

# SB-ETAS: using simulation based inference for scalable, likelihood-free inference for the ETAS model of earthquake occurrences

Samuel Stockman<sup>1\*</sup>, Daniel J. Lawson<sup>1</sup> and Maximilian J. Werner<sup>2</sup>

<sup>1</sup>School of Mathematics, University of Bristol, Woodland Rd, Bristol, BS8 1UG, United Kingdom.

<sup>2</sup>School of Earth Sciences, University of Bristol, Wills Memorial Building, Bristol, BS8 1RL, United Kingdom.

\*Corresponding author(s). E-mail(s): [sam.stockman@bristol.ac.uk](mailto:sam.stockman@bristol.ac.uk);  
Contributing authors: [dan.lawson@bristol.ac.uk](mailto:dan.lawson@bristol.ac.uk); [max.werner@bristol.ac.uk](mailto:max.werner@bristol.ac.uk);

## Abstract

Performing Bayesian inference for the Epidemic-Type Aftershock Sequence (ETAS) model of earthquakes typically requires MCMC sampling using the likelihood function or estimating the latent branching structure. These tasks have computational complexity  $\mathcal{O}(n^2)$  with the number of earthquakes and therefore do not scale well with new enhanced catalogs, which can now contain an order of  $10^6$  events. On the other hand, simulation from the ETAS model can be done more quickly  $\mathcal{O}(n \log n)$ . We present SB-ETAS: simulation-based inference for the ETAS model. This is an approximate Bayesian method which uses Sequential Neural Posterior Estimation (SNPE), a machine learning based algorithm for learning posterior distributions from simulations. SB-ETAS can successfully approximate ETAS posterior distributions on shorter catalogues where it is computationally feasible to compare with MCMC sampling. Furthermore, the scaling of SB-ETAS makes it feasible to fit to very large earthquake catalogs, such as one for Southern California dating back to 1932. SB-ETAS can find Bayesian estimates of ETAS parameters for this catalog in less than 10 hours on a standard laptop, which would have taken over 2 weeks using MCMC. Looking beyond the standard ETAS model, this simulation based framework would allow earthquake modellers to define and infer parameters for much more complex models that have intractable likelihood functions.

**Keywords:** Point Processes, Earthquakes, Approximate Bayesian Inference, Simulation Based Inference, Machine Learning, Hawkes Process

# 1 Introduction

The Epidemic Type Aftershock Sequence (ETAS) model (Ogata, 1988) has been the most dominant way of modelling seismicity in both retrospective and fully prospective forecasting experiments (e.g. Woessner et al., 2011; Taroni et al., 2018; Catania et al., 2018; Mancini et al., 2019, 2020) as well as in operational earthquake forecasting (Marzocchi et al., 2014; Rhoades et al., 2016; Field et al., 2017). The model characterises the successive triggering of earthquakes, making it effective for forecasting aftershock sequences. The parameters of the model are also used by seismologists to characterise physical phenomena of different tectonic regions such as structural heterogeneity, stress, and temperature (Utsu et al., 1995).

In recent years, an accelerated growth in the number of seismic sensors and machine learning algorithms for detecting the arrival times of earthquake phases (e.g. Zhu and Beroza, 2019), has meant that earthquake catalogs have grown to a size that has overtaken what is computationally feasible to fit an ETAS model to. In California, a deployment of a dense network of seismic sensors over the last century combined with an active tectonic regime has resulted in a comprehensive dataset of earthquakes in the region (Hutton et al., 2010). Furthermore, in more specific areas of California, through machine learning based seismic phase picking and template matching, enhanced

earthquake catalogs have been created which contain many small previously undetected earthquakes (White et al., 2019; Ross et al., 2019). It is fair to assume that these datasets will only continue to grow in the future as past continuous data is reprocessed by these new detection techniques as well as being applied to future earthquakes. Determining whether increased data size results in improved earthquake forecasts is an important question for the seismological community. This work provides some methodology for answering that question by providing a simulation based estimation procedure for the ETAS model, with a scalability that make it feasible to infer ETAS parameters for such large datasets.

Most commonly, point estimates of ETAS parameters are found through maximum likelihood estimation (MLE). From these point estimates, forecasts can be issued by simulating multiple catalogs over the forecasting horizon. Forecast uncertainty is quantified by the distribution of simulations from MLE parameter values, however, this approach fails to quantify uncertainty contained in estimating the parameters themselves. Parameter uncertainty for MLE can be estimated using the Hessian of the likelihood (Ogata, 1978; Rathbun, 1996; Wang et al., 2010), which requires a very large sample size to be effective, or through multiple runs of the MLE procedure with different initial conditions (Lombardi, 2015).

Full characterisation of the parameter uncertainty is achieved with Bayesian inference, a procedure which returns the entire probability distribution over parameters conditioned on the observed data and updated from prior knowledge. This distribution over parameters, known as the posterior, does not have a closed form expression for the temporal and spatio-temporal ETAS model and so several approaches have used Markov Chain Monte Carlo (MCMC) to obtain samples from this distribution (Vargas and Gneiting, 2012; Ross, 2021). These approaches evaluate the likelihood of the ETAS model during the procedure, which is an operation with quadratic complexity  $\mathcal{O}(n^2)$ , and therefore are only suitable for catalogs of around 10,000 events.

An alternative to MCMC is based on the Integrated Nested Laplace Approximation (INLA) method (Rue et al., 2017), which approximates the marginal posterior distributions using latent Gaussian models. The implementation of INLA for the ETAS model, named `inlabru`, by Serafini et al. (2023) demonstrated a factor 10 speed-up over the MCMC method for a catalog of 3,500 events but they do not provide results on larger catalogs. Section 2.3 provides a more in depth overview of the MCMC and INLA methods.

We propose SB-ETAS: an alternative method for Bayesian inference for the ETAS model. SB-ETAS is based on Simulation Based Inference

(SBI), a family of approximate procedures which infer posterior distributions for parameters using simulations in place of the likelihood. Simulation from the ETAS model is  $\mathcal{O}(n \log n)$ , making it a procedure which scales better than the evaluation of the likelihood,  $\mathcal{O}(n^2)$ . From a suite of SBI methods developed recently using neural density estimators (Cranmer et al., 2020), we choose Sequential Neural Posterior Estimation (SNPE) since it does not involve any MCMC. SNPE trains a neural density estimator to approximate the posterior distribution from pairs of simulations and parameters. Section 3 provides an overview of SNPE and other methods of SBI.

SBI is typically used in the case where the likelihood function is intractable, for example when the likelihood involves integrating over many unobserved latent variables. However, in our use case the likelihood function is tractable, presenting an alternative potential for SBI for problems where the simulator scales better than the likelihood. Looking beyond the example presented in this study, the likelihood-free inference procedure that we present, could be applied to other branching models for earthquakes. Such branching models need only be defined as a forward simulator and wouldn't be constrained to have a likelihood function. For example, a large problem in earthquake modeling is the time varying detection of earthquakes, particularly following very

large magnitude events (Hainzl, 2016b; Zhuang et al., 2017). These undetected events bias models such as ETAS that do not account for them (Seif et al., 2017; Stockman et al., 2023). A branching model which deletes simulated events with a time varying detection rate, does not result in a tractable likelihood but could still be fit to data using our procedure.

Other branching process models for earthquakes that incorporate physical quantities (e.g. Field et al., 2017), on the other hand, do not have tractable likelihoods and currently use parameter values that are fit independent of the physical part of the model. This work provides a framework for estimating the parameters for such models, with the hope that better estimates would lead to improved forecasts.

The remainder of this paper is structured as follows: In section 2 we give an overview of the ETAS model along with existing procedures for performing Bayesian inference; section 3 gives an overview of SBI, following which we describe the details of SB-ETAS in section 4. We present empirical results based on synthetic earthquake catalogs in section 5 and observational earthquake data from Southern California in section 6, before finishing with a discussion in section 7.

## 2 The ETAS model

The temporal Epidemic Type Aftershock Sequence (ETAS) model (Ogata, 1988) is a marked Hawkes process (Hawkes, 1971) that describes the random times of earthquakes  $t_i$  along with their magnitudes  $m_i$  in the sequence  $\mathbf{x} = \{(t_1, m_1), (t_2, m_2), \dots, (t_n, m_n)\} \in [0, T]^n \times \mathcal{M}^n \subset \mathbb{R}_+^n \times \mathcal{M}^n$ . The natural ordering of events leads to the definition of a filtration  $\mathcal{H}_t$  known as the history of the process, which contains information on all events up to time  $t$ . Marked Hawkes processes are usually specified by their conditional intensity function (Rasmussen, 2018),

$$\lambda(t, m | \mathcal{H}_t) = \quad (1)$$

$$\lim_{\Delta t, \Delta m \rightarrow 0} \frac{\mathbb{E}[N([t, t + \Delta t) \times B(m, \Delta m) | \mathcal{H}_t)]}{\Delta t |B(m, \Delta m)|}. \quad (2)$$

Where the counting measure  $N(A)$  counts the events in the set  $A \subset \mathbb{R}_+ \times \mathcal{M}$  and  $|B(m, \Delta m)|$  is the Lebesgue measure of the ball  $B(m, \Delta m)$  with radius  $\Delta m$ .

