# Samples, Symmetries and Extensions: Exploring Parameter Space in Nonlinear Problems

# **Andrew Curtis**

University of Edinburgh, School of Geosciences, EH9 3FE, Edinburgh, U.K. Andrew.Curtis@ed.ac.uk

This article has yet to be formally peer reviewed. It will be submitted to a peer-reviewed journal for publication, so subsequent versions may have slightly different content. If accepted, the final version of this manuscript will be available via the 'Peer-reviewed Publication DOI' link. Please feel free to contact the author with any comments or feedback.

# Abstract

Many scientific and technological advances require the values of a set of parameters to be constrained or estimated using recorded data. Geophysical parameters usually describe a geological process, or the state of the Earth at a particular time. It is often possible to model data that would be recorded for particular parameter values using a computable and in general nonlinear function, where its inverse does not exist as a unique-valued mapping. The space of parameter values must then be explored to find values that are consistent with measured data. This may be achieved by sampling values at a sufficiently dense set of points in parameter space, evaluating the forward function for each sample, and testing their modelled data against the measured data. However, for manyparameter (high-dimensional) systems, achieving sufficient sample density may be infeasible due to the 'curse of dimensionality' – the exponential increase in sampling required to achieve similar sample density in spaces of increasing dimensionality. This article shows that for each sample, significantly more information may be available than the forward function evaluation alone, due to symmetries and other properties of the physical system. Call this additional information the extension of a sample. In travel time tomographic imaging, the extension of every sample is shown to provide forward function values throughout continuous parameter semi-subspaces, almost for free. The dimensionality of these subspaces increases exponentially with the dimensionality of the problem, almost matching the rate of information increase required by the so-called curse of dimensionality. The use of extensions within physics-based sampling schemes may therefore contribute to increase the dimensionality of problems in which one can explore parameter space, solve inverse or inference problems, or provide distributed information about forward function values directly.

# Introduction

In many fields of science and engineering, important processes or properties are described by mathematical parameters. Geophysical parameters usually describe geological processes for example in the study of geodynamics, or the state of the Earth at a particular time as in tomographic imaging. Advances are made by estimating the family of possible parameter values from recorded data, a process referred to as Geophysical inference.

Say the values of parameters m in parameter space M must be estimated using recorded data d in data space D. It is often the case that it is possible to estimate d for any particular values of m using some known, computable, in general nonlinear function f(m), but the inverse function is unknown and does not exist as a unique-valued mapping. Since the goal is to estimate m given values of d, inverse or inference theory must then be deployed.

The most commonly used inverse methods involve linearising f(m) around an initial estimate of the parameter values, solving the resulting inverse problem of estimating values of m that best fit recorded data under that linearised (approximate) version of f, then re-linearising the function around the new parameter estimates and solving the newly linearised inverse problem. This process is iterated to convergence. Given parameterised uncertainty on data d, corresponding uncertainty in values of m may be estimated using the gradients of f linearised around the final parameter estimate (*Tarantola 2005*).

While apparently computationally efficient, linearised methods may fail to converge towards the best-fit solution when f is nonlinear, and whether or not the best-fit solution has been found remains unknown. In addition, gradient-based uncertainty estimates only account for the linearised approximation to the physics and this is only valid near the solution found. As a result, in significantly nonlinear problems, both the minimum-misfit solution found, and corresponding linearised uncertainty estimates may be highly misleading as shown by Galetti et al. (2015) and Zhang et al. (2018).

It is therefore increasingly common to solve inverse problems without linearization, using samplingbased methods. These methods explore space *M* by evaluating the consistency between parameter values *m*, recorded data *d*, and any other pertinent *a priori* (prior) information that is independent of the data, at a succession of points or samples of *M*. The samples may be chosen pseudo-randomly as in Monte Carlo methods (*Gelfand & Smith 1990; Mosegaard & Tarantola 1995; Sambridge 1999; Malinverno 2002*), from a deterministic set of values for each parameter as in grid-search methods (Lomax & Curtis 2001; Lomax et al., 2009), or using a variety of other algorithms such as particle swarms (Kennedy & Eberhart 2001), simulated annealing (Kirkpatrick et al., 1983), genetic algorithms (*Goldberg 1989; Sambridge and Gallagher 1993*) or variational methods (*Zhang & Curtis 2020a,b*). Several of these can be tuned either to converge around already-discovered models that are consistent with recorded data and prior information, or to explore space *M* more broadly.

In a third class of methods, samples can be distributed probabilistically to reflect to the state of knowledge prior to collecting any current data set. These so-called prior samples can then be used to train machine learning or other algorithms to solve inverse problems swiftly for any particular data set that may arise in future, and to provide fully nonlinear parameter uncertainty estimates (Devilee et al., 1999; Meier et al., 2007a,b; Maiti et al. 2007; Shahraeeni & Curtis 2011; Shahraeeni et al., 2012; Käufl et al., 2016; Karmakar et al., 2018; Nawaz & Curtis 2019; Earp and Curtis 2020; Earp et al., 2020). Key for these methods is the ability to sample the set of parameter values sufficiently densely to represent all possible true values by interpolation or regression through these samples.

In all of the above sampling algorithms, the function f is evaluated for each sample. These evaluations may be used to find a set of models that are consistent with measured data, or for more formal uncertainty analysis using Bayesian inference (Tarantola 2005). However, in all cases the principle holds that where the function value has not been evaluated, it remains unknown. If we wish to infer results from such explorations it is therefore important to have included samples that span all relevant parts of parameter space M (those with a significant probability of representing the true state of nature) with reasonable sample density. This is problematic because *a priori* we do not know which parts of M are relevant, so a key part of Geophysical inference is to explore M with sampling that is sufficiently dense to discover high probability regions.

The barriers to exploring space M are almost always a combination of:

- (i) *M* is too high-dimensional to be explored due to the 'curse of dimensionality' (*Curtis & Lomax 2001*).
- (ii) Prior information is too weak to constrain parameters to a sufficiently small sub-volume of *M*.
- (iii) The cost of evaluating forward function f for each sample is too high.

The first barrier reflects the so-called curse of dimensionality, which states that the number of samples required to explore M with a given sample density expands exponentially with the number of dimensions (*Curtis & Lomax 2001*). The second simply states that prior information is too weak to overcome the exponential expansion in computation required by the curse. If parameters are numerous and (i) and (ii) apply, then barrier (iii) is the least significant: one could reduce the cost of the forward evaluation almost to zero, and it would make little or no difference to our practical ability to explore high-dimensional problems by sampling. The curse of dimensionality is so extreme that without prior information to remove its exponential expansion, even storing dense sample sets on disk becomes impossible as the number of parameters gets large. Nevertheless, the third barrier remains important for low-dimensional problems as it indicates that even then the expense of computing the forward function may preclude sufficiently dense sampling.

Any attempt to overcome these barriers using smarter algorithms must take into account the '*No Free Lunch theorem*' (NFL). Loosely speaking, this states that any one optimisation algorithm performs no better than any other when averaged over all problems (Wolpert & McReady 1987). The proof of NFL covers a broad class of optimisation algorithms, and since the best sampling methods are those that select an 'optimal' set of samples according to some criterion, NFL also applies to sampling algorithms. NFL is important because it implies that more efficient sampling methods cannot even be conceived of in general: they can only be devised for limited and specific classes of problems.

Standard sampling methods therefore often follow a pragmatic strategy. They are designed to increase efficiency by exploring *M* only at a level of detail that is proportional to the probability of parameters being true. In Bayesian solutions, this consistency is embodied in a so-called *posterior* probability density function (pdf) constructed using Bayes rule, which combines prior information about parameters with information from the recorded data (shown below). A common sampling strategy is to attempt to select samples with sample-density proportionate to this pdf. This is called importance sampling, and for certain classes of problems it may be approximated using Markov chain Monte Carlo algorithms (*Green 1995; Mosegaard & Tarantola 1995; Malinverno 2002; Bodin & Sambridge 2009; Fichtner et al., 2019*) and variational methods (Liu & Wang 2016; Zhang & Curtis 2020a,b), or achieved exactly using so-called exact sampling methods (Propp & Wilson 1996; Walker & Curtis 2014).

With the exception of exact sampling methods, such algorithms only distribute their samples proportionally to the true posterior pdf in the limit of infinite samples. Indeed, their mathematical proofs assume that in the infinite limit they explore all parts of parameter space for which the posterior pdf is positive. Nevertheless, many importance sampling algorithms do not explore parameter space efficiently. Instead they tend to focus samples – indeed, this is often perceived as their strength. In practice therefore, importance sampling algorithms are approximate, use a limited number of samples, and explore in detail only the higher-probability regions of parameter space that happen to have been discovered using the existing set of samples. Any inferences made from a finite set of samples are thus potentially biased to an unknown extent because unexplored parts of M that have non-zero posterior probability are not taken into account.

If we are to believe conclusions reached using sampling methods, it is therefore important to reduce the computational power required to explore all relevant parts of M. This has been attempted previously by reducing the dimensionality of the problem to be solved (Douma et al., 1996; Grana et al., 2019), tightening *a priori* constraints on parameter values (Curtis & Wood 2004; Walker & Curtis, 2014; Nawaz & Curtis 2019; Linde et al., 2015), developing more efficient forward computations (Rawlinson & Sambridge 2005; Nissen-Meyer et al., 2014; van Driel et al., 2015; Krischer et al., 2017), approximating the forward function with an emulator that can be explored more rapidly (Das et al., 2018; Moseley et al., 2020), or improving predictions of where samples might usefully be located in unexplored parts of M (e.g., Fichtner et al., 2019; Khoshkholgh et al., 2020). However, in all such studies the same principle holds: where the parameter space M has not been sampled, the value of fis unknown.

