

RECONCILING BAYESIAN AND TOTAL VARIATION REGULARIZATION METHODS FOR BINARY INVERSION*

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. On the other hand, sparse or noisy data often demands a probabilistic approach to the reconstruction of images, to enable uncertainty quantification; the Bayesian approach to inversion is a natural framework in which to carry this out. The link between Bayesian inversion methods and perimeter regularization, however, is not fully understood. In this paper two links are studied: (i) the MAP objective function of a suitably chosen phase-field Bayesian 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 finite perimeter and to have the ability to learn about the true perimeter. Furthermore, the level set approach is shown to lead to faster algorithms for uncertainty quantification than the phase field approach.

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

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

1. Introduction.

1.1. Problem Statement. We consider the problem of recovering a function u from finite dimensional data y where

$$(1) \quad y = Ku + \varepsilon^c \eta.$$

Here $y \in \mathbb{R}^J$ denotes a finite number of observations corrupted by noise $\varepsilon^c \eta$ of size ε^c and where we assume that η is a centred Gaussian $\mathcal{N}(0, \Sigma)$. We suppose that the function $u \in BV_{\text{binary}}$ where $BV_{\text{binary}}(D) = \{\psi \in BV(D) : \psi(D) \subset \{\pm 1\}\}$. We note that this is a subset of $L^1(D)$ and we assume that the operator K is bounded and linear from $L^1(D)$ into \mathbb{R}^J . We assume that the parameter $\varepsilon \ll 1$ and we distinguish between noise which is on the same scale as the observations ($c = 0$) and small noise ($c > 0$.) Observe that from an application perspective the space BV_{binary} is a natural model for binary images.

Since the unknown observational noise η has a Gaussian distribution, and since we may also model our prior uncertainty about u through a probability distribution, it is natural to take a probabilistic approach to the recovery of u . In the Bayesian approach to inversion we impose a prior probability distribution ν_0 on the function u that we wish to reconstruct. The solution to the problem is the posterior probability

*Submitted to the editors DATE.

Funding: The research of CME was partially supported by the Royal Society via a Wolfson Research Merit Award; the work of AMS by DARPA contract contract W911NF-15-2- 0121; the work of CME and AMS by the EPSRC programme grant EQUIP; the work of MMD and AMS by AFOSR Grant FA9550-17-1-0185 and ONR Grant N00014-17-1-2079; the work of MMD by the EPSRC MASDOC Graduate Training Program; VHH gratefully acknowledges the MOE AcRF Tier 1 grant RG30/16.

[†]Computing & Mathematical Sciences, California Institute of Technology, USA, 91125 (mdunlop@caltech.edu, astuart@caltech.edu).

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

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

distribution ν^y on u given y (written $u|y$), where y is related to u via (1), and we assume that the Gaussian distribution of η is known. Bayes' theorem states that the ratio of the posterior probability on $u|y$ to the prior probability on u is proportional to the likelihood: the probability of $y|u$. Furthermore since η is assumed to be Gaussian the negative logarithm of the likelihood is proportional to a covariance weighted least squares misfit derived from $y - Ku$. We employ Bayes' theorem in a form which implies that the posterior probability measure is absolutely continuous with respect to the prior probability measure. This means that any almost sure property of the prior will be an almost sure property of the posterior; in particular this enables us to impose regularity properties on functions distributed according to the posterior by imposing those regularity properties on the prior.

1.2. Our Contribution. In interface reconstruction, classical methods have been dominated by inversion techniques which penalize the length of the perimeter between different subdomains in which the reconstruction is continuous; in particular total variation (TV) regularization has played a central role [38] and has been shown to lead to empirically effective methods which are computationally efficient. In this paper we address the question of how perimeter regularization appears within Bayesian inversion techniques for the reconstruction of binary function u . This is a notoriously difficult problem, as made transparent in the paper [36] which showed that use of discrete total variation 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 [35].

In this paper we take a different approach, trying to build connections to total variation penalization via Gaussian random fields. Two competing methodologies for the problem of representing interfaces between piecewise continuous fields are the phase-field approach [11], which introduces a length-scale over which sharp interfaces are smoothed out, and the level set method, which represents interfaces as level sets of continuous fields [41]. Both of these approaches have been used to formulate optimization approaches to inverse problems where the target unknown function is piecewise continuous [4, 18, 39]. On the other hand increase in computer power has started to render Bayesian inversion techniques tractable in some applications [33, 45, 17]. The Bayesian methodology is of interest as it enables uncertainty quantification to be performed. We study both a Bayesian phase field and a Bayesian level set approach to inversion for binary fields, making the following contributions to the understanding of Bayesian inversion and how perimeter regularization manifests in the resulting methods:

- we prove that, for appropriate choice of prior distribution, with parameters carefully scaled with respect to ε , the maximum a posteriori (MAP) estimator for a phase field Bayesian formulation has $\varepsilon \rightarrow 0$ Γ -limit which is a least squares objective function, penalized by total variation, elucidating the underlying perimeter regularization contained in the phase field Bayesian approach;
- we establish conditions under which the Bayesian level set approach leads to posterior samples with almost surely finite perimeter, and hence TV norm, demonstrating that TV penalization arises naturally out of appropriately chosen random field prior models;
- we provide numerical investigations of the properties of the two Bayesian inversion techniques demonstrating that the level set approach may be implemented quite cheaply in comparison with the phase-field approach, for

similar levels of reconstruction accuracy, and that the level set approach can learn the true perimeter.

1.3. Literature Review and Paper Overview. There are many problems in the physical sciences where piecewise constant reconstruction is of interest, for example in subsurface inversion and imaging [9, 10, 21, 29] and other problems in the physical sciences [20]; the problem of image deblurring [24] 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. We mention, however, that there is a related body of literature concerning reconstruction of QR codes from noisy observations [13, 32, 46, 44, 31]; this work exploits additional prior knowledge about the structure of the QR codes, in particular the fact that interface boundaries are aligned with known coordinate axes, which places such problems beyond the scope of our work.

There is a rich history linking probabilistic approaches to classical numerical methods [19, 43, 12, 40, 26, 5, 14]. 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 [33, 16] is then the solution of a Tikhonov-Phillips regularized least squares problem [22]. When more complex priors are used the connection between classical and Bayesian perspectives is more subtle, even for linear inverse problems [6, 25, 7, 8, 35]. Two interesting approaches to Bayesian inversion, both using thresholding as we do in this paper, may be found in [37] and [28].

In this paper we consider linear inverse problems for piecewise constant binary functions in two and three dimensions. In [section 2](#) we formulate the inverse problem of interest in a Bayesian fashion, introducing the phase-field and level set priors, both of which are non-Gaussian, and stating a well-posedness result for the resulting posterior distributions; we also discuss the properties of the length of level sets of Gaussian random fields and use this to demonstrate that the level set prior penalizes the TV norm of the binary reconstruction, for appropriately chosen parameters within the prior. [Section 3](#) characterizes the MAP estimator for the phase-field prior, demonstrating appropriate parameter scalings, in terms of assumed small noise, to obtain the desired Γ -limit for the MAP estimator in the small noise regime, using the analysis in [27]. This Γ -limit links the MAP estimator to classical TV regularization of the inverse problem. The arguments in [30] show that the MAP estimator functional for the level set method does not exist, essentially because infimizing sequences tend to zero but the point $v = 0$ does not attain the infimum. In [section 4](#) we describe numerical results which extend the foregoing discussion to full posterior exploration and demonstrate the superior efficiency of the level set representation; related one dimensional numerical results may be found in [42]. We conclude in [section 5](#), and [Appendix A](#) contains proofs of the main results.

