogr. 50, American Mathematical Society, Providence, RI, 1977.
- [33] C. L. LAWSON AND R. J. HANSON, *Solving least squares problems*, Classics Appl. Math. 15, SIAM, Philadelphia, PA, 1995.
- [34] J. MATOUŠEK AND B. GÄRTNER, *Understanding and Using Linear Programming*, Universitext, Springer, Berlin, 2007, <http://doi.org/10.1007/978-3-540-30717-4>.
- [35] P. McMULLEN, *The maximum numbers of faces of a convex polytope*, Mathematika, 17 (1970), pp. 179–184.
- [36] T. S. MOTZKIN, *Comonotone curves and polyhedra*, Bull. Amer. Math. Soc., 63 (1957).
- [37] M. NIKOLOVA, *Local strong homogeneity of a regularized estimator*, SIAM J. Appl. Math., 61 (2000), pp. 633–658.
- [38] G. PATAKI, *The geometry of semidefinite programming*, in Handbook of Semidefinite Programming: Theory, Algorithms, and Applications, Springer, Boston, MA, 2000, pp. 29–65.
- [39] E. RICHARD, F. BACH, AND J. P. VERT, *Intersecting singularities for multi-structured estimation*, in Proceedings of the 30th International Conference on Machine Learning, Proc. Mach. Learn. Res. 28, 2013, pp. 1157–1165, <http://proceedings.mlr.press/v28/>.
- [40] E. RICHARD, G. R. OBOZINSKI, AND J.-P. VERT, *Tight convex relaxations for sparse matrix factorization*, in Adv. Neural Inf. Process. Syst. 27, Curran Associates, Red Hook, NY, 2014, pp. 3284–3292, <http://papers.nips.cc/paper/5408-tight-convex-relaxations-for-sparse-matrix-factorization.pdf>.
- [41] L. I. RUDIN, S. OSHER, AND E. FATEMI, *Nonlinear total variation based noise removal algorithms*, Phys. D, 60 (1992), pp. 259–268.
- [42] B. SCHÖLKOPF, R. HERBRICH, AND A. J. SMOLA, *A generalized representer theorem*, in Computational Learning Theory, COLT 2001, Lecture Notes in Comput. Sci. 2111, Springer, 2001, pp. 416–426.
- [43] B. SCHÖLKOPF AND A. J. SMOLA, *Learning with kernels: Support vector machines, regularization, optimization, and beyond*, Adapt. Comput. Mach. Learn., MIT Press, Cambridge, MA, 2002.
- [44] A. SHAPIRO, *On duality theory of conic linear problems*, in Semi-Infinite Programming, Non-convex Optim. Appl. 57, Springer, Boston, MA 2001, pp. 135–165.
- [45] J. A. SHOHAT AND J. D. TAMARKIN, *The Problem of Moments*, Math. Surveys Monogr. 1, American Mathematical Society, Providence, RI, 1943.
- [46] G. TANG, B. N. BHASKAR, P. SHAH, AND B. RECHT, *Compressed sensing off the grid*, IEEE Trans. Inf. Theory, 59 (2013), pp. 7465–7490.
- [47] M. UNSER, J. FAGEOT, AND H. GUPTA, *Representer theorems for sparsity-promoting ℓ_1 regularization*, IEEE Trans. Inf. Theory, 62 (2016), pp. 5167–5180.
- [48] M. UNSER, J. FAGEOT, AND J. P. WARD, *Splines are universal solutions of linear inverse problems with generalized TV regularization*, SIAM Rev., 59 (2017), pp. 769–793.

- [49] H. WENDLAND, *Scattered Data Approximation*, Cambridge Monogr. Appl. Comput. Math., Cambridge University Press, 2005.
- [50] S. ZUHOVICKÍĽ, *Remarks on problems in approximation theory*, Mat. Zbirnik KDU (1948), pp. 169–183 (in Ukrainian).
- [51] S. ZUHOVICKÍĽ, *On approximation of Real Functions in the Sense of P. L. Čebyšev*, Amer. Math. Soc. Transl. Ser. 2 19, American Mathematical Society, Providence, RI, 1962, pp. 221–252
<http://dx.doi.org/10.1090/trans2/019>.