
Gaussian Processes for Probabilistic Estimates of Earthquake Ground Shaking: A 1-D Proof-of-Concept

Sam A. Scivier

Department of Earth Sciences
University of Oxford
Oxford, UK
sam.scivier@earth.ox.ac.uk

Tarje Nissen-Meyer

Department of Mathematics and Statistics
University of Exeter
Exeter, UK
t.nissen-meyer@exeter.ox.ac.uk

Paula Koelemeijer

Department of Earth Sciences
University of Oxford
Oxford, UK
paula.koelemeijer@earth.ox.ac.uk

Atılım Güneş Baydin

Department of Computer Science
University of Oxford
Oxford, UK
gunes.baydin@cs.ox.ac.uk

Abstract

Estimates of seismic wave speeds in the Earth (seismic velocity models) are key input parameters to earthquake simulations for ground motion prediction. Owing to the non-uniqueness of the seismic inverse problem, typically many velocity models exist for any given region. The arbitrary choice of which velocity model to use in earthquake simulations impacts ground motion predictions. However, current hazard analysis methods do not account for this source of uncertainty. We present a proof-of-concept ground motion prediction workflow for incorporating uncertainties arising from inconsistencies between existing seismic velocity models. Our analysis is based on the probabilistic fusion of overlapping seismic velocity models using scalable Gaussian process (GP) regression. Specifically, we fit a GP to two synthetic 1-D velocity profiles simultaneously, and show that the predictive uncertainty accounts for the differences between the models. We subsequently draw velocity model samples from the predictive distribution and estimate peak ground displacement using acoustic wave propagation through the velocity models. The resulting distribution of possible ground motion amplitudes is much wider than would be predicted by simulating shaking using only the two input velocity models. This proof-of-concept illustrates the importance of probabilistic methods for physics-based seismic hazard analysis.

1 Introduction

Seismic velocity models — estimates of the Earth’s seismic wave speeds — underpin earthquake ground motion prediction in seismic hazard analysis, as they are key inputs to wave equation solvers. They continue to be produced at different resolutions and scales, stemming from different methods (e.g., tomography [1], reflection surveys [2]). The seismic inverse problem is ill-posed as there are not enough data to constrain a unique true Earth model [3]. As such, many overlapping velocity models exist for a given region. Consequently, the choice of which velocity model to use in ground motion prediction is often arbitrary. Nevertheless, it has a significant impact on the results as different models have different structures, length scales, and amplitudes.

The key output of seismic hazard analysis is an estimate of peak ground motion, to assess potential infrastructure damage and inform earthquake engineering. Most commonly, empirical ground motion

models (GMMs) [4, 5] are used to predict the median and uncertainty of a ground motion parameter (e.g., peak ground displacement - PGD) for earthquake scenarios [6]. They make rapid predictions, but drastically simplify the underlying physical processes. Importantly, they approximate the effect of seismic velocities on ground motion, typically only using the average shear wave velocity in the uppermost 30 m [7]. GMMs are thus limited in accuracy and reliability. An alternative is to simulate earthquake scenarios in 3-D by solving the wave equation and extracting PGD estimates, requiring 3-D seismic velocity information as input. However, there are two issues: (i) simulating many earthquakes is computationally costly, and (ii) choices of input parameters are subjective, including the input velocity model. To address the first issue, recent advances in machine learning have begun accelerating wave propagation methods [8–11]. However, current physics-based hazard analysis workflows do not consider inconsistencies between velocity models. This omits a key source of uncertainty, given that predicted ground motion can be drastically impacted by velocity structure. One possible solution is to fuse different velocity models, and use the output in earthquake simulations. Unfortunately, existing methods for velocity model fusion [e.g., 12–14] typically do not produce probabilistic outputs, limiting their ability to account for differences between velocity models.

In this study, we propose a workflow to account for inconsistencies between seismic velocity models in ground motion prediction. Our method is based on the probabilistic fusion of velocity models using Gaussian processes (GPs), and estimates uncertainties owing to differences between them. We then produce probabilistic ground motion predictions with respect to these uncertainties by drawing velocity model samples from the GP predictive distribution, simulating acoustic wave propagation using each sample, and extracting the PGD predictions. We illustrate that such a probabilistic method is necessary to capture the spread of possible ground motion scenarios.