2. Bayesian Formulations of the Inverse Problem.

2.1. Orientation. Let D be the unit cube $(0, 1)^d \subset \mathbb{R}^d$. Let $X_{k,\gamma} = C_{\#}^{k,\gamma}(\bar{D}, \mathbb{R})$ denote restriction to periodic functions of the space of real-valued functions on \bar{D} whose k^{th} derivative is Hölder- γ . And let X denote the space $C(\bar{D}, \mathbb{R})$, restricted to periodic functions and let H denote $L^2(D)$. Let $K : L^1(D) \rightarrow \mathbb{R}^J$ be a bounded linear operator. We use $|\cdot|$ to denote the Euclidean norm on \mathbb{R}^J . By virtue of continuous embedding K is also a bounded linear operator on $X_{\#}^{k,\gamma}$ for any integer

$k \geq 0$. We assume that our observational data y is given by (1) where ε and c denote constants, with $\varepsilon \ll 1$ and $c > 0$ (small noise) or $c = 0$ (order one noise). We want to reconstruct u given the prior information that it takes values in BV_{binary} . The theory in this paper easily extends to globally Lipschitz nonlinear forward maps from $L^1(D)$ into \mathbb{R}^J ; the numerical methods readily extend to quite general nonlinear forward maps, for example to nonlinear forward maps from $L^\infty(D)$ into \mathbb{R}^J as arise in recovery of coefficients in divergence form elliptic PDEs.

In the phase-field approach we will impose regularity on u via the prior – functions drawn from the prior will be almost surely continuous. But we will also penalize deviations from ± 1 by means of a phase-field weighting in the prior. In the level set approach we will construct a prior on u which ensures that it is almost surely in BV_{binary} ; we will do this by writing $u = S(v)$ for $S(\cdot)$ the signum function. We will impose prior μ_0 on v , which implies almost sure continuity of v , and compute the posterior μ^y on $v|y$. The implied prior on u , ν_0 , is the pushforward of μ_0 under S ; the implied posterior on u , ν^y , is the pushforward of μ^y under S . In subsection 2.2 we formulate the phase-field Bayesian inverse problem, whilst in subsection 2.3 we formulate the level-set based Bayesian inverse problem. Both approaches are built on Gaussian random fields and in subsection 2.4 we describe the properties of, and how to generate numerically, samples from these Gaussian random fields.

Once the posterior distribution has been defined, the computational task of finding information about it remains. In this paper we consider two approaches to this task. The first, considered in section 3, is to find the point which maximizes the posterior probability distribution, known as a MAP estimator; this leads to a problem in the calculus of variations. The second, considered in section 4, is to derive correlated samples from the posterior distribution by running a Markov chain which is ergodic with respect to the posterior distribution – the Monte Carlo-Markov chain (MCMC) approach.

2.2. Phase-Field Formulation.

2.2.1. Prior. We construct a family of priors that is supported on continuous functions $X_{\#}^{0,\gamma}$ for any $\gamma < 2 - d/2$. Furthermore the prior will be designed to concentrate on functions which, for most $x \in D$, take values close to ± 1 . We achieve these properties by working with a measure absolutely continuous with respect to a Gaussian random field. Fix constants $\delta, \tau > 0$, $q \geq 0$ and $a_1, a_2, a_3 \in \mathbb{R}$. The three parameters δ, q and τ weight the contributions from the $H_{\#}^2(D)$, $H_{\#}^1(D)$ and $L^2(D)$ terms appearing the Cameron-Martin norm of the Gaussian random field; the parameters a_1, a_2 and a_3 scale these terms with respect to powers of ε . In what follows we will show how to choose a_1, a_2 and a_3 so that the MAP estimator corresponding to this prior corresponds to a phase-field relaxation of a total variation penalized least squares problem.

We define our Gaussian measure $\mu_0 = \mathbf{N}(0, C)$ on the Hilbert space H where C is the covariance operator defined implicitly by its inverse $C^{-1} : H_{\#}^4(D) \rightarrow H$ given by the identity $f = C^{-1}u \in H$ for $u \in H_{\#}^4(D)$ and

$$f = \delta\varepsilon^{-2a_1} \Delta^2 u - q\delta\varepsilon^{-2a_2} \Delta u + \tau^2 \delta\varepsilon^{-2a_3} u.$$

It follows that E , the Cameron-Martin space of μ_0 , is $H_{\#}^2(D)$ endowed with a norm

$$\|u\|_E^2 := \langle C^{-\frac{1}{2}}u, C^{-\frac{1}{2}}u \rangle = \delta \int_D (\varepsilon^{-2a_1} |\Delta u|^2 + q\varepsilon^{-2a_2} |\nabla u|^2 + \tau^2 \varepsilon^{-2a_3} u^2) dx$$

where $\langle \cdot, \cdot \rangle$ denotes the standard $L^2(D)$ inner-product. The support of this prior is characterized in [45, Lemma 6.25], [17, Theorem 2.12]. 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 the $H_{\#}^2(D)$ contribution to the Cameron-Martin norm [42]. In what follows choice of the parameters a_i will be crucial, and will be explained below; the precise values of the positive parameters δ, q are less significant. We note that $\tau > 0$ is required to make the precision C^{-1} invertible on $L_{\#}^2(D)$. This could also be addressed with $\tau = 0$ by working on spaces of functions where the mean-value is specified (for example to be zero) in situations where this prior information is natural.

Now fix constants $r, b > 0$ and define the prior probability measure ν_0 on X via the Radon-Nikodym derivative

$$(2) \quad \frac{d\nu_0}{d\mu_0} = \frac{1}{Z_0} \exp \left(-\frac{r}{\varepsilon^b} \int_D \frac{1}{4} (1 - u(x)^2)^2 dx \right).$$

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 choice of parameter b will be crucial, and will be explained below; the precise value of the positive parameter r is less significant.

