Solid Earth
Discussions

Open Access

Fully probabilistic seismic source inversion – Part 1: Efficient parameterisation

S. C. Stähler and K. Sigloch

Dept. of Earth and Environmental Sciences, Ludwig-Maximilians-Universität (LMU), München, Germany

Correspondence to: S. C. Stähler (staehler@geophysik.uni-muenchen.de)

Abstract

Seismic source inversion is a non-linear problem in seismology where not just the earthquake parameters themselves, but also estimates of their uncertainties are of great practical importance. Probabilistic source inversion (Bayesian inference) is very adapted to this challenge, provided that the parameter space can be chosen small enough to make Bayesian sampling computationally feasible. We propose a framework for PRobabilistic Inference of Source Mechanisms (PRISM) that parameterises and samples earthquake depth, moment tensor, and source time function efficiently by using information from previous non-Bayesian inversions. The source time function is expressed as a weighted sum of a small number of empirical orthogonal functions, which were derived from a catalogue of > 1000 STFs by a principal component analysis. We use a likelihood model based on the cross-correlation misfit between observed and predicted waveforms. The resulting ensemble of solutions provides full uncertainty and covariance information for the source parameters, and permits to propagate these source uncertainties into travel time estimates used for seismic tomography. The computational effort is such that routine, global estimation of earthquake mechanisms and source time functions from teleseismic broadband waveforms is feasible.

1 Introduction

Seismic source inversion is one of the primary tasks of seismology, and the need to explain devastating ground movements was at the origin of the discipline. The interest is to locate the earthquake source using seismogram recordings, and to combine this information with geological knowledge, in order to estimate the probability of further earthquakes in the same region. This purpose is served well by a variety of existing source catalogues, global and regional. Large earthquakes and those in densely instrumented areas are being studied in detail, using extended-source frameworks like finite-fault or back-projection.

Smaller earthquakes ($M_S \leq 7.5$), and especially remote events with sparse data coverage, are better parameterised by a point source. Most catalogues determine only a location and a moment tensor solution, which often allows for identification of the associated fault. But the waveform data contain additional information: for earthquakes exceeding $M_S \geq 5.5$, it is generally possible to invert for the temporal evolution of the rupture, described by a time series called the source time function (STF) (Ruff, 1989; Houston, 2001; Sigloch and Nolet, 2006). While the STF may further aid the understanding of fault characteristics and hazard, our primary motivation for estimating it is a different one: the STF convolves the broadband Green's function and strongly affects its waveform. Waveform tomography estimates three-dimensional earth structure by optimising the fit of observed to predicted waveforms, but at high frequencies (e.g., exceeding 0.1 Hz), such fits can only succeed when the source time function is incorporated into the predicted waveform (Sigloch and Nolet, 2006; Stähler et al., 2012). Hence the purpose here is to develop an automated procedure to routinely estimate broadband source time functions and point source parameters from global seismogram recordings, including a full treatment of parameter uncertainties.

A few recent catalogues now include STF estimates (Vallée et al., 2011; Garcia et al., 2013), but the treatment of parameter uncertainties is still incomplete. Uncertainties in the STF correlate most strongly with source depth estimates, especially for shallow earthquakes (Sigloch and Nolet, 2006), where surface-reflected phases (pP, sP) inevitably enter the time window for STF estimation (see Fig. 1). Inversion for the STF and the moment tensor is linear, whereas inversion for depth is inherently non-linear. Hence gradient-free optimisation techniques like Simulated Annealing (Kirkpatrick et al., 1983) or the first stage of the Neighbourhood Algorithm (Sambridge, 1999a) have become popular; Table 4 presents an overview of gradient-free source inversion algorithms from recent years. These optimisation algorithms provide only rudimentary uncertainty estimates.

A natural alternative, pursued here, is Bayesian sampling, where an ensemble of models is generated. The members of this ensemble are distributed according to the

Title Page

Abstract | Introduction

Conclusions | References

Tables | Figures

◄ | ►|

◄ | ►

Back | Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Probabilistic source inversion I – parametrization

S. C. Stähler and
K. Sigloch

Title Page

Abstract | Introduction

Conclusions | References

Tables | Figures

◄ | ►|

◄ | ►

Back | Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

posterior probability density $P(\boldsymbol{m})$, where $\boldsymbol{m}$ is the model parameter vector to estimate. Integrating over certain parameters of this joint posterior $P(\boldsymbol{m})$, or linear combinations thereof, yields marginal distributions over arbitrary individual parameters or parameter combinations. To the best of our knowledge, ensemble sampling in the context of source parameter estimation has been tried twice so far (Wéber, 2006; Debski, 2008), and has been limited to a few events in either case.

A hurdle to using sampling algorithms has been the efficient parameterisation of the source time function. We propose a parameterisation based on empirical orthogonal wavelets (Sect. 2.1), which reduces the number of free parameters to less than 12 for the STF, and to around 18 in total. We show that this makes Bayesian sampling of the entire model space computationally feasible.

A normalised moment tensor is sampled explicitly, and the scalar moment and absolute values for M_j are derived from the amplitude misfit (Sect. 2.2). Section 3 introduces Bayesian inference as a concept and explains the model space and prior assumptions. The ensemble inference is done with the Neighbourhood Algorithm (Sambridge, 1999a,b). In Sect. 4, the code is applied to a magnitude 5.7 earthquake in Virginia, 2011. Section 5 discusses aspects of our algorithm and potential alternatives. We compare to related studies by other workers in Sect. 5.4 and in the Appendix.

Our procedure is called PRISM (PRobabilistic Inference of Source Mechanisms), by applying it routinely, we plan to publish ensemble solutions for intermediate-size earthquakes in the near future. A usage of uncertainty information gained from the ensemble is demonstrated in Sect. 4.3, where the influence of source uncertainties on tomographic travel time observables is estimated. Further investigations of noise and of inter-station covariances are presented in a companion paper (Stähler et al., 2013).

2 Method

2.1 Parametrisation of the source time function

Source time function (STF) is a synonym for the moment rate $\dot{m}(t)$ of a point source, denoting a time series that describes the rupture evolution of the earthquake

$$u(t) = \sum_{j=1}^{6} G_j(\boldsymbol{r}_s, \boldsymbol{r}_r, t) \cdot s(t) \cdot M_j \tag{1}$$

where $u(t)$ is the vertical or transverse component of the displacement seismogram observed at location $\boldsymbol{r}_r$; $s(t) \equiv \dot{m}(t)$ is the STF; M_j denotes the six distinct elements of the symmetric, 3×3 moment tensor, $\boldsymbol{M}$; and $G(\boldsymbol{r}_s, \boldsymbol{r}_r, t)$ is the Green's function between the hypocenter $\boldsymbol{r}_s$ and receiver location $\boldsymbol{r}_r$.

$\boldsymbol{G}$ is not a 3-D vector because we compute either only its vertical component (for P waves) or its transverse component (for SH waves). In either case, the Green's function is a superposition of six partial Green's functions G_j, corresponding to contributions from six unique moment tensor elements M_j. The orientation of the source is considered to remain fixed during the rupture, i.e., M_j does not depend on t, so that a single time series $s(t)$ is sufficient to describe rupture evolution.

For intermediate-size earthquakes ($5.5 < M_W < 7.0$) the STF typically has a duration of several seconds, which is not short compared to the rapid sequence of P-pP-sP or S-sS pulses that shallow earthquakes produce in broadband seismograms. Most earthquakes are shallow in this sense, i.e., shallower than 50 km. In order to assemble tomography-sized data sets, it is therefore imperative to account for the source time function in any waveform fitting attempt that goes to frequencies above ≈ 0.05 Hz (Sigloch and Nolet, 2006).

Equation (1) is linear in $s(t)$, so that $s(t)$ can be determined by deconvolving $\boldsymbol{G}$ from $\boldsymbol{u}$ if M_j in considered fixed. However, $\boldsymbol{G}$ depends strongly on source depth (third component of vector $\boldsymbol{r}_s$), so that a misestimated source depth will strongly distort the shape

Title Page

Abstract | Introduction

Conclusions | References

Tables | Figures

◄ | ►◄

◄ | ►

Back | Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

of the STF, as demonstrated by Fig. 1. Another complication is present in the fact that observed seismograms $u(t)$ (as opposed to the predicted Green's functions) are time-shifted relative to each other due to 3-D heterogeneity in the earth, and should be empirically aligned before deconvolving $s(t)$.

