

Trans-dimensional matched-field geoacoustic inversion with hierarchical error models and interacting Markov chains

Jan Dettmer^{a)} and Stan E. Dosso

School of Earth and Ocean Sciences, University of Victoria, Victoria, British Columbia V8W 3P6, Canada

(Received 4 April 2012; revised 20 July 2012; accepted 23 July 2012)

This paper develops a trans-dimensional approach to matched-field geoacoustic inversion, including interacting Markov chains to improve efficiency and an autoregressive model to account for correlated errors. The trans-dimensional approach and hierarchical seabed model allows inversion without assuming any particular parametrization by relaxing model specification to a range of plausible seabed models (e.g., in this case, the number of sediment layers is an unknown parameter). Data errors are addressed by sampling statistical error-distribution parameters, including correlated errors (covariance), by applying a hierarchical autoregressive error model. The well-known difficulty of low acceptance rates for trans-dimensional jumps is addressed with interacting Markov chains, resulting in a substantial increase in efficiency. The trans-dimensional seabed model and the hierarchical error model relax the degree of prior assumptions required in the inversion, resulting in substantially improved (more realistic) uncertainty estimates and a more automated algorithm. In particular, the approach gives seabed parameter uncertainty estimates that account for uncertainty due to prior model choice (layering and data error statistics). The approach is applied to data measured on a vertical array in the Mediterranean Sea. © 2012 Acoustical Society of America. [http://dx.doi.org/10.1121/1.4746016]

PACS number(s): 43.30.Pc, 43.30.Ma, 43.60.Pt [AIT]

Pages: 2239–2250

I. INTRODUCTION

Matched-field inversion (MFI) can provide information about seabed parameters that is valuable for sonar applications and predicting sound propagation in shallow-water environments. The inversion method has been applied widely to acoustic-field data recorded on vertical- and horizontal-line arrays to obtain seabed parameter estimates (e.g., Refs. 1–3). In recent years, MFI has been addressed with probabilistic inference methods,^{4,5} which sample the posterior probability density (PPD), providing the ability to quantify parameter uncertainty and interrelationships.^{6–8} Quantifying uncertainty is of fundamental importance for inference problems because parameter variability (spatial and/or temporal heterogeneity) must be resolved in the presence of uncertainty due to incomplete and uncertain data observations.

Any probabilistic parameter inference requires the assumption of a *model*, which is defined here in the general sense to include a physics theory describing signal-system interaction, an appropriate parametrization, and a statistical representation for the data errors. It has been shown that MFI requires careful examination of the data errors, which can be strongly dependent (correlated).^{9,10} Correlated errors are generally due to both measurement and theory errors, which are inseparable. Measurement errors are typically randomly distributed and due to various instrument and noise sources. In many inference problems, theory error is due to the intrinsic inability of specifying a model that fully captures the complexity of the system-signal interaction. Theory

errors can cause significant error dependence (correlation) and are often due to insufficient model complexity (the model does not allow data predictions that sufficiently capture the structure of the signal).

Estimating parameter values and uncertainties that meaningfully describe the seabed based on incomplete, noisy acoustic-field data is challenging and strongly tied to model selection: The choice of physics theory (describing signal-sediment interaction) may not be obvious, the sediment parametrization (e.g., layering) consistent with the data resolving power is generally unknown and cannot be constrained to a single prior choice, and an appropriate statistical description of data errors is unknown. Traditional inversion approaches make subjective model assumptions, and the consequences of these assumptions are rarely investigated. However, model choice can profoundly affect the solution and, in particular, uncertainty estimates. Hence model choice should be considered as part of the inverse problem. This approach quantifies the data support for competing models and is referred to as *model selection*. One of the goals of model selection is to identify the simplest parametrization consistent with the data, and Bayes' theorem intrinsically provides such natural parsimony.⁵ Model selection has been applied to MFI with regard to choosing a single optimal parametrization using Bayesian evidence⁷ and the Bayesian information criterion.¹¹

Trans-dimensional (trans-D) inference generalizes model selection^{12,13} by relaxing the requirement of specifying a single preferred or optimal model to specifying groups of appropriate models. Hence the PPD quantifies the state of parameter knowledge for groups of models, accounting for the inability of incomplete and uncertain data to clearly discriminate between models. This results in more rigorous

^{a)}Author to whom correspondence should be addressed. Electronic mail: jand@uvic.ca

seabed parameter uncertainty estimates because the uncertainty of the model is included in the parameter uncertainty estimates. Trans-D inference has been applied to several geophysical inverse problems^{14–18} as well as seabed reflection inversion^{19,20} among other applications.

This paper develops a trans-D inference method for MFI that samples over the seabed layering (i.e., does not assume a fixed number of layers) and, in addition, applies a hierarchical autoregressive error model¹⁸ to sample data-error process parameters that account for error dependence. This rigorous approach is found to substantially improve the fit to data and results in more meaningful uncertainty estimation by significantly reducing theory error. In addition, fewer prior choices and less algorithm tuning is required, resulting in a more automated inversion algorithm.

The trans-D PPD is sampled here by simulating a Markov chain on the trans-D state space using the reversible-jump Markov-chain Monte Carlo algorithm (rjMCMC)^{13,19} which is based on a generalization of the Metropolis–Hastings (MH) algorithm.^{21,22} The trans-D state space is considered to be the union of the state spaces for each model considered in the inversion, and transitions (*jumps*) between sub-spaces are implemented with a birth-death scheme.^{23,24} A significant challenge of rjMCMC inversions is very low acceptance rates of jumps proposed between sub-spaces^{25,26} (poor chain mixing). Several strategies have been proposed to improve dimension mixing, including delayed rejection,²⁵ more efficient proposal distributions,²⁶ and interacting Markov chains.²⁷ Interacting Markov chains are a generalization of replica Monte Carlo and parallel tempering.^{28,29} Interacting Markov chains are applied here to improve cross-dimension mixing as well as mixing within dimensions. Several chains are run in parallel and distributed according to an annealing schedule. Parameter-vector exchange between trans-D chains is governed by the MH criterion for multiple chains.³⁰ In addition, delayed rejection is applied for parameter updates resulting in an adaptive proposal distribution within sub-spaces.

The trans-D MFI algorithm is applied here to acoustic-field data collected near Elba Island in the Mediterranean Sea. Parameter inference is carried out in terms of profile marginal distributions for sound velocity, density, and attenuation. In addition, ensemble residual-error analysis is carried out to examine the hierarchical error model and error dependence.

II. BAYESIAN INFERENCE

A. Bayes' rule for hierarchical trans-D models

This section reviews trans-D inference^{12,13} used here to address seabed layering in MFI. The inversion assumes that the number of unknowns in the problem is itself an unknown parameter. This parameter is a model index for a group of models chosen *a priori* and is integrated over in a hierarchical Bayesian sense. The resulting posterior density spans multiple subspaces of different dimensions. A key advantage of this formulation is that the posterior includes the uncertainty due to the limited ability to identify a single model for the analysis.

Let \mathbf{d} be a random variable of N observed data containing information about a physical system. Further, let \mathcal{I}_k denote a group of models specifying particular choices of physics theory, model parametrizations, and error statistics to explain that system, where $k \in \mathcal{K}$ and \mathcal{K} is a countable set. Let $\mathbf{m}_k \in \mathbb{R}^{M_k}$ be a random variable of the M_k free parameters representing a realization of model \mathcal{I}_k . Green¹² shows that Bayes' rule can be written for a Bayesian hierarchical model to include parameter k ,

$$P(k, \mathbf{m}_k | \mathbf{d}) = \frac{P(k)P(\mathbf{d}|k, \mathbf{m}_k)P(\mathbf{m}_k|k)}{\sum_{k' \in \mathcal{K}} \int_{\mathcal{M}} P(k')P(\mathbf{d}|k', \mathbf{m}'_{k'})P(\mathbf{m}'_{k'}|k')d\mathbf{m}'_{k'}}, \quad (1)$$

where k indexes choices of appropriate models resulting in a single hierarchical model that spans several multi-dimensional subspaces. In Eq. (1), $P(k)$ is the prior over the \mathcal{K} models considered. The state variables of the trans-D inversion are given by (k, \mathbf{m}_k) and are of dimension M_k . The state space is trans-D and given by the union of all fixed-dimensional spaces in \mathcal{K} , i.e., $\cup_{k \in \mathcal{K}} (\{k\} \times \mathbb{R}^{M_k})$. A Markov chain that samples this state space can be defined that converges to the trans-D joint posterior $P(k, \mathbf{m}_k | \mathbf{d})$.

