Reconciling Bayesian and perimeter regularization for binary inversion

Dunbar, Oliver R. A.; Dunlop, Matthew M.; Elliott, Charles M.; Hoang, Viet Ha; Stuart, Andrew M.

2020

Dunbar, O. R. A., Dunlop, M. M., Elliott, C. M., Hoang, V. H. & Stuart, A. M. (2020).
Reconciling Bayesian and perimeter regularization for binary inversion. SIAM Journal On
Scientific Computing, 42(4), A1984-A2013. <https://dx.doi.org/10.1137/18M1179559>

<https://hdl.handle.net/10356/147617>

<https://doi.org/10.1137/18M1179559>

© 2020 Society for Industrial and Applied Mathematics. All rights reserved. This paper was
published in SIAM Journal on Scientific Computing and is made available with permission
of Society for Industrial and Applied Mathematics.

Downloaded on 26 Apr 2023 08:25:54 SGT

RECONCILING BAYESIAN AND PERIMETER REGULARIZATION FOR BINARY INVERSION*

OLIVER R. A. DUNBAR[†], MATTHEW M. DUNLOP[‡], CHARLES M. ELLIOTT[§],
VIET HA HOANG[¶], AND ANDREW M. STUART^{||}

Abstract. A central theme in classical algorithms for the reconstruction of discontinuous functions from observational data is perimeter regularization via the use of total variation. On the other hand, sparse or noisy data often demand a probabilistic approach to the reconstruction of images, to enable uncertainty quantification; the Bayesian approach to inversion, which itself introduces a form of regularization, is a natural framework in which to carry this out. In this paper the link between Bayesian inversion methods and perimeter regularization is explored. In this paper two links are studied: (i) the maximum a posteriori objective function of a suitably chosen Bayesian phase-field approach is shown to be closely related to a least squares plus perimeter regularization objective; (ii) sample paths of a suitably chosen Bayesian level set formulation are shown to possess a finite perimeter and to have the ability to learn about the true perimeter.

Key words. Bayesian inversion, phase field, level set method, perimeter regularization, gamma convergence, uncertainty quantification

AMS subject classifications. 35J35, 62G08, 62M40, 94A08

DOI. 10.1137/18M1179559

1. Introduction.

1.1. Problem statement. Let D be the unit cube $(0, 1)^d \subset \mathbb{R}^d$, $d = 2, 3$. Let $K : L^1(D) \rightarrow \mathbb{R}^J$ be a bounded linear operator. We consider the problem of recovering a binary-valued function $u \in BV_{\text{binary}}(D)$, where

$$BV_{\text{binary}}(D) = \{\psi \in BV(D) : \psi(x) \in \{\pm 1\}, x \in D\}$$

from finite dimensional data $y \in \mathbb{R}^J$ satisfying

$$(1.1) \quad y = Ku + \varepsilon^c \eta.$$

Here the finite number of observations are corrupted by noise $\varepsilon^c \eta$ of size ε^c for which we assume that η is a centered Gaussian $\mathcal{N}(0, \Sigma)$ with positive definite covariance

*Submitted to the journal's Methods and Algorithms for Scientific Computing section April 9, 2018; accepted for publication (in revised form) April 9, 2020; published electronically July 6, 2020.
<https://doi.org/10.1137/18M1179559>

Funding: The work of the first author was supported by the NSF through grant AGS 1835860. The work of the third author was partially supported by the Royal Society via a Wolfson Research Merit Award. The work of the fourth author was supported by MOE AcRF Tier 1 grant RG30/16. The work of the fifth author was supported by DARPA through contract W911NF-15-2-0121. The work of the first, third, and fifth authors was supported by the EPSRC programme grant EQUIP. The work of the first, second, third, and fifth authors was supported by the EPSRC.

[†]Geological and Planetary Sciences, California Institute of Technology, Pasadena, CA 91125 (odunbar@caltech.edu).

[‡]Courant Institute of Mathematical Sciences, New York University, New York, NY 10012 (matt.dunlop@nyu.edu).

[§]Mathematics Institute, University of Warwick, CV4 7AL, UK (c.m.elliott@warwick.ac.uk).

[¶]Division of Mathematical Sciences, School of Physical and Mathematical Sciences, Nanyang Technological University, 637371 Singapore (vhhoang@ntu.edu.sg).

^{||}Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA 91125 (astuart@caltech.edu).

$\Sigma \in \mathbb{R}^{J \times J}$. Here ε and c denote constants with $\varepsilon \ll 1$ and $c > 0$ (small noise) or $c = 0$ (noise on the order of the observations). Observe that from an application perspective the space $BV_{\text{binary}}(D)$ is a natural model for binary images. The problem of determining $u \in BV_{\text{binary}}(D)$ from y given by (1.1) thus constitutes a canonical binary inverse problem for which the issue is to recover the interface between the different domains in D in which function u takes its two values. For a given u a measure of the discrepancy with the data is the following scaled misfit functional

$$(1.2) \quad J(u) := \frac{1}{2\varepsilon^{2c}} \left| \Sigma^{-\frac{1}{2}}(y - Ku) \right|^2.$$

A typical deterministic approach to recover u , based on this misfit, would be to specify a model prior space $\mathcal{P}$ for u and to regularize the misfit functional by addition of a functional $\mathcal{R}(u)$ defined on $\mathcal{P}$, and then to seek a solution to the optimization problem

$$(1.3) \quad \inf_{u \in \mathcal{P}} (J(u) + \mathcal{R}(u)).$$

An intuitive and common method of regularization for binary problems is to penalize the perimeter of the interface. In the case of $\mathcal{P} = BV_{\text{binary}}(D)$ this leads to

$$(1.4) \quad \inf_{u \in BV_{\text{binary}}(D)} \left(J(u) + \sigma \int_D |\nabla u| \right),$$

where $\int_D |\nabla u|$ denotes the total variation of u and $\sigma > 0$ is a parameter. In this way, the minimization over $BV_{\text{binary}}(D)$ identifies perimeter regularization with total variation [57] and the Mumford–Shah [50] approaches, since these regularizations coincide on binary functions.

Often perimeter regularization is relaxed to a convex regularization by allowing values of $u \in [-1, 1]$. An alternative, nonconvex, relaxation is to use a Cahn–Hilliard functional to approximate the perimeter functional. For example one might consider, for a small parameter $\tilde{\varepsilon} > 0$,

$$(1.5) \quad \inf_{W^{1,p}(D)} \left(J(u) + \sigma_W \int_D \left(\tilde{\varepsilon} |\nabla u|^2 + \frac{1}{\tilde{\varepsilon}} W(u) \right) dx \right),$$

where $W(\cdot)$ is a double well potential. The minimizers of this Cahn–Hilliard functional are known as phase fields, and this relaxation is often referred to as a phase-field regularization. In appropriate circumstances this Γ -converges to (1.4) in the limit $\tilde{\varepsilon} \rightarrow 0$.

However, since the unknown observational noise η has an assumed Gaussian distribution, it is natural to take a probabilistic approach to the recovery of u and model uncertainty about u and, hence, the interface between different domains, through a probability distribution. This leads to Bayesian formulations of the problem in which a prior probability distribution is specified on the unknown function, and the likelihood of the data is used to compute a posterior probability distribution on the unknown function, given the data. The prior probability distribution imposes a prior space $\mathcal{Q}$ where, almost surely, samples from the posterior distribution live; the mean or mode of the posterior distribution will typically live in a smoother space $\mathcal{P} \subset \mathcal{Q}$; this space $\mathcal{P}$ will be analogous to the prior space $\mathcal{P}$ described above.

We adopt two approaches. In the first, the level set method, we reformulate the inverse problem as determining smooth functions v whose zero level set defines the

interface in the unknown function u . Specifically the sign of v defines u and the prior probability distribution on v yields a space $\mathcal{Q}$ containing C^1 functions. The pushforward measure on u , defined by the sign function, has support in $BV_{\text{binary}}(D)$. In the second approach, motivated by phase-field regularization, we relax the prior measure to allow for smooth functions u with sharp interfaces near zero; the implied space $\mathcal{Q}$ for functions u contains H^s functions for any $s < 2 - d/2$, while $\mathcal{P}$ contains H^2 functions.

We define a prior distribution over the smooth function v or over u , and formulate an associated likelihood determined by $J(u)$. Bayes' theorem is then employed in a form which implies that the posterior probability measure is absolutely continuous with respect to the prior probability measure. Taking the sign of such distributions on v yields a distribution for u . By sampling the distribution one can approximate the mean. It is then interesting to make a connection between this mean and the solution to the perimeter regularization problem (1.4).

The main goals of the paper are to investigate how the connection to perimeter regularization appears for different (Bayesian) formulations of the inverse problem, and to demonstrate the performance and applicability of these formulations for both the linear inverse problem (1.1), and nonlinear generalizations.

1.2. Background. There are many problems in the physical sciences where piecewise constant reconstruction is of interest, for example, in subsurface inversion and imaging [12, 14, 29, 39, 48, 10] and other problems in the physical sciences [28]; the problem of image deblurring [34] (in particular, for barcodes and QR codes [16, 43, 65, 63, 42, 55]) is also of interest in the context of piecewise constant reconstruction. We draw our motivation from these problems and our numerical experiments are based on imaging problems possessing a variety of geometric interfaces, smooth and including edges.

A seminal paper linking probabilistic approaches to classical numerical methods is [27], and a review describing developments since then can be found in [18]. In the context of inverse problems the link between Bayesian and classical approaches is well understood in the setting of Gaussian random field priors: the Bayesian maximum a posteriori (MAP) estimator [44, 21] is then the solution of a Tikhonov–Phillips regularized least squares problem [32]. When more complex priors are used the connection between classical and Bayesian perspectives is more subtle, even for linear inverse problems [6, 36, 7, 8, 46]; see [1, 17, 35] for recent work generalizing [21] beyond the Gaussian prior setting. Two interesting approaches to Bayesian inversion, both using thresholding as we do in this paper, may be found in [51] and [38]. In the one dimensional setting an interesting construction of random functions with finite total variation (TV) norm may be found in [19]; a Poisson process is used to define points of discontinuity, with smooth processes between these discontinuity points.

In interface reconstruction, classical methods have been dominated by inversion techniques which penalize the length of the perimeter between different subdomains. Two contrasting approaches for describing interfaces are the use of a level set of a continuous function or a characteristic function which takes just two values. TV regularization has played a central role [57] and has been shown to lead to empirically effective methods which are computationally efficient. The phase-field representation of interfaces is described in [23]. The method approximates the perimeter using a scaled gradient energy and double well potential for which minimizers have diffuse interfaces of width a small length scale and which encloses a zero level set of the minimizer. In contrast, the level set approach of [60] represents interfaces as level sets of continuous fields. See [5, 26, 58] for applications of phase-field and level-set

ideas in classical, non-Bayesian, inversion for interfaces. For simplicity this paper focuses primarily on recovery of a binary function, taking two known values, with unknown interface. In the more general setting of recovering an unknown piecewise continuous function, in which the interface and the values of the function off the interface, are unknown, the classical TV and Mumford–Shah perimeter regularization methods become distinct. For an elliptic problem, [13] uses TV regularization on a level set function to penalize both perimeter length and jump discontinuities; the Mumford–Shah minimization can be written over a suitable space to jointly minimize the function and its set of discontinuities [54].