By contrast, the introduction of new analytical results has been shown to allow sampling to be avoided either partially or entirely in certain classes of spatially structured inverse problems (e.g., Nawaz & Curtis 2017, 2018). Many Geophysical problems are formulated in a spatial domain, and have additional underlying physical structure that tends to be ignored. The rest of this article therefore investigates the use of information about the structure of the physical forward function f to improve sampling algorithms.

Most of the algorithms cited above require that f has properties of continuity or smoothness, but we may have far more knowledge about its structure. While attempts have been made to include physical symmetries in sampling methods, for example in Physics (Chuang Chen et al., 2018; Ji-Yao Chen et al., 2018) and Genetics (e.g., Felsenstein 1981), that work is highly specialised to those fields and not easily transferable. This article shows how our understanding of the *geophysics* of problems to be solved may lead to highly efficient exploration of parameter space.

Symmetries of a function f are transformations of parameter values T(m) that leave f(m) unchanged, such that f(m) = f[T(m)]. The existence of symmetries thus implies that more information is available from each sample  $m_i$  than simply  $f(m_i)$  (where i indexes the sample number): we also know the value of f evaluated at all possible values of  $T(m_i)$ . Let us call this extra information the *extension* of sample  $m_i$ , defined as all information about other forward function evaluations that is available from  $f(m_i)$  without requiring any further evaluation of f.

This article shows that while some types of sample extension are already used in Geophysical applications, using them directly for sampling and exploration problems can lead to radically different algorithms. Information in an extension depends on the particular physics of the problem, so constructing extensions-based sampling algorithms circumnavigates limitations imposed by the No Free Lunch theorem.

We examine the example problem of finding maps of wave or particle speed that are consistent with measured travel time data. Extensions offer the possibility for extreme reductions in the number of samples required to explore parameter space to any desired level of detail compared to standard search methods. The most informative samples are found using a physics-based criterion which may eventually lead to methods that involve no random sampling at all. And importantly, as the number of parameters increases, information provided by these extensions increases approximately proportionately to the added volume of parameter space (i.e., exponentially). This suggests that it may be possible to explore parameter space despite the curse of dimensionality, at least in some specific classes of problems.

Similar models, data and physics to the above example occur in diverse problems of acoustic, seismic and electromagnetic travel time tomography or in positioning and location problems, and in inversions for parameters controlling fluid flow given fluid break-through times. These occur in Geophysical applications, and in other fields of science and engineering. Sample extensions may therefore significantly expand the range of scenarios in which fully nonlinear uncertainty analysis can be applied.

# Inverse Problem Formulation

Say we seek to constrain the values of a set of p parameters described by vector  $\mathbf{m} \in M_P$  where parameter space  $M_P$  is defined by parameterisation P. These will be constrained using recorded data  $\mathbf{d}$  in data space  $D_S$  which is defined as the set of all possible measured values using survey or experimental design S. Variable P defines how the mathematical model  $\mathbf{f}(\mathbf{m}): M_P \to D_S$  that relates parameters to data is parameterised since this is usually subject to choice; for example a spatial model might be represented within a variety of coordinate axis orientations and scalings, and may be parameterised using continuous basis functions of arbitrary resolution or approximated using different families of discrete cells. Variable S defines data types, and how, where and when the data are measured, all of which may also be a matter of choice.

It is common to use Bayes rule to solve such problems:

$$\mathcal{P}(\boldsymbol{m}|\boldsymbol{d}) = \frac{\mathcal{P}(\boldsymbol{d}|\boldsymbol{m}) \,\mathcal{P}(\boldsymbol{m})}{\mathcal{P}(\boldsymbol{d})} \tag{1}$$

In equation (1),  $\mathcal{P}(d|m)$  is referred to as the *likelihood* and measures the extent to which parameter values m are consistent with the measured data, and  $\mathcal{P}(m)$  is the *prior* pdf which defines the state of information about m independent of information in the data.  $\mathcal{P}(m|d)$  is called the posterior probability distribution function (pdf) and represents the final state of information about parameters m given the prior information and the data d, and is the solution to the overall inverse problem.  $\mathcal{P}(d)$  is called the *evidence* and is a measure of the consistency between data d, prior information about values m, and the mathematical model f(m) that is used to relate the two within the likelihood.

Notice that all terms in equation (1) depend on either the model or data parameterisations or both, and that such dependencies were dropped from notation. The full expression is

$$\mathcal{P}(\boldsymbol{m}|\boldsymbol{d}, \boldsymbol{P}, \boldsymbol{S}) = \frac{\mathcal{P}(\boldsymbol{d}|\boldsymbol{m}, \boldsymbol{P}, \boldsymbol{S}) \,\mathcal{P}(\boldsymbol{m}|\boldsymbol{P})}{\mathcal{P}(\boldsymbol{d}|\boldsymbol{P}, \boldsymbol{S})} \tag{2}$$

The prior pdf does not depend on S since under the Bayesian approach it represents information that exists independently of the current data set. We will use the notation in equation (1) but dependencies in equation (2) are implicit.

It is often the case that data are assumed to be mutually independent, and to have a symmetric uncertainty depending on the distance from each measured datum under an h-norm for some h. In that case the likelihood term may be assumed to take the form (Tarantola 2005)

$$\mathcal{P}(\boldsymbol{d}|\boldsymbol{m}) = c_1 \exp \sum_{i=1}^{q} \left\| \frac{d_i - f_i(\boldsymbol{m})}{\sigma_i} \right\|_h$$
(3)

where  $d_i$  is the  $i^{\text{th}}$  of q data in vector d, and  $c_1$  is a normalising constant which ensures that the right side of the equation defines a valid probability distribution (its integral over parameter space must be 1). It is common to assume a squared norm (h = 2) and Gaussian uncertainties with standard deviation  $\sigma_i$  on the  $i^{\text{th}}$  datum, whence the likelihood becomes

$$\mathcal{P}(\boldsymbol{d}|\boldsymbol{m}) = \left( (2\pi)^{n/2} \prod_{i=1}^{n} \sigma_i \right)^{-1} \exp\left(-\sum_{i=1}^{n} \left(\frac{d_i - f_i(\boldsymbol{m})}{\sigma_i}\right)^2\right)$$
(4)

The expression on the right can be rewritten  $\prod_{i=1}^{n} (2\pi\sigma_i^2)^{-\frac{1}{2}} \exp\left[-\left(\frac{d_i-f_i(\boldsymbol{m})}{\sigma_i}\right)^2\right]$  which is the product of likelihoods  $\mathcal{P}(d_i|\boldsymbol{m})$  of the sub-problem involving the single datum  $d_i$  given parameters  $\boldsymbol{m}$ . We therefore obtain  $\mathcal{P}(\mathbf{d}|\mathbf{m}) = \prod_{i=1}^{n} \mathcal{P}(d_i|\boldsymbol{m})$  and Bayes rule becomes

$$\mathcal{P}(\boldsymbol{m}|\boldsymbol{d}) = \frac{\mathcal{P}(\boldsymbol{m})}{\mathcal{P}(\boldsymbol{d})} \prod_{i=1}^{n} \mathcal{P}(d_i|\boldsymbol{m})$$
(5)

We may also be able to assume that the evidence can be written as the product of pdf's over the evidence in sub-problems involving individual data: e.g., for fixed model and data parameterisations the evidence is independent of parameter values m, and hence it is constant. It can then be divided between the various sub-problems,

$$\mathcal{P}(\boldsymbol{m}|\boldsymbol{d}) = \prod_{i=1}^{n} \frac{\mathcal{P}(\mathbf{m})}{\mathcal{P}(d_i)} \mathcal{P}(d_i|\boldsymbol{m}) = \prod_{i=1}^{n} \mathcal{P}(\boldsymbol{m}|d_i)$$
(6)

where the latter equality is obtained by Bayes rule. Thus we see that the solution to the full Bayesian problem is the product of the posterior pdf's of the set of n Bayesian sub-problems, each involving constraints from a single datum.

The expression for  $\mathcal{P}(\boldsymbol{m}|\boldsymbol{d})$  in equation (6) evaluates the pdf of any particular set of parameter values  $\boldsymbol{m}$ . Often we wish to estimate statistics of this pdf such as its mean  $\boldsymbol{\mu}$  or variance. This involves calculating integrals of the form  $\int_{\boldsymbol{m}\in M_P} m_i^r \mathcal{P}(\boldsymbol{m}|\boldsymbol{d}) d\boldsymbol{m}$  to obtain the  $r^{\text{th}}$  moment of the  $i^{\text{th}}$  parameter. In turn, this requires that we know  $\mathcal{P}(\boldsymbol{m}|\boldsymbol{d})$  over at least a representative set of samples  $\boldsymbol{m}_i$ , and from equation (4) we see that this involves evaluating the forward function  $\boldsymbol{f}(\boldsymbol{m}_i)$  for each  $\boldsymbol{m}_i$ . For example, using equation (5) the mean of parameter i is given by

$$\overline{m}_{i} = \frac{1}{P(d)} \left( (2\pi)^{n/2} \prod_{i=1}^{n} \sigma_{i} \right)^{-1} \int_{\boldsymbol{m} \in M_{P}} m_{i} P(\boldsymbol{m}) \exp\left(-\sum_{i=1}^{n} \left(\frac{d_{i}-f_{i}(\boldsymbol{m})}{\sigma_{i}}\right)^{2}\right) d\boldsymbol{m}$$
(7)

Even if we do not take a Bayesian approach to inversion, in order to estimate confidence intervals on parameters we must know  $f(m_i)$  for each of a representative set of samples  $m_i$ . It is therefore necessary to explore  $f(m_i)$  over parameter space P, and by equations (1) and (5) this exploration may be performed considering all data at once, or one datum at a time.