Our key contributions are as follows: (i) We present a workflow for probabilistic earthquake ground motion prediction that accounts for inconsistencies between seismic velocity models. (ii) We demonstrate the capability of scalable GPs for the probabilistic fusion of different estimates of the same physical parameter, through a synthetic example using 1-D seismic velocity models. The code for this work is written in Python and is available at Scivier et al. [15].

2 Gaussian processes and data fusion

GPs [16] are a class of non-parametric models for defining a distribution over function spaces. They are widely used for regression, providing robust uncertainty quantification and predictive performance. Unfortunately, exact GP regression is limited in scalability owing to a computational cost of $\mathcal{O}(n^3)$, where n is the number of data points. Despite the small datasets used in this study, seismic datasets of realistic size can have $n \sim 10^6-10^7$. To overcome this, approximate GP methods have been developed [17]. One popular method for scalable GP inference is the sparse variational Gaussian

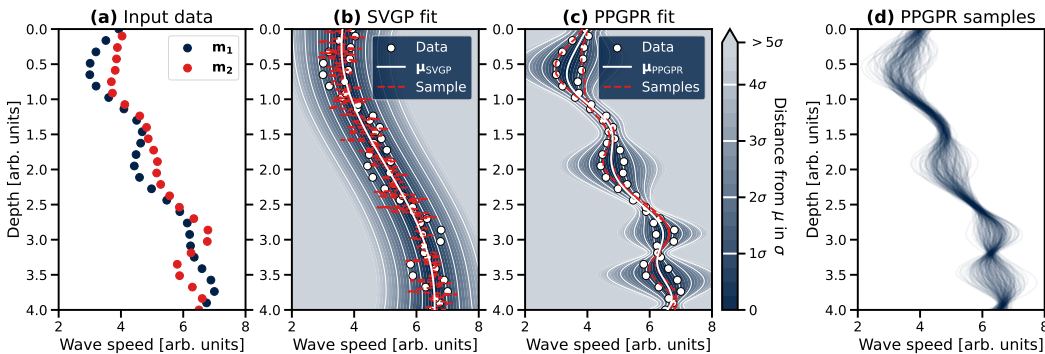


Figure 1: Comparison of SVGPR and PPGPR for the probabilistic fusion of seismic velocity models. (a) shows the input synthetic 1-D seismic velocity profiles with depth. (b) and (c) show the fusion results of SVGPR and PPGPR, and (d) shows the 200 function samples drawn from the PPGPR predictive distribution used in Section 4. The shading in (b) and (c) show the SVGPR posterior predictive distribution, $q_{\text{SVGPR}}(\mathbf{y}_*)$, and the PPGPR latent predictive distribution, $q_{\text{PPGPR}}(\mathbf{f}_*)$, respectively, in terms of distance from the predictive means in standard deviations.

process (SVGP) [18, 19], which applies variational inference to fit GPs. SVGPs introduce a set of inducing variables to approximate the full dataset using a smaller set of points $m \ll n$. Thus, the computational cost is reduced to $\mathcal{O}(nm^2 + m^3)$, making SVGPs practical to apply to large 2-D and 3-D datasets.

In this study, we use scalable GP regression for fusing seismic velocity models by fitting a GP to multiple datasets simultaneously. A key advantage of GPs is their modelling of covariance structure, which enables samples matching the spatial patterns of the input data to be drawn from the predictive distribution.

We aim to model inconsistencies between velocity estimates as uncertainty in the GP predictive distribution. Thus, we need to understand the form of the predictive variance in the SVGP model. Below we provide a brief summary of the key SVGP results to highlight the relevant context for our work. The interested reader is referred to Titsias [18], Matthews et al. [19], Jankowiak et al. [20], and Murphy [21] for full derivations and explanations. The inputs are the training set (coordinates of the input velocity models), \mathbf{X} , the inducing point locations, \mathbf{Z} , and the points at which we wish to predict, \mathbf{X}_* . Then $\mathbf{f}_\mathbf{X}$, $\mathbf{f}_\mathbf{Z}$, \mathbf{f}_* are the (unknown) velocity values that we predict at these locations; and \mathbf{y} are the observed data (input velocity values). SVGP-based methods approximate the joint posterior as $q(\mathbf{f}_*, \mathbf{f}_\mathbf{X}, \mathbf{f}_\mathbf{Z}) = p(\mathbf{f}_*, \mathbf{f}_\mathbf{X} | \mathbf{f}_\mathbf{Z}) q(\mathbf{f}_\mathbf{Z})$, where $p(\mathbf{f}_*, \mathbf{f}_\mathbf{X} | \mathbf{f}_\mathbf{Z})$ is calculated exactly [21]. The variational distribution is $q(\mathbf{f}_\mathbf{Z}) = \mathcal{N}(\mathbf{f}_\mathbf{Z} | \mathbf{m}, \mathbf{S})$, where \mathbf{m} and \mathbf{S} are (learned) variational parameters. The predictive distribution over the underlying function (at the target points) \mathbf{f}_* is given by

