



Probabilistic Inversions for Time–Distance Helioseismology

Jason Jackiewicz¹ 

Received: 6 February 2020 / Accepted: 2 July 2020 / Published online: 7 October 2020
© Springer Nature B.V. 2020

Abstract Time–distance helioseismology is a set of powerful tools to study localized features below the Sun’s surface. Inverse methods are needed to robustly interpret time–distance measurements, with many examples in the literature. However, techniques that utilize a more statistical approach to inferences, and that are broadly used in the astronomical community, are less-commonly found in helioseismology. This article aims to introduce a potentially powerful inversion scheme based on Bayesian probability theory and Monte Carlo sampling that is suitable for local helioseismology. We first describe the probabilistic method and how it is conceptually different from standard inversions used in local helioseismology. Several example calculations are carried out to compare and contrast the setup of the problems and the results that are obtained. The examples focus on two important phenomena that are currently outstanding issues in helioseismology: meridional circulation and supergranulation. Numerical models are used to compute synthetic observations, providing the added benefit of knowing the solution against which the results can be tested. For demonstration purposes, the problems are formulated in two and three dimensions, using both ray- and Born-theoretical approaches. The results seem to indicate that the probabilistic inversions not only find a better solution with much more realistic estimation of the uncertainties, but they also provide a broader view of the range of solutions possible for any given model, making the interpretation of the inversion more quantitative in nature. The probabilistic inversions are also easy to set up for a broad range of problems, and they can take advantage of software that is publicly available. Unlike the progress being made in fundamental measurement schemes in local helioseismology that image the far side of the Sun, or have detected signatures of global Rossby waves, among many others, inversions of those measurements have had significantly less success. Such statistical methods may help overcome some of these barriers to move the field forward.

Keywords Helioseismology, inverse modeling · Interior, convective zone · Oscillations, solar

✉ J. Jackiewicz
jasonj@nmsu.edu

¹ Department of Astronomy, New Mexico State University, Las Cruces, NM 88003, USA

1. Introduction

Inversions play a critical role for the interpretation of helioseismic measurements. In global helioseismology, inversions of the frequency spectrum of the Sun's low-degree modes have been used to determine its interior structure (Christensen-Dalsgaard et al., 1996), including internal differential rotation (Thompson et al., 2003; Howe, 2009). In the local framework of helioseismology, inversions of wave-packet travel times, or of ring-diagram parameters, are employed for measuring sub-surface flows (Komm et al., 2007), such as meridional circulation (Giles et al., 1997; Zhao et al., 2013; Jackiewicz, Serebryanskiy, and Kholikov, 2015; Rajaguru and Antia, 2015), supergranulation (Zhao and Kosovichev, 2003; Švanda, 2012), and velocity structures in the vicinity of sunspots (Couvidat, Birch, and Kosovichev, 2006; Gizon et al., 2009; Moradi et al., 2010).

Helioseismic inversions estimate sub-surface quantities. Two popular classes of techniques used for this estimation are Regularized Least Squares (RLS) and Optimally Localized Averages (OLA) (Gough and Thompson, 1991; Christensen-Dalsgaard, Hansen, and Thompson, 1993; Pijpers and Thompson, 1994; Schou, Christensen-Dalsgaard, and Thompson, 1994; Corbard et al., 1997; Jensen, Jacobsen, and Christensen-Dalsgaard, 1998; Jackiewicz, Gizon, and Birch, 2008; Švanda et al., 2011; Jackiewicz et al., 2012; Korda and Švanda, 2019). These methods rely on inversions of large matrices that may suffer from numerical instabilities when the matrices are ill-conditioned, which they often are. Furthermore, the cost or misfit function to be minimized may be very irregular in the parameter space, and strong regularization or smoothing often needs to be applied. Due to various tuning strategies, it is recognized that computing these inversions is sometimes as much “art” as science (Basu, 2016).

An alternative framework to interpret observational data relies on Bayesian theory and statistics. In its simplest form, Bayesian inference combines prior information on a model and its parameters with observational data to produce a posterior probability distribution function (PDF hereafter) of the model parameters. The PDF represents the complete solution to the inverse problem, and all of the information is formulated in terms of probabilities. For this to work, one must know the statistical properties of the noise in the data. For helioseismology, these properties are typically well understood (Gizon and Birch, 2004; Fournier et al., 2014).

The Bayesian computation of the PDF spans the whole model space. In the case that the PDF is Gaussian, then the inverse problem can be straightforwardly solved using the methods described above to give a reasonable “most-probable model.” However, if the nature of the data or prior information is complex, such that the PDF is not very smooth or is multi-modal, then a most-probable model has little meaning. In this case, it is important to characterize the full shape of the PDF as to provide realistic uncertainties on the estimations. The problem becomes one of sampling, rather than optimization.

This is where methods of the Markov Chain Monte Carlo (MCMC hereafter) approach come in. Modern MCMC techniques are actively being developed that efficiently and effectively sample multi-modal, multi-dimensional distribution functions of the parameter space. They work by drawing random samples that are distributed according to the properties of the PDF. Coupling these samplers to Bayesian inferences to solve problems is what will be referred to in this article as *probabilistic inversions*.

Apart from seismology of the Earth, which has a very mature MCMC inversion literature (see the references within, and the references to Sambridge and Mosegaard, 2002), global helioseismology and asteroseismology have employed probabilistic methods much more sparsely. The applications have not primarily been for standard inversions either, but for

statistical measurements of the properties of individual seismic mode parameters (frequencies, amplitudes, linewidths) (e.g. the Diamonds package of Corsaro and De Ridder, 2014). Local helioseismology has seen even less adoption. A notable exception is the current solar coronal seismology work led by Arregui (see Arregui, 2018, and the references therein). In other areas of astronomy, probabilistic inversions have proven to be a robust way to interpret astronomical observations (Sharma, 2017). Indeed, in a relatively recent article presenting a new MCMC Bayesian tool for the Python programming language, Foreman-Mackey et al. (2013) discuss its usage for general astronomical problems. That publication has over 3000 citations in ADS (as of January 2020). A couple dozen are related to asteroseismology, but none to local helioseismology.

Therefore, we feel that it could be useful to provide some examples of probabilistic inversions for local helioseismology. This article is written for people working in the field of solar physics and helioseismology who might not be very familiar with the utility of such techniques. We caution that there will be few details in the derivation of Bayesian statistics and MCMC, so that more focus can be applied to example tools and methods that can be used to solve certain classes of helioseismic problems.

The rest of the article is organized as follows: In Section 2, the basic formulations of standard linear inversions and Bayesian inferences are described, as well as how they are connected in certain cases. Section 3 provides examples of both types of inversions for two relevant problems in local helioseismology: inferring the flows of meridional circulation and those of supergranulation. This section compares in detail the results and outputs from the inversions. The final sections present a discussion of when one inversion technique might be preferable to another, and we end with a summary of the work presented. The appendices provide more details of the inversions.

2. Deterministic and Bayesian Inferences

2.1. Formulation of Standard Helioseismic Inversions

The majority of local helioseismic inversions published over the last decade or so rely on variants of the Optimally Localized Averages (OLA) method that was developed for terrestrial seismology by Backus and Gilbert (1968). The most widely used form of this class of *linear*, deterministic inversions may be the Subtractive OLA (SOLA: Pijpers and Thompson, 1994; Jackiewicz, Gizon, and Birch, 2008; Švanda et al., 2011; Jackiewicz et al., 2012; Greer, Hindman, and Toomre, 2016). However, some recent studies have begun to employ full-waveform techniques that are very promising, yet very computationally demanding. The methods are iterative in nature and do not assume linearity between the response of seismic waves and the perturbation. Hanasoge and collaborators are at the forefront of this effort (Hanasoge et al., 2011; Hanasoge, 2014; Bhattacharya and Hanasoge, 2016), which also has a mature history in terrestrial seismology.

In any case, SOLA inversions essentially provide a way to infer the perturbation one spatial location at a time. Unlike RLS-type algorithms, which try to find a best fit to the data, SOLA forms linear combinations of the data (while minimizing the errors) that spatially localize the inference. The solution can critically depend on the tuning of certain parameters. These are not model parameters, but parameters that control the type of solution one desires. There are tradeoffs in the solution, such as those between spatial resolution and noise amplification, which are tunable. There are also parameters that allow for regularizing of possible

ill-conditioned, large matrices. The choices of these parameters can be somewhat subjective and non-rigorous.

Standard derivations of the SOLA method are common in the literature (e.g. Švanda et al., 2011; Jackiewicz et al., 2012; Korda and Švanda, 2019). Here, a slightly modified version is presented that will connect to the probabilistic equations in Section 2.2. We follow closely the notation of Tarantola (2005). Where appropriate, the relationship to standard inversion terminology is given in parenthesis with italicized text.

Assume that any model can be described by $\mathbf{m}(\mathbf{r})$, where \mathbf{r} denotes space. By model, we mean the quantity that inversions are seeking, such as the flow structure of a supergranule or the sound-speed profile under sunspots. Consider a generalized discrete data set \mathbf{d} that is linearly related to the model through an integral equation

$$\mathbf{d} = g(\mathbf{m}), \quad (1)$$

where g is some functional that describes the physics of the problem. If such an equation exists, it will be called a generative model. For now, this relationship will be given as

$$\mathbf{d} = \mathbf{G}\mathbf{m}, \quad (2)$$

where \mathbf{G} is a matrix made up of vector functions (*sensitivity kernels*). The true, but as of yet unknown, model is related to some set of observed data through

$$\mathbf{d}_{\text{obs}} = \mathbf{G}\mathbf{m}_{\text{true}}, \quad (3)$$

which we consider error free for simplicity. We want to obtain a good estimate \mathbf{m}_{est} of \mathbf{m}_{true} at some location, and we therefore assume that the estimator model is linearly related to the observed data as

$$\mathbf{m}_{\text{est}} = \mathbf{w}^T \mathbf{d}_{\text{obs}}, \quad (4)$$

where the \mathbf{w} are constants (*weights*). Defining some resolution operator (*averaging kernels*) as

$$\mathbf{R} = \mathbf{w}^T \mathbf{G} \quad (5)$$

gives

$$\mathbf{m}_{\text{est}} = \mathbf{R}\mathbf{m}_{\text{true}}. \quad (6)$$

This equation implies that the estimation that will be found is a smoothed or weighted version of the true model, since with finite data \mathbf{R} will never be a delta function.

The constants $[\mathbf{w}]$ are computed by minimizing a cost function

$$\min |\mathbf{R} - \mathbf{I}|^2, \quad (7)$$

where \mathbf{I} represents a delta function, but in practice it is something more reasonable (*Gaussian target function*). Minimization with respect to the weights gives

$$\mathbf{w} = (\mathbf{G}\mathbf{G}^T)^{-1} \mathbf{G}. \quad (8)$$

This expression shows that a (usually) large matrix inversion is necessary to compute (*kernel convolution matrix*). In a standard local-helioseismic inversion, the convolution matrix can be of order $10^5 \times 10^5$ elements, although various Fourier methods can help reduce this size

(Jackiewicz et al., 2012). Additionally, in practice this matrix may contain other quantities such as the noise covariance and any regularization terms required for a smooth solution.

Finally, once the weights are obtained, the estimate is given by

$$\mathbf{m}_{\text{est}} = \mathbf{G}^T (\mathbf{G}\mathbf{G}^T)^{-1} \mathbf{d}_{\text{obs}}, \tag{9}$$

and

$$\mathbf{R} = \mathbf{G}^T (\mathbf{G}\mathbf{G}^T)^{-1} \mathbf{G}. \tag{10}$$

Notice that the observations are only involved in the last step: the calculation of \mathbf{w} is not conditioned on the data at all.

