



Blockworlds 0.1.0: A demonstration of anti-aliased geophysics for probabilistic inversions of implicit and kinematic geological models

Richard Scalzo^{1,4}, Mark Lindsay^{2,4}, Mark Jessell^{2,4}, Guillaume Pirot^{2,4}, Jeremie Giraud^{2,4},
Edward Cripps^{3,4}, and Sally Cripps^{1,4}

¹School of Mathematics and Statistics, The University of Sydney, Darlington, NSW 2008, Australia

²Centre for Exploration Targeting, School of Earth Sciences, The University of Western Australia, Crawley, WA 6009, Australia

³Department of Mathematics and Statistics, The University of Western Australia, Crawley, WA 6009, Australia

⁴ARC Industrial Transformation and Training Centre in Data Analytics for Resources and the Environment (DARE)

Correspondence: Richard Scalzo (richard.scalzo@sydney.edu.au)

Abstract. Parametric geological models such as implicit or kinematic models provide low-dimensional, interpretable representations of 3-D geological structures. Combining these models with geophysical data in a probabilistic joint inversion framework provides an opportunity to directly quantify uncertainty in geological interpretations. For best results, the projection of the geological parameter space onto the finite-resolution discrete basis of the geophysical calculation must be faithful within the power of the data to discriminate. We show that naively exporting voxelised geology as done in commonly used geological modeling tools can easily produce a poor approximation to the true geophysical likelihood, degrading posterior inference for structural parameters. We then demonstrate a numerical forward-modeling scheme for calculating anti-aliased rock properties on regular meshes for use with gravity and magnetic sensors. Finally, we explore anti-aliasing in the context of a kinematic forward model for simple tectonic histories, showing its impact on the structure of the geophysical likelihood for gravity anomaly.

10 1 Introduction

Geological modeling of subsurface structures is critical to decision-making across numerous application areas, including mining, groundwater, resource exploration, natural hazard assessment, and engineering, yet is also subject to considerable uncertainty (e.g. Pirot et al., 2015; Linde et al., 2017; Witter et al., 2019; Quigley et al., 2019; Wang et al., 2020). Uncertainty arises in numerous places within the model-building workflow, including: sparse, noisy, and/or heterogeneous geological observations and rock property measurements; the indirect nature of geophysical measurements; the non-uniqueness of inverse problem solutions; and the ambiguity of geological interpretations (Lindsay et al., 2013). Rigorous quantification of uncertainty is therefore critical to decision-making informed by models, and is becoming an increasingly active area of research in geology and geophysics (Rawlinson et al., 2014; Bond, 2015; Jessell et al., 2018; Wellmann and Caumon, 2018). The reduction of geological uncertainty through fusion of complementary types of information or data, including geophysics, has long been of comparable interest (e.g. Wegener, 1923; Eckhardt, 1940; Li and Oldenburg, 2000; Guillen et al., 2008; Jessell et al., 2010; Bosch, 2016; Lindsay et al., 2020).



Bayesian methods provide a probabilistically coherent framework for reasoning about uncertainty, particularly in the context of beliefs about latent or unobserved entities and their causal influence on observations (Tarantola and Valette, 1982; Mosegaard and Tarantola, 1995; Mosegaard and Sambridge, 2002). Although not the only way to account for uncertainty in earth science settings, Bayesian reasoning enables the natural integration of heterogeneous data and expert knowledge (Bosch, 2016; Beardsmore et al., 2016), guides the acquisition of additional data for maximum information gain (e.g. Pirot et al., 2019b), enables selection among competing conceptual models (e.g. Brunetti et al., 2019; Pirot et al., 2019a), and optimizes management of risk in decision-making over possible outcomes (e.g. Varouchakis et al., 2019).

Bayesian methods are used in both geology and geophysics, where the question of how to combine disparate types of information depends upon the framing of the problem. In the geophysics community, where quantitative inversion frameworks have been in use for decades (Backus and Gilbert, 1967, 1968, 1970), the core problem is framed as a non-parametric imaging of the subsurface. The model parameters comprise a segmentation of the subsurface along with rock properties within each segment; applying a forward geophysics operator to this model produces synthetic data that can be compared to observations. Model structures rich enough to be useful are often very high-dimensional, and variations in spatial scales and resolutions make the full covariance matrix cumbersome and impractical to evaluate (Rawlinson et al., 2014). Inverse problems are notoriously ill-posed, and priors or regularizing terms must be included alongside the likelihood to target sensible regions of parameter space. Although the discrete geophysical image representation may naturally reflect the geological interpretation — for example, in the transdimensional inversions reviewed in Sambridge et al. (2012) — this is not usually the case, and much effort goes into constraining the outputs of geophysical inversions to be geologically reasonable.

These limitations surrounding interpretability and dimensionality are common to representations of geological structure in terms of very flexible models, including those that incorporate other geological concepts. Examples include: the hierarchical Bayesian "lithologic tomography" of Bosch (1999) and Bosch et al. (2001); structure-coupled multiphysics approaches (Gallardo and Meju, 2011; de Pasquale et al., 2019); and flexible models based on Voronoi tessellations or multi-point statistics (e.g. Cordua et al., 2012; Pirot et al., 2017). Level-set methods (Osher and Sethian, 1988; Santosa, 1996; Li et al., 2017; Zhiglova et al., 2018; Giraud et al., 2021a) solve for the boundaries of discrete rock units; these methods are more parsimonious than traditional volumetric inversions, but still invert for an explicit description of geological and geophysical properties. Uncertainty in an assumed geology can be incorporated into a geophysical inversion as a prior (Giraud et al., 2017, 2018, 2019a, b; Ogarko et al., 2021; Giraud et al., 2021b), although doing so does not directly address questions about the covariance structure or interpretability of the resulting model. Giraud et al. (2020) provide some assistance to model interpretation by demonstrating how neural networks can recover geology from the geophysical model; however, this work uses a very simple neural net and the classification process does not consider prior geological uncertainty.

In contrast, the core problem in geology is to interpret observations in terms of geological histories and processes (Frodeman, 1995). This becomes important in the structural geology of ore-forming systems, where details of the history become important and flexible treatment of 3-D geological structures as random fields is inadequate (Wellmann et al., 2017). The mineral exploration community has developed 3-D geological forward models, such as implicit and kinematic models, that can capture much of the complexity of real systems (e.g. Perrouty et al., 2014). Implicit models (Houlding, 1994) represent interfaces between



geological units as isosurfaces of a scalar field defined over 3-D space, which can be conditioned directly on structural geological observations by co-kriging (Lajaunie et al., 1997); the popular industry software package 3D GeoModeller (Calcagno et al., 2008), and the open-source code GemPy (de la Varga et al., 2018), are notable examples. Kinematic models (Jessell, 1981) use a parametrized description of tectonic events to forward-model a 3-D volume; the open-source package Noddy (Jessell and Valenta, 1993) is commonly used. The recently released open-source LoopStructural package (Grose et al., 2020) combines elements of implicit and kinematic models for additional geological richness and realism.

For models that interpolate structural data with no constraints from geophysics, geological uncertainty can be quantified by generating Monte Carlo realizations (Metropolis and Ulam, 1949) of geological datasets (Pakyuz-Charrier et al., 2018a, b, 2019). In a Bayesian treatment, this amounts to drawing samples from the prior (Wellmann et al., 2017).

The automated conditioning of geological forward models on geophysical observations simplifies uncertainty quantification by sampling from the posterior distribution in the low-dimensional geological parameter space. Uncertainty can then be quantified in the volumetric basis by generating the discrete geologies associated with the sampled geological parameters. Additionally, this approach naturally obeys constraints imposed by the forward geological process, and can be used for inference over that process. Implicit and kinematic models are well-suited for this approach since they formalize geological hypotheses in terms of event histories, but have simplified parametric forms that avoid the full computational burden of dynamic process-based simulations. Performing Bayesian inference over kinematic model parameters could open the door to more general Bayesian model selection over geological histories or conceptual models.

Sampling from Bayesian posterior distributions can be complicated, and consequently probabilistic inversion workflows that incorporate forward geophysics based on forward-modelled 3-D structural geology are, to date, still uncommon. However, the uptake of recent advances in Markov chain Monte Carlo methods ("MCMC"; Metropolis et al., 1953; Hastings, 1970) by the geological and geophysical community have begun to address this issue. The *Obsidian* distributed inversion code (McCalman et al., 2014) implemented a geological model of a layered sedimentary basin with explicit unit boundaries to explore geothermal potential (Beardsmore et al., 2016); it supported multiple geophysical sensors and used a distributed parallel-tempered MCMC sampler to draw from the multi-modal posteriors that may arise in underconstrained inverse problems. *Obsidian* has been extended with new within-chain MCMC proposals inside the parallel-tempered framework (Scalzo et al., 2019) and a new sensor likelihood for field observations of surface lithostratigraphy (Olierook et al., 2020); however, the limitations of its geological model make it an unlikely engine for general-purpose inversions. Wellmann et al. (2017) present a more general workflow that uses GeoModeller to render the geology, Noddy for calculation of geophysical fields, and pymc2 (Patil et al., 2010) for MCMC sampling. GemPy's support for information other than structural geological measurements is as yet fairly limited, but includes forward modeling of gravity sensors and topology.

All of the abovementioned workflows resample the 3-D geometry of geological units onto a discrete volumetric mesh for geophysical calculations. Discontinuous changes in rock properties that commonly occur in these models across unit interfaces then become undersampled, unless the mesh is chosen to align with those interfaces (Wellmann and Caumon, 2018). Since implicit and kinematic geological models naturally represent geology volumetrically, aligning mesh cell boundaries with geological interfaces that change with model parameters is not easily done without first evaluating the model, and incurs additional



computational overhead. Flexible geophysical inversions can adjust rock properties in individual boundary voxels to partially compensate for loss of accuracy, but this is not naturally true of geological forward models. Undersampled discontinuities in forward-modeled rock properties can lead to discontinuities in the geophysical likelihood, which hinders convergence of estimation methods, degrades posterior inference over structural parameters, and precludes the use of derivative information needed by advanced optimization or sampling methods.

Obsidian addresses this issue using an anti-aliasing algorithm which averages physical rock properties in voxels with partial unit membership — though this is not mentioned in associated publications (McCalman et al., 2014; Beardsmore et al., 2016) and the implementation depends upon a specific parametrization of geological structure. Anti-aliasing to address undersampling has a long history in computer graphics (Crow, 1977; Catmull, 1978; Cook, 1986; Öztireli, 2020) and in geophysics, for example finite-difference solutions of seismic wave equations (Backus, 1962; Muir et al., 1992; Moczo et al., 2002; Capdeville et al., 2010; Koene et al., 2021).

This paper describes a generalized numerical anti-aliasing scheme for implicit and kinematic geological models discretized onto fixed three-dimensional meshes. We demonstrate the principle with a toy kinematic modeling code, Blockworlds, written in Python 3. Section 2 explains the aliasing effect, including its influence on the likelihood, and describes the algorithm used in Blockworlds to address it. Section 3 describes how Blockworlds evaluates kinematic event histories, then introduces a set of synthetic kinematic geological models we use as benchmarks. Section 4 summarizes gravity inversion experiments based on these models, demonstrating the influence of anti-aliasing on MCMC sampling. Section 5 discusses future directions, including use with other geophysical sensors and curvilinear coordinate systems, and we conclude in Section 6.

110 2 Methods

2.1 Chaining forward models in geophysics calculations from structural geology

The forward modeling of observations falls under the calculation of the likelihood $p(\mathbf{d}|\boldsymbol{\theta})$, which describes the probability of observing data \mathbf{d} given that the causal process generating the data has true parameters $\boldsymbol{\theta}$. A Bayesian inversion proceeds by combining the likelihood with a prior $p(\boldsymbol{\theta})$ that describes the strength of belief, in terms of probability, that the causal process has true parameters $\boldsymbol{\theta}$ before any data are taken into account. In non-parametric geophysical inversions the prior includes regularization terms that penalize undesirable solutions. In parametric inversions, the choice of parametrization corresponds to a strong regularization on the discretized geology, with the prior distribution providing further constraints. The posterior distribution $p(\boldsymbol{\theta}|\mathbf{d})$, describing probabilistic beliefs about $\boldsymbol{\theta}$ once the data have been taken into account, is then determined through Bayes's theorem:

$$120 \quad p(\boldsymbol{\theta}|\mathbf{d}) \propto p(\mathbf{d}|\boldsymbol{\theta})p(\boldsymbol{\theta}). \quad (1)$$

For parameter estimation, the unnormalized right-hand side can then be sampled using MCMC algorithms for full uncertainty quantification.



While some measurements, such as structural observations, can be computed directly from a continuous functional form for the geological forward model, likelihoods based on simulation of geophysical sensors may require discretization of the geology to calculate. The geophysical likelihood terms in this case each involve the composition of two forward models:

1. a mapping $g : \Theta \rightarrow \mathcal{G}$ from the parameter space Θ into a space \mathcal{G} of discrete volumetric representations (such as block models); and
2. a mapping $f : \mathcal{G} \rightarrow \mathcal{D}$ from the space of discretizations into the space of possible realizations of the data.

Structural geology is chiefly concerned with the map g , and often reckons uncertainty not in terms of parameter variance, but in terms of the properties of an ensemble of discrete voxelized models (e.g. Wellmann and Regenauer-Lieb, 2012; Lindsay et al., 2012). Inversions of geophysical sensor data such as gravity, magnetism, conductivity, or seismic are chiefly concerned with the map f .

An exact representation of the geology in discrete form is rarely possible, so that the accuracy of the inference will depend on the accuracy with which $f \circ g$ can be computed. The fidelity of discrete geological representations is important to both the geology and geophysics communities, and is reflected for example in the use of adaptive meshes that respect the strength of data constraints (Rawlinson et al., 2014) or the use of basis functions that lie along geological features (Wellmann and Caumon, 2018). Existing tools, however, frequently export to a simple fixed basis that may be adequate for visualization in an interactive workflow, but cause trouble in automated inversion workflows, producing the aliasing effects discussed in the next subsection.

