

Bayesian Detectability of Induced Polarisation in Airborne Electromagnetic Data using Reversible Jump Sequential Monte Carlo

Laurence Stuart Davies¹, Christopher Drovandi¹, Alan Yusen Ley Cooper², and Matthew Sutton¹

¹Queensland University of Technology

²Geoscience Australia

November 22, 2022

Abstract

Detection of induced polarisation (IP) effects in airborne electromagnetic (AEM) measurements does not yet have an established methodology. This contribution develops a Bayesian approach to the IP-detectability problem using decoupled transdimensional layered models, and applies an approach novel to geophysics whereby transdimensional proposals are used within the embarrassingly parallelisable and robust static Sequential Monte Carlo (SMC) class of algorithms for the simultaneous inference of parameters and models. Henceforth referring to this algorithm as Reversible Jump Sequential Monte Carlo (RJSMC), the statistical methodological contributions to the algorithm account for adaptivity considerations for multiple models and proposal types, especially surrounding particle impoverishment in unlikely models. Methodological contributions to solid Earth geophysics include the decoupled model approach and proposal of a statistic that use posterior model odds for IP detectability. A case study is included investigating detectability of IP effects in AEM data at a broad scale.

Bayesian Detectability of Induced Polarisation in Airborne Electromagnetic Data using Reversible Jump Sequential Monte Carlo

L. Davies^{1,2}, A.Y. Ley-Cooper³, M. Sutton^{1,2}, C. Drovandi^{1,2}

¹School of Mathematical Sciences, Queensland University of Technology, Brisbane QLD 4000, Australia

²Centre for Data Science, Queensland University of Technology, Brisbane QLD 4000, Australia

³Geoscience Australia, Canberra ACT 2609, Australia

Key Points:

- We introduce a reduced model space approach via decoupled layered model parameters for inference of induced polarisation (IP) models versus conductive-only models fitted to airborne electromagnetic (AEM) data.
- For model and parameter inference we develop an adaptive static sequential Monte Carlo algorithm with reversible jump Markov chain Monte Carlo proposals (RJSMC).
- We successfully apply RJSMC to perform airborne induced polarisation (AIP) detectability on data from a large survey that demonstrates spatial continuity.

Abstract

Detection of induced polarisation (IP) effects in airborne electromagnetic (AEM) measurements does not yet have an established methodology. This contribution develops a Bayesian approach to the IP-detectability problem using decoupled transdimensional layered models, and applies an approach novel to geophysics whereby transdimensional proposals are used within the embarrassingly parallelisable and robust static Sequential Monte Carlo (SMC) class of algorithms for the simultaneous inference of parameters and models. Henceforth referring to this algorithm as Reversible Jump Sequential Monte Carlo (RJSMC), the statistical methodological contributions to the algorithm account for adaptivity considerations for multiple models and proposal types, especially surrounding particle impoverishment in unlikely models. Methodological contributions to solid Earth geophysics include the decoupled model approach and proposal of a statistic that use posterior model odds for IP detectability. A case study is included investigating detectability of IP effects in AEM data at a broad scale.

Plain Language Summary

Electromagnetic models for solid-Earth geophysics often make assumptions to ease computation. These assumptions may hold in the majority of cases, however in cases where it is impossible to explain certain problematic data using such models it is necessary to either revisit the model assumptions or to consider empirical model approximations. In this research, the problematic data contains significant anomalous measurements which are hypothesised to be due to the presence of a phenomenon known as induced polarisation (IP). This phenomenon is possibly explained using an empirical extension to current epistemological physical models given sufficient statistical evidence. It is the purpose of this research to introduce a rigorous statistical methodology for detecting when it is more likely that the empirical model will explain the data versus the epistemological model. This methodology is tested on artificial and real-world data, demonstrating the applicability and usefulness of the approach.

1 Introduction

Negative or steeply-decaying measurements of magnetic flux density in a concentric-loop airborne electromagnetic (AEM) system are usually inexplicable in electromagnetic models that ignore intrinsic chargeability. Despite wide acknowledgement of the possibility that these and other anomalies are caused by induced polarisation (IP) in subsurface materials (Kratzer & Macnae, 2012), confirmation of these hypotheses for airborne data is currently an open area of research. Moreover, the influence of such IP effects is not limited to producing negative measurements; significant distortions in off-time transients can present in any number of ways that can produce incorrect conductivity values in non-IP ground models (Viezzoli & Manca, 2020).

Deterministic methods for inverting time-domain AEM data with hypothesised IP effects generally do not consider parameter or model uncertainty, and as such do not directly provide a means for model inference. Approximate methods for model selection such as the Akaike Information Criterion (AIC) (Akaike, 1974) are not robust to the pathological posteriors often found in layered-Earth geophysical models. To date, the most advanced Bayesian sampling approach applied to the Cole-Cole IP model (Cole & Cole, 1941) has been Markov Chain Monte Carlo (MCMC) for within-model parameter inference of non-airborne time and frequency domain data (Ghorbani et al., 2007). Related work in conductive-only electromagnetic models applied to airborne data has advanced as far as transdimensional inference via various implementations of reversible jump Markov chain Monte Carlo (RJMCMC) (Brodie & Sambridge, 2012; B. Minsley, 2011; Hawkins et al., 2017) and parallel-tempering RJMCMC (Blatter et al., 2018). It is a natural progression to consider Bayesian transdimensional inference methods for IP detection in airborne data in a similar manner.

Induced polarisation in AEM data, hitherto referred to as airborne induced polarisation (AIP), has attracted recent research interest posing the question of detecting such effects. This

has been driven by the fact that ground-based IP methods comprise some of the most widely used techniques in mineral exploration, not only for the discovery of many anomalous mineralisation prospects due to their chargeable response (Meju, 2002), but also for the increasing ability to identify subsurface materials and mineralogy (Merriam (2007), Qi et al. (2018), Feng et al. (2020)). Approaches using thresholds of re-parameterisations (Fiandaca et al., 2018), modelling of 3D IP effects (Nunes et al., 2019), and detection of IP using various approaches (Kang et al., 2019; Viezzoli et al., 2021) are some examples. However, Bayesian inference on the detectability of IP effects in AEM data has not been considered. This contribution will demonstrate how AIP detectability can be framed as a tractable model selection problem which can be accomplished using Bayesian methods. In addition, this work introduces the novel application of static Sequential Monte Carlo with an RJMCMC mutation kernel to parameter and model inference in AIP. First formally addressed in Del Moral et al. (2006); Jasra et al. (2008) and subsequently identified in Zhou et al. (2016) as the SMC1 algorithm, this work will expand on the implementation of adaptive considerations in such an algorithm both generally and specifically for layered Earth models. It will henceforth be referred to as the Reversible Jump Sequential Monte Carlo (RJSMC) algorithm.

Application of SMC methods to geophysics problems is still relatively new. Recent work by Amaya et al. (2021) demonstrates how adaptive static SMC can be used to an advantage when needing to sample from complex priors. Earlier work by Dettmer et al. (2011) demonstrates a completely different approach whereby in an application to sequential inference on geoacoustic survey lines the posterior distribution of a previous model in the sequence is input as the initial distribution for inference on parameters of the next model. Whilst such approaches are similar by name, the algorithms of Amaya et al. (2021) and Dettmer et al. (2011) are configured to target different statistical quantities. Our article contributes to the growing literature demonstrating that RJSMC is a viable alternative to other popular methods such as RJMCMC and parallel tempering (Swendsen & Wang, 1986) in Geophysics. In addition, we present for the first time an implementation of transdimensional proposals designed for layered-Earth models in the SMC framework.

In comparison to other particle-population methods such as Population MCMC (Jasra et al., 2007) and parallel tempering (Swendsen & Wang, 1986), SMC algorithms are based on a methodology of sequentially sampling from a sequence of probability distributions on a common space (Del Moral et al., 2006). These probability distributions are approximated using a cloud of weighted random samples, or particles, where the process of moving to the next distribution is via a combination of importance sampling, resampling, and a specifically chosen mutation or propagation kernel. SMC has several benefits: the embarrassingly-parallel mutation of N particles rather than reliance on one or few sequentially computed chains in other MCMC methods translates easily to parallel computing architectures; the algorithm is robust to high-dimensional multi-modal, often pathological posteriors (referred to in Ellis (1998) as non-uniqueness in AEM geophysics models); the availability of the particle approximation to the current target distribution allows for novel tuning possibilities; and a trivial additional output of SMC is the estimation of otherwise intractable normalising constants.

By using an MCMC kernel for particle mutation in combination with an artificial sequence of distributions (such as likelihood annealing) SMC can be readily applied to static parameter estimation problems such as AIP. While the use of SMC for model selection in static problems can take several forms as explored in Zhou et al. (2016), this contribution will focus on an implementation that returns the joint posterior of parameters and models, facilitating a straightforward application of Bayesian model selection for the detection of AIP effects.

This paper is divided into three sections following this introduction. Section 2, titled *Methodology*, describes a modelling approach that employs two decoupled layered-Earth models to reduce the size of the model space, the application of RJSMC to the IP detectability problem in AEM data, where we introduce the Bayes Factor Induced Polarisation Detectability (BFIPD) statistic. Section 3 describes the computational algorithms for SMC and RJSMC, and proposes adaptive considerations for the latter. Section 4, titled *Case Studies*, applies the described method-




Figure 1: Synthetic data produced using the Cole-Cole (Cole & Cole, 1941) parameterisation for an AEM layered-Earth model. The second layer intrinsic chargeability, m , varies from 0.0 to 0.8 to demonstrate the effect on the magnetic field response. Conductivity parameters are constant across models. Negative log values when $m \geq 0.8$ are plotted with a dashed line.

ology and RJSMC algorithm to a synthetic study and subsequently to a real data set in Colorado USA demonstrating the spatial continuity of a novel IP detection statistic (described in Section 2.3.3).

2 Methodology

For any flight location and geometry of an airborne electromagnetic (AEM) system, the response to the magnetic flux density response can be computed analytically using the 1D layered-Earth model approximation which employs a Hankel transform and propagation matrix method (Ward & Hohmann, 1988). This approximation use the quasi-static assumption to reduce the Helmholtz wave equations to more tractable diffusion equations. In doing so, the dielectric permittivity term that models chargeability is removed. To re-introduce this term, the empirical Cole-Cole equation can be formulated (following Seigel (1959)) using low and high frequency conductivity terms σ_s and σ_∞ respectively with intrinsic chargeability defined as $m = \frac{\sigma_\infty - \sigma_s}{\sigma_\infty}$. This gives the following form for complex conductivity in terms of the high-frequency conductivity σ_∞ , frequency ω , time delay constant τ and frequency dependence c :

$$\hat{\sigma}(\omega) = \sigma_\infty \left(1 - \frac{m}{1 + (i\omega\tau)^c} \right).$$

Although other formulations exist (such as that of Pelton et al. (1978)), they are not considered here since such models can be computed from the above via a bijective transformation (Tarasov & Titov, 2013). This work considers only time domain electromagnetic (TDEM) airborne data for inversion using independent 1D layered-Earth models for each AEM sounding. A depiction of models of varying chargeability and the resultant off-time TDEM synthetic responses computed using this mathematical model is in Figure 1. Since the number of layers is not known *a-priori*, a transdimensional approach is adopted to infer the number of layers. The Bayesian approach for this described in Malinverno (2002) introduced RJMCMC proposals for the birth and death of layer interfaces, each being comprised of a depth (or thickness) and one or more electromagnetic parameters. All parameters and associated prior distributions are identified in Section 2.3.1, however in the next section they will all be referred to in concatenated vector form using the symbol θ .