5   These issues can be overcome by solving iteratively for $s(t)$ and M_j (Sigloch and Nolet, 2006; Stähler et al., 2012), but the approach requires significant human interaction, which poses a challenge for the amounts of data now available for regional or global tomography. Moreover, such an optimisation approach does not provide systematic estimates of parameter uncertainties.

10   Monte Carlo sampling avoids the unstable deconvolution and permits straightforward estimation of full parameter uncertainties and covariances. However, the model space to sample grows exponentially with the number of parameters, and the STF adds a significant number of parameters. In a naive approach, this number could easily be on the order of 100, i.e., computationally prohibitive. For example, the STFs deconvolved in

15 Fig. 1 were parameterised as a time series of 25 s duration, sampled at 10 Hz, and thus yielding 250 unknowns – not efficient, since neighbouring samples are expected to be strongly correlated. This raises the question of how many independent parameters or degrees of freedom this problem actually has.

    Due to intrinsic attenuation of the earth, the highest frequencies still significantly

20 represented in teleseismic P waves are around 1 Hz. If from experience we require a duration of 25 s to render the longest possible STFs occurring for our magnitude range (Houston, 2001), then the time-bandwidth product is $1\,\text{Hz} \cdot 25\,\text{s} = 25$, and the problem cannot have more degrees of freedom than that.

    Efficient parametrisation then amounts to finding a basis of not more than 25 or-

25 thogonal functions that span the subspace of the real-world, band-limited STFs just described. In fact, we can empirically decrease the number of parameters even further. By the method of Sigloch and Nolet (2006), we have semi-automatically deconvolved more than 3000 broadband STFs while building data sets for finite-frequency tomography. Of these, we propose to use the 1000 STFs that we consider most confidently

determined as prior information for what the range of possible STFs looks like, for earthquakes of magnitude $5.5 < M_W < 7.5$. By performing a Principal Component Analysis on this large set of prior STFs, we find that only around 10 empirical orthogonal wavelets are needed to satisfactorily explain almost all of the STFs, as shown in Fig. 2.

5   In concrete terms, we applied the MATLAB function *princomp.m* to a matrix containing the 3000 prior STFs in its rows. The mean over the matrix columns (time samples) was subtracted prior to performing the decomposition, and is shown in Fig. 2a as wavelet $s_0(t)$. Principal Component Analysis then determines $s_1(t)$ as the function orthonormal to $s_0(t)$ that explains as much of the variance in the matrix rows as possible. After sub-

10   tracting (optimally weighted) $s_1(t)$ from each row, function $s_2(t)$ is determined such that it is orthonormal to $s_0(t)$ and $s_1(t)$, and explains as much as possible of the remaining variance. Each subsequent iteration generates another orthonormal s_i until $i = 256$, the number of time samples (matrix columns). The source time function can now be expressed as

15   $$s(t) = \sum_{i=1}^{255} a_i s_i(t) + s_0(t). \tag{2}$$

In this parametrisation, the new unknowns to solve for during source estimation are the a_i. Since Principal Component Analysis has sorted the a_i by their importance to explaining a typical STF, we may choose to truncate this sum at a relatively low value $N \ll 256$:

20   $$s_N(t) = \sum_{i=1}^{N} a_i s_i(t) + s_0(t), \tag{3}$$

In practice, N will be chosen based on the residual misfit between $s(t)$ and $s_N(t)$ that one is willing to tolerate. Figure 2b shows the dependence of this misfit on N. If we tolerate an average misfit of 5 % in total signal variance, $N = 12$–14 basis functions are sufficient. In the following we use $N = 12$.

1131

2.2 Parametrisation of the moment tensor

The orientation of the source can either be parametrised by a moment tensor using 6 parameters, or as a pure shear displacement source (Aki and Richards, 2002, p. 112) with strike, slip and dip (to which a term for an isotropic component might be added).

Here we want to estimate the non-double-couple content of the solutions, and hence we sample the full moment tensor. The scalar moment is fixed to 1, so that only relative M_j are estimated.

This is equivalent to sampling a hypersphere in the six-dimensional vector space $\{M_{xx}, M_{yy}, M_{zz}, M_{xy}, M_{yz}, M_{xz}\}$ with

$$M_0 = \sqrt{(M_{xx} + M_{yy} + M_{zz})^2 + 2(M_{xy} + M_{yz} + M_{xz})^2} = 1. \tag{4}$$

Uniform sampling on a n-D hypersphere can be achieved by the method of Tashiro (1977), which transforms $n-1$ uniformly distributed random variables x_i to produce n random variables r_i that are distributed uniformly on a hypersphere with $\sqrt{\sum_{i=1}^{6} r_i^2} = 1$. We identify r_i with the moment tensor components and note that the non-diagonal elements $M_{kl}, k \neq l$ appear twice in the sum (thus we actually sample an ellipsoid rather than a hypersphere). We then have

$$x_i \sim U(0,1), \quad i = 1,2,\ldots,5$$
$$Y_3 = 1; \qquad Y_2 = \sqrt{x_2}; \qquad Y_1 = Y_2 x_1$$
$$M_{xx}/M_0 = \sqrt{Y_1} \cdot \cos(2\pi x_3)$$
$$M_{yy}/M_0 = \sqrt{Y_1} \cdot \sin(2\pi x_3) \tag{5}$$
$$M_{zz}/M_0 = \sqrt{Y_2 - Y_1} \cdot \cos(2\pi x_4)$$
$$M_{xy}/M_0 = \sqrt{Y_2 - Y_1} \cdot \sin(2\pi x_4)/\sqrt{2}$$
$$M_{yz}/M_0 = \sqrt{Y_3 - Y_2} \cdot \cos(2\pi x_5)/\sqrt{2}$$
$$M_{zx}/M_0 = \sqrt{Y_3 - Y_2} \cdot \sin(2\pi x_5)/\sqrt{2}$$

Title Page

Abstract | Introduction

Conclusions | References

Tables | Figures

◄ | ►

◄ | ►

Back | Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

2.3 Forward simulation

Broadband, teleseismic Green's functions for P-pP-sP and S-sS wave trains are calculated by the WKBJ code of Chapman (1978), using IASP91 (Kennett and Engdahl, 1991) as spherically symmetric reference model for the mantle. The reference crust at the receiver site is replaced by the two-layered crust predicted by model CRUST2.0 (Bassin et al., 2000). Values for intrinsic attenuation are taken from the spherically symmetric earth model PREM (Dziewoński, 1981). The synthetic waveforms are compared to the observed seismograms in time windows that start $10\,\mathrm{s}$ before the theoretical P wave arrival time (according to IASP91) and end $41.2\,\mathrm{s}$ after.

3 Source parameter estimation by Bayesian sampling

3.1 Bayesian inversion

Bayesian Inversion is an application of Bayes' Rule

$$P(\boldsymbol{m}|\boldsymbol{d}) = \frac{P(\boldsymbol{d}|\boldsymbol{m})P(\boldsymbol{m})}{P(\boldsymbol{d})}, \tag{6}$$

where $\boldsymbol{m}$ is a vector of model parameters (in our case depth, moment tensor elements M_j and STF weights a_i), and $\boldsymbol{d}$ is a vector of data, i.e., a concatenation of P and SH waveforms. These quantities are considered to be random variables that follow Bayes' Rule. We can then identify $P(\boldsymbol{m})$ with the prior probability density of a model. This is the information on the model parameters that we have independent of the experiment. The conditional probability of $\boldsymbol{d}$ given $\boldsymbol{m}$, $P(\boldsymbol{d}|\boldsymbol{m})$, also called $\mathcal{L}(\boldsymbol{m}|\boldsymbol{d})$, is the *likelihood* of a model $\boldsymbol{m}$ to produce the data $\boldsymbol{d}$. Term $P(\boldsymbol{d})$ is constant for all models and is therefore dropped in what follows. $P(\boldsymbol{m}|\boldsymbol{d})$ is called the posterior probability density (short, "the posterior"), and denotes the probability assigned to a model $\boldsymbol{m}$ after having

Title Page

Abstract | Introduction

Conclusions | References

Tables | Figures

◄◄ | ►►

◄ | ►

Back | Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

done the experiment.

$$P(\boldsymbol{m}|\boldsymbol{d}) = P(\boldsymbol{m})\mathcal{L}(\boldsymbol{m}|\boldsymbol{d})k \tag{7}$$

Since the posterior $P(\boldsymbol{m}|\boldsymbol{d})$ may vary by orders of magnitude for different $\boldsymbol{d}$, we work with its logarithm. We introduce the quantity $\Phi(\boldsymbol{m})$ to denote some kind of data misfit such that the likelihood can be written as $\mathcal{L}(\boldsymbol{m}) = \exp[-\Phi(\boldsymbol{m}|\boldsymbol{d})]$.

$$\ln\left(P(\boldsymbol{m}|\boldsymbol{d})\right) = -\Phi(\boldsymbol{m}|\boldsymbol{d}) + \ln P(\boldsymbol{m}) - \ln k. \tag{8}$$

The normalisation constant k is

$$k = \int \exp[-\Phi(\boldsymbol{m}|\boldsymbol{d})]P(\boldsymbol{m})d\boldsymbol{m} \tag{9}$$

and calculated by the Neighbourhood Algorithm in the ensemble inference stage.