2.2 Aliasing and its effects on the posterior

To show a straightforward example, we calculate a grid scan through the posterior for inversion of forward gravity from a uniform-density spherical intrusion. An analytic solution exists, so we can compare directly to the true posterior. We model a 1 km^3 cubical volume, and fix the sphere's position at the center. Gravity anomaly data were generated from the analytic model on an evenly spaced 10×10 survey grid at the surface ($z = 0$), with Gaussian measurement noise at the level of 5% of the signal amplitude added. We used the SimPEG library (Cockett et al., 2015) to generate meshes and calculate the action of the forward model for the gravity sensor in the inversion loop. To emphasize the role of the likelihood in a scenario with vague prior constraints, we use uniform priors on the mass density ρ and radius R ($300 \pm 100 \text{ m}$) of the sphere. Given spherical symmetry, the data constrain only the total mass M of the intrusion, and so the region of high posterior probability density follows the curve

$$\rho = \frac{3M}{4\pi R^3}. \quad (2)$$

which the discrete posterior should reflect.

Figure 1 shows the results of this exercise. Four discretizations of the sphere are shown: a coarse mesh, with the rock density for each cell queried from the true geology at the center of that cell; a higher-resolution mesh, with cell sizes refined by a factor of 4 along each axis; a coarse mesh on which the rock density has been averaged throughout the cell, using the anti-aliasing scheme described in the next section; and a high-resolution mesh using the same anti-aliasing scheme. The gravity fields are

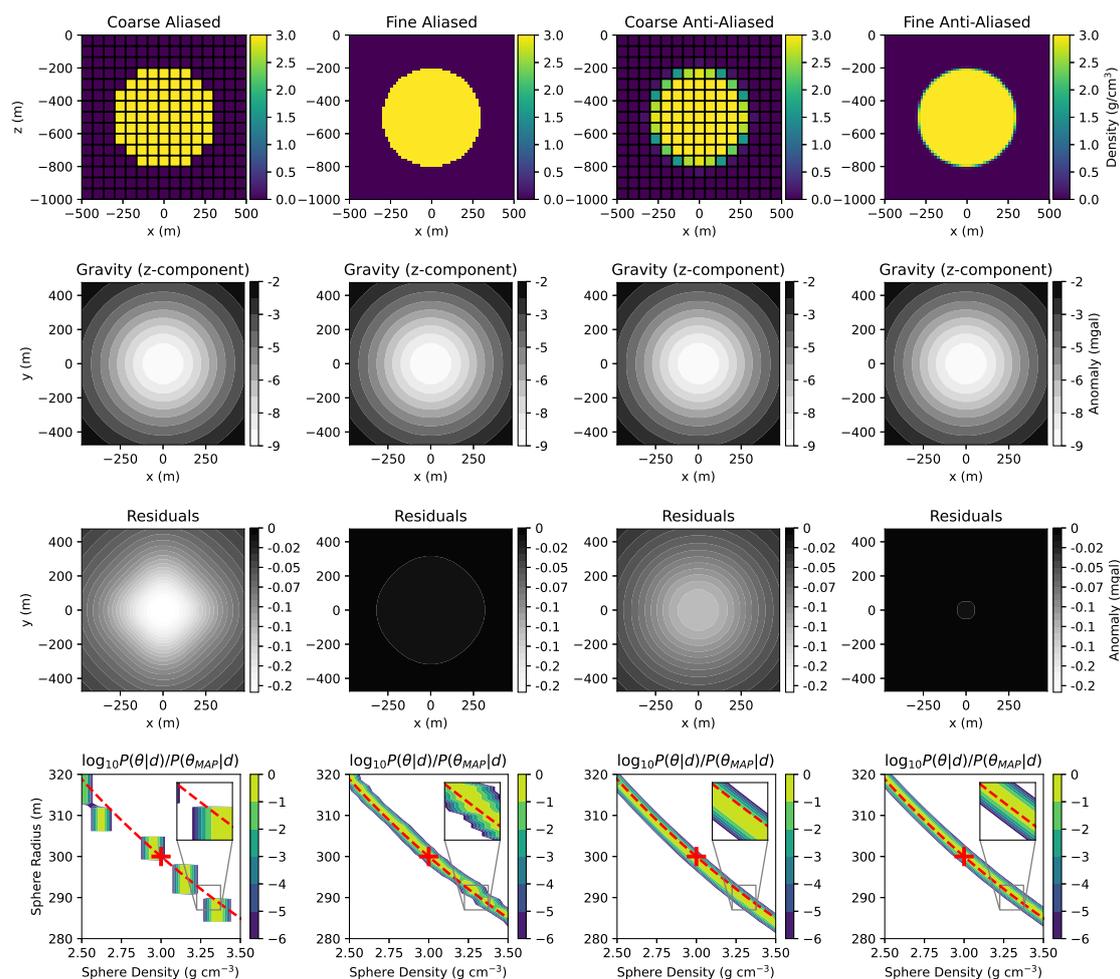


Figure 1. Calculation of posterior distribution for the radius and density of a uniform spherical inclusion from gravity inversion. Top row: cross-section of 3-D rock density field for the true model from which the data are generated. 2nd row: Simulated gravity field at surface. 3rd row: Residuals of simulated gravity field from analytic solution. Bottom row: cross-section through the posterior for an independent Gaussian likelihood. Discretization schemes shown (columns, left to right): coarse aliased mesh (15^3 cells), fine aliased mesh (60^3 cells), coarse anti-aliased mesh, fine anti-aliased mesh. The true parameters are shown by the red cross, while the red dashed line indicates the degeneracy relation in Eq. 2 that constitutes the maximum-likelihood ridge of the analytic solution.



155 visually indistinguishable for all four versions, but the posteriors look very different. When inverting on the coarse mesh with
no anti-aliasing, discontinuous changes in rock properties occur whenever the sphere radius crosses over the centre of one or
more mesh cells. This in turn induces spurious structure into the geophysical likelihood, and hence the posterior distribution
of the inversion parameters: the numerical posterior becomes a collection of isolated modes, scattered along the locus of
degeneracy between ρ and R (red dashed line) on which the true parameters (red cross) lie, but none of which appropriately
160 quantify the uncertainty in the true analytic problem. These artefacts can be suppressed by refining the mesh, though at the cost
of increasing the number of voxels (and the computation time); in this case, by a factor of $4^3 = 64$. Importantly, the underlying
posterior for this high-resolution mesh is not free of discontinuities, but takes on a terraced appearance.

The inversions with anti-aliased geology result in continuous, even smooth, posteriors that trace the analytic solution and
run at nearly the same speed. A closer look reveals that the coarse-mesh anti-aliased posterior is slightly offset toward the
165 bottom-left relative to the analytic curve, while the fine-mesh posterior is not. This is a result of the approximation used when
anti-aliasing the voxel density, which treats each interface locally as a plane surface (see Sect. 5.3). For a given gravity signal,
spheres with higher density have smaller radius and thus higher curvature at the interface; the departure from the true posterior
will grow as the sphere radius approaches the voxel size. The bias amounts to $\sim 1\%$ of the sphere's mass at the true parameter
values, which the data are just sufficient to detect.

170 The accuracy of the geophysical field calculation is not dramatically improved by anti-aliasing, but the posterior density
becomes much more navigable. This highlights another critical feature of the problem, which is that aliasing produces effects
distinct from mis-specification of the parametric geological model. Model calibration methods (e.g. Kennedy and O'Hagan,
2001) that address correlated model residuals will not remove discontinuities from aliasing.

While this example may seem artificial, widely used geological modeling tools such as Noddy and GeoModeller usually
175 export on rectangular meshes, with lithology or rock properties evaluated at the center of each mesh cell precisely as described
above. Any ongoing development or support of geological modeling tools intended for use in probabilistic workflows should
keep examples like this in mind. Similar issues may afflict geophysical inversions that use such exported geologies as priors:
although the posterior mean of the inversion output is a flexible function of space, a biased mean function in the form of an
aliased geological prior may still influence the inversion in the neighborhood of discontinuous contacts. Complex structure
180 imparted by faults, folding, dykes and sills is highly sensitive to discretization. The situation may arise where the causative
body for a strong geophysical anomaly, such as a thin and relatively magnetic or remanently magnetised dyke, is removed
from the geological prior due to overly coarse discretization parameters, while a resulting strong magnetic response remains.
Inversion schemes are not intended to address situations where the target is unintentionally removed from the data.

2.3 Fitting an anti-aliasing function

185 The anti-aliased discretization is not, in volumetric terms, as accurate a description of the underlying geology as the high-
resolution mesh. Instead, the anti-aliasing prescription captures the response of the sensor to the underlying parametrized
geology. While there could in principle be many configurations of rock that result in the same cell-averaged rock properties,
the parametric form of the geological model identifies the specific possibilities of interest and strongly constrains the inversion.



In a joint inversion of multiple sensors, different sub-mesh prescriptions or different meshes may be appropriate for the different sensors, as long as these faithfully represent the sensor response to the infinite-resolution geology and can be efficiently calculated given the geological parameters.

In the case of gravity and magnetic anomaly, calculating forward geophysics on a fixed volumetric mesh is convenient: first, it is conventional for current geology engines, and second, the sensitivities can be cached for fast forward calculations. Potential-field sensors are linear sensors that represent a spatial average of rock properties, so averaging rock properties over voxels with partial membership in multiple units is the desired prescription. We expect the true sensor response to be a continuous, perhaps even smooth, function of the geological parameters that describe the interface geometry, except perhaps in regions where a geological contact pinches a unit to zero thickness. Thus we seek a smooth interpolant to be evaluated at the centers of mesh cells as an approximation to a numerical average. This will also make the numerical posterior a smooth function of the parameters, enabling future use of advanced MCMC proposals that require derivative information, such as Hamiltonian Monte Carlo (Neal, 2011).

To fit for such an interpolating function, we approximate the interface between two units of constant density as a plane at the scale of a single voxel. Such a plane is defined by a point $\mathbf{r} = (r_x, r_y, r_z)$ on its surface, and a unit normal $\mathbf{n} = (n_x, n_y, n_z)$ (see Fig. 2). Exploiting the symmetries of the cube, we reduce these to three main predictive features:

$$u_{\perp} = (\mathbf{r} - \mathbf{r}_0) \cdot \mathbf{n} / h \quad (3)$$

$$\eta = \sqrt{1 - n_{\max}^2} / n_{\max} \quad (4)$$

$$\xi = n_{\min} / \sqrt{1 - n_{\max}^2} \quad (5)$$

where h is the voxel size, $n_{\min} = \min\{n_x, n_y, n_z\}$, $n_{\max} = \max\{n_x, n_y, n_z\}$, and \mathbf{r}_0 is the voxel center. For training, we generate a synthetic dataset of 1000 (\mathbf{r}, \mathbf{n}) pairs with which a plane interface might intersect a generic voxel; in each pair, \mathbf{r} is distributed uniformly in space, and \mathbf{n} uniformly in solid angle. We calculate the partial voxel volume beneath the plane numerically by mesh refinement to $20\times$ higher resolution.

Figure 2 shows three different possible functional forms for the interpolant: a piecewise linear function of u_{\perp} alone; a linear regression model including interaction terms up to third order, where the regressors have been transformed by a hyperbolic tangent function to maintain boundary conditions and keep fractional volumes between 0 and 1; and a Gaussian process regression. A linear regression model based on the single feature u_{\perp} — the normal distance to the plane from the voxel center — can provide smooth anti-aliasing to 1% RMS accuracy (5% worst-case); the optimized form is

$$v(u_{\perp}) = \frac{1}{2} [1 + \tanh(2.2u_{\perp} + 3.2u_{\perp}^3)] \quad (6)$$

This functional form provides negligible computational overhead in our implementation as part of a MCMC inversion loop, and has derivatives of all orders. Including orientation features in the linear regression results in negligible improvements in accuracy while increasing computation time. The Gaussian process does not provide meaningful improvement over the linear regression using u_{\perp} alone; it attains accuracy of 0.3% RMS (3% worst-case) using three features, but is many times slower to evaluate. Other models involving techniques such as functional PCA might give improved results, but all three interpolants

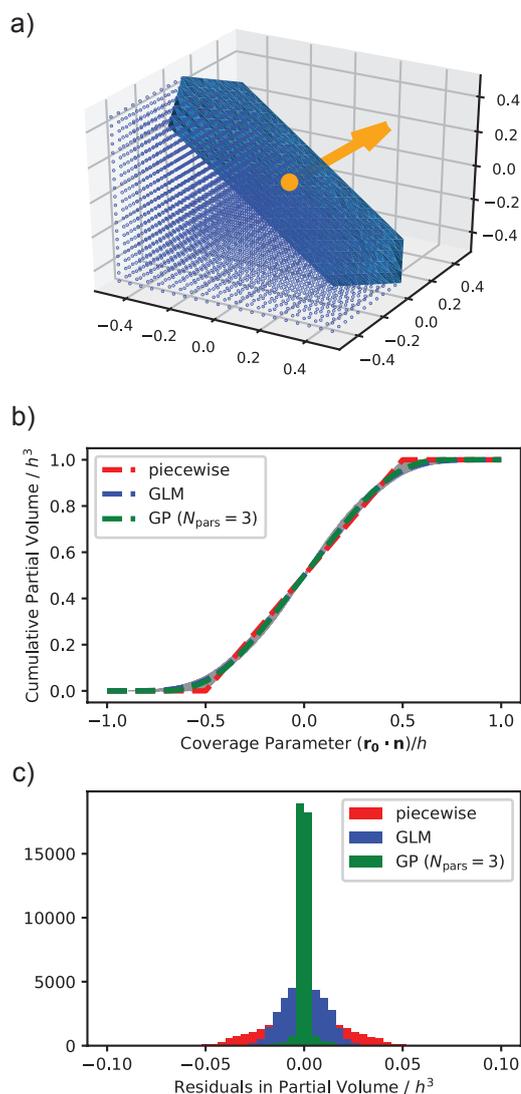


Figure 2. (a) Numerical approach to calculating partial volumes. The interface is approximated by a plane running through the voxel. A cloud of points is generated in a regular grid, and the partial volume is then the fraction of these points that satisfy $u_{\perp} = (\mathbf{r} - \mathbf{r}_0) \cdot \mathbf{n}/h = 0$, where $\mathbf{r}_0 = (r_x, r_y, r_z)_0$ is a point on the interface and $\mathbf{n} = (n_x, n_y, n_z)$ is a unit normal vector. (b) Regression models for fast evaluation of partial volumes using data generated from the approximate numerical scheme (gray curves): a piecewise linear interpolation over u_{\perp} (red), a generalized linear model with up to third-order terms (blue), and a Gaussian process that also includes features associated with the orientation of the interface (green). (c) Precision of partial volume models. The piecewise model reaches 2% RMS accuracy. A linear regression reaches 1% accuracy, as does a Gaussian process regression with a single feature u_{\perp} . A GP with three features that captures the directionality of the interface with respect to the cube boundaries reaches 0.3% accuracy.



provide satisfactory performance for our purposes. To the extent that small numerical errors are randomized across multiple interface voxels, their contributions should cancel out. We use Eq. 6 for further experiments in this paper.