$$\begin{aligned} q(\mathbf{f}_*) &= \int p(\mathbf{f}_* | \mathbf{f}_\mathbf{Z}) q(\mathbf{f}_\mathbf{Z}) d\mathbf{f}_\mathbf{Z} \\ &= \mathcal{N}(\mathbf{f}_* | \mu_*, \sigma_f(\mathbf{x}_*)^2), \end{aligned} \quad (1)$$

where,

$$\begin{aligned} \mu_* &= \mathbf{K}_{*,\mathbf{Z}} \mathbf{K}_{\mathbf{Z},\mathbf{Z}}^{-1} \mathbf{m} \\ \sigma_f(\mathbf{x}_*)^2 &= \mathbf{K}_{*,*} - \mathbf{K}_{*,\mathbf{Z}} \mathbf{K}_{\mathbf{Z},\mathbf{Z}}^{-1} (\mathbf{K}_{\mathbf{Z},\mathbf{Z}} - \mathbf{S}) \mathbf{K}_{\mathbf{Z},\mathbf{Z}}^{-1} \mathbf{K}_{\mathbf{Z},*}, \end{aligned}$$

with e.g., $\mathbf{K}_{*,\mathbf{Z}} = k(\mathbf{X}_*, \mathbf{Z})$, and $k(\cdot, \cdot)$ is the (chosen) covariance function. Assuming a Gaussian likelihood, measurements (at the target points) \mathbf{y}_* are related to the underlying function (at the target points) \mathbf{f}_* as $p(\mathbf{y}_* | \mathbf{f}_*, \sigma_y^2) = \mathcal{N}(\mathbf{y}_* | \mathbf{f}_*, \sigma_y^2 \mathbb{I}_*)$, where σ_y^2 is the (learned) observational noise variance. Thus, the predictive distribution over \mathbf{y}_* is,

$$\begin{aligned} q(\mathbf{y}_*) &= \int p(\mathbf{y}_* | \mathbf{f}_*, \sigma_y^2) q(\mathbf{f}_*) d\mathbf{f}_* \\ &= \mathcal{N}(\mathbf{y}_* | \mu_*, \sigma_f(\mathbf{x}_*)^2 + \sigma_y^2 \mathbb{I}_*). \end{aligned} \quad (2)$$

The predictive variance at a target point $\mathbf{x}_{*,i}$ is thus the sum of input-dependent variance over the underlying function, $\sigma_f(\mathbf{x}_{*,i})^2$, and observational noise, σ_y^2 : $\text{Var}(\mathbf{x}_{*,i}) = \sigma_f(\mathbf{x}_{*,i})^2 + \sigma_y^2$. Despite this symmetry in the predictive variance, Jankowiak et al. [20] highlight that the typical SVGP objective function (variational ELBO) targets only large σ_y^2 – often resulting in $\sigma_y^2 \gg \sigma_f(\mathbf{x}_{*,i})^2$, which we see from the data-fit term: $\mathcal{L}_{\text{SVGP}} \supset -\frac{1}{2\sigma_y^2} |\mathbf{y}_i - \mu_{\mathbf{X},i}|^2$. This means we would model disagreements between different velocity models as observational noise. Given the degree of disagreement varies spatially, σ_y would need to be input-dependent [e.g., 22, 23]. However, this would result in noisy samples in the predictive distribution (see Fig. 1b) - making them useless for downstream tasks.

We instead wish to model inconsistencies between velocity models as input-dependent uncertainty in the underlying physical process (i.e., $\sigma_f(\mathbf{x})$). To enable this, we use the parametric predictive GP regression (PPGPR) model [20]. PPGPR is a variation of SVGP with an objective function that directly targets the predictive distribution (Eq. (2)). Notably, the PPGPR objective encourages large $\sigma_f(\mathbf{x}_*)^2$, as seen from the data-fit term: $\mathcal{L}_{\text{PPGPR}} \supset -\frac{1}{2} \frac{1}{\sigma_y^2 + \sigma_f(\mathbf{x}_i)^2} |\mathbf{y}_i - \mu_{\mathbf{X},i}|^2$. Contrary to SVGP, this typically results in $\sigma_f(\mathbf{x}_{*,i})^2 \gg \sigma_y^2$. Therefore, we can choose to ignore σ_y^2 and use the PPGPR latent predictive distribution, $q(\mathbf{f}_*)$.