2.2.2. Likelihood. Let η be a normal random variable in \mathbb{R}^J , $\eta \sim \mathbf{N}(0, \Sigma)$ with $\Sigma \in \mathbb{R}^{J \times J}$ is the positive-definite covariance of the noise. Then the random variable $y|u$, given by (1), is distributed as the Gaussian $\mathbf{N}(Ku, \varepsilon^{2c}\Sigma)$. The negative log likelihood is then proportional to the misfit

$$\frac{1}{2} \left| \Sigma^{-\frac{1}{2}} (y - Ku) \right|^2.$$

2.2.3. Posterior. We let $\nu^y(du)$ denote the probability of the conditioned random variable $u|y$. 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

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

The following is a straightforward application of the theory in [17]:

PROPOSITION 2.1. *The posterior probability ν^y on random variable $u|y$ is a probability measure supported on $X_{\#}^{0,\gamma}$ 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. Furthermore, 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

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

2.3. Level Set Formulation.

2.3.1. Prior. We define a prior supported on functions of the form

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

where $\mathbb{1}$ denotes the characteristic function of a set, and open set D_+ has the property that $\text{leb}(\overline{D_+} \setminus D_+) = 0$. We construct this prior as follows. Define the thresholding function $S : \mathbb{R} \mapsto \{-1, 0, +1\}$ by

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

We assume that $u = S(v)$ and place the Gaussian prior $\mu_{0,\alpha} = N(0, C^{\alpha/2})$ on v for some $\alpha > \frac{d}{2}$; note that $\mu_0 = \mu_{0,2}$. This prior on v is supported on the function space $X^{k,\gamma}$ for all $\gamma \in [0, \gamma']$ where k is the largest integer such that $\gamma' := \alpha - \frac{d}{2} - k \in (0, 1]$. It induces a prior on u by push forward under S and furthermore, under this pushforward prior, $u \in \{\pm 1\}^D$ a.e., with probability 1; this is because the level sets of the Gaussian random field v have Lebesgue measure zero [30]. We work with v as our unknown for the purposes of inversion, noting that u is easily recovered by application of $S(\cdot)$. In particular draws from the induced prior on u may be created by writing

$$(3) \quad u = \mathbb{1}_{v>0} - \mathbb{1}_{v<0}$$

with $v \sim \mu_0$. It is interesting to address the question is to whether the function u is of bounded total variation. This is equivalent to asking whether the level set $v = 0$ has finite length.

LEMMA 2.2. *If function v is drawn from measure $\mu_{0,\alpha}$ with $\alpha > 1 + d/2$ then almost surely function u defined by (3) has finite total variation norm.*

Proof. If $\alpha > 1 + d/2$ then almost surely $v \sim \mu_{0,\alpha}$ will be a C^1 function. The paper [34] establishes that the level set $v = 0$ will then have finite length, almost surely. Since u is a binary function given by (3) this establishes that u will have finite total variation, almost surely. \square

2.3.2. Likelihood. Let η be a normal random variable in \mathbb{R}^J , $\eta \sim N(0, \Sigma)$, Σ as before. Using the fact that $u = S(v)$ it follows from (1) that

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

and hence that $y|v$ is distributed as the Gaussian $N(KS(v), \varepsilon^{2c}\Sigma)$. The negative log likelihood is then proportional to the misfit

$$\frac{1}{2} \left| \Sigma^{-1/2} (y - KS(v)) \right|^2.$$

2.3.3. Posterior.

PROPOSITION 2.3. *Let $\alpha > d/2$. Then the posterior probability μ^y on random variable $v|y$ is a probability measure supported on $X^{k,\gamma}$ for all $\gamma \in [0, \gamma']$ where k is the largest integer such that $\gamma' := \alpha - \frac{d}{2} - k \in (0, 1]$ and determined by*

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

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

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

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

Proof. Everything but the final statement follows from application of the theory in [30]. The final statement 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 2.2 gives the desired result. \square

2.4. Samples From the Prior. We describe how to sample numerically from the Gaussian priors $\mu_{0,\alpha} = \mathbf{N}(0, C^{\alpha/2})$ that are key to the two Bayesian inversion techniques outlined in the preceding two subsections. Using this we investigate numerically the length of the level sets of these samples. We have shown in Lemma 2.2 that choosing $\alpha > 1 + d/2$ is sufficient to ensure almost sure finite length of level sets; our numerical results will demonstrate that this is a sharp result.

Let $\{\lambda_k\}$ denote the eigenvalues of C , ordered increasing, and $\{\varphi_k\}$ the corresponding $L^2(D)$ -normalized eigenfunctions (which are Fourier modes). Then samples u from $\mu_{0,\alpha}$ may be expressed through the Karhunen-Loève expansion as

$$(4) \quad v(x, y) = \sum_{k=1}^{\infty} \lambda_k^{\alpha/4} \xi_k \varphi_k(x, y), \quad \xi_k \sim \mathbf{N}(0, 1) \text{ i.i.d.}$$

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 $\{x_i, y_j\}$ of N^2 points in D . An approximate sample on this grid is then given by

$$v^N(x_i, y_j) = \sum_{k=1}^{N^2} \lambda_k^{\alpha/4} \xi_k \varphi_k(x_i, y_j), \quad \xi_k \sim \mathbf{N}(0, 1) \text{ i.i.d.}$$

All of our numerical results in section 4 are performed on the grid which arises from this approach to generating Gaussian random fields.

We may apply the map S to a Gaussian random field generated in this way to obtain a binary grid-based function $w^N = S(v^N)$. The length of the zero level set of v may then be approximated by

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

where the operator D^N approximates the gradient on the grid $\{x_i, y_j\}$ via central differences.

In Figure 1 we confirm the above intuition about interface length for the prior distribution by observing how it scales with respect to increasing N , for a single realization of $\{\xi_k\}$ in (4), as we vary α . We use $d = 2$ for which $\alpha = 2$ is the critical value predicted by the theory. We see 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

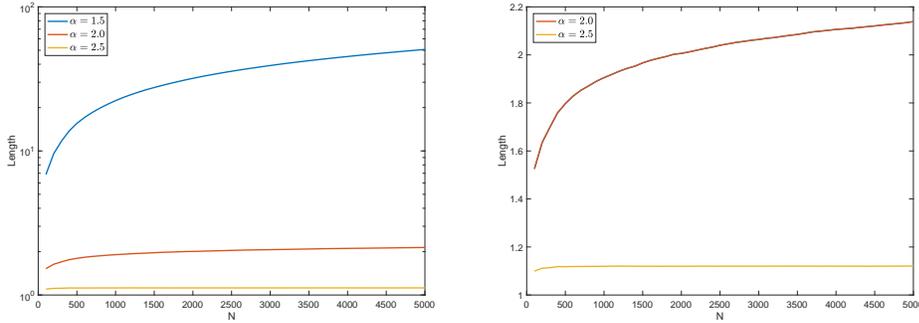


FIG. 1. The dependence of the length of the zero level set of a Gaussian sample, as a function of numerical approximation level N and of prior regularity parameter α . (Left) logarithmic axis, (right) linear axis.

left and right hand panels show this). The results, then, suggest that level sets have finite length if and only if $\alpha > 1 + d/2$.

Since almost sure properties of the prior are inherited in the posterior, this means that the phase-field level set approach will not lead to posterior samples with finite perimeter $u = 0$ in dimension $d \geq 2$. This is because the phase field formulation uses $\alpha = 2$, a choice dictated by our wish to construct a MAP estimator which, for small ε , approximately penalizes the perimeter, as we will show section 3; these is no contradiction here because MAP estimators on function space will always be smoother than draws from the measure [16]. On the other hand, in the level set method we are not constrained to choose $\alpha = 2$ and as we will show in section 4 choosing $\alpha = 3$ leads, in dimension $d = 2$ to not only finite perimeter, but posterior on the perimeter which contracts close to the true value.

