2023-06-22

Reconstructing Historical Earthquake-Induced Tsunamis: Case Study of 1852 Event Using the Adjoint Method Combined with HMC

Chelsey Noorda
Brigham Young University

Reconstructing Historical Earthquake-Induced Tsunamis: Case Study of 1852 Event Using the Adjoint Method Combined with HMC

Chelsey Noorda

A thesis submitted to the faculty of
Brigham Young University
in partial fulfillment of the requirements for the degree of

Master of Science

Jared P. Whitehead, Chair
Emily Evans
Blake Barker

Department of Mathematics

Brigham Young University

ABSTRACT

Reconstructing Historical Earthquake-Induced Tsunamis: Case Study of 1852 Event Using the Adjoint Method Combined with HMC

Chelsey Noorda
Department of Mathematics, BYU
Master of Science

Seismic hazard analysis aims to estimate human risk due to natural disasters such as earthquakes. To improve seismic hazard analysis, our group is focused on earthquake induced tsunamis and the use of statistical models to reconstruct historical earthquakes. Based on the estimated wave heights given in anecdotal historical descriptions, we created observational probability distributions to model the historically recorded observations and constructed a prior distribution on the relevant earthquake parameters based on known seismicity of a given region. Then we used the software package GeoClaw, and a Metropolis-Hastings sampler to obtain a posterior distribution of earthquake parameters that most closely matches the historically recorded tsunami.

Our method was tested on two main events that occurred in 1820 and 1852 in central and eastern Indonesia respectively. The random walk Metropolis-Hastings sampler we employed appeared to recover the causal earthquake quite well, but the computational costs were prohibitive even though both scenarios we considered were relatively simple. To improve the sampling procedure, we have focused on advanced sampling techniques such as Hamiltonian Monte Carlo (HMC) where the gradient of the forward model (Geoclaw) is required. This is problematic however as this gradient is not available computationally. To mitigate this problem, we make use of a linearized adjoint solver for the shallow water equations, and exact gradient calculations for the Okada earthquake rupture model, yielding a surrogate gradient that leads to improved sampling.

Keywords: HMC, adjoint method, applied math

Acknowledgements

Thank you to my parents and grandparents who have supported me in everything I do. They were my lending ear whenever I just needed to talk and my encouragement when I wanted to quit. I am grateful for their love and assistance in helping me to reach my goals. I do not know where I would be today without them.

I would like to thank my advisor, Dr. Jared Whitehead, for pushing me beyond my limits and helping me manage the stress of graduate school. He was very patient and never hesitated to reexplain any topic I did not understand along with answering the millions of questions that I would come up with as he reexplained. Dr. Emily Evans taught my favorite undergraduate course (Math 314) and continued to help me see my potential as a mathematician. She would meet with me outside of class and helped me make academic changes to achieve the understanding and grades I desired. Dr. Blake Barker has the kindest heart and would clarify the gaps in my understanding outside of class. He saw the potential in me I did not see in myself. It was also a privilege to work with Ron Harris and gain insight into his expertise in geology and love for Indonesia.

I would like to also thank my fellow graduate students and friends who made these past two years bearable. These are the friendships I will cherish forever and will be grateful for their help and support in accomplishing all of my classes. I would like to thank Raelynn Wonnacott, Jake Callahan, Taylor Paskett, Marja Crawford, Paul Smith, Rachel Young, and Carol Herrera for our collaboration efforts in working on this project.

Lastly, thank you to my best friend Marien Dille who was one of the only people I could talk to about my work who also partially knew what I was talking about. She was my biggest cheerleader and always encouraged me to finish my degree. I would also like to thank my roommate and friend Christine Akagi for being an example to me of the reward that comes through hard work and for the many late-night talks about life that helped me gain perspective on the future. Lastly, I would like to thank my siblings, extended family, and friends for their continued support as I move forward to the next phase of my life.

Contents

List of Figures

Chapter 1. Introduction

In 2022, the highest risk index for natural disasters in Indonesia was tsunamis [3]. This trend has continued for many years and still is one of the leading causes of deaths due to natural disasters in the world's 4th most populous country. The economic and social damage that this type of natural disaster inflicts upon residents makes it nearly impossible for locals to rebuild afterward. The end goal of this research project is to provide enough information to more accurately predict potential tsunami hazards to educate residents so they are prepared for evacuation in the incidence of a tsunami warning. Increased understanding of this natural hazard is essential to predicting the potential hazard to life and property.

To reach the end goal described above, our research group focuses on seismic hazard analysis which aims to estimate human risk due to natural disasters such as earthquakes. In order to improve seismic hazard analysis, we have focused on earthquake induced tsunamis and the use of statistical models to reconstruct historical earthquakes. The idea behind this approach is to use anecdotal records from previous earthquake induced tsunamis to build a model that can accurately post-dict the past in hopes of understanding the current and/or future hazards. In a general sense, seismologists and geologists are unable to make accurate predictions of future seismic hazards without an in-depth understanding of the seismic history of the region [4, 5, 6, 7], and hence we must have as much comprehension of the seismic history as possible to develop an adequate hazard assessment.

Previous efforts from our group have developed a statistical inverse method that has been applied to two earthquake induced tsunamis that occurred in Indonesia in 1820 (South Sulawesi) and 1852 (Banda Arc) [2]. These will be referred to as the 1820 event and the 1852 event for the remainder of this thesis. These investigations used a random walk Metropolis-Hastings sampler which appeared to recover the causal earthquake quite well, but at a very high computational cost even though both events we considered were relatively small geographically (each tsunami simulation was relatively inexpensive). This is prohibitive to

our end goal since we want to be able to apply our model to larger geographical events such as the Lisbon earthquake induced tsunami of 1755 [8] which spread all the way across the Atlantic. Our current approach would take roughly 2 hours to collect a single sample (even running on 24 cores on the supercomputing resources at Brigham Young University) and with a random walk sampler we need to collect roughly 200,000 samples to analyze the success of the trajectory predictions. That is approximately 400,000 hours to collect our desired samples which equates to 16,667 days or a little over 45 years.

Due to the long data collection time frame, we are considering alternate approaches to improve the sampling procedure. We focus on an advanced sampling technique called Hamiltonian Monte Carlo (HMC, see [9] for example) which requires the gradient of the forward model, which in our case is the tsunami simulation package Geoclaw [10]. The gradient in our case is not available computationally, so we will instead approximate the gradient using a linearized adjoint solver of the shallow water equations implemented in Geoclaw. This surrogate gradient combined with the exact gradient calculations for the Okada earthquake rupture model yield a surrogate gradient that should lead to more accurate sampling, i.e. the desired posterior distribution can be approximated with far fewer samples than when using a random walk proposal. Both mathematical innovations, the adjoint method and using HMC, will be discussed in more detail in the next section.

Chapter 2. Background

In this section the two main mathematical innovations over the method of [2] which are used to sample potential earthquake sources for the tsunami trajectories will be explained in more detail. These two methods are: Hamiltonian Monte Carlo, which is abbreviated as HMC, and the adjoint method.

2.1 Hamiltonian Monte Carlo

The method outlined and utilized in [2, 11] performs Markov Chain Monte Carlo (MCMC) sampling with Metropolis-Hastings [12] adjustment where the proposed states are selected via a random walk in parameter space. The parameter space we wish to sample over is the space of all possible earthquakes in a given geographic region. For MCMC, an initial starting point is identified in parameter space and then the selected forward model (Geoclaw in our case) uses that initial set of earthquake parameters to simulate the resultant tsunami and produce wave heights at each of the historically observed locations. A relative probability of how well this simulation matches the historical data and how well it matches geological prior information is assigned to these parameters. The random walk proposal then proposes a new earthquake by randomly moving in parameter space and obtaining a similar probability for this new event. A Metropolis-Hastings update, as outlined in Algorithm 1, [13, 14] is then calculated to either accept or reject the proposal based on the plausibility of the predicted wave heights (and wave arrival times). This method has worked well for the smaller geographical area relevant to the 1852 event [2] where the observed tsunami reached all of the relevant observation locations in under one hour, and even worked decently for the 1820 event [11] where the longest observation time was between 3 and 4 hours real time. We want to make our model applicable to situations where the arrival time is several hours such as the Lisbon 1755 tsunami. The primary issue is that random walk MCMC rejects close to 3/4 of the proposed samples, wasting nearly 75% of computational expense. To combat this disadvantage and approximate the desired distribution in a reasonable time, we need to use an algorithm that converges faster than a random walk based algorithm. This is exactly why we have decided to use Hamiltonian Monte Carlo (HMC). The following picture illustrates the convergence difference between MCMC and HMC. As seen in Figure 2.1, random walk MCMC has a higher rejection rate meaning that it will reject a lot of proposals (typically 75%) before it accepts one. On the other hand, Hamiltonian dynamics has a fairly high acceptance rate and thus converges to the desired target distribution faster than MCMC.