The seabed is parametrized with a partition model,²⁴ where layers are described in terms of interfaces in a partition extending from the water-sediment interface to a maximum depth of interest. Interface positions and geoacoustic parameters (e.g., sound velocity, density, attenuation) are sampled with the MH algorithm by accepting/rejecting perturbations to those parameters according to the MH criterion.²² Such partition models constitute a general approach to layer parametrization that allows structure in the model where needed with the ability to capture parameter gradients as well as abrupt discontinuities in the same model without requiring explicit assumptions.¹⁶

Creation and deletion of interfaces are implemented with a birth-death scheme based on the rjMCMC algorithm. Such a Markov-chain simulation [which samples the trans-D posterior in Eq. (1)] must be able to undergo dimension changes (*jumps*) in the state space (i.e., transition between subspaces for different models) while not violating the requirement for detailed balance (i.e., unbiased PPD sampling).¹⁹

Without dimension jumps, the MH algorithm is implemented using the MH acceptance for a step from the current model (k, \mathbf{m}_k) to a proposed model $(k', \mathbf{m}'_{k'})$:

$$\alpha_1 = \min \left[1, \frac{P(\mathbf{m}'_{k'})}{P(\mathbf{m}_k)} \left(\frac{P(\mathbf{d}|\mathbf{m}'_{k'})}{P(\mathbf{d}|\mathbf{m}_k)} \right)^\beta \frac{Q(\mathbf{m}_k|\mathbf{m}'_{k'})}{Q(\mathbf{m}'_{k'}|\mathbf{m}_k)} \right], \quad (2)$$

where β is an annealing parameter (see Ref. 19 for details) that can be considered unity here, and Q is the proposal distribution used to propose new states. For the trans-D case, a rjMCMC algorithm can be implemented^{12,13} similar to the MH algorithm by generalizing the acceptance rule to allow transitions between state spaces for the multiple models that are considered simultaneously. According to this rule, a step from the current model (k, \mathbf{m}_k) to a proposed model $(k', \mathbf{m}'_{k'})$ is accepted or rejected according to

$$\alpha = \min \left[1, \frac{P(k', \mathbf{m}'_k) \left(\frac{P(\mathbf{d}|k', \mathbf{m}'_k)}{P(\mathbf{d}|k, \mathbf{m}_k)} \right)^\beta \frac{Q(k, \mathbf{m}_k|k', \mathbf{m}'_k)}{Q(k', \mathbf{m}'_k|k, \mathbf{m}_k)} |\mathbf{J}| \right], \quad (3)$$

where $|\mathbf{J}|$ is the determinant of the Jacobian for the diffeomorphism from (k, \mathbf{m}_k) to (k', \mathbf{m}'_k) . The most commonly used form of rjMCMC is the birth-death scheme,^{23,24} which is applied here. A new layer is created by uniformly sampling a depth and proposing an interface at that depth. This interface splits an existing layer into two layers, one of which is randomly selected and its parameters perturbed. Such birth steps are accepted/rejected according to Eq. (3). Similarly, for a death step, an interface is selected at random and deleted. The parameters for the new (thicker) layer are randomly chosen to be the parameters of one of the two original layers, and the step is again accepted/rejected according to Eq. (3). A Gaussian proposal distribution for parameters during a birth step is applied here and the algorithm as implemented here ensures $|\mathbf{J}| = 1$.²⁴

The proposal distributions Q in Eqs. (2) and (3) cannot be arbitrarily adapted during a Markov-chain simulation. However, some adaptive scaling of Q is possible by introducing delayed rejection³¹ steps in the chain. Delayed rejection updates adapt the step size to local features in the state space by allowing additional steps to be used upon rejection as illustrated in Fig. 1. The proposal distributions Q_2 for delayed rejection steps can depend on the original proposal (standard deviations are typically chosen to be smaller) and on the rejected parameter vector. If Q_2 is independent of the rejected vector, the MH acceptance α_2 for delayed rejection in fixed dimension k is

$$\alpha_2 = \min \left[1, \frac{P(\mathbf{m}''_k) P(\mathbf{d}|\mathbf{m}''_k) Q_1(\mathbf{m}'_k|\mathbf{m}''_k) (1 - \alpha_1(\mathbf{m}'_k|\mathbf{m}''_k))}{P(\mathbf{m}_k) P(\mathbf{d}|\mathbf{m}_k) Q_1(\mathbf{m}'_k|\mathbf{m}_k) (1 - \alpha_1(\mathbf{m}'_k|\mathbf{m}_k))} \right], \quad (4)$$

where α_1 is the MH acceptance criterion for the initial step [Eq. (2)], \mathbf{m}_k the current parameter vector, \mathbf{m}'_k the rejected vector, and \mathbf{m}''_k the vector proposed using Q_2 . Equation (4) is applied here to allow some adaptability of the proposal distribution.

III. LIKELIHOOD FUNCTION AND AUTOREGRESSIVE ERROR MODEL

The conditional probability $P(\mathbf{d}|k, \mathbf{m}_k)$ in Eq. (1) describes the data-error statistics. An important concept in Bayesian inference is that once data have been observed,

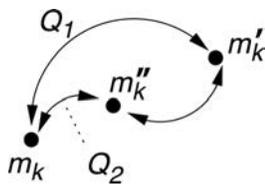


FIG. 1. Principle of delayed rejection. Original vector \mathbf{m}_k is perturbed using proposal distribution Q_1 to obtain vector \mathbf{m}'_k , which is rejected via the MH criterion, Eq. (2). A second perturbed vector \mathbf{m}''_k is then proposed using \mathbf{m}_k and Q_2 , which is evaluated according to Eq. (4).

$P(\mathbf{d}|k, \mathbf{m}_k)$ is interpreted as the likelihood function $L(k, \mathbf{m}_k)$ of the model parameter vector. Because data errors include both measurement and theory errors (which cannot generally be separated), the specific form of this distribution is often unknown. For quantitative application of the likelihood function, the form of the error distribution must be assumed and its parameters specified. Multivariate Gaussian distributions are commonly used, but the validity of this assumption should be examined *a posteriori*.^{9,32}

To formulate the MFI problem, complex acoustic pressure fields measured at a linear array of H hydrophones and F frequencies ($N = FH$) are given by $\mathbf{d} = \{\mathbf{d}_f, f = 1, F\}$. Here data errors are assumed to be independent between frequencies but can be spatially correlated between hydrophones. The assumption of independence between frequencies is justified for reasonably spaced frequencies.³³ For complex, circular-symmetric Gaussian-distributed data errors, the likelihood function is given by

$$L(k, \mathbf{m}_k) = \prod_{f=1}^F \frac{1}{\pi^H |\mathbf{C}_f|} \exp\{-[\mathbf{d}_f - \mathbf{d}_f(k, \mathbf{m}_k)]^\dagger \times \mathbf{C}_f^{-1} [\mathbf{d}_f - \mathbf{d}_f(k, \mathbf{m}_k)]\}, \quad (5)$$

where \dagger is the conjugate transpose, $\mathbf{d}_f(k, \mathbf{m}_k)$ are the predicted acoustic data for (k, \mathbf{m}_k) at frequency f , and \mathbf{C}_f are the data covariance matrices. For independent identical data errors at each frequency, Eq. (5) simplifies to

$$L(k, \mathbf{m}_k) = \frac{1}{\prod_{f=1}^F (\pi \sigma_f^2)^H} \times \exp\left\{-\sum_{f=1}^F |\mathbf{d}_f - \mathbf{d}_f(k, \mathbf{m}_k)|^2 / \sigma_f^2\right\}, \quad (6)$$

where σ_f are the standard deviations at each frequency. Because source amplitude and phase information is often unavailable in MFI, predicted data are expressed

$$\mathbf{d}_f(k, \mathbf{m}_k) = A_f \exp(i\theta_f) \mathbf{p}_f(k, \mathbf{m}_k), \quad (7)$$

where $\mathbf{p}_f(k, \mathbf{m}_k)$ are predicted acoustic pressure data for a source with zero phase and unit amplitude computed here with the normal-mode propagation model ORCA,³⁴ and A_f and θ_f represent the unknown amplitude and phase at the f th frequency. A maximum likelihood estimate for $A_f \exp(i\theta_f)$ is obtained from Eq. (5) and given by¹¹

$$A_f \exp(i\theta_f) = \frac{\mathbf{p}_f^\dagger(k, \mathbf{m}_k) \mathbf{C}_f^{-1} \mathbf{d}_f}{\mathbf{p}_f^\dagger(k, \mathbf{m}_k) \mathbf{C}_f^{-1} \mathbf{p}_f(k, \mathbf{m}_k)}. \quad (8)$$

The residual error $\mathbf{r}'_f(k, \mathbf{m}_k) = \mathbf{d}_f - \mathbf{d}_f(k, \mathbf{m}_k)$ is then given by

$$\mathbf{r}'_f(k, \mathbf{m}_k) = \mathbf{d}_f - A_f \exp(i\theta_f) \mathbf{p}_f(k, \mathbf{m}_k). \quad (9)$$