2.1 Decoupled Layered Models

A primary concern for algorithmic complexity is the size of the model space. If we desire inference on the detectability of Cole-Cole chargeability in every conductive layer in a 1D




Figure 2: (a) A single transdimensional parameter that models both conductivity and chargeability with respect to depth. Change to layer interface depths will change both conductivity and chargeability profiles because they are coupled. (b) Two de-coupled transdimensional parameters each separately modelling conductivity and chargeability with respect to depth. These are merged via a direct sum before computing the mathematical model.

model then we are implicitly using binary variables to model the inclusion of such parameters in each layer. We call this a *coupled* layered-model, where for k layers this effectively poses a model space $\mathcal{M}$ of exponential cardinality. Instead of investigating detectability of IP in this manner, we propose to *de-couple* the conductivity and Cole-Cole chargeability layer interfaces. This approach has been used in previous geophysical applications such as joint inversion of physically unrelated data in a single model (Piana Agostinetti & Bodin, 2018). In this application it permits a more flexible and much smaller model space $|\mathcal{M}_{\text{decoupled}}| = k_{\kappa} \times k_{\lambda}$ where k_{κ}, k_{λ} signify the number of independent conductivity and Cole-Cole chargeability layers respectively. A comparison of a decoupled design with a coupled design is visualised in Figure 2. The reduction of the model space cardinality is imperative when considering the inference approach we will take in the following sections whereby we begin by sampling the priors of each model.

The next subsection will introduce Bayesian inference of parameters and models, followed by a Bayesian description of the decoupled model for conductivity and chargeability. Subsequently, computational algorithms for inference on these parameters and model parameterisation will be discussed, including an introduction of the novel RJSMC sampler for which the decoupled model design is suited due to the reduced size of the model space.

2.2 Bayesian inference of parameters and models

Bayesian inference of parameters θ given data $\mathbf{y}$, likelihood $\mathcal{L}(\mathbf{y}|\theta)$ (explained in Section 2.3.2) and prior $p(\theta)$ (see Section 2.3.1) is found via application of Bayes' theorem

$$\pi(\theta|\mathbf{y}) = \frac{\mathcal{L}(\mathbf{y}|\theta)p(\theta)}{\mathcal{Z}(\mathbf{y})},$$

where the denominator term, known as the marginal likelihood or normalising constant, is the integral

$$\mathcal{Z}(\mathbf{y}) = \int_{\theta} \mathcal{L}(\mathbf{y}|\theta)\pi(\theta)d\theta,$$

which is typically computationally intractable in a non-trivial number of dimensions. This term can be used for model selection using Bayes Factor (Kass & Raftery, 1995) (a quantity we will use in later sections) where between two contending models k_1, k_2 it is the ratio

$$\mathcal{B}_{1,2} = \frac{\int p(\theta_{k_1}|k_1)\mathcal{L}(\mathbf{y}|\theta_{k_1}, k_1)d\theta_{k_1}}{\int p(\theta_{k_2}|k_2)\mathcal{L}(\mathbf{y}|\theta_{k_2}, k_2)d\theta_{k_2}} = \frac{\mathcal{Z}(\mathbf{y}|k_1)}{\mathcal{Z}(\mathbf{y}|k_2)}.$$

For inference over the joint space of models $k \in \mathcal{M}$ and parameters θ_k we extend the above approach to express the posterior as

$$\pi(\theta_k, k | \mathbf{y}) = \frac{\mathcal{L}(\mathbf{y} | \theta_k, k) p(\theta_k | k) p(k)}{\sum_{k' \in \mathcal{M}} \mathcal{Z}(\mathbf{y} | k')},$$

where the $\mathcal{Z}$ terms in the denominator will be henceforth defined as

$$\mathcal{Z}_k = \mathcal{Z}(\mathbf{y} | k) = \int_{\theta_k} \mathcal{L}(\mathbf{y} | \theta_k, k) p(\theta_k | k) p(k) d\theta_k.$$

If the above joint posterior of models and parameters is available, we can use Bayes' theorem to express Bayes Factor in terms of the ratios of posterior model marginal densities (also called the *posterior model odds*) and models priors

$$\mathcal{B}_{1,2} = \frac{\mathcal{Z}_{k_1}}{\mathcal{Z}_{k_2}} = \frac{\pi(k_1 | \mathbf{y}) \pi(k_2)}{\pi(k_2 | \mathbf{y}) \pi(k_1)}.$$

2.3 Bayesian inference of induced polarisation models

This subsection will employ the Bayesian approach to specifying IP and non-IP models and the inference of such models and parameters with respect to AEM data.

2.3.1 Parameter Priors

Using the decoupled layered model design, we identify the following parameter priors (using the indices i, j to denote the i^{th} conductive layer and j^{th} chargeable layer respectively). Note that we use the notational convenience $\phi = \log_{10} \sigma$.

$$\begin{aligned} \text{Number of conductive layers : } \kappa &\sim \mathcal{U}\{0, \dots, \kappa_{max}\}, \\ \text{Number of chargeable layers : } \lambda &\sim \mathcal{U}\{0, \dots, \lambda_{max}\}, \\ \text{Background Log10 conductivity : } \theta_{\phi,b} &\sim \mathcal{U}(-4, 2), \\ i^{th}\text{-layer Log10 conductivity : } \theta_{\phi,i} &\sim \mathcal{U}(-4, 2), \\ i^{th}\text{-layer Interface Depth : } \theta_{z_{\sigma},i} &\sim \frac{\kappa!}{z_{max}^{\kappa}}, \\ \text{Background Intrinsic Chargeability : } \theta_{m,b} &= 0, \\ j^{th}\text{-layer Intrinsic Chargeability : } \theta_{m,j} &\sim \mathcal{U}(0, 1), \\ j^{th}\text{-layer Interface Depth : } \theta_{z_m,j} &\sim \frac{\lambda!}{z_{max}^{\lambda}}, \\ \text{Time Constant : } \theta_{\tau} &\sim \mathcal{U}(0, 1), \\ \text{Frequency Dependence : } \theta_c &\sim \mathcal{U}(0, 1). \end{aligned}$$

Using bold notation θ_{ϕ} to denote all parameters $\{\theta_{\phi,i}\}_{i=1}^{\kappa}$ (and similarly for $\theta_{z_{\sigma}}$, θ_m , and θ_{z_m}) the full prior distribution becomes

$$p(\kappa, \lambda, \theta_{\phi,b}, \theta_{\phi}, \theta_{z_{\sigma}}, \theta_m, \theta_{z_m}, \theta_{\tau}, \theta_c) = p(\kappa) p(\theta_{\phi}, \theta_{z_{\sigma}} | \kappa) p(\lambda) p(\theta_m, \theta_{z_m} | \lambda) p(\theta_{\phi,b}) p(\theta_{\tau}) p(\theta_c).$$

The parameter vector in full, using the pair κ, λ to identify the model, is

$$\theta_{\kappa,\lambda} = \{\theta_{\phi,b}, \theta_{\phi,1}, \dots, \theta_{\phi,\kappa}, \theta_{z_{\sigma},1}, \dots, \theta_{z_{\sigma},\kappa}, \theta_{m,1}, \dots, \theta_{m,\lambda}, \theta_{z_m,1}, \dots, \theta_{z_m,\lambda}, \theta_{\tau}, \theta_c\}.$$

Note that θ_{τ} and θ_c Cole-Cole parameters are global for all chargeable layers. This could feasibly mean that there theoretically exist chargeable models that poorly fit the above parameterisation; however, in practical examples this has not yet been an issue.

2.3.2 The Likelihood

The data $\mathbf{y}$ is represented by an time-series array of response values, usually the time rate-of-change of the magnetic field ($\frac{dB}{dt}$) with units in picoteslas per second (pT/s), which for conductive only models are usually positive. Acquisition of data occurs during the off-time phase of the periodic waveform for the transmitted current. For example, the VTEM ET system (Eadie et al., 2018) begins data acquisition from 5 microseconds after the start of the off-time period and continues sampling for up to an additional 15 milliseconds. Data is then down-sampled to a time series of gates which are typically exponentially-spaced with respect to increasing time. Gates are indexed in the below using the square bracket notation $y[i]$. Visualisations of synthetic data were introduced at the beginning of this section on methodology using various example θ parameterisations in Figure 1.

The mathematical *forward model* for an AEM system is a non-injective (and often non-surjective) map from the 1D layered-Earth model parameterisation of conductivity and chargeability parameters, θ , to the data space of off-time magnetic flux density responses, denoted $\mathbf{y}^\theta$. The likelihood $\mathcal{L}(\theta|\mathbf{y})$ is a multivariate Gaussian $\mathcal{N}(\mathbf{y}; \mathbf{y}^\theta, \Sigma_{\mathbf{y}^\theta})$ computed in the synthetic data space. Typical methods for determining $\Sigma_{\mathbf{y}^\theta}$, i.e. the system noise, use an empirical approximation from high-altitude measurements where the effects of the ground are not present. This research employs the empirical model of A. Green and Lane (2003) where $\Sigma_{\mathbf{y}^\theta}$ is a diagonal covariance matrix where each diagonal element is the summation in quadrature of additive noise ϵ_{AN} for that time window and multiplicative noise ϵ_{MN} , that is

$$\Sigma_{\mathbf{y}^\theta}^{(i,j)} = \begin{cases} (\epsilon_{AN}[i])^2 + (\epsilon_{MN} \times y^\theta[i])^2, & i = j, \\ 0, & i \neq j. \end{cases}$$

Such an approach requires high-altitude calibration lines to be flown immediately prior to data acquisition, and such calibration determines a per-time-window value for ϵ_{AN} and an overall value for ϵ_{MN} , where typical values for ϵ_{MN} are in the order of 5%. The empirical additive and multiplicative noise model is adequate for data where only conductivity parameters are of interest. However, in data with one or several significant zero-crossings in middle to late time windows, the additive noise will become the dominant influence in the likelihood. In this research, it was found that “small” additive noise resulted in poor estimation of the parameter posteriors; this observation was consistent with other Bayesian research in transdimensional geophysics methods (Bodin et al., 2012) which notes that the variance of data noise strongly affects the shape of the posterior. Since the AEM system noise models were not necessarily constructed with Bayesian sampling methods for IP in mind, the additive noise was kept above 10^{-3} pT/s to ensure that reasonable estimates of the posterior were feasible. It would be the subject of future research to determine a noise model that is parametric such as that in Bodin et al. (2012).

2.3.3 A statistic for Bayesian detectability of induced polarisation

Inference of detectability of IP effects in this work is via Bayes factor using the expected probabilities of chargeable models and non-chargeable models. Denoting the chargeable versus non-chargeable estimate of Bayes factor as $\hat{\mathcal{B}}_{\kappa\lambda}$, where $k = \kappa$ is a non-chargeable model and $k = \lambda$ is a chargeable model, we define the Bayes Factor Induced Polarisation Detectability (BFIPD) statistic to be the grouped Bayes Factor

$$\hat{\mathcal{B}}_{\kappa\lambda} = \frac{|\mathcal{M}_\kappa| \sum_{\lambda \in \mathcal{M}_\lambda} \hat{\pi}(\lambda|\mathbf{y})/p(\lambda)}{|\mathcal{M}_\lambda| \sum_{\kappa \in \mathcal{M}_\kappa} \hat{\pi}(\kappa|\mathbf{y})/p(\kappa)}.$$

Using $\log \hat{\mathcal{B}}_{\kappa\lambda}$, a chargeable model is more likely when this value is greater than zero. Applied to real data, the third section in Figure 5 demonstrates how this quantity can be used to detect chargeability in Earth materials assuming no other anomalous effects.

3 Computational algorithms for Bayesian model inference

The canonical approach for inference of $\pi(\theta_k, k|\mathbf{y})$ is to implement a RJMCMC algorithm that proposes on the space of k as well as θ . Whilst such algorithms asymptotically converge to the posterior, there are several concerns. One is the difficulty of designing a well-mixing RJMCMC proposal, another is the lack of parallelisability due to the dependent nature of the algorithm. To solve the first, either the various adaptive proposals that conform to rules of diminishing adaptation (Haario et al., 2001) or proposals that use local derivative information (Roberts and Stramer (2002), Girolami and Calderhead (2011)) have shown success for within model MCMC, but these do not always translate well for RJMCMC proposal design. For the second concern, frameworks such as parallel tempering (Swendsen & Wang, 1986) can be parallelised to a certain degree, but do not upscale well to take advantage of increasingly common very wide computing architectures.