Computer power has started to render Bayesian inversion techniques tractable in some applications [44, 64, 22], enabling uncertainty quantification. In this paper we address the question of how perimeter regularization appears within Bayesian inversion techniques for the reconstruction of a binary function u . This is a notoriously difficult problem, as made transparent in the paper [47] which showed that use of discrete TV regularization, in a Bayesian setting, does not lead to a meaningful problem in the continuum limit; this work led to the development of new Besov priors in [46], and the combination of TV and Gaussian priors [67]. Other approaches, with demonstrable numerical success, make use of the introduction of hyperparameters [48, 9]. Instead of approximating the TV regularization, one can derive Bayesian approaches based on the Mumford–Shah functional [36]; these methods can be extended to higher order functionals such as Blake–Zisserman [11]. Our probabilistic level set based method generalizes to a hierarchical Bayesian approach that learns the unknown continuous function off the unknown interfaces, as well as the interface itself; the interface may be viewed as a nonparametric hyperparameter.

1.3. Our contribution. In detail our contributions are as follows:

1. We formulate a Bayesian level set approach and establish conditions under which this leads to posterior samples with almost surely finite perimeter and, hence, almost surely a finite TV norm. This demonstrates that TV regularization arises naturally out of appropriately chosen Bayesian formulations of inversion.
2. We formulate a Bayesian based phase-field approach and establish a link with perimeter regularization through its MAP estimator. We prove, for an appropriate choice of prior distribution and parameters carefully scaled with respect to ε (and so not strictly Bayesian), that the MAP estimator for this phase-field approach has a Γ -limit as $\varepsilon \rightarrow 0$. This limit is exactly the perimeter (TV) regularization of the least squares fidelity objective function.
3. For a linear inverse problem we provide numerical investigations of these approaches; we also compare with (widely used) Gaussian process regression which is a natural method in this linear setting. These investigations demonstrate that the level set approach may be implemented quite cheaply in comparison with the phase-field approach, for similar levels of reconstruction accuracy. Also it is demonstrated that the level set approach can learn the true perimeter. Gaussian process regression also performs well at a low computational cost for the linear problem, but is not readily extended to nonlinear problems.
4. We provide numerical evidence for the flexibility of the Bayesian level set approach, by showing an application to a nonlinear inverse problem arising from the eikonal equation. Within this context, we also show that hyperparameters contained in the statistical model may also be efficiently learned.

1.4. Some notation. We use $|\cdot|$ to denote the Euclidean norm on $\mathbb{R}^J$. Let $C_{\#}^{k,\gamma}(\bar{D})$, $k \geq 0$ denote the space of real valued continuous periodic functions on $\bar{D}$ whose derivatives up to the k th order derivative are Hölder continuous with exponent γ . By virtue of continuous embedding K is a bounded linear operator on $C_{\#}^{k,\gamma}(\bar{D})$ for any integer $k \geq 0$. Also let $H_{\#}^k(D)$, $k \geq 0$ denote the restriction to periodic functions of the Sobolev space of $H^k(D)$ of k -times weakly differentiable real valued functions on D . These Sobolev spaces are readily characterized as weighted ℓ^2 spaces on Fourier coefficients [56]. Let X denote the space $C(\bar{D})$, restricted to periodic functions, and let H denote $L^2(D)$.

Fix constants $\delta > 0$, $\tau > 0$, $q \geq 0$, and $a_i \geq 0$, $i = 1, 2, 3$. Denote $\tilde{\delta} = (\delta, \tau, q)$ and $\tilde{a} = (a_1, a_2, a_3)$. We define a covariance operator $\mathcal{C}_{\varepsilon, \tilde{\delta}, \tilde{a}}$ implicitly as the solution operator corresponding to the weak formulation of the following elliptic boundary value problem: given $f \in H$ find $\eta \in H_{\#}^2(D)$ so that

$$(1.6) \quad \delta \varepsilon^{-2a_1} \Delta^2 \eta - q \delta \varepsilon^{-2a_2} \Delta \eta + \tau^2 \delta \varepsilon^{-2a_3} \eta = f.$$

Elliptic regularity gives $\eta \in H_{\#}^4(D)$ and so we may define $(\mathcal{C}_{\varepsilon, \tilde{\delta}, \tilde{a}})^{-1} : H_{\#}^4(D) \rightarrow H$ by $(\mathcal{C}_{\varepsilon, \tilde{\delta}, \tilde{a}})^{-1} \eta = f$ for $f \in H$. The Hilbert space $\mathcal{E}_{\varepsilon, \tilde{\delta}, \tilde{a}}$, is defined to be $H_{\#}^2(D)$ endowed with the norm

$$\|\xi\|_{\mathcal{E}_{\varepsilon, \tilde{\delta}, \tilde{a}}}^2 := \langle (\mathcal{C}_{\varepsilon, \tilde{\delta}, \tilde{a}})^{-\frac{1}{2}} \xi, (\mathcal{C}_{\varepsilon, \tilde{\delta}, \tilde{a}})^{-\frac{1}{2}} \xi \rangle = \delta \int_D (\varepsilon^{-2a_1} |\Delta \xi|^2 + q \varepsilon^{-2a_2} |\nabla \xi|^2 + \tau^2 \varepsilon^{-2a_3} \xi^2) dx,$$

where $\langle \cdot, \cdot \rangle$ denotes the standard $L^2(D)$ inner product. By polarization an inner product is defined on $\mathcal{E}_{\varepsilon, \tilde{\delta}, \tilde{a}}$. The three parameters δ, q , and τ weight the contributions of the $H_{\#}^2(D)$, $H_{\#}^1(D)$, and $L^2(D)$ terms appearing in the Hilbert space $\mathcal{E}_{\varepsilon, \tilde{\delta}, \tilde{a}}$ norm whereas the parameters a_1, a_2 , and a_3 scale these terms with respect to powers of ε . The Hilbert space $\mathcal{E}_{\varepsilon, \tilde{\delta}, \tilde{a}}$ is exactly the Cameron–Martin space for the Gaussian $\mathcal{N}(0, \mathcal{C}_{\varepsilon, \tilde{\delta}, \tilde{a}})$ [22, Definition 6.26]. In the following we write $\mathcal{C}$ and $\mathcal{E}$, with the dependence on the parameters being understood. In the computations it is made clear which values of the parameters are chosen.

Remark 1. We note that including an $H_{\#}^2(D)$ contribution in the Cameron–Martin norm is required in dimensions $d = 2, 3$ in order to ensure that the underlying Gaussian is supported on continuous functions. In dimension $d = 1$ it is possible to remove this requirement [61].

Remark 2. The requirement that $\tau > 0$ is made to ensure that $(\mathcal{C}_{\varepsilon, \tilde{\delta}, \tilde{a}})^{-1}$ is invertible on $L_{\#}^2(D)$. This could also be addressed when $\tau = 0$ by working on spaces of functions where the mean value is zero.

1.5. Outline of the paper. In subsection 2.1 we formulate three inversion approaches for the linear inverse problem (1.1), a Bayesian level set based approach, a Bayesian based phase-field regularization, and Gaussian process regression. We introduce level set and phase-field priors, both of which are non-Gaussian and state a well-posedness result for the resulting posterior distributions, as well as some relevant properties. Section 3 characterizes the MAP estimator for the phase-field prior. Under appropriate parameter scalings, we obtain perimeter regularization as a Γ -limit for the MAP estimator in the small noise regime, using the analysis in [37]. This Γ -limit links the MAP estimator to classical deterministic perimeter regularization. In section 4 we describe testing the approaches with numerical experiments based on Markov chain Monte Carlo (MCMC). We also discuss the properties of the length of level sets of

Gaussian random fields, and of random fields whose law has density with respect to a Gaussian random field, and use this to demonstrate that the level set approach *learns* the perimeter. These numerical results establish interesting links between Bayesian level set inversion and perimeter regularization. Within section 5 we go beyond linear inverse problems, showing that the Bayesian level set approach readily extends both to a nonlinear inverse problem arising from the eikonal equation, and to the learning of unknown hyperparameters from the prior; for hyperparameter learning we use algorithms introduced in [31] and further developed in [15]. Appendix A contains proofs of the main results relating to MAP estimators for the phase-field approach.

2. Inversion approaches.

2.1. A Bayesian level set based approach to the inverse problem.

2.1.1. Prior and likelihood. Let $\mathbb{1}_\cdot$ denote the characteristic function of a set. Our model prior for u is that

$$u = \mathbb{1}_{D_+} - \mathbb{1}_{D_+^c},$$

where $D_+ \subset D$ is a random set defined in such a way that $\text{leb}(\overline{D_+} \setminus D_+) = 0$, almost surely; this ensures that u is almost surely in BV_{binary} . This is achieved by working with an auxiliary variable v and recovering u by an application of a thresholding (sign) function $S : \mathbb{R} \mapsto \{-1, 0, +1\}$ defined by

$$S(v) = 1, v > 0, \quad S(0) = 0, \quad \text{and} \quad S(v) = -1, v < 0.$$

The prior on v is chosen to be $\mu_{0,\alpha} = \mathcal{N}(0, C^{\frac{\alpha}{2}})$ with the power satisfying $\alpha > d/2$. This implies almost sure continuity of v . The prior for u is then defined by the pushforward of $\mu_{0,\alpha}$ under S :

$$(2.1) \quad \mathcal{E}^S := \{u = S(v) \mid v \in \mathcal{E}\}.$$

This is justified since the level sets of the Gaussian random field v have Lebesgue measure zero [41, Proposition 2.5]. Also by the lemma which follows, it holds almost surely that if $\alpha > 1 + d/2$ then $u \in \{\pm 1\}^D$ a.e. and has bounded TV.

LEMMA 2.1. *If function v is drawn from measure $\mu_{0,\alpha}$ with $\alpha > 1 + d/2$ then almost surely function u defined by (2.1) has finite TV norm.*

Proof. If $\alpha > 1 + d/2$ then almost surely $v \sim \mu_{0,\alpha}$ will be a C^1 function. The paper [45] establishes that the level set $v = 0$ will then have finite length, almost surely. Since u is a binary function given by (2.1) this establishes that u will have finite TV, almost surely. $\square$

If we set $u = S(v)$ it follows that model (1.1) becomes

$$y = KS(v) + \varepsilon^c \eta$$

and hence that $y|v$ is distributed as the Gaussian $\mathcal{N}(KS(v), \varepsilon^{2c}\Sigma)$. The likelihood is given by the Gaussian density proportional to

$$\exp(-J(S(v))) = \exp\left(-\frac{1}{2\varepsilon^{2c}} \left| \Sigma^{-\frac{1}{2}} (y - KS(v)) \right|^2\right).$$

$\Phi(v; y) = J \circ S(v)$ is the negative log-likelihood function.

2.1.2. Posterior. In the following proposition we establish a relationship between the (posterior) distribution μ^y of the random variable $v|y$, the prior on v , and the likelihood on $y|v$, by means of an infinite dimensional Bayes' theorem [22].

