Geophys. J. Int. (2024) **???**, 1–**??**

A concept for the global assessment of tomographic resolution and uncertainty

Roman Freissler[1,*], Bernhard S.A. Schuberth[1], Christophe Zaroli[2]

[1] *Dept. of Earth and Environmental Sciences, Geophysics, Ludwig-Maximilians-Universität München, Theresienstr. 41, 80333 München, Germany*

[2] *Institut Terre et Environnement de Strasbourg, UMR 7063, Université de Strasbourg, EOST/CNRS, 67084 Strasbourg CEDEX, France*

[*] *Corresponding author, E-mail: Roman.Freissler@lmu.de*

17 May 2024

SUMMARY

A major challenge in seismic tomography consists in quantifying and representing model resolution and uncertainty, particularly at global scales. This information is crucial for interpretations of tomographic images and their technical application in geodynamics. However, due to large computational costs, there have been only few attempts so far to coherently analyse the spatially varying resolving power for a complete set of model parameters. Here, we present a concept for an effective evaluation and global representation of the 3-D resolution information contained in a full set of averaging kernels. In our case, these kernels are constructed using the 'Subtractive Optimally Localized Averages' (SOLA) method, a variant of classic Backus-Gilbert inversion suitable for global tomography. Our assessment strategy incorporates the following steps: 1) a 3-D Gaussian function is fitted to each averaging kernel to measure resolution lengths in different directions; 2) we define a classification scheme for the quality of the averaging kernels based on their focus with respect to the estimated 3-D Gaussian, allowing us to reliably identify whether the inferred resolution lengths are robust. This strategy is not restricted to SOLA inversions, but can, for example, be applied in all cases where point-spread functions are computed in other tomographic frameworks.

Together with model uncertainty estimates that are derived from error propagation in the SOLA method, our concept reveals at which locations, resolution lengths and interpretations of model values are actually meaningful. We finally illustrate how the complete information from our analysis can be used to calibrate the SOLA inversion parameters —locally tunable target resolution kernels and trade-off parameters— without the need for visual inspection of the individual resulting averaging kernels. Instead, our global representations provide a tool for designing tomographic models with specific resolution-uncertainty properties that are useful in geodynamic applications, especially for linking seismic inversions to models of mantle flow.

Key words: Seismic tomography - Inverse theory - Body waves - Structure of the Earth

1 INTRODUCTION

Global seismic tomography is the primary technique for revealing the physical structure of the deep Earth. The first tomographic models of the Earth's mantle, developed more than four decades ago, have mainly been concerned with mapping seismic heterogeneity at spherical harmonic degrees of 6-8; that is, at wavelengths of thousands of kilometres (Sengupta & Toksöz 1976; Aki et al. 1977; Dziewonski et al. 1977). Over the years, the resolution of global tomographic images has steadily been improving by the general increase in data coverage, by exploitation of datasets with complementary sensitivity, as well as through advanced forward and inverse modelling techniques (e.g. Ritsema et al. 2011; Schaeffer & Lebedev 2013; Zaroli et al. 2015; French & Romanowicz 2015; Koelemeijer et al. 2016; Fichtner et al. 2018; Lu et al. 2019; Hosseini et al. 2020; Lei et al. 2020).

Still, in many regions there is only little consensus on the seismic heterogeneity at shorter length scales of ∼300-500 km and less. Not only the exact geographic distribution, but in particular the magnitudes of heterogeneity, are difficult to constrain with tomographic methods. For example, regularization with damping and smoothing constraints is typically needed to counteract the ill-posed nature of the problem, but this inevitably biases the recovered model amplitudes (e.g. Ritsema et al. 2007; Nolet 2008; Schuberth et al. 2009a). Further intricacies arise from complex non-linear wavefield effects and trade-offs between physical parameters (e.g Hung et al. 2001; Favier et al. 2004; Zhang & Shen 2008; Mercerat & Nolet 2012; Schuberth et al. 2015; Koroni et al. 2022). The continuing desire in tomographic studies to increase resolution beyond the current limits is not an end in itself, but for global applications motivated by the estimated thickness of the thermal boundary layers of the mantle and the associated expected size of slabs and plumes. In light of such geodynamic considerations, an accurate retrieval of heterogeneity at spatial scales of $\sim$100 km and less is crucial for subsequent quantitative inferences in studies of the lower mantle (e.g. Schuberth et al. 2009b; Koelemeijer et al. 2018; Choblet et al. 2023; Richards et al. 2023), reconstructed time evolution of mantle flow (e.g. Bunge et al. 2003; Spasojevic et al. 2009; Shephard et al. 2010; Horbach et al. 2014; Colli et al. 2018; Ghelichkhan et al. 2021) and surface dynamic topography (e.g. Davies et al. 2019, 2023). In addition to the dynamically inherent size of thermal anomalies, variations in mineral phase assemblage and chemical composition likely occur on even shorter scales (e.g. Stixrude & Lithgow-Bertelloni 2007; Papanagnou et al. 2022).

Despite the great progress in global seismic tomography, relatively few studies addressed explicitly the problem of quantifying the spatially variable resolving power of a given inversion (e.g. Boschi 2003; Ritsema et al. 2004; Soldati et al. 2006; Ritsema et al. 2011; Koelemeijer et al. 2016; Simmons et al. 2019). Those studies have in common that their tomographic systems are based on a linearization of the (weakly) non-linear problem, such that the quantification of resolution is straightforward from a theoretical point of view. This requires the computation of the resolution matrix, which allows for a complete characterization of the underlying effects of imperfect data coverage and regularization. Moreover, the linear nature of the solution lends itself to practical applications where limited resolution is a critical aspect of quantitative model interpretation. An example for this is tomographic filtering of geodynamic Earth models, which is a necessary step to obtain fair and consistent comparisons between these independent theoretical predictions of present-day seismic heterogeneity and the tomographically imaged structures (e.g. Mégnin et al. 1997; Ritsema et al. 2007; Schuberth et al. 2009a; Nerlich et al. 2016; Simmons et al. 2019; Freissler et al. 2020). Together with the posterior covariance matrix, which includes the variances and correlations of model parameters, the non-uniqueness and quality of the tomographic solution can be fully appraised (e.g. Nolet 2008; Simmons et al. 2019).

For non-linear systems, in contrast, a complete formal quantification of resolution and uncertainty is often not viable in practice. In full waveform inversions based on numerical wavefield simulations and adjoint techniques (e.g. Igel et al. 1996; Pratt 1999; Fichtner et al. 2009; Tape et al. 2009; Colli et al. 2013; Krischer et al. 2018; Ma et al. 2022; Rodgers et al. 2022), there is still a gap between the relative wealth of information in the data and the available tools to assess the general non-uniqueness, especially in global models. One possibility to approach this problem is Bayesian inference (Tarantola 2005) in order to elegantly deal with the non-linearity. However, probabilistic approaches that rely on sheer random sampling of the posterior probability density function are out of reach for global scale applications due to the high computational costs of repeatedly solving the forward problem in such analyses. Several other strategies have thus been proposed to address this issue, which mostly involve an approximation of the Hessian matrix for the misfit function in the vicinity of the 'optimal' model. This is motivated by the insight that the Hessian matrix may be exploited in a local sense as the inverse of the posterior model covariance (e.g. Tarantola 2005; Bui-Thanh et al. 2013; Liu et al. 2021). It allows for some practical inferences, such as in extremal bounds analysis (Fichtner 2010), or for an efficient exploration of the model nullspace (e.g. Deal & Nolet 1996; Liu & Peter 2020; Fichtner et al. 2021). Hessian-vector products may also be used to compute local point-spread functions (e.g. Fichtner & Trampert 2011; Fichtner et al. 2013) that are equivalent to the columns of the resolution matrix in a linear framework. However, except for point-spread function tests for a few individual locations in the full-waveform models GLAD-M15 (Bozdağ et al. 2016) and GLAD-M25 (Lei et al. 2020), these sophisticated methods have found only limited usage in global scale applications so far.

It must be noted that even in the linear case, computing formal resolution and uncertainty is a formidable challenge (e.g. Rawlinson et al. 2014). Stochastic techniques may yield specific characteristics of the resolution matrix, such as depth-dependent average resolution lengths (Trampert et al. 2013) or the main diagonal elements (MacCarthy et al. 2011). The diagonal entries give an indication of the resolvability at the parameter location of interest, while resolution lengths characterize the impact range of off-diagonal entries representing inter-parameter trade-offs. More detailed information can be extracted, for example by a statistical resolution matrix (An 2012), or a stochastic estimation of point-spread function parameters, which can in turn be applied to both linear and non-linear problems (Fichtner & Leeuwen 2015). It is also possible in large-scale problems to use direct approaches for computing the resolution matrix that involve efficient numerical strategies (Boschi 2003; Soldati et al. 2006; Bogiatzis et al. 2016). Alternatively, practical approximations can be made to estimate both the resolution matrix and the posterior covariance (e.g. Nolet et al. 1999; Simmons et al. 2019).

A straightforward method to compute directly the resolution as well as uncertainty can be found in the seminal work by Backus & Gilbert (1967, 1968, 1970). In Backus–Gilbert theory, the estimates of individual model parameters can be interpreted as localized spatial averages around a given target location. In contrast to the more commonly used linear methods in tomography, which often employ Tikhonov regularization for norm damping, no a priori constraints on model values need to be prescribed that may bias model amplitudes. Instead, a certain control can be exerted on the trade-off between a favourable spatial structure of averaging kernels (that determine resolution) and the amount of data errors propagating into the averages as model uncertainties (Backus & Gilbert 1970).

A variant of the Backus–Gilbert method, called Subtractive Optimally Localized Averages (SOLA), was introduced to global seismic tomography by Zaroli (2016). Originally formulated and termed SOLA by Pijpers & Thompson (1992, 1994) for 1-D inversions in helioseismology, the method may have been discovered independently by several authors in different contexts (e.g. Oldenburg 1981; Louis & Maass 1990). In geophysics, it was further adapted for solving discrete and continuous 2-D and 3-D large-scale tomographic problems (Zaroli et al. 2017; Zaroli 2019) and has since been applied to surface wave tomography (Latallerie et al. 2022; Amiri et al. 2023), normal modes (Restelli et al. 2024), and modelling the radial magnetic field at the core-mantle boundary (Hammer & Finlay 2019). The great advantage of the SOLA method compared to the classic Backus–Gilbert formulation arises from the implementation of target kernels with prescribed finite size, which specify the volume around the specific parameter location in which the inversion shall provide the spatial average of the model values. The target kernels make it possible to provide a priori information on the expected local resolution length scales (i.e. they allow for potentially exerting a rather direct control on the final resolution), while the so-called trade-off parameter moderates error propagation. Moreover, the SOLA method enables perfectly parallel computations of the model values as well as of the averaging kernels and propagated uncertainty. So far, however, there is no definite method or recipe for the automatic selection of the SOLA inversion parameters, namely the size of the individual target kernels and the particular choice for the trade-off parameters, throughout a complete model. In global seismic tomography, such a strategy would be particularly helpful due to the highly inhomogeneous data coverage, leading to locally different quality of the tomographic solution. The remaining issue in that regard is the lack of tools for assessing the entire set of averaging kernels in a 3-D setting. Furthermore, even if one can compute with SOLA, or any other tomographic method, a complete set of averaging kernels (or point-spread functions), one will never be able to visually calibrate and analyse each one individually. In other words, it is still a challenge in itself to effectively represent and communicate the resolution information (Trampert 1998).

The objective of this paper is therefore twofold: First, we want to explore for a previously employed tomographic dataset, how different inversion parameter choices in the SOLA method applied to global S-wave tomography lead to different local resolving power and model uncertainties. To this end, we systematically test several combinations of target kernels and trade-off parameters spanning the range from low-resolution to high-resolution inversions, each with varying degrees of resulting model uncertainty. Second, in order to effectively analyse the results from different inversion parameter combinations, we develop a combined analysis of the resolution length scales in 3-D together with a specific test of the adequateness of the method for estimating these lengths. This allows us to represent the practically relevant information on resolution in the SOLA averaging kernels on a global scale, which can then be inspected alongside the uncertainty propagating into the model.

We start with a brief review of the SOLA method and describe the general tomographic system that we employ in Section 2 and Appendix A. Computational aspects regarding the efficient solution of the linear SOLA system are described in Appendix B. Section 3 provides examples of typical SOLA averaging kernels for different inversion choices and motivates the development of a strategy for estimating resolution lengths with a Gaussian approximation in Section 4. Important for this analysis will be to test this Gaussian approximation, for which we introduce the concept of 'focus' that allows us to define different quality categories for the averaging kernels. In Section 5, we provide global maps of tomographic resolution lengths in specific but globally coherent directions, estimated in a robust manner using the combined concepts of resolution and focus, together with the propagated uncertainty. Finally, we discuss the impact of the different possible choices of inversion parameters in the SOLA method in light of possible optimal design towards practical applications.

2 TOMOGRAPHIC METHODOLOGY

2.1 Review of the SOLA Backus–Gilbert method

The main insight of Backus–Gilbert theory is relatively straightforward: given a finite amount of data, one can generally not retrieve exact point estimates of the Earth model parameters $m(\mathbf{r})$ of interest. Nevertheless, it is often possible to infer at a model target location $\mathbf{r}^{(k)}$ a unique weighted average $\hat{m}^{(k)}$, such that

$$\hat{m}^{(k)} = \int_V A^{(k)}(\mathbf{r})\, m(\mathbf{r})\; d^3\mathbf{r}\,, \tag{1}$$

where $A^{(k)}(\mathbf{r})$ is the averaging or resolving kernel (Backus & Gilbert 1968, 1970). Classically, the objective is to obtain an optimally localized averaging kernel that approximates a delta peak at $\mathbf{r}^{(k)}$, constructed from a linear combination of N data(-sensitivity) kernels $K_i(\mathbf{r})$. In the linear case, the data kernels $K_i(\mathbf{r})$ relate model parameters $m(\mathbf{r})$ to the measured data d_i in the form of

$$d_i = \int_V K_i(\mathbf{r})\, m(\mathbf{r})\, d^3\mathbf{r} + n_i\,, \qquad i = 1, ..., N\,, \tag{2}$$

where the data d_i include an error (or noise) component n_i that is assumed here to be independent and normally distributed with zero mean and variance $\sigma^2_{d_i}$. To compute global sets of averaging kernels $A^{(k)}(\mathbf{r})$, we employ the SOLA method (Zaroli 2016) that solves the following optimization problem:

$$\min_{\mathbf{x}^{(k)}} \int_V \left(A^{(k)}(\mathbf{r}) - T^{(k)}(\mathbf{r})\right)^2 d^3\mathbf{r} \;+\; \eta^2\; \sigma^2_{\hat{m}^{(k)}} \tag{3}$$

$$\text{subject to} \quad \int_V A^{(k)}(\mathbf{r})\, d^3\mathbf{r} = 1\,, \tag{4}$$

where $T^{(k)}(\mathbf{r})$ is a target (resolution) kernel, η the trade-off parameter and $\sigma^2_{\hat{m}^{(k)}}$ the model variance from error propagation. The solution of eq. (3)+(4) yields a set of coefficients $\mathbf{x}^{(k)}$ that can be interpreted as (the k-th row of) a generalized inverse operator and determines the estimated average $\hat{m}^{(k)}$, the averaging kernel $A^{(k)}(\mathbf{r})$ and the model uncertainty from error propagation $\sigma_{\hat{m}^{(k)}}$ (plus covariance if desired):

$$\mathbf{x}^{(k)} \Longrightarrow \begin{cases} \sum_{i=1}^N x_i^{(k)} d_i & \longrightarrow \quad \hat{m}^{(k)} \\ \sum_{i=1}^N x_i^{(k)} K_i(\mathbf{r}) & \longrightarrow \quad A^{(k)}(\mathbf{r}) \\ (\sum_{i=1}^N (x_i^{(k)})^2 (\sigma_{d_i})^2)^{1/2} & \longrightarrow \quad \sigma_{\hat{m}_k}\,. \end{cases} \tag{5}$$

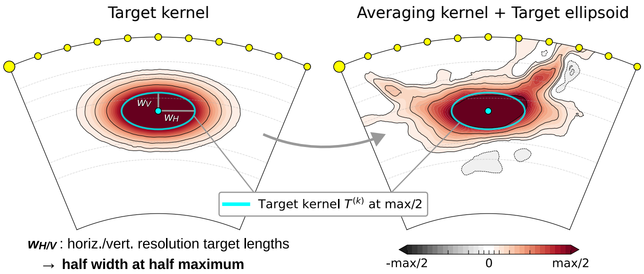

Figure 1. Left: Cross-section of a target kernel with horizontal and vertical half widths $w_{\mathrm{H/V}} = 600/300$ km. The cyan ellipse marks the contour line at half maximum. Right: Cross-section for an averaging kernel at the same target location. To get an idea of the fit to the target kernel, we also plot the target ellipse at half maximum.

The unimodular constraint in eq. (4), which is also implied by the classical Backus–Gilbert theory, ensures that $\hat{m}^{(k)}$ represents a physical volumetric average. In the hypothetical presence of a constant model value around the target location $\mathbf{r}^{(k)}$, the estimated model amplitude would thus not be artificially scaled. Most important, the introduction of a target kernel $T^{(k)}(\mathbf{r})$ in the SOLA method means that, instead of an ideal delta peak, as in the original Backus–Gilbert formulation, a practically more relevant spatial function is underlying the construction of the averaging kernels. Different choices of $T^{(k)}(\mathbf{r})$ then allow us to use information on the expected local resolution by varying the shape and weighting present in the target kernels. At the same time, the trade-off parameter η ensures that the fit to the target kernel and the variance of the propagated errors in the inferred averages can be controlled. Both $T^{(k)}(\mathbf{r})$ and η may be selected with a subjective preference, however with the benefit that the inversion results include complete information on resolution and uncertainty. Note also that each choice of η and $T^{(k)}(\mathbf{r})$ has a specific impact on all local model properties (for details on the SOLA inversions employed here, including computational aspects to solve the system efficiently for all model parameters, see Appendix A and B).