The SMC family of algorithms can be adapted for static parameter inference (Chopin (2002), Del Moral et al. (2006)), with the added benefit of providing an estimate of the marginal likelihood (Del Moral et al., 2006). For parameter estimation in a single non-linear model, static SMC has been demonstrated to have advantages when posteriors are pathological. To date, applications of static SMC in solid Earth geophysics are currently sparse, however recent work by Amaya et al. (2021) demonstrates static SMC for within model parameter inference as well as the use of estimates of the marginal likelihood for model selection. In the following section, a static SMC algorithm will be introduced and followed by an implementation which uses transdimensional proposals as well as within-model proposals, which we call Reversible Jump Sequential Monte Carlo (RJSMC).

3.1 Static Sequential Monte Carlo

In this section we will briefly state the algorithm in the common configuration where an MCMC kernel is used for mutation (the reader may be interested in the work of Dai et al. (2020) for a more in-depth and up-to-date review of SMC samplers). Such a configuration ensures that the computation of incremental particle weights can be evaluated with linear time complexity without the need to compute or approximate an expensive integral (Del Moral et al., 2006). The first implementation of this approach was called Iterated Batch Importance Sampling (Chopin, 2002), which used a *data annealing* schedule for successive target distributions. For brevity, the below description does not include commonly employed adaptive schemes such as those employed in Fearnhead and Taylor (2013), however such approaches are addressed in the following section for RJSMC.

Following from Bayes Theorem in section 2.2, the posterior of parameters θ is proportional to the likelihood $\mathcal{L}$ and prior $p(\cdot)$, that is

$$\pi(\theta|\mathbf{y}) \propto \mathcal{L}(\mathbf{y}|\theta)p(\theta).$$

In a static SMC algorithm, a sequence of distributions π_t , $t = 0, \dots, T$ is specified that “smoothly” transitions from a starting distribution, most commonly the prior $p(\cdot)$, to the target posterior distribution $\pi(\cdot|\mathbf{y})$. For inference using sparse data (such as AEM time-domain data) a common choice for this sequence is *likelihood annealing*, where a monotonic sequence $\{\gamma_t\}_{t=0}^{t=T}$ with $\gamma_0 = 0$ ascending to $\gamma_T = 1$ defines the sequence of target distributions

$$\pi_t(\theta|\mathbf{y}) = \frac{\mathcal{L}(\mathbf{y}|\theta)^{\gamma_t} p(\theta)}{\mathcal{Z}_t}.$$

The proportionality constant $\mathcal{Z}_t$ usually cannot be evaluated, so we instead define a tractable term η_t as

$$\eta_t(\theta) = \mathcal{Z}_t \pi_t(\theta|\mathbf{y}) = \mathcal{L}(\mathbf{y}|\theta)^{\gamma_t} p(\theta).$$

The target distributions π_t are approximately represented by a set of N particles. During initialisation, the particles are drawn from the prior $\pi_0 = p(\cdot)$. Following this, for each tem-

perature γ_t , a three-step procedure of *importance-sampling*, *resampling*, and *mutation* is evaluated. A typical mutation kernel choice for static parameter inference is a target-invariant MCMC kernel, resulting in importance weights taking the incremental form

$$w_t^{(i)} = W_{t-1}^{(i)} \frac{\eta_t(\boldsymbol{\theta}^{(i)})}{\eta_{t-1}(\boldsymbol{\theta}^{(i)})} = W_{t-1}^{(i)} \mathcal{L}(\mathbf{y}|\boldsymbol{\theta})^{\gamma_t - \gamma_{t-1}},$$

where $\boldsymbol{\theta}^{(i)}$ is the parameter vector for the i^{th} particle, and the term $W_{t-1}^{(i)}$ represent normalised weights from the previous target, computed via

$$W_t^{(i)} = \frac{w_t^{(i)}}{\sum_{j=1}^N w_t^{(j)}},$$

224 where $W_0^{(i)} = \frac{1}{N}$ for all i particles. The resampling step simply draws N new particles $\boldsymbol{\theta}^{(i)*}$
 225 from the weighted representation of particles $\{W_t^{(i)}, \boldsymbol{\theta}^{(i)}\}_{i=1}^N$ and resets the normalised weights
 226 to $W_t^{(i)} = \frac{1}{N}$ for all i particles.

As stated earlier, the mutation step uses a target-invariant MCMC kernel such as several iterations of random-walk Metropolis-Hastings (Hastings, 1970) to perturb each particle. The resulting particle will then be used as input for the next target π_{t+1} . Denoting particles for targets at steps t and $t+1$ as $\boldsymbol{\theta}_t^{(i)}$ and $\boldsymbol{\theta}_{t+1}^{(i)}$ respectively, we write the mutation step as

$$\boldsymbol{\theta}_{t+1}^{(i)} \sim K(\cdot | \boldsymbol{\theta}_t^{(i)*}).$$

This procedure is summarised in Algorithm 2 in Appendix A. A by-product of this algorithm is the marginal likelihood estimate as found by the following trivial computation:

$$\hat{\mathcal{Z}} = \prod_{t=1}^T \sum_{i=1}^N w_t^{(i)}.$$

227 Some of the concerns mentioned at the end of Section 2.2 are addressed by using an SMC al-
 228 gorithm. By virtue of the particle approximation of π_t , proposals can be designed that take
 229 advantage of this information without violating any rules of diminishing adaptation. Secondly,
 230 the computation of $\mathcal{L}(\mathbf{y}|\boldsymbol{\theta})$ for N particles can be performed independently in parallel, thus
 231 lending itself well to a distributed computing architecture.

232 Another advantageous difference between the static SMC framework and traditional Markov
 233 Chain sampling methods is the implicit stopping condition that an adaptive SMC algorithm
 234 provides. Rather than determining the number of MCMC iterations *a-prior*, or relying on con-
 235 vergence criteria, SMC terminates after traversing the sequence of target distributions.

236 Whilst static SMC does produce an unbiased estimate of the normalising constant, it does
 237 require running an instance of the algorithm for each model. This general approach was rec-
 238 ommended by Zhou et al. (2016), but it ignores efficiencies that can be leveraged from exist-
 239 ing research in Bayesian model selection for particular problems. In the case of layered-Earth
 240 models, the RJMCMC proposals first introduced by Malinverno (2002) and subsequently de-
 241 veloped in Dosso et al. (2014) are (assuming non-pathological likelihood conditions) capable
 242 of efficiently traversing a medium-sized model space and sampling high likelihood models more
 243 often than low likelihood models, thereby implicitly introducing a sampling efficiency that in
 244 practice yields posterior model probabilities with low variability. We intend to leverage this
 245 efficiency in the implementation of static SMC with RJMCMC proposals discussed in the next
 246 section, and apply this approach to a model space with cardinality approaching 10^2 . Such a
 247 broad prior of models is commonplace in exploration geophysics problems where there is usu-
 248 ally a wide variation of possible posteriors with very little informative prior knowledge.

249 3.2 Reversible Jump Sequential Monte Carlo

The target density we wish to consider is the joint posterior of models and parameters

$$\pi(\boldsymbol{\theta}_k, k | \mathbf{y}) \propto \mathcal{L}(\mathbf{y} | \boldsymbol{\theta}_k, k) p(\boldsymbol{\theta}_k | k) p(k).$$

Based on the SMC1 archetype identified in Zhou et al. (2016), we specify a static SMC algorithm using a likelihood-annealed sequence of target distributions where the annealing exponent sequence γ_t is monotonically increasing on $[0, 1]$. The below sequence of target distributions forms the basis of the RJSMC algorithm

$$\pi_t(\boldsymbol{\theta}_k, k|\mathbf{y}) \propto \mathcal{L}(\mathbf{y}|\boldsymbol{\theta}_k, k)^{\gamma_t} p(\boldsymbol{\theta}_k|k) p(k), \quad t = 0, \dots, T.$$

If we denote the normalising constant defined in Section 2.2 for target density π_t as $\mathcal{Z}_{t,k}$, we note that the above is proportional to $\sum_{k \in \mathcal{M}} \mathcal{Z}_{t,k}$. For convenience of notation, we introduce k terms $\eta_{t,k}$ for each unnormalised target density

$$\eta_{t,k}(\boldsymbol{\theta}_k) = \mathcal{Z}_{t,k} \pi_t(\boldsymbol{\theta}_k, k|\mathbf{y}) = \mathcal{L}(\mathbf{y}|\boldsymbol{\theta}_k, k)^{\gamma_t} p(\boldsymbol{\theta}_k|k) p(k).$$

Using this, and assuming an RJMCMC kernel is used for particle dynamics, we again choose to approximate the posterior of models and parameters π_t with N particles, and as such we define the importance weights for each particle, indexed by i and for convenience the model k , as

$$w_{t,k}^{(i)} = W_{t-1,k}^{(i)} \frac{\eta_{t,k}(\boldsymbol{\theta}_k^{(i)})}{\eta_k(\boldsymbol{\theta}_k^{(i)})} = W_{t-1,k}^{(i)} \mathcal{L}(\mathbf{y}|\boldsymbol{\theta}_k)^{\gamma_t - \gamma_{t-1}},$$

where we denote the number of particles representing $\pi_{t,k}(\boldsymbol{\theta}|k)$ by $N_{t,k}$. Note that since the algorithm is using particle approximations to such conditional densities, the total number of particles $N = \sum_{k \in \mathcal{M}} N_{t,k}$ should be set such that $N \ll |\mathcal{M}|$. The weights are then normalised such that they approximately represent the conditional density $\pi_t(\boldsymbol{\theta}_{t,k}|k, \mathbf{y})$. These normalised weights are given by

$$W_{t,k}^{(i)} = \frac{w_{t,k}^{(i)}}{\sum_{j=1}^{N_{t,k}} w_{t,k}^{(j)}}.$$

Following the notation from the previous section on single-model SMC, parameters for the i^{th} particle associated with target k at steps t and $t+1$ will henceforth be denoted $\boldsymbol{\theta}_{t,k}^{(i)}$ and $\boldsymbol{\theta}_{t+1,k}^{(i)}$ respectively. At initialisation, after sampling from the joint prior of model and parameters, the normalised weights are uniform, i.e. $W_{t,k}^{(i)} = 1/N_{t,k}$ for all particles $i = 1, \dots, N_{t,k}$ in model k . The same uniform initialisation is applied after the mutation step on the (now different) set of particles that represent the subsequent conditional target density $\pi_{t+1}(\boldsymbol{\theta}_{t+1,k}|k, \mathbf{y})$.

Since we are using an RJMCMC kernel, $N_{t,k}$ will change at the mutation step and as such it is not fixed for all t in $\pi_t(\boldsymbol{\theta}_{t,k}|k, \mathbf{y})$. It is a design choice of this algorithm to keep $N_{t,k}$ constant in the resample step, and a natural way to achieve this is to sample $\boldsymbol{\theta}_{t,k}^{(i)*} \sim \hat{\pi}_t(\cdot|k)$ where $\hat{\pi}_t(\cdot|k)$ is represented by the weighted particles $\{W_{t,k}^{(i)}, \boldsymbol{\theta}_{t,k}^{(i)}\}_{i=1}^{N_{t,k}}$ from model k . We refer to this constraint as *within-model re-sampling*, and by observing this constraint a typical scheme such as multinomial or systematic resampling can be used. The complete adaptive algorithm is summarised in Algorithm 1.

3.2.1 Adaptive Considerations

The configuration of a static SMC sampler over the joint posterior of models and parameters is not well-researched in terms of the sequence of target densities for more than one model and the particle counts between models. In this subsection we will discuss the adaptation of the sequence of targets π_t , and in subsequent subsections we will address the determination of the number of MCMC mutation steps, and overcoming particle impoverishment for the later target densities of unlikely models.

