This paper is non-peer reviewed EarthArXiv preprint submitted to GJI

Constraining Earth Properties using Deterministic Backus-Gilbert Inferences

Adrian M. Mag[1], Christophe Zaroli[2], Paula Koelemeijer[1]

[1] *Department of Earth Sciences, University of Oxford, Oxford, United Kingdom*

[2] *Institut Terre et Environnement de Strasbourg, Université de Strasbourg, EOST, CNRS, UMR 7063,*

5 rue Descartes, Strasbourg F-67084, France

Received date; in original form date

SUMMARY

Seismic tomography is routinely used to image the Earth's interior using seismic data. However, in practice, data limitations lead to discretised inversions or the use of regularisations, which complicates tomographic model interpretations. In contrast, Backus-Gilbert methods make it possible to infer properties of the true Earth, providing useful insights into the internal structure of our planet. Here, we show how in the absence of data uncertainties the Backus-Gilbert-based Subtractive Optimally Localized Averages (SOLA) method can be utilised as a deterministic linear inference (DLI) method to derive a new method: SOLA-DLI. SOLA-DLI enables us to interpret results through the target kernels, rather than the imperfect resolving kernels. This also allows us to build families of models, rather than just properties. These advantages are illustrated in this contribution using three case studies. In the first, we illustrate how properties such as various local averages and gradients can be obtained, including associated bounds on these properties and resolution information. Our second case study shows how trade-offs between physical parameters are naturally included in SOLA-DLI, which is particularly relevant to inferences of Earth structure using normal-mode data. Using our final

case study, we demonstrate that SOLA-DLI can be utilised to obtain the coefficients of basis function expansions, which leads to discretised models with specific advantages compared to classical least-squares solutions. This publication is accompanied by a SOLA-DLI software package that allows the interested reader to reproduce our results and to utilise the method for their own research.

Key words: Inverse theory, Seismology, Seismic tomography, Structure of the Earth, Surface waves and free oscillations.

1 INTRODUCTION

Seismic tomography relies on mathematical inversions (Rawlinson et al. 2010; Nolet 2008) to model Earth's interior from collected data. Resulting tomography models highlight persistent features, such as subducted plates, rising plumes, and large scale velocity anomalies, believed to mirror real Earth characteristics (Ritsema & Lekić 2020). Improving these models often involves the development of new models with different data or methods to enhance the resolution of certain features or to reduce uncertainties. However, seismic inversions encounter a major challenge: data scarcity. This leads to non-uniqueness in solutions (e.g. Backus & Gilbert 1970), which often is mitigated using regularisation. Yet, such prior information might inadvertently impose unrealistic constraints or introduce artefacts in the models (e.g. Nolet 2008; Zaroli et al. 2017). While incorporating such new information is not inherently wrong, it must be accurate and well-understood to avoid misinterpretations of the resulting seismic tomography models.

Due to the inherent non-uniqueness of seismic tomography, the chosen data often provide insufficient constraints for the entire model. They may however offer useful constraints for certain parts or properties of the model. The approach of seeking specific properties rather than inverting for the entire model falls under the domain of mathematical inference (Valentine & Sambridge 2023; Tsai 2023). If all properties are known, the complete model is also known. As this is typically not possible, the data limitations lead to compromises. Rather than inverting for a complete model, we can aim to directly constrain a set of properties. This means we utilise all information provided

by the data in an optimal way, which is the core idea behind the Backus-Gilbert SOLA (Subtractive Optimally Localized Averages) linear inference method (Pijpers & Thompson 1994; Zaroli 2016).

The Backus-Gilbert method was originally used to only compute local averages of the unknown model from the data (Backus & Gilbert 1967, 1968, 1970). The issue of non-uniqueness is addressed by averaging, removing the need for regularisation. In the 1970s, Backus extended this to compute more general "properties" of the unknown model by fitting linear combinations of the sensitivity kernels to some other predefined kernel (later termed "Target kernel" by some authors) (Backus 1970a,b), but this earlier work does not delve deeper into the kind of properties one could analyse. The ideas of Backus and Gilbert have been popularised and practically applied in different areas, including helioseismology (e.g. Pijpers & Thompson 1994, 1992, who also coined the name SOLA), deconvolution (Oldenburg 1981), geomagnetism (Hammer & Finlay 2019; Hammer et al. 2021), and (1D, 2D and 3D) seismic tomography (Nolet 1985; Masters & Gubbins 2003; Zaroli 2016, 2019; Lau & Romanowicz 2021; Latallerie et al. 2022; Amiri et al. 2023; Restelli et al. 2024). However, these studies have mostly concentrated on extracting local average information about the unknown model. A notable exception is the paper of Pijpers & Thompson (1994) who used Gaussian derivatives to extract information about gradients. Past studies have also overlooked the importance of placing a norm bound on the model space, as data alone cannot sufficiently constrain any property (Backus 1970a,b; Backus & Gilbert 1970). Recently, Al-Attar (2021) has generalised the original methods of Backus and Gilbert by developing a rigorous mathematical framework for Backus-Gilbert inferences and by highlighting the importance of the model norm bound for obtaining model properties.

More issues arise in practical applications of BG SOLA theory. First of all, results obtained using the classic method must be interpreted using the resolving kernels, which are almost always distinct from one location to another, and sometimes may not resemble the target kernels at all. Furthermore, with classic BG-SOLA we obtain properties of the true model rather then the true model itself (or an approximation of it). In situations where our data depend on multiple physical parameters (e.g. when using normal modes), the sensitivity to parameters not of interest has typically been accounted for by additional noise (Masters & Gubbins 2003; Restelli et al. 2024), but

this means prior information is assumed in the inference problem. Furthermore, given that Backus-Gilbert methods provide model properties, it is thought that the method cannot be used to develop models that cover the full spatial domain (Valentine & Sambridge 2023), making it difficult to compute the fit to data and to utilise the models in forward simulations. These issues, together with the developments on general deterministic linear inferences by Al-Attar (2021) motivate this work.

The paper is structured as follows. In Section 2, we present the theoretical framework, including an overview of the general theory of deterministic linear inferences (Section 2.1), an overview of the classic BG SOLA method (Section 2.2), followed by a description of the modified SOLA-DLI method that incorporates model norm bounds in a more general mathematical framework (Section 2.3). Furthermore, we show how SOLA-DLI can be expanded by considering other types of target kernels (Section 2.4) and how families of discretised models can be obtained through an appropriate choice of target kernels (Section 2.5). In Section 3, we illustrate several aspects of the SOLA-DLI method through three case studies. Section 3.1 serves to illustrate how different properties are retrieved in a fully synthetic example, Section 3.2 shows how a simple resolution and trade-off analysis can be performed using synthetic data with real sensitivity kernels, and Section 3.3 demonstrates how discretised models can be obtained with SOLA-DLI. Finally, in Section 4 we discuss the advantages and limitations of SOLA-DLI and summarise our findings.

2 THEORY

2.1 Deterministic Linear Inferences

Relationships between data and model parameters are often linear in seismology, or can be linearised using perturbation theory (Tromp et al. 2005), provided that some assumptions about the size of the perturbations apply. Therefore, let us assume that we have some error-free data d that are related to some model m via a linear relation G:

$$G(m) = d. \tag{1}$$

We will call such model-data relationships "deterministic data constraints", because we assume the data to be known exactly. The model is part of the model space $\mathcal{M}$, and the data are part of the data space $\mathcal{D}$. In inversions, we attempt to find the solution model from the data by inverting the forward relation (Equation 1). However, in most cases, the forward relation cannot be inverted due to insufficient or inadequate data. For continuous models, this leads to either no solutions, or infinitely many solutions (Backus & Gilbert 1967). In the absence of measurement errors, the only case when there are no solutions, is when the data are outside the range of the forward operator and thus incompatible with the physical law. Typically, in such a situation we use a different forward relation. Throughout this paper, we will therefore assume that the data are compatible with the forward relation. Thus, we have infinitely many solutions (see Fig. 1a), which we will denote by $\{m\}$. Inversions can be performed by adding constraints (regularisations) to the model space $\mathcal{M}$ until a single model $\tilde{m}$ is "selected". For example, one may choose the model that has the smallest average gradient, which corresponds in some way to the smoothest model. However, if the assumptions built into the chosen regularisation are not correct, the model will not be a good representation of reality.

We often seek specific properties of the true model $\bar{m}$ rather than the entire model itself. These properties, for example the average structure over some volume within the Earth or the depths of discontinuities, belong to a distinct space known as the property space $\mathcal{P}$, following the work of Al-Attar (2021). Therefore, we can define a new (inference) problem as:

Given that:

$$G(\bar{m}) = d \tag{2}$$

Find:

$$\mathcal{T}(\bar{m}) = \bar{p} \tag{3}$$

where $\mathcal{T}$ (the property mapping) is a linear relation that extracts a property of any model, and $\bar{p} \in \mathcal{P}$ represents the value extracted by $\mathcal{T}$ (property) when applied to the true model $\bar{m}$. It can be shown that in most practical situations the desired property $\bar{p}$ can be anything given a finite number of deterministic data constraints (Al-Attar 2021; Backus 1970a). In other words, given the

data constraints, $\bar{p}$ may take any value from the property space $\mathcal{P}$ (see Fig. 1a), leaving us unable to say anything about the property of the true model $\bar{m}$. Backus (1970a) and subsequently Al-Attar (2021) showed that this problem can be overcome by introducing a norm bound M in the model space:

$$\|m\|_{\mathcal{M}} \leq M \tag{4}$$

This constraint, applied on the model space, differs from the constraints implicitly applied during regularisations because it does not aim to single out a particular model. However, the model norm bound constraint does not lead to an invertible system, but rather limits the possible solutions to a bounded subset of $\mathcal{M}$. If the set of models that respect the norm bound is denoted U_M (Equation 4), then the set of solutions that respect the norm bound constraint and the data constraint is $U_M \cap \{m\}$, which is a bounded subset (Al-Attar 2021). Al-Attar (2021) further showed that such a constraint leads to the true property $\bar{p}$ being constrained in a bounded subset $\{p\} \subset \mathcal{P}$ as long as the norm of the true model is smaller than the chosen norm bound (see Fig. 1b for a visual representation of these concepts). Without any other prior information, all properties found in $\{p\}$ are equally likely to represent the true property $\bar{p}$. If we choose the appropriate property mapping $\mathcal{T}$ and prior model norm bound, we will be able to constrain the set of possible properties sufficiently for us to infer useful information about $\bar{p}$. Choosing the appropriate $\mathcal{T}$ can be seen as finding the observable that can be best constrained by the data. For a more detailed and mathematical description of these ideas, the reader is referred to Section 2 of Al-Attar (2021).

2.2 Classic Backus-Gilbert SOLA Theory

Equation 3 describes the general linear inference problem. In realistic applications, we will always have a finite number of data N and a finite number of physical parameters N'. In the special case of classic BG SOLA, the relationship between model and data is assumed to be:

$$d_i = G(m^j) = \sum_{j=1}^{N'} \int_{\Omega_j} K_i^j m^j d\Omega_j \tag{5}$$

where d_i are the (error-free) data, m^j are the various physical parameters that the data depend on (such as density, shear and compressional wave speeds, locations of discontinuities, etc.), and

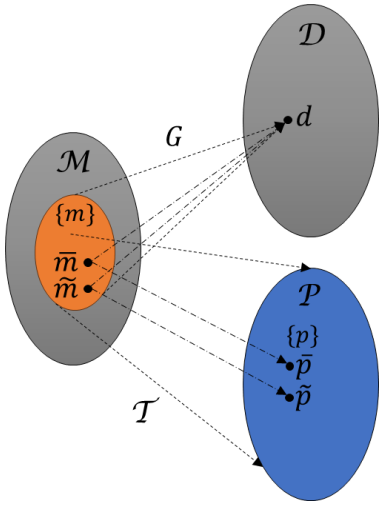

(a) No norm bound applied on the model space.

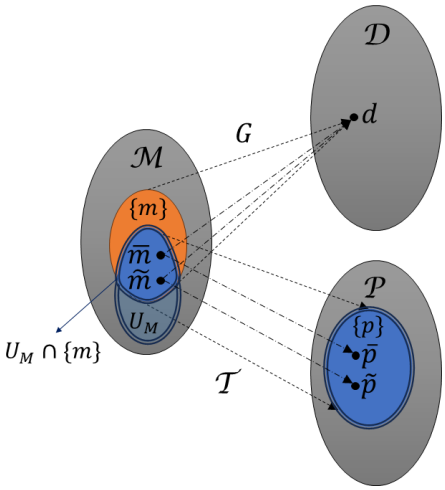

(b) With a norm bound applied on the model space.

Figure 1. Schematic of general linear inference problems, illustrating the effect of bounding the model space (see Table A1 for symbols definitions). Sets with thin lines for margins represent unbounded sets, while sets with double line margins represent bounded sets. The true model is denoted by $\bar{m}$ and the least norm model solution is denoted by $\tilde{m}$. When no bounds are imposed, the property of a model that respects the data constraint may take any value in the property space, which is an unbounded set (see Al-Attar 2021, Theorem 2.2). In other words, $\mathcal{T}(\{m\}) = \{p\} = \mathcal{P}$. Applying the norm bound on the model space leads to the intersection between $\{m\}$ and U_M to be bounded, which gets mapped under $\mathcal{T}$ to a bounded subset of $\mathcal{P}$. In other words, $\mathcal{T}(U_M \bigcap \{m\}) = \{p\} \subset \mathcal{P}$

.

K_i^j are all the sensitivity kernels associated with m^j (e.g. Tromp et al. 2005; Liu & Tromp 2008; Fichtner et al. 2006). In the continuous BG SOLA formalism, we consider the physical parameters m^j and the sensitivity kernels K_i^j to be functions of space. The sensitivity kernels and the corresponding physical parameters are usually defined on some regions Ω_j of the Earth, which may be 1D, 2D, or 3D and they live in function spaces $\mathcal{M}_j$. The data are always real, and the integrals are computed over the entire spatial domains over which the physical parameters are defined.

The Backus-Gilbert SOLA (BG SOLA) theory developed by Pijpers & Thompson (1994), Zaroli (2016, 2019), and others (Oldenburg 1981), has been used to find approximate local averages of some physical parameter. Let $\tilde{p}^{l,(k)}$ be an approximate local average centered around the spatial location $\mathbf{r}^{(k)}$ of the l^{th} physical parameter m^l. This approximation is obtained only from linear combinations of the data (Equation 5):

$$\tilde{p}^{l,(k)} := \sum_{i=1}^{N} x_i^{(k)} d_i = \sum_{j=1}^{N'} \int_{\Omega_j} \sum_{i=1}^{N} x_i^{(k)} K_i^j(\mathbf{r}) m^j(\mathbf{r}) d\Omega_j \tag{6}$$

where $x_i^{(k)}$ are some unknown scalars that optimise the sensitivity to m^l, $\mathbf{r}$ is the position vector, and $d\Omega_j$ is the volume/surface/line element of the domain on which m^j is defined. We further define $A^{j,(k)}$:

$$A^{j,(k)}(\mathbf{r}) := \sum_{i=1}^{N} x_i^{(k)} K_i^j(\mathbf{r}) \tag{7}$$

to be the "resolving kernel" (Zaroli 2019; Pijpers & Thompson 1994) of the physical parameter m^j at the spatial location $\mathbf{r}^{(k)}$.