In case of Gaussian-distributed noise on the data with a covariance matrix $\mathbf{S}_D$

$$\boldsymbol{d} = g(\boldsymbol{m}) + \epsilon, \quad \epsilon \sim \mathcal{N}(0, \mathbf{S}_D) \tag{10}$$

where $g(\boldsymbol{m})$ is the data predicted by model $\boldsymbol{m}$, we would obtain the familiar expression

$$\Phi(\boldsymbol{m}|\boldsymbol{d}) = k'\left(\frac{1}{2}(\boldsymbol{d} - g(\boldsymbol{m}))^T\mathbf{S}_D^{-1}(\boldsymbol{d} - g(\boldsymbol{m}))\right). \tag{11}$$

This term is usually called Mahalanobis distance or ℓ^2 misfit. We do not choose this sample-wise difference between observed and predicted waveforms as our measure of misfit, since for tomography we do not use it either. There are questions about the Gaussian noise assumption for real data, but mainly we consider there to be a measure that is more robust and adapted to our purpose, the cross-correlation (mis-)fit between data and synthetics (Sigloch and Nolet, 2006; Nolet, 2008), which essentially quantifies phase misfit. In the optimisation-based, linearised approach to tomography,

fitting the phase shift between two waveforms remains a near-linear problem in a wider range around the reference model than fitting the waveforms sample-wise. The cross-correlation fit is defined as:

$$CC(\Delta T_i) = \int_t \left(u_i^c(t - \Delta T_i) \cdot u_i(t) dt \right) \tag{12}$$

5 where $u_i(t)$ is the measured and $u_i^c(t)$ is the synthetic waveform for a model **m** at station i. In general, CC is a function of the time lag ΔT_i for which we compare the observed and predicted waveforms, but here we imply that ΔT_i has already been chosen such as to maximise $CC(\Delta T_i)$. (This value of ΔT_i that maximises the cross-correlation is called the "finite-frequency travel time anomaly" of waveform $u_i(t)$, and represents the

10 most important observable for finite-frequency tomography (Nolet, 2008; Sigloch and Nolet, 2006). Section 4.3, which discusses error propagation from source inversion into tomographic observables, should clarify this motivation of the cross-correlation criterion further.)

Correlation $CC(\Delta T_i)$ measures goodness of fit, so we choose decorrelation $D_i = $

15 $1 - CC(\Delta T_i)$ as our measure of misfit (one scalar per wavepath i). From the large set of pre-existing deterministic source solutions described in Sect. 2.1, we estimated the distribution of this misfit D_i, based on our reference data set of about 1000 very confidently deconvolved STF solutions. For this large and highly quality-controlled set of earthquakes, we empirically find that the decorrelation D_i of its associated seismo-

20 grams $u_i(t) - u_i^c(t)$ follows a log-normal distribution in the presence of the actual noise and modelling errors. The statistics of this finding are discussed further in the companion paper (Stähler et al., 2013), but here we use it to state our likelihood function $\mathcal{L}$, which is the multivariate log-normal distribution:

$$\mathcal{L} = \frac{\exp\left(-\frac{1}{2}(\ln(\mathbf{D}) - \boldsymbol{\mu})^T \mathbf{S}_D^{-1} (\ln(\mathbf{D}) - \boldsymbol{\mu})\right)}{(2\pi)^{\frac{n}{2}} \sqrt{|\det(\mathbf{S}_D)|}} \tag{13}$$

Title Page

Abstract | Introduction

Conclusions | References

Tables | Figures

◄ | ►

◄ | ►

Back | Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

D is the decorrelation vector into which n decorrelation coefficients D_i are gathered. Each D_i was measured on a pair of observed/predicted broadband waveforms that contained either a P or an SH arrival. The parameters of this multivariate log-normal distribution are its mean vector $\boldsymbol{\mu}$ containing n means μ_i and its covariance matrix $\mathbf{S}_D$.

5 Empirically we find that the μ_i and the standard deviations σ_i (diagonal elements of $\mathbf{S}_D$) depend mainly on the signal-to-noise-ratio (SNR) of waveform u_i. The data covariance between two stations i and j (off-diagonal elements in $\mathbf{S}_D$) is predominantly a function of the distance between station i and station j. We estimate their values from the data set of the 1000 trustworthy STF solutions, i.e., from prior information, and proceed to

10 use these $\boldsymbol{\mu}$ and $\mathbf{S}_D$ in our Bayesian source inversions.

It follows from Eq. (13) that the misfit Φ is

$$\Phi = \frac{1}{2}\left(\sum_i^n \sum_j^m \left(\ln(D_j) - \mu_j\right)^T S_{D,ij}^{-1}\left(\ln(D_j) - \mu_j\right)\right) + \frac{1}{2}\ln\left((2\pi)^n |\det(\mathbf{S}_D)|\right) \tag{14}$$

3.2 Construction of the prior probability density

A crucial step in Bayesian inference is the selection of prior probabilities $P(\boldsymbol{m})$ on the

15 model parameters $\boldsymbol{m}$. Our model parameters are:

- m_1: source depth. We assume a uniform prior based on the assumed depth of the event in the NEIC catalogue. If the event is shallow according to the ISC catalogue ($< 30\,\text{km}$), we draw from depths between $0\,\text{km}$ and $50\,\text{km}$, i.e., $m_1 \sim \mathcal{U}(0,50)$. For deeper events, we draw from depths between $20\,\text{km}$ and $100\,\text{km}$. Events deeper

20   than $100\,\text{km}$ have to be treated separately.

- $m_2, \ldots, m_{13} = a_1, \ldots, a_{12}$, the weights of the source time function (Eq. 3). The samples are chosen from uniform distributions with ranges shown in Table 1, but are subjected to a prior π_{STF} (see below).

- $m_{14}, \ldots, m_{18} = x_1, \ldots x_5$, the constructor variables for the moment tensor (Eq. 5).

25    $x_i \sim \mathcal{U}(0,1)$, but they are subjected to two priors π_{iso} and π_{CLVD} (see below).

An earthquake is caused by the release of stress that had built up along a fault, driven by shear motion in the underlying, viscously flowing mantle. Hence the rupture is expected to proceed in only one direction, the direction that releases the stress. The source time function is defined as the time-derivative of the moment, $s(t) = \dot{m}(t)$. The moment is proportional to the stress and thus monotonous, and hence $s(t)$ should be non-negative. In practice, an estimated STF is often not completely non-negative (unless this characteristic was strictly enforced). The reason for smaller amounts of "negative energy" (time samples with negative values) in the STF include reverberations at heterogeneities close to the source, which produce systematic oscillations that are present in most or all of the observed seismograms. Motivated by waveform tomography, our primary aim is to fit predicted to observed waveforms. If a moderately non-negative STF produces better-fitting synthetics, then our pragmatic approach is to accept it, since we are not interested in source physics per se. However, we still need to moderately penalise non-negative samples in the STF, because otherwise they creep in unduly when the problem is underconstrained, due to poor azimuthal receiver coverage. In such cases, severely negative STFs often produce marginally better fits by fitting the noise.

Our approach is to punish slightly non-negative STF estimates only slightly, but to severely increase the penalty once the fraction of "negative energy" I exceeds a certain threshold I_0. To quantify this, we define I as the squared negative part of the STF divided by the entire STF squared:

$$I = \frac{\int_0^T s_N(t)^2 \cdot \Theta(-s_N(t))\mathrm{d}t}{\int_0^T s_N(t)^2}, \quad \text{where} \tag{15}$$

$$s_N = s_0(t) + \sum_{i=1}^{N} a_i s_i(t) \tag{16}$$

and Θ is the Heaviside function. Based on I, we define a prior π_{STF}:

$$\pi_{\mathrm{STF}}(m_2, \ldots, m_{13}) = \exp\left[-\left(\frac{I}{I_0}\right)^3\right], \tag{17}$$

where the third power and $I_0 = 0.1$ have been found to work best. In other words, we do not mind if up to 10 % of STF variance is contributed by negative samples (mostly oscillations), but we do not tolerate more negative samples than that.