2.2 Resolution length and choice of target kernels

To obtain robust resolution information from the SOLA averaging kernels, Pijpers & Thompson (1994) calibrated a practical threshold value based on visual inspection of their target fit. This way, they were able to distinguish well-localized from non-localized 1-D averaging kernels and could use the prescribed target kernel peak width to infer resolution lengths. This is however not directly applicable in our case. In 3-D tomography, the fit of $A^{(k)}$ to the target resolution might be good in a specific direction and poor in any other one. Therefore, no simple choice of target kernel $T^{(k)}$ and trade-off parameter η is reducing equally well the kernel difference in all regions. Instead, additional information on possible shifts and varying length scales in different directions is required to assess the averaging kernels (see Section 4). However, with the choice of a specific target $T^{(k)}$, one can still promote a desired shape and size of the $A^{(k)}$ prior to the inversion.

For this purpose, an intuitive parametrization for the target kernel $T^{(k)}$ is given by 3-D Gaussian functions that have been used previously in seismology for analysing tomographic resolution (see e.g. Fichtner & Trampert 2011; An 2012). The kernels are centred at the target locations $\mathbf{r}^{(k)}$, which correspond to the points in a tomographic grid (see Section 2.3) By using a local Cartesian frame with origin at $\mathbf{r}'^{(k)}$ (the prime indicates the change of basis), we can align the principal axes of the Gaussian function along the horizontal and vertical directions; i.e. they are oriented along tangents in the latitudinal and longitudinal direction and along the radius in the vertical direction, with respect to $\mathbf{r}'^{(k)}$. Our 3-D Gaussian target kernels are then given by

$$T^{(k)}(x', y', z') = \frac{a^3}{\sqrt{(2\pi)^3 \cdot w_{x'} \cdot w_{y'} \cdot w_{z'}}} \times \exp\Big[-\frac{a^2}{2}\left(\frac{x'^2}{w_{x'}^2} + \frac{y'^2}{w_{y'}^2} + \frac{z'^2}{w_{z'}^2}$$

where $w_{x',y',z'}$ are the half widths at half maximum and x', y' and z' refer to the axes in the local basis. Examples for a target kernel $T^{(k)}$ and a resulting averaging kernel $A^{(k)}$ are visualized in Fig. 1. We specify the $w_{x',y'}$ in horizontal (w_{H}) and $w_{z'}$ in vertical direction (w_{V}) of $T^{(k)}$ as target resolution lengths. As a remark, their relation to the standard deviation is $w = a \cdot \sigma$, where $a = \sqrt{2\ln(2)} \approx 1.17$. Using w is particularly useful because it allows one to relate the kernel width to its maximum at the peak, which facilitates comparisons of the volumetric change of various kernels with different peak amplitudes. We directly evaluate eq. (8) on the discrete tomographic grid (see Section 2.3). In contrast to Zaroli (2016), the $T^{(k)}$ are here not strictly normalized, i.e. we do not enforce $\int_V T^{(k)}(\mathbf{r})\, d^3\mathbf{r} = 1$. This is done deliberately in order to preserve the Gaussian shape of $T^{(k)}$ within the finite volume V of the grid.

For the SOLA inversions in this study, we take 3 variations of $T^{(k)}$ using Gaussian functions of progressively larger horizontal and vertical target resolution lengths ($w_{\mathrm{H/V}} = 300/200$, $600/300$, $900/400$ km). Along with this, we test 3 different values for the trade-off parameter η ($\eta_1 = 5$, $\eta_2 = 10$, $\eta_3 = 30$; see also Table 1). Increasing values for η generally promote smaller model uncertainty $\sigma_{\hat{m}^{(k)}}$ while deteriorating the fit of the averaging kernel $A^{(k)}$ to $T^{(k)}$. The range of values of η was chosen empirically from a few SOLA inversions for all $T^{(k)}$ in order to cover a range of tomographically relevant levels for $\sigma_{\hat{m}^{(k)}}$, resolution length scales and target fits.

2.3 Tomographic grid and dataset

Since we are interested in quantifying the impact of different inversion parameter choices, we choose a global grid that can well represent the shape of averaging kernels down to all target length scales; that is, it covers at least the smallest target half widths $w_{\mathrm{H/V}} = 300/200$ km used here. In general, we follow the parametrization strategy of Zaroli (2016), where grid nodes are the upper vertices of triangular prisms based on a spherical Delaunay triangulation for several distinct depth layers across the entire mantle. For details, the reader is referred to Zaroli (2010, 2016). In radial direction, we take the 18 depth layers from SOLA-Z16 (Zaroli 2016) that are between 100-200 km thick, whereas in lateral direction, we use approximately equidistant spherical Fibonacci

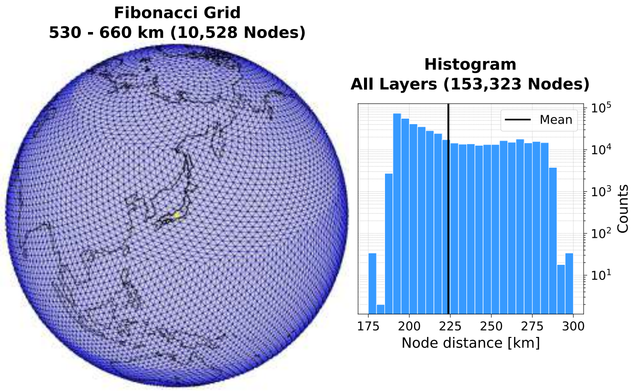

Figure 2. Left: Fibonacci grid on a sphere for the 530-660 km depth layer using the algorithm by Swinbank & Purser (2006). The nodes are the upper vertices of spherical triangular prisms that are constructed by Delaunay triangulation. Right: Histogram of internode distances for all 18 depth layers. The total amount of unique lateral node connections (i.e. the number of triangle edges) covered in the histogram is 447,567. The average distance is ∼224 km.

grids following Swinbank & Purser (2006). In Fig. 2 the grid nodes of the layer at 530-660 km depth are shown as an example. To create a suitable realization of the Fibonacci grid, we empirically determine a specific amount of points for each layer such that the maximum distance of neighbours from the spherical Delaunay triangulation is less than 300 km. The distribution of the node distances across all layers is shown in a histogram in Fig. 2. Minimum and average neighbour distances are about 176 km and 224 km, respectively. In total, the grid includes 153,323 grid nodes.

As tomographic dataset, we use the source-receiver configurations from the SOLA-Z16 model (Zaroli 2016) that include 79,765 cross-correlation traveltime measurements for S and SS seismic phases at 22.5 s central period (Zaroli 2010). Sensitivities to shear-wave velocity perturbations are calculated using paraxial finite-frequency kernels (Dahlen et al. 2000). The coverage of the dataset is particularly suitable for investigating the velocity structure at depths between 400 and 1710 km (Zaroli 2016). However, there are mainly two relevant regions that have different characteristic behaviour for resolution. At depths of ≈400–810 km the finite-frequency kernels for teleseismic S-waves have not bottomed out yet, thus finding suitable linear combinations of the data to enhance resolution locally can be more difficult there. In contrast, at depths below ≈810 km, larger horizontal volumes are covered by the finite-frequency kernels, leading to a higher chance for crossing volumes of sensitivity. To effectively investigate the resulting averaging kernels, we therefore focus on two particular depth layers in this study (at 530–660 km, midpoint at 595 km and 1110– 1310 km, midpoint at 1210 km depth) representing each situation.

3 SOLA AVERAGING KERNELS

To first get a rough appreciation of the behaviour of the SOLA averaging kernels, we visually inspect some examples before we proceed with the global analysis of all kernels in Section 4. To this end, we show kernel cross-sections for the different combinations (see Table 1) of trade-off parameters η and target kernels $T^{(k)}$ that we introduced in Section 2.2.

Table 1. SOLA inversion parameters used in this study. To facilitate discussion, a high resolution(HR) and a low resolution (LR) setup are introduced. Horizontal and vertical target resolutions $w_{\mathrm{H/V}}$ correspond to the half widths at half maximum of the Gaussian target kernels (see eq. 6).

Target resolution $w_{\mathrm{H/V}}$ (horizontal/vertical)	Trade-off parameter η
300/200 km	$\eta_1 = 5$
600/300 km	$\eta_2 = 10$
900/400 km	$\eta_3 = 30$

High resolution (HR) setup : 300/200 km, η_1
Low resolution (LR) setup: 900/400 km, η_3

3.1 Influence of data coverage

A compilation of different averaging kernels is shown in Fig. 3 for two end-member combinations of $T^{(k)}$ and η. They represent a high-resolution (HR) inversion with η_1 (= 5) and a target kernel size of $w_{\mathrm{H/V}} = 300/200$ km (Figs 3a and d), and the contrary case of a low-resolution (LR) inversion using η_3 (= 30) and a target size of $w_{\mathrm{H/V}} = 900/400$ km (Figs 3b and c). Dashed ellipses defined by horizontal and vertical semi-axes w_{H} and w_{V} indicate the size of $T^{(k)}$ for the particular inversion (as shown in Fig. 1). Data coverage is particularly good in the Northern Hemisphere (top row in Fig. 3), where one can see that both HR and LR inversion setups lead to averaging kernels that are overall well localized and fit the shape of the target kernels. In Fig. 3(b) using the LR setup for an averaging kernel centred below North America, one can still observe stronger positive amplitudes south of Hawaii outside the broader target region. Poor data coverage in the Southern Hemisphere (bottom row in Fig. 3) typically leads to increased power in side lobes, reflecting the difficulty of fitting the target kernel with a locally incomplete dataset. For example, the averaging kernel in Fig. 3(c) for the LR setup is hardly centred and apparently dominated by a subset of unidirectional data kernels reaching to the surface between Kerguelen Islands and Australia. On the other hand, the $A^{(k)}$ in Fig. 3(d) combines poor data coverage with the HR setup, leading again to more pronounced side lobes, but also to a good fit to the target for the bulk of the averaging sensitivity. Fig. 3 already suggests that the target kernel has a strong control on the resolution properties of $A^{(k)}$ (as expected), and that poor data coverage, although making the averaging kernels more prone to oscillatory behaviour, is not necessarily preventing one from fitting the target resolution. Although being barely visible here, negative values in the averaging kernels do exist, but their amplitudes are mainly located outside the target region and are generally small. A more quantitative analysis taking this into account is presented in the classification scheme that we develop in Section 4.2.

3.2 Variable target kernel and trade-off parameter

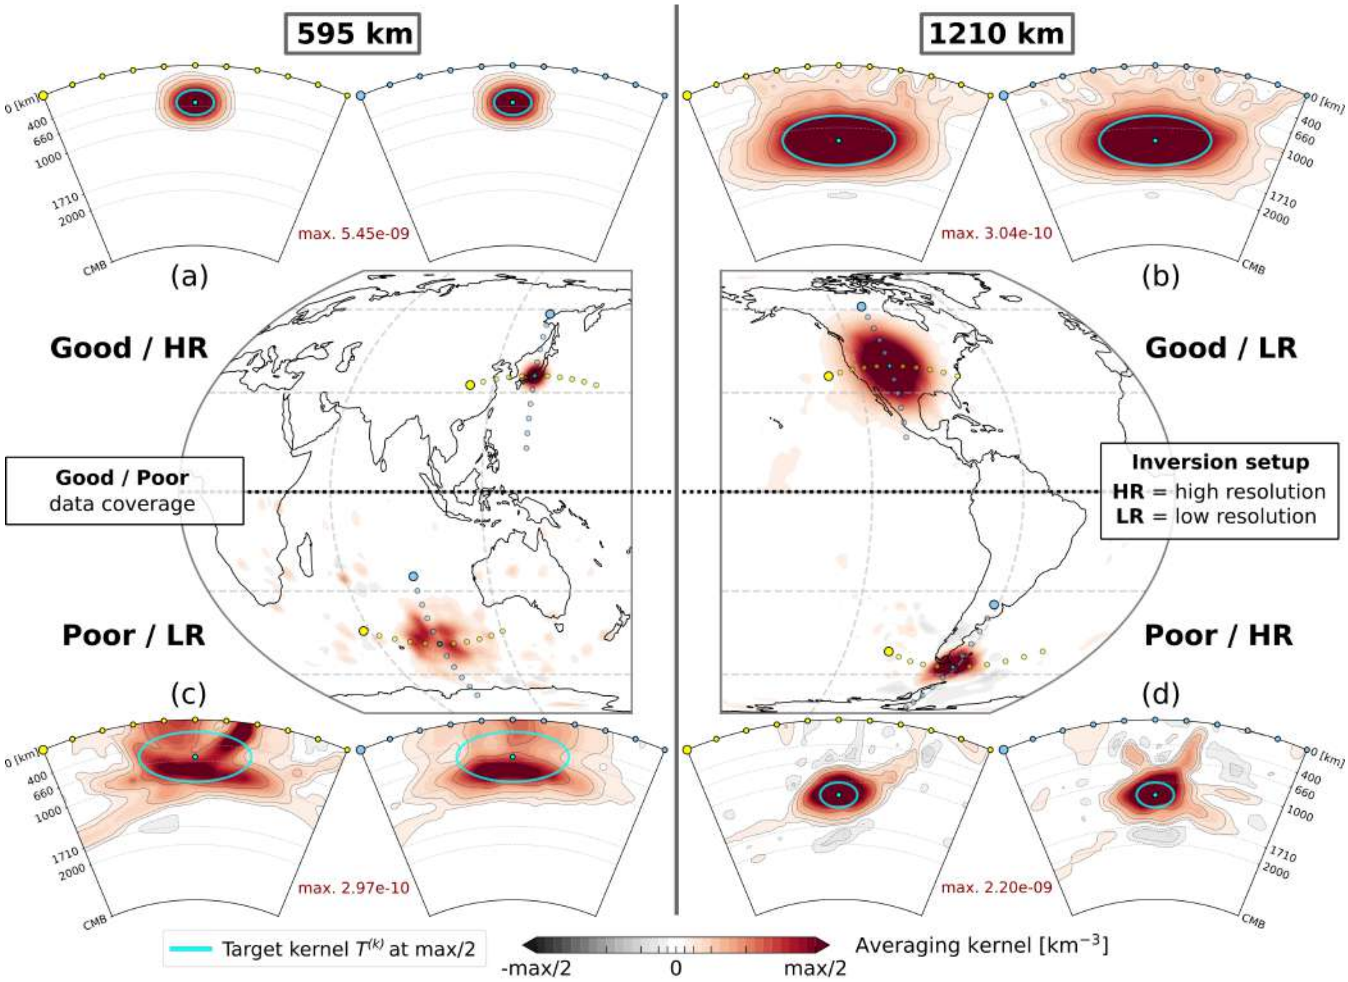

Figure 3. Lateral and vertical cross-sections for four different averaging kernels in regions of good and poor data coverage in the Northern and Southern Hemisphere, respectively. Left and right panels show target locations at 595 km and 1210 km depth, respectively. a) and d): 'High-resolution' setup (HR), $A^{(k)}$ for a trade-off parameter of $\eta_1 = 5$ and target kernel $T^{(k)}$ with horizontal/vertical extent of $w_{\mathrm{H/V}}$ = 300/200 km. b) and c): 'Low-resolution setup (LR)', $A^{(k)}$ for $\eta_3 = 30$ and $T^{(k)}$ with $w_{\mathrm{H/V}}$ = 900/400 km. The low-resolution setup leads to a preference of lower uncertainty $\sigma_{\hat{m}^{(k)}}$ by using a larger trade-off parameter η_3 at the cost of a worse fit to the target resolution. The latter is marked by dashed ellipses (drawn at the half widths at half maximum $w_{\mathrm{H/V}}$ of each $T^{(k)}$). Values of the averaging kernels are normalized with their respective maximum. Because our tomographic grid is fine enough for the target length scales, it follows that for an almost ideal fit, the darkest red values at the cut-off of the colour scale would all be within the target ellipses. Black contour lines in the vertical cross-sections further indicate iso-surfaces at 1/2, 1/4, 1/8, 1/16 and 1/32 of the peak value.

While Fig. 3 visualizes the consequences of variable data coverage, we can also systematically exploit the capabilities of the SOLA method and investigate how different choices of target kernel and trade-off parameter affect $A^{(k)}$ and the propagated uncertainty $\sigma_{\hat{m}^{(k)}}$. Fig. 4 features such a set of alternatively possible averaging kernels, corresponding to the target location and the low-resolution example of Fig. 3(c) in the Southern Hemisphere. Here, we focus on vertical cross-sections in west-east direction; complementary figures for all cross-sections and kernels can be found in the online supplementary material. The different averaging kernels are organized in a matrix layout, going from smaller to larger target kernel sizes from top to bottom and increasing values for the trade-off parameter η from left to right.