The ETAS model typically has the form,

$$\lambda(t, m | \mathcal{H}_t) = \left( \mu + \sum_{i: t_i < t} g(t - t_i, m_i) \right) f_{GR}(m), \quad (3)$$

where  $\mu$  is a constant background rate of events and  $g(t, m)$  is a non-negative excitation kernel which describes how past events contribute to the likelihood of future events. For the ETAS model

this triggering kernel factorises the contribution from the magnitude and the time,

$$g(t, m) = k(m; K, \alpha)h(t; c, p) \quad (4)$$

$$k(m; K, \alpha) = Ke^{\alpha(m-M_0)} : m \geq M_0 \quad (5)$$

$$h(t; c, p) = c^{p-1}(p-1)(t+c)^{-p} : t \geq 0 \quad (6)$$

where the  $k(m; K, \alpha)$  is known as the Utsu law of productivity (Utsu, 1970) and  $h(t; c, p)$  is a power law known as the Omori decay (Utsu et al., 1995). The magnitudes are said to be “unpredictable” since they do not depend on previous events and are distributed according to the Gutenberg-Richter law for magnitudes (Gutenberg and Richter, 1936) with probability density  $f_{GR}(m) = \beta e^{\beta(m-M_0)}$  on the support  $m : m \geq M_0$ .

## 2.1 Branching Process Formulation

An alternative way of formulating the ETAS model is as a Poisson cluster process (Rasmussen, 2013). In this formulation a set of immigrants  $I$  are realisations of a Poisson process with rate  $\mu$ . Each immigrant  $t_i \in I$  has a magnitude  $m_i$  with probability density  $f_{GR}$  and generates offspring  $S_i$  from an independent non-homogeneous Poisson process, with rate  $g(t-t_i, m_i)$ . Each offspring  $t_j \in S_i$  also has magnitude  $m_j$  with probability density  $f_{GR}$  and generate offspring  $S_j$  of their own. This

process is repeated over generations until a generation with no offspring in time interval  $[0, T]$  is produced. Although it is not observed in the data, this process is accompanied by latent branching variables  $B = \{B_1, \dots, B_n\}$  which define the branching structure of the process,

$$B_i = \begin{cases} 0 & \text{if } t_i \in I \text{ (i.e. } i \text{ is a background event)} \\ j & \text{if } t_i \in S_j \text{ (i.e. } i \text{ is an offspring of } j) \end{cases} \quad (7)$$

Both the intensity function formulation as well as the branching formulation define different methods for simulating as well as inferring parameters for the Hawkes process. We now give a brief overview of some of these methods, however, we direct the reader to Reinhart (2018) for a more detailed review.

## 2.2 Simulation

A simulation algorithm based on the conditional intensity function was proposed by Ogata (1998). This algorithm requires generating events sequentially using a thinning procedure. Simulating forward from an event  $t_i$ , the time to the next event  $\tau$  is proposed from a Poisson process with rate  $\lambda(t_i | \mathcal{H}_{t_i})$ . The proposed event  $t_{i+1} = t_i + \tau$  is then rejected with probability  $1 - \frac{\lambda(t_i | \mathcal{H}_{t_i})}{\lambda(t_i + \tau | \mathcal{H}_{t_i})}$ . This procedure is then repeated from the newest simulated event until a proposed event falls outside a predetermined time window  $[0, T]$ .

This procedure requires evaluating the intensity function  $\lambda(t|\mathcal{H}(t))$  at least once for each one of  $n$  events that are simulated. Evaluating the intensity function requires a summation over all events before time  $t$ , thus giving this simulation procedure time complexity  $\mathcal{O}(n^2)$ .

Another algorithm which instead simulates using the branching process formulation of the ETAS model was proposed by Zhuang et al. (2004). This algorithm simulates events over generations  $G^{(i)}, i = 0, \dots$ , until no more events fall within the interval  $[0, T]$ .

---

**Algorithm 1** Hawkes Branching Process Simulation in the interval  $[0, T]$

---

1. Generate events  $(t_i, m_i) \in G^{(0)}$  from a stationary Poisson process with intensity  $\mu$  in  $[0, T]$ .
  2.  $l = 0$ .
  3. For each  $(t_i, m_i) \in G^{(l)}$  simulate its  $N^{(i)}$  offspring, where  $N^{(i)} \sim \text{Poisson}(k(m_i))$ . Each offspring's time is generated from  $h(t)$  and magnitude from  $f_{GR}$  and are labelled  $O_i^{(l)}$ .
  4.  $G^{(l+1)} = \bigcup_{i \in G^{(l)}} O_i^{(l)}$
  5. If  $G^{(l)}$  is not empty,  $l = l + 1$  and return to step 3.
  6. Sort and return  $S = \bigcup_{j=0}^l G^{(j)}$  as the set of all simulated events.
- 

This procedure has time complexity  $\mathcal{O}(n)$  for steps 1-5, since there is only a single pass over all events. Additional time constraint is added in step 6. where the whole set of events are sorted chronologically, which is at best  $\mathcal{O}(n \log n)$ .

## 2.3 Bayesian Inference

Given we observe the sequence  $\mathbf{x}_{\text{obs}} = \{(t_1, m_1), (t_2, m_2), \dots, (t_n, m_n)\}$  in the interval  $[0, T]$ , we are interested in the posterior probability  $p(\theta|\mathbf{x}_{\text{obs}})$  for the parameters  $\theta = (\mu, K, \alpha, c, p)$  of the ETAS model defined in (3)-(6), updated from some prior probability  $p(\theta)$ . The posterior distribution, expressed in Bayes rule,

$$p(\theta|\mathbf{x}_{\text{obs}}) \propto p(\mathbf{x}_{\text{obs}}|\theta)p(\theta), \quad (8)$$

is known up to a constant of proportionality through the product of the prior  $p(\theta)$  and the likelihood  $p(\mathbf{x}_{\text{obs}}|\theta)$ , where,

$$\log p(\mathbf{x}_{\text{obs}}|\theta) = \sum_{i=1}^n \log \left[ \mu + \sum_{j=1}^{i-1} h(t; c, p)k(m; K, \alpha) \right] \quad (9)$$

$$- \mu T - \sum_{i=1}^n k(m; K, \alpha)H(T - t_i; c, p). \quad (10)$$

Here,  $H(t; c, p) = \int_0^t h(s; c, p)ds$ , denotes the integral of the Omori decay kernel.

Vargas and Gneiting (2012) draw samples from the posterior  $p(\theta|\mathbf{x}_{\text{obs}})$  through independent random walk Markov Chain Monte Carlo (MCMC) with Metropolis-Hastings rejection of proposed samples. This approach, however, can suffer from slow convergence due parameters of the ETAS model having high correlation (Ross, 2021).

Ross (2021) developed an MCMC sampling scheme, **bayesianETAS**, which conditions on the latent branching variables  $B = \{B_1, \dots, B_n\}$ . The scheme iteratively estimates the branching structure,

$$\mathbb{P}(B_i^{(k)} = j | \mathbf{x}_{\text{obs}}, \theta^{(k-1)}) = \quad (11)$$

$$\begin{cases} \frac{\mu^{(k-1)}}{\mu^{(k-1)} + \sum_{j=1}^{i-1} k(m_j)h(t_i - t_j)} & : j = 0 \\ \frac{k(m_j)h(t_i - t_j)}{\mu^{(k-1)} + \sum_{j=1}^{i-1} k(m_j)h(t_i - t_j)} & : j = 1, 2, \dots, i-1 \end{cases} \quad (12)$$

and then samples parameters,  $\theta^{(k)} = (\mu^{(k)}, K^{(k)}, \alpha^{(k)}, c^{(k)}, p^{(k)})$ , from the conditional likelihood,

$$\log p(\mathbf{x}_{\text{obs}} | \theta, B^{(k)}) = \quad (13)$$

$$|S_0^{(k)}| \log \mu - \mu T + \sum_{j=1}^n \left( -k(m_j; K, \alpha) H(T - t_j; c, p) \right. \quad (14)$$

$$\left. + |S_j^{(k)}| \log k(m_j; K, \alpha) + \sum_{t_i \in S_j^{(k)}} \log h(t_i - t_j; c, p) \right) \quad (15)$$

Where  $|S_j|$  denotes the number of events that were triggered by the event at  $t_j$ .

By conditioning on the branching structure, the dependence between parameters  $(K, \alpha)$  and  $(c, p)$  is reduced, decreasing the time it takes for the sampling scheme to converge. We can see from equation (11) that estimating the branching

structure from the data is a procedure that is  $\mathcal{O}(n^2)$ . Since for every event  $i = 1, \dots, n$ , to estimate its parent we must sum over  $j = 1, \dots, i-1$ . For truncated version of the time kernel  $h(t)$ , this operation can be streamlined to  $\mathcal{O}(n)$ , however, this is not the case in general.

