

LINEAR INVERSE PROBLEMS WITH NONNEGATIVITY CONSTRAINTS: SINGULARITY OF OPTIMISERS

CAMILLE POUCHOL AND OLIVIER VERDIER

ABSTRACT. We look at continuum solutions in optimisation problems associated to linear inverse problems $y = Ax$ with non-negativity constraint $x \geq 0$. We focus on the case where the noise model leads to maximum likelihood estimation through general divergences, which cover a wide range of common noise statistics such as Gaussian and Poisson. Considering x as a Radon measure over the domain on which the reconstruction is taking place, we show a general singularity result. In the high noise regime corresponding to $y \notin \{Ax \mid x \geq 0\}$ and under a key assumption on the divergence as well as on the operator A , any optimiser has a singular part with respect to the Lebesgue measure. We hence provide an explanation as to why any possible algorithm successfully solving the optimisation problem will lead to undesirably spiky-looking images when the image resolution gets finer, a phenomenon well documented in the literature. We illustrate these results with several numerical examples inspired by medical imaging.

1. INTRODUCTION

We consider linear inverse problems of the type $Ax = y$, where A is a linear operator, and y lies in a suitable linear space. These problems are often endowed with constraints stemming from the model at hand. One of the most common such constraints is nonnegativity of the unknown, i.e., $x \geq 0$. This happens in various applications, in particular in medical imaging, where x is the *activity*, which is nonnegative for physical reasons. One important example is that of Positron Emission Tomography (PET) [26]. In this setting, the unknown x lies in \mathbb{R}^r , the data y lies in \mathbb{R}^m , and the operator A lies in $\mathbb{R}^{m \times r}$, where m is the number of data points, r is the number of voxels. Deconvolution problems often also incorporate such constraints [15].

Depending on the noise model, the corresponding (negative) log-likelihood problem is typically of the form

$$\min_{x \geq 0} D(y, Ax),$$

where D is some divergence functional. If the noise model is Gaussian, for instance, then D is simply the Euclidean distance, whereas if the noise model is Poisson, D is the Kullback–Leibler divergence.

We analyse the effect of an ever increasing resolution, which leads us to regard the unknown x as a function (henceforth denoted μ) in some functional space X of functions over a compact set $K \subset \mathbb{R}^p$, and the operator A is now a linear mapping from the space X to \mathbb{R}^m . This leads to the optimisation problem

$$(1) \quad \min_{\mu \geq 0} D(y, A\mu).$$

The non-negativity constraint has been shown to cause sparsity in various contexts in optimisation and optimal control [8]. As a result, the proper functional space X

to be considered appears to be that of Radon measures, where discrete measures are regarded as sparse, with the operator A defined for all $\mu \in X, \mu \geq 0$ by

$$(A\mu)_i := \int a_i d\mu, \quad a_i \geq 0.$$

Here, the assumption $a_i \geq 0$ reflects our interest in applications where the measured data is nonnegative.

This sparsity phenomenon has been observed in optimal control [9, 18], as well as for the optimisation problem (1) when the divergence D is specifically the Kullback–Leibler divergence [22]. This divergence and the framework of Radon measures can even be natural from the physical point of view, as is the case in PET where the underlying model is based on Poisson point processes [26, 23].

In the examples above, sparsity (which can arise in the form of Dirac masses) is undesirable as the sought-for image is expected to be at least piecewise smooth. In other contexts, sparsity of the signal must be enforced, as is the case for instance in sparse super-resolution where Dirac masses are the desired outcome [12, 11].

In the latter case, the situation is completely different as the unknown signal is *known* to be a sum of (nonnegative) Dirac masses and the aim is to make sure that the optimisation problem has *at least one minimiser* which is a sum of Dirac masses, thereby hoping to recover the number and support of the unknown discrete measure. A very general treatment of the existence of such sparse optimisers has recently been given in [2, 3].

The goal of the present work is to establish general conditions for the operator A and the divergence D to fulfill, so that, when the measured data y is not in the cone $\{A\mu \mid \mu \geq 0\}$, *all minimisers* are singular in the sense defined below. We use Lebesgue’s decomposition theorem with respect to the Lebesgue measure: any Radon measure μ over K decomposes into the sum of an absolutely continuous measure (its *absolutely continuous part*) and a singular measure (its *singular part*).

Definition 1.1. A measure μ will be said to be *singular* if its singular part is non-zero, and *completely singular* if it is singular and its absolutely continuous part is zero.

The paper [22] established that optimal measures are completely singular in the specific case where

- (i) the functions a_i are linearly independent and (real) analytic,
- (ii) D is the Kullback–Leibler divergence.

The first assumption (i) is restrictive since many practical problems are such that all the functions a_i are compactly supported inside the interior of K . This makes analyticity irreconcilable with the condition $a_i \geq 0$. The second assumption (ii) also restricts the possible noise models to the single, albeit important, case of Poisson measurements.

In fact, the so-called β -divergences have attracted interest recently in non-negative matrix factorisation [13], and are now also advocated for in some medical imaging contexts, such as in PET [5]. This is a family of divergences depending on a real parameter β . It has the appealing property of interpolating between the Kullback–Leibler divergence ($\beta = 1$) and the Euclidean distance ($\beta = 2$) [7], which correspond to different noise models (see §2.4).

In this article, we first generalise the above result (when (i) holds) to very general divergences.

Then, we proceed to treat the case of more general detector functions a_i relaxing assumption (i). This forces us to restrict the class of divergences. Informally at this stage, we hence consider

- (C) detector functions a_i which are *locally* linearly independent and *piecewise* analytic,
- (H) divergences which are sufficiently averse to zero values whenever the data is non-zero (see the precise statement in equation (H) in §2.2).

Most functions we are aware of are expected to satisfy (C), such as those defining PET (although local linear independence might be difficult to prove for each particular case). Condition (H) limits the class of divergences, as it for instance excludes the Euclidean distance. It still covers a wide class of β -divergences, f-divergences and Bregman divergences, in particular those which arise in contexts featuring nonnegativity constraints.

Our main result may then be stated as follows (see Theorem 3.6 for a precise statement).

Theorem 1.2. *If $y \notin \{A\mu, \mu \geq 0\}$ with A satisfying (C) and D satisfying (H), then all optimal solutions μ^* to (1) are singular.*

In other words, if the data is not in the image of the cone of nonnegative measures $\{\mu \geq 0\}$ under the operator A , any optimiser μ^* is singular, even if the data has been generated from an image that is absolutely continuous with respect to the Lebesgue measure.

The condition that the data y does not belong to the image cone $\{A\mu \mid \mu \geq 0\}$ can be interpreted as a condition on the level of noise: the more noise there is, the more likely it is that this condition be fulfilled (see [22] and §2.4). Undesirable singular-looking images will hence arise in the high-noise regime.

Our results show that the optimisation problem itself leads to singular results. Thus, any algorithm successfully solving (1) will inevitably lead to spiky-looking images as one keeps iterating. In the context of medical imaging, this has been observed when using the Maximum-Likelihood-Expectation-Maximisation (ML-EM, also called the Richardson–Lucy algorithm) for solving (1), and has been referred to in the literature as the “night-sky” or the “draughtsboard” effect [28].

Consequently, the same kind of artefacts will be observed for other likelihoods, hinting at the necessity of either early stopping when solving (1) (see [24], [20]), or adding appropriate regularisation terms of the form

$$(2) \quad \min_{\mu \geq 0} D(y, A\mu) + \lambda R(\mu),$$

with regularisation parameter λ that is sufficiently large to alleviate the issue [5].

In this work, we also provide numerical examples coming from PET where, when solving (1) for some commonly used divergences, and applying a sufficient amount of noise, reconstructions exhibit the night-sky effect (§5). As our theoretical results suggest, this effect should be more and more prominent as one keeps iterating a convergent algorithm for solving (1), or a convergent algorithm for solving (2) with a sufficiently small regularisation parameter λ .

Outline of the paper. The paper is organised as follows. In §2, we define the inverse problem by setting the functional analytic framework as well as the family of divergences being considered, leading to the corresponding maximum likelihood problem. In §3, we proceed to study the resulting optimisation problem and prove our main result Theorem 1.2. Section §4 is devoted to discussing the assumption (C) on the detector functions in more detail, in the case of a toy example and for the 2-dimensional PET operator on a regular polygon. We finally illustrate our results about singularity by numerical simulations in §5. They feature different examples with and without regularisation.

2. INVERSE PROBLEM SETUP

2.1. Linear inverse problem. We aim at reconstructing an image μ defined on a non-empty compact $K \subset \mathbb{R}^p$, $p \geq 1$, where K is the closure of a bounded Lipschitz connected open set.

2.1.1. Measure-theoretic background. The unknown image μ is an element of the space of Radon measures, denoted $\mathcal{M}(K)$, which is the topological dual space of continuous functions over the compact, denoted $\mathcal{C}(K)$. We endow $\mathcal{M}(K)$ with the weak-* topology, making $\mathcal{C}(K)$ its dual space. We also recall that the bounded sets of $\mathcal{M}(K)$ are relatively compact in the weak-* topology, by virtue of the Banach–Alaoglu Theorem [4].

We denote the dual pairing between a function $\mu \in \mathcal{M}(K)$ and a function $f \in \mathcal{C}(K)$ by $\langle \mu, f \rangle$, and $\mathcal{M}_+(K)$ stands for the set of nonnegative Radon measures. Whenever the context is clear, we shall drop the reference to the compact K and write \mathcal{M} , \mathcal{M}_+ and \mathcal{C} .

