

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

Richard Scalzo^{1,7}, Mark Lindsay^{2,4,7}, Mark Jessell^{2,7}, Guillaume Pirot^{2,7}, Jeremie Giraud^{2,5,7}, Edward Cripps^{3,7}, and Sally Cripps^{1,6,7}

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

⁴CSIRO Mineral Resources, Kensington, WA 6151, Australia

⁵RING Team, GeoResources, Université de Lorraine, 54000, Nancy, France

⁶CSIRO Data61, Eveleigh, NSW 2015, Australia

⁷ARC 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, care must be taken with the intermediate step of rendering parametric geology in a finite-resolution discrete basis for the geophysical calculation. Calculating geophysics from naively voxelised geology, as exported from commonly used geological modeling tools, can produce a poor approximation to the true likelihood, degrading posterior inference for structural parameters. We develop a simple integrated Bayesian inversion code, called Blockworlds, showcasing a numerical scheme to calculate anti-aliased rock properties over regular meshes for use with gravity and magnetic sensors. We use Blockworlds to demonstrate anti-aliasing in the context of an implicit model with kinematic action for simple tectonic histories, showing its impact on the structure of the likelihood for gravity anomaly.

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).

20 Bayesian methods provide a probabilistically coherent framework for reasoning about uncertainty (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 (Jessell et al., 2010; 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., 25 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, but with different framing for inference problems. 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 directly define discretized spatial fields of rock properties, from which forward geophysics are calculated to reproduce observations. Model 30 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). 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 developing models with priors or regularizing terms to constrain the outputs of geophysical inversions to be geologically reasonable. Examples include: the 35 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; Zheglova 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 a flexible, high-dimensional description of static 3-D 40 properties.

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 that can capture much of the complexity of real systems 45 (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. The popular industry software package 3D GeoModeller (Calcagno et al., 2008), and the open-source code GemPy (de la Varga et al., 2018), interpolate this scalar field directly from structural geological observations by co-kriging (Lajaunie et al., 1997). Kinematic models (Jessell, 1981) use a parametrized description of volume deformations produced by tectonic events; the open-source package Noddy (Jessell and Valenta, 1993) is commonly used. The recently 50 released open-source package LoopStructural (Grose et al., 2020) combines elements of co-kriging and kinematic implicit models for additional geological richness and realism. For models 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).

Conditioning geological forward models on geophysical observations makes it tractable to perform full posterior inference in an interpretable parameter space of reduced dimension. This can be done using sampling methods such as Markov chain Monte Carlo methods ("MCMC"; Metropolis et al., 1953; Hastings, 1970), which are becoming increasingly sophisticated and widely used in geological and geophysical modeling (Mosegaard and Tarantola, 1995; Mosegaard and Sambridge, 2002; Sambridge et al., 2012). Additionally, parametric geology naturally obeys constraints imposed by the forward geological process, and can be used for inference over that process, unlike non-parametric geophysical inversions that incorporate a static, voxelized geological model as a prior mean (Giraud et al., 2017, 2018, 2019a, b; Ogarko et al., 2021; Giraud et al., 2021b). Kinematic models are especially well-suited for this approach since they represent geological hypotheses in terms of structural observations and event histories, but 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.

Probabilistic inversion workflows that incorporate forward geophysics based on forward-modelled 3-D structural geology are, to date, still uncommon, owing in part to computational challenges faced by sampling algorithms. The *pyNoddy* code (Wellmann et al., 2016) provided a wrapper to Noddy to enable Monte Carlo sampling of 3-D block models and potential fields, but not posterior sampling of an inversion conditioned on gravity or magnetic data. 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 (Beardmore et al., 2016); it featured multiple geophysical sensors and a parallel-tempered MCMC sampler for complete exploration of multi-modal posteriors. *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); but the limitations of its geological model make it an unlikely engine for general-purpose inversions. Wellmann et al. (2017) present a workflow using *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.

This paper describes a simple Bayesian inversion framework, called *Blockworlds*, which to our knowledge is the first to perform full posterior inference based on forward geophysics from an implicit geological model with kinematic elements. Although not yet intended to serve as a production-ready geological modeling package, *Blockworlds*' design addresses an obstacle for inversion workflows that resample parametrized 3-D geometry of geological units onto a discrete volumetric mesh for geophysical calculations. Discontinuous changes in rock properties across unit interfaces may 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, this alignment is not easily made without first evaluating the model, incurring additional computational overhead. Flexible geophysical inversions can adjust rock properties in individual boundary voxels to partially compensate for loss of accuracy, but models that condition rock properties directly on geological parameters lose much of this freedom. This can result in artefacts in the likelihood that hinder convergence of estimation methods, degrade posterior inference over structural parameters, and preclude the use of posterior derivative information.

To address the issue of undersampled interfaces, Blockworlds incorporates numerical anti-aliasing into its discretization step. 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). *Obsidian* anti-aliases physical rock properties using volume averages 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 the Obsidian parametrization. Blockworlds’ prescription is more general, operating on the scalar field values that define interface positions in all implicit models, and thus can be straightforwardly implemented for implicit models using other types of interpolation schemes such as co-kriging.

The paper is organized as follows. 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.

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’ theorem:

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

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 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

120 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 . The accuracy of a parametric inversion will depend on the accuracy with which $f \circ g$ can be computed.

125 Although discretization methods such as adaptive meshes (Rawlinson et al., 2014) or basis functions aligned with geological features (Wellmann and Caumon, 2018) can improve the fidelity of g , these mappings are not guaranteed to be suitable for efficient and accurate computation of f . Furthermore, existing geological engines 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.

130 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 20×20 survey grid at the surface ($z = 0$), with Gaussian measurement noise at the level of 10%
135 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

$$140 \quad \rho = \frac{3M}{4\pi R^3} \quad (2)$$

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
145 visually indistinguishable for all four versions, but the posteriors look very different.

When inverting on the coarse mesh with no anti-aliasing, the density in a given cell changes only when the sphere radius crosses over the centre of one or more mesh cells, and then it changes discontinuously. This in turn induces spurious structure into the 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
150 (red cross) lie, but none of which appropriately 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,




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.