The covariance matrices \mathbf{C}_f in Eq. (5) are often unknown due to poor understanding of error sources, in particular, theory errors. In some cases, \mathbf{C}_f can be estimated using multiple data measurements to obtain ensemble averages; however,

care is required to appropriately sample over all error processes affecting the data. Therefore it can be advantageous to estimate error distribution parameters from residual error analysis by non-parametric approaches. Non-parametric approaches⁹ use an optimal (maximum-likelihood) estimate of the parameter vector to provide residual errors that provide estimates of variances and covariances (under the assumption of ergodicity). However, non-parametric estimates are problematic for trans-D problems because a specific model is assumed to evaluate the estimate. Which model to choose for point-estimate residual errors may not be clear *a priori*, and the approach may be inappropriate when significant posterior probability is found for more than one model. In addition, non-parametric approaches may require windowing, smoothing, or damping to ensure non-singular matrices, which relies on ad-hoc tuning parameters requiring trial and error techniques.³²

This paper uses a hierarchical autoregressive (AR) error model to account for error dependence.¹⁸ Such models can capture general covariance structure with a small number of parameters and quantify the effect of uncertainties in data error statistics on parameter uncertainty estimates. Autoregressive models³⁵ are given by

$$\hat{d}_t = \sum_{i=1}^p a_i d_{t-i} + \epsilon_t, \quad (10)$$

where t indexes the data, p is the order of the AR model, a_i are the AR parameters, and ϵ_t is an uncorrelated Gaussian noise term. The model is abbreviated as AR(p). In the case of the AR(1) model, a closed form expression for the autocovariance is given by

$$A_l = \frac{\sigma^2}{1 - a_1^2} a_1^{|l|}, \quad (11)$$

where l is the lag and σ the standard deviation of ϵ_t . Hence, the data errors for an AR(1) process are Gaussian distributed with a Toeplitz covariance matrix in which off-diagonal terms fall off exponentially. In inversion of complex data, two AR(p) processes are applied (for real and imaginary parts) using one a_p parameter for both parts. The a_p are sampled with MH updates; this accounts for uncertainty in the AR estimate in the environmental model.

The complete hierarchical error model at frequency f is given by $\boldsymbol{\sigma}_f = (\sigma, \mathbf{a})_f$, where $\mathbf{a}_f = \{a_{ij}, i = 1, p\}$ and σ is the standard deviation. After applying the AR model, the likelihood function is given by

$$L(k, \mathbf{m}_k) = \frac{1}{\prod_{f=1}^F (\pi \sigma_f^2)^H} \times \exp \left\{ - \sum_{f=1}^F |\mathbf{d}_f - A_f \exp(i\theta_f) \mathbf{p}_f(k, \mathbf{m}_k) - \hat{\mathbf{d}}(\mathbf{a}_f)|^2 / \sigma_f^2 \right\}. \quad (12)$$

The AR predictions $\hat{\mathbf{d}}(\mathbf{a}_f)$ are obtained by applying Eq. (10) to the residual error [Eq. (9)] The *total data residuals* are

given by $\mathbf{r}_f = \mathbf{d}_f - \mathbf{d}_f(k, \mathbf{m}_k) - \hat{\mathbf{d}}(\mathbf{a}_f)$ and are assumed to be uncorrelated Gaussian distributed with standard deviations σ_f . The assumptions about the residuals can be examined by applying statistical tests to the total residual sample from a large number of models drawn randomly from the posterior (residual samples are saved along with model parameters during the inversion). In practice, Eq. (12) means that the observed data are fit with the predicted data $\mathbf{d}_f(k, \mathbf{m}_k)$ due to a physics process (e.g., acoustic-seabed interaction) and with an additional predicted AR process $\hat{\mathbf{d}}(\mathbf{a}_f)$ to account for serial correlations of the data errors. Autoregressive parameters \mathbf{a}_f and standard deviation σ_f are unknowns in the inversion. This formulation has the advantage of providing a general sampling scheme that can include complex error models with no requirements to invert a covariance matrix or compute its determinant at each step of the inversion. Hierarchical-error models are particularly useful in trans-D inverse problems where point estimates (limited to one fixed-dimensional subspace of the trans-D state space) may not be representative of all subspaces.

Here maximum-likelihood estimates of σ_f are used to implicitly sample³⁶ the error standard deviations, in which case the likelihood function is given by

$$L(k, \mathbf{m}_k) = -H \sum_{f=1}^F \log_e |\mathbf{r}_f(k, \mathbf{m}_k, \mathbf{a}_f)|^2. \quad (13)$$

Estimating error statistics from the data has the intrinsic challenge that resolving physical model structure is traded off with estimating error-distribution parameters. However, the scale of data features giving rise to interpretation of seabed structure is generally large compared to features of error processes. Reasonable parametrizations of physics models typically only predict a limited range of data features, while AR processes can fit general features. Hence observations for many inverse problems exhibit dominant features that are most straightforwardly fit by the physics model. This allows the AR process to fit error-process features not explained by typical geoacoustic models. In practice, the AR model should have the lowest order that sufficiently removes correlations, and the magnitude of AR predictions can be constrained appropriately by prior information. To ensure that the AR process does not remove meaningful geoacoustic structure from the PPD, inversions should be carried out and compared with and without hierarchical AR models.

IV. INTERACTING MARKOV CHAINS

Markov-chain Monte Carlo (MCMC) is based on the simulation of a Markov chain (“sampling”) that is guaranteed to converge to the PPD (the target distribution) in the limit of infinitely long chain simulations. In practice, the simulation is carried out for a limited number of MCMC steps until the chain has converged sufficiently to the target distribution. Testing convergence is difficult, and no definitive method to establish convergence exists.³⁷ To examine chain convergence, the sampling is typically run for large numbers of steps, and results are monitored continuously

until moments of the PPD cease to change significantly. However, such tests can be misleading, and it is advisable to continue sampling for some time after apparent convergence.

The rate of convergence depends critically on how efficiently the chain can sample the high-probability regions of the state space. The efficiency is governed by the proposal distribution in Eq. (3), which must be specified *a priori* and cannot be adapted during sampling. Defining efficient proposal distributions is an important but potentially difficult aspect of applying MCMC sampling to Bayesian inference problems. The sampling is particularly challenging when multiple modes exist in the target distribution that are separated by regions of low probability, and the chain is required to jump freely back and forth (“mix”) between modes to converge in reasonable time. One approach to increasing efficiency for multi-modal problems are annealed or tempered methods,^{29,38} which relax the target distribution by raising the likelihood function to the power of an annealing parameter β [referred to as the inverse “temperature,” also see Eq. (3)]. Such approaches typically require increased computational effort that is balanced by a much more rapid rate of convergence. More generally, these methods can be described as population-based or interacting MCMC.^{27,30,39} Formally, interacting MCMC is based on the concept that to sample from the target distribution $\pi = P(\mathbf{m}|\mathbf{d})$, a new target π^* is defined

$$\pi^* = \prod_{i=1}^N \pi_i, \quad (14)$$

where at least one $\pi_i \equiv \pi$. The additional distributions are defined such that they are simpler to sample and are introduced to aid in the sampling of $P(\mathbf{m}|\mathbf{d})$. The sequence of π_i can be chosen in any meaningful way and often includes an annealed sequence $\{\pi_i = \pi^{\beta_i}, \beta_i \in [0, 1]\}$. Each distribution in the sequence is then targeted by corresponding MCMC chains that are simulated in parallel. Such annealed sequences allow large, efficient jumps (rapid mixing) of the chains in the parameter space for low β values while the PPD target is given when $\beta_i = 1$. To take advantage of the rapid mixing of the chains that target low- β π_i , chains can exchange information by way of *exchange updates*. Exchange updates are based on randomly selecting pairs of chains from the sequence and proposing to swap the parameter vectors of the chains. For two chains n and q with parameter vectors \mathbf{m}_n and \mathbf{m}_q , respectively, updates are accepted/rejected based on the MH criterion for exchange updates given by

$$\hat{\alpha} = \min \left[1, \frac{\pi_n(\mathbf{m}_q)\pi_q(\mathbf{m}_n)}{\pi_n(\mathbf{m}_n)\pi_q(\mathbf{m}_q)} \right]. \quad (15)$$

Note that the proposal distribution is symmetric if chains are randomly selected from the population, resulting in a unit proposal ratio.

For trans-D problems, jumps between subspaces for different parametrizations and dimensions pose a substantial challenge for inter-model mixing. Improving trans-D mixing has been the focus of considerable research.^{25–27,40} Here, an

annealed sequence of interacting trans-D Markov chains is simulated in parallel to improve dimension mixing. The sequence is used to target $\pi = P(k, \mathbf{m}|\mathbf{d})$, and exchange updates are proposed periodically throughout the chain simulation. The β sequence is chosen to accept exchange updates $\sim 50\%$ of the time. Each chain is updated according to the MH criterion in Eq. (3). Hence some chains explore the trans-D parameter space much more freely than others and can exchange information according to Eq. (15), and annealed chains show significantly increased dimension mixing. Because the π_i are trans-D, exchange moves are also trans-D and allow parameter vectors of different dimension to be exchanged, improving dimension mixing of individual chains. In addition to improved dimension mixing, chain mixing within dimensions is also addressed by the population of chains.