3. MAP Estimators for Phase-Field Posterior. Recall that the MAP estimator of a Bayesian posterior distribution maximizes the posterior probability; we make this precise below. Here we explain constraints on the parameters a_1, a_2 and a_3 that are needed in order to obtain an $\varepsilon \rightarrow 0$ Γ -limit for the MAP estimator associated with the phase-field posterior; this limit is total variation penalized least squares. We do not address choice of the non-negative parameters δ, r, q and τ as, other than requiring strict positivity, their choices do not play a big role and optimizing them will be very case specific.

Recall the Cameron-Martin space E of the Gaussian measure μ_0 on X is $H_{\#}^2(D)$ with the norm given by given by

$$\|u\|_E^2 = \frac{\delta}{\varepsilon^{2a_1}} \|\Delta u\|_{L^2(D)}^2 + \frac{\delta q}{\varepsilon^{2a_2}} \|\nabla u\|_{L^2(D)}^2 + \frac{\delta \tau^2}{\varepsilon^{2a_3}} \|u\|_{L^2(D)}^2.$$

Now define $\Psi : X \rightarrow \mathbb{R}^+$ by

$$(5) \quad \Psi(u) = \frac{r}{\varepsilon^b} \int_D \frac{1}{4} (1 - u(x)^2)^2 dx$$

and $\Phi : X \times \mathbb{R}^J \rightarrow \mathbb{R}^+$ by

$$(6) \quad \Phi(u, y) = \Psi(u) + \frac{1}{2\varepsilon^{2c}} |\Sigma^{-1/2}(y - Ku)|^2.$$

We define the Onsager-Machlup functional, J^ε , associated with the measure ν^y by

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

For $\rho > 0$, let $B^\rho(z)$ be the ball centred at $z \in X$ with radius ρ and define

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

Following Dashti et al. [16], we define a MAP estimator as follows. Intuitively this definition captures the idea that the MAP estimator locates points in X at which arbitrarily small balls will have maximal probability.

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

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

Then we have the following result demonstrating the role of the Onsager-Machlup functional from [16, 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 .*

This establishes a direct connection between Bayesian inversion and classical regularization. We now study the Γ -limit of the MAP estimator as $\varepsilon \rightarrow 0$. The functional $J^\varepsilon(u)$ can be written as

$$(7) \quad J^\varepsilon(u) = \varepsilon^{-2a_1-3} I^\varepsilon(u),$$

where

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

As $\varepsilon > 0$ the critical points of J^ε and I^ε coincide, therefore in what follows we will consider I^ε . In order to obtain the desired Γ -limit for I^ε we are forced to consider the case where

$$(8) \quad a_2 - a_1 = 1, \quad 3 + 2a_1 - b = -1, \quad 3 + 2a_1 - 2c = 0, \quad 3 + 2(a_1 - a_3) = a > 0.$$

Thinking of c as given (defining the unit of small noise) these provide the values for a_1, a_2 and b , and a constraint on a_3 , necessary to obtain our desired Γ -limit. Note, in addition, that the constraint $c > 0$ (the small noise assumption) means that $2a_1 + 3 = 2c > 0$ and that $b = 2a_1 + 4 = 2c + 1 > 0$.

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

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

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

We study the Γ -convergence of the functional I^ε , basing our analysis on the work of Hilhorst et al [27]. We define the functional

$$e^\delta(U) = \int_{-\infty}^{\infty} \left(\frac{1}{2} \delta(U''(t))^2 + \frac{q}{2} \delta(U'(t))^2 + \frac{r}{4} (1 - U(t)^2)^2 \right) dt;$$

and the constant

$$P^\delta = \inf_{U \text{ odd}} e^\delta(U).$$

We then have the following theorem:

THEOREM 3.3. *Define*

$$I_0^\delta = \frac{1}{2} \int_D P^\delta |\nabla u| dx + \frac{1}{2} |\Sigma^{-1/2}(y - Ku)|^2, \quad \text{if } u \in BV_{\text{binary}}(D),$$

where $BV_{\text{binary}}(D) = \{\psi \in BV(D) : \psi(D) \subset \{\pm 1\}\}$. Then

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

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

We present the proof in [Appendix A](#). The theorem shows that the MAP estimator is, for small observational noise ε , close to a total variation penalized least squares minimization over piecewise constant binary functions. Furthermore, since $2a_1 + 3 > 0$, (7) 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 next, support this conjecture.

4. Numerical Experiments. We describe numerical results which illustrate: (i) the quality of reconstructions; (ii) the relative efficiency of the phase field and level set simulations; (iii) the ability to learn the perimeter length. For benchmark we also use a Gaussian process (GP) regression approach to the problem, although we note that this method does not generalize to nonlinear forward maps and so is not as widely applicable as the phase field and level set approaches. All of our numerical results are performed using the grid which arises from the approach to generating Gaussian random fields which is outlined in subsection 2.4, with $N = 2^{14}$. The integral required to evaluate the phase-field likelihood is approximated using quadrature on the grid.

Markov Chain Monte Carlo (MCMC) simulations may be used to sample the measures ν^y, μ^y defined above. Here we simply *assume* that the resulting Markov chain $\{u^{(m)}\}$ is ergodic. The theory in [23] 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. We thus expect that we may make the approximation

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

where the error e_M is Gaussian with variance c_g/M , and c_g the integrated autocorrelation of $g(u^{(m)})$. A similar expression also holds for expectations under μ^y using samples $\{v^{(m)}\}$. These samples can then be used to produce point estimates for the unknown 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 both the phase-field and level set formulations. 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 K , and study the convergence to equilibrium of quantities, such as the acceptance probability of the chain, as a function of K . Preliminary numerical results for the one dimensional analogue of the problem studied here may be found in [42].

Our overall conclusion is that the Bayesian level set approach achieves comparable reconstruction accuracy to the phase-field approach, but at much lower cost; GP regression performs similarly to Bayesian level set and is very cheap for the linear inverse problem studied here, but unlike the Bayesian level set it does not readily extend to nonlinear forward maps. We also study the behaviour of the total variation norm of the reconstructions arising from the Bayesian level set approach demonstrating that, for appropriately chosen Gaussian prior, the Bayesian level set approach penalizes the perimeter of interfaces in the solution, and can learn about its true value from observations. In summary, the Bayesian level set method provides a viable approach to interface inference problems and indirectly imposes a form of total variation regularization on the reconstructed function.

4.1. Setup of Numerical Experiments. Apart from in subsection 4.2.4, we we make the choice $\alpha = 2$ in all the experiments, so that $\mu_{0,\alpha} = \mu_0$. We now detail the other parameter settings and algorithmic choices.

4.1.1. MCMC for Phase-Field and Level Set Formulations. We employ the preconditioned Crank-Nicolson (pCN) algorithm [1, 15] which may be used to sample any measure σ of the form

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

Both of our posterior measures can be written in this form. We perform numerical experiments in both the small noise and order one noise regimes.

The proposal variance parameter $\beta \in (0, 1]$ will be referred to later. 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. 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-Loeve expansion described in subsection 2.4 [23]. Furthermore derivatives of $A(u)$ are not required to implement the method.

