

Combining Multiple Earthquake Models in Real Time for Earthquake Early Warning

by Sarah E. Minson, Stephen Wu, James L. Beck, and Thomas H. Heaton

Abstract The ultimate goal of earthquake early warning (EEW) is to provide local shaking information to users before the strong shaking from an earthquake reaches their location. This is accomplished by operating one or more real-time analyses that attempt to predict shaking intensity, often by estimating the earthquake's location and magnitude and then predicting the ground motion from that point source. Other EEW algorithms use finite rupture models or may directly estimate ground motion without first solving for an earthquake source. EEW performance could be improved if the information from these diverse and independent prediction models could be combined into one unified, ground-motion prediction. In this article, we set the forecast shaking at each location as the common ground to combine all these predictions and introduce a Bayesian approach to creating better ground-motion predictions. We also describe how this methodology could be used to build a new generation of EEW systems that provide optimal decisions customized for each user based on the user's individual false-alarm tolerance and the time necessary for that user to react.

Electronic Supplement: Animations of the waveform envelope fits and predicted shaking intensity for both the 2014 M_w 6.0 Napa earthquake and the 1 July 2015 false alarm, along with the details of all reports from all earthquake early warning (EEW) algorithms for both events, as well as for the 2014 M_w 6.8 off Cape Mendocino earthquake.

Introduction

The goal of earthquake early warning (EEW) is to provide users with an estimate of the ground motion that they will feel and the time that shaking will occur at their location. With this information, people and automated systems can take action to minimize the impact of the earthquake. Many countries and regions around the world, including Japan, Mexico, Taiwan, and the United States, are operating, or are in the process of building, EEW systems. Sometimes, these systems employ multiple real-time analyses for predicting local ground shaking; in these cases, a more robust warning system could be built if the information from these different analyses could be combined into one unified shaking forecast. In this article, we outline a strategy to accomplish precisely that.

Our examples focus on ShakeAlert, the EEW system for the west coast of the United States. ShakeAlert is being built by the U.S. Geological Survey in cooperation with the States of California, Oregon, and Washington, as well as the California Institute of Technology, University of California (Berkeley), University of Oregon, and University of Washington. ShakeAlert's design is to run multiple analyses simultaneously and then to synthesize information from these

independent algorithms to issue one unified forecast of local shaking (Böse *et al.*, 2014).

Currently, ShakeAlert receives warnings from three algorithms: (1) Onsite (e.g., Kanamori, 2005; Böse *et al.*, 2009), Earthquake Alarm Systems (ElarmS; e.g., Allen and Kanamori, 2003; Kuyuk *et al.*, 2013), and Virtual Seismologist (VS; e.g., Cua and Heaton, 2007; Cua *et al.*, 2009). Each of these algorithms independently uses real-time measurements of seismic data to detect an earthquake and determine its point-source description (magnitude, location, and origin time [OT]). Collectively, these three algorithms provide a spectrum of behavior, from faster response (at the potential cost of less accurate source parameters) to more accurate source information (at the potential cost of less timeliness). The location and magnitude information from these algorithms are used as input to ground-motion prediction equations (GMPEs) that are used to predict the ground motion that a user at a particular location may expect. However, if the earthquake is larger than what can be reasonably modeled as a point source, the expected shaking intensity (SI) at a certain station will also be dependent on the distance to that fault plane, which can be different from the hypocentral distance. Thus, these point-

source models may underpredict ground motions for large earthquakes with finite rupture extents. (Of course, other sources of ground-motion variability may cause under- or overprediction of expected ground motion as well.) Therefore, in addition to these three point-source algorithms, it is anticipated that future generations of the ShakeAlert system will include real-time finite-fault source models as well. For example, currently under development and testing are a finite seismic line source solution (FinDer; Böse *et al.*, 2012, 2015) and a number of finite-fault source models based on real-time, high-rate Global Positioning System data (e.g., Grapenthin *et al.*, 2014; Minson *et al.*, 2014; Crowell *et al.*, 2016).

In the current ShakeAlert prototype production system, the outputs from the three point-source algorithms (Onsite, ElarmS, and VS) are averaged together to obtain an average latitude and longitude for the earthquake's epicenter, average magnitude, and average OT; these averaged parameters are then input into a GMPE for predicting each user's expected ground motion. This framework is not optimal because it is not known how to map uncertainties in the separate predictions of the source parameters to uncertainties in the ground-motion predictions; more importantly, there is no way to objectively suppress false alarms. Further, this framework cannot treat non-point-source algorithms: it is nonphysical to attempt to average the source properties of three point sources, one line source, and a collection of various distributed slip models, each of which is built upon a different fault geometry. Additionally, we anticipate that in the future there will be a need to combine earthquake source models with algorithms that forecast future ground motion directly from current ground motion without an assumed earthquake source (e.g., Hoshiaba, 2013; Hsu *et al.*, 2013; Hoshiaba and Aoki, 2015). In this article, we propose a new approach that allows us to probabilistically combine information from multiple rupture models in real time (regardless of how each rupture model is parameterized) to provide a single unified and accurate ground-motion forecast, along with associated uncertainties and better suppression of false alarms.

Although different early warning algorithms may use very different source parameterizations (or no source parameterization), which makes it impossible to combine their source models, we can still combine their ground-motion predictions. We accomplish this by using Bayesian analysis to obtain what is known as the posterior hyper-robust predictive (PHRP) probability density function (PDF; e.g., Beck, 2010). This PDF is so named because it accounts not only for uncertainties in the model parameters that any EEW algorithm may use (such as earthquake magnitude, location, and OT) but also for uncertainties in the choice of which EEW algorithm to use to predict ground motion. Further, when we evaluate this PDF, we also include an extra virtual EEW algorithm that predicts no ground motion. (We refer to this as the "No Event" algorithm because it is the algorithm that always thinks there is no earthquake happening, and thus always predicts zero ground motion.) This allows us to calculate the probability that the output from the EEW algorithm(s) is a false alarm by comparing to the observed ground motion both the ground-motion predictions from the

EEW algorithm(s) and the prediction of zero ground motion that we would expect if there actually were no event.

We envision that this analysis would be done in a module that accepts input from EEW algorithms and then sends appropriate ground-motion warnings to users. We call this module the central decision module (CDM). Each EEW algorithm would send its output to the CDM, which would then construct the PHRP PDF from all reporting algorithms and send users a probabilistic description of the expected ground motion at their locations.

Methodology

A detailed derivation of the mathematical underpinnings of our methodology is given in the [Theoretical Foundation for the Method](#) section. Readers who are particularly interested in probabilistic inference may wish to read that section now. Other readers may prefer simply to read this section and skip the mathematical details in the [Appendix](#).

An EEW system such as ShakeAlert may operate any number of EEW algorithms, each of which produces an independent prediction of ground motion. Our goal is to combine these ground-motion predictions into a single unified shaking forecast. We accomplish this by estimating the probability that each EEW algorithm is correct, and then use these probabilities to determine how to appropriately blend the ground-motion predictions from each algorithm. The resulting probability distribution describing the expected future ground motion is known, as stated earlier, as the PHRP PDF (Beck, 2010). Posterior (as opposed to prior) simply means that the probability distributions have been updated based on observations of ground motions. Robust means that the ground-motion prediction is based on all potential values of an algorithm's model parameters (e.g., earthquake magnitude, location, and OT) rather than a single estimate. Finally, hyper-robust means that it also accounts for uncertainty concerning which EEW algorithm to use for predicting ground motion. (This terminology originates in the fact that the space of ground-motion predictions from all EEW algorithms is a hyperspace encompassing the space of ground-motion predictions for all possible values of the model parameters of a single EEW algorithm.)

The PHRP PDF for a quantity y , given the predictions from all of N EEW algorithms, follows directly from the total probability theorem:

$$p(y|\mathcal{D}) = \sum_{i=1}^N p(y|\mathcal{D}, A_i)P(A_i|\mathcal{D}) \quad (1)$$

(Beck, 2010), in which \mathcal{D} is the observed seismic data, and A_i is the i th EEW algorithm issuing a report for this potential earthquake detection ($i = 1, \dots, N$). The PDF $p(y|\mathcal{D}, A_i)$ is the robust posterior prediction for y from the algorithm A_i , and $P(A_i|\mathcal{D})$ is the posterior probability of algorithm A_i , based on data \mathcal{D} . ($\sum_{i=1}^N P(A_i|\mathcal{D}) = 1$, by definition, because probabilities must sum to one.)