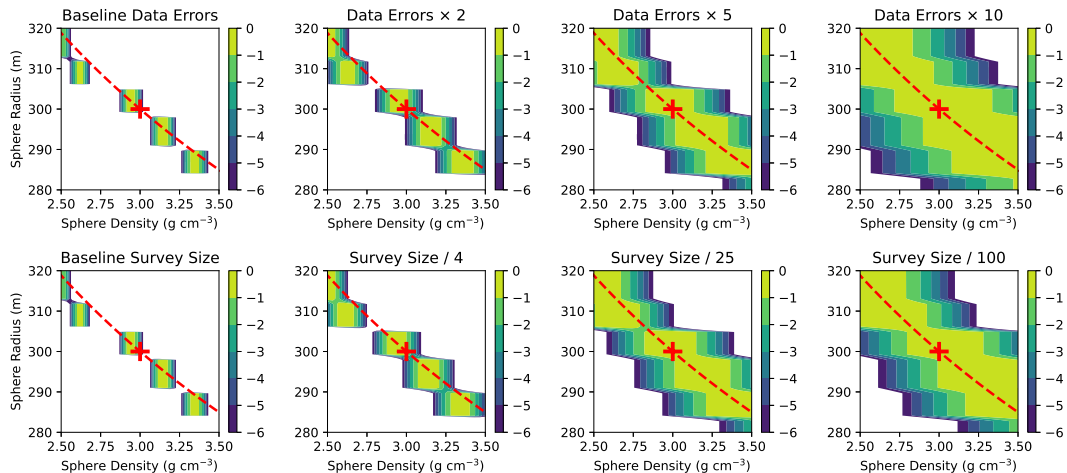


Figure 2. Influence of data error on the posterior distribution for the sphere inversion problem. The panels from left to right show the effects of multiplying the standard deviation of Gaussian noise on gravity observations by factors of 1 (baseline), 2, 5, and 10.

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
 155 run at nearly the same speed. A closer look reveals that the coarse-mesh anti-aliased posterior is slightly offset toward the bottom-left relative to the analytic curve, while the fine-mesh posterior is not. This is a product of low mesh resolution, which the anti-aliasing approximation only partially mitigates (see Sect. 2.4). Each anti-aliased interface is treated locally as a plane surface. 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 $\sim 1\%$
 160 of the sphere’s mass at the true parameter values, which the data are just sufficient to detect. Thus, while anti-aliasing may incidentally improve the accuracy of geophysical field calculations, its primary benefit is to reproduce the continuous structure of the underlying likelihood.

We will later directly test the influence of aliasing on the performance of MCMC, since the capacity for full posterior inference is one of the major advantages of parametric geological inversions. However, aliasing will cause problems for a
 165 broad variety of estimation algorithms. Even for fine mesh resolution, the likelihood is piecewise constant along spatially aliased geological parameters, and so components of the likelihood gradient in those directions are zero almost everywhere, in a measure-theoretic sense. Methods such as Nelder-Mead (for optimization) and Metropolis random walk (for sampling) that do not rely on posterior gradient information will seem to work with a fine enough mesh, although at greater computational cost and with no warning of the extent of aliasing effects unless run multiple times from different starting points. Gradient-based
 170 methods, such as stochastic gradient descent or Hamiltonian Monte Carlo, will fail catastrophically, since the zero gradient of an aliased likelihood will not reflect its geometry. The covariance matrix used to estimate local uncertainty in the fit from

an optimization algorithm will be singular, since the components of the Hessian along aliased directions will be zero. Finally, while the gravity inverse problems we consider in this paper are linear, an attempt to invert for parametric geology based on a nonlinear sensor would face discontinuous changes in the sensitivity kernel with parameters, with unpredictable effects on parameter estimation.

One could argue that the likelihood might be less badly aliased if the data constraints were weaker. Inflating the data errors can be viewed as a form of tempering, which we would expect to merge the multiple aliased modes. Figure 2 shows the results of inflating the errors in the "coarse aliased" discretization of the sphere problem by various factors until the isolated modes in the original problem have merged. The distribution becomes fully navigable only for data errors a factor of 5–10 times larger than our original problem — that is, for data errors comparable to the amplitude of the signal itself. Furthermore, the blockiness or terracing is a function of the discretization scale only, so that any estimation methods relying on gradients will break down along the aliased spatial parameters in the same way as before.

This kind of catastrophic breakdown occurs because of the restrictive conditioning of rock properties on parametric geology, so we do not expect the same kind of effect to arise in flexible geophysical inversions that derive their priors from discretized geological models (e.g. Giraud et al., 2021b). The geophysical sensor model is a smooth function of the *rock properties*, and so when the rock properties are the primary model parameters, the likelihood remains smooth. Highly restrictive model error terms that are inconsistent with the data may still hamper performance or inference. While it would be interesting to examine non-zero voxelization model errors for the kind of parametric inversion we investigate in this paper, this is a technically complex issue that we defer to future work. For now we note simply that aliasing still arises even in a parametric model that is in other respects perfectly specified, and that treating it explicitly will remove aliasing contributions to model errors in future parametric inversions.

While our example may seem artificial, widely used geological modeling tools such as Noddy and GeoModeller usually 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. Complex structure 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

The main idea behind anti-aliasing is to produce rock property values in each mesh cell such that the action of a sensor on the anti-aliased model best approximates the action of the same sensor on the same geology exported to a higher-resolution mesh. In this section we will develop an anti-aliasing prescription for gravity anomaly, which responds linearly to mass density. We frame the problem as a regression that uses the position of a geological interface within a mesh cell to predict what the mean mass density would be for that cell in a high-resolution model.

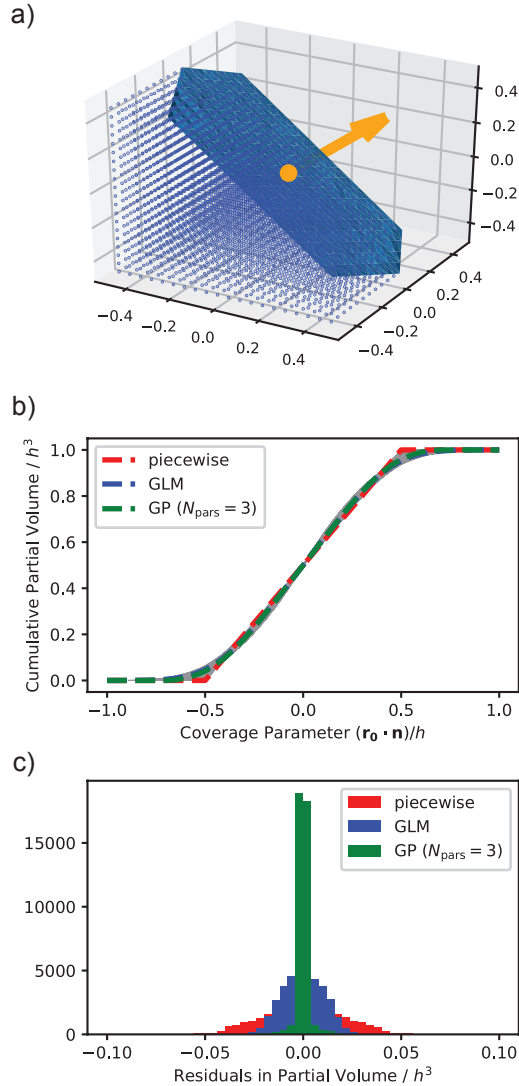


Figure 3. (a) Numerical approach to calculating partial volumes. The interface is approximated by a plane running through the voxel. A cloud of points is generated in a regular grid, and the partial volume is then the fraction of these points that satisfy $u_{\perp} = (\mathbf{r} - \mathbf{r}_0) \cdot \mathbf{n}/h = 0$, 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.