Finally, we recall that by Lebesgue’s decomposition theorem, any measure $\mu \in \mathcal{M}_+(K)$ can uniquely be written $\mu = \mu_1 + \mu_2$, where $\mu_1 \in \mathcal{M}_+(K)$ is absolutely continuous, and $\mu_2 \in \mathcal{M}_+(K)$ is singular, where absolute continuity and singularity are meant with respect to the Lebesgue measure. The measure μ is then said to have a *singular part* if $\mu_2 \neq 0$.

2.1.2. Operator. A data point y is a vector of m scalar nonnegative measurements, that is,

$$y \in \mathbb{R}_+^m.$$

This vector itself typically is the realisation of some random variable with mean $A\mu$, where $\mu \in \mathcal{M}_+(K)$ is the image to be reconstructed, and A is a linear mapping $A: \mathcal{M}(K) \rightarrow \mathbb{R}^m$.

The only assumption we make on A is that it is continuous in the weak-* topology. From [4, Proposition 3.14], this implies that A is of the form

$$(A\mu)_i = \langle \mu, a_i \rangle, \quad i = 1, \dots, m,$$

where the *detector functions* a_i are elements of $\mathcal{C}(K)$. This covers the case of PET [19, 22] and more generally the setting of kernel operators: if the underlying operator in infinite dimension is of the form

$$\mu \mapsto \int_K k(\cdot, y) d\mu(y),$$

for some smooth kernel $k \in \mathcal{C}(K \times K)$, the operator A typically is obtained from a sampling for m points $x_i \in K$ or integrating the kernel over some subdomains $\Omega_i \subset K$, namely

$$a_i = k(x_i, \cdot), \quad \text{or} \quad a_i = \int_{\Omega_i} k(x, \cdot) dx, \quad i = 1, \dots, m.$$

We will make the assumption that A maps $\mathcal{M}_+(K)$ into the set of (componentwise) nonnegative vectors denoted \mathbb{R}_+^m which of course is equivalent to

$$a_i \geq 0, \quad i = 1, \dots, m.$$

Finally, we assume $a_i \neq 0$ for all $i = 1, \dots, m$, as well as

$$(3) \quad \sum_{i=1}^m a_i > 0.$$

Indeed, at any point x such that $\sum_{i=1}^m a_i(x) = 0$, we would have no information on the unknown μ .

Another consequence of the simple continuity assumption on A is that its adjoint $A^* : \mathbb{R}^m \rightarrow \mathcal{C}(K)$ is simply defined as

$$A^* \lambda = \sum_{i=1}^m \lambda_i a_i, \quad \lambda \in \mathbb{R}^m.$$

We shall sometimes need to know when A^* is injective. This is of course equivalent to the linear independence of the family $(a_i)_{i=1, \dots, m}$. Note that since the codomain of A is finite dimensional, we have

$$A^* \text{ is injective} \iff A \text{ is surjective.}$$

In order to solve the inverse problem with the nonnegativity constraint, we aim at solving the optimisation problem

$$(4) \quad \min_{\mu \in \mathcal{M}_+(K)} D(y, A\mu).$$

Here, D stands for some divergence over \mathbb{R}_+^m , and we will use the notation

$$(5) \quad \ell(\mu) := D(y, A\mu).$$

Finally, let us define the notion of support for the various relevant cases (all these cases can be covered in one single definition, but we prefer separating them for clarity):

- for a vector $w \in \mathbb{R}_+^m$,

$$\text{supp}(w) = \{i \in \{1, \dots, m\} \mid w_i > 0\},$$

- for a nonnegative function $f \in \mathcal{C}(K)$,

$$\text{supp}(f) = \overline{\{x \in K \mid f(x) > 0\}},$$

where \overline{F} stands for the closure of a set $F \subset K$,

- for a nonnegative measure $\mu \in \mathcal{M}_+(K)$,

$$\text{supp}(\mu) := \{x \in K \mid \mu(N) > 0, \forall N \in \mathcal{N}(x)\},$$

where $\mathcal{N}(x)$ is the set of all open neighbourhoods of x .

Notice that with these notations in place, we may now rewrite assumption (3) as

$$\bigcup_{i=1}^m \text{supp}(a_i) = K.$$

2.2. Divergences. We shall assume that D is a *separable* divergence, i.e.,

$$D(u, v) = \sum_{i=1}^m d(u_i, v_i) \quad u, v \in \mathbb{R}_+^m.$$

where $d: \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}_+ \cup \{+\infty\}$ is a (scalar) divergence. When needed, we will implicitly extend the function d to the whole \mathbb{R}^2 by $+\infty$ for $u < 0$ or $v < 0$. Throughout, we will assume that d satisfies the following basic properties

Separation:

$$d(u, v) = 0 \iff u = v.$$

Convexity: for $u \in \mathbb{R}_+$, $v \mapsto d(u, v)$ is convex on \mathbb{R}_+ .

Regularity: for $u \in \mathbb{R}_+$, the mapping $v \mapsto d(u, v)$ is lower-semicontinuous on \mathbb{R}_+ ,

Coercivity: for $u \in \mathbb{R}_+$, the mapping $v \mapsto d(u, v)$ is coercive, i.e.,

$$\forall u \in \mathbb{R}_+, \quad \lim_{v \rightarrow +\infty} d(u, v) = +\infty.$$




FIGURE 1. An example of map $v \mapsto d(1, v)$ which fulfills the assumptions above. In particular, the derivative tends to $-\infty$ when v approaches zero, and the divergence is zero only at the point $v = 1$. Note that this is a plot of the function $v \mapsto d_\beta(1, v)$ (see § 2.3.1) for $\beta = 1.2$.

The above assumptions ensure that, for a fixed $u \in \mathbb{R}_+$, $v \mapsto d(u, v)$ is subdifferentiable at every $v > 0$. For our results about singularity of optimisers, we will also sometimes consider a subclass of such divergences for which we must specify (non-)subdifferentiability at $v = 0$ as follows:

(H) $v \mapsto d(u, v)$ is not subdifferentiable at 0 if $u > 0$.

From standard convex analysis, this is equivalent to the directional derivative of $d(u, \cdot)$ being $-\infty$ at 0 for all $u > 0$, i.e.,

$$(H) \iff \forall u > 0, \lim_{t \rightarrow 0, t > 0} \frac{d(u, t) - d(u, 0)}{t} = -\infty.$$

Finally, we will need one very weak assumption involving both the operator A and the divergence D :

(6) there exists $\mu \in \mathcal{M}_+$ such that $D(y, \cdot)$ is finite and continuous at $A\mu$.

One sufficient and simpler assumption for the above condition to hold is the existence of some point $x \in K$ such that the divergences $d(y_i, \cdot)$ are continuous at $a_i(x)$ for all $i \in \{1, \dots, m\}$, which follows from choosing $\mu = \delta_x$.

2.3. Examples.

2.3.1. *Beta Divergences.* For $u > 0$, $v > 0$ scalar variables, $\beta \in (1, 2]$, we define

$$d_\beta(u, v) := \frac{1}{\beta(\beta - 1)} (u^\beta + (\beta - 1)v^\beta - \beta uv^{\beta-1}),$$

which for $\beta = 2$ gives the Euclidean distance

$$d_2(u, v) = \frac{1}{2}(u - v)^2,$$

and by continuity for $\beta = 1$ the Kullback–Leibler divergence

$$d_1(u, v) = u \ln \left(\frac{u}{v} \right) - u + v.$$

The corresponding divergence over \mathbb{R}_+^m will be denoted D_β .

More precisely, using the convention $0/0 = 0$, $0 \ln 0 = 0$, d_1 is defined for nonnegative scalars $u \geq 0$, $v \geq 0$ as follows:

$$d_1(u, v) = \begin{cases} v & \text{if } u = 0, v \geq 0, \\ +\infty & \text{if } u > 0, v = 0, \\ u \ln\left(\frac{u}{v}\right) - u + v & \text{if } u > 0, v > 0. \end{cases}$$

Note that the β -divergences d_β satisfy Hypothesis (H) for all $\beta \in [1, 2)$, but d_2 does not.

2.3.2. Reverse f - and Bregman divergences. Given a convex function $F: \mathbb{R}_+ \rightarrow \mathbb{R}_+ \cup \{+\infty\}$ with $F(1) = 0$, one can define the following “reverse f -divergence”

$$d_F^f(u, v) := F(v/u)u, \quad u > 0 \quad v \geq 0,$$

Similarly, with a convex function $F: \mathbb{R} \rightarrow \mathbb{R} \cup \{+\infty\}$ such that F is differentiable, we define the “reverse Bregman divergence”

$$d_F^B(u, v) := F(v) - F(u) - F'(u)(v - u), \quad u > 0 \quad v \geq 0.$$

Both fulfill our assumptions as long as

- F is lower semicontinuous,
- $F'(0) = -\infty$,
- they are suitably extended to the case $u = 0$ in such a way that $d(0, \cdot)$ is a nonnegative convex, lower semicontinuous and coercive function.

Note that β -divergences are not reverse f -divergences nor reverse Bregman divergences for $\beta \in (1, 2)$, the Kullback–Leibler divergence d_1 is a reverse f -divergence but not a reverse Bregman divergence, and the Euclidean distance d_2 is a reverse Bregman divergence but not a reverse f -divergence.

2.4. The noise model. In general, y is drawn according to a distribution parameterised by $w = A\mu$ as well as some additional dispersion parameter $\phi > 0$ controlling the noise level. For a better understanding of why the condition $y \notin A(\mathcal{M}_+)$ typically arises in the high noise case, we review the underlying statistical model in the specific case of β -divergences.