For EEW, the quantity being predicted y will typically be peak ground acceleration (PGA), peak ground velocity (PGV), or SI, such as modified Mercalli intensity (MMI). Some users

may be interested in response spectra or computing a particular spectral acceleration (SA; e.g., [Convertito et al., 2008](#)). Thus, $p(y|\mathcal{D}, A_i)$, the PDF describing PGA (or PGV or SI or SA) predicted from each algorithm can be obtained by inputting each algorithm's earthquake source model into a GMPE.

Equation (1) shows that the PHRP PDF of the ground-motion parameter (e.g., PGA or PGV or SI or SA) y is simply the combination of the predicted value of y for each algorithm $p(y|\mathcal{D}, A_i)$, weighted by the posterior probability of each algorithm $P(A_i|\mathcal{D})$. This approach properly propagates the uncertainties from the predictions of all algorithms A_i into the PHRP PDF. The PHRP PDF is robust to parameter uncertainty within a model class (i.e., uncertainty in the model parameters estimated by an individual EEW algorithm, typically earthquake magnitude, location, and OT) and hyper-robust to model class uncertainty (i.e., uncertainty concerning which A_i to choose to make probabilistic predictions).

According to Bayes' theorem

$$P(A_i|\mathcal{D}) \propto p(\mathcal{D}|A_i)P(A_i), \quad (2)$$

in which $p(\mathcal{D}|A_i)$ is known as the marginal likelihood or evidence in favor of the model class (i.e., algorithm) A_i , and the prior probability $P(A_i)$ is a measure of the plausibility of algorithm A_i as a predictor of the quantity y . If we take an unbiased stance in which all algorithms are considered equally plausible before observing any data (i.e., $P(A_i) = \text{constant}$), then

$$P(A_i|\mathcal{D}) \propto p(\mathcal{D}|A_i). \quad (3)$$

Let \mathcal{D} be the observed waveform data and θ_i represent the source parameters for EEW algorithm A_i . For example, for On-site, ElarmS, and VS, $\theta_i = \{\text{magnitude, location, OT}\}$, whereas for FinDer, $\theta_i = \{\text{location, OT, fault strike, rupture length}\}$. Then, applying the total probability theorem

$$p(\mathcal{D}|A_i) = \int p(\mathcal{D}|\theta_i, A_i)p(\theta_i|A_i)d\theta_i. \quad (4)$$

Each EEW algorithm A_i should report its marginal likelihood $p(\mathcal{D}|A_i)$ via equation (4) as part of its real-time earthquake analysis. However, because none of the current EEW algorithms in ShakeAlert employ Bayesian analysis or compute their formal probabilities, this information is unavailable. We tested several methods of estimating the marginal likelihood $p(\mathcal{D}|A_i)$, which yields the probability of each algorithm $P(A_i|\mathcal{D})$ through equation (3). In the end, we chose to approximate $P(A_i|\mathcal{D})$ using a maximum likelihood estimate that we obtained by a constrained least-squares (LSQ) fit of the observed waveforms to the predictions from all reporting EEW algorithms. We view this as an effective but somewhat *ad hoc* fix to overcome the current inability of the EEW algorithms to provide formal posterior probabilities computed via equations (3) and (4). However, we can justify our choice by noting that using this approach ensures that the mean of the PHRP PDF, and thus the mean of the CDM's predicted ground motion, matches the observed ground motions. (See the [Appendix](#) for proof.)

The framework in equation (1) is completely general and applies to any quantity, y , which we might want to robustly predict and any PDF describing an individual EEW algorithm's prediction of y , $p(y|\mathcal{D}, A_i)$. For EEW, the quantities that we want to predict are measures of ground motion. As an example, we consider estimating the expected PGA. The most likely PGA estimate from each EEW algorithm is obtained using a GMPE. However, ground-motion variability about the mean predicted by the GMPE is quite large and a significant contributor to uncertainty in EEW predictions of expected ground motion (e.g., [Allen, 2007](#); [Iervolino et al., 2009](#)). Observed ground-motion variability about the log of the PGA predictions from GMPEs is typically modeled as Gaussian-distributed with a standard deviation σ of about a factor of 2. (Specifically, $\sigma = 1.55 - 1.8$ for \mathbf{M} 8 to \mathbf{M} 5 earthquakes in [Chiou and Youngs, 2014](#); $\sigma = 2.2$ for $\mathbf{M} < 4$ and 1.8 for $\mathbf{M} > 5.5$ in [Boore et al., 2014](#); σ ranges from 2.3 at \mathbf{M} 3 to 1.8 at \mathbf{M} 8 in [Campbell and Bozorgnia, 2008](#); and σ goes from 2.45 at \mathbf{M} 3 to 1.9 at \mathbf{M} 8 in [Abrahamson et al., 2014](#).) Thus, the PDF describing the ground-motion prediction from a single EEW algorithm A_i is

$$\mathcal{N}(\log \text{PGA} | \log \text{PGA}_i(A_i), \log \sigma), \quad (5)$$

in which $\log \text{PGA}_i(A_i)$ is the log of the GMPE-based ground-motion prediction from the i th EEW algorithm A_i , and $\mathcal{N}(\log \text{PGA} | \log \text{PGA}_i(A_i), \log \sigma)$ denotes a Gaussian distribution with mean $\log \text{PGA}_i(A_i)$ and standard deviation $\log \sigma$. Equivalently, we can write

$$\begin{aligned} p(y|\mathcal{D}, A_i) &= p(\text{PGA}|\mathcal{D}, A_i) \\ &= \log \mathcal{N}(\text{PGA} | \log \text{PGA}_i(A_i), \log \sigma), \end{aligned} \quad (6)$$

in which $\log \mathcal{N}(\text{PGA} | \log \text{PGA}_i(A_i), \log \sigma)$ denotes the (base 10) lognormal distribution, such that $\log \text{PGA}$ is Gaussian distributed with mean $\log \text{PGA}_i(A_i)$ and standard deviation $\log \sigma$.

For simplicity, we assume $\sigma = 2$. Then, the PHRP PDF for PGA is

$$\begin{aligned} p(\text{PGA}|\mathcal{D}) &= [P_{\text{No event}} \times \delta(\text{PGA})] \\ &+ \sum_{i=2}^N \{\log \mathcal{N}(\text{PGA} | \log \text{PGA}_i(A_i), \log 2) P(A_i|\mathcal{D})\}, \end{aligned} \quad (7)$$

in which we take A_1 as the No Event virtual algorithm mentioned in the [Introduction](#) (i.e., the algorithm that always believes there is no earthquake happening and so always predicts there is zero expected ground motion), and δ is the Dirac delta function (equal to one at $\text{PGA} = 0$ and zero otherwise).

The cumulative density function (CDF) associated with the PDF in equation (7) is

$$\begin{aligned} P(\text{PGA} < x|\mathcal{D}) &= P_{\text{No event}} \\ &+ \sum_{i=2}^N \{P(A_i|\mathcal{D}) \int_0^x \log \mathcal{N}(\text{PGA} | \log \text{PGA}_i(A_i), \log 2) d\text{PGA}\}. \end{aligned} \quad (8)$$