The Neighbourhood Algorithm supports only uniform distributions on parameters. The introduction of π_{STF} defined by Eq. (17) leads to a certain inefficiency, in that parts of the model space are sampled that are essentially ruled out by the prior. We carefully selected the ranges of the a_i by examining their distributions for the 3000 catalogue solutions. A test was to count which fraction of random models were consistent with $I < 0.1$. For the ranges given in Table 1, we found that roughly 10 % of the random STF estimates had $I < 0.1$.

A second prior constraint is that earthquakes caused by stress release on a fault should involve no volume change, meaning that the isotropic component $M_{\mathrm{iso}} = M_{xx} + M_{yy} + M_{zz}$ of the moment tensor should vanish. Hence we introduce another prior constraint

$$\pi_{\mathrm{iso}}(m_{14}, \ldots, m_{18}) = \exp\left[-\left(\frac{M_{\mathrm{iso}}/M_0}{\sigma_{\mathrm{iso}}}\right)^3\right] \tag{18}$$

where M_0 is the scalar moment, and $\sigma_{\mathrm{iso}} = 0.1$ is chosen empirically.

Third, we also want to encourage the source to be double-couple-like. A suitable prior is defined on the Compensated Linear Vector Dipole (CLVD) content, which is the ratio $\epsilon = |\lambda_3|/|\lambda_1|$ between smallest and largest deviatoric eigenvalues of the moment tensor:

$$\pi_{\mathrm{CLVD}}(m_{14}, \ldots, m_{18}) = \exp\left[-\left(\frac{\epsilon}{\sigma_{\mathrm{CLVD}}}\right)^3\right] \tag{19}$$

A moment tensor with $\epsilon = 0.5$ corresponds to a purely CLVD source, while $\epsilon = 0$ is a pure DC source. Again we have to decide on a sensible value for the characteristic constant σ_{CLVD}. For intermediate-sized earthquakes of the kind we are interested in, $\epsilon \leq 0.2$ seems to be a reasonable value (Kuge and Lay, 1994).

The total prior probability density is then

$$P(\boldsymbol{m}) = \pi_{\mathrm{STF}}(m_2, \ldots, m_{13}) \qquad\qquad (20)$$
$$+ \pi_{\mathrm{iso}}(m_{14}, \ldots, m_{18}) + \pi_{\mathrm{CLVD}}(m_{14}, \ldots, m_{18})$$

3.3 Sampling with the Neighbourhood Algorithm

Our efficient wavelet parametrisation of the STF reduces the total number of model parameters to around 18, but sampling this space remains non-trivial. A 18-dimensional model space with a multimodal fitness landscape is poorly suited to the popular Metropolis Hastings-Algorithm (Hastings, 1970), which has large autocorrelation times and a tendency to get stuck in local minima. Better suited and faster is a Gibbs sampler, but it needs to know the conditional distribution $p(x_j | x_1, \ldots x_{j-1}, x_{j+1}, x_n)$ along parameter x_j in the n-dimensional model space (Geman and Geman, 1984). This conditional distribution is usually not available, especially not for nonlinear inverse problems.

To overcome the problem of navigation in complex high-dimensional model spaces, the Neighbourhood Algorithm uses Voronoi cells (Sambridge, 1998) to approximate a map of the misfit landscape (Sambridge, 1999a, first stage), followed by a Gibbs sampler to appraise an ensemble based on this map (Sambridge, 1999b, second stage).

In order to point the map-making first stage of the NA into the direction of a priori allowed models, we use a pre-calculated set of starting models. For that, the NA is run without forward simulations and without calculating the likelihood, so that only a map of the prior landscape is produced, from 32 768 samples (Fig. 3a). The resulting 32 768 Voronoi cells are used as a starting set to produce a map of the posterior landscape. This means that from the start, the map will be more detailed in a priori favourable

Full Screen / Esc

regions, and avoids that the algorithm will waste too much time refining the map in regions that are essentially ruled out by the prior.

Next, the prior landscape is loaded and a forward simulation is run for each member in order to evaluate its posterior probability. Then this map is further refined by 512 forward simulations around the 128 best models. This is repeated until a total of 65 536 models have been evaluated.

In the second stage of the NA, which is the sampling stage, 400 000 ensemble members are drawn according to the posterior landscape from the first step. This process runs on a 16 core Xeon machine and takes around 2 h in total per earthquake.

4   A fully worked example

4.1   23 August 2011 Virginia earthquake

In the following we present a fully worked example for a Bayesian source inversion, by applying our software to the M_W 5.7 earthquake that occurred in Central Virginia on 23 August 2011 (Figs. 4 and 5, also compare to Fig. 1). While not considered a typical earthquake region, events from this area have nevertheless been recorded since the early days of quantitative seismology (Taber, 1913). Due to its occurrence in a relatively unusual but densely populated area, this relatively small earthquake was studied in considerable detail, affording us the opportunity to compare to results of other workers. Moderate-sized events of this kind are typical for our targeted application of assembling a large catalogue. The greatest abundance of suitable events is found just below magnitude 6; toward smaller magnitudes, the teleseismic signal-to-noise ratio quickly deteriorates below the usable level.

For the inversion, we used a set of 41 P waveforms and 17 SH waveforms recorded by broadband stations at teleseismic distances (Fig. 4). For waveform modelling, the crustal structure of model CRUST 2.0 (Bassin et al., 2000) was assumed around the source region; values are given in Table 2. The algorithm ran 65 536 forward simu-

lations to generate a map of the posterior landscape, and produced an ensemble of 400 000 members in the second step. From this ensemble, the source parameters were estimated. Table 3 shows the estimated credible intervals and the mode of the depth and the moment tensor. These quantiles represent only a tiny part of the information contained in the ensemble, i.e., two statistics of 1-dimensional marginals derived from a 16-dimensional probability density function. Some credible intervals are large, for example we cannot constrain the depth to a range of less than 10 km with 90 % credibility. Using such credible interval estimates, routine production runs of our software should be able to clarify whether depth uncertainties in existing catalogues tend to be overly optimistic or pessimistic. The complete marginal distribution of the source depth estimate is shown in Fig. 3, bottom left.

We aim for additional, informative ways of summarising and conveying the resulting ensemble. Figure 5 is what we call a "Bayesian beach ball": an overlay of 1024 focal mechanisms drawn from the ensemble at random. The thrust faulting character of the event is unambiguous, but the direction of slip is seen to be less well constrained. The estimate of the source time function and its uncertainty are displayed in Fig. 4, bottom right. Within their frequency limits, our teleseismic data prefer a single-pulsed rupture of roughly 3 s duration, with a certain probability of a much smaller foreshock immediately preceding the main event. Smaller aftershocks are possible, but not constrained by our inversion.

4.2 Comparison to source estimates of other workers

Our solution is consistent with the solution from the SCARDEC catalogue (Vallée, 2012), which puts the depth of this event at 9 km, and its STF duration at 2.5 s. Chapman (2013) studied the source process of the 2011 Virginia event in great detail. He argues for three sub-events having occurred within 1.6 s at a depth of 7–8 km, and spaced less than 2 km apart. This is compatible with our solution: since teleseismic waveforms contain little energy above frequencies of 1 Hz, we would not expect to resolve three pulses within 1.6 s with the method presented here. Chapman (2013) used

both local and teleseismic recordings, and was therefore able to exploit high frequencies recorded close to the source. His local crustal model featured an upper crustal velocity that was 50 % higher than ours, which may explain why he estimates the source 1–2 km deeper than our most probable depth of 5.9 km (Fig. 4, bottom left).

4.3 Uncertainty propagation into tomographic observables

We are interested in source estimation primarily because we want to account for the prominent signature of the source wavelet in the broadband waveforms that we use for waveform tomography. Input data for the inversion, primarily traveltime anomalies ΔT_i, where i is the station index, are generated by cross-correlating observed seismograms with predicted ones. A predicted waveform consists of the convolution of a synthetic Green's function with an estimated source time function (Eq. 1). Thus uncertainty in the STF estimate propagates into the cross-correlation measurements that generate our input data for tomography. Previous experience has led us to believe that the source model plays a large role in the uncertainty of ΔT_i. The probabilistic approach presented here permits to quantify this influence by calculating $\Delta T_{i,j}$ for each ensemble member j. From all values for one station, the ensemble mean $\overline{\Delta T_i}$ and its standard deviation σ_i can then be used as input data for the tomographic inversion. Thus we obtain a new and robust observable: Bayesian traveltime anomalies with full uncertainty information.