An increase in η leads to an increase in model uncertainty $\sigma_{\hat{m}^{(k)}}$, which is indicated in each cross-section in Fig. 4, ranging from 0.47 % for the HR setup, down to 0.03 % $\mathrm{dln}(v_S)$ for the LR setup. Likewise, at constant η and growing target kernel sizes, model uncertainties also decrease. None of the averaging kernels exactly attains the maximum amplitude of their underlying $T^{(k)}$. At a target size of $w_{\mathrm{H/V}} = 300/200$ km, for example, the maxima of the corresponding $A^{(k)}$ are off by more than a factor of 2. For larger target lengths, this difference is less severe. Overall, an unsatisfactory visual fit of the averaging kernels to their respective target kernels is observed, which can be ascribed here to a single finite-frequency kernel that seems to dominate the averaging. A different situation is shown in Fig. 5. At this location, the data coverage is excellent, which leads to a suite of averaging kernels that are Gaussian-shaped for all parameter combinations tested. Peak amplitudes of the averaging kernels for the lowest value of the trade-off parameter are close to their respective target kernel value and only moderately drop as η is increased. Even at lower

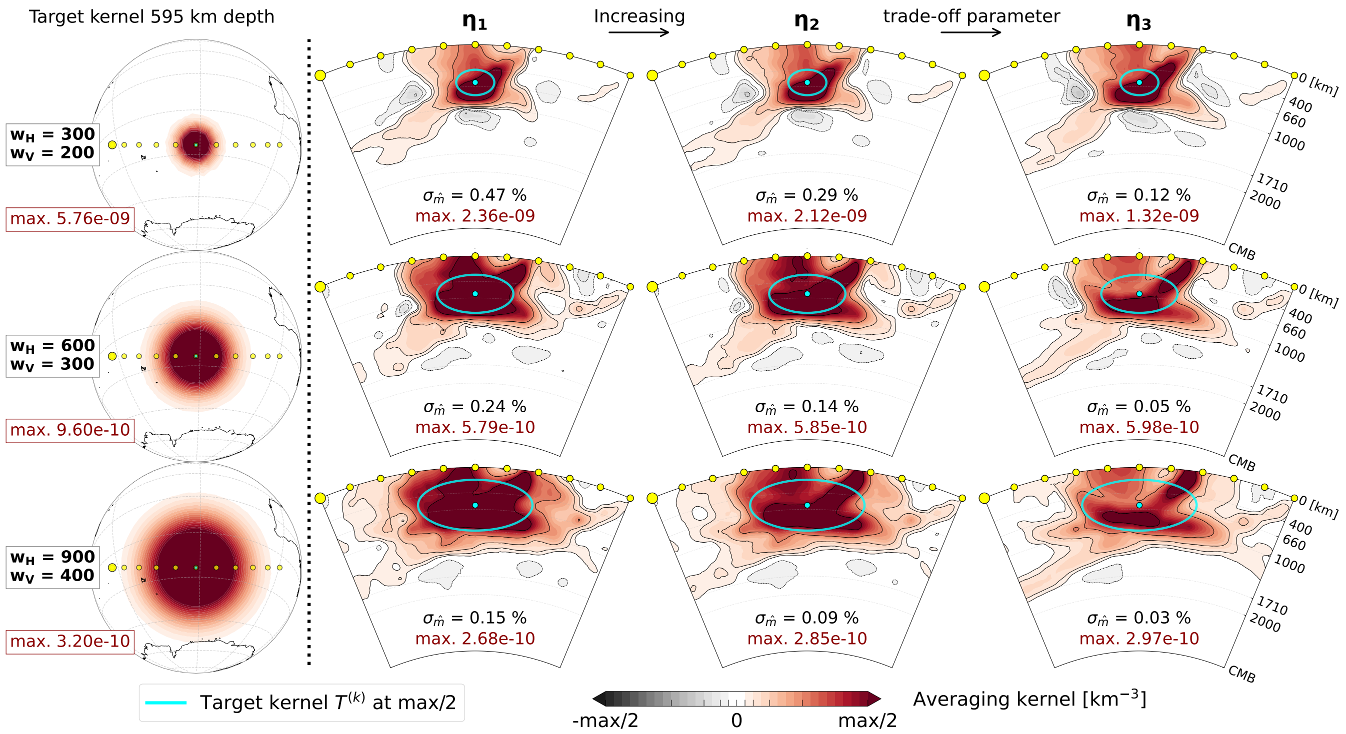

Figure 4. Influence of increasing the trade-off parameter and target kernel size on the averaging kernels for a region of relatively poor data coverage in the southern Indian Ocean, east of the Kerguelen Islands at 595 km depth. Left: lateral cross-sections of the target kernel $T^{(k)}$ with horizontal and vertical extent w_{H} and w_{V}, respectively, used for the inversion. Right: vertical cross-sections (West-East) of the resulting averaging kernels $A^{(k)}$ for all combinations of η ($\eta_{1,2,3}$ = 5, 10, 30) and $T^{(k)}$. Dashed ellipses indicate the size of $T^{(k)}$ with semi-axes of length w_{H} and w_{V}. The kernel values are normalized with their respective half maximum to facilitate comparisons. Model uncertainties $\sigma_{\hat{m}^{(k)}}$ are given in % dln(v_S).

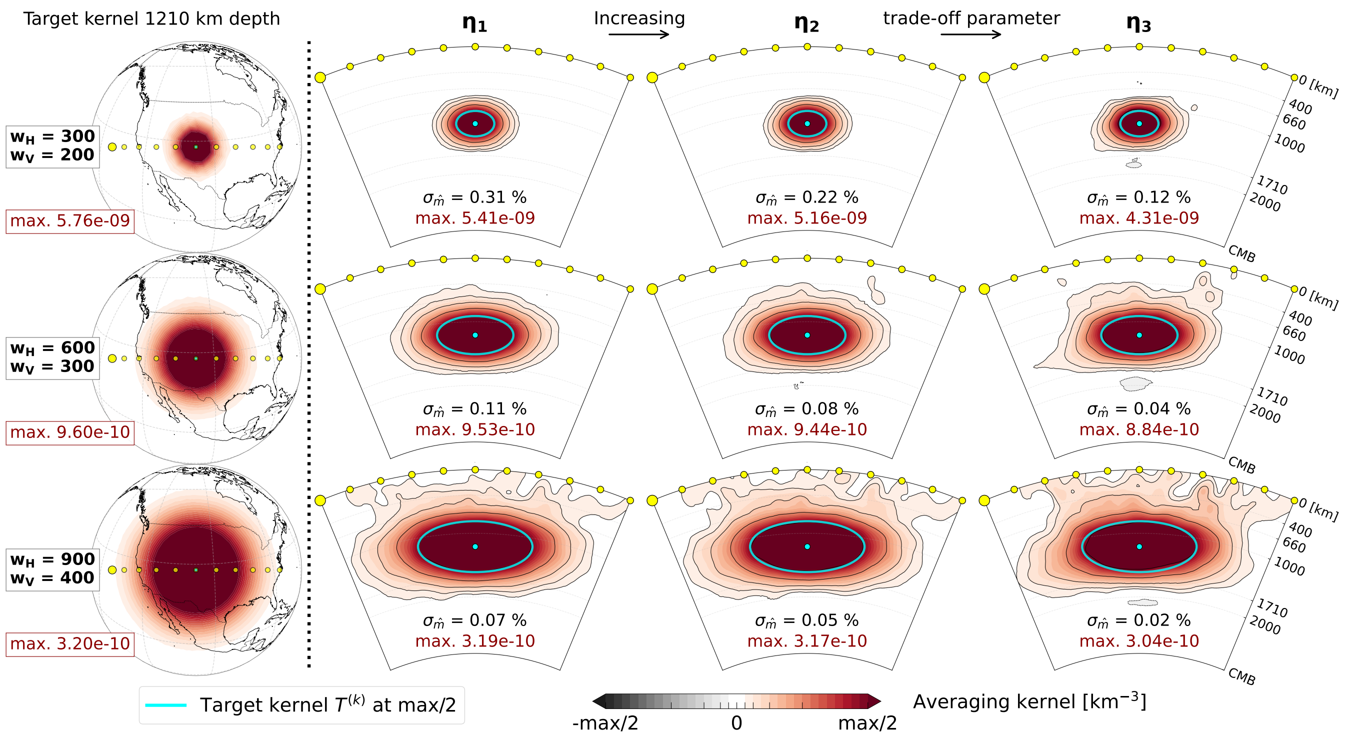

Figure 5. Influence of increasing the trade-off parameter and target kernel size for a region of excellent data coverage in the central US at 1210 km depth. Same layout as Fig. 4.

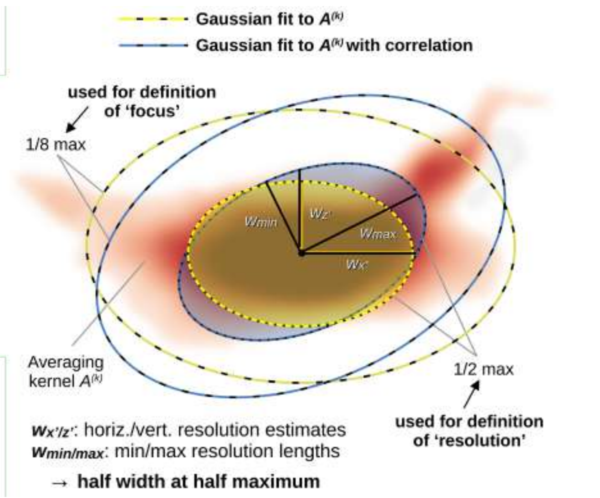

Figure 6. Concept for our resolution analysis and definition of the 'focus' for an averaging kernel $A^{(k)}$ parametrized with a 3-D Gaussian. Yellow and blue indicate an uncorrelated and correlated version of the Gaussian, respectively. Filled ellipses at the half widths at half maximum ($w_{x'/y'/z'}$) define estimated resolution lengths, while the larger, unfilled ellipses at the half width at one-eighth maximum define the region used for computing the focus ξ (see Section 4.2). In contrast to the simpler, uncorrelated Gaussian $\hat{g}^{(k)}$, the Gaussian $\tilde{g}^{(k)}$ with correlation includes rotation with respect to the axes of the local Cartesian frame centred at the target location. This could be used for estimating minimum/maximum resolution lengths $w_{\min/\max}$, which however, would not lie along the same directions for all kernels.

model uncertainties $\sigma_{\hat{m}^{(k)}}$, the kernels also mostly retain localization and recover the target shape.
It is clear that a visual inspection of the averaging kernels might be insufficient to judge on their quality and cannot be performed individually for every model parameter. Also, it does not provide information on the length scales of resolution that need to be quantified for a more thorough analysis. This raises the question of how one could assess the quality and nature of the $A^{(k)}$ in a consistent manner, such that sensible estimates and comparisons of resolution lengths can at best be made for the entire set of averaging kernels all together.

4 QUANTIFYING RESOLUTION

A general problem for inferring meaningful information on resolution in seismic tomography is that one cannot guarantee that every individual averaging kernel is localized and reasonably Gaussian-shaped for a given selection of inversion parameters. Therefore, we cannot reliably use the prescribed size of the target kernels $T^{(k)}$ to examine resolution on a global scale.Instead, resolution lengths should be quantified consistently for all averaging kernels and independent of the respective target form of $T^{(k)}$. In addition, it needs to be determined whether the resulting length estimates are meaningful in the given context, while bearing in mind the possible complexity of the averaging kernels. A robust quantification and interpretation of the resolution information contained in tomographic inversions thus requires two individual tools: the concept of resolution lengths together with a classification scheme for the quality of the averaging kernels.

4.1 Gaussian estimates and resolution lengths

Owing to our choice of target kernels in the form of 3-D Gaussian functions, it is useful to independently determine also a best-fitting 3-D Gaussian for each averaging kernel in order to quantify resolution lengths. A general 3-D Gaussian parametrization is given by (see also Fichtner & Trampert 2011)

$$\tilde{g}^{(k)}(\mathbf{r}') = \frac{N^* a^3}{\sqrt{(2\pi)^3 \det \mathbf{C}}} \times \exp\left[-\frac{a^2}{2}(\mathbf{r}'-\boldsymbol{\mu}')^{\mathrm{T}}\mathbf{C}^{-1}(\mathbf{r}'-\boldsymbol{\mu}')\right], \tag{7}$$

where the position $\mathbf{r}'$ and mean vector $\boldsymbol{\mu}'$ are defined in the same local Cartesian coordinate system used for the target kernel $T^{(k)}$ in eq. (6). Any non-zero mean location $\boldsymbol{\mu}'$ implies that the centre of mass of the averaging kernel is displaced from the target location (this is called 'distortion' by Fichtner & Trampert 2011). The scaling factor N^* represents the total mass of the Gaussian. We include N^* to facilitate finding suitable resolution length estimates (given by the half widths $w_{x',y',z'}$) that better reflect the averaging volume around the main peak of $A^{(k)}$. The covariance matrix $\mathbf{C}$ describes the spatial extent and rotation/tilting of the Gaussian in the local framework and incorporates the half widths $w_{x',y',z'}$ and parameters for correlation $\rho_{x'y'}$, $\rho_{x'z'}$ and $\rho_{y'z'}$. The complete set of parameters for estimation would thus be given by $\tilde{\mathbf{p}} = (N^*, \mu_{x'}, \mu_{y'}, \mu_{z'}, w_{x'}, w_{y'}, w_{z'}, \rho_{x'y'}, \rho_{x'z'}, \rho_{y'z'})$. By using the full set of correlation parameters, one could therefore also extract from the averaging kernels minimum/maximum resolution lengths in any necessary direction.

For global comparisons of averaging length scales with respect to the uncorrelated target kernels that we have chosen, it might in fact be simpler for interpretation, but not less informative, to focus only on specific directions. To this end, we can neglect the correlations and simplify eq. (7) to get a Gaussian parametrization in the following form:

$$\hat{g}^{(k)}(x', y', z') = \frac{N^* a^3}{\sqrt{(2\pi)^3 \cdot w_{x'} \cdot w_{y'} \cdot w_{z'}}} \times \exp\left[-\frac{a^2}{2}\left(\frac{(x'-\mu_{x'})^2}{w_{x'}^2} + \frac{(y'-\mu_{y'})^2}{w_{y'}^2} + \frac{(z'-\mu_{z'})^2}{w_{z'}^2}\right)\right]. \tag{8}$$

The reduced set of parameters for estimation is given by $\mathbf{p} = (N^*, \mu_{x'}, \mu_{y'}, \mu_{z'}, w_{x'}, w_{y'}, w_{z'})$, from which we obtain information on the shift of the centre of mass $\boldsymbol{\mu}'$ of $A^{(k)}$ away from the target location and resolution lengths $w_{x',y',z'}$ in the three directions defined by the local Cartesian frame. Fitting is performed by using the Levenberg—Marquardt algorithm for solving the weighted least-squares problem

$$\arg\min_{\mathbf{p}} \sum_j \mathcal{V}_j \left(A_j^{(k)} - \hat{g}^{(k)}(\mathbf{r}'_j; \mathbf{p})\right)^2, \tag{9}$$

using the volume $\mathcal{V}_j$ associated with each j-th grid node. For a practical initial guess, we use the specifications of the corresponding target kernel $T^{(k)}$ at the location of interest. As a remark, other parametrizations are possible that could approximate even more accurately the shape of averaging kernels. Fichtner & Trampert (2011) pointed out that Gram-Charlier expansions can be used to that end. These make it possible to approximate functions, or more strictly distributions, like the averaging kernels from a series of their cumulants (e.g. mean, variance, third central moment and more complicated quantities at higher orders). While this can lead to more accurate approximations of the complete function, Fichtner & Trampert (2011) suggested that from a practical point of view, these expansions might not always offer more physically interpretable information than the plain Gaussian approximation itself. However, these and other suitable parametrizations may be required for analysing functions that are vastly different across the model domain, such as kernels for the trade-off between different physical model parameters (which is not the case for the $A^{(k)}$ with Gaussian targets $T^{(k)}$ in this study).

The difference between the two Gaussian parametrizations $\hat{g}^{(k)}$ and $\tilde{g}^{(k)}$ is visualized schematically in Fig. 6. The semi-axes of the ellipse at half the maximum are the estimated resolution lengths $w_{x',z'}$ ($+w_{y'}$ in the actual 3-D ellipsoid), that can be compared to the target widths $w_{\mathrm{H/V}}$. These estimates, as well as the remaining parameters in $\mathbf{p}$, may differ slightly between the two parametrizations. The largest mean absolute deviation between $\hat{g}^{(k)}$ and $\tilde{g}^{(k)}$ we found for any length estimate and given parameter combination is $\sim$16 km. However, in case of kernels with reasonably good target fits these differences are small, and we noticed that on global scales, the choice between uncorrelated and correlated Gaussian has only a minor effect (see supporting material). A remaining problem is the identification of complex averaging kernels that cannot be well approximated by either of the Gaussian parametrizations. We therefore propose a practical strategy to deal with this issue in the following section.

4.2 Kernel classification based on the 'focus'

For a robust interpretation of the resolution information, one needs to know for every target location whether the Gaussian-based estimates of resolution length actually relate well to the shape and mass of the averaging kernel (i.e. the integral over $A^{(k)}$ for a given volume). Since the centre of $\hat{g}^{(k)}$ is an estimate of the centre of mass in $A^{(k)}$, the (main) peak regions of both functions should by design overlap. If we can measure, by comparison of the masses of the two kernels, if this is actually true for a given target location, it would allow us to quickly identify for which choices of the trade-off parameter η and target resolution the resulting averaging kernels deviate significantly from a Gaussian.