4.1.2. Gaussian Process Regression. Since the forward map is linear, we may perform inversion using Gaussian Process (GP) regression. We look at the particular case of this methodology that arises from setting $r = 0$ in the phase-field formulation. Then $\nu_0 = \mu_0$ and $\nu^y = \mathbf{N}(m_y, C_u)$ is Gaussian and determined by the formulae

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

This distribution may be sampled directly, without the need for MCMC. The mean m_y is also the MAP estimator in this case, and so may be characterized as the unique minimizer of the functional

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

We may additionally consider the thresholded field $S(u)$, $u \sim \mathbf{N}(m_y, C_y)$, whose statistics may be deduced from those of the Gaussian field.

4.1.3. Small Observational Noise. Throughout we set $c = 3/2$ and $\varepsilon = 0.01$. The implied standard deviation of the observational noise is thus 0.001.

For the phase-field formulation, we make a choice of parameters in the prior covariance operator C such that the relations (8) hold; the MAP estimator for ν^y then approximates the minimizer of I_0^δ as given in [Theorem 3.3](#), and we expect the posterior mass to concentrate fairly close to this MAP estimator. Specifically, we make the choices $a_1 = 0$, $a_2 = 1$, $a_3 = 0$, $b = 4$, $\delta = 0.01$, $q = 0.1$, $\tau = 1$ and $r = 1$.

For the level set formulation we make the same choices of prior parameters as for the phase-field formulation, except we set $\delta = 1$, $q = 0$ and $\tau = 50$. Note that in general we need not insist on the parameters being related via (8) 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 level set method.

4.1.4. Order One Observational Noise. We set $c = 0.0$. Note that now ε does not enter the observational noise; it is simply a parameter that enters the prior. 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 $\delta = 100$, $q = 0.1$, $\tau = 1$ 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. Results of Numerical Experiments. 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 images BV_{binary} and are illustrated in the obvious way in [Figure 2](#). 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 with standard deviation $\varepsilon^c = 0.001$, as in equation (1).¹ MCMC sampling and GP regression is performed on a mesh of 2^{14} points. The ratio of the mesh-scale to ε is thus $\mathcal{O}(1)$. In all MCMC runs we generate 10^6 samples, and discard the first 5×10^5 samples as burn-in when computing means. In the notation of [15] for the pCN algorithm, depending on the noise model and number of observations, we take the proposal standard deviation parameter β between 0.002 and 0.02 for phase-field simulations, and β between 0.02 and 0.1 for level set 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. In order to avoid an inverse crime [33], 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; the inversion then attempts to recover approximations of these highly resolved fields on the coarser grids used in the algorithm.

4.2.1. Small Noise Reconstructions. In [Figure 3](#) we present sample means associated with small-noise observations for the phase-field, level set and GP regression models, both with and without thresholding by S . Note that the phase-field and GP

¹Pointwise 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 unchanged in any substantive way.

regression models attempt to fit the un-thresholded field to the datapoints, whereas the level set method attempts to fit the thresholded field; the un-thresholded field for the level set method is hence on a different scale to 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 Truth C 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. 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.

4.2.2. Order One Noise Reconstructions. In [Figure 4](#) 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 .

4.2.3. 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 simulations performed here, evaluation of $A(u)$ is approximately twice as expensive for the phase-field model than 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 4.1.2](#), with the cost arising from the matrix multiplications and inversion involved. However, the GP approach does not generalize to nonlinear forward maps.

For the phase-field and level set approaches, the most significant discrepancy in computational cost arises from the statistical properties of the chains. In [Figure 5](#) we show the evolution of the local acceptance rates of proposed states for Truth A with small observational noise; the evolutions are similar for the other datasets and so are not presented for brevity. These figures suggest 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 6](#), which shows a selection of samples early in the chains for Truth B with small observational noise. After 10000 samples the three inclusions have already been identified by the level set chain, however after 30000 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.

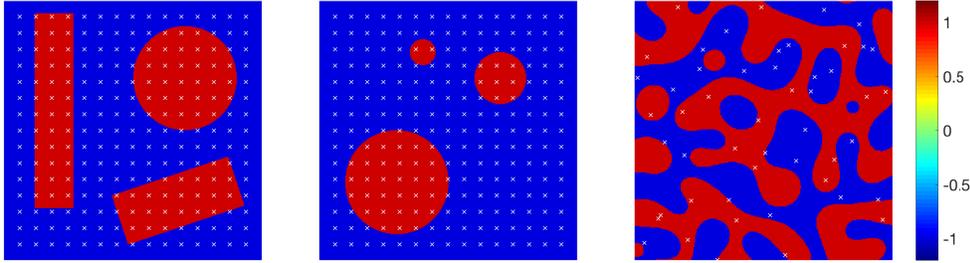


FIG. 2. The three true fields used for inversion; the field on the left will be referred to as Truth A, the field in the middle as Truth B and the field on the right as Truth C. The sets of observation points are shown in each figure.

Another observation to make from Figure 5 is that the acceptance rates for the phase-field chains are much lower than those for the level set chains, despite the proposal standard deviation parameter β being one tenth of the size. To understand why this is the case, note that the measure ν^y can informally be thought of as having Lebesgue density proportional to $\exp(-J^\varepsilon(u)) = \exp(-\varepsilon^{-2\alpha_1-3}I^\varepsilon(u))$. 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 .

4.2.4. Perimeter Penalization. Here we study perimeter learning for the Bayesian level set method. The material in subsection 2.4 suggests that if we wish to compare the perimeter distribution between the prior and posterior, we must choose $\alpha > 1+d/2$ to ensure the length of the zero level set is well-defined. Working in dimension $d = 2$ we thus choose $\alpha = 3$ and compute the prior and posterior distributions on the zero level set resulting from Truth B in the small noise setting. The results are shown in Figure 7. Whilst the perimeter still retains some variation under the posterior, the variation is much lower and, in contrast to the prior, it is concentrated around near to the true value of the perimeter. Thus the Bayesian level set approach has the ability to estimate the perimeter, and quantify uncertainty in the estimation.

5. Conclusions. We have studied Bayesian inversion for unknown functions which are known to be piecewise constant on a bounded open set. The methods we study do not explicitly penalize the perimeter. We study the manner in which perimeter regularization nonetheless arises in the methods studied. To be concrete we consider binary functions taking values in $\{\pm 1\}$ but the ideas are readily generalizable.

Method 1 is a Bayesian approach which, rather than thresholding, uses a non-Gaussian prior which penalizes deviations from $|u| = 1$; this prior is defined by using a Ginzburg-Landau penalty as density with respect to a Gaussian measure with Cameron-Martin norm equivalent to the norm on H^2 . Method 2 defines a non-Gaussian prior by thresholding a smooth function on which a Gaussian prior is placed, with Cameron-Martin norm equivalent to the norm on H^α ; different α are considered.