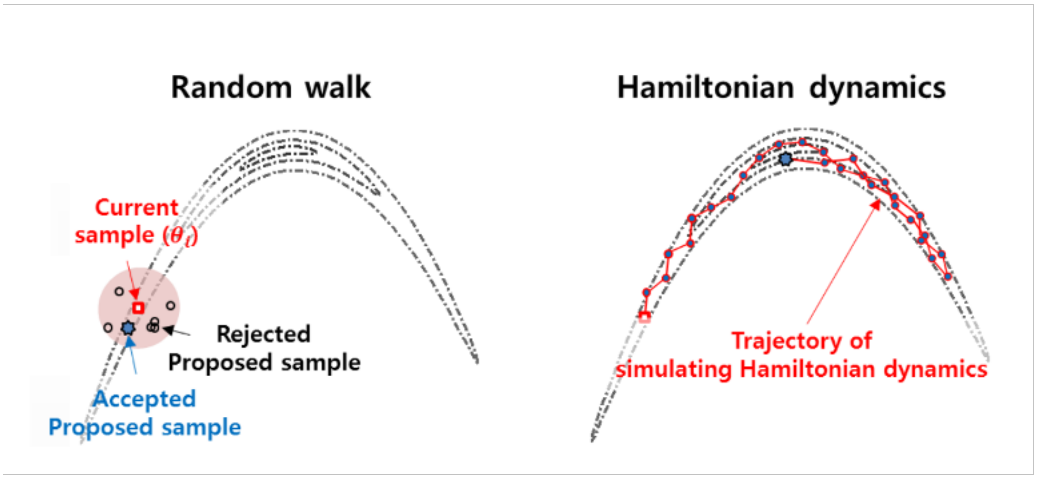

Figure 2.1: A cartoon illustration of the difference between random walk MCMC and HMC. The target distribution is depicted via contours of the probability density function. The random walk (left) is allowed to step in a random direction within an exponentially decaying radius (the step size is normally distributed), but most proposed parameters are immediately rejected, and even those accepted proposals are close to the original sample. HMC on the other hand is able to move along the target distribution and propose values that are far away from the original sample and yet still yield reasonable probable values. [1]

In order to understand how HMC and Hamiltonian dynamics work together, we are going to consider by analogy the canonical distributions, the distributions of the possible states that are in thermal equilibrium, and how their energy changes for a given conserved physical system. We will represent the probability density function of a tsunami induced earthquake as $P(x)$ with a corresponding energy function $E(x)$ of the physical state x (parameters of the earthquake in our case). This canonical distribution with temperature T and normalizing constant Z has the following probability density function (pdf):

$$P(x) = \frac{1}{Z} \exp\left(-E(x)/T\right). \tag{2.1}$$

To simplify, we will assume the total temperature of our system is 1, i.e. $T = 1$ and we will define the energy function to be the Hamiltonian of a conserved system with position q (below we will substitute $x = q$) and the artificially produced momentum p. This makes $E(q, p) = H(q, p) = U(q) + K(p)$ where $U(q)$ is the potential energy and $K(p)$ is the kinetic

energy. The joint distribution of q and p then satisfies

$$\begin{aligned} P(q,p) &= \frac{1}{Z}\exp\left(-H(q,p)\right) \\ &= \frac{1}{Z}\exp -(U(q)+K(p)) \\ &= \frac{1}{Z}\exp\left(-U(q)\right)\exp\left(-K(p)\right). \end{aligned} \tag{2.2}$$

We will assume that both q and p are independent and let q represent our earthquake parameters and p represent an artificially introduced momentum. Hamiltonian systems such as this will conserve the energy of the system, that is $H(q,p)$ will remain fixed for all time. As the momentum p is artificially introduced to move the sample parameters q across the target distribution, we do not focus on the definition of the kinetic energy $K(p) = \frac{1}{2}p^T M^{-1} p$ where M is a mass matrix (assumed to be diagonal in all of our calculations). For Monte Carlo sampling, the potential energy function for q is as follows:

$$U(q) = -\log[\pi(q)L(q|D)] \tag{2.3}$$

where $\pi(q)$ is the prior distribution and $L(q|D)$ is the likelihood function given the data. Hence the potential energy of the system is a constant value away from the logarithm of the posterior distribution (the target distribution for Bayesian inverse problems, see below for a more thorough discussion on Bayes' Rule). The rest of the HMC calculations, including the acceptance update step, are depicted in Algorithm 2.

Algorithm 1 Metropolis-Hastings

for t in 1, 2, ... **do**
  Draw $x \sim q(.|\theta_t)$
  Set $\alpha = \min(1, \frac{p(x|d)p(x)}{p(\theta_t|d)p(\theta_t)} * \frac{q(\theta_t|x)}{q(x|\theta_t)})$
  Draw $u \sim \text{unif}(0,1)$
  if $u \leq \alpha$ **then**
    $\theta_{t+1} = x$
  else
    $\theta_{t+1} = \theta_t$
  end if
end for

The basic idea is that the parameters q (position) is updated according to Hamilton's equations:

$$q' = \frac{\partial H}{\partial p}, \quad p' = -\frac{\partial H}{\partial q}, \tag{2.4}$$

where the artificial momentum p is evolved to ensure that the system is conserved, i.e. H remains constant. The system (2.1) is evolved forward for a limited amount of time via a Leapfrog time-stepping mechanism with the final value of q being the new proposal in the sampler. The acceptance rate is dictated by how far off the numerical discretization gets from conserving the Hamiltonian $H(q, p)$.

Algorithm 2 HMC

```
q = current_q
p = norm(len(q),0,1)                          ▷ independent standard normal variates
current_p = p
p = p - epsilon * grad_U(q) / 2       ▷ Make a half step for momentum at the beginning

# Alternate full steps for position and momentum
for i in 1, 2, ..., L do
    q = q + epsilon * p                                  ▷ Make a full step for the position
    if (i != L) then     ▷ Make a full step for the momentum, except at end of trajectory
        p = p - epsilon * grad_U(q)
    end if
end for

p = p - epsilon * grad_U(q) / 2                ▷ Make a half step for momentum at the end
p = -p          ▷ Negate momentum at end of trajectory to make the proposal symmetric

# Evaluate potential and kinetic energies at start and end of trajectory
current_U = U(current_q)
current_K = sum(current_p ** 2) / 2
proposed_U = U(q)
proposed_K = sum(p ** 2) / 2

# Accept or reject the state at end of trajectory
if (runif(1) < exp(current_U-proposed_U+current_K-proposed_K)) then
    return q                          ▷ Return the position at the end of trajectory (accept)
else
    return current_q                                     ▷ Return initial position (reject)
end if
```

Since we are using Hamiltonian dynamics, this method requires the calculation of

$$\frac{\partial U}{\partial y_i} \tag{2.5}$$

This may seem like a setback because computing derivatives is computationally expensive, but this is why we are going to use the adjoint method to make this calculation less computationally expensive. A thorough description of the adjoint and its usefulness in approximating the gradient is given below.

2.2 The Adjoint Method and Surrogate Gradients

The main contribution to this project is the use of the linearized adjoint to approximate the gradient of Geoclaw with respect to its initial parameters. Before getting into the specifics of the adjoint method for this application, we will present an explanation of adjoint methods using linear algebra and investigate a one dimensional case.

We begin by defining two finite dimensional inner product spaces W and V. We will define $A : V \Longrightarrow W$ to be a linear transformation. The inner product in W and V will be the same as the standard inner product in $\mathbb{R}^n$, i.e. if B is a matrix representing the linear transformation A (below we will often not distinguish between the linear transformation A and its matrix representation) then $\langle Bx, y\rangle = x^T B^T y$.

Definition 2.1. The adjoint of a linear operator A, denoted as A^* is defined such that the following equation is satisfied for all $x \in V$ and $y \in W$:

$$\langle Ax, y\rangle = \langle y, A^* x\rangle. \tag{2.6}$$

Example 2.2. When $V = \mathbb{R}^n$ and $W = \mathbb{R}^m$ this would be the same as finding a matrix A^* such that

$$x^T Ay = yA^* x$$

is satisfied, which occurs exactly when $A^* = A^T$, the transpose of the matrix representation of the linear transformation.

A finite dimensional example. For an illustration of how the adjoint is used, consider the linear ordinary differential equation (ODE) in $\mathbb{R}^n$:

$$x' = Ax, \quad x(0) = \tilde{x}, \tag{2.7}$$

where $x \in \mathbb{R}^n$ and A is an $n \times n$ real valued matrix. The solution of (2.7) is $x(t) = e^{At}\tilde{x}$ where the matrix exponential is defined via a standard Taylor series similar to the scalar exponential.

Now suppose that we wish to determine the effect that the initial conditions $\tilde{x}$ have on the solution at some future time $t = t_f$. We will focus on a single entry of the solution, that is we want to determine the influence of the initial conditions on the kth entry of $x(t)$ which we denote as

$$x_k(t) = \langle x(t), e_k \rangle, \tag{2.8}$$

where e_k is the vector of all zeros with a one in the kth position. We want to calculate the gradient of $x_k(t)$ with respect to $\tilde{x}$. For this particular example that isn't all that difficult, but using the adjoint operator we can illustrate how we can do this even for the infinite dimensional setting in which (2.7) becomes a partial differential equation rather than an ODE.

First, we note that the adjoint of the exponential of a matrix is simply the exponential of the adjoint of the matrix, that is

$$\left(e^A\right)^* = e^{A^*}. \tag{2.9}$$

Hence we can write

$$x_k(t_f) = \langle x(t_f), e_k \rangle \tag{2.10}$$

$$= \langle e^{At_f}\tilde{x}, e_k \rangle \tag{2.11}$$

$$= \langle \tilde{x}, e^{A^*t_f} e_k \rangle. \tag{2.12}$$

It follows that if we want to compute the derivative of $x_k(t_f)$ with respect to the jth entry

of $\tilde{x}$ (which we denote as $\tilde{x}_j$) then we need to compute

$$\frac{\partial x_k(t_f)}{\partial \tilde{x}_j} = \langle e_j, e^{A^* t_f} e_k \rangle, \tag{2.13}$$

and all the other entries of the desired gradient can be found by varying j.

As pointed out above, the use of the adjoint for this simple, finite-dimensional example, was rather unnecessary. The true value of using the adjoint to calculate the gradient is only evident in the infinite-dimensional case where the explicit gradient is otherwise entirely unavailable. This example does illustrate some very key points however, the first of which is that we can note that the final calculation in (2.13) is really given by

$$\frac{\partial x_k(t_f)}{\partial \tilde{x}_j} = \langle e_j, y(t_f) \rangle, \tag{2.14}$$

where $y(t)$ solves the adjoint equation 'backward' in time from the desired endpoint e_k, that is $y(t)$ solves

$$y' = A^* y, \quad y(0) = e_k. \tag{2.15}$$

Hence, to calculate the gradient via the adjoint method we need to solve the adjoint equation *backward* from the endpoint we are interested in (e_k) at time $t = 0$ to the final time $t = t_f$, and then the inner product of this solution with the desired location will yield the derivative. An analogous calculation will work even for the infinite dimensional (PDE) setting so long as the underlying system is linear.