More recently Serafini et al. (2023) have constructed an approximate method of Bayesian inference for the ETAS model based on an Integrated Nested Laplace Approximation (INLA) implemented in the R-package **inlabru** as well as a linear approximation of the likelihood. This approach expresses the log-likelihood as 3 terms,

$$\log p(\mathbf{x}_o | \theta) = \quad (16)$$

$$- \Lambda_0(\mathbf{x}_{\text{obs}}, \theta) - \sum_{i=1}^n \Lambda_i(\mathbf{x}_{\text{obs}}, \theta) + \sum_{i=1}^n \log \lambda(t_i | \mathcal{H}_{t_i}), \quad (17)$$

where,

$$\begin{aligned} \Lambda_0(\mathbf{x}_{\text{obs}}, \theta) &= \int_0^T \mu dt, \\ \Lambda_i(\mathbf{x}_{\text{obs}}, \theta) &= \sum_{h=1}^{B_i} \int_{b_{h,i}}^{B_i} g(t - t_i, m_i) dt \\ &= \sum_{h=1}^{B_i} \Lambda_i(\mathbf{x}_{\text{obs}}, \theta, b_{h,i}). \end{aligned}$$

Where  $b_{1,i}, \dots, b_{B_i,i}$  are chosen to partition the interval  $[0, T]$ . The log-likelihood is then linearly approximated with a first order Taylor expansion

with respect to the posterior mode  $\theta^*$ ,

$$\begin{aligned}\widehat{\log p(\mathbf{x}_{\text{obs}}|\theta; \theta^*)} &= -\widehat{\Lambda}_0(\mathbf{x}_{\text{obs}}, \theta, \theta^*) \\ &\quad - \sum_{i=1}^n \sum_{h=1}^{B_i} \widehat{\Lambda}_i(\mathbf{x}_{\text{obs}}, \theta, b_{h,i}; \theta^*) + \sum_{i=1}^n \widehat{\log \lambda}(\mathbf{x}_{\text{obs}}, \theta; \theta^*) \\ &= -\exp\{\overline{\log \Lambda}_0(\mathbf{x}_{\text{obs}}, \theta, \theta^*)\} \\ &\quad - \sum_{i=1}^n \sum_{h=1}^{B_i} \exp\{\overline{\log \Lambda}_i(\mathbf{x}_{\text{obs}}, \theta, b_{h,i}; \theta^*)\} \\ &\quad + \sum_{i=1}^n \overline{\log \lambda}(\mathbf{x}_{\text{obs}}, \theta; \theta^*),\end{aligned}$$

where the notation,  $\widehat{\Lambda}$ , denotes the approximation of  $\Lambda$  and  $\overline{\log \Lambda}(\cdot; \theta^*)$  denotes the first order Taylor expansion of  $\log \Lambda$  about the point  $\theta^*$ .

The posterior mode  $\theta^*$  is found through a Quasi-Newton optimisation method and the final posterior densities are found using INLA, which approximates the marginal posteriors  $p(\theta_i|\mathbf{x}_{\text{obs}})$  using a latent Gaussian model.

This approach speeds up computation of the posterior densities, since it only requires evaluation of the likelihood function during the search for the posterior mode. However, the approximation of the likelihood requires partitioning the space into a number of bins, which the authors recommend choosing as greater than 3 per observation. This results in the approximate likelihood having complexity  $\mathcal{O}(n^2)$ .

### 3 Simulation Based Inference

A family of Bayesian Inference methods have evolved from application settings in science, economics or engineering where stochastic models are used to describe complex phenomena. In this setting, the model may simulate data from a given set of input parameters, however, the likelihood of observing data given parameters is intractable. The task in this setting is to approximate the posterior  $p(\theta|\mathbf{x}_{\text{obs}}) \propto p(\mathbf{x}_{\text{obs}}|\theta)p(\theta)$ , with the restriction that we cannot evaluate  $p(\mathbf{x}|\theta)$  but we have access to the likelihood implicitly through samples  $\mathbf{x}_r \sim p(\mathbf{x}|\theta_r)$  from a simulator, for  $r = 1, \dots, R$  and where  $\theta_r \sim p(\theta)$ . This approach is commonly referred to as Simulation Based Inference (SBI) or likelihood-free inference.

Until recently, the most predominant approach for SBI was Approximate Bayesian Computation (Beaumont et al., 2002). In its simplest form, parameters are chosen from the prior  $\theta_r \sim p(\theta)$ ,  $r = 1, \dots, R$ , the simulator then generates samples  $\mathbf{x}_r \sim p(\mathbf{x}|\theta_r)$ ,  $r = 1, \dots, R$ , and each sample is kept if it is within some tolerance  $\epsilon$  of the observed data, i.e.  $d(\mathbf{x}_r, \mathbf{x}_{\text{obs}}) < \epsilon$ .

This approach, although exact when  $\epsilon \rightarrow 0$ , is inefficient with the use of simulations. Sufficiently small  $\epsilon$  requires simulating an impractical number of times, and this issue scales poorly with the dimension of  $\mathbf{x}$ . In light of this an MCMC



approach to ABC makes proposals for new simulator parameters  $\theta_r \sim q(\cdot | \theta_{r-1})$  using a Metropolis-Hastings kernel (Beaumont et al., 2002; Marjoram et al., 2003). This leads to a far higher acceptance of proposed simulator parameters but still scales poorly with the dimension of  $\mathbf{x}$ .

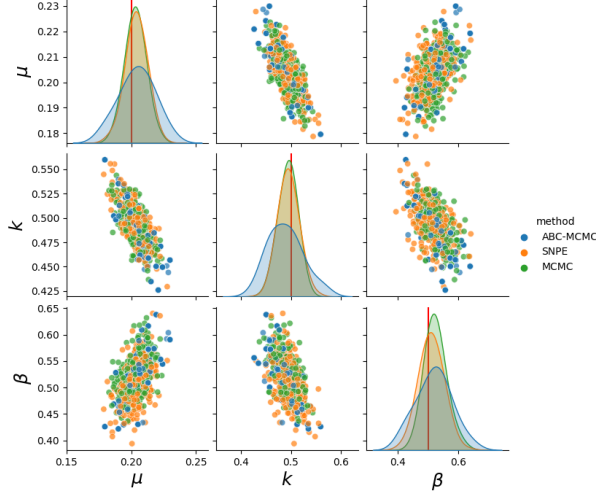
In order to cope with high dimensional simulator outputs  $\mathbf{x} \in \mathbb{R}^n$ , summary statistics  $S(\mathbf{x}) \in \mathbb{R}^d$  are chosen to reduce the dimension of the sample whilst still retaining as much information as possible. These are often chosen from domain knowledge or can be learnt as part of the inference procedure (Prangle et al., 2014). Summary statistics  $S(\mathbf{x})$  are then used in place of  $\mathbf{x}$  in any of the described methods for SBI.

### 3.1 Neural Density Estimation

Recently, SBI has seen more rapid development as a result of neural network based density estimators, which seek to approximate the density  $p(x)$  given samples of points  $x \sim p(x)$ . A popular method for neural density estimation is normalising flows (Rezende and Mohamed, 2015), in which a neural network parameterizes an invertible transformation  $x = g_\phi(u)$ , of a variable  $u$  from a simple base distribution  $p(u)$  into the target distribution of interest. In practice, the transformation is typically composed of a stack of invertible transformations, which allows it to learn the complex target density. The parameters

of the transformation are trained through maximising the likelihood of observing  $p_g(x)$ , which is given by the change of variables formula. Since  $x$  is expressed as a transformation of a simple distribution  $u \sim p(u)$ , samples from the learnt distribution  $p_g(x)$  can be generated by sampling from  $p(u)$  and passing the samples through the transformation. Neural density estimators may also be generalised to learn conditional densities  $p(\mathbf{x}|\mathbf{y})$  by conditioning the transformation  $g_\phi$  on the variable  $y$  (Papamakarios et al., 2017).

In the task of SBI, a neural density estimator can be trained on pairs of samples  $\theta_r \sim p(\theta)$ ,  $\mathbf{x}_r \sim p(\mathbf{x}|\theta_r)$  to approximate either the likelihood  $p(\mathbf{x}_{\text{obs}}|\theta)$  or the posterior density  $p(\theta|\mathbf{x}_{\text{obs}})$ , from which posterior samples can be obtained. If the posterior density is estimated, in a procedure known as Neural Posterior Estimation (NPE) (Lueckmann et al., 2017), then samples can be drawn from the normalising flow. If the likelihood is estimated, known as Neural Likelihood estimation (NLE) (Papamakarios et al., 2019), then the approximate likelihood can be used in place of the true likelihood in a MCMC sampling algorithm to obtain posterior samples. These density estimation techniques generally outperform ABC-based methods since they are able to efficiently interpolate between different simulations (Figure 1, Lueckmann et al. (2021)). Other neural network methods exist for SBI such as ratio estimation



**Fig. 1** Posterior densities for a univariate Hawkes process with exponential kernel. The ‘observed’ data contains 4806 events and was simulated from parameters indicated in red on the diagonal plots. In blue are posterior samples found using the ABC-MCMC method for SBI using 300,000 simulations. In orange are posterior samples from SNPE using the same summary statistics as ABC-MCMC but only 10,000 simulations. In green are posterior samples found using MCMC sampling with likelihood function. A Uniform( $[0.05, 0, 0]$ ,  $[0.85, 0.9, 3]$ ) prior was used for all three methods.

(Izbicki et al., 2014) or score matching (Geffner et al., 2022; Sharrock et al., 2022), however, we direct the reader to (Cranmer et al., 2020) for a more comprehensive review of modern SBI.

