

ESTIMATING AND USING INFORMATION IN INVERSE PROBLEMS

WOLFGANG BANGERTH*, CHRIS R. JOHNSON†, DENNIS K. NJERU‡, AND BART VAN BLOEMEN WAANDERS§

Abstract. For inverse problems one attempts to infer spatially variable functions from indirect measurements of a system. To practitioners of inverse problems, the concept of “information” is familiar when discussing key questions such as which parts of the function can be inferred accurately and which cannot. For example, it is generally understood that we can identify system parameters accurately only close to detectors, or along ray paths between sources and detectors, because we have “the most information” for these places.

Although referenced in many publications, the “information” that is invoked in such contexts is not a well understood and clearly defined quantity. Herein, we present a definition of *information density* that is based on the variance of coefficients as derived from a Bayesian reformulation of the inverse problem. We then discuss three areas in which this information density can be useful in practical algorithms for the solution of inverse problems, and illustrate the usefulness in one of these areas – how to choose the discretization mesh for the function to be reconstructed – using numerical experiments.

AMS subject classifications. 65N21, 35R30, 94A17

1. Introduction. Inverse problems – i.e., determining distributed internal parameters of a system from measurements of its state – are frequently ill-posed. Mathematically, this ill-posedness is often described as the lack of a continuous mapping from the space of measurements to the corresponding parameters reconstructed from a measurement. A consequence of ill-posedness is that a small measurement error can result in a significantly different reconstructed parameter unless the problem is *regularized* in some way.

Concretely, let us consider that we want to identify a spatially varying parameter $q(\mathbf{x})$, for instance the density and elastic moduli of the earth in seismology, or the absorption and scattering properties of the human body in biomedical imaging. This will require using measurements z of some part of the state of the system under interrogation, e.g. the time-dependent displacement at a seismometer station, or the light intensity at the surface of the body as recorded by the pixels of a camera. If z is corrupted by noise of level ε , we will get two reconstructions (q_1, q_2) for two measurements (z_1, z_2) that differ only by the realization of the measurement noise. Ideally, we would be able to show that

$$\|q_1(\mathbf{x}) - q_2(\mathbf{x})\| \leq C \|z_1 - z_2\| \quad (1.1)$$

for some appropriate choice of norms and a constant C of moderate size. The problem is “ill-posed” if such an estimate does not exist. Many inverse problems fall into this category of ill-posedness.

On the other hand, a pragmatic view of inverse problems is that the ill-posedness of the problem is simply the result of a *lack of information*. Some inverse problems can

*Department of Mathematics, Department of Geosciences, Colorado State University, Fort Collins, CO 80523-1874, USA bangerth@colostate.edu.

†Scientific Computing and Imaging Institute, University of Utah, Salt Lake City, UT 84112, USA crj@sci.utah.edu.

‡Scientific Computing and Imaging Institute, University of Utah, Salt Lake City, UT 84112, USA dnjeru@sci.utah.edu.

§Sandia National Laboratories, MS 0370, PO Box 5800, Albuquerque, NM 87185, USA bartv@sandia.gov.

achieve well-posedness by obtaining different kinds of measurements from the system under consideration, but even if that is not the case, it would already be advantageous to simply improve the *degree of ill-posedness*. In either case, we may ask whether it is possible to derive estimates of the kind

$$\|J(\mathbf{x})(q_1(\mathbf{x}) - q_2(\mathbf{x}))\| \leq C\|z_1 - z_2\|, \quad (1.2)$$

again with an appropriate choice of norms. We will call $J(\mathbf{x}) \geq 0$, or a related quantity such as its square root, the *information density*. Equation (1.2) is motivated by the observation that at places where $J(\mathbf{x})$ is large, we can accurately determine the value of the coefficient $q(\mathbf{x})$ we are looking for (i.e., $q_1(\mathbf{x}) - q_2(\mathbf{x})$ must be small to satisfy the inequality). Conversely, the places where $J(\mathbf{x})$ is small coincide with those locations where we have little control over the coefficient, and even small amounts of noise in z may lead to large variations $q_1(\mathbf{x}) - q_2(\mathbf{x})$. In the extreme case when the problem is truly ill-posed, $J(\mathbf{x})$ would not be bounded away from zero and consequently $\|J(\mathbf{x})\varphi(\mathbf{x})\|$ would not be a norm of φ .¹

It is unlikely that for practical problems one can find meaningful expressions for J that give rise to provable estimates of the form (1.2). This is because for many inverse problems, what can and cannot be recovered stably is often not about where in space we are, but about which *modes in feature space* (for example low- versus high-frequency components of a function $q(\mathbf{x})$) are identifiable. In our discussions below, we therefore consider estimates such as (1.2) *aspirational*: We will instead seek statements such as

$$\|j(\mathbf{x})(q_1(\mathbf{x}) - q_2(\mathbf{x}))\| \simeq C\|z_1 - z_2\|, \quad (1.3)$$

where $j(\mathbf{x})$ takes on the role of the information density, and where \simeq expresses a relationship of the form “behaves conceptually like, but possibly only when spatial discretization is used”. While we will not be able to (and one likely cannot) prove that any choice of j in (1.3) implies an estimate of the form (1.2), the conceptual approach of seeking a function $j(\mathbf{x})$ that expresses the idea of an information density of how much we know about q at different points in space will turn out to be useful in practice – as we will demonstrate in Sections 4 and 5.

References to information in inverse problems in the research literature. The notion of information density is not new, in particular in applications where $q(\mathbf{x})$ is replaced by finite-dimensional parameter vectors $(q_i)_{i=1}^N$. Indeed, similar notions can be found in many areas of inverse and parameter estimation problems in various forms, and among practitioners of inverse problems, there is a degree of “knowledge” that *information* is a key concept. At the same time, practitioners do not appear to have a clear understanding of what information actually means, and uses of this concept in the literature appear to be vaguely defined and disconnected.

In our review of the existing literature, we have come across many publications that touch on the concept of information in inverse problems. The most obvious application of information concepts to inverse problems is in optimal experimental design where the goal of the design of schemes is to measure data about the system so as to minimize the uncertainty (that is, to maximize the information) in the parameters

¹In many inverse problems, for example in imaging, the ill-posedness manifests not by there being *locations* \mathbf{x} at which $q(\mathbf{x})$ is not identifiable. Rather, it is the *high-frequency (Fourier) content* of q that is often not identifiable without regularization. In this case, one may want to consider an equation like (1.2) with \mathbf{x} replaced by the wave number \mathbf{k} . Regardless of this obvious difference, let us for the moment move forward with the derivation as stated.

we wish to recover [5, 10, 23]. This relation to uncertainty is most clearly articulated in the Bayesian setting of optimal experimental design where the *information gain* of the posterior probability distribution over the prior is maximized. However information, generally defined in less concise terms, is also a topic discussed in other contexts. For example, considering concrete applications, [37] presents Fisher information for a single-particle system and proposes a new uncertainty relationship based on Fisher information. Similarly, [12, 47] discuss the use of a resolution matrix in seismic tomography (see also [36]); related concepts of “resolution”, “resolution length scales”, “event kernels”, “sensitivity kernels”, or “point spread functions” also appear in both seismic imaging and a number of other fields, see for example [21, 29–31, 35, 43]. In many other cases, the literature references the Fischer information matrix that, together with the Cramér-Rao bound, quantifies how accurately we know what the inverse problem seeks to identify [28]; examples include [11], which uses this approach for estimating diffusion in a single particle tracking process; [25], which compares Fisher matrices to the Hessian calculation in boundary value inversion problem using the heat equation; and [33] presents a preconditioning and regularization scheme based on Fisher information.

These publications and several others make connections between the accuracy of measurements and the uncertainty in the recovered parameters of the inverse problem, but these studies have not been undertaken to specifically *identify the role of information in the spatially variable ability to recover parameters in inverse problems*. Indeed, a common feature of the many definitions of resolution, adjoints, controllability, and identifiability that can be found in the literature, is that most of these notions originate in, and were developed for, *deterministic* inverse problems. On the other hand, “information” is probably best understood as a statistical concept, and a useful definition will therefore be rooted in statistical reformulations of the inverse problems. It is the connection between the Fisher information matrix and the variance of reconstructed parameters, via the Cramér-Rao bound, that we will utilize to derive information densities in Section 2.2 below.

The differences in the concepts mentioned above, and the lack of a common language to describe them, then presents the background of this work. Our goals are as follows:

- To introduce the notion of an *information density* based on a statistical interpretation of the inverse problem, the Fisher information matrix, and an application of the Cramér-Rao bound.
- To outline a number of applications for which we believe that an information density can be usefully employed.
- To practically evaluate our concepts in a concrete application, namely the choice of mesh on which to discretize an inverse source-identification problem.

Herein, we will pursue these goals by first considering a finite-dimensional, linear model problem in Section 2 that we use to provide a conceptual overview of what we are trying to achieve, followed by the extension of this model problem to the infinite-dimensional case in Section 3. Having so set the stage, we provide “vignettes” for three ways in which we believe information densities can be used in practice in Section 4. Section 5 then explores one of these – the choice of mesh for discretizing an infinite-dimensional inverse problem – in detail and with numerical and quantitative results. We conclude in Section 6.