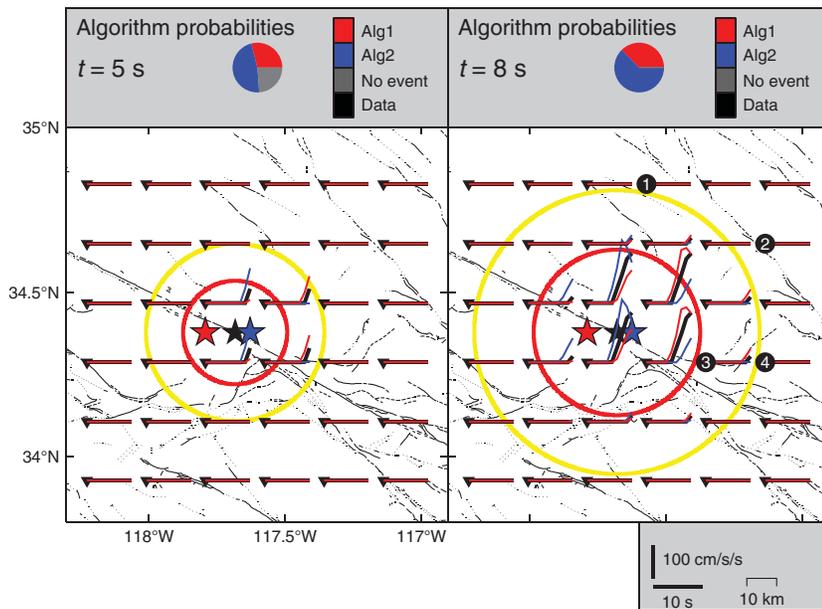


Figure 1. Waveform comparisons for a synthetic test of two algorithms that accurately estimate the earthquake’s magnitude but mislocate the event. The envelopes of acceleration for the scenario source are plotted in black. The envelopes of acceleration predicted by the two algorithms are shown in red and blue. Triangles are station locations. The black star is the location of the input source, whereas the blue and red stars are locations estimated by the algorithms. The yellow and red circles are the spatial extent after 5 s (left) and 8 s (right) of the P and S waves, assuming velocities of 6 and 3.5 km/s, respectively. Numbers in circles identify locations for which probability density functions (PDFs) of expected ground motion are plotted in Figure 2. A least-squares (LSQ) fit between the observations and the algorithms’ predictions yields the probabilities for each algorithm shown in the pie chart at the top of each subplot.

The median predicted ground motion and associated confidence intervals can be determined from the CDF (equation 8). For example, the median and upper and lower 95% confidence bounds are found by solving for x , such that $P(\text{PGA} < x|\mathcal{D})$ equals 0.5, 0.975, and 0.025, respectively.

Examples

In the following examples, at each one second increment of time, we follow these analysis steps: first, we take as our data \mathcal{D} envelopes of acceleration (Cua, 2005) observed over the last five minutes at all regional strong-motion stations contributing to ShakeAlert. Second, we generate predicted envelopes of acceleration from each reporting EEW algorithm A_i (including the predictions of zero ground motion from the No Event virtual algorithm that always thinks there is no earthquake happening) using the Cua (2005) relations between ground-motion envelopes and the distance and magnitude of an earthquake. Third, in the absence of probabilities provided by the EEW algorithms, we estimate the probability of each algorithm (including the No Event algorithm) $P(A_i|\mathcal{D})$, using a constrained LSQ fit between the observed and predicted ground-motion envelopes. (In this way, we obtain not only the probability associated with each EEW algorithm but also $P_{\text{No event}}$, the probability that the earthquake report is a

false alarm.) Finally, we obtain the PHRP prediction of ground motion by combining the ground-motion predictions of each algorithm (including the No Event algorithm) according to equation (1).

The basis for our comparison between observed and predicted ground motion is the Cua (2005) envelope relationships. Following Cua (2005), the observed envelopes of ground motion are defined as the maximum of the absolute value of the ground motion in a 1 s window. Cua (2005) used 30,000 three-component seismograms from stations within 200 km of 70 M_w 2.0–7.3 southern California earthquakes to determine relationships for the expected waveform envelope as a function of source magnitude and distance from the source. These relationships allow us to compare the time evolution of the observed acceleration waveforms to those predicted by each algorithm’s source model using only the point-source location and magnitude information that the EEW algorithms provide. Further, although in this article we have restricted ourselves to considering strong-motion acceleration records, the Cua (2005) relationships will allow us to, in the future, utilize velocity and filtered displacement seismograms as well. Alternatively, these envelope

relations could someday be replaced by a more advanced set of envelope relationships that include a PDF describing the uncertainty in the ground-motion envelopes themselves.

We present a simple synthetic test plus two real examples from ShakeAlert. In our synthetic test, we assume that two algorithms (Alg1 and Alg2) are operating and issue an alert 3.5 s after the initiation of an earthquake and that data are available for analysis within 1 s, making the first solution available at 5 s after origin (in integer seconds). The earthquake is located on the southern San Andreas fault and is surrounded by stations spaced 20 km apart. Both algorithms produce correct magnitude estimates, but the locations from the two algorithms have errors of 15 and 5 km, respectively. Although a case in which two EEW algorithms produce errors up to 15 km and no magnitude error is somewhat extreme, these are not implausible values. In a retrospective test of ShakeAlert EEW algorithms using large historic events, >75% of events have epicentral errors less than 20 km, and >80% of events have magnitude errors less than 0.5 units (E. Cochran, personal comm., 2017).

Figure 1 shows a comparison of the observed waveform envelopes of acceleration to the corresponding waveform envelopes predicted from the two algorithms. This comparison allows us to observe that the Alg2 solution (which has the smaller location error) is more probable than the Alg1 solution, and that both are more probable than the possibility

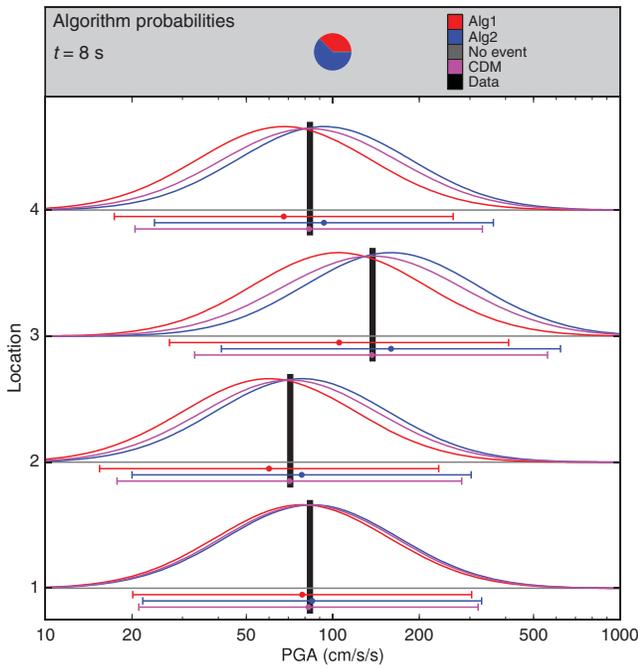


Figure 2. PDFs of expected peak ground acceleration (PGA) at the four locations marked in Figure 1. The red and blue lines are the PDFs of expected ground motion for the two earthquake early warning (EEW) algorithms, assuming that the distribution of expected ground motion is given by a lognormal distribution with mean given by the ground-motion prediction equation output for that algorithm's source model and standard deviation $\sigma = 2$ (equation 5 or, equivalently, equation 6). The magenta line is the posterior hyper-robust predictive (PHRP) PDF calculated by combining these PDFs (along with a prediction of zero ground motion), weighted by the probabilities given in the pie chart at top. Medians and 95% confidence bounds are given by the circles and brackets below the PDFs. Thick black bars show the actual PGA that will eventually arrive at each location. This figure shows how the central decision module (CDM) combines the predictions from the two EEW algorithms to produce a PGA forecast (magenta circle) that is closer to the PGA that will be observed (black line) than either of the original EEW algorithms (red and blue circles).

that this is a false alarm. As time goes on and more ground motions are available for comparison with the predictions from the algorithms, this divergence increases, with the probability that this is a false alarm going to 0% and the solution increasingly favoring Alg2 over Alg1.

Figure 2 shows the process by which the CDM creates its PDF of predicted ground motion from the probabilities computed in Figure 1. The source parameters from Alg1 and Alg2 are each used as input into a GMPE and each GMPE output is used to define the center of a lognormal PDF (red and blue lines in Fig. 2). These PDFs, along with a prediction of zero ground motion from the No Event algorithm, are weighted by the probabilities in Figure 1 and then summed, yielding the CDM prediction of ground motion (magenta line in Fig. 2). This process successfully combines the predictions from Alg1 and Alg2 to create a new PDF that better predicts the peak ground motion that will eventually be observed in these synthetic seismograms.