Adjoint method for the shallow water equations. To implement HMC for the historical earthquake-tsunami inversion problem described in more detail below, we need to be able to compute the gradient of the modeled tsunami wave height at a specific location and time with respect to the initial wave disturbance, i.e. we seek to compute

$$\frac{\partial \eta(x_k, t_k)}{\partial h_0(x)}, \tag{2.16}$$

where $\eta(x, t)$ is the sea surface perturbation from a quiescent mean value, $h_0(x)$ is the initial disturbance of η, x_k is the physical location of interest and t_f is the time at which the observation takes place. The forward evolution of the sea surface height η is simulated using the Geoclaw [10] package which uses a dynamically adaptive mesh in a finite volume

framework to solve the nonlinear shallow water equations. Ideally we would implement automatic differentiation into Geoclaw to calculate $\frac{\partial \eta}{\partial h_0}$, but the complexity of the source code precludes this solution in a reasonable time frame. As the tsunami propagation follows the nonlinear shallow water equations, a full adjoint approach is not feasible (adjoint methods on nonlinear systems require frequent linearization about the current state), and so we will instead rely on a linearized adjoint first introduced in [15] for adaptive mesh refinement, that provides a surrogate gradient.

In two spatial dimensions the shallow water equations are given by

$$\frac{\partial h}{\partial t} = \mathcal{L}_1(u, v, h), \tag{2.17}$$

$$\frac{\partial (hu)}{\partial t} = \mathcal{L}_2(u, v, h), \tag{2.18}$$

$$\frac{\partial (hv)}{\partial t} = \mathcal{L}_3(u, v, h), \tag{2.19}$$

where $h(x, y, t)$ is the water height, and $u(x, y, t)$ and $v(x, y, t)$ are the horizontal velocities (hu and hv are the corresponding components of the momentum). Each of the $\mathcal{L}_j$ are spatial differential operators. Specifically

$$\mathcal{L}_1 = -\partial_x(hu) - \partial_y(hv), \tag{2.20}$$

$$\mathcal{L}_2 = -\partial_x \left(hu^2 + \frac{g}{2} h^2 \right) + \partial_y(huv) - ghB_x, \tag{2.21}$$

$$\mathcal{L}_3 = -\partial_x(huv) - \partial_y \left(hv^2 + \frac{g}{2} h^2 \right) - ghB_y, \tag{2.22}$$

where g is the gravitational constant, and $B(x, y)$ is the bathymetric seafloor. We are interested in the evolution of the sea surface height $\eta(x, y, t) = h(x, y, t) + B(x, y)$, and in particular approximating the gradient of this function with respect to the initial perturbed state. Ideally we would linearize (2.20)-(2.22) repeatedly about the current solution, but for tsunami propagation we actually have a more appealing procedure.

In the deep open ocean the tsunami wave is a very small perturbation relative to the depth of the ocean, and so the linearized shallow water equations are valid. It is only near the shoreline, or immediately following the uplift due to the earthquake that the linearized equations fail to adequately capture the dynamics. Hence, we will investigate the adjoint

of the linearized shallow water equations, not linearized about individual (time-varying) solutions, but linearized about the deep ocean approximation.

Following [15] we let $\mu = hu$ and $\gamma = hv$, and linearize the shallow water equations about a flat surface with zero velocities $\overline{u} = \overline{v} = 0$ with $\overline{h} = \overline{\eta} - B(x, y)$ (we suppose that the full velocity is $u = \overline{u} + u$ where $\tilde{u}$ is small for example). This leads to the linearized shallow water equations:

$$\frac{\partial \tilde{\eta}}{\partial t} = L_1, \tag{2.23}$$

$$\frac{\partial \tilde{\mu}}{\partial t} = L_2, \tag{2.24}$$

$$\frac{\partial \tilde{\gamma}}{\partial t} = L_3, \tag{2.25}$$

where as before

$$L_1 = -\tilde{\mu}_x - \tilde{\gamma}_x, \tag{2.26}$$

$$L_2 = -g\overline{h}\tilde{\eta}_x, \tag{2.27}$$

$$L_3 = -g\overline{h}\tilde{\eta}_y. \tag{2.28}$$

To re-establish this in a notation familiar to the previous simple example of a finite-dimensional system (and dropping the $\tilde{\cdot}$ going forward), let

$$q(x, y, t) = \begin{pmatrix} \eta \\ \mu \\ \gamma \end{pmatrix}, \tag{2.29}$$

so that q satisfies

$$q_t(x, y, t) + F(x, y)q_x(x, y, t) + G(x, y)q_y(x, y, t) = 0, \tag{2.30}$$

where

$$F(x, y) = \begin{pmatrix} 0 & 1 & 0 \\ g\overline{h} & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad G(x, y) = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ g\overline{h} & 0 & 0 \end{pmatrix}. \tag{2.31}$$

It follows that the adjoint equation for q is given by:

$$\hat{q}_t + \left(F^T(x,y)\hat{q}\right)_x + \left(G^T(x,y)\hat{q}\right)_y = 0, \tag{2.32}$$

where the adjoint is derived using inner product induced by the integral L^2 norm.

Let Ω refer to the rectangular region in (x, y) space wherein the tsunami propagates, that is Ω is the set of all latitude-longitude coordinates that are simulated via Geoclaw for each simulation. Following [15] we enforce the boundary conditions $\hat{q}^T(x,y)F(x,y) = 0$ and $\hat{q}^T(x,y)G(x,y) = 0$ for $(x,y) \in \partial\Omega$. As mentioned above, we are interested in the wave height at specific locations (x_k, y_k) and times t_k when the wave reaches those locations. Rather than focusing on the sea surface perturbation $\eta(x,y,t)$ (which can easily be computed from the bathymetry) we consider h, i.e.

$$h(x_k, y_k, t_k) = \int_\Omega q(x,y,t_k)^T \hat{q}(x,y,0;x_k,y_k)dxdy = \langle q(x,y,t_k), \hat{q}(x,y,0;x_k,y_k)\rangle_{L^2}, \tag{2.33}$$

where we will specify that

$$\hat{q}(x,y,0;x_k,y_k) = \begin{pmatrix} \delta(x-x_k, y-y_k) \\ 0 \\ 0 \end{pmatrix}, \tag{2.34}$$

i.e. $\hat{q}(x,y,t=0;x_k,y_k)$ is the Dirac delta function centered at the desired location (x_k, y_k).

Following the same derivation as above for the finite dimensional linear system, we will approximate the gradient of h via

$$\frac{\partial h(x_k,y_k,t_k)}{\partial h_0(\tilde{x},\tilde{y})} \approx \langle q_0(x,y;\tilde{x},\tilde{y}), \hat{q}(x,y,t_k;x_k,y_k)\rangle_{L^2}, \tag{2.35}$$

where $\hat{q}(x,y,t;x_k,y_k)$ satisfies the adjoint equation with initial condition given above in (2.34) and

$$q_0(x,y;\tilde{x},\tilde{y}) = \begin{pmatrix} \delta(x-\tilde{x}, y-\tilde{y}) \\ 0 \\ 0 \end{pmatrix}. \tag{2.36}$$

The calculation of (2.35) leads to the gradient of the wave height at specific locations (and times) with respect to certain points in the initial sea surface perturbation, and only requires

the calculation of the linearized adjoint solution at the specific time t_k with the specified initial condition (2.34).

Specifying the Dirac delta function as the initial condition is computationally impossible, so following [15] we instead specify $\hat{q}(t = 0)$ as a narrow Gaussian centered at (x_k, y_k). In fact, this is done for all of the observation locations of interest, i.e. $k = 1, 2, \ldots, 9$ (see [2] for more detail) at the same time, as each observation location is sufficiently spatially separated from each other to not invoke any overlap in the initial condition. The solution of the linearized adjoint problem for these initial conditions has already been computed in [2] to use in defining the adaptivity of the underlying spatial mesh. The solution field $\hat{q}$ is saved at a series of discrete times t_n. To utilize this computation, we identify the value of n so that $t_k \approx t_n$ as close as possible, and read in the corresponding data saved in an appropriate adjoint file. Geoclaw by default saves the $\hat{h}$, $\hat{h}\hat{u}$, $\hat{h}\hat{v}$ and $\hat{\eta} = \hat{h} + B$ fields. As the bathymetry $B(x, y)$ is unchanged throughout the simulation (in both the forward model and the linearized adjoint model) then

$$\frac{\partial h}{\partial h_0} = \frac{\partial \eta}{\partial h_0}, \tag{2.37}$$

so we need only read in the $\hat{\eta}$ field from the adjoint solution for the selected value of t_n.

It is worth noting at this point that (2.35) yields a single value as the derivative of the wave height at a specific location (and time) with respect to the initial wave perturbation at a single location. This means that the desired gradient

$$\frac{\partial h(x_k, y_k, t_k)}{\partial h_0(x, y)} \tag{2.38}$$

is actually a function of x and y. In practice this is an array with values specified at each latitude-longitude value of (x, y). Even so, we emphasize the form of the derivative shown in (2.35) as this illustrates how to utilize this gradient in applications of the chain rule as discussed below.

