Bayesian Seismic Monitoring

David A. Moore and Stuart J. Russell

Second Bayesian Young Statisticians Meeting (BAYSM 2014)
Vienna, September 18–19, 2014

UC Berkeley, Computer Science Division, Berkeley, USA
<dmoore,russell>@cs.berkeley.edu

Abstract

Nuclear monitoring is the problem of detecting nuclear weapons tests through analysis of seismic signals recorded by a global network of sensors. We present a Bayesian approach using a generative probabilistic model of natural and human-caused seismic events, the physics of wave propagation, and raw seismic waveforms. Our system combines insights from multiple techniques in current seismological practice (pick-based localization, array beamforming, waveform matching) into a unified model with principled handling of uncertainty.

Keywords: signals; open universe models; Gaussian processes; MCMC

1 Introduction

The *Comprehensive Test Ban Treaty (CTBT)* bans all testing of nuclear weapons by its signatories. To ensure that this ban is credible, the treaty provides for an *International Monitoring System (IMS)* to detect potential nuclear tests using data from a worldwide network of sensors including 170 seismic stations. The recorded signals are complex, unassociated, and often jumbled, so forming accurate inferences requires combining noisy statistical evidence from observations at multiple sensors. We present a system, called SIG-VISA (Signal-based Vertically Integrated Seismic Monitoring), that applies Bayesian reasoning to a very rich representation – raw observed seismic signals – to infer a set of seismic events, including potential nuclear tests. This extends previous work, NET-VISA [1], which conditions only on a set of discrete station 'detections'.

2 Generative Model

Our approach is conceptually straightforward: we define a prior distribution on the locations, magnitudes, and occurrence rates of nuclear explosions and natural earthquakes, along with a forward model that specifies the probability of observing a given set of signals conditioned on a hypothesized set of events. The prior on seismic event frequency is a Poisson process, with event locations sampled according to historical seismicity (for natural earthquakes) or a uniform prior (nuclear explosions), and magnitudes from an exponential (Gutenberg-Richter) distribution.

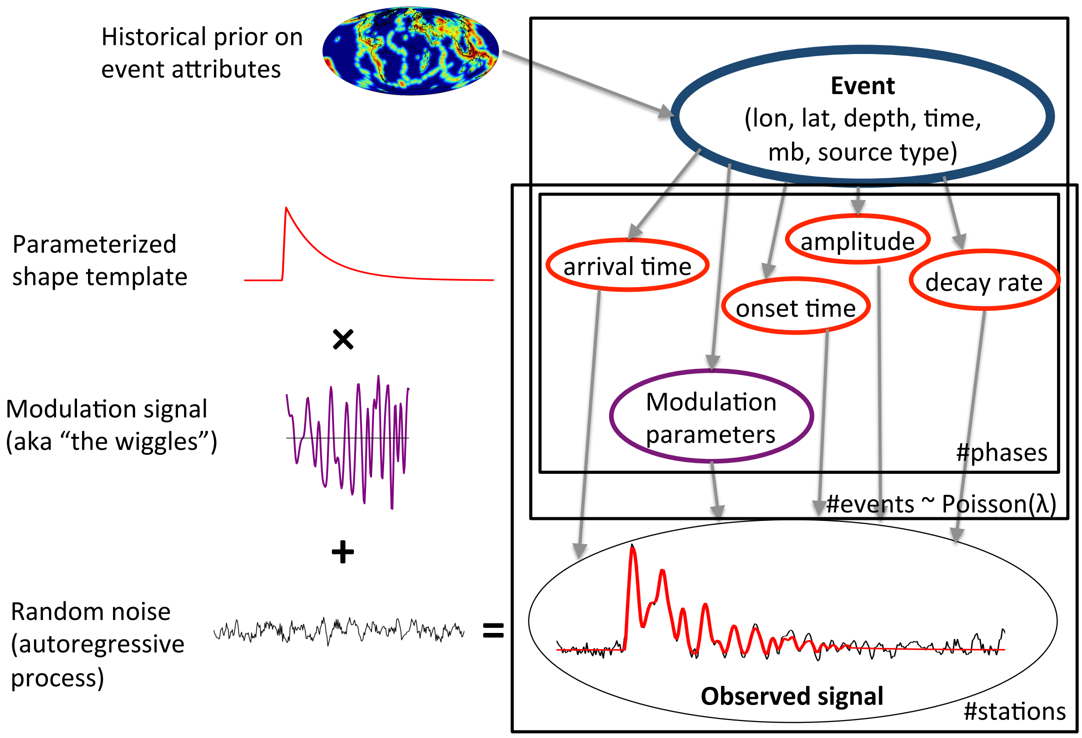

Figure 1: Generative model of seismic signals.