3 Fusing synthetic seismic velocity models

We present a proof-of-concept demonstrating the applicability of PPGPR for the probabilistic fusion of two synthetic 1-D seismic velocity models. We note that we do not consider uncertainties attached to input velocity models in this study (i.e., the input models themselves are not probability distributions). Two datasets, \mathbf{s}_1 and \mathbf{s}_2 , are sampled ($n = 25$ data points, each) from a GP prior, using radial basis function (RBF) kernels with different length scales. The samples have different coordinates in an overlapping region. The first model is set as the first sample, $\mathbf{m}_1 = \mathbf{s}_1$. The second model is a weighted superposition of the two samples, $\mathbf{m}_2 = \frac{2}{3}\mathbf{s}_1 + \frac{1}{3}\mathbf{s}_2$, to create larger-scale similarities and smaller-scale differences - typical of different seismic velocity models. Fig. 1a shows the input velocity models, which are 1-D profiles with respect to depth.

At training time, the input models are concatenated and used to condition the GP as a single dataset. For regression, we employ both PPGPR and SVGPR for comparison, using scaled RBF kernels, $m = 20$ inducing points (with learned locations), and Gaussian likelihoods. The models are trained using the Adam optimiser [24] and identical hyperparameters (i.e., learning rate and number of iterations). Hyperparameters are chosen through trial-and-error, and full hyperparameter details are provided in the code [15]. Training the models takes one minute for each method, on a laptop using an NVIDIA T500 2GB GDDR6 GPU.

Fig. 1b and c show the results of SVGPR and PPGPR, respectively, on the two velocity models. As discussed in Section 2, optimising the SVGP objective results in $\sigma_{\text{obs}}^2 \gg \sigma_{\mathbf{f}(\mathbf{x})}^2$, making it unsuitable for this task owing to a lack of input-dependence on the predictive variance and noisy function samples. On the other hand, PPGPR performs well, with predictive samples appearing to reflect the spatial patterns of the input models. In an ideal case of maximum likelihood estimation for fitting a univariate Gaussian distribution to two observations, y_1 and y_2 , the resulting distribution is $\mathcal{N}\left(\mu = \frac{1}{2}(y_1 + y_2), \sigma^2 = \left(\frac{y_1 - y_2}{2}\right)^2\right)$ [25]. In our example, we therefore expect the $\pm 1\sigma$ contours to approximately follow each of the input velocity models. We interpolated \mathbf{m}_1 and \mathbf{m}_2 at the test points using cubic splines and calculated the root mean square error (RMSE) of the SVGP and PPGPR predictive means and variances with respect to the above ideal result. The RMSEs on μ_{SVGP} and σ_{SVGP}^2 were 0.243 and 0.098 (in wave speed units), respectively, while for μ_{PPGPR} and σ_{PPGPR}^2 the RMSEs were 0.045 and 0.012 (in wave speed units). The PPGPR predictive distribution thus appropriately quantifies the uncertainty on the knowledge of seismic velocities in the region. Most importantly for our application, the covariance structure of the data is modelled. This enables the drawing of samples from the predictive distribution that match the spatial patterns of the input data. Despite only fusing two velocity models here, our approach is generally applicable for fusing any number of input datasets.