The top panel of Figure 3 shows the geometry of the problem. For simplicity, we assume constant rock density within each unit, which reduces the problem to predicting the fraction of the cell’s volume lying to either side of the interface. We further assume that the interface is flat at the spatial scale of a single mesh cell, so that it can be approximated by a plane surface defined by a point $\mathbf{r} = (r_x, r_y, r_z)$ and a unit normal $\mathbf{n} = (n_x, n_y, n_z)$. Due to the translational symmetry of a plane, the influence of \mathbf{r} can be reduced to the projected distance u_\perp of the plane from the voxel centre \mathbf{r}_0 . The unit normal \mathbf{n} can be described in terms of two angles in spherical coordinates; due to the rotational symmetries of a cubical voxel, it is enough to specify the polar angle with respect to the normal to the cube face that a ray drawn from \mathbf{r}_0 along \mathbf{n} intersects, and the azimuthal angle with respect to a coordinate system projected onto that cube face. Expressing the angles in terms of ratios of components of \mathbf{n} results in 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.

To train the regression, 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 3 shows three different possible functional forms for the regression: 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. We use Eq. 6 for further experiments in this paper.

2.4 Potential 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. The approximation giving rise to Eq. 3 and Eq. 6 will break down when interface curvature becomes significant. Common sense suggests that no length scale parameter in the inference problem should ever be less than the voxel size h . These situations can be difficult to detect from inside the inversion. Spatial scale parameters, such as fold wavelengths,

may produce high-curvature interfaces only in some regions of parameter space. Curvature measured using the Hessian adds computational overhead and may miss structures beneath the Nyquist limit if evaluated on the same mesh as the inversion itself. 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.

Another drawback in anti-aliasing is that it ignores the relative orientations of interfaces in voxels spanning multiple interfaces. If only one interface passes through a voxel, Eq. 6 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. 6 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; we expect it to be more pronounced in models with multiply-faulted interfaces in regions where the geophysical sensitivity is the highest, for example near the surface. Such regions of complexity can be identified by 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. In the next section we will discuss metrics to evaluate the global agreement of posteriors based on different mesh sizes, allowing quantitative evaluation, online or offline, of the value to the user of running at a given mesh resolution.

255 3 Experiments with kinematic models

To illustrate the impact of anti-aliasing on more realistic geologies, Blockworlds implements 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 kinematic calculations. We construct a range of 3-D models and visualize slices through each Bayesian 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, $g_0(\mathbf{r}|\rho_0) = \rho_0$. Each subsequent event is parametrized by a collection of parameters θ_i , with $\Theta_i = \{\theta_0, \theta_1, \dots, \theta_i\}$ denoting the collection of parameters for 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 between two geological units with rock properties g_+ and g_- , at projected distance u_{\perp} from the voxel centre, is given by

$$A[u_{\perp}; g_+, g_-] = g_- + (g_+ - g_-)v(u_{\perp}) \quad (7)$$

where u_{\perp} and $v(u_{\perp})$ are given by Eq. 3 and Eq. 6. The anti-aliasing must be applied in the co-moving coordinate frame of each event.

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

- **Stratigraphic layer** — $\theta_i = (\Delta z_i, \rho_i)$ for a layer with thickness Δz_i and mass density ρ_i :

$$g_i(\mathbf{r}|\Theta_i) = A[\mathbf{r} \cdot \mathbf{u}_z - \Delta z_i; g_{i-1}(\mathbf{r} - \Delta z_i \mathbf{u}_z | \Theta_{i-1}), \rho_i] \quad (8)$$

- **Fault** — $\theta_i = (\mathbf{r}_{0,i}, \mathbf{n}_i, s_i)$ for a fault passing through a point $\mathbf{r}_{0,i}$ with unit normal \mathbf{n}_i and dip-slip displacement s_i . Our parametrization constrains $\mathbf{r} = (x, y, 0)$ to lie on the surface, and uses a polar representation for \mathbf{n} derived from dip and dip direction (Pakyuz-Charrier et al., 2018b). Given these elements, we calculate $\mathbf{v} = \frac{(\mathbf{u}_z \times \mathbf{n}) \times \mathbf{n}}{|(\mathbf{u}_z \times \mathbf{n}) \times \mathbf{n}|}$ resulting in a vector pointing in the dip-slip direction parallel to the fault plane, with no strike-slip component. Then:

$$g_i(\mathbf{r}|\Theta_i) = A[(\mathbf{r} - \mathbf{r}_{0,i}) \cdot \mathbf{n}_i; g_{i-1}(\mathbf{r} + s\mathbf{v} | \Theta_{i-1}), g_{i-1}(\mathbf{r} | \Theta_{i-1})] \quad (9)$$

- **Fold** — $\theta_i = (\mathbf{n}_i, \psi_i, \phi_i, L_i, B_i)$ for a fold with axis defined by a unit vector \mathbf{n}_i , pitch angle ψ_i , phase angle ϕ_i , wavelength L_i and amplitude B_i . Given \mathbf{n}_i , we generate two unit vectors $\mathbf{v}_0 = \frac{\mathbf{u}_z \times \mathbf{n}}{|\mathbf{u}_z \times \mathbf{n}|}$ and $\mathbf{v}_1 = \frac{(\mathbf{u}_z \times \mathbf{n}) \times \mathbf{n}}{|(\mathbf{u}_z \times \mathbf{n}) \times \mathbf{n}|}$, defining an orthonormal structural frame $\{\mathbf{v}_0, \mathbf{v}_1, \mathbf{n}_i\}$. Rock elements are displaced laterally by a vector field

$$\begin{aligned} \Delta \mathbf{r} &= B_i \sin(2\pi(\mathbf{r} \cdot \mathbf{n}_i)/L + \phi) \\ &\times (\sin(\psi_i)\mathbf{v}_0 + \cos(\psi_i)\mathbf{v}_1) \end{aligned} \quad (10)$$

with the background geology displaced according to

$$g_i(\mathbf{r}|\Theta_i) = g_{i-1}(\mathbf{r} + \Delta \mathbf{r} | \Theta_{i-1}) \quad (11)$$

- **Spherical intrusion** — $\theta_i = (\mathbf{r}_{0,i}, R_i, \rho_i)$ for a sphere of radius R_i and uniform density ρ_i centered at $\mathbf{r}_{0,i}$:

$$g_i(\mathbf{r}|\Theta_i) = A[|\mathbf{r} - \mathbf{r}_{0,i}| - R_i; \rho_i, g_{i-1}(\mathbf{r} | \Theta_{i-1})] \quad (12)$$

Figure 4 shows 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

290 We construct a series of fifteen kinematic models, with true configurations shown in Fig. 5. 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 additional tectonic events, varying in order and positioning.