2.4 Anti-aliasing for implicit models

225 For implicit geological models, each interface is defined as the level set $\Phi(\mathbf{r}) = \Phi_0$ of some scalar field Φ . In contemporary geological modeling suites this field represents an interpolation of structural measurements (Lajaunie et al., 1997; Calcagno et al., 2008). The value of Φ corresponds roughly to depth or to geological time, but has no intrinsic physical meaning except to ensure that the geologic series of interfaces it defines are conformable; the specific value Φ_0 representing an interface is fixed by observations (from surface measurements or boreholes). In this case, u_{\perp} can be calculated easily by locally calibrating the
230 scalar field to represent a physical distance in the neighborhood of an interface: for a voxel centered at \mathbf{r}_0 , a Taylor expansion of Φ gives

$$\Phi(\mathbf{r}) = \Phi(\mathbf{r}_0) + (\mathbf{r} - \mathbf{r}_0) \cdot \nabla\Phi + O(\mathbf{r} - \mathbf{r}_0)^2 = \Phi_0, \quad (7)$$

which, since $\mathbf{n} = \nabla\Phi/|\nabla\Phi|$, can be rearranged to give

$$u_{\perp} = (\Phi_0 - \Phi(\mathbf{r}_0))/|\mathbf{J}\nabla\Phi|. \quad (8)$$

235 In this way, each voxel can be anti-aliased simply by evaluating the field and its gradient at the voxel center.

3 Experiments with kinematic models

To illustrate the impact of anti-aliasing on more realistic geologies, we implement anti-aliased geology within a simplified kinematic model, patterned after Noddy (Jessell, 1981; Jessell and Valenta, 1993). We chose to write our own demonstration code rather than modifying Noddy both for ease of prototyping, and to demonstrate how anti-aliasing interacts with the kine-
240 matic calculations. We then construct a range of Bayesian 3-D models and visualize slices through each model posterior to demonstrate the influence of anti-aliasing on the calculations.

3.1 Kinematic model events

The calculation begins with a basement layer of uniform density, $g_0(\mathbf{r}|\rho_0) = \rho_0$. For non-trivial geologies, we parametrize each subsequent event by a collection of parameters θ_i , and let $\Theta_i = \{\theta_0, \theta_1, \dots, \theta_i\}$ denote the collection of parameters for
245 all events up to event i . Each event operates recursively on the geologies that came before it. The action of the anti-aliasing operator A for an interface $\Phi(\mathbf{r}) = \Phi_0$ between two geological units with rock properties g_+ and g_- is given by

$$A[\Phi(\mathbf{r}); g_+, g_-] = g_- + (g_+ - g_-)v(u_{\perp}) \quad (9)$$

where $v(u_{\perp})$ and u_{\perp} are given by Eq. 6 and Eq. 8 with voxel size h . The anti-aliasing must be applied in the co-moving frame of each event, to ensure that voxels that are modified by multiple transformations are averaged sequentially for each event.



250 The Blockworlds code covers four elementary event types, each parametrized by its own scalar field that takes a zero value at the relevant interface:

– **Stratigraphic layer** — $\theta_i = (\Delta z_i, \rho_i)$:

$$g_i(\mathbf{r}|\Theta_i) = A[\mathbf{r} \cdot \mathbf{u}_z - \Delta z_i; g_{i-1}(\mathbf{r} - \Delta z_i \mathbf{u}_z | \Theta_{i-1}), \rho_i] \quad (10)$$

– **Fault** — $\theta_i = (\mathbf{r}_{0,i}, \mathbf{n}_i, s_i)$:

255
$$g_i(\mathbf{r}|\Theta_i) = A[(\mathbf{r} - \mathbf{r}_{0,i}) \cdot \mathbf{n}_i; g_{i-1}(\mathbf{r} + s\mathbf{v} | \Theta_{i-1}), g_{i-1}(\mathbf{r} | \Theta_{i-1})] \quad (11)$$

where \mathbf{r}_0 is a point on the fault, \mathbf{n} is a unit normal direction to the fault, and \mathbf{v} is a slip vector. In our parametrization we constrain \mathbf{r} to lie at zero depth, we parametrize \mathbf{n} in polar representation derived from dip and dip direction (Pakyuz-Charrier et al., 2018b), and we fix $v = \frac{(\mathbf{u}_z \times \mathbf{n}) \times \mathbf{n}}{|(\mathbf{u}_z \times \mathbf{n}) \times \mathbf{n}|}$ so that s describes dip-slip displacement only with no strike-slip component.

260 – **Fold** — $\theta_i = (\mathbf{n}_i, \psi_i, \phi_i, L_i, B_i)$:

$$g_i(\mathbf{r}|\Theta_i) = g_{i-1}(\mathbf{r} + \Delta \mathbf{r} | \Theta_{i-1}) \quad (12)$$

where $\{\mathbf{v}_0, \mathbf{v}_1, \mathbf{n}_i\}$ form an orthonormal frame including the fold axis \mathbf{n} , and

$$\begin{aligned} \Delta \mathbf{r} &= B_i \sin(2\pi(\mathbf{r} \cdot \mathbf{n}_i)/L + \phi) \\ &\times (\sin(\psi_i)\mathbf{v}_0 + \cos(\psi_i)\mathbf{v}_1) \end{aligned} \quad (13)$$

265 – **Spherical intrusion** — $\theta_i = (\mathbf{r}_{0,i}, R_i, \rho_i)$:

$$g_i(\mathbf{r}|\Theta_i) = A[|\mathbf{r} - \mathbf{r}_{0,i}| - R_i; \rho_i, g_{i-1}(\mathbf{r} | \Theta_{i-1})] \quad (14)$$

Figures 3 and 4 show the action of the kinematic elements of the model alongside an anti-aliased version of the voxelized density volume, calculated over a 1 km^3 rectilinear mesh $15 \times 15 \times 15$ voxels on a side.

3.2 Setup of specific 3-D models

270 We construct a series of eight kinematic models, with true configurations shown in Fig. 3 and Fig. 4. Each model occupies a 1 km^3 cubical volume, and is discretized on a rectilinear mesh with cubical voxels. All have three stratigraphic layers and two planar dip-slip faults, but the relative positioning of the faults differs.

The final configuration of Model 1 represents a typical graben structure with two intersecting faults, both exhibiting normal movement (see Fig. 3). The final configuration of Model 2 represents a negative flower structure resulting from trans-tension,
 275 where both a sub-vertical fault and high-angle fault exhibit normal movement. The final configuration of Model 3 represents



Table 1. Parameter true values and prior distributions for each kinematic model, labeled M1 to M8. The prior mean is set to the true value for each parameter.

Parameter	M1	M2	M3	M4	M5	M6	M7	M8	Prior
Basement Density (g cm^{-3})	3.0	3.0	3.0	3.0	3.0	3.0	3.0	3.0	Lognormal ($\sigma = 0.1 \text{ g cm}^{-3}$)
Layer 1 Thickness (m)	350	350	350	350	350	350	350	350	Lognormal ($\sigma = 50 \text{ m}$)
Layer 1 Density (g cm^{-3})	2.5	2.5	2.5	2.5	2.5	2.5	2.5	2.5	Lognormal ($\sigma = 0.1 \text{ g cm}^{-3}$)
Layer 2 Thickness (m)	190	190	190	190	190	190	190	190	Lognormal ($\sigma = 50 \text{ m}$)
Layer 2 Density (g cm^{-3})	2.0	2.0	2.0	2.0	2.0	2.0	2.0	2.0	Lognormal ($\sigma = 0.1 \text{ g cm}^{-3}$)
Fold 1 Axis Elevation (deg)	—	—	—	—	—	—	—	0	vMF ($\kappa = 25$)
Fold 1 Axis Azimuth (deg)	—	—	—	—	—	—	—	0	vMF ($\kappa = 25$)
Fold 1 Pitch Angle (deg)	—	—	—	—	—	—	—	0	Normal ($\sigma = 1 \text{ deg}$)
Fold 1 Phase Angle (deg)	—	—	—	—	—	—	—	0	Normal ($\sigma = 1 \text{ deg}$)
Fold 1 Wavelength (m)	—	—	—	—	—	—	—	1000	Lognormal ($\sigma = 100 \text{ m}$)
Fold 1 Amplitude (m)	—	—	—	—	—	—	—	100	Lognormal ($\sigma = 15 \text{ m}$)
Fault 1 Contact X Position (m)	-400	-450	-50	-250	-450	-400	-400	250	Normal ($\sigma = 1 \text{ m}$)
Fault 1 Contact Y Position (m)	0	0	0	0	0	0	0	0	Normal ($\sigma = 1 \text{ m}$)
Fault 1 Polar Elevation (deg)	45	45	-20	0	30	20	20	-205	vMF ($\kappa = 25$)
Fault 1 Polar Azimuth (deg)	0	0	0	0	0	0	0	0	vMF ($\kappa = 25$)
Fault 1 Dip-Direction Slip (m)	-220	-220	-220	-220	220	220	140	200	Normal ($\sigma = 150 \text{ m}$)
Fault 2 Contact X Position (m)	400	50	50	250	50	-300	-300	—	Normal ($\sigma = 1 \text{ m}$)
Fault 2 Contact Y Position (m)	0	0	0	0	0	0	0	—	Normal ($\sigma = 1 \text{ m}$)
Fault 2 Polar Elevation (deg)	-45	20	20	0	10	40	40	—	vMF ($\kappa = 25$)
Fault 2 Polar Azimuth (deg)	0	0	0	0	0	0	0	—	vMF ($\kappa = 25$)
Fault 2 Dip-Direction Slip (m)	220	220	220	220	-220	-220	80	—	Normal ($\sigma = 150 \text{ m}$)

two reverse faults dipping at the same angle and away from each other. The final configuration of Model 4 represents two parallel sub-vertical faults of opposing displacement: the left fault with left-side-up movement, the right fault with right-side-up movement. The final configuration of Model 5 represents a positive flower structure resulting from transpression, where both a sub-vertical fault and high-angle fault exhibit reverse movement (see Fig. 4). The final configuration of Model 6 represents two high-angle faults, the left displaying thrust movement and the right displaying normal movement. The final configuration of Model 7 represents a similar scenario as Model 6, except the movement for both faults is reverse. The final configuration of Model 8 diverges from the two-fault scheme of the previous models by including a folding event. The stratigraphic history remains the same, however is subject to upright and open folding of 500 m wavelength. A west-dipping, high-angle reverse fault offsetting the fold limbs completes the scenario.



285 We also choose prior distributions for the parameters to simulate the realistic incorporation of structural geological knowl-
 edge, as shown in Table 1. We use maximum-entropy distributions — normal, lognormal (for non-negative quantities), and
 von-Mises-Fisher (for angular quantities) — for simplicity and interpretability in terms of scale parameters, which reflect the
 precision of available structural information. The positions of the anchor points lying on the fault planes have very narrow priors
 ($\sigma = 1$ m), appropriate for surface observation localized by GPS. The fault slip parameters reflect rough estimates ($\sigma = 150$ m),
 290 an extent comparable to the true slip value for each fault. The polar representations of the fault directions are constrained by a
 von Mises-Fisher (vMF) distribution with $\kappa = 25$, corresponding to a full width at half maximum of about 16 degrees (Pakyuz-
 Charrier et al., 2018b). In each case the prior mode rests at the true parameter values, to ensure that mis-specification of the
 prior does not complicate the aliasing effects we set out to examine.

For each model, we examine four regimes of resolution and aliasing. We consider a "coarse" $15 \times 15 \times 15$ mesh ($h = 66.6$ m)
 295 as well as a "fine" $75 \times 75 \times 75$ mesh ($h = 13.3$ m), both with and without anti-aliasing. The anti-aliased model on the fine
 mesh is taken to be the true model for the purposes of calculating a common synthetic forward gravity dataset against which
 to evaluate the likelihood.