In this setting, the minimisation problem (4) is (up to constants) the corresponding (negative-log) likelihood maximum problem. We write the statistical model for y and w as scalar variables, as the full statistical model is straightforwardly obtained by assuming independent components.

A general way to write the noise model giving rise to β -divergences is to use the so-called Tweedie distributions [27]. The β -divergences are a special case of such distributions, as the corresponding Tweedie distribution is given for w fixed by

$$y \mapsto H_\beta(y, \phi) \exp\left(-\frac{1}{\phi} d_\beta(y, w)\right),$$

where H_β is a normalisation factor for the above function to integrate to 1. Observing the data y to estimate w , minimising the negative log-likelihood problem indeed is equivalent to minimising $w \mapsto d_\beta(y, w)$. We refer to [27] for more details.

The underlying density is not always tractable (this is the case if $1 < \beta < 2$), making the noise model unclear. In other cases, the noise model can be further identified as follows.

Case $\beta = 2$: the noise model is Gaussian, i.e.,

$$y \sim \mathcal{N}(w, \phi).$$

Case $1 < \beta < 2$: to the best of our knowledge, no explicit model is known.

Case $\beta = 1$: the noise model is Poisson, i.e.,

$$y \sim \phi \mathcal{P}\left(\frac{1}{\phi} w\right).$$

Furthermore, y has mean w and variance $\phi^{2-\beta}w$: y concentrates around w as the noise level ϕ vanishes. Hence, the less noise there is, the more likely it is that $y \in A(\mathcal{M}_+)$. For a more quantitative version of this statement in the case of $d = d_1$, i.e., Poisson distributed measurements, see [22].

Remark 2.1. Note that the Gaussian case $\beta = 2$ is the only one which does not (necessarily) lead to nonnegative data $y \in \mathbb{R}_+^m$. Yet, we will always make this assumption throughout. Indeed, the Gaussian noise model is frequently used as an approximate basic model, even for inverse problems where the data satisfies $y \geq 0$ (for physical reasons). The problem is then typically solved using nonnegative least-squares, hence it is worth studying this case as well.

3. OPTIMISATION PROBLEM

We now investigate the optimisation problem (4), starting with a related optimisation problem and its dual.

3.1. Related optimisation problem and its dual. We consider the cone

$$A(\mathcal{M}_+) = \{ A\mu \mid \mu \in \mathcal{M}_+ \}.$$

We notice that $A(\mathcal{M}_+)$ is a closed convex set:

Lemma 3.1. *The cone $A(\mathcal{M}_+)$ is closed.*

Proof. Pick a sequence $w_n = A\mu_n$ in $A(\mathcal{M}_+)$ converging to some $w \in \mathbb{R}_+^m$. Appealing to (3), we choose $c > 0$ such that $\sum_{i=1}^m a_i \geq c$ and write

$$\mu_n(K) \leq \frac{1}{c} \int_K \left(\sum_{i=1}^m a_i \right) d\mu_n = \frac{1}{c} \sum_{i=1}^m (w_n)_i,$$

where the last quantity is bounded since (w_n) converges. Hence, the sequence (μ_n) is bounded, it has a weak-* converging subsequence, say to some $\mu \in \mathcal{M}_+$. Along the subsequence, $(A\mu_n)$ converges to $A\mu$ by weak-* continuity of A , which proves that $w = A\mu \in A(\mathcal{M}_+)$. \square

The original optimisation problem (4) is related to the following optimisation problem:

$$(7) \quad \min_{w \in A(\mathcal{M}_+)} D(y, w).$$

More precisely, for any w^* optimal for the above problem, any measure $\mu^* \in \mathcal{M}_+$ such that $A\mu^* = w^*$ is optimal for the original problem.

Lemma 3.2. *The minimum in (4) is attained.*

Proof. We show that the optimal value in the minimisation problem (7) is attained. The function $w \mapsto D(y, w)$ is coercive and lower semicontinuous. Since the (non-empty) cone $A(\mathcal{M}_+)$ is closed (Lemma 3.1), there exists an optimal $w^* \in A(\mathcal{M}_+)$ for the problem (7). Any $\mu^* \in \mathcal{M}_+$ such that $A\mu^* = w^*$ then provides a minimiser for the original problem (4). \square

We now compute the (Lagrange) dual problem to the problem (7). These computations turn out to be crucial in analysing results from numerical simulations, determining whether we should expect singularity or not for some given $y \in \mathbb{R}_+^m$ from appropriately defined *singularity certificates*.

We define the cone $A(\mathcal{M}_+)^* := \{ \lambda \in \mathbb{R}^m \mid \langle \lambda, w \rangle \geq 0, \forall w \in A(\mathcal{M}_+) \}$ dual to $A(\mathcal{M}_+)$, which can be characterised as in [14] by

$$A(\mathcal{M}_+)^* = \{ \lambda \in \mathbb{R}^m \mid A^* \lambda \geq 0 \text{ on } K \}.$$

The dual problem writes

$$\max_{\lambda \in A(\mathcal{M}_+)^*} g(\lambda),$$

where the function $g: \mathbb{R}^m \rightarrow \mathbb{R}$ is defined for $\lambda \in A(\mathcal{M}_+)^*$ by

$$(8) \quad g(\lambda) := \min_{w \in \mathbb{R}_+^m} D(y, w) - \langle \lambda, w \rangle.$$

Note that strong duality holds since the problem is convex and Slater's condition is obviously satisfied since the minimisation occurs over a cone in finite dimension.

The point of taking the Lagrange dual is that, since D decomposes, so does g and we find

$$g(\lambda) = \sum_{i=1}^m \min_{w_i \geq 0} (d(y_i, w_i) - \lambda_i w_i).$$

As a result, defining

$$(9) \quad h(y, \lambda) := \min_{w \geq 0} (d(y, w) - \lambda w)$$

for a scalar $y \geq 0$, the resulting dual function will take the form

$$g(\lambda) = \sum_{i=1}^m h(y_i, \lambda_i).$$

The explicit computation of the function h in (9) is carried out for the case of β -divergences in Appendix A.

3.2. Optimality conditions. We now use Fenchel duality to compute the optimality conditions for problem (4). The assumption that the operator A is continuous in the weak-* topology plays a crucial role.

Proposition 3.3. *Let μ^* be an optimal measure for (4). Then, there exists $\lambda^* \in \partial D(y, \cdot)(A\mu^*)$ such that*

$$(10) \quad A^* \lambda^* \geq 0 \text{ on } K, \quad A^* \lambda^* = 0 \text{ on } \text{supp}(\mu^*).$$

Proof. We use the Fenchel–Rockafellar Theorem [25], with the optimisation problem rewritten as

$$\min_{\mu \in \mathcal{M}} D(y, A\mu) + \delta_{\mathcal{M}_+}(\mu),$$

where $\delta_{\mathcal{M}_+}$ is the indicator function in the convex analytic sense, i.e. $\delta_{\mathcal{M}_+}(\mu) = 0$ if $\mu \in \mathcal{M}_+$ and $+\infty$ otherwise.

In the context of the Fenchel–Rockafellar Theorem we need *paired spaces*. The natural choice is \mathcal{C} (endowed with its strong topology) and \mathcal{M} (endowed with its weak-* topology). We may apply the theorem thanks to the hypothesis (6) and we obtain

$$\mu^* \text{ is optimal} \iff 0 \in A^* \partial D(y, \cdot)(A\mu^*) + N_{\mathcal{M}_+}(\mu^*),$$

where $N_{\mathcal{M}_+}(\mu)$ is the normal cone of \mathcal{M}_+ at μ , defined by

$$N_{\mathcal{M}_+}(\mu) := \{ f \in \mathcal{C} \mid \forall \nu \in \mathcal{M}_+, \langle \mu - \nu, f \rangle \geq 0 \}.$$

The normal cone can be identified as

$$N_{\mathcal{M}_+}(\mu) = \{ f \in \mathcal{C} \mid f \leq 0 \text{ on } K, f = 0 \text{ on } \text{supp}(\mu) \},$$

see [22]. Picking $\lambda^* \in \partial D(y, \cdot)(A\mu^*)$ such that $A^* \lambda^* \in -N_{\mathcal{M}_+}(\mu^*)$, we exactly obtain (10). \square

Note that for μ^* optimal and with the notations of the above result, the separable form of D ensures

$$(11) \quad \lambda^* \in \partial D(y, \cdot)(A\mu^*) \iff \forall i \in \{1, \dots, m\}, \lambda_i^* \in \partial d(y_i, \cdot)((A\mu^*)_i).$$

Remark 3.4. Let us then emphasise the following important (yet straightforward) property: if μ^* is optimal, then any $\lambda^* \in \partial D(y, \cdot)(A\mu^*)$ is such that

$$y_i \neq (A\mu^*)_i \implies \lambda_i^* \neq 0.$$

Indeed, assume $\lambda_i^* = 0$. Then the inclusion $\lambda_i^* = 0 \in \partial d(y_i, \cdot)((A\mu^*)_i)$ is equivalent to $(A\mu^*)_i$ minimising $d(y_i, \cdot)$, which by the separation property enforces $y_i = (A\mu^*)_i$.

3.3. Singularity theorems. We now address the following question

If the data y is not in the image cone $A(\mathcal{M}_+)$, when do the optimality conditions (10) lead to singular, or completely singular measures?

By singular, we recall that we mean measures that have a singular part with respect to the Lebesgue measure, where completely singular measures are singular measures which furthermore have a zero absolutely continuous part, as per Definition 1.1. In order to prove that a given measure μ is completely singular, our approach will be to show that $\mu \neq 0$ and that $\text{supp}(\mu)$ has zero Lebesgue measure.