Using the adjoint in the chain rule. The previous subsection focused on calculating the derivative of modeled tsunami wave heights with respect to initial perturbations of the sea surface, but this isn't the entire story. The initial sea surface perturbation is not the

function we are interested in, but instead we are interested in determining a parameterization of the earthquake that generated that sea surface perturbation. To map from earthquake parameters to sea surface perturbation, Geoclaw takes advantage of the Okada model [16, 17] which parameterizes an earthquake with the following parameters:

- latitude-longitude of the earthquake centroid (center of the earthquake rupture), denoted as lat and $long$ here.
- depth of the earthquake centroid below the surface of the earth in kilometers, denoted as $depth$ or d below.
- length and width (in kilometers) of the rectangular rupture zone of the specified earthquake, denoted l and w respectively.
- slip (in meters) is the distance that the rectangular rupture zone instantaneously jumps in this model, denoted as s.
- strike is the orientation angle of the fault with respect to due north (in degrees, not radians unfortunately).
- dip is the angle which the fault makes with the vertical direction (also in degrees).
- rake is the angle relative to the strike angle that the actual slip occurs along (in degrees).

If we refer to each of these earthquake parameters as the Okada parameters and denote them as the vector valued variable z, then what we really want to compute is

$$\frac{\partial h(x_k, y_k, t_k)}{\partial z_i}, \tag{2.39}$$

where each value of the vector z is denoted as z_i.

To identify how this gradient is performed, let $\mathcal{O}$ denote the Okada model, that is we allow $\mathcal{O}(z) = h_0(x, y)$ be the function (array in practice) that represents the initial sea surface perturbation induced by the Okada parameters z. It follows that $\frac{\partial \mathcal{O}}{\partial z_i}$ is a function

of (x, y) for every value of i. We want to combine this via the chain rule with $\frac{\partial h(x_k, y_k, t_k)}{\partial h_0(x,y)}$ which was discussed above, however we note that although both of these terms are functions of (x, y) (or arrays across the discrete grid in the numerical implementation), $\frac{\partial h(x_k, y_k, t_k)}{\partial z_i}$ is a scalar, i.e. a single value. This is precisely why we need to utilize the notation used in (2.35) when computing the gradient.

Note that (2.35) specifies the gradient of $h(x_k, y_k, t_k)$ with respect to the initial data at a specific point $(\tilde{x}, \tilde{y})$. The partial derivative of the initial condition at that specific point with respect to any of the Okada parameters would then be $\frac{\partial h_0(\tilde{x}, \tilde{y})}{\partial z_i}$. We however need to compute the derivative of the entire initial condition with respect to each Okada parameter which implies that we need to compute (via the chain rule):

$$\frac{\partial h(x_k, y_k, t_k)}{\partial z_i} = \sum_{(\tilde{x}, \tilde{y})} \frac{\partial h(x_k, y_k, t_k)}{\partial h_0(\tilde{x}, \tilde{y})} \frac{\partial h_0(\tilde{x}, \tilde{y})}{\partial z_i}, \tag{2.40}$$

if there was a finite number of points $(\tilde{x}, \tilde{y})$ in the domain. Instead, we want to consider this summation over a continuous set of points $(\tilde{x}, \tilde{y}) \in \Omega$ so that the derivative instead becomes:

$$\frac{\partial h(x_k, y_k, t_k)}{\partial z_i} = \int_\Omega \frac{\partial h(x_k, y_k, t_k)}{\partial h_0(\tilde{x}, \tilde{y})} \frac{\partial h_0(\tilde{x}, \tilde{y})}{\partial z_i} d\tilde{x} d\tilde{y} = \left\langle \hat{\eta}(x, y, t_k; x_k, y_k), \frac{\partial \mathcal{O}}{\partial z_i}(x, y) \right\rangle_{L^2}, \tag{2.41}$$

where we have used the earlier definitions of the Okada model and the derivation of the adjoint-based gradient calculation. Note that we would have arrived at the same formula if we had considered the formula in (2.35) as an operator which acted via the chain rule on $\frac{\partial \mathcal{O}}{\partial z_i}$, the derivative of the Okada model with respect to the earthquake parameter z_i.

The Okada derivative. The Okada model detailed in [16, 17] is based on the Green's function solution of a linear, elastic PDE on a semi-infinite half-space (the boundary being the surface of the earth). Ignoring the curvature of the earth, there is an explicit formula for the resultant Green's function which is then integrated over the specified rupture rectangle specified by the length and width, providing an explicit, analytic formula for the seafloor displacement (which in turn yields an exact formula for $h_0(x, y)$). Although there is an exact formula for the seafloor displacement, it is quite messy and requires a rather involved several lines of code to implement so that calculating $\frac{\partial O}{\partial z_i}$ is not a straightforward nor simple process.

To address this issue, R. Molen (an ugrad student at BYU) converted the Okada model implemented in Geoclaw to the Jax library. Automatic differentiation in jax then allowed for a straightforward calculation of the desired gradient. Unfortunately even with the optimized methods introduced in jax, the computational cost of computing these derivatives is extremely prohibitive (the derivative is calculated at every single grid point on the domain), particularly on a CPU (forward propagation of the Geoclaw model is really only feasible on a CPU cluster, not on GPU's yet). As we desire to calculate these derivatives several times for each proposed sample, we implemented the required derivatives for subsets of the full domain Ω, i.e. $\frac{\partial \mathcal{O}}{\partial z_i}$ can be computed on $\omega_l \subset \Omega$ for any rectangularly defined sub-region ω_l. Eventually this will allow for parallelization across the subsets ω_l, but for now we only use this to eliminate parts of the full domain Ω where the adjoint is near zero.

Sample parameters. As discussed in more detail in [2] and more particularly in [18], the Okada parameters z_i are NOT the parameters we search over in MCMC. These earthquake parameters are not independent and some are more correlated/dependent than others. For instance the specific example described below of an earthquake in the Banda Sea has the most likely source of the known subduction zone along the outer Banda Arc meaning that each latitude-longitude value will map to a likely earthquake depth, dip, strike, and rake as dictated by the geometry of the fault itself and provided by existing data [19]. The length, width, and slip are also interdependent for recorded earthquakes so that searching over each of these parameters independently is not logical. For this reason, we actually search over a set of *sample parameters*:

- c_{lat} and c_{lon} the centroid latitude and longitude coordinates.
- m, the magnitude of the earthquake (defined as a logarithmic scaling of the product of the length, width and slip).
- Δw and Δl, the logarithmic deviation of the length and width from a data fitted regression with magnitude.

- d_0 the depth offset from a mapped depth value.

Hence the Okada parameters z are really a function of these sample parameters (which we denote as ζ) with a function we refer to as the *map to Okada* so that $z = M2O(\zeta)$. Hence the actual gradient approximation that we use is

$$\frac{\partial h(x_k, y_k, t_k)}{\partial \zeta_j} = \left\langle \hat{\eta}(x, y, t_k; x_k, y_k), \sum_i \frac{\partial \mathcal{O}}{\partial z_i}(x, y) \frac{\partial z_i}{\partial \zeta_j} \right\rangle_{L^2}, \tag{2.42}$$

where the derivative z_i is the ith entry of the map $M2O(\zeta)$ and is detailed more closely in [18].

2.3 Gradient of the Prior

Calculation of the gradient described above is only part of the problem when addressing implementation of HMC. The potential energy in the formulation of HMC requires a derivative of the full posterior distribution which will also necessitate calculating a derivative of the prior distribution (all of the previous discussion applies only to the likelihood). The gradient of the prior is much more straightforward however, as precise formulae are available for all of the different components of the prior distribution, or at least a finite difference approach is quite viable for the depth-map that arises from the data provided in SLAB2 [19]. Details for this calculation are contained in [18].

Chapter 3. Tsunami inversion

As stated earlier, the initial goal of this thesis was to develop an improved sampling strategy for the inverse problem of identifying historical earthquakes from sparse and anecdotal evidence of tsunami observations. This section first describes the previous efforts using a random walk Monte Carlo method for a specific historical event, explaining previous results and clarifying the need to consider better sampling strategies. We then describe the actual implementation of HMC using a surrogate linearized adjoint gradient to the same problem

and summarize the advances and difficulties posed by using the adjoint based surrogate gradient.

3.1 Previous Work

To employ Bayesian inference to the tsunami inversion problem, we needed to identify the observations and describe them probabilistically. We started this process by collecting anecdotal records from historical accounts including the recently translated Wichmann catalog [20, 21, 6]. These records included newspaper clippings, journal articles, and any other written form that described the wave heights and arrival times of the tsunami. Below is the English translation of an account from the Dutch colonial soldiers stationed on the island of Banda Neira:

> *1852, November 26, 7:40.*
>
> At Banda Neira, barely had the ground been calm for a quarter of an hour when the flood wave crashed in...The water rose to the roofs of the store-houses...and reached the base of the hill on which Fort Belgica is built on.

From this description, we constructed an observational probability distribution that assigned maximum probability to the wave height that we most closely associate with this observational record. For instance, standard construction at the time were homes and store houses that were approximately $2.5m$, and most often built on stilts (to avoid storm surge and flooding). In addition, on the morning of November 26, 1852 the Banda Sea was under an extremely (seasonally and daily) low tide so that for a tsunami to reach the 'roofs of the store-houses' then we estimate that the wave must have been approximately $6.5m$. Due partially to the level of uncertainty in this historical record, and to the lack of a definitive measurement we identify the wave height observational distribution at Banda Neira as a normal distribution centered at $6.5m$ with a variance of $2m$. The observational distribution on the arrival time is similarly constructed as a normal distribution centered on 15 minutes,

but with a reasonably wide variance to allow for likely arrival times between $10-20$ minutes.

Similar historical accounts to that described above are identified for nine different physical locations for the 1852 earthquake and tsunami, two of which include an observation of arrival time of the wave. All of the constructed observational probability distributions are depicted in Fig. 3.1 for the 1852 event. The distributions for the observation at Banda Neira are depicted in green.