2. A finite-dimensional, linear model problem. We first consider a simple, finite-dimensional problem. Specifically, let us consider a problem in which a state

vector \mathbf{u} is related to source terms via the relationship

$$A\mathbf{u} = \sum_k q_k \mathbf{s}_k, \quad (2.1)$$

where \mathbf{s}_k are possible source vectors and q_k are their relative strengths. We assume that the system matrix A is invertible although it may be ill-conditioned. In the inverse problem, we are then interested in recovering unknown source strengths q_k through a number of (linear, noisy) measurements

$$z_\ell = \mathbf{m}_\ell^T \mathbf{u} + \varepsilon_\ell, \quad (2.2)$$

where ε_ℓ is measurement noise. We assume that we have a guess σ_ℓ for the magnitude of the noise ε_ℓ .

For convenience, let us collect the quantities $q_k, \mathbf{s}_k, z_\ell, \mathbf{m}_\ell$ into vectors and matrices $\mathbf{q}, S, \mathbf{z}, M$, where the \mathbf{s}_k form the columns of S and the \mathbf{m}_ℓ the rows of M . Then, we can state the source strength recovery problem we will consider here as

$$\min_{\mathbf{q}, \mathbf{u}} \frac{1}{2} \|M\mathbf{u} - \mathbf{z}\|_{\Sigma^{-2}}^2 + \frac{\beta}{2} \|R\mathbf{q}\|^2, \quad (2.3)$$

such that $A\mathbf{u} = S\mathbf{q}$.

Here, we have used the weighted norm $\|M\mathbf{u} - \mathbf{z}\|_{\Sigma^{-2}}^2 = \sum_\ell \frac{1}{\sigma_\ell^2} |\mathbf{m}_\ell^T \mathbf{u} - z_\ell|^2$, where the diagonal matrix $\Sigma_{\ell\ell} = \sigma_\ell$ weighs measurements according to the assumed *certainty* $\frac{1}{\sigma_\ell}$ we have of the ℓ th measurement. We have also added a Tikhonov-type regularization term where β is the regularization parameter and R a matrix that amplifies the undesirable components of \mathbf{q} .

By eliminating the state variable using the state equation (2.1), we can re-state this problem as an unconstrained, quadratic minimization problem:

$$\min_{\mathbf{q}} \mathcal{J}_{\text{red}}(\mathbf{q}) := \frac{1}{2} \|MA^{-1}S\mathbf{q} - \mathbf{z}\|_{\Sigma^{-2}}^2 + \frac{\beta}{2} \|R\mathbf{q}\|^2. \quad (2.4)$$

It is then not difficult to show that the minimizer \mathbf{q} of this problem satisfies

$$\underbrace{(S^T A^{-T} M^T \Sigma^{-2} M A^{-1} S + \beta R^T R)}_Q \mathbf{q} = S^T A^{-T} M^T \Sigma^{-2} \mathbf{z}. \quad (2.5)$$

If we consider the noise to be random, we can ask how the solution \mathbf{q} depends on concrete measurements. Specifically, if we have two, presumably nearby, measurements $\mathbf{z}_1, \mathbf{z}_2$, then the following relationship holds for the corresponding solutions $\mathbf{q}_1, \mathbf{q}_2$:

$$Q(\mathbf{q}_1 - \mathbf{q}_2) = S^T A^{-T} M^T \Sigma^{-2} (\mathbf{z}_1 - \mathbf{z}_2). \quad (2.6)$$

For the following discussions, it is important to point out that the structure of Q guarantees that it is a symmetric, positive, semidefinite matrix; that is, all of its eigenvalues are non-negative. We will assume that the user has chosen β and R in such a way that Q is positive definite, although some of the eigenvalues may be small.

2.1. Defining an “information content” for components of the solution vector. The question we want to investigate herein is if the relationship (2.6) between $\mathbf{q}_1 - \mathbf{q}_2$ and $\mathbf{z}_1 - \mathbf{z}_2$ allows us to define stability bounds such as those outlined in (1.2) or (1.3) above. In the context of this finite-dimensional situation, such a bound would have the form

$$\|\mathbf{j} \odot (\mathbf{q}_1 - \mathbf{q}_2)\| \simeq C \|\mathbf{z}_1 - \mathbf{z}_2\|, \quad (2.7)$$

with a constant C of possibly unknown size, and where \odot indicates the Hadamard product that scales each entry of the vector $\mathbf{q}_1 - \mathbf{q}_2$ by the corresponding entry of the “information vector” \mathbf{j} . (Alternatively, we can interpret $\mathbf{j} \odot (\mathbf{q}_1 - \mathbf{q}_2)$ as $\text{diag}(\mathbf{j})(\mathbf{q}_1 - \mathbf{q}_2)$ where $\text{diag}(\mathbf{j})$ is a diagonal matrix with diagonal entries j_k .)

A *meaningful* statement² such as (2.7) will not always follow from (2.5) unless either the action of the matrix $Q = S^T A^{-T} M^T \Sigma^{-2} M A^{-1} S + \beta R^T R$ can somehow be approximated from below by a diagonal matrix, or Q^{-1} be approximated from above by a diagonal matrix. Indeed, we could choose \mathbf{j} to be a vector whose elements are all equal to $j_k = \lambda_{\min}(Q) = [\lambda_{\max}(Q^{-1})]^{-1}$. If, in addition, $C = \|S^T A^{-T} M^T \Sigma^{-2}\|$, then (2.7) holds true. This approach works as long as the regularization is chosen so that all eigenvalues of Q are reasonably large, i.e., that the problem is well-posed; in practice, however, this choice may over-regularize the problem.

At the same time, the choice $j_k = \lambda_{\min}(Q)$ is not interesting since it does not give us any insight into *which elements of \mathbf{q} can be estimated accurately and which cannot*. Furthermore, if the problem is indeed ill-posed, as the regularization parameter β is reduced, *all* elements j_k will become small, despite the fact that the unregularized problem may only imply that *some* components of \mathbf{q} cannot be stably recovered. To address this problem, we will appeal to a stochastic (Bayesian) interpretation of the inverse problem [27, 42]. From this perspective, we assume that our measurements \mathbf{z} are stochastic because they are corrupted by noise, and that consequently our recovered coefficients \mathbf{q} are also stochastic variables whose joint probability distribution we would like to infer. If we assume that the components of \mathbf{z} are distributed according to $\mathbf{z} = \hat{\mathbf{z}} + N(0, \Sigma^2)$ (that is, we assume that the noise is Gaussian and that our guessed noise levels σ_ℓ are indeed correct), then the desired probability distribution for \mathbf{q} will be of the form

$$p(\mathbf{q}|\mathbf{z}) = \kappa e^{-\frac{1}{2} \mathcal{J}_{\text{red}}(\mathbf{q})} \quad (2.8)$$

where κ is a normalization constant whose concrete value is not of importance to us.

Given this interpretation, the question of how much we know about the individual components of \mathbf{q} can be related to the *uncertainty* under $p(\mathbf{q}|\mathbf{z})$ – namely, we should choose the information weights j_k as the *inverse of the standard deviation of q_k* , that is,

$$j_k = \frac{1}{\sqrt{\text{var}_p(\mathbf{q})_k}}, \quad \text{where} \quad \text{var}_p(\mathbf{q})_k = \int (q_k - \mathbb{E}_p[\mathbf{q}]_k)^2 p(\mathbf{q}|\mathbf{z}) \, d\mathbf{q}. \quad (2.9)$$

Here, $\mathbb{E}_p[\mathbf{q}] = \int \mathbf{q} p(\mathbf{q}|\mathbf{z}) \, d\mathbf{q}$. Because we are considering a linear problem (2.1) and because the objective function \mathcal{J}_{red} is quadratic, this expectation value $\mathbb{E}_p(q_k)$

²Equation (2.5) implies that $\|\mathbf{q}_1 - \mathbf{q}_2\| \leq C \|\mathbf{z}_1 - \mathbf{z}_2\|$ with $C = \|Q^{-1} S^T A^{-T} M^T \Sigma^{-2}\|$, which corresponds to choosing \mathbf{j} in (2.7) as a vector of ones. At the same time, this estimate reflects no specifics of the problem and we do not consider this choice useful because it does not help us identify *which components of \mathbf{q} can be identified accurately and which cannot*.

is equal to the solution of the original, deterministic problem (2.5). This choice of \mathbf{j} has the pleasant property of making (2.7) dimensionally correct also for cases where the components of \mathbf{q} have different physical units. We note that the inverse of the variance is often called the “precision” with which a parameter is known [22]; j_k is then the square root of the k th parameter’s precision.

2.2. Estimating the information content for components of the solution vector. This definition of an information content j_k based on the inverse of the variance in the stochastic inverse problem makes intuitive sense. The question remains whether these weights j_k can be computed efficiently or at least estimated.