It is interesting to point out that the model estimate is likely not the true model (again using a finite amount of data). So it is reasonable to postulate that the true model may have a form in which it is related to the estimated model, plus some arbitrary, properly scaled model \mathbf{m}_0 of similar smoothness

$$\mathbf{m} = \mathbf{m}_{\text{est}} + (\mathbf{I} - \mathbf{R}) \mathbf{m}_0. \tag{11}$$

This expression serves as a general solution to the inverse problem (Backus and Gilbert, 1968). Since $\mathbf{G}(\mathbf{I} - \mathbf{R})\mathbf{m}_0 = 0$, the same operation on Equation 11 gives

$$\mathbf{G}\mathbf{m} = \mathbf{G}\mathbf{m}_{\text{est}} = \mathbf{G}\mathbf{R}\mathbf{m}_{\text{true}} = \mathbf{G}\mathbf{m}_{\text{true}} = \mathbf{d}_{\text{obs}}, \tag{12}$$

which was shown in Equation 3.

2.2. Background to Bayesian Inferences

Full discussions of MCMC in the general context of Bayesian theory and astronomy applications can be found in many places (e.g. Sharma, 2017; Hilbe, de Souza, and Ishida, 2017). A particularly useful pedagogical treatment is given by Hogg and Foreman-Mackey (2018). Here a simple overview is provided to guide the later discussion and examples.

Imagine we have N measurements of some observable comprising a data set $\mathbf{d} = \{d_i \mid i = 1, \dots, N\}$, and each measurement has an uncertainty σ_i , which is considered independent and normally distributed for simplicity. Now assume that we possess a generative model that can, in principle, make predictions of the data through the operation $g(\mathbf{m})$, as in Equation 1. \mathbf{m} is a model made up of M parameters $\mathbf{m} = \{m_i \mid i = 1, \dots, M\}$. If many repeated measurements are made, then the expected frequency distribution (*probability*) of datum d_j is

$$p(d_j \mid \mathbf{m}, \sigma_j) = \frac{1}{\sqrt{2\pi\sigma_j^2}} \exp \left[-\frac{(d_j - g_j(\mathbf{m}))^2}{2\sigma_j^2} \right]. \tag{13}$$

The vertical bar \mid is read as “given,” so this expression is the probability of the datum d_j given the model and the uncertainty on d_j . Clearly, if the operation

$$g_j(\mathbf{m}) \equiv \sum_{i=1}^M g_j(m_i) \tag{14}$$

gives a number far from d_j , the resulting probability will be small. One wishes to maximize the probability, not just of one data point, but of the entire set of observations. This is usually

referred to as the likelihood function $[L]$, which is a product of individual probabilities

$$L(\mathbf{d}|\mathbf{m}, \boldsymbol{\sigma}) = \prod_{j=1}^N p(d_j|\mathbf{m}, \sigma_j). \quad (15)$$

In practice, one may stop here and find the parameters that maximize the likelihood function, or, more conveniently, minimize the negative logarithm of it. The problem then reduces to least-squares fitting. The resulting model, identified from the probability of the *data given the parameters*, is interpreted, however, as the likelihood of the *parameters given the data*. This interpretation presents a formal inconsistency.

Bayes' theorem can be easily derived from sum and product rules of probability theory. The result has four quantities: One quantity is the likelihood function in Equation 15. Another is any prior information $[I]$ that we possess on the model and the uncertainties, which will be denoted $\rho(\mathbf{m}, \sigma|I)$. The third is the evidence $[p(\mathbf{d}|I)]$, which is effectively a normalization term and will not be important for our discussion. The final ingredient is the posterior probability distribution function, which is computed as

$$\text{PDF}(\mathbf{m}|\mathbf{d}, \sigma, I) = \frac{L(\mathbf{d}|\mathbf{m}, \boldsymbol{\sigma})\rho(\mathbf{m}, \sigma|I)}{p(\mathbf{d}|I)}, \quad (16)$$

and defines Bayes' theorem. This important quantity is the statistical probability of the model *given* the data, uncertainties, and any prior knowledge.

The model PDF is therefore related to the likelihood function, and it will closely resemble it if the priors are not very specific or informative. In this case the interpretation stated above is not so fatal. However, some of the power of the Bayes framework is that if the prior knowledge of model parameters is non-trivial, then the PDF is too, and its complexity requires more sophisticated inference methods to be applied. Priors also restrict the parameter space to a smaller region than a likelihood function alone can.

In general, and in the examples below, the likelihood function is a multivariate normal distribution

$$L(\mathbf{d}|\mathbf{m}, \boldsymbol{\Sigma}) = \frac{1}{\sqrt{(2\pi)^k |\boldsymbol{\Sigma}|}} \exp \left[-\frac{1}{2} (\mathbf{d} - \mathbf{g}(\mathbf{m}))^T \boldsymbol{\Sigma}^{-1} (\mathbf{d} - \mathbf{g}(\mathbf{m})) \right], \quad (17)$$

where $\boldsymbol{\Sigma}$ is the data covariance matrix, $|\boldsymbol{\Sigma}|$ is its determinant, and k is the dimension of the problem (length of \mathbf{d}).

To summarize, probabilistic inversions use Equation 16 to compute the posterior PDF – the joint probability distribution of parameters that is consistent with the data. The PDF will rarely have an analytic form, and it is not necessarily well-behaved or uni-modal. The goal is not to optimize (maximize) the PDF, or find its peak, but to know the whole distribution and sufficiently sample it. One can imagine one strategy, which would be looping over a (uniform) grid of parameter values and computing the resulting PDF. However, for high-dimensional problems, this would be extremely expensive and inefficient. Too many low-probability realizations would be calculated. Fortunately, there are alternative approaches. There is a robust literature of different probability-distribution sampling methods, but most modern ones rely on MCMC techniques. The basic difference among these methods is how the sampler “moves” through the parameter space; i.e. how an algorithm decides to choose trial parameter values, so that it hopefully spends more time in high-probability space. MCMC uses random numbers to drive the process. Metropolis–Hastings is one of the simplest and well-known algorithms (Press et al., 2007).

A recently developed MCMC algorithm, affine-invariant sampling (Goodman and Weare, 2010), is what we adopt in this work. This method is in a class of ensemble MCMC, since multiple chains, called “walkers,” are all in execution simultaneously as they explore the parameter space. The walkers can therefore be run in parallel, but they are allowed to interact in certain ways to adapt the proposal densities and maintain their Markov properties. It is a promising tool for sampling PDFs that are not extremely complex (Foreman-Mackey et al., 2013).

The connection of the SOLA inversion described in Section 2.1 to the probabilistic language where priors and covariances are considered is useful, and it can be made quite easily. Consider a linear least-squares problem. Including model priors, one could construct a cost function (or least-squares function, or χ^2 -function) called S as

$$2S(\mathbf{m}) = (\mathbf{d}_{\text{obs}} - \mathbf{G}\mathbf{m})^T \mathbf{C}_D^{-1} (\mathbf{d}_{\text{obs}} - \mathbf{G}\mathbf{m}) + \tag{18}$$

$$+ (\mathbf{m} - \mathbf{m}_{\text{prior}})^T \mathbf{C}_M^{-1} (\mathbf{m} - \mathbf{m}_{\text{prior}}). \tag{19}$$

\mathbf{C}_D and \mathbf{C}_M are the covariance matrices of the data and model priors (if known), respectively. As outlined above, the Gaussian posterior PDF can be computed from the cost function S that has a form $\sim \exp(-S(\mathbf{m}))$. The center of the distribution, i.e. the most likely model of the Gaussian PDF (the model that minimizes the cost function) $\tilde{\mathbf{m}}$ and its covariance $\tilde{\mathbf{C}}_M$ can be computed by differentiation and shown to be

$$\tilde{\mathbf{m}} = \mathbf{m}_{\text{prior}} + \mathbf{C}_M \mathbf{G}^T (\mathbf{G} \mathbf{C}_M \mathbf{G}^T + \mathbf{C}_D)^{-1} (\mathbf{d}_{\text{obs}} - \mathbf{G} \mathbf{m}_{\text{prior}}), \tag{20}$$

$$\tilde{\mathbf{C}}_M = \mathbf{C}_M - \mathbf{C}_M \mathbf{G}^T (\mathbf{G} \mathbf{C}_M \mathbf{G}^T + \mathbf{C}_D)^{-1} \mathbf{G} \mathbf{C}_M. \tag{21}$$

In the SOLA method, there are no priors in the model space. In the probabilistic language, this implies white noise with no correlations and a (possibly) infinite variance of the priors:

$$\mathbf{C}_M \approx k \mathbf{I} \quad (k \rightarrow \infty). \tag{22}$$

Using the definitions in Section 2.1, the posterior centers then reduce to

$$\begin{aligned} \tilde{\mathbf{m}} &= \mathbf{G}' (\mathbf{G} \mathbf{G}')^{-1} \mathbf{d}_{\text{obs}} + (\mathbf{I} - \mathbf{R}) \mathbf{m}_{\text{prior}}, \\ &= \mathbf{R} \mathbf{m}_{\text{true}} + (\mathbf{I} - \mathbf{R}) \mathbf{m}_{\text{prior}}, \end{aligned} \tag{23}$$

$$\tilde{\mathbf{C}}_M = (\mathbf{I} - \mathbf{R}) \mathbf{C}_M. \tag{24}$$

Equation 23 is precisely Equation 11, a general solution of a SOLA inversion with the prior information replacing the arbitrary model \mathbf{m}_0 . In the SOLA language, if $\mathbf{R} \approx \mathbf{I}$ (a δ -function like *averaging kernel*), then $\tilde{\mathbf{m}} = \mathbf{m}_{\text{est}} \approx \mathbf{m}_{\text{true}}$. In the probabilistic language, this implies there are no uncertainties in the posterior solution, so $\tilde{\mathbf{C}}_M \approx \mathbf{0}$.

These two conclusions are identical, showing that, in principle, the methods can arrive at similar results, yet only in ideal circumstances. What is hopefully demonstrated throughout the rest of this article is that the probabilistic method, in practice, is robust, practical, and gives more realistic uncertainties.

3. Examples for Time–Distance Local Helioseismology

The forward problem in time–distance helioseismology is symbolically formulated as (e.g. Kosovichev and Duvall, 1997; Gizon and Birch, 2002, 2004)

$$\delta\tau = \int_{\odot} K \delta q \, dr, \quad (25)$$

where the travel-time shifts $[\delta\tau]$ between two surface locations are caused by some (small) interior perturbation δq . The sensitivity kernels $[K]$ mediate this relationship, which is considered to be linear. Any inversion consists of using the observed surface $\delta\tau$ and computed K to find the unknown δq . In SOLA methods, δq is inferred at each spatial location, or at least one depth at a time. In probabilistic inversions, δq must first be parametrized by some number of free parameters. The parameters are estimated using Bayes' theorem and MCMC, and then δq can be studied over the whole domain.

We elucidate two rather simple examples of inversions based on common research areas in local helioseismology. We compare the probabilistic inversions with the SOLA method and contrast the computational particulars. We will only consider examples of flows, and therefore travel-time differences are the important observables.

It is important to keep in mind that in what follows we are not solving any real problem. In one case, we are only inverting synthetic observations that are computed in the forward sense from Equation 25. This does not tell us anything about the accuracy of the sensitivity kernels. They could be completely wrong. It only tells us about the inverse process, which is the goal here. In the other case, inversions of a realistic numerical model are shown. Most helioseismic studies employ two ways of modeling the interaction of seismic waves with inhomogeneities: ray theory or Born theory. Our examples span these two cases.