V. APPLICATION TO MEDITERRANEAN-SEA DATA

A. Experiment and data

The PROSIM’97 shallow-water geoacoustic experiment was carried out by the SACLANT (now NATO) Undersea Research Centre in the Mediterranean Sea off the west coast of Italy near Elba Island. The experiment and data have been considered previously^{11,41} and are described only briefly here. The experiment consisted of recording acoustic signals from a transducer towed at approximately 12-m depth over a track with nearly range-independent bathymetry (water depth ~ 132 m). The source emitted a linear frequency-modulated chirp signal that swept over the frequency band 400–800 Hz in 0.5 s. The signals were received at a bottom-moored vertical line array (VLA) of 48 hydrophones that spanned from 26- to 120-m depth with a 2-m spacing. The acoustic pressure at each hydrophone was sampled at 3 kHz, transmitted to the ship, and recorded. The ocean sound-velocity profile (SVP) was measured during the experiment and consisted of a weakly downward-refracting gradient that decreased from about 1520 to 1510 m/s over the water column. The data analyzed here were recorded for a source-receiver range of approximately 3.85 km.

The experiment and the environmental and geometric parameters included in the model \mathbf{m} are illustrated in Fig. 2. The

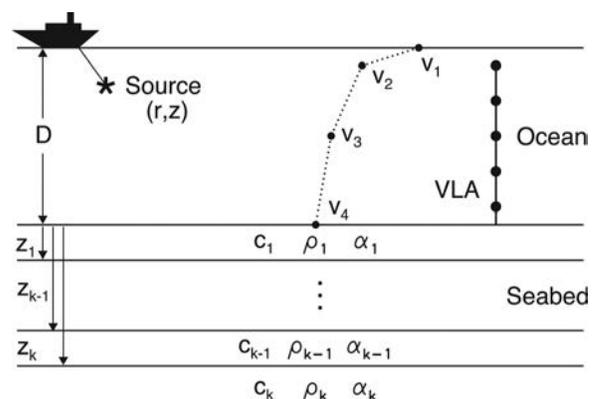


FIG. 2. Experiment geometry and parametrization of matched field experiment.

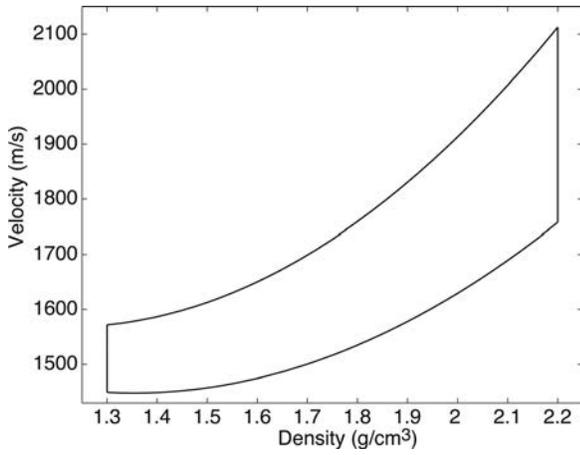


FIG. 3. Prior bounds for density-velocity relationship.

acoustic source is at depth z and range r from the VLA in water of depth D . The SVP is represented by four unknown sound velocities v_1-v_4 at depths of 0, 10, 50, and D m. For the purposes of extracting seabed information, the SVP and geometric unknowns represent nuisance parameters that are not themselves of interest but are included in the model to account for the effect of their uncertainties on the estimated geoacoustic parameters. The seabed is parametrized in terms of a partition model over the sub-bottom depths of interest with k interfaces, where intervals between interfaces are homogeneous layers represented by sound velocity c , density ρ , and attenuation coefficient α . Prior distributions for sound velocity and density are constrained by uniform joint prior distributions that are based on empirical data.^{42,43} For a given density ρ , lower (c_{low}) and upper (c_{up}) values for velocity are constrained by the empirical expressions (Fig. 3)

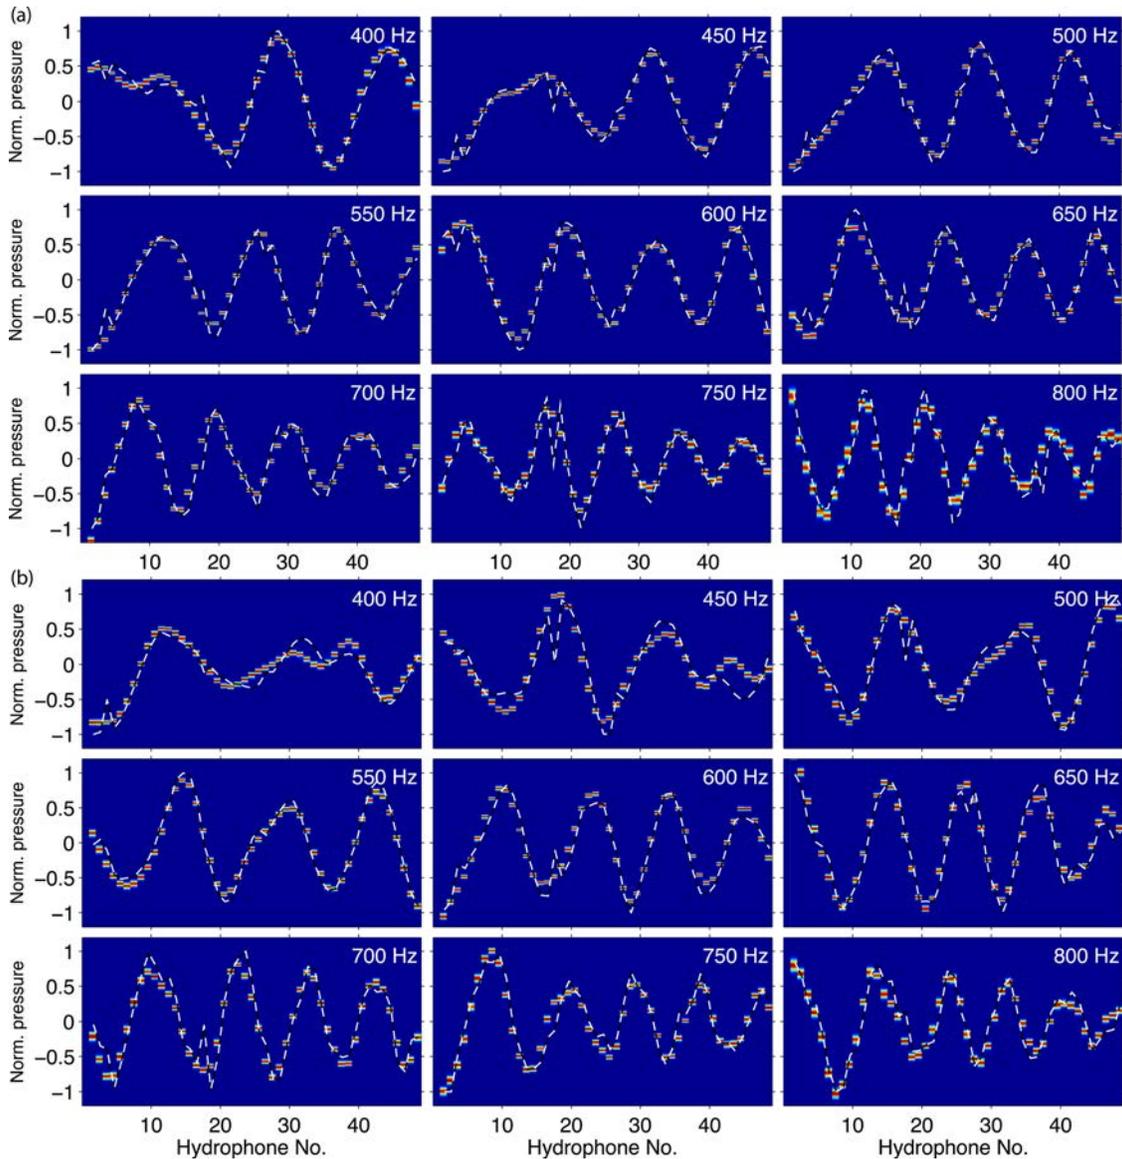


FIG. 4. (Color online) Measured data (dashed lines) in terms of normalized pressure as a function of hydrophone number and frequency [(a) real, (b) imaginary parts]. The shading indicates the density of the data predictions from an ensemble of models (a random subset of the PPD).