## 4 SB-ETAS

We now present SB-ETAS, our scalable method for Bayesian inference for the ETAS model. The method avoids having to evaluate the computationally expensive likelihood function and instead leverages fast simulation from the ETAS branching process. The inference method uses Sequential Neural Posterior Estimation (SNPE) (Papamakarios and Murray, 2016; Lueckmann et al., 2017),

---

### Algorithm 2 SB-ETAS

---

**Input:** observed data  $\mathbf{x}_{\text{obs}}$ , summary statistic  $S(\mathbf{x})$ , estimator  $q_\phi(\theta|\mathbf{x})$ , prior  $p(\theta)$ , number of rounds  $K$ , simulations per round  $L$ .

```

 $q_\phi^0(\theta|\mathbf{x}_{\text{obs}}) = p(\theta)$  and  $\mathcal{D} = \{\}$ 
for  $k = 1 : K$  do
    for  $l = 1 : L$  do
        sample  $\theta_l \sim q_\phi^{k-1}(\theta|\mathbf{x}_{\text{obs}})$ 
        simulate  $\mathbf{x}_l \sim p(\mathbf{x}|\theta_l)$  using Algorithm 1
        add  $(\theta_l, S(\mathbf{x}_l))$  to  $\mathcal{D}$ 
    end for
    train  $q_\phi^k(\theta|\mathbf{x}_{\text{obs}})$  on  $\mathcal{D}$ 
end for

```

**Output:** approximate posterior

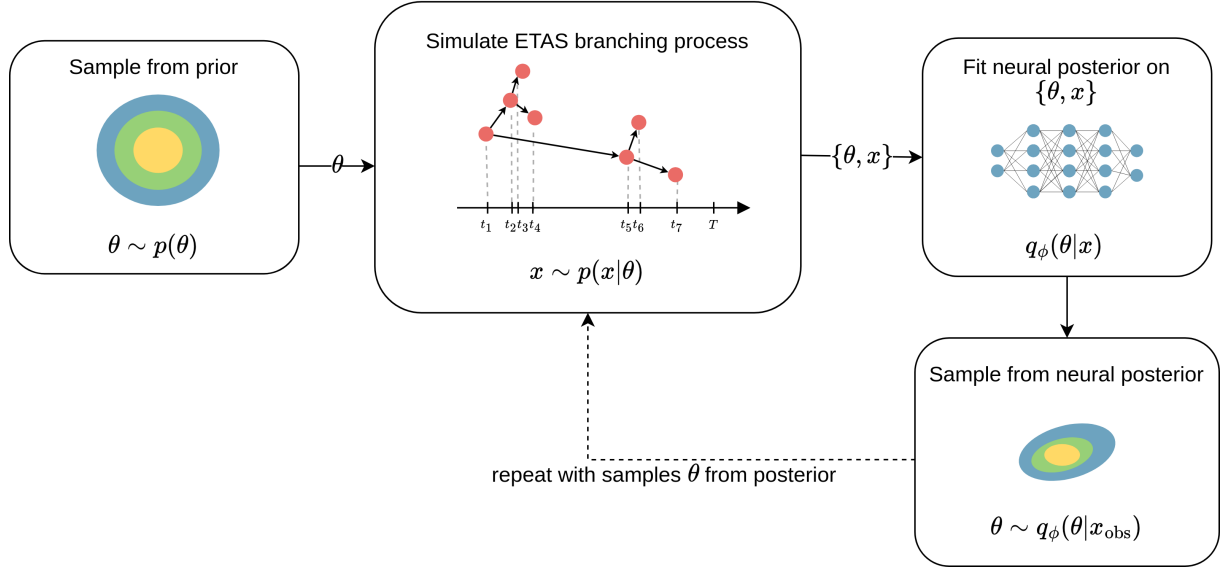
$$\hat{p}(\theta|\mathbf{x}_{\text{obs}}) = \hat{q}_\phi^K(\theta|\mathbf{x}_{\text{obs}})$$


---

a modified version of NPE which performs inference over rounds. Each round, an estimate of the posterior proposes new samples for the simulator, a neural density estimator is trained on those samples and the estimated posterior is updated (Figure 2, Algorithm 2). SNPE was chosen over other methods of Neural-SBI, as it avoids the need to perform MCMC sampling, a slow procedure. Instead sampling from the posterior is fast since the approximate posterior is a normalising flow.

### 4.1 Summary Statistics

What differentiates Hawkes process models from other simulator models used in Neural-SBI is that the output of the simulator  $\mathbf{x} = (t_1, m_1), \dots, (t_n, m_n)$  itself has random dimension. For a specified time interval over which to simulate earthquakes  $[0, T]$ , one particular parameter  $\theta_1$  will generate different numbers of events if simulated repeatedly. This is problematic for neural



**Fig. 2** An outline of the SB-ETAS inference procedure. Samples from the prior distribution are used to simulate many ETAS sequences. A neural density estimator is then trained on the parameters and simulator outputs to approximate the posterior distribution. Samples from the posterior given the observed earthquake sequence can then be used to improve the estimate over rounds or are returned as the final posterior samples.

density estimators since even though they are successful over high dimensional data, they require a fixed dimensional input. For this reason, we fix the dimension of simulator output through calculating summary statistics of the data  $S(\mathbf{x})$ .

Works to perform ABC on the univariate Hawkes process with exponential decay kernel have found summary statistics that perform well in that setting. Ertekin et al. (2015) use the histogram of inter-event times as summary statistics as well as the number of events. Deutsch and Ross (2021) extend these summary statistics by adding Ripley’s K statistic (Ripley, 1977), which is a popular choice of summary statistic in spatial point processes (Møller and Waagepetersen, 2003). Figure 1 shows the performance of the ABC-MCMC

method developed by Deutsch and Ross (2021), using the aforementioned summary statistics. Using SNPE on the same summary statistics yields a more confident estimation of the “true” posterior which is found through MCMC sampling using the likelihood and requires far fewer simulations (10,000 versus 300,000).

The ETAS model is more complex than a univariate Hawkes process since it is both marked (i.e. it contains earthquake magnitudes) and contains a power law decay kernel which decays much more slowly than exponential making it harder to estimate (Bacry et al., 2016). For SB-ETAS we borrow similar summary statistics to Ertekin et al. (2015), namely  $S_1(\mathbf{x}) = \log(\# \text{ events})$ ,  $S_2, \dots, S_4(\mathbf{x}) = 20\text{th}, 50\text{th and } 90\text{th quantiles of}$

the inter-event time histogram. Similar to [Deutsch and Ross \(2021\)](#), we use another statistic  $S_5(\mathbf{x})$  = the ratio of the mean and median of the inter-event time histogram.

#### 4.1.1 Ripley's K Statistic

For the remaining summary statistics, we develop upon the introduction of Ripley's K statistic by ([Deutsch and Ross, 2021](#)). For a univariate point process  $\mathbf{x} = (t_1, \dots, t_n)$ , Ripley's K statistic is ([Dixon, 2001](#)),

$$K(\mathbf{x}, w) = \quad (18)$$

$$\frac{1}{\lambda} \mathbb{E}(\# \text{ of events within } w \text{ of a random event}). \quad (19)$$

Here,  $\lambda$  is the unconditional rate of events in the time window  $[0, T]$ . An estimator for the K-statistic is derived in [Diggle \(1985\)](#),

$$\hat{K}(\mathbf{x}, w) = \frac{T}{n^2} \sum_{i=1}^n \sum_{j \neq i}^n \mathbb{I}(0 < t_j - t_i \leq w). \quad (20)$$

Despite containing a double-sum, computation of this estimator has complexity  $\mathcal{O}(n)$  since  $\{t_i\}_{i=1}^n$  is an ordered sequence, i.e.  $(t_3 - t_1 < w) \Rightarrow (t_3 - t_2 < w)$ . Calculation of Ripley K-statistic therefore satisfies the complexity requirement of our procedure if the number of windows  $w$  for which we evaluate  $\hat{K}(\mathbf{x}, w)$  is less than  $\log n$ . In fact, our results suggest that less than 20 are required.

The use of Ripley's K-statistic for non-marked

Hawkes data is motivated by ([Bacry and Muzy, 2014](#)), where the authors show that second-order properties fully characterise a Hawkes process and can be used to estimate a non-parametric triggering kernel. [Bacry et al. \(2016\)](#) go on to give a recommendation for a binning strategy to estimate slow decay kernels such as a power law, using a combination of linear and log-scaling. It therefore seems reasonable to define  $S_6(\mathbf{x}) \dots S_{23}(\mathbf{x}) = \hat{K}(\mathbf{x}, w)$ , where  $w$  scales logarithmically between  $[0, 1]$  and linearly above 1.

We modify Ripley's K statistic to account for the particular interaction between marks and points in the ETAS model. Namely, the magnitude of an earthquake directly affects the clustering that occurs following it, expressed in the productivity relationship (5). In light of this, a quantity of interest is a magnitude thresholded Ripley K-statistic,

$$K_T(\mathbf{x}, w, M_T) = \quad (21)$$

$$\frac{1}{\lambda_T} \mathbb{E}(\# \text{ events within } w \text{ of an event } m_i \geq M_T). \quad (22)$$

Where  $\lambda_T$  is the unconditional rate of events above  $M_T$ . One can see that  $K_T(\mathbf{x}, w, M_0) = K(\mathbf{x}, w)$ . We estimate  $K_T$  with,


$$\hat{K}_T(\mathbf{x}, w, M_T) = \frac{T}{\omega^2} \sum_{i: m_i \geq M_T} \sum_{j \neq i} \mathbb{I}(0 < t_i - t_j \leq w), \quad (23)$$

where  $\omega$  is the number of events above magnitude threshold  $M_T$ . For general  $M_T$ , we lose the  $\mathcal{O}(n)$  complexity that the previous statistic has, instead it is  $\mathcal{O}(\omega n)$ . However, if the threshold is chosen to be large enough, evaluation of this estimator is fast. In our experiments,  $M_T$  is chosen to be  $(4.5, 5, 5.5, 6)$ , with  $w = (0.2, 0.5, 1, 3)$ . This defines the remaining statistics  $S_{24}(\mathbf{x}), \dots, S_{39}(\mathbf{x})$ .