In the single-model static SMC formulation of Schäfer and Chopin (2011), next target density π_{t+1} is specified adaptively using an estimate of the Effective Sample Size (ESS) (Kish, 1965), where a predetermined threshold α is used to find γ_{t+1} such that the resampling phase

Algorithm 1: Adaptive Reversible Jump Sequential Monte Carlo**Input:** TESS threshold proportion α_{TESS} , c probability of mutation**Output:** $\{\theta_{T,k}^{(i)}\}_{i=1}^{N_{t,k}}$ and $\widehat{\mathcal{Z}}_{T,k}$ for $k \in \mathcal{M}$ Initialise particles from the prior for $i = 1, \dots, N, k \sim \pi_0(k) = p(k)$,

$$\theta_{0,k}^{(i)} \sim \pi_0(\cdot|k) = p(\cdot|k).$$

Initialise the particle weights for $i = 1, \dots, N_{t,k}, k \in \mathcal{M}$, $W_{0,k}^{(i)} \leftarrow N_{t,k}^{-1}$.Set $t \leftarrow 0$, $\gamma_0 \leftarrow 0$, and $\widehat{\mathcal{Z}}_{0,k} = 1$ for $k \in \mathcal{M}$.**while** $\gamma_t \neq 1$ **do** Set $t \leftarrow t + 1$ Compute γ_t using a bisection method such that $\widehat{\text{TESS}}_t \approx \alpha_{\text{TESS}} N$ **Re-weight** particles with normalised weights for $i = 1, \dots, N_{t,k}, k \in \mathcal{M}$

$$w_{t,k}^{(i)} \leftarrow W_{t-1,k}^{(i)} \mathcal{L}(\mathbf{y}|\theta_{t,k}^{(i)})^{\gamma_t - \gamma_{t-1}}$$

$$W_{t,k}^{(i)} \leftarrow \frac{w_{t,k}^{(i)}}{\sum_{i=1}^{N_{t,k}} w_{t,k}^{(i)}}$$

$$\widehat{\mathcal{Z}}_{t,k} \leftarrow \widehat{\mathcal{Z}}_{t-1,k} \sum_{i=1}^{N_{t,k}} w_{t,k}^{(i)}$$

Re-sample $N_{t,k}$ particles $\theta_{t,k}^{(i)*} \sim \widehat{\pi}_t(\cdot|k)$ where $\widehat{\pi}_t(\cdot|k)$ is represented by the weighted particles $\{W_{t,k}^{(i)}, \theta_{t,k}^{(i)}\}_{i=1}^{N_{t,k}}$ for $k \in \mathcal{M}$ **Adapt** MCMC and RJMCMC proposals $q_{k \rightarrow k'}$ using weighted power posterior $\{\{\nu_{s,k}^{(i)}, \theta_{s,k}^{(i)}\}_{i=1}^{N_{t,k}}\}_{s=0}^t$ where

$$\nu_{t,s,k}^{(i)} = \frac{\mathcal{L}(\mathbf{y}|\theta_{s,k}^{(i)})^{\gamma_t} p(\theta_{s,k}^{(i)})}{\frac{1}{t} \sum_{l=0}^t \mathcal{L}(\mathbf{y}|\theta_{s,k}^{(i)})^{\gamma_l} p(\theta_{s,k}^{(i)}) (\widehat{\mathcal{Z}}_{l,k})^{-1}} \quad \text{for } k \in \mathcal{M}$$

 Compute $R_t = \left\lceil \frac{\log c}{\log(1-p_{\text{acc}}^{\min})} \right\rceil$ using trial mutations **Mutate** particles R_t times using an RJMCMC kernel with the n_q adapted proposals

$$k' \sim q_{\mathcal{M}}(\cdot|k),$$

$$\theta_{t+1,k'}^{(i)} \sim q_{k \rightarrow k'}(\cdot|\theta_{t,k}^{(i)*}, k')$$

end

273 is initiated when the ESS of π_{t+1} falls below αN . Such a scheme can be implemented naïvely
 274 by choosing a small step size δ (where $\gamma_{t+1} = \gamma_t + \delta$) and simply incrementing t until the
 275 threshold is reached. Alternatively, a bisection method can be used to determine the next γ_{t+1}
 276 Jasra et al. (2011).

In the presence of multiple models, the ESS estimate for the conditioned density $\pi_t(\cdot|k)$
 is

$$\widehat{\text{ESS}}_{t,k} = \frac{1}{\sum_{i=1}^{N_{t,k}} (W_{t,k}^{(i)})^2}.$$

If there are K models where $K > 1$, a new condition is required for when to stop and re-sample/mutate since there are now K ESS estimates. Given a threshold $0 < \alpha < 1$, such a condition can be met by taking a statistic of $\widehat{\text{ESS}}_{t,k}$. A naïve approach would be to take the threshold condition $\min_k (\widehat{\text{ESS}}_{t,k}/N_{t,k}) < \alpha$, but this exposes the algorithm to high variability of $\widehat{\text{ESS}}_{t,k}$ in unlikely models. For this reason, we choose to have the threshold condition be dominated by the more likely models. First, we need to define the *normalised* effective sample size (NESS) for the conditional target density $\pi_t(\boldsymbol{\theta}|\mathbf{y}, k)$ as

$$\text{NESS}_{t,k} = \frac{1}{N_{t,k}} \text{ESS}_{t,k}.$$

Using this quantity, we choose our threshold condition to use what we define as the Total Effective Sample Size (TESS), which expressed in terms of $\text{NESS}_{t,k}$ is

$$\mathbb{E}[\text{TESS}_t] := N \cdot \mathbb{E}[\text{NESS}_{t,k}].$$

It is shown in Appendix C that an estimate of the TESS is simply

$$\widehat{\text{TESS}}_t = \sum_{k \in \mathcal{M}} \widehat{\text{ESS}}_{t,k}.$$

277 The algorithm will resample when $\widehat{\text{TESS}}_t < \alpha N$. Intuitively, one can reason that in low-
 278 likelihood models the variability of $\text{ESS}_{t,k}$ would increase, but since such models are repre-
 279 sented by proportionally fewer particles this variability does not translate to increased vari-
 280 ability in the overall TESS. As such, this approach is in practice robust to low particle counts
 281 in the presence of unlikely models.

An estimate of the ratio of normalising constants between successive target distributions in an SMC algorithm with an MCMC kernel is given by the sum of the weights. When considering the same quantity for an RJSMC algorithm, we show in Appendix D that the normalising constant estimate $\widehat{\mathcal{Z}}_{T,k}$ reduces to the same form considering particle weights for model k . This gives us

$$\widehat{\mathcal{Z}}_{T,k} = \prod_{t=0}^T \frac{\widehat{\mathcal{Z}}_{t,k}}{\mathcal{Z}_{t-1,k}} = \prod_{t=0}^T \sum_{i=1}^{N_{t,k}} w_{t,k}^{(i)},$$

noting that $\mathcal{Z}_{0,k} = 1$. From here it is possible to compute estimates of the ratios of normalising constants between models using

$$\widehat{\mathcal{B}}_{1,2} = \frac{\widehat{\mathcal{Z}}_{T,k_1}}{\widehat{\mathcal{Z}}_{T,k_2}}.$$

Since we are using an RJMCMC kernel, another approach (which is adopted in this work) to estimate ratios of normalising constants is to take ratios of the empirical posterior model marginal density weighted by prior densities

$$\widehat{\mathcal{B}}_{1,2} = \frac{\widehat{\pi}_T(k_1|\mathbf{y}) p(k_2)}{\widehat{\pi}_T(k_2|\mathbf{y}) p(k_1)},$$

where

$$\widehat{\pi}_T(k|\mathbf{y}) = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(k^{(i)} = k) = \frac{N_{T,k}}{N},$$

282 noting $k^{(i)}$ is the model indicator for particle i at step T . Other precise quantities relating to
 283 ratios of normalising constants derived from RJMCMC posteriors are explored in Bartolucci
 284 et al. (2006). The motivating reason for adopting these latter approaches is to, if possible, re-
 285 duce the posterior model-odds variability. This places the onus of reducing such variability
 286 on the RJMCMC proposal performance rather than the performance of the SMC sampler as
 287 a whole, which is a topic that would require further research as it is not central to this arti-
 288 cle. The form of the MCMC and RJMCMC proposals for the AIP inference application are
 289 discussed in the next subsection.

3.2.2 Transdimensional and within-model adaptive proposals

For this work, within-model MCMC proposals are an adaptive component-wise random-walk Gaussian proposal distribution which utilises the availability of particles approximating $\pi_{t,k}$ to adapt the search direction and step size.

Typical component-wise proposals treat the parameter vector as independent components and randomly choose a component and then sample a value for that component. This approach relies on well-conditioned posterior for good performance. In a layered Earth model of electrical properties, the parameter space generally yields a pathological posterior, hence a naïve approach is inefficient. A simple approximate fit of a multivariate Gaussian $\mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$ to the particles $\boldsymbol{\theta}^{(i)}$, $i = 1, \dots, N$ approximating π_t opens the ability to use principle components as search directions. Taking the eigenvalue decomposition $\boldsymbol{\Sigma}_t = \mathbf{U}\boldsymbol{\Lambda}\mathbf{U}^T$ where $\boldsymbol{\Lambda}$ is diagonal, our component-wise proposal selects a random component/column $\mathbf{U}_r$ for the search direction and scales the step size by the corresponding eigenvalue $\boldsymbol{\Lambda}_r$. A bisection algorithm on trial proposals is then used to further scale proposals to target a desired acceptance rate. In this work we used 0.44 since it is accepted as the optimal acceptance rate for component-wise sampling of a standard multivariate Gaussian target (Roberts & Rosenthal, 2001).

The cross-dimensional proposals use the RJMCMC framework, first proposed by P. J. Green (1995) and developed for the 1D layered-Earth model by Malinverno (2002). The form of the proposals used in this work extend the above using a simple adaptive design constructed in Appendix D. Two separate proposals are used, each mapping auxiliary variables to Earth property parameters $\theta_{\phi,i}$ and $\theta_{m,j}$ respectively.

3.2.3 Adapting the number of MCMC steps for multiple proposals

There is no set procedure for determining the number of MCMC iterations for effective mutation. Taking the approach introduced by Drovandi and Pettitt (2011), the number of mutation steps can be determined simply as a function of the acceptance rate $R(p_{\text{acc}})$ for a single proposal and a tuning parameter c . However, for RJSMC there are multiple proposal types as a result of using a birth/death RJMCMC pair of proposals and a within-model MCMC proposal. Fortunately, by simple extension of the approach taken in Drovandi and Pettitt (2011), it is shown in Appendix B that for n_q proposals, the minimum acceptance rate

$$p_{\text{acc}}^{\min} = \min_{j \in n_q} p_{\text{acc}}^{(j)}$$

will determine the minimum number of mutations required to ensure that all particles mutate with probability $1 - c$. This results in the formula for determining the minimum number of mutation steps to be

$$R = \left\lceil \frac{\log c}{\log(1 - p_{\text{acc}}^{\min})} \right\rceil.$$

Other more recent work by Bon et al. (2021) adopts a generalised approach which adaptively chooses a proposal step size based on a target expected squared jumping distance (ESJD). Such an approach could be examined in the context of RJSMC samplers in future research.

3.2.4 Overcoming particle impoverishment in unlikely models

Cross-dimensional proposals that target $\pi_t(\boldsymbol{\theta}_k, k|\mathbf{y})$ will favour models with higher probability. Therefore, in a particle approximation of π_t where N is fixed, examples will occur where unlikely models are represented by few or no particles. In the case where naïve (RJ)MCMC proposals (i.e. those that do not adapt to π_t) are used, this is not necessarily an issue. However, in order to take advantage of the availability of π_t for proposal tuning (see section 3.2.2 and Appendix D for examples in this application), particle impoverishment becomes an important issue.




Figure 3: Shown are four case studies employing a three-layered conductivity model, each with varying chargeability and depth of the second layer, and the other layers having zero chargeability. First and second layer conductivities are 0.01 S/m and 0.1 S/m respectively, second layer thickness is 20m, and the conductivities of the the halfspace/basement layer (shown above each heatmap) reflect typical geological examples. BFIPD values above 1 indicate higher likelihood of chargeability. Within each study the heatmaps show how depth and magnitude of the chargeable layer affects the respective BFIPD statistic for each model posterior. Between studies the heatmaps show how increasing halfspace conductivity relative to the chargeable layer reduces BFIPD statistics.

One approach to alleviate particle impoverishment is to re-use particles from previous target densities. Since we are considering a likelihood-annealed sequence of distributions, it is natural to consider the deterministic mixture recycling approach employed by Nguyen et al. (2015) and further implemented in SMC by South et al. (2019).