Figure 6 shows the standard deviation σ_i of P wave ΔT_i at all stations. Comparison to the signal-to-noise ratios of Fig. 6 shows no overall correlation, except for South American stations, where a higher noise level is correlated with a somewhat larger uncertainty on ΔT_i. By contrast, European stations all have good SNR, but uncertainties in the travel times are large nonetheless, because source uncertainty happens to propagate into the estimates of ΔT_i more severely in this geographical region. This information would not have been available in a deterministic source inversion and could strongly affect the results of seismic tomography.

Probabilistic source inversion I – parametrization

S. C. Stähler and
K. Sigloch

Title Page

Abstract | Introduction

Conclusions | References

Tables | Figures

◄ | ►

◄ | ►

Back | Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

5 Discussion

5.1 Performance of the empirical orthogonal basis for STF parametrisation

We choose to parametrise the source time function in terms of empirical orthogonal functions (eofs), which by design is the most efficient parametrisation if the character-istics of the STFs are well known. We think that they are, having semi-automatically deconvolved thousands of STFs in prior work (Sigloch and Nolet, 2006; Sigloch, 2011) and compared them with other studies (Tanioka and Ruff, 1997; Houston, 2001; Tocheport et al., 2007). The flip side of this tailored basis is that it might quickly turn inefficient when atypical STFs are encountered. From the appearance of the eofs in Fig. 2a, it is for example obvious that STFs longer than 20 s could not be well expressed well as a weighted combination of only 10 eofs. Hence the STFs of the strongest earth-quakes considered (around M_W 7.5) might not be fit quite as well as the bulk of smaller events, which contributed more weight to defining the eof base. For our tomography ap-plication, this behaviour is acceptable and even desirable, since the largest events are no more valuable than smaller ones (often quite the opposite, since the point source approximation starts to break down for large events).

At first glance it might seem unintuitive that the basis functions have oscillatory char-acter and thus negative parts, rather than resembling a set of non-negative basis func-tions (a set of triangles would be one such set). Remember however that the training collection to which the principal components analysis was applied did consist of pre-dominantly non-negative functions, which by construction are then represented par-ticularly efficiently, even if the eofs may not give this appearance. On top of this, we explicitly encourage non-negativity of the solution via the prior π_{STF} (Eq. 17). A rough estimation showed that roughly 90 % of the model space are "forbidden" by the condi-tion that the source should have a vanishing negative part.

We wanted to know how many basis functions of a more generic basis (e.g., wavelets) would be required in order to approximate the STF collection equally well as with the eofs. A trial with a basis of sinc-wavelets showed that 16 basis functions

were needed to achieve the same residual misfit as delivered by our optimised basis of only 10 eofs. Since the size of the model space grows exponentially with the number of parameters, avoiding 6 additional parameters makes a big difference in terms of sampling efficiency.

5.2 Moment tensor parametrisation

The parametrisation of the moment tensor is a technically non-trivial point. We discuss the pros and cons of possible alternatives to our chosen solution:

- Parameterisation in terms of strike ϕ_f, slip λ and dip δ is problematic for sampling. Strike and dip describe the orientation of the fault plane; an equivalent description would be the unit normal vector $\boldsymbol{n}$ on the fault.

$$\boldsymbol{n} = \begin{pmatrix} c - \sin\delta\sin\phi_f \\ -\sin\delta\cos\phi_f \\ \cos\delta \end{pmatrix} \tag{21}$$

All possible normal vectors form a unit sphere. In order to sample uniformly on this unit sphere, samples have to be drawn from a uniform volumetric density (Tarantola, 2005, Sect. 6.1). Since the Neighbourhood Algorithm (and most other sampling algorithms) implicitly assume Cartesian coordinates in the model space, the prior density has to be multiplied by the Jacobian of the transformation into the actual coordinate system, in our case $1/\sin\delta$. To our knowledge, this consideration is neglected in most model space studies, but it would be more severe in ensemble sampling than in gradient-based optimisation.

- A different issue with strike-dip parametrisation is the following: the Euclidean distances applied to $\{\phi_f, \lambda, \delta\}$ by the NA and similar, Cartesian-based algorithms are in fact a rather poor measure of the similarity of two double couple sources. A more suitable measure of misfit is the Kagan angle (Kagan, 1991) or the Tape measure (Tape and Tape, 2012), which is the smallest angle required to rotate

the principal axes of one double couple into the corresponding principal axes of the other.

This is an issue in model optimisation with the first stage of the Neighbourhood Algorithm (Kennett et al., 2000; Sambridge and Kennett, 2001; Vallée et al., 2011). Wathelet (2008) has introduced complex boundaries to the NA, but unfortunately no periodic ones.

– An alternative would be to sample $\{M_{xx}, M_{yy}, M_{zz}, M_{xy}, M_{yz}, M_{xz}\}$ independently, but this is inefficient because the range of physically sensible parameters spans several orders of magnitude.

– Finally, one might choose not to sample the moment tensor at all. Instead, one might sample only from the $\{S_i, d\}$ model space, followed by direct, linear inversion of the six moment tensor elements corresponding to each sample. This would speed up the sampling considerably since the dimensionality of the model space would be reduced from 16 to 10. Moment tensor inversion is a linear problem (Eq. 1), and hence we would not lose much information about uncertainties. In a potential downside, moment tensor inversion can be unstable in presence of noise or bad stations, but from our experience with supervised, linear inversions, this is typically not a severe problem in practice. Therefore we are considering this pragmatic approach of reduced dimensionality for production runs.

5.3 Neighbourhood Algorithm

The Neighbourhood Algorithm avoids some of the pitfalls of other sampling algorithms. It does not suffer from the same huge autocorrelation times as the Metropolis-Hastings Algorithm, and it uses fewer tuning parameters than Genetic Algorithms or Simulated Annealing. (Note that Genetic Algorithms and Simulated Annealing are not sampling algorithms. Like the first stage of the NA, they only explore the model space for a best-fitting solution. Their results however might be used as input for the second stage of

Title Page

Abstract | Introduction

Conclusions | References

Tables | Figures

◀ | ▶

◀ | ▶

Back | Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

the NA. We are not aware of any examples.) The NA uses only two tuning parameters, which control (a) how many new models are generated in each step and (b) in how many different cells these models are generated. As in every optimisation algorithm, they control the tradeoff between exploration of the model space, and exploitation of the region around the best models.

There is no hard and fast rule for choosing values for these tuning parameters. Since we do not want to optimise for only one "best" solution, we tend towards an explorative strategy and try to map large parts of the model space. Compared to other source inversion schemes, we are explicitly interested in local minima in the misfit landscape. Local minima are often seen as nuisance, especially in the rather aggressive iterative optimisation frameworks, but in our view, they contain valuable information. What may appear as a local minimum to the specific data that we are using for inversion, might turn out to be the preferred solution of another source inversion method (e.g., surface waves, GPS or InSAR).

However, an ensemble that does not resolve the best fitting model is equally useless. The posterior of all models gets normalised after all forward simulations have been done (see Eq. 9). If one peak (the best solution) is missing, the normalisation constant Z will be too small, and therefore $P(\mathbf{m}|\mathbf{d})$ will be too high for all models, meaning that the credibility bounds will be too large. It is possible that other sampling schemes, such as *parallel tempering*, might find better compromises between exploration and exploitation, which could be a topic of further study.

5.4 Comparison with other source inversion schemes

Table 4 shows a list of other point source inversion algorithms proposed and applied over the past 15 yr. Most widely used is probably the Global Centroid Moment Tensor (CMT) catalogue (Dziewoński et al., 1981; Ekström et al., 2012), which is mostly based on intermediate-period (> 40 s) waveforms to determine a centroid moment tensor solution. Its results are less applicable to short-period body wave studies, since waveforms in the latter are dominated by the hypocenter, which may differ significantly

from the centroid. Another classical catalogue is the ISC bulletin (Bondár and Storchak, 2011), which goes back as long as 1960. The ISC catalogue focuses on estimating event times and locations, neither of which are the topic of this study. The ISC recently adopted a global search scheme based on the first stage of the NA, similar to Sambridge and Kennett (2001), followed by an attempt to refine the result by linearised inversion, including inter-station covariances. Garcia et al. (2013) and Tocheport et al. (2007) use Simulated Annealing to infer depth and moment tensor. A STF is estimated from the P waveforms. By neglecting all crustal contributions and reducing the forward simulation to mantle attenuation, this approach is very efficient.

Similarly, Kolář (2000) used a combination of Simulated Annealing and bootstrapping to estimate uncertainties of the moment tensor, depth and a source time function. The study was limited to two earthquakes and was never extended.