We begin with a result ensuring that all optimal measures are completely singular (and not only singular), but which holds only under the restrictive assumption that the detector functions a_i be analytic. This generalises the result of [22] to general divergences.

Proposition 3.5. *Assume that the functions $(a_i)_{i=1, \dots, m}$ are linearly independent in $\mathcal{C}(K)$, analytic in $\text{int}(K)$.*

Then, if $y \notin A(\mathcal{M}_+)$, any optimal measure, if it is non zero, is completely singular.

Note that condition (H) is one sufficient condition that D can satisfy to ensure that all optimal measures μ^* satisfy $\mu^* \neq 0$, as the proof of Theorem 3.6 shows.

Proof. Let us pick some optimal measure $\mu^* \neq 0$. The optimality conditions (10) provide λ^* such that $\lambda^* \in \partial D(y, \cdot)(A\mu^*)$ and $A^*\lambda^* = 0$ on $\text{supp}(\mu^*)$. In other words, $\text{supp}(\mu^*)$ is contained in the zeros of the continuous function $\varphi^* := A^*\lambda^* = \sum_{i=1}^m \lambda_i^* a_i$.

Since $y \notin A(\mathcal{M}_+)$, there exists i such that $y_i \neq (A\mu^*)_i$. Using Remark 3.4, this implies $\lambda_i^* \neq 0$. Now, the linear independence of the functions a_i shows that φ^* is not identically zero. Since the function is by assumption analytic on $\text{int}(K)$, its zero set intersected with $\text{int}(K)$ is of zero Lebesgue measure. Note that ∂K has zero measure since $\text{int}(K)$ is assumed to be a Lipschitz bounded connected open set. Hence the support of μ^* has zero Lebesgue measure, which concludes the proof. \square

The main reason why this theorem is not satisfactory is that in many applications, the detector functions a_i are compactly supported inside K , even though they may be analytic (or piecewise analytic) on their support. This is why we relax the analyticity assumption and consider now detector functions a_i which are only *piecewise* analytic. Dealing with this more complicated case also requires restricting our attention to divergences which satisfy the property (H).

We make some assumptions on the detector functions a_i , which require some notations. First, we define

$$J_i := \{ j \in \{1, \dots, m\} \mid \text{supp}(a_j) \cap \text{supp}(a_i) \neq \emptyset \} \quad i \in \{1, \dots, m\}.$$

It is the set storing which detector functions a_j are active on $\text{supp}(a_i)$. In particular, we have $i \in J_i$.

We consider the following condition, which combines *local linear independence* and *piecewise analyticity*:

$$(C) \quad \text{there exists a partition } K = \bigcup_{k=1}^r \bar{\Omega}_k \text{ such that}$$

- the sets Ω_k , $k \in \{1, \dots, r\}$ are Lipschitz open connected sets,
- the detector functions a_i , $i \in \{1, \dots, m\}$ are piecewise analytic on the partition, i.e., for each $i \in \{1, \dots, m\}$, $k \in \{1, \dots, r\}$, a_i is analytic on Ω_k ,
- for all $i \in \{1, \dots, m\}$, the family $(a_j)_{j \in J_i}$ is linearly independent in $\mathcal{C}(\Omega_k)$ for any k such that $\Omega_k \cap \text{supp}(a_i) \neq \emptyset$.

Then, the following holds.

Theorem 3.6. *Assume that the divergence D satisfies (H), and that the functions $(a_i)_{i=1, \dots, m}$ satisfy (C). Then, if $y \notin A(\mathcal{M}_+)$, any μ^* optimal is singular.*

More precisely, for any $i_0 \in \{1, \dots, m\}$ such that $y_{i_0} \neq (A\mu^)_{i_0}$, $\mu^*_{|\text{supp}(a_{i_0})}$ is completely singular.*

Proof. Let μ^* be optimal. We may write the optimality conditions (10), namely

$$A^* \lambda^* \geq 0 \text{ on } K, \quad A^* \lambda^* = 0 \text{ on } \text{supp}(\mu^*),$$

where $\lambda_i^* \in \partial d(y_i, \cdot)((A\mu^*)_i)$ for all i from (11).

We shall prove that the following key property holds:

$$(P) \quad \forall i \in \{1, \dots, m\}, \quad y_i \neq (A\mu^*)_i \implies (A\mu^*)_i > 0.$$

Indeed, let us pick i_0 such that $(A\mu^*)_{i_0} = 0$; we need to show that $y_{i_0} = (A\mu^*)_{i_0} = 0$. If we had $y_{i_0} > 0$, we would find

$$\lambda_{i_0} \in \partial d(y_{i_0}, \cdot)((A\mu^*)_{i_0}) = \partial d(y_{i_0}, \cdot)(0),$$

contradicting the emptiness of the subdifferential as given by (H), and hence proving (P).

Continuing with $i_0 \in \{1, \dots, m\}$ such that $y_{i_0} \neq (A\mu^*)_{i_0}$ and hence such that $\lambda_{i_0}^* \neq 0$ as explained in Remark 3.4, Property (P) entails

$$(12) \quad \text{supp}(\mu^*) \cap \text{supp}(a_{i_0}) \neq \emptyset,$$

which in particular shows that $\mu^*_{|\text{supp}(a_{i_0})} \neq 0$. The optimality of μ^* also ensures

$$\text{supp}(\mu^*) \subset \left\{ x \in K \mid \sum_{j=1}^m \lambda_j^* a_j(x) = 0 \right\}.$$

Hence, we have

$$\text{supp}(\mu^*) \cap \text{supp}(a_{i_0}) \subset \left\{ x \in \text{supp}(a_{i_0}) \mid \sum_{j \in J_{i_0}} \lambda_j^* a_j(x) = 0 \right\},$$

and the set on the right-hand side cannot be empty owing to (12). We denote it K_{i_0} .

Now let us have a closer look at this set. Suppose that for some k , the set $\bar{\Omega}_k$ intersects $\text{supp}(a_{i_0})$. The set K_{i_0} may contain parts of the edges $\partial\Omega_k$. Other than that, its intersection with Ω_k cannot have positive Lebesgue measure. If it were the case, analyticity would imply $\sum_{j \in J_{i_0}} \lambda_j^* a_j = 0$ identically on Ω_k . Then, linear independence would enforce $\lambda_j^* = 0$ for all $j \in J_{i_0}$, contradicting $\lambda_{i_0}^* \neq 0$. Hence, the support of $\mu^*_{|\text{supp}(a_{i_0})}$ is included in a finite union of zero-measure sets (lying in the interior of the partition pieces) and of $(p-1)$ -dimensional edges (lying at their boundary). As a result, the support of $\mu^*_{|\text{supp}(a_{i_0})}$ is of measure zero, which concludes the proof. \square

Remark 3.7. We reached the conclusion that if $y \notin A(\mathcal{M}_+)$, any optimal measure μ^* will be such that $\mu^*_{|\text{supp}(a_{i_0})}$ is not the zero measure and has support of Lebesgue measure zero for at least some $i_0 \in \{1, \dots, m\}$. In fact, one can be a bit more specific and assert that $\mu^*_{|\text{supp}(a_{i_0})}$ has support included in a union of subvarieties, each of them having dimension at most $p-1$ [16, Theorem 6.3.3].

Remark 3.8. We emphasise that Theorem 3.6 does use the fact that $\lambda^* \in \partial D(y, \cdot)(A\mu^*)$, the optimality condition $A^*\lambda^* = 0$ over $\text{supp}(\mu^*)$, but it does not use the optimality condition $A^*\lambda^* \geq 0$ over K . The second optimality condition combined with the first one indeed yields

$$\text{supp}(\mu^*) \subset \arg \min \left(\sum_{j=1}^m \lambda_j^* a_j \right).$$

This in turn can only make measures “more singular”. For instance, one has for all $k \in \{1, \dots, r\}$

$$\text{supp}(\mu^*) \cap \Omega_k \subset \left\{ x \in \Omega_k \mid \sum_{j=1}^m \lambda_j^* \nabla a_j(x) = 0 \right\}.$$

4. DETECTOR FUNCTIONS a_i LEADING TO SINGULARITY

We investigate a few examples of detector functions and discuss whether they satisfy Property (C) or not. We start with a toy example and then consider a situation which more closely matches functions encountered in practice (in the case of PET).

4.1. Shifted polynomials in dimension 1. Let $\mathbb{R}_k[X]$ denote the vector space of real polynomials of degree at most $k \geq 0$ over \mathbb{R} .

Lemma 4.1. *Let $P \in \mathbb{R}_k[X]$ with $\deg(P) = k$, and z_1, \dots, z_r be distinct in \mathbb{R} . Then*

the family $(P(\cdot + z_1), \dots, P(\cdot + z_r))$ is linearly independent in $\mathbb{R}_k[X]$ if $r \leq k + 1$.