PROPOSITION 2.2. *Let $\alpha > d/2$. Then the posterior probability μ^y on random variable $v|y$ is a probability measure supported on $C_{\#}^{k,\gamma}(\bar{D})$ for all $\gamma \in [0, \gamma')$, where $k \in \{0, 1, 2, \dots\}$ is chosen so that $\gamma' := \alpha - \frac{d}{2} - k \in (0, 1]$, and is determined by*

$$\frac{d\mu^y}{d\mu_{0,\alpha}} = \frac{1}{Z} \exp \left(-\frac{1}{2\varepsilon^{2c}} \left| \Sigma^{-\frac{1}{2}}(y - KS(v)) \right|^2 \right);$$

here $Z \in (0, \infty)$ is the normalization constant that makes μ^y a probability measure.

Proof. The proof follows from an application of the theory in [41]. $\square$

We state some beneficial properties of the posterior, making it fit for purpose in this application. Specifically, the posterior μ^y has a continuous dependence on y , and the pushforward under S defines an implied posterior ν^y , whose samples have finite TV.

Recall the Hellinger distance between measures μ and μ' , defined with respect to any common reference measure μ_0 (but independent of it) and given by

$$d_{\text{hell}}(\mu, \mu') = \sqrt{\left(\int_X \left(\sqrt{\frac{d\mu}{d\mu_0}} - \sqrt{\frac{d\mu'}{d\mu_0}} \right)^2 d\mu_0 \right)}.$$

PROPOSITION 2.3. *Given the setting of Proposition 2.2, (i) the posterior measure μ^y is locally Lipschitz continuous with respect to $y \in \mathbb{R}^J$; more precisely, if $|y| < \rho$ and $|y'| < \rho$ for a constant $\rho > 0$ then there is a constant $C = C(\rho)$ such that*

$$d_{\text{hell}}(\mu^y, \mu^{y'}) \leq C(\rho)|y - y'|;$$

(ii) if $\alpha > 1 + d/2$ then $u = S(v)$ with $v \sim \mu^y$, has finite TV norm, almost surely.

Proof. (i) follows from application of the theory in [41]. (ii) follows by noting that, since μ^y has density with respect to $\mu_{0,\alpha}$, anything which holds almost surely under $\mu_{0,\alpha}$, will also hold almost surely under μ^y . Application of Lemma 2.1 gives the desired result. $\square$

2.2. A phase-field regularization based Bayesian approach to the inverse problem. The Bayesian level set approach of section 2.1 is formulated in terms of a prior on a smooth variable v whose pushforward under the thresholding map gives a function u taking values in $\{-1, 1\}$, respecting the fact that the data take values of the form $\{-1, 1\} + \text{noise}$. Here we describe a different approach, one in which the ε -dependent prior on u may take values anywhere in $\mathbb{R}$, but concentrates close to $\{-1, 1\}$ when $\varepsilon \ll 1$. This leads to a connection with phase-field regularization.

2.2.1. Prior and likelihood. Fixing constants $r, b > 0$ define $\Psi : X \rightarrow \mathbb{R}^+$ by

$$(2.2) \quad \Psi(u) = \frac{r}{\varepsilon^b} \int_D \frac{1}{4} (1 - u(x)^2)^2 dx.$$

We define the prior probability measure ν_0 on X via the Radon–Nikodym derivative

$$(2.3) \quad \frac{d\nu_0}{d\mu_0} = \frac{1}{Z_0} \exp(-\Psi(u)),$$

where μ_0 is the Gaussian measure $\mu_0 = \mathcal{N}(0, C)$ on the Hilbert space H . The normalization Z_0 is chosen so that ν_0 is a probability measure. Since the Gaussian measure μ_0 is supported on continuous functions in dimensions 2 and 3, so is the non-Gaussian measure ν_0 . Furthermore, since $r, b > 0$ and $\varepsilon \ll 1$, this measure will concentrate on functions taking values close to ± 1 . In what follows the choice of parameter b will be crucial, and will be explained below; the precise value of the positive parameter r is less significant.

The random variable $y|u$, given by (1.1), has a Gaussian distribution $\mathcal{N}(Ku, \varepsilon^{2c}\Sigma)$ and the likelihood is the Gaussian density proportional to

$$\exp(-J(u)) = \exp\left(-\frac{1}{2\varepsilon^{2c}} \left| \Sigma^{-\frac{1}{2}}(y - Ku) \right|^2\right).$$

$\Phi(u; y) = J(u)$ is the negative log-likelihood function.

2.2.2. Posterior. We let $\nu^y(du)$ denote the probability of the conditioned random variable $u|y$. The following propositions are the analogue of Propositions 2.2 and 2.3. They are proved by a straightforward application of the theory in [22, 41].

PROPOSITION 2.4. *The posterior probability ν^y on random variable $u|y$ is a probability measure supported on $C_{\#}^{k,\gamma}(\bar{D})$ for any $\gamma < 2 - d/2$ and determined by*

$$\frac{d\nu^y}{d\nu_0} = \frac{1}{Z} \exp\left(-\frac{1}{2\varepsilon^{2c}} \left| \Sigma^{-\frac{1}{2}}(y - Ku) \right|^2\right),$$

where $Z \in (0, \infty)$ is the normalization constant that makes ν^y a probability measure.

PROPOSITION 2.5. *In the setting of Proposition 2.4, the posterior measure ν^y is locally Lipschitz continuous with respect to $y \in \mathbb{R}^J$; more precisely, if $|y| < \rho$ and $|y'| < \rho$ for a constant $\rho > 0$ then there is a constant $C = C(\rho)$ such that*

$$d_{\text{hell}}(\nu^y, \nu^{y'}) \leq C(\rho)|y - y'|.$$

Remark 3. For computations, it is convenient to draw samples from the Gaussian prior μ_0 , not the prior ν_0 ; to this end we note that the posterior ν^y may be written as

$$\frac{d\nu^y}{d\mu_0} = \frac{1}{Z_0 Z} \exp\left(-\frac{r}{\varepsilon^b} \int_D \frac{1}{4} (1 - u(x)^2)^2 dx - \frac{1}{2\varepsilon^{2c}} \left| \Sigma^{-\frac{1}{2}}(y - Ku) \right|^2\right)$$

for normalization constants $Z_0, Z \in (0, \infty)$.

2.3. A Gaussian process regression on the inverse problem. A popular approach for linear inverse problems is to use a Gaussian process (GP) regression to find posterior parameter distributions [52, 66]; this may be combined with thresholding to perform classification [66], an approach we adapt here to learning a binary function.

A GP is a collection of random variables, with all finite subsets being described by a joint Gaussian distribution. Adapted to our specific inverse problem, GP regression proceeds by imposing a Gaussian prior on the unknown function, and then

conditioning this on y given by (1.1). We take as prior μ_0 the Gaussian $\mathcal{N}(0, C)$. The posterior is completely described as $\nu^y = \mathcal{N}(m_y, C_y)$ with mean m_y and covariance function C_y . This is attractive as with a closed form for m_y and C_y , we may sample the posterior directly without the need for MCMC, thus it is very computationally efficient. Closed forms are found to be

$$m_y = CK^*(\varepsilon^{2c}\Sigma + KCK^*)^{-1}y, \quad C_y = C - CK^*(\varepsilon^{2c}\Sigma + KCK^*)^{-1}KC.$$

In the linear setting m_y is the MAP estimator of the posterior, thus is the unique minimizer of the functional

$$(2.4) \quad J(u) + \mathbf{R}(u) = \frac{1}{2\varepsilon^{2c}} \left| \Sigma^{-\frac{1}{2}}(y - Ku) \right|^2 + \frac{1}{2} \|u\|_{\mathcal{E}}^2.$$

This approach is unnatural from a modeling point of view, as the difference $y - Ku$ appearing in (2.4) contains a comparison between data produced from a binary field and the forward map of a nonbinary field. It is therefore difficult to interpret the results from this approach. Nonetheless we proceed in a fashion standard in machine learning, namely, to threshold the solution of the regression to obtain a classifier; in our particular setting this corresponds to application of S to m_y , or to samples from the Gaussian posterior distribution. We also note that the GP regression methodology is very specific to the linear inverse problems and does not extend directly to nonlinear forward mappings.

3. MAP estimators, phase field regularization, and Γ -convergence. A MAP estimator of a Bayesian posterior distribution maximizes the posterior probability. Intuitively, the MAP estimator locates points in X at which arbitrarily small balls will have maximal probability. It is defined as follows [21].

DEFINITION 3.1. *A point $\bar{z} \in X$ is a MAP estimator for the posterior measure ν^y if*

$$\lim_{\rho \rightarrow 0} \frac{\nu^y(B^\rho(\bar{z}))}{\nu^y(B^\rho(z^\rho))} = 1,$$

where

$$z^\rho = \operatorname{argmax}_{z \in X} \nu^y(B^\rho(z))$$

and $B^\rho(z), \rho > 0$ is the ball centered at $z \in X$ with radius ρ .

We explore this concept in the context of phase-field regularization. Set $\Phi : X \times \mathbb{R}^J \rightarrow \mathbb{R}^+$ to be the sum

$$(3.1) \quad \Phi(u; y) = \Psi(u) + \frac{1}{2\varepsilon^{2c}} |\Sigma^{-\frac{1}{2}}(y - Ku)|^2,$$

where Ψ is defined in (2.2). We define the Onsager–Machlup functional J^ε , associated with the measure ν^y by

$$J^\varepsilon(u) = \begin{cases} \frac{1}{2} \|u\|_{\mathcal{E}}^2 + \Phi(u; y) & \text{if } u \in \mathcal{E}, \\ \infty & \text{if } u \notin \mathcal{E}. \end{cases}$$

Recall $(\mathcal{E}, \|\cdot\|_{\mathcal{E}})$ is the Hilbert space and corresponding norm defined in subsection 1.4. From a probabilistic perspective, $(\mathcal{E}, \|\cdot\|_{\mathcal{E}})$ is the Cameron–Martin space associated with the Gaussian measure $\mathcal{N}(0, C)$ [22, Definition 6.26]. We have the following result demonstrating the role of the Onsager–Machlup functional defined on the Cameron–Martin space from [21, Theorem 3.5].

PROPOSITION 3.2. *There exists a MAP estimator for the posterior measure ν^y which is a minimizer of the functional J^ε .*

The functional $J^\varepsilon(u)$ can be written as

$$(3.2) \quad J^\varepsilon(u) = \varepsilon^{-2a_1-3} I^\varepsilon(u),$$

where

$$\begin{aligned} I^\varepsilon(u) = & \frac{1}{2} \delta \varepsilon^3 \|\Delta u\|_{L^2(D)}^2 + \frac{1}{2} \delta q \varepsilon^{3+2(a_1-a_2)} \|\nabla u\|_{L^2(D)}^2 + \frac{1}{2} \delta \tau^2 \varepsilon^{3+2(a_1-a_3)} \|u\|_{L^2(D)}^2 \\ & + r \varepsilon^{3+2a_1-b} \int_D \frac{1}{4} (1 - u(x)^2)^2 dx + \frac{1}{2} \varepsilon^{3+2a_1-2c} |\Sigma^{-\frac{1}{2}}(y - Ku)|^2. \end{aligned}$$

We consider the case where

$$(3.3) \quad a_2 - a_1 = 1, \quad 3 + 2a_1 = b - 1 = 2c, \quad 3 + 2(a_1 - a_3) = a > 0.$$

With these parameter constraints the functional $I^\varepsilon(u)$ becomes, for $u \in H_\#^2(D)$,

$$\begin{aligned} I^\varepsilon(u) = \int_D & \left(\frac{1}{2} \delta \varepsilon^3 |\Delta u|^2 + \frac{1}{2} \delta q \varepsilon |\nabla u|^2 + \frac{r}{4\varepsilon} (1 - u(x)^2)^2 + \delta \tau^2 \varepsilon^a u(x)^2 \right) dx \\ & + \frac{1}{2} |\Sigma^{-\frac{1}{2}}(y - Ku)|^2 \end{aligned}$$

and $I^\varepsilon(u) = +\infty$ when $u \in H \setminus H_\#^2(D)$.