We find the following outcomes. In Method 1 the specific choice of Cameron-Martin norm means that draws from both the prior and posterior u will not have a finite length zero level set – hence their TV norm is infinite. However the penalization set-up provides a link between the Bayesian approach and phase field regularization made rigorously via the MAP estimator. Furthermore this can then be linked rig-

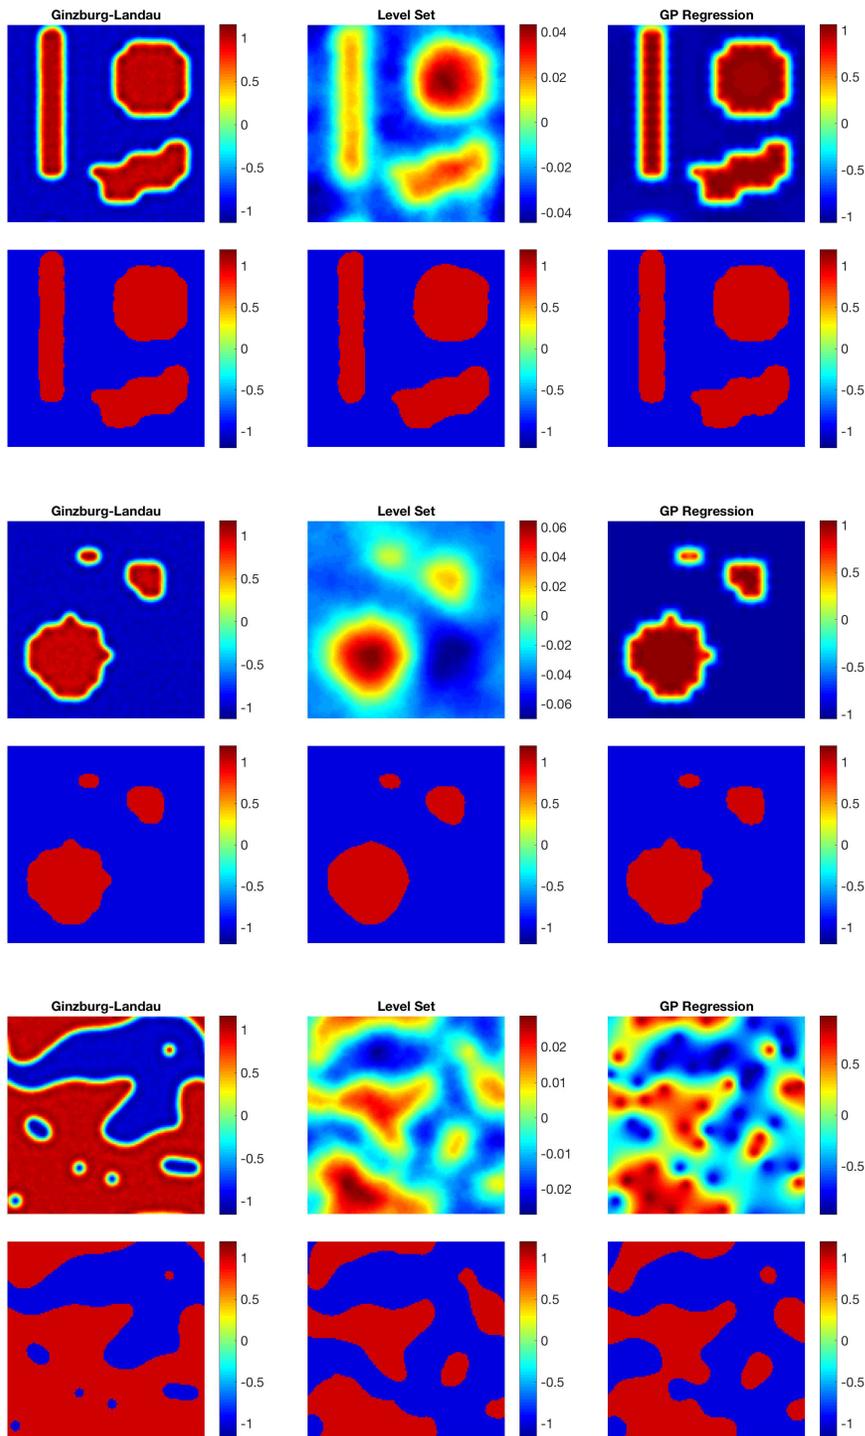


FIG. 3. Sample means for Truth A (top block), Truth B (middle block) and Truth C (bottom block) with small observational noise. To top row of each block shows Monte Carlo approximations to $\mathbb{E}^{\nu^y}(u)$, the underlying continuous fields, and the bottom row in each block shows Monte Carlo approximations to $S(\mathbb{E}^{\mu^y}(u))$, the thresholded fields.