The functions that we call resolving kernels have typically been referred to as *averaging kernels* in past studies, but we prefer to use the more general term *resolving kernel*. Our reasons for this will become more apparent in Section 2.4. If one finds the coefficients $x_i^{(k)}$ such that $A^{l,(k)}$ resembles a local averaging function $T^{(k)}$ for physical parameter m^l, then $\tilde{p}^{l,(k)}$ will represent the approximate local average of the true physical parameter $\bar{m}^l$ around the point $\mathbf{r}^{(k)}$. $T^{(k)}$ are called target kernels and represent the perfect resolving kernels we aim to obtain.

In practice, it is impossible to obtain $A^{l,(k)} = T^{(k)}$, but we may obtain very good approximations $A^{l,(k)} \approx T^{(k)}$ when we have sufficient data coverage. If we lack sufficient data, then $A^{l,(k)}$ may be significantly different from $T^{(k)}$, because the target kernels are typically linearly inde-

pendent of the sensitivity kernels. Furthermore, the resolving kernels $A^{j\neq l,(k)}$ for other physical parameters that affect the data will typically not be 0. Thus, $\tilde{p}^{l,(k)}$ will also contain contributions from physical parameters that are not of interest to us. This "contamination" is the mathematical manifestation of the trade-offs that exist between various physical parameters, such as volumetric heterogeneity and topography on internal boundaries inside the Earth.

In general, we therefore find

$$\text{Approximate property of true physical parameter } \bar{m}^l$$

$$\tilde{p}^{l,(k)} = \int_{\Omega_l} A^{l,(k)}(\mathbf{r})\bar{m}^l(\mathbf{r})d\Omega_l + \sum_{j\neq l}^{N'} \int_{\Omega_j} A^{j,(k)}(\mathbf{r})\bar{m}^j(\mathbf{r})d\Omega_j \tag{8}$$

when we really want to find

$$\text{True property of true physical parameter } \bar{m}^l$$

$$\bar{p}^{l,(k)} = \int_{\Omega_l} T^{(k)}(\mathbf{r})\bar{m}^l(\mathbf{r})d\Omega_l \tag{9}$$

and

$$\tilde{p}^{l,(k)} \neq \bar{p}^{l,(k)} \neq \bar{m}^l(\mathbf{r}^{(k)}). \tag{10}$$

In practical applications, for example inversions of normal mode splitting, the contaminant terms in Equation 2.2 have typically been accounted for either by scaling the contaminant physical parameters and adding them into the term containing the desired physical parameter (i.e. scaled sensitivity kernels Ritsema et al. 1999; Moulik & Ekström 2014) or by using the so called "3D noise" approach" (Restelli et al. 2024; Masters & Gubbins 2003; Lau & Romanowicz 2021). On the other hand, the complications arising from imperfect resolving kernels have been treated mostly by being cautious about the meaning of the values of $\tilde{p}^{l,(k)}$.

2.3 SOLA-DLI: Backus-Gilbert SOLA theory through linear inferences

Here, we present several modifications of the classic Backus-Gilbert (BG) SOLA method by applying the general theory of deterministic linear inferences (see Section 2.1 and Al-Attar (2021))

to the specific problem that SOLA is concerned with, leading to a new method that we call SOLA-DLI. By doing this, we will alleviate the issues of imperfect resolving kernels and contaminant terms at the cost of imposing a weak prior model constraint. For this, it is useful to first introduce some more general concepts.

An element of the model space $\mathcal{M}$ is a tuple containing all the physical parameters considered $m = (m^0, m^1, ...)$ (for example $m = (\rho, v_s, v_p)$). These parameters m^j are continuous variables, typically functions of space defined over some domains Ω_j. Each physical parameter lives in some space $\mathcal{M}_j$ and for our purposes it suffices to assume that the spaces $\mathcal{M}_j$ are Hilbert spaces. The total model space $\mathcal{M}$ is defined as the direct sum of all constituent model spaces $\mathcal{M}_j$, which is also a Hilbert space. Similarly, the sensitivity kernels are also continuous variables of the same type as the physical parameters they are associated with. We can thus generalise Equation 5 to:

$$d_i = G(m) = \sum_j \left\langle K_i^j, m^j \right\rangle_{\mathcal{M}_j} \tag{11}$$

where $\langle \cdot, \cdot \rangle_{\mathcal{M}_j}$ are inner products defined on each space $\mathcal{M}_j$. We have dropped the summation limits here for brevity, and in the remainder of the manuscript we will assume that the indices of our data and physical parameters run over the following sets $i \in \{1, 2, \ldots, N\}, j \in \{1, 2, \ldots, N'\}$, respectively. Most often the spaces $\mathcal{M}_j$ will be function spaces with inner products defined by:

$$\langle f, g \rangle := \int_\Omega f g \, d\Omega. \tag{12}$$

In some cases the physical parameter and its sensitivity kernels are represented by real numbers in which case the inner product becomes a real number multiplication (i.e. for sensitivities at internal boundaries within a 1D Earth model (Lau & Romanowicz 2021)). In this notation, a property of a model $m = (m^0, m^1, ...)$ is defined using target kernels as:

$$p^{(k)} = \mathcal{T}(m) := \sum_j \left\langle T^{j,(k)}, m^j \right\rangle_{\mathcal{M}_j} \tag{13}$$

where a target kernel $T^{j,(k)}$ must be defined for each physical parameter and $\mathcal{T}$ is the property mapping. Note that unlike in the previous subsection, we do not use the l index in $\tilde{p}^{l,(k)}$. We reserve that index for the cases when only one physical parameter (denoted m^l) has non-zero target kernels associated with it.

With our generalised notation, the inference problem (Equation 11) can be rewritten as:

Given

$$d_i = \sum_j \left\langle K_i^j, \bar{m}^j \right\rangle_{\mathcal{M}_j} \tag{14}$$

Find

$$\bar{p}^{(k)} = \sum_j \left\langle T^{j,(k)}, \bar{m}^j \right\rangle_{\mathcal{M}_j}. \tag{15}$$

Classic BG SOLA methods (assuming no data errors) would find the coefficients $x_i^{(k)}$ or simply X in matrix form (see Equation 6) such that the following cost function is minimised for all target kernels (Lau & Romanowicz 2021; Zaroli et al. 2017):

$$\Phi = \sum_j \left\| T^{j,(k)} - A^{j,(k)} \right\|^2_{\mathcal{M}_j}. \tag{16}$$

The true properties $\bar{p}^{(k)}$ are estimated by $\tilde{p}^{(k)}$, which are linear combinations of the data d_i weighted by the coefficients $x_i^{(k)}$ (Equation 6). The interpretation of these estimates is tied to their corresponding resolving kernels $A^{j,(k)}$. However, these resolving kernels often show ringing effects that make their interpretation as averaging weight functions difficult (Masters & Gubbins 2003). Furthermore, if multiple properties with identical target kernels but different spatial locations are computed, their resolving kernels will usually also be different. This leads to non-uniform interpretations of the properties.

These difficulties can be eliminated if we interpret the BG SOLA results through the target kernels, rather than the resolving kernels, as these remain unchanged. This makes us revisit the distinction between the true property $\bar{p}^{(k)}$ and the property estimate $\tilde{p}^{(k)}$. The true property has a clear meaning that is defined by the target kernels, but it cannot be found in practice. The estimate, on the other hand, can be found, but it is often difficult to interpret, since its meaning is defined by the resolving kernel. As discussed in Section 2.1, imposing a prior norm bound on the model space enables us to find an interval in which $\bar{p}^{(k)}$ is contained (Al-Attar 2021; Backus 1970a). By doing this, we effectively pay the price of imposing prior information and loose the convenience of having a single estimate for the true property, but we gain exact interpretability of our results.

In this manuscript, we use these concepts to modify BG SOLA giving rise to the new SOLA-DLI method.

Following Al-Attar (2021), we assume that the norm of each true physical property is bounded:

$$\left\|\bar{m}^j\right\|_{\mathcal{M}_j} \leq M^j \tag{17}$$

which leads to a norm bound for the true model $\bar{m}$:

$$\|\bar{m}\|_{\mathcal{M}} \leq \sum_j M^j, \tag{18}$$

where the norm on the direct sum $\mathcal{M} = \bigoplus_j \mathcal{M}_j$ is defined by $\|m\|_{\mathcal{M}} = \sum_j \|m^j\|_{\mathcal{M}_j}$ for any $m \in \mathcal{M}$. Examples of such model norm bounds will be given in Section 3.

Imposing a norm bound leads to the following solution estimate for the true property vector $\bar{p}$ (see Appendix A and Al-Attar (2021)):

Solution

$$\bar{p}^{(k)} \in \left[\tilde{p}^{(k)} - \epsilon^{(k)}, \tilde{p}^{(k)} + \epsilon^{(k)}\right] \tag{19}$$

Definitions

$$\tilde{p} = \mathcal{A}(\bar{m}) = Xd = \mathcal{T}(\tilde{m}) \tag{20}$$

$$\tilde{m} = G^*(GG^*)^{-1}d \tag{21}$$

$$\epsilon^{(k)} = \sqrt{\left(M^2 - \|\tilde{m}\|^2_{\mathcal{M}}\right)\mathcal{H}_{kk}} \tag{22}$$

where $\mathcal{A}$ is the "approximate mapping" that maps the true model through the resolving kernels to the property space (see equation A.14), $\mathcal{H}$ effectively quantifies the difference between $\mathcal{A}$ and $\mathcal{T}$ (see Figure A1), $\tilde{m}$ is the least norm solution and G^* is the adjoint of G. The solution estimate is a set of intervals that define the values that our property can take. For example, if the property that we search is a local average at 10 spatial locations, then the solution estimate will be one interval for each of the 10 local averages. Any average value within such an interval is equally likely to be the real average value. Furthermore, each interval is centered on the property of the least norm solution $\tilde{p}^{(k)}$. The solution intervals are dependent on the prior norm bound and the difference between the target kernels and the resolving kernels as quantified by the matrix $\mathcal{H}$. The

larger their differences, the larger the intervals of our solution. For an overview of the derivations of these equations, see appendix A.

The introduction of the prior norm bound allows us to shift the interpretation from the resolving kernels to the target kernels. In addition, it presents an alternative to the 3D noise method for estimating the contaminant terms (Equation 2.2) that does not require an estimate of errors in the physical parameters.

2.4 Choice of Target Kernels

In this section, we will discuss how different information about the unknown model can be extracted by choosing appropriate target kernels. For this, we assume the model to be a triplet of piece-wise continuous and bounded functions $m = (m^1, m^2, m^3)$ defined on the interval $[0, 1]$. This leads to a 1D inference problem, i.e. $\mathbf{r}^{(k)} = r^{(k)}$. However, the results can be easily generalised. The true model is assumed to be known and is plotted in Fig. 2. This model is arbitrary and has no physical significance.

2.4.1 Local Average Targets

Previous studies have primarily used the box car function as a target kernel for its simplicity and ease of interpretation - it gives a uniform local average (Restelli et al. 2024; Masters & Gubbins 2003). However, many other types of target kernels could be used to obtain local averages. Here, we use three different averaging target kernels:

Uniform Local Average (for reference):

$$T_U^{(k)}(r) := \begin{cases} C & r \in V_k \\ 0 & \text{else} \end{cases} \tag{23}$$

Gaussian Local Average:

$$T_G^{(k)}(r) := C \exp\left[-\frac{\left\| r - r^k \right\|_2^2}{2\sigma^2} \right] \quad r \in \Omega \tag{24}$$

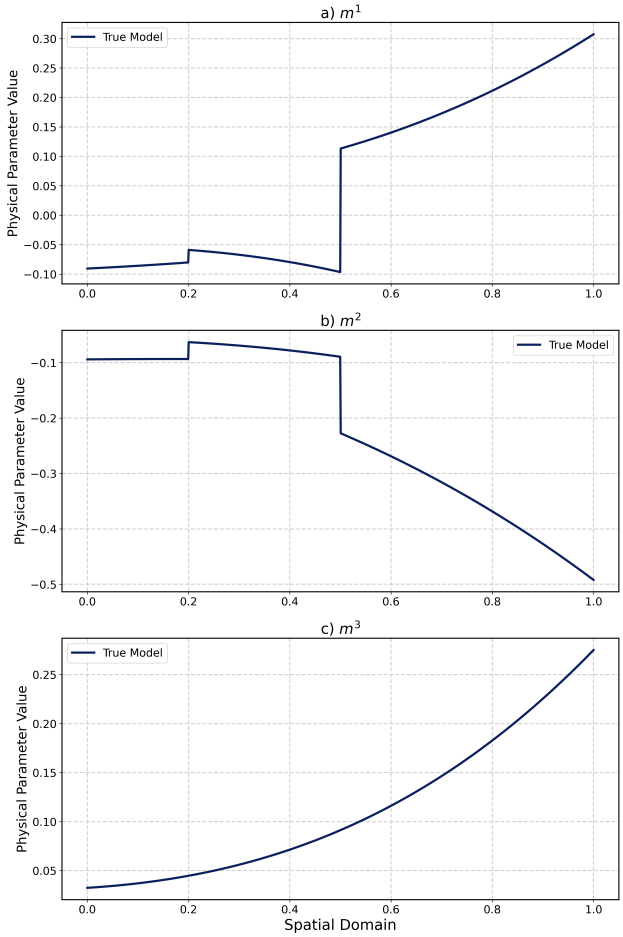

Figure 2. An arbitrary synthetic "true model". m^j denotes the physical parameters of the model.