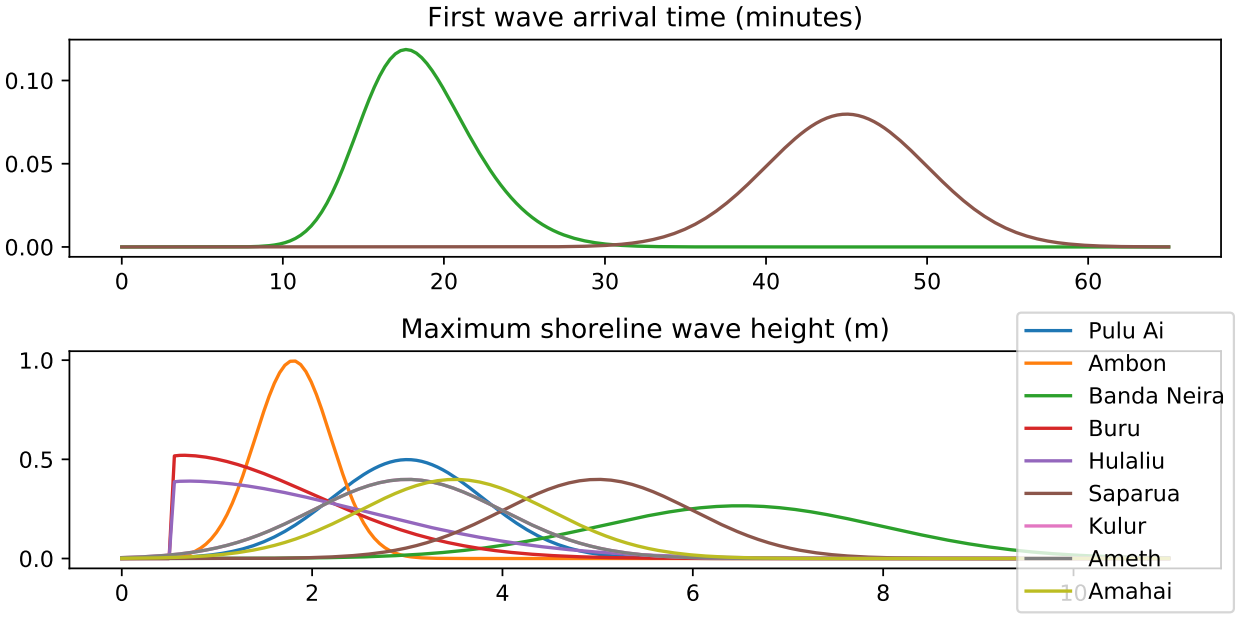

Figure 3.1: All likelihood densities for the observations at each observation location for the 1852 event. Each density represents an interpretation of an account from the Wichmann catalog.

After careful construction of a prior distribution on all of the sample parameters, [2] used the random walk Metropolis Hastings algorithm to search over parameter space to identify the most probable parameters and approximate the relevant posterior distribution. [2] ran 168,000 total tsunami simulations, eventually identifying a posterior distribution that was seismically realistic and matched the observational data and prior distribution adequately. The sampled posterior distribution from [2] on the earthquake centroid location is depicted in Fig. 3.2.

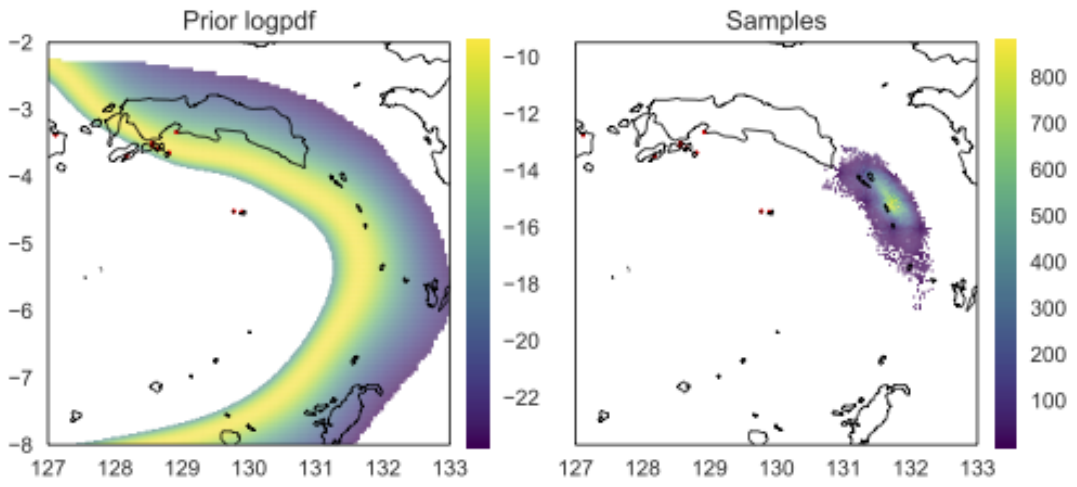

Figure 3.2: The sampled posterior distribution (right) for latitude-longitude coordinates of the earthquake centroid location for the 1852 event compared to the log-prior distribution of the centroid location (left).

The collection of the samples necessary to create Fig. 3.2 representing the posterior for the 1852 event took approximately two weeks of real time computational resources, running 14 separate chains with each chain running on 24 cores at a time. Using the standard computational resources available through BYU's supercomputing center, each forward tsunami simulation required approximately 1.5-2 minutes of compute time (parallelized across 24 cores). This computational time corresponds to a 1 hour simulation time evolution of the tsunami (the wave typically reached the farthest observation point in approximately 45 minutes). In [22], a different event from south Sulawesi in 1820 is investigated that had simulation times of 5-6 hours leading to a computational time of 10-12 minutes for each forward run of the tsunami model. As a result this 1820 event took over two months of computational time to generate the desired posterior distribution.

The 1852 Banda Sea and 1820 south Sulawesi events are just the first of several events we intend to investigate, but the computational time required for these additional events is

even higher than it was for the 1820 event. As an example, in 1755 there was an earthquake in the eastern Atlantic Ocean near Lisbon, Portugal [8]. This earthquake and the resultant tsunami devastated the city and the entire Portugese empire. The tsunami generated by this earthquake impacted not only the Portugese coastline, but was also observed in northern Africa, the British Isles, and all the way across the Atlantic in Martinique and Barbados. Recreating this event will necessarily require simulating a tsunami that crosses the entire Atlantic Ocean which would take approximately 2 hours of compute time on our current resources for a single simulation. Even running multiple chains simultaneously this represents a prohibitively costly investigation that would take several years of real time to generate the necessary 150,000+ simulations to adequately resolve the posterior distribution using a random walk proposal.

This leads to the proposed work in this thesis. As described earlier, we seek an improved sampling strategy that will resolve the posterior with a significantly reduced number of necessary simulations, and hence the current consideration of surrogate gradient HMC methods. We employ this approach for the 1852 Banda Sea event here, as the posterior is already sufficiently resolved from [2] so that we can make comparisons between the results obtained from the surrogate HMC approach outlined above. In the following, we use the same prior distribution and observational probability distributions used in [2]. In addition, all settings for the forward model Geoclaw are identical and we use the same bathymetry files etc..

3.2 Current HMC results

3.2.1 Preliminaries. Before running the surrogate HMC chains, we needed to run the adjoint problem with initial sea surface perturbation specified at each observation location as described above. The benefit of the current approach is that we need to only solve the linearized adjoint problem one time, and upon saving the relevant output files, can use that solution for every forward sample run of Geoclaw. Hence the computational temporal cost

of using the adjoint is negligible, except that the spatial complexity (storage space for the output files) and the input/output operations for reading those files in can be prohibitive. We determined the number of adjoint files we would save by first examining the longest arrival time recorded in our anecdotal records. The longest arrival time recorded was around 45 minutes so we decided to double that time to account for outliers, though we do not expect any sample to produce an arrival time that long. This means that our final time would be 5400 seconds. We saved the linearized adjoint output every 60 seconds, resulting in 90 output files. The specific output file needed to calculate the surrogate gradient was then identified as the arrival time of the wave rounded to the nearest minute.

Once the adjoint files were collected, we ran into the issue of spatial complexity because the files were so large. As stated previously, the adjoint file consists of the value of the adjoint at a series of specific sub-grids of the entire grid with these sub-grids ranging in width from 36 to 38 grid points and in length from 44 to 46. The variance in the sub-grid size is due to the use of ghost cells for sub-regions near the boundary of the overall region which had to be removed to simplify the output. There is a stored set of values on each sub-grid calculated for the following fields: $\hat{h}$, $\hat{h}\hat{u}$, $\hat{h}\hat{v}$ and $\hat{\eta} = \hat{h} + B$. Since we only needed the data for the displacement of water, the sub-grids for $\hat{\eta}$ were all that would be needed. After much finagling with Geoclaw's bounds and the output file settings, smaller files which only contained the desired $\hat{\eta}$ sub-grids were created and used throughout the remainder of the code, and the ghost cells were removed so that each sub-grid was a standard size. This cut down computational times significantly as the grid sizes were smaller and already in the correct format for their desired use.

For the initial set of collected samples, we relied on the linearized adjoint for a surrogate gradient. However the current and proposed parameters in HMC were orders of magnitude different representing a drastic change in the momentum, but almost no movement in the position (sample values). This caused the acceptance probability to almost universally yield a NaN value so that every sample was rejected. This caused us to re-evaluate our method and

determine if there was something we were missing. While the linearized adjoint gives very good approximations of the gradient of the maximal wave height at specific locations, it yields no information on the gradient of the arrival time for the wave (either the maximal wave or the 'first' wave). Essentially we had ignored the effects that the arrival times would have on the gradient which led to highly un-physical proposals that were automatically rejected. Luckily for the 1852 event there were only two observations of arrival time at Banda Neira and Saparua, so we only needed to adjust these two observations in the surrogate gradient to capture this information.