To get a diagnostic tool for this specific problem, we develop in the following a global classification scheme for the averaging kernels, measuring their quality with respect to the best-fitting Gaussian $\hat{g}^{(k)}$. First, it is necessary to define the actual volume of the peak region of $\hat{g}^{(k)}$, for which we can draw inspiration from the simpler 1-D case. For a 1-D Gaussian function, 76 per cent of the total area lies within $\pm w$, the half width at half maximum (our measure for resolution length), around the mean. In higher dimensions, however, this well-known concept does not hold. In fact, the integrated mass of $\hat{g}^{(k)}$ within the 3-D ellipsoid given by the three half widths at half maximum $w_{x',y',z'}$ (respectively the inner ellipse in Fig. 6) defined by

$$\left\{ \left(\frac{x'-\mu_{x'}}{w_{x'}}\right)^2 + \left(\frac{y'-\mu_{y'}}{w_{y'}}\right)^2 + \left(\frac{z'-\mu_{z'}}{w_{z'}}\right)^2 \leq 1 \right\} \quad (10)$$

covers merely $\approx$29 per cent (if $w_{x'} = w_{y'} = w_{z'}$) of the total mass of the Gaussian. The exact value also depends on the specific half widths in each direction, and thus may vary for different $\hat{g}^{(k)}$. Since the SOLA kernels determine average values in volumetric regions, we consider 29 per cent of the total mass to be insufficient to properly describe the characteristics of $A^{(k)}$. Therefore, we aim to reproduce the 1-D convention, with roughly 76 per cent of the total mass, for the 3-D situation here as well. This can be achieved by scaling up the ellipsoid in eq. (10) and replacing the $w_{x',y',z'}$ with the corresponding half widths at one-eighth maximum from the best-fitting Gaussian $\hat{g}^{(k)}$ (the outer ellipses in Fig. 6). In practice, we determine the ellipsoid by evaluating $\hat{g}^{(k)}$ directly on the tomographic grid and including all nodes that hold values larger than one-eighth of the maximum (taken from the continuous function). The absolute mass $\hat{g}^{(k)}_{in}$ inside this spatial domain $E^{(k)}$ then approximately represents 76 per cent of the total mass ($\hat{g}^{(k)}_{in} + \hat{g}^{(k)}_{out}$). The exact numeric value, however, may still differ by a few per cent of the total mass depending on the size and location as well as on the error from evaluation on a discrete grid. Having formulated an expectation on the volume and mass $\hat{g}^{(k)}_{in}$ of the peak region, we can then separate the mass contribution of the associated averaging kernel $A^{(k)}$ within and outside the ellipsoid $E^{(k)}$:

$$\underbrace{\int_{\mathbf{r}'\in E^{(k)}} A^{(k)}(\mathbf{r})\ d^3\mathbf{r}}_{A^{(k)}_{in}} + \underbrace{\int_{\mathbf{r}'\notin E^{(k)}} A^{(k)}(\mathbf{r})\ d^3\mathbf{r}}_{A^{(k)}_{out}} = 1\,. \quad (11)$$

Note again that the total mass of $A^{(k)}$ is equal to 1 in the SOLA method owing to the unimodular condition in eq. (4). Generally, both $A^{(k)}_{in}$ and $A^{(k)}_{out}$ include positive as well as negative kernel values. However, the negative contributions to $A^{(k)}_{in}$ were found to not exceed 2–3 per cent for any averaging kernel that we computed. For more than $\sim$80 per cent of all averaging kernels, negative values in $A^{(k)}_{in}$ are $\leq$1 per cent and vanished completely for $\sim$8-25 per cent (depending on the specific $T^{(k)}$ and η). We then define the 'focus' ξ of the averaging kernel based on the mass ratio of $A^{(k)}_{in}$ and $\hat{g}^{(k)}_{in}$ as

$$\xi = \left(\frac{A^{(k)}_{in}}{A^{(k)}_{in} + A^{(k)}_{out}}\right) \Big/ \left(\frac{\hat{g}^{(k)}_{in}}{\hat{g}^{(k)}_{in} + \hat{g}^{(k)}_{out}}\right) = \frac{A^{(k)}_{in}(\hat{g}^{(k)}_{in} + \hat{g}^{(k)}_{out})}{\hat{g}^{(k)}_{in}}\,. \quad (12)$$

The normalization with the total mass ($\hat{g}^{(k)}_{in} + \hat{g}^{(k)}_{out}$) of $\hat{g}^{(k)}$ within the model domain allows one to take into account possible errors through discretization (for the absolute values of $\hat{g}^{(k)}_{in}$) and ellipsoids $E^{(k)}$ that are intersected by the surface. Using these relative mass contributions is helpful because it makes ξ a uniform metric for all $A^{(k)}$. In case that the unimodular condition for $A^{(k)}$ (eq. 4) is not fulfilled, e.g. when using other kernels that are not derived by a Backus–Gilbert type inversion, one should modify the focus ξ and normalize $A^{(k)}_{in}$ also with the total mass of the averaging kernel (i.e. one should use the more general expression of ξ in eq. 12). To finally classify the quality of the averaging kernels, we define 5 categories for ξ:

- $\xi < 0.5$, *'Not Focused'*:

$A^{(k)}$ is hardly focused and likely degraded by multiple peaks or strong side lobes, often due to either individual finite-frequency kernels dominating the averages or strong sensitivity in the upper mantle;

- $\xi \geq 0.75$, *'Sufficient (Suff.)'*:

Threshold between acceptable and unacceptable kernel fit with respect to $\hat{g}^{(k)}$; ensuring that all sufficient $A_{in}^{(k)}$ are at least more concentrated inside the Gaussian ellipsoid $E^{(k)}$ rather than outside of it;

- $0.9 \leq \xi < 1.1$, *'Good'*:

Deviation of $A_{in}^{(k)}$ from the ideal case is less than 10 per cent;

- $\xi = 1$, *'Ideal'*:

Relative mass contributions of $A^{(k)}$ and $\hat{g}^{(k)}$ within ellipsoid $E^{(k)}$ are equal;

- $\xi \geq 1.1$, *'Highly Focused'*:

$A^{(k)}$ is more peaked, i.e. has significantly more mass within the ellipsoid $E^{(k)}$ than $\hat{g}^{(k)}$.

This classification with the focus can be seen as a heuristic tool for deciding on the quality of the kernels, since a relative redistribution of the mass of $A^{(k)}$ to outside the peak region is clearly associated with a decrease in ξ. Therefore, although these focus categories do not fully characterize the detailed shape of a specific averaging kernel, they yet provide a basic way to test whether the Gaussian approximation is locally plausible. The classification scheme also gives an indication for the fit to $T^{(k)}$, if $\hat{g}^{(k)}$ is not shifted or broadly stretched beyond the target resolution length. This is because averaging kernels that agree well with $\hat{g}^{(k)}$ tend to be also centred.

4.3 Examples for resolution quantification and classification

Before we apply the previously introduced methods on a global scale (see Section 4.4), we demonstrate more explicitly how they act together to describe resolution. Therefore, instead of averaging kernels that are obviously Gaussian-shaped (see Fig. 5), we consider two less-intuitive examples in Fig. 7. Again, we employ a high-resolution and a low-resolution setup (Figs 7a and 7b, respectively). Unlike the averaging kernels, the functions $\hat{g}^{(k)}$ are defined beyond the mantle domain V, and the associated Gaussian ellipsoids for the focus ξ (at max/8) may thus come close to or extend through the surface. As a consequence, the best-fitting Gaussian $\hat{g}^{(k)}$ in Fig. 7(a) can yield a vertical half width $w_{z'}$ of 322 km that is considerably larger than the target length of $w_{\mathrm{V}} = 200$ km. The estimated resolution lengths in horizontal direction of 370 km and 340 km (North-South and West-East, respectively) are as well somewhat larger than the actual target length of $w_{\mathrm{H}} = 300$ km.

The averaging kernel for the low-resolution case in Fig. 7(b) is even more complex. It exhibits localized smearing of uppermost mantle structure from individual data kernels, and several distinct local maxima in and outside the target region. Resolution lengths estimated by $\hat{g}^{(k)}$ are in horizontal direction 890 km and 864 km (North-South and West-East) and in vertical direction 459 km. They are thus horizontally narrower but vertically larger than the target lengths of $w_{\mathrm{H/V}} = 900/400$ km. There is also a considerable shift of the centre of $\hat{g}^{(k)}$ from the target location of $\boldsymbol{\mu}' = (-66,-97,-106)$ km, which might not be expected by merely inspecting the cross-sections at these azimuths. Overall, it is at least debatable whether the best-fitting Gaussian $\hat{g}^{(k)}$ is an adequate approximation in this case. For the $A^{(k)}$ in Fig. 7(a) we obtain a focus value $\xi = 0.89$ (i.e., certainly 'sufficient' and almost in the range of 'good' kernels), meaning that although the target shape is not fully matched, the best-fitting Gaussian $\hat{g}^{(k)}$ can be used with confidence for describing local resolution lengths. The averaging kernel in Fig. 7(b), however, is classified as insufficient, with a focus value $\xi = 0.62$ suggesting that maxima present outside the ellipsoid $E^{(k)}$ may have a significant impact on the corresponding model estimate for the average value. Further resolution estimates and focus values ξ for the averaging kernels in Fig. 3 can be found in the supplementary material.

4.4 Global classification of resolution

The results of our kernel classification on a global scale, for all previously used combinations of the trade-off parameter η and the target kernels $T^{(k)}$, are visualized in Figs 8 and 9. As expected, we find overall a larger number of Gaussian-shaped averaging kernels at smaller values for η. The percentage of acceptable $A^{(k)}$ ($\xi \geq 0.75$) is shown above each map, varying between 8-79 per cent at 595 km and 34-91 per cent at 1210 km depth. While there is a strict trend of fewer acceptable averaging kernels towards higher η, this is not necessarily the case when increasing the size of $T^{(k)}$. The largest number of acceptable kernels is actually obtained for the intermediate target kernel size with $w_{\mathrm{H/V}} = 600/300$ km using η_1 (left column in Fig. 8). Consistently fewer acceptable kernels are found for the target resolution of $w_{\mathrm{H/V}} = 900/400$ km. Note that this difference between the target sizes does not imply that model averages are better constrained at the smaller scales rather than at large ones as we are not investigating the resolving power itself here, but the appropriateness of the Gaussian approximation. Instead, the classification maps suggest that one could lower the trade-off parameter η even further and obtain a larger number of Gaussian-shaped averaging kernels also for the larger target sizes. 'Highly focused' kernels (i.e. $\xi \geq 1.1$), are only observed for the HR inversion (Fig. 8, top left panel) at a few target locations east of Hawaii. This category suggests that narrower Gaussian estimates $\hat{g}^{(k)}$ may be possible in those cases and the associated resolution lengths could then be regarded as upper bounds of the size of the corresponding $A^{(k)}$. Alternatively, increased focus values could in some cases be explained by an overshooting local maximum at the peak combined with pronounced negative kernel values outside the peak region . The trend in Fig. 8 then also suggests that a larger value for η, with consequently lower model uncertainty, can be chosen locally if such a highly focused $A^{(k)}$ is not desirable (note again that the trade-off parameter can in principle be chosen for each target location individually). Given all the combinations of $T^{(k)}$ and η that we tested, a consistent classification of 'sufficient' and better is only possible in a few regions of excellent data coverage in the Northern Hemisphere.

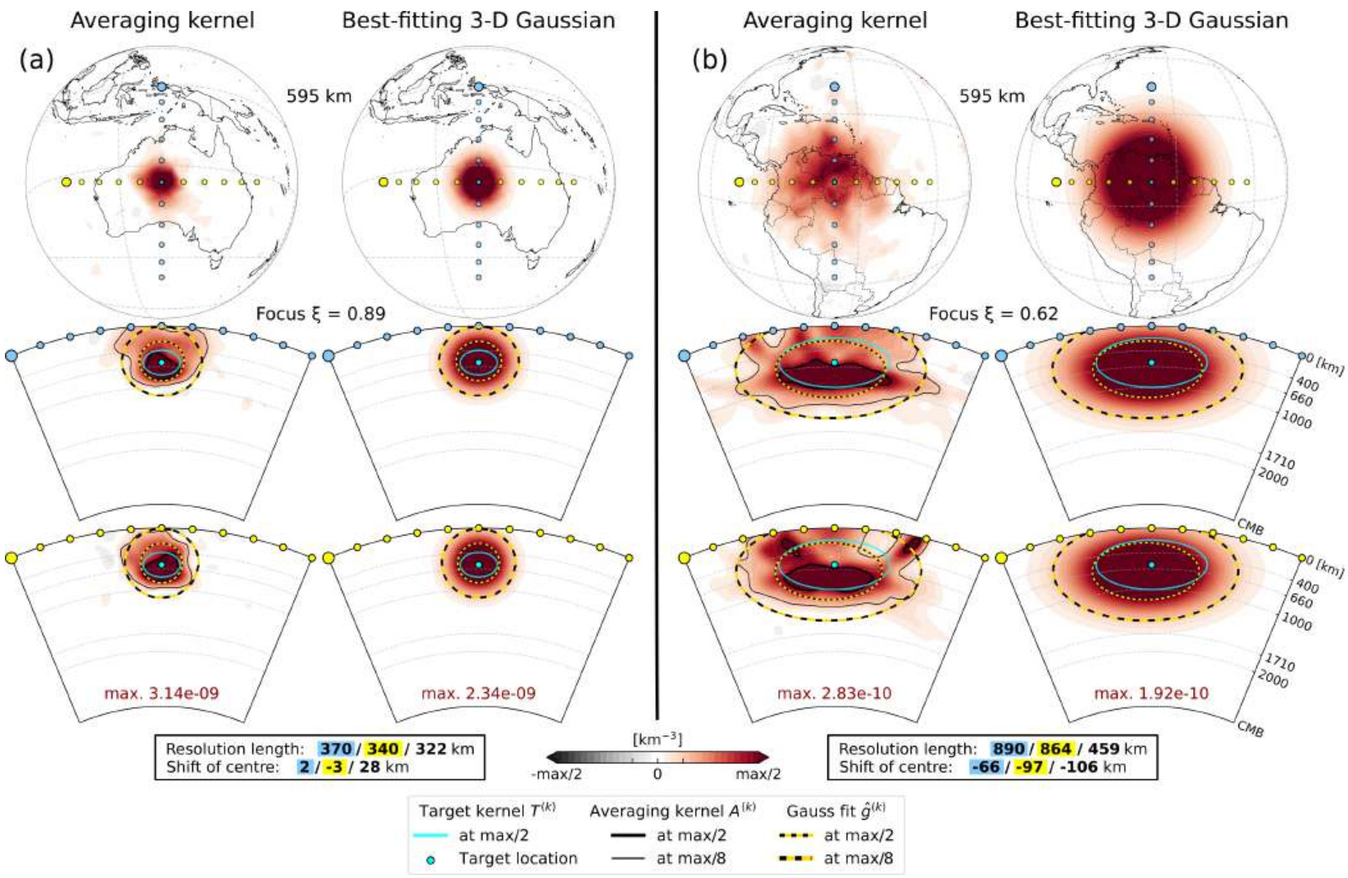

Figure 7. Averaging kernels $A^{(k)}$ and their best-fitting Gaussians $\hat{g}^{(k)}$ using (a) the high-resolution setup (η_1, $T^{(k)}$ with horizontal and vertical extent $w_{\mathrm{H/V}} = 300/200$ km), below Western Australia and (b) the low-resolution setup (η_3, $T^{(k)}$ with $w_{\mathrm{H/V}} = 900/400$ km), below Venezuela. Both target locations and the corresponding lateral cross-sections are at 595 km depth. Vertical cross-sections are marked by blue and yellow dots. Larger dots specify the negative direction for the local coordinate system in which resolution lengths and shifts of the centre are quantified. Additionally, the value for the focus ξ is shown, which serves as a metric for quantifying the appropriateness of $\hat{g}^{(k)}$ (see Section 4.2).

5 GLOBAL RESOLUTION & UNCERTAINTY MAPS

The tools presented in the previous section allow us to inspect the resolution lengths for the varying tomographic parameter setups, in connection with a basic test of the validity of the Gaussian approximation. On a global scale, this has the power to reveal concisely the impact of the different inversion parameters on resolution. In combination with the propagated model uncertainties $\sigma_{\hat{m}^{(k)}}$, which are straightforwardly calculated with the SOLA method (see eq. 5), this also makes it possible to uncover the locally varying trade-offs between resolution and uncertainty. To illustrate this, we show global maps of the estimated resolution lengths for all inversion parameter combinations in Figs 10–13. We apply our classification scheme to mask all 'insufficiently' focused averaging kernels in these maps. In addition, the mean and standard deviation of all resolution lengths, given by averaging kernels classified as 'sufficient' and better, are specified above each map.