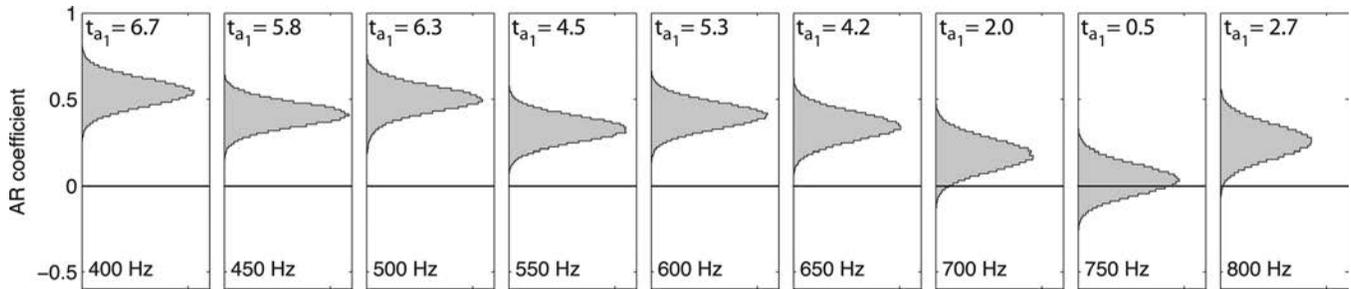


FIG. 5. Marginal probability distributions of AR(1) parameters for each frequency.

$$c_{\text{low}} = (1.54 - 0.907\rho + 0.3659\rho^{1.88})1500.4,$$

$$c_{\text{up}} = (1.60 - 0.907\rho + 0.3695\rho^{2.01})1501.4. \quad (16)$$

These joint prior bounds allow a wide range of physically meaningful c - ρ combinations to be considered, while excluding physically unreasonable combinations.

Bounds for SVP and geometric parameters (v_i , D , r , z) are chosen to be relatively small, representing measurement uncertainty (e.g., instrument calibration) and possible effects of lateral and/or temporal variability. Unless otherwise mentioned, prior bounds are represented by the plot bounds in the figures that follow. The goal of MFI is to estimate model parameters that predict data that match measured acoustic

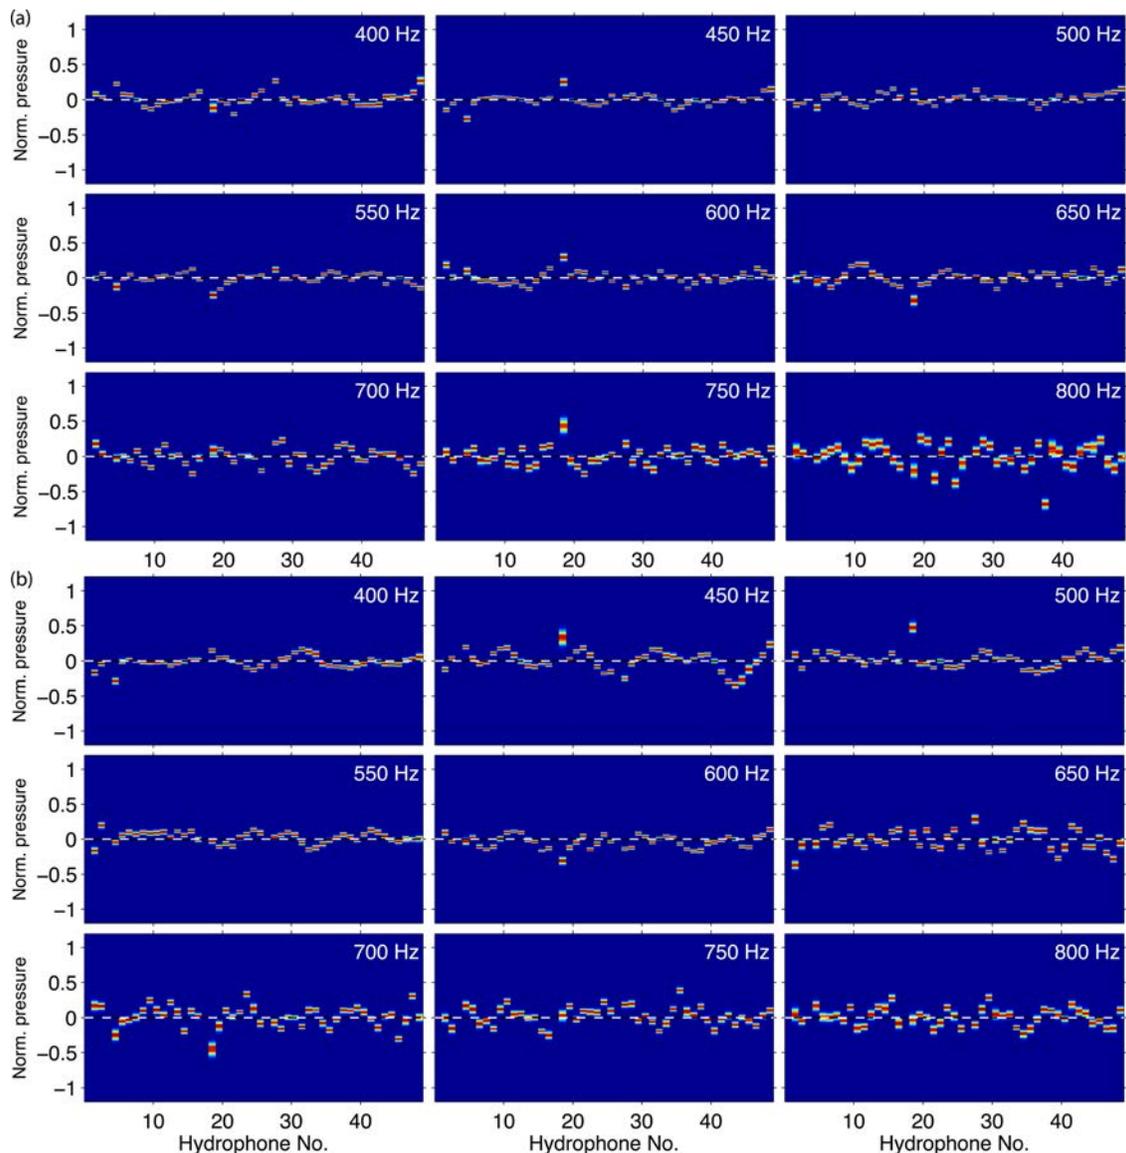


FIG. 6. (Color online) Total data residual errors [(a) real, (b) imaginary part] for the inversion results with an AR(1) error model.

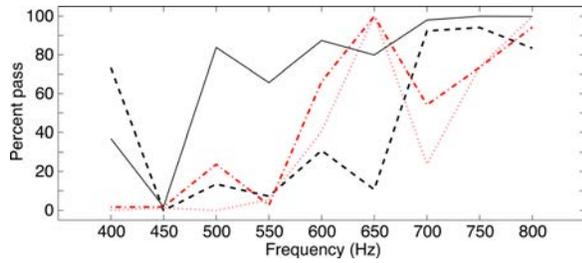


FIG. 7. (Color online) Runs-test results for real (solid) and imaginary (dashed) data for the inversion with an AR(1) error model and for real (dashed-dotted) and imaginary (dotted) data for the inversion without an AR model.

fields. The measured data indicate¹¹ that the acoustic sensors at depths of 32, 58, and 76 m had gain problems, and these data were excluded from the analysis. The data considered for the inversions in this paper consist of acoustic fields at nine frequencies at 50-Hz intervals from 400 to 800 Hz.

B. Posterior parameter inference

Figure 4 shows the measured acoustic-field data (real and imaginary parts) and the range of data predictions (excluding AR predictions) from the inversion, which represents an ensemble inference constituting the density of predicted fields for a random sub-sample of the PPD. The predicted-field density is considered here as a more rigorous approach to examine data fit compared to considering only the prediction for a single best-fit model vector. Figure 4 shows that the model fits the data well. In fact, the fit in this figure is significantly better than in previous inversions of the same data.¹¹ This likely indicates that previous inversions did not consider an amount of seabed structure consistent with data information (models were under-parametrized).

In addition to the hierarchical partition model for seabed layers, an AR(1) model was included in the inversions to address error dependence. The data errors at each frequency were modeled by one AR parameter, which was applied to both the real and imaginary parts. The marginal probability distributions for the AR(1) parameters are shown in Fig. 5

and indicate that the data show significant correlation at several frequencies. Low frequencies generally require higher AR(1) parameter values than high frequencies. In particular, data at 700 and 750 Hz, which show the least amount of correlation in Fig. 4, also have marginal distributions in Fig. 5 that are close to zero. The mean values of the distributions are generally consistent with the degree of error dependence seen in Fig. 4. The t -test⁴⁴ can be applied to examine the significance of including an AR parameter in the error model:

$$t_{a_{\max}} = \hat{a}_{\max} / \text{stderr}(a_{\max}) \gtrsim 2, \quad (17)$$

where a_{\max} is the highest-order coefficient included in the analysis and stderr represents standard error. The t -test results are given for each frequency in Fig. 5 and indicate that at eight of the nine frequencies, the a_1 coefficients are significant. Only data at 750 Hz show an insignificant t -test value of 0.5. The inversion was also carried out for an AR(2) model, and the results (not shown here) indicate that a second-order ($a_{\max} = a_2$) AR model is not justified at any frequency. Therefore the AR(1) model is taken to be necessary and sufficient for the error correlations in this inversion.