DEFINITION 3.3. *Define the following two functionals and constant:*

$$\begin{aligned} I_0^\delta &= \frac{1}{2} \int_D P^\delta |\nabla u| dx + \frac{1}{2} |\Sigma^{-\frac{1}{2}}(y - Ku)|^2, \quad \text{if } u \in BV_{\text{binary}}(D), \\ e^\delta(U) &= \int_{-\infty}^{\infty} \left(\frac{1}{2} \delta (U''(t))^2 + \frac{q}{2} \delta (U'(t))^2 + \frac{r}{4} (1 - U(t)^2)^2 \right) dt, \\ P^\delta &= \inf_{U_{\text{odd}}} e^\delta(U). \end{aligned}$$

Based on the work of Hilhorst, Peletier, and Schätzle [37], we have the following theorem for Γ -convergence of the functional I^ε .

THEOREM 3.4. *Then*

$$I_0^\delta = \lim_{\varepsilon \rightarrow 0} I^\varepsilon$$

in the sense of Γ -convergence in the strong $L^1(D)$ topology.

Proof. See Appendix A for the proof. $\square$

This theorem shows that the MAP estimator is, for small observational noise $\varepsilon^c \eta$, close to a perimeter regularization. Furthermore, since $2a_1 + 3 > 0$, (3.2) together with the preceding Γ -limit theorem suggest that, when $\varepsilon \ll 1$, the measure will approximately concentrate on a single point close to a minimizer of I_0^δ . Our numerical results, presented in the next section, support this conjecture.

Remark 4. This establishes a link with perimeter regularization for the Bayesian phase-field approach at the level of the MAP estimator. This is not available for the Bayesian level set method because MAP estimators do not exist [41]. Conversely, the Bayesian level set approach has a link to perimeter length at the level of samples

FIG. 1. The three true fields used for inversion; the field on the left will be referred to as Truth A, the field in the middle as Truth B, and the field on the right as Truth C. The sets of observation points are shown in each figure.

from the posterior (see Lemma 2.1), which is not true for the phase-field approach. We explain why this is the case. Recall that almost sure properties of the prior are inherited in the posterior. Prior samples are drawn from the centered Gaussian with covariance C , which corresponds to choosing $\alpha = 2$. In dimension $d \geq 2$ we thus do not have $\alpha > 1 + d/2$ and we cannot deduce that samples have finite perimeter almost surely. (Numerical results, reported in subsection 4.5, demonstrate that the lower bound $\alpha = 1 + d/2$ is sharp and that random draws beneath this value do not have finite perimeter almost surely). MAP estimators, on the other hand, will live in the Cameron–Martin space of the underlying Gaussian reference measure and are necessarily smoother than draws from the measure itself [21]. Thus there is no contradiction between the fact that the MAP estimator, for small ε , approximately penalizes the perimeter, while samples from the posterior may have infinite perimeter.

Remark 5. The phase-field approach has the (undesirable) property that, for $c > 0$, the prior construction depends on the noise level ε^c through (3.3) meaning that it is not strictly Bayesian.

4. Numerical simulations.

4.1. Test problems. We test the three inversion techniques on three images referred to as Truth A, Truth B, and Truth C. These are three fields $u^\dagger$ lying in the set of BV_{binary} images and are illustrated in the obvious way in Figure 1. Truth A and Truth B are observed on a uniform grid of 15×15 points, Truth C is observed at 50 uniformly distributed points, and all of these observations are corrupted by additive Gaussian noise as in (1.1). Pointwise observation does not fit our theory as we assume K is linear on $L^1(D)$; however, mollification can be used to address this and leads to results which are not different in any substantive way, as noted in [40].

In order to avoid an inverse crime [44], Truth A and Truth B are generated on a square mesh of 2^{16} points, and Truth C is generated on a square mesh of 320^2 points, but the Gaussian random fields are constructed over a mesh of 2^{14} points ($N = 2^7$ in subsection 4.2). We perform numerical experiments in both the small noise and order one noise regimes.

4.1.1. Small observational noise setup. We set $c = 3/2$ and $\varepsilon = 0.01$. The implied standard deviation of the observational noise is thus 0.001. We make the choices of parameters in the prior covariance operator $C(=C_{\varepsilon, \tilde{\delta}, \tilde{a}})$: $a_1 = 0$, $a_2 = 1$, $a_3 = 0$, $b = 4$, and $r = 1$ for both the Bayesian level set approach and the phase-field approach. For the Bayesian level set approach we set $\delta = 1$, $q = 0$, $\tau = 50$,

and $\alpha = 3$, whereas for the phase-field approach $\delta = 0.01$, $q = 0.1$, $\tau = 1$, and $\alpha = 2$. Thus we ensure that the relations (3.3) hold so that the phase-field MAP estimator for ν^y approximates the minimizer of I_0^δ as given in Theorem 3.4 and we expect the posterior mass to concentrate fairly close to this MAP estimator. Note that in general we need not insist on the parameters being related via (3.3) for the level set formulation; this is because, unlike the phase-field formulation, there is no MAP estimator whose properties we are seeking to control via parameter choices. For these small noise experiments the GP regression used the same parameters as for the Bayesian level set method.

4.1.2. Order one observational noise setup. We set $c = 0$. Note that now ε does not enter the observational noise; it is simply a parameter that enters the prior and so the phase-field formulation is truly Bayesian. With this choice of c we require, for the phase-field formulation, $a_1 = -3/2$, $a_2 = -1/2$, $a_3 = -1$, $b = 1$. We also set $\delta = 100$, $q = 0.1$, $\tau = 1$, $\alpha = 2$, and $r = 1$. For the level set formulation we retain the same choice of parameters as for the small noise case above. For the GP regression we use the same parameters as for the phase-field approach for these order one noise experiments.

4.2. Sampling from the Gaussian prior space. We describe how to sample numerically from Gaussian priors $\zeta_0 = \mathcal{N}(0, \mathbf{C})$ that are key to the techniques outlined in the preceding two subsections. Here $\mathbf{C}$ is either $C^{\frac{\alpha}{2}}$ or C . We consider the case that D is the unit square $(0, 1)^2$. Let $\{\lambda_k\}$ denote the eigenvalues of $\mathbf{C}$ in increasing order with corresponding $L^2(D)$ -normalized eigenfunctions (which are Fourier modes) $\{\varphi_k\}$. Then samples z from ζ_0 may be expressed through the Karhunen–Loève expansion as

$$(4.1) \quad z(x) = \sum_{k=1}^{\infty} \lambda_k^{\frac{1}{2}} \xi_k \varphi_k(x), \quad \xi_k \sim \mathcal{N}(0, 1) \text{ independent and identically distributed (i.i.d.).}$$

We implement an approximation to this by jointly approximating the field via spectral truncation and evaluation on a discrete grid of points. Such an approximation may be efficiently implemented using the fast Fourier transform. We work on a uniformly spaced grid $\{x_i\}$ of N^d points in D . An approximate sample on this grid is then given by

$$z^N(x_i) = \sum_{k=1}^{N^d} \lambda_k^{\frac{1}{2}} \xi_k \varphi_k(x_i), \quad \xi_k \sim \mathcal{N}(0, 1) \text{ i.i.d.}$$

All of our numerical results are performed on the two dimensional grid which arises from this approach to generating Gaussian random fields in dimension $d = 2$. In practice we choose $N = 2^7$, and so the discrete grid for our inversion is 2^{14} points.

4.3. MCMC simulation. MCMC simulations may be used to sample measures ν^y .

In all MCMC runs we generate 10^6 samples and, when computing means, discard the first 5×10^5 samples as burn-in.

We employ the preconditioned Crank–Nicolson (pCN) algorithm [2, 20] which may be used to sample any measure σ of the form

$$\frac{d\sigma}{d\sigma_0}(z) = \frac{1}{Z} \exp(-A(z)), \quad \sigma_0 = \mathcal{N}(0, C),$$

without computing derivatives of $A(\cdot)$. Both of our posterior measures can be written in this form. For the Bayesian level set approach we have

$$\frac{d\nu^y}{d\mu_{0,\alpha}} = \frac{1}{Z} \exp \left(-\frac{1}{2\varepsilon^{2c}} \left| \Sigma^{-\frac{1}{2}}(y - KS(v)) \right|^2 \right),$$

whereas for the phase-field approach we have

$$\frac{d\nu^y}{d\mu_0} = \frac{1}{Z} \exp \left(-\frac{r}{\varepsilon^b} \int_D \frac{1}{4} (1 - u(x)^2)^2 dx - \frac{1}{2\varepsilon^{2c}} \left| \Sigma^{-\frac{1}{2}}(y - Ku) \right|^2 \right),$$

where $\mu_0, \mu_{0,\alpha}$ are Gaussian measures, and with the appropriate normalization constants Z . The pCN method has the advantage that, unlike the standard random walk Metropolis MCMC algorithm, its rate of convergence to equilibrium can be bounded below independently of the number of terms used in the truncated Karhunen–Loève expansion described in subsection 4.2 [33].

In the notation of [20] for the pCN algorithm, $\beta \in (0, 1]$ denotes the proposal variance parameter. Note that larger β tends to lead to smaller acceptance probability, but to greater exploration of state space when steps are accepted; the optimal β is a trade-off between these two competing effects. Depending on the noise model and number of observations, we take the proposal standard deviation parameter β between 0.02 and 0.1 for level set simulations and β between 0.002 and 0.02 for phase-field simulations. These choices are made in order to balance acceptance rate and size of proposed move with a view to optimizing the convergence rate of the Markov chain.

We simply *assume* that the resulting Markov chains $\{\eta^{(m)}\}$ are ergodic and make the approximation for an associated measure ν^y that

$$\mathbb{E}^{\nu^y} g(\eta) \approx \frac{1}{M} \sum_{m=1}^M g(\eta^{(m)}) + e_M,$$

where the error e_M is Gaussian with variance c_g/M , and c_g is the integrated autocorrelation of $g(\eta^{(m)})$. We do not impose specific stopping criteria on the Markov chains, rather we will examine the approximation qualities derived from the chains, as a function of M , and study the convergence to equilibrium of quantities of interest; in particular, in subsection 4.3.1 we compare the acceptance probability of the chain, as a function of M , for the level set and phase-field approaches. The samples up to step M can then be used to produce point estimates for the fields, by calculating, for example, their mean or the sign of their mean. We compare the cost of sampling versus the quality of reconstruction with these point estimates, for differing formulations.