3.1. Meridional Circulation in a Ray-Theory Approach

3.1.1. The Toy Problem

We use a simple, single-cell, meridional-flow model first described by van Ballegoijen and Choudhuri (1988) and later utilized by Dikpati and Charbonneau (1999), among others. The parametric model is given by Equations 57–61 of van Ballegoijen and Choudhuri (1988) and will not be reproduced here. It is computed in a polar (r, θ) meridional plane. For our purposes, the meridional profile has effectively three free parameters, which will be denoted p_1 , p_2 , and p_3 . p_1 controls the flow amplitude, while p_2 and p_3 control the latitudinal and radial (depth) dependence of the flow structure, respectively. The model provides two-dimensional flows in the radial and latitudinal directions $\mathbf{v}(r, \theta) = v_\theta(r, \theta)\hat{\theta} + v_r(r, \theta)\hat{r}$ that satisfy mass conservation in the 2D domain: $\nabla \cdot \rho \mathbf{v} = 0$. The density is a function that scales as $\rho \sim r^{-1.52}$, similar to van Ballegoijen and Choudhuri (1988), but slightly modified to match Model S (Christensen-Dalsgaard et al., 1996) in the region of interest. The input values of the three parameters are such that the poleward surface flow reverses direction at $r \approx 0.79 R_\odot$. We use a grid that has 150 points in latitude and 100 points in radius, covering $\theta = \pm 90^\circ$ and from $r = 0.68 R_\odot$ to $r = R_\odot$.

Ray kernels are computed for a set of latitudes and distances that sample the model relatively well (although by no means exhaustively). We consider ten skip distances from 2° to 45° . The central latitude range is $\pm 77^\circ$, resulting in a total of 122 ray kernels. Figure 1 shows the given circulation model with all ray paths overplotted. The weaker radial flows of the model are not shown here.

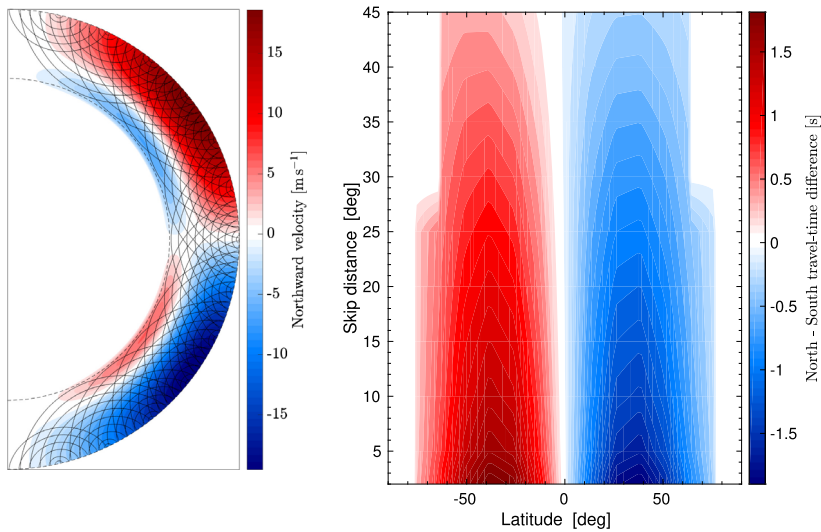


Figure 1 Left: Input model latitudinal-flow profile and ray paths. The color scale shows the northward velocity of a model computed with parameters $\mathbf{p} = \{5000, 1.0, 0.5\}$. The solid curves are 122 ray paths used in the analysis to compute flow kernels. The dashed half circle represents the radius $r = 0.7 R_{\odot}$. Right: Forward (noiseless) travel times computed from ray kernels for each distance and latitude.

Synthetic forward travel-time differences are then computed from the flow model and kernels, shown on the right of Figure 1. To these travel times, artificial, normally distributed random noise is added at two different levels: $\mathcal{N}(0, \sigma_1^2)$ and $\mathcal{N}(0, \sigma_2^2)$. In the low-noise case, $\sigma_1 = 0.016$ seconds is about 2% of the rms of the travel-time differences (0.8 seconds), and about 20% in the high-noise case ($\sigma_2 = 0.16$ seconds). These noise levels roughly correspond to typical meridional-flow measurements made over three years and one month, respectively (Braun and Birch, 2009).

3.1.2. SOLA Solution to the Problem

We first demonstrate the standard inversion method described in Section 2.1. It is the SOLA inversion applied by Jackiewicz, Serebryanskiy, and Kholikov (2015) and other recent studies. The synthetic travel-time differences are considered to be uncorrelated, and thus the noise covariance matrix used in the inversion is diagonal. No mass-conserving constraint is imposed, and therefore it is hopeless to try to recover the small radial velocity in this inversion, which is about 10% of the amplitude of the latitudinal flows.

The SOLA inversion estimates the velocities at specific *target* locations. In this example, there are 110 target locations, 10 in depth and 11 in latitude. At each location, a 2D Gaussian target function was computed with a full-width-half-maximum (FWHM) in the radial direction of $0.08 R_{\odot}$ and in the latitudinal direction of 10° . The target function replaces the unrealistic δ -function given in Equation 7, and it gives a measure of the spatial resolution of the inversion results.

The results of the SOLA inversion are shown in Figure 2 after inverting the low-noise and the high-noise travel times. To aid in comparison with the known model, the retrieved flows at the 110 spatial locations have been interpolated onto the model grid. The recovered flows generally follow the pattern of the model. The deeper return flow is not reliably found




Figure 2 SOLA inversion results for travel times of different noise levels. The left two columns show the flows from the inversion and the difference with the known model. The middle column is the inferred noise at each inversion location. The fourth column is the misfit value at each inversion location. The last column is an example averaging kernel from an inversion at a target location at $(r, \theta) = (0.9R_{\odot}, -15^{\circ})$. Top row: SOLA inversion for the low-noise case. The overall median noise is about 0.4 m s^{-1} . Bottom row: SOLA inversion for the high-noise case. The overall median noise is about 0.5 m s^{-1} .

in either case. Since the SOLA inversion always returns a flow pattern that is a smoothed version of the real one (see Jackiewicz et al., 2012; Švanda, 2012), the amplitude is underestimated. On average, the underestimation is about 2 m s^{-1} in the low-noise case and about 4.5 m s^{-1} in the high-noise case, but in some locations up to 10 m s^{-1} .

The inferred noise is too small and not consistent with the errors. Specifically, the retrieved velocity is $\approx 10\sigma$ away from the input in the high-noise case. In other words, if the true answer were not known and we surmised that our result is within 1 or 2σ away from the truth, we would make an error of one order of magnitude. Another feature of the inversions reveals the expected less-localized averaging kernel for the case of the noisier travel times. Note that one can tune the trade-off parameters to obtain different results (smoother/less noisy, more localized/noisier, etc.), making the interpretation of the validity of the inferences challenging.

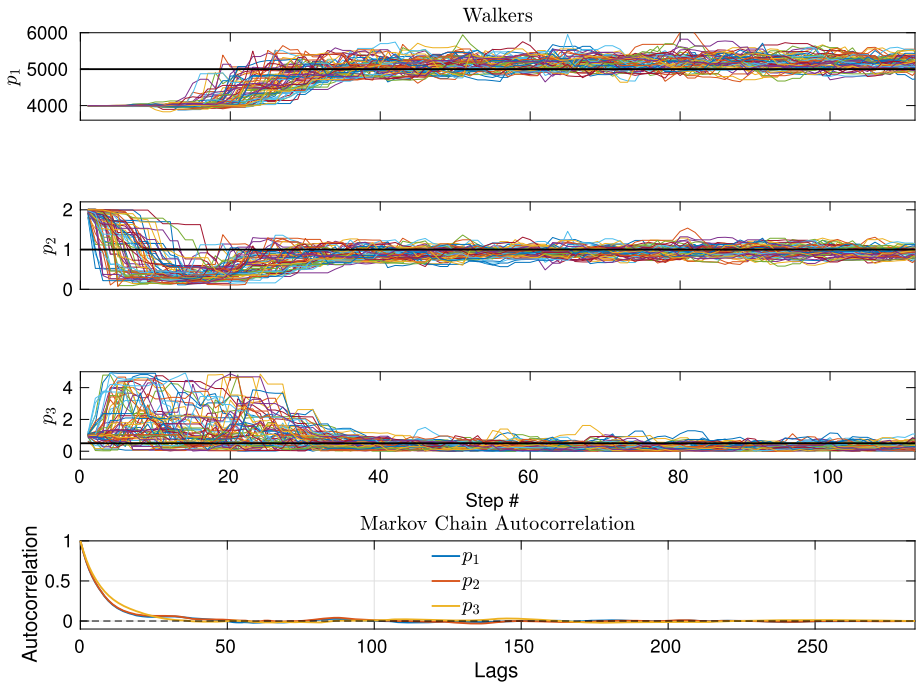


Figure 3 Probabilistic inversion diagnostics. The top three panels show how each of the 60 walkers of each parameter traverses parameter space. Each walker is a different color. The black horizontal lines are the input (known) values. The y-intercepts are the starting values of the walkers. Only the first 1/3 of the steps in the run are shown. The bottom panel shows the autocorrelation of each parameter’s walkers as a function of the step lags. The burn-in phase was discarded before the calculation. The effective sample size is 1168.

3.1.3. Probabilistic Solution to the Problem: Parameter Posteriors

Before showing the results of the Bayesian MCMC inversion in a standard way, it is important to explore the output at the level of the walkers and the multi-dimensional PDF of the parameters. In this example, the total number of steps (iterations) was chosen to be 10^5 . Each of the three free parameters was assigned 60 walkers (chains). The sampling of each walker was every five steps, which is a “thinning” procedure, whereby only the fifth step is stored. The PDF was therefore sampled $10^5 / (60 \times 5) = 333$ times per walker. The choice of the standard deviation of Gaussian likelihood function is $\sigma = 0.5$ second. Since the measurements are assumed to be uncorrelated, $(\Sigma)_{ij} = \sigma_i^2 \delta_{ij}$, the likelihood function in Equation 17 reduces to

$$L(\mathbf{d}|\mathbf{m}, \sigma_i) = \sum_i \frac{1}{\sigma_i \sqrt{2\pi}} \exp \left[-\frac{1}{2} \left(\frac{d_i - g(m_i)}{\sigma_i} \right)^2 \right]. \tag{26}$$

The priors are taken as flat and rather wide, for demonstration purposes, as if we did not have a good idea of their values. These are known as “uninformative” priors.

