



Heriot-Watt University
Research Gateway

Revisiting Maximum-A-Posteriori Estimation in Log-Concave Models

Citation for published version:

Pereyra, M 2019, 'Revisiting Maximum-A-Posteriori Estimation in Log-Concave Models', *SIAM Journal on Imaging Sciences*, vol. 12, no. 1, pp. 650-670. <https://doi.org/10.1137/18M1174076>

Digital Object Identifier (DOI):

[10.1137/18M1174076](https://doi.org/10.1137/18M1174076)

Link:

[Link to publication record in Heriot-Watt Research Portal](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

SIAM Journal on Imaging Sciences

Publisher Rights Statement:

© 2019, Society for Industrial and Applied Mathematics

General rights

Copyright for the publications made accessible via Heriot-Watt Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

Heriot-Watt University has made every reasonable effort to ensure that the content in Heriot-Watt Research Portal complies with UK legislation. If you believe that the public display of this file breaches copyright please contact open.access@hw.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Revisiting Maximum-A-Posteriori Estimation in Log-Concave Models*

Marcelo Pereyra[†]

Abstract. Maximum-a-posteriori (MAP) estimation is the main Bayesian estimation methodology in imaging sciences, where high dimensionality is often addressed by using Bayesian models that are log-concave and whose posterior mode can be computed efficiently by convex optimization. However, despite its success and wide adoption, MAP estimation is not theoretically well understood yet. In particular, the prevalent view in the community is that MAP estimation is not proper Bayesian estimation in the sense of Bayesian decision theory because it does not minimize a meaningful expected loss function (unlike the minimum mean squared error (MMSE) estimator that minimizes the mean squared loss). This paper addresses this theoretical gap by presenting a general decision-theoretic derivation of MAP estimation in Bayesian models that are log-concave. A main novelty is that our analysis is based on differential geometry and proceeds as follows. First, we use the underlying convex geometry of the Bayesian model to induce a Riemannian geometry on the parameter space. We then use differential geometry to identify the so-called natural or canonical loss function to perform Bayesian point estimation in that Riemannian manifold. For log-concave models, this canonical loss coincides with the Bregman divergence associated with the negative log posterior density. Following on from this, we show that the MAP estimator is the only Bayesian estimator that minimizes the expected canonical loss, and that the posterior mean or MMSE estimator minimizes the dual canonical loss. We then study the question of MAP and MMSE estimation performance in high dimensions. Precisely, we establish a universal bound on the expected canonical error as a function of image dimension, providing new insights on the good empirical performance observed in convex problems. Together, these results provide a new understanding of MAP and MMSE estimation in log-concave settings, and of the multiple beneficial roles that convex geometry plays in imaging problems. Finally, we illustrate this new theory by analyzing the regularization-by-denoising Bayesian models, a class of state-of-the-art imaging models where priors are defined implicitly through image denoising algorithms, and an image denoising model with a wavelet shrinkage prior.

Key words. Bayesian inference, maximum-a-posteriori estimation, decision theory, mathematical imaging, inverse problems, differential geometry, convex optimization

AMS subject classifications. 62C10, 60E05, 53B21, 94A08, 68U10

DOI. 10.1137/18M1174076

1. Introduction. We consider the estimation of an unknown image $x \in \mathbb{R}^n$ from some data y , related to x by a statistical model with likelihood $p(y|x)$. Adopting a Bayesian statistical approach, we postulate a prior distribution $p(x)$ modeling the prior knowledge available and

*Received by the editors March 6, 2018; accepted for publication (in revised form) January 16, 2019; published electronically March 28, 2019.

<http://www.siam.org/journals/siims/12-1/M117407.html>

Funding: The work of the author was funded through a Marie Curie Intra-European Research Fellowship for Career Development by the European Research Council. The author conducted the work partly at the University of Bristol, partly at Heriot Watt University, and partly when he was a visiting professor at the Institut Henri Poincaré in France.

[†]School of Mathematical and Computer Sciences, Heriot Watt University, Edinburgh, EH14 4AS, UK (m.pereyra@hw.ac.uk).

base our inferences on the posterior distribution [32]

$$p(x|y) = \frac{p(y|x)p(x)}{\int_{\mathbb{R}^n} p(y|x)p(x)dx},$$

which models our knowledge about x after observing y . In this paper we focus on the case where $p(x|y)$ belongs to the class of log-concave distribution, i.e.,

$$(1) \quad p(x|y) = \frac{\exp\{-\phi(x)\}}{\int_{\mathbb{R}^n} \exp\{-\phi(s)\}ds},$$

for some proper convex function $\phi : \mathbb{R}^n \rightarrow (-\infty, \infty)$ [9], and where we note that for notational convenience we do not write explicitly the dependence of ϕ on y .

Log-concave models (1) are ubiquitous in modern imaging sciences. For instance, many imaging methods to solve linear imaging inverse problems involving additive Gaussian noise use models of the form $\phi(x) = \|y - Ax\|^2/2\sigma^2 + \psi(Bx) + \mathbf{1}_{\mathcal{S}}(x)$ for some linear operators A and B , convex regularizer ψ , and convex set \mathcal{S} (see [1, 6, 5] for examples related to image deblurring, inpainting, compressive sensing, superresolution, and tomographic reconstruction, with total-variation and wavelet priors). Similar log-concave Bayesian models can be considered for problems involving other observation noise models, such as Poisson noise [26], and for phase retrieval problems [11]. Log-concave models (1) are also used extensively in other areas of data science such as machine learning [41].

Because drawing conclusions directly from $p(x|y)$ is difficult, Bayesian imaging methods deliver summaries of $p(x|y)$ —namely Bayes point estimators—that summarize $p(x|y)$ optimally in a single value \hat{x} . This estimated value is optimal in the following decision-theoretic sense [38].

Definition 1. Let $L : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}_0^+$ be a loss function that quantifies the difference between two points in \mathbb{R}^n . A Bayes estimator associated with L is any estimator that minimizes the posterior expected loss, i.e.,

$$\hat{x}_L = \operatorname{argmin}_{u \in \mathbb{R}^n} E_{x|y}[L(u, x)].$$

Recall that the posterior expectation $E_{x|y}[L(u, x)] \triangleq \int_{\mathbb{R}^n} L(u, x)p(x|y)dx$. Sensible loss functions L usually verify the following three desiderata [38]:

- $L(u, x) \geq 0 \forall u, x \in \mathbb{R}^n$,
- $L(u, x) = 0 \iff u = x$,
- L is strictly convex w.r.t. its first argument (to guarantee estimator uniqueness).

Estimator uniqueness is important because it implies admissibility (i.e., Bayesian estimator \hat{x}_L is not dominated by any other estimator) [38]. Observe that L is not necessarily symmetric; that is, $L(u, x) \neq L(x, u)$. We do not enforce symmetry because the arguments of L have clearly different roles in the decision problem.

In a purely theoretical Bayesian exercise, L should be chosen based on specific aspects of the problem and application considered. This is particularly important in imaging problems that are ill-posed or ill-conditioned, as the choice of L can significantly impact estimation

results. However, specifying a bespoke loss function for high dimensional problems is not easy. Consequently, most methods in the imaging literature use default losses and estimators.

In particular, Bayesian imaging methods have traditionally used the minimum mean squared error (MMSE) estimator, given by the posterior mean $\hat{x}_{\text{MMSE}} = \int_{\mathbb{R}^n} p(x|y) x dx$. This estimator is widely regarded as a gold standard in the field, in part because of its good empirical performance and favorable theoretical properties, and also perhaps in part because of cultural heritage. From Bayesian decision theory, MMSE estimation is optimal with respect to the entire class of Euclidean or Mahalanobis squared distances, given by quadratic loss functions of the form $L(u, x) = (u - x)^\top Q(u - x)$ with $Q \in \mathcal{S}_{++}^n$ (i.e., the set of $n \times n$ positive definite matrices), which includes the popular mean square loss $L(u, x) = \|u - x\|_2^2$ when $Q = I_n$ [38]. This gives \hat{x}_{MMSE} a straightforward geometric interpretation. Moreover, MSE estimation is optimal w.r.t. a more general class of functions [7] that provides a second order approximation to any strongly convex loss function; hence \hat{x}_{MMSE} can act as a universal proxy for other Bayesian estimators in this sense.

Unfortunately, calculating \hat{x}_{MMSE} in high dimensional models can be very difficult because it requires solving integrals that are often too computationally expensive. This has stimulated much research on the topic, from fast Monte Carlo simulation methods to efficient approximations with deterministic algorithms [37, 25]. But with ever increasingly large problems and datasets, many imaging methods have focused on alternatives to MMSE estimation.

In particular, modern methods rely strongly on maximum-a-posteriori (MAP) estimation,

$$\begin{aligned} \hat{x}_{\text{MAP}} &= \operatorname{argmax}_{x \in \mathbb{R}^n} p(x|y), \\ &= \operatorname{argmin}_{x \in \mathbb{R}^n} \phi(x), \end{aligned}$$

whose calculation is a convex problem that can often be solved very efficiently, even in very high dimensions (e.g., $n > 10^6$), by using proximal convex optimization techniques [19, 13, 21, 27]. Modern nonstatistical imaging methods also predominately solve problems by convex optimization, and their solutions are often equivalent to performing MAP estimation for some implicit Bayesian model. The precise sense in which these solutions are equivalent to MAP estimators is an interesting discussion topic that is beyond the scope of this paper.

Following a decade of intensive activity, there is now abundant evidence that MAP estimation delivers accurate results for a wide range of imaging problems. However, from a theoretical viewpoint MAP estimation is not well understood. Currently the predominant view is that MAP estimation is not formal Bayesian estimation in the decision-theoretic sense postulated by Definition 1 because it does not generally minimize a known expected loss. The prevailing interpretation is that MAP estimation is in fact an approximation arising from the degenerate loss $L_\epsilon(u, x) = \mathbf{1}_{\|x-u\| < \epsilon}$ with $\epsilon \rightarrow 0$ [38] (this derivation holds for all log-concave models but is not generally true [8]). However, this asymptotic derivation does not lead to a proper Bayesian estimator. More importantly, the resulting loss is very difficult to motivate for inference problems in continuous domains such as \mathbb{R}^n and does not help explain the good empirical performance reported in the literature.