To approximate the arrival time of the wave, we relied on the fact that velocity is equal to distance divided by time traveled. For the linearized shallow water equations over a deep ocean of average depth H, the phase speed of the wave is $u = \sqrt{gH}$ where g is the gravitational constant, and so the approximate arrival time of the wave is $T = \frac{d}{\sqrt{gH}}$ where d is the distance traveled. To calculate d we used the Haversine function hav which computes the greatest circle distance between two points on the spherical approximation to the earth. Hence the arrival time of a wave observed at the point $(lat, long)$ that originated from (x_0, y_0) is approximated as

$$T(lat, long, x_0, y_0) = \frac{hav(lat, long, x_0, y_0)}{\sqrt{gH}}, \tag{3.1}$$

where $hav(lat, long, x_0, y_0)$ calculates the greatest circle distance from the origin (x_0, y_0) to $(lat, long)$. In the end, we don't really want to calculate $T(lat, long, x_0, y_0)$ but we want to calculate the derivative of this with respect to x_0 and y_0 which required calculating the derivative of the Haversine formula which we coded up using a previously coded version of hav from [2].

Lastly, we needed to multiply $\frac{\partial T}{\partial x_0}$ or $\frac{\partial T}{\partial y_0}$ by the derivative of observational probability for the arrival time. The observational distribution for arrival time at Saparua was a normal distribution centered at 45 minutes with a standard deviation of 5 minutes and for Banda Neira the arrival time was modeled as a skew normal distribution centered at 15 minutes with a standard deviation of 5 minutes and a skew normal constant of 2. Since it is extremely

difficult to calculate the derivative of a skew-normal distribution, we instead estimated the arrival time to be a normal distribution with the same mean and standard deviation and used the derivative of the normal distribution instead (the skew parameter plays a very minor role, particularly in evaluating the derivative of the distribution). Since the observational distribution was not skewed by that much, we feel this is a fair generalization to make. After combining both of these gradient calculations, the new calculation of the surrogate derivative of the likelihood with respect to latitude and longitude is as follows

$$\frac{\partial L}{\partial y_0} = \frac{\frac{\partial hav}{\partial y_0}}{\sqrt{gH}} * \frac{\partial obs}{\partial T} \tag{3.2}$$

$$\frac{\partial L}{\partial x_0} = \frac{\frac{\partial hav}{x_0}}{\sqrt{gH}} * \frac{\partial obs}{\partial T} \tag{3.3}$$

where L is the likelihood and *obs* represents the observational distributions which are standard normal distributions for both locations. From the latitude-longitude coordinates of the earthquake centroid (x_0, y_0) and observation location $(lat, long)$ we compute an estimated arrival time, but more importantly approximate the derivative of that arrival time with respect to the latitude-longitude location of the earthquake centroid. Potential concerns with this surrogate gradient of the wave arrival time includes:

(i) We have assumed a constant velocity of the tsunami wave which is clearly invalid over the region of propagation. The water depth is certainly not constant throughout the region of interest.

(ii) The phase speed $u = \sqrt{gH}$ is an excellent approximation for tsunami wave propagation of a single wave that acts linearly, an assumption that is also not born out in practice as earthquake induced tsunamis have multiple waves that nonlinearly interact with each other.

(iii) The approach outlined above ignores the effect of varying coastlines. For instance, the observation location at Banda Neira is on the west side of the island while we anticipate that the earthquake originated on the east of the island. The method discussed here for

approximating arrival times ignores this and acts as if the tsunami will travel straight across the island.

(iv) We have treated the tsunami as propagating from a point source at the earthquake centroid. In reality earthquakes create seafloor deformation that spreads across a significantly large area which spans several hundreds of kilometers, i.e. the initial wave disturbance is a region on the sea surface and not a single point. To correct this in our approximate calculation above, we would need to identify the point on the initial seafloor deformation closest to the observation point, and estimate the arrival time from that position.

3.2.2 Actual results. The sampling efficiency of HMC depends largely on the hyper parameters ϵ (the time step), and the total number of time steps taken for each proposal. If ϵ is too large, the leapfrog scheme is numerically unstable and the Hamiltonian structure is violated which leads to a low acceptance rate because there will be a higher integration error. On the other hand, if ϵ is too small with the same number of time steps, this leads to computational problems as it will take a long time to generate a single proposal. In addition if ϵ is too small we risk not moving the parameters at all i.e. the proposal will be severely restricted to be close to the current sample. If there are a small number of time steps, then the samples appear to show a high auto-correlation between iterations, i.e. they barely move in parameter space. If the number of time steps is large, then parameters might start to return to their original state because they have a large trajectory length [23], however more time steps also allows for more numerical error to accumulate which can lead to frequent rejections. To investigate the utility of HMC, we selected two different values of ϵ: 0.0001 and 0.00001, and three numbers of time steps: 2, 4, and 6, and then we evaluated the acceptance ratio of HMC after collecting 200 samples for each. The results of our hyper parameter search with different values of ϵ and the number of time steps with the same initial conditions can be seen in Fig. 3.3.

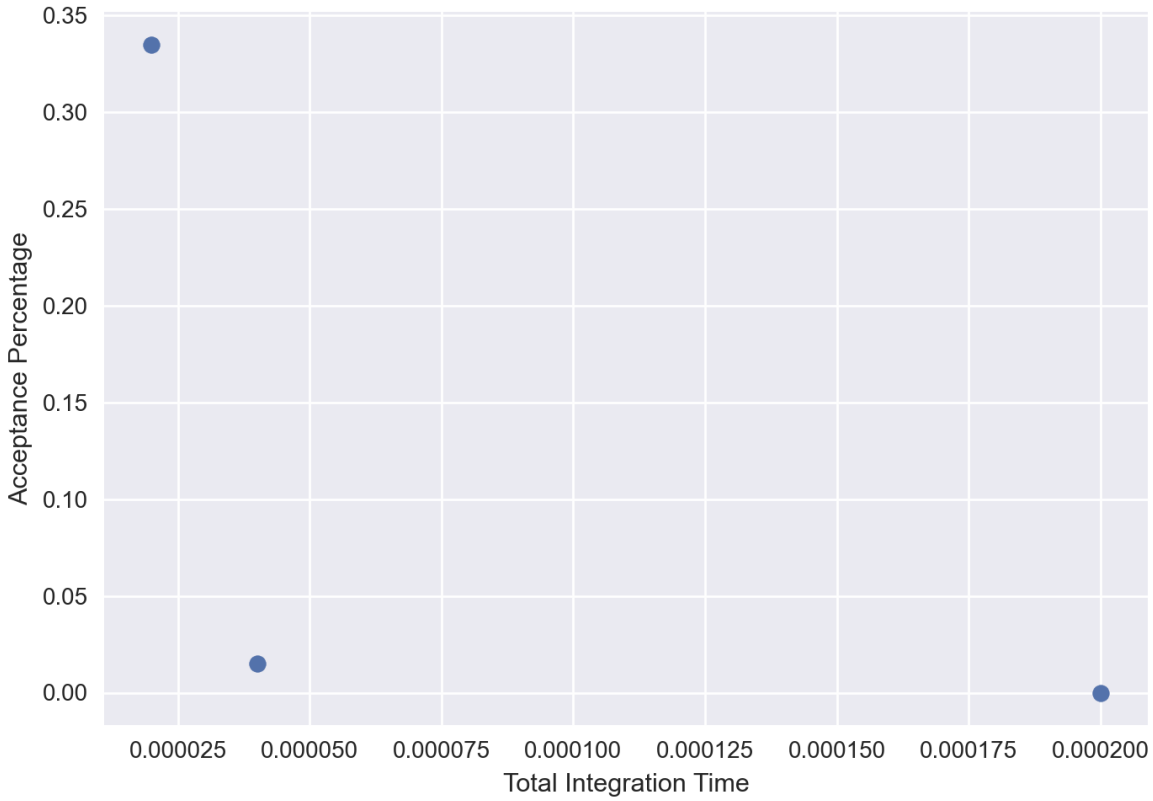

Figure 3.3: The acceptance rate versus the total integration time (we found that the total integration time rather than distinct values of ϵ or total number of time steps were the determining factor dictating the acceptance rate).

Figure 3.3 demonstrates the real problem we had with the surrogate gradient for arrival time. The highest acceptance ratio we achieved of 34% was only achieved using a very small integration time and corresponding value of ϵ. Pushing the integration time longer (whether with increased ϵ or number of time steps) led to a very low acceptance rate, due to the disparity between the true gradient and the surrogate described above. The problem with the lowest integration time depicted in Fig. 3.3 is that although the acceptance rate is significantly improved from the expected 23.4% for random walk Monte Carlo, such a short integration time restricts the proposal to limited exploration of the parameter space, i.e. the proposal is not that dissimilar from the current sample meaning that exploration of the parameter space is still restricted. In addition (not shown in Fig. 3.3) several other

chains that we initialized with different initial sample values and various other values of ϵ and the number of time steps, led to rejection of all of the proposed samples, i.e. the HMC integration was so poor that numerical differences between the true gradient and the surrogate overwhelmed the Metropolis-Hastings adjustment.

As described earlier, the approach outlined above for the surrogate gradient relies on a rather costly integration of the linearized adjoint data against the jax calculated gradient of the Okada model. Although there are likely parallelization methods we can use to speed up this calculation, in its current form, generation of each proposed sample requires several minutes of computational time on the supercomputer (as opposed to 1-2 minutes in [2]). This means that we were unable to collect a large number of samples, particularly with such a poor acceptance rate in HMC. Reconstructed histograms of the observations are shown in Fig. 3.4 for those chains we ran that successfully obtained a nonzero acceptance rate. We note that due to the severe limitation on the number of total samples we collected, the histograms are not very filled out, particularly when compared with Fig. 3.5 which is the identical figure created using all 168,000 samples generated via random walk.

Although the data represented in Fig. 3.4 is much more sparse than that in Fig. 3.5, it is clear that the linearized adjoint surrogate HMC method generates data in the 'correct' region of sample space, i.e. the simulated observations are a subset of those identified in [2]. Unfortunately the prohibitive cost of using the adjoint calculation in the surrogate, and the poor approximation provided by our model of arrival time appears to make this surrogate HMC impractical.