Bump Function Average:

$$T_B^{(k)}(r) := \begin{cases} C \exp\left[\frac{w^2}{2(r-r^{(k)})^2 - w^2}\right] & r \in V_k \\ 0 & \text{else} \end{cases} \tag{25}$$

where

$$V_k = \left[r^{(k)} - \frac{w}{2}, r^{(k)} + \frac{w}{2}\right] \tag{26}$$

is the compact support of the boxcar and bump function with width w, σ is the standard deviation of the Gaussian, and C is in each case an appropriate normalisation constant that ensures the unimodularity of each target kernel. Examples of these averaging kernels are plotted in Fig. 3. Using these target kernels we can define, for example, the following property mappings for physical

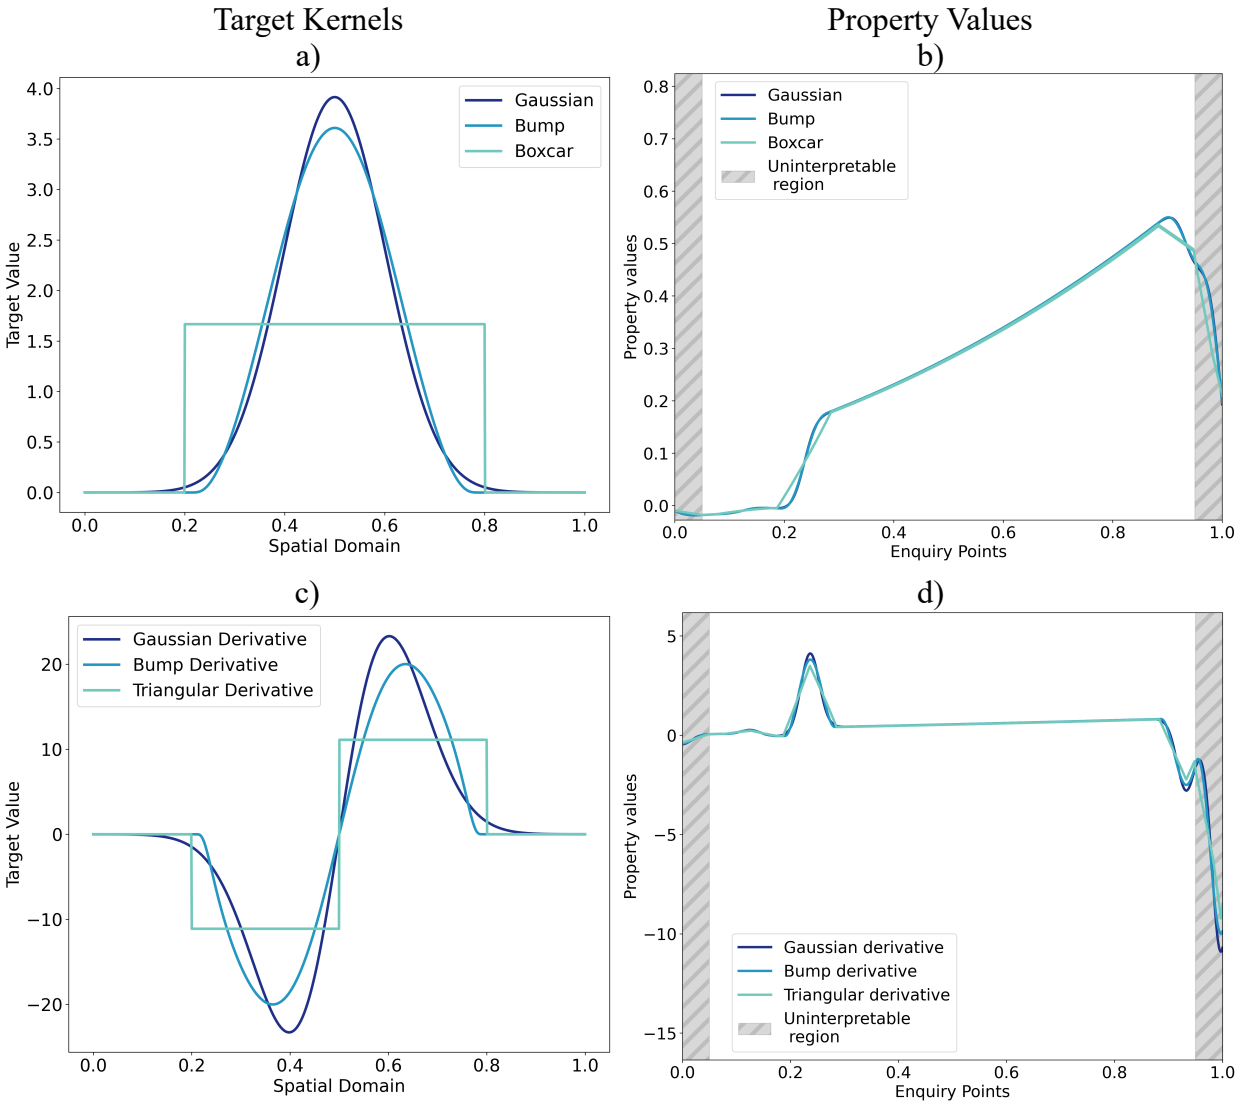

Figure 3. Examples of averaging (a) and gradient (c) target kernels and corresponding properties (b and d) obtained using these target kernels. First column: averaging and gradient target kernels with width 0.6. Second column: property values as a function of the 1000 enquiry points for different local averages and gradients of the true model using target kernels with width 0.1. The grey hatched regions represent parts of the domain where the target kernels are clipped (half-width of the target kernels).

parameter m^1:

$$\bar{p}^{1,(k)}_{U/G/B} = \mathcal{T}_{U/G/B}(m) = \int_0^1 T^{1,(k)}_{U/G/B}(r) m^1(r) dr. \tag{27}$$

The property vector extracted by each such property mapping (uniform/Gaussian/bump) is a vector of local averages centered at a set of points $\{r^{(k)}\}$ that we call "enquiry points". In Fig. 3 we also plot the property vectors $\bar{p}^{1,(k)}_{U/G/B}$ for 1000 evenly spaced enquiry points (right column). The grey hatched regions are parts of the domain where the target kernels $T^{1,(k)}_{U/G/B}$ are clipped and therefore

uninterpretable. The target kernels for physical parameters m^2, m^3 are 0 since we are not interested in these parameters.

Each of these target kernels has advantages and disadvantages. The boxcar function (Restelli et al. 2024) is simple and has compact support, providing a clear interpretation of the resolution of these kernels. However, most sensitivity kernels used in seismology are smooth, giving rise to poor resolving kernels that do not resemble boxcar functions. Consequently, the property error bounds are large.

Gaussian targets are often better reconstructed and lead to better constrained property values. They are, however, not defined on a compact domain and restricting a Gaussian to a compact domain leads to clipping. A clipped Gaussian is no longer a Gaussian, and different centering of the Gaussian leads to different clipping and therefore a "non-uniform" interpretation of the property values. If most of the Gaussian is located well within the bounds of the model domain, then the errors introduced by clipping can be negligible, but not readily quantifiable. For such target kernels we define a "width" that contains some large and arbitrary percentage of the function's weight (such as 90%) and pretend as if the entire weight of the function is concentrated in this region.

Bump functions are both smooth and defined on a compact support, therefore offering the advantages of both boxcars and Gaussian targets. The one shown here is just one example of a family of functions with similar characteristics.

2.4.2 Local Gradient Targets

If we want to obtain some local estimate of the gradient of a 1D physical parameter such as:

$$\bar{p}^{l,(k)} = \int_0^1 T^{l,(k)} \frac{\mathrm{d}m^l}{\mathrm{d}r} dr \tag{28}$$

we can use integration by parts to obtain:

$$\bar{p}^{l,(k)} = -\int_0^1 \frac{\mathrm{d}T^{l,(k)}}{\mathrm{d}r} m^l dr + \left[T^{l,(k)} m^l\right]_0^1. \tag{29}$$

For target kernels $T^{l,(k)}$ with compact support Equation 29 will reduce to:

$$\bar{p}^{l,(k)} = -\int_0^1 \frac{\mathrm{d}T^{l,(k)}}{\mathrm{d}r} m^l dr \tag{30}$$

in the interval where the target kernels are not clipped. For any other target kernel the second term will not necessarily be 0. However, it will be very close to zero if the target kernel is centered well within the bounds of the domain. Therefore, we can use Equation 30 to define new target kernels that extract local gradients of the true model.

In the bottom row of Fig 3 we plot the derivative of a Gaussian, Bump, and a Triangular function. The derivative of a Boxcar gives a sum of Dirac delta distributions. These cannot be used in our framework as they do not belong to a useful Hilbert space. Instead we have opted for the derivative of a triangular function, which yields a Haar function. The true property values obtained using the different gradient target kernels are shown in row 2, column 2 of the same figure.

The idea of using different target kernels to extract different types of information about the unknown model has been applied previously in helioseismology by Pijpers & Thompson (1994). They used Gaussian target kernels for extracting average information, and derivatives of the Gaussian to extract first and higher order derivatives of the model. However, as far as we are aware, this approach has not yet been used in seismic tomography. More importantly, the work of Pijpers & Thompson (1994) regard this approach as an inversion, whereas we believe it should be considered as an inference problem instead.

2.5 Obtaining Discretised Models through Target Kernels

One perceivable downside of linear inferences, such as SOLA-DLI, is the seeming impossibility of obtaining models that cover the full spatial domain (Valentine & Sambridge 2023). We illustrate here how SOLA-DLI can in fact be used to obtain discretised models by choosing appropriate target kernels.

Consider a model $m \in \mathcal{M}$ related to some data $d \in \mathcal{D}$ by:

$$d_i = G(m) = \langle K_i, m \rangle_{\mathcal{M}} \,. \tag{31}$$

For the reasons already mentioned in 2.1 an inversion of this equation is ill-posed leading to an

infinite set of solutions $\{m\}$. A common method for fixing this issue, besides regularisation, is discretisation. Typically, a set of orthonormal basis functions $\{B_l\} \in \mathcal{M}$ is chosen and any model in $\mathcal{M}$ is projected on the subspace formed by the span of this set, leading to a parallel $m^{\parallel}$ and perpendicular $m^{\perp}$ component of the model (i.e. $m^{\parallel}$ is the component that can be expressed with $\{B_l\}$ and $m^{\perp}$ is the residual term):

$$m = m^{\parallel} + m^{\perp} \tag{32}$$

$$m^{\parallel} = \sum_l p_l B_l \;\rightarrow\; \text{projection} \tag{33}$$

$$m^{\perp} = m - m^{\parallel} \tag{34}$$

where p_l are the coefficients given by the projection of m onto the basis functions:

$$p_l = \langle B_l, m \rangle_{\mathcal{M}} \tag{35}$$

We can then reformulate the initial inverse problem as:

$$\text{Find } \{p_l\} \text{ s.t.}$$

$$G\left(\sum_l p_l B_l\right) = d_i - G(\bar{m}^{\perp}) \tag{36}$$

The data correction term $G(\bar{m}^{\perp})$ subtracts from the original data the component corresponding to the part of the true model that is not within the span of the basis functions. In real applications, this term can never be computed since we do not know the true model $\bar{m}$, nor how much of it is outside the span of $\{B_l\}$. This term is therefore typically omitted and the equation solved in practice is just given by:

$$G\left(\sum_l p_l B_l\right) = d_i \tag{37}$$

which, combined with Equation 31, leads to the discretised inverse problem:

$$d_i = \sum_l \langle K_i, B_l \rangle \, p_l \tag{38}$$

In the seismic tomography literature, the matrix $\langle K_i, B_l \rangle$ is often denoted by G. However, we will not use that notation here since we already have a distinct (but related) use of the letter G.

When the number of coefficients p_l is chosen to be smaller than the number of data, such that

Equation 38 is overdetermined, it is often solved in a least square (or regularised least square) manner to produce the coefficients $\{\hat{p}_l\}$. These are systematically different from the true coefficients $\{\bar{p}_l\}$, because the correction term $G(\bar{m}^{\perp})$ is ignored. Including more data while keeping the same basis functions $\{B_l\}$ will not eliminate the systematic error caused by omitting the correction term. In order to converge to the true solution $\{\bar{p}_l\}$, one has to increase the number of data and the number of basis functions in the expansion. Increasing the number of basis functions shrinks the space in which $\bar{m}^{\perp}$ resides and thus decreases the size of the correction term $G(\bar{m}^{\perp})$.

SOLA-DLI can be used to tackle the same inverse problem while providing some notable advantages. The crucial step required to transform the inverse problem into a SOLA-DLI problem is to note that the coefficients $\{p_l\}$ can be viewed as values of a property p that resides in a property space $\mathcal{P}$. The property mapping $\mathcal{T}(m)$ is defined by choosing the target kernels to be the basis functions $\{B_l\}$:

$$\mathcal{T}(m) = \langle B_l, m \rangle_{\mathcal{M}} \tag{39}$$

Subsequently, we can introduce a prior model norm bound and apply the solution presented in Equation 19 to obtain a range of possible values for the coefficients $\{p_l\}$:

$$[\tilde{p}_l - \epsilon_l, \tilde{p}_l + \epsilon_l] \tag{40}$$

where $\tilde{p}_l$ are the coefficients of the least norm solution to Equation 31. The set of possible coefficients from Equation 40 can be inserted into Equation 33 to form sets of possible solutions for the projection of the true model on the chosen basis span.

Assuming that the prior norm bound is larger than the true model norm, the property bounds will contain the value of the true property, i.e. coefficients $\{\bar{p}_l\}$. By adding more data, the center of the property bounds (the least norm solution $\tilde{p}$) approaches the true property and the widths of the bounds decrease. SOLA-DLI thus provides a way to converge to the true property while maintaining the same number of basis functions. In fact, SOLA-DLI can even be used on a single coefficient - for example a particular spherical harmonic coefficient - and it will converge to the true value as more independent data are included. This contrasts with the classical inverse solution,

which, in the absence of the data correction term, requires an increase in both data and basis functions to converge to the true coefficients.

3 APPLICATIONS

We use three case studies to showcase the advantages and capabilities of the SOLA-DLI method. In Case 1 (Subsection 3.1), we show the effect of the prior model norm bound and target widths on the solution as well as how different types of local averages can be constrained. In Case 2 (Subsection 3.2), we illustrate how to perform a simple resolution and trade-off analysis with SOLA-DLI. In Case 3 (Subsection 3.3), we demonstrate how discretised model solutions can be obtained using SOLA-DLI, comparing the results with a least square inversion solution.

3.1 Case 1: General Multiparameter Model

In this completely synthetic case study, we show how SOLA-DLI is used to obtain estimates for three types of local averages, finding that some averages can be better constrained than others. We also illustrate how the prior information and the desired resolution change the local average estimates.

3.1.1 Setup

We consider a 1D model space containing three physical parameters m^1, m^2, m^3, all of which are piece-wise continuous functions defined on the interval $[0, 1]$. The synthetic true model (Fig. 4) is generated quasi-randomly and has no physical meaning.