To do so, recall that the variances $\text{var}_p(\mathbf{q})_k = \text{cov}_p(\mathbf{q})_{kk}$ are the diagonal entries of the covariance matrix associated with $p(\mathbf{q}|\mathbf{z})$, where the covariance matrix is defined as

$$\text{cov}_p(\mathbf{q})_{kl} = \int (q_k - \mathbb{E}_p[\mathbf{q}]_k) (q_l - \mathbb{E}_p[\mathbf{q}]_l) p(\mathbf{q}|\mathbf{z}) d\mathbf{q}. \quad (2.10)$$

This matrix is not computable via the integral in an efficient way. However, we can find estimates of its elements via the Cramér-Rao bound that states that

$$\text{cov}_p(\mathbf{q}) \geq I_p^{-1}$$

in the sense that $[\text{cov}_p(\mathbf{q}) - I_p^{-1}]$ is a positive semidefinite matrix. Here, I_p is the Fisher information matrix defined by

$$(I_p)_{kl} = -\mathbb{E} \left[\frac{\partial^2}{\partial q_k \partial q_l} \ln p(\mathbf{q}|\mathbf{z}) \right], \quad (2.11)$$

which for our choice of $p(\mathbf{q}|\mathbf{z})$ and $\mathcal{J}_{\text{red}}(\mathbf{q})$ evaluates to

$$(I_p)_{kl} = \mathbb{E} \left[\frac{\partial^2}{\partial q_k \partial q_l} \left(-\ln \kappa + \frac{1}{2} \mathcal{J}_{\text{red}}(\mathbf{q}) \right) \right] = \frac{1}{2} \mathbb{E} \left[\frac{\partial^2}{\partial q_k \partial q_l} \mathcal{J}_{\text{red}}(\mathbf{q}) \right] = Q_{kl}.$$

Moreover, the following inequality holds [17]:

$$\text{var}_p(\mathbf{q})_k = \text{cov}_p(\mathbf{q})_{kk} \geq [I_p^{-1}]_{kk} \geq [(I_p)_{kk}]^{-1}.$$

These statements then provide us with an efficient way to estimate $\text{var}_p(\mathbf{q})$:

$$\text{var}_p(\mathbf{q})_k \geq [(I_p)_{kk}]^{-1} = [Q_{kk}]^{-1}. \quad (2.12)$$

In the spirit of the transition from (1.2) to (1.3), let us then define the information content of the k th parameter as

$$j_k := \sqrt{Q_{kk}}. \quad (2.13)$$

REMARK 1. Based on the definition $Q = S^T A^{-T} M^T \Sigma^{-2} M A^{-1} S + \beta R^T R$, the elements $j_k = \sqrt{Q_{kk}}$ can be computed in different ways by setting parentheses in the defining expression. The first way computes

$$\begin{aligned} Q_{kk} &= \mathbf{e}_k Q \mathbf{e}_k \\ &= (\Sigma^{-1} M A^{-1} \mathbf{S} \mathbf{e}_k)^T (\Sigma^{-1} M A^{-1} \mathbf{S} \mathbf{e}_k) + \beta (\mathbf{R} \mathbf{e}_k)^T (\mathbf{R} \mathbf{e}_k) \\ &= (\Sigma^{-1} M A^{-1} \mathbf{s}_k)^T (\Sigma^{-1} M A^{-1} \mathbf{s}_k) + \beta (\mathbf{R} \mathbf{e}_k)^T (\mathbf{R} \mathbf{e}_k) \\ &= (\Sigma^{-1} M \mathbf{h}_k)^T (\Sigma^{-1} M \mathbf{h}_k) + \beta \mathbf{r}_k^T \mathbf{r}_k \\ &= \sum_{\ell} \frac{1}{\sigma_{\ell}^2} (\mathbf{m}_{\ell}^T \mathbf{h}_k)^2 + \beta \mathbf{r}_k^T \mathbf{r}_k, \end{aligned}$$

where \mathbf{e}_k is the k th unit vector and

$$\mathbf{h}_k = A^{-1}S\mathbf{e}_k = A^{-1}\mathbf{s}_k, \quad \mathbf{r}_k = R\mathbf{e}_k.$$

That is, computing the information content vector \mathbf{j} requires one solve with the forward operator A for each of the source terms, plus a few matrix vector products.

An alternative way involves computing $MA^{-1} = (A^{-T}M^T)^T$ first. Because the vectors \mathbf{m}_ℓ form the rows of M (and so the columns of M^T), we can compute vectors

$$\mathbf{h}_\ell^* = A^{-T}\mathbf{m}_\ell,$$

and then recognize that

$$Q_{kk} = \sum_{\ell} \frac{1}{\sigma_{\ell}^2} ((\mathbf{h}_{\ell}^*)^T \mathbf{s}_k)^2 + \beta \mathbf{r}_k^T \mathbf{r}_k.$$

This approach requires solving a linear system with A^T for each measurement.

Which of the two ways of computing Q_{kk} is more efficient depends on whether there are more measurements than source terms, or the other way around.³

Regardless of the way Q_{kk} (and consequently $j_k = \sqrt{Q_{kk}}$) is computed, it can be interpreted as having contributions from all measurements (through the sum over ℓ) and from regularization. The scalar product $\mathbf{m}_{\ell}^T \mathbf{h}_k$ can be considered as the influence of the forward propagated sources (\mathbf{h}_k) on measurements. On the other hand, the equivalent term $(\mathbf{h}_{\ell}^*)^T \mathbf{s}_k$ corresponds to a view where we first compute an adjoint solution \mathbf{h}_{ℓ}^* that indicates which possible source terms affect a measurement functional, and then take the dot product with a concrete source \mathbf{s}_k . Both views represent the sensitivity of measurement functionals to sources.

Interestingly, the formula expresses the intuitive concept that *information is additive*: If there are no measurements and no regularization, then $j = 0$; each measurement in turn adds a non-negative contribution. Finally, because the measurement contributions to Q_{kk} are proportional to $\frac{1}{\sigma_{\ell}^2}$, we have the pleasant and reasonable property that *information is inversely proportional to the measurement uncertainty*.

3. Extension to infinite-dimensional inverse problems. We can extend the reasoning of the previous section to infinite-dimensional inverse problems. Specifically, let us consider the linear source identification problem

$$\mathcal{L}u(\mathbf{x}) = \sum_k q_k s_k(\mathbf{x}), \quad \forall \mathbf{x} \in \Omega, \quad (3.1)$$

where \mathcal{L} is a differential operator acting on functions defined on a domain $\Omega \subset \mathbb{R}^d$, and the equations are augmented by appropriate boundary conditions on $\partial\Omega$ whose details we will skip for the moment. As before, $s_k(\mathbf{x})$ are possible source vectors and q_k are their relative strengths. We again seek to identify source strengths q_k . Importantly, we will assume that the source terms are all of the form

$$s_k(\mathbf{x}) = \chi_{\omega_k}(\mathbf{x}),$$

where χ_{ω_k} is the characteristic function of a subdomain ω_k , and we assume that $\omega_k \cap \omega_l = \emptyset$ for $k \neq l$ and $\bigcup_k \bar{\omega}_k = \bar{\Omega}$. In other words, we seek to identify a source term that is a step function defined on a partition of the domain Ω .

³Clearly, both ways are expensive for real-world cases with many parameters and many measurements. We will come back to this in our conclusions and outlook, Section 6.

We will infer the source strengths q_k through (linear, noisy) measurements

$$z_\ell = \langle m_\ell, u \rangle + \varepsilon_\ell = \int_{\Omega} m_\ell(\mathbf{x})u(\mathbf{x}) \, dx + \varepsilon_\ell, \quad (3.2)$$

where ε_ℓ is again measurement noise assumed to have magnitude σ_ℓ . The formalism we will develop will allow us to assign an *information content* to each q_k . Because the source strengths q_k correspond to characteristic functions s_k of subdomains ω_k , the information content j_k divided by the volume $|\omega_k|$ will define *information density* $j(\mathbf{x})$, for which we can consider the limit case $|\omega_k| \rightarrow 0$. This limit is not computable, but we can hope to use finite subdivisions into regions ω_k that allow us to approximate it with reasonable accuracy.

To make these concepts concrete, in Section 5 we will consider this model where \mathcal{L} is an advection-diffusion operator, $\mathcal{L} = -D\Delta + \mathbf{b} \cdot \nabla$, and where z_ℓ correspond to weighted point measurements of $u(\mathbf{x})$ (or a well-defined approximation of point measurements if the solution u is not guaranteed to be a continuous function). This example is motivated by a desire to identify sources of air pollution from sparse measurements at a finite number of points.

3.1. Definition of the inverse problem. The inverse problem we have described in words above then has the following mathematical formulation, where we also include an L_2 regularization term:

$$\begin{aligned} \min_{\mathbf{q}, u} \mathcal{J}(\mathbf{q}, u) &= \frac{1}{2} \sum_{\ell} \frac{1}{\sigma_\ell^2} |\langle m_\ell, u \rangle - z_\ell|^2 + \frac{\beta}{2} \left\| \sum_k q_k s_k \right\|_{L_2(\Omega)}^2, \\ \text{such that } \mathcal{L}u &= \sum_k q_k s_k. \end{aligned} \quad (3.3)$$

In the same spirit as in the previous section, we can define a reduced objective function

$$\begin{aligned} \mathcal{J}_{\text{red}}(\mathbf{q}) &= \mathcal{J} \left(\mathbf{q}, \mathcal{L}^{-1} \sum_k q_k s_k \right) \\ &= \frac{1}{2} \sum_{\ell} \frac{1}{\sigma_\ell^2} \left| \left\langle m_\ell, \mathcal{L}^{-1} \sum_k q_k s_k \right\rangle - z_\ell \right|^2 + \frac{\beta}{2} \left\| \sum_k q_k s_k \right\|_{L_2(\Omega)}^2, \end{aligned} \quad (3.4)$$

which gives rise to a related stochastic inverse problem with a probability density $p(\mathbf{q}|\mathbf{z})$ defined as in (2.8).