Kennett et al. (2000) used the first stage of the NA to optimise for hypocenter depth, moment tensor, and the duration of a trapezoidal STF, using essentially the same kind of data as the present study, and an advanced reflectivity code for forward modelling. However, no uncertainties were estimated.

Debṣki (2008) is one of the only two studies we are aware of that did Bayesian inference of the source time function. He studied magnitude 3 events in a copper mine in Poland. By using the Empirical Green's Functions (EGF) method, it was not necessary to do an explicit forward simulation. The study was limited to inverting for the STF, which he parametrised sample-wise. This was possible since the forward problem was computationally very inexpensive to solve.

The second sampling study is Wéber (2006), which used an Octree importance sampling algorithm to infer probability density functions for depth and moment tensor rate function. The resulting ensemble was decomposed into focal mechanisms and source time functions, a nontrivial and nonunique problem (Wéber, 2009). With this algorithm, a catalogue of Hungarian seismicity was produced until 2010, but apparently this promising work was not extended to a global context.

Probabilistic source inversion I – parametrization

S. C. Stähler and
K. Sigloch

Title Page

Abstract | Introduction

Conclusions | References

Tables | Figures

◄ | ►►

◄ | ►

Back | Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

The most recent global source catalogue is the SCARDEC method by Vallée et al. (2011). It uses the first stage of the Neighbourhood Algorithm to optimise the parameters source depth, strike, dip and rake. For each model and each station, a *Relative Source Time Function* is calculated. The misfit is comprised of a waveform misfit and ⁵ the differences between the RSTF at different stations. Uncertainties of the parameters are estimated by the variation of the misfit along different parameters.

The PRISM algorithm as presented here is the first to enable Bayesian inference of seismic source parameters on a global scale and in a flexible framework. It allows for sampling of the source time function by a set of optimised, wavelet-like basis functions. ¹⁰ By producing a whole ensemble of solutions, arbitrary parameters, like the uncertainty of travel time misfits can be estimated from the ensemble afterwards, at little additional cost.

6 Conclusions

We showed that routine Bayesian inference of source parameters from teleseismic ¹⁵ body waves is possible and provides valuable insights. From clearly stated a priori assumptions, followed by data assimilation, we obtain rigorous uncertainty estimates of the model parameters. The resulting ensemble of a posteriori plausible solutions permits to estimate the propagation of uncertainties from the source inversion to other observables of practical interest to us, such as travel time anomalies for seismic to- ²⁰ mography.

were done on the computational facilities of the Sektion Geophysik at LMU and the Leibniz-Rechenzentrum in Garching.

References

Aki, K. and Richards, P. G.: Quantitative Seismology, vol. II, University Science Books, 2002. 1132

Bassin, C., Laske, G., and Masters, G.: The Current Limits of Resolution for Surface Wave Tomography in North America, in: EOS Trans AGU, vol. 81, F897 pp., 2000. 1133, 1140

Bondár, I. and Storchak, D. A.: Improved location procedures at the International Seismological Centre, Geophys. J. Int., 186, 1220–1244, 2011. 1147, 1156

Chapman, C. H.: A new method for computing synthetic seismograms, Geophys. J. Roy. Astr. S., 54, 481–518, 1978. 1133

Chapman, M. C.: On the rupture process of the 23 August 2011 Virginia earthquake, B. Seismol. Soc. Am., 103, 613–628, 2013. 1141

Debski, W.: Estimating the earthquake source time function by Markov Chain Monte Carlo Sampling, Pure Appl. Geophys., 165, 1263–1287, 2008. 1128, 1147, 1156

Dziewoński, A. M.: Preliminary reference Earth model, Phys. Earth Planet. Int., 25, 297–356, 1981. 1133

Dziewoński, A. M., Chou, T.-A., and Woodhouse, J. H.: Determination of earthquake source parameters from waveform data for studies of global and regional seismicity, J. Geophys. Res., 86, 2825–2852, 1981. 1146

Ekström, G., Nettles, M., and Dziewoński, A. M.: The global CMT project 2004–2010: centroid-moment tensors for 13,017 earthquakes, Phys. Earth Planet. Int., 200–201, 1–9, 2012. 1146, 1156

Garcia, R. F., Schardong, L., and Chevrot, S.: A nonlinear method to estimate source parameters, amplitude, and travel times of teleseismic body waves, B. Seismol. Soc. Am., 103, 268–282, 2013. 1127, 1147, 1156

Geman, S. and Geman, D.: Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images, IEEE T. Pattern Anal., 6, 721–741, 1984. 1139

Hastings, W.: Monte Carlo sampling methods using Markov Chains and their applications, Biometrika, 57, 97–109, 1970. 1139

Houston, H.: Influence of depth, focal mechanism, and tectonic setting on the shape and duration of earthquake source time functions, J. Geophys. Res., 106, 11137–11150, 2001. 1127, 1130, 1143

Kagan, Y.: 3-D rotation of double-couple earthquake sources, Geophys. J. Int., 106, 709–716, 1991. 1144

Kennett, B. L. N. and Engdahl, E. R.: Traveltimes for global earthquake location and phase identification, Geophys. J. Int., 105, 429–465, 1991. 1133

Kennett, B. L. N., Marson-Pidgeon, K., and Sambridge, M.: Seismic source characterization using a neighbourhood algorithm, Geophys. Res. Lett., 27, 3401–3404 doi:10.1029/2000GL011559, 2000. 1145, 1147

Kirkpatrick, S., Gelatt, C. D., and Vecchi, M. P.: Optimization by simulated annealing, Science, 220, 671–680, 1983. 1127, 1156

Kolář, P.: Two attempts of study of seismic source from teleseismic data by simulated annealing non-linear inversion, J. Seismol., 4, 197–213, 2000. 1147, 1156

Kuge, K. and Lay, T.: Data-dependent non-double-couple components of shallow earthquake source mechanisms: effects of waveform inversion instability, Geophys. Res. Lett., 21, 9–12, 1994. 1139

Marson-Pidgeon, K. and Kennett, B. L. N.: Source depth and mechanism inversion at teleseismic distances using a neighborhood algorithm, B. Seismol. Soc. Am., 90, 1369–1383, 2000. 1156

Nolet, G.: A Breviary of Seismic Tomography: Imaging the Interior of the Earth and Sun, Cambridge University Press, 2008. 1134, 1135

Ruff, L.: Multi-trace deconvolution with unknown trace scale factors: omnilinear inversion of P and S waves for source time functions, Geophys. Res. Lett., 16, 1043–1046, 1989. 1127

Sambridge, M.: Exploring multidimensional landscapes without a map, Inverse Probl., 14, 427–440 doi:10.1088/0266-5611/14/3/005, 1998. 1139

Sambridge, M.: Geophysical inversion with a neighbourhood algorithm – I. Searching a parameter space, Geophys. J. Int., 138, 479–494,1999a. 1127, 1128, 1139

Sambridge, M.: Geophysical inversion with a neighbourhood algorithm – II. Appraising the ensemble, Geophys. J. Int., 138, 727–746, 1999b. 1128, 1139, 1156

Sambridge, M. and Kennett, B. L. N.: Seismic event location: nonlinear inversion using a neighbourhood algorithm, Pure Appl. Geophys., 158, 241–257, 2001. 1145, 1147, 1156

Probabilistic source inversion I – parametrization

S. C. Stähler and K. Sigloch

Title Page

Abstract | Introduction

Conclusions | References

Tables | Figures

◀ | ▶

◄ | ►

Back | Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Sigloch, K.: Mantle provinces under North America from multifrequency *P* wave tomography, Geochem. Geophy. Geosy., 12, Q02W08, doi:10.1029/2010GC003421, 2011. 1143

Sigloch, K. and Nolet, G.: Measuring finite-frequency body-wave amplitudes and traveltimes, Geophys. J. Int., 167, 271–287, 2006. 1127, 1129, 1130, 1134, 1135, 1143, 1156, 1157

Stähler, S. C., Sigloch, K., and Nissen-Meyer, T.: Triplicated P-wave measurements for waveform tomography of the mantle transition zone, Solid Earth, 3, 339–354, doi:10.5194/se-3-339-2012, 2012. 1127, 1130

Stähler, S. C., Sigloch, K., and Zhang, R.: Probabilistic seismic source inversion II: data misfits and covariances, Solid Earth, in preparation, 2013. 1128, 1135

Taber, S.: Earthquakes in Buckingham County, Virginia, B. Seismol. Soc. Am., 3, 124–133, 1913. 1140