4 Probabilistic ground motion prediction

We propose a proof-of-concept workflow for propagating the predicted uncertainty on seismic velocities through simulations of the acoustic wave equation, to produce probabilistic ground motion predictions. First, we draw 200 function samples from the PPGPR latent predictive distribution (Fig. 1d). Then for each sample, we simulate the 1-D acoustic wave equation for displacement, \mathbf{u} , with a Ricker wavelet as the earthquake source, using a finite difference scheme (6 s for 200 simulations on a laptop - vectorised over velocity models). At the surface, we implement a free-surface boundary condition (i.e., the acoustic pressure $\mathbf{p} = 0$). At depth, we implement an absorbing boundary layer according to Chern [26]. Fig. 2a-f shows snapshots of the displacement field at various time steps in one of the simulations. For each simulation, we record the peak ground displacement at the surface (i.e., PGD; depth = 0), producing one PGD estimate per simulation (i.e., per sample velocity model). Fig. 2g shows a histogram of the recorded PGD measurements from the simulations. Additionally, we ran simulations using interpolated versions of \mathbf{m}_1 and \mathbf{m}_2 and marked the resulting PGD measurements in Fig. 2g, to investigate how much information is gained by running simulations for many velocity model samples. Clearly, the PGD measurements for \mathbf{m}_1 and \mathbf{m}_2 do not account for the spread of possible ground motions, given the degree of knowledge of seismic velocities in this example. Despite being 1-D, our work already shows that it is not possible to approximate the full distribution of possible ground motions using only two velocity models.

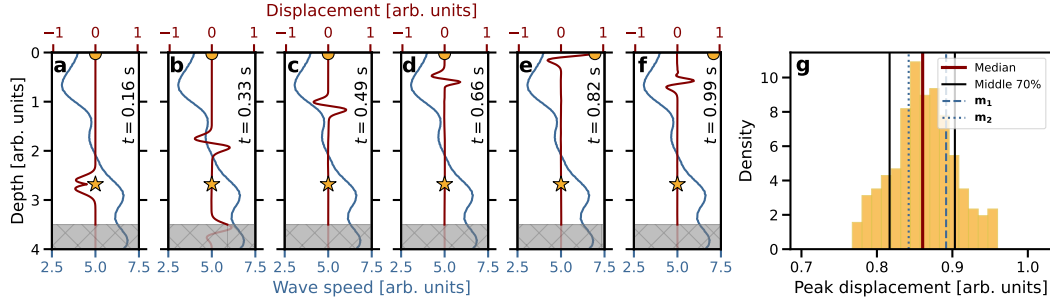


Figure 2: Wavefield snapshots and probabilistic ground motion prediction. (a)–(f) shows wavefield displacement snapshots at increasing time steps for one simulation. Each panel includes the source location (yellow star), the maximum PGD up to that time step (yellow dot), and the underlying velocity model of the simulation. The shaded region indicates where an absorbing boundary layer is applied [26]. (g) shows a histogram of the PGD measurements from the simulations, and highlights the median and middle 70% of predictions. Also shown are the PGD measurements resulting from simulations using just m_1 or m_2 as input.

5 Limitations

This work is a proof-of-concept and can be extended in several ways. For example, we do not account for data with varying length scales or structure, or address kernel design or choice, which would be required for dealing with real seismic datasets. For real-world applicability, it will also be important to extend our workflow from 1-D to 2-D and 3-D, and to solve the elastic wave equation instead of the acoustic wave equation. Working with 3-D velocity models would add complexity, but in principle it would consist of changing the GP coordinate space from 1-D to 3-D. There are many optimised 3-D seismic wave propagation codes that could then be used for the simulation component of the workflow [e.g., 11, 27–29]. If the input velocity models had different spatial densities of data points, the objective function would be weighted towards one of them, and the result would be skewed. This can be readily solved by weighting their contributions to the objective. Additionally in this work, the synthetic input velocity models do not have uncertainties, which we plan to incorporate in the future (i.e., the input velocity models would themselves be probability distributions).

In this study, we were unable to compare our method with existing methods for velocity model fusion. Current methods are generally designed for enhancing larger-scale velocity models using smaller-scale models - and are thus not applicable when the models occupy the same domain and/or have similar spatial data density, as in our case. Additionally, existing methods typically do not produce probabilistic outputs, meaning it is not possible to compare them in the ground motion prediction component of the study (Section 4).

6 Conclusion

Seismic velocity models underpin predictions of earthquake ground motion. Current methods for physics-based seismic hazard analysis do not account for inconsistencies between existing velocity models. In this study, we present a proof-of-concept workflow for probabilistic ground motion prediction that takes this source of uncertainty into account. Firstly, we demonstrate the applicability of scalable GP regression to the probabilistic fusion of input velocity models, showing that inconsistencies between velocity models can be modelled as predictive uncertainty. This provides access to any number of plausible velocity models for the region by drawing samples from the predictive distribution. Secondly, we build up a distribution of possible ground motion scenarios for the family of possible velocity models according to the GP predictive distribution. Our results show a much wider spread of possible peak ground motions than would be predicted by simulating earthquake scenarios using just the input velocity models themselves. We thus highlight the value of using probabilistic methods, such as the one presented here, in physics-based seismic hazard analysis to account for differences between velocity models.