Furthermore, most other theoretical results for MAP estimation only hold for very specific models, or have been derived by adopting analyses that are extrinsic to the Bayesian decision

theory framework (e.g., by analyzing MAP estimation as constrained or regularized least-squares regression; see, for example, [18, 17]). As a trivial example of results that only hold for specific models, when $p(x|y)$ is symmetric we have $\hat{x}_{\text{MAP}} = \hat{x}_{\text{MMSE}}$, and thus MAP estimation inherits the favorable properties of MMSE estimation. This result has been partially extended to some denoising models of the form $p(x|y) \propto \exp\{\|y - x\|^2/2\sigma^2 + \lambda h(x)\}$ in [28], where it is shown that MAP estimation coincides with MMSE estimation with a different model $\tilde{p}(x|y) \propto \exp\{-\|y - x\|^2/2\sigma^2 - \tilde{\lambda}\tilde{h}(x)\}$ involving a different prior distribution. It follows that for these models MAP estimation is decision-theoretic Bayesian estimation w.r.t. the weighted loss $L(u, x) = \|u - x\| \exp\{\tilde{\lambda}\tilde{h}(x) - \lambda h(x)\}$. This is of course a post-hoc loss, but the result is interesting in that it highlights that a single estimator may have a plurality of origins. More importantly, it raises the question of whether MAP estimation is merely a computational proxy for MMSE estimation, which unlike \hat{x}_{MAP} has an appealing theoretical underpinning. This question was recently answered in Burger and Lucka [15]: MAP estimation is proper decision-theoretic Bayesian estimation for all models of the form $p(x|y) \propto \exp\{-\|y - Ax\|_{\Sigma^{-1}}^2/2 - \lambda h(x)\}$, with known linear operator A and noise covariance Σ , and where h is convex and Lipschitz continuous. More precisely, that paper shows that for this class of models MAP estimation is optimal w.r.t. the loss $L(u, x) = \|A(u - x)\|_{\Sigma^{-1}}^2 + 2\lambda D_h(u, x)$, where $D_h(x) = h(u) - h(x) - \nabla h(x)^\top (u - x)$ is the h -Bregman divergence [9]. The paper also shows that \hat{x}_{MAP} outperforms \hat{x}_{MMSE} w.r.t. the expected Bregman error $E_{x|y}\{D_h(u, x)\}$, an error measure that grasps some distinctive features of x (e.g., sparsity, regularity, smoothness, etc.). It may appear that the loss identified in [15] is rather artificial and difficult to analyze and motivate; however, the new results presented in section 3 show that it is a specific instance of a more general loss that stems directly from the consideration of the geometry of the Bayesian model.

It is also worth mentioning that several recent works have studied MAP estimation in infinite-dimensional settings, which is important for our understanding of how the technique behaves in increasingly large problems. An important advance in this area is the connection of the topological description of the MAP estimate to a variational problem, developed in [23] for nonlinear inverse problems in a Gaussian framework, and subsequently extended to non-Gaussian settings in [31]. Agapiou et al. made another key contribution in this area by studying infinite-dimensional MAP estimation with Besov priors, which are very relevant to imaging sciences because they promote sparsity and preservation of edges [2]. We also mention the recent work [33], which further improves our understanding of modes in infinite dimensions.

In order to better understand MAP estimation, in this paper we first revisit the choice of the loss function for Bayesian point estimation in the context of models that are log-concave, where MAP is a convex problem (we limit our analysis to finite-dimensional problems). A main novelty of our analysis is that, instead of specifying the loss directly, we use differential geometry to derive the loss from the geometry of the model. Precisely, we show that under some regularity assumptions, the log-concavity of $p(x|y)$ induces a specific Riemannian differential geometry on the parameter space, and that taking into account this space geometry naturally leads to an intrinsic or canonical loss function to perform Bayesian point estimation in that space. Following on from this, we establish that the canonical loss for the parameter space is the Bregman divergence associated with $\phi(x) = -\log p(x|y)$, and that the Bayesian

estimator w.r.t. this loss is the MAP estimator. We then show that the MMSE estimator is the Bayesian estimator associated with the dual canonical loss, and we propose universal estimation performance guarantees for MAP and MMSE estimation in log-concave models. We conclude by illustrating our theory with an application to linear inverse problems with sparsity-promoting wavelet priors, and an analysis of the regularization-by-denoising models proposed recently in [39].

The remainder of the paper is organized as follows: section 2 introduces the elements of differential geometry that are essential to our analysis. In section 3 we present our main theoretical results: a decision-theoretic and differential-geometric derivation of MAP and MMSE estimation, with universal bounds on the estimation error involved. Section 5 discusses the impact of relaxing the regularity assumptions adopted in section 3. Conclusions are finally reported in section 6. Proofs are presented in the appendix.

2. Riemannian geometry and the canonical divergence function. In this section we recall some elements of differential geometry that are necessary for our analysis. For a detailed introduction to this topic we refer the reader to [3].

An n -dimensional Riemannian manifold (\mathbb{R}^n, g) , with metric $g : \mathbb{R}^n \rightarrow \mathcal{S}_{++}^n$ and global coordinate system x , is a vector space that behaves locally as a Euclidean space.¹ More precisely, at any point $x \in \mathbb{R}^n$, we have a tangent space $\mathcal{T}_x\mathbb{R}^n$ with inner product $\langle u, x \rangle = u^\top g(x)x$ and norm $\|x\| = \sqrt{x^\top g(x)x}$. This tangent space describes how the manifold (\mathbb{R}^n, g) behaves locally at x . The geometry is local and may vary smoothly from $\mathcal{T}_x\mathbb{R}^n$ to neighboring tangent spaces (i.e., the inner product and norm used are local properties and vary spatially). The variations are encoded in the affine connection Γ , with coefficients given by $\Gamma_{ij,k}(x) = \partial_k g_{i,j}(x)$ describing the spatial evolution of the metric g .

A crucial property of (\mathbb{R}^n, g) is that, similarly to Euclidean spaces, manifolds support divergence functions.

Definition 2. A function $D : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ is a divergence function on \mathbb{R}^n if the following conditions hold for any $u, x \in \mathbb{R}^n$:

- $D(u, x) \geq 0 \forall u, x \in \mathbb{R}^n$,
- $D(u, x) = 0 \iff x = u$,
- $D(u, x)$ is strongly convex w.r.t. u , and \mathcal{C}^2 w.r.t. u and x .

Observe that the class of divergence functions is equivalent to that of loss functions for Bayesian point estimation specified in section 1, with some mild additional regularity conditions. This suggests that divergence functions are sensible losses to define estimators. Moreover, divergence functions also provide a link to the differential geometry of the space, which allows relating space geometry and Bayesian decision theory. This relationship has been used previously to analyze Bayesian decision problems from a Riemannian geometric viewpoint, leading to the so-called decision geometry framework [24]. Here we adopt an opposite perspective: we start by considering a Riemannian manifold (\mathbb{R}^n, g) and then use the relationship to identify the divergence functions that arise naturally in that space. In particular, we focus on the so-called canonical divergence on (\mathbb{R}^n, g) , which generalizes the Euclidean squared distance to this kind of manifold [4].

¹Recall that \mathcal{S}_{++}^n is the set of $n \times n$ positive definite matrices.

Definition 3 (Canonical divergence [4]). For any two points $u, x \in \mathbb{R}^n$, the (\mathbb{R}^n, g) -canonical divergence is given by

$$(2) \quad D(u, x) = \int_0^1 t \dot{\gamma}_t^\top g(\gamma_t) \dot{\gamma}_t dt,$$

where γ_t is the Γ -geodesic from u to x and $\dot{\gamma}_t = d/dt \gamma_t$.

The reason D is the (\mathbb{R}^n, g) -canonical divergence is that it fully specifies the geometry of (\mathbb{R}^n, g) ; i.e., there is a one-to-one relationship between D and the metric g .

Observe that D has connections to the length of the Γ -geodesic between u and x . Precisely, by noting that the squared length of a curve $\zeta_t : [0, 1] \rightarrow \mathbb{R}^n$ on the manifold (\mathbb{R}^n, g) is given by $\int_0^1 \dot{\zeta}_t^\top g(\zeta_t) \dot{\zeta}_t dt$, we observe that $D(u, x)$ is essentially the squared length of the Γ -geodesic γ_t weighted linearly along the path from u to x . This weighting in (2) is important because it guarantees that $D(u, x)$ is convex in u , a necessary condition to define a divergence function (the weighting also leads to other important properties such as linearity w.r.t. g ; see section 3). Note that the weighting also introduces an asymmetry, i.e., generally $D(u, x) \neq D(x, u)$, which will have deep implications for Bayesian estimation.

Finally, it is easy to check that (2) reduces to the Euclidean squared distance $D(u, x) = \frac{1}{2}(u-x)^\top g(u-x)$ when (\mathbb{R}^n, g) is the Euclidean space with product $\langle u, x \rangle = u^\top gx$.² More generally, D is always consistent with the local Euclidean geometry of the manifold (\mathbb{R}^n, g) . That is, for any point $x + dx$ in the neighborhood of x we have $D(x + dx, x) = \|dx\|^2/2 + o(\|dx\|^2)$, where $\|\cdot\|$ is the Euclidean norm of the tangent space $\mathcal{T}_x \mathbb{R}^n$ (a higher order approximation of $D(x + dx, x)$ is also possible by using the connection Γ [3]). And because D is the canonical divergence, if we use the decision geometry framework [24] to derive the Riemannian geometry induced by D on \mathbb{R}^n , we obtain

$$g_{i,j}^{(D)}(x) \triangleq \partial_i \partial_j D(x, x) = g_{i,j}(x),$$

$$\Gamma_{i,j,k}^{(D)}(x) \triangleq \partial_i \partial_j \partial'_k D(x, x) = \Gamma_{i,j,k}(x)$$

(here ∂ and ∂' denote differentiation w.r.t. the first and second components of D , respectively), indicating that D fully specifies the geometry of (\mathbb{R}^n, g) , and vice versa.

3. A differential-geometric derivation of MAP and MMSE estimation.

3.1. Canonical Bayesian estimation: From differential geometry to decision theory.

In this section we use differential geometry to relate the geometry of $p(x|y)$ to the loss functions used for Bayesian estimation of x . Precisely, we exploit the log-concavity of $p(x|y)$ to induce a Riemannian geometry on the solutions space. This in turn defines a canonical loss for that space and two Bayesian estimators: a primal estimator related to $D(u, x)$ and a dual estimator related to the dual divergence $D_\phi^*(u, x) = D_\phi(x, u)$. We focus on the case where $p(x|y)$ is smooth and strongly log-concave, and later we analyze the effect of relaxing these assumptions.

²In the Euclidean case we have that case g is constant, and the Γ -geodesic is $\gamma_t = u + t(x - u)$, so $D(u, x) = \int_0^1 t(u-x)^\top g(u-x) dt = \int_0^1 t dt (x-u)^\top g(x-u) = \frac{1}{2}(x-u)^\top g(x-u)$.