3.2. Defining the information content. In the same spirit as in Section 2.2, we can again identify the *information content* associated with each parameter q_k via the precision, i.e., inverse of the variance $\text{var}_p(\mathbf{q})_k = \text{cov}_p(\mathbf{q})_{kk}$, and the estimate we have for the variance based on the Fisher information matrix.

In the finite-dimensional case, the Fisher information matrix I_p could be computed by solving one forward problem for each source vector \mathbf{s}_k . The same is true for the current infinite-dimensional situation:

PROPOSITION 3.1. *For the model problem defined above, the Fisher information matrix I_p defined in (2.11) has the following form:*

$$(I_p)_{kl} = Q_{kl} \quad (3.5)$$

where

$$Q_{kl} = \sum_{\ell} \frac{1}{\sigma_{\ell}^2} \langle m_{\ell}, h_k \rangle \langle m_{\ell}, h_l \rangle + \beta \int_{\Omega} s_k s_l, \quad (3.6)$$

and where h_k satisfies the equation

$$\mathcal{L}h_k(\mathbf{x}) = s_k(\mathbf{x}) \quad \forall \mathbf{x} \in \Omega, \quad (3.7)$$

again augmented by appropriate boundary conditions for h_k .

Proof. Recall that

$$(I_p)_{kl} = -\mathbb{E} \left[\frac{\partial^2}{\partial q_k \partial q_l} \ln p(\mathbf{q}|\mathbf{z}) \right], \quad \text{with} \quad p(\mathbf{q}|\mathbf{z}) = \kappa e^{-\frac{1}{2} \mathcal{J}_{\text{red}}(\mathbf{q})}.$$

Based on the definition of $\mathcal{J}_{\text{red}}(\mathbf{q})$ and the linearity of \mathcal{L} , we then obtain

$$\begin{aligned} (I_p)_{kl} &= \frac{\partial^2}{\partial q_k \partial q_l} \mathcal{J}(\mathbf{q}) \\ &= \frac{\partial^2}{\partial q_k \partial q_l} \left[\frac{1}{2} \sum_{\ell} \frac{1}{\sigma_{\ell}^2} \left\langle m_{\ell}, \mathcal{L}^{-1} \sum_r q_r s_r \right\rangle^2 + \frac{\beta}{2} \left\| \sum_r q_r s_r \right\|_{L_2(\Omega)}^2 \right] \\ &= \sum_{\ell} \frac{1}{\sigma_{\ell}^2} \langle m_{\ell}, \mathcal{L}^{-1} s_k \rangle \langle m_{\ell}, \mathcal{L}^{-1} s_l \rangle + \beta \int_{\Omega} s_k s_l, \end{aligned}$$

as claimed when using $h_k := \mathcal{L}^{-1} s_k$. \square

PROPOSITION 3.2. *The matrix Q that appears in the definition of the Fisher matrix in (3.6) can alternatively be expressed through the following formula:*

$$Q_{kl} = \sum_{\ell} \frac{1}{\sigma_{\ell}^2} \langle h_{\ell}^*, s_k \rangle \langle h_{\ell}^*, s_l \rangle + \beta \int_{\Omega} s_k s_l, \quad (3.8)$$

and where h_{ℓ}^* satisfies the equation

$$\mathcal{L}^* h_{\ell}^*(\mathbf{x}) = m_{\ell}(\mathbf{x}) \quad \forall \mathbf{x} \in \Omega, \quad (3.9)$$

where \mathcal{L}^* is the adjoint operator to \mathcal{L} , and with appropriate boundary conditions for h_{ℓ}^* .

Proof. The proposition follows from the observation that

$$\langle m_{\ell}, h_k \rangle = \langle m_{\ell}, \mathcal{L}^{-1} s_k \rangle = \langle \mathcal{L}^{-*} m_{\ell}, s_k \rangle = \langle h_{\ell}^*, s_k \rangle.$$

\square

REMARK 2. *As in the finite-dimensional case, the Fisher information matrix (3.5) is easy to compute for problems with either not too many parameters (using (3.6)) or not too many measurements (then using (3.8)). In either case, the functions h_k or h_{ℓ}^* can be computed independently in parallel. Which of the two forms is more efficient depends on whether there are more measurements than source terms or the other way around. That said, in the discussions below, we will want to let $|\omega_k| \rightarrow 0$ and consequently make the number of source terms infinite, and in that case the adjoint formulation in (3.8) provides the more useful perspective.*

The Fisher information matrix approximates the inverse of the covariance matrix, and the diagonal elements of the Fisher matrix I_p therefore provide a means to estimate the certainty in the corresponding parameters q_k . In the same way as for the finite-dimensional case in (2.13), we can then define an information content for the parameter k via

$$j_k := \sqrt{Q_{kk}}, \quad (3.10)$$

where now

$$Q_{kk} = \sum_{\ell} \frac{1}{\sigma_{\ell}^2} \langle m_{\ell}, h_k \rangle^2 + \beta \int_{\Omega} s_k^2 = \sum_{\ell} \frac{1}{\sigma_{\ell}^2} \langle h_{\ell}^*, s_k \rangle^2 + \beta \int_{\Omega} s_k^2.$$

3.3. Defining the information density. The discussions in the previous section did not make use of any particular properties of the source basis functions s_k . Let us now come back to the special case where we aim to identify a piecewise constant source function, i.e., where

$$s_k(\mathbf{x}) = \chi_{\omega_k}(\mathbf{x}).$$

In that case, we have that the information content for the parameter q_k associated with area ω_k is

$$j_k = \sqrt{Q_{kk}} = \sqrt{\sum_{\ell} \frac{1}{\sigma_{\ell}^2} \left(\int_{\omega_k} h_{\ell}^* \right)^2 + \beta |\omega_k|}. \quad (3.11)$$

Since this quantity scales with the size of the subdomains ω_k , it is reasonable to define a piecewise constant *information density* as

$$\begin{aligned} j(\mathbf{x})|_{\omega_k} &= \frac{1}{|\omega_k|} j_k = \frac{1}{|\omega_k|} \sqrt{Q_{kk}} \\ &= \sqrt{\sum_{\ell} \frac{1}{\sigma_{\ell}^2} \left(\frac{1}{|\omega_k|} \int_{\omega_k} h_{\ell}^* \right)^2 + \beta \frac{1}{|\omega_k|}} \end{aligned} \quad (3.12)$$

$$\approx \sqrt{\sum_{\ell} \frac{1}{\sigma_{\ell}^2} h_{\ell}^*(\mathbf{x})^2 + \beta \frac{1}{|\omega_k|}}. \quad (3.13)$$

We can make a number of observations based on these definitions, analogous to the finite-dimensional case of the previous section:

- Because the definition of h_{ℓ}^* is independent of the choice of ω_k , the formulas shown above can be interpreted as saying that the information density has a component that results from the measurements ℓ (and, in particular, grows monotonically with the number of measurements), and a component that results from regularization.
- As before, the information density is inversely proportional to the measurement uncertainties σ_{ℓ} .
- Regularization bounds the amount of information from below: $j(\mathbf{x})|_K \geq \sqrt{\beta/|\omega_k|}$. This dependence on the square root of the regularization parameter is well known [18].

REMARK 3. We end this section by remarking that the definitions of information content and information density above depend only on the forward operator and the measurement functionals, but not on concrete measured values z_ℓ . Consequently, and as anticipated, information quantities can be computed before measurements, and they are independent of the specific noise in measurements later obtained. We will come back to this point in Sections 5.2 and 5.3.1.

REMARK 4. We end this section by remarking that the idea of using variances Q_{kk} as a spatially variable measure of certainty is not new. For example, [35, Section 3.1] illustrates the spatially variable variance for a seismic imaging problem. Yet, the authors’ definition is unclear regarding the role of regularization, misses the square root, and is then discarded as not very useful.

4. Using information densities. Having shown a way to define an information density $j(\mathbf{x})$, the question is whether it is useful for anything. Indeed, there are numerous questions related to the practical solution of inverse problems for which information densities could be useful. In the following subsections, we therefore first outline three “vignettes” of situations in which the information density could be useful. In Section 5, we then illustrate one of these ideas using a concrete numerical example.

In the examples below, we will concretely consider the situation where we have discretized the source term $q(\mathbf{x}) = \sum_k q_k s_k(\mathbf{x})$ in (3.1) on a “mesh” \mathbb{T} , as is common in the finite element method. Because there are no differentiability requirements on $q(\mathbf{x})$, it is common to identify the source term as a piecewise constant function on this mesh, and in this case, the source functions s_k are the characteristic functions of the cells K of \mathbb{T} . By identifying the index k with a cell K , (3.10) and (3.11) then define an “information content” j_K for each cell K of the mesh.

4.1. Using information densities for regularization. As a first example of where we believe that information densities could be used, let us consider the regularization of inverse problems. In a large number of practical applications, one regularizes inverse problems by adding a penalty term to the misfit function, in hopes of penalizing undesirable aspects of the recovered function. For example, in our definition of the source identification problem in Section 3.1 (see also equation (3.3)), we have penalized the *magnitude* of the source term to be identified. The strength of this penalization is provided by the factor β .