## 5 Experiments and Results

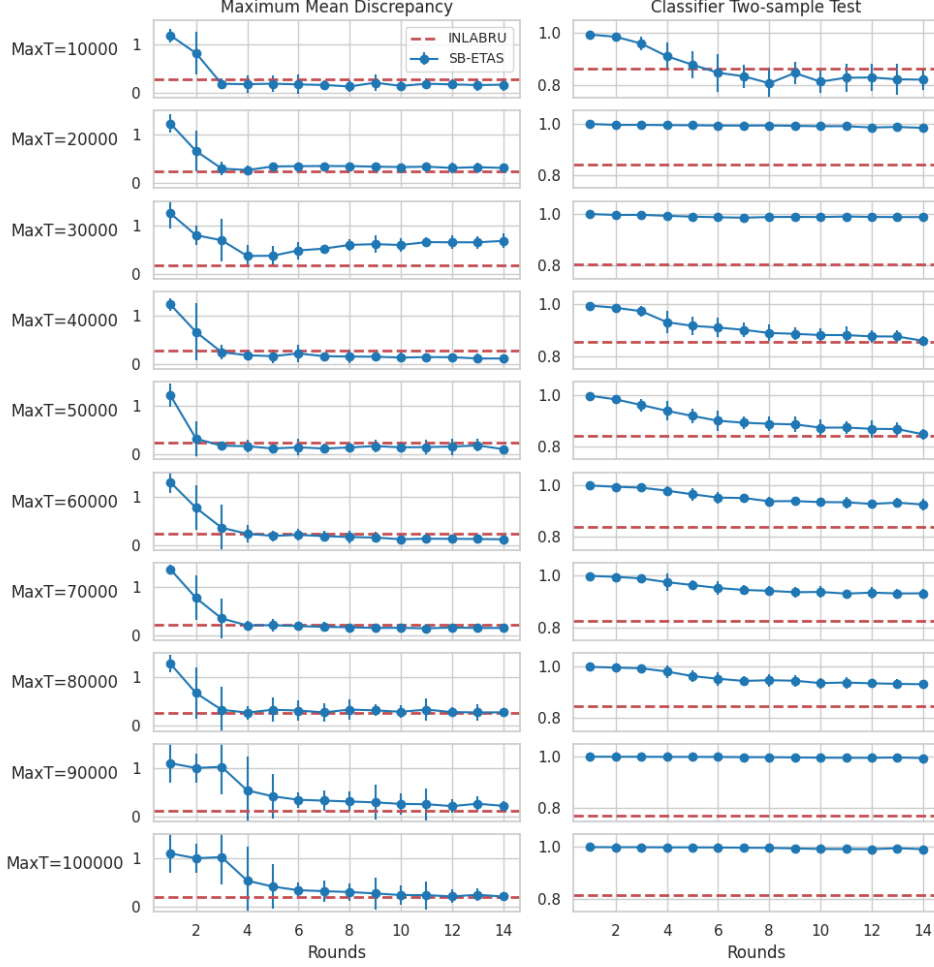
To evaluate the performance of SB-ETAS, we conduct inference experiments on a series of synthetic ETAS catalogs. On each simulated catalog we seek to obtain 5000 samples from the posterior distribution of ETAS parameters. The latent variable MCMC inference procedure, **bayesianETAS**, will be used as a reference model in our experiments since it uses the ETAS likelihood without making any approximations. We compare samples from this exact method with samples from approximate methods, **inlabru** and SB-ETAS.

Multiple catalogs are simulated from a fixed set of ETAS parameters,  $(\mu, k, \alpha, c, p) = (0.2, 0.2, 1.5, 0.5, 2)$  with magnitude of completeness  $M_0 = 3$  and Gutenberg-Richter distribution parameter  $\beta = 2.4$ . Each new catalog is simulated in a time window  $[0, T]$ , where  $T \in (10, 20, 30, 40, 50, 60, 70, 80, 90, 100, 250, 500, 1000) \times 10^3$ . This allows us to vary the number of events in each simulated catalog. Figure 3 shows the run-



**Fig. 3** The runtime for parameter inference versus the catalog size for SB-ETAS, **inlabru** and **bayesianETAS**. Separate ETAS catalogs were generated with the same intensity function parameters but for varying size time-windows. The runtime in hours and the number of events are plotted in log-log space.

time of each inference method as a function of the number of events in each catalog. Each method was run on a high-performance computing node with 8 2.4 GHz Intel E5-2680 v4 (Broadwell) CPUs, which is equivalent to what is commonly available on a standard laptop. On the catalogs with up to 100,000 events, **inlabru** is the fastest inference method, around ten times quicker on a catalog of 20,000 events. However, the superior scaling of SB-ETAS allows it to be run on the catalog of  $\sim 500,000$  events, which was unfeasible for **inlabru** given the same computational resources i.e. it exceeded a 336 hour time limit. The gradient of 2 for both **bayesianETAS** and **inlabru** in log-log space confirm the  $\mathcal{O}(n^2)$  time complexity of both inference methods. SB-ETAS, on the other hand, has gradient  $\frac{2}{3}$  which suggests



**Fig. 4** Maximum Mean Discrepancy and Classifier Two-Sample Test score for samples from each round of simulations in SB-ETAS. Each row corresponds to a different simulated catalog simulated with identical intensity function parameters but over a different length time-window. In red is the performance metric evaluated for samples from `inlabru`. 95% confidence intervals are plotted for SB-ETAS across 10 different initial seeds.

that the theoretical  $\mathcal{O}(n \log n)$  time complexity is a conservative upper-bound.

The prior distribution of parameters for each implementation are not identical since each has its own requirements. Priors are chosen to replicate the fixed implementation in the `bayesianETAS` package,

$$\mu \sim \text{Gamma}(0.1, 0.1) \quad (24)$$

$$K, \alpha, c \sim \text{Unif}(0, 10) \quad (25)$$

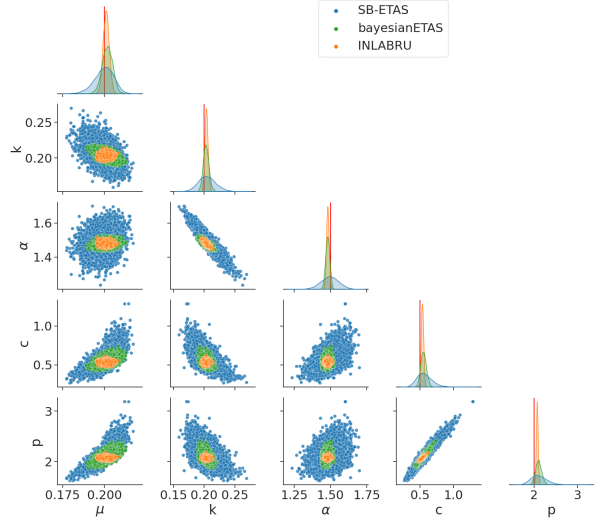
$$p \sim \text{Unif}(1, 10). \quad (26)$$

The implementation of `inlabru` uses a transformation  $K_b = \frac{K(p-1)}{c}$ , with prior  $K_b \sim \text{LogN}(-1, 2.03)$  chosen by matching 1% and 99%

quantiles with the `bayesianETAS` prior for  $K$ . SB-ETAS uses a  $\mu \sim \text{Unif}(0.05, 0.3)$  prior in place of the gamma prior as well as enforcing a non-critical parameter region  $K\beta < \beta - \alpha$  (Zhuang et al., 2012). Both the uniform prior and the restriction on  $K$  and  $\alpha$  stop unnecessarily long or infinite simulations.

Once samples are obtained from SB-ETAS and `inlabru`, we measure their (dis)similarity with samples from the exact method `bayesianETAS` using the Maximum Mean Discrepancy (MMD) (Gretton et al., 2012) and the Classifier Two-Sample Test (C2ST) (Lehmann and Romano, 2005; Lopez-Paz and Oquab, 2016). Figure 4 shows the value of these performance metrics for samples from each of 15 rounds of simulations in SB-ETAS compared with the performance of `inlabru`. Since SB-ETAS involves random sampling in the procedure, we repeat it for 10 different initial seeds and plot a 95% confidence intervals. Across the 10 synthetic catalogs, SB-ETAS performs best in terms of MMD and `inlabru` performs best in terms of C2ST. Figure 5 shows samples from the  $T = 60,000$  catalog where the performance metrics rank the inference methods differently. Samples from `inlabru` are overconfident with respect to the `bayesianETAS` samples, whereas SB-ETAS samples are more conservative. This phenomenon is shared across the samples

from all the simulated catalogs and we speculate that it accounts for the metrics ranking the methods differently.



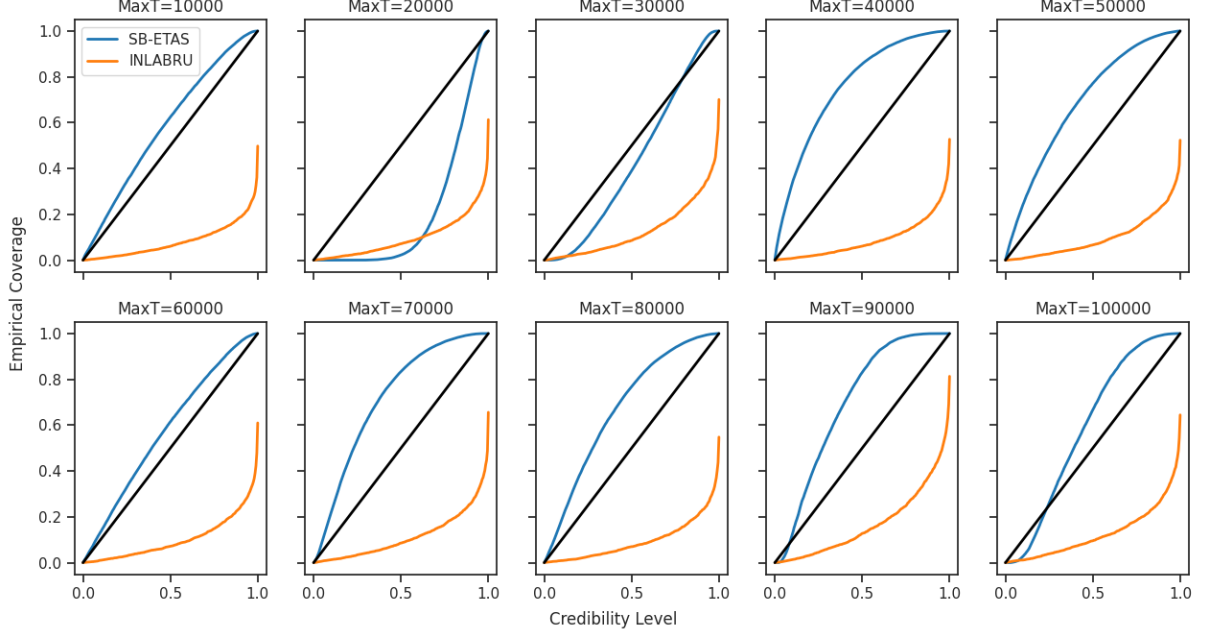
**Fig. 5** Samples from the posterior distribution of ETAS parameters for the simulated catalog with  $T = 60,000$ , for `bayesianETAS`, `inlabru` and SB-ETAS. The data generating parameters are marked in red in the diagonal plots.