Theorem 4 (Canonical Bayesian estimators). *Suppose that $\phi(x) = -\log p(x|y)$ is strongly convex, continuous, and \mathcal{C}^3 on \mathbb{R}^n . Let (\mathbb{R}^n, g) denote the Riemannian manifold induced by ϕ , with metric coefficients $g_{i,j}(x) = \partial_i \partial_j \phi(x)$. Then, the canonical divergence on (\mathbb{R}^n, g) is the ϕ -Bregman divergence, i.e.,*

$$D_\phi(u, x) = \phi(u) - \phi(x) - \nabla \phi(x)(u - x).$$

In addition, the Bayesian estimator associated with $D_\phi(u, x)$ is unique and is given by the MAP estimator

$$\begin{aligned} \hat{x}_{D_\phi} &\triangleq \operatorname{argmin}_{u \in \mathbb{R}^n} \mathbb{E}_{x|y}[D_\phi(u, x)] \\ &= \operatorname{argmin}_{x \in \mathbb{R}^n} \phi(x) \\ &= \hat{x}_{\text{MAP}}. \end{aligned}$$

The Bayesian estimator associated with the dual canonical divergence $D_\phi^(u, x) = D_\phi(x, u)$ is also unique and is given by the MMSE estimator*

$$\begin{aligned} \hat{x}_{D_\phi^*} &\triangleq \operatorname{argmin}_{u \in \mathbb{R}^n} \mathbb{E}_{x|y}[D_\phi^*(u, x)] \\ &= \int_{\mathbb{R}^n} x p(x|y) dx \\ &= \hat{x}_{\text{MMSE}}. \end{aligned}$$

The proof is reported in the appendix.

Theorem 4 provides several interesting new insights into MAP and MMSE estimation in log-concave models. First and foremost, MAP estimation stems from Bayesian decision theory, and hence it stands on the same theoretical footing as the core Bayesian methodologies such as MMSE estimation (albeit w.r.t. a different class of loss functions). The MAP loss, $D_\phi(u, x)$, is a generalization of the Euclidean squared distance that arises naturally from the consideration of the geometry of $p(x|y)$. Consequently, the conventional definition of the MAP estimator as the maximizer $\hat{x}_{\text{MAP}} = \operatorname{argmax}_{x \in \mathbb{R}^n} p(x|y)$ is mainly algorithmic for these models and useful to highlight that these estimators take the form of a convex optimization problem. (Of course, this is a key computational advantage over other Bayesian point estimators because it allows computation of \hat{x}_{MAP} by using modern proximal convex optimization algorithms that scale very efficiently to high dimensions; see, e.g., [19] for details.) Moreover, Theorem 4 also reveals a surprising form of duality between MAP and MMSE estimation, with the two estimators intimately related to each other by the (asymmetry of the) canonical divergence function that $p(x|y)$ induced on the solutions space. Note that Gaussian models are special because (\mathbb{R}^n, g) is Euclidean in that case, which is a self-dual space; consequently $D_\phi(u, x) = D_\phi(x, u) = \frac{1}{2} \|u - x\|_{\Sigma^{-1}}^2$ and the primal and dual canonical estimators coincide as a result. Finally, Theorem 4 also shows that, under log-concavity and smoothness assumptions, the posterior mode is a global property of $p(x|y)$, similarly to the posterior mean.

The way in which the Bregman divergence $D_\phi(u, x)$ measures the similarity between u and x is directly related to the geometry of $p(x|y)$. Precisely, because $\phi(x) = -\log p(x|y)$ is

strongly convex, then $\phi(u) > \phi(x) - \nabla\phi(x)(u - x)$ for any $u \neq x$. The divergence $D_\phi(u, x)$ essentially quantifies this gap, which, as mentioned previously, is directly related to the length of the affine geodesic from u to x (and hence not only to the relative position of u and x but also to the space geometry induced by $p(x|y)$). Moreover, this geometry can depend on the value of y ; however, for the important class of models of the $p(x|y) \propto \exp\{-\|y - Ax\|_{\Sigma^{-1}}^2/2 - \lambda h(x)\}$ the geometry is completely specified by Σ and λh independently of y . Furthermore, observe that because D_ϕ is linear w.r.t. ϕ , for any decomposition $\phi = \alpha\phi_1 + \beta\phi_2$ based on two convex functions ϕ_1 and ϕ_2 and $\alpha, \beta \in \mathbb{R}$, we obtain $D_\phi = \alpha D_{\phi_1} + \beta D_{\phi_2}$. It follows that for the specific case of Gaussian linear observation models, the canonical divergence D_ϕ is equivalent to the specific loss identified in [15].

We also mention at this point that for Gaussian denoising models, i.e., $\phi(x) = \|y - x\|_2^2/2\sigma^2 + \psi(x)$, the estimator $\hat{x}_{D_\phi} = \hat{x}_{\text{MAP}}$ results from the computation of the proximal operator $\text{prox}_{\sigma^2\psi}(y) = \text{argmin}_{x \in \mathbb{R}^n} \|y - x\|_2^2/2\sigma^2 + \psi(x)$ [19]. This is equivalent to a gradient step on the Moreau–Yoshida regularization of ψ ; i.e., $\hat{x}_{\text{MAP}} = y + \sigma^2 \nabla \tilde{\phi}(y)$, with $\tilde{\phi}(y) = \inf_{x \in \mathbb{R}^n} \|y - x\|_2^2/2\sigma^2 + \psi(x)$. Likewise, $\hat{x}_{D_\phi^*} = \hat{x}_{\text{MMSE}}$ can be expressed as the gradient step $\hat{x}_{\text{MMSE}} = y + \sigma^2 \nabla \bar{\phi}(y)$, where $\bar{\phi} = \log \int \exp\{-\|y - x\|_2^2/2\sigma^2 - \psi(x)\} dx$ is a different smooth approximation of ψ (please see [36] for details).

Also note that a different Bregman divergence, namely the KL divergence $KL(u, x) = \int \log \left[\frac{p(y|x)}{p(y|u)} \right] p(y|x) dy$, is often used in Bayesian point estimation to define an estimator that is independent of the parametrization of the likelihood [38]. This estimator is particularly relevant when the object of interest is $p(y|x)$, as opposed to the value of x itself, for example in prediction problems. This estimator is not often used in imaging sciences.

Finally, we notice that because $D_\phi(u, x)$ is derived from $p(x|y)$ it may depend on the value of y , which is controversial in some lines of Bayesian thinking because it implies that the decision problem underpinning the estimator is defined a posteriori. This can happen, for example, in problems involving non-Gaussian observation models. Our view on this matter is that although decision problems are generally defined a priori, the case of Bayesian estimators is special because the decision involved is precisely how to summarize $p(x|y)$, and this decision can be considered a posteriori if this allows delivery of an estimator with favorable accuracy or computational properties. Of course, loss functions that do not depend on the model considered also have advantages, namely the mean squared error loss that also leads to an estimator with good properties (albeit often very expensive to compute). In any case, it is fundamental that one understands how the estimator that one uses summarizes $p(x|y)$, and the aim of this work is to improve our understanding of the widely used MAP estimator.

3.2. Error bounds for MAP and MMSE estimation. Theorem 4 establishes that under certain conditions \hat{x}_{MAP} is a proper Bayesian estimator. Following on from this, it is natural to study the accuracy of \hat{x}_{MAP} as a Bayesian estimator. The Bayesian approach to this question is to infer the accuracy of \hat{x}_{MAP} according to the posterior distribution $p(x|y)$. For \hat{x}_{MMSE} this generally corresponds to computing the expected mean squared error loss, related to the posterior covariance. This type of analysis can be useful, for example, for identifying high dimensional stability conditions (i.e., conditions under which the error grows linearly with $n = \dim(x)$).

Here we perform this type of analysis for \hat{x}_{MAP} w.r.t. the canonical loss. Precisely, we establish universal estimation error bounds w.r.t. the dual error function $D_\phi^*(u, x)$ for MAP and MMSE estimation, for which we have the following result.

Proposition 5 (Expected error bound). *Suppose that $\phi(x) = -\log p(x|y)$ is convex on \mathbb{R}^n and $\phi \in \mathcal{C}^1$. Then,*

$$\mathbb{E}_{x|y} [D_\phi^*(\hat{x}_{\text{MMSE}}, x)] \leq \mathbb{E}_{x|y} [D_\phi^*(\hat{x}_{\text{MAP}}, x)] \leq n.$$

Proof. The proof is reported in the appendix. ■

We read Proposition 5 as a high dimensional stability result for MAP and MMSE estimation, stating that the expected estimation error, as measured by the dual loss D_ϕ^* , grows at most linearly with the number of image pixels. Therefore, even if the likelihood $p(y|x)$ is poorly identifiable because $\dim(y) \ll \dim(x)$, or because the linear operator A is very rank deficient, or because y is corrupted by Poison noise, in smooth log-concave settings the expected error cannot grow polynomially or with a linear constant greater than 1.

To formally study the expected error as n increases we consider a generic log-concave stochastic process $\mathbb{X} = \{x^{(n)}, n \in \mathbb{N}\}$, where, for each $n \in \mathbb{N}$, the random vector $x^{(n)} = (x_1, \dots, x_n) \in \mathbb{R}^n$ has marginal distribution $p_n(x^{(n)}|y) \propto \exp\{-\phi_n(x^{(n)})\}$ for some convex function $\phi_n : \mathbb{R}^n \rightarrow (-\infty, \infty]$. We also assume that the entropy rate of \mathbb{X} is finite; i.e., $\lim_{n \rightarrow \infty} \mathbb{E}_{x^{(n+1)}|y}[\phi_{n+1}(x^{(n+1)})] - \mathbb{E}_{x^{(n)}|y}[\phi_n(x^{(n)})] < \infty$ [22]. This limit captures the asymptotic information gain per pixel and characterizes global statistical features of the image, particularly correlations at any rage. In log-concave settings, this condition holds, for example, when $\lim_{n \rightarrow \infty} \phi_n(\hat{x}_{\text{MAP}}^{(n)})/n < \infty$; it also holds when \mathbb{X} is strongly stationary [12]. By analyzing Proposition 5 in this setting we see that

$$\mathbb{E}_{x^{(n)}|y} [D_{\phi_n}^*(\hat{x}_{\text{MMSE}}^{(n)}, x^{(n)})] \leq \mathbb{E}_{x^{(n)}|y} [D_{\phi_n}^*(\hat{x}_{\text{MAP}}^{(n)}, x^{(n)})] \leq n.$$

Then, because the entropic rate of \mathbb{X} is finite, $\lim_{n \rightarrow \infty} \mathbb{E}_{x^{(n)}|y}[\phi_n(x^{(n)})]/n < \infty$ [22], and hence the dimension-normalized expected errors verify

$$\lim_{n \rightarrow \infty} \mathbb{E}_{x^{(n)}|y} [D_{\phi_n}^*(\hat{x}_{\text{MMSE}}^{(n)}, x^{(n)})/n] \leq \lim_{n \rightarrow \infty} \mathbb{E}_{x^{(n)}|y} [D_{\phi_n}^*(\hat{x}_{\text{MAP}}^{(n)}, x^{(n)})/n] \leq 1.$$

We emphasize again that this form of dimension stability is not generally available in estimation problems, and is a direct consequence of the log-concavity of the model and its relationship with the MAP and MMSE estimators. Finally, observe that the above error bounds are tight; e.g., the trivial independent and identically distributed (i.i.d.) process $x_i|y \sim \text{Exp}(\lambda_y)$, for $i \geq 1$, $\lambda_y \in \mathbb{R}^+$, attains the bound. We conjecture that this bound can be improved for specific subclasses of log-concave models by using entropy rate results from the probability literature; future work will investigate this.