For vertical resolution estimates at 595 km depth, one can observe a strong variability for the case with target lengths of $w_{\mathrm{H/V}} = 300/200$ km (Fig. 10, top row). At this target size, the mean vertical resolution length ranges between 316–329 km across all trade-off values and with increasing values of η, one can see a clear progression towards larger vertical extent of the averaging kernels. Minimum-maximum values of vertical resolution for each target size $w_{\mathrm{V}} = 200$, 300, 400 km are 211–501 km, 321–521 km and 416–534 km, respectively. The maps therefore show that the vertical target length only is approached in regions of high data coverage, but overall cannot be reached by any 'sufficient' averaging kernel. Numerous $A^{(k)}$ deviate strongly (>100 km) from the vertical target resolution in Fig. 10, but they do not all necessarily fall in the category of 'insufficient' averaging kernels (e.g. beneath the North Pacific). This means that a useful Gaussian $\hat{g}^{(k)}$ was obtained, although a more accurate fit to the target kernel $T^{(k)}$ could not be achieved at these locations. We chose to additionally indicate the non-Gaussian $A^{(k)}$ as shaded areas to roughly analyse the range of estimated lengths there, even if they are less reliable. The vertical resolution lengths in those regions are in fact often in line with surrounding acceptable $A^{(k)}$, but may as well be anomalously low or high, quite strikingly for instance around the East Pacific Rise (we analyse an example for this region in more detail in Section 6.1). An opposite trend can be observed in the mid-mantle at 1210 km depth in Fig. 11, where the prescribed target lengths are overall

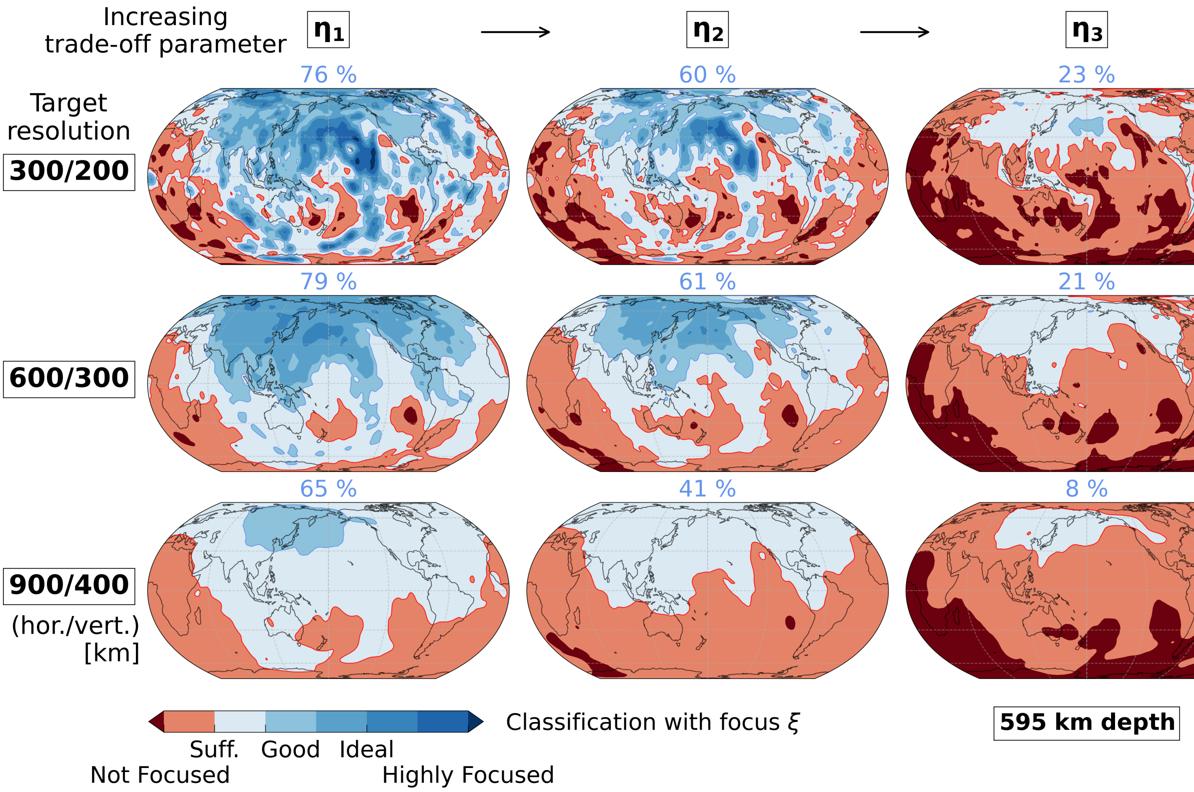

Figure 8. Classification of averaging kernels based on the focus ξ at a depth of 595 km depth for different target resolution lengths and trade-off parameters η. Blue colours mark averaging kernels that are sufficiently ('Suff.') well approximated by the best-fitting Gaussian $\hat{g}^{(k)}$, meaning that resolution length can be robustly extracted from $\hat{g}^{(k)}$. 'Good' averaging kernels are subdivided in at levels of $\xi =$ 0.9, 0.95, 1.0, 1.05 to reveal their deviation from the ideal case in more detail. Red colours accordingly indicate more complex averaging kernels $A^{(k)}$ that are often affected by strong kernel side lobes. Percentages above the maps display the relative amount of acceptable $A^{(k)}$.

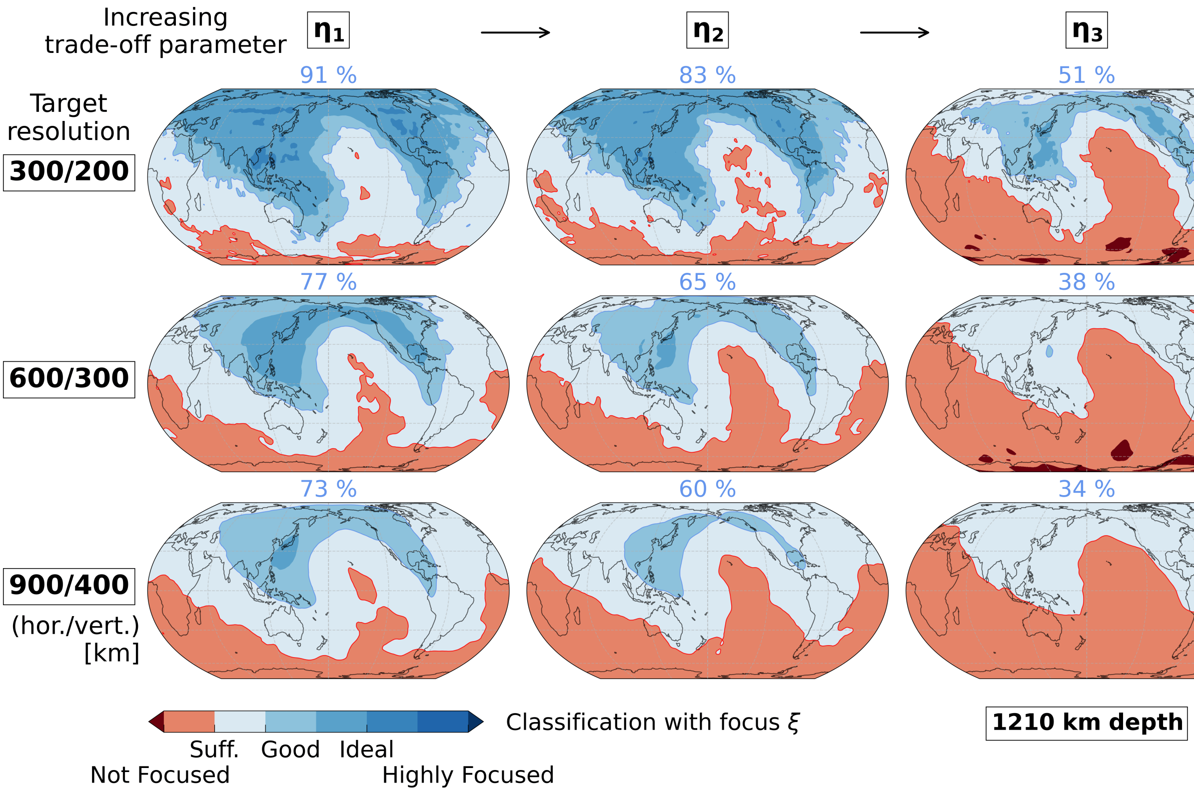

Figure 9. Classification of averaging kernels based on the focus ξ at a depth of 1210 km.

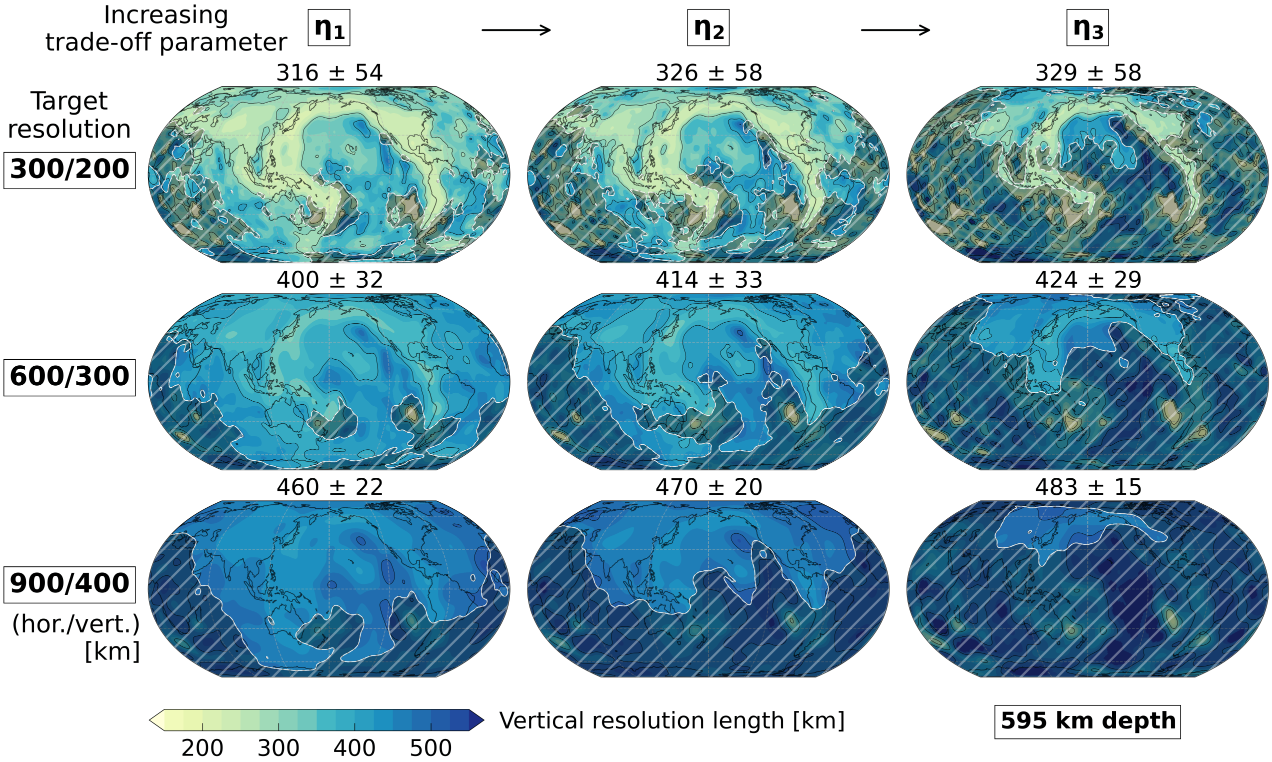

Figure 10. Vertical resolution lengths of the averaging kernels at 595 km depth, as estimated from the best-fitting Gaussian $\hat{g}^{(k)}$. Shaded areas mark the regions where the classification from Fig. 8 indicates that the $A^{(k)}$ are insufficiently Gaussian-shaped and that resolution length must be interpreted with caution there. Mean values and standard deviations of the resolution lengths of all 'sufficient' averaging kernels are given above each map. Contour lines are drawn at 200, 300, 400 and 500 km.

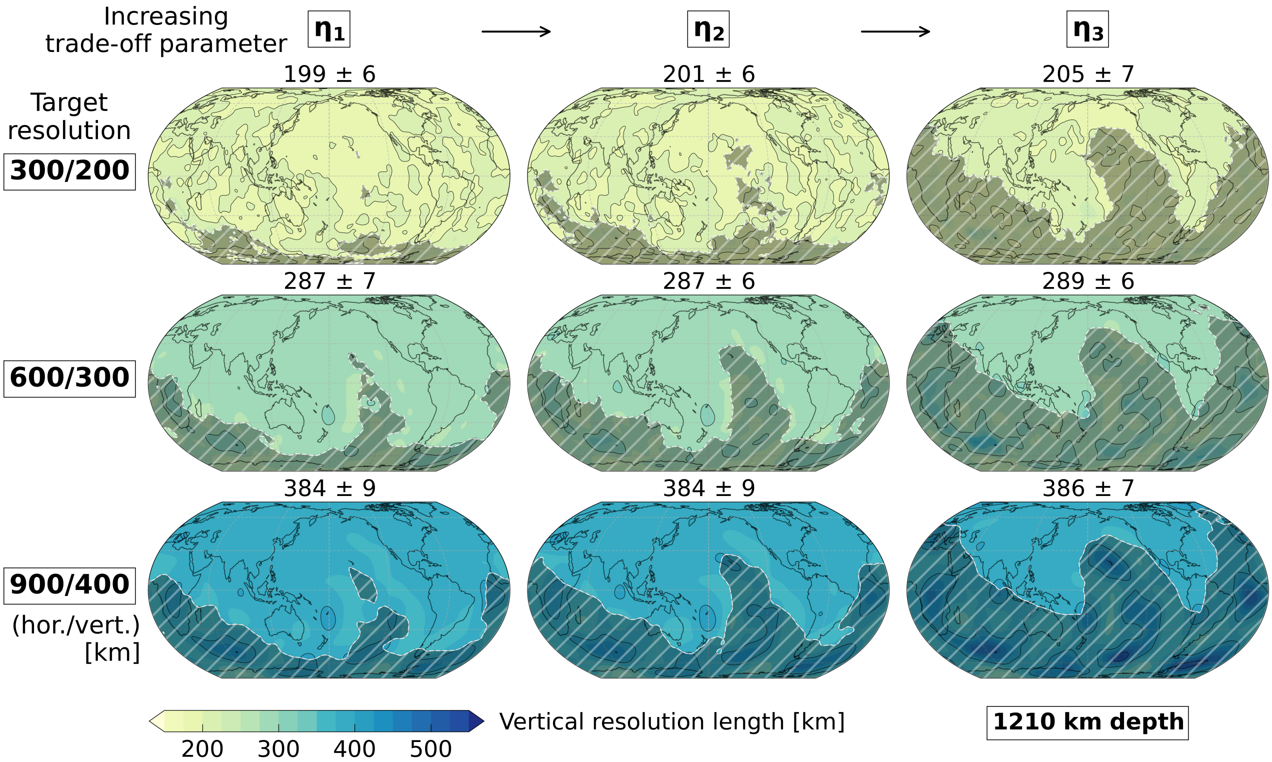

Figure 11. Vertical resolution lengths of the averaging kernels at 1210 km depth. Same as Fig. 10 with shaded areas given by Fig. 9.

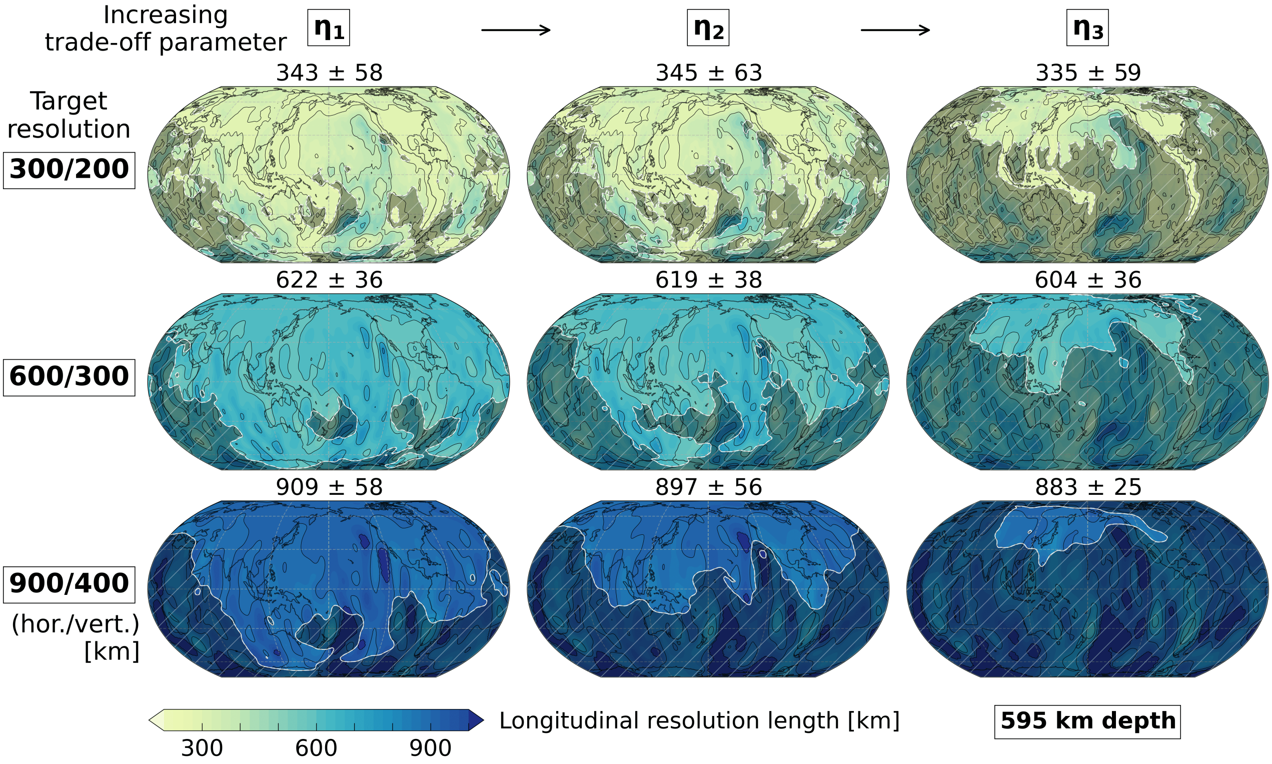

Figure 12. Longitudinal resolution lengths of the averaging kernels at 595 km depth, as estimated from the best-fitting Gaussian $\hat{g}^{(k)}$. Same layout as Fig. 10.