Almost all standard Bayesian sampling methods to solve nonlinear inverse problems consider all of the data at once and explore the solution using equation (1). This follows from an assumption that the more data constraints placed on parameter vector m, the more limited will be high-likelihood areas

of parameter space, hence the more tractable will become the sampling problem. Alternatively, one might choose to solve single-datum travel time tomography problems to obtain  $\mathcal{P}(d_i|\mathbf{m})$ : as shown below, extensions to single-datum problems may be easier to describe and evaluate, and may span large hyper sub-spaces of *P*. Single-datum solutions can then be combined using equation (5).

# Sample Extensions

Define the *i*th sample of parameter space to be  $\mathbf{m}_i = [m_{i,1}, ..., m_{i,p}] \in M_P$  where  $m_{i,j}$  is the *j*th of *p* parameters of the mathematical model. For each sample we assume that the forward function  $\mathbf{f}(\mathbf{m}_i)$  has been evaluated. The total state of information derived from the *i*th sample is then

$$\{\boldsymbol{m}_i, \boldsymbol{f}(\boldsymbol{m}_i; \boldsymbol{M}_P, \boldsymbol{D}_S), \boldsymbol{X}_i\}$$
(8)

Here,  $f(\mathbf{m}_i, M_P, D_S)$  is the forward function (noting explicitly that the model and data space affect the form of f), and  $X_i$  represents the *extension* of sample i. The extension  $X_i$  is defined to be all information about other forward function evaluations that is obtainable from  $\{\mathbf{m}_i, f(\mathbf{m}_i, M_P, D_S)\}$  without performing further forward function evaluations. Extension  $X_i$  is a set: it may contain information about the value of f for parameter values other than  $\mathbf{m}_i$ , or values of different forward functions  $\hat{f}(\hat{\mathbf{m}}_i, \hat{M}, \hat{D})$  which map samples  $\hat{\mathbf{m}}_i$  in space  $\hat{M}$  to space  $\hat{D}$ , where spaces  $\hat{M}$  and  $\hat{D}$  may differ from  $M_P$  and  $D_S$ , respectively. By analogy, we define the *second order sample extension*  $X_{i,j}$  of samples  $\mathbf{m}_i$  and  $\mathbf{m}_j$  to be all information about other forward function evaluations that is obtainable from the set  $\{\mathbf{m}_i, f_i(\mathbf{m}_i; M_{P,i}, D_{S,i})\}$ ;  $\mathbf{m}_j, f_j(\mathbf{m}_j; M_{P,j}, D_{S,j})\}$  without performing further forward functions) might differ between parameter values  $\mathbf{m}_i$  and  $\mathbf{m}_j$ , hence the additional subscripts i and j throughout. Higher order extensions are defined analogously. In this article we refer to the first order sample extension simply as *the extension* of  $\mathbf{m}_i$ , and to simplify notation the forward function  $f(\mathbf{m}_i, M_P, D_S)$  is denoted  $f(\mathbf{m}_i)$  except where we wish to focus on the parameter and data space dependencies.

Symmetries **T** always provide sample extensions. By definition,  $f(m_i) = f[T(m_i)]$ , so if  $f(m_i)$  has already been evaluated then  $f[T(m_i)]$  is a member of the extension of  $m_i$  that is available at the cost of evaluating  $T(m_i)$ . In linear problems, the particular set of symmetries for which the parameters best-fit recorded data are referred to as the null space and in theory can be found using matrix methods (Snieder & Trampert 2000). In nonlinear problems they must be found using sampling methods (Fichtner & Zunino 2018). Symmetries extend the concept of a nonlinear null space, as they provide information about parameters throughout  $M_P$ , irrespective of fit to a particular data set.

In addition to those derived from symmetries, a variety of other types of extension may also be available. Extensions generally fall into four main types: *function value, data space, parameter space* and *interpretational extensions* of  $f(m_i)$ , defined as follows.

**Function Value Extensions**: These provide the values of f(m) for parameter values other than  $m_i$ . Functional symmetries lead to function value extensions as explained above. Other types of information may also exist within these extensions, such as function values for numerical combinations of already-evaluated parameter values.

**Parameter Space Extensions**: These provide forward evaluations from spaces that differ from  $M_P$ . They typically arise by changing parameterisation P in ways that affect the function value predictably. This may involve adding, removing or otherwise changing parameters. Parameter space extensions may be particularly useful in problems involving many sets of parameter values with different parameterisations, including model selection problems (Box & Draper 1987; Linde 2014), reversiblejump Monte Carlo methods (*Green 1995; Malinverno 2002; Bodin & Sambridge 2009*) and interrogation problems (Arnold & Curtis 2018).

**Data Space Extensions**: These provide forward evaluations to spaces that differ from  $D_S$ . Data space extensions are already used in various contexts. They are particularly important for experimental design problems where we attempt to define an optimal suite of data types, locations and recording times (all defined by design S) in order to obtain particular target information (Curtis 2004; Maurer et al., 2010; Bloem et al., 2020). In these studies, the ability to test different designs without having to evaluate different functions f increases efficiency.

**Interpretational Extensions**: These provide alternative interpretations of the same parameter, data or function values. They are important in a variety of settings, and such extensions may differ substantially in their nature.

Extensions are particularly important when they can be evaluated at substantially lower cost than evaluating f. In such cases, extensions of previous samples might be used to guide more efficient future sampling strategies, or may answer questions of interest directly. The importance of each type of extension becomes clear when discussing tangible cases, so we now introduce a specific application that we follow through the rest of this article.

# Travel Time Tomography

Travel time tomography is an inverse problem that uses the measured travel times of energy or particles travelling between source-receiver pairs, to infer the speed or slowness (the reciprocal of speed) of propagation across the spatial domain through which the energy travelled. The parameterisation consists of any means to define the slowness map across the domain, so the parameter space is the set of all parameter values that define a valid slowness map. The data space is the set of possible travel times between each source-receiver pair considered in the problem. One set of travel times is recorded experimentally, which constitutes the measured data. Prior information usually exists about physical parameters, and must be combined with information from the data to find the posterior uncertainty in the parameter values. The posterior pdf that describes this uncertainty is the solution to the inverse problem.



Figure 1 (a) Map of slowness of wave propagation (colour) and ray (green) of first-arriving energy at receiver (triangle) from the source (star). (b) Corresponding map of time taken to travel from the source (star) to each grid cell (colours) and rays (black) to two of several receivers (triangles).

A specific example is introduced in Figure 1. Figure 1a depicts a gridded (cellular) slowness structure of a 2-dimensional isotropic medium. This is referred to as a slowness map, and represents sample  $m_i$  where each element of vector  $m_i$  is the slowness in one cell. A ray path between a source and receiver depicts the fastest route between the two points. For an isotropic source, this is the path taken by the first energy to arrive at the receiver, and we define its travel time to be  $d_i$ . The forward function in this case calculates  $d_i$  given map  $m_i$  and the source and receiver locations. It is now standard practice to calculate travel times by solving the eikonal equation across the grid using a fast marching method (e.g., Rawlinson & Sambridge 2005). This computation is the forward function  $d_i = f(m_i)$ . If desired, rays can then be calculated by tracing the negative gradient of travel time from each receiver back to the source.

#### Function Value Extensions

We obtain a first example of a function value extension by invoking the physics of refracting rays. For easy reference, we define this extension using the following Lemma:

Lemma 1: Given travel times d = f(m) through slowness map  $m \in M_P$ , the corresponding travel times  $f(\lambda m)$  for map  $\lambda m$  are given by  $\lambda d$  for any positive scalar  $\lambda$ .

*Proof*: Ray geometries are defined by angles of incidence and refraction at each cell boundary, which in turn are governed by Snell's law:  $s_1 \cos \theta_1 = s_2 \cos \theta_2$  where angles  $\theta_1$  and  $\theta_2$  and slownesses  $s_1$  and  $s_2$  are defined for one cell boundary in Figure 2a. If the slownesses are multiplied by positive scalar value  $\lambda$ , Snell's law remains identical:  $\lambda$  appears on both sides and cancels, so the ray path remains in its original position. The travel time of any packet of energy travelling along a ray is the length of path in each cell multiplied by the cell slowness, summed over all cells. Since the ray and hence these path lengths do not change, the travel time along the ray in slowness map  $\lambda m$  is the travel time in map m multiplied by the same scalar  $\lambda$  (Figure 2b). End Proof.



Figure 2 (a) Slownesses  $s_1$  and  $s_2$  in two cells leads to angles of incidence and refraction  $\theta_1$  and  $\theta_2$ , respectively, measured from the cell boundary. (b) As for (a) but with slownesses multiplied by  $\lambda$ .