A practical question is how large this factor β should be. Many criteria have been proposed in the literature [24], but, in practice many studies do not use any of these automatic criteria and instead choose values of β that yield “reasonable” results based on trial and error.

Moreover, it is clear to many practitioners that regularization may not be necessary to the same degree in all parts of the domain. For example, if measurements are available only in parts of the domain (say, on the boundary), then intuitively “more information” is available to identify source strengths close to the boundary than deep in the interior of the domain. A particularly obvious example is in seismic imaging: There, we can only accurately identify properties of the Earth in those places that are crossed by ray paths from earthquake sources (predominantly located at plate boundaries) to seismometer stations (predominantly located on land), but not in the rest of the Earth [6, 13, 30, 35, 38]. In the definition of the information density $j(\mathbf{x})$ in (3.13), this would imply that the first term under the square root would be large along these ray paths, but small elsewhere. In cases such as this, a reasonable approach would be to make the regularization parameter spatially variable: large where little information is available, and small where regularization is not as important, for example so that

$j(\mathbf{x}) \geq j_0$. This could be achieved by replacing the regularization term in (3.3) by

$$\frac{1}{2} \left\| \sqrt{\beta(\mathbf{x})} \left(\sum_k q_k s_k(\mathbf{x}) \right) \right\|_{L_2(\Omega)}^2,$$

and defining $\beta(\mathbf{x})$ in some appropriate way.

Using spatially variable regularization is not a new idea (see, for example, [4, 6, 13, 16, 34, 41, 45, 46]), though we are not aware of any references that would provide an overarching, systematic framework for choosing $\beta(\mathbf{x})$. In contrast, it is clear that the connection between information density $j(\mathbf{x})$ and $\beta(\mathbf{x})$ in (3.13) has the potential to provide such a systematic approach. A scheme based on this observation also satisfies other considerations that appear reasonable. For example, increasing the number of measurements, or decreasing the measurement error, leads to a larger information density and therefore to a smaller regularization term to satisfy $j(\mathbf{x}) \geq j_0$.

4.2. Using information densities to guide the discretization of an inverse problem. In actual practice, inverse problems are solved by discretization. In our derivation above, we have done so by choosing finitely many source functions s_k that we have assumed are the characteristic functions of “cells” K of some kind of mesh or subdivision of the domain Ω on which $q(\mathbf{x})$ is defined, and then expanded the function we seek as

$$q(\mathbf{x}) = \sum_k q_k s_k(\mathbf{x}).$$

A practical question is how to choose this subdivision into cells K . Oftentimes, the subdivision is chosen fine enough to resolve the features of interest but coarse enough to keep the computational cost in check. Regularization is frequently used to ensure that an overly fine mesh does not lead to unwanted oscillations in the recovered coefficients – in other words, to keep the problem reasonably well-posed.

Most often in the literature, the mesh for the inverse problem is either uniform or at least chosen a priori through insight into the problem (for approaches in the latter direction, see for example, [6, 14, 30]). On the other hand, discretization is a form of regularization, and it is reasonable to choose the mesh finer where more information is available – say, close to a measurement device – but coarser where our measurements have little information to offer. This idea has been used as a heuristic in the past [6], or at least mentioned (see Section 2.3 of [35] and the references therein), but, as with regularization, no overarching scheme is available to guide this choice of the mesh.

At the same time, an information density can provide such a guide to determine optimal cell sizes. First, we might conjecture that meshes should be graded in such a way that the information content of each cell (i.e., roughly the information density times the measure of a cell) is approximately equal among all cells. We will explore in detail how well this works in practice in Section 5. Second, in mathematical research, mesh refinement cycles are frequently terminated whenever we run out of memory, out of patience, or both, whereas in applications, mesh refinement is stopped whenever an expert deems the solution sufficiently accurate. Either approach is unsatisfactory, and the amount of information available per cell might provide a more rational criterion to stop mesh refinement.

4.3. Using information densities for experimental design. As a final example, let us consider optimal experimental design, that is, the question of what, how,

or where to measure so as to minimize the uncertainty in recovered parameters given a certain noise level in measurements.

Optimal experimental design for inverse problems is more difficult than for finite-dimensional parameter estimation problems because it is not entirely clear what the objective function should be when minimizing or maximizing by varying the specifics of measuring. For finite-dimensional problems, objective functions include the A -, C -, D -, E -, and T -optimality criteria, plus many variations [5].

For the infinite-dimensional case (or discretized versions thereof), the choice might be to maximize the information content in all of Ω , or a subset $\omega \subset \Omega$:

$$\phi(\{\mathbf{m}_k\}) = \int_{\omega} j(\mathbf{x}) \, dx,$$

where $\{\mathbf{m}_k\}$ denotes the set of measurements to be performed and optimization will typically happen over a set of implementable such measurements.

If $\omega = \Omega$, then the integral above reduces to

$$\phi(\{\mathbf{m}_k\}) = \sum_k \sqrt{Q_{kk}},$$

based on the definitions in (3.11) and (3.12). Recalling that the matrix Q is the Fisher information matrix, we recognize that the criterion ϕ above is similar to – but distinct from – the T -optimality criterion that maximizes the sum of diagonal entries of Q (i.e., the *trace* of Q).

5. Numerical examples of using information densities. The previous section provided three “vignettes” of how we imagine information densities could be used for practical computations. That said, exploring all of these ideas through numerical examples exceeds the reasonable length of a single publication, and as a consequence we will focus on only one of these applications: namely, mesh refinement.

In the following subsections, we will first lay out the inverse problem we will use as a test case. We will then show some numerical results that illustrate the use of information densities as applied to this problem.

All numerical results were obtained with a program that is based on the open-source finite element library `deal.II` [2, 3]. This program is itself also available under an open source license as part of the `deal.II` code gallery at <https://dealii.org/developer/doxygen/deal.II/CodeGallery.html> under the name “Information density-based mesh refinement”.

5.1. The test case. Let us consider the following question: Given an advection-diffusion problem for a concentration $u(\mathbf{x})$, can we identify the sources $q(\mathbf{x})$ of the concentration field from point measurements of u at points ξ_ℓ ? This kind of problem is widely considered in environmental monitoring of pollution sources [26, 32, 39, 44], and also when trying to identify the sources of nuclear radiation.

Mathematically, we will assume the concentration field satisfies the stationary advection-diffusion equation

$$\mathcal{L}u(\mathbf{x}) \equiv \mathbf{b}(\mathbf{x}) \cdot \nabla u(\mathbf{x}) - D\Delta u(\mathbf{x}) = q(\mathbf{x}) \quad \text{in } \Omega, \quad (5.1)$$

where \mathbf{b} is a (known) wind field and D is the (known) diffusion constant. For simplicity, we will assume homogenous Dirichlet boundary conditions

$$u = 0 \quad \text{on } \partial\Omega. \quad (5.2)$$




FIG. 5.1. *Left: The solution $u^*(\mathbf{x})$ of the forward problem from which we generate “synthetic” measurements z_ℓ via (5.3). Right: The source term $q^*(\mathbf{x})$ from which we compute synthetic measurements is constant and nonzero only in the solid red circle offset from the center; the detector locations $\xi_\ell, \ell = 1, \dots, L = 100$ are marked by dots.*

Concretely, for our computations, we will assume that $\Omega = (-1, 1)^2 \subset \mathbb{R}^2$ is a square, $D = 1$, and $\mathbf{b} = (100, 0)^T$. These choices lead to a Péclet number of 200; that is, the problem is advection dominated.

For the inverse problem, we ask whether we can recover the function $q(\mathbf{x})$ (or a discretized version of it) from measurements z_ℓ at a number of points $\xi_\ell \in \Omega, \ell = 1, \dots, L$. That is, we consider (3.2) with $m_\ell(\mathbf{x}) = \delta(\mathbf{x} - \xi_\ell)$ where σ_ℓ is the assumed noise level for the measurement at location ξ_ℓ (see below). We choose these points ξ_ℓ equally distributed around two concentric circles of radius 0.2 and 0.6, centered at the origin, with 50 points on each of the circles, for a total of $L = 100$ measurement points.

For our experiments, we will consider a situation where the data we have, z_ℓ , has been obtained by solving the forward problem with the finite element method, using a known source distribution $q^*(\mathbf{x})$ that is equal to one in a circle of radius 0.2 centered at $(-0.25, 0)^T$. A solution $u^*(\mathbf{x})$ can then be evaluated at the points ξ_ℓ to obtain “synthetic” measurements z_ℓ via

$$z_\ell = \langle m_\ell, u^* \rangle + \varepsilon_\ell = u^*(\xi_\ell) + \varepsilon_\ell. \quad (5.3)$$

We choose Gaussian noise $\varepsilon_\ell = N(0, \sigma_\ell^2)$ and set $\sigma_\ell = 0.1 \max_{\mathbf{x} \in \Omega} |u^*(\mathbf{x})|$.

To avoid an inverse crime, we solve for u^* on a mesh that is different from the meshes used for all other computations. The solution of this forward problem so computed to obtain synthetic measurements is shown in Fig. 5.1, along with the locations of the source term and the detector locations.