The first set of models focus on planar dip-slip faults, with final configurations as follows.

- **Model 1:** a typical graben structure with two intersecting faults, both exhibiting normal movement.




Figure 4. Calculation of forward geology for Models 1 and 8 through a sequence of tectonic events: the addition of two stratigraphic layers and two faults. Top: geology; bottom: anti-aliased, voxelised rendering.

- 295
- **Model 2:** a negative flower structure resulting from trans-tension, where both a sub-vertical fault and high-angle fault exhibit normal movement.
 - **Model 3:** two reverse faults dipping at the same angle and away from each other.
 - **Model 4:** two parallel sub-vertical faults of opposing displacement: the left fault with left-side-up movement, the right fault with right-side-up movement.
- 300
- **Model 5:** a positive flower structure resulting from transpression, where both a sub-vertical fault and high-angle fault exhibit reverse movement.




Figure 5. Voxelized final states of all models in suite, displaying 3-D geology in cross section.

- **Model 6:** two high-angle faults, the left displaying thrust movement and the right displaying normal movement.
- **Model 7:** a similar scenario to Model 6, except the movement for both faults is reverse.

We also include some additional folded configurations, selected to test the anticipated limitations of anti-aliasing described in Section 2.4. These use the same stratigraphic history as the two-fault events except when parameter changes are explicitly noted.

- **Model 8:** an upright, open fold of 500 m wavelength, with a west-dipping, high-angle reverse fault offsetting the fold limbs.
- **Model 9:** a tighter fold of wavelength 300 m, cut across by a low-angle reverse fault to create near-surface voxels in the block model traversed by multiple interfaces.
- **Model 10:** a more extreme version of Model 9 created by reducing the fold wavelength to 150 m, near the Nyquist sampling limit for the coarse voxel size.
- **Model 11:** similar to Model 10, but with a deeper surface layer to create aliased voxels farther beneath the surface.

Table 1. Parameter true values and prior distributions for fault-focused kinematic models, labeled M1 to M7. The prior mean is set to the true value for each parameter.

Parameter	M1	M2	M3	M4	M5	M6	M7	Prior
Basement Density (g cm^{-3})	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	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	Lognormal ($\sigma = 0.1 \text{ g cm}^{-3}$)
Layer 2 Thickness (m)	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	Lognormal ($\sigma = 0.1 \text{ g cm}^{-3}$)
Fault 1 Contact X Position (m)	-400	-450	-50	-250	-450	-400	-400	Normal ($\sigma = 1 \text{ m}$)
Fault 1 Contact Y Position (m)	0	0	0	0	0	0	0	Normal ($\sigma = 1 \text{ m}$)
Fault 1 Polar Elevation (deg)	45	45	-20	0	30	20	20	vMF ($\kappa = 25$)
Fault 1 Polar Azimuth (deg)	0	0	0	0	0	0	0	vMF ($\kappa = 25$)
Fault 1 Dip-Direction Slip (m)	-220	-220	-220	-220	220	220	140	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 ($\sigma = 150 \text{ m}$)

315

– **Model 12:** a folded fold, with two folds of wavelength 300 m and with the second upright fold at a 70-degree azimuthal with respect to the first.

– **Model 13:** two upright folds of different wavelengths (300 m and 500 m) with axes aligned.

– **Model 14:** a version of Model 13 with thinner surface stratigraphic layers (75 m each) to create near-surface voxels with multiple parallel interfaces.

320

– **Model 15:** a stratigraphy with a thinner surface layer (150 m) and a low-angle reverse fault that has been folded, again to create a final state with high complexity in near-surface voxels.

325

We choose prior distributions for the parameters to simulate the realistic incorporation of structural geological knowledge, as shown in Table 1 and Table 2. 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 ($\sigma = 1 \text{ m}$), appropriate for surface observation localized by GPS. The fault slip parameters reflect rough estimates ($\sigma = 150 \text{ m}$), 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-

Table 2. Parameter true values and prior distributions for fold-focused kinematic models, labeled M8 to M15.

* In M15, the stratigraphy is faulted first, then folded.

Parameter	M8	M9	M10	M11	M12	M13	M14	M15*	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	75	150	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	390	190	190	75	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	0	0	0	0	0	0	0	vMF ($\kappa = 25$)
Fold 1 Axis Azimuth (deg)	0	0	0	0	0	0	0	0	vMF ($\kappa = 25$)
Fold 1 Pitch Angle (deg)	0	0	0	0	0	0	0	0	Normal ($\sigma = 1 \text{ deg}$)
Fold 1 Phase Angle (deg)	0	0	0	0	0	0	0	0	Normal ($\sigma = 1 \text{ deg}$)
Fold 1 Wavelength (m)	1000	300	150	150	300	300	300	300	Lognormal ($\sigma = 100 \text{ m}$)
Fold 1 Amplitude (m)	100	100	100	100	100	100	200	200	Lognormal ($\sigma = 15 \text{ m}$)
Fold 2 Axis Elevation (deg)	—	—	—	—	0	0	0	—	vMF ($\kappa = 25$)
Fold 2 Axis Azimuth (deg)	—	—	—	—	70	0	0	—	vMF ($\kappa = 25$)
Fold 2 Pitch Angle (deg)	—	—	—	—	0	0	0	—	Normal ($\sigma = 1 \text{ deg}$)
Fold 2 Phase Angle (deg)	—	—	—	—	0	0	0	—	Normal ($\sigma = 1 \text{ deg}$)
Fold 2 Wavelength (m)	—	—	—	—	300	500	500	—	Lognormal ($\sigma = 100 \text{ m}$)
Fold 2 Amplitude (m)	—	—	—	—	100	100	200	—	Lognormal ($\sigma = 15 \text{ m}$)
Fault 1 Contact X Position (m)	250	450	450	450	—	—	—	450	Normal ($\sigma = 1 \text{ m}$)
Fault 1 Contact Y Position (m)	0	0	0	0	—	—	—	0	Normal ($\sigma = 1 \text{ m}$)
Fault 1 Polar Elevation (deg)	-205	-255	-255	-245	—	—	—	-245	vMF ($\kappa = 25$)
Fault 1 Polar Azimuth (deg)	0	0	0	0	—	—	—	0	vMF ($\kappa = 25$)
Fault 1 Dip-Direction Slip (m)	200	200	200	200	—	—	—	200	Normal ($\sigma = 150 \text{ m}$)

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.