A common measure for the appropriateness of a prediction’s uncertainty is the coverage property (Prangle et al., 2014; Xing et al., 2019). The coverage of an approximate posterior assesses the quality of its credible regions  $\hat{C}_{\mathbf{x}_{obs}}$  which satisfy,

$$\gamma = \mathbb{E}_{\hat{p}(\theta|\mathbf{x}_{obs})} \left( \mathbb{I}\{\hat{\theta} \in \hat{C}_{\mathbf{x}_{obs}}\} \right) \quad (27)$$

An approximate posterior has perfect coverage if its operational coverage,

$$b(\mathbf{x}_{obs}) = \mathbb{E}_{p(\theta|\mathbf{x}_{obs})} \left( \mathbb{I}\{\hat{\theta} \in \hat{C}_{\mathbf{x}_{obs}}\} \right) \quad (28)$$



**Fig. 6** Empirical estimates of the coverage of both SB-ETAS and `inlabru`. Coverage below the black line  $y = x$  indicates an overconfident approximation, whereas coverage above  $y = x$  indicates a conservative approximation.

is equal to the credibility level  $\gamma$ . The approximation is conservative if it has operational coverage  $b(\mathbf{x}_{obs}) > \gamma$  and is overconfident if  $b(\mathbf{x}_{obs}) < \gamma$  (Hermans et al., 2021). Expectations in equations (27)-(28) cannot be computed exactly and so are replaced with Monte Carlo averages, resulting in empirical coverage  $c(\mathbf{x}_{obs})$ . Figure 6 shows the empirical coverage for both SB-ETAS, averaged across the 10 initial seeds, along with `inlabru` on the 10 synthetic catalogs. `inlabru` consistently gives overconfident approximations, as the empirical coverage lies well below the credibility level. SB-ETAS has empirical coverage that indicates conservative estimates, but that is generally closer to the credibility level.

## 6 SCEDC Catalog

We now continue our evaluation of the performance of SB-ETAS on some observational data from Southern California. The Southern California Seismic Network has produced an earthquake catalog for Southern California since 1932 Hutton et al. (2010). This catalog contains many famous large earthquakes such as the 1952  $M_W$  7.5 Kern County, 1992  $M_W$  7.3 Landers, 1999  $M_W$  7.1 Hector Mine and  $M_W$  7.1 Ridgecrest sequences. The dataset contains  $N = 43,537$  events from 01/01/1981 - 31/12/2021 with earthquake magnitudes  $\geq M_W$  2.5 and can be downloaded from the Southern California Earthquake Data Center <https://service.scedc.caltech>.



[edu/ftp/catalogs/SCSN/](http://edu/ftp/catalogs/SCSN/). This size of catalog contains too many events to find ETAS posteriors using `bayesianETAS` (i.e. it would take longer than 2 weeks). Therefore we run only SB-ETAS and `inlabru` on the entire catalog and validate their performance by comparing the compensator,  $\Lambda^*(t; \theta) = \int_0^t \lambda^*(s; \theta) ds$ , with the observed cumulative number of events in the catalog  $N(t)$ .  $\Lambda^*(t; \theta)$  gives the expected number of events at time  $t$ , and therefore a model and its parameters are consistent with the observed data if  $\Lambda^*(t; \theta) = N(t)$ .

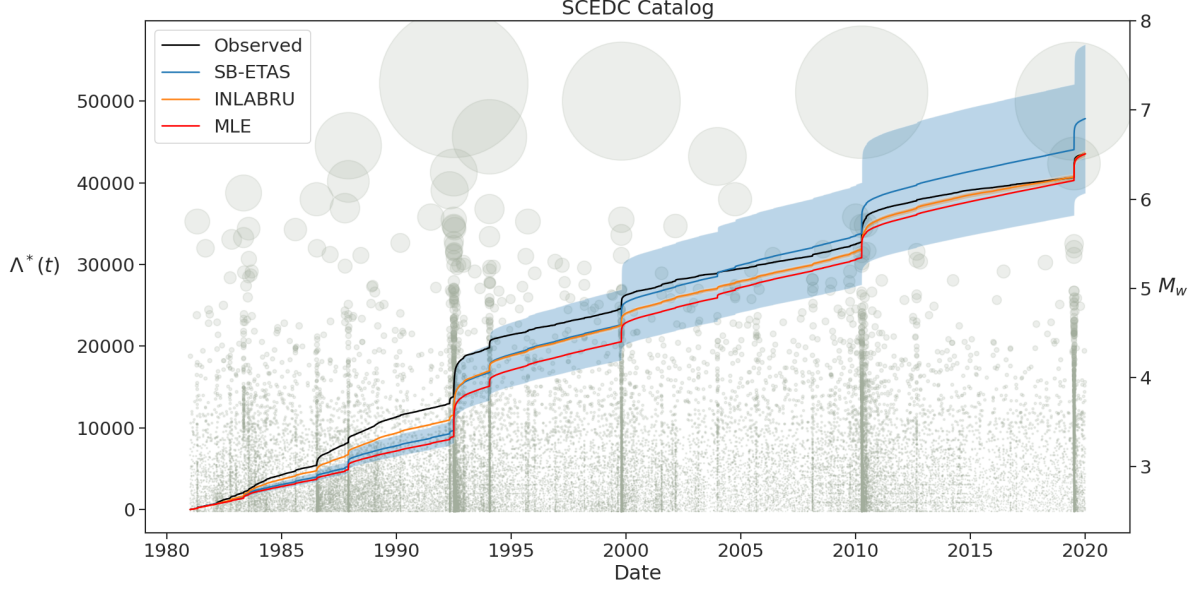
We generate 5,000 samples using SB-ETAS and `inlabru` and use each sample to generate a compensator curve  $\Lambda^*(t; \theta)$ . We display 95% confidence intervals of these curves in Figure 7, along with a curve for the Maximum Likelihood Estimate (MLE). Consistent with the synthetic experiments, we find that SB-ETAS gives a conservative estimate of the cumulative number of events across the catalog, whereas `inlabru` is overconfident and does not contain the observed number of events within its very narrow confidence interval. Both `inlabru` and the MLE fit to the total observed number of events in the catalog, since this value,  $\Lambda^*(T)$ , is a dominant term in each of their loss functions (the likelihood) during estimation.

We fix the  $\alpha$  parameter equal to the  $\beta$  parameter of the Gutenberg-Richter law  $f_{GR(m)}$ , a result

that is consistent with other temporal only studies of Southern California (Felzer et al., 2004; Helmstetter et al., 2005), and which reproduces Båth’s law for aftershocks (Felzer et al., 2002). We were unable to successfully fix  $\alpha$  for `inlabru` and therefore use the 5 parameter implementation of ETAS. Posterior distributions are displayed in Figures 8 and 9 including  $\alpha = \beta$  and free  $\alpha$  implementations of the MLE. Although the modes of the marginal distributions do not match the MLE, the SB-ETAS posteriors contain the MLE parameters within their wider confidence ranges. Since `inlabru` has much narrower confidence, although the distributions are relatively close to the MLE, the confidence ranges do not contain MLE parameters.

## 7 Discussion and Conclusion


The growing size of earthquake catalogs generated through machine learning based phase picking and an increased density of seismic networks has overtaken the ability of our current forecasting models to fit to such large datasets. Although improvements have been made to reduce the computational time of performing Bayesian inference for the ETAS model, first with `bayesianETAS` followed by `inlabru`, both of these approaches do not improve upon the scalability of inference. The implication of this is that as catalogs continue to grow in size, these methods become less feasible



**Fig. 7** The compensator  $\Lambda^*(t)$  found from estimating the ETAS posterior distribution on the SCEDC catalog (events displayed in background). 5,000 Samples from the posterior using both SB-ETAS and `inlabru` were used to generate a mean and 95% confidence interval. The compensator is compared against the observed cumulative number of events in the catalog along with the MLE.

to use. On experiments where we give SB-ETAS, `bayesianETAS` and `inlabru` access to the same 8 CPUs, only SB-ETAS could be used to fit a catalog of 500,000 events and was the fastest method for catalogs above 100,000 events. Both `inlabru` and SB-ETAS are parallelized methods and would therefore see a reduction in runtime if given access to more CPUs. This is unlike `bayesianETAS` which is not parallelized in its current implementation. It is also worth noting that although SB-ETAS and `inlabru` were given the same CPUs, `inlabru` required over 4 times the amount of memory than SB-ETAS with catalogs over 100,000 events, (Figure A1).


Comparing samples from both approximate methods to samples from `bayesianETAS` demonstrates the quality of each inference method. Our general finding is that `inlabru` provides overconfident estimates of posteriors whereas SB-ETAS gives conservative estimates. Although it might seem reasonable to judge an approximate posterior by its closeness to the exact posterior, for practical use, overconfident estimates should be penalised more than under-confident ones. Bayesian inference for the ETAS model seeks to identify a range of parameter values which are then used to give confidence over a range of forecasts. However, failure to identify regions of the



**Fig. 8** The posterior distribution of ETAS parameters found on the SCEDC catalog using SB-ETAS. This implementation of ETAS fixes  $\alpha = \beta$ . MLE parameters are plotted for comparison.

parameter space that give likely parameters would result in omission of a range of likely forecasts.