Chapter 4. Sensitivity Analysis

As described earlier, the linearized adjoint approach utilized here is not only useful for a surrogate HMC method, but also provides an estimate of the gradient of the wave heights at specific locations with respect to any/all of the Okada earthquake parameters. This gives us a method to estimate the sensitivity of the forward model with respect to the Okada model

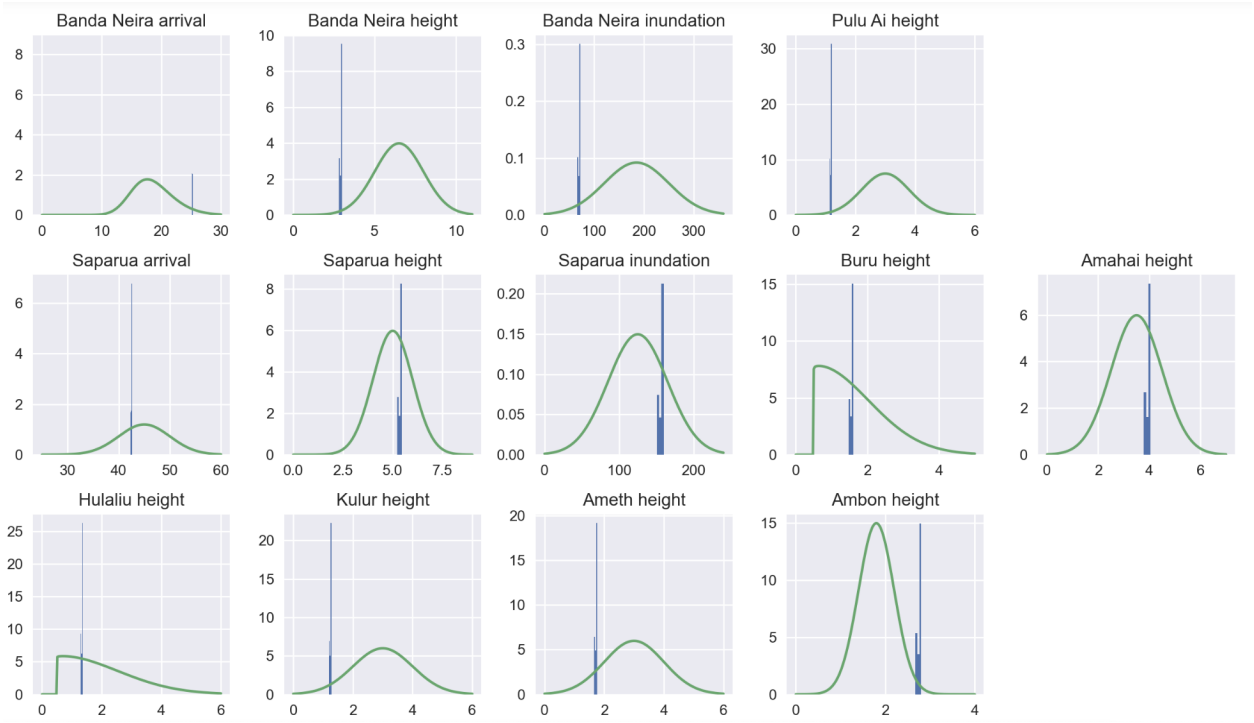

Figure 3.4: Histograms of the observations generated from the successful HMC chains. The green curves are the observational probability distributions that were created in [2] and used in the calculation of the likelihood.

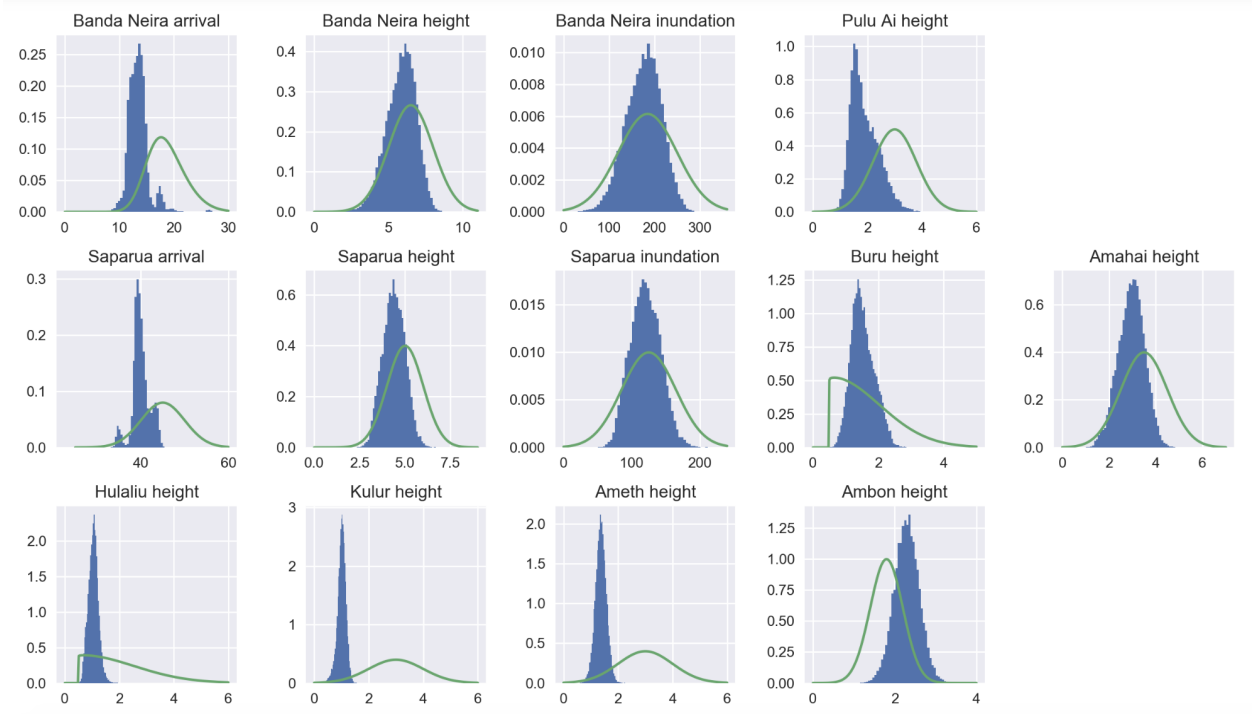

Figure 3.5: The same as Fig. 3.4 but with the data collected from [2] with all 168,000 samples.

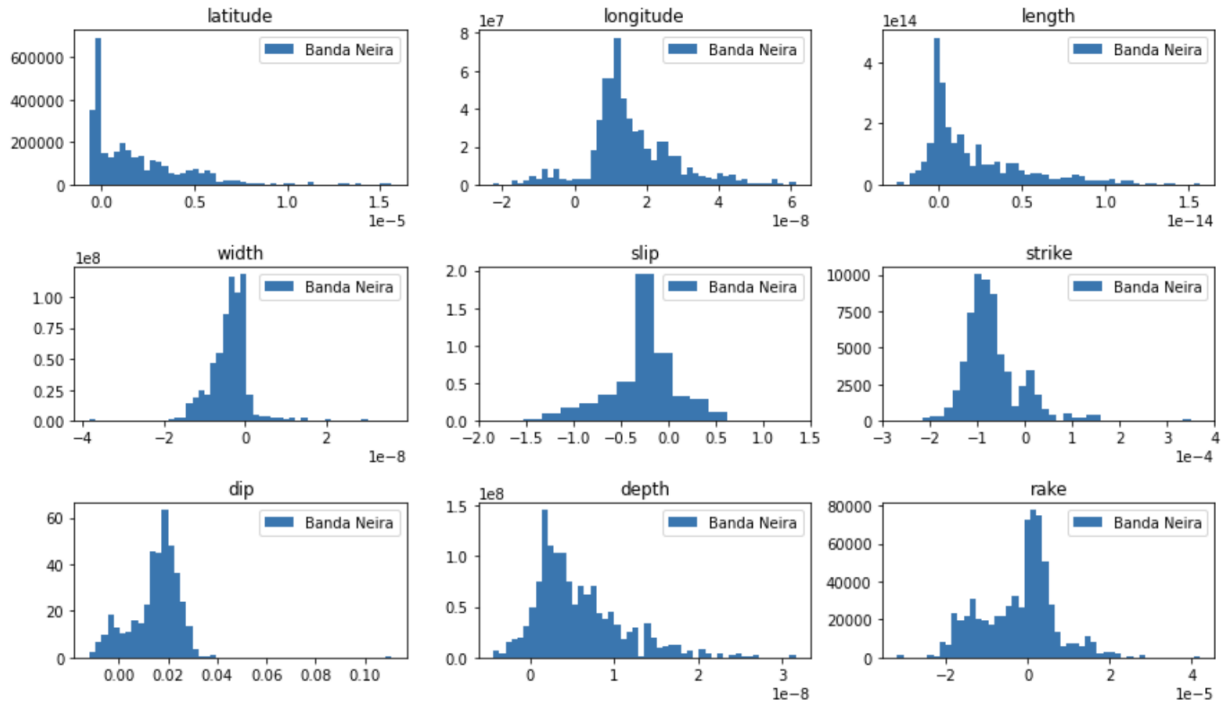

Figure 4.1: The value of the relative gradient for each earthquake parameter taken from 900 samples of the posterior distribution from [2]. The horizontal axis is the value of the gradient while the height of the histogram is the frequency of that gradient value. Note in particular the different magnitudes of the horizontal axis for each subplot.

inputs. To compute these sensitivities we analyzed the surrogate gradient for a subset of the posterior distribution from [2]. The estimated relative gradient values (the calculated gradient divided by the actual value of the parameter) for each of the Okada model inputs can be seen in Fig. 4.1 where the gradient is calculated on the maximal wave height at Banda Neira.

There are several key observations that can be made from this sensitivity analysis which are geologically intuitive:

(i) The relative gradient of slip is clearly much larger than it is for any other earthquake parameter. This is anticipated because larger slip values will lead to more pronounced initial waves that propagate to the observational location.