Tanioka, Y. and Ruff, L. J.: Source time functions, Seismol. Res. Lett., 68, 386–400, 1997. 1143

Tape, W. and Tape, C.: Angle between principal axis triples, Geophys. J. Int., 191, 813–831, 2012. 1144

Tarantola, A.: Inverse Problem Theory and Methods for Model Parameter Estimation, SIAM, Philadelphia, 2005. 1144

Tashiro, Y.: On methods for generating uniform random points on the surface of a sphere, Ann. I. Stat. Math., 29, 295–300, 1977. 1132

Tocheport, A., Rivera, L., and Chevrot, S.: A systematic study of source time functions and moment tensors of intermediate and deep earthquakes, J. Geophys. Res., 112, 1–22, 2007. 1143, 1147, 1156

Vallée, M.: SCARDEC solution for the 23/08/2011 Virginia earthquake, available at: http://www.geoazur.net/scardec/Results/Previous_events_of_year_2011/20110823_175103_VIRGINIA/carte.jpg, (last access: 23 July 2013), 2012. 1141

Vallée, M., Charléty, J., Ferreira, A. M. G., Delouis, B., and Vergoz, J.: SCARDEC: a new technique for the rapid determination of seismic moment magnitude, focal mechanism and source time functions for large earthquakes using body-wave deconvolution, Geophys. J. Int., 184, 338–358, 2011. 1127, 1145, 1148, 1156

Wathelet, M.: An improved neighborhood algorithm: parameter conditions and dynamic scaling, Geophys. Res. Lett., 35, 1–5, 2008. 1145

Wéber, Z.: Probabilistic local waveform inversion for moment tensor and hypocentral location, Geophys. J. Int., 165, 607–621, 2006. 1128, 1147, 1156

Probabilistic source inversion I – parametrization

S. C. Stähler and K. Sigloch

Title Page

Abstract | Introduction

Conclusions | References

Tables | Figures

◀ | ▶

◀ | ▶

Back | Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Wéber, Z.: Estimating source time function and moment tensor from moment tensor rate functions by constrained L 1 norm minimization, Geophys. J. Int., 178, 889–900, 2009. 1147

Probabilistic source inversion I – parametrization

S. C. Stähler and
K. Sigloch

S. C. Stähler and
K. Sigloch

Table 1. Sampling of the prior probability distribution: range of STF weights a_i that are permitted in the first stage of the Neighbourhood Algorithm.

i	range	i	range	i	range
1	±1.5	6	±0.7	11	±0.4
2	±1.0	7	±0.7	12	±0.4
3	±0.9	8	±0.6	13	±0.3
4	±0.8	9	±0.5	14	±0.3
5	±0.8	10	±0.5	15	±0.3

Table 2. Crustal model assumed for the source region of the 2011 Virginia earthquake (CRUST2.0).

	V_P	V_S	ρ	depth
Upper crust:	$4.10\,\mathrm{km\,s^{-1}}$	$2.15\,\mathrm{km\,s^{-1}}$	$2.51\,\mathrm{Mg\,m^{-3}}$	$10.5\,\mathrm{km}$
Lower crust:	$6.89\,\mathrm{km\,s^{-1}}$	$3.84\,\mathrm{km\,s^{-1}}$	$2.98\,\mathrm{Mg\,m^{-3}}$	$24.5\,\mathrm{km}$

Table 3. Credible intervals for source parameters of the Virginia earthquake. The moment tensor components M_{kl} need to be multiplied by 10^{16} Nm.

Credible intervals for source parameters:			
	1st decile	median	9th decile
depth	1.8	5.9	11
M_{W}	5.57	5.67	5.74
M_{tt}	−0.233	1.38	2.54
M_{tp}	−1.99	−0.955	−0.165
M_{rt}	−2.7	−0.325	2.72
M_{pp}	−9.4	−4.74	−2.7
M_{rp}	−3.25	−0.563	1.87
M_{rr}	3.16	4.42	7.84

Table 4. Overview of similar source inversion algorithms

Name	characteristics				Inversion parameters					algorithms	
	Focus	Probabilistic	Depth range	Data	Catalogue	Depth	Location	Moment tensor	STF	Forward algorithm and model	Inversion algorithm
PRISM (this paper)	global	yes	full	waveforms, P, SH, teleseismic	yes	yes	no	yes	yes	reflectivity, IASP91 + crust	NA[a], both stages
Tocheport et al. (2007)	global	no	> 100 km	waveforms, P, teleseismic	no	no	no	yes	yes	none	SA[b]
Garcia et al. (2013)	global	no	full	waveforms, P, teleseismic	no	no	no	yes	yes	none	SA
Marson-Pidgeon and Kennett (2000)	global	no	full	waveforms, P, SH, SV	no	yes	no	yes	duration	reflectivity, ak135 + crust	NA, first stage
Sambridge and Kennett (2001)	global	no	full	traveltimes, P, S	no	yes	yes	no	no	ak135	NA, first stage
ISC (Bondár and Storchak, 2011)	global	no	full	traveltimes, all phases	yes	yes[c]	yes	no	no	ak135	NA, first stage
Global CMT, Ekström et al. (2012)	global	no	full	waveforms, P, S + surface	yes	yes	yes	yes	no	normal modes	
Sigloch and Nolet (2006)	global	no	full	waveforms, P, teleseismic	no	yes	no	yes	yes	WKBJ, IASP91	LSQR, iterative
Kolář (2000)	global	uncertainties	full	waveforms, P, SH, teleseismic	no	yes	yes	strike, slip, dip	yes	reflectivity (?), local model	SA + bootstrapping
SCARDEC, Vallée et al. (2011)	global	uncertainties	full	waveforms, P, SH, teleseismic	yes, fast	yes	no	strike, slip, dip	RSTF[d]	reflectivity, IASP91 + crust	NA, first stage
Wéber (2006)	local	yes	shallow	waveforms, P, local	no	yes	yes	yes	MTRF[e]	reflectivity, local model	octree importance sampling
Debski (2008)	local	yes	shallow	waveforms, P, local	no	no	no	no	yes	EGF[f]	Metropolis-Hastings

[a] Neighbourhood Algorithm Sambridge (1999b).
[b] Simulated Annealing Kirkpatrick et al. (1983).
[c] Binning allowed.
[d] Relative STF, one STF per station.
[e] Moment tensor rate function, one STF per MT component.
[f] Empirical Green's functions.