The difference between observed and predicted data are referred to as residual errors and are considered here in terms of ensemble-density distributions (Fig. 6) for real and imaginary parts. Figure 6 shows the total residual error, including the AR(1) model, which should be an uncorrelated random process. The residuals at most frequencies appear largely uncorrelated and increasingly so with increasing frequency. To quantify these results, runs tests⁴⁴ were applied to the full ensemble of models from the PPD with the results shown in Fig. 7. In particular, Fig. 7 compares the runs-test results of the inversion both including and not including the AR(1) model. Generally, including the AR(1) model improves the runs-test results with a larger percentage of the models in the PPD passing the test at the 0.05 significance level. Hence inclusion of the AR(1) model in the inversion reduces the error dependence. However, some dependence is still apparent in the residuals at low to mid frequencies.

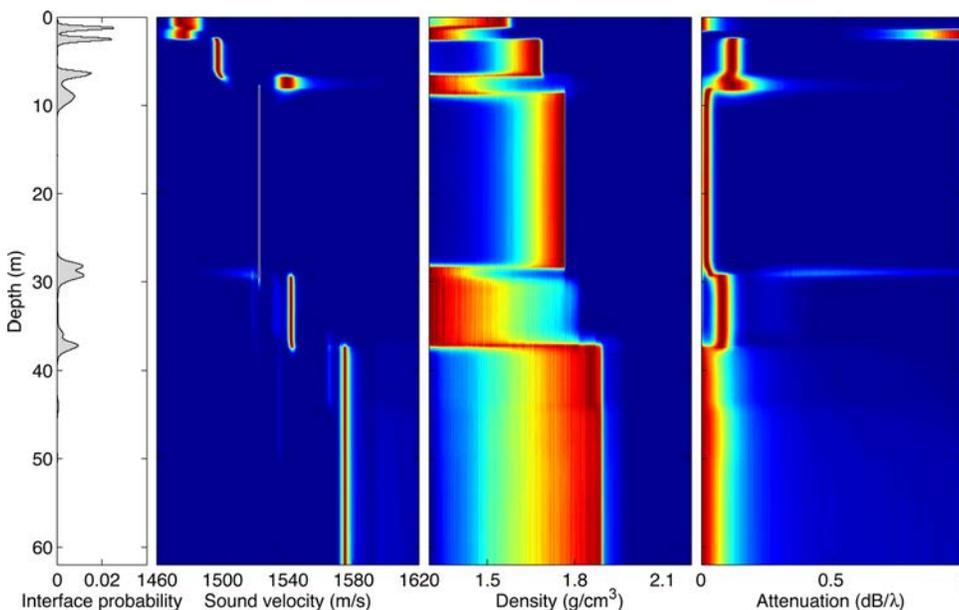


FIG. 8. (Color online) Trans-D marginal profile distributions for sound velocity, density, attenuation, and interface marginal probability as a function of depth.

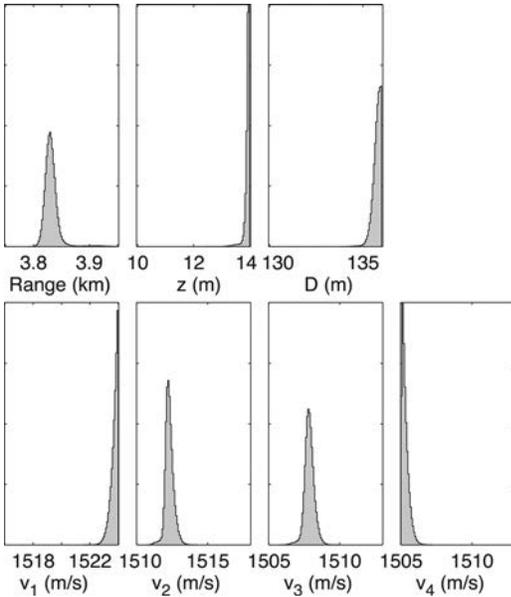


FIG. 9. One-dimensional marginal probability distributions for water-column and geometric parameters.

The main inversion result for the PROSIM data is shown in Fig. 8 in terms of profile marginal distributions for interface probability, velocity, density, and attenuation. Note that the marginal profile distributions are normalized at each depth individually for display purposes. The results show that the MFI inversion is largely sensitive to velocity and resolves structure consistent with other previous work.¹¹ The uncertainties of the velocity distribution are generally larger for thin layers than for thicker layers. In particular, over the depths from ~ 10 to 30 m, velocity uncertainty is very small, which was previously observed. Density appears to be largely unconstrained, which is common for MFI. Attenuation values are mostly low, which is consistent with the relatively large depth to which the data appear to resolve

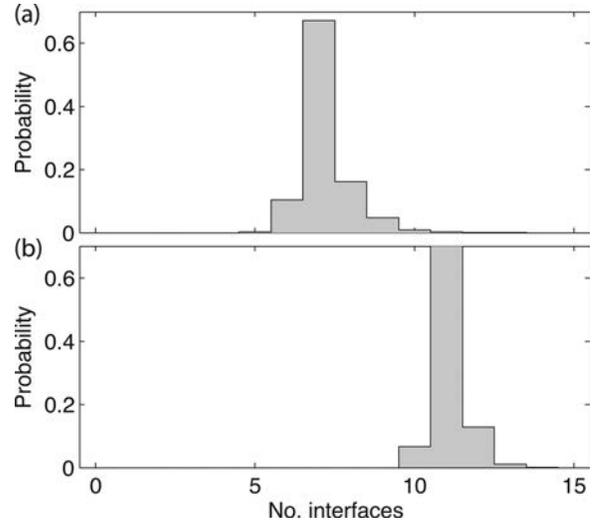


FIG. 10. Marginal probability distribution of the number of interfaces in the sediment stack: (a) including the AR(1) model and (b) no error model.

layering (~ 40 m). Figure 8 also shows interface probability as a function of depth, which indicates approximately six major interfaces with the uncertainty of interface positions generally increasing with depth. The water-column and geometric parameters are shown in Fig. 9. The prior bounds for all parameters in Fig. 9 are given by the plot boundaries. It can be seen that several parameters are constrained by the prior bounds that are chosen based on experiment values. The peaking of these distributions at the boundaries implies that the parameters are not resolved well. Allowing wide bounds could potentially bias other parameter results.

Figure 10 shows a comparison of the marginal distribution of the number of interfaces supported by the data for the two inversions. The distribution for the inversion including the AR(1) model shows significantly fewer layers (peak at 7) than the inversion ignoring correlated errors (peak at 11 layers). Including the AR model effectively reduces the

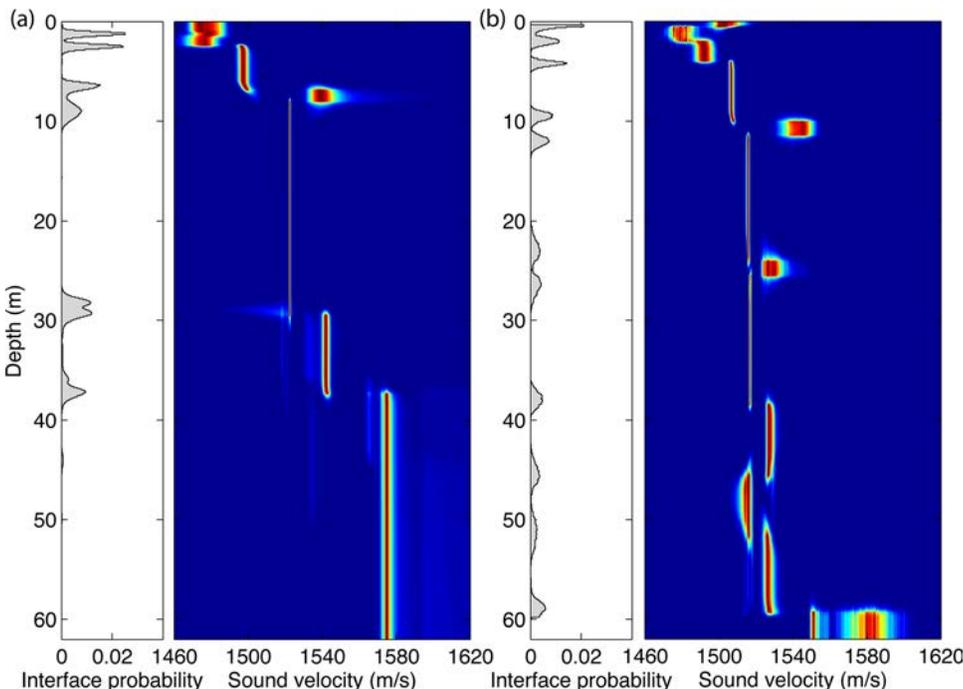


FIG. 11. (Color online) Comparison of interface probability and velocity marginal profile plots for (a) inversion including an AR error model and (b) ignoring correlated errors.