SB-ETAS provides an efficient procedure for performing Bayesian Inference for a temporal ETAS model on large earthquake catalogs. However, Operational earthquake forecasting requires spatial forecasts in addition to temporal forecasts. The spatio-temporal form of the ETAS model extends the temporal model used in the study by modeling earthquake spatial interactions with an isotropic Gaussian spatial triggering kernel (Ogata, 1998). This spatio-temporal ETAS model is also defined as a branching process and would therefore retain the  $\mathcal{O}(n \log n)$  complexity of simulation. This study has illustrated that the Ripley K-statistic is an informative summary



**Fig. 9** The posterior distribution of ETAS parameters found on the SCEDC catalog using inlabru. This implementation of ETAS has a free  $\alpha$  parameter. MLE parameters are plotted for comparison.

statistic for the triggering parameters of the temporal ETAS model. It seems fair to assume that the spatio-temporal Ripley K-statistic,

$$\hat{K}(\mathbf{x}, w_t, w_s) = \frac{AT}{n^2} \sum_{i=1}^n \sum_{j \neq i}^n \mathbb{I}(0 < t_j - t_i \leq w_t) \mathbb{I}(\|s_j - s_i\|_2 \leq w_s).$$

where  $A$  is the area of the study region, would be a reasonable choice for the spatio-temporal form of SB-ETAS. This statistic loses the  $\mathcal{O}(n)$  efficiency that the purely temporal one benefits from. Instead Wang et al. (2020) have developed a distributed procedure for calculating this statistic with  $\mathcal{O}(n \log n)$  complexity that would retain the overall time complexity that SB-ETAS has.

SB-ETAS demonstrates that inference for parameters of a clustered temporal point process can be achieved without having to evaluate an expensive likelihood function. However, generally simulation based inference is used in cases where the likelihood function is intractable. For example, consider the same ETAS branching process used in this study, but instead events are deleted with a time varying probability  $h(t)$ . The induced likelihood of this process,

$$p(\mathbf{x}|\theta) \propto \int p(\mathbf{x}, \mathbf{x}_u|\theta) \prod_{t_i \in \mathbf{x}} h(t_i) \prod_{t_j \in \mathbf{x}_u} (1 - h(t_j)) d\mathbf{x}_u, \quad (29)$$

is intractable since it involves integrating over the set of unobserved events  $\mathbf{x}_u$  (Deutsch and Ross, 2021). This time-varying detection of events is a huge source of bias in earthquake models and current methods to deal with it typically estimate the true earthquake rate from the apparent earthquake rate assuming no contribution from undetected events (Hainzl, 2016a). There are biases from ignoring such triggering (Sornette and Werner, 2005) which could be avoided by using this simulation based random deletion model with intractable likelihood Eq. 29.

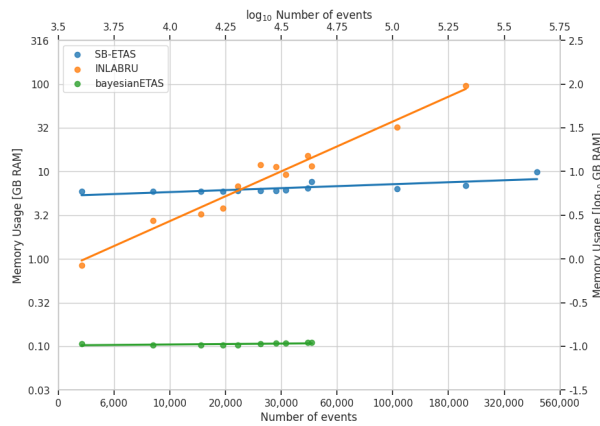
Another application of this simulation based framework is to earthquake branching models that

include complex physical dependencies. One possible example would be to calibrate the UCERF3-ETAS model, a unified model for fault rupture and ETAS earthquake clustering (Field et al., 2017). This model extends the standard ETAS model by explicitly modeling fault ruptures in California and includes a variable magnitude distribution which significantly affects the triggering probabilities of large earthquakes. This model is only defined as a simulator and uses ETAS parameters found independently to the joint ETAS and fault model. In fact, Page and van der Elst (2018) validate the models performance through a comparison of summary statistics from the outputs of the model. This validation could be extended to comprise part of the inference procedure for model parameters using the same simulation based framework as SB-ETAS.

## Appendix A Memory

### References

- Bacry, E., Jaisson, T., Muzy, J.-F.: Estimation of slowly decreasing hawkes kernels: application to high-frequency order book dynamics. *Quantitative Finance* **16**(8), 1179–1201 (2016)
- Bacry, E., Muzy, J.-F.: Second order statistics characterization of hawkes processes and non-parametric estimation. *arXiv preprint arXiv:1401.0903* (2014)



**Fig. A1** The memory usage for parameter inference versus the catalog size for SB-ETAS, `inlabru` and `bayesianETAS`. Separate ETAS catalogs were generated with the same intensity function parameters but for varying size time-windows. The memory usage in GB RAM and the number of events are plotted in log-log space.

Beaumont, M.A., Zhang, W., Balding, D.J.: Approximate bayesian computation in population genetics. *Genetics* **162**(4), 2025–2035 (2002)

Cranmer, K., Brehmer, J., Louppe, G.: The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences* **117**(48), 30055–30062 (2020)

Cattania, C., Werner, M.J., Marzocchi, W., Hainzl, S., Rhoades, D., Gerstenberger, M., Liukis, M., Savran, W., Christophersen, A., Helmstetter, A., *et al.*: The forecasting skill of physics-based seismicity models during the 2010–2012 canterbury, new zealand, earthquake sequence. *Seismological Research Letters* **89**(4), 1238–1250 (2018)

Diggle, P.: A kernel method for smoothing point

process data. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **34**(2), 138–147 (1985)

Dixon, P.: Ripley’s k function (2001)

Deutsch, I., Ross, G.J.: ABC Learning of Hawkes Processes with Missing or Noisy Event Times (2021)

Ertekin, Ş., Rudin, C., McCormick, T.H.: Reactive point processes: A new approach to predicting power failures in underground electrical systems (2015)

Felzer, K., Abercrombie, R., Ekstrom, G.: A common origin for aftershocks, foreshocks, and multiplets. *Bulletin of The Seismological Society of America - BULL SEISMOL SOC AMER* **94**, 88–98 (2004) <https://doi.org/10.1785/0120030069>

Felzer, K.R., Becker, T.W., Abercrombie, R.E., Ekström, G., Rice, J.R.: Triggering of the 1999 mw 7.1 hector mine earthquake by aftershocks of the 1992 mw 7.3 landers earthquake. *Journal of Geophysical Research: Solid Earth* **107**(B9), 6 (2002)

Field, E.H., Jordan, T.H., Page, M.T., Milner, K.R., Shaw, B.E., Dawson, T.E., Biasi, G.P., Parsons, T., Hardebeck, J.L., Michael, A.J., *et*

- al.*: A synoptic view of the third uniform california earthquake rupture forecast (ucerf3). *Seismological Research Letters* **88**(5), 1259–1267 (2017)
- Field, E.H., Milner, K.R., Hardebeck, J.L., Page, M.T., Elst, N., Jordan, T.H., Michael, A.J., Shaw, B.E., Werner, M.J.: A spatiotemporal clustering model for the third uniform california earthquake rupture forecast (ucerf3-etas): Toward an operational earthquake forecast. *Bulletin of the Seismological Society of America* **107**(3), 1049–1081 (2017)
- Gretton, A., Borgwardt, K.M., Rasch, M.J., Schölkopf, B., Smola, A.: A kernel two-sample test. *The Journal of Machine Learning Research* **13**(1), 723–773 (2012)
- Geffner, T., Papamakarios, G., Mnih, A.: Score modeling for simulation-based inference. In: *NeurIPS 2022 Workshop on Score-Based Methods* (2022)
- Gutenberg, B., Richter, C.F.: Magnitude and energy of earthquakes. *Science* **83**(2147), 183–185 (1936)
- Hainzl, S.: Apparent triggering function of aftershocks resulting from rate-dependent incompleteness of earthquake catalogs. *Journal of Geophysical Research: Solid Earth* **121**(9), 6499–6509 (2016)
- Hainzl, S.: Rate-dependent incompleteness of earthquake catalogs. *Seismological Research Letters* **87**(2A), 337–344 (2016)
- Hawkes, A.G.: Spectra of some self-exciting and mutually exciting point processes. *Biometrika* **58**(1), 83–90 (1971)
- Hermans, J., Delaunoy, A., Rozet, F., Wehenkel, A., Begy, V., Louppe, G.: A trust crisis in simulation-based inference? your posterior approximations can be unfaithful. *arXiv preprint arXiv:2110.06581* (2021)
- Helmstetter, A., Kagan, Y.Y., Jackson, D.D.: Importance of small earthquakes for stress transfers and earthquake triggering. *Journal of Geophysical Research: Solid Earth* **110**(B5) (2005) <https://doi.org/10.1029/2004JB003286> <https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2004JB003286>
- Hutton, K., Woessner, J., Hauksson, E.: Earthquake monitoring in southern california for seventy-seven years (1932–2008). *Bulletin of the Seismological Society of America* **100**(2), 423–446 (2010)
- Izbicki, R., Lee, A., Schafer, C.: High-dimensional density ratio estimation with extensions to approximate likelihood computation. In: *Artificial Intelligence and Statistics*, pp. 420–429 (2014). PMLR