Title Page

Abstract | Introduction

Conclusions | References

Tables | Figures

◄◄ | ►►

◄ | ►

Back | Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Full Screen / Esc
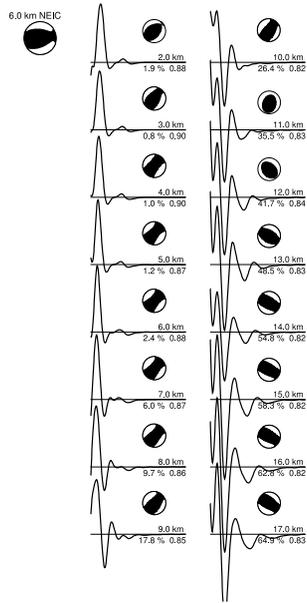
Fig. 1. Source time function solutions for a M_W5.7 earthquake in Virginia, USA (23 August 2011) obtained from joint inversion for STF and moment tensor **M**, using the iterative linearised optimisation algorithm of Sigloch and Nolet (2006). Trial source depths ranged from 2 km to 17 km, in increments of 1 km, and each deconvolution was based on the same 86 broadband, teleseismic P waveforms. Note the strong changes in STF and moment tensor as a function of depth. Top left shows the moment tensor solution from the NEIC catalogue for comparison. For every candidate solution, the percentage of "non-negative" energy is given, a proxy for how oscillatory (and thus inherently non-physical) the solution is. The third number gives the average cross-correlation coefficient between observed and predicted waveforms achieved by each solution. At depths between 2 and 7 km, the STF is pulse-like, simple, non-negative, and waveform cross-correlation attains its maximum, signalling the most likely depth range for this event. The present study offers an approach to quantify these qualitative tradeoffs and judgements.
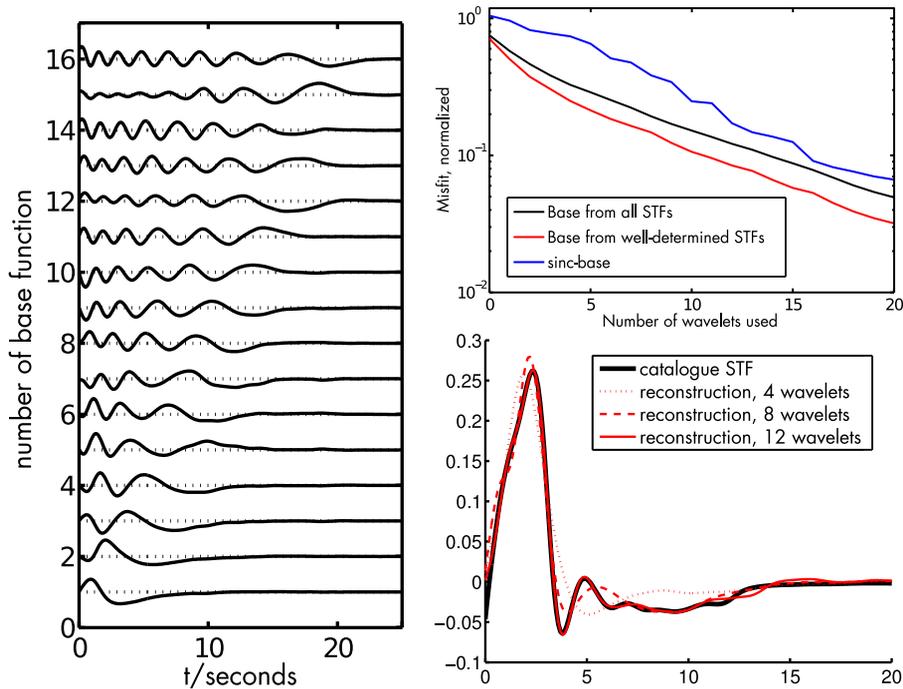
Fig. 2. Efficient parametrisation of the STF in terms of empirical orthogonal functions, computed from a large set of manually deconvolved STFs that effectively serve as prior information. **(a)** First 16 members of the basis of empirical orthogonal functions. **(b)** Median RMS misfit between members of the prior STF catalogue and their projection on a subspace of the model space spanned by the first wavelet basis functions. **(c)** A typical STF from the catalogue, and its projection onto several subspaces spanned by the first few basis functions ($N = [4, 8, 12]$).

Title Page

Abstract | Introduction

Conclusions | References

Tables | Figures

|◄ | ►|

◄ | ►

Back | Close

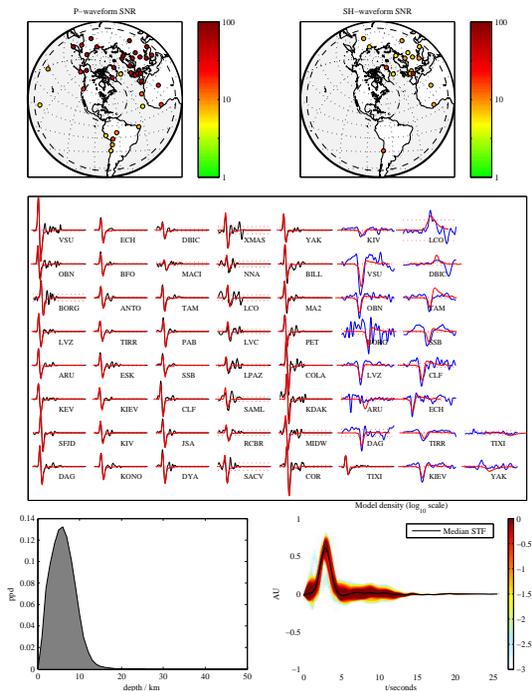Full Screen / Esc

Printer-friendly Version

Interactive Discussion
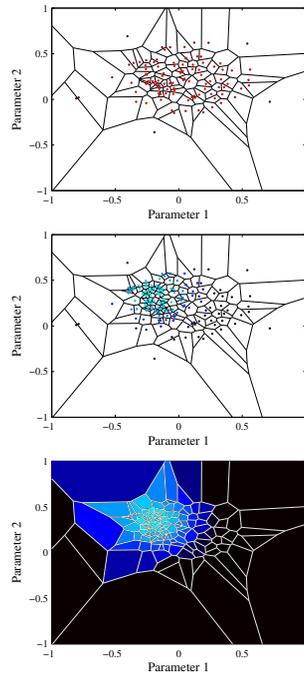
Fig. 3. Principle of the Neighbourhood Algorithm, demonstrated for a two-dimensional toy problem (the underlying distributions are fictional and chosen for demonstration purposes). Top: in the pre-mapping stage, only the prior distribution is evaluated, resulting in a map of starting models that cluster in regions of high prior probability (marked by lighter shades of red). Middle: next, the NA loads this map, evaluates the posterior probability for every sample, and refines the map only in the best fitting Voronoi cells. Lighter shades of blue correspond to a higher posterior probability. Bottom: in the sampling or appraisal stage, the value of the posterior is interpolated to the whole Voronoi cell. The Gibbs sampler uses this map to produce an ensemble. This ensemble can be used to calculate integrals over the model space, like the mean or mode of selected parameters.

Fig. 4. Waveform data and source estimates for the 23 August 2011 Virginia earthquake (M_W 5.7). Top row: distribution of 41 and 17 teleseismic broadband stations that recorded P and S waveforms, respectively. Station colour corresponds to the signal-to-noise ratio in the relevant waveform window. Middle row: synthetic broadband waveforms (red), compared to the data for the best fitting model. Black waveforms are P seismograms, blue waveforms are SH seismograms. The time windows are 51.2 s long and start 5 s before the theoretical phase arrival time. The amplitudes of all P and SH waveforms have been normalised. Bottom left: posterior marginal distribution of estimated source depth. Bottom right: posterior marginal distribution of the source time function. Probability densities are marked by colour and are highest in the areas shaded red.

Title Page

Abstract | Introduction

Conclusions | References

Tables | Figures

|◄ | ►|

◄ | ►

Back | Close

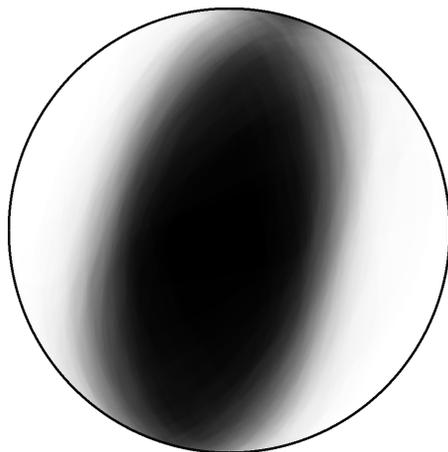Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Fig. 5. Bayesian beach ball: probabilistic display of focal mechanism solutions for the 2011 Virginia earthquake.

Title Page

Abstract | Introduction

Conclusions | References

Tables | Figures

◀ | ▶|

◀ | ▶

Back | Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

P–wave dT standard deviation (s)
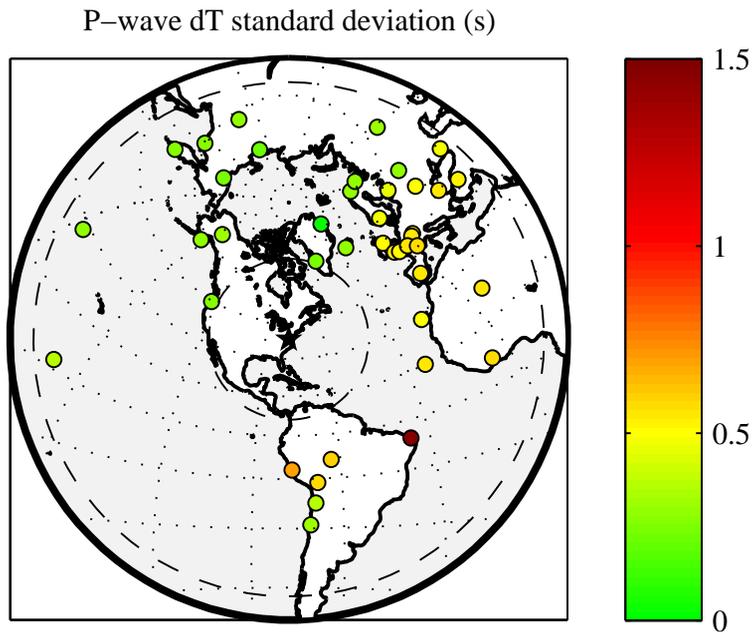
Fig. 6. Standard deviations σ_i of P wave travel times ΔT_i, as calculated from the ensemble of solutions. The travel time estimates are by-products of using waveform cross-correlation as the measure for goodness of fit, and they represent our main input data for tomographic inversions. The unit on the colour scale is seconds.