Proof. Since P has degree k , the set of vectors

$$e_\alpha := \frac{D^\alpha P}{\alpha!} \quad 0 \leq \alpha \leq k$$

is a basis of $\mathbb{R}_k[X]$. Now, using the Taylor formula for polynomials, we obtain

$$P(X + z) = \sum_{\alpha=0}^k z^\alpha \frac{D^\alpha P}{\alpha!}(X) = \sum_{\alpha=0}^k z^\alpha e_\alpha(X) \quad z \in \mathbb{R}.$$

If we express $(P(\cdot + z_1), \dots, P(\cdot + z_r))$ in the basis e_α , we thus obtain the Vandermonde matrix

$$\begin{pmatrix} 1 & 1 & \dots & 1 \\ z_1 & z_2 & \dots & z_r \\ z_1^2 & z_2^2 & \dots & z_r^2 \\ \vdots & \vdots & \dots & \vdots \\ z_1^k & z_2^k & \dots & z_r^k \end{pmatrix}.$$

Hence, since the shifts z_1, \dots, z_r are distinct, this matrix is of rank r as soon as $r \leq k + 1$, which concludes the proof. \square

Now, consider some univariate polynomial $P \in \mathbb{R}_k[X]$ such that $P > 0$ on some interval (a, b) , with $P(a) = P(b) = 0$. Let $a_0 = P\mathbf{1}_{(a,b)}$ and define a countable family of shifted functions by

$$a_0(\cdot - lh), \quad l \in \mathbb{Z},$$

where $h > 0$ is the grid spacing. Suppose K is some compact interval in \mathbb{R} , and m stands for the number the above functions whose support intersects the interior of K , which we denote a_1, \dots, a_m . For an example, see Figure 2.




FIGURE 2. Example of shifts with $P(x) := 1 - x^4$, $a = -1$, $b = 1$, $h = \frac{2}{5}$, shown over the compact $K = [0, 3]$. The resulting functions satisfy Assumption (C), as per Corollary 4.2.

Corollary 4.2. Let $r := \lceil \frac{b-a}{h} \rceil$ and assume $r \leq k+1$. Then the functions $(a_i)_{1 \leq i \leq m}$ satisfy Assumption (C).

Proof. The functions $(a_i)_{1 \leq i \leq m}$ are clearly piecewise analytic. Let us tackle the issue of local linear independence in more detail. We let $\Omega \subset K$ be any open interval such that all functions a_i are analytic on Ω . It is easily checked that there are at most r functions active on each such set Ω . Hence, if i is such that $\text{supp}(a_i)$ intersects Ω , a linear combination of $(a_j)_{j \in J_i}$ writes

$$\sum_{\lambda_j \in J_i} \lambda_j P(\cdot + z_j),$$

for some distinct shifts z_j (which are multiples of h). This polynomial vanishes on Ω , and hence is the zero polynomial. We may then use Lemma 4.1 to infer that $\lambda_j = 0$ for all $j \in J_i$, hence the local linear independence. \square

4.2. PET functions on a regular polygon. By *PET functions*, we informally refer to the functions associated to a continuous-discrete PET forward operator. In view of illustrating why PET functions are expected to satisfy Assumption (C), we focus on a representative example of PET, the 2D case with detectors regularly placed on the unit circle.

Let us review the basics of PET physics: $a_i(x)$ stands for the probability that an emission that occurred at position x leads to a detection by the i th pair of detectors. More precisely, a positron emitted at position x interacts with an electron, almost instantly and hence at a position which can be approximated to be x . This leads to the emission of two photons in opposite directions, where the direction is uniformly chosen at random. These two photons are then (almost) simultaneously detected by a pair of detectors.

4.2.1. Modelling for PET functions. Let $\theta = \frac{2\pi}{N}$ for some $N \in \mathbb{N}^*$. Working with complex notations, we denote $z_k = e^{ik\theta}$, $k = 0, \dots, N-1$. Define P_N to be the regular polygon associated to the points z_k , which lies inside the unit disk. Throughout this section, $T(u, v, w)$ stands for the (closed) triangle defined by three non-aligned complex numbers u, v, w , and $\text{int}(B)$ denotes the interior of a set $B \subset \mathbb{C}$.

For convenience, we use two indices j, k to index functions, letting $a_{j,k}$ be the function associated to the pair of detectors formed by the line segments $[z_j, z_{j+1}]$ and $[z_k, z_{k+1}]$ for $0 \leq j < k \leq N-1$. There are $m = \frac{1}{2}N(N-1)$ such functions.

Given two such segments and $z \in P_N$, $a_{j,k}(z)$ equals the probability that a line passing through z intersects both line segments, when the line orientation is chosen uniformly at random. This definition is ambiguous at the vertices z_k since these points belong to two detectors. In fact, one can decide for a convention at these points, but whatever the choice, some of the functions will not be continuous there, see also Figure 3.

In practice, however, the object to be imaged will not reach the detectors: hence the compact K to be considered will typically satisfy $K \subset \text{int}(P_N)$. This in turn will ensure that we are in the framework of functions in $\mathcal{C}(K)$. In what follows, we shall thus define the functions $a_{j,k}$ over $\text{int}(P_N)$. We note, however, that all the definitions given below are valid on $P_N \setminus \{z_0, \dots, z_{N-1}\}$.

4.2.2. Formulae for PET functions. We let $0 \leq j < k \leq N-1$ be fixed. Given the probabilistic model chosen for the PET functions, computing $a_{j,k}(z)$ boils down to computing angles (after proper normalisation).

Adjacent detectors. First assume that the two segments have one point in common. In other words, we are interested in the functions $a_{j,j+1}$, $0 \leq j \leq N-2$, and $a_{0,N-1}$.

The functions $a_{j,j+1}$, $0 \leq j \leq N-2$ vanish outside of $T(z_j, z_{j+1}, z_{j+2})$. Inside the triangle and for $z \neq z_j, z_{j+2}$, the angle we are looking for is the (non-oriented) angle $(\overrightarrow{z_j z}, \overrightarrow{z z_{j+2}})$. Hence, after normalisation and using the notation Arg for the principal argument, we find in complex notations

$$\forall z \in T(z_j, z_{j+1}, z_{j+2}) \cap \text{int}(P_N), \quad a_{j,j+1}(z) = \frac{1}{\pi} \text{Arg} \left(\frac{z_{j+2} - z}{z - z_j} \right).$$

These functions are continuous on $\text{int}(P_N)$ and analytic on $\text{int}(T(z_j, z_{j+1}, z_{j+2}))$.

The above reasoning also covers the function $a_{0,N-1}$, which vanishes outside of $T(z_{N-1}, z_0, z_1)$, and

$$\forall z \in T(z_{N-1}, z_0, z_1) \cap \text{int}(P_N), \quad a_{0,N-1}(z) = \frac{1}{\pi} \text{Arg} \left(\frac{z_{N-1} - z}{z - z_1} \right).$$

Non-adjacent detectors. Now, consider the remaining functions $a_{j,k}$, i.e., when $2 \leq k-j \leq N-2$. Then, $[z_j, z_{j+1}]$ and $[z_k, z_{k+1}]$ have no point in common and the lines (z_{j+1}, z_k) and (z_j, z_{k+1}) are easily shown to be parallel. Again, $a_{j,k}$ is zero outside of the closed trapezium defined by $z_j, z_{j+1}, z_k, z_{k+1}$.

With $c_{j,k}$ denoting the point at which the diagonals of the trapezium intersect, we find by reasoning as in the adjacent case

- If $z \in T(z_j, z_{j+1}, c_{j,k}) \cap \text{int}(P_N)$,

$$(13) \quad a_{j,k}(z) = \frac{1}{\pi} (\overrightarrow{z_k z}, \overrightarrow{z z_{k+1}}) = \frac{1}{\pi} \text{Arg} \left(\frac{z - z_{k+1}}{z - z_k} \right).$$

- If $z \in T(z_k, z_{k+1}, c_{j,k}) \cap \text{int}(P_N)$,

$$(14) \quad a_{j,k}(z) = \frac{1}{\pi} (\overrightarrow{z_j z}, \overrightarrow{z z_{j+1}}) = \frac{1}{\pi} \text{Arg} \left(\frac{z - z_{j+1}}{z - z_j} \right).$$

- If $z \in T(z_{j+1}, z_k, c_{j,k}) \cap \text{int}(P_N)$,

$$(15) \quad a_{j,k}(z) = \frac{1}{\pi} (\overrightarrow{z_k z}, \overrightarrow{z z_{j+1}}) = \frac{1}{\pi} \text{Arg} \left(\frac{z_{j+1} - z}{z - z_k} \right).$$




FIGURE 3. Two views of the functions described in § 4.2: here the two detectors are the line segment $[u, v]$ with $u = 1 + 2i$, $v \sim -0.7 + 1.4i$, and the symmetric line segment $[\bar{u}, \bar{v}]$ with respect to the real axis. We only plot the function on half of its domain since it is symmetric with respect to the real axis. As can be seen, the functions cannot be continuously extended up to the vertices.

- If $z \in T(z_{k+1}, z_j, c_{j,k}) \cap \text{int}(P_N)$,

$$(16) \quad a_{j,k}(z) = \frac{1}{\pi} \frac{\overrightarrow{z_j z} \cdot \overrightarrow{z z_{k+1}}}{|\overrightarrow{z_j z}| |\overrightarrow{z z_{k+1}}|} = \frac{1}{\pi} \text{Arg} \left(\frac{z_{k+1} - z}{z - z_j} \right).$$

One can check that these definitions are coherent where the triangles intersect. We note that these functions are continuous on $\text{int}(P_N)$. Furthermore, they are analytic on all four open subtriangles $\text{int}(T(z_j, z_{j+1}, c_{j,k}))$, $\text{int}(T(z_k, z_{k+1}, c_{j,k}))$, $\text{int}(T(z_{j+1}, z_k, c_{j,k}))$ and $\text{int}(T(z_{k+1}, z_j, c_{j,k}))$.

Remark 4.3. In fact, the formulae (13)-(14)-(15)-(16) derived above apply to any pair of detectors given by two line segments that do not intersect.

The situation is illustrated by Figure 3 which shows the plot of a typical PET function between two given segments. Figure 4 shows the supports of two PET functions in the case where $N = 12$.