- Lueckmann, J.-M., Boelts, J., Greenberg, D., Goncalves, P., Macke, J.: Benchmarking simulation-based inference. In: International Conference on Artificial Intelligence and Statistics, pp. 343–351 (2021). PMLR
- Lueckmann, J.-M., Goncalves, P.J., Bassetto, G., Öcal, K., Nonnenmacher, M., Macke, J.H.: Flexible statistical inference for mechanistic models of neural dynamics. *Advances in neural information processing systems* **30** (2017)
- Lombardi, A.M.: Estimation of the parameters of etas models by simulated annealing. *Scientific reports* **5**(1), 8417 (2015)
- Lopez-Paz, D., Oquab, M.: Revisiting classifier two-sample tests. *arXiv preprint arXiv:1610.06545* (2016)
- Lehmann, E.L., Romano, J.P.: *Testing Statistical Hypotheses*, 3rd edn. Springer Texts in Statistics, p. 784. Springer, New York (2005)
- Marzocchi, W., Lombardi, A.M., Casarotti, E.: The establishment of an operational earthquake forecasting system in Italy. *Seismological Research Letters* **85**(5), 961–969 (2014)
- Marjoram, P., Molitor, J., Plagnol, V., Tavaré, S.: Markov chain monte carlo without likelihoods. *Proceedings of the National Academy of Sciences* **100**(26), 15324–15328 (2003)
- Mancini, S., Segou, M., Werner, M., Cattania, C.: Improving physics-based aftershock forecasts during the 2016–2017 central Italy earthquake cascade. *Journal of Geophysical Research: Solid Earth* **124**(8), 8626–8643 (2019)
- Mancini, S., Segou, M., Werner, M.J., Parsons, T.: The predictive skills of elastic coulomb rate-and-state aftershock forecasts during the 2019 Ridgecrest, California, earthquake sequence. *Bulletin of the Seismological Society of America* **110**(4), 1736–1751 (2020)
- Møller, J., Waagepetersen, R.P.: An introduction to simulation-based inference for spatial point processes. In: *Spatial Statistics and Computational Methods*, pp. 143–198. Springer, ??? (2003)
- Ogata, Y.: Estimators for stationary point processes. *Ann. Inst. Statist. Math* **30**(Part A), 243–261 (1978)
- Ogata, Y.: Statistical models for earthquake occurrences and residual analysis for point processes. *Journal of the American Statistical Association* **83**(401), 9–27 (1988)
- Ogata, Y.: Space-time point-process models for earthquake occurrences. *Annals of the Institute of Statistical Mathematics* **50**(2), 379–402 (1998)
- Prangle, D., Blum, M.G., Popovic, G., Sisson,

- S.: Diagnostic tools for approximate bayesian computation using the coverage property. *Australian & New Zealand Journal of Statistics* **56**(4), 309–329 (2014)
- Prangle, D., Fearnhead, P., Cox, M.P., Biggs, P.J., French, N.P.: Semi-automatic selection of summary statistics for abc model choice. *Statistical applications in genetics and molecular biology* **13**(1), 67–82 (2014)
- Papamakarios, G., Murray, I.: Fast  $\varepsilon$ -free inference of simulation models with bayesian conditional density estimation. *Advances in neural information processing systems* **29** (2016)
- Papamakarios, G., Pavlakou, T., Murray, I.: Masked autoregressive flow for density estimation. *Advances in neural information processing systems* **30** (2017)
- Papamakarios, G., Sterratt, D., Murray, I.: Sequential neural likelihood: Fast likelihood-free inference with autoregressive flows. In: *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 837–848 (2019). PMLR
- Page, M.T., Elst, N.J.: Turing-style tests for ucerf3 synthetic catalogs. *Bulletin of the Seismological Society of America* **108**(2), 729–741 (2018)
- Rasmussen, J.G.: Bayesian inference for hawkes processes. *Methodology and Computing in Applied Probability* **15**, 623–642 (2013)
- Rasmussen, J.G.: Lecture notes: Temporal point processes and the conditional intensity function. arXiv preprint arXiv:1806.00221 (2018)
- Rathbun, S.L.: Asymptotic properties of the maximum likelihood estimator for spatio-temporal point processes. *Journal of Statistical Planning and Inference* **51**(1), 55–74 (1996)
- Reinhart, A.: A review of self-exciting spatio-temporal point processes and their applications. *Statistical Science* **33**(3), 299–318 (2018)
- Ripley, B.D.: Modelling spatial patterns. *Journal of the Royal Statistical Society: Series B (Methodological)* **39**(2), 172–192 (1977)
- Rhoades, D., Liukis, M., Christophersen, A., Gerstenberger, M.: Retrospective tests of hybrid operational earthquake forecasting models for canterbury. *Geophysical Journal International* **204**(1), 440–456 (2016)
- Rezende, D., Mohamed, S.: Variational inference with normalizing flows. In: *International Conference on Machine Learning*, pp. 1530–1538 (2015). PMLR
- Ross, G.J.: Bayesian estimation of the etas model for earthquake occurrences. *Bulletin of the Seismological Society of America* **111**(3), 1473–1480



- (2021)
- Rue, H., Riebler, A., Sørbye, S.H., Illian, J.B., Simpson, D.P., Lindgren, F.K.: Bayesian computing with inla: a review. *Annual Review of Statistics and Its Application* **4**, 395–421 (2017)
- Ross, Z.E., Trugman, D.T., Hauksson, E., Shearer, P.M.: Searching for hidden earthquakes in southern california. *Science* **364**(6442), 767–771 (2019)
- Serafini, F., Lindgren, F., Naylor, M.: Approximation of bayesian hawkes process with inlabru. *Environmetrics*, 2798 (2023)
- Stockman, S., Lawson, D.J., Werner, M.J.: Forecasting the 2016–2017 central apennines earthquake sequence with a neural point process. *Earth’s Future* **11**(9), 2023–003777 (2023) <https://doi.org/10.1029/2023EF003777>. e2023EF003777 2023EF003777
- Seif, S., Mignan, A., Zechar, J.D., Werner, M.J., Wiemer, S.: Estimating etas: The effects of truncation, missing data, and model assumptions. *Journal of Geophysical Research: Solid Earth* **122**(1), 449–469 (2017)
- Sharrock, L., Simons, J., Liu, S., Beaumont, M.: Sequential neural score estimation: Likelihood-free inference with conditional score based diffusion models. *arXiv preprint arXiv:2210.04872* (2022)
- Sornette, D., Werner, M.J.: Apparent clustering and apparent background earthquakes biased by undetected seismicity. *Journal of Geophysical Research: Solid Earth* **110**(B9) (2005)
- Taroni, M., Marzocchi, W., Schorlemmer, D., Werner, M.J., Wiemer, S., Zechar, J.D., Heiniger, L., Euchner, F.: Prospective csep evaluation of 1-day, 3-month, and 5-yr earthquake forecasts for italy. *Seismological Research Letters* **89**(4), 1251–1261 (2018)
- Utsu, T., Ogata, Y., *et al.*: The centenary of the omori formula for a decay law of aftershock activity. *Journal of Physics of the Earth* **43**(1), 1–33 (1995)
- Utsu, T.: Aftershocks and earthquake statistics (1): Some parameters which characterize an aftershock sequence and their interrelations. *Journal of the Faculty of Science, Hokkaido University. Series 7, Geophysics* **3**(3), 129–195 (1970)
- Vargas, N., Gneiting, T.: Bayesian point process modelling of earthquake occurrences. Technical report, Technical Report, Ruprecht-Karls University Heidelberg, Heidelberg, Germany ... (2012)
- White, M.C., Ben-Zion, Y., Vernon, F.L.: A detailed earthquake catalog for the san jacinto

- fault-zone region in southern california. *Journal of Geophysical Research: Solid Earth* **124**(7), 6908–6930 (2019)
- Wang, Y., Gui, Z., Wu, H., Peng, D., Wu, J., Cui, Z.: Optimizing and accelerating space-time ripley’s k function based on apache spark for distributed spatiotemporal point pattern analysis. *Future Generation Computer Systems* **105**, 96–118 (2020)
- Woessner, J., Hainzl, S., Marzocchi, W., Werner, M., Lombardi, A., Catalli, F., Enescu, B., Cocco, M., Gerstenberger, M., Wiemer, S.: A retrospective comparative forecast test on the 1992 landers sequence. *Journal of Geophysical Research: Solid Earth* **116**(B5) (2011)
- Wang, Q., Schoenberg, F.P., Jackson, D.D.: Standard errors of parameter estimates in the etas model. *Bulletin of the Seismological Society of America* **100**(5A), 1989–2001 (2010)
- Xing, H., Nicholls, G., Lee, J.: Calibrated approximate bayesian inference. In: *International Conference on Machine Learning*, pp. 6912–6920 (2019). PMLR
- Zhu, W., Beroza, G.C.: Phasenet: a deep-neural-network-based seismic arrival-time picking method. *Geophysical Journal International* **216**(1), 261–273 (2019)
- Zhuang, J., Harte, D.S., Werner, M.J., Hainzl, S., Zhou, S.: Basic models of seismicity: Temporal models (2012)
- Zhuang, J., Ogata, Y., Vere-Jones, D.: Analyzing earthquake clustering features by using stochastic reconstruction. *Journal of Geophysical Research: Solid Earth* **109**(B5) (2004)
- Zhuang, J., Ogata, Y., Wang, T.: Data completeness of the kumamoto earthquake sequence in the jma catalog and its influence on the estimation of the etas parameters. *Earth, Planets and Space* **69**(1), 1–12 (2017)