Say we evaluate the travel time function  $d = f(\mathbf{m}_i)$  for any slowness map  $\mathbf{m}_i$  to obtain a sourcereceiver travel time  $d_i$ . By Lemma 1, we obtain the corresponding travel times at an infinite number of other slownesses  $\lambda \mathbf{m}_i$  essentially for free by using the relation  $\lambda d_i = f(\lambda \mathbf{m}_i)$ . All such models in the set { $\lambda \mathbf{m}_i$ :  $\forall \lambda \in \mathbb{R}_{>0}$ } are therefore in the extension of  $\mathbf{m}_i$ , where  $\mathbb{R}_{>0}$  is the set of strictly positive real numbers.

It is useful to visualise the extension of models in parameter space. Figure 3a plots a conceptual slowness map  $m_i$  in a graph of its first three slowness parameters (labelled  $s_1$  to  $s_3$ ), and we evaluate  $d_i = f(m_i)$ . The extension  $\{\lambda m_i: \forall \lambda \in \mathbb{R}_{>0}\}$  is depicted by the red line in that figure, showing that we immediately know the function value at an infinite number of other models. A similar line, generally at different trajectories, exists in all parameter sub-spaces spanned by slownesses in  $m_i$ .



Figure 3 Parameter space for 3 cells containing slownesses s<sub>1</sub>, s<sub>2</sub> and s<sub>3</sub>. (a) Sample (circle) and extension (line). (b) & (c) Similar to (a) but extensions are planar (red) if cell 1 (b) or cell 2 (c) does not lie on the ray. (d) Extension is a solid volume (red) if neither cells 1 nor 2 lie on the ray.

Function value extensions can also be derived using the following result:

Lemma 2: Let  $r_1$  be a vector that describes a shortest travel-time path between a fixed source and receiver, through a medium with slowness map m. Increasing the slowness in any part of m that does not lie along ray  $r_1$  does not alter the smallest travel time source-to-receiver path  $r_1$ , nor the travel time  $t_1$  of the first-arriving energy at the receiver.

*Proof*: Consider any perturbation dm to parameter values m such that  $dm_i \ge 0 \quad \forall i$ , and  $dm_j = 0$  for cells j intersected by ray  $r_1$ . By definition of dm the travel time along path  $r_1$  does not change in map m + dm. Consider any second path  $r_2 \ne r_1$  between the source and receiver. By definition of  $r_1$ , energy travelling along path  $r_2$  in map m has travel time  $t_2 \ge t_1$ . And by definition, dm cannot increase speed (decrease slowness) anywhere on path  $r_2$ . Thus path  $r_2$  has travel time  $t_2 \ge t_1$  in map m + dm. Thus path  $r_1$  remains a shortest travel time path, and the travel time remains  $t_1$ . End Proof.

Lemma 2 defines symmetry  $T_1(m) = m + dm$  with dm defined as above, since  $f(m) = f[T_1(m)]$ . This provides travel times for a family of other maps without further evaluation of f. The power of this symmetry can be observed by applying Lemma 2 to any of the models  $\lambda m_i$  in the extensions of  $m_i$ . Assume that cell k in the model does not lie on the particular source-receiver ray path for which  $f(m_i)$  was evaluated. Lemma 2 states that slowness  $s_k$  can be increased arbitrarily without changing travel time  $d_i$ . Thus we observe that if  $dm_k = [0, ..., 0, 1, 0, ..., 0]$  is the unit p-vector with kth element equal to 1, then the set of models  $\{m_i + \mu dm_k: \forall \mu \in \mathbb{R}_{>0}\}$  lies in the extension of  $m_i$ . Since there exists a similar extension to every model  $\lambda m_i$  along the extension defining the line in Figure 3a, all models in the set  $\{\lambda \boldsymbol{m}_i + \mu \boldsymbol{dm}_k: \forall \lambda, \mu \in \mathbb{R}_{>0}\}$  lie in the extension of  $\boldsymbol{m}_i$ . This set is the plane shown in Figure 3b for the case in which cell 1 does not lie on the source-receiver ray, and in Figure 3c if cell 2 does not lie on the ray. If neither cell 1 nor cell 2 lie on the ray then every model in the extension in the  $s_1$  direction (Figure 3b) can also be extended in the  $s_2$  direction, so the solid wedge of models in Figure 3d also lies in the extension of  $\boldsymbol{m}_i$ . Thus it is clear that if q cells do not lie on the ray, there is an extension of the line  $\{\lambda \boldsymbol{m}_i: \forall \lambda \in \mathbb{R}_{>0}\}$  into q dimensions of parameter space; this creates a hyperwedge in q + 1 dimensions of the p-dimensional parameter space  $M_p$  within which the forward function values are obtainable given a single function evaluation for parameters  $\boldsymbol{m}_i$ .



Figure 4 Same map as in Figure 1(a), but highlighting 226 cells (outlined in yellow) that are not intersected by the ray.

Figure 4 shows that for the slowness maps used in examples above, 226 cells do not lie on the sourcereceiver ray. The extension of the model shown in the figure is therefore a solid hyper-wedge in 227 dimensions. This example shows the power of sample extensions: a space of such dimensionality is impossible to search with Monte Carlo or grid search methods (having a sample density of only 2 samples per dimension in 227 dimensional space requires that we evaluate the forward function  $2^{227}$ = 2x10<sup>68</sup> times. This is currently infeasible on any computer for any forward function, either to evaluate the function or to store the results. Yet sample extensions make this information available almost for free anywhere inside the hyper-wedge, at the cost of a single forward model evaluation and tracing a ray. Clearly this extension type may be important when the goal is to explore the values of the function across parameter space  $M_P$ .

Notice that the extension defined using Lemma 2 is exactly half of the null space that would be calculated using linearised methods if the recorded travel time was equal to the calculated value for any sample. Such methods define the set of linear combinations of slownesses that can be added to a particular solution  $\mathbf{m}^*$  of an inverse problem without affecting the data values predicted by the linearised forward function  $\mathbf{f}(\mathbf{m} + d\mathbf{m}) = \mathbf{f}(\mathbf{m}^*) + d\mathbf{m}^T[\partial \mathbf{f}(\mathbf{m}^*)/\partial \mathbf{m}]$  where the *l*th row of the matrix in square brackets is  $[\partial f(\mathbf{m}^*)/\partial m_l]^T$  assuming column vectors. Standard practice is to evaluate singular values of the matrix of partial derivatives, identify parameter space singular vectors with associated singular values equal to zero, and define the linearised null space by adding any multiple of those vectors to the particular solution  $\mathbf{m}^*$ . In the above example, some of these vectors would span all off-ray cells, and so the null space would define part of the extension of sample  $\mathbf{m}^*$  under the approximate function  $\mathbf{f}(\mathbf{m})/\partial \mathbf{m}$  are close to  $\partial f(\mathbf{m}^*)/\partial \mathbf{m}$  for all  $\mathbf{m}$ . This is only guaranteed in the vicinity of  $\mathbf{m}^*$  for forward functions with smoothly varying first derivatives. Smoothness is not

guaranteed if we decrease the slowness of cells that are not intersected by the ray. The fastest ray may then jump to a different location, changing first derivatives discontinuously. This explains why Lemma 2 restricts consideration to linearised null-space vectors that give positive perturbations to slownesses away from the ray: this guarantees that the ray remains constant, and that the corresponding extension exists for the true, nonlinear forward function f(m).

A similar type of extension exists for some cases where the data consists of recordings of waveforms of arriving wave energy. In waveform tomography, the recording of the first-arriving wave (rather than its travel time) may be used as data to invert for the slowness map. In that case the forward function involves either solving the wave equation to estimate wavefields between each source and receiver from which the first-arriving waves are extracted (e.g., Zhou et al., 2018), or estimation of first arriving waveforms directly (e.g., Lomax 1994; Snieder & Lomax 1996).

It will usually be the case that slownesses across much of the map have no influence on the first arriving waveform, and for such cells we can invoke Lemma 2 to define similar extensions to those above. For a cell to have influence, energy must be able to travel from the source to that cell, and from that cell back to the receiver, before the end of the first arriving waveform. Say we evaluate two travel time calculations across the grid similar to that in Figure 1 (these are typically far cheaper than calculating waveforms): the first has a source at the true source location, the second has a source at the receiver location. These are sufficient to calculate the travel time along the fastest path from the source to any cell and then onward to the receiver. Any cell for which this total time is larger than that of the latest energy in the first arriving wave cannot affect that waveform. We may therefore increase the slowness in such cells without inducing a change in those waveforms, and all such perturbed slowness structures lie in the extension of the unperturbed model.

#### Data Space Extensions

Data space extensions are often available when the evaluation of f at existing or proposed spatiotemporal data-recording points and types defined by experimental design S, necessarily also produces synthetic data for other recording points and types. Such data do not lie in data space  $D_S$ , hence they extend sample information to other spaces. Function evaluations at these other data points and types are often obtained almost for free.