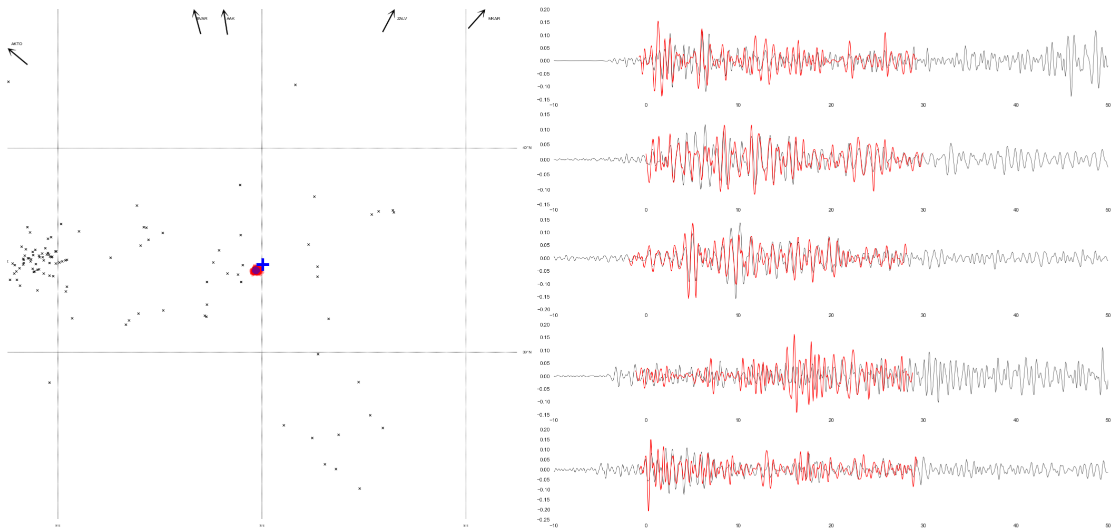

Figure 2: Posterior on seismic event location from five stations, from observed waveforms (right side, in black) showing model-predicted waveforms overlayed (red). The posterior mean (purple circle) is 4.3km from the catalog location (blue cross); training events are shown as black X's.

The forward model (Figure 1) begins by modeling, for each arriving phase, a *shape template* representing the envelope of the resulting signal. This is multiplied by a *modulation signal*, modeled with a wavelet parameterization, which captures the predictable 'wiggles' resulting from distortion of the original event impulse along the path to the station. Finally, the observed waveform at each station is taken to be the sum of all arriving phases, plus station background noise generated by an autoregressive process. All of the shape parameters, as well as the modulation signal, are moduled using Gaussian process regression conditioned on event location, trained from historical data.

This model unifies several existing approaches to seismic event detection and localization. The envelope shape component generalizes traditional 'pick'-based localization, in which point estimates of arrival times are used to solve for event locations. The joint model of signals at nearby stations recovers *array beamforming*, in which signals from multiple nearby stations are used to reduce noise and estimate the azimuth of arriving signals. The Gaussian process model of waveform modulation allows for a *waveform matching* effect: energy from nearby events travels along similar paths to each station, generating highly correlated waveforms; this can be exploited to localize an event from a single station, and to detect events that would otherwise fall below the noise threshold [2, 3]. Our model exploits these advantages while smoothly degrading to pick-based inferences in regions where no historical waveform data is available.

3 Inference

We use Metropolis-Hastings to sample from the posterior distribution on event hypotheses, using reversible-jump birth and death moves to search over hypotheses containing varying numbers of events; new phase arrivals are proposed with probability proportional to signal amplitude, and new event locations are proposed from a Hough transform on all currently-unassociated phase arrivals. We also use a number of custom moves that jointly propose new envelope parameters along with a new modulation signal, to maintain the fit to the observed waveform.

4 Results

We have preliminary experiments showing promising results for event localization from small networks (Figure 2). We are working to scale up inference and evaluate the model against existing approaches on a full global dataset.

References

[1] Nimar Arora, Stuart Russell, Paul Kidwell, and Erik B. Sudderth. Global seismic monitoring as probabilistic inference. In *Advances in Neural Information Processing Systems*, page 7381, 2010.

[2] Steven J. Gibbons and Frode Ringdal. The detection of low magnitude seismic events using array-based waveform correlation. *Geophysical Journal International*, 165(1):149166, 2006.

[3] M. Slinkard, Stephen Heck, D. Carr, Regina Eckert, and C. Young. Broad area event detection using waveform correlation and distributed computing. *Proceedings of 2012 Monitoring Research Review: Ground-Based Nuclear Explosion Monitoring Technologies*, page 806815, 2012.