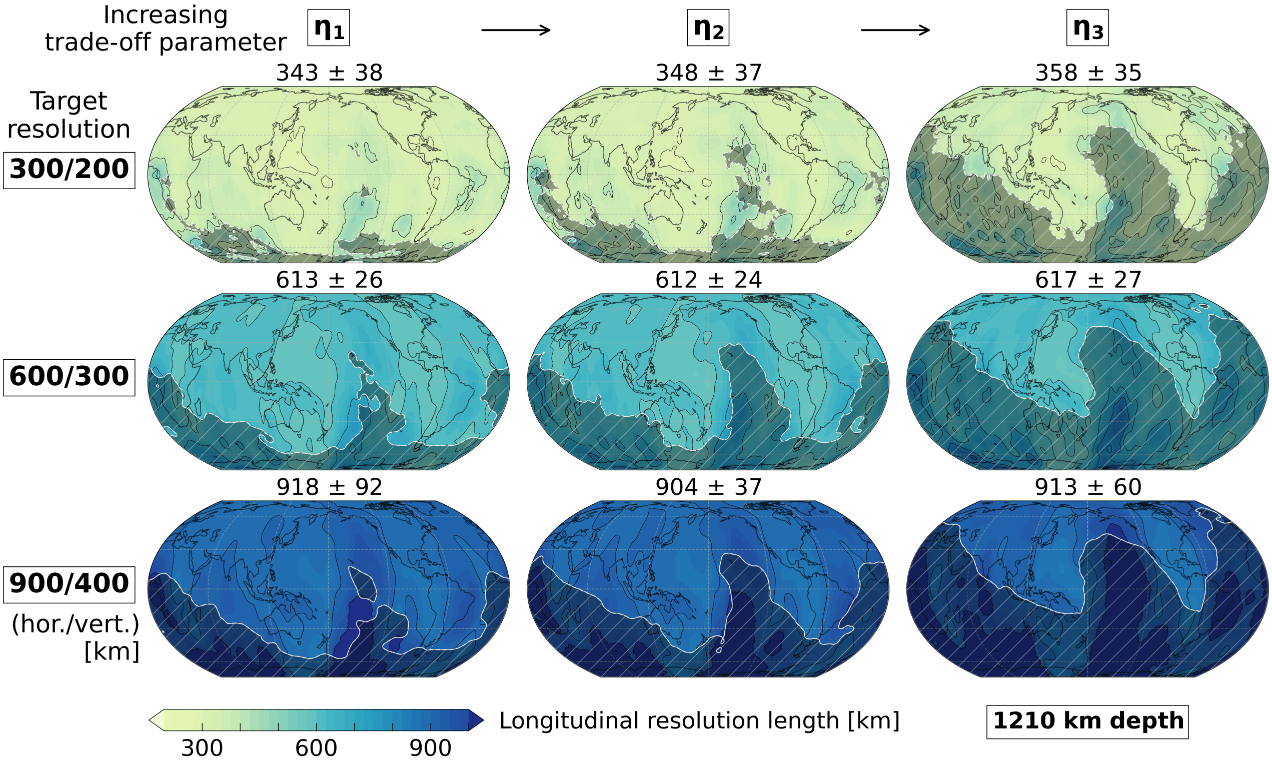

Figure 13. Longitudinal resolution lengths of the averaging kernels at 1210 km depth. Same layout as Fig. 11.

fitted well on a global scale. At this depth, the mean values of the estimated vertical resolution lengths are however often smaller compared to their target, and the full range of estimates is 176–239 km, 261–336 km and 351–423 km for $w_{\mathrm{V}} = 200$, 300, 400 km, respectively.

Maps of longitudinal resolution lengths for the same depths are shown in Figs 12 and 13. The global mean values for each selection of $w_{\mathrm{H/V}}$ and η as well as the associated distributions of longitudinal resolution lengths are quite similar at both depths. The respective target lengths are well approached in many regions. At 595 km depth, minimum-maximum values of longitudinal resolution for the different target resolutions $w_{\mathrm{H}} = 300$, 600, 900 km (again across all values of η and for all 'sufficient' and better averaging kernels) are 253–673 km, 495–826 km and 778–1082 km, respectively. At 1210 km depth, a larger number of 'sufficient' averaging kernels relative to the total number per layer is found. Still, the corresponding ranges are slightly narrower with minimum-maximum values of 293–581 km, 560–778 km and 828–1104 km. Fairly similar maps and trends can be found for the estimated latitudinal resolution lengths (shown in Appendix C, Figs C1 and C2).

To complete the results for our tests of SOLA inversion parameters, we show the associated global maps of propagated model uncertainty $\sigma_{\hat{m}^{(k)}}$ in Figs 14 and 15. As intuitively expected, the uncertainty increases systematically with lower values of η, but also with smaller target sizes. A possible reason for this could be that fewer finite-frequency kernels interact within the smaller target kernel volume in this case. The highest global mean and largest variability thus correspond to the HR inversion setup (η_1, $w_{\mathrm{H/V}} = 300/200$ km) with $\sigma_{\hat{m}^{(k)}} = 0.668 \pm 0.138$ % at 595 km depth, and 0.478 ± 0.112 % at 1210 km depth. For the largest target kernels with $w_{\mathrm{H/V}} = 900/400$ km, the model uncertainties are vanishingly small. This indicates that lower values for η should be chosen at larger target scales compared to the best choice of η at smaller target resolution. Overall, there is no strict pattern emerging for the variability of the model uncertainty with respect to the associated resolution length estimates. Instead, one can observe regions of reduced $\sigma_{\hat{m}^{(k)}}$ both where the potential for overall high resolution is observed (e.g. around Japan, with 'ideal' focus), but also where our classification tool suggested a poor Gaussian approximation of the averaging kernels.

6 DISCUSSION

6.1 Role of the Gaussian approximation

As shown in Section 4, an approximation of the averaging kernels $A^{(k)}$ by some parametric function is crucial for determining their resolving power shown in our global resolution maps. In this respect, Gaussian functions are a convenient choice and serve as a tool for both estimating resolution lengths and for identifying kernels of more complex shape and consequently low focus. Earlier studies have also relied on the Gaussian approximation in order to parametrize the resolution matrix (An 2012) or the point-spread functions based on the Hessian matrix (Fichtner & Trampert 2011) prior to the actual inversion step. In contrast to this earlier work, we computed entire sets of averaging kernels without additional assumptions on their structure, and tested subsequently if a reduction of their complexity in the form of a Gaussian parametrization is warranted. Although using the SOLA method and targeting Gaussian-shaped averaging kernels here, it was not possible to obtain robust resolution proxies from the kernels at every location and at reasonable levels of uncertainty (highlighted by the classification in Figs 8, 9 and uncertainty maps in Figs 14, 15). We therefore suggest that our concept should be applied to new datasets, including hypothetical ones, also to investigate potential improvements in resolution resulting from additional seismic stations at new locations.

The reliability of our resolution length analysis, including the computation of the kernel focus ξ, clearly depends on the quality of the individual Gaussian approximations. For example, an irregular feature was discovered by the classification scheme at a depth of 595 km (Fig. 8), suggesting strongly non-Gaussian averaging kernels around the East-Pacific Rise. Across this region, the kernels are insufficiently focused and show notably low estimates of vertical resolution lengths for all inversion parameter combinations (Fig. 10). In Fig. 16, we provide an example of such a kernel using η_1 and target lengths $w_{\mathrm{H/V}} = 300/200$ km (i.e. the 'high-resolution' scenario). It is strongly shifted to greater depths with respect to the target location by about 190 km, as estimated from the centre of the best-fitting Gaussian. While the Gaussian appears to fit this kernel well at the level of the half width at half maximum, the low value for the focus of $\xi = 0.25$, however, suggests a poor approximation. This is mainly due to elevated averaging values and oscillations of the kernel outside the plane of the cross-section (indicated for example by the local maximum visible in the lateral cross-section in Fig. 16). The estimated resolution lengths may thus still be reasonable even for these complex kernels, while the focus successfully points out their inadequacy for describing comprehensively the resolution characteristics. Again, this demonstrates why the computation of resolution lengths from 2-D cross-sections can be potentially misleading. A more robust assessment of resolution lengths for these complex scenarios instead requires some combination of accurate estimates, based for example on the 3-D Gaussian parametrization, and a classification that we realize using the concept of 'focus' here. Our analysis illustrates that both together are a useful way to extract the relevant information on resolution reliably from a large set of averaging kernels.

6.2 Implications for SOLA tomography

A key result of Zaroli (2016) was that choosing a constant value for the trade-off parameter η per layer can produce coherent SOLA tomographic images (i.e. showing geodynamically interpretable large-scale features) with bounded uncertainty on a global scale. In contrast to Zaroli (2016), in which spheroidal, constant target functions of variable size adapted to an irregular data-driven grid were used, we chose a laterally (almost) equidistant model parametrization and tested a homogeneous Gaussian target kernel size consistently at all model locations. Our results show that, for a given selection of inversion parameters η and $T^{(k)}$, the model uncertainties $\sigma_{\hat{m}^{(k)}}$ from propagated data errors may have a low variability across all target locations. This sup-

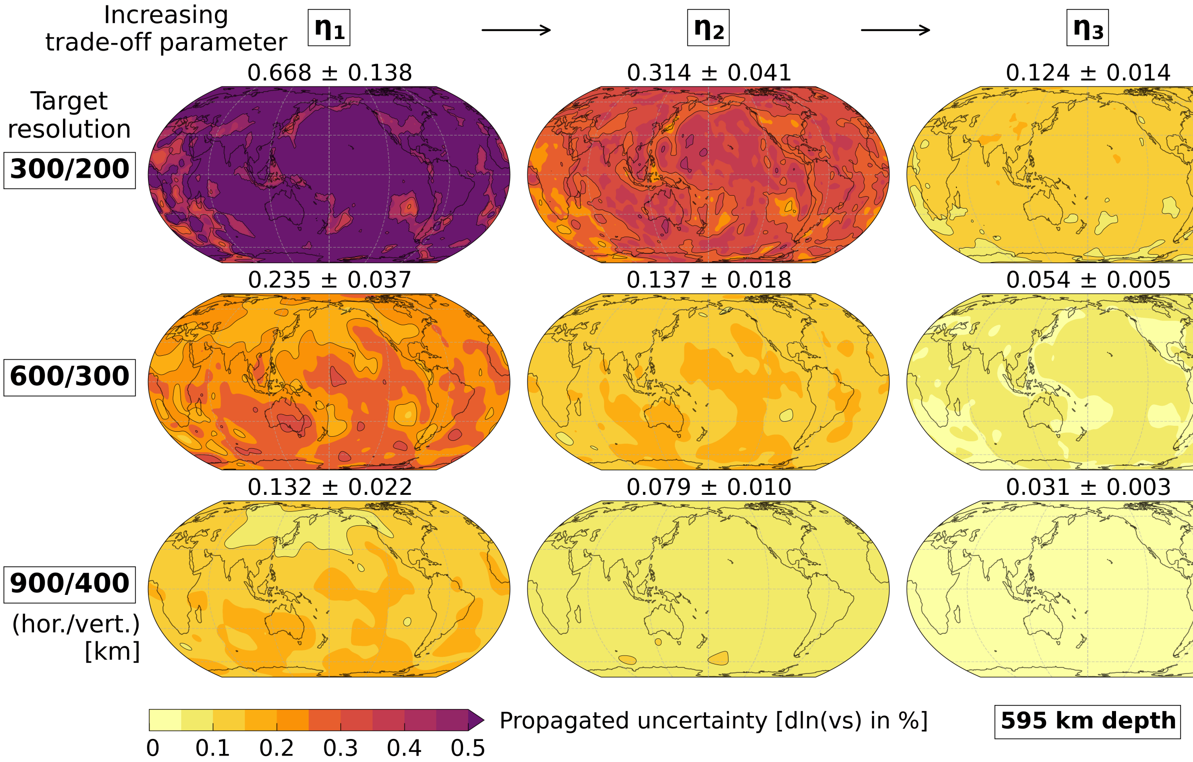

Figure 14. Propagated model uncertainty $\sigma_{\hat{m}^{(k)}}$ at 595 km depth for different target kernel sizes and trade-off parameters. Dashed contour lines are only drawn at the values shown in the colour bar. Values above each map show the global mean ± standard deviation.

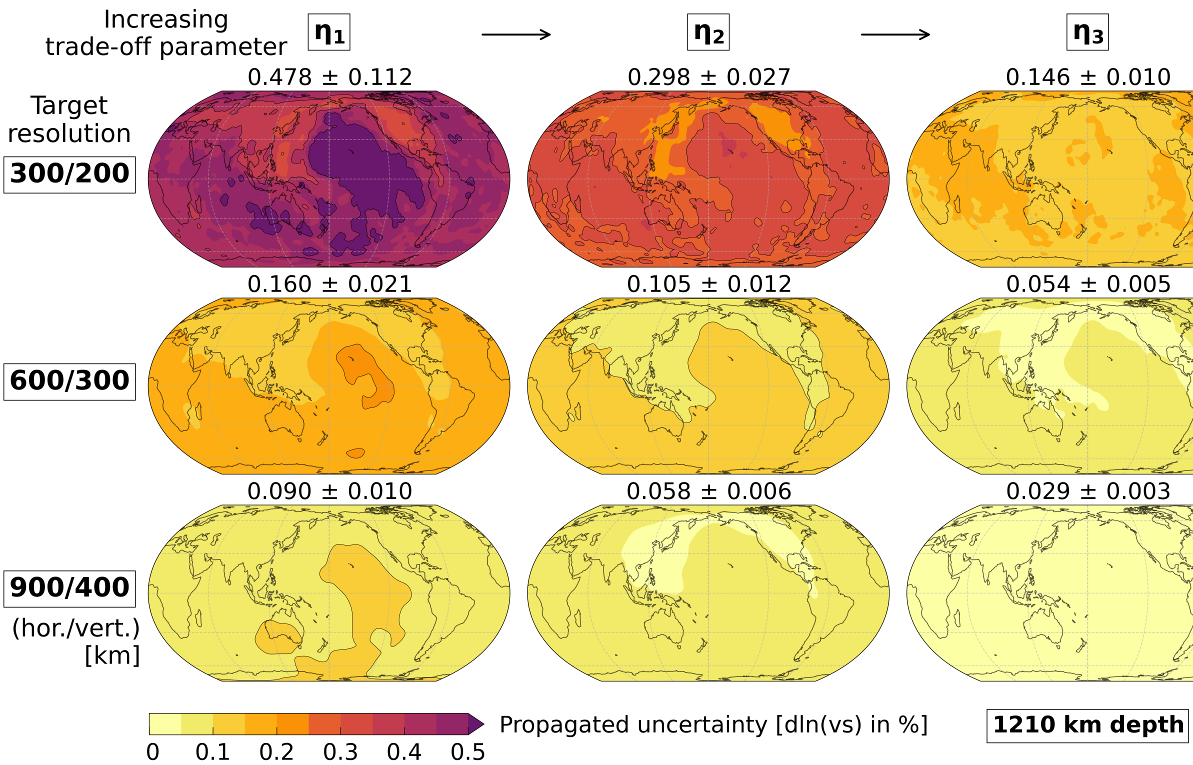

Figure 15. Same as Fig. 14 (model uncertainty $\sigma_{\hat{m}^{(k)}}$) at a depth of 1210 km.

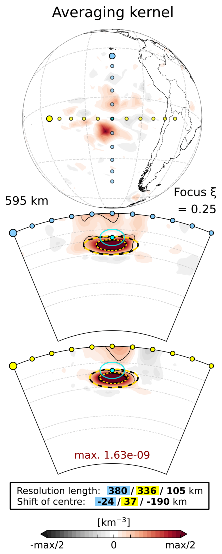

Figure 16. Example of a poorly centred averaging kernel $A^{(k)}$ in the vicinity of the East-Pacific Rise for the high-resolution inversion setup (η_1, target resolution $w_{\mathrm{H}}/w_{\mathrm{V}} = 300/200$ km). The lateral cross-section at 595 km (top plot) does not reveal the main peak due to its vertical shift of about -190 km and the limited vertical half width of 105 km. The kernel is insufficiently focused ($\xi < 0.75$), which implies that its best-fitting Gaussian is not representative for the complete averaging volume. Ellipses and contour lines are explained in Fig. 7.

ports the notion that, from the perspective of error propagation, it is indeed viable to use one single trade-off parameter (e.g., here η_2 with the current dataset) per tomographic layer. In previous applications of the SOLA method it has so far not been clear, however, whether the averaging kernels also reasonably fit the target kernels $T^{(k)}$ at a given level of uncertainty. The classification scheme developed here offers additional guidance in this regard by making it possible to assess the quality of the averaging kernels with respect to a best-fitting 3-D Gaussian instead of the target kernel. For certain research questions it might be further on desirable to obtain global tomographic models with different characteristics, favouring either high-resolution, uniformly good 'focusing' of averaging or resolving kernels, low uncertainty $\sigma_{\hat{m}^{(k)}}$, or a regionally variable mix thereof. It is, for example, not clear whether geodynamic inversions that aim at retrodicting past mantle evolution would benefit from either homogeneous tomographic resolution or from homogeneous model uncertainty (see e.g. Colli et al. 2020). This in turn requires the joint adaptation of target kernels $T^{(k)}$ and trade-off values $\eta^{(k)}$ for every individual parameter location. Such an analysis is beyond the scope of the current study, but can readily be tackled with the tools developed here. For the given shape of the target kernels in terms of a Gaussian, we have demonstrated that it becomes clear from just a few inversions, which range of values for $\eta^{(k)}$ is practically relevant and whether another size or shape for the target function $T^{(k)}$ should be employed. This empirical procedure is necessary because no automatic criterion or rule exists to determine the 'ideal' SOLA inversion parameters at every target location without excessive testing. For global applications, this would be computationally prohibitive. The joint analysis of model uncertainties $\sigma_{\hat{m}^{(k)}}$ and the focus ξ, which we used to set up the classification scheme for the averaging kernels, in comparison, is practically viable and provides essential information for selecting useful local combinations of $T^{(k)}$ and $\eta^{(k)}$, globally.

7 CONCLUSIONS

We have presented a practical concept and its application for the assessment of resolution and uncertainty of tomographic images on a global scale. It is based on: 1) explicitly available averaging kernels and uncertainties, here enabled by the SOLA method, 2) a 3-D Gaussian parametrization of the averaging kernels for estimating resolution lengths consistently, and 3) a classification scheme for identifying regions where the Gaussian approximation may not accurately represent the spatial averaging. Through this combination, it is possible not only to investigate and visualize the resolution information for all the averaging kernels together, but also to indicate in a straightforward fashion in which regions the results can be reliably interpreted. At the same time, it shows in which locations specific care must be taken, for example where individual analysis of the local resolving power of the given dataset may be advisable.