A typical case occurs when evaluating f requires a computer simulation of a physical process that evolves dynamically through space or time until it spans observation points in design S. The full simulation then also represents values of the underlying process at other space-time points, perhaps also for different data types, the set of which define an experimental design S'. These synthetic data would be members of a different data space  $D_{S'}$ , and hence are values of a forward function  $\hat{f}(\boldsymbol{m}_i, \boldsymbol{M}_P, \boldsymbol{D}_{S'})$  where  $\hat{f}: \boldsymbol{M}_P \to D_{S'}$ . Such information would be included in extension  $X_i$  because it is available without requiring further forward evaluations – it merely requires that we store the dynamic evolution at other space-time points as we evaluate f.

Figure 1 depicts a specific example. Figure 1a shows a slowness map  $m_i$ , and a ray path between a source and an existing receiver that depicts the fastest route between the two points. Say the travel time along that ray is  $d_i$ . Simulating  $d_i$  using the fast-marching method constitutes forward function  $d_i = f(m_i)$  and also provides travel times to every point on the grid (Figure 1b). So as a by-product of this calculation we obtain travel times to any other receiver location. Such travel times lie in a different data space, and hence in the data space extension. As another example, the derivatives of f in space can be calculated by applying finite-difference operators to the values of f in Figure 1b. If spatial derivatives of travel times are measured, as for example in stereo slope tomography (Tavakoli F. et al., 2017), then this represents a change in data type and hence experimental design. All such

information is included in the data-space extension of slowness  $m_i$ , and is available at low computational cost given the original calculation of  $f(m_i)$ .

# Parameter Space Extensions

Common examples of parameter space extension are referred to as grid refinement, grid coarsening and effective medium theories. In all of these, a newly gridded parameterisation P' produces a sample  $\hat{m}_i$  from a new parameter space  $M_{P'}$ . This sample has known function value  $\hat{f}(\hat{m}_i)$  where  $\hat{f}: M_{P'} \rightarrow D_S$ , which is obtained without any additional function evaluations, and hence is in the parameter space extension.

Say parameterisation P defines a gridded model where each grid-cell represents a region that has certain parameterised properties as in Figure 1, and evaluation of f requires a hypothesised physical process (here, travel time) to be simulated across this grid. Any grid cell may usually be replaced by a partition of smaller grid cells with properties that exactly match those of the original cell, without changing the value of f since the old and new grids represent exactly the same distribution of properties in the underling physical space or process. This refined parameterisation increases the dimensionality of  $M_P$ , as depicted in Figure 5 for the same travel time calculation as in Figure 1.



Figure 5 Same map as in Figure 1(a), but with one cell divided into a partition of four sub-cells (magnified on right).

Grid coarsening (sometimes called up-scaling) is the opposite process, and may be applied for example where neighbouring cells contain identical properties, or affect the physical process as though they were a single cell with a single 'effective' property: such cells can be replaced by a single, larger cell containing the effective property without changing the evolution of the underlying physical process. For example, if all of the red cells in Figure 5 had the same slowness, any adjacent subset of them could be replaced by a single cell with the same slowness without changing either the ray path or the travel time of the wave travelling between the source and receiver. This change in parameterisation reduces the dimensionality of  $M_P$ .



Figure 6 Top-left: same map as in Figure 1(a) with magnified portion shown top-right. Lower plots are similar but the magnified cells have had all of their slownesses replaced by their maximum value.

Above we showed that Lemma 2 defines a symmetry  $T_1(m_i)$  that provides travel times for a family of other models of the same dimensionalities. In fact it does so also for models of different dimensionalities, found for example by combining cells, as depicted in Figure 6. The slowness map  $m_i$ shown in the upper left panel is the same as that in Figure 1. Four cells shown in the upper-right inset have differing slownesses, and none of those cells lie on the source-receiver ray path. Lemma 1 states that if we increase the slownesses of any such cells, the source-receiver travel time remains unchanged. This is illustrated in the lower plots where three cells have increased slowness to match the maximum of the original four slownesses. Since the four cells now have equal slowness they can be combined into a single cell as shown lower-right, and thus we obtain the slowness distribution in the lower-left panel. This has lower dimensionality 252 rather than the original 255, and the travel time for this model is known without further evaluations of  $f(m_i)$ , and is equal to that in map  $m_i$ . Any sets of neighbouring cells that do not lie on the ray path can be combined in a similar manner to provide models of significantly reduced dimensionality and with known travel time, all of which lie in the extension of the original sample  $m_i$ .

The above are simple examples of effective medium theories. These essentially define symmetries across parameterisations of different spatial scales, and therefore also imply parameter space extensions. They are used to find properties of parameters at a coarser scale, which produce exactly the same data as media with (in general) different properties at finer scales. They exist not only for travel time data, but for data sets consisting of amplitudes, anisotropy, or even full seismic waveforms (Capdeville 2010; Capdeville & Cance 2015).

#### Interpretational Extensions

This type of extension involves no changes in parameter or data spaces, nor in the values of function f, but extends the *interpretation* of the model, data or function values. This may be important in a variety of settings, between which such extensions may differ substantially in both nature and form.

One philosophical example is that under certain conditions on parameter and data uncertainty distributions, it can be shown that a linearised inverse problem solution can be formulated that has exactly the same mathematical form under a Bayesian or non-Bayesian interpretation of the parameter-data system (e.g., Snieder & Trampert 2000, equations 44-48). Since numerical values of the data, and the mathematical operations are identical, the numerical values of parameters in the solution are identical. Solving the problem within either the Bayesian or the frequentist philosophy thus also provides the solution to the problem interpreted within the other philosophy, for free.

Interpretational extensions of a different nature are often available from an evaluation of a forward function at one scale in a physical domain, or when the function is evaluated in dimensionless units. Extensions to any other scaling in the physical domain may then be represented by identical values in parameter space. For example, the scalar wave equation  $\frac{d^2u}{dx^2} = \frac{1}{c^2} \frac{d^2u}{dt^2}$  remains identical under the coupled transformations  $\bar{x} = \lambda x$ ,  $\bar{c} = \lambda c$ :  $\forall \lambda \in \mathbb{R}_{>0}$ , where c(x) is the local wave phase speed at location x. So, say the forward function involves solving the acoustic wave equation to predict recorded waveforms for speed map parameterised at one spatial scale (or in dimensionless form). Then exactly the same waveforms would be obtained in a spatial domain that is  $\lambda$  times as large with a speed map that is  $\lambda$  times as fast – and we obtain that solution for free through a simple reinterpretation of the original solution. Thus, the solutions to infinitely many other problems may be interpreted from a single forward calculation, and all of these alternative interpretations lie in its extension.

#### Example: 3-parameter travel time tomography

We now define a particular travel time tomography problem that has only three parameters so that parameter space, samples and extensions can be represented fully in three coordinate axes. Equation (5) shows that the Bayesian solution can be found provided that the likelihood can be integrated through parameter space, and equation (4) shows that this typically requires that the forward function is known across this space. So our goal here is to define the forward function  $f(m_i)$  at a dense set of points  $m_i$  distributed throughout  $M_P$ .



Figure 7 (a) Slowness map comprising 3 cells of constant slowness (lighter grey-shades are slower). Source and receiver locations are marked and the fastest ray is shown in green. (b) In parameter space the red sample corresponds to map in (a); blue line and green surface are extensions described in the main text.

In this example, parameter vectors  $m_i \in M_P$  consist of three slowness values that define the full medium considered. Figure 7(a) shows a simple 3-cell model for the slowness structure around a source and receiver pair between which a first-arrival travel time is measured. The cells have laterally-constant wave slowness (grey shade) and fixed geometry, leaving only the 3 cell slowness parameters ( $m_k$  = slowness in cell k; k = 1,2,3) variable.

The function f(m) that predicts travel time data given a set of parameters m comprises components

$$f_j(\boldsymbol{m}) = \int_{\boldsymbol{x} \in R_j(\boldsymbol{m})} s(\boldsymbol{x}, \boldsymbol{m}_i) \, d\boldsymbol{x}$$
(9)

Here,  $s(x, m_i)$  is the slowness at location x in the medium which is specified by parameters  $m_i$ ,  $R_j$  is the ray path between the  $j^{\text{th}}$  source and receiver pair, and  $f_j$  is the travel time along that path. While this integral looks linear provided that s is linear in  $m_i$ , in fact f is always nonlinear since ray  $R_j(m_i)$  is the path along which energy takes least time to travel from source to receiver. This is also a function of the slowness distribution in the medium, hence ray  $R_j$  also depends on parameters  $m_i$ , and the problem is nonlinear.

The forward problem consists of calculating the travel time of the earliest arriving wave travelling between the source and receiver, in this case found by calculating the fastest ray path and integrating the slowness along its length. Figure 7(a) shows an example slowness structure and corresponding source-receiver ray path, for a particular parameter-space sample shown as a red spot in Figure 7(b). Lemma 1 showed that if the absolute magnitudes of the slownesses are scaled by any constant then the ray does not change, and the corresponding travel time is given by the same scaling. Hence all scaled structures lie in the extension of the sample, and these are represented by the blue line through the origin in Figure 7(b).

The ray does not pass through cell 3. Lemma 2 proves that making cell 3 slower than its current value can not change the shortest travel time; hence, all parameter sets constructed from those within the blue extension by making cell 3 slower, are also in the extension of this sample. These are represented by the green slice in Figure 7(b). Thus, it can be seen that any slowness map m that lies in a significant part of a 2-dimensional subspace now has known travel time. Given a measured datum, the likelihood in equation (5) is also directly calculable at every such map m, and this information is available given only a single forward function evaluation.