We generate synthetic geophysics data based on the true model parameters, with measurements spaced evenly in a 20×20
 grid at the surface (grid spacing of 50 m), and add independent Gaussian noise with a standard deviation σ_0 of 5% of the
 300 sample standard deviation of the data across the survey. Each data point y_j is therefore generated according to

$$p(y_j|\boldsymbol{\theta}, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{(y_j - f_j(\boldsymbol{\theta}))^2}{2\sigma^2}\right]. \quad (15)$$

where $f_j(\boldsymbol{\theta})$ is the forward model for the geophysical measurement y_j given the geological parameters $\boldsymbol{\theta}$. In a realistic situation,
 we may not know the noise variance σ^2 exactly, but we can account for our uncertain knowledge over it by including a
 hierarchical prior for σ^2 . If we choose this prior to be an inverse gamma distribution,

$$305 \quad p(\sigma^2|\alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} \sigma^{-2(\alpha+1)} \exp\left(-\frac{\beta}{\sigma^2}\right), \quad (16)$$

the integral over σ^2 can be solved analytically, saving the computational expense of sampling over σ^2 by MCMC (see McCal-
 man et al., 2014; Scalzo et al., 2019):

$$\begin{aligned} p(y_j|\boldsymbol{\theta}, \alpha, \beta) &= \int_0^\infty p(y_j|\boldsymbol{\theta}, \sigma^2) p(\sigma^2|\alpha, \beta) d\sigma^2 \\ &= \frac{\Gamma(\alpha + \frac{1}{2})}{\Gamma(\alpha)\sqrt{2\pi\beta}} \left[\frac{(y_j - f_j(\boldsymbol{\theta}))^2}{2\beta} + 1 \right]^{-(\alpha + \frac{1}{2})} \end{aligned} \quad (17)$$

310 so that the likelihood has the form of a t -distribution with $\nu = 2\alpha$ degrees of freedom and scale parameter $\sqrt{\beta/\alpha}$. For our
 experiments we choose $\alpha = 2.5$, $\beta = 2.5 \times \sigma_0^2$, a weakly informative prior with mean and mode close to the true variance used
 to generate the data. The full likelihood is then the product of the likelihoods for independent data points,

$$p(\mathbf{d}|\boldsymbol{\theta}, \alpha, \beta) = \prod_{j=1}^N p(y_j|\boldsymbol{\theta}, \alpha, \beta). \quad (18)$$



As with the earlier synthetic examples for a spherical intrusion, we used SimPEG (Cockett et al., 2015) for calculation of forward gravity. The SimPEG `Simulation3DIntegral` class uses a closed-form integral expression for the contribution to the gravitational field from a rectangular prism of constant density (Okabe, 1979; Li and Oldenburg, 1998): for example, for the z -component,

$$g_z(\mathbf{r}) = -G\rho_0 \sum_{i=1}^2 \sum_{j=1}^2 \sum_{k=1}^2 \mu_{ijk} \times \left[x'_i \ln(y'_i + r'_{ijk}) + y'_j \ln(y'_j + r'_{ijk}) + 2z'_k \arctan\left(\frac{x'_i y'_j}{z'_k r'_{ijk}}\right) \right] \quad (19)$$

where $\mathbf{r} = (x, y, z)$ is the location of the gravity sensor, $\mathbf{r}_{ijk} = (x_i, y_j, z_k)$ runs over each corner of the prism, and $\mathbf{r}' = \mathbf{r} - \mathbf{r}_{ijk}$ with $r' = |\mathbf{r}'|$. This form is widely used in geophysics and is one of the benefits of working over rectangular meshes. The total gravity signal is then the sum of the contributions for each mesh cell. Since the expression is linear in the density for each mesh cell, the forward model sensitivities can be cached for fast likelihood evaluation.

3.3 Posterior slices and MCMC sampling

To illustrate some of the effects caused by aliasing and some of the limits of anti-aliasing, we visualize a series of two-dimensional slices through the posterior distribution. Pairs of free parameters are scanned on a regular 30×30 grid centered on their true values, with all other parameters held fixed at their true values. This produces a set of figures similar to Fig. 1 for the kinematic models.

We also perform MCMC over the parameters of the kinematic models to demonstrate the impact of aliasing on chain mixing. Since the aliased posterior is effectively piecewise constant, it has no useful derivative information, so we are limited to random walk proposals. Due to the strong covariances seen in the conditional posterior slices, we choose the adaptive Metropolis random walk algorithm of Haario et al. (2001), which starts from an initial guess θ_0 and proposes a new value θ' at time step t from a multivariate normal distribution centered on the current state θ_t :

$$q(\theta' | \theta_t) \sim N(\theta_t, \eta_d \Sigma_t), \quad (20)$$

where η_d is a global step size parameter depending only on the dimension of the space, and

$$\Sigma_t = \begin{cases} \Sigma_0 & \text{for } t \leq t_0 \\ \text{cov}\{\theta_i\} & \text{for } t > t_0 \end{cases} \quad (21)$$

The covariance matrix of the proposal is estimated from the current chain history, so that after an initial non-adaptive random walk period of length t_0 steps, the walk gradually transitions from an initial-guess covariance Σ_0 to match the estimated posterior covariance. Although this process depends upon the past chain history and thus is not strictly Markov, the preconditioning of the random walk steps approaches a fixed value as the number of steps increases. A more rigorous definition of "vanishing adaptation" (Andrieu and Thoms, 2008), which this method satisfies, guarantees convergence to the true posterior distribution.



Since this method is still a Metropolis-Hastings method, it uses the usual acceptance criterion $\theta_{t+1} = \theta'$ with probability

$$\alpha = \min \left\{ 1, \frac{p(\mathbf{d}|\theta')p(\theta')q(\theta|\theta')}{p(\mathbf{d}|\theta)p(\theta)q(\theta'|\theta)} \right\} \quad (22)$$

345 and $\theta_{t+1} = \theta_t$ with probability $1 - \alpha$. This proposal can become maladapted for posteriors with locally varying covariance structure (Scalzo et al., 2019); we use it here since it is commonly employed and also for comparison with contemporary studies such as Wellmann et al. (2017). We use an adaptation length of $t_0 = 1000$ and an initial diagonal covariance matrix of step sizes set to 20% of the prior width for each parameter (0.02 g cm^{-3} for densities, 10 m for layer thicknesses, 30 m for fault slips and 3 degrees for angular variables).

350 We use the coarse $15 \times 15 \times 15$ mesh for calculations, since this will give much faster results for MCMC and highlight the impact of aliasing on chain mixing. We run $M = 4$ chains for $N = 10^5$ samples from the posterior of each model, with and without anti-aliasing. We initialize each chain with an independent prior draw, and discard the first 20% of samples as burn-in. Each chain took 6 minutes to run on an Intel Core i7 2.6 GHz processor.

We use the integrated autocorrelation time to measure MCMC efficiency within each chain, and the potential scale reduction factor (PSRF; Gelman and Rubin, 1992) to estimate the extent of convergence to the true posterior. The autocorrelation time τ is the approximate number of chain iterates required to produce a single independent sample from the posterior, so that the chain length N divided by τ becomes one possible measure of effective sample size. The potential scale reduction factor \hat{R} , formed from multiple chains, is the factor by which the posterior variance estimated from those chains could be further reduced by continuing to sample. A PSRF near 1 suggests that multiple chains from different starting conditions have achieved similar means and variances, and thus are sampling from a common stationary distribution; a large PSRF shows a residual dependence on starting conditions. The PSRF is usually given for each parameter in a multi-dimensional chain; some dimensions may mix quickly while others take a long time to converge.

4 Results

4.1 Posterior cross sections

365 Two-dimensional slices through the posterior distribution for selected variable pairs are shown for Model 1 (in Fig. 5) and Model 6 (in Fig. 6). Similar sets of plots for all eight models are available in the Supplementary Material.

Model 1 is relatively well-behaved; the aliased posterior has a single dominant mode near the true value. This is a situation in which the aliasing effects function mainly to make the likelihood appear blocky and terraced. Although proposals that require derivatives of the likelihood will fail on this posterior, it can still be navigated by appropriately tuned random walks, or by the discontinuous Hamiltonian Monte Carlo sampler of Nishimura et al. (2019).

370 Nevertheless, the benefits of anti-aliasing are still clear. Merely increasing the mesh resolution of the coarse aliased model by a factor of 5 (and the computational cost by over $100\times$) is insufficient to completely suppress aliasing artefacts. The coarse aliased model fails to capture the full extent of the very strong correlation between the dip-slips of the two faults, though the



375 fine aliased model does so reasonably well. The posteriors of the two anti-aliased models deviate only negligibly from each other for all parameters.

380 Model 6 presents a more challenging case where a coarse mesh does not fully resolve structures close to the surface, resulting in bias for the low-resolution models. The layer thicknesses of the coarse aliased model are explicitly multi-modal as the interfaces hover between depths, and the modes for several variables are significantly offset from their true values. These values become sharp modes in a broader parameter space, easily missed by any inappropriately scaled MCMC proposal. The fine aliased model reproduces the overall posterior shapes, with some distortions relative to the fine anti-aliased reference model. The coarse anti-aliased model produces a smooth posterior and recovers the correct overall correlations between parameters, but are somewhat offset from the true values.

385 The biases caused by aliasing seem to be strongest in angular variables such as fault directions. For Model 6, the best-fit dip-slip angles are offset from the true values by more than 10 degrees, corresponding to the extent through which the scan must sweep to obtain any difference in rock property for a few near-surface voxels in the coarse aliased model. Anti-aliasing the coarse-mesh model largely, but not completely, mitigates this bias.

4.2 MCMC sampling

390 Table 2 shows the average, best-case, and worst-case values of τ and \hat{R} . The posterior is challenging to navigate and all chains have autocorrelation times of order hundreds to thousands of samples, comparable to the performance of *Obsidian* for the Moomba models (Scalzo et al., 2019). Chains for anti-aliased models have shorter autocorrelation times by a factor of 2–10 than chains for aliased models, owing to the smoother, more Gaussian structure of the posterior. Most dramatically, the aliased chains still have large PSRF for several variables even after 10^5 samples. The anti-aliased chains consistently converge to the same distribution, with final PSRF values closer to 1 for all variables.

395 Models 5 and 6 produce challenging exceptions to these trends on which the anti-aliased chain also has trouble mixing fully. An example trace plot is shown in Fig. 7. One of the chains in the anti-aliased version has fallen into a different mode from the other three, characterized by shifts in the thicknesses and densities of stratigraphic layers as well as a shift in the dip angle of the fault. In contrast, none of the four traces in the aliased model seems to be sampling from the same distribution; each chain has learned a different proposal scaling, and chains occasionally jump between modes at adjacent parameter values characteristic of the ones seen in lower-dimensional projections. Each chain, considered on its own, would give a very different impression of the uncertainty in the inversion. This qualitative behavior is characteristic of sampling for the other models.

400 To fully characterize uncertainty in these poorly-constrained models, parallel tempering should be used to sample from all modes in proportion to their posterior probability. However, without anti-aliasing, the low-temperature chain in the tempering scheme will mix poorly and will rely more heavily on swap proposals to explore the posterior. The modes that do appear in anti-aliased posteriors are also more likely to represent distinct interpretations, rather than poorly-resolved strong correlations between variables.



Table 2. Performance metrics for each MCMC run, including mean (best-case, worst-case) integrated autocorrelation time τ in samples ($\times 10^3$), and mean (best-case, worst-case) potential scale reduction factor \hat{R} calculated across 4 runs starting from independent prior draws. Runs marked with 'AA' are anti-aliased, while those marked with only a single-digit model number have anti-aliasing turned off.

Model	$\tau_{\text{mean}} (\tau_{\text{min}}, \tau_{\text{max}})$	$\hat{R}_{\text{mean}} (\hat{R}_{\text{min}}, \hat{R}_{\text{max}})$
1	$0.6 (0.3, 0.8) \times 10^3$	1.34 (1.00, 3.53)
1-AA	$0.4 (0.2, 0.6) \times 10^3$	1.01 (1.00, 1.04)
2	$2.3 (0.6, 5.3) \times 10^3$	1.59 (1.00, 5.83)
2-AA	$0.2 (0.2, 0.3) \times 10^3$	1.00 (1.00, 1.01)
3	$2.5 (0.3, 6.2) \times 10^3$	2.65 (1.00, 15.41)
3-AA	$1.0 (0.2, 3.2) \times 10^3$	1.11 (1.00, 1.47)
4	$3.5 (0.6, 6.9) \times 10^3$	3.03 (1.01, 21.61)
4-AA	$0.6 (0.1, 1.9) \times 10^3$	1.01 (1.00, 1.06)
5	$3.5 (0.8, 6.6) \times 10^3$	214.88 (1.04, 2719.72)
5-AA	$0.3 (0.2, 0.3) \times 10^3$	19.91 (1.00, 120.74)
6	$4.1 (0.4, 6.6) \times 10^3$	6.92 (1.01, 26.97)
6-AA	$0.6 (0.2, 1.2) \times 10^3$	12.12 (1.00, 72.07)
7	$2.0 (0.6, 5.6) \times 10^3$	2.39 (1.00, 8.48)
7-AA	$0.3 (0.2, 0.4) \times 10^3$	1.00 (1.00, 1.01)
8	$2.5 (0.3, 6.7) \times 10^3$	5.52 (1.00, 36.93)
8-AA	$0.1 (0.1, 0.1) \times 10^3$	1.00 (1.00, 1.00)

5 Discussion

5.1 Anti-aliasing in curvilinear coordinates

The development of the anti-aliasing mappings in Sections 2.3 and 2.4 assumes that each voxel is a cube. However, mesh cells with varying aspect ratios are common in geophysical inversions, for example in dealing with varying sensitivity with depth or with boundary conditions where high resolution is less important. It may also be useful for some problems to work in curvilinear coordinates.

Although Blockworlds does not implement this functionality, a change of variable should be enough to extend the anti-aliasing prescription to more general coordinate systems. Let $\mathbf{J} = \partial \mathbf{r} / \partial \mathbf{u}$ be the Jacobian of the transformation from dimensionless coordinates \mathbf{u} on a unit voxel to physical coordinates \mathbf{r} , so that to first order $\mathbf{u} = \mathbf{J}^{-1}(\mathbf{r} - \mathbf{r}_0)$. Equation 6, defining the projection of a reference point in the voxel onto an interface running through the voxel, then becomes

$$u_{\perp} = \mathbf{u} \cdot \frac{1}{|\mathbf{J}\mathbf{n}|} \mathbf{J}\mathbf{n}, \quad (23)$$



and Equation 8, defining u_{\perp} in terms of scalar field values for an implicit model, becomes

$$u_{\perp} = (\Phi_0 - \Phi(\mathbf{r}_0)) / |\mathbf{J}\nabla\Phi|. \quad (24)$$

5.2 Anti-aliasing for other geophysical sensors

420 In this paper we have focused on gravity since it is among the most widely used geophysical sensors in an exploration context, and since its linear response makes an anti-aliasing treatment straightforward. All of our results should extend immediately to magnetic sensors given the strong mathematical equivalence, and may also be appropriate for linear forward models of other sensors such as thermal or electric conductivity, or for the slowness in travel-time tomography.

Useful schemes to anti-alias forward-modeled geology may exist for other sensors, as long as the sensor action at scales
425 beneath the mesh spacing can be usefully approximated by a computationally simple function of the rock properties and the interface geometry. Frequency-dependent sensors that probe a range of physical scales may require a frequency-dependent anti-aliasing function, and similarly for sensors with anisotropic interactions with interfaces. While there is nothing fundamental that prevents anti-aliasing for sensors that respond non-linearly to rock properties, mesh refinement may be more important in these cases to ensure the numerical accuracy, as well as the continuity, of the posterior. The framework in this paper treats only
430 quantities defined at mesh cell centers; while it may be possible to derive similar relations for quantities defined on mesh cell faces and edges, as in finite-volume treatments of electromagnetic sensors, this represents future work.