The particle mixture weights representing the power-posterior at step t for particles $i = 1, \dots, N_{t,k}$, steps $s = 0, \dots, t$, and models $k \in \mathcal{M}$ are

$$\nu_{t,s,k}^{(i)} = \frac{\mathcal{L}(\mathbf{y}|\boldsymbol{\theta}_{s,k}^{(i)})^{\gamma_t} p(\boldsymbol{\theta}_{s,k}^{(i)})}{\frac{1}{t} \sum_{l=0}^t \mathcal{L}(\mathbf{y}|\boldsymbol{\theta}_{s,k}^{(i)})^{\gamma_l} p(\boldsymbol{\theta}_{s,k}^{(i)}) (\hat{\mathcal{Z}}_{l,k})^{-1}}$$

Since the normalising constant $\mathcal{Z}_{t,k}$ for the power posterior π_t is not available, we use the RJSMC estimate $\hat{\mathcal{Z}}_{t,k}$.

It is important to note that the above form for recycled particle weights is conditional on the model k , meaning that these weighted particles are suited only to fitting such conditional proposal densities and are not valid for representing the joint posterior of parameters and models.

Whilst particle recycling has been used offline for estimation of posterior statistics, an online implementation for estimation of intermediate densities for proposal tuning is not known to the authors, and is thus proposed here as a novel step to mitigate particle impoverishment during the course of the algorithm.

4 Case Studies

This section will demonstrate the application of parameter and model inference via the RJSMC algorithm as applied to various synthetic and real-data examples. It will begin with a comprehensive synthetic case study demonstrating the performance of Bayes Factor detectability of chargeability, followed by inference for IP detectability in a 2D ground section of AEM line data.

4.1 Synthetic Studies

A comprehensive series of synthetic studies were designed to demonstrate targeted quantities that are generated by the application of RJSMC to AEM data with IP effects. The quantities of interest are detectability of chargeability, recoverability of model parameters, and goodness of fit.

For investigation of IP detectability using the BFIPD statistic from Section 2.3.3, a three-layered model with chargeability in only the middle layer was selected as basis of four almost identical case studies where only the basement conductivity varied across each of the four studies. The basement conductivities were chosen to simulate four common scenarios: a strongly resistive igneous/metamorphic basement (0.001 S/m), a hard sedimentary basement (0.01 S/m), a moderately-conductive sedimentary basement (0.032 S/m), and a conductive basement (0.1 S/m). The conductivity of the upper two layers were held constant in all cases, being 0.01 S/m and 0.1 S/m for the first and second layer respectively. The time constant and frequency dependence parameters were set to $\theta_\tau = 4.07 \times 10^{-4}$ and $\theta_c = 1.0$ respectively. The data was generated using the same VTEM ET AEM system configuration as used in Zamudio et al. (2021) and the data noise was simulated using the noise model discussed in Section 2.3.2.

Within each case study, both the depth of the upper interface and the chargeability of the second layer were jointly varied such that the BFIPD statistic could be examined as a function of the interaction of these two parameters. The BFIPD statistics were visualised as a heatmap in Figure 3. As would be expected the BFIPD statistic decreases with depth of the second layer and increases with the magnitude of second layer chargeability. Also, it can be seen that across the four studies, variation in the basement conductivity significantly affects the BFIPD statistic, where strongly resistive basement materials admit high BFIPD statistics, ranging down to low BFIPD statistics for conductive sedimentary basement materials such as shales, clays, or aquifers.

The recoverability of model parameters is well-known to be confounded by a phenomenon geophysicists term “non-uniqueness” (Ellis, 1998). In the language of Bayesian statistics, this translates to the situation where the *maximum a posteriori* (MAP) model is not necessarily close to the *data-generating* model. This is frequently seen in conductivity-only inversion of AEM data and thus it is expected to be present in Bayesian AIP inference. For this investigation, we compared a selection of the model-averaged posterior densities from the four case studies outlined above against the data-generating parameters. Figure 4 shows three selected data sets from the above detectability study. The posterior is summarised in two model-averaged plots showing depth versus conductivity and depth versus intrinsic chargeability, and the “true” data-generating model is shown on each posterior plot with a dark-red dashed line. It can be seen that the shallow layers are generally well-recovered, however the deeper layers and layer interfaces of the data-generating model are not usually reflected in the high-probability regions of the posterior. This phenomenon is well-understood for conductive-only models (Ellis, 1998) but the extent to which it is present in models with chargeability is not well-quantified and should be the subject of further research.

A canonical Bayesian form of goodness of fit uses the *posterior predictive distribution* (PPD) (Gelman et al., 2013), which has the mathematical form

$$p(\mathbf{y}|\tilde{\mathbf{y}}) = \int p(\tilde{\mathbf{y}}|\boldsymbol{\theta})\pi(\mathbf{y}|\boldsymbol{\theta})d\boldsymbol{\theta}.$$

A visual inspection of the PPD constitutes a *posterior predictive check* (PPC). PPCs for posteriors generated from three synthetic data sets are shown in Figure 4. This approach is straightforward for individual soundings; however, it can be cumbersome when examining PPDs of the thousands of soundings in a single AEM line. In such cases, it is more feasible to display summary statistics of the PPD, such as the sample mean and variance, and this approach can be seen for the Colorado study in Figure 5.




Figure 4: Three figures showing data selected from the detectability synthetic study (Figure 3) to represent posterior predictive checks and parameter posteriors from data generating models with various intrinsic chargeability and depth parameters. *Left:* Posterior predictive check plots where posterior samples are drawn from particles in the final target density of each inversion. *Right:* Joint posterior density of the inversion of synthetic data. The true (data-generating) model parameters are identified by the dark-red dashed line in each plot. In the third figure showing the posterior of a conductive-only model the recoverability of conductivity parameters at greater depths is visibly more successful than when chargeability is present in the data-generating model.

4.1.1 Depth of Investigation

A Bayesian approach to computing the level of information of the data with respect to depth, known as the depth of investigation (DOI), was developed by Blatter et al. (2018) using entropy to measure information gain. By evaluating the continuous conductivity or chargeability parameters at discrete intervals with respect to depth, the depth of investigation divergences are computed at 1-metre intervals. It is necessary to choose a suitable divergence threshold to mask the posterior in regions that are considered to be occluded; for this work a threshold of 1 was chosen to work well in most scenarios. Figure 4 shows the DOI for the intrinsic chargeability vs depth averaged model posterior in each case. The DOI for conductivity was below the a priori maximum interface depth, indicating that in these cases the resolution of conductivity was greater than that of chargeability.

4.2 Colorado Study

The above methods were applied to parameter and model inference on one flight-line of AEM survey data from Colorado, USA (B. J. Minsley and Ball (2018), Zamudio et al. (2021)). Notably, there was a relatively low particle count (9600 particles per sounding) required to sufficiently permit model and parameter inference on the resultant posteriors. The selected data visibly contained negative magnetic flux density measurements that cannot be modelled using non-chargeable physics. Data was spatially decimated so that soundings are spaced by approximately 30 metres and sections of the mean summaries are shown in Figure 5. Although there is no published ground truth survey for this region, it is expected that the continuous section of non-zero intrinsic chargeability in flight line 11810 is due to the frequent occurrence of disseminated pyrite present in the Mancos shale formations consistent with the surrounding geology (Vanderwilt, 1937).

This flight line of data was also the subject of an earlier study by Viezzoli et al. (2019). In their study, the authors compare Cole-Cole IP ground parameters, recovered via a Tikhonov regression inversion, to a geology map of the region to demonstrate spatial consistency with known geological structures in the area. Notably, their findings for MAP parameters were consistent with the results in this work, however their study did not take into account parameter and model uncertainty and did not attempt to discriminate data on the basis of the presence of induced polarisation effects.

5 Conclusions

In this study we demonstrated the effectiveness of Bayesian parameter inference and model inference, the latter specifically for inference on the likelihood of non-zero intrinsic chargeability when the number of layers in conductivity and chargeability depth profiles are unknown a-priori. We have developed an SMC algorithm for inference of parameters and model probabilities that can exploit practically any parallel computing architecture, independent of the configuration of the model space and the number of particles, and is robust in the presence of pathological posteriors. In synthetic studies we have shown how well chargeability can be detected via the BFIPD statistic for a given AEM system configuration. In real data examples we have demonstrated how the BFIPD statistic is spatially consistent in a line of AEM data.

Future research stemming from this work could take many avenues. Geophysics practitioners may delve into the parameterisation and/or the assumptions of the existing noise model that was discussed in Section 2.3.2. It would also be of interest to investigate the computational efficiency of cross-dimensional proposals considering that such proposal design can be enhanced with the availability of particles approximating π_t . There are potentially other efficiencies that could be implemented in the SMC algorithm itself which could quickly make it computationally competitive with well-established Bayesian inference software for geophysics. The Python 3 code that was developed for this research is part of a larger framework under




Figure 5: *Top:* Mean posterior summaries of each sounding for conductivity and chargeability parameter posteriors given VTEM ET data from the Colorado USA AEM survey. Each mean posterior summary is computed using the mean of the model-averaged parameter posterior versus depth. The third panel shows the Bayes Factor Induced Polarisation Detectability (BFIPD) statistic, which is the expected non-chargeable model probability versus the expected chargeable model probability. The fourth panel displays a summary of each data set standardised to the respective posterior predictive distribution (PPD) such that any significant deviations of the data at each gate time with respect to the PPD can be easily identified. Values between -1 and 1 fall within the standard deviations of the PPD. A selected sounding is highlighted with a red rectangle for further inspection in the bottom figure. *Bottom:* Posterior predictive check and posterior vs depth plot of conductivity and chargeability for data highlighted in the red rectangle in the top figure.

development and will be released in the near future, accessible via <http://www.terrascope.com.au/> and as a repository at <https://github.com/daviesl/>.

Acknowledgments

This work was financially supported by the Australian Government Research Training Program, the Queensland University of Technology Centre for Data Science, and an Australian Research Council Discovery Project (DP200102101). Computing resources were provided by Queensland University of Technology. The authors would like to give special thanks to Dr Andrea Viezzoli of Aarhus Geophysics for crucial advice concerning the Colorado data, and to Dr Ross Brodie of Geoscience Australia for providing the forward model code for IP in a TDEM system (see the development branch in Brodie (2016)).

AEM Data from the Colorado case study was first presented in (B. J. Minsley & Ball, 2018) and will be available under Zamudio et al. (2021) as an entry in the data repository website <http://www.sciencebase.gov>.

References

- Akaike, H. (1974, December). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6), 716–723. (Conference Name: IEEE Transactions on Automatic Control) doi: 10.1109/TAC.1974.1100705
- Amaya, M., Linde, N., & Laloy, E. (2021, August). Adaptive sequential Monte Carlo for posterior inference and model selection among complex geological priors. *Geophysical Journal International*, 226(2), 1220–1238. Retrieved 2021-05-21, from <https://doi.org/10.1093/gji/ggab170> doi: 10.1093/gji/ggab170
- Bartolucci, F., Scaccia, L., & Mira, A. (2006). Efficient Bayes Factor Estimation from the Reversible Jump Output. *Biometrika*, 93(1), 41–52. Retrieved 2020-07-13, from <https://www.jstor.org/stable/20441258> (Publisher: [Oxford University Press, Biometrika Trust])
- Blatter, D., Key, K., Ray, A., Foley, N., Tulaczyk, S., & Auken, E. (2018, July). Trans-dimensional Bayesian inversion of airborne transient EM data from Taylor Glacier, Antarctica. *Geophysical Journal International*, 214(3), 1919–1936. Retrieved 2021-04-13, from <https://academic.oup.com/gji/article/214/3/1919/5043225> doi: 10.1093/gji/ggy255
- Bodin, T., Sambridge, M., Tkalčić, H., Arroucau, P., Gallagher, K., & Rawlinson, N. (2012). Transdimensional inversion of receiver functions and surface wave dispersion. *Journal of Geophysical Research: Solid Earth*, 117(B2). Retrieved 2020-09-23, from <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2011JB008560> (.eprint: <https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2011JB008560>) doi: 10.1029/2011JB008560
- Bon, J. J., Lee, A., & Drovandi, C. (2021, August). Accelerating sequential Monte Carlo with surrogate likelihoods. *Statistics and Computing*, 31(5), 62. Retrieved 2021-08-19, from <https://doi.org/10.1007/s11222-021-10036-4> doi: 10.1007/s11222-021-10036-4
- Brodie, R. C. (2016). *Geoscience Australia Airborne Electromagnetics Programs*. Geoscience Australia. Retrieved from <https://github.com/GeoscienceAustralia/ga-aem>
- Brodie, R. C., & Sambridge, M. (2012, December). Transdimensional Monte Carlo Inversion of AEM Data. *ASEG Extended Abstracts*, 2012(1), 1–4. Retrieved 2020-05-21, from <https://www.tandfonline.com/doi/full/10.1071/ASEG2012ab095> doi: 10.1071/ASEG2012ab095
- Chopin, N. (2002, August). A sequential particle filter method for static models. *Biometrika*, 89(3), 539–552. Retrieved 2019-03-28, from <https://academic.oup.com/>