3.3. Connections to other works. We conclude this section by discussing some connections between this paper and other theoretical works related to MAP estimation. As previously explained, Theorem 4 directly builds on [15], which considered the class of log-concave models $p(x|y) \propto \exp\{-\|y - Ax\|_{\Sigma^{-1}}^2/2 - \lambda h(x)\}$ with Gaussian likelihood $y \sim \mathcal{N}(Ax, \Sigma)$ and prior

$p(x) \propto \exp\{\lambda h(x)\}$, and establishes that in this case \hat{x}_{MAP} is the Bayesian estimator for the Bregman loss $L(u, x) = \|A(u - x)\|_{\Sigma^{-1}}^2 + 2\lambda D_h(u, x)$. Theorem 4 generalizes this result to a larger class of posterior distributions and provides motivation for this unconventional loss function by showing that it stems directly from the consideration of the geometry of the parameter space. Proposition 5 provides further motivation for this loss by establishing explicit bounds on the expected estimation error.

It is worth mentioning that the generalization of [15] to other log-concave models was developed simultaneously and independently in Burger, Dong, and Sciacchitano [14] (see [14, Theorem 3.2] for MAP estimation and [14, Theorem 4.3] for MMSE estimation, which is also closely related to [7, Proposition 1]). Moreover, that work also analyzes the expected estimation error involved in MAP and MMSE estimation but w.r.t. other divergence functions. More precisely, Burger, Dong, and Sciacchitano [14] use the Bregman divergence D_h related to the regularizer or negative log-prior, whereas we use the canonical Bregman divergence D_ϕ related to the negative log-posterior. As mentioned previously, D_h grasps important features of x (e.g., sparsity, regularity, smoothness) and is always independent of the observed data y , whereas D_ϕ is independent of y only in special cases (e.g., Gaussian linear observation models).

Moreover, Burger, Dong, and Sciacchitano also show that $\mathbb{E}_{x|y} [D_h^*(\hat{x}_{\text{MMSE}}, x)] \leq \mathbb{E}_{x|y} [D_h^*(\hat{x}_{\text{MAP}}, x)]$ and conclude that \hat{x}_{MMSE} outperforms \hat{x}_{MAP} when the estimation error is measured in this way, which is independent of y . To analyze how these expected errors behave as dimensionality increases, we combine this result with Proposition 5 and obtain the following bound.

Corollary 6. *Suppose that $h(x) = -\log p(x)$ is convex on \mathbb{R}^n and $\phi \in \mathcal{C}^1$. Then,*

$$\mathbb{E}_{x|y} [D_h^*(\hat{x}_{\text{MMSE}}, x)] \leq \mathbb{E}_{x|y} [D_h^*(\hat{x}_{\text{MAP}}, x)] \leq n.$$

Proof. The proof follows directly from combining [14, Theorem 5.1] with Proposition 5 and the fact that $D_h^* \leq D_\phi^*$ for any splitting $\phi = h + f$ where h and f are convex functions. This result can also be derived from the integration by parts argument in [14]. ■

Again, we read Corollary 6 as a high dimensional stability result for MAP and MMSE estimation, stating that the expected estimation error, measured in this case by the dual loss D_h^* , grows at most linearly with the number of image pixels. Polynomial growth or faster linear growth is not possible within the class of smooth log-concave models, even if the problem is very ill-conditioned. At the same time, this linear rate cannot be further improved, as any i.i.d. process $p(x|y) \prod_{i=1}^n p(x_i|y)$ will have an error that grows linearly with n .

4. Illustrative examples. As a way of illustrating our theory, we now analyze the geometry of a simple image denoising model in the wavelet domain, and of the regularization-by-denoising (RED) Bayesian models recently proposed in [39].

4.1. Wavelet image denoising model. In this example we analyze the behavior of MAP estimation in linear inverse problems with sparsity-promoting or shrinkage priors. Without loss of generality, we first consider a simple additive noise observation model $y = x + w$, with noise $w \sim \mathcal{N}(0, \sigma^2 I_n)$ with variance $\sigma^2 \in \mathbb{R}^+$, that allows a detailed analysis. To recover x we put a shrinkage prior on a wavelet representation $z = Wx$ of x , where W is some orthogonal

wavelet transform. More precisely, we use the smoothed Laplace prior

$$(3) \quad p(z) \propto \exp \left\{ -\lambda \sum_{i=1}^n \sqrt{z_i^2 + \alpha^2} \right\},$$

where $\lambda \in \mathbb{R}^+$ and $\alpha \in \mathbb{R}^+$ are, respectively, scale and shape regularization parameters; this prior is also known as the pseudo-Huber, Hardy, or Charbonnier prior [10, 20]. The likelihood is $p(y|z) \propto \exp\{-\frac{1}{2\sigma^2}\|y - W^\top z\|_2^2\}$, and hence the posterior for the wavelet coefficients is

$$p(z|y) \propto \exp \left\{ -\frac{1}{2\sigma^2}\|y - W^\top z\|_2^2 - \lambda \sum_{i=1}^n \sqrt{z_i^2 + \alpha^2} \right\}.$$

To check that Theorem 4 and Proposition 5 apply, we note that $\phi(z) = -\frac{1}{2\sigma^2}\|y - W^\top z\|_2^2 - \lambda \sum_{i=1}^n \sqrt{z_i^2 + \alpha^2}$ belongs to \mathcal{C}^∞ and has a diagonal Hessian matrix with elements given by

$$\frac{\partial^2}{\partial z_i^2} \phi(z) = \frac{1}{\sigma^2} + \frac{\lambda \alpha^2}{(\alpha^2 + z_i^2)^{3/2}}.$$

Noticing that the elements $\frac{\partial^2}{\partial z_i^2} \phi(z)$ take values in $[\frac{1}{\sigma^2}, \frac{1}{\sigma^2} + \lambda]$ for all $z \in \mathbb{R}^n$, we conclude that $\phi(z)$ is strongly convex. Notice that, similarly to the previous example, the geometry of the manifold $\{\mathbb{R}^n, g\}$ does not depend on the value of the observation y , and hence the canonical divergences are independent of y , too.

Because of the action of the shrinkage prior (3), the Bayesian model $p(z|y)$ will promote solutions that have some large wavelet coefficients and some coefficients close to zero. This behavior is controlled by the regularization parameter λ and also by the choice of the Bayesian estimator used to summarize $z|y$. In particular, MAP estimation may significantly accentuate shrinkage. This can be theoretically analyzed in different ways, and in particular by using Theorem 4. Accordingly, \hat{z}_{MAP} minimizes the expected canonical divergence on $\{\mathbb{R}^n, g\}$, given by the ϕ -Bregman divergence

$$\begin{aligned} D_\phi(u, z) &= \phi(u) - \phi(z) - \nabla \phi(z)^\top (u - z) \\ &= \frac{1}{2\sigma^2} \|W^\top u - W^\top z\|_2^2 + \lambda \sum_{i=1}^n \left[\sqrt{u_i^2 + \alpha^2} - \sqrt{z_i^2 + \alpha^2} + \frac{z_i^2 - z_i u_i}{\sqrt{z_i^2 + \alpha^2}} \right]. \end{aligned}$$

Because $WW^\top = I_n$ we have that D_ϕ is fully separable, i.e., $D_\phi(u, z) = \sum_{i=1}^n D_\psi(u_i, z_i)$ with

$$D_\psi(u_i, z_i) = \frac{1}{2\sigma^2} (u_i - z_i)^2 + \lambda \frac{\sqrt{z_i^2 + \alpha^2} \sqrt{u_i^2 + \alpha^2} - z_i u_i - \alpha^2}{\sqrt{z_i^2 + \alpha^2}}.$$

Because D_ψ is a divergence it promotes values of u_i that are close to z_i . To develop an intuition for D_ψ it is useful to analyze its behavior when z_i is small and when it is large relative to

α . Observe that D_ϕ has a quadratic term related to the likelihood and a nonquadratic term related to the shrinkage prior. When $z_i \gg \alpha$ the nonquadratic term vanishes, and hence

$$D_\psi(u_i, z_i) \approx \frac{1}{2\sigma^2}(u_i - z_i)^2.$$

As a result, if the observed data is such that the posterior distribution for $z_i|y$ has most of its mass in large values of z_i , the MAP estimate for z_i will essentially coincide with the MMSE estimate given by the posterior mean of $z_i|y$. In this case there is no additional shrinkage from the estimator. Conversely, when $z_i \ll \alpha$ the estimator will significantly boost shrinkage. More precisely, when $z_i \ll \alpha$, the nonquadratic term behaves as

$$\begin{aligned} D_\psi(u_i, z_i) &\approx \frac{1}{2\sigma^2}(u_i - z_i)^2 + \lambda|u_i| \\ &\approx \frac{1}{2\sigma^2}u_i^2 + \lambda|u_i| \end{aligned}$$

for $u_i \gg \alpha$, and for $u_i \ll \alpha$ it behaves as

$$\begin{aligned} D_\psi(u_i, z_i) &\approx \frac{1}{2\sigma^2}(u_i - z_i)^2 + \lambda \left[\frac{u_i^2}{2\alpha} + \frac{z_i^2}{2\alpha} - \frac{-z_i u_i}{\alpha} \right] = \left(\frac{1}{2\sigma^2} + \frac{\lambda}{2\alpha} \right) (u_i - z_i)_2^2 \\ &\approx \left(\frac{1}{2\sigma^2} + \frac{\lambda}{2\alpha} \right) u_i^2. \end{aligned}$$

In these two cases D_ψ strongly promotes u_i values that are close to zero, either explicitly via the shrinkage term $\lambda|u_i|$, or by amplifying the convexity constant of the quadratic loss from $1/2\sigma^2$ to $1/2\sigma^2 + \lambda/2\alpha$. As a result, if the posterior distribution for $z_i|y$ has mass in small values of z_i , then the MAP estimate will intensify the shrinkage effect of the prior. This additional shrinkage is not observed with other loss functions, e.g., MMSE, and is consistent with the empirical observation that MAP estimation performs well with shrinkage priors.