Remark 6. The theory in [33] demonstrates ergodicity for problems similar to those arising in the phase-field formulation. Developing an analogous theory for the level set formulation is an open and interesting research direction; however, our numerics do suggest that ergodicity holds in this case too.

Remark 7. Preliminary numerical results for the one dimensional analogue of the problem studied here may be found in [61]. They are consistent with what we report here in dimension two.

4.3.1. Computational cost. For MCMC sampling, which we use for both the phase-field and level set approaches, every set of the Markov chain requires an evaluation of $A(u)$. Due to the presence of an extra integral term, this evaluation will typically be more expensive for the phase-field model than the level set model; for the

FIG. 2. The evolution of the acceptance rate of proposals for the phase-field (left) and level set (right) MCMC chains for Truth A with small observational noise. Acceptance rates are calculated over a moving window of 1000 samples.

simulations performed here, evaluation of $A(u)$ is approximately twice as expensive for the phase-field model as for the level set model. For the GP regression simulations no sampling is required and so the computational cost is significantly cheaper; the means were calculated from the expression in subsection 2.3, with the cost arising from the matrix multiplications and inversion involved.

For the phase-field and level set approaches, the most significant discrepancy in computational cost arises from the statistical properties of the Markov chain used to sample the posterior approximately. In Figure 2 we show the evolution of the local acceptance rates of proposed states for Truth A with small observational noise, as a function of M . The evolutions are similar for the other datasets and so are not presented for brevity. The parameter β , the proposal variance, is chosen so that the acceptance probability is neither close to one nor zero; recall that this results in an order of magnitude smaller value of β for the phase-field method in comparison with level set, meaning that the former method makes a much slower exploration of the posterior distribution. Figure 2 suggests that the phase-field chains have not reached equilibrium until after at least 5×10^5 samples, whereas the level set chains converge much earlier. This is illustrated in Figure 3, which shows a selection of samples for $M = \mathcal{O}(10^4)$ for Truth B with small observational noise. With $M = 10^4$ samples, the three inclusions have already been identified by the level set chain; however, after $M = 3 \times 10^4$ samples the phase-field chain has only started to identify a second inclusion. Thus, even though for both models we produced the same number of samples, it would have sufficed to terminate the level set chains much earlier, significantly reducing the computational cost.

The fact that the acceptance rates for the phase-field chains are lower than those for the level set chains, despite the proposal standard deviation parameter β being one tenth of the size can be understood as follows. Note that the measure ν^y can informally be thought of as having Lebesgue density proportional to $\exp(-J^\varepsilon(u)) = \exp(-\varepsilon^{-2a_1-3}I^\varepsilon(u))$. Thus for small ε the probability mass is concentrated in a small neighborhood of critical points of $I^\varepsilon \approx I_\delta^0$. The MCMC simulations for ν^y could hence be viewed as a form of derivative-free optimization for the functional J^ε .

4.4. Reconstruction of the means.

4.4.1. Small observational noise. In Figure 4 we present sample means associated with small-noise observations for the phase-field, level set, and GP regression

FIG. 3. Examples of samples near the start of chains for Truth B with small observational noise. Sample numbers 10000, 20000, 30000, 40000, and 50000 are shown from left-to-right for the phase-field chain (top) and the level set chain (bottom).

models, both with and without thresholding by S . Note that the phase-field and GP regression models attempt to fit the unthresholded field to the data points, whereas the level set method attempts to fit the thresholded field; the unthresholded field for the level set method is hence on a different scale than the other two models.

For Truth A and Truth B, the general quality of the reconstruction is similar for all three models after thresholding, though the level set method does not overfit to the datapoints as significantly as the other two methods; this overfitting for the phase field and GP regression is manifest in a boundary for the largest inclusion in Truth B which has variations on the scale of the observational noise. Another noticeable effect in the quality of the phase field and GP regression, manifest in Truth A, is that the edges of the circular inclusion are rendered approximately piecewise linear; this might be ameliorated by using a small mesh increment to ε ratio. The level set method has no small length scale to resolve and, hence, does not suffer from this effect.

For all three models reconstruction of Truth C is fairly inaccurate as the sparse observation network does not resolve the length scale on which the true field varies. The level set and GP regression models perform similarly, whereas the phase-field model places much more mass into the positive class; it is likely that this reflects a lack of convergence of the Markov chain for the phase-field model.

4.4.2. Order one noise reconstructions. In Figure 5 the sample means associated with order one observational noise are shown. As would be expected, reconstruction quality is generally poorer than for the small noise observations, though overfitting to the observational noise is no longer an issue for the phase-field and GP approaches. The three models perform similarly, though there seems to be an increased amount of penalization on the length of the interface from left-to-right. Without thresholding, the GP regression means provide poor estimates of the truth in terms of scale, due to the far weaker influence of the likelihood and lack of prior information enforcing values close to ± 1 .

4.5. Perimeter learning. Here we study perimeter learning for the Bayesian level set method. The length of the zero level set of η^N may then be approximated by using the discrete variation of $w^N := S(\eta^N)$,

$$\ell(N) = \frac{1}{2N^2} \sum_{i,j=1}^N |D^N w^N(x_i, y_j)| \approx \frac{1}{2} \int_D |\nabla w^N(x, y)|,$$

FIG. 4. Sample means for Truth A (top block), Truth B (middle block), and Truth C (bottom block) with small observational noise. The top row of each block shows Monte Carlo approximations to $\mathbb{E}^{\nu^y}(v)$, the underlying continuous fields, and the bottom row in each block shows Monte Carlo approximations to $S(\mathbb{E}^{\nu^y}(v))$, the thresholded fields.

FIG. 5. Sample means for Truth A (top block), Truth B (middle block), and Truth C (bottom block) with order one observational noise. The top row of each block shows Monte Carlo approximations to $\mathbb{E}^{\nu^y}(v)$, the underlying continuous fields, and the bottom row in each block shows Monte Carlo approximations to $S(\mathbb{E}^{\nu^y}(v))$, the thresholded fields.

FIG. 6. The dependence of the length of the zero level set of a Gaussian sample, as a function of numerical approximation level N and of prior regularity parameter α . (Left) logarithmic axis, (right) linear axis.

where the operator D^N approximates the gradient on the grid $\{x_i, y_j\}$ via central differences. Using this we investigate numerically the length of the level sets of these samples. We have shown in Lemma 2.1 that choosing $\alpha > 1 + d/2$ is sufficient to ensure the almost sure finite length of level sets. Numerical experiments indicate that this is a sharp result. In Figure 6 interface lengths for prior distributions approximated as described above for a single realization of $\{\xi_k\}$ in (4.1) are displayed as a function of N for varying values of α . We use $d = 2$ and observe that for $\alpha < 2$ the length of the interface diverges algebraically with N (left-hand panel), for $\alpha = 2$ it diverges logarithmically (right-hand panel shows this best), and for $\alpha > 2$ it converges to a constant (both left- and right-hand panels show this). The results, then, suggest that level sets have finite length almost surely if and only if $\alpha > 1 + d/2$.

In order to compare the perimeter distribution between the prior and posterior, the choice $\alpha > 1 + d/2$ ensures that the length of the zero level set is well-defined so in two dimensions we take $\alpha = 3$. The results for recovery of Truth B are shown in Figure 7. While the perimeter still retains some variation under the posterior, the variation is much lower and, in contrast to the prior, it is concentrated around the true value of the perimeter. We see that though the Bayesian level set approach does not explicitly penalize the perimeter, it has the ability to estimate the perimeter, and quantify uncertainty in the estimation.

5. Eikonal equation. In this section we build on what we have learned so far for the linear inverse problem defined by (1.1) and use it to study a nonlinear inverse problem from the eikonal equation. GP regression is fast to implement, and appears to give qualitatively comparable accuracy to the Bayesian level set method; but it does not generalize to nonlinear problems and so we do not consider it further. The experiments in the previous section, set-up in subsection 4.1, suggest that the Bayesian level set approach to binary recovery has two advantages over the Bayesian phase-field formulation, for the linear inverse problem considered: the level set method is faster and draws from the posterior contain information about the true perimeter. Thus we focus attention purely on the Bayesian level set method. Within this context we also demonstrate the benefits of hierarchical Bayesian inversion.

5.1. The forward equation. Let $x_0 \in D$ be a wave emitting source, and define the first arrival time of the wave at $x \in D$ as $T(x)$. The wave passes through a medium which adjusts the wave speed according to a scalar function $u: D \rightarrow \mathbb{R}_+$ known as

FIG. 7. The distribution of the perimeter under the prior and posterior distributions. The vertical dashed line indicates the perimeter of the true field.

the *slowness*. It is shown in [49, 62] that $T(x)$ may be viewed as the solution of a stationary Hamilton–Jacobi equation, namely, the following eikonal equation:

$$(5.1) \quad |\nabla T(x)| = u(x) \quad \forall x \in D \setminus \{x_0\},$$

$$(5.2) \quad T(x_0) = 0,$$

$$(5.3) \quad \nabla T(y) \cdot n(y) \geq 0 \quad \forall y \in \partial D,$$

where n is the outward pointing unit normal. The Soner boundary condition (5.3) ensures ray paths terminate at ∂D [62]. The recovery of the slowness function u from observations of arrival times is known as first arrival traveltime tomography. Extensive discussion of the well-posedness of the forward problem can be found in [24, 25, 30]. For our application, we consider a binary slowness function $u: D \rightarrow \{u_{\min}, u_{\max}\}$, where $0 < u_{\min} \leq u_{\max}$.

We define the solution map G mapping the slowness u to the travel times T via solution of the eikonal equation with source x_0 . Because we are interested in binary slowness functions we also introduce

$$(5.4) \quad S(v) = S(v) \cdot (u_{\max} - u_{\min})/2 + (u_{\max} + u_{\min})/2;$$

here $S(v)$ is the sign function defined in subsection 2.1.

To solve the forward problem we first discretize using an upwind finite difference scheme. We then use a fast marching procedure (see [59]) to solve the discrete eikonal equation. A formulation of the discretization and marching algorithm, along with a proof of numerical convergence, is found in [25].