The model-data relationship for d_i with $i \in \{1, 2, .., N\}$ is given by:

$$d_i = G(m^1, m^2, m^3) = \int_0^1 K_i^1(r)m^1(r)dr + \int_0^1 K_i^2(r)m^2(r)dr + \int_0^1 K_i^3(r)m^3(r)dr \tag{41}$$

For each physical parameter, the sensitivity kernels are also produced quasi-randomly using the equation:

$$K_i^j(r) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(r-\mu_{i,j})^2}{2\sigma^2}\right) \sin(\omega r) \sum_q c_q (r - r_q)^2 \tag{42}$$

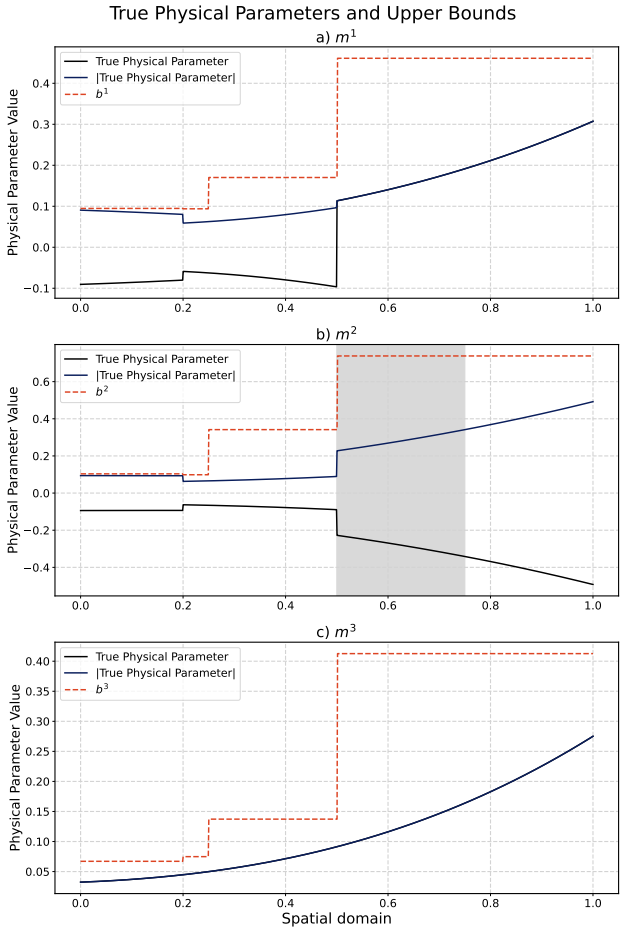

Figure 4. Case 1: True model and model norm bounds. Panels a–c show the synthetic quasi-randomly generated true model (comprised of three physical parameters) and some arbitrary piece-wise upper bound functions (b_i) used for computing the norm bound. In each panel we present both the physical parameter (black), and the absolute value of the physical parameter (blue).

where $\mu_{i,j}, c_q, r_q, \omega$ are randomly generated (see Fig. 5). We choose to use $N = 150$ (e.g. 150 observations) with the sensitivity kernels computed for each physical parameter. To simulate the lack of data sensitivity to a particular region (e.g. no S-wave sensitivity in the Earth's outer core), we manually set the sensitivity kernels for m^2 to zero in the interval $[0.5, 0.75]$. The synthetic (error-free) data are then produced using Equation 41 combined with the synthetic sensitivity kernels and the synthetic true model. As target kernels we use those defined in Equations 23, 24 and 25 with a width of 0.2.

The least norm solution to this problem (Equation 41) is given by the Moore-Penrose right-

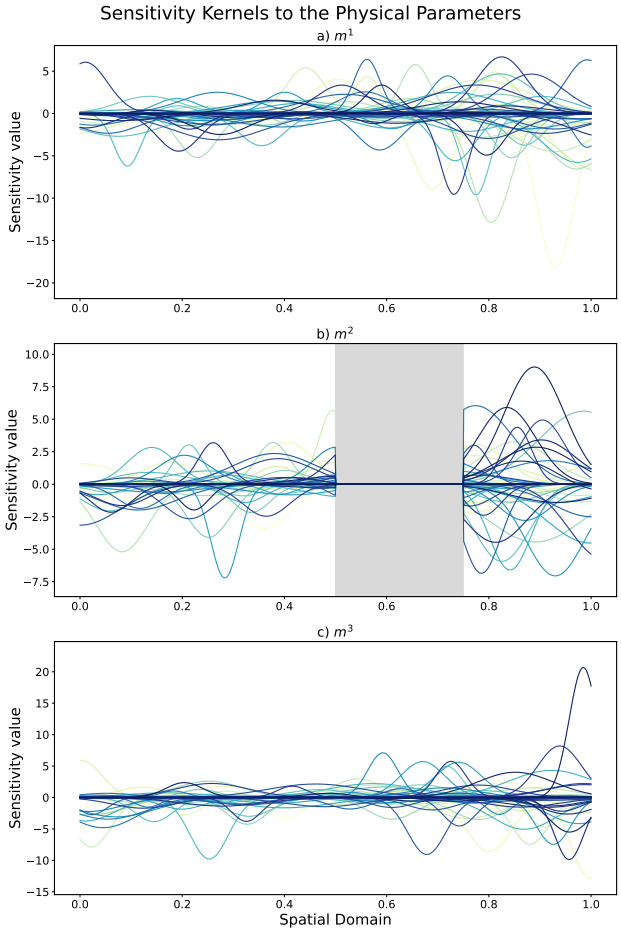

Figure 5. Case 1: Sensitivity kernels. Panels a–c show the synthetic quasi-randomly generated sensitivity kernels for physical parameters m^1, m^2, m^3. The region with no sensitivity to m^2 is shaded in gray.

inverse:

$$\tilde{m} = G^*(GG^*)^{-1}d \tag{43}$$

and shown in Fig. 6. This is a regularized inverse solution obtained by selecting the solution with the least norm from the set of all possible solutions $\{m\}$. The lack of sensitivity in the grey region is immediately apparent in this solution, as the least norm results in a zero model value here.

To solve the SOLA-DLI problem, upper bound functions b^i are chosen arbitrarily (Fig. 4) such that:

$$\left|m^j(r)\right| \leq b^j(r) \quad \forall r \in [0,1], \tag{44}$$

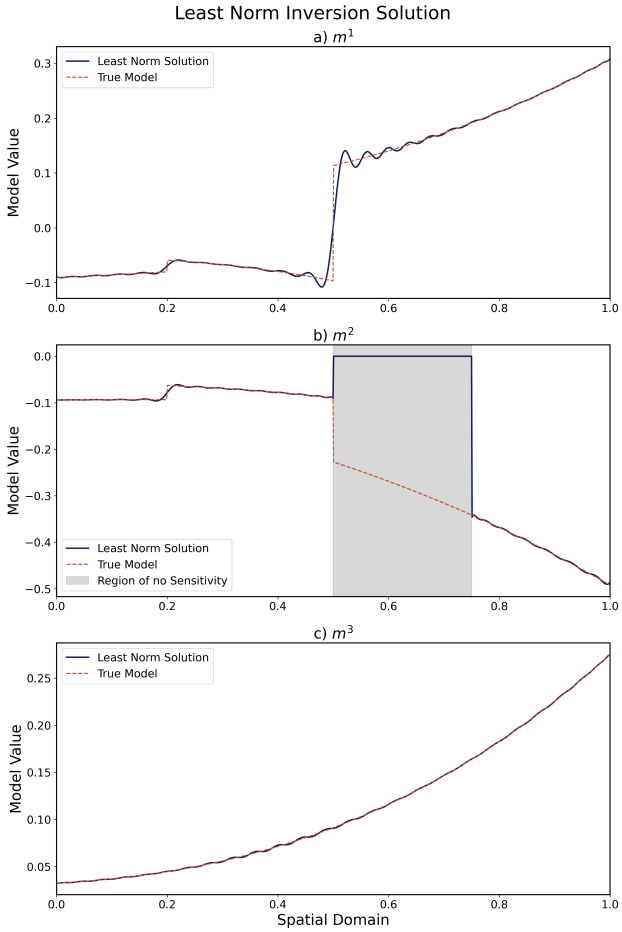

Figure 6. Case 1: Least norm solution for (m^1, m^2, m^3) obtained using Equation 41.

which leads to the following upper bound on the model norm:

$$\left\|m^j\right\|_{\mathcal{M}^j} = \sqrt{\int_0^1 (m^j)^2 dr} \leq \sqrt{\int_0^1 (b^j)^2 dr} = M^j \tag{45}$$

$$\|m\|_{\mathcal{M}} \leq M = M^1 + M^2 + M^3 \tag{46}$$

In real applications, the upper bound functions b^j should be chosen carefully based on physical arguments, for example using constraints from mineral physics.

3.1.2 Local averages and resolving kernels

In this case study, we consider three types of local averages: uniform local averages, Gaussian averages, and bump averages. Specifically, we are interested in obtaining these local averages for parameter m^2 at 100 equally spaced inquiry points in the spatial domain, with the results plotted in Fig. 7.

For each type of average, at each of the 100 enquiry points, the solution (Equation 19) gives an upper and a lower bound. Fig. 7 shows that the uniform local average is the least constrained property, while the Gaussian average is the best constrained property. This is not surprising considering that the sensitivity kernels are Gaussians modulated by polynomial and sinusoidal functions. If the sensitivity kernels were more similar to boxcar functions, then we may expect the uniform local averages to be better constrained. The region with no sensitivity is poorly constrained since the only constraint comes from the model norm bound.

In Fig. 7 we also plot, for each type of property, one set of target and resolving kernels. For Gaussian averages we observe that the resolving kernels are visually identical to the target kernels, which is the reason for the tight property bounds. In contrast, we retrieve visibly imperfect resolving kernels for the uniform averages, leading to large property bounds. Such large bounds ultimately make it impossible to infer any useful information about the true uniform local averages.

We have to be careful in comparing our SOLA-DLI results of Fig. 7 to the result of a regularised inversion (e.g. the least norm solution shown in Fig. 6). The regularised solution is a member of the model space $\mathcal{M}$, which in this case is a tuple of piece-wise continuous and bounded functions. In contrast, the solutions obtained using SOLA-DLI are members of the property spaces $\mathcal{P}$ (members of $\mathbb{R}^{N_p}$ in this case, where N_p is the number of enquiry points). At each enquiry point in Fig. 7 we have a range of possible values for the desired local average of the true model solution, while at each spatial point in Fig. 6 we have the tuple of values of the least norm solution. The regularised solution offers us a model spanning the full spatial domain, but it assumes that the true model has a minimal norm. The SOLA-DLI method offers bounds for several local averages of the true model solution and assumes only that the true model has a norm smaller than our norm bound.

3.1.3 Effect of the prior model norm bound

When we change the norm bound prior information, only the property error bounds are affected (see Equation 19). This is illustrated in Fig. 8, where we show results for three different upper bounds on the true model. Bound 3 is the most conservative, assuming a constant function three

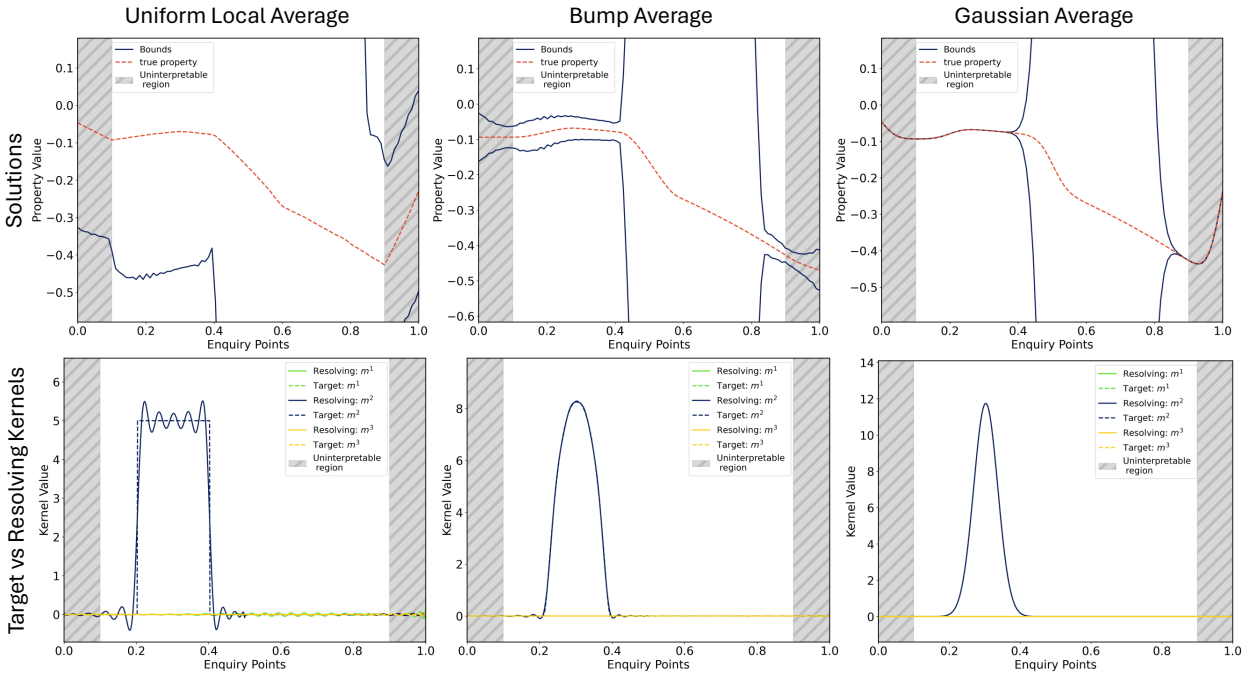

Figure 7. Case 1: SOLA-DLI solutions for three different types of local average properties. Top row: solution bounds for three types of local averages of the physical parameter m^2 at 100 evenly spaced enquiry points. Bottom row: target and resolving kernels for each type of property at the enquiry point located at $r^k = 0.3$ with width 0.2.

times larger than the maximum of the true physical parameter. Bounds 1 and 2 are tighter and therefore assume more prior knowledge. The bottom panel of the same figure illustrates that, as expected, tighter norm bounds lead to tighter property bounds. In all cases, the range remains centered on the true property. It is interesting to note that restricting the bounding function b^j in some local region does not lead to a tighter property bounds at an enquiry point in the same region, but rather it will lead to a uniform decrease of the property bounds at all enquiry points.