Figure 3 shows the time series of all of the walkers during the run. Upon inspection, the first thing to point out is that the initial $\approx 30 - 40$ steps are when the sampling “burns-in.” This essentially means that the chains take a few steps to wander towards and reach a

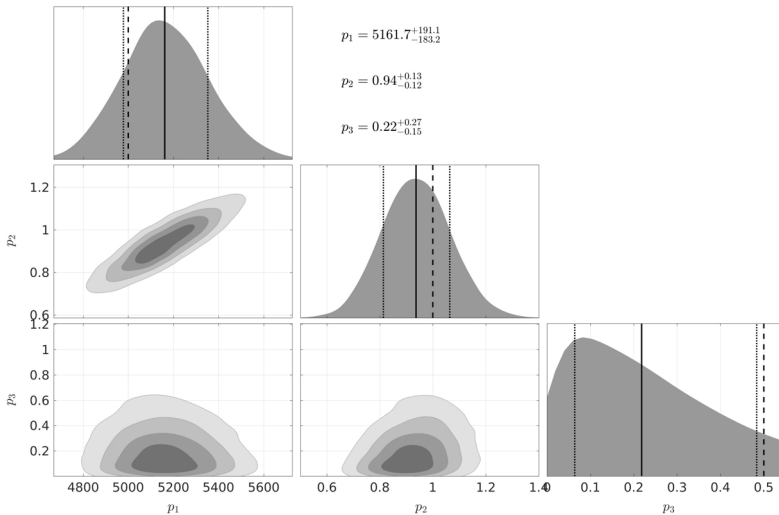


Figure 4 Corner plot showing the marginalized PDFs of the three parameters. The marginalized distribution for each parameter independently is shown in the histograms along the diagonal, and the marginalized 2D distributions as contour plots in the other panels. For each 1D histogram, the median of the PDF is the solid black line, and the dotted lines enclose the 68% confidence interval. The numerical values are given at the top. The dashed-black lines are the known input-parameter values. The contour levels of the 2D joint probability densities are at 20%, 40%, 60%, and 80% confidence intervals.

high-probability region, since the starting values typically might be far from such regions, as in this example (by choice). There is endless debate about burn-in validity in the literature (e.g. Hogg and Foreman-Mackey, 2018, Section 7) into which we will not delve. In any case, the walker behavior is acceptable, in that once burnt-in, the space of the PDF is fully explored. The acceptance rate of the proposed steps is about 30% – a good value for MCMC algorithms.

Sample draws can be correlated in MCMC algorithms due to noise or other factors. If each draw were completely independent, then the variance would decrease as more and more samples are drawn. It is critical to know if independent samples are drawn from the PDF so that the parameter estimation is not biased, and reliable estimates of the mean/median and variance can be computed. The standard way of determining this is by calculating the autocorrelation of the walkers of each parameter. When and if the autocorrelation approaches zero, one can be confident that the walkers “lose their memory” of where they started and reach some state of equilibrium. The autocorrelation can be computed empirically from the time series of walkers, and it is shown in Figure 3 in the right panel. In this case, the rate of convergence is quite rapid compared to the length of the run, and even fewer total steps could have been chosen. The effective sample size (ESS) is another concept to understand how many independent samples were drawn in the walker time series, and it is found from the autocorrelation (Sokal, 1997). In this example, the ESS is 1168. For a standard deviation σ of the PDF of a given parameter, the Monte Carlo standard error goes as $\sigma/\sqrt{\text{ESS}}$. This means that we are able to measure the median of a parameter with about a 3% error compared to the overall uncertainty σ .

The PDF in this example is three-dimensional. The “corner” plot matrix in Figure 4 shows histograms of all of the one- and two-dimensional projections of the PDF of the parameters. The marginalized 1D PDFs are along the diagonal, and correlations between




Figure 5 Results for both components of the meridional circulation using the probabilistic inversion. The inferred flows use the median of the parameter PDFs. The panels show the difference of the two flow components inferred from the inversion with the model, for both levels of noise. The color scale extends to the limits of the data in each panel. The factor necessary to multiply the lower-amplitude radial velocities to achieve this is shown at the top.

parameters are given in the off-diagonal elements (marginalized 2D PDFs). In this case, the PDFs are not multi-modal, which can be an indication that the model is parametrized well. This should not be surprising since the input model follows the same parametrization as the forward model.

The power of the projected model parameter PDFs in Figure 4 is that one immediately sees the distribution widths, as well as any correlations between model parameters. In this example, the PDFs bracket the known input value of the parameters within the 16 and 84 percentiles, except parameter p_3 , which is just beyond that range. This parameter also has the least Gaussian PDF.

3.1.4. Probabilistic Solution to the Problem: Model Space

The median of the PDFs of the inversion for the three parameters are used to generate a flow-circulation profile. This provides a way to visualize the results in the space of the model, similar to what was shown earlier for SOLA. The resulting profiles are very comparable to the input model, so much so that in Figure 5 only the differences with the model are shown. The differences are significantly less than the SOLA example described in Section 3.1.2. In the low-noise example, the inversion very slightly overestimates the poleward flow amplitude (parameter p_1), and therefore the equatorward flow is weakly underestimated. This affects the radial-velocity differences in the manner shown. In the case of noisier travel times, p_1 is again overestimated, but the other two parameters are slightly underestimated, leading to some small-scale differences in the relative flows.

3.1.5. Probabilistic Solution to the Problem: Data Space

The MCMC method allows for a visualization of the results in data space too, much more naturally and quickly than the SOLA method. Figure 6 presents the data-space solutions for

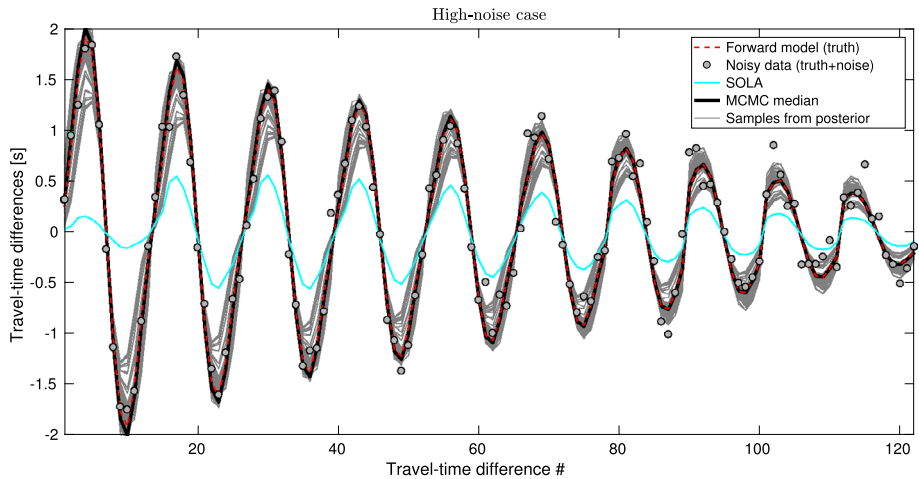


Figure 6 Inversion solutions in the data (travel-time) space for the high-noise case. The dashed-red line shows the input travel times, while the filled-gray circles represent the random noise addition. The forward-modeled travel times from the SOLA inversion result are in cyan. The thick-black line is computed from the median of the parameter PDF. Finally, the thin-gray lines are forward travel times computed from 100 realizations of the posterior distribution. The travel times are plotted such that each “oscillation” is a different skip distance (smaller to larger from left to right), and the points within each oscillation correspond to each latitude (see Böning et al., 2017, for similar plots).

both inversion methods using noisier travel times. One immediately sees the manifestation of the underestimated velocity in the SOLA inversion in the smaller-amplitude travel times. The near-surface region (left of the figure at smaller skip distances) is particularly evident. On the other hand, the travel times generated from the median of the PDF are highly consistent with the input ones. Also shown are 100 random realizations of the PDF, which quickly gives a picture of the statistical uncertainties in the data space.

3.2. Supergranulation in the Born Theory

The Sun’s supergranulation is an important component of near-surface convection-zone dynamics and a strong source of advection of magnetic fields. While its sub-surface flow structure has not yet been faithfully determined by helioseismology, it may have a simple enough form to be parametrized by a model. It is therefore another suitable test case for a probabilistic inversion. For this example, we consider five new aspects that add complexity and richness to the demonstration:

- i) The problem is set up in three dimensions (rather than two).
- ii) Born sensitivity kernels are used (instead of ray kernels).
- iii) The observations are from a 3D numerical simulation with stochastic, realistic noise properties (not synthetic forward-modeled observations).
- iv) A proper noise covariance matrix is computed and used in the inversions (not just diagonal variances).
- v) The supergranule model in the simulation is *different* from the model and parameters used to estimate the PDF.

Regarding the last point, this means that the “true” values of the parameters used to simulate the supergranule are essentially unknown, unlike the meridional-flow example where the

input p_i could be directly compared to the posteriors. We briefly describe the problem setup before studying the results.

3.2.1. The Models

The supergranulation model is taken from Dombroski et al. (2013). In that work, realistic wave propagation using the SPARC code (Hanasoge et al., 2006) was simulated through a single, kinematic supergranule flow pattern to quantify the effects on seismic waves. The mass-conserving flow structure was modeled using seven parameters. Two control the horizontal extent of the divergent flow, three control the depth dependence and strength of the outflow, and two more parameters control the depth dependence of the boundary inflow.

The model supergranule has a radial extent of about 30 Mm at the surface, where the maximum horizontal speed and the vertical speed are 250 m s^{-1} and 20 m s^{-1} , respectively. The outflow switches to an inflow at a depth of $\approx 10 \text{ Mm}$. The maximum vertical speed is about 100 m s^{-1} , peaked around 4 Mm below the photosphere. For purposes later, this will be referred to as the reference (“ref”) model.

As a notable aside, while this model is reasonable and at least consistent with surface observations (Duvall and Birch, 2010; Rieutord and Rincon, 2010), supergranulation has proven very difficult to fully understand. There are even questions about whether models that have separable flows (in horizontal and vertical directions) are appropriate for supergranulation (Ferret, 2019; Dhruv, Bhattacharya, and Hanasoge, 2019). Addressing such issues is beyond the scope of this article.

The model that we use in the probabilistic inversion is instead from Duvall and Hanasoge (2012). Also employing a separable, mass-conserving flow, this model has five free parameters. In fact, two of the parameters are equivalent between the models, those that control the horizontal diverging flow as

$$\mathbf{g}(r) = J_1(kr) \exp(-r/R) \hat{\mathbf{r}}, \tag{27}$$

where J_1 is an order-one Bessel function, k is a wavenumber, and R represents a decay length in the distance coordinate from the origin $[r]$. The values from Dombroski et al. (2013) are $k = 2\pi/30 \text{ rad Mm}^{-1}$ and $R = 15 \text{ Mm}$, identical to those in Duvall and Hanasoge (2012). In addition, this model has a Gaussian depth dependence of the velocities, determined by three additional parameters: a peak amplitude $[v_0]$, a peak flow location $[z_0]$, and a Gaussian width $[\sigma_z]$, leading to the function

$$u(z) = \frac{v_0}{k} \exp\left(-\frac{(z - z_0)^2}{2\sigma_z^2}\right). \tag{28}$$

Once \mathbf{g} is computed, the model vertical flows are constructed first as $v_z(r, z) = u(z)\nabla_h \cdot \mathbf{g}$. Then the horizontal flows are $\mathbf{v}_h(r, z) = -f(z)\mathbf{g}(r)$, where f is obtained from applying the continuity equation. We compute this model in three spatial Cartesian dimensions (x, y, z) for illustration sake, even though it is axisymmetric and the problem can be solved in only two. This will be referred to as the “trial” model.

Apart from the two common free parameters, each model is derived differently enough that the priors on the other three free parameters are not well known. We take uniform priors that are kept identical in each of the probabilistic inversions discussed below. The priors are given in Table 1.