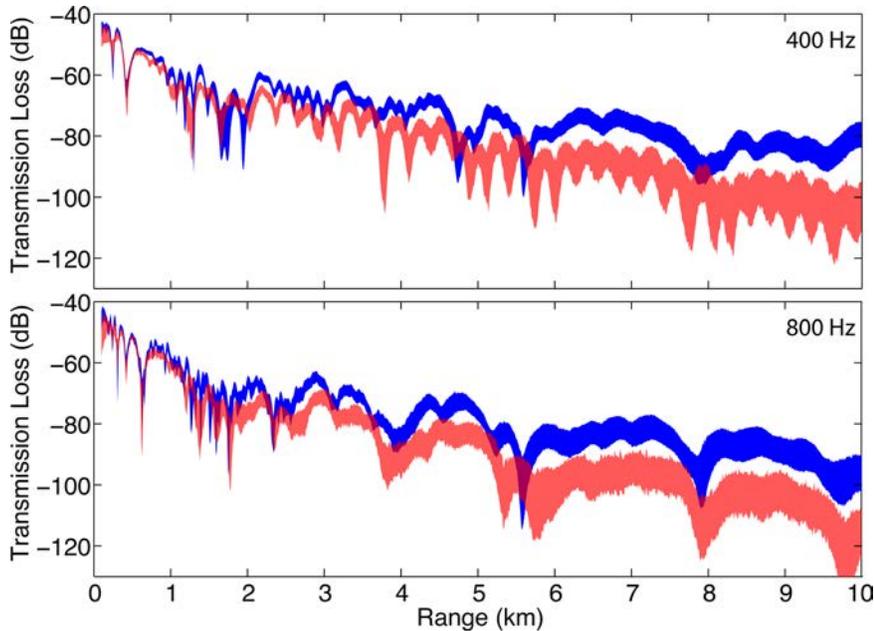


FIG. 12. (Color online) Impact of including the AR error model in MFI on TL predictions for 400 and 800 Hz: 95% highest probability density credibility intervals for inversion including AR model (solid) and ignoring correlated errors (transparent).

amount of seabed structure supported by the data by accounting for correlated errors. This is an important result because in some cases, data features that actually represent error processes can be fit by additional seabed layers (given the model); this may lead to unjustified structure. The AR model appears to balance the tendency of including unreasonable structure by allowing error dependence.

Figure 11 compares velocity and interface probability marginal profiles for both inversions. Similar to Fig. 10, unconstrained structure is evident for the inversion that ignores correlated data errors. This is particularly prominent for the interface probabilities, where neglecting error dependence leads to many additional interfaces with small velocity contrasts. In addition, the velocity-profile distribution indicates unconstrained structure with velocities jumping back and forth from layer to layer. Further, the maximum partition depth in the inversion was set to 60 m, and strong interfaces appear to this depth, indicating unconstrained structure. Finally, Fig. 11(b) shows a bimodal distribution for the basement velocity that is not observed in Fig. 11(a).

To illustrate the significance of a well-formulated geoacoustic inversion to acoustic propagation modeling and sonar performance prediction applications, Fig. 12 compares transmission loss (TL) predictions for both inversions with 20-m source depth and 50-m receiver depth out to 10-km range. The lower (400 Hz) and upper (800 Hz) ends of the frequency band used in the inversion are considered. For both frequencies, TL predictions are significantly different for the two inversion results, which illustrates the impact of model complexity on TL predictions. The more complex seabed model that exhibits unjustified layering structure causes significantly more structure in the TL predictions in Fig. 12.

C. Interacting chains and chain mixing

All inversions reported in this paper were carried out with a population of seven chains with logarithmically spaced $\beta = [1., 1., 0.87, 0.756, 0.658, 0.571, 0.497]$. Note

that two chains are simulated for $\beta_i = 1$, resulting in the seven chains yielding PPD samples at 3.5 times the computational effort compared to running a single $\beta_i = 1$ chain. This section considers the mixing of the chain targeting $\pi = \pi^{\beta_i}$ to illustrate the increase in across-model mixing and parameter mixing within a model. Figure 13 compares across-model mixing for inversions with and without interacting chains. During 400 algorithm steps, interacting chains change model dimension 116 times, while dimension changes only 5 times without interacting chains. In addition, note that model jumps for interacting chains can change k by more than 1. Figure 14 shows a mixing comparison for the velocity parameter of the first layer. Both chains and chain-autocorrelation functions are shown and indicate a substantial increase in mixing. In particular, the autocorrelation-function peak for interacting chains is much narrower; this indicates less sample-to-sample dependence. Similar analysis for other parameters gives similar results (not shown). This illustrates how interacting chains can significantly improve mixing for high-dimensional trans-D inverse problems while only moderately increasing computational effort (3.5 times).

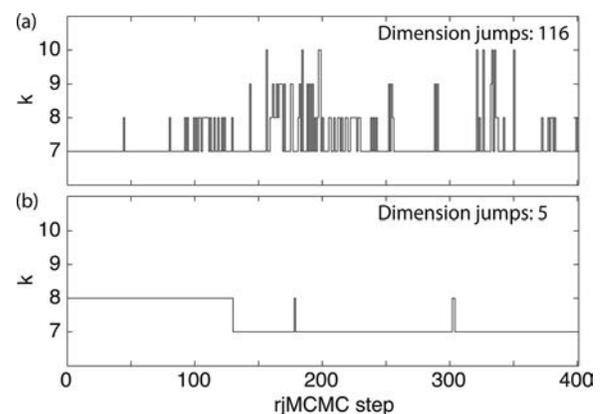


FIG. 13. Model-index mixing for inversions using (a) interacting chains and (b) a single chain. A small subset of 400 rjMCMC steps is shown.

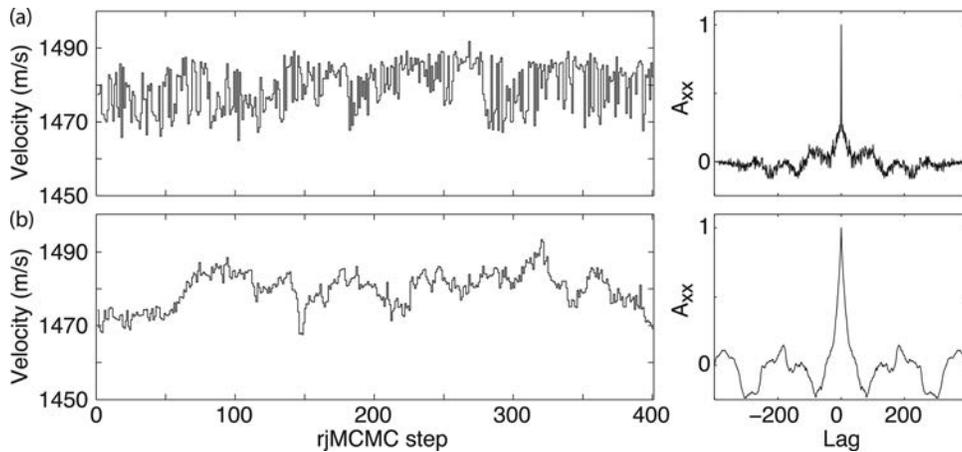


FIG. 14. Mixing of the velocity parameter of the first layer and the associated autocorrelation functions for (a) interacting chains and (b) single chain. A small subset of 400 rjMCMC steps is shown.

VI. SUMMARY AND DISCUSSION

A trans-D inversion with a hierarchical data-error model was developed and applied here to matched field inversion. The trans-D algorithm was implemented using interacting Markov chains that significantly improve performance of the algorithm. In particular, jumping between models was substantially improved by using interacting chains; this addresses an important difficulty in trans-D inversion.

The algorithm was applied to the PROSIM data recorded in the Mediterranean Sea. Results are largely sensitive to seabed velocities and resolve more detailed structure than previous inversions. The trans-D seabed model allows simultaneous consideration of many seabed parametrizations in the inversion and effectively addresses model selection with respect to layering as part of the inverse problem. The PPD estimate resulting from this inversion jointly quantifies the state of knowledge about seabed parameters given data and prior information for all models included in the inversion. Hence trans-D geoacoustic uncertainties include the uncertainty due to the intrinsically limited ability to specify a single model for the unknown seabed. Therefore prior assumptions with regard to model choice are relaxed to a group of models, resulting in more realistic results.

Error dependence has been shown to be an issue for MFI and is addressed here by a hierarchical error model using AR processes to model error dependence. The AR model was shown to result in more reasonable structure, while ignoring correlations resulted in an unconstrained velocity profile. Hierarchical error models further relax the requirement of specifying aspects of the model *a priori*.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the support of Office of Naval Research Grant No. 322AO. The computational work was carried out on a parallel high-performance computing cluster operated by the authors at the University of Victoria funded by the Natural Sciences and Engineering Research Council of Canada and the Office of Naval Research. The authors thank the NATO Undersea Research Centre and Dr. Peter Nielsen for at-sea data collection and two anonymous reviewers for valuable comments that improved the quality of the manuscript.