The first of two real examples from the ShakeAlert EEW system is the 2014 M_w 6 South Napa earthquake. For this earthquake, ElarmS produced a timely and accurate source model. Onsite had identified the stations closest to the hypocenter as having high background noise, and accordingly set a high triggering threshold for these stations. Because of this, the earliest Onsite triggers came from stations located significantly south of the earthquake, resulting in an Onsite solution that was issued late and that inaccurately located the earthquake's epicenter. VS did not issue a relevant warning for this earthquake.

One of the advantages of our CDM analysis is that it can combine the strengths and help overcome the weaknesses of each EEW algorithm, without having to know the details of how each EEW algorithm operates. It is neither necessary to know why VS did not issue a warning nor necessary to know the particulars of ElarmS and Onsite's triggering approaches or their relative effectiveness overall. (In this case, Onsite's higher triggering threshold was a weakness, but perhaps in other cases it has successfully suppressed false alarms.) What is important is to identify that, for this earthquake, the ElarmS ground-motion prediction is more probable than the Onsite ground-motion prediction, as well as more probable than the prediction that there is no earthquake, so that we can then use these probabilities to produce a probabilistic prediction of the expected ground motion.

We now apply our CDM analysis to the output from the ShakeAlert EEW algorithms for the 2014 South Napa earthquake. As time after origin increases, N in equation (1) ranges from 1 (No Event) to 2 (No Event, ElarmS) to 3 (No Event, ElarmS, Onsite). Until the first EEW algorithm reports, the CDM has no choice but to issue a 100% probability of No Event and thus zero predicted ground motion. The ideal CDM behavior would be to identify, as soon as the ElarmS solution is released, that the ElarmS solution is a good predictor of the observed ground motion and immediately push the probabilities to heavily favor the ElarmS model over No Event and to estimate ground motions in accordance with the ElarmS predictions. Then, when the Onsite solution is issued, the ideal CDM would recognize that it is a poor solution and severely downweight both Onsite and No Event in favor of the ElarmS solution. Replaying our proposed CDM analysis on the waveform data for the Napa earthquake and the evolution of alerts issued by ElarmS and Onsite produces exactly this behavior (Figs. 3–5). At 6 s after the OT, the first second for which a solution from ElarmS exists, comparison with the observed envelopes of ground motion yields a probability of $\sim 97.6\%$ in favor of the ElarmS report to 2.4% probability of a false alarm (Fig. 3). The first Onsite report is available at 12 s after OT but fits the data so poorly that it received 0% probability, with 98.4% probability to ElarmS and 1.6% probability to No Event. This means that the resulting PHRP PDFs for PGA, PGV, and SI (which are based on combining the predictions of ElarmS, Onsite, and No Event with those probabilities) is heavily weighted toward the ElarmS solution,

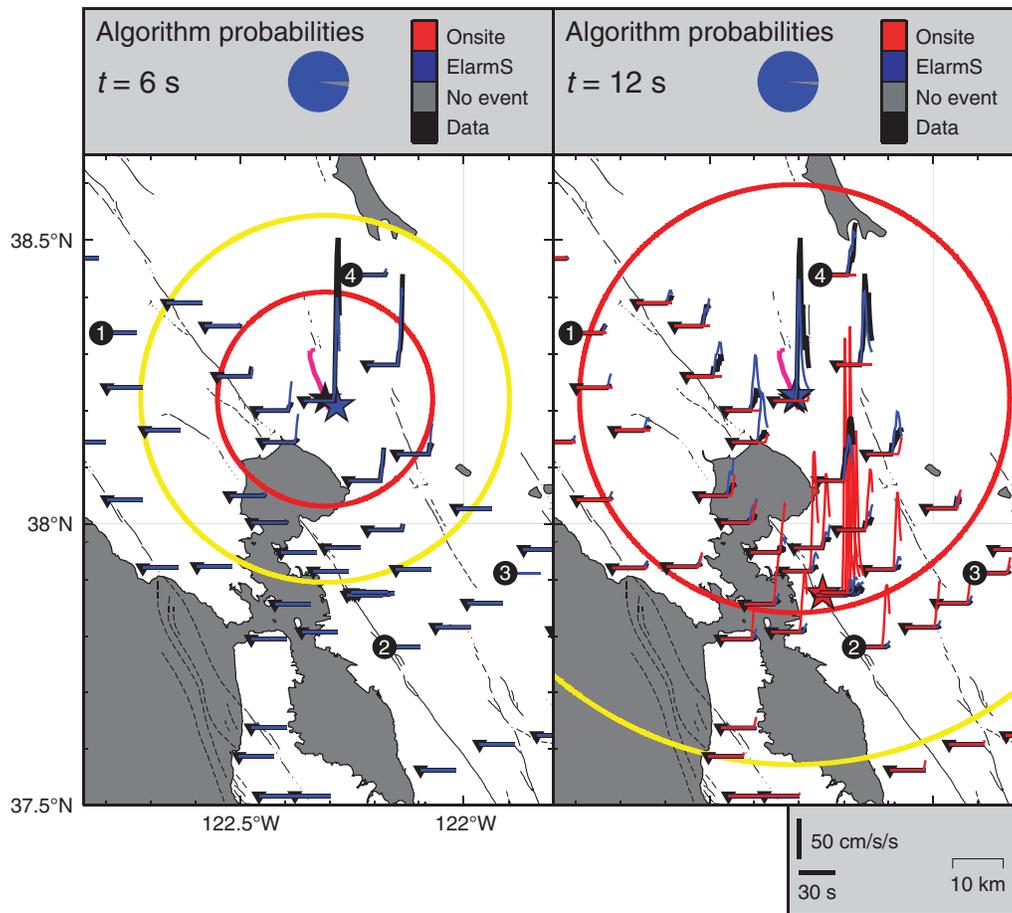


Figure 3. Waveform comparisons for the 2014 M_w 6 South Napa earthquake (left) at 6 s after origin time (OT; the time of the first Earthquake Alarm Systems [ElarmS] report) and (right) 12 s after OT (the time of the first Onsite report). The observed envelopes of vertical acceleration are shown in black. The predictions from the most recent ElarmS report are shown in blue. Onsite predictions are shown in red. The black star is the location of the earthquake (38.22° N, 122.31° W), and the pink line shows the observed surface rupture. Yellow and red circles are the farthest reach of P and S waves after 6 (left) and 12 (right) s, assuming velocities of 6 and 3.5 km/s, respectively. The blue and red stars are the current estimated epicenters from ElarmS and Onsite, respectively. The probabilities for each algorithm calculated from these envelope comparisons are shown in the pie charts. Numbers in circles identify locations for which PDFs of expected ground motion are plotted in Figure 4.

and the CDM's PDF of predicted ground motion is almost identical to the ElarmS PDF (Fig. 4). However, ElarmS does not entirely predict the ground motion observed so far. This results in the CDM ground-motion prediction gaining uncertainty relative to the original ElarmS prediction, and thus the 95% confidence bounds for the CDM ground-motion prediction are slightly broader than the ElarmS confidence bounds. It follows that the forecast ShakeMap of MMI SI computed from the CDM predictions for PGA (Fig. 4) and PGV is very similar to the ShakeMap computed from the ElarmS source model and does an excellent job of predicting the observed final ground-motion intensities (Fig. 5).