4.2.3. *PET functions and Assumption (C).* The angle interpretation shows that we obviously have $a_{j,k} \geq 0$ for all j, k , as well as the relation

$$\forall z \in \text{int}(P_N), \quad \sum_{0 \leq j < k \leq N-1} a_{j,k}(z) = 1,$$

hence PET functions satisfy the positivity assumption (3) over any compact $K \subset \text{int}(P_N)$. By symmetry, we also have for any $j < k$, $j' < k'$ such that $k' - j' = k - j = r$

$$\forall z \in \text{int}(P_N), \quad a_{j',k'}(z) = a_{j,k}(e^{-ir\theta} z),$$

meaning that all functions can be obtained from $a_{0,j}$, $j = 1, \dots, N-1$, up to appropriate rotations, which is reminiscent of the previous toy example with functions all equal up to translation. Here, there is not 1 but N "base" functions from which

all the others are deduced by rotation. Finally, we note that further (mirror) symmetries can be exploited, with respect to line segments linking a given vertex and its diametric opposite vertex when N is even (or a given vertex and the middle point of the diametrically opposed detector when N is odd).

From the above discussion, PET functions are in $\mathcal{C}(K)$ for any compact $K \subset \text{int}(P_N)$. Furthermore, we may build a partition ensuring that they are piecewise analytic in the sense defined by Assumption (C) in any compact $K \subset \text{int}(P_N)$.

A partition that ensures piecewise analyticity is easily built by drawing all line segments $[z_j, z_k]$ (and then intersecting with K). This construction naturally isolates the four subtriangles associated to non-adjacent detectors.

The issue of local linear independence (associated to such partitions) is more subtle. We conjecture that the PET functions are locally linearly independent, for any value of N and any (reasonably large) compact $K \subset \text{int}(P_N)$. To support this conjecture, let us consider the case of $N = 4$ with $K \subset \text{int}(P_4)$ any (reasonable) compact.

Proposition 4.4. *For $N = 4$, and any compact $K \subset \text{int}(P_4)$ containing at least a neighbourhood of 0, the PET functions are locally linearly independent over K .*

Proof. There are 6 PET functions for $N = 4$, and the partition described above is obtained by considering the four quadrants intersected with K . By symmetry, we will be done if we prove the result for only one of the four: we focus on the first quadrant, i.e.

$$\Omega_1 := \text{int}(K \cap \{z \in \mathbb{C}, \text{Re}(z) > 0, \text{Im}(z) > 0\}).$$

Over Ω_1 , there are 4 active PET functions, namely $a_{0,1}$, $a_{0,2}$, $a_{0,3}$ and $a_{1,3}$. Consider a vanishing linear combination, i.e. $\lambda_{0,1}$, $\lambda_{0,2}$, $\lambda_{0,3}$ and $\lambda_{1,3}$ real numbers such that

$$\forall z \in \Omega_1, \quad \lambda_{0,1}a_{0,1}(z) + \lambda_{0,2}a_{0,2}(z) + \lambda_{0,3}a_{0,3}(z) + \lambda_{1,3}a_{1,3}(z) = 0.$$

In fact, this equality extends to the boundary of Ω_1 by continuity of the functions on K .

We successively evaluate the equality at $z = 0$, any $z \neq 0$ on the real axis, any $z \neq 0$ on the imaginary axis and any z on the first diagonal.

For $z = 0$, $a_{0,1}(z) = a_{0,3}(z) = 0$ and by symmetry $a_{0,2}(z) = a_{1,3}(z) \neq 0$, hence we find $\lambda_{0,2} + \lambda_{1,3} = 0$. For $z \neq 0, z \in \Omega_1$ on the real axis, $a_{0,1}(z) = 0$, and by symmetry $a_{0,2}(z) = a_{1,3}(z) \neq 0$, hence $\lambda_{0,3} = 0$ since $a_{0,3}(z) \neq 0$. Similarly with $z \neq 0, z \in \Omega_1$ on the imaginary axis, we uncover $\lambda_{0,1} = 0$. Finally, one can check that the function $t \mapsto a_{0,2}(z(t)) - a_{1,3}(z(t))$ with $z(t) = (1+i)t$, i.e., on the first diagonal, has a positive derivative at $t = 0$. Hence $a_{0,2}(z) > a_{1,3}(z)$ for any $z \in \Omega_1, z \neq 0$ on the first diagonal sufficiently close to 0. Picking such an element z leads to $\lambda_{0,2}a_{0,2}(z) + \lambda_{1,3}a_{1,3}(z) = 0$ which, together with $\lambda_{0,2} + \lambda_{1,3} = 0$, implies $\lambda_{0,2} = \lambda_{1,3} = 0$ and concludes the proof. \square

5. NUMERICAL EXPERIMENTS

We here consider some simulations of algorithms solving (4) in different contexts, where the results exhibit singularity as expected from the theoretical results.

5.1. Singularity Certificates. As evidenced by our results, under the Assumptions (H) and Assumption (C), the relevant criterion for singularity of optimisers of (4) is independent of the specific divergence D , as the question reduces to:

$$\text{do we have } y \in A(\mathcal{M}_+)?$$

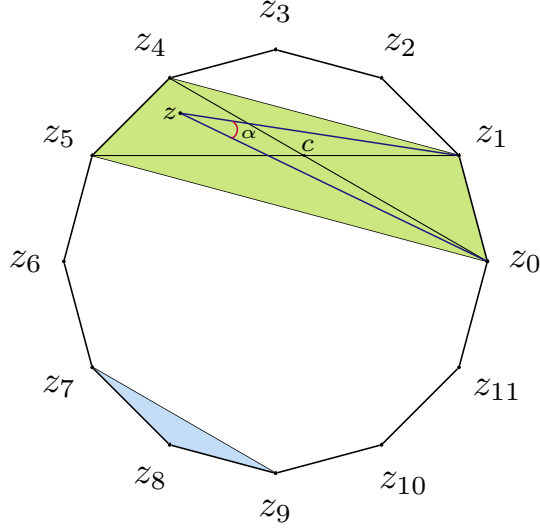


FIGURE 4. Example of PET functions for the regular polygon P_N , with $N = 12$. Outside of the pale blue region, $a_{7,8}$ vanishes. Outside of the pale green region (a trapezium whose diagonals intersect at the point denoted $c = c_{0,4}$), the function $a_{0,4}$ vanishes. At the point z , $a_{0,4}(z)$ is given by $\frac{\alpha}{\pi}$.

In practice, as one wants to solve (4) (or possibly a regularised version thereof) in the form of some iterative algorithm defined by iterates of the form

$$\mu_{k+1} = G_k(\mu_k),$$

we are looking for methods allowing us to *guarantee* that $y \notin A(\mathcal{M}_+)$. Our strategy is to devise a method to prove that $y \notin A(\mathcal{M}_+)$ which writes as a function of μ_k and that gets better as $k \rightarrow +\infty$.

One approach towards this is to make use of duality: by weak duality (and recalling the definition of g in (8)) we always have

$$\forall \mu \in \mathcal{M}_+, \forall \lambda \in A(\mathcal{M}_+)^*, \quad \ell(\mu) \geq g(\lambda)$$

where we recall that ℓ refers to the functional to be minimised $\ell(\mu) = D(y, A\mu)$ as per definition (5). Since $y \in A(\mathcal{M}_+)$ if and only if $\min_{\mu \in \mathcal{M}_+} \ell(\mu) = 0$, this entails the following straightforward result:

$$(\exists \lambda \in A(\mathcal{M}_+)^*, g(\lambda) > 0) \implies y \notin A(\mathcal{M}_+).$$

We will call a vector $\lambda \in A(\mathcal{M}_+)^*$ such that $g(\lambda) > 0$ a *dual certificate of singularity*.

Remark 5.1. We note that checking whether $\lambda \in A(\mathcal{M}_+)^* \iff A^*\lambda \geq 0$ for a given $\lambda \in \mathbb{R}^m$ can easily be done at the discrete level, i.e., for the discretised version of the operator A . However, it does not mean that $A^*\lambda \geq 0$ holds at the continuous level; one would need to control discretisation errors in order to ensure such an inequality. When performing numerical experiments, we did not account for them. Hence, whenever we will have claimed to have found a dual certificate, this will be abusive and will refer to discrete certificates.

This provides a natural method when it comes to establishing singularity: assume we have a convergent algorithm for solving (4), in the sense that each subsequence

of $(\mu_k)_{k \in \mathbb{N}}$ has subsequence that converges (in the weak-* sense) to some minimiser μ^* of (4).

Ofentimes, one can derive an explicit link between primal and dual variables from the relation

$$\lambda^* \in \partial D(y, \cdot)(w^*),$$

which we know hold for optimal w^* under the hypotheses of Theorem 3.6. For instance, in the case of β -divergences, an optimal dual variable $\lambda^* \in A(\mathcal{M}_+)^*$ is related to a primal variable w^* by $\lambda^* := (w^*)^{\beta-2}(w^* - y)$ where powers and multiplications (or divisions) are to be understood componentwise. Then, a candidate of choice for a dual certificate is given by

$$\lambda_k := (A\mu_k)^{\beta-2}(A\mu_k - y).$$

In particular, λ_k will converge to λ^* along subsequences, hence the convergence of $g(\lambda_k)$ to $\max_{\lambda \in A(\mathcal{M}_+)^*} g(\lambda) = \min_{\mu \geq 0} \ell(\mu)$, the last equality being valid since strong duality obtains. Hence, if $y \notin A(\mathcal{M}_+)$, we should have $\lim g(\lambda_k) > 0$ as $k \rightarrow +\infty$.