Table 1 Table of (uniform) priors for the probabilistic inversions using the “trial” model. They are ordered according to how the results are presented.

| Parameter | Label | Min. | Max. | Unit |
|------------|-------|------|---------|----------------------|
| v_0 | p_1 | 50 | 400 | m s^{-1} |
| σ_z | p_2 | 0 | 12 | Mm |
| z_0 | p_3 | 0 | -12^a | Mm |
| k | p_4 | 20 | 40 | rad Mm^{-1} |
| R | p_5 | 10 | 25 | Mm |

^aNegative values denote sub-surface

3.2.2. Setup of the Problem

Helioseismic measurements were computed from the numerical simulation using a time series of the vertical velocity, which is sampled every one minute at 200 km above the model photosphere over a total of 24 hours. The horizontal spatial domain extends 100 Mm and is sampled every 1/3 Mm. The vertical velocity is first filtered to isolate different ridges (radial orders n), including the f -mode (n_0) and the first two acoustic-mode ridges (n_1, n_2) using standard methods (Braun and Birch, 2008b; Gizon et al., 2009; DeGrave, Jackiewicz, and Rempel, 2014). Cross correlations were measured in center-to-annulus and center-to-quadrant geometries for 15 different travel distances, ranging from 6 Mm to 20 Mm. For each ridge and each distance, three travel-time difference maps are computed (at the same spatial resolution) across 50 Mm of the simulation domain: “out–in” [$\delta\tau_{oi}$], “west–east” [$\delta\tau_{we}$], and “north–south” [$\delta\tau_{ns}$]. Such geometries are sensitive to flows (Duvall et al., 1997). This results in 135 unique travel-time maps, which are very comparable to the ones computed using helioseismic holography by Dombroski et al. (2013). Only a fraction of these measurements are used in the sample inversions.

Dombroski et al. (2013) computed a second simulation without a background supergranule. We use these data (split into 12 two-hour cubes) to estimate the noise covariance in the travel times according to the noise model of Gizon and Birch (2004). The exact same measurement procedure explained above is carried out on these cubes to determine the covariances $\text{Cov}[\delta\tau_i, \delta\tau_j]$.

The linear forward problem (Gizon and Birch, 2002) in this example can be written

$$\delta\tau_i^\alpha(x, y) = \iiint \mathbf{K}_i(x - x', y - y', z) \cdot \mathbf{v}^\alpha(x', y', z) \, dx' \, dy' \, dz, \tag{29}$$

where flows are \mathbf{v} , the sensitivity kernels \mathbf{K}_i are vector-valued, and each index i corresponds to a given ridge, geometry, and travel distance. Born-approximation kernels are computed from Birch and Gizon (2007) in a point-to-point fashion, and then averaged over annuli to be consistent with the travel-time geometries (oi, we, ns). The α -superscripts refer to the particular measurement source or supergranule model under consideration.

We can be precise about this example. The SOLA inversion is looking to infer \mathbf{v}^{ref} in Equation 29 given the kernels and the measurements $\delta\tau^{\text{ref}}$ of $v_z^{\text{ref}}(x, y, z = 200 \text{ km})$, and so $\alpha = \text{ref}$. Let us call the estimate \mathbf{v}^{SOLA} . The probabilistic inversion is using Equation 29 with $\mathbf{v}^{\text{trial}}$ and the same kernels to forward compute $\delta\tau^{\text{trial}}$, thus $\alpha = \text{trial}$ in that case. The $\delta\tau^{\text{trial}}$ are used in the computation of the likelihood function along with $\delta\tau^{\text{ref}}$. In model space, the probabilistic inversion is seeking to estimate suitable values of parameters such that $\mathbf{v}^{\text{trial}}$ will resemble \mathbf{v}^{ref} . Further details in the setup of the inversions are mentioned in Appendix A.




Figure 7 Comparison of supergranulation inversion results at 3 Mm beneath the model photosphere. The left (right) panels are for the v_x (v_z) inversion. The top rows, from left to right, are the flow fields for the simulation, the SOLA inversion, and the probabilistic inversion computed from the median of the PDF. Darker shading corresponds to positive velocities (to the right for v_x , and out of the page for v_z), and the scale is the same in each set. The bottom panels are cuts through the supergranule at $y = 0$, shown by the dashed line in the top panels. Twenty random samples of the PDF are drawn and computed in the model space. A few representative points from the SOLA inversion are given with uncertainties. The SOLA velocities are scaled by the factor indicated.

3.2.3. Results

SOLA inversions in 3D are not very convenient to compare with the probabilistic inversions, neither in model space nor data space. The reason is that, typically, the flows in (x, y) are inferred one depth at a time, and these depths are usually few. Furthermore, the prescribed “resolution” in both directions can vary from depth to depth. In this example, the SOLA inversions were carried out at three target depths $z_0 = (0, -3, -4.5)$ Mm. Each target depth had a different target width in the horizontal and vertical directions. The actual depth at which the inversion is most sensitive can also be distant from z_0 due to non-localized averaging kernels. By contrast, and by construction, the probabilistic inversions provide parameters that allow one to estimate flows over the full (or any) spatial domain.

To make meaningful comparisons between inversion results and the reference model, several steps need to be taken. Since the SOLA results are rather coarse in depth and smooth in the horizontal direction, we decide to adapt everything else to them. Firstly, for a given target depth, the reference-model velocities are convolved with the target function of the inversion [the \mathbf{I} in Equation 7, which in this case is not a δ -function but a 3D Gaussian sphere]. This process is represented by Equation 6, whereby the estimated flows are a smoothed version of the true ones: $\mathbf{v}^{\text{SOLA}} = \mathbf{R}\mathbf{v}^{\text{ref}}$. The result is then integrated over depth, giving a 2D flow map that can be compared to the SOLA inferences.

For the probabilistic inversion, we use draws of the model PDF parameters (median or otherwise) and compute the flow model $\mathbf{v}^{\text{trial}}$ on the same spatial grid as \mathbf{v}^{ref} and the \mathbf{v}^{SOLA} . It is also appropriately smoothed by the SOLA inversion target function and integrated over depth in the same manner. This process results in three sets of flow maps at three nominal target depths for three flow components, although we restrict comparisons to v_x and v_z . Only results in model (velocity) space will be presented.

Figure 7 shows a comparison of these two flow components at 3 Mm beneath the photosphere, where the horizontally divergent structure is apparent. In general, we find the SOLA inversions severely underestimate horizontal velocities (note the scaling factor), while the

probabilistic inversions weakly overestimate them. The bottom panel of Figure 7 shows a cut through the models at $y = 0$. The noise in the SOLA inversion, even using covariance matrices, is highly underestimated. The horizontal error bars show the FWHM of the target function, which were quite wide to get sensible results. On the other hand, a random sample of the parameter PDF from the probabilistic inversion gives a reasonable spread of solutions in model space.

At the same depth, the inferences on the weaker vertical velocity are also shown in Figure 7 on the right. In this case, the SOLA flow inferences are marginal at best. At inversions just below this depth, the SOLA flows are anticorrelated with the reference flows, as in Figure 11 in Appendix B. Dombroski et al. (2013) found the same result in their inversions of this model, and they demonstrated that the culprit was the “cross talk” between vertical and horizontal flows that the sensitivity kernels, and inversions, are unable to disentangle. In our SOLA inversion, an explicit cross-talk term is included (Švanda et al., 2011), and even still the problem persists. The sensitivity kernels are not completely accurate. This can be verified by comparing measured and forward-modeled travel-time differences, and as Dombroski et al. (2013) showed (and we verified) there are anomalies in some of the travel-time maps. However, since both inversions use the same kernel functions, the relative comparisons are meaningful. Examples at other depths are given in Appendix B.

Figure 9 in Appendix B shows the corner plot for the probabilistic inversion. The only “known” parameters are $p_4 = R$ and $p_5 = k$, so comparisons between the input ones and inferred ones cannot be made due to the differing models. The probabilistic inversion overestimates the flow speeds, and is mainly due to the estimation of the $p_2 = \sigma_z$ and $p_3 = z_0$ parameters (Section 3.2.4 gives more evidence of this). These control the location of the peak of the vertical-velocity profile and its width. There are (at least) two reasons for the poor estimation of p_2 and p_3 . The first, as the corner plot shows in the 2D marginalized PDFs, is that these two parameters are not highly correlated with the others, but more so with themselves. There must be some correlations due to the continuity-equation constraint, but it is a weak one. This could indicate a poor parametrization of this particular model for supergranulation¹. The second reason has to do with the sensitivity functions used here. They have very little sensitivity below 8 Mm, while the depth of the profile extends to a depth of about 12 Mm. The likelihood is thus not informative, and the 1D PDF for parameter p_3 is not very Gaussian.

3.2.4. Does Additional Data Bring New Information?

For researchers who have experience computing inversions in local helioseismology, it can be non-trivial to understand how the addition of extra observations will (positively or adversely) affect the results. Indeed, a brief discussion regarding this point is presented by Dombroski et al. (2013) in their results section. For instance, consider one set of measurements using particular seismic waves. Now, consider another set of measurements using the same seismic waves but where the only difference is different travel distances. Will including the second set with the first improve the inversion, just add unwanted noise, or improve the noise? Just doing this experiment may not answer the question either, since the differences may be subtle, and SOLA or RLS inversions can be very sensitive to any outlier measurement points.

¹Indeed, the conclusions from earlier work using this model (Duvall and Hanasoge, 2012; Duvall, Hanasoge, and Chakraborty, 2014) showed a very shallow supergranule that is quite different from what other studies in the literature have found. While the analysis may very well be correct, more checks on the model need to be carried out, which is beyond the scope of this article




Figure 8 Metrics for 11 example probabilistic inversions. Each inversion comprises N (left-most numbers) travel-time maps, consisting of the configuration shown by the next eight columns of the matrix. The n_i denote the mode radial order, then the annulus geometry, and then the travel distances [Mm]. Filled circles indicate inclusion in the inversion. Beyond the vertical dashed line, the next five columns represent the variance reduction in the PDF of the parameters compared to the priors (given by the gray scale). The final two columns give the correlation coefficient C_τ between the inversion results and the measurements in data space. Also given is the slope [m_τ] of a fit to the correlated data sets.

The probabilistic inversion provides a way to study this question more quantitatively. We design a simple demonstration experiment, and leave a full analysis to another article. Eleven probabilistic inversions are computed, each one having different combinations or different numbers of input data sets. Everything else is kept fixed.

Two metrics are then calculated to assess the results. We compare the variance of the priors to the variance of the posteriors. Imagine the worst-case scenario, when the variance is not reduced at all. This would imply that the addition of data has provided no new information on the model parameters. The goal of any inversion is to reduce the variance of the parameter estimation. The variance reduction metric is computed as

$$\frac{\text{var}[\text{PDF}(\mathbf{m})] - \text{var}[\rho(\mathbf{m})]}{\text{var}[\rho(\mathbf{m})]} \times 100, \tag{30}$$

where $\rho(\mathbf{m})$ is the distribution of the model priors (see Equation 16) whose range of values is in Table 1. The other metric is the simple correlation coefficient between the travel-time measurements and the forward measurements provided using the median of the PDF.

The results are provided in Figure 8. To understand what is shown, consider, for example, the second row of the matrix. $N = 2$ means there are two travel-time maps used, $\delta\tau_{oi}$ and $\delta\tau_{we}$ for the f -mode at a travel distance of 6 Mm. The black circles for these quantities are filled. To the right of the dashed line, the variance reduction of the five parameters (as a percentage) are given by the gray scale. To be specific, the values for p_1 through p_5 in row 2 are [74.7, 34.2, 1.6, 97.0, 55.5]%. The data have not provided much information on p_3 at all, as suspected. After that, C_τ is the correlation coefficient, and m_τ is the slope of a simple linear fit between the travel-time vectors. In almost all trials, the inferred data have larger amplitudes ($m_\tau \gtrsim 1$).

The third row is an inversion with only one change: the $\delta\tau_{we}$ measurements are removed and an extra travel distance is added. The precise values of the variance reduction are now [82.5, 27.6, 0.4, 97.2, 68.2]%. The first parameter has gone from dark gray to black, over the 80% level, p_2 and p_3 are marginally worse, and p_5 is marginally better. Finally, the fourth row also uses two sets of measurements, with only one annulus geometry and one distance, but now two ridges (n_0 and n_1). This results in a better variance reduction of p_2 than the other cases, and brings the slope closer to one. One might conclude, for this scenario, that, given a very limited number of measurements, it is best to use more ridges than adding distances or anything else.