- biomet/article/89/3/539/251804 doi: 10.1093/biomet/89.3.539
- Cole, K. S., & Cole, R. H. (1941, April). Dispersion and Absorption in Dielectrics I. Alternating Current Characteristics. *The Journal of Chemical Physics*, 9(4), 341–351. Retrieved 2020-06-29, from <http://aip.scitation.org/doi/10.1063/1.1750906> doi: 10.1063/1.1750906
- Dai, C., Heng, J., Jacob, P. E., & Whiteley, N. (2020, July). An invitation to sequential Monte Carlo samplers. *arXiv:2007.11936 [stat]*. Retrieved 2021-04-13, from <http://arxiv.org/abs/2007.11936> (arXiv: 2007.11936)
- Del Moral, P., Doucet, A., & Jasra, A. (2006). Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(3), 411–436. Retrieved 2019-04-04, from <https://rss.onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-9868.2006.00553.x> doi: 10.1111/j.1467-9868.2006.00553.x
- Dettmer, J., Dosso, S. E., & Holland, C. W. (2011, April). Sequential trans-dimensional Monte Carlo for range-dependent geoacoustic inversion. *The Journal of the Acoustical Society of America*, 129(4), 1794–1806. Retrieved 2020-12-11, from <http://asa.scitation.org/doi/10.1121/1.3557052> doi: 10.1121/1.3557052
- Dosso, S., Dettmer, J., Steininger, G., & Holland, C. (2014, October). Efficient trans-dimensional Bayesian inversion for geoacoustic profile estimation. *Inverse Problems*, 30, 114018. doi: 10.1088/0266-5611/30/11/114018
- Drovandi, C. C., & Pettitt, A. N. (2011, March). Estimation of Parameters for Macroparasite Population Evolution Using Approximate Bayesian Computation. *Biometrics*, 67(1), 225–233. Retrieved 2019-08-12, from <http://doi.wiley.com/10.1111/j.1541-0420.2010.01410.x> doi: 10.1111/j.1541-0420.2010.01410.x
- Eadie, T., Legault, J. M., Plastow, G., Prikhodko, A., & Tishin, P. (2018, December). VTEM ET: An improved helicopter time-domain EM system for near surface applications. *ASEG Extended Abstracts*, 2018(1), 1–5. Retrieved 2021-08-26, from https://doi.org/10.1071/ASEG2018abW9_3H (Publisher: Taylor & Francis eprint: https://doi.org/10.1071/ASEG2018abW9_3H) doi: 10.1071/ASEG2018abW9_3H
- Ellis, R. G. (1998, March). Inversion of airborne electromagnetic data. *Exploration Geophysics*, 29(1-2), 121–127. Retrieved 2020-06-18, from <https://www.tandfonline.com/doi/full/10.1071/EG998121> doi: 10.1071/EG998121
- Fearnhead, P., & Taylor, B. M. (2013, June). An Adaptive Sequential Monte Carlo Sampler. *Bayesian Analysis*, 8(2), 411–438. Retrieved 2020-07-21, from <https://projecteuclid.org/euclid.ba/1369407558> (Publisher: International Society for Bayesian Analysis) doi: 10.1214/13-BA814
- Feng, L., Li, Q., Cameron, S. D., He, K., Colby, R., Walker, K. M., ... Ertaş, D. (2020, February). Quantifying Induced Polarization of Conductive Inclusions in Porous Media and Implications for Geophysical Measurements. *Scientific Reports*, 10(1), 1669. Retrieved 2020-06-26, from <https://www.nature.com/articles/s41598-020-58390-z> (Number: 1 Publisher: Nature Publishing Group) doi: 10.1038/s41598-020-58390-z
- Fiandaca, G., Madsen, L. M., & Maurya, P. K. (2018). Re-parameterisations of the Cole–Cole model for improved spectral inversion of induced polarization data. *Near Surface Geophysics*, 16(4), 385–399. Retrieved 2020-07-02, from <https://onlinelibrary.wiley.com/doi/abs/10.3997/1873-0604.2017065> (eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.3997/1873-0604.2017065>) doi: 10.3997/1873-0604.2017065
- Gelman, A., Carlin, J., Stern, H., Dunson, D., Vehtari, A., & Rubin, D. (2013). *Bayesian Data Analysis, Third Edition*. Taylor & Francis. Retrieved from <https://books.google.com.au/books?id=ZXL6AQAAQBAJ>
- Ghorbani, A., Camerlynck, C., Florsch, N., & Cosenza, P. (2007, July). Bayesian inference of the Cole–Cole parameters from time-and frequency-domain induced polarization. *Geophysical Prospecting*, 55, 589–605. doi: 10.1111/j.1365-2478.2007.00627.x
- Girolami, M., & Calderhead, B. (2011). Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Se-*

- ries *B (Statistical Methodology)*, 73(2), 123–214. Retrieved 2021-06-11, from <https://rss.onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-9868.2010.00765.x> (eprint: <https://rss.onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-9868.2010.00765.x>) doi: 10.1111/j.1467-9868.2010.00765.x
- Green, A., & Lane, R. (2003, August). Estimating Noise Levels in AEM Data. *ASEG Extended Abstracts*, 2003(2), 1–5. Retrieved 2020-05-21, from <https://www.tandfonline.com/doi/full/10.1071/ASEG2003ab093> doi: 10.1071/ASEG2003ab093
- Green, P. J. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika*, 22.
- Haario, H., Saksman, E., & Tamminen, J. (2001). An Adaptive Metropolis Algorithm. *Bernoulli*, 7(2), 223–242. Retrieved 2021-06-13, from <https://www.jstor.org/stable/3318737> (Publisher: International Statistical Institute (ISI) and Bernoulli Society for Mathematical Statistics and Probability) doi: 10.2307/3318737
- Hastings, W. K. (1970, April). Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1), 97–109. Retrieved 2020-08-07, from <https://academic.oup.com/biomet/article/57/1/97/284580> (Publisher: Oxford Academic) doi: 10.1093/biomet/57.1.97
- Hawkins, R., Brodie, R. C., & Sambridge, M. (2017, February). Trans-dimensional Bayesian inversion of airborne electromagnetic data for 2D conductivity profiles. *Exploration Geophysics*, 49(2), 134–147. Retrieved 2020-06-18, from <https://www.publish.csiro.au/eg/EG16139> (Publisher: CSIRO PUBLISHING) doi: 10.1071/EG16139
- Jasra, A., Doucet, A., Stephens, D. A., & Holmes, C. C. (2008, January). Interacting sequential Monte Carlo samplers for trans-dimensional simulation. *Computational Statistics & Data Analysis*, 52(4), 1765–1791. Retrieved 2019-03-12, from <https://linkinghub.elsevier.com/retrieve/pii/S0167947307003398> doi: 10.1016/j.csda.2007.09.009
- Jasra, A., Stephens, D. A., Doucet, A., & Tsagaris, T. (2011, March). Inference for Lévy-Driven Stochastic Volatility Models via Adaptive Sequential Monte Carlo: Lévy-driven stochastic volatility. *Scandinavian Journal of Statistics*, 38(1), 1–22. Retrieved 2021-08-26, from <https://onlinelibrary.wiley.com/doi/10.1111/j.1467-9469.2010.00723.x> doi: 10.1111/j.1467-9469.2010.00723.x
- Jasra, A., Stephens, D. A., & Holmes, C. C. (2007, August). Population-Based Reversible Jump Markov Chain Monte Carlo. *Biometrika*, 94(4), 787–807. Retrieved 2019-02-25, from <https://academic.oup.com/biomet/article-lookup/doi/10.1093/biomet/asm069> doi: 10.1093/biomet/asm069
- Kang, S., Oldenburg, D. W., & Heagy, L. J. (2019, September). Detecting induced polarization effects in time-domain data: a modeling study using stretched exponentials. *arXiv:1909.12993 [physics]*. Retrieved 2021-04-13, from <http://arxiv.org/abs/1909.12993> (arXiv: 1909.12993)
- Kass, R. E., & Raftery, A. E. (1995, June). Bayes Factors. *Journal of the American Statistical Association*, 90(430), 773. Retrieved 2019-03-28, from <https://www.jstor.org/stable/2291091?origin=crossref> doi: 10.2307/2291091
- Kish, L. (1965). *Survey sampling*. New York: John Wiley & Sons. (OCLC: 256017)
- Kratzer, T., & Macnae, J. C. (2012, September). Induced polarization in airborne EM. *GEO-PHYSICS*, 77(5), E317–E327. Retrieved 2020-06-29, from <http://library.seg.org/doi/10.1190/geo2011-0492.1> doi: 10.1190/geo2011-0492.1
- Malinverno, A. (2002, December). Parsimonious Bayesian Markov chain Monte Carlo inversion in a nonlinear geophysical problem. *Geophysical Journal International*, 151(3), 675–688. Retrieved 2020-06-18, from <https://academic.oup.com/gji/article-lookup/doi/10.1046/j.1365-246X.2002.01847.x> doi: 10.1046/j.1365-246X.2002.01847.x
- Meju, M. A. (2002, March). Geoelectromagnetic Exploration For Natural Resources: Models, Case Studies And Challenges. *Surveys in Geophysics*, 23(2), 133–206. Retrieved