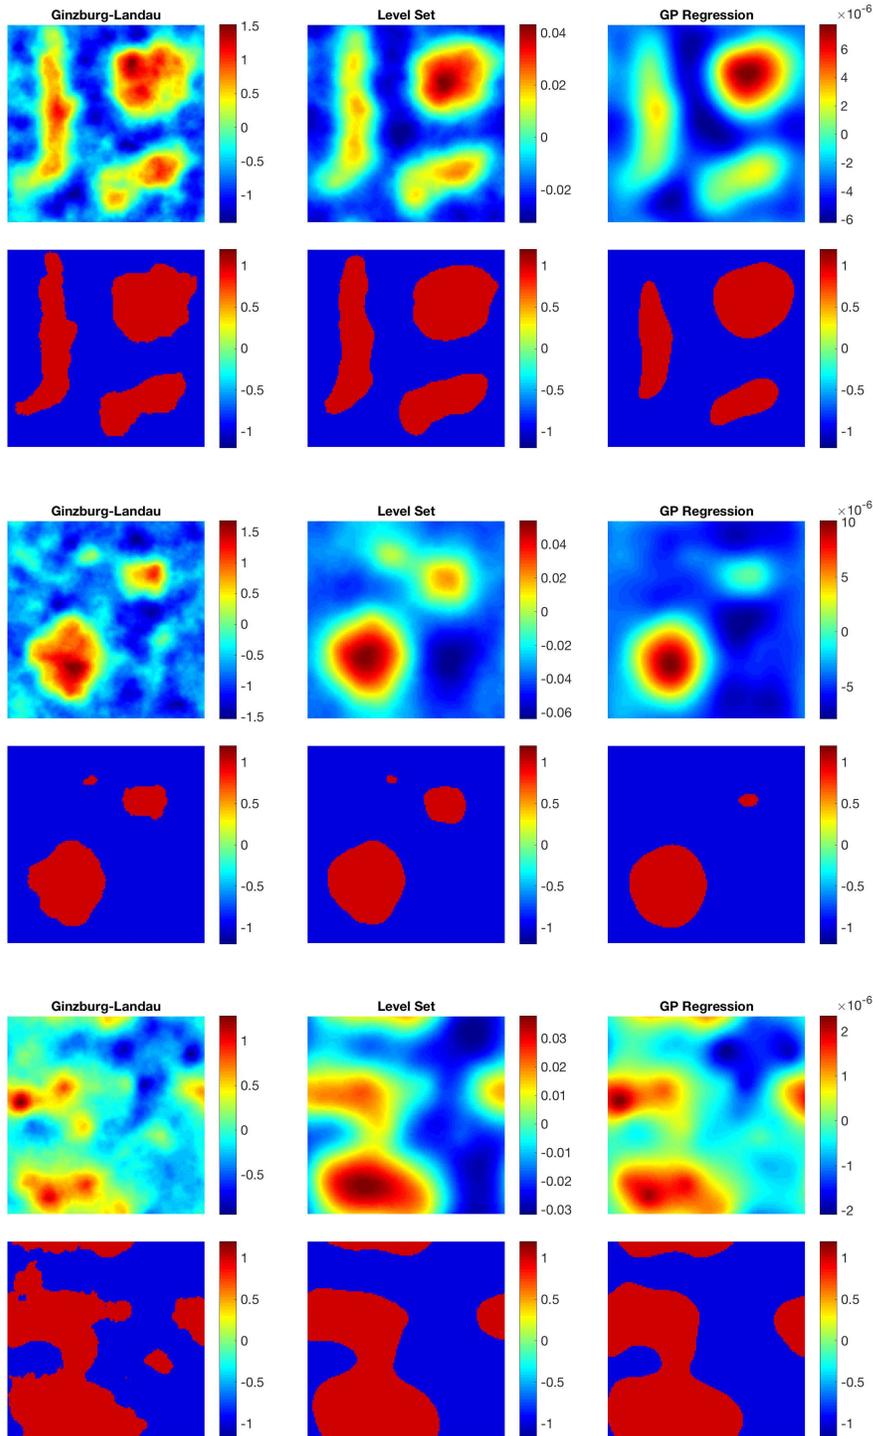


FIG. 4. Sample means for Truth A (top block), Truth B (middle block) and Truth C (bottom block) with order one observational noise. To top row of each block shows Monte Carlo approximations to $\mathbb{E}^y(u)$, the underlying continuous fields, and the bottom row in each block shows Monte Carlo approximations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

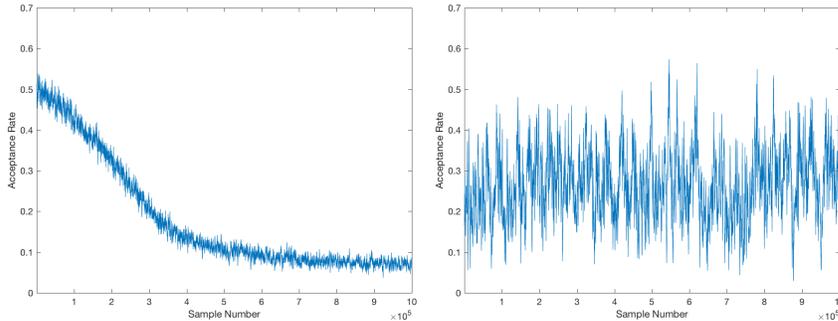


FIG. 5. The evolution of the acceptance rate of proposals for the phase-field (left) and level set (right) MCMC chains, for Truth A with small observational noise. Acceptance rates are calculated over a moving window of 1000 samples.

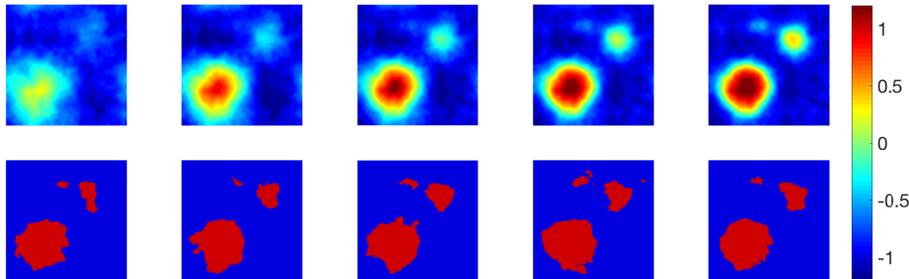


FIG. 6. Examples of samples near the start of chains for Truth B with small observational noise. Sample numbers 10000, 20000, 30000, 40000 and 50000 are shown from left-to-right for the phase-field chain (top) and the level set chain (bottom).

ously to perimeter regularization via Γ -convergence. In Method 2 the underlying smoothness allows the signum function $S(v)$ of samples to have a finite TV norm for appropriately chosen α , a topic studied in this paper. Figure 7 shows the remarkable improvement of the perimeter estimation, under the prior and then under the posterior after data fitting. The realization of Method 2 is computationally more efficient than Method 1. Both Method 1 and 2 demonstrate how the Bayesian approach implicitly induces perimeter regularization, but in differing ways. For Method 1 via the MAP estimator result; for Method 2 by the Figure 7.

The ideas in this paper can be combined in different ways: other smoothed thresholding functions could be contemplated within Method 2, such as the double-obstacle approximation to the signum function [2, 3], and penalties other than Ginzburg-Landau quartic potential could be used within Method 1, again including the double-obstacle approximation. The key conclusion is that perimeter penalization is possible within the Bayesian approach to inversion, but care is required to elucidate the manner in which this penalization manifests: within samples (Method 2), or only at the MAP point (Method 1).

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 [16], we need to show

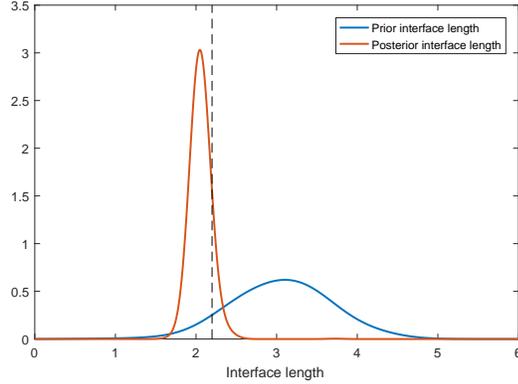


FIG. 7. The distribution of the perimeter under the prior and posterior distribution. The vertical dashed line indicates the perimeter of the true field.

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 non-negative 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 Lipschitzness, we have

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

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 3.3. We adapt the proof of Hilhorst et al. to allow for periodic boundary conditions and the additional L^2 norm appearing in the functional to be infimized. From Hilhorst et al., we have that if $u^\varepsilon \rightarrow u$ in $L^1(D)$ then

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

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

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

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 neighbourhood 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

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

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 et al.. 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 et al. 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 neighbourhood 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 et al., 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. □