Figure 8(a) shows 15 random samples { $m_i$ : i = 1, ..., 15} generated from a Uniform prior distribution in the range (0,1] in each slowness. An average of 2 samples per dimension would require  $2^3 = 8$ samples, whereas 3 samples per dimension would require  $3^3 = 27$  samples, so 15 samples represents a density of between 2 and 3 samples per dimension. It can clearly be seen that much of parameter space has not been sampled: throughout the region of white space neither the forward function value (travel time) nor the likelihood is known, so we can not know whether or not we have sampled the high probability regions.

Samples in this figure have been colour coded to represent the ray that gives the fastest path from source to receiver. If the ray passes through cells 1 and 2 similarly to that in Figure 7 the sample is green, if the ray goes directly from source to receiver within cell 1 the sample is blue, whereas if the ray goes through cells 1 and 3 the sample is red. Similarly to Figure 7, samples whose rays go through cells 1 and 2 have both absolute scaling extensions (Lemma 1), and extensions in the direction of larger slowness in cell 3 (Lemma 2), and these are represented for each sample by green semi-planes in Figure 8(b). The symmetric situation is where rays pass through cells 1 and 3 and hence have extensions towards increasing slowness in cell 2, represented by red semi-planes in Figure 8(b). Rays

that only pass through cell 1 have extensions towards larger values of slowness in cells 2 or 3: this creates volumetric 3D pyramid-shaped extensions, which are represented in Figure 8(b) as sets of multi-coloured, closely spaced planes so that extensions of one sample that lies within the extension of another can also be seen.



*Figure 8 (a) 15 randomly selected samples plotted in parameter space. Green samples do not intersect cell 3; red samples do not intersect cell 2; blue samples pass only through cell 1. (b) Extensions corresponding to samples in (a).* 

A comparison of Figure 8(a) and (b) illustrates why extensions are important. They provide continuous, volumetric information about both forward function and likelihood values, potentially across the full dimensionality of parameter space (the pyramid structures above is three dimensional), or across lower-dimensional sub-spaces. Without using extensions, sampling provides no continuous information across parameter space, leaving open the possibility that high probability regions lie between samples and so remain undetected. Of course, regions of parameter space remain unsampled in Figure 8(b), but these are dramatically reduced in volume compared to that in Figure 8(a).

Figure 8(b) also shows that extensions render some of the random samples entirely redundant – specifically those that lie within the extension of another sample. These can be seen as pyramids within the largest pyramidal extension in that figure. The latter pyramid spans around 1/6<sup>th</sup> of parameter space (and below it is shown that this could have been up to around 1/4), so a large proportion of forward model evaluations could be considered to be wasted if selected randomly without reference to extensions. Indeed, the situation might be significantly worsened had we deployed typical Markov chain Monte Carlo (McMC) sampling schemes: the strength of those methods is to focus future samples around previously discovered areas of high probability. If such a method discovers a high probability region within the largest pyramid, it would tend to focus future samples there too, with the possibility that much of this computation would be redundant.

Figure 9(a) shows an alternative set of 13 non-random samples. These have been chosen in a systematic manner: six closely spaced pairs of red and green samples lie either side of a line (not plotted) rising diagonally in the  $m_1 = 1$  plane from the point (1,0,0), terminating in a single blue sample in the  $m_3 = 1$  plane. Their importance is apparent in Figure 8(b) which shows their extensions: parameter space is now relatively densely sampled, with known travel time values across regularly distributed portions of continuous subspaces in either 2 or 3 dimensions. There are no redundant samples (the origin and two coordinate axes appear to be repeated by several samples but are

precluded on physical grounds since zero slowness is impossible), and the distribution of known travel times and hence likelihoods is far more informative in Figure 8 than in Figure 7a or 7b, despite using fewer samples.



Figure 9 Similar to Figure 8 but for deterministically distributed samples.

The samples in Figure 9 have the following characteristics. The central point between each pair of red and green samples describes a slowness map which exhibits multi-pathing along two fastest rays (let us call this double-pathing). In fact, any map that lies exactly on the line passing between the sample pairs has two rays that give the same shortest travel time between source and receiver – one path going through cell 2, the other through cell 3. Moving the sample even infinitesimally in the direction of increasing slowness in cell 2 ensures that the path through cell 3 becomes uniquely fastest, and vice versa, thus producing each pair of samples shown. It can therefore be understood that the red-green pairs of samples are merely for illustration, since their extensions are both given from the single double-pathed point that lies in between: that point produces *both* the red and green extensions. Using double-pathed samples instead of sample pairs reduces the total number of samples required to produce this set of extensions to 7 – fewer than 2 samples per dimension, and fewer than half the number used in Figure 8.

The entire pyramid-shaped extension is generated by the single sample in blue. The sample has a straight ray between source and receiver within cell 1, and varying the absolute scaling of this sample produces the line between the sample and origin; the pyramid consists of all models that can be derived from those on that line by increasing the slowness in cells 2 or 3.

The blue sample is clearly the most important from the point of view of exploring parameter space since its extension spans almost a quarter of the total volume. It is in fact the single point of triple-pathing: rays could take either a direct path, a path through cell 2, or a path through cell 3 through the corresponding slowness map, and arrive at the receiver in the same shortest possible time. Graphically, if we were to place another pair of samples just to the left of the blue point they would have red and green 2D extensions similar to those of the other sample pairs, each 2D extension bounding the pyramid on one side. However, similarly to above, this implies that we only need the single, triple-pathed sample to define all three extensions: by increasing the slowness value of any two cells infinitesimally (such that the travel time through those cells increases by an infinitesimally small amount) we obtain the extension of the path that traverses the third.

It is possible to reduce sampling further while also increasing information density if we invoke more physics about rays, and allow a certain degree of approximation. *Fermat's principle* states that travel time varies to first order with changes in slowness, and to second order with changes in ray path. In other words, if we compare small changes in ray paths versus small changes in cell slowness along the ray, the latter have far greater effect on travel times. This is often used to justify linearising tomographic problems by fixing ray paths in their prior best estimate locations, and solving for slownesses along those rays that are consistent with data. By contrast, here we use Fermat's theorem only to infer that travel time varies slowly with changes in ray geometry, compared to changes in slownesses, leading to a different linearization as follows.

All sample pairs in Figure 9 describe maps with rays that pass through the same cells. However, rays in different sample pairs vary in their exact geometry. Provided ray paths do not change too much, Fermat's theorem states that it is reasonable to interpolate the travel time between the top and bottom point pair. This implies that to within a certain approximation, the travel time of *any* double-pathed slowness structure can be inferred from only two samples: the lower-left double-pathed sample (**A** in Figure 9), and the upper-right triple-pathed sample (**B** in Figure 9) since the latter is also double-pathed with similar rays. The approximate travel time of any intermediate sample is obtained by interpolation, and since Lemmas 1 and 2 apply to that sample, the approximate travel time throughout the entire continuous parameter space can then be estimated – in this case using only two samples. This is an example of a second-order (approximate) extension: information in the extension is obtained from pairs of samples rather than individual samples alone.

Whether 2 samples are used to obtain approximate information about the forward function continuously throughout parameter space, or 7 samples are used to obtain exact information throughout the coloured regions in Figure 9 (which can be interpolated to every point in that space), reasonably dense information about f(m) is now available throughout parameter space. Similar physical principles hold for any source and receiver geometry, so a similarly dense set of information could be obtained for each travel time datum providing information about f(m). Equation (4) can therefore be evaluated to produce an estimate of the posterior pdf at all points in parameter space, allowing integrals of moments in equation (8) to be calculated efficiently. Thus we solve the tomographic inverse problem.

# Discussion

In the absence of strong prior information, the curse of dimensionality indicates that to explore solutions to high-dimensional problems in a computationally tractable manner we must obtain far more information than simply probabilities of point samples. The additional information contributed by the extension of each sample depends on the physics of the problem at hand. In the examples above, the ratio of the dimensionality of extensions to the dimensionality of parameter space approaches 1 as the density of cells in slowness maps (hence the number of parameters) increases. This exponential increase in the volume of information provide]d by extensions almost matches the order of increase in volume that must be searched due to the curse of dimensionality. This indicates that physics-based sampling methods may be tractable in high-dimensional systems even with weak prior information, at least for some classes of problem.

The results in the simple tomographic problem analysed above suggest that the information within a set of extensions is likely to be maximised using multi-pathed samples, for which the shortest source-receiver travel time occurs along multiple rays simultaneously. This is also true in more general

tomographic problems: for an *n*-pathed sample we need only increase the slowness anywhere along exactly n - 1 of those rays in order to produce another sample for which (a) we know that the single remaining ray produces the first arrival, (b) we know its travel time, and hence (c) we know its extension. Systematically perturbing the slownesses on each subset of n - 1 rays produces all n single-pathed extensions, so all such extensions are included in the extension of the original n-pathed sample.

This implies that to use travel-time tomographic extensions efficiently we might develop a procedure to find multi-pathed slowness maps. For example, considering *m* samples at a time that have *m* different rays, one could attempt to generate a single sample that produces the same first arrival time along a subset or all of the rays. For example, the bifocal design algorithms of Winterfors & Curtis (2008; 2012) focus on travel times from two samples at a time, and it is possible that a similar approach could be taken for the task above. The forward function of the original *m* samples might then have to be evaluated, and these samples may have useful extensions in their own right. However, the single multi-pathed sample could contain additional information. For example, compare the combined extensions of any red and a green sample in Figure 7 with the single multipathed example lying between any sample pair in Figure 8: the latter produces both red and green extensions that are always larger than the former combination.