For illustration, Figure 1 shows an experiment with the House image of size 256×256 pixels. Figure 1(a) shows a corrupted observation $y = x + w$ with noise $w \sim \mathcal{N}(0, \sigma)$ with $\sigma = 0.08$, which has a signal-to-noise ratio of 17.1dB. The restored imaged obtained by MAP estimation is displayed in Figure 1(b); this estimate has a signal-to-noise ratio of 22.3dB (we used a Haar wavelet decomposition with four levels and $\lambda = 12$ and $\alpha = 0.01$ for all scales except the coarse scale, for which we used a Jeffreys prior $p(z_i) \propto 1$ to avoid excessively biasing the estimates). For comparison, we also report \hat{x}_{MMSE} , which in this experiment performs poorly (signal-to-noise ratio of 17.7dB). The MAP estimator obtained with a conventional Laplace or ℓ_1 prior (i.e., with $\alpha \rightarrow 0$) has a worse signal-to-noise ratio (21.4dB, not displayed).

Because $p(z|y)$ is fully separable, i.e., $p(z|y) \prod_{i=1}^n p(z_i|y)$, and thus $D_\psi = \sum_{i=1}^n D_\psi(u_i, z_i)$, the action of this estimator can be clearly visualized by plotting the estimation function that performs the denoising of the wavelet coefficients: for MAP estimation this is given by $\hat{z}_{i,MAP}(y) : y \rightarrow \operatorname{argmin}_{u_i} \mathbb{E}_{z_i} [D_\psi(u_i, z_i)|w_i^\top y]$, where $w_i^\top y$ is the i th wavelet coefficient of y ; and for MMSE estimation it is the marginal posterior mean $\hat{z}_{i,MMSE}(y) : y \rightarrow \int z_i p(z_i|y) dz_i$. These functions, displayed in Figure 1(d), clearly show the importance of the choice of the loss used to summarize $p(z|y)$.

We emphasize at this point that this experiment has been selected to highlight the additional shrinkage obtained by using MAP estimation instead of MMSE estimation. However, there are other models where, because of the likelihood and the choice of the wavelet

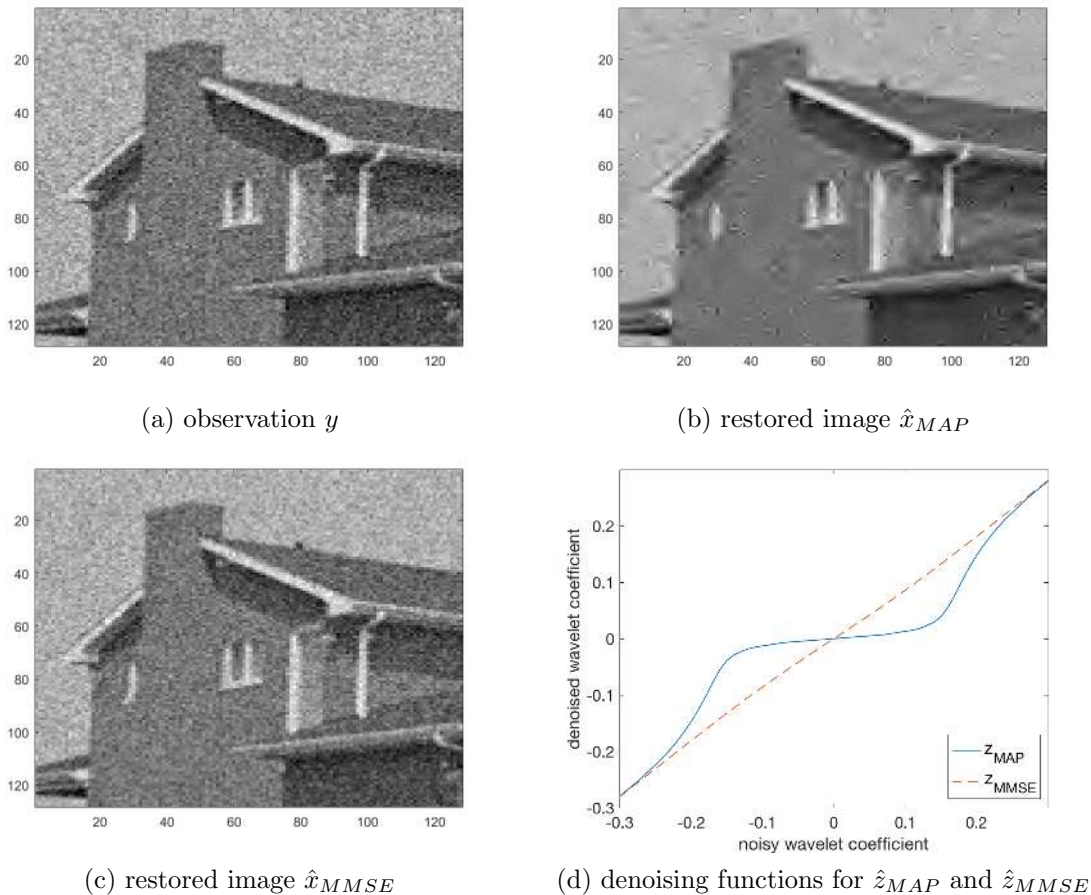


Figure 1. Wavelet denoising experiment with the House image with the smoothed Laplace prior (3).

representation and the parameters used, shrinkage arises mainly from the prior. In that case MAP and MMSE estimation perform equally well. To illustrate this point, Figure 2 below shows the reconstruction results obtained in [16] with MAP and MMSE estimation for a radio-interferometric imaging problem with a very similar model of the form $p(z|y) \propto \exp\{-\frac{1}{2\sigma^2}\|y - AW^T z\|_2^2 - \lambda \sum_{i=1}^n |z_i|\}$, where the difference is that the likelihood term involves a linear operator A modeling the radio-telescope system (see [16] for more details about the model and the algorithms used to compute the estimates). Observe that in this case both MAP and MMSE estimation deliver excellent and remarkably similar results. A similar empirical observation is reported in [15] for a sparse tomography experiment using the Besov wavelet model of [30], which is closely related to the model considered here.

To conclude, shrinkage priors promote solutions that are sparse or approximately sparse through two mechanisms: directly through the definition of the Bayesian model $p(z|y)$, and indirectly through the loss function used to summarize $z|y$. In the case of MAP estimation, this loss function is a Bregman divergence that can significantly amplify shrinkage. In some cases this may lead to better estimation performance. Therefore, when designing Bayesian

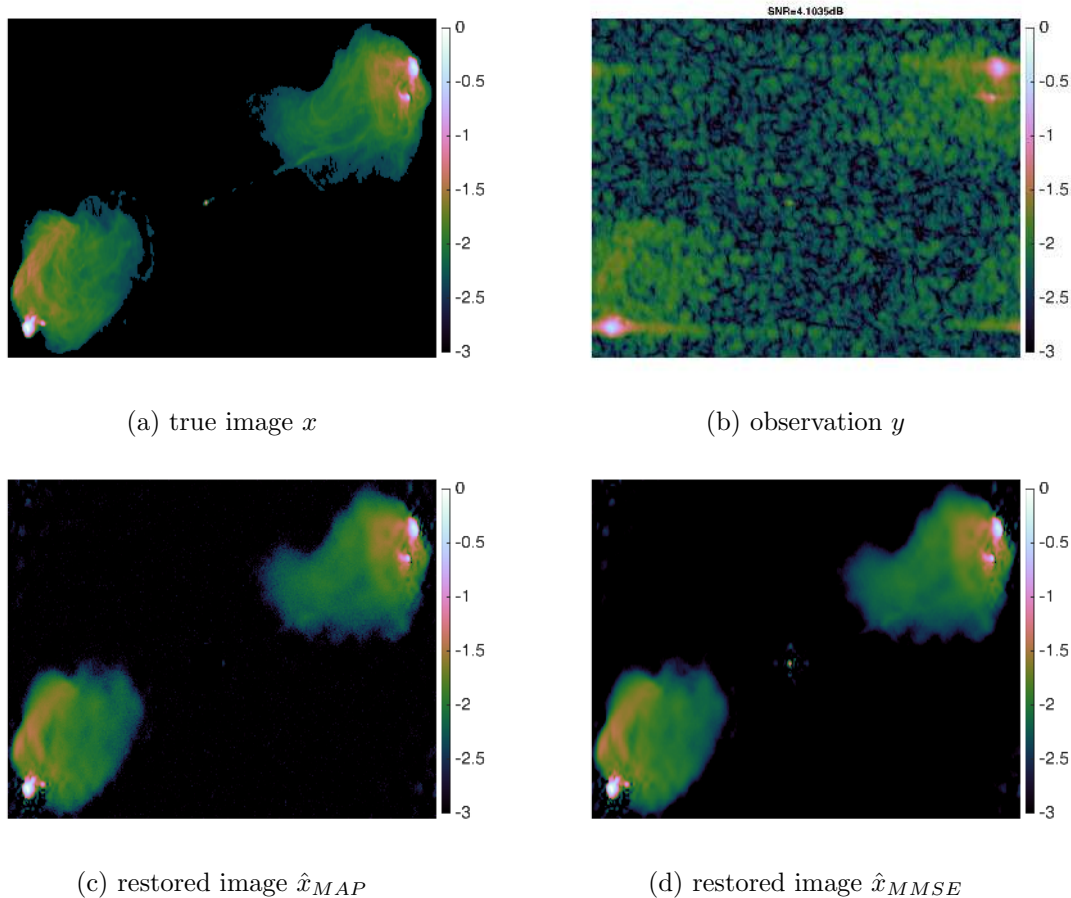


Figure 2. Bayesian radio-interferometric imaging experiment with the *Cygnus A* radio galaxy (size 256×512 pixels). See [16] for more details.

procedures for imaging problems, it is important to carefully consider both the model and the Bayesian estimator used to summarize the information provided by the model.

4.2. Regularization-by-denoising (RED) Bayesian models. As a second illustrative example we analyze the geometry of the RED Bayesian models [39]. In this class of models the prior $p(x)$ is defined implicitly through an image denoising algorithm. Precisely, starting from some image denoising filter $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$, we posit the prior

$$(4) \quad p(x) \propto \exp\left\{-\frac{\lambda}{2} x^\top [x - f(x)]\right\},$$

which promotes values of x that are approximately invariant to filtering by f (i.e., for which $f(x) \approx x$), the rationale being that these are values that f considers to be realistic images. Note that this approach has close connections to plug-and-play priors, which are also defined through denoising algorithms [40].