Acknowledgments and Disclosure of Funding

Acknowledgments The authors thank the five anonymous reviewers for their time and insightful comments, which have improved the quality of the paper. The authors thank Adrian Marin Mag for fruitful discussions and comments on the paper, as well as Fatima Ramadan and Franck Latalerie for useful feedback.

Reproducibility The code used to produce the results and figures in this study is available at Scivier et al. [15].

Software The code for this work was written in Python, and used the open-source software libraries JUPYTER NOTEBOOKS [30], BINDER [31], NUMPY v2.1.3 [32], SCIPY v1.14.1 [33], MATPLOTLIB v3.9.2 [34], PYTORCH v2.5.1 [35], and GPYTORCH v1.13 [36].

Funding SAS is funded by a UKRI NERC DTP Award (NE/S007474/1) and gratefully acknowledges their support. PK acknowledges financial support from a Royal Society University Research Fellowship (UR/R1180377). AGB was supported by grants from NVIDIA and the Munich Institute for Astro-, Particle and BioPhysics (MIAPbP), funded by the Deutsche Forschungsgemeinschaft under Germany’s Excellence Strategy – EXC-2094 – 390783311.

Competing interests The authors have no competing interests to declare.

References

- [1] Carl Tape, Qinya Liu, Alessia Maggi, and Jeroen Tromp. Seismic tomography of the southern California crust based on spectral-element and adjoint methods. *Geophysical Journal International*, 180(1):433–462, 01 2010. ISSN 0956-540X. doi: 10.1111/j.1365-246X.2009.04429.x. URL <https://doi.org/10.1111/j.1365-246X.2009.04429.x>.
- [2] P.S. Schultz. Seismic velocity estimation. *Proceedings of the IEEE*, 72(10):1330–1339, 1984. doi: 10.1109/PROC.1984.13021.
- [3] N. Rawlinson, S. Pozgay, and S. Fishwick. Seismic tomography: A window into deep Earth. *Physics of the Earth and Planetary Interiors*, 178(3):101–135, 2010. ISSN 0031-9201. doi: <https://doi.org/10.1016/j.pepi.2009.10.002>. URL <https://www.sciencedirect.com/science/article/pii/S0031920109002106>.
- [4] Carlo Meletti, Warner Marzocchi, Vera D’Amico, Giovanni Lanzano, Lucia Luzi, Francesco Martinelli, B. Pace, Andrea Rovida, Francesco Visini, MPS Group, Giovanni Barreca, Carmelo Monaco, and Junio Iervolino. The new Italian Seismic Hazard Model (MPS19). *Annals of Geophysics*, 64(1), 05 2021. doi: 10.4401/ag-8579.
- [5] Barbara Motnikar, Polona Zupančič, Mladen Živčić, J. Atanackov, Petra Jamšek Rupnik, Martina Čarman, Laurentiu Danciu, and Andrej Gosar. The 2021 seismic hazard model for Slovenia (SHMS21): overview and results. *Bulletin of Earthquake Engineering*, 20:1–30, 08 2022. doi: 10.1007/s10518-022-01399-8.
- [6] Norman A. Abrahamson, Nicolas M. Kuehn, Melanie Walling, and Niels Landwehr. Probabilistic Seismic Hazard Analysis in California Using Nonergodic Ground-Motion Models. *Bulletin of the Seismological Society of America*, 109(4):1235–1249, 07 2019. ISSN 0037-1106. doi: 10.1785/0120190030. URL <https://doi.org/10.1785/0120190030>.
- [7] Robert Graves, Thomas Jordan, Scott Callaghan, Ewa Deelman, Edward Field, Gideon Juve, Carl Kesselman, Philip Maechling, Gaurang Mehta, Kevin Milner, David Okaya, Patrick Small, and Karan Vahi. CyberShake: A Physics-Based Seismic Hazard Model for Southern California. *Pure and Applied Geophysics*, 168:367–381, 03 2010. doi: 10.1007/s00024-010-0161-6.
- [8] Yan Yang, Angela F. Gao, Jorge C. Castellanos, Zachary E. Ross, Kamyar Azizzadenesheli, and Robert W. Clayton. Seismic wave propagation and inversion with Neural Operators, 2021.
- [9] Ben Moseley, Andrew Markham, and Tarje Nissen-Meyer. Finite basis physics-informed neural networks (FBPINNs): a scalable domain decomposition approach for solving differential equations. *Advances in Computational Mathematics*, 49, 07 2023. doi: 10.1007/s10444-023-10065-9.