Figure 10 Cartoon of a slowness map comprising slow (white) and fast (blue) cells. Two rays are shown (yellow) between a source (star) and receiver (triangle).

A more systematic algorithm might generate multi-pathed slowness maps directly. Figure 10 shows a map that has two fast paths between the source and receiver. This map could be adjusted in a number of ways: given the travel time along a ray that follows the lower path, slownesses along the upper path can be directly assigned to match that travel time, thus generating a double-pathed model. Similarly a triple pathed model could be generated by suitably increasing the speed of cells that lie on the straight line between source and receiver. The rays can then also be shifted by decreasing the speed in blue cells and increasing the speed of their neighbours. All such maps will have high-dimensional extensions in parameter space. This suggests that an algorithm might be conceived that systematically samples sets of multi-pathed rays; maps that embody those rays could be generated automatically.

Given travel time data between multiple source-receiver pairs, generally different parameter values will produce multi-pathed slowness maps for each datum. This might suggest that one should develop sampling algorithms that focus on multiple data at a time. However, above it was shown that interpretational extensions often allow the forward evaluation for one source-receiver pair to be converted to a similar evaluation for another source-receiver pair, by a simple rescaling. If the system can also be rotated and translated in space without affecting the forward function, then the forward evaluation for any source-receiver pair also provides a related forward evaluation for all other such pairs. The entire system of extensions found for any one datum may therefore be applicable to all others, by rescaling, rotation and translation, so single-datum extensions may remain important even in multiple-data scenarios.

The approximate extension proposed in the three parameter tomographic example demonstrates the increased efficiency with which information can be obtained. Not only are the number of samples required reduced to two, but the information obtained from their first and second order extensions spans all of parameter space continuously. While for the region outside of the pyramidal extension this information is approximate, it is nevertheless useful for estimating approximate moments of the posterior pdf using equation (8), or for identifying regions of parameter space that are more likely to provide good data fits and hence high posterior probabilities. As a result this information could be used to guide sampling algorithms that invoke proposal distributions for each new sample selected, such as McMC algorithms (Mosegaard & Tarantola 1995). Proposal distributions are used to select samples that are expected to have a good chance of fitting the data *before* their forward function has been evaluated, and improving proposal distributions has been shown to dramatically improve the performance of McMC (Khoshkholgh et al. 2020). While exact extensions provide posterior information directly, approximate extensions could guide McMC towards better areas to sample.

McMC and other importance sampling algorithms are usually used for integration, for example to estimate the expected value of any function over a posterior pdf, including moments of that pdf as in equation (8). Sampling apparently redundant areas within previously discovered extensions may then have value: if samples are distributed according to the posterior pdf then integrals such as equation (8) can be evaluated by simply averaging values of that function over the set of samples (Mosegaard & Tarantola 1995). Such algorithms might be made significantly more efficient if instead of evaluating computationally intensive forward functions for every sample, they used known values from previously discovered extensions. Indeed, McMC integration could be deployed with no further function evaluations at all, by estimating likelihoods for each new sample by interpolating previously discovered likelihoods between the nearest neighbour extensions.

The redundancy of samples chosen in Figure 8, the deterministic (non-random) nature of highlyinformative samples in Figure 9, and the method of selecting samples using Figure 10 all suggest that efficient sampling of parameter space for travel time tomography may require a departure from existing random sampling methods, and indeed that it might be achieved deterministically. However achieved, the arguments presented herein show that physics-based sampling algorithms avoid limitations that might otherwise be imposed by the No Free Lunch theorem, can be particularly efficient in terms of number of samples evaluated, may add continuous and widespread information throughout parameter subspaces for each sample evaluated, avoid redundancy that may be rife in randomly-selected samples, and crucially may provide information that grows exponentially with the number of dimensions. The latter property provides hope that sampling algorithms may in future be tractable even in high-dimensional problems.

# Acknowledgements

The author wishes to thank Claire Allmark for producing travel time maps used in some of the figures.

# References

Arnold, R. & Curtis, A., 2018. Interrogation Theory, *Geophys. J. Int.*, **214**, 1830-1846. doi: 10.1093/gji/ggy248

van den Berg, J., Curtis, A. & Trampert, J., 2003. Optimal nonlinear Bayesian experimental design: an application to amplitude versus offset experiments, *Geophys. J. Int.*, **155**, 411-421

Bloem, H., Curtis, A. & Maurer, H., 2020. Experimental Design for Fully Nonlinear Source Location Problems: Which Method Should I Choose? *Geophysics*, **223**, 944-958, doi: 10.1093/gji/ggaa358

Bodin, T. & Sambridge, M., 2009. Seismic tomography with the reversible jump algorithm. *Geophys. J. Int.*, **178**, 1411-1436, doi:10.1111/j.1365-246X.2009.04226.x

Box, G.E. & Draper, N.R., 1987. *Empirical Model-Building and Response Surfaces*, Wiley series in probability and mathematical statistics. John Wiley & Sons

Capdeville, Y. & Cance, P., 2015. Residual homogenization for elastic wave propagation in complex media, *Geophys. J. Int.*, **200**, 986–999, https://doi.org/10.1093/gji/ggu452

Capdeville Y., Guillot, L. & Marigo, J.J., 2010. 1-D non periodic homogenization for the wave equation, *Geophys. J. Int.*, **181**, 897-910

Chuang Chen, Xu, X.Y., Liu, J., Batrouni, G., Scalettar, R. & Meng, Z.Y., 2018. Symmetry-enforced self-learning Monte Carlo method applied to the Holstein model, *Phys. Rev. B*, **98**, 041102(R)

Ji-Yao Chen, Capponi, S. & Poilblanc, D., 2018. Discrete lattice symmetry breaking in a twodimensional frustrated spin-1 Heisenberg model, *Phys. Rev. B*, **98**, 045106

Curtis, A., 2004. Theory of model-based geophysical survey and experimental design Part B; Nonlinear Problems, *The Leading Edge*, **23**(10), 1112-1117

Curtis, A. & Lomax, A., 2001. Prior information, sampling distributions and the curse of dimensionality, *Geophysics*, **66**, 372-378.

Das, S., Chen, X., Hobson, M.P., Phadke, S., van Beest, B., Goudswaard, J. & Hohl, D., 2018. Surrogate regression modelling for fast seismogram generation and detection of microseismic events in heterogeneous velocity models, *Geophys. J. Int.*, **215**(2), 1257–1290, https://doi.org/10.1093/gji/ggy283

Devilee, R., Curtis A. & Roy-Chowdhurry, K., 1999. An efficient, probabilistic, neural network approach to solving inverse problems: Inverting surface wave velocities for Eurasian crustal thickness, *J. Geophys. Res.*, **104**(B12), 28,841-28,856

Douma, H., Snieder, R., & Lomax, A., 1996. Ensemble inference in terms of Empirical Orthogonal Functions, *Geophys. J. Int.*, **127**, 363-378

Earp S., & Curtis, A., 2020. Probabilistic Neural-Network Based 2D Travel Time Tomography, *Neural Computing & Applications*, **32**, 17077–17095, doi:10.1007/s00521-020-04921-8

Earp, S., Curtis, A., Zhang, X. & Hansteen, F., 2020. Probabilistic Neural Network Tomography across Grane field (North Sea) from Surface Wave Dispersion Data, *Geophys. J. Int.*, **223**, 1741-1757, doi: 10.1093/gji/ggaa328

Felsenstein, J., 1981. Evolutionary trees from DNA sequences: A maximum likelihood approach, *J. Mol. Evol.*, **17**, 368–376, https://doi.org/10.1007/BF01734359

Fichtner, A. & Zunino, A., 2018. Hamiltonian nullspace shuttles, *Geophys. Res. Lett.*, **46**(2), 644-651, https://doi.org/10.1029/2018GL080931

Fichtner, A., Zunino, A. and Gebraad, L., 2019. Hamiltonian Monte Carlo solution of tomographic inverse problems, *Geophys. J. Int.*, **216**, 1344–1363, https://doi.org/10.1093/gji/ggy496

Galetti, E., Curtis, A., Meles G. & Baptie, B., 2015. Uncertainty loops in travel-time tomography from nonlinear wave physics, *Phys. Rev. Lett.*, **114**, 148501, doi:10.1103/PhysRevLett.114.148501

Gelfand, A.E. & Smith, A.F.M., 1990. Sampling based approaches to calculating marginal densities, *J. Am. statist. Ass.*, **85**, 398–409.

Goldberg, D.E., 1989. *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley, Reading, MA.

Grana, D., Passos de Figueiredo, L. & Azevedo, L., 2019. Uncertainty quantification in Bayesian inverse problems with model and data dimension reduction, *Geophysics*, **84**(6), 1ND-Z34, https://doi.org/10.1190/geo2019-0222.1

van Driel, M., Krischer, L., Stähler, S. C., Hosseini, K., and Nissen-Meyer, T., 2015. Instaseis: instant global seismograms based on a broadband waveform database], *Solid Earth*, **6**, 701-717, http://dx.doi.org/10.5194/se-6-701-2015

Green, P., 1995. Reversible jump MCMC computation and Bayesian model selection, *Biometrika*, **82**, 711–732.