One can continue this way for the other cases to find interesting trends. Inspecting the matrix as a whole, a few things stand out. p_1 and p_4 are the best “resolved” parameters, and p_2 and p_3 are the least. A quick glance at the PDFs in the corner plot in Figure 9 confirms this. The last row in the matrix is from an inversion using 18 different travel-time maps, and the variance reduction for p_2 and p_3 are 62% and 40%, respectively, the best in the set. The correlation between maps is consistently high, and the slope fluctuates a bit, but is overall acceptable.

To answer the question posed in this subsection – yes, at least in this particular example. While the addition of new data might not visually and qualitatively improve the comparison in data space or model space (as demonstrated by the unchanging correlations), the variance, or uncertainties of the parameters, generally does improve.

In principle, such an analysis could also be achieved from the SOLA inversions, but it would be more cumbersome. Minimizing the cost function properly, i.e. calculating a good averaging kernel, becomes much more difficult with fewer and fewer kernels. Then the trade-off parameters change and some of the results would not be sensible. However, in the probabilistic framework, this is entirely reasonable and instructive.

4. Discussion

The previous sections have contrasted two linear methods for interpreting helioseismic measurements. For comparison, we label the class of inversions similar to SOLA as Method 1, and the class of statistical and probabilistic inversions as Method 2. There are several similarities between these two approaches. Both methods require a type of forward equation relating the unknowns and measurements. Several such equations are provided (Equations 1, 25, 29). Both methods can be run numerically using parallelization when formulated appropriately.

The practical differences outnumber the similarities. Method 2 needs more than a forward equation; it requires a generative model that can be parametrized with a manageable number of parameters. Otherwise, the computational cost may become prohibitive. It would be difficult to utilize Method 2 to make synoptic flow maps of the Sun that contain many different convective structures and size scales, as standard “pipeline” inversions now do for local helioseismic data (Zhao et al., 2012). It would require hundreds of parameters, with very little prior information. Similarly, there is no pre-determined form of the solution when Method 1 is used, and as such it cannot incorporate priors like Method 2. Method 1 does not use the data until the last step, whereby it combines the measurements in an “optimal” way based on how the sensitivity functions were combined (it does use the noise covariance, however, in the computation of the large matrix). Method 1 requires ways to deal with computing the inverse of a large, usually ill-conditioned matrix. Method 2 provides a statistical interpretation of the solution, while Method 1 is forced to provide a “best model.”

Beyond similarities and differences, inversion methods need to be validated. Numerous helioseismic studies over the past decade have employed numerical models for validation purposes. This is a powerful strategy, since one can quantitatively test inversion results on the known answer from the model. The results of these studies provide very consistent conclusions. On the one hand, there are those that use Method 1 and measurements of (non-magnetic) numerical models that do not have realistic noise, although usually some form of noise is added to the measurements after the fact. The findings are generally encouraging (e.g. example 1 in this article; Hartlep et al., 2013; Jackiewicz, Serebryanskiy, and Kholikov, 2015; Korda and Švanda, 2019). This would seem to indicate rather persuasively that Method 1, as well as the sensitivity functions (either ray or Born), can be used to accurately solve problems. On the other hand, when more realistic simulation models are studied in the same way, the results are *somewhat* in agreement near the surface, but quickly diverge below ≈ 3 Mm (example 2 in this article; Zhao et al., 2007; Dombroski et al., 2013; DeGrave, Jackiewicz, and Rempel, 2014; DeGrave et al., 2018). The solar-like realization noise in these models is a significant barrier, which brings skepticism to *any* inversion results using actual solar data and Method 1 (as commented on by Braun and Birch, 2008a; Švanda, 2015; Korda, Švanda, and Zhao, 2019).

Despite heroic efforts and substantial progress, there unfortunately have not been as many significant advancements as one would expect in our understanding of the Sun from explicit inversions of local helioseismic data (Gizon, Birch, and Spruit, 2010). The two examples in this work are cases in point, where still no consensus has been established regarding supergranulation and meridional circulation (Giles et al., 1997; Zhao et al., 2013; Rajaguru and Antia, 2015; Liang and Chou, 2015; Jackiewicz, Serebryanskiy, and Kholikov, 2015; Duvall and Hanasoge, 2012; Hathaway, 2012; Greer, Hindman, and Toomre, 2016).

Indeed, most of the fundamental breakthroughs in local helioseismology have come from the observations alone, rather than the formal interpretation of them. Examples include far-side imaging from acoustic holography and time–distance (Lindsey and Braun, 2000; Zhao, 2007), direct imaging of large-scale flows (Woodard, 2002), acoustic absorption by sunspots (Braun, Duvall, and Labonte, 1987), flared-induced sunquakes (Kosovichev and Zharkova, 1998), and the recent detection of solar Rossby waves using different helioseismic measurement strategies (Löptien et al., 2018; Hanasoge and Mandal, 2019; Proxauf et al., 2020), among many others.

The potential issues that inhibit a full helioseismic analysis of certain outstanding problems include systematics and realization noise inherent in measurements, the theoretical treatment of seismic wave scattering from solar perturbations, and the inverse methods applied. There are many ways for dealing with each of these factors at various levels, and this work provides a possible avenue forward for exploration of the inversion component.

5. Summary and Outlook

In this article we described a probabilistic inversion scheme for time–distance helioseismology that uses Bayesian statistics and Monte Carlo sampling. A few simple examples were carried out and compared with the commonly used SOLA technique. The examples used synthetic data where the known answer was the target of the inversions. Given that the input sets of measurements and sensitivity functions were rather minimal, the goal was not to solve these problems completely (i.e. infer the flows as well as possible), but to demonstrate some of the strengths and weaknesses of these two approaches.

While the examples were highly idealized, the intercomparison consistently showed that the SOLA inversions systematically underestimate the flow speeds and the noise levels, compared to the other method. This may not be too surprising given that Method 2 exploits a generative model with relatively few parameters. It is not surprising either that the solutions using Method 2 are always smooth, since they are constructed as such. However, the probabilistic inversions also crucially provide informative posterior probability distribution functions on the model parameters that are more consistent with the known answer. This was the case even when using uninformative priors. One may question the need to use Method 2 for (likely) highly linear problems such as meridional circulation. In some of the example cases, however, the posteriors are not Gaussian, which could be a possible reason why SOLA or least-squares methods are not optimal. At the very least, the probabilistic method could be used to explore which helioseismic problems have complex PDFs and demonstrate why other inverse methods get trapped in local minima.

SOLA inversions can be tuned at some level to obtain different properties of the solution. A particular advantage of the probabilistic inversion scheme is that many realizations of the solution (in data space and/or model space) are computed automatically, allowing for a broad view of any particular model and its relative probability given the data.

Future work in this area should concentrate on developing well-parametrized models of solar structures amenable to helioseismic investigation. For example, the recent meridional-circulation model of Liang et al. (2018) is much more flexible than the one presented here. Good models may also increase computational efficiency. Finally, one could imagine forward (generative) models that compute other observables than travel times, such as the more fundamental and information-laden cross correlations. This would be another move in the direction of full-waveform inversions.

While this work is focused on time–distance helioseismology, application to ring-diagram analysis, helioseismic holography, or direct modeling is straightforward. For those interested in similar applications, the affine-invariant ensemble MCMC algorithm has been made available in Python (`emcee`: github.com/dfm/emcee) and Matlab (`GWMCMC`: github.com/grinsted/gwmcnc) and can be adapted to many types of problems.

Acknowledgements This article is dedicated to the life and work of the late Michael J. Thompson, who unselfishly taught me a significant amount of inverse theory. The author wishes to thank Aaron Birch, Aleczander Herczeg, Jon Holtzman, and Shukur Kholikov for fruitful discussions, as well as Doug Braun for making the simulation data publicly available. The article made use of the `GWMCMC` code written by Aslak Grinsted. Partial funding support is acknowledged from the National Science Foundation under Grant Number 1351311 and from NASA under Award Number 80NSSC18K0672.

Disclosure of Potential Conflicts of Interest The author declares that he has no conflicts of interest.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Appendix A: Supergranulation Inversion Details

This article presents a comparison of six SOLA inversions (two flow components at three target depths) to the probabilistic one for a simulated supergranule model. For completeness, details of the inversions and how they are compared are itemized here:

- i) The simulation domain is 100 Mm on a side horizontally, sampled evenly at 1/3 Mm. It extends to -25 Mm at the bottom boundary. Cross correlations and travel times were measured in an area 50 Mm on a side at the same sampling. Each map is therefore 150×150 points.

- ii) Born-sensitivity kernels are computed on the same horizontal grid as the travel times and they extend to 15 Mm below the surface using 55 points in depth. The kernels use input model power spectra that have been Fourier filtered to separately retain the first three radial orders, including the f -mode. Everything is computed to match the details of the travel times as discussed in Section 3.2.2.
- iii) The SOLA inversion code is the faster formulation in Fourier space that can be run in parallel (Jackiewicz et al., 2012). It includes a cross-talk parameter as in Švanda et al. (2011). Trade-off curves (L-curves) were studied to find the best regularization parameters. To obtain sensible results in terms of noise and misfit, the 3D Gaussian target functions have FWHM on the order of 10 Mm and 3 Mm in the horizontal and vertical directions, respectively. In general, the SOLA inversion results are very similar to the ones published by Dombroski et al. (2013) using an RLS inversion scheme.
- iv) The probabilistic inversion has 2×10^5 steps. Each of the five parameters was assigned 120 walkers. With thinning every five steps, the PDF was sampled $2 \times 10^5 / (120 \times 5) = 333$ times per walker, as before. More walkers were used in this inversion than the meridional example because the PDFs are more complicated, and it is suggested to use more walkers for good sampling. These values are sufficient to give good autocorrelation properties of the walker time series.
- v) The data vector \mathbf{d} is composed of only a small subset (12 of the 135) of the travel-time measurement maps. To speed up the computation, each of the 150×150 pixel maps was binned down to only 19×19 pixels. The data vector is therefore 4332 measurements. To match the size of the data vector, the forward vector resulting from the operation $g(\mathbf{m})$ uses kernels that are also binned down. This is noteworthy, in that the SOLA inversion uses the full maps. The probabilistic inversion is quite powerful with substantially less data. Of course, the 4332 measurements are not all completely independent, as described by the covariance matrix.
- vi) The full noise covariance matrix used in the likelihood function (Equation 17) was constructed from smaller pairs of covariances $\text{Cov}[\delta\tau_i, \delta\tau_j]$ and ordered to correctly match the travel-time data vector.
- vii) The uniform priors on the model are given in Table 1. As is common (Foreman-Mackey et al., 2013), the walkers are initialized with values in a small Gaussian ball, centered somewhere within the prior bounds. The initial guess is not critical, as the walkers soon quickly explore the full parameter space after the burn-in stage.
- viii) The entire run on a two-core desktop machine (thus minimal parallelization) was about 80 minutes. The most expensive task is the computation of the forward travel times in Equation 29, which takes about 0.2 seconds each.

Appendix B: Supergranulation Flow Inversions

A few more results are presented here to add to those in Section 3.2.3. The corner plot for the probabilistic inversion is given in Figure 9.