Instead of interpreting model values, we employed the approach presented here for testing various combinations of SOLA inversion parameters in terms of their effects on final resolution and uncertainty, as there is no formal way to determine any 'ideal' setup. Our analysis shows that only a few inversions are required for a given realization of the target kernels to pinpoint whether their size or shape needs to be adapted and which range of values for the trade-off parameter is useful. This in turn is important for effectively selecting the proper combinations of these parameters at each target location in case of specific requirements (e.g. tuning towards more homogeneous resolution or more homogeneous model uncertainty). Finally, we emphasize once more that the general analysis performed here as well as the classification scheme are not restricted to the setup based on SOLA. They can also be applied to the closely-related resolution concept for point-spread functions, in case they are explicitly available.

In summary, the analysis with our tomographic framework revealed conclusively that resolution lengths from the SOLA method can be primarily tuned by the choice of target kernel size, and only to a lesser extent by the trade-off parameter. Despite this design control, a good focusing of the averaging kernels (i.e. adequacy of the approximation with a Gaussian) cannot be guaranteed on a global scale with the data and possible inversion setups employed here (especially for a target resolution going down to ~300 km horizontally and ~200 km vertically). Most notably, as expected with body waves, one has less control on vertical than on horizontal resolution length, especially at shallower depths in the mantle. However, a laterally homogeneous distribution of uncertainty is generally possible by choosing locally varying trade-off parameters. Additionally, the SOLA uncertainties remind us that even if averaging kernels or point-spread functions with high resolution could be obtained everywhere with specific regularization choices, the estimated model may be highly speculative in regions of poor data coverage. It is for these reasons that a proper and coherent quantification of resolution and model uncertainty is of critical importance, since this is a prerequisite to better inform independent geophysical studies that rely on global tomographic images.

ACKNOWLEDGMENTS

We thank the editor Carl Tape, as well as Nathan Simmons and an anonymous reviewer for their constructive comments that improved the clarity of our manuscript. This work was supported by the Deutsche Forschungsgemeinschaft (DFG) under grant SCHU 2914/7-1 (project ID 456788150). We are grateful to the High-Performance Computing Centre of the University of Strasbourg and the Leibniz-Rechenzentrum (LRZ) for providing support and access to computing resources. We also want to thank Jens Oeser for an excellent in-house computational infrastructure at LMU Munich. This work was partly supported by the Programme National de Planétologie (PNP) of CNRS/INSU, co-funded by CNES.

DATA AVAILABILITY

The data underlying this article will be shared on reasonable request to the corresponding author.

REFERENCES

Aki, K., Christoffersson, A., & Husebye, E. S., 1977. Determination of the three-dimensional seismic structure of the lithosphere, *Journal of Geophysical Research (1896-1977)*, **82**(2), 277–296.

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).

An, M., 2012. A simple method for determining the spatial resolution of a general inverse problem, *Geophysical Journal International*, **191**(2), 849–864.

Backus, G. E. & Gilbert, J. F., 1967. Numerical Applications of a Formalism for Geophysical Inverse Problems, *Geophys. J. Int.*, **13**(1-3), 247–276.

Backus, G. E. & Gilbert, J. F., 1968. The Resolving Power of Gross Earth Data, *Geophys. J. Int.*, **16**(2), 169–205.

Backus, G. E. & Gilbert, J. F., 1970. Uniqueness in the Inversion of Inaccurate Gross Earth Data, *Phil. Trans. R. Soc. Lond. Series A*, **266**, 123–192.

Bogiatzis, P., Miaki, I., & Davis, T. A., 2016. Towards using direct methods in seismic tomography: computation of the full resolution matrix using high-performance computing and sparse QR factorization, *Geophys. J. Int.*, **205**(2), 830–836.

Boschi, L., 2003. Measures of resolution in global body wave tomography, *Geophys. Res. Lett.*, **30**(19).

Bozdağ, E., Peter, D., Lefebvre, M., Komatitsch, D., Tromp, J., Hill, J., Podhorszki, N., & Pugmire, D., 2016. Global adjoint tomography: first-generation model, *Geophys. J. Int.*, **207**(3), 1739–1766.

Bui-Thanh, T., Ghattas, O., Martin, J., & Stadler, G., 2013. A computational framework for infinite-dimensional Bayesian inverse problems part I: The linearized case, with application to global seismic inversion, *SIAM Journal on Scientific Computing*, **35**(6), A2494–A2523.

Bunge, H.-P., Hagelberg, C. R., & Travis, B. J., 2003. Mantle circulation models with variational data-assimilation: Inferring past mantle flow and structure from plate motion histories and seismic tomography, *Geophys. J. Int.*, **2**(152), 280–301.

Choblet, G., Deschamps, F., Amit, H., & Lasbleis, M., 2023. Inferring the relationship between core-mantle heat flux and seismic tomography from mantle convection simulations, *Physics of the Earth and Planetary Interiors*, **342**, 107072.

Colli, L., Fichtner, A., & Bunge, H.-P., 2013. Full waveform tomography of the upper mantle in the South Atlantic region: evidence for pressure-driven westward flow in a shallow asthenosphere, *Tectonophysics*, **604**(0), 26–40.

Colli, L., Ghelichkhan, S., Bunge, H.-P., & Oeser, J., 2018. Retrodictions of Mid Paleogene mantle flow and dynamic topography in the Atlantic region from compressible high resolution adjoint mantle convection models: Sensitivity to deep mantle viscosity and tomographic input model, *Gondwana Research*, **53**, 252–272, Rifting to Passive Margins.

Colli, L., Bunge, H.-P., & Oeser, J., 2020. Impact of model inconsistencies on reconstructions of past mantle flow obtained using the adjoint method, *Geophysical Journal International*, **221**(1), 617–639.

Dahlen, F. A., Hung, S. H., & Nolet, G., 2000. Fréchet kernels for finite-frequency traveltimes – I. Theory, *Geophys. J. Int.*, **141**(1), 157–174.

Davies, D. R., Valentine, A. P., Kramer, S. C., Rawlinson, N., Hoggard, M. J., & Eakin, C. M.and Wilson, C. R., 2019. Earth's multi-scale topographic response to global mantle flow, *Nature Geoscience*, **12**(10), 845–850.

Davies, D. R., Ghelichkhan, S., Hoggard, M., Valentine, A., & Richards, F., 2023. Chapter 11 - Observations and models of dynamic topography: Current status and future directions, in *Dynamics of Plate Tectonics and Mantle Convection*, pp. 223–269, ed. Duarte, J. C., Elsevier.

Deal, M. & Nolet, G., 1996. Nullspace shuttles, *Geophys. J. Int.*, **124**, 372–380.

Dziewonski, A. M., Hager, B. H., & O'Connell, R. J., 1977. Large-scale heterogeneities in the lower mantle, *J. Geophys. Res.*, **82**(2), 239—255.

Favier, N., Chevrot, S., & Komatitsch, D., 2004. Near-field influences on shear wave splitting and traveltime sensitivity kernels, *Geophys. J. Int.*, **156**, 467–482.

Fichtner, A., 2010. *Full seismic waveform modelling and inversion*, Springer Science & Business Media.

Fichtner, A. & Leeuwen, T. v., 2015. Resolution analysis by random probing, *Journal of Geophysical Research: Solid Earth*, **120**(8), 5549–5573.

Fichtner, A. & Trampert, J., 2011. Resolution analysis in full waveform inversion, *Geophys. J. Int.*, **187**(3), 1604–1624.

Fichtner, A., Kennett, B. L. N., Igel, H., & Bunge, H.-P., 2009. Full seismic waveform tomography for upper-mantle structure in the Australasian region using adjoint methods, *Geophys. J. Int.*, **179**(3), 1703–1725.

Fichtner, A., Trampert, J., Cupillard, P., Saygin, E., Taymaz, T., Capdeville, Y., & Villaseñor, A., 2013. Multiscale full waveform inversion, *Geophysical Journal International*, **194**(1), 534–556.

Fichtner, A., van Herwaarden, D.-P., Afanasiev, M., Simutė, S., Krischer, L., Çubuk Sabuncu, Y., Taymaz, T., Colli, L., Saygin, E., Villaseñor, A., Trampert, J., Cupillard, P., Bunge, H.-P., & Igel, H., 2018. The Collaborative Seismic Earth Model: Generation 1, *Geophysical Research Letters*, **45**(9), 4007–4016.

Fichtner, A., Zunino, A., Gebraad, L., & Boehm, C., 2021. Autotuning Hamiltonian Monte Carlo for efficient generalized nullspace exploration, *Geophysical Journal International*, **227**(2), 941–968.

Fong, D. C.-L. & Saunders, M., 2011. LSMR: An iterative algorithm for sparse least-squares problems, *SIAM Journal on Scientific Computing*, **33**(5), 2950–2971.

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.

French, S. W. & Romanowicz, B., 2015. Broad plumes rooted at the base of the Earth's mantle beneath major hotspots, *Nature*, **525**, 95–9.

Ghelichkhan, S., Bunge, H.-P., & Oeser, J., 2021. Global mantle flow retrodictions for the early Cenozoic using an adjoint method: evolving dynamic topographies, deep mantle structures, flow trajectories and sublithospheric stresses, *Geophysical Journal International*, **226**(2), 1432–1460.

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.

Horbach, A., Bunge, H.-P., & Oeser, J., 2014. The adjoint method in geodynamics: derivation from a general operator formulation and application to the initial condition problem in a high resolution mantle circulation model, *GEM - Int. J. Geomath.*, **5**(2), 163–194.

Hosseini, K., Sigloch, K., Tsekhmistrenko, M., Zaheri, A., Nissen-Meyer, T., & Igel, H., 2020. Global mantle structure from multifrequency tomography using P, PP and P-diffracted waves, *Geophysical Journal International*, **220**(1), 96–141.

Huang, H., Dennis, J. M., Wang, L., & Chen, P., 2013. A scalable parallel LSQR algorithm for solving large-scale linear system for tomographic problems: A case study in seismic tomography, *Procedia Computer Science*, **18**, 581–590, 2013 International Conference on Computational Science.

Hung, S.-H., Dahlen, F. A., & Nolet, G., 2001. Wavefront healing: a banana-doughnut perspective, *Geophys. J. Int.*, **146**(2), 289–312.

Igel, H., Djikpéssé, H., & Tarantola, A., 1996. Waveform inversion of marine reflection seismograms for P impedance and Poisson's ratio, *Geophys. J. Int.*, **124**, 363–371.

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.

Koelemeijer, P., Schuberth, B., Davies, D., Deuss, A., & Ritsema, J., 2018. Constraints on the presence of post-perovskite in Earth's lowermost mantle from tomographic-geodynamic model comparisons, *Earth Planet. Sci. Lett.*, **494**, 226–238.

Koroni, M., Borgeaud, A., Fichtner, A., & Deschamps, F., 2022. An analysis of core–mantle boundary related seismic waves using full-waveform modelling and adjoint methods, *Geophysical Journal International*, **232**(2), 1259–1275.

Krischer, L., Fichtner, A., Boehm, C., & Igel, H., 2018. Automated large-scale full seismic waveform inversion for North America and the North Atlantic, *Journal of Geophysical Research: Solid Earth*, **123**(7), 5902–5928.

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.

Lee, E.-J., Huang, H., Dennis, J. M., Chen, P., & Wang, L., 2013. An optimized parallel LSQR algorithm for seismic tomography, *Computers & Geosciences*, **61**, 184–197.

Lei, W., Ruan, Y., Bozdağ, E., Peter, D., Lefebvre, M., Komatitsch, D., Tromp, J., Hill, J., Podhorszki, N., & Pugmire, D., 2020. Global adjoint tomography—model GLAD-M25, *Geophysical Journal International*, **223**(1), 1–21.

Liu, Q. & Peter, D., 2020. Square-root variable metric-based nullspace shuttle: A characterization of the nonuniqueness in elastic full-waveform inversion, *Journal of Geophysical Research: Solid Earth*, **125**(2), e2019JB018687.

Liu, Q., Beller, S., Lei, W., Peter, D., & Tromp, J., 2021. Preconditioned BFGS-based uncertainty quantification in elastic full-waveform inversion, *Geophysical Journal International*, **228**(2), 796–815.

Louis, A. K. & Maass, P., 1990. A mollifier method for linear operator equations of the first kind, *Inverse Problems*, **6**(3), 427.

Lu, C., Grand, S. P., Lai, H., & Garnero, E. J., 2019. TX2019slab: A new P and S tomography model incorporating subducting slabs, *Journal of Geophysical Research: Solid Earth*, **124**(11), 11549–11567.

Ma, J., Bunge, H.-P., Thrastarson, S., Fichtner, A., Herwaarden, D.-P. v., Tian, Y., Chang, S.-J., & Liu, T., 2022. Seismic full-waveform inversion of the crust-mantle structure beneath China and adjacent regions, *Journal of Geophysical Research: Solid Earth*, **127**(9), e2022JB024957, e2022JB024957 2022JB024957.

MacCarthy, J. K., Borchers, B., & Aster, R. C., 2011. Efficient stochastic estimation of the model resolution matrix diagonal and generalized cross–validation for large geophysical inverse problems, *Journal of Geophysical Research: Solid Earth*, **116**(B10).

Mégnin, C., Bunge, H.-P., Romanowicz, B., & Richards, M., 1997. Imaging 3-D spherical convection models: What can seismic tomography tell us about mantle dynamics?, *Geophys. Res. Lett.*, **24**(11), 1299–1302.

Mercerat, E. & Nolet, G., 2012. Comparison of ray-based and adjoint-based sensitivity kernels for body-wave seismic tomography, *Geophys. Res. Lett.*, **39**, L12301.

Montoison, A. & Orban, D., 2023. Krylov.jl: A Julia basket of hand-picked Krylov methods, *Journal of Open Source Software*, **8**(89), 5187.

Nerlich, R., Colli, L., Ghelichkhan, S., Schuberth, B., & Bunge, H.-P., 2016. Constraining central Neo-Tethys Ocean reconstructions with mantle convection models, *Geophys. Res. Lett.*, **43**(18), 9595—-9603.

Nolet, G., 2008. *A Breviary of Seismic Tomography*, Cambridge Univ. Press, New York, ISBN 978-0-521-88244-6.

Nolet, G., Montelli, R., & Virieux, J., 1999. Explicit, approximate expressions for the resolution and a posteriori covariance of massive tomographic systems, *Geophysical Journal International*, **138**(1), 36–44.

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

Papanagnou, I., Schuberth, B. S. A., & Thomas, C., 2022. Geodynamic predictions of seismic structure and discontinuity topography of the mantle transition zone, *Geophysical Journal International*, **234**(1), 355–378.

Pijpers, F. & Thompson, M., 1992. Faster formulations of the Optimally Localized Averages method for helioseismic inversions, *Astronomy and Astrophysics*, **262**, L33–L36.

Pijpers, F. & Thompson, M., 1994. The SOLA method for helioseismic inversion, *Astronomy and Astrophysics*, **281**, 231–240.

Pratt, R. G., 1999. Seismic waveform inversion in the frequency domain, Part 1: Theory and verification in a physical scale model, *Geophysics*, **64**(3), 888–901.

Rawlinson, N., Fichtner, A., Sambridge, M., & Young, M. K., 2014. Seismic tomography and the assessment of uncertainty, *Advances in Geophysics*, **55**, 1–76.

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.

Richards, F. D., Hoggard, M. J., Ghelichkhan, S., Koelemeijer, P., & Lau, H. C., 2023. Geodynamic, geodetic, and seismic constraints favour deflated and dense-cored LLVPs, *Earth and Planetary Science Letters*, **602**, 117964.