3.1.4 Effect of target kernel width

Changing the width of the target kernels can be interpreted as changing the resolution of the property evaluated at a given enquiry point. To investigate this, we have varied the target width between 1% and 100% of the domain width and computed the relative error bounds for all the enquiry points and widths. The results are plotted in Fig. 9. The relative error bound shown in the

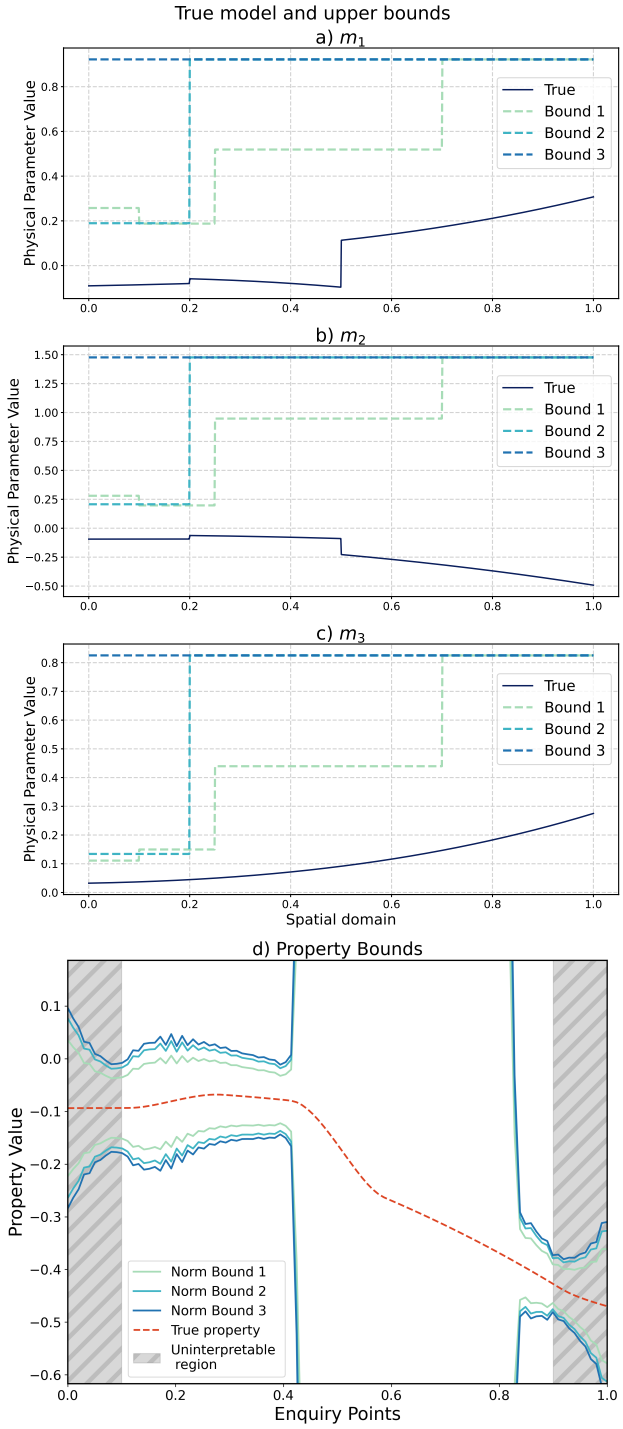

Figure 8. Case 1: Effect of prior norm bound on the property bounds. a–c) indicate the levels of three different upper bounds on the true model. Our choice of norm bound functions results in the following prior norm bounds (M^i): 2.44 (Bound 1), 3.89 (Bound 2) and 4.34 (Bound 3), which are all larger than the true model norm of 0.32. d) Solutions corresponding to the three different model upper bounds. Tighter norm bounds lead to tighter constraints on the desired properties.

first column is defined as:

$$e^{(k)} = \frac{\epsilon^{(k)}}{\max(\tilde{p}) - \min(\tilde{p})} \tag{47}$$

where $\tilde{p}$ is the property of the least norm model solution $\tilde{m}$ (see Equation 20). This metric has been chosen as the absolute error ϵ is not a good metric for determining whether a property is well constrained, while the classic relative error defined as $\epsilon/\bar{p}$ cannot be computed without knowing the true property $\bar{p}$. While there is no quantitative rule for what constitutes an unacceptable high relative property error bound, we believe any relative error higher than 100% is "certainly too high", and relative errors less than 10% are "generally good".

In general, we find that for all target kernel types the relative error bounds increase when we decrease the width of the target kernel (i.e. increase resolution). In addition, regions with no sensitivity always lead to large relative errors. As expected, the width of the uninterpretable regions at the edge of the domain increases with larger target kernel width (decreasing resolution) as the half-width of the kernel increases. Finally, in this setup, we find that particular properties, e.g. Gaussian averages, are constrained better (i.e. lower relative error bounds) than uniform or bump local averages, for all enquiry points and for target kernel widths (i.e. all resolutions), likely due to the Gaussian-like sensitivity kernels used.

This case study illustrates the general notion that we typically use inference methods to answer specific questions about a true model rather than finding the entire model itself. In SOLA-DLI, these questions are encoded in our chosen target kernels. The differing extent to which we are able to retrieve the target kernels effectively shows that our data can answer some questions better than others.

3.2 Case 2: Quasi Synthetic Normal Mode Application

In this quasi-synthetic case study, we show how to conduct a simple resolution analysis without real data or model values nor any prior information, based solely on the sensitivity kernels of the data set. We also illustrate how the results of such a resolution analysis can be linked to trade-offs between physical parameters.

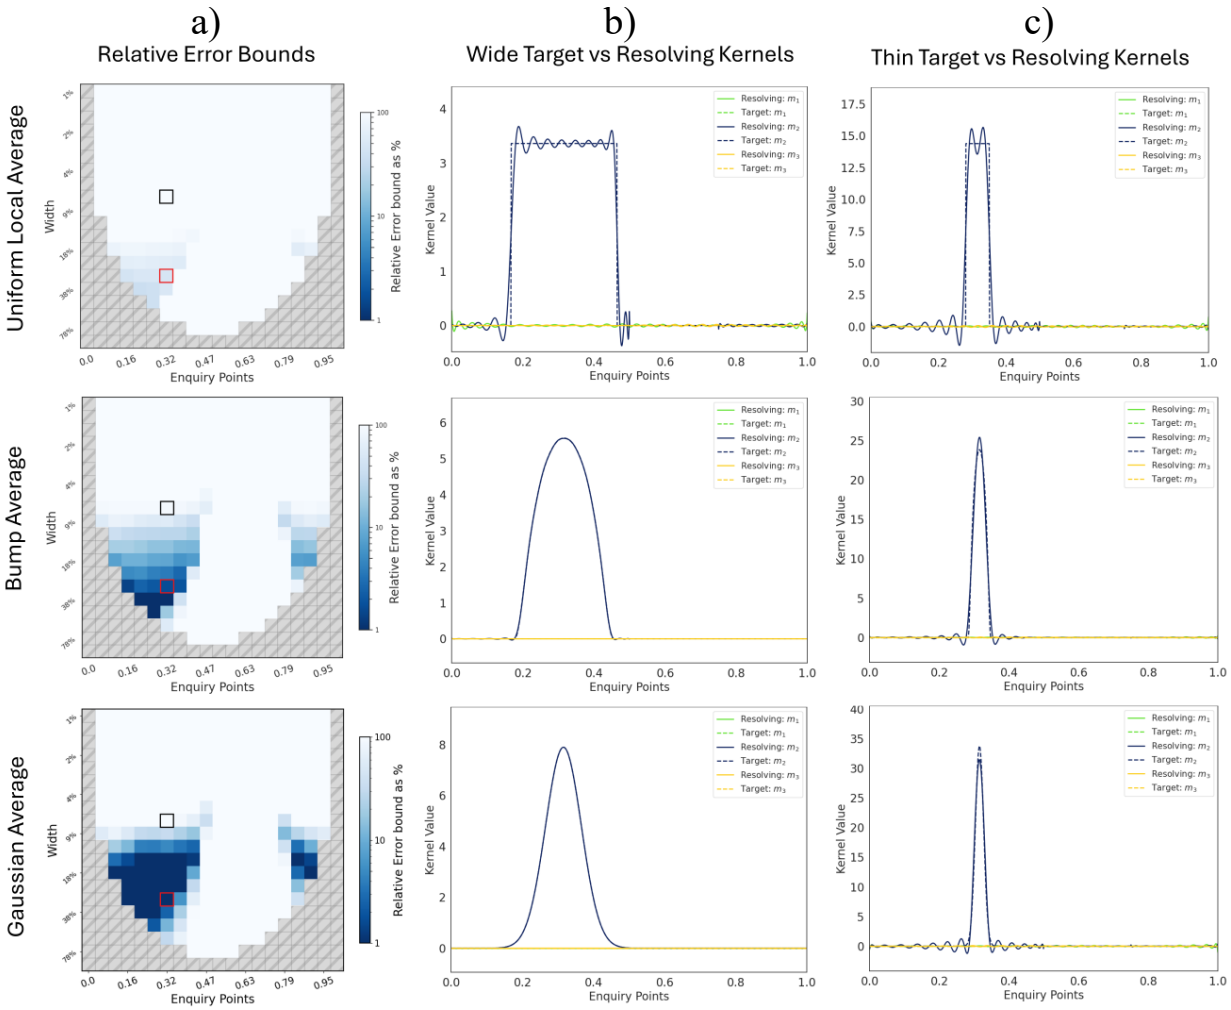

Figure 9. Case 1: Relative error bounds with examples of resolving kernels compared to their target kernels. The different rows correspond to different type of target kernels, e.g. uniform local average (top), bump average (middle) and Gaussian average (bottom). In the three columns we show the relative error bounds e (a); examples for wide target and resolving kernels, corresponding to the red squares (b) and examples for narrow target and resolving kernels (c), corresponding to the black squares.

3.2.1 Setup

Here, we consider a model formed by the triplet $m = (\delta \ln(v_s), \delta \ln(v_p), \delta \ln(\rho))$, where v_s is shear wave speed, v_p is compressional wave speed, and ρ is density (Fig. 10). Each physical parameter is assumed to be a piece-wise continuous function defined over the interval $[0, R_E]$ where R_E is Earth's radius (approximately 6371 km). We aim to constrain Gaussian averages and gradients of this synthetic true model using realistic normal mode sensitivity kernels (Woodhouse & Dahlen 1978). Specifically, we select the same modes as in the SP12RTS dataset (Koelemeijer et al. 2016;

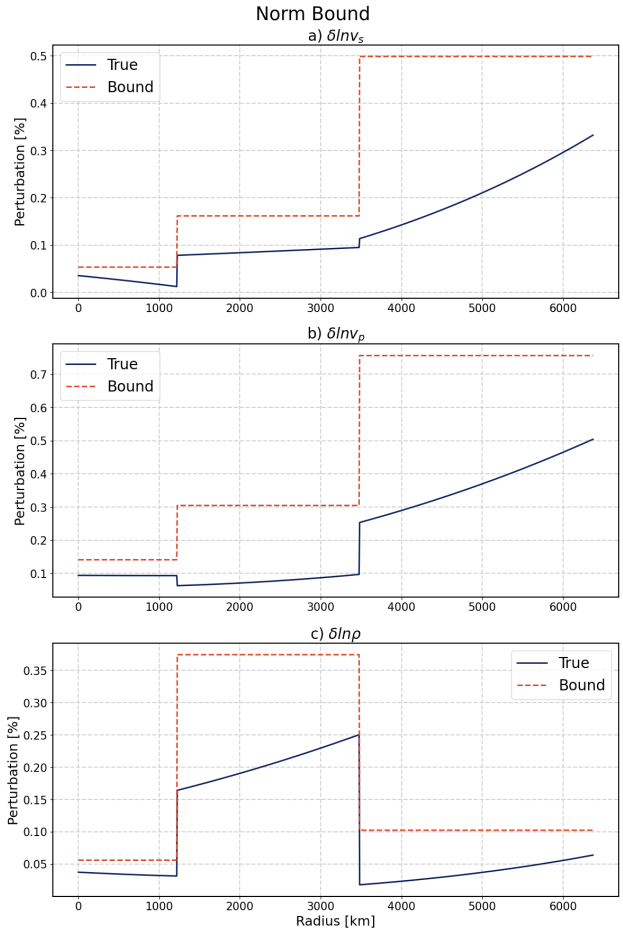

Figure 10. Case 2: Arbitrary quasi-random synthetic true model and the upper bound functions used to compute the prior upper bound norm.

Restelli et al. 2024), i.e. 143 modes with their sensitivity to $\delta \ln(v_s), \delta \ln(v_p)$ and $\delta \ln(\rho)$ concentrated mostly in the mantle (see Fig. 11).

3.2.2 Resolution Analysis

Before introducing any data or model values, we are able to perform a simple resolution analysis to investigate where and at what resolution our data contain information regarding the Earth model. While the SOLA-DLI solution depends on the model norm bound via M (see Eq. 19), indirectly on the data via $\|\tilde{m}\|_{\mathcal{M}}$ (see Eq. 21), and on the relationships between the target kernels and sensitivity kernels via $\mathcal{H}$ (see Eq. A.14–A.23), the resolving kernels only depend on the data geometry, i.e. data sensitivity kernels.

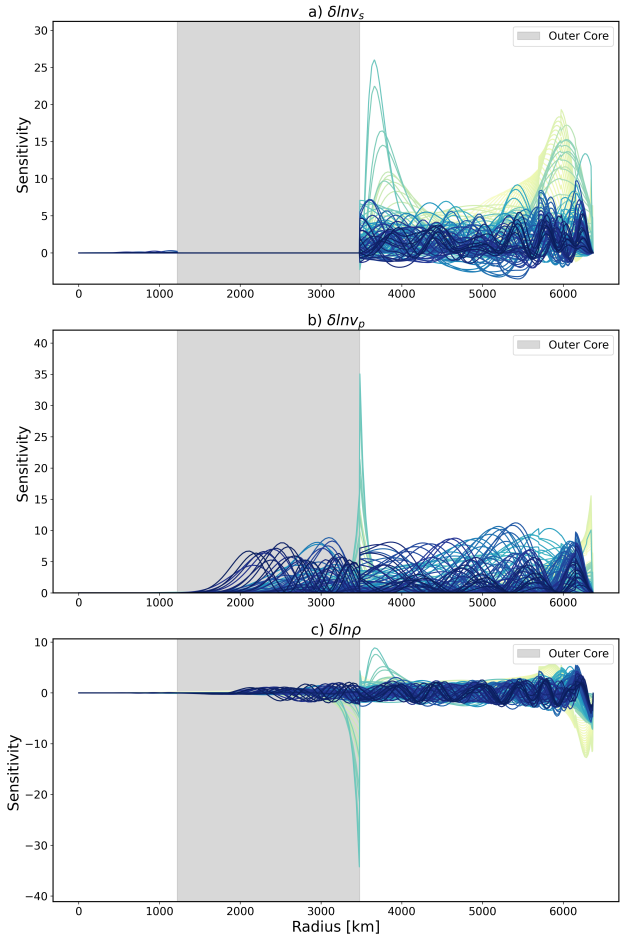

Figure 11. Case 2: Normal mode sensitivity kernels for a) $\delta \ln(v_s)$, b) $\delta \ln(v_p)$ and c) $\delta \ln(\rho)$, obtained using a modified version of OBANI. The shaded region indicates the depth range of the outer core, where the sensitivity to v_s is zero.