5.2. The inverse problem. The inverse problem we seek to solve is the identification of the source term $q(\mathbf{x})$ (which we approximate via a finite-dimensional expansion $\sum_k q_k s_k(\mathbf{x})$) in (5.1), based on the measurements described by (5.3). We approach this problem by reformulating it in the form of the constrained optimization problem (3.3), where we set the regularization parameter to $\beta = 10^4$. This problem is then solved by introducing a Lagrangian

$$\mathfrak{L}(u, q, \lambda) = \mathcal{J}(u, q) + \int_{\Omega} \lambda(\mathbf{x}) (\mathcal{L}u(\mathbf{x}) - q(\mathbf{x})) dx,$$

and then solving the linear system of partial differential equations that results by setting the derivatives of \mathfrak{L} to zero (that is, the optimality conditions). In strong

form, these optimality conditions read

$$\begin{aligned}\mathcal{L}u(\mathbf{x}) &= q(\mathbf{x}), \\ \mathcal{L}^*\lambda(\mathbf{x}) &= -\sum_{\ell} \frac{1}{\sigma_{\ell}^2} (u(\xi_{\ell}) - z_{\ell}) \delta(\mathbf{x} - \xi_{\ell}), \\ \beta q - \lambda &= 0.\end{aligned}\tag{5.4}$$

The solution of this system of equations is facilitated by discretizing on a finite element mesh. We use continuous, piecewise cubic elements for u and λ , and discontinuous, piecewise constants elements for q .⁴

The three components u, q, λ of this solution, computed on a very fine mesh with $256 \times 256 = 65\,536$ cells, for which the coupled problem has 1 248 258 unknowns, are shown in Fig. 5.2. The maximal value of the recovered source is less than half the maximal size of the “true” source, owing to the effect of the L_2 regularization term. As a consequence, the forward solution u is also too small. Furthermore the inverse problem places the source in a broader region than where it really is, but this is not surprising: In an advection-dominated problem, it is only possible to say with accuracy that the source is *upstream* of a detector, but not where in the upstream region it actually is unless another detector further upstream indicates that it must be downstream from the latter.

The adjoint variable λ clearly illustrates the effect that the adjoint operator \mathcal{L}^* transports information in the opposite direction $-\mathbf{b}$ of the forward operator \mathcal{L} , and that the sources of the adjoint equation are the residuals $-(u(\xi_{\ell}) - z_{\ell})\delta(\mathbf{x} - \xi_{\ell})$; here $u(\xi_{\ell}) - z_{\ell}$ reflects the measurement error and, based on our choice of noise above, is Gaussian distributed with both positive and negative values.

The bottom right panel of the figure also shows the information density $j(\mathbf{x})$ that corresponds to this problem, as defined in (3.13). It illustrates that, given the location of detectors and the nature of the equation, information is primarily available upstream of detector locations. Notably, and as mentioned in Remark 3, the information density is based solely on the operator \mathcal{L} and the measurement functionals m_{ℓ} , but not on the actual measurements z_{ℓ} (or the noise that is part of z_{ℓ}).

5.3. Choice of mesh for the inverse problem. The question of interest then is how we can use information densities for mesh refinement. To this end, we have repeated the computations discussed above, but instead of using a uniformly refined mesh, we have used a sequence of meshes in which we refine cells hierarchically so as to equilibrate the information content j_k of each cell ω_k , see (3.10), by always refining those cells that have the largest information content. The reconstructions and the sequence of meshes they are computed on are shown in Fig. 5.3. For comparison with the computations mentioned above and shown in Fig. 5.2, the rightmost mesh has 1642 cells and the coupled problem solved on it has 32 582 unknowns.

5.3.1. Comparison with other mesh refinement criteria. The relevant question to ask is whether this mesh is better suited to the task than any other mesh we could come up with. Answering this question is notoriously difficult in inverse problems because, in general, the exact solution of the problem is unknown if

⁴This choice of higher-order finite element spaces for u and λ is akin to solving for the forward and adjoint variables on a finer mesh than the source terms we seek to identify. As a result, we need not worry about satisfying discrete stability properties for the resulting saddle point problem. We can also, in essence, consider the forward and adjoint equation to be solved nearly exactly, with the majority of the discretization error resulting from the discretization of the source term $q(\mathbf{x})$.




FIG. 5.2. The solution of the inverse problem (5.4), computed on a very fine finite element mesh. Top left: The recovered primal variable $u(\mathbf{x})$, shown with the same scale for color and isocontours as in Fig. 5.1. Top right: The recovered sources $q(\mathbf{x})$. Bottom left: The recovered adjoint variable $\lambda(\mathbf{x})$. Bottom right: The information density $j(\mathbf{x})$ associated with this problem, as defined in (3.13).

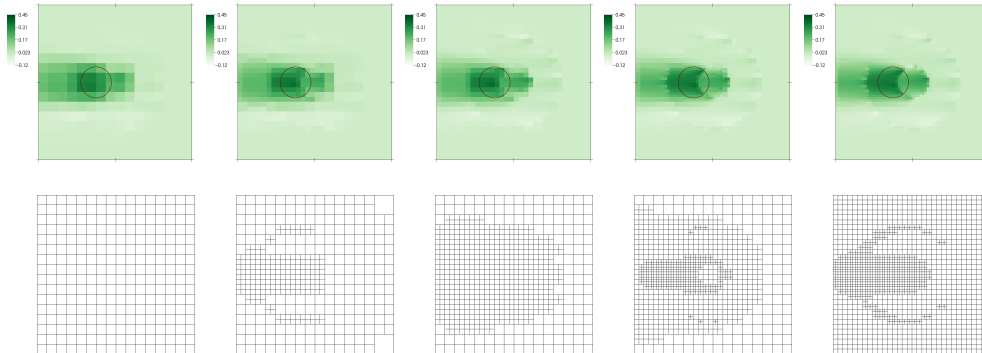


FIG. 5.3. Reconstructions (top) on a sequence of meshes (bottom) refined based on the information content of each cell of the mesh.

only finitely many measurements are available and if regularization is used. As a consequence, it is difficult to answer the question through comparison of convergence rates of different methods, for example.

However we can sometimes make intuitive comparisons based on experience on “how a good mesh should look”, even though for problems like the one under consideration, it is generally difficult to create such meshes by hand a priori. To this end, Fig. 5.4 shows the meshes generated by always refining those cells K that have the largest “refinement indicators” η_K defined in two different ways. In the top row of the figure, this indicator is the cell-wise norm of the residual of the third equation of (5.4) and is thus an a posteriori error indicator that can be derived in a way similar to that shown in [7–9]:

$$\eta_K = \|\beta q - \lambda\|_{L_1(K)}.$$




FIG. 5.4. Sequences of meshes generated by different mesh refinement criteria. Top: Mesh refinement is driven by an a posteriori error indicator. Bottom: Mesh refinement is driven by a smoothness indicator.

We will consequently refer to this quantity as the “error estimator”.

In the bottom row of the figure, we show meshes generated by evaluating a finite difference approximation $\nabla_h q(\mathbf{x})$ on each cell by comparing the values of $q|_K$ with the values of q on neighboring cells, and then computing

$$\eta_K = h_K \|\nabla_h q\|_{L_2(K)}.$$

This quantity is proportional to the interpolation error of a continuous function when approximated by a piecewise constant finite element function (as we do here); the indicator therefore measures where the piecewise constant approximation is likely poor. We will refer to this criterion as the “smoothness indicator”. It is also used in [38], for example.

As outlined in Section 4.2, the literature contains discussions of many other ways to refine meshes for the inverse problem, but we consider the two mentioned above as representative mesh refinement criteria to compare our approach against.

The meshes generated by these two criteria and shown in Fig. 5.4 are structurally similar to those generated based on the information content and shown in Fig. 5.3. However, they lack the top-bottom symmetry of the ones in Fig. 5.3 and look generally “less organized”, owing to the fact that *they are based on the solution of the inverse problem, which is subject to the noise in the measurements*, whereas the information density reflects only how much we *know* about the solution at a specific point in the domain – that is, a quantity that is *independent of the concrete realization of the noise that is part of the measurements*, see also Remark 3. Conceptually, the best mesh should be independent of the concrete realization of noise, although dependent on *what* is being measured. The refinement by information content allows us to construct the mesh even before solving the inverse problem because it does not depend on the solution of the inverse problem.

If the individual measurements m_ℓ had had differently sized measurement errors, then this would also have affected the information density-based mesh refinement and led to smaller cells where more accurate measurements are available. In contrast, there is no such direct dependence for the other refinement methods; rather, for those methods, variable noise levels only affect mesh refinement because λ indirectly depends on the error level.

5.3.2. Quantitative evaluation: Condition numbers of matrices. A more concrete comparison between meshes would be to measure the degree of ill-posedness of the problem. Of course, we use regularization to make the problem well-posed, but a well-chosen mesh results in a matrix after discretization that has a better condition number than a poorly chosen mesh, and for which the reconstruction is consequently less sensitive to noise. In practice, the condition number is a poor indicator since it considers only the largest and smallest eigenvalues; we hypothesize that a better criterion would be to ask how many “large” eigenvalues there are, and it is this criterion that we will consider below.