5.2. The Bayesian inverse problem. Let $\eta \sim \mathcal{N}(0, \Sigma)$ be a normal random variable in $\mathbb{R}^J$. Fix $x_0 \in \bar{D}$. Defining the observation map K taking traveltimes to $\mathbb{R}^J$ we define the inverse problem of finding v , given observed data y satisfying

$$(5.5) \quad y = K \circ G \circ S(v) + \varepsilon^c \eta = K \circ G(u) + \varepsilon^c \eta.$$

We will assume that K is defined so that the data y is a set of observed first hitting times at fixed known receiver locations $\{z_j\}_{j=1}^J \in \bar{D}$. The random variable $y|v \sim \mathcal{N}(K \circ G \circ S(v), \varepsilon^{2c} \Sigma)$, leading to the negative log-likelihood defined, up to an additive

constant, by

$$\Phi(v; y) = \frac{1}{2\varepsilon^{2c}} \left| \Sigma^{-\frac{1}{2}} (y - K \circ G \circ S(v)) \right|^2.$$

We will treat problems of multiple sources $\{x_0^m\}_{m=0}^M$ as multiple experiments, with solution maps G^m each producing data $y^m \in \mathbb{R}^J$. The natural extension is to consider the following negative log-likelihood,

$$(5.6) \quad \Phi(v; y) = \frac{1}{2\varepsilon^{2c}} \sum_{m=1}^M \left| \Sigma^{-\frac{1}{2}} (y^m - K \circ G^m \circ S(v)) \right|^2$$

for the m observations $y = (y^1, \dots, y^m)$. We assume that we are given prior measure $\zeta_0 = \mathcal{N}(0, C)$ and let $\zeta^y(dv)$ denote the probability distribution of the conditioned random variable $v|y$. Then using Bayes' theorem (see Proposition 2.2), we deduce that ζ^y is a probability measure supported on continuous functions, determined by

$$\frac{d\zeta^y}{d\zeta_0} = \frac{1}{Z} \exp \left(-\frac{1}{2\varepsilon^{2c}} \sum_{m=1}^M \left| \Sigma^{-\frac{1}{2}} (y^m - K \circ G^m \circ S(v)) \right|^2 \right)$$

with normalization constant Z . This problem has been formulated for a piecewise constant slowness function with Ginzburg–Landau type regularization in the deterministic setting [30], where simulations and proofs of convergence of numerical schemes can be found. We choose here to use a level set formulation for the reasons discussed at the start of the section.

5.2.1. Hierarchical inference for inverse length scale. In hierarchical Bayesian inference we employ a Gaussian prior measure $\zeta_0 = \mathcal{N}(0, C(\tau))$ in which the covariance C depends on parameter $\tau > 0$ which we interpret as an additional unknown to be learned during the inversion process. In particular we will work with settings in which τ has an interpretation as an inverse length scale. To this end, define the hierarchical prior ζ_0 by decomposing as follows:

$$\zeta_0(dv, d\tau) = \zeta_0(dv|\tau) \pi_0(\tau) d\tau.$$

We call $\pi_0(\tau)$ the hyperprior. Generalizing the derivation of the posterior in the preceding subsection, we now find that the distribution of $v, \tau|y$ is determined by probability measure $\zeta^y(dv, d\tau)$ defined by

$$(5.7) \quad \zeta^y(dv, d\tau) = \frac{1}{Z} \exp(-\Phi(v; y)) \zeta_0(dv|\tau) \pi_0(\tau) d\tau$$

for normalization constant Z . For reasons discussed in [53, 68] it can be advantageous to reparameterize the hierarchical inverse problem in terms of (ξ, τ) , rather than (v, τ) , where ξ is a Gaussian white noise distributed as $\mathcal{N}(0, I)$, so that $v = \sqrt{C(\tau)}\xi$; the underlying latent Gaussian white noise ξ may be identified with the collection of i.i.d. unit Gaussians $\{\xi_k\}$ used to construct prior samples in subsection 4.2. Abusing notation we may write the posterior distribution ζ^y , now for the variable (ξ, τ) , as

$$(5.8) \quad \zeta^y(d\xi, d\tau) = \frac{1}{Z} \exp(-\Phi(\sqrt{C(\tau)}\xi; y)) \zeta_0(d\xi) \pi_0(\tau) d\tau.$$

Working with variables (v, τ) as in (5.7) is referred to as the *centered* problem formulation; using variables (ξ, τ) as in (5.8) is referred to as the *noncentered* problem formulation.

FIG. 8. The sign of fields v , where $S(v)$ (for appropriate $u_{\min}$, $u_{\max}$) will be used as a truth for inversion; the field on the left will be referred to as Truth D and the field on the right as Truth E. The sets of observation ($\times$) and source ($+$) points are shown in each figure.

Numerical evidence described in [31, 15] demonstrates that, for level set based inverse problems, use of the noncentered formulation in (5.8) confers considerable advantages in terms of speed of convergence of MCMC. We thus employ the noncentered formulation. To sample from this distribution we use [15, Algorithm 6.1], known as the noncentered pCN-within-Gibbs. This algorithm updates ξ and τ in separate substeps of a pCN sampling method, and we perform this in a random order each iteration.

5.2.2. Hierarchical inference for contrast. We also consider a model hyperparameter that originates in the application (whereas τ appears when regularizing through the prior). We investigate the contrast between the binary materials, and rewriting the relationship between u and v , we obtain

$$S(v) = \frac{\kappa}{2}(S(v) + 1) + u_{\min},$$

where $\kappa = (u_{\max} - u_{\min}) > 0$ is the contrast. Knowing the parameter $u_{\min}$, we consider the parameter κ as a positive random variable to be learned from the data. This increases flexibility of techniques in application as we can apply our methods to scenarios where slowness contrast is uncertain. We may include this in a noncentered pCN-within-Gibbs algorithm as above, by exploiting the form

$$S(v) = \frac{\kappa}{2}(S(\sqrt{C(\tau)}\xi) + 1) + u_{\min},$$

substituting this into the likelihood (5.6). We update each of ξ , τ , κ separately and in random order in each iteration.

5.2.3. Numerical results: Length scale hyperparameter only. For this first test we wish to demonstrate recovery of hyperparameters with the nonlinear eikonal forward model as described in subsection 5.2.1. Our domain is given by $D = [0, 1] \times [0, 1]$. We choose Truth D as seen in Figure 8; here the truth has been produced by applying the slowness function $S(v^*)$ for an instance $v^* \sim \mu_0(\tau^*)$ with hyperparameter $\tau^* = \exp(6.5)$. We take $u_{\min} = 1.0$, and a known contrast $\kappa = 0.2$ (thus $u_{\max} = 1.2$).

FIG. 9. Sample means for truth D. The left shows the Monte Carlo approximation of $\mathbb{E}^{\mu^y}(v)$, the underlying continuous field. The right shows the Monte Carlo approximation of $S(\mathbb{E}^{\mu^y}(v))$, the thresholded and scaled field.

The discretization uses an equidistributed mesh with grid spacing $h = 128^{-1}$. To avoid an inverse crime we produce the data on a numerical mesh with spacing $h/2$. We take 4 sources $\{x_0^m\}_{m=1}^4$ and 100 receivers uniformly distributed in the (coarser) discrete domain.

We work with the same parameters as taken in subsection 4.1.1 with the exception of setting $c = 1$ and, of course, viewing τ as an unknown. For consistency across different source–receiver combinations we additionally scale the noise by the range of the traveltimes observations, so the effective observational noise is 10^{-2} . We take $\alpha = 3$ to ensure finite perimeter.

For the hyperparameter τ , we choose a lognormal prior to ensure positivity, we take prior $\pi(\log \tau) = \mathcal{N}(5, 2.5)$, and initialize the Markov chain at $\log(3)$. We use the noncentered pCN-within-Gibbs algorithm [15, Algorithm 6.1] described in subsection 5.2.1 with a random walk Metropolis proposal for the hyperparameter step. The step sizes were chosen to achieve 20%–25% acceptance rates for ξ and τ .

The result of the recovery of truth D after 2×10^5 iterations run is displayed in Figure 9. The recovery of the field is fairly faithful; we note that information can only be learned between source and receiver pairs; see Truth D in Figure 8 for their random distribution. We additionally provide the distribution of the perimeter in Figure 10 with the true perimeter marked, and one can see that this falls well within the posterior distribution with high probability mass given to a small neighborhood of the truth. In Figure 11 we see the results of learning the hyperparameter distribution. We see the marked truth and the prior and posterior distribution and, again, we see the posterior peaks near the true value and gives a large mass to a small neighborhood of the truth. We note that the parameter $\log \tau$ is sampling close to its posterior within 10^4 iterations in this example.

5.3. Numerical results: Length scale and contrast hyperparameters.

For this test we demonstrate recovery of the contrast in the medium. We assume that we are in the situation of performing the inverse eikonal problem where we do not have exact information on the contrast between the binary phases. We choose Truth E as seen in Figure 8, comprising four circles; three of diameter 0.1 and one of diameter 0.15. In the figure we see the choice of four sources, and take 15^2 equally

FIG. 10. The distribution of the perimeter under the prior and posterior distributions for truth D . The vertical dashed line indicates the perimeter of the true field.

FIG. 11. The distribution of the logarithm of the hyperparameter τ under the prior and posterior distributions for truth D . The vertical dashed line indicates the hyperparameter value $\log \tau = 6.5$ of the true field.

spaced receivers over the domain. We take $u_{\min} = 1.0$, and again we work with the parameters as detailed in subsection 4.1.1 with the exception of setting $c = 1$, and viewing τ as unknown and, as in the previous subsection, we set $\alpha = 3$. We assume the contrast is lognormal with prior $\pi(\log \kappa) = \mathcal{N}(\log(0.2), 0.3)$, as we require a positive prior, we also treat τ as a random unknown quantity, and take a lognormal $\pi(\log \tau) = \mathcal{N}(5, 2.5)$. We initialize the MCMC method to sample these variables at $\log \tau = 8$ and $\log \kappa = \log(0.1)$.

Our recovery of Truth E after 2×10^5 iterations is recorded in Figure 12. We see an excellent recovery of the simple geometry, in particular, the color scale shows the approximation of the mean recovered contrast to the true contrast, where $u_{\min} = 1$, $u_{\max} = 1.2$. This recovery is demonstrated in Figure 13 where we see the prior and posterior densities for the interfacial length. The posterior places most weight in a small neighborhood of the truth, and peaks nearby.

FIG. 12. Sample means for truth E . The left shows the Monte Carlo approximation of $\mathbb{E}^{\mu^y}(v)$, the underlying continuous field. The right shows the Monte Carlo approximation of $S(\mathbb{E}^{\mu^y}(v))$, the thresholded and scaled field.

FIG. 13. The distribution of the perimeter under the prior and posterior distributions for truth E . The vertical dashed line indicates the perimeter of the true field.

The hyperparameter recovery is displayed in Figures 14 and 15. We see the profiles for prior and posterior for $\log \tau$ in Figure 14, demonstrating that the posterior concentrates within a sensible range; note that there is no true τ for this example. In Figure 15 we see the contrast recovery, and observe that the posterior places large weight close to the true value. It again takes less than 10^4 iterations for both hyperparameters to draw approximately from their posterior distributions.

6. Conclusions. The paper investigates the reconciliation of perimeter and Bayesian regularization for the reconstruction of functions with interfaces, from direct or indirect noisy measurements. Three approaches are studied: Formulation 1 is based on the Bayesian level method; Formulation 2 is based on the Bayesian phase-field regularization; and Formulation 3 is based on GP regression and classification.

By studying a class of linear inverse problems we show that Formulation 2 exhibits perimeter regularization in the context of its MAP estimator, but not at the level of

FIG. 14. The distribution of the logarithm of τ under the prior and posterior distributions for truth E .

FIG. 15. The distribution of the logarithm of κ under the prior and posterior distributions for truth E . The vertical dashed line indicates the true value $\log \kappa = \log(0.2)$.

samples from the posterior distribution. Formulation 1 exhibits perimeter regularization at the level of individual samples from the posterior; there is no MAP estimator in this context. Both Formulations 1 and 2 require careful choices of constants in construction of the prior, but Formulation 2 is far more constrained in this regard. Furthermore, as a consequence of these constraints, Formulation 2 exhibits a measure concentration phenomenon meaning that MCMC based algorithms using Formulation 1 are considerably faster than those based on Formulation 2. Formulation 3 is competitive with Formulation 1 in terms of both sample properties and speed, but does not generalize beyond linear problems. We study Formulation 1 for a nonlinear inverse problem, demonstrating that it is effective in this context and, in addition, showing how hierarchical methods may be used to learn model hyperparameters appearing in the prior.