The diagonal elements of the matrix $\mathcal{H}$ can be shown to equal:

$$\mathcal{H}_{kk} = \sum_j \left\| T^{j,(k)} - A^{j,(k)} \right\|^2_{\mathcal{M}_j} \tag{48}$$

which essentially quantifies the cumulative difference between our target and resolving kernels. Using $\mathcal{H}$ we can also define the resolving misfit as a more useful metric:

$$R_k = \frac{\sqrt{\mathcal{H}_{kk}}}{\sum_j \left\| T^{j,(k)} \right\|_{\mathcal{M}_j}} = \frac{\sqrt{\sum_j \left\| T^{j,(k)} - A^{j,(k)} \right\|^2_{\mathcal{M}_j}}}{\sum_j \left\| T^{j,(k)} \right\|_{\mathcal{M}_j}} \tag{49}$$

which is a generalisation of the "resolution misfit" as defined in Restelli et al. (2024). The resolving misfit is 0 when all the resolving kernels associated with some property evaluated at $r^{(k)}$ are equal to the corresponding target kernels. This would mean that our data contain exact information about the desired property and the property error bounds are 0. On the other hand, the resolving misfit

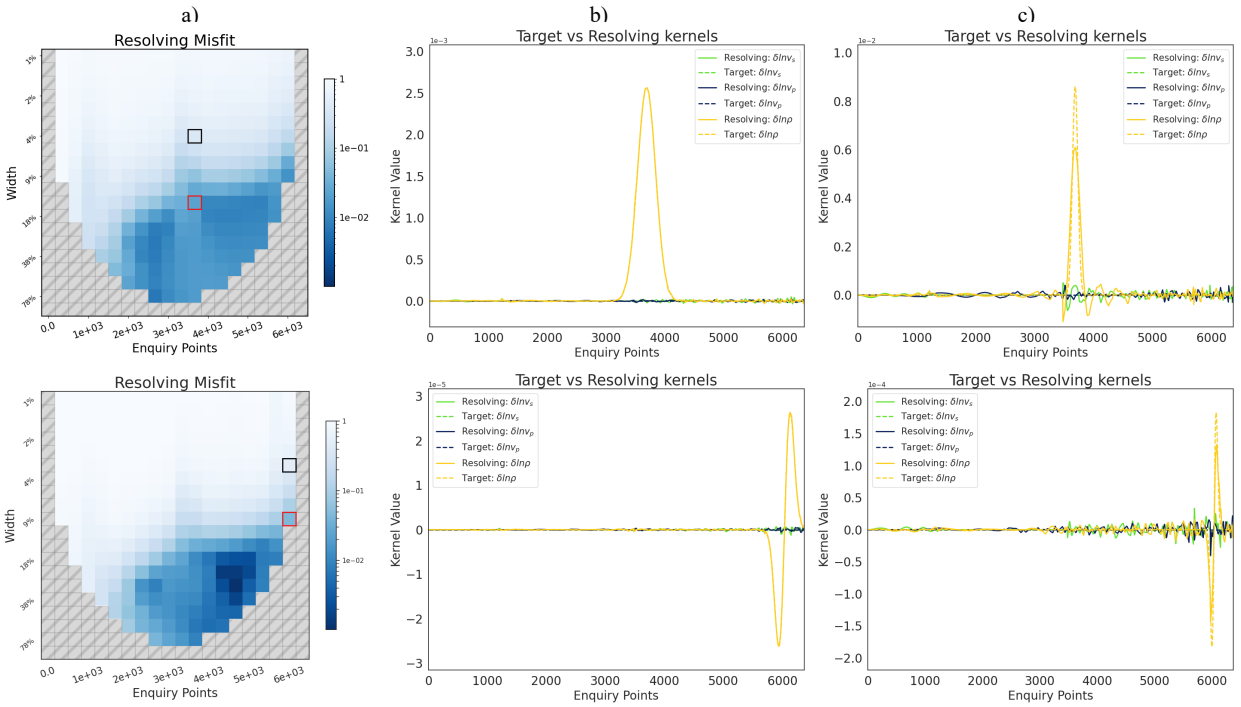

Figure 12. Case 2: Resolution analysis for a Gaussian average (top) and gradient (bottom) target for $\delta \ln(rho)$ using realistic mode sensitivity kernels. The resolving misfit (left) can be computed without the need for data or any prior norm bound information. The middle and right panels illustrate the target and resolving kernels for a wide and thin target, including the contaminant kernels that indicate trade-offs between parameters.

is equal to 1 when our resolving kernels are zero, which would correspond to a complete lack of sensitivity of our data to the desired property. It is important to note that the computation of the resolving misfit does not use the data vector d nor any prior model information, it only uses the "geometry of the data set", i.e. the sensitivity kernels (Latallerie et al. 2024). Fig. 12 illustrates the information that is provided by the resolving misfit (left column). As indicated by a low resolving misfit (darker shades of blue), our data mostly contain information in the mantle, as expected for this set of sensitivity kernels. The resolving misfit is also typically low for wide target kernels (over 18% domain width, or over 1000 km). Wide gradient kernels can be better recovered in the mantle, while wide averaging kernels can be better recovered in the lower outer core, indicating that our choice of target (i.e. property) is important.

3.2.3 Trade-offs between physical parameters

When our data are sensitive to two or more physical parameters, it may become difficult or impossible to obtain properties of a single parameter in isolation from the others. These trade-offs between physical parameters pose problems for interpretations, particularly in regions such as the lower mantle where the sensitivity of normal modes to seismic velocities and density is similar.

Our setup with SOLA-DLI, where we explicitly set the target kernels for parameters not of interest to zero, enables us to easily visualise and consider model parameter trade-offs. Suppose we are interested in some local property of $\delta \ln(\rho)$, for example the Gaussian local average in the deep mantle or Gaussian gradient in the mantle transition zone (Fig. 12). If we choose low resolution (wide) target kernels (middle column), we find that the resolving kernels for $\delta \ln(\rho)$ match the target kernels well. Furthermore, the resolving kernels for $\delta \ln(v_s)$ and $\delta \ln(v_p)$ also match their respective target kernels, which are just zero. Such zero or near zero resolving kernels indicate that the trade-off between the physical parameter of interest and the other physical parameters is small. However, if we choose higher resolution (thin) target kernels (right column), we notice that the resolving kernels are struggling to match their respective target kernels. The resolving kernels for $\delta \ln(v_p)$ and particularly $\delta \ln(v_s)$ are far from zero, indicating significant trade-offs with the desired property of $\delta \ln(\rho)$, which are regarded as *contaminants*. Such trade-offs between physical parameters are naturally taken into account by SOLA-DLI and typically result in higher error bounds on the property. If instead we would account for the sensitivity to $\delta \ln(v_p)$ and $\delta \ln(v_s)$ by scaling the sensitivity kernels, we would obtain tighter bounds, at the expense of assuming more prior information.

3.3 Case 3: Discretised inversions using continuous SOLA-DLI

This case study serves to illustrate how we can obtain a family of discretised model solutions using SOLA-DLI, which we compare with a typical least-squares inversion model solution.

3.3.1 Setup

Here, we consider a model m with only one physical parameter, denoted also m (see the true model in Fig. 13 a)). Our model space $\mathcal{M}$ is $PCb[0,1]$ and the data are given by:

$$d_i = \langle K_i, m \rangle_{\mathcal{M}} \tag{50}$$

where K_i are some quasi-randomly functions, generated again using Eq. 42 (see Fig. 13 b)).

In this setup, we choose to discretise the model using a Fourier expansion. The resulting basis functions are (see Fig. 13 c)):

$$B_l(r) = \begin{cases} 1, & l = 0 \\ \sqrt{2}\sin\big(2\pi\frac{l+1}{2}r\big), & l \text{ odd} \\ \sqrt{2}\cos\big(2\pi\frac{l}{2}r\big), & l \text{ even} \end{cases} \tag{51}$$

and a possible model expansion with Fourier coefficients p_l is given by:

$$m(r) \approx \sum_l p_l B_l(r). \tag{52}$$

The discretised model–data relation used for the least-squares inversion is:

$$d_i = \sum_l \langle K_i, B_l \rangle_{\mathcal{M}} p_l = \sum_l \Gamma^*_{il} p_l, \tag{53}$$

where (see also Appendix A2):

$$\Gamma^*_{ij} = \langle K_i, B_l \rangle_{\mathcal{M}}. \tag{54}$$

This leads to the following least-squares solution for p_l:

$$\hat{p} = (\Gamma\Gamma^*)^{-1}\Gamma d \tag{55}$$

Using the least-squares solution $\{\hat{p}_l\}$, we can thus find the corresponding model solution by using the Fourier expansion:

$$\hat{m} = \sum_l \hat{p}_l B_l. \tag{56}$$

To obtain the SOLA-DLI solution, we consider the Fourier coefficients p_l to be elements of a property vector obtained from the property mapping:

$$p_l = \mathcal{T}(m) = \langle B_l, m \rangle_{\mathcal{M}}. \tag{57}$$

We also introduce a prior model norm bound (see Fig. 13 a)). This leads to the following SOLA-DLI problem:

Given

$$d_i = \langle K_i, m \rangle_{\mathcal{M}} \tag{58}$$

Find

$$p_l = \langle B_l, m \rangle_{\mathcal{M}} \,. \tag{59}$$

This problem is readily solved using Eq. 19 to obtain upper and lower bounds for the possible values of the Fourier coefficients:

$$\bar{p}_l \in [\tilde{p}_l - \epsilon_l, \tilde{p}_l + \epsilon_l]. \tag{60}$$

In this case $\tilde{p}_l$ are the Fourier coefficients of the least norm solution to Eq. 50, which are not to be confused with the least-squares solution $\hat{p}_l$. The property bounds obtained from SOLA-DLI (Equation 60 offer a family of solutions that can be sampled.

3.3.2 Discretised least-squares vs SOLA-DLI solution

We compute both the least-squares and discretised SOLA-DLI solution using different number of data points (50, 70 or 100), solving for 29 Fourier coefficients. SOLA-DLI provides bounds on the Fourier coefficients, and we therefore have to draw samples from the distribution for each Fourier coefficient distribution to obtain a model solution. Fig. 14 shows the results for both the Fourier coefficients and the resulting model solutions, including two random samples for the SOLA-DLI solution. When using few data (Fig. **??**, the least-squares inversion generally struggles to retrieve Fourier coefficients close to the true ones, while the bounds of the SOLA-DLI solution always encompass the true property. However, for certain Fourier coefficients, and in certain parts of the model, the least-squares solution appears to approach the true property (Fourier coefficients) and the true model better than the SOLA-DLI solution. Increasing the number of data to 70 leads to a better least-squares solution, especially for the first 10 Fourier coefficients, and tighter bounds of the SOLA-DLI solution. However, it now becomes clear that the SOLA-DLI bounds offer more accurate information, always encompassing the true Fourier coefficients and better resembling the

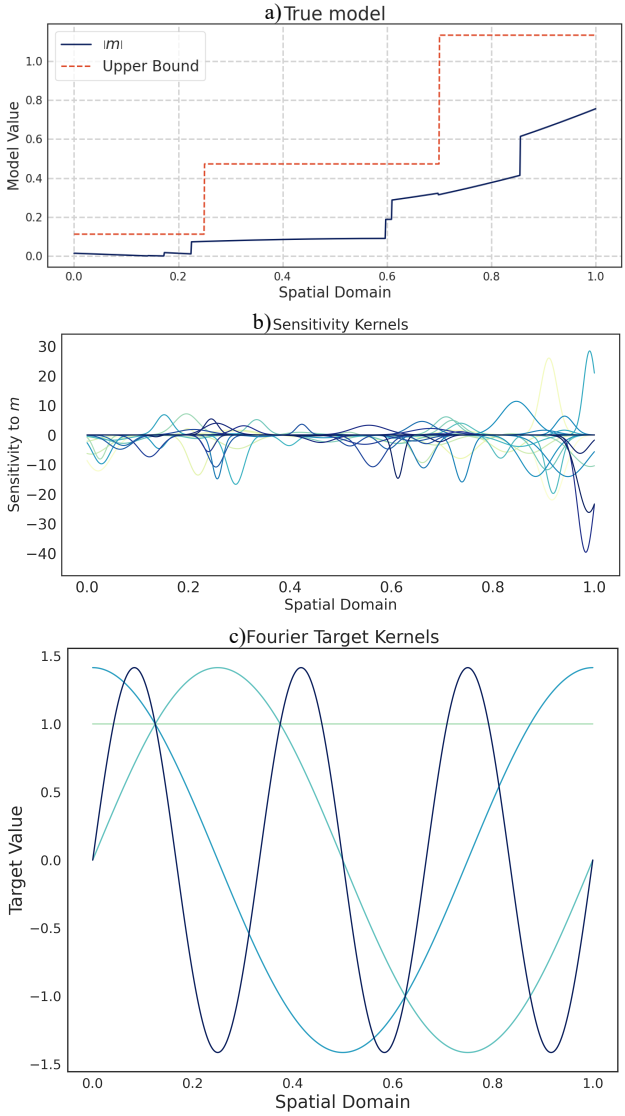

Figure 13. Case 3: True model m with the upper bound function used to compute the prior model norm bound.Case 3: Synthetic quasi-random sensitivity kernels.Case 3: Examples of four Fourier basis functions.

true model. When we further increase the number of data points to 100 (see Figs. **??** and **??**), we note that the SOLA-DLI solution converges to the true Fourier coefficients and model, while the least-squares inversion systematically deviates.

In our synthetic setup, it is possible to compute the data correction term, that captures the components of the true model that are not within the span of the basis functions (see Section 2.5, Eq. 2.5). When we correct the data, using our knowledge of the true model, we find that the least-squares inversion solution converges to the true property (Fourier coefficients), even for few data (Fig. 15). This demonstrates the equivalence of the discretised least-squares and SOLA-DLI solu-

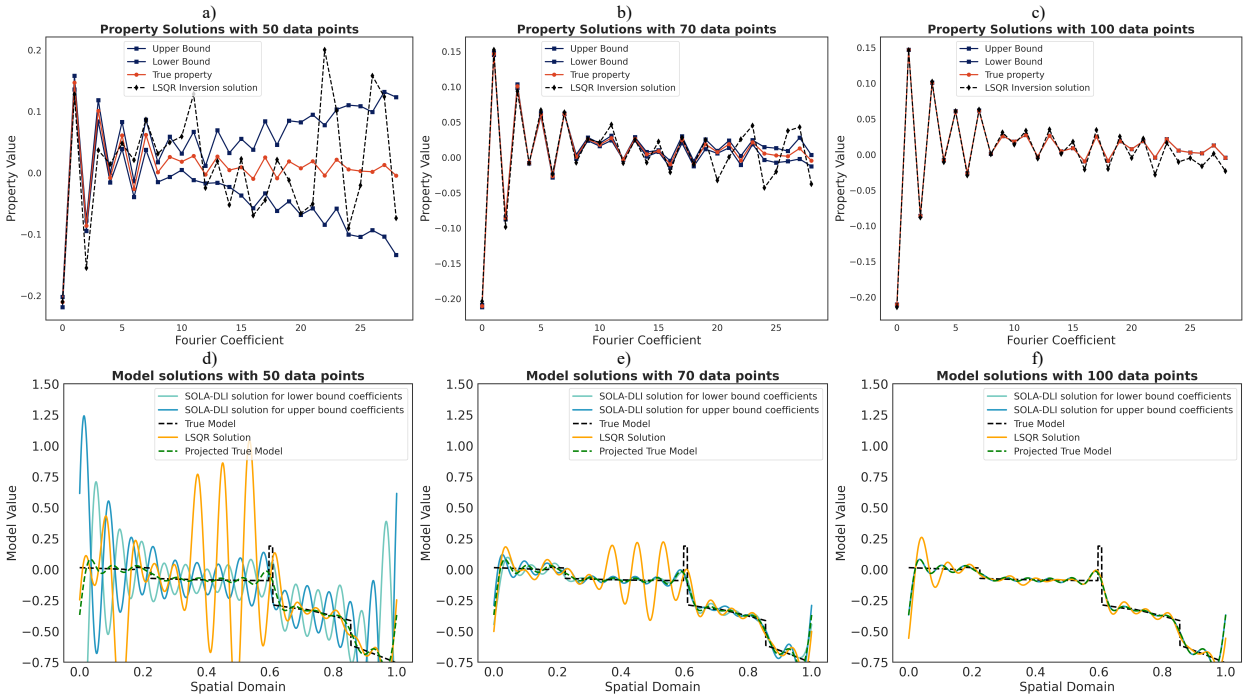

Figure 14. Case 3: Comparison between least-squares and SOLA-DLI solution for a model discretised using Fourier basis functions. a–c) Fourier coefficients from discretised least-squares inversion and SOLA-DLI using a) 50 data points, b) 70 data points and c) 100 data points. d–f) Discretised model solution from discretised least-squares and SOLA-DLI inversion using d) 50 data points, e) 70 data points and f)100 data points

tions. However, in real world applications, when the true model is unknown, this data correction term cannot be computed. Consequently, the SOLA-DLI solution is preferred over the discretised least-squares inversion method.