330 For each model, we examine four regimes of resolution and aliasing. We consider a "coarse" $15 \times 15 \times 15$ mesh ($h = 66.6 \text{ m}$) as well as a "fine" $75 \times 75 \times 75$ mesh ($h = 13.3 \text{ m}$), 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.

335 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

sample standard deviation of the data across the survey. Each data point y_j is thus generated according to

$$p(y_j|\boldsymbol{\theta}, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{(y_j - f_j(\boldsymbol{\theta}))^2}{2\sigma^2}\right], \quad (13)$$

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
 340 hierarchical prior for σ^2 . If we choose this prior to be an inverse gamma distribution,

$$p(\sigma^2|\alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} \sigma^{-2(\alpha+1)} \exp\left(-\frac{\beta}{\sigma^2}\right), \quad (14)$$

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

$$\begin{aligned} p(y_j|\boldsymbol{\theta}, \alpha, \beta) &= \int_0^\infty p(y_j|\boldsymbol{\theta}, \sigma^2) p(\sigma^2|\alpha, \beta) d\sigma^2 \\ &= \frac{\Gamma(\alpha + \frac{1}{2})}{\Gamma(\alpha)\sqrt{2\pi\beta}} \left[\frac{(y_j - f_j(\boldsymbol{\theta}))^2}{2\beta} + 1\right]^{-(\alpha + \frac{1}{2})} \end{aligned} \quad (15)$$

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,

$$p(\mathbf{d}|\boldsymbol{\theta}, \alpha, \beta) = \prod_{j=1}^N p(y_j|\boldsymbol{\theta}, \alpha, \beta). \quad (16)$$

350 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,

$$\begin{aligned} g_z(\mathbf{r}) &= -G\rho_0 \sum_{i=1}^2 \sum_{j=1}^2 \sum_{k=1}^2 \mu_{ijk} \\ &\times \left[x'_i \ln(y'_i + r'_{ijk}) + y'_j \ln(y'_j + r'_{ijk}) \right. \\ &\quad \left. + 2z'_k \arctan\left(\frac{x'_i y'_j}{z'_k r'_{ijk}}\right) \right] \end{aligned} \quad (17)$$

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
 360 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 anti-aliasing on chain mixing. The aliased posterior has no useful derivative information for structural parameters, so we are limited to random walk proposals. For our tests we choose the adaptive Metropolis random walk algorithm of Haario et al. (2001), which is used in contemporary studies such as Wellmann et al. (2017) and is the simplest algorithm likely to give good performance on posteriors with the kind of strong covariances seen in the conditional posterior slices. This algorithm 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 :

$$q(\theta'|\theta_t) \sim N(\theta_t, \eta_d \Sigma_t), \quad (18)$$

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

$$\Sigma_t = \begin{cases} \Sigma_0 & \text{for } t \leq t_0 \\ \text{cov}\{\theta_i\} & \text{for } t > t_0 \end{cases} \quad (19)$$

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. States are accepted according to the usual Metropolis-Hastings criterion, setting $\theta_{t+1} = \theta'$ with probability

$$\alpha = \min \left\{ 1, \frac{p(\mathbf{d}|\theta')p(\theta')q(\theta|\theta')}{p(\mathbf{d}|\theta)p(\theta)q(\theta'|\theta)} \right\} \quad (20)$$

and $\theta_{t+1} = \theta_t$ with probability $1 - \alpha$. 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).

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.

Finally, to quantify the overall accuracy of the posterior density calculation, we use the Kullback-Liebler divergence

$$D_{\text{KL}}(p||q) = \int (\log p(\boldsymbol{\theta}) - \log q(\boldsymbol{\theta})) p(\boldsymbol{\theta}) d\boldsymbol{\theta} \quad (21)$$

where $p(\boldsymbol{\theta}) = p_{\text{lo}}(\boldsymbol{\theta}|\mathbf{y})$ represents the posterior based on a low-resolution mesh (with or without anti-aliasing), and $q(\boldsymbol{\theta}) = p_{\text{hi}}(\boldsymbol{\theta}|\mathbf{y})$ represents the high-resolution, anti-aliased posterior. D_{KL} is a global measure sensitive to biases and covariances across all parameters. Since it requires integration over the posterior, this metric is not available to check results of single-point optimization algorithms. We use the MCMC chain itself to calculate the integral, representing it as the difference in log probability density averaged over samples from the low-resolution posterior.

4 Results

4.1 Posterior cross sections

Two-dimensional slices through the posterior distribution for selected variable pairs are shown for Model 1 (in Fig. 6) and Model 6 (in Fig. 7). Similar sets of plots for all models can be regenerated automatically by scripts in the Blockworlds repository.

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).

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

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.

Models 10 and 11 show yet more dramatic differences between the low-resolution and high-resolution models, even with
425 anti-aliasing. This is expected, since the true fold wavelength for these models is close to the Nyquist limit for the coarse mesh
scale, and sub-Nyquist wavelengths are not excluded by the priors. Model 9 shows that when the fold wavelength is relaxed
from $\sim 2h$ to $\sim 5h$, anti-aliasing gives much better results. In Model 11, where the undersampled structure is positioned farther
beneath the surface, the apparent effects are less apparent than in Model 10, but still appear in directly linked variables such as
the fold wavelength.

430 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. In Model 10 the layer thicknesses are the only model
parameters that can be robustly estimated without refining the mesh to resolve structure.

435 These results confirm that our algorithm can reproduce the continuous behavior in an underlying posterior distribution, and
can enable significant computational savings even for complex models as long as the interfaces are smooth on the mesh scale.
It cannot — and is not intended to — recover more complex structures beneath the mesh scale that violate the core assumptions
under which it was derived.

4.2 MCMC sampling

440 Table 3 shows the average, best-case, and worst-case values of τ and \hat{R} across all model parameters, as well as the K-L
divergence D_{KL} . 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
445 after 10^5 samples. The anti-aliased chains converge more consistently to the same distribution, with final PSRF values close
to 1 for all variables for most models, and much closer to 1 than the aliased chains even on the more challenging, adversarial
model configurations.

Models 5, 6, 10, 14, and 15 are challenging models for which the anti-aliased chain also has trouble mixing fully. An
example trace plot from Model 5 is shown in Fig. 8. One of the chains in the anti-aliased version has fallen into a different
450 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.