5.3 Limitations of anti-aliasing

Since the anti-aliasing approximation amounts to a strong prior on sub-mesh structure, understanding its limitations is critical in practical modeling. Common sense suggests that no length scale parameter in the inference problem should ever be less than
435 the voxel size h . However, as shown in Fig. 1, departures from the true posterior can occur even at higher resolution, depending upon the resolving power of the data and the structure of the likelihood.

The approximation giving rise to the first-order anti-aliasing scheme described in Eq. 6 and Eq. 8 will break down when higher-order terms in the expansion become significant. This will manifest in a departure of the position of the best-fit flat plane from the expected position, and in curvature of local interfaces. We can use the second-order term of the expansion to
440 estimate this, so that our particular scheme assumes

$$\max_{|\delta\mathbf{u}| < 1/2} \frac{\delta\mathbf{u}^T \mathbf{J}^T \mathbf{H}_{\Phi} \mathbf{J} \delta\mathbf{u}}{2|\mathbf{J}\nabla\Phi|} \ll 1, \quad (25)$$

where \mathbf{H}_{Φ} is the Hessian of Φ evaluated at the cell center \mathbf{r}_0 .

This leads us towards a criterion we can use to evaluate whether our model is under-resolved within a given mesh cell. Although we would expect the Hessian to be expensive to calculate at every mesh cell, it should be accessible to implicit
445 models based on co-kriging, and could be spot-checked only for cells lying near interfaces. Furthermore, within the context of MCMC, it might suffice to perform such a spot check about once per autocorrelation time, when the Markov chain migrates to a new, independent region of parameter space. Carrying out the maximization explicitly is also probably not needed; the



behavior will depend on the largest singular value of \mathbf{H}_Φ , but a simpler criterion to evaluate is

$$\delta v = \frac{\|\mathbf{J}^T \mathbf{H}_\Phi \mathbf{J}\|_{\max}}{|\mathbf{J} \nabla \Phi|} \ll 1. \quad (26)$$

450 An example for the sphere follows: setting $\Phi(\mathbf{r}) = |\mathbf{r}|^2$, we have $\nabla \Phi = 2\mathbf{r}$ and $\mathbf{H}_\Phi = 2\mathbf{I}$. Equation 26 then holds when $h/|\mathbf{r}| \gg 1$, and we expect the aliasing to hold near the true sphere radius R when $h \ll R$, as we might expect.

A more practical calibration can be given, again, by working out a limiting case. For a sphere with radius equal to the mesh cell size, Eq. 8 will overestimate its volume by a factor between 8% (if $\nabla \Phi$ points normal to a voxel face) to 20% (if it points along the voxel diagonal instead). The sphere is an extreme case in that it has positive curvature everywhere,
455 while a sinusoidally folded interface has no intrinsic curvature. Analytic calculations for other special cases give deviations proportional to δv , with estimated partial volume errors around $0.1\delta v$. For the sphere in Sect. 2.2, $R/h = 4.5$, so that $0.1\delta v = 2\%$. Since about half of the sphere's voxels lie at the interface, this is enough to explain the bias in inference shown in Fig. 1.

Another drawback in anti-aliasing is that it ignores the relative orientations of interfaces in voxels spanning multiple inter-
460 faces. If only one interface passes through a voxel, Eq. 26 should result in a faithful representation of its density. If, however, that voxel is later faulted or brought up against an unconformity, the true mean density in the voxel depends on which of the two formations is preferentially replaced by the new obtruding unit; Eq. 26 will only give some mean value roughly correct when the two interfaces intersect at right angles. Anti-aliasing will thus become increasingly inaccurate with each successive application to a voxel. This effect can cause biases even in geologies with exactly flat interfaces, but should remain small as long as the affected voxels are not too numerous or too close to the surface. Such regions of complexity can be identified by
465 flagging voxels within a projected distance $u_\perp < h$ of an interface, and the additional uncertainty in the lithology can be treated explicitly — for example, by assigning latent variables in the statistical model to account for it.

Once an inversion has been performed, the best practice to assess bias will still be to repeat it with a different cell size. However, using these metrics to assess an ensemble of models drawn from the prior, or tracking affected voxels while a trial inversion is running, could still be useful in determining a suitable cell size for the final model.

470 6 Conclusions

Our experiments show the potential pitfalls of using oversimplified projections of 3-D structural geology onto a volumetric basis for calculation of synthetic geophysics, and demonstrate an intuitive, efficient solution. Anti-aliasing reproduces the smooth behavior of the underlying posterior with respect to the geological parameters, to enhance convergence of optimization or mixing of sampling methods and to enable the use of derivative information in these methods. Anti-aliasing enables a fixed
475 volumetric basis to more faithfully represent sensor response to latent discrete geology with interfaces that are flat on the scale of a mesh cell, reducing the computational burden of forward geophysics in an inversion loop. Our algorithm is also coordinate-invariant and can be combined with curvilinear meshes, and with mesh refinement techniques such as octrees (e.g. Haber and Heldmann, 2007), for even stronger results.



We have focused on calculations over 3-D volumetric meshes because implicit and kinematic models are naturally volumetric, and because existing geological modeling codes already export rock properties onto meshes as block models. Introducing anti-aliasing into these geological models is a minimally invasive modification to enhance their use in MCMC-based Bayesian inversions. Other methods, such as the Cauchy surface method for calculating gravity and magnetics (Zhdanov and Liu, 2013; Cai and Zhdanov, 2015), require many fewer discrete elements than a 3-D volumetric mesh to achieve a given accuracy, since they follow an explicit lower-dimensional discretization of each interface. These methods are most suitable for geological models that already parametrize unit interfaces explicitly. Using them with implicit models, or with kinematic models that work by deforming and transporting volumetric elements, would introduce additional computational overhead in tracking the positions of interfaces and updating an associated lower-dimensional mesh. The sensor response coefficients would then also have to be updated each time the model parameters are updated. There may still be geological use cases where tracking interfaces during the sampling process is more efficient than a finite volume solution involving every voxel, but we defer examination of such cases to future work.

The accurate and efficient projection of these simple geological models onto meshes for geophysics calculations is a prerequisite to inversion for structural and kinematic parameters in more realistic situations. We can now take clear next steps towards MCMC sampling of kinematic histories for richer, higher-dimensional models. In addition, we can now move towards sampling of hierarchical geophysical inversions that use a parametric structural model as a mean function. This will enable voxel-space inversions for which the geological prior is expressed in terms of uncertain interpretable parameters, or inversions for geology that include uncertainty due to residual rock property variations constrained by geophysics. These would represent more complete probabilistic treatments of geological uncertainty in light of available constraints from geophysics.

Code and data availability. The version of the Blockworlds model code used in this paper is available at <https://github.com/rscalzo/blockworlds/releases/tag/v0.1.0-beta.2>. The datasets used in the inversions are synthetic and can be reproduced exactly by running a set of scripts from the command line that

fix the random number seed used to generate those datasets, as described in the package manual.

Appendix A: MCMC performance metrics

In the definitions to follow, we follow notation introduced in Scalzo et al. (2019). Consider M Markov chains, each of length N and with d -dimensional parameter vectors, are run independently or in parallel. Let $\theta_{ki}^{[j]}$ denote parameter k of d , drawn at iteration $[j]$ of N in chain i of M . Let

$$\hat{\theta}_{ki} = \frac{1}{N} \sum_{j=1}^N \theta_{ki}^{[j]} \quad (\text{A1})$$

denote the sample mean of parameter k over the N iterates of chain i , and

$$s_{ki} = \frac{1}{M-1} \sum_{j=1}^N (\theta_{ki}^{[j]} - \hat{\theta}_{ki}) \quad (\text{A2})$$



denote the sample variance of parameter k over chain i . Let $\tilde{\theta}_k = \sum 1/M \sum_{i=1}^M \hat{\theta}_{ki}$ denote the sample mean of $\hat{\theta}_{ki}$ across all chains. To summarize variation in parameter estimates across chains, let

$$510 \quad B_k = \frac{1}{M-1} \sum_{i=1}^M (\theta_{ki}^{[j]} - \hat{\theta}_{ki}) \quad (A3)$$

denote the sample variance in $\hat{\theta}_{ki}$ across chains, and

$$W_k = \frac{1}{M} \sum_{i=1}^M s_{ki}^2 \quad (A4)$$

515 denote the sample mean of s_{ki}^2 across chains. Thus B_k is a measure of the variance in parameter estimates made from the history of samples in any single chain on its own, and W_k is a measure of the overall variance of a parameter within any single chain.

A1 Integrated autocorrelation time

The autocorrelation function, measuring the correlation between parameter draws separated by a lag l when treating each chain as a time series, is

$$\rho_{lki} = \frac{1}{(N-l)W_k} (\theta_{ki}^{[j]} - \hat{\theta}_{ki})(\theta_{ki}^{[j-l]} - \hat{\theta}_{ki}). \quad (A5)$$

520 The integrated autocorrelation time (IACT), the sum of the autocorrelation function over lags l gives an estimate of the number of chain samples required to obtain an independent draw from the stationary distribution:

$$\tau_{ki} = 1 + 2 \sum_{l=1}^N \left(1 - \frac{l}{N}\right) \rho_{lki}. \quad (A6)$$

A2 Potential scale reduction factor

525 The potential scale reduction factor (PSRF) was introduced as a metric for chain convergence by Gelman and Rubin (1992). It measures the extent to which uncertainty in a parameter could be reduced by continuing to sample beyond the nominal chain length N . A simple version often implemented is

$$\frac{\hat{V}_k}{W_k} = \frac{N-1}{N} + \frac{M+1}{MN} \frac{B_k}{W_k} \quad (A7)$$

530 The behavior of the metric is driven by the ratio B_k/W_k , or the variance in parameter means from different chains as a fraction of the overall marginal variance in that parameter. This can be large because of small number statistics in a unimodal problem, but it can also be large because of poor mixing between modes in a multimodal problem, which increases the autocorrelation between samples. Gelman and Rubin (1992) advocate searching for multiple modes in advance and drawing starting points for MCMC from an overdispersed approximation to the posterior. They also account for sampling variability, including correlation between samples, via the modified metric

$$\hat{R} = \frac{\hat{V}}{W} \frac{\nu}{\nu-2}, \quad (A8)$$



535 where $\nu = 2\hat{V}^2/\hat{\text{var}}(\hat{V})$ represents the number of degrees of freedom in a t -distribution for $\tilde{\theta}_k$, with

$$\begin{aligned}\hat{\text{var}}(\hat{V}_k) &= \left(\frac{N-1}{N}\right)^2 \frac{1}{M} \hat{\text{var}}(s_{ki}^2) \\ &+ \left(\frac{M+1}{MN}\right)^2 \frac{2}{M-1} B^2 \\ &+ 2 \frac{(M+1)(N-1)}{MN^2} \frac{N}{M} \\ &\times \left[\hat{\text{cov}}(s_{ki}^2, \hat{\theta}_{ki}^2) - 2\tilde{\theta}_k \hat{\text{cov}}(s_{ki}^2, \hat{\theta}_{ki}) \right].\end{aligned}\tag{A9}$$

540 We use the metric \hat{R} for our experiments, given that autocorrelation times can be quite long for our chains.

Author contributions. The study was conceptualized by SC, who with MJ provided funding and resources. RS was responsible for project administration, designed the methodology, wrote the Blockworlds code resulting in v0.1, and carried out the main investigation and formal analysis under supervision from SC, EC, and MJ. RS and ML validated and visualized the results. RS wrote the original draft text with contributions from ML, for which all coauthors provided a review and critical evaluation.

545 *Competing interests.* The authors declare that they have no conflict of interest.

Acknowledgements. Parts of this research were conducted by the Australian Research Council Industrial Transformation Training Centre in Data Analytics for Resources and the Environment (DARE), through project number IC190100031. The work has been supported by the Mineral Exploration Cooperative Research Centre whose activities are funded by the Australian Government's Cooperative Research Centre Programme. This is MinEx CRC Document 2021/34. MDL acknowledges funding from ARC DECRA DE190100431 and MinEx CRC. We

550 thank Alistair Reid and Malcolm Sambridge for useful discussions.