The RED framework [39] assumes that f verifies three regularity conditions that are necessary to make inference with (4) analytically and computationally tractable. First, f is

smooth—at least \mathcal{C}^2 . Second, the Jacobian matrix $J_f(x)$ is symmetric and has all its eigenvalues in $[-1, 1]$ for all $x \in \mathbb{R}^n$. Third, f is locally homogeneous, i.e., $\lim_{\epsilon \rightarrow 0} \epsilon^{-1}[f(x + \epsilon x) - f(x)] = f(x) \forall x \in \mathbb{R}^n$; this property implies that $f(x) = J_f(x)x$. Under these assumptions, it is possible to express (4) in the following pseudo-quadratic form:

$$(5) \quad p(x) \propto \exp\{-\frac{\lambda}{2}x^\top \Lambda_f(x)x\},$$

where $\Lambda_f(x) = I_n - J_f(x)$ plays the role of an image-adapted graph Laplacian operator, highlighting the connection between f and the model geometry [39].

Notice that these regularity assumptions imply the regularizer or negative log-prior $h(x) = -\log p(x)$ is \mathcal{C}^3 and convex, which is important, for example, for the efficient computation of x_{MAP} by optimization. To show that $h(x) \in \mathcal{C}^3$ we first use the fact that $\nabla h(x) = \nabla \log p(x) = -\lambda[x - f(x)/2 - J_f(x)^\top x/2] = -\lambda[x - f(x)]$, where we have used the symmetry $J_f(x)^\top = J_f(x)$ and the identity $f(x) = J_f(x)x$ related to the local homogeneity of f . Therefore, the Hessian matrix of $h(x)$ has elements given by $\partial_i \partial_j h(x) = -\Lambda_f(x)_{i,j} = J_f(x) - I_n$, which are also continuously differentiable because $f \in \mathcal{C}^2$, and hence $h(x) \in \mathcal{C}^3$. Moreover, the Hessian matrix of $h(x)$, given by $-\Lambda_f(x) = J_f(x) - I_n$, is negative-semidefinite because $J_f(x)$ has all its eigenvalues in $[-1, 1] \forall x \in \mathbb{R}^n$, and consequently $h(x)$ is convex. As a result, if the negative log-likelihood $-\log p(y|x)$ is \mathcal{C}^3 and convex w.r.t. x , then $\phi(x) = -\log p(x|y)$ is also \mathcal{C}^3 and log-concave, and Theorem 4 and Proposition 5 apply.

As an illustrative example, consider linear inverse problems of the form $y = Ax + w$, where A is a known linear operator, and $w \sim \mathcal{N}(0, \sigma^2 I_n)$ with noise variance $\sigma^2 \in \mathbb{R}^+$. The resulting RED Bayesian model has posterior density³

$$(6) \quad p(x|y) \propto \exp\{-\frac{1}{2\sigma^2}\|y - Ax\|_2^2 - \frac{\lambda}{2}x^\top \Lambda(x)x\}.$$

This distribution is strongly log-concave and \mathcal{C}^3 , and hence Theorem 4 and Proposition 5 apply. More precisely, we have $\phi(x) = \frac{1}{2\sigma^2}\|y - Ax\|_2^2 + \frac{\lambda}{2}x^\top \Lambda(x)x$, which induces the metric

$$(7) \quad g_{i,j}(x) = \partial_i \partial_j \phi(x) = \{\sigma^{-2}A^\top A + \lambda \Lambda(x)\}_{i,j}.$$

Observe that (7) combines a Euclidean geometry term $A^\top A$ from the Gaussian likelihood and a non-Euclidean term from the Laplacian $\Lambda(x)$. Again, note that for this class of models the geometry of the manifold $\{\mathbb{R}^n, g\}$ does not depend on the value of the observation y .

Moreover, from Theorem 4, the estimator \hat{x}_{MAP} is the Bayesian estimator associated with the canonical divergence on $\{\mathbb{R}^n, g\}$, given by the ϕ -Bregman divergence

$$\begin{aligned} D_\phi(u, x) &= \phi(u) - \phi(x) - \nabla \phi(x)^\top (u - x) \\ &= \frac{1}{2\sigma^2}\|y - Au\|_2^2 + \frac{\lambda}{2}u^\top \Lambda(u)u - \frac{1}{2\sigma^2}\|y - Ax\|_2^2 + \frac{\lambda}{2}x^\top \Lambda(x)x \\ &\quad - [\sigma^{-2}A^\top (Ax - y) + \lambda x - \lambda f(x)]^\top (u - x) \\ &= \sigma^{-2}[u^\top A^\top Au + x^\top A^\top Ax - 2u^\top A^\top Ax] + \lambda[u^\top \Lambda(u)u + x^\top \Lambda(x)x - 2u^\top \Lambda(x)x] \\ &= \sigma^{-2}D_{A^\top A}(u, x) + \lambda D_\Lambda(u, x), \end{aligned}$$

³Because $\Lambda(x)$ is potentially rank deficient, to guarantee that $p(x|y)$ is a proper probability density function we further assume that $\ker(A^\top A) \cap \ker(\Lambda(x)) = \{0\}$ for all $x \in \mathbb{R}^n$.

where the Mahalanobis (Euclidean) distance

$$\begin{aligned} D_{A^\top A}(u, x) &= \|u - x\|_{A^\top A} \\ &= u^\top A^\top A u + x^\top A^\top A x - 2u^\top A^\top A x \end{aligned}$$

is a measure of prediction of mean squared error related to the Gaussian likelihood, and

$$D_\Lambda(u, x) = u^\top \Lambda(u)u + x^\top \Lambda(x)x - 2u^\top \Lambda(x)x$$

is related to the Laplacian $\Lambda(x)$, which encodes the geometry of the prior (to compute D_Λ we used the local homogeneity property $f(x) = J_f(x)x$ of the denoiser; see [39] for details).

Finally, observe that D_Λ is very similar to the Euclidean loss $D_{A^\top A}$ in that it measures the difference between the squared norms of u and x and the projection of u on x , with the only difference being that for $D_\Lambda(u, x)$ these norms and projections are measured on the tangent spaces $\mathcal{T}_u \mathbb{R}^n$ and $\mathcal{T}_x \mathbb{R}^n$ of the manifold $\{\mathbb{R}^n, \Lambda\}$, with inner products specified by Λ .

5. Relaxation of regularity conditions. We now examine the effect of relaxing the regularity assumptions of Theorem 4. We consider three main cases: lack of smoothness, lack of strong convexity, and lack of continuity.

5.1. Nonsmooth models. The results of Theorem 4 hold for nonsmooth models with the following modifications. First, assume that ϕ is *almost everywhere* \mathcal{C}^3 on \mathbb{R}^n ; i.e., the set of points of \mathbb{R}^n where ϕ is not smooth has dimension $n - 1$ and hence zero Lebesgue measure. To check that ϕ is \mathcal{C}^3 almost everywhere it is necessary to analyze the regularity of the second order derivatives $\partial_i \partial_j \phi(x)$ (e.g., if the second derivatives are Lipschitz continuous, then ϕ is *almost everywhere* \mathcal{C}^3 by Rademacher's theorem [34]). Because in models that are almost everywhere smooth the set of nondifferentiable points has no probability mass, this set can be simply omitted in the computation of expectations. Second, because the nondifferentiable points do not have Euclidean tangent spaces, instead of a global manifold we need to consider the collection local manifolds associated with the regions of \mathbb{R}^n where $p(x|y)$ is \mathcal{C}^3 . Each one of these regions has a local canonical divergence given by the Bregman divergence $D(u, x) = D_\phi(u, x) = \phi(u) - \phi(x) - \nabla \phi(x)^\top (u - x)$. Therefore, for these models we need to posit $D_\phi(u, x)$ as the global loss function for any $(u, x) \in \mathbb{R}^n \times \mathbb{R}^n$ (technically the global loss is the generalized Bregman divergence $D_\phi(u, x) = \phi(u) - \phi(x) - q_x^\top (u - x)$, where q_x belongs to the subdifferential set of ϕ at x [9]; however, the expectation $\mathbb{E}_{x|y}[D_\phi(u, x)]$ is taken over the points where ϕ is \mathcal{C}^3 and hence $q_x = \nabla \phi(x)$). By calculating the primal and dual Bayesian estimators related to this global loss, we obtain that $\hat{x}_{\text{MAP}} = \operatorname{argmin}_{u \in \mathbb{R}^n} \mathbb{E}_{x|y}[D_\phi(u, x)]$ and $\hat{x}_{\text{MMSE}} = \operatorname{argmin}_{u \in \mathbb{R}^n} \mathbb{E}_{x|y}[D_\phi^*(u, x)]$, similarly to Theorem 4. Observe that these modifications do not affect the fact that \hat{x}_{MAP} and \hat{x}_{MMSE} can correspond to nondifferentiable points. Also note that despite not being a global canonical divergence, $D_\phi(u, x)$ is still consistent with the space's Riemannian geometry which is local.

Many imaging models involve nonsmooth norms such as the ℓ_1 and the nuclear norm, or the total-variation pseudo-norm, which are almost everywhere \mathcal{C}^1 but not \mathcal{C}^3 . More generally, all Lipschitz continuous functions are almost everywhere \mathcal{C}^1 . In this case, only the second and third parts of Theorem 4 hold. That is, we posit $D_\phi(u, x)$ as the loss function for any $u \in \mathbb{R}^n$ and any $x \in \mathbb{R}^n$, excluding nondifferentiable points, and obtain that

$\hat{x}_{\text{MAP}} = \operatorname{argmin}_{u \in \mathbb{R}^n} \mathbb{E}_{x|y}[D_\phi(u, x)]$ and $\hat{x}_{\text{MMSE}} = \operatorname{argmin}_{u \in \mathbb{R}^n} \mathbb{E}_{x|y}[D_\phi^*(u, x)]$ by removing nondifferentiable points from the calculation of the expectations. To recover the differential geometric derivation of D_ϕ it is necessary to use a smooth approximation of the model, i.e., the smoothed L1 norm $|s| = \sqrt{s^2 + \alpha^2}$ for some arbitrarily small $\alpha > 0$.

Finally, also note that the bound $\mathbb{E}_{x|y}[D_\phi^*(\hat{x}_{\text{MAP}}, x)] \leq n$ in Proposition 5 is straightforwardly extended to nonsmooth models by using the generalized dual Bregman divergence $D_\phi^*(u, x) = \phi(x) - \phi(u) - q_u^\top(x - u)$ with subgradient $q_{\hat{x}_{\text{MAP}}} = 0$. Conversely, the other bound $\mathbb{E}_{x|y}[D_\phi^*(\hat{x}_{\text{MMSE}}, x)] \leq n$ is lost (see the appendix for details).