In our second example, we consider the case of 1 July 2015 in which a timing offset at several stations along California's central coast was misidentified by ElarmS as an M_w 8.2 earthquake. Neither Onsite nor VS issued a false alarm. (So, in this example N moves from 1 [No Event] to 2 [No Event, ElarmS].) The ideal CDM behavior for this

example would be to identify, as quickly as possible, that the ElarmS solution is not consistent with the observed records, make No Event highly probable, and thus push the weighted combination of predicted ground motion toward as little predicted ground motion as possible. Again, this is exactly how the CDM performs (Figs. 6–8). When ElarmS issued its first alert (which it calculated as being 34 s after the OT of the earthquake it had supposedly identified), the analysis assigned the ElarmS solution a probability of just 2.5%, and gives 97.5% probability that it is false alarm. Because at all times the probability of “No Event” is $>95\%$, the median and 95% confidence bounds of the CDM's PDF are all identically zero-predicted ground motion (Fig. 7), and the CDM's predicted ShakeMap is always zero SI (Fig. 8). For the first ElarmS report (Fig. 7, left) when there is 2.5% probability that the ElarmS solution is correct, a very slight deflection of the CDM PDF (magenta) above the zero prediction from the No Event algorithm (gray) can be seen at the location of

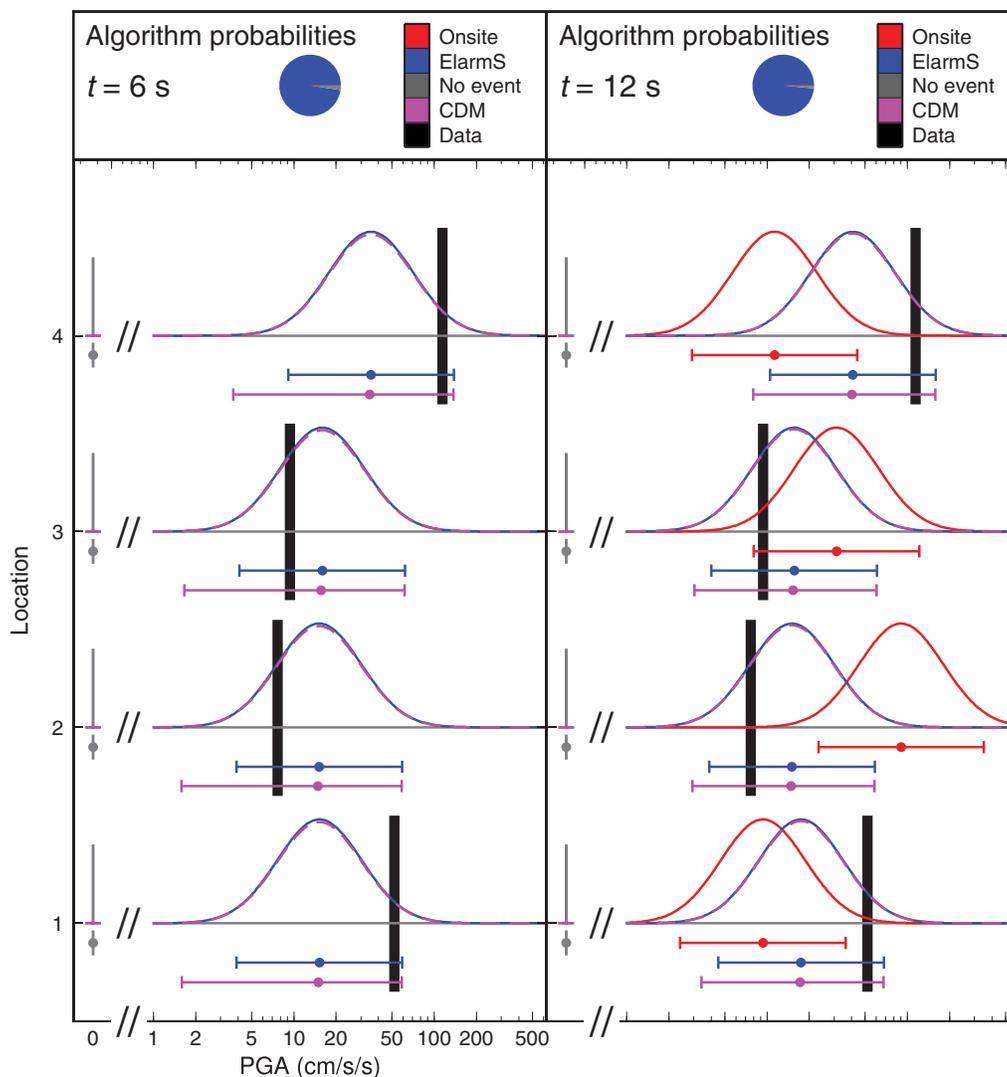


Figure 4. PDFs of expected PGA for the 2014 M_w 6 South Napa earthquake at the four locations marked in Figure 3. The red and blue lines are the PDFs of expected ground motion based on the Onsite and ElarmS source parameters, respectively. The magenta line is the PHRP PDF calculated by combining these PDFs (along with a prediction of zero ground motion), weighted by the probabilities given in the pie chart at top. Medians and 95% confidence bounds are given by the circles and brackets below the PDFs. Thick black bars show the actual PGA that will eventually arrive at each location. The probability associated with the No Event algorithm, that is, the probability that the user should expect zero ground motion because an earthquake is not actually happening, is shown by the delta function to the left of the break in scale.

the peak of the ElarmS PDF (blue). Later, when more data are available and the probability of the ElarmS solution has decreased to just 0.5% (Fig. 7, right), the CDM PDF is almost identically a prediction of zero ground motion.

Although our methodology can determine the most probable ground-motion prediction from anywhere in the space of ground-motion predictions bounded by the EEW algorithms' output and No Event, it cannot predict ground motions from outside that space. By including No Event as a ground-motion model in addition to earthquake sources from the EEW algorithms, the CDM can successfully suppress spurious ground-motion predictions from a false alarm, as we have just shown for the 1 July 2015 false alarm. However, our approach has no way to introduce additional ground motion that the EEW algorithms have failed to predict. The

effects of this can be seen in Figure 5. ElarmS underpredicts the ground motion on the northern end of the rupture, probably because the ElarmS source model is a point source located near the southern end of the rupture. Because there are no EEW algorithms contributing predictions of strong ground motion on the northern part of the rupture, the CDM has no opportunity to include that shaking in the PHRP PDF. Therefore, the CDM's ground-motion prediction in Figure 5d,h shares the underprediction of ground motion in the north exhibited by both ElarmS and Onsite.

Discussion and Concluding Remarks

The CDM framework described in this article provides an effective way to evaluate the quality of real-time earthquake