- 2021-08-30, from <https://doi.org/10.1023/A:1015052419222> doi: 10.1023/A:1015052419222
- Merriam, J. B. (2007, July). Induced polarization and surface electrochemistry. *GEOPHYSICS*, 72(4), F157–F166. Retrieved 2021-08-30, from <https://library.seg.org/doi/10.1190/1.2732554> (Publisher: Society of Exploration Geophysicists) doi: 10.1190/1.2732554
- Minsley, B. (2011, August). A trans-dimensional Bayesian Markov chain Monte Carlo algorithm for model assesment using frequency-domain electromagnetic data. *Geophysical Journal International*, 187, 252–272. doi: 10.1111/j.1365-246X.2011.05165.x
- Minsley, B. J., & Ball, L. B. (2018). Airborne geophysical characterization of geologic structure in a mountain headwater system, upper East River, Colorado.. Retrieved 2021-04-12, from <https://pubs.er.usgs.gov/publication/70217684>
- Nguyen, T., Septier, F., Peters, G., & Yves, D. (2015, April). Efficient Sequential Monte-Carlo Samplers for Bayesian Inference. *IEEE Transactions on Signal Processing*, 64. doi: 10.1109/TSP.2015.2504342
- Nunes, C. M. B., R´egis, C., & de Lima, O. A. L. (2019, August). Detectability of induced polarization effects on the frequency domain fields of a large loop. In *SEG Technical Program Expanded Abstracts 2019* (pp. 1159–1163). San Antonio, Texas: Society of Exploration Geophysicists. Retrieved 2020-06-29, from <https://library.seg.org/doi/10.1190/segam2019-3215024.1> doi: 10.1190/segam2019-3215024.1
- Pelton, W. H., Ward, S. H., Hallof, P. G., Sill, W. R., & Nelson, P. H. (1978, April). MINERAL DISCRIMINATION AND REMOVAL OF INDUCTIVE COUPLING WITH MULTIFREQUENCY IP. *GEOPHYSICS*, 43(3), 588–609. Retrieved 2020-07-02, from <http://library.seg.org/doi/10.1190/1.1440839> doi: 10.1190/1.1440839
- Piana Agostinetti, N., & Bodin, T. (2018, October). Flexible Coupling in Joint Inversions: A Bayesian Structure Decoupling Algorithm. *Journal of Geophysical Research. Solid Earth*, 123(10), 8798–8826. Retrieved 2020-12-30, from <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6282997/> doi: 10.1029/2018JB016079
- Qi, Y., Soueid Ahmed, A., Revil, A., Ghorbani, A., Abdulsamad, F., Florsch, N., & Bonnenfant, J. (2018, September). Induced polarization response of porous media with metallic particles — Part 7: Detection and quantification of buried slag heaps. *GEOPHYSICS*, 83(5), E277–E291. Retrieved 2021-08-30, from <https://library.seg.org/doi/10.1190/geo2017-0760.1> (Publisher: Society of Exploration Geophysicists) doi: 10.1190/geo2017-0760.1
- Roberts, G. O., & Rosenthal, J. S. (2001). Optimal Scaling for Various Metropolis-Hastings Algorithms. *Statistical Science*, 16(4), 351–367. Retrieved 2021-05-04, from <https://www.jstor.org/stable/3182776> (Publisher: Institute of Mathematical Statistics)
- Roberts, G. O., & Stramer, O. (2002, December). Langevin Diffusions and Metropolis-Hastings Algorithms. *Methodology And Computing In Applied Probability*, 4(4), 337–357. Retrieved 2021-07-30, from <https://doi.org/10.1023/A:1023562417138> doi: 10.1023/A:1023562417138
- Schäfer, C., & Chopin, N. (2011, November). Sequential Monte Carlo on large binary sampling spaces. *arXiv:1101.6037 [math, stat]*. Retrieved 2020-07-21, from <http://arxiv.org/abs/1101.6037> (arXiv: 1101.6037)
- Seigel, H. O. (1959, July). Mathematical formulation and type curves for induced polarization. *GEOPHYSICS*, 24(3), 547–565. Retrieved 2020-06-18, from <https://library.seg.org/doi/10.1190/1.1438625> (Publisher: Society of Exploration Geophysicists) doi: 10.1190/1.1438625
- South, L. F., Pettitt, A. N., & Drovandi, C. C. (2019, September). Sequential Monte Carlo Samplers with Independent Markov Chain Monte Carlo Proposals. *Bayesian Analysis*, 14(3), 753–776. Retrieved 2020-09-14, from <https://projecteuclid.org/euclid.ba/1560240027> doi: 10.1214/18-BA1129

- Swendsen, R., & Wang, J.-S. (1986, December). Replica Monte Carlo Simulation of Spin-Glasses. *Physical review letters*, 57, 2607–2609. doi: 10.1103/PhysRevLett.57.2607
- Tarasov, A., & Titov, K. (2013, October). On the use of the Cole–Cole equations in spectral induced polarization. *Geophysical Journal International*, 195(1), 352–356. Retrieved 2020-07-01, from <https://academic.oup.com/gji/article/195/1/352/608470> (Publisher: Oxford Academic) doi: 10.1093/gji/ggt251
- Vanderwilt, J. W. (1937). *Geology and mineral deposits of the Snowmass Mountain area, Gunnison County, Colorado* (Bulletin No. 884-886). United States Geological Survey. Retrieved 2021-08-26, from <http://pubs.er.usgs.gov/publication/b884> doi: 10.3133/b884
- Viezzoli, A., Dauti, F., & Wijns, C. (2021, January). Robust scanning of AEM data for IP effects. *Exploration Geophysics*, 0(0), 1–12. Retrieved 2021-02-19, from <https://doi.org/10.1080/08123985.2020.1856624> (Publisher: Taylor & Francis _eprint: <https://doi.org/10.1080/08123985.2020.1856624>) doi: 10.1080/08123985.2020.1856624
- Viezzoli, A., & Manca, G. (2020, January). On airborne IP effects in standard AEM systems: tightening model space with data space. *Exploration Geophysics*, 51(1), 155–169. Retrieved 2021-03-16, from <https://doi.org/10.1080/08123985.2019.1681895> (Publisher: Taylor & Francis _eprint: <https://doi.org/10.1080/08123985.2019.1681895>) doi: 10.1080/08123985.2019.1681895
- Viezzoli, A., Minsley, B., & Menghini, A. (2019, December). Insight from AIP modelling of VTEM ET data from Colorado. *ASEG Extended Abstracts*, 2019(1), 1–6. Retrieved 2021-08-26, from <https://www.tandfonline.com/doi/full/10.1080/22020586.2019.12073178> doi: 10.1080/22020586.2019.12073178
- Ward, S. H., & Hohmann, G. W. (1988, January). 4. Electromagnetic Theory for Geophysical Applications. In *Electromagnetic Methods in Applied Geophysics: Volume 1, Theory* (pp. 130–311). Society of Exploration Geophysicists. Retrieved 2020-07-07, from <https://library.seg.org/doi/abs/10.1190/1.9781560802631.ch4> doi: 10.1190/1.9781560802631.ch4
- Zamudio, K., Minsley, B., & Ball, L. (2021). *Airborne electromagnetic, magnetic, and radiometric survey, upper East River and surrounding watersheds near Crested Butte, Colorado, 2017*. U.S. Geological Survey. Retrieved from <https://doi.org/10.5066/P949ZCZ8>. (Type: dataset) doi: 10.5066/P949ZCZ8.
- Zhou, Y., Johansen, A. M., & Aston, J. A. D. (2016, July). Toward Automatic Model Comparison: An Adaptive Sequential Monte Carlo Approach. *Journal of Computational and Graphical Statistics*, 25(3), 701–726. Retrieved 2020-07-21, from <https://doi.org/10.1080/10618600.2015.1060885> (Publisher: Taylor & Francis _eprint: <https://doi.org/10.1080/10618600.2015.1060885>) doi: 10.1080/10618600.2015.1060885

691

Appendix A Static Sequential Monte Carlo Algorithm

Algorithm 2: Static Sequential Monte Carlo with an MCMC Kernel

Output: $\{\theta_T^{(i)}\}_{i=1}^N$ and $\hat{Z}_T$

Initialise particles from the prior for $i = 1, \dots, N$, $\theta_0^{(i)} \sim \pi_0(\cdot) = p(\cdot)$.

Initialise the particle weights for $i = 1, \dots, N$, $W_0^{(i)} \leftarrow N^{-1}$.

Set $t \leftarrow 0$, $\gamma_0 \leftarrow 0$, and $\hat{Z}_0 = 1$.

for $t \leftarrow 1, \dots, T$ **do**

Re-weight particles with normalised weights for $i = 1, \dots, N$.

$$w_t^{(i)} \leftarrow W_{t-1}^{(i)} \mathcal{L}(\mathbf{y} | \theta_t^{(i)})^{\gamma_t - \gamma_{t-1}}$$

$$W_t^{(i)} \leftarrow \frac{w_t^{(i)}}{\sum_{i=1}^{N_t} w_t^{(i)}}$$

$$\hat{Z}_t \leftarrow \hat{Z}_{t-1} \sum_{i=1}^{N_t} w_t^{(i)}$$

Re-sample N_t particles $\theta_t^{(i)*} \sim \hat{\pi}_t(\cdot)$ where $\hat{\pi}_t(\cdot)$ is represented by the weighted particles $\{W_t^{(i)}, \theta_t^{(i)}\}_{i=1}^{N_t}$

Mutate particles using an MCMC kernel with proposal $q(\cdot)$

$$\theta_{t+1}^{(i)} \sim q(\cdot | \theta_t^{(i)*}).$$

end

693

Appendix B Adapting the number of MCMC moves for multiple proposals

694

If we require at least one accepted proposal with a high probability greater than a tuning parameter $c' = 1 - c$, say $c' = 0.99$,

695

$$P(X \geq 1) \geq 1 - c \tag{B1}$$

696

Evaluate B1 by taking the complement

$$P(X \geq 1) = 1 - P(X = 0) \geq 1 - c$$

$$P(X = 0) \leq c$$

Taking the Binomial expansion at $X = 0$

$$(1 - p_{acc})^R \leq c$$

$$R \log 1 - p_{acc} \leq \log c$$

$$R \geq \frac{\log c}{\log 1 - p_{acc}}$$

697

Take R as the least integer upper bound

$$R = \left\lceil \frac{\log c}{\log 1 - p_{acc}} \right\rceil \tag{B2}$$

For more than one proposal type and many parameters, the note the number of accepted proposals $\mathbf{X} \sim \text{Multinomial}(R, \mathbf{a})$ where $\mathbf{X} = \{X_1, \dots, X_n\}$ and $\mathbf{a} = \{a_1, \dots, a_x\}$ for all parameters $i \in \{1, \dots, n\}$. Evaluating using the same tuning parameter $1 - c$ as above,

$$P(\mathbf{X} \geq 1) \geq 1 - c \quad (\text{B3})$$

Evaluate B3 by taking the complement

$$P(\mathbf{X} \geq 1) = 1 - P\left(\bigcup_{i=1}^n X_i = 0\right) \geq 1 - c$$

$$P\left(\bigcup_{i=1}^n X_i = 0\right) \leq c$$

Instead of evaluating the union of sets, it is sufficient to note that the union $\bigcup_{i=1}^n (X_i = 0)$ bounded below by any of the events $\{X_i = 0\}$ and that the marginal distribution of the multinomial is a binomial distribution. Hence enforcing

$$\max \{P(X_i = 0)\} = (1 - p_{acc}^{min})^R \leq P\left(\bigcup_{i=1}^n X_i = 0\right) \leq c \quad (\text{B4})$$

will satisfy B3 and as such B4 results in the same form as B2 using p_{acc}^{min} , that is

$$R = \left\lceil \frac{\log c}{\log 1 - p_{acc}^{min}} \right\rceil \quad (\text{B5})$$

698

Appendix C Total Effective Sample Size Estimator

First we define the Normalised Effective Sample Size (NESS) as a value between 0 and 1 giving the proportional quantity of representative samples of a distribution $\pi_t(\boldsymbol{\theta}|\mathbf{y}, k)$ that can be approximated with

$$\widehat{\text{NESS}}_{t,k} := \frac{\widehat{\text{ESS}}_{t,k}}{N_{t,k}}.$$

We then define the Total Effective Sample Size (TESS) as a value between 0 and N in terms of the expected NESS of all models conditional on the marginal probability of the models, namely

$$\begin{aligned} \mathbb{E}[\text{TESS}_t] &:= N \cdot \mathbb{E}[\text{NESS}_{t,k}] \\ &= N \cdot \mathbb{E}[\mathbb{E}[\text{NESS}_{t,k}|k]] \quad (\text{Law of Total Expectation}) \\ &= N \sum_{k \in \mathcal{M}} \mathbb{E}[\text{NESS}_{t,k}|k] \pi(k|\mathbf{y}) \end{aligned}$$

Using the estimator $\widehat{\text{NESS}}_{t,k}$ from above, and noting that $\hat{\pi}(k|\mathbf{y}) = \frac{N_{t,k}}{N}$, we have the estimate

$$\widehat{\text{TESS}}_t = N \sum_{k \in \mathcal{M}} \frac{\widehat{\text{ESS}}_{t,k}}{N_{t,k}} \frac{N_{t,k}}{N} = \sum_{k \in \mathcal{M}} \widehat{\text{ESS}}_{t,k}$$

699

Appendix D Ratio of Normalising Constants Between Successive Target Densities

This appendix demonstrates that the ratio of normalising constants between successive target densities in RJSMC is derived in a similar way to those for single model static SMC. Here we will use the notation $p_t(\boldsymbol{\theta}|k)$ to represent the conditional (normalised) target distribution. Noting $\eta_{t,k}(\boldsymbol{\theta}_k|k) = Z_{t,k} p_t(\boldsymbol{\theta}|k)$, the ratio of normalising constants can be found as follows

$$\mathcal{Z}_{t,k} = \int_{\boldsymbol{\theta}_k} \eta_{t,k}(\boldsymbol{\theta}_k|k) d\boldsymbol{\theta}$$

$$\begin{aligned}
 &= \int_{\boldsymbol{\theta}} \frac{\eta_{t,k}(\boldsymbol{\theta}_k|k)}{p_{t-1}(\boldsymbol{\theta}_k|k)} p_{t-1}(\boldsymbol{\theta}_k|k) d\boldsymbol{\theta} \\
 &= \int_{\boldsymbol{\theta}} \mathcal{Z}_{t-1,k} \frac{\eta_{t,k}(\boldsymbol{\theta}_k|k)}{\eta_{t-1,k}(\boldsymbol{\theta}_k|k)} p_{t-1}(\boldsymbol{\theta}_k|k) d\boldsymbol{\theta} \\
 \frac{\mathcal{Z}_{t,k}}{\mathcal{Z}_{t-1,k}} &= \int_{\boldsymbol{\theta}} \frac{\eta_{t,k}(\boldsymbol{\theta}_k|k)}{\eta_{t-1,k}(\boldsymbol{\theta}_k|k)} p_{t-1}(\boldsymbol{\theta}_k|k) d\boldsymbol{\theta}.
 \end{aligned}$$