455 Complementing the trace plots and MCMC performance metrics as a global measure of posterior inference *accuracy*, the
K-L divergence shows that anti-aliasing usually improves correspondence with the high-resolution posterior by a large margin.
Anti-aliased models with large values of D_{KL} correspond well with sub-Nyquist structures or with regions of complexity close

to the surface, confirming our intuition that inference in these models benefits from finer mesh resolution. The aliased posteriors show differences with the high-resolution posteriors spanning many orders of magnitude in probability density, averaged over
 460 those modes the chain was able to reach. Since the aliased chains did not in general mix well, their values of D_{KL} should be taken as indicative, whereas the anti-aliased chains with low \hat{R} represent more robust estimates of the true divergence.

To fully characterize uncertainty in truly multi-modal problems, 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
 465 posteriors are also more likely to represent distinct interpretations, rather than poorly-resolved strong correlations between variables.

5 Discussion and Future Work

We have experimentally demonstrated the use and benefit of anti-aliasing only for kinematic models acting on a regular rectangular mesh. However, we see several natural extensions to this work that extend the range and impact of anti-aliasing for
 470 parametric geological models. We describe these future directions in the following section.

5.1 Anti-aliasing in terms of the implicit scalar field

Although Blockworlds forward-models geology through the action of parametrized tectonic events, it is still an implicit model in the sense that the unit interfaces are defined to be level sets of a scalar field, which we parametrize simply in terms of the distance to the interface. The mathematics of our anti-aliasing prescription generalizes to any differentiable scalar field,
 475 including those used in co-kriging models that interpolate structural measurements (Lajaunie et al., 1997; Calcagno et al., 2008). These models will have different posterior geometries than kinematic models, since the modification of structural parameters produces more localized changes in the block models. However, we see no intrinsic reason why the following treatment for general scalar fields should not work just as well as for other types of implicit models, provided the interpolated geological structures have minimal curvature on the mesh scale.

480 For implicit geological models, each interface is defined as the level set $\Phi(\mathbf{r}) = \Phi_0$ of some scalar field Φ . 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. In interpolation models, 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 scalar field to represent a physical distance in the neighborhood of an interface: for a voxel centered at \mathbf{r}_0 , a first-order
 485 approximation of Φ gives

$$\Phi(\mathbf{r}) = \Phi(\mathbf{r}_0) + (\mathbf{r} - \mathbf{r}_0) \cdot \nabla\Phi + O(\mathbf{r} - \mathbf{r}_0)^2 = \Phi_0, \quad (22)$$

which, using Equation 3 and writing the interface unit normal as $\mathbf{n} = \nabla\Phi/|\nabla\Phi|$, can be rearranged to give

$$u_{\perp} = (\Phi_0 - \Phi(\mathbf{r}_0))/|h\nabla\Phi|. \quad (23)$$

In this way, each voxel can be anti-aliased simply by evaluating the scalar field and its gradient at the voxel center. For co-kriging models, the scalar field itself must be evaluated at the voxel centers to render it, and the field gradient can be evaluated straightforward using the same trained model by differentiating the kernel.

5.2 Anti-aliasing in curvilinear coordinates

The development of the anti-aliasing mapping in Section 2.3 and its generalization in 5.1 assume 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.

Another useful avenue of future work would be to extend Equation 23 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 3, defining the projection of a reference point in the voxel onto an interface running through the voxel, then becomes

$$u_{\perp} = \mathbf{u} \cdot \frac{1}{|\mathbf{J}\mathbf{n}|} \mathbf{J}\mathbf{n}, \quad (24)$$

and Equation 3, defining u_{\perp} in terms of scalar field values for an implicit model, becomes

$$u_{\perp} = (\Phi_0 - \Phi(\mathbf{r}_0)) / |\mathbf{J}\nabla\Phi|. \quad (25)$$

5.3 Anti-aliasing for other geophysical sensors

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 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 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 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.

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 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. Finally, the K-L divergence between versions of the log posterior using different mesh resolutions can be calculated using MCMC samples, either online or in post-processing as a diagnostic, allowing the user to determine how much the problem would benefit from additional mesh refinement.

We have focused on calculations over 3-D Cartesian 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 discretization schemes built to align with geological interfaces, such as pillar grids (Wellmann and Caumon, 2018), may mitigate aliasing along the pillar, but still suffer aliasing of features cutting across the pillar axis and are more limited in their representational power. Alternate methods for calculating gravity and magnetics, such as the Cauchy surface method (Zhdanov and Liu, 2013; Cai and Zhdanov, 2015) or finite-volume methods on tetrahedral meshes Götze and Lahmeyer (1988); Schmidt et al. (2020) 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.

550 *Code and data availability.* The version of the Blockworlds model code used in this paper is archived with Zenodo (DOI: 10.5281/zenodo.5759225). 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

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

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

$$B_k = \frac{1}{M-1} \sum_{i=1}^M (\hat{\theta}_{ki} - \tilde{\theta}_k) \quad (\text{A3})$$

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

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

565 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

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

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:

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

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

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

580 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
585 between samples, via the modified metric

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

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

$$\begin{aligned} \hat{\text{var}}(\hat{V}_k) &= \left(\frac{N-1}{N}\right)^2 \frac{1}{M} \hat{\text{var}}(s_{ki}^2) \\ &+ \left(\frac{M+1}{MN}\right)^2 \frac{2}{M-1} B^2 \\ 590 &+ 2 \frac{(M+1)(N-1)}{MN^2} \frac{N}{M} \\ &\times \left[\hat{\text{cov}}(s_{ki}^2, \hat{\theta}_{ki}^2) - 2\tilde{\theta}_k \hat{\text{cov}}(s_{ki}^2, \hat{\theta}_{ki}) \right]. \end{aligned} \quad (\text{A9})$$

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
595 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.

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

600 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 thank Alistair Reid and Malcolm Sambridge for useful discussions.

References

- Backus, G. and Gilbert, F.: The Resolving Power of Gross Earth Data, *Geophysical Journal of the Royal Astronomical Society*, 16, 169–205, 605 <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* 610 *Astronomical Society*, 13, 247–276, <https://doi.org/10.1111/j.1365-246X.1967.tb02159.x>, 1967.
- Beardsmore, 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.
- 615 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.
- 620 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 625 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*, 630 *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, 635 <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.