4 DISCUSSION

Inference methods such as Backus-Gilbert and SOLA-DLI are centered around specific questions that we wish to answer, in contrast to inversion methods that typically aim to obtain the entire model. The questions we wish to answer guide our choices of specific target kernels. We have illustrated several possibilities in this contribution (e.g. uniform averages, Gaussian averages or gradients), but there are many other useful choices of target kernels. An obvious next step is the implementation of 2D and 3D target kernels that extend beyond the classic 2D and 3D uniform averaging kernels (e.g. Zaroli 2016; Latallerie et al. 2022; Freissler et al. 2024). Utilising ideas

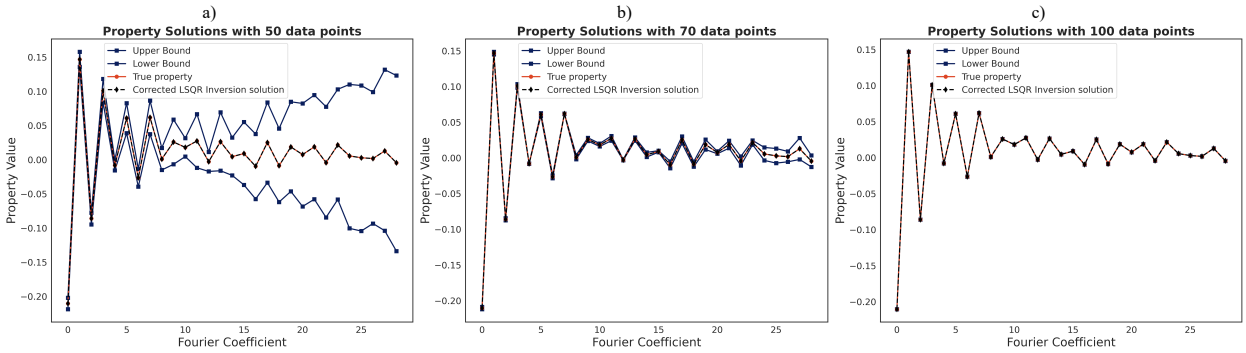

Figure 15. Case 3: Comparison between Fourier coefficients obtained using SOLA-DLI and the discretised least-squares method with a data correction term, using a) 50 data points, b) 70 data points and c) 100 data points. A comparison with the least-squares solutions in Fig. 14 indicates that the data correction leads to the systematic error in the Fourier coefficients.

from signal processing, we can think of targets intended to detect 3D spatial gradients, spherical harmonic coefficients, or unique target kernels to detect shape specific features inside the Earth. We also foresee applications of the theory in other geophysical disciplines, e.g. geomagnetic and geodetic imaging studies.

While the choice of target kernels depends in first instance on the question we wish to answer, it is also important to consider the geometry of our data (i.e. sensitivity kernels). Certain types of local averages can be better constrained depending on the details of our sensitivity kernels, e.g. uniform local averages are difficult to constrain with smooth normal mode kernels (Section 3.2). It then depends on the specifics of the question we want to answer, whether we may be able to find slight variations of the question whose associated target kernels are more easily reconstructed from linear combinations of the sensitivity kernels. While this might lead to more mathematically complex target kernels, the interpretation of the results will always remain well defined. One could even use families of target kernels that contain hyperparameters found by optimising the ability of our sensitivity kernels to reconstruct them. Such approaches can be extended and lead to applications in observable optimisation - finding the observable best constrained by our data.

The introduction of prior model information via the model norm bound is of great importance in the SOLA-DLI method. The model norm bound (L_2 norm) chosen here is the most common due to its mathematical simplicity, but as pointed out by Al-Attar (2021), there might be better prior

constraints. Other norms for our model space can be thought off that allow to place bounds on the maximum point-value of the true model, or its gradients. Such modifications necessitates the use of more general spaces than Hilbert spaces, which add significant theoretical complications. We refer the interested reader to Al-Attar (2021) for the theoretical modifications needed.

An important aspect that has not been used or addressed in this contribution is the presence of data errors, i.e. an assumption of perfect (true) data with zero uncertainty. In such a situation, the data provide deterministic constraints, hence the name "DLI". In contrast, if there exists a probability measure defined on the data space and the true data are uncertain, then we would be in a probabilistic constraint case. Past applications have included data measurement errors using various variations of classic Backus-Gilbert methods (Backus & Gilbert 1970; Zaroli 2016, 2019; Freissler et al. 2020; Latallerie et al. 2022; Restelli et al. 2024). However, tackling them requires added mathematical complexity which, we believe, would have distracted from the main points of this paper. We believe that the main ideas presented here will not change by the introduction of data errors, but the output of SOLA-DLI problems will be affected. Notably, the property bounds will increase or might be replaced by probability density functions, as they will contain the effects of imperfect resolving kernels and propagated data errors. Future work will focus on the modifications to the theory that are required for dealing with real data, whose uncertainties cannot be neglected.

5 CONCLUSION

In this contribution, we have presented the theory and possible applications of SOLA-DLI, which combines the Backus-Gilbert SOLA method with deterministic linear inferences. Through three synthetic case studies with error-free data, we have illustrated the abilities and advantages of SOLA-DLI over other inverse and inference methods. Specifically, we have shown how the method can be used to constrain different properties of an unknown model, defined by different target kernels. By using a norm bound on the model space, the values of these properties can be constrained and interpreted through the lens of the target kernels, rather than the resolving kernels. Using a set of realistic normal-mode data sensitivity kernels, we have also shown how a resolution analysis can

be performed using SOLA-DLI and how trade-offs between physical parameters can be visualised and quantified. Finally, by selecting a set of basis function coefficients as our target properties, we have demonstrated that we can infer robust bounds on the coefficients of the unknown model's projection onto these basis functions. This enables us to generate families of discretised model solutions to an equivalent inversion problem, showing that the development of Earth models spanning the full spatial domain can be obtained using Backus-Gilbert methods. While the presented work assumes error-free data, data errors can be included into the framework in several ways and are expected to primarily lead to an increase in the property error bounds.

6 ACKNOWLEDGEMENTS

AMM is funded by a NERC DTP studentship NE/S007474/1 and gratefully acknowledges their support. AMM also received funding from Royal Society grant RF\ERE\210182 awarded to PK. CZ acknowledges financial support from ITES (Institut Terre et Environnement de Strasbourg, UMR 7063) for a research visit to Oxford. PK acknowledges financial support from a Royal Society University Research Fellowship (URF\1\180377). Normal mode sensitivity kernels were computed using a modified version of OBANI, which is based on the work of Woodhouse & Dahlen (1978). The authors thank Sam Scivier and Franck Latallerie for fruitful theoretical discussions. The following Python packages were used extensively for producing and plotting the synthetic data: Scipy (Virtanen et al. 2020), Numpy (Harris et al. 2020) and Matplotlib (Hunter 2007).

7 DATA AVAILABILITY

The sensitivity kernels and all the codes used to produce the figures in this paper can be found at `https://github.com/Adrian-Mag/SOLA_DLI` (Mag et al. 2024).

APPENDIX A: DERIVATION OF THE SOLA-DLI SOLUTION

In this appendix we derive the equations and equalities needed for the solution of the SOLA-DLI method presented in Section 2.3.

A1 Data-model relationships

Given a model $m = (m^1, m^2, ...)$, the relationship between the model and data is defined by:

$$d_i = G(m) = \sum_j \left\langle K_i^j, m^j \right\rangle_{\mathcal{M}_j}. \tag{A.1}$$

It is useful to define:

$$G^j(m) = \left\langle K_i^j, m^j \right\rangle_{\mathcal{M}_j} \tag{A.2}$$

$$G = \sum_j G^j. \tag{A.3}$$

The adjoint of G is defined by:

$$\langle G(m), d' \rangle_{\mathcal{D}} = \langle m, G^*(d') \rangle_{\mathcal{M}} \tag{A.4}$$

for all $m \in \mathcal{M}$ and $d' \in \mathcal{D}$. We expand the LHS:

$$\sum_i \sum_j \left\langle K_i^j, m^j \right\rangle_{\mathcal{M}_j} d_i' = \langle m, G^*(d') \rangle_{\mathcal{M}} \tag{A.5}$$

For the RHS, we use the formula for the inner product in the direct sum space $\mathcal{M}$:

$$\langle a, b \rangle_{\mathcal{M}} = \sum_j \left\langle a^j, b^j \right\rangle_{\mathcal{M}_j} \tag{A.6}$$

where a, b are some members of $\mathcal{M}$. Therefore, we write

$$\sum_i \sum_j \left\langle K_i^j, m^j \right\rangle_{\mathcal{M}_j} d_i' = \sum_j \left\langle m^j, G^{j*}(d') \right\rangle_{\mathcal{M}_j} \tag{A.7}$$

Taking the sum over i inside, we can also write:

$$\sum_j \left\langle \sum_i d_i' K_i^j, m^j \right\rangle_{\mathcal{M}_j} = \sum_j \left\langle m^j, G^{j*}(d') \right\rangle_{\mathcal{M}_j} \tag{A.8}$$

and we identify:

$$G^{j*}(d') = \sum_i d'_i K_i^j \tag{A.9}$$

$$G^* = (G^{1*}, G^{2*}, ...) \tag{A.10}$$

G^* maps elements from the data space to elements (tuples) in the model space. A similar approach shows that the adjoint of the property mapping $\mathcal{T}$ is given by:

$$\mathcal{T}^{j*}(p) = \sum_k p^{(k)} T^{j,(k)} \tag{A.11}$$

$$\mathcal{T}^* = (\mathcal{T}^{1*}, \mathcal{T}^{2*}, ...) \tag{A.12}$$

A2 H matrix

The $\mathcal{H}$ matrix introduced in Section 2.3 quantifies the difference between the target and resolving kernels. It is defined by Al-Attar (2021, see Equation 2.84) as:

$$\mathcal{H} = HH^* \tag{A.13}$$

$$H = \mathcal{T} - \mathcal{A} \tag{A.14}$$

where $\mathcal{A}$ is the "approximate mapping", given by:

$$\mathcal{A} = \mathcal{T}G^*(GG^*)^{-1}G \tag{A.15}$$

This mapping takes any model $m \in \mathcal{M}$ into the data space $d \in \mathcal{D}$, then finds the least norm solution to $G(m) = d$ and maps this least norm solution into the property space. When applied to one of the possible model solutions $U_M \bigcap \{m\}$, it gives the property of the model solution that has the smallest norm. Combining (A.13), (A.14), and (A.15) we obtain:

$$\mathcal{H} = (\mathcal{T} - \mathcal{A})(\mathcal{T} - \mathcal{A})^* \tag{A.16}$$

$$\mathcal{H} = \mathcal{T}\mathcal{T}^* - \mathcal{T}G^*(GG^*)^{-1}G\mathcal{T}^* \tag{A.17}$$

Let us denote

$$\Lambda := GG^*. \tag{A.18}$$

Using a simple application of Equation (A.1) and (A.10), we can then easily obtain:

$$\Lambda_{iq} = \sum_j \left\langle K_i^j, K_q^j \right\rangle_{\mathcal{M}_j} \tag{A.19}$$

Similarly, we denote:

$$\chi := \mathcal{T}\mathcal{T}^* \tag{A.20}$$

$$\chi_{k,l} = \sum_j \left\langle T^{j,(k)}, T^{j,(l)} \right\rangle_{\mathcal{M}_j} \tag{A.21}$$

and

$$\Gamma := \mathcal{T}G^* \tag{A.22}$$

$$\Gamma_{ki} = \sum_j \left\langle T^{j,(k)}, K_i^j \right\rangle_{\mathcal{M}_j} \tag{A.23}$$

Using the definitions of Λ, χ, Γ we can write (A.17) as:

$$\mathcal{H} = \chi - \Gamma\Lambda^{-1}\Gamma^T. \tag{A.24}$$

Figure A1 provides a visualisation of the ellipse in the property space as determined by $\mathcal{H}$ when only two properties are considered.

A3 Error Bounds

The error bounds defined in Equations (19) and (22) are derived from the property bounds defined by Al-Attar (2021, see Equation 2.84) as:

$$\left\langle \mathcal{H}^{-1}(p - \tilde{p}), p - \tilde{p} \right\rangle_{\mathcal{P}} \leq M^2 - \|\tilde{m}\|^2_{\mathcal{M}} \tag{A.25}$$

Equation (A.25) describes a hyperellipsoid centered on $\tilde{p}$ with major axes given by the eigenvalues of $\mathcal{H}^{-1}$ scaled by $\sqrt{M^2 - \|\tilde{m}\|^2_{\mathcal{M}}}$. If the matrix $\mathcal{H}$ is diagonal, then the inverse is trivial to find and the hyperellipsoid has its major axes aligned with the coordinate axes of the property space. In all other cases, the hyperellipsoid will have some arbitrary orientation and $\mathcal{H}$ will be difficult to invert numerically.