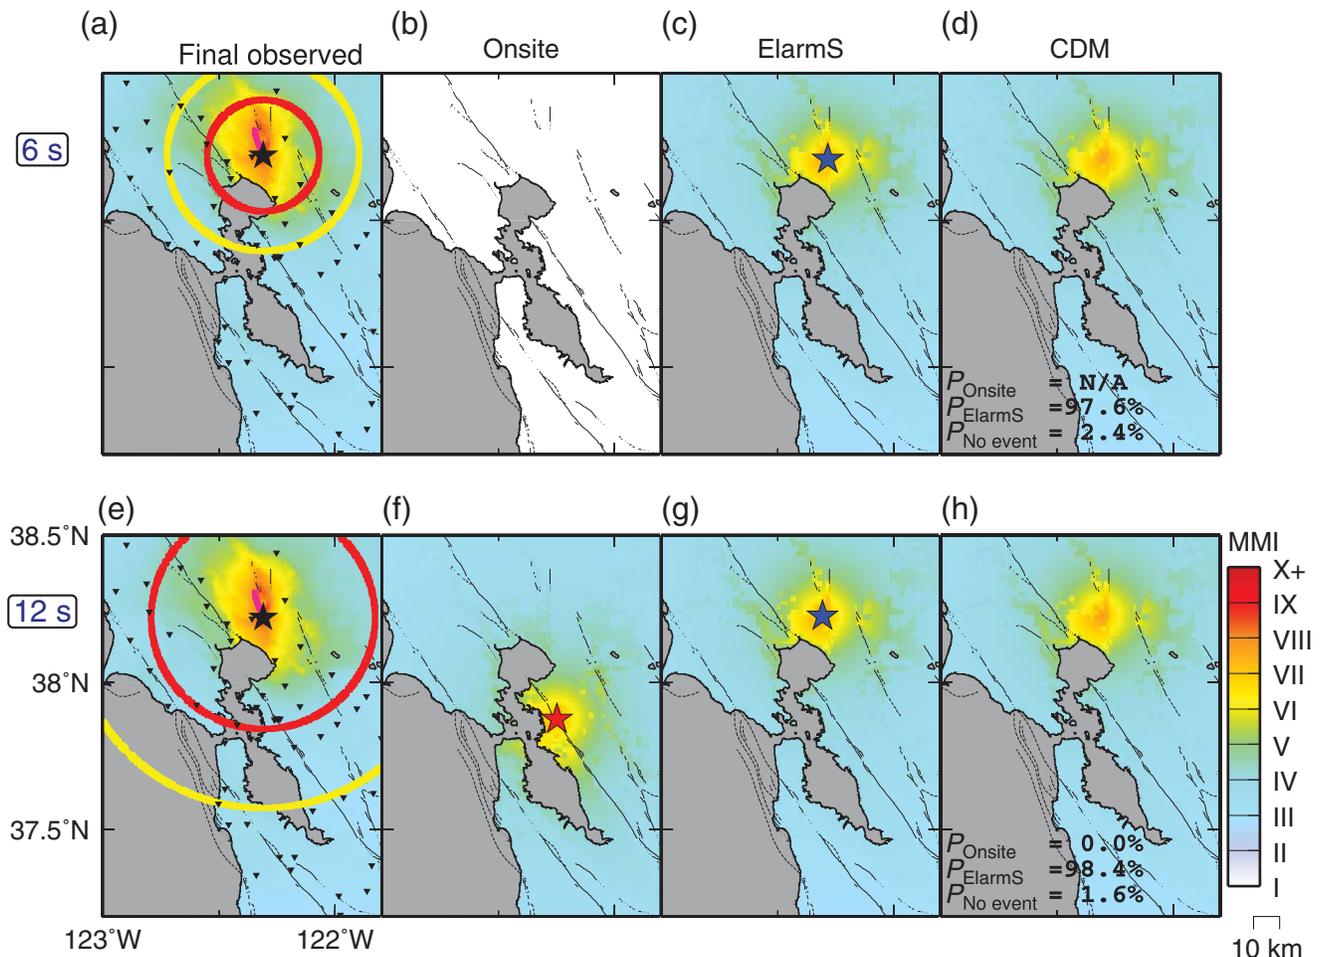


Figure 5. (a) Final observed peak horizontal shaking intensity for the 2014 M_w 6 South Napa earthquake. The triangles are the locations of stations used in this analysis. All other symbols are the same as in Figure 3. (b,c) Hypocenter location (star) and predicted distribution of shaking from the most recent Onsite and ElarmS solutions, respectively. (d) Posterior hyper-robust prediction of ground motion from the CDM. The probabilities used in the PHRP PDF were obtained via a maximum-likelihood estimate. (These are the probabilities that were computed from an LSQ fit between the observed and predicted waveform envelopes in Fig. 3 and shown in the previous pie charts.) (a) and (e) are the final observed ShakeMap for the Napa earthquake, with the P - and S -wave positions plotted at 6 and 12 s after OT, respectively. (b)–(d) uses data and EEW algorithm reports from 6 s after OT. This is the time of the first ElarmS report, and Onsite has not yet reported. Note that the CDM immediately confirms the accuracy of the ElarmS report and produces a shaking prediction that is very close to the ElarmS solution and that accurately predicts the final observed shaking intensities. (f)–(h) Solution at 12 s after OT. This is when the first Onsite report is released. Note that the CDM correctly identifies that the ElarmS solution is preferable to the Onsite solution and creates a combined solution that favors the ElarmS shaking predictions and thus matches the observed shaking values.

information. It provides one unified prediction of the most probable expected ground motion by combining shaking predictions from multiple EEW algorithms, even if those algorithms have physically incompatible source models. Our framework also provides a path to combining information from earthquake-source-based algorithms, shaking forecast models, and on-site methods of EEW. Using this approach, EEW systems can, in real time, combine information from independent EEW algorithms and assess whether the report of an earthquake from these algorithms is a false alarm. In each of our three examples, the CDM framework was able to construct an accurate ground-motion prediction as soon as the first EEW algorithm issued an alert. This suggests that the CDM analysis can provide improved ground-motion prediction and

false alert suppression without any additional latency on top of whatever duration of observation is required by the original EEW algorithm to issue an alert. (A more detailed discussion of data latencies is reserved for the [Appendix](#).)

In our framework, despite the fact that we calculate the probability that an earthquake report is a false alarm, we do not classify individual results as real earthquakes and false alarms. Instead, the CDM produces the most probable ground-motion predictions based on available data. Thus, successful false alarm suppression is expressed as producing a PDF of predicted ground motion that is tightly peaked near a prediction of zero ground motion. For the 1 July 2015 false alarm, the most probable ground motion is a combination of 97.5% no ground motion with 2.5% of the ground motion

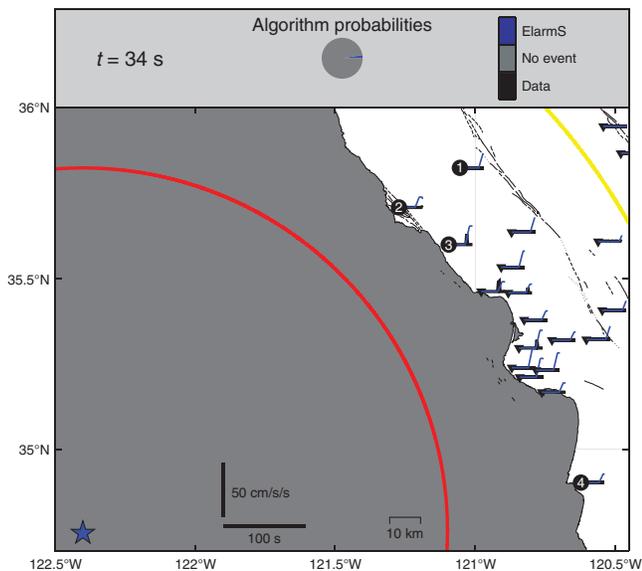


Figure 6. Same as Figure 3 for the 1 July 2015 M_w 8.2 ElarmS false alarm at the time of the first ElarmS report (34 s after the OT reported by ElarmS). The P - and S -wave positions are drawn relative to the location and OT of the ElarmS false earthquake source. ElarmS predicts large motions at stations that have no apparent motion, guiding the CDM to assign an almost 97.5% probability that this is a false alarm. Numbers in circles identify locations for which PDFs of expected ground motion are plotted in Figure 7.

predicted by ElarmS. This pushes the CDM’s predicted ground motion to zero with greater than 95% confidence.

It might seem counterintuitive to not categorize earthquake reports into true earthquakes, missed events, and false alarms, but, although it is tempting to think of an EEW system’s output as discretely and exclusively a caught earthquake, missed event, or false alarm, this does not describe how EEW actually functions. Users receive location-specific predictions for ground motion and, if such a prediction exceeds a user’s threshold for taking action, the user will react. The warning is only effective for that user if the system correctly predicts shaking above (or below) that threshold, and then the actual shaking at that location is also above (or below) that threshold. Although an EEW system that issues a report when no earthquake has occurred is obviously providing its users with a false alert, it is equally a false alarm for one or more users if, in the event of an actual earthquake, they receive a warning for shaking above their action threshold and then experience ground motions lower than their threshold. Similarly, even if an EEW system detects an earthquake and issues a shaking forecast, it will be seen by a user as a missed event if the predicted level of shaking is below the threshold for taking action, but the actual shaking exceeds the threshold. Because different users have different thresholds for actions, an identical prediction of shaking for different users may be seen variously as a correct warning, a missed warning, or a false warning, depending on how the predicted and observed ground motions compare with each user’s threshold for taking action.

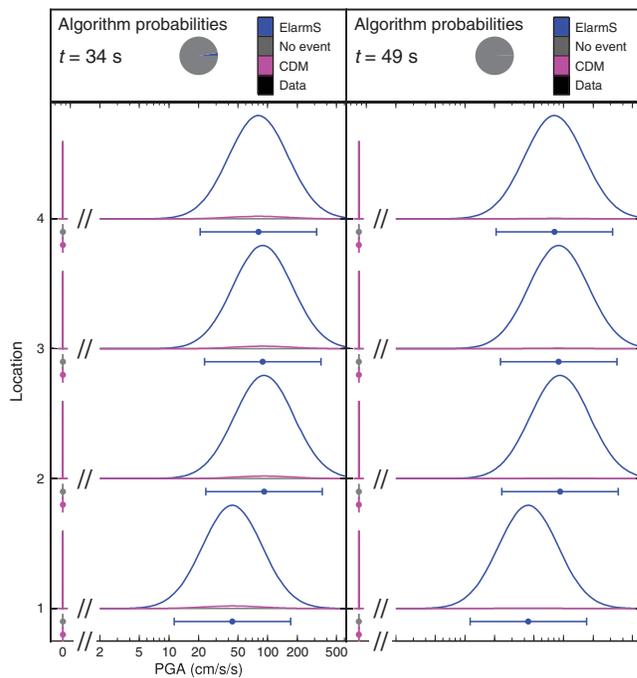


Figure 7. PDFs of expected PGA for the 1 July 2015 M_w 8.2 ElarmS false alarm at the four locations marked in Figure 6. The blue line is the PDF of expected ground motion based on the ElarmS source parameters. The magenta line is the PHRP PDF calculated by combining the ElarmS PDF and a prediction of zero ground motion, weighted by the probabilities given in the pie chart at top. Medians and 95% confidence bounds are given by the circles and brackets below the PDFs. The probability associated with the No Event algorithm, that is, the probability that the user should expect zero ground motion because an earthquake is not actually happening, is shown by the delta function to the left of the break in the scale. Because the probability of No Event is >95% at all times, the median and 95% confidence bounds of the CDM’s PDF are all identically zero-predicted ground motion.