To test this hypothesis, let us consider the discretized version of (5.4). If we collect the degrees of freedom of a finite element discretization of u into a vector U , and similarly those of λ into a vector Λ and those of q into the vector P (a symbol chosen to avoid confusion with the matrix Q of Section 2), then (5.4) corresponds to the following system of linear equations after discretization by the finite element method:

$$\begin{aligned} AU &= BP, \\ A^T \Lambda &= -C(U - Z), \\ \beta MP - B^T \Lambda &= 0. \end{aligned} \tag{5.5}$$

Here, the matrix A corresponds to the discretized operator $\mathcal{L} = \mathbf{b} \cdot \nabla - D\Delta$ acting on the finite element space chosen to discretize the state and adjoint variables, and M is the mass matrix on the finite element space chosen for the source q – that is, on the set of piecewise constant functions $s_k(\mathbf{x})$ associated with the cells of the mesh. The matrix B results from the product $B_{ik} = (\varphi_i, s_k)_\Omega$ between the shape functions for u and q , and C corresponds to terms of the form $C_{ij} = \sum_\ell \frac{1}{\sigma_\ell^2} \varphi_i(\xi_\ell) \varphi_j(\xi_\ell)$. By noting that the matrices A and M are invertible, we can reduce this system of equations to an equation for P by repeated substitution to

$$HP = B^T A^{-1} CZ, \tag{5.6}$$

where the matrix H is the Schur complement,

$$H = B^T A^{-T} C A^{-1} B + \beta M. \tag{5.7}$$

The matrix H , which is symmetric and at least positive semidefinite, thus relates the vector of measurements Z to the vector of coefficients H we would like to recover. H can be thought of as the discretized counterpart to the matrix Q in (2.5). Each eigenvalue of H then corresponds to an eigenvector (“mode”) of the coefficient $q(\mathbf{x})$ we would like to recover. Moreover, large eigenvalues correspond to modes that are insensitive to noise, whereas small eigenvalues correspond to modes that are strongly affected by noise. As a consequence, we would like to aim for discretizations that result in many large and few small eigenvalues.

Fig. 5.5 provides a numerical evaluation of this perspective. It shows that when refining the mesh using the information content criterion, the eigenvalues of H are further to the “top right” – in other words, there are more large eigenvalues than when using refinement by the error estimator or the smoothness indicator. This pattern persists after both three refinement cycles (the left part of the figure) and six refinement cycles (the right part).

The stair-step structure of the figure results from the fact that mesh refinement turns one large cell into four small ones. Consequently in general one large eigenvalue

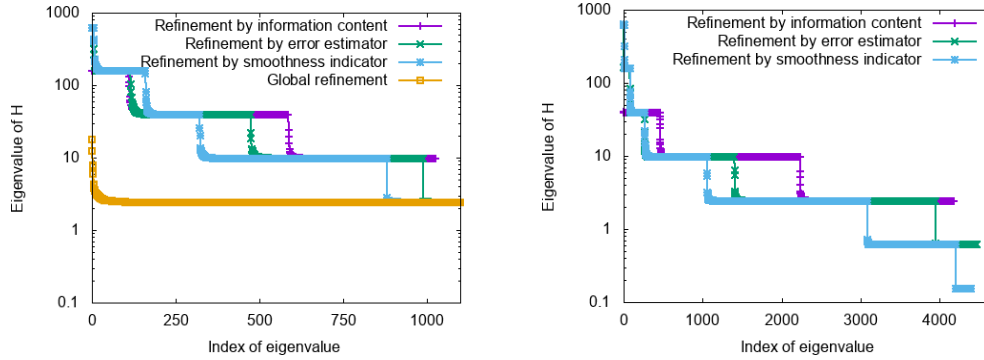


FIG. 5.5. Comparison of the eigenvalues of the resolution matrix H for different ways of refining meshes. Left: After three refinement cycles, yielding problems with approximately 1,000 parameters to identify for the adaptive refinement criteria, and 16,384 for global refinement. Right: After six refinement cycles, resulting in problems with approximately 4,000 parameters for the adaptive refinement criteria. (Global refinement would have resulted in more than one million parameters; the eigenvalues of this matrix could not be computed.)

turns into four smaller eigenvalues. By counting the number of derivatives present in the operators that enter into H , we can conjecture that the conditioning of the problem scales with the mesh size h squared; indeed, the levels in the plots confirm that each mesh refinement step reduces the size of the smallest eigenvalues by approximately a factor of $(h_{\text{large}}/h_{\text{small}})^2 = 4$.

The left part of the figure also shows the eigenvalues of H for meshes constructed via global refinement, where the globally refined mesh is chosen so that it has the same finest resolution as the adaptively refined ones. Global refinement results in vastly more unknowns than for the adaptively refined meshes, with a large majority of eigenvalues small. These eigenvalues all correspond to small cells to the right of the array of detectors where very little information is available. Given the size of the problem, we were not able to compute the eigenvalues of H following six global refinement steps; the corresponding data are therefore omitted in the right panel.

The comparison shown in the figure confirms the suspicion we have laid out at the beginning of the section: namely, that refining the mesh based on information contents leads to an inverse problem that is better posed than those that result from any of the other refinement criteria we have compared with, in the sense that our approach leads to more large eigenvalues.

6. Conclusions and outlook. In this paper, we have used a statistical approach to define how much “information” we have about the parameter that is recovered in an inverse problem. More concretely, we have defined a density $j(\mathbf{x})$ that corresponds to *how much we know* about the solution at a given point \mathbf{x} , and derived an explicit expression for it that can be computed. We have then outlined a number of ways in which we believe that this information density can be used, via three “vignettes”. Finally, we have assessed one of these application areas numerically and showed that basing mesh refinement for inverse problems on information densities indeed leads to meshes that not only visually look more suited to the task than other criteria, but also quantitatively lead to better-posed discrete problems.

At the same time, this paper did not address many areas that would make for very natural next questions, including the following:

- In our work, we have chosen a simple L_2 regularization term $\frac{\beta}{2} \|q\|_{L_2}^2$, see

(3.3). In practice, however, regularization terms would typically be used that penalize oscillations, for example using terms of the form $\frac{\beta}{2} \|\nabla^\alpha q\|_{L_2}^2$, $\alpha \geq 1$. How this would affect the definition of the information density would be interesting to ask in a future study. In order to consider the limit of infinite-dimensional inverse problems, we need to ensure that the covariance operators that result from regularization are of trace class (see, for example, [15, 40]), which is also likely a precondition for the definition of information densities.

- For many inverse problems – such as ultrasound or seismic imaging, or electrical impedance tomography – the quantity we would like to identify is not a right-hand source term, but a coefficient in the operator on the left side of the equation. In these cases, the definition of the information density will have to be linearized around the solution of the inverse problem, which then may make the definition of $j(\mathbf{x})$ dependent on actual noise values. How this affects the usefulness of the information density is a priori unclear, but it is clear that it would at least require solving the inverse problem before we can compute $j(\mathbf{x})$.
- The computation of $j(\mathbf{x})$ requires the solution of a number of forward or adjoint problems, which is expensive, especially for three-dimensional inverse problems, even though these problems are all independent of each other and can be computed in parallel. At the same time, although we have shown in Sections 2 and 3 what is necessary to form the complete matrix Q , we only need the diagonal entries of this matrix, see (3.12). It is conceivable that one can compute approximations to the entries Q_{kk} more cheaply, for example by random projections, low-rank approximations, or hierarchical low-rank approximations. Such ideas can, for example, be found in [19, 20], and in [15, Section 5] and [1] for the closely related reduced Hessian matrix.

We leave the exploration of these topics to future work.

Acknowledgments. WB would like to thank Dustin Steinhauer for an introduction to the Fisher information matrix many years ago. WB’s work was partially supported by the National Science Foundation under award OAC-1835673; by award DMS-1821210; and by award EAR-1925595. CRJ’s work was partially supported by the Intel Graphics and Visualization Institutes of XeLLENCE, the National Institutes of Health under award R24 GM136986, the Department of Energy under grant number DE-FE0031880, and the Utah Office of Energy Development.

This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy’s National Nuclear Security Administration under contract DE-NA-0003525. SAND2022-11142 O.

Bibliography.

REFERENCES

- [1] I. AMBARTSUMYAN, W. BOUKARAM, T. BUI-THANH, O. GHATTAS, D. KEYES, G. STADLER, G. TURKIYYAH, AND S. ZAMPINI, *Hierarchical matrix approximations of Hessians arising in inverse problems governed by PDEs*, SIAM J. Sci. Comput., 42 (2020), pp. A3397–A3426.