References

- Andrieu, C. and Thoms, J.: A tutorial on adaptive MCMC, *Statistics and Computing*, 18, 343–373, <https://doi.org/10.1007/s11222-008-9110-y>, <http://link.springer.com/10.1007/s11222-008-9110-y>, 2008.
- Backus, G. and Gilbert, F.: The Resolving Power of Gross Earth Data, *Geophysical Journal of the Royal Astronomical Society*, 16, 169–205, <https://doi.org/10.1111/j.1365-246X.1968.tb00216.x>, 1968.
- Backus, G. and Gilbert, F.: Uniqueness in the inversion of inaccurate gross Earth data, *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 266, 123–192, <https://doi.org/10.1111/j.1365-246X.1968.tb00216.x>, 1970.
- Backus, G. E.: Long-wave elastic anisotropy produced by horizontal layering, *Journal of Geophysical Research*, 67, 4427–4440, 1962.
- Backus, G. E. and Gilbert, J. F.: Numerical Applications of a Formalism for Geophysical Inverse Problems, *Geophysical Journal of the Royal Astronomical Society*, 13, 247–276, <https://doi.org/10.1111/j.1365-246X.1967.tb02159.x>, 1967.
- Beardmore, G., Durrant-Whyte, H., and Callaghan, S. O.: A Bayesian inference tool for geophysical joint inversions, *ASEG Extended Abstracts 2016.1* (2016): 1–10., pp. 1–10, 2016.
- Bond, C. E.: Uncertainty in structural interpretation: Lessons to be learnt, *Journal of Structural Geology*, 74, 185–200, <https://doi.org/10.1016/j.jsg.2015.03.003>, 2015.
- Bosch, M.: Lithologic tomography: From plural geophysical data to lithology estimation, *Journal of Geophysical Research: Solid Earth*, 104, 749–766, <https://doi.org/10.1029/1998JB900014>, 1999.
- Bosch, M.: Inference Networks in Earth Models with Multiple Components and Data, in: *Geophysical Monograph Series*, edited by Moorkamp, M., Lelièvre, P. G., Linde, N., and Khan, A., pp. 29–47, John Wiley & Sons, Inc, Hoboken, NJ, <https://doi.org/10.1002/9781118929063.ch3>, 2016.
- Bosch, M., Guillen, A., and Ledru, P.: Lithologic tomography: an application to geophysical data from the Cadomian belt of northern Brittany, France, *Tectonophysics*, 331, 197–227, [https://doi.org/10.1016/S0040-1951\(00\)00243-2](https://doi.org/10.1016/S0040-1951(00)00243-2), 2001.
- Brunetti, C., Bianchi, M., Pirot, G., and Linde, N.: Hydrogeological Model Selection Among Complex Spatial Priors, *Water Resources Research*, 55, 6729–6753, <https://doi.org/10.1029/2019WR024840>, 2019.
- Cai, H. and Zhdanov, M.: Application of Cauchy-type integrals in developing effective methods for depth-to-basement inversion of gravity and gravity gradiometry data, *GEOPHYSICS*, 80, G81–G94, <https://doi.org/10.1190/geo2014-0332.1>, <http://library.seg.org/doi/10.1190/geo2014-0332.1>, 2015.
- Calcagno, P., Chilès, J., Courrioux, G., and Guillen, A.: Geological modelling from field data and geological knowledge, *Physics of the Earth and Planetary Interiors*, 171, 147–157, <https://doi.org/10.1016/j.pepi.2008.06.013>, 2008.
- Capdeville, Y., Guillot, L., and Marigo, J.-J.: 1-D non-periodic homogenization for the seismic wave equation. *Geophysical Journal International*, *Geophysical Journal International*, 181, 897–910, 2010.
- Catmull, E.: A hidden-surface algorithm with anti-aliasing, in: *Proceedings of the 5th Annual Conference on Computer Graphics and Interactive Techniques*, pp. 6–11, 1978.
- Cockett, R., Kang, S., Heagy, L. J., Pidlisecky, A., and Oldenburg, D. W.: SimPEG: An open source framework for simulation and gradient based parameter estimation in geophysical applications, *Computers & Geosciences*, 85, 142–154, <https://doi.org/10.1016/j.cageo.2015.09.015>, 2015.
- Cook, R. L.: Stochastic sampling in computer graphics, *ACM Transactions on Graphics*, 5, 51–72, 1986.



- Cordua, K. S., Hansen, T. M., and Mosegaard, K.: Monte Carlo full-waveform inversion of crosshole GPR data using multiple-point geostatistical a priori information, *GEOPHYSICS*, 77, H19–H31, <https://doi.org/10.1190/geo2011-0170.1>, 2012.
- Crow, F. C.: The aliasing problem in computer-generated shaded images, *Communications of the ACM*, 20, 799–805, 1977.
- 590 de la Varga, M., Schaaf, A., and Wellmann, F.: GemPy 1.0: open-source stochastic geological modeling and inversion, *Geoscience Model Development Discussions*, submitted, 50, 2018.
- de Pasquale, G., Linde, N., Doetsch, J., and Holbrook, W. S.: Probabilistic inference of subsurface heterogeneity and interface geometry using geophysical data, *Geophysical Journal International*, 217, 816–831, <https://doi.org/10.1093/gji/ggz055>, 2019.
- Eckhardt, E. A.: Partnership between geology and geophysics in prospecting for oil, *GEOPHYSICS*, 5, 209–214,
595 <https://doi.org/10.1190/1.1441804>, 1940.
- Frodeman, R.: Geological reasoning: Geology as an interpretive and historical science, *Geological Society of America Bulletin*, 107, 960–9968, [https://doi.org/10.1130/0016-7606\(1995\)107<0960:GRGAAI>2.3.CO;2](https://doi.org/10.1130/0016-7606(1995)107<0960:GRGAAI>2.3.CO;2), 1995.
- Gallardo, L. A. and Meju, M. A.: Structure-coupled multiphysics imaging in geophysical sciences, *Reviews of Geophysics*, 49, 2011.
- Gelman, A. and Rubin, D.: Inference from Iterative Simulation Using Multiple Sequences, *Statistical Science*, 7, 457–472, 1992.
- 600 Giraud, J., Pakyuz-Charrier, E., Jessell, M., Lindsay, M., Martin, R., and Ogarko, V.: Uncertainty reduction through geologically conditioned petrophysical constraints in joint inversion, *GEOPHYSICS*, 82, ID19–ID34, <https://doi.org/10.1190/geo2016-0615.1>, 2017.
- Giraud, J., Pakyuz-Charrier, E., Ogarko, V., Jessell, M., Lindsay, M., and Martin, R.: Impact of uncertain geology in constrained geophysical inversion, *ASEG Extended Abstracts*, 2018, 1, https://doi.org/10.1071/ASEG2018abM1_2F, 2018.
- Giraud, J., Lindsay, M., Ogarko, V., Jessell, M., Martin, R., and Pakyuz-Charrier, E.: Integration of geoscientific uncertainty into geophysical
605 inversion by means of local gradient regularization, *Solid Earth*, 10, 193–210, <https://doi.org/10.5194/se-10-193-2019>, 2019a.
- Giraud, J., Ogarko, V., Lindsay, M., Pakyuz-Charrier, E., Jessell, M., and Martin, R.: Sensitivity of constrained joint inversions to geological and petrophysical input data uncertainties with posterior geological analysis, *Geophysical Journal International*, 218, 666–688, <https://doi.org/10.1093/gji/ggz152>, 2019b.
- Giraud, J., Lindsay, M., Jessell, M., and Ogarko, V.: Towards plausible lithological classification from geophysical inversion: honouring
610 geological principles in subsurface imaging, *Solid Earth*, 11, 419–436, <https://doi.org/10.5194/se-11-419-2020>, 2020.
- Giraud, J., Lindsay, M., and Jessell, M.: Generalization of level-set inversion to an arbitrary number of geologic units in a regularized least-squares framework, *GEOPHYSICS*, 86, R623–R637, <https://doi.org/10.1190/geo2020-0263.1>, <https://library.seg.org/doi/10.1190/geo2020-0263.1>, 2021a.
- Giraud, J., Ogarko, V., Martin, R., Jessell, M., and Lindsay, M.: Structural, petrophysical and geological constraints in potential field inversion
615 using the Tomofast-x v1.0 open-source code, *Geoscientific Model Development Discussions*, 2021, 1–51, <https://doi.org/10.5194/gmd-2021-14>, <https://gmd.copernicus.org/preprints/gmd-2021-14/>, 2021b.
- Grose, L., Ailleres, L., Laurent, G., and Jessell, M.: LoopStructural 1.0: Time aware geological modelling, *Geoscientific Model Development Discussions*, 2020, 1–31, <https://doi.org/10.5194/gmd-2020-336>, <https://gmd.copernicus.org/preprints/gmd-2020-336/>, 2020.
- Guillen, A., Calcagno, P., Courrioux, G., Joly, A., and Ledru, P.: Geological modelling from field data and geological knowledge Part
620 II. Modelling validation using gravity and magnetic data inversion, *Physics of the Earth and Planetary Interiors*, 171, 158–169, <https://doi.org/10.1016/j.pepi.2008.06.014>, 2008.
- Haario, H., Saksman, E., and Tamminen, J.: An Adaptive Metropolis Algorithm, *Bernoulli*, 7, 223, <https://doi.org/10.2307/3318737>, ISBN: 1350-7265, 2001.



- Haber, E. and Heldmann, S.: An octree multigrid method for quasi-static Maxwell's equations with highly discontinuous coefficients, *Journal of Computational Physics*, 223, 783–796, <https://doi.org/https://doi.org/10.1016/j.jcp.2006.10.012>, <https://www.sciencedirect.com/science/article/pii/S0021999106004748>, 2007.
- Hastings, W. K.: Monte Carlo sampling methods using Markov chains and their applications, *Biometrika*, 57, 97–109, <https://doi.org/10.1093/biomet/57.1.97>, 1970.
- Houlding, S. W.: 3D Geoscience modeling: computer techniques for geological characterization, Springer Verlag, 46, 85–90, 1994.
- 630 Jessell, M., Pakyuz-Charrier, E., Lindsay, M., Giraud, J., and de Kemp, E.: Assessing and Mitigating Uncertainty in Three-Dimensional Geologic Models in Contrasting Geologic Scenarios, in: *Metals, Minerals, and Society*, Society of Economic Geologists (SEG), <https://doi.org/10.5382/SP.21.04>, 2018.
- Jessell, M. W.: “Noddy” – An interactive Map creation Package, Master's thesis, Imperial College of Science and Technology, London, UK, 1981.
- 635 Jessell, M. W. and Valenta, R. K.: Structural geophysics: Integrated structural and geophysical modelling, in: *Structural Geology and Personal Computers*, edited by De Paor, D. G., pp. 303–324, Elsevier, Oxford, UK, 1993.
- Jessell, M. W., Ailleres, L., and de Kemp, E. A.: Towards an integrated inversion of geoscientific data: What price of geology?, *Tectonophysics*, 490, 294–306, <https://doi.org/10.1016/j.tecto.2010.05.020>, 2010.
- Kennedy, M. C. and O'Hagan, A.: Bayesian calibration of computer models, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63, 425–464, <https://doi.org/10.1111/1467-9868.00294>, 2001.
- 640 Koene, E. F. M., Wittsten, J., and Robertsson, J. O. A.: Finite-difference modeling of 2-D wave propagation in the vicinity of dipping interfaces: a comparison of anti-aliasing and equivalent medium approaches, 2021.
- Lajaunie, C., Courrioux, G., and Manuel, L.: Foliation fields and 3D cartography in geology: Principles of a method based on potential interpolation, *Mathematical Geology*, 29, 571–584, <https://doi.org/10.1007/BF02775087>, 1997.
- 645 Li, W., Lu, W., Qian, J., and Li, Y.: A multiple level-set method for 3D inversion of magnetic data, *Geophysics*, 82, J61–J81, <https://doi.org/10.1190/GEO2016-0530.1>, 2017.
- Li, Y. and Oldenburg, D. W.: 3-D inversion of gravity data, *Geophysics*, 63, 109–119, <https://doi.org/10.1190/1.1444302>, 1998.
- Li, Y. and Oldenburg, D. W.: Incorporating geological dip information into geophysical inversions, *GEOPHYSICS*, 65, 148–157, <https://doi.org/10.1190/1.1444705>, 2000.
- 650 Linde, N., Ginsbourger, D., Irving, J., Nobile, F., and Doucet, A.: On uncertainty quantification in hydrogeology and hydrogeophysics, *Advances in Water Resources*, 110, 166 – 181, <https://doi.org/10.1016/j.advwatres.2017.10.014>, 2017.
- Lindsay, M., Jessell, M., Ailleres, L., Perrouy, S., de Kemp, E., and Betts, P.: Geodiversity: Exploration of 3D geological model space, *Tectonophysics*, 594, 27–37, <https://doi.org/https://doi.org/10.1016/j.tecto.2013.03.013>, <https://www.sciencedirect.com/science/article/pii/S004019511300187X>, 2013.
- 655 Lindsay, M. D., Aillères, L., Jessell, M. W., de Kemp, E. A., and Betts, P. G.: Locating and quantifying geological uncertainty in three-dimensional models: Analysis of the Gippsland Basin, southeastern Australia, *Tectonophysics*, 546–547, 10–27, <https://doi.org/https://doi.org/10.1016/j.tecto.2012.04.007>, <https://www.sciencedirect.com/science/article/pii/S0040195112002004>, 2012.
- Lindsay, M. D., Occhipinti, S., Laflamme, C., Aitken, A., and Ramos, L.: Mapping undercover: integrated geoscientific interpretation and 3D modelling of a Proterozoic basin, *Solid Earth*, 11, 1053–1077, <https://doi.org/10.5194/se-11-1053-2020>, <https://se.copernicus.org/articles/11/1053/2020/>, 2020.
- 660