The flow-inversion comparison near the top of the simulation domain is shown in Figure 10. For the horizontal flows, the SOLA inversion underestimates the (smoothed) flow amplitude by about a factor of two, while the probabilistic inversion overestimates the amplitude by 15%. The profile structure is primarily determined by parameters p_4 and p_5 . The peak vertical velocity is underestimated by 35% in the SOLA inversion, and it is overestimated by 60% in the probabilistic inversion. The f -mode is generally the most sensitive at

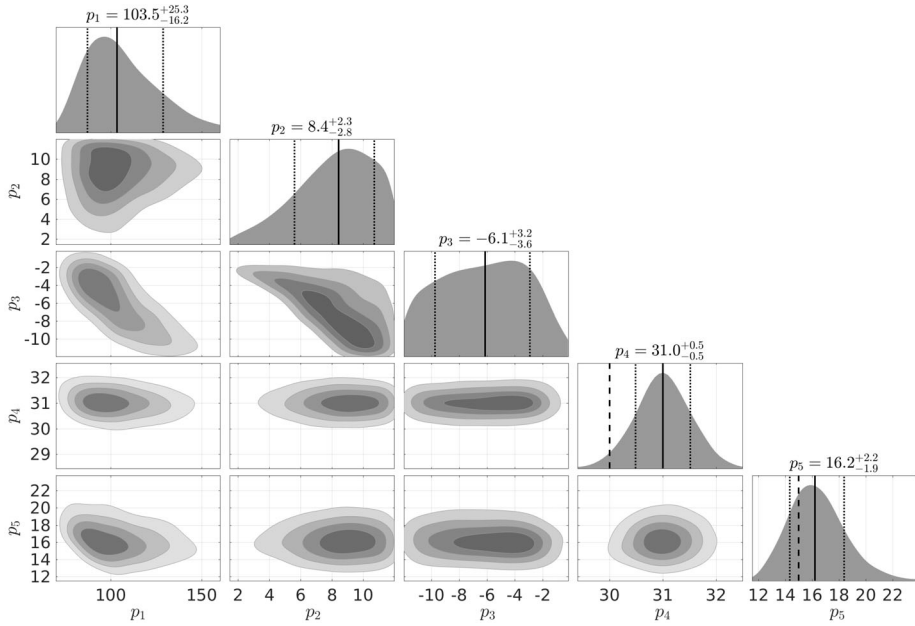


Figure 9 Corner plot showing the marginalized PDFs of the five parameters in the supergranulation example. The marginalized distribution for each parameter independently is shown in the histograms along the diagonal, and the marginalized 2D distributions as contour plots in the other panels. For each 1D histogram, the median of the PDF is the solid-black line, and the dotted lines give the 68% confidence interval, whose numerical values are provided at the top of each panel. The dashed-black lines for p_4 and p_5 are the known input parameter values for the two in common. The contour levels of the 2D joint probability densities are at 20%, 40%, 60%, and 80% confidence intervals.




Figure 10 Comparison of supergranulation inversion results at the model photosphere. The left (right) panels are for the v_x (v_z) inversion. The top rows are the flow fields for the simulation, the SOLA inversion, and the probabilistic inversion computed from the median of the PDF. Darker shading corresponds to positive velocities, and the scale is the same in each set. The bottom panels are cuts through the supergranule at $y = 0$, given by the dashed line in the top panels. Twenty random samples of the PDF are drawn and computed in the model space. A few representative points from the SOLA inversion are given with uncertainties.




Figure 11 Comparison of supergranulation inversion results near about 5 Mm beneath the model photosphere. The left (right) panels are for the v_x (v_z) inversion. The top rows are the flow fields for the simulation, the SOLA inversion, and the probabilistic inversion computed from the median of the PDF. Darker shading corresponds to positive velocities, and the scale is the same in each set. The bottom panels are cuts through the supergranule at $y = 0$, given by the dashed line in the top panels. Twenty random samples of the PDF are drawn and computed in the model space. A few representative points from the SOLA inversion are given with uncertainties.

these layers, but its sensitivity to vertical flows is weak. This affects both inversions. v_z^{SOLA} also has artifacts in its flow structure.

A comparison near 5 Mm beneath the surface is given in Figure 11. As is common in standard helioseismic inversions using typical data sets, the inferences at this depth are uninformative (Zhao et al., 2007). The vertical velocity inferred is weak and is the wrong sign, which was also noticed in Dombroski et al. (2013). The probabilistic inversion provides a good estimate of the flow structure at this depth.

References

- Arregui, I.: 2018, Bayesian coronal seismology. *Adv. Space Res.* **61**, 655. DOI. ADS.
- Backus, G.E., Gilbert, J.F.: 1968, The resolving power of gross Earth data. *Geophys. J. Roy. Astron. Soc.* **16**, 169. DOI. ADS.
- Basu, S.: 2016, Global seismology of the Sun. *Liv. Rev. Solar Phys.* **13**, 2. DOI. ADS.
- Bhattacharya, J., Hanasoge, S.M.: 2016, Strategies in seismic inference of supergranular flows on the Sun. *Astrophys. J.* **826**, 105. DOI. ADS.
- Birch, A.C., Gizon, L.: 2007, Linear sensitivity of helioseismic travel times to local flows. *Astron. Nachr.* **328**, 228. DOI. ADS.
- Böning, V.G.A., Roth, M., Jackiewicz, J., Kholikov, S.: 2017, Validation of spherical born approximation sensitivity functions for measuring deep solar meridional flow. *Astrophys. J.* **838**, 53. DOI. ADS.
- Braun, D.C., Birch, A.C.: 2008a, Prospects for the detection of the deep solar meridional circulation. *Astrophys. J. Lett.* **689**, L161. DOI. ADS.
- Braun, D.C., Birch, A.C.: 2008b, Surface-focused seismic holography of sunspots: I. Observations. *Solar Phys.* **251**, 267. DOI. ADS.
- Braun, D.C., Birch, A.C.: 2009, How much data do we need to detect the deep solar meridional circulation? In: Dikpati, M., Arentoft, T., González Hernández, I., Lindsey, C., Hill, F. (eds.) *Solar-Stellar Dynamics as Revealed by Helio- and Asteroseismology: GONG 2008/SOHO 21*, CS-416, Astron. Soc. Pacific, San Francisco, 131. ADS.
- Braun, D.C., Duvall, T.L. Jr., Labonte, B.J.: 1987, Acoustic absorption by sunspots. *Astrophys. J. Lett.* **319**, L27. DOI. ADS.
- Christensen-Dalsgaard, J., Hansen, P.C., Thompson, M.J.: 1993, Generalized singular value decomposition analysis of helioseismic inversions. *Mon. Not. Roy. Astron. Soc.* **264**, 541. DOI. ADS.

- Christensen-Dalsgaard, J., Dappen, W., Ajukov, S.V., Anderson, E.R., Antia, H.M., Basu, S., Baturin, V.A., Berthomieu, G., Chaboyer, B., Chitre, S.M., Cox, A.N., Demarque, P., Donatowicz, J., Dziembowski, W.A., Gabriel, M., Gough, D.O., Guenther, D.B., Guzik, J.A., Harvey, J.W., Hill, F., Houdek, G., Iglesias, C.A., Kosovichev, A.G., Leibacher, J.W., Morel, P., Proffitt, C.R., Provost, J., Reiter, J., Rhodes, E.J. Jr., Rogers, F.J., Roxburgh, I.W., Thompson, M.J., Ulrich, R.K.: 1996, The current state of solar modeling. *Science* **272**, 1286. [ADS](#).
- Corbard, T., Berthomieu, G., Morel, P., Provost, J., Schou, J., Tomczyk, S.: 1997, Solar internal rotation from LOWL data. A 2D regularized least-squares inversion using B-splines. *Astron. Astrophys.* **324**, 298. [ADS](#).
- Corsaro, E., De Ridder, J.: 2014, DIAMONDS: a new Bayesian nested sampling tool. Application to peak bagging of solar-like oscillations. *Astron. Astrophys.* **571**, A71. [DOI](#). [ADS](#).
- Couvidat, S., Birch, A.C., Kosovichev, A.G.: 2006, Three-dimensional inversion of sound speed below a sunspot in the born approximation. *Astrophys. J.* **640**, 516. [DOI](#). [ADS](#).
- DeGrave, K., Jackiewicz, J., Rempel, M.: 2014, Validating time-distance helioseismology with realistic quiet-Sun simulations. *Astrophys. J.* **788**, 127. [DOI](#). [ADS](#).
- DeGrave, K., Braun, D.C., Birch, A.C., Crouch, A.D., Javornik, B.: 2018, Validating forward modeling and inversions of helioseismic holography measurements. *Astrophys. J.* **863**, 34. [DOI](#). [ADS](#).
- Dhruv, V., Bhattacharya, J., Hanasoge, S.M.: 2019, Validating time-distance helioseismic inversions for non-separable subsurface profiles of an average supergranule. *Astrophys. J.* **883**, 136. [DOI](#). [ADS](#).
- Dikpati, M., Charbonneau, P.: 1999, A Babcock–Leighton flux transport dynamo with solar-like differential rotation. *Astrophys. J.* **518**, 508. [DOI](#). [ADS](#).
- Dombroski, D.E., Birch, A.C., Braun, D.C., Hanasoge, S.M.: 2013, Testing helioseismic-holography inversions for supergranular flows using synthetic data. *Solar Phys.* **282**, 361. [DOI](#). [ADS](#).
- Duvall, T.L. Jr., Birch, A.C.: 2010, The vertical component of the supergranular motion. *Astrophys. J. Lett.* **725**, L47. [DOI](#). [ADS](#).
- Duvall, T.L. Jr., Hanasoge, S.M.: 2012, Subsurface supergranular vertical flows as measured using large distance separations in time-distance helioseismology. *Solar Phys.*, 136. [DOI](#). [ADS](#).
- Duvall, T.L. Jr., Hanasoge, S.M., Chakraborty, S.: 2014, Additional evidence supporting a model of shallow, high-speed supergranulation. *Solar Phys.* **289**, 3421. [DOI](#). [ADS](#).
- Duvall, T.L. Jr., Kosovichev, A.G., Scherrer, P.H., Bogart, R.S., Bush, R.I., de Forest, C., Hoeksema, J.T., Schou, J., Saba, J.L.R., Tarbell, T.D., Title, A.M., Wolfson, C.J., Milford, P.N.: 1997, Time-distance helioseismology with the MDI instrument: initial results. *Solar Phys.* **170**, 63. [ADS](#).
- Ferret, R.Z.: 2019, SDO/HMI observations of the average supergranule are not compatible with separable flow models. *Astron. Astrophys.* **623**, A98. [DOI](#). [ADS](#).
- Foreman-Mackey, D., Hogg, D.W., Lang, D., Goodman, J.: 2013, emcee: the MCMC hammer. *Publ. Astron. Soc. Pac.* **125**, 306. [DOI](#). [ADS](#).
- Fournier, D., Gizon, L., Hohage, T., Birch, A.C.: 2014, Generalization of the noise model for time-distance helioseismology. *Astron. Astrophys.* **567**, A137. [DOI](#). [ADS](#).
- Giles, P.M., Duvall, T.L. Jr., Scherrer, P.H., Bogart, R.S.: 1997, A subsurface flow of material from the Sun's equator to its poles. *Nature* **390**, 52. [DOI](#). [ADS](#).
- Gizon, L., Birch, A.C.: 2002, Time-distance helioseismology: the forward problem for random distributed sources. *Astrophys. J.* **571**, 966. [DOI](#). [ADS](#).
- Gizon, L., Birch, A.C.: 2004, Time–distance helioseismology: noise estimation. *Astrophys. J.* **614**, 472. [DOI](#). [ADS](#).
- Gizon, L., Birch, A.C., Spruit, H.C.: 2010, Local helioseismology: three-dimensional imaging of the solar interior. *Annu. Rev. Astron. Astrophys.* **48**, 289. [DOI](#). [ADS](#).
- Gizon, L., Schunker, H., Baldner, C.S., Basu, S., Birch, A.C., Bogart, R.S., Braun, D.C., Cameron, R., Duvall, T.L. Jr., Hanasoge, S.M., Jackiewicz, J., Roth, M., Stahn, T., Thompson, M.J., Zharkov, S.: 2009, Helioseismology of sunspots: a case study of NOAA region 9787. *Space Sci. Rev.* **144**, 249. [DOI](#). [ADS](#).
- Goodman, J., Weare, J.: 2010, Ensemble samplers with affine invariance. *Commun. Appl. Math. Comput. Sci.* **5**, 65. [DOI](#). [ADS](#).
- Gough, D.O., Thompson, M.J.: 1991,. In: Cox, A.N., Livingston, W.C., Matthews, M. S. (eds.) *The Inversion Problem*, 519. [ADS](#).
- Greer, B.J., Hindman, B.W., Toomre, J.: 2016, Helioseismic imaging of supergranulation throughout the Sun's near-surface shear layer. *Astrophys. J.* **824**, 128. [DOI](#). [ADS](#).
- Hanasoge, S.M.: 2014, Full waveform inversion of solar interior flows. *Astrophys. J.* **797**, 23. [DOI](#). [ADS](#).
- Hanasoge, S., Mandal, K.: 2019, Detection of Rossby waves in the Sun using normal-mode coupling. arXiv e-prints. [arXiv](#). [ADS](#).
- Hanasoge, S.M., Larsen, R.M., Duvall, T.L. Jr., DeRosa, M.L., Hurlburt, N.E., Schou, J., Roth, M., Christensen-Dalsgaard, J., Lele, S.K.: 2006, Computational acoustics in spherical geometry: steps toward validating helioseismology. *Astrophys. J.* **648**, 1268. [DOI](#). [ADS](#).