5.2. Strictly log-concave models. For models that are strictly but not strongly log-concave, only the second and third results of Theorem 4 remain true. It is easy to check that the Bayesian estimators w.r.t. $D_\phi = \phi(u) - \phi(x) - \nabla\phi(x)^\top(u - x)$ are still $\hat{x}_{\text{MAP}} = \operatorname{argmin}_{u \in \mathbb{R}^n} \mathbb{E}_{x|y}[D_\phi(u, x)]$ and $\hat{x}_{\text{MMSE}} = \operatorname{argmin}_{u \in \mathbb{R}^n} \mathbb{E}_{x|y}[D_\phi^*(u, x)]$, similarly to strongly log-concave models (see in the appendix that strong log-concavity is not required to prove the second and third parts of Theorem 4). Thus, the decision-theoretic derivation of \hat{x}_{MAP} remains valid, and \hat{x}_{MAP} and \hat{x}_{MMSE} remain dual to each other. The high dimensional performance guarantees of Proposition 5 also hold because ϕ is convex. However, without strong convexity, g becomes semi-positive definite and (\mathbb{R}^n, g) becomes a singular manifold. Currently, the validity of the interpretation of D_ϕ as a canonical divergence in singular manifolds is not clear. The generalization of canonical divergences and of Theorem 4 to singular manifolds is currently under investigation. In any case, without strong convexity D_ϕ is no longer a divergence in the sense of Definition 2 as $D_\phi(x, u) = 0$ does not imply $x = u$, which is an important desired property for loss functions.

5.3. Models involving constraints. Finally, in cases where $x|y$ is constrained to a convex set $\mathcal{S} \subset \mathbb{R}^n$, only the first and the third results of Theorem 4 hold. Proceeding similarly to the proof of Theorem 4, it is easy to show that D_ϕ is the canonical divergence of the manifold (\mathcal{S}, g) , and that the Bayesian estimator related to the dual divergence is $\hat{x}_{\text{MMSE}} = \operatorname{argmin}_{u \in \mathcal{S}} \mathbb{E}_{x|y}[D_\phi^*(u, x)]$. However, the Bayesian estimator that minimizes the canonical divergence is now a shifted or tilted MAP estimator,

$$\hat{x}_{D_\phi} = \operatorname{argmin}_{u \in \mathcal{S}} D_\phi(u, \hat{x}_{\text{MAP}}) + u^\top \mathbb{E}_{x|y}[\nabla\phi(x)],$$

where generally $\mathbb{E}_{x|y}[\nabla\phi(x)] \neq 0$ (see the appendix for details). It is not clear at this point under what conditions $\hat{x}_{\text{MAP}} \approx \hat{x}_{D_\phi}$. Nevertheless, the high dimensional guarantees of Proposition 5 still hold for \hat{x}_{MAP} , providing some theoretical justification for using this estimator.

5.4. Models with heavy-tails. We conclude this section by discussing the difficulties of extending our results to models that are heavy-tailed and hence not log-concave, such as imaging models involving heavy-tailed priors related to compressible distributions [29]. Unfortunately, extending our results to heavy-tailed settings is extremely challenging for several reasons. First, the Hessian matrix of ϕ does not define a Riemannian metric because there are regions of the space where it has negative eigenvalues. Also, directly postulating $D_\phi = \phi(u) - \phi(x) - \nabla\phi(x)^\top(u - x)$ as a loss function is not appropriate either because D_ϕ can take negative values. The analysis is further complicated by the fact that $p(x|y)$ may have an infinite number of maximizers in disconnected areas of the parameter space. As mentioned

previously, the derivation of MAP estimation as an approximation arising from the degenerate loss $L_\epsilon(u, x) = \mathbf{1}_{\|x-u\|<\epsilon}$ with $\epsilon \rightarrow 0$ also fails in this case [8]. Interestingly, MMSE estimation may also struggle here given that models in this class may not have a posterior mean [38].

6. Conclusion. MAP estimation is one of the the most successful Bayesian estimation methodologies in imaging science, with a track record of accurate results across a wide range of challenging imaging problems. Our aim here has been to contribute to the theoretical understanding of this widely used methodology, particularly by placing it in the Bayesian decision theory framework that underpins the core Bayesian inference methodologies.

In order to analyze MAP estimators, we have adopted an entirely new approach: we allowed the model to specify the loss function, or equivalently the Bayesian estimator, that is used to summarize the information that the model represents. This was achieved by using the connections between model log-concavity, Riemannian geometry, and divergence functions. We first established that if $p(x|y)$ is strongly log-concave, continuous, and \mathcal{C}^3 on \mathbb{R}^n , then $\phi(x) = -\log p(x|y)$ induces a dually flat Riemannian structure on the parameter space, where the canonical divergence is the Bregman divergence associated with ϕ , and where the MAP estimator is the unique Bayesian estimator w.r.t. to this loss function. We also established that the MMSE estimator is the Bayesian estimator w.r.t. the dual canonical loss, and that both estimators enjoy favorable stability properties in high dimensions. We then examined the effect of relaxing these assumptions to models with weaker regularity conditions.

The theoretical results presented in this work provide several valuable new insights into MAP and MMSE estimation. In particular, both estimators stem from Bayesian decision theory and from the consideration of the geometry of the parameter space, and exhibit an interesting form of duality. Also, the expected estimation error—as measured by the canonical loss—is stable in high dimensions; this is in agreement with the remarkable empirical performance observed imaging and other large scale settings. The fact that MAP estimators are available as solutions to convex problems is a fundamental practical advantage. However, our results also show that the predominant view of MAP estimators as hastily approximate inferences, motivated only by computational efficiency, is fundamentally incorrect. We hope that these results will provide some clarity to imaging scientists using MAP estimators, and that they stimulate further research into the theory of this powerful Bayesian methodology.

Appendix A. Proofs of Theorem 4 and Proposition 5.

Proof of Theorem 4. The first part of Theorem 4 follows directly from differential geometry and from the regularity properties of ϕ (see [3] for an introduction to differential geometry). From differential geometry, under the conditions of Theorem 4, ϕ induces a Riemannian metric on \mathbb{R}^n with coefficients

$$g_{i,j}(x) = \partial_i \partial_j \phi(x),$$

and where we note that $g(x)$ is positive definite from the strong convexity of ϕ . Similarly, we have the affine connection coefficients $\Gamma_{ij,k} = \partial_i \partial_j \partial_k \phi(x)$.

Moreover, because ϕ is convex, it endows (\mathbb{R}^n, g) with a dual affine coordinate system η , related to the primal coordinate system by the duality $\eta_x = \nabla \phi(x)$ and $x_\eta = \nabla \phi^*(\eta)$, where $\phi^*(\eta) = \max_{x \in \mathbb{R}^n} x^\top \eta - \phi(x)$ is the convex conjugate of ϕ [3, Ch. 3]. As a result, we have a

dual Riemannian metric g^* w.r.t. η , with coefficients given by

$$g_{i,j}^*(\eta) = \partial_i \partial_j \phi^*(\eta),$$

and a dual affine connection Γ^* with coefficients given by

$$\Gamma_{i,j,k}^*(\eta) = \partial_i \partial_j \partial_k \phi^*(\eta).$$

Finally, it is easy to check that x and η are mutually dual w.r.t. g . That is, for all $x \in \mathbb{R}^n$

$$g^*(\eta_x) = g(x)^{-1},$$

which implies that $(\mathbb{R}^n, g, \Gamma, \Gamma^*)$ is a dually flat Riemannian manifold [3, Ch. 3]. Please see [35, section 2] for an excellent introduction to dually flat structures and their main properties.

From [4], in such manifolds the Γ -geodesic connecting $u \rightarrow x$ in (2) is given by $\gamma_t = u + t(x - u)$, and $\dot{\gamma}_t = x - u$. The proof is then concluded by integration by parts of (2) to obtain the Bregman divergence $D_\phi(u, x) = \phi(u) - \phi^*(\eta_x) - \eta_x^\top u$, which also admits the more familiar expression $D_\phi(u, x) = \phi(u) - \phi(x) - \nabla \phi(x)(u - x)$.

To prove the second part of Theorem 4, we use the linearity property of the expectation operator to express the definition $\hat{x}_{D_\phi} = \operatorname{argmin}_{u \in \mathbb{R}^n} \mathbb{E}_{x|y}[D_\phi(u, x)]$ as follows:

$$\begin{aligned} \hat{x}_{D_\phi} &= \operatorname{argmin}_{u \in \mathbb{R}^n} \phi(u) + \mathbb{E}_{x|y}[\phi(x)] - u^\top \mathbb{E}_{x|y}[\nabla \phi(x)] - x^\top \mathbb{E}_{x|y}[\nabla \phi(x)] \\ &= \operatorname{argmin}_{u \in \mathbb{R}^n} \phi(u) - u^\top \mathbb{E}_{x|y}[\nabla \phi(x)]. \end{aligned}$$

In a manner akin to [15], the proof is concluded by using the divergence theorem, together with the fact that $p(x|y)$ is continuous and vanishes at least exponentially as $\|x\| \rightarrow 0$, to show that $\mathbb{E}_{x|y}[\nabla \phi(x)] = \int_{\mathbb{R}^n} \nabla p(x|y) dx = 0$. Hence,

$$\begin{aligned} \hat{x}_{D_\phi} &= \operatorname{argmin}_{u \in \mathbb{R}^n} \phi(u) \\ &= \hat{x}_{\text{MAP}}. \end{aligned}$$

Note that in the case where $p(x|y)$ involves hard constraints on the parameter space, generally $\mathbb{E}_{x|y}[\nabla \phi(x)] \neq 0$, and we have $\hat{x}_{D_\phi} = \operatorname{argmin}_{u \in \mathbb{R}^n} D_\phi(u, \hat{x}_{\text{MAP}}) - u^\top \mathbb{E}_{x|y}[\nabla \phi(x)]$ generally different from \hat{x}_{MAP} .

Finally, the proof of the third part of Theorem 4 follows directly from [7, Proposition 1], which for completeness we detail below:

$$\begin{aligned} \hat{x}_{D_\phi^*} &= \operatorname{argmin}_{u \in \mathbb{R}^n} \mathbb{E}_{x|y}[D_\phi^*(u, x)] \\ &= \operatorname{argmin}_{u \in \mathbb{R}^n} \mathbb{E}_{x|y}[D_\phi(x, u)] \\ &= \operatorname{argmin}_{u \in \mathbb{R}^n} \mathbb{E}_{x|y}[D_\phi(x, u)] - \mathbb{E}_{x|y}[D_\phi(x, \hat{x}_{\text{MMSE}})] \\ &= \operatorname{argmin}_{u \in \mathbb{R}^n} \phi(\hat{x}_{\text{MMSE}}) - \phi(u) - (\hat{x}_{\text{MMSE}} - u)^\top \nabla \phi(u) \\ &= \operatorname{argmin}_{u \in \mathbb{R}^n} D_\phi(\hat{x}_{\text{MMSE}}, u) \\ &= \hat{x}_{\text{MMSE}}. \end{aligned}$$

Notice that strict log-concavity suffices to prove the second and third parts of Theorem 4.