- McCalman, L., O'Callaghan, S. T., Reid, A., Shen, D., Carter, S., Krieger, L., Beardsmore, G. R., Bonilla, E. V., and Ramos, F. T.: Distributed Bayesian geophysical inversions, Proceedings of the Thirty-Ninth Workshop on Geothermal Reservoir Engineering, Stanford University, pp. 1–11, 2014.
- Metropolis, N. and Ulam, S.: The Monte Carlo Method, *Journal of the American Statistical Association*, 44, 335–341, 1949.
- 665 Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E.: Equation of State Calculations by Fast Computing Machines, *The Journal of Chemical Physics*, 21, 1087–1092, <https://doi.org/10.1063/1.1699114>, 1953.
- Moczó, P., Kristek, J., Vavryčuk, V., Archuleta, R. J., and Halada, L.: 3d heterogeneous staggered-grid finite-difference modeling of seismic motion with volume harmonic and arithmetic averaging of elastic moduli and densities., *Bulletin of the Seismological Society of America*, 92, 3042–3066, 2002.
- 670 Mosegaard, K. and Sambridge, M.: Monte Carlo analysis of inverse problems, *Inverse Problems*, 18, R29–R54, <https://doi.org/10.1088/0266-5611/18/3/201>, 2002.
- Mosegaard, K. and Tarantola, A.: Monte Carlo sampling of solutions to inverse problems, *Journal of Geophysical Research: Solid Earth*, 100, 12 431–12 447, <https://doi.org/10.1029/94JB03097>, 1995.
- Muir, F., Dellinger, J., Etgen, J., and Nichols, D.: Modeling elastic fields across irregular boundaries, *Geophysics*, 57, 1189–1193, 1992.
- 675 Neal, R. M.: MCMC using Hamiltonian dynamics, *Handbook of Markov Chain Monte Carlo*, pp. 113–162, <https://doi.org/doi:10.1201/b10905-6>, arXiv: 1206.1901 ISBN: 9781420079418, 2011.
- Nishimura, A., Dunson, D., and Lu, J.: Discontinuous Hamiltonian Monte Carlo for discrete parameters and discontinuous likelihoods, arXiv:1705.08510 [stat], <http://arxiv.org/abs/1705.08510>, arXiv: 1705.08510, 2019.
- Ogarko, V., Giraud, J., Martin, R., and Jessell, M.: Disjoint interval bound constraints using the alternating direction method of multipliers
680 for geologically constrained inversion: Application to gravity data, *GEOPHYSICS*, 86, G1–G11, <https://doi.org/10.1190/geo2019-0633.1>, 2021.
- Okabe, M.: Analytical expressions for gravity anomalies due to homogeneous polyhedral bodies and translations into magnetic anomalies, *Geophysics*, 44, 730, <https://doi.org/10.1190/1.1440973>, 1979.
- Olierook, H. K., Scalzo, R., Kohn, D., Chandra, R., Farahbakhsh, E., Clark, C., Reddy, S. M., and Müller, R. D.: Bayesian geological and geophysical data fusion for the construction and uncertainty quantification of 3D geological models, *Geoscience Frontiers*,
685 <https://doi.org/10.1016/j.gsf.2020.04.015>, 2020.
- Osher, S. and Sethian, J. A.: Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations, *Journal of Computational Physics*, 79, 12–49, [https://doi.org/10.1016/0021-9991\(88\)90002-2](https://doi.org/10.1016/0021-9991(88)90002-2), 1988.
- Öztireli, A. C.: A Comprehensive Theory and Variational Framework for Anti-aliasing Sampling Patterns, *Computer Graphics Forum*, 39,
690 133–148, 2020.
- Pakyuz-Charrier, E., Giraud, J., Ogarko, V., Lindsay, M., and Jessell, M.: Drillhole uncertainty propagation for three-dimensional geological modeling using Monte Carlo, *Tectonophysics*, 747–748, 16–39, <https://doi.org/10.1016/j.tecto.2018.09.005>, 2018a.
- Pakyuz-Charrier, E., Lindsay, M., Ogarko, V., Giraud, J., and Jessell, M.: Monte Carlo simulation for uncertainty estimation on structural data in implicit 3-D geological modeling, a guide for disturbance distribution selection and parameterization, *Solid Earth*, 9, 385–402,
695 <https://doi.org/10.5194/se-9-385-2018>, 2018b.
- Pakyuz-Charrier, E., Jessell, M., Giraud, J., Lindsay, M., and Ogarko, V.: Topological analysis in Monte Carlo simulation for uncertainty propagation, *Solid Earth*, 10, 1663–1684, <https://doi.org/10.5194/se-10-1663-2019>, 2019.
- Patil, A., Huard, D., and Fonnesbeck, C. J.: PyMC: Bayesian stochastic modelling in Python, *J. Stat. Softw.*, pp. 1–81, 2010.



- 700 Perrouty, S., Lindsay, M., Jessell, M., Aillères, L., Martin, R., and Bourassa, Y.: 3D modeling of the Ashanti Belt, southwest Ghana: Evidence for a litho-stratigraphic control on gold occurrences within the Birimian Sefwi Group, *Ore Geology Reviews*, 63, 252 – 264, <https://doi.org/10.1016/j.oregeorev.2014.05.011>, 2014.
- Pirot, G., Renard, P., Huber, E., Straubhaar, J., and Huggenberger, P.: Influence of conceptual model uncertainty on contaminant transport forecasting in braided river aquifers, *Journal of Hydrology*, 531, 124 – 141, <https://doi.org/10.1016/j.jhydrol.2015.07.036>, groundwater flow and transport in aquifers: Insights from modeling and characterization at the field scale, 2015.
- 705 Pirot, G., Linde, N., Mariethoz, G., and Bradford, J. H.: Probabilistic inversion with graph cuts: Application to the Boise Hydrogeophysical Research Site, *Water Resources Research*, 53, 1231–1250, <https://doi.org/10.1002/2016WR019347>, 2017.
- Pirot, G., Huber, E., Irving, J., and Linde, N.: Reduction of conceptual model uncertainty using ground-penetrating radar profiles: Field-demonstration for a braided-river aquifer, *Journal of Hydrology*, 571, 254 – 264, <https://doi.org/10.1016/j.jhydrol.2019.01.047>, 2019a.
- Pirot, G., Krityakierne, T., Ginsbourger, D., and Renard, P.: Contaminant source localization via Bayesian global optimization, *Hydrology and Earth System Sciences*, 23, 351–369, <https://doi.org/10.5194/hess-23-351-2019>, 2019b.
- 710 Quigley, M. C., Bennetts, L. G., Durance, P., Kuhnert, P. M., Lindsay, M. D., Pembleton, K. G., Roberts, M. E., and White, C. J.: The provision and utility of science and uncertainty to decision-makers: earth science case studies, *Environment Systems and Decisions*, 39, 307–348, <https://doi.org/10.1007/s10669-019-09728-0>, 2019.
- Rawlinson, N., Fichtner, A., Sambridge, M., and Young, M. K.: Seismic Tomography and the Assessment of Uncertainty, in: *Advances in Geophysics*, vol. 55, pp. 1–76, Elsevier, <https://doi.org/10.1016/bs.agph.2014.08.001>, 2014.
- 715 Sambridge, M., Bodin, T., Gallagher, K., and Tkalcic, H.: Transdimensional inference in the geosciences, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 371, 20110 547–20110 547, <https://doi.org/10.1098/rsta.2011.0547>, ISBN: 1364503X (ISSN), 2012.
- Santosa, F.: A level-set approach for inverse problems involving obstacles, *ESAIM: Control, Optimisation and Calculus of Variations*, 1, 17–33, <https://doi.org/10.1051/cocv:1996101>, 1996.
- 720 Scalzo, R., Kohn, D., Olierook, H., Houseman, G., Chandra, R., Girolami, M., and Cripps, S.: Efficiency and robustness in Monte Carlo sampling for 3-D geophysical inversions with Obsidian v0.1.2: setting up for success, *Geoscientific Model Development*, 12, 2941–2960, <https://doi.org/10.5194/gmd-12-2941-2019>, 2019.
- Tarantola, A. and Valette, B.: Generalized nonlinear inverse problems solved using the least squares criterion, *Reviews of Geophysics*, 20, 219–232, <https://doi.org/10.1029/RG020i002p00219>, 1982.
- 725 Varouchakis, E. A., Yetilmezsoy, K., and Karatzas, G. P.: A decision-making framework for sustainable management of groundwater resources under uncertainty: combination of Bayesian risk approach and statistical tools, *Water Policy*, 21, 602–622, <https://doi.org/10.2166/wp.2019.128>, 2019.
- Wang, Z., Yin, Z., Caers, J., and Zuo, R.: A Monte Carlo-based framework for risk-return analysis in mineral prospectivity mapping, *Geoscience Frontiers*, 11, 2297 – 2308, <https://doi.org/10.1016/j.gsf.2020.02.010>, 2020.
- 730 Wegener, A.: Die Entstehung der Kontinente und Ozeane, *The Geographical Journal*, 61, 304, <https://doi.org/10.2307/1781265>, 1923.
- Wellmann, F. and Caumon, G.: 3-D Structural geological models: Concepts, methods, and uncertainties, in: *Advances in Geophysics*, vol. 59, pp. 1–121, Elsevier, <https://doi.org/10.1016/bs.agph.2018.09.001>, 2018.
- Wellmann, J. F. and Regenauer-Lieb, K.: Uncertainties have a meaning: Information entropy as a quality measure for 3-D geological models, *Tectonophysics*, 526-529, 207–216, <https://doi.org/10.1016/j.tecto.2011.05.001>, 2012.
- 735



- Wellmann, J. F., de la Varga, M., Murdie, R. E., Gessner, K., and Jessell, M.: Uncertainty estimation for a geological model of the Sandstone greenstone belt, Western Australia – insights from integrated geological and geophysical inversion in a Bayesian inference framework, Geological Society, London, Special Publications, 453, 41–56, <https://doi.org/10.1144/SP453.12>, 2017.
- 740 Witter, J. B., Trainor-Guitton, W. J., and Siler, D. L.: Uncertainty and risk evaluation during the exploration stage of geothermal development: A review, *Geothermics*, 78, 233 – 242, <https://doi.org/10.1016/j.geothermics.2018.12.011>, 2019.
- Zhdanov, M. S. and Liu, X.: 3-D Cauchy-type integrals for terrain correction of gravity and gravity gradiometry data, *Geophysical Journal International*, 194, 249–268, <https://doi.org/10.1093/gji/ggt120>, <http://academic.oup.com/gji/article/194/1/249/650016/3D-Cauchytype-integrals-for-terrain-correction-of>, 2013.
- 745 Zheglava, P., Lelievre, P. G., and Farquharson, C. G.: Multiple level-set joint inversion of traveltimes and gravity data with application to ore delineation: A synthetic study, *Geophysics*, 83, R13–R30, <https://doi.org/10.1190/GEO2016-0675.1>, 2018.

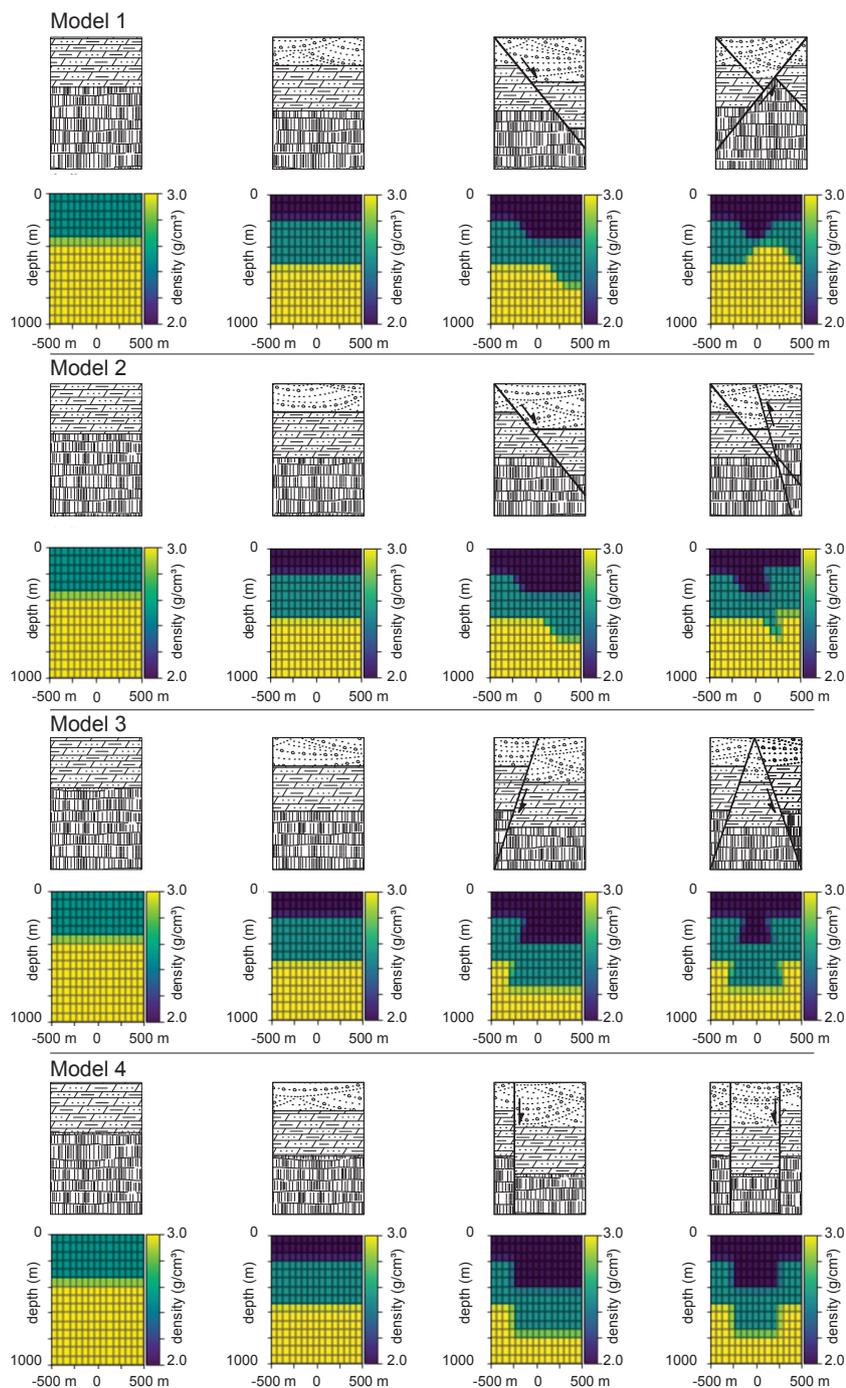


Figure 3. Calculation of forward geology for Models 1–4 through a sequence of tectonic events: the addition of two stratigraphic layers and two faults. Top: geology; bottom: anti-aliased, voxelised rendering.