- Hanasoge, S.M., Birch, A., Gizon, L., Tromp, J.: 2011, The adjoint method applied to time-distance helioseismology. *Astrophys. J.* **738**, 100. DOI ADS.
- Hartlep, T., Zhao, J., Kosovichev, A.G., Mansour, N.N.: 2013, Solar wave-field simulation for testing prospects of helioseismic measurements of deep meridional flows. *Astrophys. J.* **762**, 132. DOI ADS.
- Hathaway, D.H.: 2012, Supergranules as probes of solar convection zone dynamics. *Astrophys. J. Lett.* **749**, L13. DOI ADS.
- Hilbe, J.M., de Souza, R.S., Ishida, E.E.O.: 2017, *Bayesian Models for Astrophysical Data Using R, JAGS, Python, and Stan*. DOI ADS.
- Hogg, D.W., Foreman-Mackey, D.: 2018, Data analysis recipes: using Markov chain Monte Carlo. *Astrophys. J. Suppl.* **236**, 11. DOI ADS.
- Howe, R.: 2009, Solar interior rotation and its variation. *Liv. Rev. Solar Phys.* **6**, 1. ADS.
- Jackiewicz, J., Gizon, L., Birch, A.C.: 2008, High-resolution mapping of flows in the solar interior: fully consistent OLA inversion of helioseismic travel times. *Solar Phys.* **251**, 381. DOI ADS.
- Jackiewicz, J., Serebryanskiy, A., Kholikov, S.: 2015, Meridional flow in the solar convection zone. II. Helioseismic inversions of GONG data. *Astrophys. J.* **805**, 133. DOI ADS.
- Jackiewicz, J., Birch, A.C., Gizon, L., Hanasoge, S.M., Hohage, T., Ruffio, J.-B., Švanda, M.: 2012, Multi-channel three-dimensional SOLA inversion for local helioseismology. *Solar Phys.* **276**, 19. DOI ADS.
- Jensen, J.M., Jacobsen, B.H., Christensen-Dalsgaard, J.: 1998, MCD inversion for sound speed using time-distance data. In: Korzennik, S. (ed.) *Structure and Dynamics of the Interior of the Sun and Sun-like Stars*, SP-418, ESA, Noordwijk, 635. ADS.
- Komm, R., Howe, R., Hill, F., Miesch, M., Haber, D., Hindman, B.: 2007, Divergence and vorticity of solar subsurface flows derived from ring-diagram analysis of MDI and GONG data. *Astrophys. J.* **667**, 571. DOI ADS.
- Korda, D., Švanda, M.: 2019, Combined helioseismic inversions for 3D vector flows and sound-speed perturbations. *Astron. Astrophys.* **622**, A163. DOI ADS.
- Korda, D., Švanda, M., Zhao, J.: 2019, Comparison of time-distance inversion methods applied to SDO/HMI dopplergrams. *Astron. Astrophys.* **629**, A55. DOI ADS.
- Kosovichev, A.G., Duvall, T.L. Jr.: 1997, Acoustic tomography of solar convective flows and structures. In: Pijpers, F.P., Christensen-Dalsgaard, J., Rosenthal, C.S. (eds.) *ASSL 225: SCORE '96: Solar Convection and Oscillations and Their Relationship*, Kluwer, Dordrecht, 241. ADS.
- Kosovichev, A.G., Zharkova, V.V.: 1998, X-ray flare sparks quake inside Sun. *Nature* **393**, 317. DOI ADS.
- Liang, Z.-C., Chou, D.-Y.: 2015, Effects of solar surface magnetic fields on the time-distance analysis of solar subsurface meridional flows. *Astrophys. J.* **805**, 165. DOI ADS.
- Liang, Z.-C., Gizon, L., Birch, A.C., Duvall, T.L. Jr., Rajaguru, S.P.: 2018, Solar meridional circulation from twenty-one years of SOHO/MDI and SDO/HMI observations. Helioseismic travel times and forward modeling in the ray approximation. *Astron. Astrophys.* **619**, A99. DOI ADS.
- Lindsey, C., Braun, D.C.: 2000, Basic principles of solar acoustic holography. *Solar Phys.* **192**, 261. ADS.
- Löptien, B., Gizon, L., Birch, A.C., Schou, J., Proxauf, B., Duvall, T.L. Jr., Bogart, R.S., Christensen, U.R.: 2018, Global-scale equatorial Rossby waves as an essential component of solar internal dynamics. *Nat. Astron.* **2**, 568. DOI ADS.
- Moradi, H., Baldner, C., Birch, A.C., Braun, D.C., Cameron, R.H., Duvall, T.L. Jr., Gizon, L., Haber, D., Hanasoge, S.M., Hindman, B.W., Jackiewicz, J., Khomenko, E., Komm, R., Rajaguru, P., Rempel, M., Roth, M., Schlichenmaier, R., Schunker, H., Spruit, H.C., Strassmeier, K.G., Thompson, M.J., Zharkov, S.: 2010, Modeling the subsurface structure of sunspots. *Solar Phys.*, 171. DOI ADS.
- Pijpers, F.P., Thompson, M.J.: 1994, The SOLA method for helioseismic inversion. *Astron. Astrophys.* **281**, 231. ADS.
- Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P.: 2007, *Numerical Recipes 3rd Edition: The Art of Scientific Computing*, 3rd edn. Cambridge University Press, Cambridge. ISBN: 0521880688.
- Proxauf, B., Gizon, L., Löptien, B., Schou, J., Birch, A.C., Bogart, R.S.: 2020, Exploring the latitude and depth dependence of solar Rossby waves using ring-diagram analysis. *Astron. Astrophys.* **634**, A44. DOI arXiv ADS.
- Rajaguru, S.P., Antia, H.M.: 2015, Meridional circulation in the solar convection zone: time-distance helioseismic inferences from four years of HMI/SDO observations. *Astrophys. J.* **813**, 114. DOI ADS.
- Rieutord, M., Rincon, F.: 2010, The Sun's supergranulation. *Liv. Rev. Solar Phys.* **7**, 2. ADS.
- Sambridge, M., Mosegaard, K.: 2002, Monte Carlo methods in geophysical inverse problems. *Rev. Geophys.* **40**, 1009. DOI ADS.
- Schou, J., Christensen-Dalsgaard, J., Thompson, M.J.: 1994, On comparing helioseismic two-dimensional inversion methods. *Astrophys. J.* **433**, 389. DOI ADS.
- Sharma, S.: 2017, Markov chain Monte Carlo methods for Bayesian data analysis in astronomy. *Annu. Rev. Astron. Astrophys.* **55**, 213. DOI ADS.

- Sokal, A.: 1997. In: DeWitt-Morette, C., Cartier, P., Folacci, A. (eds.) *Monte Carlo Methods in Statistical Mechanics: Foundations and New Algorithms*, Springer, Boston, 131. [DOI](#).
- Švanda, M.: 2012, Inversions for average supergranular flows using finite-frequency kernels. *Astrophys. J. Lett.* **759**, L29. [DOI](#). [ADS](#).
- Švanda, M.: 2015, Issues with time-distance inversions for supergranular flows. *Astron. Astrophys.* **575**, A122. [DOI](#). [ADS](#).
- Švanda, M., Gizon, L., Hanasoge, S.M., Ustyugov, S.D.: 2011, Validated helioseismic inversions for 3D vector flows. *Astron. Astrophys.* **530**, A148. [DOI](#). [ADS](#).
- Tarantola, A.: 2005, *Inverse Problem Theory and Methods for Model Parameter Estimation*, Society for Industrial and Applied Mathematics, Philadelphia.
- Thompson, M.J., Christensen-Dalsgaard, J., Miesch, M.S., Toomre, J.: 2003, The internal rotation of the Sun. *Annu. Rev. Astron. Astrophys.* **41**, 599. [DOI](#). [ADS](#).
- van Ballegoijen, A.A., Choudhuri, A.R.: 1988, The possible role of meridional flows in suppressing magnetic buoyancy. *Astrophys. J.* **333**, 965. [DOI](#). [ADS](#).
- Woodard, M.F.: 2002, Solar subsurface flow inferred directly from frequency-wavenumber correlations in the seismic velocity field. *Astrophys. J.* **565**, 634. [DOI](#). [ADS](#).
- Zhao, J.: 2007, Time-distance imaging of solar far-side active regions. *Astrophys. J. Lett.* **664**, L139. [DOI](#). [ADS](#).
- Zhao, J., Kosovichev, A.G.: 2003, On the inference of supergranular flows by time-distance helioseismology. In: Sawaya-Lacoste, H. (ed.) *GONG+ 2002. Local and Global Helioseismology: The Present and Future*, *SP-517*, ESA, Noordwijk, 417. [ADS](#).
- Zhao, J., Georgobiani, D., Kosovichev, A.G., Benson, D., Stein, R.F., Nordlund, Å.: 2007, Validation of time-distance helioseismology by use of realistic simulations of solar convection. *Astrophys. J.* **659**, 848. [DOI](#). [ADS](#).
- Zhao, J., Couvidat, S., Bogart, R.S., Parchevsky, K.V., Birch, A.C., Duvall, T.L. Jr., Beck, J.G., Kosovichev, A.G., Scherrer, P.H.: 2012, Time-distance helioseismology data-analysis pipeline for helioseismic and magnetic imager onboard Solar Dynamics Observatory (SDO/HMI) and its initial results. *Solar Phys.* **275**, 375. [DOI](#). [ADS](#).
- Zhao, J., Bogart, R.S., Kosovichev, A.G., Duvall, T.L. Jr., Hartlep, T.: 2013, Detection of equatorward meridional flow and evidence of double-cell meridional circulation inside the Sun. *Astrophys. J. Lett.* **774**, L29. [DOI](#). [ADS](#).