The ideas in this paper can be combined in different ways: the methodology may be extended beyond binary-valued functions to a variety of piecewise continuous problems; or other limiting functionals could be contemplated, such as Mumford–Shah

[36]; and other smoothed thresholding functions could be contemplated within the level set method, such as the double-obstacle approximation to the signum function [3, 4]. The success of the Bayesian level set method suggests that further analysis of it, as well as its deployment in new application domains, would be very valuable.

Appendix A. Proofs of main results.

Proof of Proposition 3.2. Throughout this proof C is a universal constant whose value may change between occurrences. To apply Theorem 4.12 from [21], we need to show that the function $\Phi(\cdot, y)$ is bounded from below, is locally bounded from above, and is locally Lipschitz. We note that $\Phi(\cdot, y)$ is always nonnegative so is bounded from below. If $\|u\|_X = \max_{x \in \bar{D}} |u(x)| \leq \rho$ then we may bound $|\Phi(u, y)|$ by a constant depending on ρ , i.e., $\Phi(\cdot, y)$ is locally bounded. For the local Lipschitz continuity, we have

$$\begin{aligned} \Phi(u, y) - \Phi(v, y) &= \frac{r}{4\epsilon^b} \int_D (2 - u(x)^2 - v(x)^2)(u(x) + v(x))(u(x) - v(x)) \, dx \\ &\quad + \frac{1}{2\epsilon^{2c}} \langle \Sigma^{-\frac{1}{2}}(2y - Ku - Kv), \Sigma^{-\frac{1}{2}}K(v - u) \rangle. \end{aligned}$$

Assume that $\|u\|_X \leq \rho$ and $\|v\|_X \leq \rho$. Then, since K is a bounded linear operator on $L^1(D)$,

$$\begin{aligned} |\Phi(u, y) - \Phi(v, y)| &\leq C \int_D |u(x) - v(x)| \, dx + C|K(v - u)| \\ &\leq C\|u - v\|_{L^1(D)} \leq C|D|^{1/2}\|u - v\|_{L^2(D)} \leq C\|u - v\|_X. \end{aligned}$$

The desired result follows. $\square$

Proof of Theorem 3.4. We adapt the proof of Hilhorst, Peletier, and Schätzle [37] to allow for periodic boundary conditions and the additional L^2 norm appearing in the functional to be infimized. From Hilhorst, Peletier, and Schätzle [37] we have that if $u^\epsilon \rightarrow u$ in $L^1(D)$ then

$$\begin{aligned} \liminf_{\epsilon \rightarrow 0} I^\epsilon(u^\epsilon) &\geq \liminf_{\epsilon \rightarrow 0} \int_D \left(\frac{1}{2} \delta \epsilon^3 |\Delta u^\epsilon|^2 + \frac{1}{2} \delta q \epsilon |\nabla u^\epsilon|^2 + \frac{r}{4\epsilon} (1 - u^\epsilon(x)^2)^2 \right) \, dx \\ &\quad + \frac{1}{2} |\Sigma^{-\frac{1}{2}}(y - Ku^\epsilon)|^2 \\ &\geq I_0^\delta(u). \end{aligned}$$

Now we show that for each $u \in L^1(D)$, there is a sequence $\{u^\epsilon\} \subset H_\#^2(D)$ which converges strongly to u in $L^1(D)$ such that $\limsup_{\epsilon \rightarrow 0} I^\epsilon(u^\epsilon) \leq I_0^\delta(u)$. We first review the main points in the proof of Hilhorst, Peletier, and Schätzle [37] for functions $u \in H^2(D)$. Considering the case $I^\delta(u) < \infty$, without loss of generality, we assume that

$$u = \mathbb{1}_Q - \mathbb{1}_{\mathbb{R}^d \setminus Q},$$

where Q is a bounded domain, with $\partial Q \in C^\infty$ and $Q \subset \subset D$. The sign distance function is defined as

$$d(x) = \begin{cases} +\inf_{y \in \partial Q} |x - y| & \text{if } x \in Q, \\ -\inf_{y \in \partial Q} |x - y| & \text{if } x \notin Q. \end{cases}$$

Let N_h be an h neighborhood of ∂Q (we choose h so that h is less than the distance between ∂Q and ∂D). We choose a function $\eta \in C^2(\bar{D})$ such that $\eta(x) = d(x)$ for

$x \in N_h$, $\eta(x) \geq h$ when $x \in Q \setminus N_h$ and $\eta(x) \leq -h$ when $x \in D \setminus (Q \cup N_h)$. Let U be an odd minimizer of the functional $e^\delta(U)$ with $\lim_{t \rightarrow \infty} U(t) = 1$ and $\lim_{t \rightarrow -\infty} U(t) = -1$. We let

$$u^\varepsilon = U\left(\frac{\eta(x)}{\varepsilon}\right).$$

We note that $u^\varepsilon(x)$ is uniformly bounded pointwise and $u^\varepsilon(x) \rightarrow u(x)$ for all $x \in D$. From the Lebesgue dominated convergence theorem, $u^\varepsilon \rightarrow u$ in $L^1(D)$ and in $L^2(D)$. Thus

$$\lim_{\varepsilon \rightarrow 0} |\Sigma^{-\frac{1}{2}}(y - Ku^\varepsilon)|^2 = |\Sigma^{-\frac{1}{2}}(y - Ku)|^2$$

and, since $a > 0$,

$$\lim_{\varepsilon \rightarrow 0} \int_D \delta \tau^2 \varepsilon^a u^\varepsilon(x)^2 dx = 0.$$

To show that $\lim_{\varepsilon \rightarrow 0} I^\varepsilon(u) = I_0^\delta(u)$, we follow the approach of Hilhorst, Peletier, and Schätzle [37]. The integral

$$\int_D \left(\frac{1}{2} \delta \varepsilon^3 |\Delta u^\varepsilon|^2 + \frac{1}{2} \delta q \varepsilon |\nabla u^\varepsilon|^2 + \frac{r}{4\varepsilon} (1 - u^\varepsilon(x)^2)^2 \right) dx$$

is written as

$$\begin{aligned} & \int_{D \setminus N_h} \left(\frac{1}{2} \delta \varepsilon^3 |\Delta u^\varepsilon|^2 + \frac{1}{2} \delta q \varepsilon |\nabla u^\varepsilon|^2 + \frac{r}{4\varepsilon} (1 - u^\varepsilon(x)^2)^2 \right) dx \\ & + \int_{N_h} \left(\frac{1}{2} \delta \varepsilon^3 |\Delta u^\varepsilon|^2 + \frac{1}{2} \delta q \varepsilon |\nabla u^\varepsilon|^2 + \frac{r}{4\varepsilon} (1 - u^\varepsilon(x)^2)^2 \right) dx. \end{aligned}$$

Using the exponential decay of U, U' , and U'' at ∞ and $-\infty$, we deduce that the integral over $D \setminus N_h$ goes to 0 when $\varepsilon \rightarrow 0$ (note that $|\eta(x)|/\varepsilon > h/\varepsilon$ which goes to ∞ when $\varepsilon \rightarrow 0$ for $x \in D \setminus N_h$). The integral over N_h is shown to converge to $I^\delta(u)$ when $\varepsilon \rightarrow 0$.

To adapt this proof of Hilhorst, Peletier, and Schätzle [37] to functions with periodic boundary condition on D , we only need to choose the function η so that η is periodic and $\eta(x) \geq h$ when $x \in Q \setminus N_h$ and $\eta(x) \leq -h$ when $x \in D \setminus (Q \cup N_h)$. Such a function can be constructed as follows. Let $\psi(x) \in C_0^\infty(D)$ be such that $\psi(x) = 1$ when x is in a neighborhood of $Q \cup N_h$, and $0 \leq \psi(x) \leq 1$ for all $x \in D$. Let $\eta_1(x)$ be a smooth periodic function with $\eta_1(x) \leq -h$ for all $x \in D$. Using the function η of Hilhorst, Peletier, and Schätzle [37], we define a new function

$$\bar{\eta}(x) = \psi(x)\eta(x) + (1 - \psi(x))\eta_1(x).$$

The function $\bar{\eta}(x)$ satisfies the requirement. $\square$

REFERENCES

- [1] 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.
- [2] A. BESKOS, G. O. ROBERTS, A. M. STUART, AND J. VOSS, *MCMC methods for diffusion bridges*, Stoch. Dyn., 8 (2008), pp. 319–350.
- [3] J. BLOWEY AND C. ELLIOTT, *Curvature dependent phase boundary motion and parabolic double obstacle problems*, in Degenerate Diffusions, Springer, New York, 1993, pp. 19–60.