There are further complications in EEW system performance. For example, different amounts of time are required to take different actions. (Taking cover requires little time, but redirecting airport traffic may be a much slower process.) So the identical warning that might be timely for some users might be too late for others. Additionally, some users are more false-alarm tolerant than others. (Shutting down a nuclear reactor is very costly, but opening elevator doors to allow people to escape has rather insignificant consequences if the earthquake warning turns out to be a false alarm.) So, some users that are false-alarm-tolerant but also require more time for taking action may prefer an earlier warning even if it is more likely that it will later be revealed to be a false alert. Other users might prefer to wait to take action until the shaking prediction is more certain.

An advantage of the CDM analysis is that its output is a PDF describing the relative plausibility of different ground-motion predictions rather than a single earthquake source model or even a single ground-motion estimate. Only a probabilistic approach can address the complex interactions

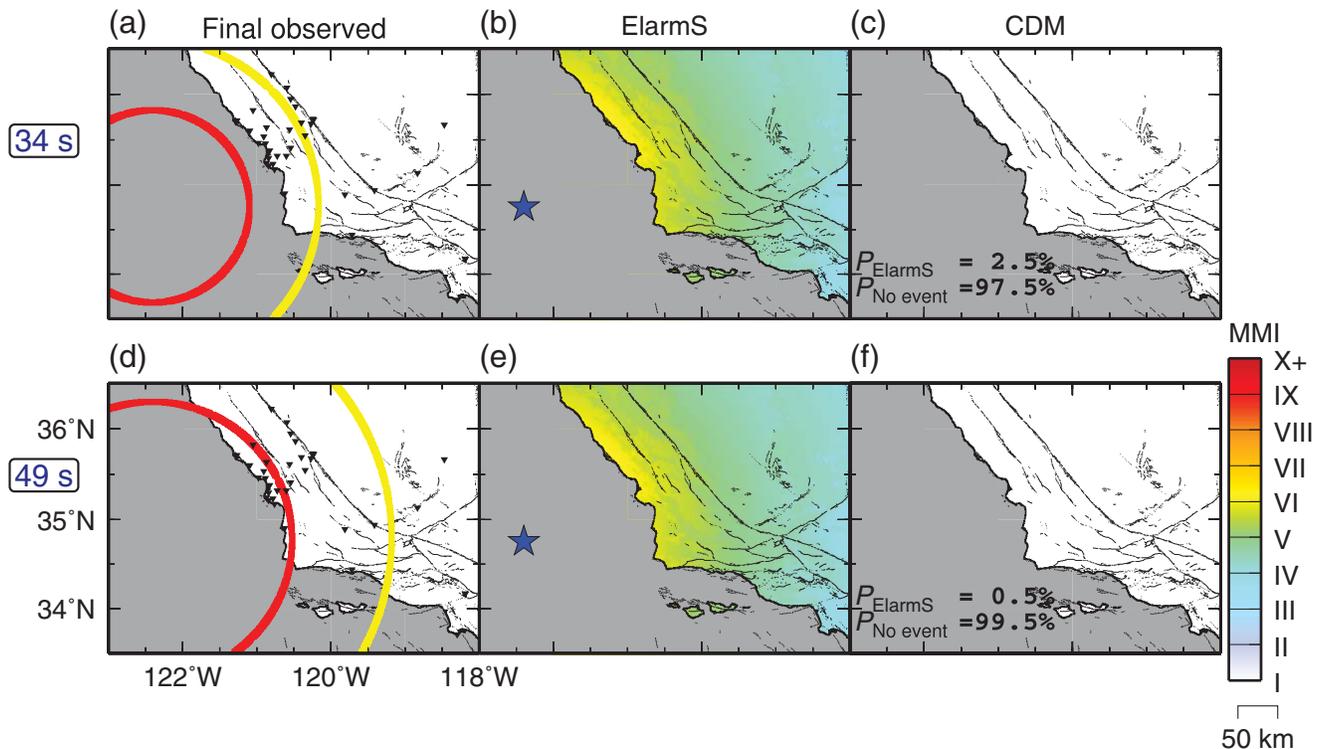


Figure 8. Same as Figure 5 for the 1 July 2015 M_w 8.2 ElarmS false alarm. (a)–(c) Snapshot of CDM performance at the time of the first ElarmS report (34 s after the OT reported by ElarmS). The CDM’s estimate of the algorithm probabilities immediately identifies that it is highly probable that this is not an earthquake and that, in turn, correctly pushes the most probable prediction to one of no shaking. As time progresses and more data become available, the probabilities even more strongly support this alert as a false alarm (d–f).

between the uncertainties of the ground-motion predictions, the warning time available to a specific user to decide whether and when to act, and that user’s individual potential costs and benefits of taking action. Ideally, each user should receive a site-specific PDF describing the expected ground motion and associated uncertainties, allowing each user to decide whether the probability that there will be ground motion exceeding the user’s threshold is high enough to warrant action or instead to wait for more information.

The user decision-making process may sound complex, but it need not be. Unsophisticated users who do not wish to deal with probabilities can use the most probable predicted ground motion from the CDM exactly as they would use a point estimate of predicted ground motion from any existing EEW system. Even users who wish to take advantage of the full probabilistic ground-motion prediction to make optimal and customized decisions could simply utilize a black box algorithm that does the probabilistic decision-making for them.

Wu *et al.* (2013) describe such a user-side algorithm called ePAD that gives each user an optimal decision as to when or if to take action. This decision is based on asking each user in advance what that user’s cost is for a false alarm, the cost of a missed event, and how long the user requires to take action. However, ePAD analyses require the EEW system to provide real-time updates of the PDFs describing the expected ground motion at the user’s location, so that the user-side module can analyze whether it is appropriate for

the user to take action based on the predicted ground motion and its uncertainty. The PHRP PDF described in this article is exactly the EEW-system-side counterpart needed to power user-specific robust decision-making systems, such as the one in Wu *et al.* (2013). Thus, not only does the CDM framework proposed here formally provide the most probable shaking forecast to all users but also it could power a next-generation EEW system that optimizes its warnings for each user by providing sufficient information for individual users to make robust customized decisions.

Data and Resources

Acceleration waveforms were obtained from the Incorporated Research Institutions for Seismology (IRIS) Data Management Center BREQ_FAST request service at http://ds.iris.edu/ds/nodes/dmc/manuals/breq_fast/ (last accessed May 2016). Details of earthquake early warning (EEW) algorithm reports are available in the  electronic supplement to this article.

Acknowledgments

The authors would like to thank Annemarie Baltay and T. C. Hanks who provided internal reviews, constructive criticisms, and useful discussions. Annemarie Baltay further provided background research on ground-motion uncertainties. Thanks are also due to Elizabeth Cochran who provided ShakeAlert performance statistics. The authors would also like to acknowledge

the journal editors and reviewers for helping revise and improve the article.

Any use of trade, firm, or product names is for descriptive purposes only and does not imply endorsement by the U.S. Government.