- 640 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.
- Frodeman, R.: Geological reasoning: Geology as an interpretive and historical science, *Geological Society of America Bulletin*, 107, 960–
645 0968, [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.
- 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.
- 650 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 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 geo-
655 logical 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., 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.
- 660 Giraud, J., Ogarko, V., Martin, R., Jessell, M., and Lindsay, M.: Structural, petrophysical and geological constraints in potential field inversion 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.
- Götze, H. and Lahmeyer, B.: Application of three-dimensional interactive modeling in gravity and magnetics, *Geophysics*, 53, 1096–1108, <https://doi.org/10.1190/1.1442546>, 1988.
- 665 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.
- 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.
- 670 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.
- 675 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.
- 680 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.
- Koene, E. F. M., Wittsten, J., and Robertsson, J. O. A.: Finite-difference modeling of 2-D wave propagation in the vicinity of dipping
685 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.
- 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.
- 690 Li, Y. and Oldenburg, D. W.: 3-D inversion of gravity data, *Geophysics*, 63, 109–119, <https://doi.org/10.1190/1.1444302>, 1998.
- 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., Perrouty, 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.
- 695 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.
- 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
700 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.
- 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.
- 705 Moczo, P., Kristek, J., Vavrycuk, 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.
- 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.
- 710 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.
- 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.

- 715 Ogarko, V., Giraud, J., Martin, R., and Jessell, M.: Disjoint interval bound constraints using the alternating direction method of multipliers 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.
- 720 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*, <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.
- 725 Öztireli, A. C.: A Comprehensive Theory and Variational Framework for Anti-aliasing Sampling Patterns, *Computer Graphics Forum*, 39, 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, <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.
- 735 Perrouy, 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.
- 740 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.
- 745 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.
- 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.
- 750 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.

- 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, 20110547–20110547, <https://doi.org/10.1098/rsta.2011.0547>, ISBN: 1364503X (ISSN), 2012.
- 755 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.
- 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.
- 760 Schmidt, S., Anikiev, D., Götze, H.-J., Gomez Garcia, A., Gomez Dacal, M. L., Meessen, C., Plonka, C., Rodriguez Piceda, C., Spooner, C., and Scheck-Wenderoth, M.: IGMAS+ – a tool for interdisciplinary 3D potential field modelling of complex geological structures, in: *EGU General Assembly 2020*, pp. EGU2020–8383, <https://doi.org/10.5194/egusphere-egu2020-8383>, <https://meetingorganizer.copernicus.org/EGU2020/EGU2020-8383.html>, 2020.
- 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.
- 765 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.
- 770 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.
- 775 Wellmann, J. F., Thiele, S. T., Lindsay, M. D., and Jessell, M. W.: pynoddy 1.0: an experimental platform for automated 3-D kinematic and potential field modelling, *Geoscientific Model Development*, 9, 1019–1035, <https://doi.org/10.5194/gmd-9-1019-2016>, 2016.
- 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.
- 780 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.
- 785 Zhiglova, 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.

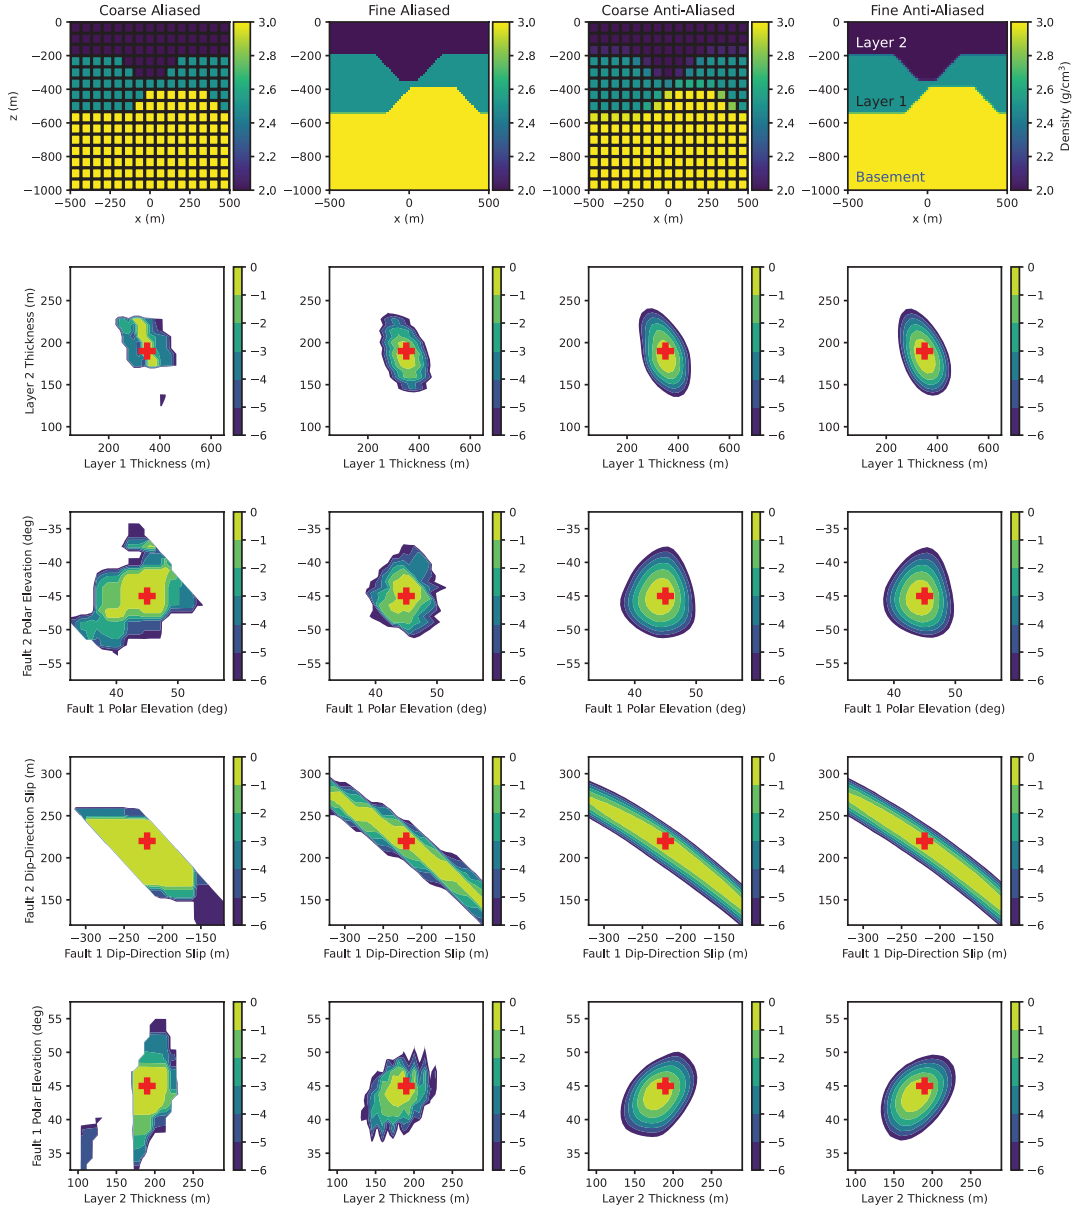


Figure 6. The log posterior distribution of pairs of parameters, denoted by θ , in kinematic model 1, relative to the maximal value $\log_{10}[p(\theta|\mathbf{d})/p(\hat{\theta}_{\text{MAP}}|\mathbf{d})]$, where $\hat{\theta}_{\text{MAP}} = \text{argmax}_{\theta} p(\theta|\mathbf{d})$. In each slice, 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 7. The log posterior distribution of pairs of parameters for Model 6. All plot properties are the same as in Fig. 6.