A caveat with our choice is that we should only have dual admissibility $\lambda_k \in A(\mathcal{M}_+)^*$ at the limit $k \rightarrow +\infty$, and not for a fixed iteration number k . In practice, we take k large and if $\lambda_k \notin A(\mathcal{M}_+)^*$, we set $\tilde{\lambda}_k := \lambda_k + c$ where $c > 0$ is the smallest constant restoring dual admissibility, estimated by bisection. In other words, we choose the minimal $c > 0$ such that $A^*\tilde{\lambda}_k \geq 0$, i.e., $\tilde{\lambda}_k \in A(\mathcal{M}_+)^*$. Such a constant c exists because $A^*(\lambda_k + c) = A^*\lambda_k + cA^*1$ and $A^*1 = \sum_{i=1}^m a_i > 0$ over K by assumption (3).

5.2. Emission Tomography Example. We first look at an example from PET, where the aim is to solve (4). We are investigating whether optimisers are singular depending on the noise level. Indeed, our conjecture that PET functions satisfy hypothesis (C) means that singular measures should be obtained, at least for a sufficient amount of noise.

When the divergence used is a β -divergence D_β with $\beta \in [1, 2]$, a common way to solve the optimisation problem (4) is to use the following iterates, called *multiplicative* [17, 13]. Starting from some $\mu_0 \in \mathcal{M}_+$ (typically with a positive constant density over the domain), the iterates write

$$\mu_{k+1} = \mu_k \frac{A^*((A\mu_k)^{\beta-2}y)}{A^*((A\mu_k)^{\beta-1})},$$

assuming that one can prove $\frac{A^*((A\mu_k)^{\beta-2}y)}{A^*((A\mu_k)^{\beta-1})} \in \mathcal{C}$ along iterates.

These iterates have the ML-EM algorithm ($\beta = 1$) and the Iterative Image Space Reconstruction ($\beta = 2$) as particular cases [10], and proofs of convergence for these algorithms with any $\beta \in [1, 2]$ can be found in [29], in the finite-dimensional case. One advantage of these algorithms is the decrease of the functional ℓ along iterates, see the proof in [13] in the finite-dimensional setting.

We now present the results of applying the algorithm in the case of a 2D PET operator A with 90 views and 64 tangential positions (hence, $m = 90 \times 64 = 5760$). We run the simulations using the Operator Discretization Library [1], a Python wrapper around the ASTRA toolbox [21].

The image μ is the Derenzo phantom, denoted μ_r . The data is obtained by (re-scaled) Poisson draws, with a time-variable (or dose-variable) t which accounts for the level of noise. In other words, $y \sim \frac{1}{t}\mathcal{P}(tA\mu_r)$, and the higher t , the lower the noise. In order to approach the infinite-dimensional setting of our work, we increase the resolution to 512×512 pixels.

Finally, we take $\beta = 1.2$, on purpose not quite matching the noise statistics, as Poisson noise should lead one to take $\beta = 1$. We hence mimic the situation of not knowing the exact noise statistics.

In Figure 5 and Figure 6, we display the evolution of the loss function ℓ along iterates, i.e. $k \mapsto \ell(\mu_k)$, starting from $\mu_0 = 1$, for noise levels $t = 1$ and $t = 10^{-1}$ respectively. As expected, the function decreases. We also plot the maximum attained for each reconstruction, namely $k \mapsto \max(\mu_k)$, which tends to increase. Finally, we show the reconstruction after $k = 100$ and $k = 1000$ iterates.

In the noisier case $t = 1$, some pixels clearly take over as one keeps iterating. Moreover, we can guarantee that we should indeed expect singularity, as we may provide a dual certificate proving that the data y is not in the image cone $A(\mathcal{M}_+)$. This is also suggested by the fact that $k \mapsto \ell(\mu_k)$ seems to converge to a positive value rather than to zero.




FIGURE 5. Case $t = 1$. (A) Divergence along iterates. (B) Maximum of reconstruction along iterates. (C) Reconstruction after 100 iterates. (D) Reconstruction after 1000 iterates.

In the less noisy case $t = 10^{-1}$, it seems that the loss $\ell(\mu_k)$ is not converging to zero, which may be a hint that we should also expect singularity. However, in this case we are not able to guarantee it with dual certificates. Note that if singularity were to be true in this case as well, the reason could be that many more iterates are needed to ascertain its presence.

In fact, we conjecture that singularity does arise at this noise level. But, as our theoretical results suggest, singularity is expected on specific parts of the domain.

This is what seem to be observed for medium noise: although some few pixels take larger values along iterates, the rest of the image remains rather smooth. As a result, cropping the images to some value is a good practical solution to mitigating the night-sky effect.




FIGURE 6. Case $t = 10^{-1}$. (A) Divergence along iterates. (B) Maximum of reconstruction along iterates. (C) Reconstruction after 100 iterates. (D) Reconstruction after 1000 iterates.

5.3. Examples with $\beta = 2$ and Regularisation.

5.3.1. *Toy Example.* We illustrate singularity with a toy example. There will just be two detectors a_1 and a_2 , so $m = 2$. We choose specifically

$$K = [0, 1], \quad a_1 = 1, \quad a_2(x) = x.$$

With such analytic a_i functions, Proposition 3.5 applies.

We also choose $\beta = 2$. Recall from § 2.3.1 that the divergence reduces to the Euclidean distance in this case. We can now compute the singular solutions explicitly depending on the parameter $y \in \mathbb{R}^2$, as shown in Figure 7. For $y \in \mathbb{R}_+^2$ outside of the cone

$$A(\mathcal{M}_+) = \{(y_1, y_2) \in \mathbb{R}^2, 0 \leq y_1 \leq y_2\},$$

the solution is of the form $\xi\delta_1$ with $\xi \geq 0$ a varying parameter. For completeness, we also display the solution for $y \notin \mathbb{R}_+^2$ (which is either 0 or $\xi\delta_0$ with $\xi \geq 0$ a varying parameter).




FIGURE 7. The singular solutions of the problem (4) with two detectors, $a_1 = 1$, $a_2(x) = x$ on the interval $K = [0, 1]$. There are three distinct regions outside the cone AM_+ , but only one which intersects the first quadrant $y \in \mathbb{R}_+^2$, where the optimal solution is a Dirac, given by $\xi\delta_1$ for some $\xi \geq 0$.




FIGURE 8. The solutions for $y = (0, 1)$ of the problem in Figure 7 with total variation regularisation. Each curve corresponds to a different regularisation parameter, labelled by its base-10 logarithm. We see that when the regularisation parameter goes to zero, the computed solution converges to the expected exact solution, which, we see from Figure 7, is $\mu = .5\delta_1$, depicted by a blue circle.

We also look at the effect of regularisation. In this case, following standard practice in many image reconstruction problems, we use total variation regularisation, that is, we solve

$$\min_{\mu \geq 0} \ell(\mu) + \rho \text{TV}(\mu),$$

where $\text{TV}(\mu)$ is the total variation of the derivative of the measure μ and $\rho > 0$ is a regularisation parameter. In this discretised, one-dimensional setting, this is simply $\text{TV}(\mu) = \sum_i |\mu^{i+1} - \mu^i|$, where μ^i is the value of the discretised measure at pixel i . We then compute the minimum using a primal-dual hybrid gradient method [6]. We plot the resulting minima for various values of the regularisation parameter ρ

in Figure 8, in the case $y = (0, 1)$. As ρ goes to 0, the solution approaches the expected singular limit $\frac{1}{2}\delta_1$.

5.3.2. Tomography example. We finally look at a more realistic 2D example taken from tomography, where the unknown μ equals the usual Shepp-Logan phantom used as a benchmark in CT tomography. The example of Figure 9 features an image resolution of 127×127 , and there are 285 angles and 183 tangential coordinates.

We consider the case of the Euclidean distance d_2 . The data y is obtained by Gaussian draws with negative values clipped to 0, i.e., $y \sim \max(\mathcal{N}(A\mu, \sigma^2), 0)$, which ensures $y \in \mathbb{R}_+^m$. We then solve the corresponding TV-regularised problem

$$\min_{\mu \geq 0} D_2(y, A\mu) + \rho \text{TV}(\mu),$$


where $\rho > 0$, by a primal-dual hybrid gradient method. Here, TV refers to isotropic TV.

As Figure 9 shows, the maximum of the reconstruction gets bigger as ρ tends to 0. In fact, the reconstruction for small ρ clearly exhibits the night-sky effect. The fact that $\rho = 0$ does not lead to higher maximum values is due to the resolution, which acts as a regulariser for the expected singular measures.


Acknowledgments. The authors are indebted to an anonymous referee whose remarks greatly improved the first versions of this paper.