Proof of Proposition 5. Assume that $\phi(x) = -\log p(x|y)$ is convex on \mathbb{R}^n and C^1 . From the optimality condition of \hat{x}_{MAP} , $\nabla\phi(\hat{x}_{\text{MAP}}) = 0$, and hence the dual divergence

$$D_{\phi}^*(\hat{x}_{\text{MAP}}, x) = \phi(x) - \phi(\hat{x}_{\text{MAP}}).$$

Noting that $E_{x|y}[\phi(x)]$ is the entropy of $x|y$, we use Proposition I.2 of [12] and obtain

$$E_{x|y}[D_{\phi}^*(\hat{x}_{\text{MAP}}, x)] = E_{x|y}[\phi(x)] - \phi(\hat{x}_{\text{MAP}}) \leq n.$$

Finally, using that \hat{x}_{MMSE} minimizes the posterior expectation of $D_{\phi}^*(\hat{x}_{\text{MMSE}}, x)$, we obtain

$$E_{x|y}[D_{\phi}^*(\hat{x}_{\text{MMSE}}, x)] \leq E_{x|y}[D_{\phi}^*(\hat{x}_{\text{MAP}}, x)] \leq n,$$

concluding the proof.

Acknowledgments. The author is grateful to Yoann Altmann, Gavin Gibson, Peter Green, Abderrahim Halimi, Bernd Schroers, Jonty Rougier, and Ben Powell for useful discussions.

REFERENCES

- [1] M. AFONSO, J. M. BIOCAS-DIAS, AND M. A. T. FIGUEIREDO, *An augmented Lagrangian approach to the constrained optimization formulation of imaging inverse problems*, IEEE Trans. Image Process., 20 (2011), pp. 681–695.
- [2] S. AGAPIOU, M. BURGER, M. DASHTI, AND T. HELIN, *Sparsity-promoting and edge-preserving maximum a posteriori estimators in non-parametric Bayesian inverse problems*, Inverse Problems, 34 (2018), 045002, <https://doi.org/10.1088/1361-6420/aaacac>.
- [3] S.-I. AMARI AND H. NAGAOKA, *Methods of Information Geometry*, Transl. Math. Monogr. 191, AMS, Providence, RI, 2000.
- [4] N. AY AND S.-I. AMARI, *A novel approach to canonical divergences within information geometry*, Entropy, 17 (2015), pp. 8111–8129.
- [5] S. D. BABACAN, R. MOLINA, AND A. KATSAGGELOS, *Variational Bayesian super resolution*, IEEE Trans. Image Process., 20 (2011), pp. 984–999; available online at <http://decsai.ugr.es/vip/files/journals/2011SR.BMK.pdf>.
- [6] S. D. BABACAN, R. MOLINA, AND A. K. KATSAGGELOS, *Bayesian compressive sensing using Laplace priors*, IEEE Trans. Image Process., 19 (2010), pp. 53–63.
- [7] A. BANERJEE, S. MERUGU, I. S. DHILLON, AND J. GHOSH, *Clustering with Bregman divergences*, J. Mach. Learn. Res., 6 (2005), pp. 1705–1749.
- [8] R. BASSETT AND J. DERIDE, *Maximum A Posteriori Estimators as a Limit of Bayes Estimators*, preprint, <https://arxiv.org/abs/1611.05917>, 2016.
- [9] H. H. BAUSCHKE AND P. L. COMBETTES, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*, Springer, New York, 2011.
- [10] J. BIOCAS-DIAS, *Bayesian wavelet-based image deconvolution: A GEM algorithm exploiting a class of heavy-tailed priors*, IEEE Trans. Image Process., 15 (2006), pp. 937–951.
- [11] J. M. BIOCAS-DIAS AND G. VALADAO, *Phase unwrapping via graph cuts*, IEEE Trans. Image Process., 16 (2007), pp. 698–709.
- [12] S. BOBKOV AND M. MADIMAN, *The entropy per coordinate of a random vector is highly constrained under convexity conditions*, IEEE Trans. Inform. Theory, 57 (2011), pp. 4940–4954.
- [13] S. BOYD AND L. VANDENBERGHE, *Convex Optimization*, Cambridge University Press, Cambridge, UK, 2004.
- [14] M. BURGER, Y. DONG, AND F. SCIACCHITANO, *Bregman Cost for Non-Gaussian Noise*, preprint, <https://arxiv.org/abs/1608.07483>, 2016.
- [15] M. BURGER AND F. LUCKA, *Maximum a posteriori estimates in linear inverse problems with log-concave priors are proper Bayes estimators*, Inverse Problems, 30 (2014), 114004.

- [16] X. CAI, M. PEREYRA, AND J. D. MCEWEN, *Uncertainty quantification for radio interferometric imaging: I. Proximal MCMC methods*, Monthly Notices Roy. Astronom. Soc., 480 (2018), pp. 4154–4169, <https://doi.org/10.1093/mnras/sty2004>.
- [17] E. J. CANDÈS AND B. RECHT, *Exact matrix completion via convex optimization*, Found. Comput. Math., 9 (2009), pp. 717–772.
- [18] E. J. CANDÈS, J. ROMBERG, AND T. TAO, *Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information*, IEEE Trans. Inform. Theory, 52 (2006), pp. 489–509.
- [19] A. CHAMBOLLE AND T. POCK, *An introduction to continuous optimization for imaging*, Acta Numer., 25 (2016), pp. 161–319.
- [20] P. CHARBONNIER, L. BLANC-FERAUD, G. AUBERT, AND M. BARLAUD, *Deterministic edge-preserving regularization in computed imaging*, IEEE Trans. Image Process., 6 (1997), pp. 298–311.
- [21] P. L. COMBETTES AND J.-C. PESQUET, *Proximal splitting methods in signal processing*, in Fixed-Point Algorithms for Inverse Problems in Science and Engineering, Springer, New York, 2011, pp. 185–212.
- [22] T. M. COVER AND J. A. THOMAS, *Elements of Information Theory*, 2nd ed., Wiley, New York, 2006.
- [23] M. DASHTI, K. J. H. LAW, A. M. STUART, AND J. VOSS, *MAP estimators and their consistency in Bayesian nonparametric inverse problems*, Inverse Problems, 29 (2013), 095017, <https://doi.org/10.1088/0266-5611/29/9/095017>.
- [24] A. P. DAWID, *The geometry of proper scoring rules*, Ann. Inst. Stat. Math., 59 (2007), pp. 77–93.
- [25] A. DURMUS, É. MOULINES, AND M. PEREYRA, *Efficient Bayesian computation by proximal Markov chain Monte Carlo: When Langevin meets Moreau*, SIAM J. Imaging Sci., 11 (2018), pp. 473–506, <https://doi.org/10.1137/16M1108340>.
- [26] M. A. T. FIGUEIREDO AND J. M. BIOCAS-DIAS, *Restoration of Poissonian images using alternating direction optimization*, IEEE Trans. Image Process., 19 (2010), pp. 3133–3145.
- [27] P. J. GREEN, K. LATUSZYŃSKI, M. PEREYRA, AND C. P. ROBERT, *Bayesian computation: A summary of the current state, and samples backwards and forwards*, Stat. Comput., 25 (2015), pp. 835–862.
- [28] R. GRIBONVAL, *Should penalized least squares regression be interpreted as maximum a posteriori estimation?*, IEEE Trans. Signal Process., 59 (2011), pp. 2405–2410.
- [29] R. GRIBONVAL, V. CEVHER, AND M. E. DAVIES, *Compressible distributions for high-dimensional statistics*, IEEE Trans. Inform. Theory, 58 (2012), pp. 5016–5034.
- [30] K. HÄMÄLÄINEN, A. KALLONEN, V. KOLEHMAINEN, M. LASSAS, K. NIINIMÄKI, AND S. SILTANEN, *Sparse tomography*, SIAM J. Sci. Comput., 35 (2013), pp. B644–B665, <https://doi.org/10.1137/120876277>.
- [31] T. HELIN AND M. BURGER, *Maximum a posteriori probability estimates in infinite-dimensional Bayesian inverse problems*, Inverse Problems, 31 (2015), 085009.
- [32] J. KAIPIO AND E. SOMERSALO, *Statistical and Computational Inverse Problems*, Springer, New York, 2005.
- [33] H. C. LIE AND T. J. SULLIVAN, *Equivalence of weak and strong modes of measures on topological vector spaces*, Inverse Problems, 34 (2018), 115013.
- [34] C. NICULESCU AND L. E. PERSSON, *Convex Functions and Their Applications*, Springer, New York, 2018.
- [35] F. NIELSEN AND G. HADJERES, *Monte Carlo Information Geometry: The Dually Flat Case*, preprint, <https://arxiv.org/abs/1803.07225>, 2018.
- [36] F. ONG, P. MILANFAR, AND P. GETREUER, *Local Kernels that Approximate Bayesian Regularization and Proximal Operators*, preprint, <https://arxiv.org/abs/1803.03711>, 2018.
- [37] M. PEREYRA, P. SCHNITER, E. CHOUZENOUX, J.-C. PESQUET, J.-Y. TOURNERET, A. HERO, AND S. MCLAUGHLIN, *A survey of stochastic simulation and optimization methods in signal processing*, IEEE J. Selected Topics Signal Process., 10 (2016), pp. 224–241.
- [38] C. P. ROBERT, *The Bayesian Choice*, 2nd ed., Springer-Verlag, New York, 2001.
- [39] Y. ROMANO, M. ELAD, AND P. MILANFAR, *The little engine that could: Regularization by denoising (RED)*, SIAM J. Imaging Sci., 10 (2017), pp. 1804–1844, <https://doi.org/10.1137/16M1102884>.
- [40] S. SREEHARI, S. V. VENKATAKRISHNAN, B. WOHLBERG, L. F. DRUMMY, J. P. SIMMONS, AND C. A. BOUMAN, *Plug-and-Play Priors for Bright Field Electron Tomography and Sparse Interpolation*, preprint, <https://arxiv.org/abs/1512.07331>, 2015.
- [41] S. THEODORIDIS, *Machine Learning, a Bayesian and Optimization Perspective*, Academic Press, Cambridge, UK, 2015.