Table 3. Performance metrics for each MCMC run: mean (best-case, worst-case) integrated autocorrelation time τ in samples ($\times 10^3$); mean (best-case, worst-case) potential scale reduction factor \hat{R} calculated across 4 runs starting from independent prior draws; and sample Kullback-Liebler divergence D_{KL} between the given model and the fine-mesh anti-aliased version, calculated by combining samples from all 4 runs thinned by a factor of 1000.

Runs marked with 'AA' are coarse-mesh anti-aliased, while those marked with only a single-digit model number have anti-aliasing turned

Model	$\tau_{\text{mean}} (\tau_{\text{min}}, \tau_{\text{max}})$	$\hat{R}_{\text{mean}} (\hat{R}_{\text{min}}, \hat{R}_{\text{max}})$	D_{KL}
1	$0.6 (0.3, 0.8) \times 10^3$	1.34 (1.00, 3.53)	14.78
1-AA	$0.4 (0.2, 0.6) \times 10^3$	1.01 (1.00, 1.04)	0.12
2	$2.3 (0.6, 5.3) \times 10^3$	1.59 (1.00, 5.83)	17.63
2-AA	$0.2 (0.2, 0.3) \times 10^3$	1.00 (1.00, 1.01)	-0.08
3	$2.5 (0.3, 6.2) \times 10^3$	2.65 (1.00, 15.4)	21.65
3-AA	$1.0 (0.2, 3.2) \times 10^3$	1.11 (1.00, 1.47)	0.01
4	$3.5 (0.6, 6.9) \times 10^3$	3.03 (1.01, 21.6)	19.27
4-AA	$0.6 (0.1, 1.9) \times 10^3$	1.01 (1.00, 1.06)	-0.08
5	$3.5 (0.8, 6.6) \times 10^3$	214 (1.04, 2720)	184.5
5-AA	$0.3 (0.2, 0.3) \times 10^3$	19.9 (1.00, 120)	-3.14
6	$4.1 (0.4, 6.6) \times 10^3$	6.92 (1.01, 27.0)	27.94
6-AA	$0.6 (0.2, 1.2) \times 10^3$	12.1 (1.00, 72.1)	-8.71
7	$2.0 (0.6, 5.6) \times 10^3$	2.39 (1.00, 8.48)	29.81
7-AA	$0.3 (0.2, 0.4) \times 10^3$	1.00 (1.00, 1.01)	-0.45
off. 8	$2.5 (0.3, 6.7) \times 10^3$	5.52 (1.00, 36.9)	29.76
8-AA	$0.1 (0.1, 0.1) \times 10^3$	1.00 (1.00, 1.00)	-1.74
9	$1.9 (0.5, 4.0) \times 10^3$	13.0 (1.05, 89.5)	12.18
9-AA	$0.3 (0.1, 0.6) \times 10^3$	4.97 (1.00, 25.0)	0.53
10	$3.6 (1.3, 5.4) \times 10^3$	1.60 (1.02, 5.47)	-2.78
10-AA	$2.6 (2.1, 3.3) \times 10^3$	1.09 (1.00, 1.26)	-10.76
11	$0.2 (0.2, 0.2) \times 10^3$	1.00 (1.00, 1.01)	1.72
11-AA	$0.1 (0.1, 0.1) \times 10^3$	1.00 (1.00, 1.01)	1.14
12	$1.8 (0.4, 5.6) \times 10^3$	75.4 (1.03, 659)	4.33
12-AA	$0.1 (0.1, 0.2) \times 10^3$	1.00 (1.00, 1.01)	-5.38
13	$2.4 (1.0, 3.7) \times 10^3$	17.0 (1.01, 89.9)	11.85
13-AA	$0.2 (0.1, 0.3) \times 10^3$	1.00 (1.00, 1.02)	-8.11
14	$5.9 (1.5, 9.5) \times 10^3$	160 (1.37, 2350)	101.4
14-AA	$1.1 (0.2, 1.7) \times 10^3$	16.5 (1.00, 133)	8.70
15	$3.7 (0.5, 8.2) \times 10^3$	176 (1.20, 2180)	47.16
15-AA	$2.0 (0.1, 6.2) \times 10^3$	1.73 (1.01, 7.09)	-11.97

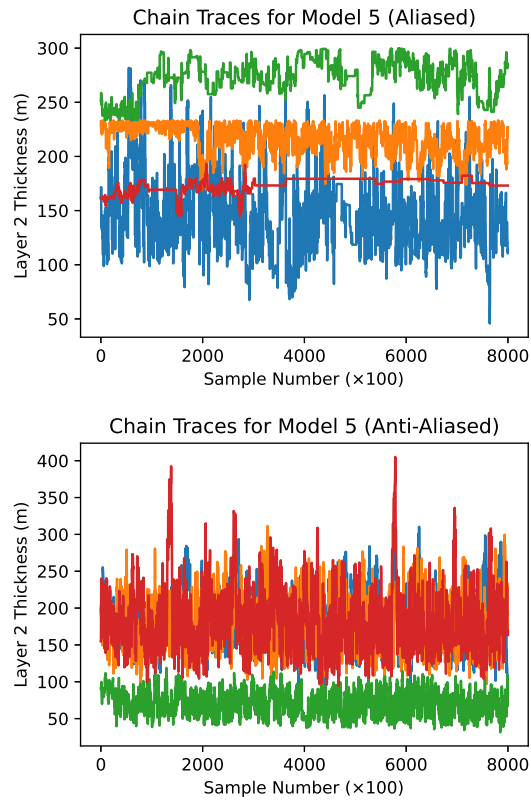


Figure 8. Trace plots for Fault 1 basement density during MCMC sampling of Model 5 (coarse mesh). The four colors in each plot represent four separate chains started from different prior draws. Traces begin at the end of burn-in and are thinned by a factor of 100 (every 100th trace point is shown). Top: aliased model. Bottom: anti-aliased model.