- [10] Fatme Ramadan, Bill Fry, and Tarje Nissen-Meyer. Rapid Computation of Physics-Based Ground Motions in the Spectral Domain using Neural Networks. In *EGU General Assembly Conference Abstracts*, EGU General Assembly Conference Abstracts, page 18444, April 2024. doi: 10.5194/egusphere-egu24-18444.
- [11] Fanny Lehmann, Filippo Gatti, Michaël Bertin, and Didier Clouteau. 3D elastic wave propagation with a Factorized Fourier Neural Operator (F-FNO). *Computer Methods in Applied Mechanics and Engineering*, 420:116718, 2024. ISSN 0045-7825. doi: <https://doi.org/10.1016/j.cma.2023.116718>. URL <https://www.sciencedirect.com/science/article/pii/S0045782523008411>.
- [12] Andreas Fichtner, Dirk-Philip van Herwaarden, Michael Afanasiev, Saulé Simutè, Lion Krischer, Yeşim Çubuk Sabuncu, Tuncay Taymaz, Lorenzo Colli, Erdinc Saygin, Antonio Villaseñor, Jeannot Trampert, Paul Cupillard, Hans-Peter Bunge, and Heiner Igel. The Collaborative Seismic Earth Model: Generation 1. *Geophysical Research Letters*, 45(9):4007–4016, 2018. doi: 10.1029/2018GL077338. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2018GL077338>.
- [13] R. Ajala and P. Persaud. Effect of Merging Multiscale Models on Seismic Wavefield Predictions Near the Southern San Andreas Fault. *Journal of Geophysical Research: Solid Earth*, 126(10), 2021. doi: 10.1029/2021JB021915. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2021JB021915>.
- [14] Hao Zhang and Yehuda Ben-Zion. Enhancing Regional Seismic Velocity Models With Higher-Resolution Local Results Using Sparse Dictionary Learning. *Journal of Geophysical Research: Solid Earth*, 129(1):e2023JB027016, 2024. doi: <https://doi.org/10.1029/2023JB027016>. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2023JB027016>. e2023JB027016 2023JB027016.
- [15] Sam A. Scivier, Tarje Nissen-Meyer, Paula Koelemeijer, and Atılım Güneş Baydin. Gaussian Processes for Probabilistic Estimates of Earthquake Ground Shaking: A 1-D Proof-of-Concept, November 2024. URL <https://doi.org/10.5281/zenodo.14246055>.
- [16] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 11 2005. ISBN 9780262256834. doi: 10.7551/mitpress/3206.001.0001. URL <https://doi.org/10.7551/mitpress/3206.001.0001>.
- [17] Haitao Liu, Yew-Soon Ong, Xiaobo Shen, and Jianfei Cai. When Gaussian Process Meets Big Data: A Review of Scalable GPs. *IEEE Transactions on Neural Networks and Learning Systems*, 31(11):4405–4423, 2020. doi: 10.1109/TNNLS.2019.2957109.
- [18] Michalis Titsias. Variational Learning of Inducing Variables in Sparse Gaussian Processes. In David van Dyk and Max Welling, editors, *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics*, volume 5 of *Proceedings of Machine Learning Research*, pages 567–574, Hilton Clearwater Beach Resort, Clearwater Beach, Florida USA, 16–18 Apr 2009. PMLR. URL <https://proceedings.mlr.press/v5/titsias09a.html>.
- [19] Alexander G. de G. Matthews, James Hensman, Richard Turner, and Zoubin Ghahramani. On Sparse Variational Methods and the Kullback-Leibler Divergence between Stochastic Processes. In Arthur Gretton and Christian C. Robert, editors, *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *Proceedings of Machine Learning Research*, pages 231–239, Cadiz, Spain, 09–11 May 2016. PMLR. URL <https://proceedings.mlr.press/v51/matthews16.html>.
- [20] Martin Jankowiak, Geoff Pleiss, and Jacob Gardner. Parametric Gaussian process regressors. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 4702–4712. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/jankowiak20a.html>.
- [21] Kevin P. Murphy. *Probabilistic Machine Learning: Advanced Topics*. MIT Press, 2023. URL <http://probml.github.io/book2>.
- [22] Haitao Liu, Yew-Soon Ong, and Jianfei Cai. Large-scale Heteroscedastic Regression via Gaussian Process, 2020. URL <https://arxiv.org/abs/1811.01179>.
- [23] Kristian Kersting, Christian Plagemann, Patrick Pfaff, and Wolfram Burgard. Most likely heteroscedastic Gaussian process regression. In *Proceedings of the 24th International Conference on Machine Learning*, ICML '07, page 393–400, New York, NY, USA, 2007. Association for Computing Machinery. ISBN 9781595937933. doi: 10.1145/1273496.1273546. URL <https://doi.org/10.1145/1273496.1273546>.
- [24] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization, 2017. URL <https://arxiv.org/abs/1412.6980>.