Ritsema, J., van Heijst, H. J., & Woodhouse, J. H., 2004. Global transition zone tomography, *J. Geophys. Res.*, **109**(B2),

B02302.
Ritsema, J., McNamara, A. K., & Bull, A. L., 2007. Tomographic filtering of geodynamic models: Implications for model interpretation and large-scale mantle structure, *J. Geophys. Res.*, **112**(B1), B01303.
Ritsema, J., Deuss, A., van Heijst, H. J., & Woodhouse, J. H., 2011. S40RTS: a degree-40 shear-velocity model for the mantle from new Rayleigh wave dispersion, teleseismic traveltime and normal-mode splitting function measurements, *Geophys. J. Int.*, **184**(3), 1223–1236.
Rodgers, A., Krischer, L., Afanasiev, M., Boehm, C., Doody, C., Chiang, A., & Simmons, N., 2022. WUS256: An adjoint waveform tomography model of the crust and upper mantle of the Western United States for improved waveform simulations, *Journal of Geophysical Research: Solid Earth*, **127**(7), e2022JB024549, e2022JB024549 2022JB024549.
Schaeffer, A. J. & Lebedev, S., 2013. Global shear speed structure of the upper mantle and transition zone, *Geophysical Journal International*, **194**(1), 417–449.
Schuberth, B. S. A., Bunge, H.-P., & Ritsema, J., 2009a. Tomographic filtering of high-resolution mantle circulation models: Can seismic heterogeneity be explained by temperature alone?, *Geochem. Geophys. Geosyst.*, **10**(5), Q05W03.
Schuberth, B. S. A., Bunge, H.-P., Steinle-Neumann, G., Moder, C., & Oeser, J., 2009b. Thermal versus elastic heterogeneity in high-resolution mantle circulation models with pyrolite composition: High plume excess temperatures in the lowermost mantle, *Geochem. Geophys. Geosyst.*, **10**(1), Q01W01.
Schuberth, B. S. A., Zaroli, C., & Nolet, G., 2015. Traveltime dispersion in an isotropic elastic mantle: strong lower-mantle signal in differential-frequency residuals, *Geophys. J. Int.*, **203**(3), 2099–2118.
Sengupta, M. K. & Toksöz, M. N., 1976. Three dimensional model of seismic velocity variation in the Earth's mantle, *Geophysical Research Letters*, **3**(2), 84–86.
Shephard, G. E., Müller, R. D., Liu, L., & Gurnis, M., 2010. Miocene drainage reversal of the Amazon River driven by plate-mantle interaction, *Nature Geoscience*, **3**(12), 870–875.
Simmons, N. A., Schuberth, B. S. A., Myers, S. C., & Knapp, D. R., 2019. Resolution and covariance of the LLNL-G3D-JPS global seismic tomography model: Applications to travel time uncertainty and tomographic filtering of geodynamic models, *Geophys. J. Int.*, **217**(3), 1543–1557.
Soldati, G., Boschi, L., & Piersanti, A., 2006. Global seismic tomography and modern parallel computers, *Annals Of Geophysics*, **49**(4-5), 977–986.
Spasojevic, S., Liu, L., & Gurnis, M., 2009. Adjoint models of mantle convection with seismic, plate motion, and stratigraphic constraints: North America since the Late Cretaceous, *Geochemistry, Geophysics, Geosystems*, **10**(5).
Stixrude, L. & Lithgow-Bertelloni, C., 2007. Influence of phase transformations on lateral heterogeneity and dynamics in Earth's mantle, *Earth Planet. Sci. Lett.*, **263**, 45–55.
Swinbank, R. & Purser, J. R., 2006. Fibonacci grids: A novel approach to global modelling, *Quarterly Journal of the Royal Meteorological Society*, **132**(619), 1769–1793.
Tape, C., Liu, Q., Maggi, A., & Tromp, J., 2009. Adjoint tomography of the Southern California Crust, *Science*, **325**(5943), 988–992.
Tarantola, A., 2005. *Inverse problem theory and methods for model parameter estimation*, SIAM.
Trampert, J., 1998. Global seismic tomography: the inverse problem and beyond, *Inverse Problems*, **14**(3), 371.
Trampert, J., Fichtner, A., & Ritsema, J., 2013. Resolution tests revisited: the power of random numbers, *Geophys. J. Int.*, **192**(2), 676–680.
Zaroli, C., 2010. *Global Multiple-Frequency S-wave Tomography of the Earth's Mantle*, Ph.D. thesis, EOST/IPGS, Université de Strasbourg.
Zaroli, C., 2016. Global seismic tomography using Backus–Gilbert inversion, *Geophys. J. Int.*, **207**(2), 876–888.
Zaroli, C., 2019. Seismic tomography using parameter-free Backus–Gilbert inversion, *Geophys. J. Int.*, **218**(1), 619–630.
Zaroli, C., Lambotte, S., & Lévêque, J.-J., 2015. Joint inversion of normal-mode and finite-frequency *S*-wave data using an irregular tomographic grid, *Geophys. J. Int.*, **203**(3), 1665–1681.
Zaroli, C., Koelemeijer, P., & Lambotte, S., 2017. Toward seeing the Earth's interior through unbiased tomographic lenses, *Geophys. Res. Lett.*, **44**(22), 11,399–11,408.
Zhang, Z. & Shen, Y., 2008. Cross-dependence of finite-frequency compressional waveforms to shear seismic wave speeds, *Geophys. J. Int.*, **174**(3), 941–948.

APPENDIX A: USING BACKUS-GILBERT THEORY FOR TOMOGRAPHY

In the original theory, Backus and Gilbert propose to construct optimal kernels $A^{(k)}(\mathbf{r})$ by approximation of a delta peak $\delta(\mathbf{r}-\mathbf{r}^{(k)})$ in terms of a specific deltaness criterion (Backus & Gilbert 1968), for example by minimizing the so-called spread (Backus & Gilbert 1970). For a perfect delta peak, the linear averaging in eq. (1) would simply show that $\hat{m}^{(k)} \approx m(\mathbf{r}^{(k)})$, which is however only possible if the data were complete and free of errors. Instead, one will typically need to accept $A^{(k)}(\mathbf{r})$ that deviate from $\delta(\mathbf{r}-\mathbf{r}^{(k)})$, due to the available set of data kernels $K_i(\mathbf{r})$, and in order to moderate the propagation of data errors into the inferred averages. A tomographic model would then consist in a collection of M local averages, $\hat{m}^{(k)}$, for $k = 1, ..., M$. Rewriting eq. (1) in a discrete notation with volumetric weights $\mathcal{V}_j$ appropriate for each j-th grid node as

$$\hat{m}^{(k)} = \sum_j A_j^{(k)} \mathcal{V}_j m_j = \sum_j \mathcal{R}_{kj} m_j \,, \tag{A.1}$$

one can see that the averaging kernels projected onto a discrete model parametrization represent a single row of the resolution operator $\mathcal{R}_{k\cdot}$. The operator $\mathcal{R}$ can then be used to retrieve information on the resolution for a specific model parametrization, while $A^{(k)}(\mathbf{r})$ essentially refers to the local resolution for the continuous model $m(\mathbf{r}^{(k)})$ (Trampert 1998). Theoretically, however there is no need to discretise the model, as the Backus–Gilbert approach essentially solves a continuous inverse problem. This can be exploited fully for example by the 'parameter-free' SOLA approach (Zaroli 2019), while in the present paper we still make use of the discrete formulation as described in Zaroli (2016). A noteworthy limitation is that the theory does not guarantee that the collection of averages together actually explains the data. We tested this for our class of models that use only body waves at the moment, and we observed that the global misfit reduction can actually be comparable to and sometimes even be better than in classic damped-least squares inversion with model norm damping. However, the question of data misfits in SOLA tomographies, especially for different inversion parameter choices, might be a relevant subject for future investigations.

APPENDIX B: EFFICIENTLY SOLVING THE SOLA B-G SYSTEM

Zaroli (2016) explained how (the discrete version of) the SOLA optimization problem in eq. (3) can be solved together with the unimodular constraint eq. (4) using a least-squares approach. Therefore, the SOLA system can be rewritten in the following fashion (see Appendix A1, Zaroli 2016):

$$\hat{\mathbf{x}}^{(k,\eta)} = \underset{\hat{\mathbf{x}}^{(k)}}{\arg\min} \left\| \begin{bmatrix} \mathbf{Q}^{(\eta)} \\ \eta\,\mathbf{I}_{N-1} \end{bmatrix} \hat{\mathbf{x}}^{(k)} - \begin{bmatrix} \mathbf{y}^{(k,\eta)} \\ \mathbf{0}_{N-1} \end{bmatrix} \right\|_2^2 . \qquad \text{(B.1)}$$

Both the matrix $\mathbf{Q}^{(\eta)}$, of size $(M+1) \times (N\text{-}1)$, and the right-hand side $\mathbf{y}^{(k,\eta)}$, which incorporates the target kernel, depend on the choice of a particular trade-off parameter η. The complete solution $\mathbf{x}^{(k)}$ can finally be recovered from $\hat{\mathbf{x}}^{(k)}$ (the intermediate solution in eq. B.1). Considering our tomographic grid and dataset, in each inversion we need to solve for the Backus–Gilbert coefficients $\mathbf{x}^{(k)}$ with a SOLA system matrix $\mathbf{Q}^{(\eta)}$ of size 153,324 × 79,764 (∼2 per cent non-zero elements, ∼2 GB). As one needs to perform a single inversion for every grid node, computational costs for empirically testing inversion parameters can thus quickly become prohibitive. The SOLA method has the computational advantage that the left-hand side of the corresponding linear system in eq. (B.1) is independent of the target location $\mathbf{r}^{(k)}$, i.e. for a given η the SOLA system matrix $\mathbf{Q}^{(\eta)}$ does not change. This enables perfectly parallel computation over all M model parameters. To this end, one could simply increase the number of processors P at the cost of having to store $2 \cdot P$ times the SOLA system matrix $\mathbf{Q}$. This obviously becomes problematic if the available computing system is limited in RAM, especially if the tomographic systems become even larger than the ones considered in this study. Alternatively, using improved parallel solvers based on LSQR (Huang et al. 2013; Lee et al. 2013) or using instead efficient direct methods (Bogiatzis et al. 2016) are other possible options that we considered and list here for documentation. We decided to solve eq. (B.1) by using a GPU version of the algorithm LSMR (Fong & Saunders 2011). We use LSMR as implemented by Krylov.jl (a package of selected Krylov methods written in the programming language *Julia*, see Montoison & Orban 2023), where, conveniently, no significant code changes are required compared to the CPU version. Once the GPU compute kernel is compiled, different left-hand sides $\mathbf{y}^{(k,\eta)}$ can be asynchronously copied from a CPU to the GPU and solution vectors can be efficiently recovered for each run without additional time spent for data transfer or solver setup. This way, performing the inversion for a single model parameter in our computations was about 50–100 times faster on the GPU (using double precision) compared to a single CPU. Also, this only requires that $\mathbf{Q}^{(\eta)}$, a large but highly sparse matrix, fits twice into GPU memory. As a remark, LSMR is recommended to be used over LSQR by Fong & Saunders (2011) if iterations have to be stopped early. This could theoretically be exploited to keep the solver time limited in case a reasonable maximum number of iterations for all inversions globally is found. However, we have not drawn on this here and relied on the stopping criteria suggested by Fong & Saunders (2011). For a given trade-off parameter η, we found that the time to solution was practically constant for all model parameters and overall increased roughly by a factor of 2 for an equal decrease in η (i.e. solutions for $\eta_1 = 5$ took about twice as long as for $\eta_2 = 10$). The use of GPUs therefore proves to be very useful for SOLA inversions with a least-squares approach and can greatly reduce the time required for computing a complete model with M parameters. As a final note, we used 2 NVidia RTX A5000 in this study, but modern GPU clusters and supercomputers often feature many more units. From a computational point of view, it should thus be straightforward to use larger grids and datasets with the SOLA method than presented here.

APPENDIX C: ADDITIONAL GLOBAL RESOLUTION MAPS

To complete the global analysis, we show additionally the estimated latitudinal resolution lengths in Figs C1 and C2. Plots for the horizontal and vertical shifts of the centres of mass of the averaging kernels can be found in the supplementary (online) material, as well as comparisons between the uncorrelated and correlated (i.e. including rotation) 3-D Gaussian estimates.

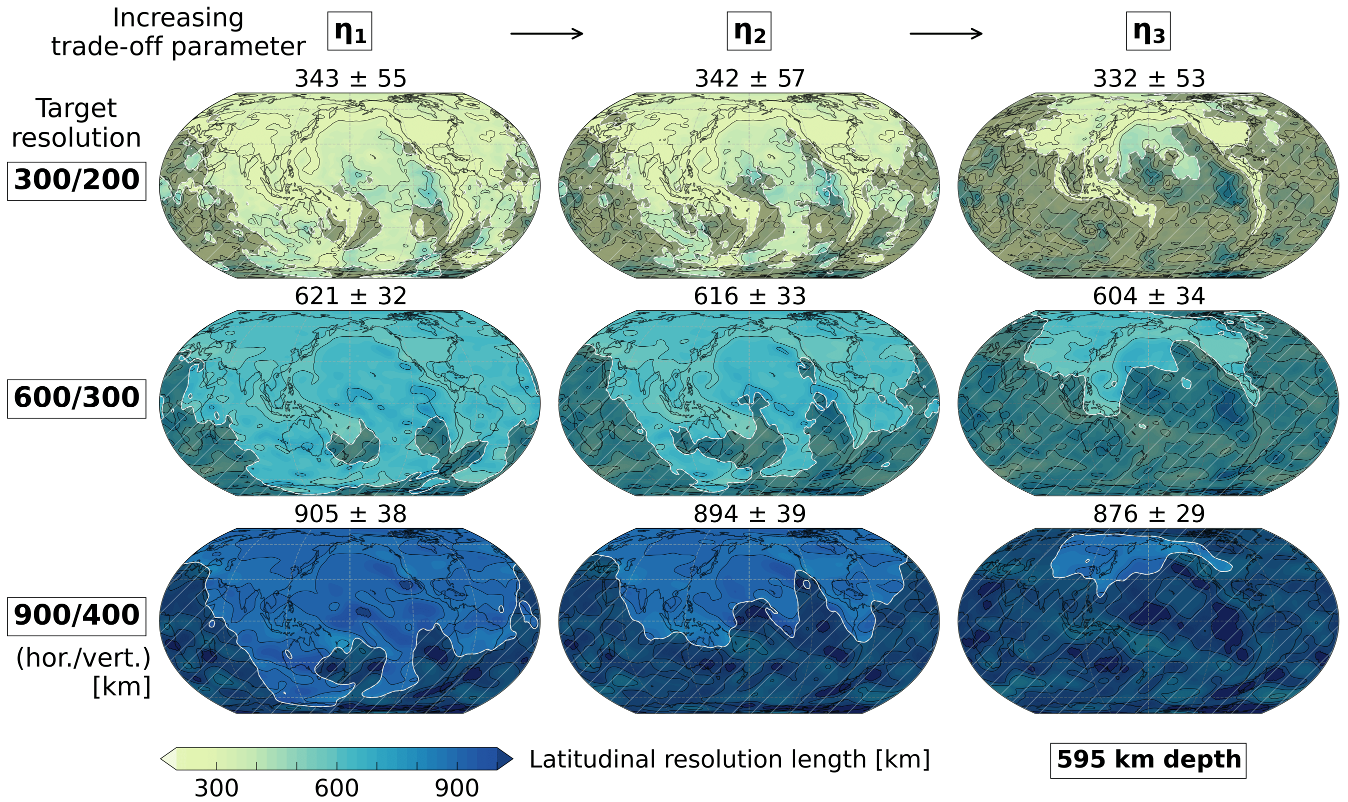

Figure C1. Latitudinal resolution lengths of the averaging kernels at 595 km depth, as estimated from the best-fitting Gaussian $\hat{g}^{(k)}$. Same layout as Fig. 10. As shown in the main text, shaded areas highlight regions of low 'focus' (i.e. the Gaussian is inadequate to reliably estimate resolution length for the averaging kernel).

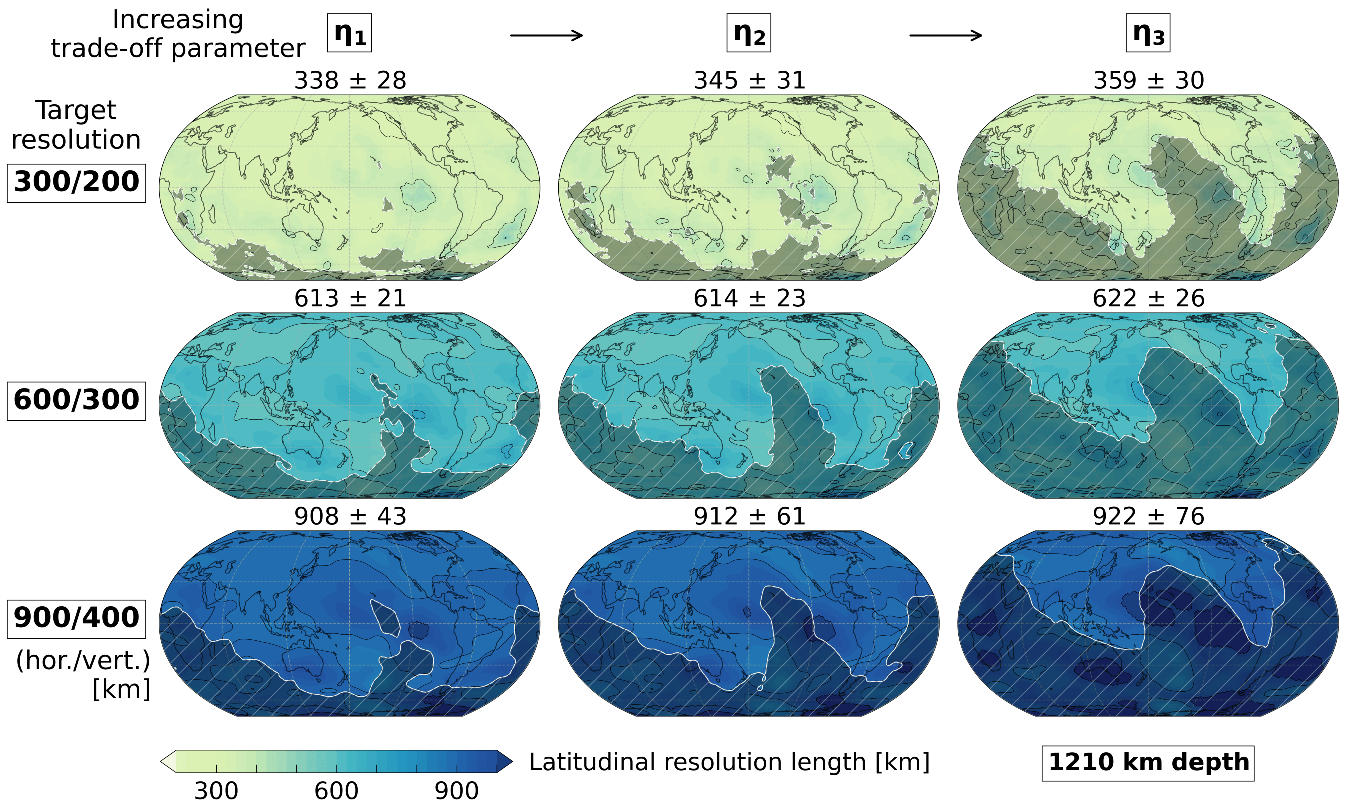

Figure C2. Latitudinal resolution lengths of the averaging kernels at 1210 km depth. Same layout as Fig. 11.