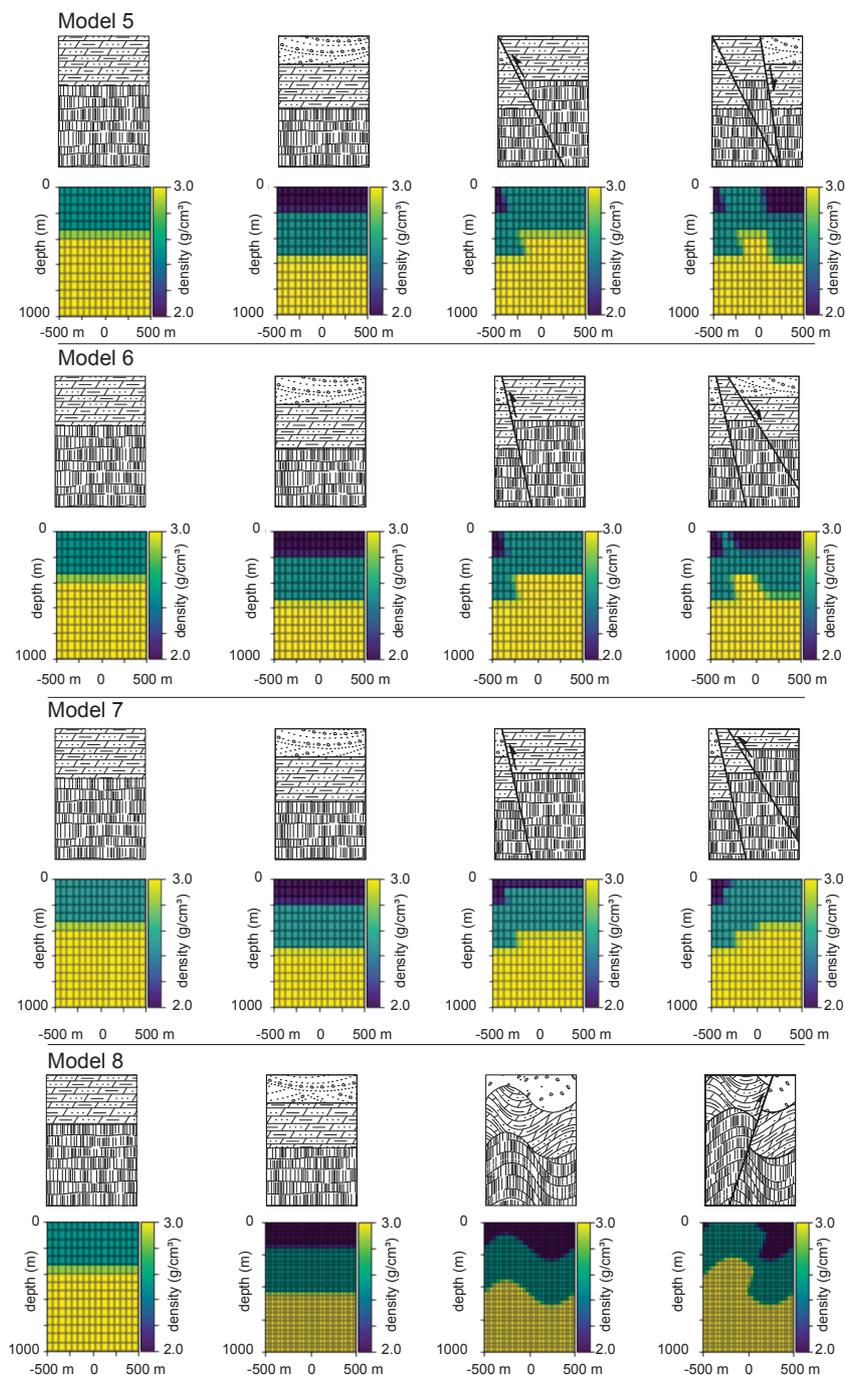


Figure 4. Calculation of forward geology for Models 5–8, each with basement and two additional stratigraphic layers. Models 5, 6, and 7 each involve two faults as for Models 1–4 in Fig. 3, while Model 8 involves a fold as well as a single fault. Top: geology; bottom: anti-aliased, voxelised rendering.

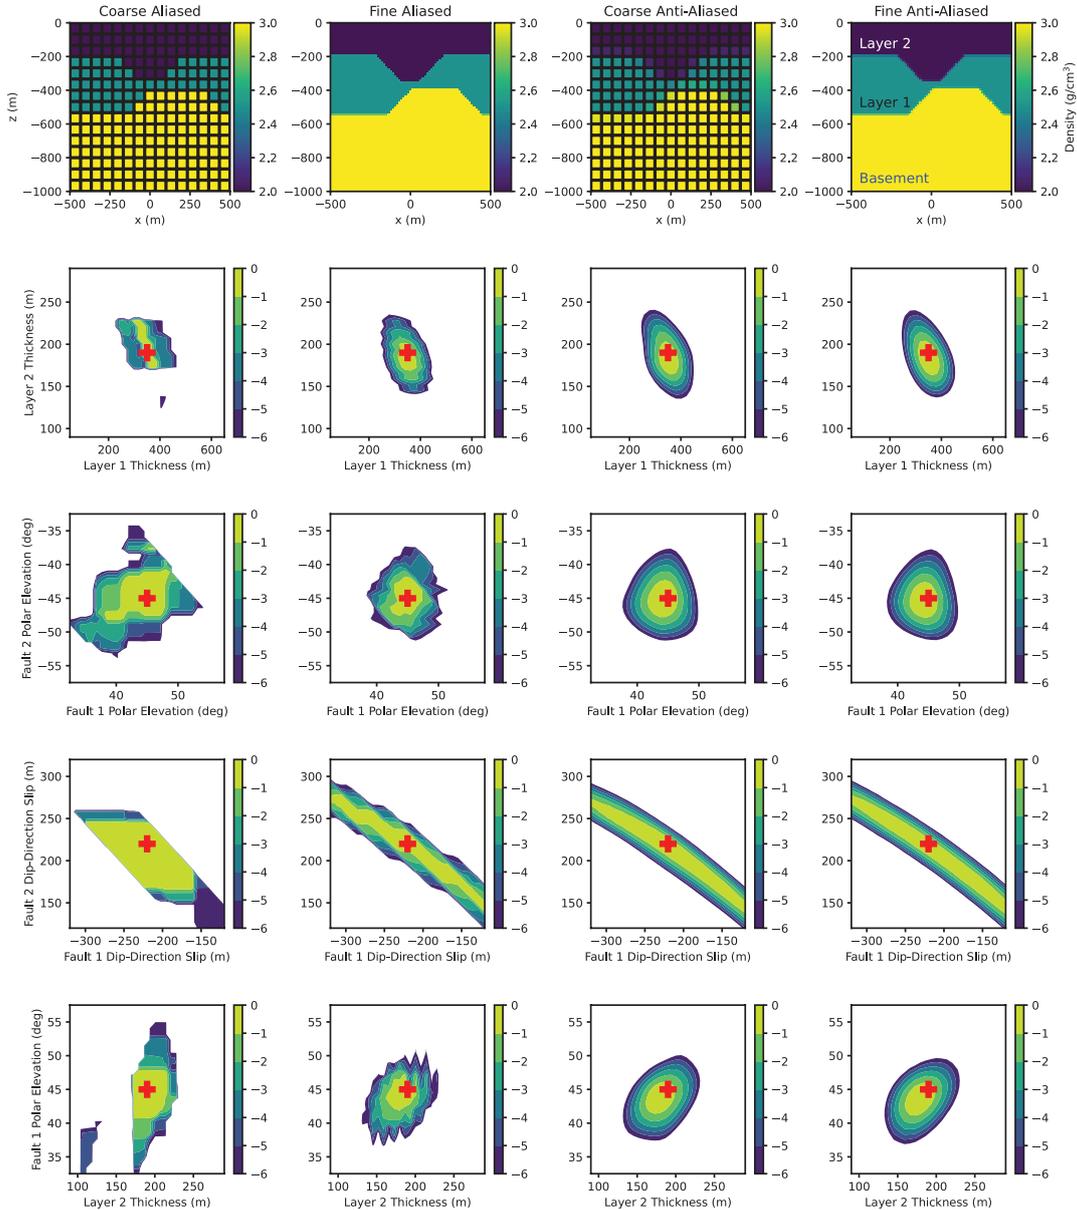


Figure 5. The log posterior distribution of pairs of parameters, denoted by θ , in kinematic model 1, relative to the maximal value $\log_{10}[p(\theta|\mathbf{d})/p(\hat{\theta}_{\text{MAP}}|\mathbf{d})]$, where $\hat{\theta}_{\text{MAP}} = \text{argmax}_{\theta} p(\theta|\mathbf{d})$. In each slice, the true values of gridded parameters are marked with red crosses, and parameters not plotted are fixed at their true values. Columns correspond to different discretization schemes, from left to right: coarse mesh (15^3) with no anti-aliasing; fine mesh (75^3) with no anti-aliasing; coarse mesh with anti-aliasing; fine mesh with anti-aliasing. The top row shows a vertical cross-section at $y = 0$ of the 3-D rock density field under each discretization scheme. Rows 2-5 correspond to different pairs of parameters in θ ; For row 2, θ consists of the two non-basement layer thicknesses; for row 3, θ consists of the dip direction angles of the two faults; for row 4, θ consists of the fault slips; for row 5, θ consists of the dip of the first fault and the thickness of the top layer.

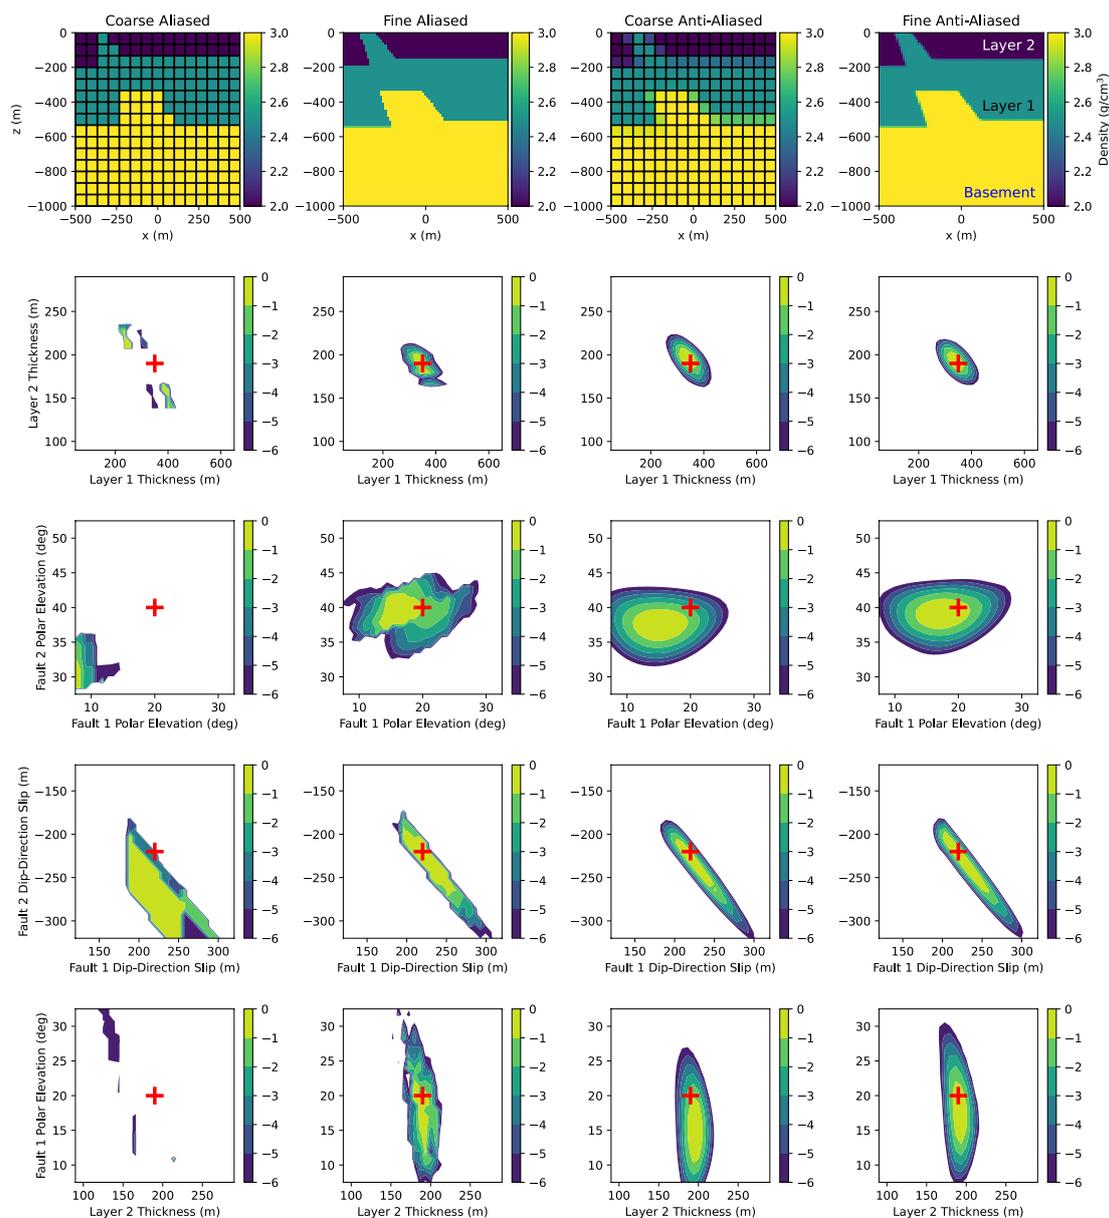


Figure 6. The log posterior distribution of pairs of parameters for Model 6. All plot properties are the same as in Fig. 5.

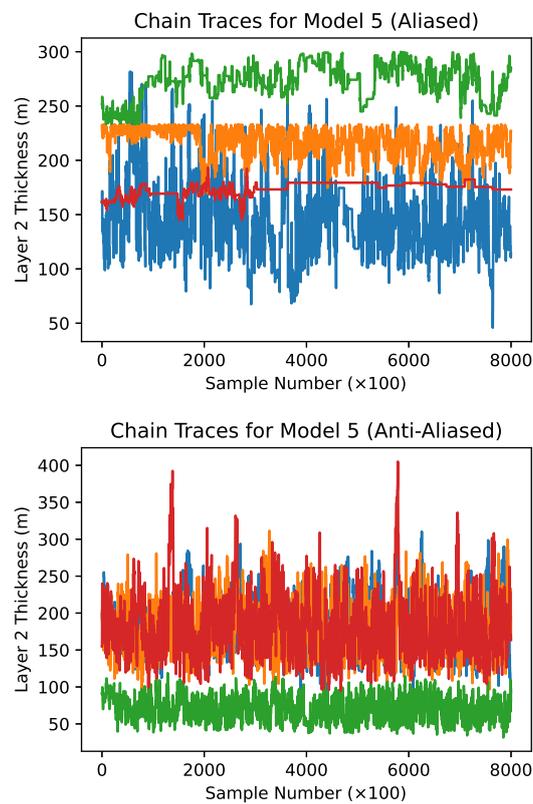


Figure 7. Trace plots for Fault 1 basement density during MCMC sampling of Model 5 (coarse mesh). The four colors in each plot represent four separate chains started from different prior draws. Traces begin at the end of burn-in and are thinned by a factor of 100 (every 100th trace point is shown). Top: aliased model. Bottom: anti-aliased model.