Now using the weighted sample $\{W_{t-1}^{(i,k)}, \boldsymbol{\theta}_{t-1}^{(i,k)}\}_{i=1}^{N_{t,k}}$ from $p_{t-1}(\boldsymbol{\theta}_k|k)$ we obtain the following Monte Carlo estimate

$$\frac{\widehat{\mathcal{Z}_{t,k}}}{\mathcal{Z}_{t-1,k}} = \sum_{i=1}^{N_{t,k}} W_{t-1}^{(i,k)} \frac{\eta_{t,k}(\boldsymbol{\theta}^{(i)}|k)}{\eta_{t-1,k}(\boldsymbol{\theta}^{(i)}|k)}.$$

In the case where we are using a RJMCMC kernel, the term $\eta_{t,k}(\boldsymbol{\theta}^{(i)}|k)/\eta_{t-1,k}(\boldsymbol{\theta}^{(i)}|k)$ is the incremental weight. In this instance, the estimate of the ratio of normalising constants is

$$\frac{\widehat{\mathcal{Z}_{t,k}}}{\mathcal{Z}_{t-1,k}} = \sum_{i=1}^{N_{t,k}} w_{t,k}^{(i)}.$$

Appendix E Reversible jump MCMC proposals

Using $\pi(\boldsymbol{\theta}_k, k|\mathbf{y})$ as the target of the state space $\boldsymbol{\theta} = \bigcup_{k \in \mathcal{K}} (\{k\} \times \mathcal{R}^{N_{t,k}})$, RJMCMC proposal construction requires the following design choices:

1. **Dimension Match:** Given $n_k = |M_k|$, $n_{k'} = |M_{k'}|$, draw random variables $\mathbf{u}_k \sim g_k(\cdot)$, $\mathbf{u}_{k'} \sim g_{k'}(\cdot)$ of length w_k and $w_{k'}$ such that

$$n_k + w_k = n_{k'} + w_{k'}.$$

2. **Bijective Map:** $h_{k \rightarrow k'} : \mathcal{R}^{n_k} \times \mathcal{R}^{w_k} \rightarrow \mathcal{R}^{n_{k'}} \times \mathcal{R}^{w_{k'}}$ is chosen to map

$$\boldsymbol{\theta}'_{k'} = h_{k \rightarrow k'}(\boldsymbol{\theta}_k, \mathbf{u})$$

The acceptance ratio that satisfies detailed balance is

$$\alpha[(k, \boldsymbol{\theta}_k), (k', \boldsymbol{\theta}'_{k'})] = 1 \wedge \frac{\pi(k', \boldsymbol{\theta}'_{k'}|\mathbf{y}) q_{k' \rightarrow k}(\boldsymbol{\theta}_k, k|\boldsymbol{\theta}'_{k'}, k') g_{k' \rightarrow k}(\mathbf{u}_{k'})}{\pi(k, \boldsymbol{\theta}_k|\mathbf{y}) q_{k \rightarrow k'}(\boldsymbol{\theta}'_{k'}, k'|\boldsymbol{\theta}_k, k) g_{k \rightarrow k'}(\mathbf{u}_k)} \left| \frac{\partial h_{k \rightarrow k'}(\boldsymbol{\theta}_k, \mathbf{u})}{\partial(\boldsymbol{\theta}_k, \mathbf{u})} \right|.$$

E1 RJMCMC Proposals for the 1D layered Earth model

(Malinverno, 2002) introduced a proposal for the 1D layered model based on the change-point model from (P. J. Green, 1995). This model essentially specifies the model space as an arbitrary number of identically-distributed change-points representing layer interfaces, each with associated properties. The original definition used the so-called “grid trick” to enable birth-death proposals of layer interfaces. This was formalised instead by (Dosso et al., 2014) to use a Dirichlet-type prior on the layers themselves, visualised as the homogeneous space between the layer interfaces. The priors are

$$\begin{aligned}
 \text{Number of layers : } P(k) &= \frac{1}{k_{\max} - k_{\min}} \\
 i^{th}\text{-layer associated physical property } j : P(\beta_{i,j}|k) &= \frac{1}{\beta_{\max} - \beta_{\min}} \\
 \text{Layer thickness : } P(z_k|k) &= \frac{k!}{z_{\max}^k} \\
 i^{th}\text{-layer parameter } \boldsymbol{\theta}_{i,k} &= \{\beta_{i,0}, \dots, \beta_{i,J}, z_i\}
 \end{aligned}$$

The combined prior is given by

$$P(k)P(\boldsymbol{\theta}_k|k) = \frac{1}{k_{\max} - k_{\min}} \frac{P(\mathbf{z}_k|k)}{\prod_{i=1}^J (\beta_{\max} - \beta_{\min})^{k+1}}$$

The birth acceptance term is comprised of the prior and proposal ratios. The prior ratio is

$$\frac{P(k')P(\boldsymbol{\theta}_{k'}|k')}{P(k)P(\boldsymbol{\theta}_k|k)} = \frac{k+1}{z_{\max} \prod_{i=1}^J (\beta_{\max} - \beta_{\min})}$$

The proposal ratio for inserting a new layer in the i^{th} position (or a uniform probability of inserting a layer in the interval $(0, z_{\max})$) combined with the reverse step (death of the proposed layer, with probability $1/(k+1)$) is

$$\frac{Q(k, \boldsymbol{\theta}_k|k', \boldsymbol{\theta}_{k'})}{Q(k', \boldsymbol{\theta}_{k'}|k, \boldsymbol{\theta}_k)} = \frac{1}{k+1} \frac{z_{\max}}{Q_{\beta}(\boldsymbol{\beta}'|\boldsymbol{\beta})}$$

The acceptance term for layer birth, combining all terms, is

$$A_{\text{birth}}[(k, \boldsymbol{\theta}_k), (k', \boldsymbol{\theta}_{k'})] = 1 \wedge \frac{1}{Q_{\beta}(\boldsymbol{\beta}'|\boldsymbol{\beta})} \frac{1}{\prod_{i=1}^J (\beta_{\max} - \beta_{\min})} \frac{\mathcal{L}(\mathbf{y}|k', \boldsymbol{\theta}_{k'})}{\mathcal{L}(\mathbf{y}|k, \boldsymbol{\theta}_k)}.$$

If $Q_{\beta}(\boldsymbol{\beta}'|\boldsymbol{\beta})$ is chosen to be simply drawing from the uniform prior for β_i for the i^{th} layer being inserted, the acceptance ratio reduces to the likelihood ratio

$$A_{\text{Naive}}[(k, \boldsymbol{\theta}_k), (k', \boldsymbol{\theta}_{k'})] = 1 \wedge \frac{\mathcal{L}(\mathbf{y}|k', \boldsymbol{\theta}_{k'})}{\mathcal{L}(\mathbf{y}|k, \boldsymbol{\theta}_k)}.$$

E2 An adaptive reversible jump MCMC proposal for the 1D layered Earth model

A common ailment of RJMCMC proposals is the poor mixing phenomenon identified by very low acceptance rates. The design of adaptive proposals is a vast field of research and is naturally problem specific. In this appendix we will identify one possible construction for an adaptive RJMCMC proposal which takes advantage of the availability of π_t in an SMC algorithm without guarantees for performance in any particular context.

This proposal design follows that of the birth/death design in Appendix E1 where the vector of auxilliary variables in the birth move is $\mathbf{u} = [u_d, u_p]$, and the depth of the new layer u_d is first drawn independently and uniformly over the range of the allowed depths. Note that since layer interfaces are sorted in order of depth, we can find the index i of the layer interface being inserted. Following this, we consider that the second auxilliary variable is drawn as from a standard Gaussian and then transformed via a bijective map. Thus we consider below the construction of $h_{k \rightarrow k'}(\boldsymbol{\theta}_k, \mathbf{u})$.

If it can be assumed that $\pi_t(\boldsymbol{\theta}_t|k)$ is unimodal and approximately Gaussian (an assumption that rarely holds), its variance can be approximated via the sample variance

$$\boldsymbol{\Sigma}_{k,t} = \frac{1}{N_{t,k} - 1} \sum_{i=1}^{N_{t,k}} (\boldsymbol{\theta}_{k,t}^{(i)} - \mu_{k,t})(\boldsymbol{\theta}_{k,t}^{(i)} - \mu_{k,t})^T$$

We choose the bijective transform for the birth of one layer to be

$$\boldsymbol{\theta}_{k+1,t}[j] = \boldsymbol{\theta}_{k,t}[h(j)] + \boldsymbol{\Sigma}_{k+1,t}^{\frac{1}{2}}[i, j]u$$

where j indexes each parameter component independently, $h(j)$ maps the indices of $\boldsymbol{\theta}_k$ to $\boldsymbol{\theta}_{k+1}$, and $\boldsymbol{\Sigma}_{k+1,t}^{\frac{1}{2}}$ is the ZCA colouring matrix (square root) of $\boldsymbol{\Sigma}_{k+1,t}$. Also note that the row index i is essentially arbitrary, however a common choice is the index of the layer being inserted into the model.

To form the Jacobian, we take the partial derivatives

$$\begin{aligned}\frac{\partial \boldsymbol{\theta}_{k+1}[j]}{\partial \boldsymbol{\theta}_k[h(j)]} &= 1 \\ \frac{\partial \boldsymbol{\theta}_{k+1}[j]}{\partial u} &= \boldsymbol{\Sigma}_{k+1,t}^{\frac{1}{2}}[i, j]\end{aligned}$$

The Jacobian becomes a mostly triangular matrix with one row replicating the non-zero element locations. It can be re-arranged to a block triangular matrix

$$\mathbf{J} = \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ 0 & \mathbf{D} \end{bmatrix}$$

where $\mathbf{A}$ is diagonal and comprised of $\frac{\partial \boldsymbol{\theta}_{k+1}[j]}{\partial \boldsymbol{\theta}_k[h(j)]} = 1$ entries, and

$$\mathbf{D} = \begin{bmatrix} 1 & \frac{\partial \boldsymbol{\theta}_{k+1}[h^{-1}(j)]}{\partial u_p} \\ 1 & \frac{\partial \boldsymbol{\theta}_{k+1}[j]}{\partial u_p} \end{bmatrix}.$$

Hence the determinant becomes

$$\begin{aligned}\det \mathbf{J} &= \det \mathbf{A} \det \mathbf{D} \\ &= \frac{\partial \boldsymbol{\theta}_{k+1}[j]}{\partial u_p} - \frac{\partial \boldsymbol{\theta}_{k+1}[h^{-1}(j)]}{\partial u_p} \text{ as } \det \mathbf{A} = 1 \\ &= \boldsymbol{\Sigma}_{k+1,t}^{\frac{1}{2}}[i, j] - \boldsymbol{\Sigma}_{k+1,t}^{\frac{1}{2}}[i, h(j)].\end{aligned}$$

For the death move, we need to solve for u_p . Using the knowledge that elements $\boldsymbol{\theta}_k[j]$ and $\boldsymbol{\theta}_k[h(j)]$ are equal for $j \neq h(j)$, we derive

$$u_p = \frac{\boldsymbol{\theta}_{k+1}[j] - \boldsymbol{\theta}_{k+1}[h(j)]}{\boldsymbol{\Sigma}_{k+1,t}^{\frac{1}{2}}[i, j] - \boldsymbol{\Sigma}_{k+1,t}^{\frac{1}{2}}[i, h(j)]}.$$