¹M. D. Collins, W. A. Kuperman, and H. Schmidt, "Nonlinear inversion for ocean-bottom properties," *J. Acoust. Soc. Am.* **93**, 2770–2783 (1992).

²S. E. Dosso, M. L. Y Jeremy, J. M. O zard, and N. R. Chapman, "Estimation of ocean-bottom properties by matched-field inversion of acoustic field data," *IEEE J. Ocean. Eng.* **18**, 232–239 (1993).

³P. Gerstoft, "Inversion of seismoacoustic data using genetic algorithms and a posteriori probability distributions," *J. Acoust. Soc. Am.* **95**, 770–782 (1994).

⁴J. S. Liu, *Monte Carlo Strategies in Scientific Computing* (Springer, New York, 2001), pp. 1–343.

⁵D. J. C. MacKay, in *Information Theory, Inference, and Learning Algorithms* (Cambridge University Press, Cambridge, UK, 2003), pp. 343–386.

⁶S. E. Dosso, "Quantifying uncertainty in geoacoustic inversion. I. A fast Gibbs sampler approach," *J. Acoust. Soc. Am.* **111**, 129–142 (2002).

⁷D. J. Battle, P. Gerstoft, W. S. Hodgkiss, W. A. Kuperman, and P. L. Nielsen, "Bayesian model selection applied to self-noise geoacoustic inversion," *J. Acoust. Soc. Am.* **116**, 2043–2056 (2004).

⁸Y. Jiang, N. R. Chapman, and H. A. DeFerrari, "Geoacoustic inversion of broadband data by matched beam processing," *J. Acoust. Soc. Am.* **119**, 3707–3716 (2006).

⁹S. E. Dosso, P. L. Nielsen, and M. J. Wilmut, "Data error covariance in matched-field geoacoustic inversion," *J. Acoust. Soc. Am.* **119**, 208–219 (2006).

¹⁰C.-F. Huang, P. Gerstoft, and W. S. Hodgkiss, "On the effect of error correlation on matched-field geoacoustic inversion," *J. Acoust. Soc. Am.* **121**, EL64–EL69 (2007).

¹¹S. Dosso and J. Dettmer, "Bayesian matched-field geoacoustic inversion," *Inverse Probl.* **27**, 055009 (2011).

¹²P. J. Green, "Reversible jump Markov chain Monte Carlo computation and Bayesian model determination," *Biometrika* **82**, 711–732 (1995).

¹³P. J. Green, "Trans-dimensional Markov chain Monte Carlo" in *Highly Structured Stochastic Systems, Oxford Statistical Science Series* (Oxford University Press, Oxford, UK, 2003), pp. 179–198.

¹⁴A. Malinverno, "Parsimonious Bayesian Markov chain Monte Carlo inversion in a non-linear geophysical problem," *Geophys. J. Int.* **151**, 675–688 (2002).

¹⁵M. Sambridge, K. Gallagher, A. Jackson, and P. Rickwood, "Trans-dimensional inverse problems, model comparison and the evidence," *Geophys. J. Int.* **167**, 528–542 (2006).

¹⁶T. Bodin and M. Sambridge, "Seismic tomography with the reversible jump algorithm," *Geophys. J. Int.* **178**, 1411–1436 (2009).

¹⁷N. P. Agostinetti and A. Malinverno, "Receiver function inversion by trans-dimensional Monte Carlo sampling," *Geophys. J. Int.* **181**, 858–872 (2010).

¹⁸J. Dettmer, S. Molnar, G. A. M. W. Steininger, S. E. Dosso, and J. F. Cassidy, "Transdimensional inversion of microtremor array dispersion data with hierarchical autoregressive error models," *Geophys. J. Int.* **188**, 719–734 (2012).

¹⁹J. Dettmer, S. E. Dosso, and C. W. Holland, "Trans-dimensional geoacoustic inversion," *J. Acoust. Soc. Am.* **128**, 3393–3405 (2010).

²⁰J. Dettmer, S. E. Dosso, and C. W. Holland, "Sequential trans-dimensional Monte Carlo for range-dependent geoacoustic inversion," *J. Acoust. Soc. Am.* **129**, 1794–1806 (2011).

- ²¹N. Metropolis, A. Rosenbluth, M. Rosenbluth, and A. T. A. E. Teller, "Equations of state calculations by fast computing machines," *J. Chem. Phys.* **21**, 1087–1092 (1953).
- ²²W. K. Hastings, "Monte Carlo sampling methods using Markov chains and their applications," *Biometrika* **57**, 97–109 (1970).
- ²³C. Geyer and J. Moller, "Simulation procedures and likelihood inference for spatial point processes," *Scand. J. Stat.* **21**, 359–373 (1994).
- ²⁴D. G. T. Denison, C. C. Holmes, B. K. Mallick, and A. F. M. Smith, *Bayesian Methods for Nonlinear Classification and Regression* (Wiley, New York, 2002), pp. 1–227.
- ²⁵P. J. Green and A. Mira, "Delayed rejection in reversible jump Metropolis–Hastings," *Biometrika* **88**, 1035–1053 (2001).
- ²⁶S. P. Brooks and G. O. R. P. Giudici, "Efficient construction of reversible jump Markov chain Monte Carlo proposal distributions," *J. R. Stat. Soc.* **65**, 3–39 (2003).
- ²⁷A. Jasra, D. A. Stephens, and C. C. Holmes, "Population-based reversible jump Markov chain Monte Carlo," *Biometrika* **94**, 787–807 (2007).
- ²⁸R. H. Swendsen and J. S. Wang, "Replica Monte Carlo simulation of spin-glasses," *Phys. Rev. Lett.* **57**, 2607–2609 (1986).
- ²⁹C. J. Geyer, "Markov chain Monte Carlo maximum likelihood," in *Computing Science and Statistics: Proceedings of the 23rd Symposium on the Interface* (Interface Foundation, Fairfax Station, VA, 1991), pp. 156–163.
- ³⁰A. Jasra, D. A. Stephens, and C. C. Holmes, "On population-based simulation for static inference," *Stat. Comp.* **17**, 263–279 (2007).
- ³¹L. Tierney and A. Mira, "Some adaptive Monte Carlo methods for Bayesian inference," *Stat. Med.* **18**, 2507–2515 (1999).
- ³²C. W. Holland, J. Dettmer, and S. E. Dosso, "Remote sensing of sediment density and velocity gradients in the transition layer," *J. Acoust. Soc. Am.* **118**, 163–177 (2005).
- ³³S. E. Dosso and P. Nielsen, "Quantifying uncertainty in geoacoustic inversion. II. Application to broadband, shallow-water data," *J. Acoust. Soc. Am.* **111**, 143–159 (2002).
- ³⁴E. K. Westwood, C. T. Tindle, and N. R. Chapman, "A normal mode model for acoustoelastic ocean environments," *J. Acoust. Soc. Am.* **100**, 3631–3645 (1996).
- ³⁵R. Shumway and D. Stoffer, in *Time Series Analysis and Its Applications* (Springer, New York, 2000), pp. 89–212.
- ³⁶S. E. Dosso and M. J. Wilmut, "Data uncertainty estimation in matched-field geoacoustic inversion," *IEEE J. Ocean. Eng.* **31**, 470–479 (2006).
- ³⁷C. J. Geyer, "Introduction to Markov chain Monte Carlo" in *Handbook of Markov Chain Monte Carlo* (Springer, New York, 2011), pp. 3–47.
- ³⁸R. M. Neal, "Annealed importance sampling," *Stat. Comp.* **11**, 125–139 (2001).
- ³⁹P. D. Moral, A. Doucet, and A. Jasra, "Sequential Monte Carlo samplers," *J. R. Stat. Soc.* **68**, 411–436 (2006).
- ⁴⁰A. Jasra, A. Doucet, D. A. Stephens, and C. C. Holmes, "Interacting sequential Monte Carlo samplers for trans-dimensional inference," *Comput. Stat. Data Anal.* **52**, 1765–1791 (2008).
- ⁴¹M. R. Fallat and S. E. Dosso, "Geoacoustic inversion via local, global, and hybrid algorithms," *J. Acoust. Soc. Am.* **105**, 3219–3230 (1999).
- ⁴²D. R. Jackson and M. D. Richardson, "Geoacoustic properties" in *High-Frequency Seafloor Acoustics* (Springer, New York, 2007), pp. 123–170.
- ⁴³J. E. Quijano, S. E. Dosso, J. Dettmer, L. M. Zurk, M. Siderius, and C. Harrison, "Bayesian geoacoustic inversion using wind-driven ambient noise," *J. Acoust. Soc. Am.* **131**, 2658–2667 (2012).
- ⁴⁴D. C. Montgomery and E. A. Peck, in *Introduction to Linear Regression Analysis* (Wiley, New York, 1992), pp. 115–136.