REFERENCES


- [1] ADLER, J., KOHR, H., AND ÖKTEM, O. ODL-a Python framework for rapid prototyping in inverse problems. *Royal Institute of Technology* (2017).
- [2] BOYER, C., CHAMBOLLE, A., CASTRO, Y. D., DUVAL, V., DE GOURNAY, F., AND WEISS, P. On representer theorems and convex regularization. *SIAM Journal on Optimization* 29, 2 (2019), 1260–1281.
- [3] BREDIES, K., AND CARIONI, M. Sparsity of solutions for variational inverse problems with finite-dimensional data. *Calculus of Variations and Partial Differential Equations* 59, 1 (2020), 1–26.
- [4] BREZIS, H. *Functional analysis, Sobolev spaces and partial differential equations*. Springer Science & Business Media, 2010.
- [5] CAVALCANTI, Y. C., OBERLIN, T., DOBIGEON, N., FÉVOTTE, C., STUTE, S., RIBEIRO, M.-J., AND TAUBER, C. Factor analysis of dynamic PET images: beyond Gaussian noise. *IEEE transactions on medical imaging* (2019).
- [6] CHAMBOLLE, A., AND POCK, T. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision* 40, 1 (2010), 120–145.
- [7] CICHOCKI, A., AND AMARI, S.-I. Families of alpha-beta-and gamma-divergences: Flexible and robust measures of similarities. *Entropy* 12, 6 (2010), 1532–1568.
- [8] CLASON, C., KALTENBACHER, B., AND RESMERITA, E. Regularization of ill-posed problems with non-negative solutions. In *Splitting Algorithms, Modern Operator Theory, and Applications*. Springer, 2019, pp. 113–135.
- [9] CLASON, C., AND SCHIELA, A. Optimal control of elliptic equations with positive measures. *ESAIM: Control, Optimisation and Calculus of Variations* 23, 1 (2017), 217–240.
- [10] DE PIERRO, A. R. On the relation between the ISRA and the EM algorithm for positron emission tomography. *IEEE transactions on Medical Imaging* 12, 2 (1993), 328–333.
- [11] DEBARRE, T., DENOYELLE, Q., AND FAGEOT, J. On the uniqueness of solutions for the basis pursuit in the continuum. *ArXiv preprint arXiv:2009.11855* (2020).
- [12] DENOYELLE, Q., DUVAL, V., AND PEYRÉ, G. Support recovery for sparse super-resolution of positive measures. *Journal of Fourier Analysis and Applications* 23, 5 (2017), 1153–1194.
- [13] FÉVOTTE, C., AND IDIER, J. Algorithms for nonnegative matrix factorization with the β -divergence. *Neural computation* 23, 9 (2011), 2421–2456.
- [14] GEORGIU, T. T. Solution of the general moment problem via a one-parameter imbedding. *IEEE transactions on automatic control* 50, 6 (2005), 811–826.
- [15] HENROT, S., SOUSSEN, C., AND BRIE, D. Fast positive deconvolution of hyperspectral images. *IEEE Transactions on Image Processing* 22, 2 (2012), 828–833.



(A) The maximum of the reconstruction for various regularisation parameters ρ



(B) Reconstruction with no regularisation ($\rho = 0$) (C) Reconstruction with small regularisation ($\rho = 0.4$)



(D) Reconstruction with medium regularisation ($\rho = 0.8$) (E) Reconstruction with high regularisation ($\rho = 1.2$)

FIGURE 9. Relation between regularisation parameter and the reconstruction of the Shepp–Logan phantom μ . Here, the data is given by $y \sim \max(\mathcal{N}(A\mu, \sigma^2), 0)$, with $\sigma = 10$.

- [16] KRANTZ, S. G., AND PARKS, H. R. *A primer of real analytic functions*. Springer Science & Business Media, 2002.
- [17] LEE, D. D., AND SEUNG, H. S. Algorithms for non-negative matrix factorization. In *Advances in neural information processing systems* (2001), pp. 556–562.
- [18] LOHÉAC, J., TRÉLAT, E., AND ZUAZUA, E. Minimal controllability time for the heat equation under unilateral state or control constraints. *Mathematical Models and Methods in Applied Sciences* 27, 09 (2017), 1587–1644.
- [19] MAIR, B., RAO, M., AND ANDERSON, J. Positron emission tomography, Borel measures and weak convergence. *Inverse Problems* 12, 6 (1996), 965.
- [20] ÖKTEM, O., POUCHOL, C., AND VERDIER, O. Spatiotemporal PET reconstruction using ML-EM with learned diffeomorphic deformation. In *International Workshop on Machine Learning for Medical Image Reconstruction* (2019), Springer, pp. 151–162.
- [21] PALENSTIJN, W. J. Astra toolbox, Jan. 2012.
- [22] POUCHOL, C., AND VERDIER, O. The ML-EM algorithm in continuum: sparse measure solutions. *Inverse Problems* 36, 3 (2020).
- [23] POUCHOL, C., AND VERDIER, O. Statistical model and ML-EM algorithm for emission tomography with known movement. *Journal of Mathematical Imaging and Vision* (2021), 1–14.
- [24] RESMERITA, E., ENGL, H. W., AND IUSEM, A. N. The expectation-maximization algorithm for ill-posed integral equations: a convergence analysis. *Inverse Problems* 23, 6 (2007), 2575.
- [25] ROCKAFELLAR, R. T. Extension of Fenchel’s duality theorem for convex functions. *Duke Mathematical Journal* 33 (1966), 81–89.
- [26] SHEPP, L. A., AND VARDI, Y. Maximum likelihood reconstruction for emission tomography. *IEEE transactions on medical imaging* 1, 2 (1982), 113–122.
- [27] SIMSEKLI, U., CEMGIL, A. T., AND YILMAZ, Y. K. Learning the beta-divergence in tweedie compound poisson matrix factorization models. In *International Conference on Machine Learning* (2013), pp. 1409–1417.
- [28] VARDI, Y., SHEPP, L., AND KAUFMAN, L. A statistical model for positron emission tomography. *Journal of the American statistical Association* 80, 389 (1985), 8–20.
- [29] YANG, Z., AND OJA, E. Unified development of multiplicative algorithms for linear and quadratic nonnegative matrix factorization. *IEEE transactions on neural networks* 22, 12 (2011), 1878–1891.

APPENDIX A. COMPUTATION OF THE DUAL

Let us denote

$$\psi_y(w) := d_\beta(y, w) - \lambda w.$$

The goal is to compute the function h defined by (9).

A.1. Case $\beta = 1$. We have $\psi_y(w) \sim (1 - \lambda)w$ as $w \rightarrow +\infty$, hence $h(y, \lambda) = -\infty$ if $\lambda > 1$. If $\lambda = 1$ and $y > 0$, $\psi_y(w) \sim -y \ln(w)$ as $w \rightarrow 0$: the function tends to $-\infty$ as $w \rightarrow 0$ and $h(y, \lambda) = -\infty$. If $y = 0$, the function equals 0 identically and its minimum is 0.

We now focus on the case $\lambda < 1$. We still have $\psi_y(w) \sim (1 - \lambda)w$ as $w \rightarrow +\infty$. If $y > 0$, $\psi_y(w) \sim -y \ln(w)$ as $w \rightarrow 0$, thus the function tends to $+\infty$ at both ends. Since ψ_y is strictly convex in this case, it has a unique minimum for $w > 0$, which we again denote $w(y, \lambda)$, solving

$$w(y, \lambda)^{-1}(w(y, \lambda) - y) = \lambda \iff w(y, \lambda) = \frac{y}{1 - \lambda}.$$

If $y = 0$, the function $\psi_y(w) = w - \lambda w$ is minimised at $w = 0$, with value 0.

We may also gather the cases $y = 0$ and $y > 0$ whenever $\lambda \leq 1$, since the formula for $w(y, \lambda)$ shows that it vanishes with y .

Summing up, we find

$$h(y, \lambda) = \begin{cases} -\infty & \text{if } \lambda > 1 \\ d_1(y, w(y, \lambda)) - \lambda w(y, \lambda) & \text{if } \lambda \leq 1 \end{cases}$$

In the last case, further computations lead to

$$d_1(y, w(y, \lambda)) - \lambda w(y, \lambda) = y \ln(1 - \lambda).$$

A.2. Case $1 < \beta < 2$. If $y > 0$, $\psi_y(w) \sim \frac{w^\beta}{\beta}$ as $w \rightarrow +\infty$, and $\psi_y(w) \sim -\frac{y}{\beta-1}w^{\beta-1}$ as $w \rightarrow 0$. The derivative ψ'_y satisfies $\psi'_y(w) \sim -yw^{\beta-2}$ as $w \rightarrow 0$, and $\psi'_y(w) \sim w^{\beta-1}$ as $w \rightarrow +\infty$. Thus the function ψ' increases (as ψ is convex) from $-\infty$ to $+\infty$. As a consequence, it has a unique minimum $w > 0$, which we again denote $w(y, \lambda)$, solving

$$w(y, \lambda)^{\beta-2}(w(y, \lambda) - y) = \lambda.$$

If $y = 0$ and $\lambda \leq 0$, it is easily seen that the function $\psi_y(w) = \frac{1}{\beta}w^\beta - \lambda w$ is minimised at $w = 0$, with value 0, whereas if $\lambda > 0$, it has a unique minimum (also defined by the equation for $w(y, \lambda)$). Summing up, we find

$$h(y, \lambda) = \begin{cases} 0 & \text{if } y = 0, \lambda \leq 0 \\ d_\beta(y, w(y, \lambda)) - \lambda w(y, \lambda) & \text{otherwise.} \end{cases}$$

A.3. Case $\beta = 2$. In this case, ψ_y has a unique minimum w given by $w = \lambda + y$, which gives the explicit formula

$$h(y, \lambda) = -\frac{1}{2}(\lambda + y)^2 + \frac{1}{2}y^2.$$

UNIVERSITÉ PARIS CITÉ, FP2M, CNRS FR 2036, MAP5 UMR 8145, F-75006 PARIS, FRANCE.
E-mail address: `camille.pouchol@u-paris.fr`

DEPARTMENT OF COMPUTING, ELECTRICAL ENGINEERING AND MATHEMATICAL SCIENCES, WESTERN NORWAY UNIVERSITY OF APPLIED SCIENCES, BERGEN, NORWAY.
E-mail address: `olivier.verdier@hvl.no`