(ii) The latitude and longitude of the earthquake centroid appear to be less significant in determining the wave height at Banda Neira (magnitudes of the relative gradient of

latitude are mostly concentrated around 10^{-8} except for a few outliers that make it appear to be 10^{-5}). This is because tsunamis dissipate quite slowly so that a wave traveling a marginally longer distance will be nearly the same height when it reaches the observational location. On the other hand, we anticipate that the arrival time of the wave will critically depend on the earthquake location, and expect that this indicates why the surrogate developed for the arrival time did not work well for HMC.

(iii) The relative gradient for the length of the earthquake is nearly the same level as machine epsilon. We expect this is unique to this particular observational location, and earthquake as the rupture of the 1852 earthquake runs primarily north-south, and Banda Neira lies nearly directly west. According to the Okada model, variations in the length of the earthquake would extend that rupture further north or south, but would likely have very little effect on the wave that is observed at Banda Neira.

(iv) The width of the earthquake has a more pronounced effect on the wave height at Banda Neira (relatively the same effect as the centroid location), which as explained above, fits because the width of the rupture zone will push the edge of the initial disturbance either closer to Banda Neira, or farther away.

(v) The effect of the strike was surprisingly significant relative to many of the other Okada parameters. The strike dictates the angle from due north that the fault line is aligned. The significance of this parameter is likely because a strike tilted more to the west will place the earthquake and initial disturbance closer to Banda Neira which would influence the wave height there significantly.

(vi) The most surprising (in our opinion) observation from Fig. 4.1 is the effect of the dip on the wave height at Banda Neira. The outlier at 0.10 in this plot do skew our perception slightly, but nevertheless the dip is clearly the most significant parameter after the slip. The dip of the earthquake is the angle at which the fault dips horizontally relative to being completely flat. Significant changes in the dip will certainly modify the amount

of seafloor deformation (and hence the initial wave) produced by the earthquake, but we are honestly surprised that the effect appears as significant as it does.

(vii) The depth of the earthquake is the depth of the centroid below the surface of the earth. This is important because althought the slip predominantly dictates the overall seafloor deformation, some of that motion can be dissipated over a sufficiently large depth, i.e. the deeper earthquakes will not result in large surface displacements even for extremely large slip values. Changes in depth do effect the initial wave, but the changes are very slight so that the magnitude of the relative gradient depicted here is not surprising.

(viii) The rake defines the angle relative to the strike along with the earthquake ruptures. This means that a rake of 90° means the slip motion is completely perpendicular to the fault, leading to the greatest amount of motion transferred from the earthquake to the seafloor displacement (and hence tsunami wave). For this reason we anticipated that the rake would have a significant impact on the wave heights of the tsunami, although not as significant as the slip which is verified quantitatively in Fig. 4.1.

In summary, most of the observations taken from Fig. 4.1 are not surprising, and instead reconfirm geologic intuition (the slip is the most dominant parameter etc.), however the predicted influence of the dip parameter is surprising, and we do not yet have a good answer for why this is the case. What is more, this is the first time (to our knowledge) that such geologic intuition has been put on a quantitative footing, i.e. these results although approximate through the use of the linearized adjoint give the first quantitative analysis of the sensitivity of tsunami wave heights with respect to the generating Okada earthquake parameters. Further investigation of these effects on other events with different geometries and different types of earthquakes will lead to more quantitative methods to determine the sensitivity of Geoclaw (and to actual tsunamis) to the geometry of the causal earthquake. Unfortunately such investigations are computationally expensive (each sample in Fig. 4.1

required several minutes of compute time), although we will consider alternative methods to speed up that process.

Chapter 5. Conclusions and Future Work

The understanding of Hamiltonian Monte Carlo and the adjoint method combined to give new insights into the Bayesian estimation of past earthquakes. Our investigation with using different values of ϵ and different numbers of time steps indicated that the real problem was not with tuning of hyperparameters, but in the surrogate for arrival times. There are several avenues to improve the adjoint-surrogate HMC method that will be explored in future work, but the primary avenue for improvement is in the surrogate model used for the arrival times. As outlined earlier, there are several steps that will be taken to improve this part of the surrogate gradient, and there is still some potential improvements that can be made to speed the computation up. Overall however, it appears that the linearized adjoint approach toward a surrogate gradient is not very viable for implementing HMC for this problem.

On the other hand, the sensitivity analysis for the wave height at a specific location with respect to the Okada parameters yielded quantitative information about the effect of the geometry, location, and magnitude of the causal earthquake on the tsunami wave's height. This not only confirmed some geological intuition for the generation of tsunamis, but also led to some potential interesting questions. We will continue to utilize the linearized adjoint in this sense to explore the effects of the earthquake on the tsunami wave's propagation. This is particularly exciting as no previous efforts (to our knowledge) have arrived at quantitative estimates for these same parameters.

Bibliography

[1] Heekun Jung Hyung-Jo. Jin, Seung-Seop Ju. Adaptive markov chain monte carlo algorithms for bayesian inference: recent advances and comparative study. structure and infrastructure engineering. 2019.

[2] Hayden Ringer, Jared Whitehead, Justin Krometis, Harris Ronald A, Nathan Glatt-Holtz, Spencer Giddens, Claire Ashcraft, Garret Carver, Adam Robertson, McKay Harward, Joshua Fullwood, Kameron Lightheart, Ryan Hilton, Ashley Avery, Cody Kesler, Martha Morrise, and Michael Hunter Klein. Methodological reconstruction of historical seismic events from anecdotal accounts of destructive tsunamis: a case study for the great 1852 banda arc mega-thrust earthquake and tsunami. 2021.

[3] Natural Disaster Risk Index. `https://www.statista.com/statistics/920857/indonesia-risk-index-for-natural-disasters`. Accessed: 2023-04-24.

[4] Kerry Sieh, Danny H Natawidjaja, Aron J Meltzner, Chuan-Chou Shen, Hai Cheng, Kuei-Shu Li, Bambang W Suwargadi, John Galetzka, Belle Philibosian, and R Lawrence Edwards. Earthquake supercycles inferred from sea-level changes recorded in the corals of west Sumatra. *Science*, 322(5908):1674–1678, 2008.

[5] Z. Y. C. Liu and R. A. Harris. Discovery of possible mega-thrust earthquake along the Seram Trough from records of 1629 tsunami in eastern Indonesian region. *Natural Hazards*, 72:1311–1328, 2014.

[6] Ron Harris and Jonathan Major. Waves of destruction in the East Indies: the Wichmann catalogue of earthquakes and tsunami in the Indonesian region from 1538 to 1877. *Geological Society, London, Special Publications*, 441:SP441–2, 2016.

[7] TszMan L Fisher and Ron A Harris. Reconstruction of 1852 Banda Arc megathrust earthquake and tsunami. *Natural Hazards*, 83(1):667–689, 2016.

[8] David K Chester. The 1755 lisbon earthquake. *Progress in Physical Geography*, 25(3):363–383, 2001.

[9] Radford M Neal et al. MCMC using hamiltonian dynamics. *Handbook of markov chain monte carlo*, 2(11):2, 2011.

[10] Marsha J Berger, David L George, Randall J LeVeque, and Kyle T Mandli. The geoclaw software for depth-averaged flows with adaptive refinement. *Advances in Water Resources*, 34(9):1195–1206, 2011.

[11] A tale of two faults: Statistical reconstruction of the 1820 flores sea earthquake using tsunami observations alone.

[12] Jari Kaipio and Erkki Somersalo. *Statistical and computational inverse problems*, volume 160 of *Applied Mathematical Sciences*. Springer Science & Business Media, 2005.

[13] Nicholas Metropolis, Arianna W Rosenbluth, Marshall N Rosenbluth, Augusta H Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953.

[14] W Keith Hastings. Monte carlo sampling methods using markov chains and their applications. 1970.

[15] B Davis and R LeVeque. Adjoint methods for guiding adaptive mesh refinement in tsunami modeling. *Pure & Applied Geophysics*, 173(12), 2016.

[16] Yoshimitsu Okada. Surface deformation due to shear and tensile faults in a half-space. *Bulletin of the Seismological Society of America*, 75(4):1135–1154, 1985.

[17] Yoshimitsu Okada. Internal deformation due to shear and tensile faults in a half-space. *Bulletin of the Seismological Society of America*, 82(2):1018–1040, 1992.

[18] J. Callahan. Hamiltonian Monte Carlo for Reconstructing Historical Earthquake-induced Tsunamis. *Masters Thesis*, 2023.

[19] Gavin P. Hayes, Ginevra L. Moore, Daniel E. Portner, Mike Hearne, Hanna Flamme, Maria Furtney, and Gregory M. Smoczyk. Slab2, a comprehensive subduction zone geometry model. *Science*, 362(6410):58–61, 2018. Slab2 GitHub: `https://github.com/usgs/slab2`, data download: `https://www.sciencebase.gov/catalog/item/5aa1b00ee4b0b1c392e86467`.

[20] A Wichmann. The earthquakes of the Indian archipelago to 1857. *Verhandl. Koninkl. Akad. van Wetenschappen, 2nd sec*, 20(4):1–193, 1918.

[21] A Wichmann. The earthquakes of the Indian archipelago from 1858 to 1877. *Verhandl. Koninkl. Akad. van Wetenschappen, 2nd sec*, 22(5):1–209, 1922.

[22] T Paskett, JP Whitehead, RA Harris, C Ashcroft, JA Krometis, I Sorensen, and R Wonnacott. A tale of two faults: Statistical reconstruction of the 1820 flores sea earthquake using tsunami observations alone. *arXiv preprint arXiv:2305.01718*, 2023.

[23] Motohide Nishio and Aisaku Arakawa. Performance of Hamiltonian Monte Carlo and No-U-Turn Sampler for estimating genetic parameters and breeding values. 2019.