Linde, N., 2014. Falsification and corroboration of conceptual hydrological models using geophysical data, *Wiley Interdiscip. Rev. Water*, **1**(2), 151-171, 10.1002/wat2.1011

Karmakar, M., Maiti, S., Singh, A., Ojha, M. & Maity, B.S., 2018. Mapping of rock types using a joint approach by combining the multivariate statistics, self-or]ganizing map and Bayesian neural networks: an example from IODP 323 site, *Mar. Geophys. Res.*, **39**, 407–419, DOI 10.1007/s11001-017-9327-2

Käufl, P., Valentine, A.P., de Wit, R.W.L. & Trampert, J., 2016. Solving probabilistic inverse problems rapidly with prior samples, *Geophys. J. Int.*, **205**, 1710-1728, doi:10.1093/gji/ggw108

Kennedy, J. & Eberhart, R.C., 2001. Swarm Intelligence. Morgan Kaufmann. ISBN 978-1-55860-595-4.

Khoshkholgh, S., Zunino, A. & Mosegaard, K., 2020. Informed Proposal Monte Carlo, *arXiv preprint* arXiv:2005.14398

Kirkpatrick, S.C., Gelatt, D. & Vecchi, M.P., 1983, Optimization by simulated annealing, *Science*, **220**, 671–680.

Krischer, L., Hutko, A. R., Driel, M. van, Stähler, S., Trabant, C. & Nissen-Meyer, T., 2017. On-demand custom broadband synthetic seismograms. *Seismol. Res. Lett.*, **88**(4), 1127–1140, https://doi.org/10.1785/0220160210

Linde, N., P. Renard, T. Mukerji & J. Caers, 2015. Geological realism in hydrogeological and geophysical inverse modeling: a review. Advances in Water Resources, 86, pp.86-101, doi: 10.1016/j.advwatres.2015.09.019

Liu Q. & Wang, D., 2016. Stein variational gradient descent: A general purpose Bayesian inference algorithm, in *Advances in neural information processing systems*, **29**, pp.2378-2386

Lomax, A., 1994. The wavelength-smoothng method for approximating broad-band wave propagation through complicated velocity structures, *Geophys. J. Int.*, **117**, 313-334

Lomax, A. & Curtis, A., 2001. Fast, probabilistic earthquake location in 3D models using oct-tree importance sampling, in abstracts of *European Geophysical Society*, Nice.

Lomax, A., Michelini, A. & Curtis, A., 2009. Earthquake location, Direct, Global-Search Methods, in *Encyclopedia of Complexity and System Science*. Springer.

Maiti, S. & Tiwari, R. K., 2014. A comparative study of artificial neural networks, Bayesian neural networks and adaptive neuro-fuzzy inference system in groundwater level prediction, *Environmental Earth Sciences*, **71**, 3147-3160

Maiti, S., Tiwari, R.K. & Kümpel, H.-J., 2007. Neural network modelling and classification of lithofacies using well log data: A case study from KTB borehole site, *Geophys. J. Int.*, **169**, 733-746, DOI:10.1111/j.1365-246X.2007.03342.x

Malinverno, A. & Briggs, V.A., 2004. Expanded uncertainty quantification in inverse problems: Hierarchical Bayes and empirical Bayes, *Geophysics*, **69**, 1005-1016, doi:10.1190/1.1778243.

Malinverno, A., 2002. Parsimonious Bayesian Markov chain Monte Carlo inversion in a nonlinear geophysical problem, *Geophys. J. Int.*, **151**, 675-688, doi:10.1046/j.1365-246X.2002.01847.x.

Maurer, H., Curtis, A. & Boerner, D., 2010. Recent advances in optimized geophysical survey design. *Geophysics*, 75(5), 75A177–75A194, doi: 10.1190/1.3484194

Meier, U., Curtis, A. & Trampert, J., 2007. A global crustal model constrained by non-linearised inversion of fundamental mode surface waves, *Geophys. Res. Lett.*, **34**, L16304, doi:10.1029/2007GL030989

Meier, U., Curtis, A. & Trampert, J., 2007. Global crustal thickness from neural network inversion of surface wave data, *Geophys. J. Int.*, **169**, 706-722

Mosegaard, K. & Tarantola, A., 1995. Monte Carlo sampling of solutions to inverse problems, *J. Geophys. Res.*, **100**(B7), 12,431–12,447.

Moseley, B., Nissen-Meyer, T. & Markham, A., 2020. Deep learning for fast simulation of seismic waves in complex media, *Solid Earth*, **11**, 1527–1549, https://doi.org/10.5194/se-11-1527-2020

Nawaz, A., & Curtis, A., 2017. Bayesian Inversion of Seismic Attributes for Geological Facies using a Hidden Markov Model, *Geophys. J. Int.*, **208**, 1184–1200, doi: 10.1093/gji/ggw411

Nawaz, A., & Curtis, A., 2018. Variational Bayesian Inversion of Seismically Derived Quasi-Localized Rock Properties for the Spatial Distribution of Geological Facies, *Geophys. J. Int.*, **214**, 845–875, doi: 10.1093/gji/ggy163

Nawaz, A., & Curtis, A., 2019. Rapid Discriminative Variational Bayesian Inversion of Geophysical Data for the Spatial Distribution of Geological Properties, *J. Geophys. Res. Solid Earth*, **124**(6), 5867-5887, doi: 10.1029/2018JB016652

Nissen-Meyer, T., van Driel, M., Stähler, S. C., Hosseini, K., Hempel, S., Auer, L., Colombi, A. & Fournier, A., 2014. AxiSEM: broadband 3-D seismic wavefields in axisymmetric media, *Solid Earth*, **5**, 425–445, https://doi.org/10.5194/se-5-425-2014.

*Propp,* J.G. & Wilson, D.B., 1996. Exact Sampling with Coupled Markov Chains and Applications to Statistical Mechanics, Random Structures and Algorithms, **9**, 223-252

Rawlinson, N. & Sambridge M., 2005. The fast marching method: An effective tool for tomographic imaging and tracking multiple phases in complex layered media, *Explor. Geophys.*, **36**, 341-350.

Sambridge, M., 1999. Geophysical inversion with a neighbourhood algorithm—I. Searching a parameter space, *Geophys. J. Int.*, **138**(2), 479–494, <u>https://doi.org/10.1046/j.1365-246X.1999.00876.x</u>

Sambridge, M.S., & Gallagher, K., 1993. Hypocentre location using genetic algorithms, *Bull. seism. Soc. Am.*, **83**(5), 1467-1491

Shahraeeni M. & Curtis, A., 2011. Fast probabilistic nonlinear petrophysical inversion, *Geophysics*, **76**(2), E45–E58, doi: 10.1190/1.3540628

Shahraeeni M., Curtis, A., & Chao, G., 2012. Fast probabilistic petrophysical mapping of reservoirs from 3D seismic data, *Geophysics*, **77**(3), O1–O19, doi:10.1190/GEO2011-0340.1

Snieder, R. & Lomax, A., 1996. Wavefield smoothing and the effect of rough velocity perturbations on arrival times and amplitudes, *Geophys. J. Int.*, **125**, 796-812.

Snieder, R. & Trampert, J., 2000. Linear and nonlinear inverse problems, in *Geomatic methods for the analysis of data in the earth sciences*, Eds. A. Dermanis, A. Grun and F. Sanso, Springer, Berlin, 93-164

Tarantola, A., 2005. *Inverse Problem Theory and Methods for Model Parameter Estimation*. SIAM, ISBN: 978-0-89871-572-9.

Tavakoli B.F., Operto, S., Ribodetti, A. & Virieux, J., 2017. Slope tomography based on eikonal solvers and the adjoint-state method, *Geophys. J. Int.*, **209**, 1629–1647, doi: 10.1093/gji/ggx111

Winterfors, E. & Curtis, A., 2008. Numerical detection and reduction of non-uniqueness in nonlinear inverse problems. *Inverse Problems*, **24**(2), 025016, doi:10.1088/0266-5611/24/2/025016 (14pp).

Winterfors, E. & Curtis, A., 2012. A Bifocal Measure of Expected Ambiguity in Bayesian Nonlinear Parameter Estimation, *Technometrics*, **54**(2), 179-190, doi: 10.1080/00401706.2012.676953

Wolpert, D.H. & Macready, W.G., 1997. No Free Lunch Theorems for Optimization, *IEEE Transactions* on Evolutionary Computation 1(67), 67-82

Zhang, X., Curtis, A., Galetti, E. & de Ridder, S., 2018. 3D Monte Carlo Surface Wave Tomography, *Geophys. J. Int.*, **215**(3), 1644-1658, https://doi.org/10.1093/gji/ggy362

Zhang X. & Curtis, A., 2020a. Seismic tomography using variational inference methods, *J. Geophys. Res.: Solid Earth*, **125**, e2019JB018589

Zhang X. & Curtis, A., 2020b. Variational Full Waveform Inversion, *Geophys. J. Int.*, **222**, 406-411, doi: 10.1093/gji/ggaa170

Zhou, W., Brossier, R., Operto, S., Virieux J. & Yang, P., 2018. Velocity model building by waveform inversion of early arrivals and reflections: A 2D case study with gas-cloud effects, *Geophysics*, **83**(2), 1942-2156.