- [4] J. BLOWEY AND C. ELLIOTT, *A phase-field model with a double obstacle potential*, in Motion by Mean Curvature and Related Topics (Trento, 1992), de Gruyter, Berlin, 1994, pp. 1–22.
- [5] C. BRETT, C. M. ELLIOTT, AND A. S. DEDNER, *Phase field methods for binary recovery*, in Optimization with PDE Constraints, Springer, Cham, Switzerland, 2014, pp. 25–63.
- [6] 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.
- [7] D. CALVETTI AND E. SOMERSALO, *A Gaussian hypermodel to recover blocky objects*, Inverse Problems, 23 (2007), pp. 733–754.
- [8] D. CALVETTI AND E. SOMERSALO, *Hypermodels in the Bayesian imaging framework*, Inverse Problems, 24 (2008), 034013.
- [9] D. CALVETTI, E. SOMERSALO, AND A. STRANG, *Hierarchical Bayesian models and sparsity: ℓ_2 -magic*, Inverse Problems, 35 (2019), 035003.
- [10] M. CARDIFF AND P. KITANIDIS, *Bayesian inversion for facies detection: An extensible level set framework*, Water Resour. Res., 45 (2009), WR007675.
- [11] M. CARRIERO, A. LEACI, AND F. TOMARELLI, *A survey on the Blake–Zisserman functional*, Milan J. Math., 83 (2015), pp. 397–420.
- [12] J. N. CARTER AND D. A. WHITE, *History matching on the Imperial College fault model using parallel tempering*, Comput. Geosci., 17 (2013), pp. 43–65.
- [13] T. F. CHAN AND X.-C. TAI, *Level set and total variation regularization for elliptic inverse problems with discontinuous coefficients*, J. Comput. Phys., 193 (2004), pp. 40–66.
- [14] H. CHANG, D. ZHANG, AND Z. LU, *History matching of facies distribution with the ENKF and level set parameterization*, J. Comput. Phys., 229 (2010), pp. 8011–8030, <https://doi.org/10.1016/j.jcp.2010.07.005>.
- [15] V. CHEN, M. M. DUNLOP, O. PAPASPILIOPOULOS, AND A. M. STUART, *Dimension-Robust MCMC in Bayesian Inverse Problems*, manuscript, <https://arxiv.org/abs/1803.03344>, 2019.
- [16] R. CHOKSI, Y. VAN GENNIP, AND A. OBERMAN, *Anisotropic Total Variation Regularized L^1 -Approximation and Denoising/Deblurring of 2d Bar Codes*, preprint, <https://arxiv.org/abs/1007.1035>, 2010.
- [17] C. CLASON, T. HELIN, R. KRETSCHMANN, AND P. PIROINEN, *Generalized modes in Bayesian inverse problems*, SIAM/ASA J. Uncertain. Quantif., 7 (2019), pp. 652–684.
- [18] J. COCKAYNE, C. J. OATES, T. J. SULLIVAN, AND M. GIROLAMI, *Bayesian probabilistic numerical methods*, SIAM Rev., 61 (2019), pp. 756–789.
- [19] A. COHEN AND J.-P. D’ALES, *Nonlinear approximation of random functions*, SIAM J. Appl. Math., 57 (1997), pp. 518–540.
- [20] S. L. COTTER, G. O. ROBERTS, A. M. STUART, D. WHITE, ET AL., *MCMC methods for functions: Modifying old algorithms to make them faster*, Statist. Sci., 28 (2013), pp. 424–446.
- [21] M. DASHTI, K. J. LAW, A. M. STUART, AND J. VOSS, *MAP estimators and their consistency in Bayesian nonparametric inverse problems*, Inverse Problems, 29 (2013), 095017.
- [22] M. DASHTI AND A. M. STUART, *The Bayesian approach to inverse problems*, in Handbook of Uncertainty Quantification, Springer, Cham, Switzerland, 2017, pp. 311–428.
- [23] K. DECKELNICK, G. DZIUK, AND C. M. ELLIOTT, *Computation of geometric partial differential equations and mean curvature flow*, Acta Numer., 14 (2005), pp. 139–232.
- [24] K. DECKELNICK AND C. M. ELLIOTT, *Uniqueness and error analysis for Hamilton–Jacobi equations with discontinuities*, Interfaces Free Bound., 6 (2004), pp. 329–349.
- [25] K. DECKELNICK, C. M. ELLIOTT, AND V. STYLES, *Numerical analysis of an inverse problem for the Eikonal equation*, Numer. Math., 119 (2011), 245.
- [26] K. DECKELNICK, C. M. ELLIOTT, AND V. STYLES, *Double obstacle phase field approach to an inverse problem for a discontinuous diffusion coefficient*, Inverse Problems, 32 (2016), 045008.
- [27] P. DIACONIS, *Bayesian numerical analysis*, in Statistical Decision Theory and Related Topics IV, Springer, New York, 1988, pp. 163–175.
- [28] O. DORN AND D. LESSELIER, *Level set methods for inverse scattering—some recent developments*, Inverse Problems, 25 (2009), 125001, <http://stacks.iop.org/0266-5611/25/i=12/a=125001>.
- [29] O. DORN AND R. VILLEGAS, *History matching of petroleum reservoirs using a level set technique*, Inverse Problems, 24 (2008), 035015, <http://stacks.iop.org/0266-5611/24/i=3/a=035015>.
- [30] O. DUNBAR AND C. M. ELLIOTT, *Binary recovery via phase field regularization for first-arrival traveltime tomography*, Inverse Problems, 35 (2019), 095004.
- [31] M. M. DUNLOP, M. A. IGLESIAS, AND A. M. STUART, *Hierarchical Bayesian level set inversion*, Statist. Comput., 27 (2017), pp. 1555–1584.
- [32] H. W. ENGL, M. HANKE, AND A. NEUBAUER, *Regularization of Inverse Problems*, Math. Appl. 375, Kluwer, Dordrecht, The Netherlands, 1996.

- [33] M. HAIRER, A. M. STUART, AND S. J. VOLLMER, *Spectral gaps for Metropolis-Hastings algorithms in infinite dimensions*, Ann. Appl. Probab., 24 (2014), pp. 2455–290, <https://doi.org/10.1214/13-AAP982>.
- [34] P. C. HANSEN, J. G. NAGY, AND D. P. O’LEARY, *Deblurring Images: Matrices, Spectra, and Filtering*, Fundam. Algorithms, SIAM, Philadelphia, 2006.
- [35] T. HELIN AND M. BURGER, *Maximum a posteriori probability estimates in infinite-dimensional Bayesian inverse problems*, Inverse Problems, 31 (2015), 085009.
- [36] T. HELIN AND M. LASSAS, *Hierarchical models in statistical inverse problems and the Mumford–Shah functional*, Inverse Problems, 27 (2011), 015008.
- [37] D. HILHORST, L. A. PELETIER, AND R. SCHÄTZLE, *γ -limit for the extended Fisher–Kolmogorov equation*, Proc. Roy. Soc. Edinburgh Sect. A, 132 (2002), pp. 141–162.
- [38] B. HOSSEINI, *Well-posed Bayesian inverse problems with infinitely divisible and heavy-tailed prior measures*, SIAM/ASA J. Uncertain. Quantif., 5 (2017), pp. 1024–1060.
- [39] M. IGLESIAS, K. LIN, AND A. STUART, *Well-posed Bayesian geometric inverse problems arising in subsurface flow*, Inverse Problems, 30 (2014), 114001, <https://doi.org/doi:10.1088/0266-5611/30/11/114001>.
- [40] M. A. IGLESIAS, K. J. LAW, AND A. M. STUART, *Evaluation of Gaussian approximations for data assimilation in reservoir models*, Comput. Geosci., 17 (2013), pp. 851–885.
- [41] M. A. IGLESIAS, Y. LU, AND A. M. STUART, *A Bayesian level set method for geometric inverse problems*, Interfaces Free Bound., 18 (2016), pp. 181–217.
- [42] M. A. IWEN, F. SANTOSA, AND R. WARD, *A symbol-based algorithm for decoding bar codes*, SIAM J. Imaging Sci., 6 (2013), pp. 56–77.
- [43] R. JIN, S. ZHAO, X. XU, E. SONG, AND C.-C. HUNG, *Super-resolving barcode images with an edge-preserving variational Bayesian framework*, J. Electron. Imaging, 25 (2016), 033016.
- [44] J. KAIPIO AND E. SOMERSALO, *Statistical and computational inverse problems*, Appl. Math. Sci. 160, Springer, New York, 2006.
- [45] M. F. KRATZ, *Level crossings and other level functionals of stationary Gaussian processes*, Probab. Surv., 3 (2006), pp. 230–288.
- [46] M. LASSAS, E. SAKSMAN, AND S. SILTANEN, *Discretization-invariant Bayesian inversion and Besov space priors*, Inverse Probl. Imaging, 3 (2009), pp. 87–122.
- [47] M. LASSAS AND S. SILTANEN, *Can one use total variation prior for edge-preserving Bayesian inversion?*, Inverse Problems, 20 (2004), pp. 1537–1564.
- [48] J. LEE AND P. KITANIDIS, *Bayesian inversion with total variation prior for discrete geologic structure identification*, Water Res. Res., 49 (2013), pp. 7658–7669.
- [49] P.-L. LIONS, *Generalized Solutions of Hamilton-Jacobi Equations*, Res. Notes Math. 69, Pitman, Boston, MA, 1982.
- [50] D. MUMFORD AND J. SHAH, *Optimal approximations by piecewise smooth functions and associated variational problems*, Comm. Pure Appl. Math., 42 (1989), pp. 577–685.
- [51] E. NIEMI, M. LASSAS, A. KALLONEN, L. HARHANEN, K. HÄMÄLÄINEN, AND S. SILTANEN, *Dynamic multi-source x-ray tomography using a spacetime level set method*, J. Comput. Phys., 291 (2015), pp. 218–237.
- [52] H. OWHADI AND C. SCOVEL, *Operator-Adapted Wavelets, Fast Solvers, and Numerical Homogenization: From a Game Theoretic Approach to Numerical Approximation and Algorithm Design*, Cambridge Monogr. Appl. Comput. Math. 35, Cambridge University Press, Cambridge, 2019.
- [53] O. PAPASPILIOPOULOS, G. O. ROBERTS, AND M. SKÖLD, *A general framework for the parametrization of hierarchical models*, Statist. Sci. 22, (2007), pp. 59–73.
- [54] R. RAMLAU AND W. RING, *Regularization of ill-posed Mumford–Shah models with perimeter penalization*, Inverse Problems, 26 (2010), 115001.
- [55] G. RIOUX, C. SCARVELIS, R. CHOKSI, T. HOHEISEL, AND P. MARECHAL, *Blind deblurring of barcodes via Kullback-Leibler divergence*, IEEE Trans. Pattern Anal. Mach. Intell., to appear.
- [56] J. C. ROBINSON, *Infinite-dimensional dynamical systems: An introduction to dissipative parabolic PDEs and the theory of global attractors*, Cambridge Texts Appl. Math. 28, Cambridge University Press, Cambridge, 2001.
- [57] L. I. RUDIN, S. OSHER, AND E. FATEMI, *Nonlinear total variation based noise removal algorithms*, Phys. D, 60 (1992), pp. 259–268.
- [58] F. SANTOSA, *A level-set approach for inverse problems involving obstacles*, ESAIM Control Optim. Calc. Var., 1 (1996), pp. 17–33.
- [59] J. A. SETHIAN, *Fast marching methods*, SIAM Rev., 41 (1999), pp. 199–235.
- [60] J. A. SETHIAN, *Level Set Methods and Fast Marching Methods: Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision, and Materials Science*, Cambridge Monogr. Appl. Comput. Math. 3, Cambridge University Press, Cambridge, 1999.

- [61] I. SIVAK, *Bayesian Reconstruction of Piecewise Constant Signals*, MSc. Dissertation, Warwick University, Coventry, England, (2014).
- [62] H. M. SONER, *Optimal control with state-space constraint I*, SIAM J. Control Optim., 24 (1986), pp. 552–561.
- [63] G. SÖRÖS, S. SEMMLER, L. HUMAIR, AND O. HILLIGES, *Fast blur removal for wearable QR code scanners*, in Proceedings of the 2015 ACM International Symposium on Wearable Computers, ACM, New York, 2015, pp. 117–124.
- [64] A. M. STUART, *Inverse problems: A Bayesian perspective*, Acta Numer., 19 (2010), pp. 451–559.
- [65] Y. VAN GENNIP, P. ATHAVALÉ, J. GILLES, AND R. CHOKSI, *A regularization approach to blind deblurring and denoising of QR barcodes*, IEEE Trans. Image Process., 24 (2015), pp. 2864–2873.
- [66] C. K. WILLIAMS AND C. E. RASMUSSEN, *Gaussian Processes for Machine Learning*, Vol. 2, MIT Press Cambridge, MA, 2006.
- [67] Z. YAO, Z. HU, AND J. LI, *A TV-Gaussian prior for infinite-dimensional Bayesian inverse problems and its numerical implementations*, Inverse Problems, 32 (2016), 075006.
- [68] Y. YU AND X.-L. MENG, *To center or not to center: that is not the question – an ancillarity-sufficiency interweaving strategy (ASIS) for boosting MCMC efficiency*, J. Comput. Graph. Statist., 20 (2011), pp. 531–570.