- [25] Kevin P. Murphy. *Probabilistic Machine Learning: An introduction*. MIT Press, 2022. URL probml.ai.
- [26] Albert Chern. A reflectionless discrete perfectly matched layer. *Journal of Computational Physics*, 381: 91–109, March 2019. ISSN 0021-9991. doi: 10.1016/j.jcp.2018.12.026. URL <http://dx.doi.org/10.1016/j.jcp.2018.12.026>.
- [27] Daniel Peter, Dimitri Komatitsch, Yang Luo, Roland Martin, Nicolas Le Goff, Emanuele Casarotti, Pieyre Le Loher, Federica Magnoni, Qinya Liu, Celine Blitz, Tarje Nissen-Meyer, Piero Basini, and Jeroen Tromp. Forward and adjoint simulations of seismic wave propagation on fully unstructured hexahedral meshes. *Geophys. J. Int.*, 186(2):721–739, 2011. doi: 10.1111/j.1365-246X.2011.05044.x.
- [28] Takuto Maeda, Shunsuke Takemura, and Takashi Furumura. OpenSWPC: an open-source integrated parallel simulation code for modeling seismic wave propagation in 3D heterogeneous viscoelastic media. Technical report, Springer, 2017.
- [29] Kuangdai Leng, Tarje Nissen-Meyer, Martin van Driel, Kasra Hosseini, and David Al-Attar. AxiSEM3D: broad-band seismic wavefields in 3-D global earth models with undulating discontinuities. *Geophysical Journal International*, 217(3):2125–2146, 02 2019. ISSN 0956-540X. doi: 10.1093/gji/ggz092. URL <https://doi.org/10.1093/gji/ggz092>.
- [30] Thomas Kluyver, Benjamin Ragan-Kelley, Fernando Pérez, Brian Granger, Matthias Bussonnier, Jonathan Frederic, Kyle Kelley, Jessica Hamrick, Jason Grout, Sylvain Corlay, Paul Ivanov, Damián Avila, Safia Abdalla, Carol Willing, and Jupyter development team. Jupyter Notebooks - a publishing format for reproducible computational workflows. In Fernando Loizides and Birgit Schmidt, editors, *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, pages 87–90. IOS Press, 2016. URL <https://eprints.soton.ac.uk/403913/>.
- [31] Project Jupyter, Matthias Bussonnier, Jessica Forde, Jeremy Freeman, Brian Granger, Tim Head, Chris Holdgraf, Kyle Kelley, Gladys Nalvarte, Andrew Osheroff, M Pacer, Yuvi Panda, Fernando Perez, Benjamin Ragan Kelley, and Carol Willing. Binder 2.0 - Reproducible, interactive, sharable environments for science at scale. In Fatih Akici, David Lippa, Dillon Niederhut, and M Pacer, editors, *Proceedings of the 17th Python in Science Conference*, pages 113 – 120, 2018. doi: 10.25080/Majora-4af1f417-011.
- [32] Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming with NumPy. *Nature*, 585(7825): 357–362, September 2020. doi: 10.1038/s41586-020-2649-2. URL <https://doi.org/10.1038/s41586-020-2649-2>.
- [33] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2.
- [34] J. D. Hunter. Matplotlib: A 2D graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007. doi: 10.1109/MCSE.2007.55.
- [35] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library, 2019. URL <https://arxiv.org/abs/1912.01703>.
- [36] Jacob R Gardner, Geoff Pleiss, David Bindel, Kilian Q Weinberger, and Andrew Gordon Wilson. GPyTorch: Blackbox Matrix-Matrix Gaussian Process Inference with GPU Acceleration. In *Advances in Neural Information Processing Systems*, 2018.