- [2] D. ARNDT, W. BANGERTH, B. BLAIS, M. FEHLING, R. GASSMÖLLER, T. HEISTER, L. HELTAI, U. KÖCHER, M. KRONBICHLER, M. MAIER, P. MUNCH, J.-P. PELTERET, S. PROELL, K. SIMON, B. TURCK SIN, D. WELLS, AND J. ZHANG, *The deal.II library, version 9.3*, Journal of Numerical Mathematics, 29 (2021), pp. 171–186.
- [3] D. ARNDT, W. BANGERTH, D. DAVYDOV, T. HEISTER, L. HELTAI, M. KRONBICHLER, M. MAIER, J.-P. PELTERET, B. TURCK SIN, AND D. WELLS, *The deal.II finite element library: design, features, and insights*, Computers & Mathematics with Applications, 81 (2021), pp. 407–422.
- [4] S. R. ARRIDGE AND M. SCHWEIGER, *Inverse methods for optical tomography*, in Information Processing in Medical Imaging, Harrison H. Barrett and A. F. Gmitro, eds., Berlin, Heidelberg, 1993, Springer Berlin Heidelberg, pp. 259–277.
- [5] A. C. ATKINSON AND A. N. DONEV, *Optimum Experimental Design*, Clarendon Press, Oxford, 1992.
- [6] L. AUER, L. BOSCHI, T. W. BECKER, T. NISSEN-MEYER, AND D. GIARDINI, *Savani: A variable resolution whole-mantle model of anisotropic shear velocity variations based on multiple data sets*, J. Geophys. Res. Solid Earth, 119 (2014), pp. 3006–3034.
- [7] W. BANGERTH, *Adaptive Finite Element Methods for the Identification of Distributed Parameters in Partial Differential Equations*, PhD thesis, University of Heidelberg, 2002.
- [8] W. BANGERTH AND A. JOSHI, *Adaptive finite element methods for the solution of inverse problems in optical tomography*, Inverse Problems, 24 (2008), p. 034011.
- [9] W. BANGERTH AND R. RANNACHER, *Adaptive Finite Element Methods for Differential Equations*, Birkhäuser Verlag, 2003.
- [10] H. T. BANKS, K. HOLM, AND F. KAPPEL, *Comparison of optimal design methods in inverse problems*, Inverse Problems, 27 (2011), p. 075002.
- [11] A. J. BERGLUND, *Statistics of camera-based single-particle tracking*, Physical Review E, 82 (2010).
- [12] J. G. BERRYMAN, *Analysis of approximate inverses in tomography I. Resolution analysis of common inverses*, Optimization and Engineering, 1 (2000), pp. 87–115.
- [13] R. BONADIO, S. LEBEDEV, T. MEIER, P. ARROUCAU, A. J. SCHAEFFER, A. LICCIARDI, M. R. AGIUS, C. HORAN, L. COLLINS, B. M. O’REILLY, P. W. READMAN, AND IRELAND ARRAY WORKING GROUP, *Optimal resolution tomography with error tracking and the structure of the crust and upper mantle beneath Ireland and Britain*, Geophysical Journal International, 226 (2021), pp. 2158–2188.
- [14] L. BORCEA, V. DRUSKIN, AND A. V. MAMONOV, *Circular resistor networks for electrical impedance tomography with partial boundary measurements*, Inverse Problems, 26 (2010), pp. 045010/1–32.
- [15] T. BUI-THANH, O. GHATTAS, J. MARTIN, AND G. STADLER, *A computational framework for infinite-dimensional bayesian inverse problems part i: The linearized case, with application to global seismic inversion*, SIAM J. Sci. Comput., 35 (2013), pp. A2494–A2523.
- [16] C. COHEN-BACRIE, Y. GOUSSARD, AND R. GUARDO, *Regularized reconstruction in electrical impedance tomography using a variance uniformization constraint*, IEEE Trans. Med. Imag., 16 (1997), pp. 562–571.
- [17] T. M. COVER AND J. A. THOMAS, *Elements of Information Theory*, Wiley, 2006.
- [18] D. R. COX AND D. V. HINKLEY, *Theoretical Statistics*, Chapman, 1974.
- [19] S. DAS, *Efficient calculation of Fisher information matrix: Monte Carlo approach using prior information*, PhD thesis, Johns Hopkins University, 2007.
- [20] S. DAS, J. C. SPALL, AND R. GHANEM, *Efficient monte carlo computation of fisher information matrix using prior information*, Computational Statistics & Data Analysis, 54 (2010), pp. 272–289.
- [21] A. DATTA-GUPTA, J. XIE, N. GUPTA, M. J. KING, AND W. J. LEE, *Radius of investigation and its generalization to unconventional reservoirs*, J. Petroleum Techn., (2011), pp. 52–55.
- [22] M. H. DEGROOT, *Optimal statistical decisions*, Wiley-Interscience, 1970.
- [23] A. F. EMERY AND A. V. NENAROKOMOV, *Optimal experiment design*, Measurement Science and Technology, 9 (1998), pp. 864–876.
- [24] H. W. ENGL, M. HANKE, AND A. NEUBAUER, *Regularization of Inverse Problems*, Kluwer Academic Publishers, Dordrecht, 1996.
- [25] Y. FAVENNEC, *Hessian and Fisher matrices for error analysis in inverse heat conduction problems*, Numerical Heat Transfer, Part B: Fundamentals, 52 (2007), pp. 323–340.
- [26] A. JAMSHIDI, J. M. V. SAMANI, H. M. V. SAMANI, A. ZANINI, M. G. TANDA, AND M. MAZAHERI, *Solving inverse problems of unknown contaminant source in groundwater-river integrated systems using a surrogate transport model based optimization*, Water, 12 (2020).
- [27] J. KAIPIO AND E. SOMERSALO, *Statistical and Computational Inverse Problems*, Springer, 2006.

- [28] S. M. KAY, *Fundamentals of Statistical Signal Processing, Volume I: Estimation Theory*, Prentice Hall, 1993.
- [29] F. J. KUCHUK, *Radius of investigation for reserve estimation from pressure transient well tests*, SPE, (2009), pp. 120515/1–22.
- [30] Q. LIU AND Y. J. GU, *Seismic imaging: From classical to adjoint tomography*, Tectonophysics, 566–567 (2012), pp. 31–66.
- [31] R. MONTELLI, G. NOLET, F. A. DAHLEN, AND G. MASTERS, *A catalogue of deep mantle plumes: New results from finite-frequency tomography*, Geochemistry, Geophysics, Geosystems, 7 (2006), pp. n/a–n/a.
- [32] R. M. NEUPAUER, B. BORCHERS, AND J. L. WILSON, *Comparison of inverse methods for reconstructing the release history of a groundwater contamination source*, Water Resources Research, 36 (2000), pp. 2469–2475.
- [33] S. NORDEBO AND M. GUSTAVSSON, *A priori modeling for gradient based inverse scattering algorithms*, Progress in Electromagnetics Research B, 16 (2009), pp. 407–432.
- [34] B. W. POGUE, T. O. MCBRIDE, J. PREWITT, U. L. ÖSTERBERG, AND K. D. PAULSEN, *Spatially variant regularization improves optical tomography*, Appl. Optics, 38 (1999), pp. 2950–2961.
- [35] N. RAWLINSON, A. FICHTNER, M. SAMBRIDGE, AND M. K. YOUNG, *Seismic tomography and the assessment of uncertainty*, in Advances in Geophysics, Elsevier, 2014, pp. 1–76.
- [36] J. RITSEMA, A. DEUSS, H. J. VAN HEIJST, AND J. H. WOODHOUSE, *S4ORTS: a degree-40 shear-velocity model for the mantle from new rayleigh wave dispersion, teleseismic traveltimes and normal-mode splitting function measurements*, Geophysical Journal International, 184 (2010), pp. 1223–1236.
- [37] E. ROMERA, P. SÁNCHEZ-MORENO, AND J. S. DEHESA, *The Fisher information of single-particle systems with a central potential*, Chemical Physics Letters, 414 (2005), pp. 468–472.
- [38] M. SAMBRIDGE AND N. RAWLINSON, *Seismic tomography with irregular meshes*, in Seismic earth: Array analysis of broadband seismograms, A. Levander and G. Nolet, eds., American Geophysical Union, 2005, pp. 49–65.
- [39] T. H. SKAGGS AND Z. J. KABALA, *Recovering the release history of a groundwater contaminant*, Water Resources Research, 30 (1994), pp. 71–79.
- [40] A. M. STUART, *Inverse problems: A Bayesian perspective*, Acta Numerica, 19 (2010), pp. 451–559.
- [41] L. TANG, G. HAMARNEH, AND R. ABUGHARBIH, *Reliability-driven, spatially-adaptive regularization for deformable registration*, in Biomedical Image Registration, 4th International Workshop, WBIR 2010, Lübeck, Germany, July 11–13, 2010. Proceedings, B. Fischer, B. M. Dawant, and C. Lorenz, eds., vol. 6204 of Lecture Notes in Computer Science, Springer, 2010, pp. 173–185.
- [42] A. TARANTOLA, *Inverse Problem Theory and Methods for Model Parameter Estimation*, SIAM, 2004.
- [43] D. W. VASCO, H. KEERS, AND K. KARASAKI, *Estimation of reservoir properties using transient pressure data: An asymptotic approach*, Water Resources Res., 36 (2000), pp. 3447–3465.
- [44] J. A. VRUGT, P. H. STAUFFER, T. WÖHLING, B. A. ROBINSON, AND V. V. VESSELINOV, *Inverse Modeling of Subsurface Flow and Transport Properties: A Review with New Developments*, Vadose Zone Journal, 7 (2008), pp. 843–864.
- [45] M. W. WOOLRICH, T. E. J. BEHRENS, C. F. BECKMANN, AND S. M. SMITH, *Mixture models with adaptive spatial regularization for segmentation with an application to MRI*, IEEE Trans. Med. Imag., 24 (2005), pp. 1–11.
- [46] Y. YIN, E. A. HOFFMAN, AND C.-L. LIN, *Lung lobar slippage assessed with the aid of image registration*, Med. Image. Comput. Comput. Assist. Interv., 13 (2010), pp. 578–585.
- [47] H. ZHANG AND C. H. THURBER, *Estimating the model resolution matrix for large seismic tomography problems based on lanczos bidiagonalization with partial reorthogonalization*, Geophysical Journal International, 170 (2007), pp. 337–345.