- [1] A. BESKOS, G. O. ROBERTS, A. M. STUART, AND J. VOSS, *MCMC methods for diffusion bridges*, Stochastics and Dynamics, 8 (2008), pp. 319–350.
- [2] J. BLOWEY AND C. ELLIOTT, *Curvature dependent phase boundary motion and parabolic double obstacle problems*, in Degenerate Diffusions, Springer, 1993, pp. 19–60.
- [3] J. BLOWEY AND C. ELLIOTT, *A phase-field model with a double obstacle potential*, Motion by mean curvature and related topics (Trento, 1992), (1994), pp. 1–22.
- [4] C. BRETT, C. M. ELLIOTT, AND A. S. DEDNER, *Phase field methods for binary recovery*, in Optimization with PDE Constraints, Springer, 2014, pp. 25–63.
- [5] F.-X. BRIOL, C. J. OATES, M. GIROLAMI, M. A. OSBORNE, AND D. SEJDINOVIC, *Probabilistic integration: A role for statisticians in numerical analysis?*, arXiv preprint arXiv:1512.00933.
- [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), p. 114004.
- [7] D. CALVETTI AND E. SOMERSALO, *A Gaussian hypermodel to recover blocky objects*, Inverse Problems, 23 (2007), p. 733.
- [8] D. CALVETTI AND E. SOMERSALO, *Hypermodels in the Bayesian imaging framework*, Inverse Problems, 24 (2008), p. 034013.
- [9] J. N. CARTER AND D. A. WHITE, *History matching on the Imperial College fault model using parallel tempering*, Computational Geosciences, 17 (2013), pp. 43–65.
- [10] 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>, <http://dx.doi.org/10.1016/j.jcp.2010.07.005>.
- [11] L.-Q. CHEN, *Phase-field models for microstructure evolution*, Annual review of materials research, 32 (2002), pp. 113–140.
- [12] O. A. CHKREBTII, D. A. CAMPBELL, B. CALDERHEAD, AND M. A. GIROLAMI, *Bayesian solution uncertainty quantification for differential equations*, Bayesian Anal., 11 (2016), pp. 1239–1267, <https://doi.org/10.1214/16-BA1017>, <http://dx.doi.org/10.1214/16-BA1017>.
- [13] R. CHOKSI, Y. VAN GENNIP, AND A. OBERMAN, *Anisotropic total variation regularized l^1 -approximation and denoising/deblurring of 2d bar codes*, arXiv preprint arXiv:1007.1035, (2010).
- [14] J. COCKAYNE, C. OATES, T. SULLIVAN, AND M. GIROLAMI, *Bayesian probabilistic numerical methods*, arXiv preprint arXiv:1702.03673, (2017).
- [15] S. L. COTTER, G. O. ROBERTS, A. M. STUART, D. WHITE, ET AL., *MCMC methods for functions: modifying old algorithms to make them faster*, Statistical Science, 28 (2013), pp. 424–446.
- [16] 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), p. 095017.
- [17] M. DASHTI AND A. M. STUART, *The Bayesian approach to inverse problems*, arXiv preprint arXiv:1302.6989, (2013).
- [18] 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), p. 045008.
- [19] P. DIACONIS, *Bayesian numerical analysis*, Statistical Decision Theory and Related Topics IV, 1 (1988), pp. 163–175.
- [20] O. DORN AND D. LESSELIER, *Level set methods for inverse scattering—some recent developments*, Inverse Problems, 25 (2009), p. 125001, <http://stacks.iop.org/0266-5611/25/i=12/a=125001>.
- [21] O. DORN AND R. VILLEGAS, *History matching of petroleum reservoirs using a level set technique*, Inverse Problems, 24 (2008), p. 035015, <http://stacks.iop.org/0266-5611/24/i=3/a=035015>.
- [22] H. W. ENGL, M. HANKE, AND A. NEUBAUER, *Regularization of inverse problems*, vol. 375, Springer Science & Business Media, 1996.
- [23] M. HAIRER, A. M. STUART, AND S. J. VOLLMER, *Spectral gaps for Metropolis-Hastings algorithms in infinite dimensions*, The Annals of Applied Probability, 24 (2014), pp. 2455–290, <https://doi.org/10.1214/13-AAP982>.
- [24] P. C. HANSEN, J. G. NAGY, AND D. P. O’LEARY, *Deblurring images: matrices, spectra, and filtering*, SIAM, 2006.
- [25] T. HELIN AND M. LASSAS, *Hierarchical models in statistical inverse problems and the Mumford–Shah functional*, Inverse problems, 27 (2011), p. 015008.
- [26] P. HENNIG, *Probabilistic interpretation of linear solvers*, SIAM Journal on Optimization, 25 (2015), pp. 234–260.
- [27] D. HILHORST, L. A. PELETIER, AND R. SCHÄTZLE, *γ -limit for the extended Fisher–Kolmogorov equation*, Proceedings of the Royal Society of Edinburgh Section A: Mathematics, 132

- (2002), pp. 141–162.
- [28] B. HOSSEINI, *Well-posed bayesian inverse problems with infinitely divisible and heavy-tailed prior measures*, SIAM/ASA Journal on Uncertainty Quantification, 5 (2017), pp. 1024–1060.
- [29] M. IGLESIAS, K. LIN, AND A. STUART, *Well-posed bayesian geometric inverse problems arising in subsurface flow*, Inverse Problems, 30 (2014), p. 114001, <https://doi.org/doi:10.1088/0266-5611/30/11/114001>.
- [30] M. A. IGLESIAS, Y. LU, AND A. M. STUART, *A Bayesian level set method for geometric inverse problems*, Interfaces and Free Boundary Problems, (2016).
- [31] M. A. IWEN, F. SANTOSA, AND R. WARD, *A symbol-based algorithm for decoding bar codes*, SIAM Journal on Imaging Sciences, 6 (2013), pp. 56–77.
- [32] R. JIN, S. ZHAO, X. XU, E. SONG, AND C.-C. HUNG, *Super-resolving barcode images with an edge-preserving variational bayesian framework*, Journal of Electronic Imaging, 25 (2016), pp. 033016–033016.
- [33] J. KAIPPIO AND E. SOMERSALO, *Statistical and computational inverse problems*, vol. 160, Springer Science & Business Media, 2006.
- [34] M. F. KRATZ, *Level crossings and other level functionals of stationary gaussian processes*, Probability Surveys, 3 (2006).
- [35] M. LASSAS, E. SAKSMAN, AND S. SILTANEN, *Discretization-invariant Bayesian inversion and Besov space priors*, Inverse Problems and Imaging, 3 (2009), pp. 87–122.
- [36] M. LASSAS AND S. SILTANEN, *Can one use total variation prior for edge-preserving bayesian inversion?*, Inverse Problems, 20 (2004), p. 1537.
- [37] 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*, Journal of Computational Physics, 291 (2015), pp. 218–237.
- [38] L. I. RUDIN, S. OSHER, AND E. FATEMI, *Nonlinear total variation based noise removal algorithms*, Physica D: Nonlinear Phenomena, 60 (1992), pp. 259–268.
- [39] F. SANTOSA, *A level-set approach for inverse problems involving obstacles fadil santosa*, ESAIM: Control, Optimisation and Calculus of Variations, 1 (1996), pp. 17–33.
- [40] M. SCHÖBER, D. K. DUVENAUD, AND P. HENNIG, *Probabilistic ODE solvers with Runge-Kutta means*, in Advances in Neural Information Processing Systems, 2014, pp. 739–747.
- [41] J. A. SETHIAN, *Level set methods and fast marching methods: evolving interfaces in computational geometry, fluid mechanics, computer vision, and materials science*, vol. 3, Cambridge university press, 1999.
- [42] I. SIVAK, *Bayesian reconstruction of piecewise constant signals*, MSc. Dissertation, Warwick University, (2014).
- [43] J. SKILLING, *Bayesian solution of ordinary differential equations*, in Maximum Entropy and Bayesian Methods, Springer, 1992, pp. 23–37.
- [44] 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, 2015, pp. 117–124.
- [45] A. M. STUART, *Inverse problems: a Bayesian perspective*, Acta Numerica, 19 (2010), pp. 451–559.
- [46] Y. VAN GENNIP, P. ATHAVALE, J. GILLES, AND R. CHOKSI, *A regularization approach to blind deblurring and denoising of QR barcodes*, IEEE Transactions on Image Processing, 24 (2015), pp. 2864–2873.