To avoid numerical complications, we use here a different, more relaxed approximation for the

error bounds given in the form:

$$\|p - \tilde{p}\|^2_{\mathcal{P}} \leq (M^2 - \|\tilde{m}\|^2_{\mathcal{M}}) diag(\mathcal{H}) \tag{A.26}$$

where $diag(\mathcal{H})$ is the diagonal of $\mathcal{H}$. We can also write this in component form:

$$\left\|p^{(k)} - \tilde{p}^{(k)}\right\|^2_{\mathcal{P}} \leq (M^2 - \|\tilde{m}\|^2_{\mathcal{M}}) \mathcal{H}_{kk} \tag{A.27}$$

Inequality (A.27) describes a hyperparallelepiped that contains the error bounds of (A.25) with sides parallel to the coordinate axes of $\mathcal{P}$ (see Fig. A1). As this approximation does not require the inversion of the $\mathcal{H}$ matrix, it is computationally advantageous. Visually, the hyperellipsoid fits "perfectly" inside the hyperparallelepiped (Fig. A1), but the error bounds of the hyperparallelepiped are easier to visualise in a static plot (see for example first row of Fig. 7). The hyperellipsoid encodes the correlations between the error bounds of the various components of the property vector (such as the correlation between the error bounds of two different local averages). Plotting the bounds for each component of the property vector simultaneously would therefore be very difficult, since the error bounds of each property component would depend on the values of the bounds on all other property components. The hyperparallelepiped ignores these correlations, simplifying thus the plotting. However, it overestimates the property bounds, which will likely make it more difficult to interpret the property values.

To show how (A.27) arises from (A.25), we need to prove the following:

Given that

$$x^T A^{-1} x \leq b \tag{A.28}$$

Show that

$$x_k^2 \leq b A_{kk} \tag{A.29}$$

where $x = \tilde{p} - \epsilon$, $A = \mathcal{H}$, and $b = M^2 - \|\tilde{m}\|^2_{\mathcal{M}}$. To prove this, we start by finding the maximum extent of the hyperellipsoid (A.28) along the k^{th} coordinate axis, which can be described

mathematically as:

$$\text{Find } \max(c^T x) \tag{A.30}$$

$$\text{Given that } x^T A^{-1} x \leq b. \tag{A.31}$$

where c^T will be chosen later to be a vector with all entries 0 except the k^{th} one. We shall use the Lagrangian approach to solve this problem. We introduce the slack constant s and use it to transform the inequality A.31 into an equality (slack constraint):

$$x^T A^{-1} x - b + s^2 = 0 \tag{A.32}$$

Let λ be a Lagrange multiplier. The problem then becomes finding the extremum points of the Lagrangian:

$$f(x, \lambda) = c^T x + \lambda(b - x^T A^{-1} x - s^2). \tag{A.33}$$

Differentiating f with respect to x and setting the result to zero leads to:

$$c - 2\lambda A^{-1} x = 0. \tag{A.34}$$

Notice that $\lambda = 0$ leads to $c = 0$, which is a contradiction. Therefore we must have $\lambda \neq 0$ and $s^2 = 0$, which means that our constraint is active. Assuming A^{-1} to be invertible, we obtain:

$$x = \frac{Ac}{2\lambda}. \tag{A.35}$$

We next differentiate f with respect to λ (using $s^2 = 0$ since we have shown the constraint to be active) to obtain the second Lagrange equation. Setting the result equal to zero leads to:

$$b - x^T A^{-1} x = 0. \tag{A.36}$$

Substituting (A.35) into (A.36) and rearranging for b, we obtain:

$$b = \left(\frac{Ac}{2\lambda}\right)^T A^{-1} \frac{Ac}{2\lambda}. \tag{A.37}$$

Since A is symmetric this leads to:

$$\lambda^2 = \frac{c^T A c}{4b}. \tag{A.38}$$

Assuming that A is positive definite (its eigenvalues give the lengths of the hyperellipsoids' major

axes), we must have

$$\lambda = \frac{\sqrt{c^T Ac}}{2\sqrt{b}}. \tag{A.39}$$

Finally, using (A.38) and (A.35) the optimal vector solution x can be expressed for any vector c as:

$$x = \frac{Ac}{2\lambda} = \frac{Ac}{\frac{\sqrt{c^T Ac}}{\sqrt{b}}} = \frac{\sqrt{b}Ac}{\sqrt{c^T Ac}}, \tag{A.40}$$

and the maximal value of $c^T x$ is thus:

$$\sqrt{bc^T Ac}. \tag{A.41}$$

Now, we consider a fixed index k between 1 and N, and we define c to be the following vector:

$$c := (\delta_{ik})_{1 \leq i \leq N}, \tag{A.42}$$

where δ_{ik} is 1 if $i = k$, and 0 if $i \neq k$. Substituting this for c in (A.41), we obtain:

$$\max(x_k) = \sqrt{bA_{kk}} \tag{A.43}$$

or equivalently:

$$x_k \leq \sqrt{bA_{kk}} \tag{A.44}$$

If instead we choose $c = -\delta_{ik}$, then we have:

$$x_k \geq -\sqrt{bA_{kk}} \tag{A.45}$$

These two inequalities can be summarised in the final answer:

$$(x_k)^2 \leq bA_{kk} \tag{A.46}$$

REFERENCES

Al-Attar, D., 2021. Linear inference problems with deterministic constraints, *arXiv preprint arXiv:2104.12256*.

Amiri, S., Maggi, A., Tatar, M., Zigone, D., & Zaroli, C., 2023. Rayleigh wave group velocities in north-west iran: Sola backus-gilbert vs. fast marching tomographic methods, *Seismica*, **2**(2).

Backus, G., 1970a. Inference from inadequate and inaccurate data, i, *Proceedings of the National Academy of Sciences*, **65**(1), 1–7.

Table A1. Table summarising the main mathematical symbols used in the manuscript. Elements are grouped on columns and rows depending on the relationships between them. For example, $\bar{m}$ is part of the model space $\mathcal{M}$ and is related to $\bar{d}$ through G, which is determined by K. Similarly, $\tilde{p}$ is related to $\bar{m}$ through $\mathcal{A}$ and to $\tilde{m}$ through $\mathcal{T}$.

$\mathcal{M}$	$\mathcal{D}$	$\mathcal{P}$		
Model space	Data space	Property space	Mapping	Kernels
$\bar{m}$: TRUE model	$\bar{d}$: TRUE data		G: Forward mapping	K: Sensitivity kernels
$\bar{m}$: TRUE model		$\bar{p}$: TRUE property	$\mathcal{T}$: Property mapping	T: Target kernels
$\bar{m}$: TRUE model		$\tilde{p}$: Aproximate property	$\mathcal{A}$: Approximate mapping	A: Resolving kernels
$\tilde{m}$: Least norm model		$\tilde{p}$: Aproximate property	$\mathcal{T}$: Property mapping	T: Target kernels
		ϵ: Property error	$\mathcal{H}$: Hyperellipsoid matrix	
j: Physical parameter index	i: Data index	k: Property index		
N': Number of physical parameters	N: number of data			

Backus, G., 1970b. Inference from inadequate and inaccurate data, iii, *Proceedings of the National Academy of Sciences*, **67**(1), 282–289.

Backus, G. & Gilbert, F., 1968. The resolving power of gross earth data, *Geophysical Journal International*, **16**(2), 169–205.

Backus, G. & Gilbert, F., 1970. Uniqueness in the inversion of inaccurate gross earth data, *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*, **266**(1173), 123–192.

Backus, G. E. & Gilbert, J. F., 1967. Numerical applications of a formalism for geophysical inverse problems, *Geophysical Journal International*, **13**(1-3), 247–276.

Fichtner, A., Bunge, H.-P., & Igel, H., 2006. The adjoint method in seismology: I. theory, *Physics of the Earth and Planetary Interiors*, **157**(1-2), 86–104.

Freissler, R., Zaroli, C., Lambotte, S., & Schuberth, B. S. A., 2020. Tomographic filtering via the generalized inverse: a way to account for seismic data uncertainty, *Geophysical Journal International*, **223**(1), 254–269.

Freissler, R., Schuberth, B. S. A., & Zaroli, C., 2024. A concept for the global assessment of tomographic resolution and uncertainty, *Geophysical Journal International*, p. ggae178.

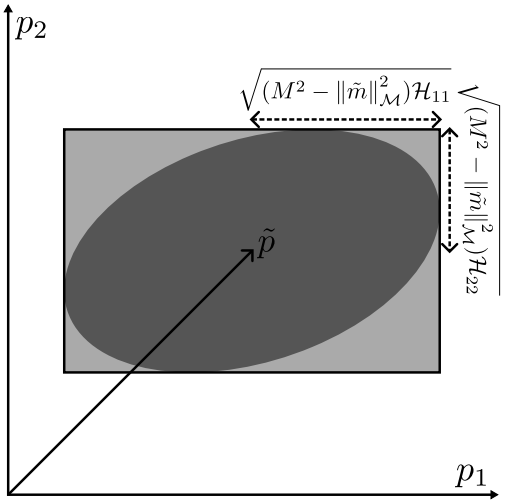

Figure A1. Illustration of the relationship between the hyperellipsoid defined in Equation A.25 and the hyperparallelepiped defined in Equation A.27 for the case when the property space is two-dimensional. p_1 and p_2 could represent, for example, two local averages at two different spatial locations. The dark shaded ellipse contains all the possible combinations of these two properties given by the tighter inequality of Equation A.25. The lighter gray shaded rectangle contains all the possible combinations of these two properties under the simplified inequality in Equation A.27.

Hammer, M. D. & Finlay, C. C., 2019. Local averages of the core–mantle boundary magnetic field from satellite observations, *Geophysical Journal International*, **216**(3), 1901–1918.

Hammer, M. D., Finlay, C. C., & Olsen, N., 2021. Applications for cryosat-2 satellite magnetic data in studies of earth's core field variations, *Earth, Planets and Space*, **73**, 1–22.

Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., Haldane, A., del Río, J. F., Wiebe, M., Peterson, P., Gérard-Marchant, P., Sheppard, K., Reddy, T., Weckesser, W., Abbasi, H., Gohlke, C., & Oliphant, T. E., 2020. Array programming with NumPy, *Nature*, **585**(7825), 357–362.

Hunter, J. D., 2007. Matplotlib: A 2d graphics environment, *Computing in Science & Engineering*, **9**(3), 90–95.

Koelemeijer, P., Ritsema, J., Deuss, A., & Van Heijst, H.-J., 2016. SP12RTS: a degree-12 model of shear- and compressional-wave velocity for Earth's mantle, *Geophys. J. Int.* , **204**(2), 1024–1039.

Latallerie, F., Zaroli, C., Lambotte, S., & Maggi, A., 2022. Analysis of tomographic models using resolution and uncertainties: a surface wave example from the pacific, *Geophysical Journal International*, **230**(2), 893–907.

Latallerie, F., Zaroli, C., Lambotte, S., Maggi, A., Walker, A., & Koelemeijer, P., 2024. Surface-wave tomography with 3d resolution and uncertainty, *Seismica*.

Lau, H. C. & Romanowicz, B., 2021. Constraining jumps in density and elastic properties at the 660 km discontinuity using normal mode data via the backus-gilbert method, *Geophysical Research Letters*, **48**(9), e2020GL092217.

Liu, Q. & Tromp, J., 2008. Finite-frequency sensitivity kernels for global seismic wave propagation based upon adjoint methods, *Geophysical Journal International*, **174**(1), 265–286.

Mag, A., Zaroli, C., & Koelemeijer, P., 2024. Adrian-Mag/SOLA_DLI: SOLA-DLI.

Masters, G. & Gubbins, D., 2003. On the resolution of density within the Earth, *Phys. Earth Planet. Inter.* , **140**(1-3), 159–167.

Moulik, P. & Ekström, G., 2014. An anisotropic shear velocity model of the earth's mantle using normal modes, body waves, surface waves and long-period waveforms, *Geophysical Journal International*, **199**(3), 1713–1738.

Nolet, G., 1985. Solving or resolving inadequate and noisy tomographic systems, *Journal of Computational Physics*, **61**(3), 463–482.

Nolet, G., 2008. A breviary of seismic tomography, *A Breviary of Seismic Tomography, by Guust Nolet, Cambridge, UK: Cambridge University Press, 2008*.

Oldenburg, D., 1981. A comprehensive solution to the linear deconvolution problem, *Geophysical Journal International*, **65**(2), 331–357.

Pijpers, F. & Thompson, M., 1992. Faster formulations of the optimally localized averages method for helioseismic inversions, *Astronomy and Astrophysics (ISSN 0004-6361), vol. 262, no. 2, p. L33-L36.*, **262**, L33–L36.

Pijpers, F. & Thompson, M., 1994. The sola method for helioseismic inversion, *Astronomy and Astrophysics (ISSN 0004-6361), vol. 281, no. 1, p. 231-240*, **281**, 231–240.

Rawlinson, N., Pozgay, S., & Fishwick, S., 2010. Seismic tomography: a window into deep earth, *Physics of the Earth and Planetary Interiors*, **178**(3-4), 101–135.

Restelli, F., Zaroli, C., & Koelemeijer, P., 2024. Robust estimates of the ratio between s- and p-wave velocity anomalies in the earth's mantle using normal modes, *Physics of the Earth and Planetary Interiors*, **347**, 107135.

Ritsema, J. & Lekić, V., 2020. Heterogeneity of seismic wave velocity in earth's mantle, *Annual Review of Earth and Planetary Sciences*, **48**, 377–401.

Ritsema, J., Heijst, H. J. v., & Woodhouse, J. H., 1999. Complex shear wave velocity structure imaged beneath africa and iceland, *Science*, **286**(5446), 1925–1928.

Tromp, J., Tape, C., & Liu, Q., 2005. Seismic tomography, adjoint methods, time reversal and banana-doughnut kernels, *Geophysical Journal International*, **160**(1), 195–216.

Tsai, V. C., 2023. The future of earth imaging, *Seismological Research Letters*, **94**(5), 2119–2128.

Valentine, A. P. & Sambridge, M., 2023. Emerging directions in geophysical inversion, *Applications of*

Data Assimilation and Inverse Problems in the Earth Sciences, **5**(9).

Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C. J., Polat, İ., Feng, Y., Moore, E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P., & SciPy 1.0 Contributors, 2020. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python, *Nature Methods*, **17**, 261–272.

Woodhouse, J. H. & Dahlen, F. A., 1978. The effect of a general aspherical perturbation on the free oscillations of the earth, *Geophysical Journal International*, **53**(2), 335–354.

Zaroli, C., 2016. Global seismic tomography using Backus–Gilbert inversion, *Geophysical Journal International*, **207**(2), 876–888.

Zaroli, C., 2019. Seismic tomography using parameter-free backus–gilbert inversion, *Geophysical Journal International*, **218**(1), 619–630.

Zaroli, C., Koelemeijer, P., & Lambotte, S., 2017. Toward seeing the earth's interior through unbiased tomographic lenses, *Geophysical Research Letters*, **44**(22), 11–399.