References

- Abrahamson, N. A., W. J. Silva, and R. Kamai (2014). Summary of the ASK14 ground motion relation for active crustal regions, *Earthq. Spectra* **30**, no. 3, 1025–1055.
- Akaike, H. (1974). A new look at the statistical model identification, *IEEE Trans. Automat. Control* **19**, no. 6, 716–723.
- Akaike, H. (1976). On entropy maximization principle, in *Application of Statistics*, P. Krishnaiah (Editor), North-Holland, Amsterdam, The Netherlands, 27–41.
- Allen, R. M. (2007). The ElarmS earthquake early warning methodology and application across California, in *Earthquake Early Warning Systems*, P. Gasparini, G. Manfredi, and J. Zschau (Editors), Springer-Verlag, Berlin/Heidelberg, Germany, 21–43.
- Allen, R. M., and H. Kanamori (2003). The potential for earthquake early warning in southern California, *Science* **300**, no. 5620, 786–789.
- Beck, J. L. (2010). Bayesian system identification based on probability logic, *Struct. Control Health Monit.* **17**, no. 7, 825–847.
- Beck, J. L., and L. S. Katafygiotis (1998). Updating models and their uncertainties. I: Bayesian statistical framework, *J. Eng. Mech.* **124**, no. 4, 455–461.
- Beck, J. L., and K.-V. Yuen (2004). Model selection using response measurements: Bayesian probabilistic approach, *J. Eng. Mech.* **130**, no. 2, 192–203.
- Boore, D. M., J. P. Stewart, E. Seyhan, and G. M. Atkinson (2014). NGA-West2 equations for predicting PGA, PGV, and 5% damped PSA for shallow crustal earthquakes, *Earthq. Spectra* **30**, no. 3, 1057–1085.
- Böse, M., R. Allen, H. Brown, G. Gua, M. Fischer, E. Hauksson, T. Heaton, M. Hellweg, M. Liukis, D. Neuhauser, *et al.* (2014). CISN ShakeAlert: An earthquake early warning demonstration system for California, in *Early Warning for Geological Disasters: Scientific Methods and Current Practice*, F. Wenzel and J. Zschau (Editors), Springer-Verlag, Berlin/Heidelberg, Germany, 49–69.
- Böse, M., C. Felizardo, and T. Heaton (2015). Finite-fault rupture detector (FinDer): Going real-time in Californian ShakeAlert warning system, *Seismol. Res. Lett.* **86**, no. 6, 1692–1704.
- Böse, M., E. Hauksson, K. Solanki, H. Kanamori, and T. Heaton (2009). Real-time testing of the on-site warning algorithm in southern California and its performance during the July 29 2008 M_w 5.4 Chino Hills earthquake, *Geophys. Res. Lett.* **36**, L00B03, doi: [10.1029/2008GL036366](https://doi.org/10.1029/2008GL036366).
- Böse, M., T. H. Heaton, and E. Hauksson (2012). Real-time finite fault rupture detector (FinDer) for large earthquakes, *Geophys. J. Int.* **191**, no. 2, 803–812.
- Campbell, K. W., and Y. Bozorgnia (2008). NGA ground motion model for the geometric mean horizontal component of PGA, PGV, PGD and 5% damped linear elastic response spectra for periods ranging from 0.01 to 10 s, *Earthq. Spectra* **24**, no. 1, 139–171.
- Chiou, B. S.-J., and R. R. Youngs (2014). Update of the Chiou and Youngs NGA model for the average horizontal component of peak ground motion and response spectra, *Earthq. Spectra* **30**, no. 3, 1117–1153.
- Convertito, V., I. Iervolino, A. Zollo, and G. Manfredi (2008). Prediction of response spectra via real-time earthquake measurements, *Soil Dynam. Earthq. Eng.* **28**, no. 6, 492–505.
- Crowell, B. W., D. A. Schmidt, P. Bodin, J. E. Vidale, J. Gombert, J. R. Hartog, V. C. Kress, T. I. Melbourne, M. Santillan, S. E. Minson, *et al.* (2016). Demonstration of the Cascadia G-FAST geodetic earthquake early warning system for the Nisqually, Washington earthquake, *Seismol. Res. Lett.* **87**, no. 4, 930–943, doi: [10.1785/0220150255](https://doi.org/10.1785/0220150255).
- Cua, G., and T. Heaton (2007). The Virtual Seismologist (VS) method: A Bayesian approach to earthquake early warning, in *Earthquake Early Warning Systems*, P. Gasparini, G. Manfredi, and J. Zschau (Editors), Springer-Verlag, Berlin/Heidelberg, Germany, 97–132.
- Cua, G., M. Fischer, T. Heaton, and S. Wiemer (2009). Real-time performance of the Virtual Seismologist earthquake early warning algorithm in southern California, *Seismol. Res. Lett.* **80**, no. 5, 740–747.
- Cua, G. B. (2005). Creating the Virtual Seismologist: Developments in ground motion characterization and seismic early warning, *Ph.D. Thesis*, California Institute of Technology.
- Grapenthin, R., I. A. Johanson, and R. M. Allen (2014). Operational real-time GPS-enhanced earthquake early warning, *J. Geophys. Res.* **119**, no. 10, 7944–7965.
- Hoshiba, M. (2013). Real-time prediction of ground motion by Kirchhoff–Fresnel boundary integral equation method: Extended front detection method for earthquake early warning, *J. Geophys. Res.* **118**, no. 3, 1038–1050.
- Hoshiba, M., and S. Aoki (2015). Numerical shake prediction for earthquake early warning: Data assimilation, real-time shake mapping, and simulation of wave propagation, *Bull. Seismol. Soc. Am.* **105**, no. 3, 1324–1338.
- Hsu, T.-Y., S.-K. Huang, Y.-W. Chang, C.-H. Kuo, C.-M. Lin, T.-M. Chang, K.-L. Wen, and C.-H. Loh (2013). Rapid on-site peak ground acceleration estimation based on support vector regression and P-wave features in Taiwan, *Soil Dynam. Earthq. Eng.* **49**, 210–217.
- Iervolino, I., M. Giorgio, C. Galasso, and G. Manfredi (2009). Uncertainty in early warning predictions of engineering ground motion parameters: What really matters? *Geophys. Res. Lett.* **36**, L00B06, doi: [10.1029/2008GL036644](https://doi.org/10.1029/2008GL036644).
- Kanamori, H. (2005). Real-time seismology and earthquake damage mitigation, *Annu. Rev. Earth Planet. Sci.* **33**, 195–214.
- Kuyuk, H. S., R. M. Allen, H. Brown, M. Hellweg, I. Henson, and D. Neuhauser (2013). Designing a network-based earthquake early warning algorithm for California: ElarmS-2, *Bull. Seismol. Soc. Am.* **104**, no. 1, doi: [10.1785/0120130146](https://doi.org/10.1785/0120130146).
- Minson, S. E., J. R. Murray, J. O. Langbein, and J. S. Gombert (2014). Real-time inversions for finite fault slip models and rupture geometry based on high-rate GPS data, *J. Geophys. Res.* **119**, 3201–3231, doi: [10.1002/2013JB010622](https://doi.org/10.1002/2013JB010622).
- Muto, M., and J. L. Beck (2008). Bayesian updating and model class selection for hysteretic structural models using stochastic simulation, *J. Vib. Control* **14**, nos. 1/2, 7–34.
- Oh, C. K., J. L. Beck, and M. Yamada (2008). Bayesian learning using automatic relevance determination prior with an application to earthquake early warning, *J. Eng. Mech.* **134**, no. 12, 1013–1020.
- Schwarz, G. (1978). Estimating the dimension of a model, *Ann. Stat.* **6**, no. 2, 461–464.
- Wu, S., J. L. Beck, and T. H. Heaton (2013). ePAD: Earthquake probability-based automated decision-making framework for earthquake early warning, *Comput. Aided Civ. Infrastruct. Eng.* **28**, no. 10, 737–752.

Appendix

Note on Data Latencies

In this section, we explore how quickly we might expect the central decision module (CDM) framework to produce robust ground-motion predictions. Data latencies are not routinely archived, making it impossible to replay past real-time performance exactly. For this article, we assumed a data latency of 1 s for our synthetic test and ignored data latencies for our examples using real earthquakes. In reality, there would be fewer data available to the CDM analysis at each instant of time because some data would not yet have arrived at the analysis center. However, the CDM will never be without data because, if there were no data available, then there would be no reports from any earthquake early warning

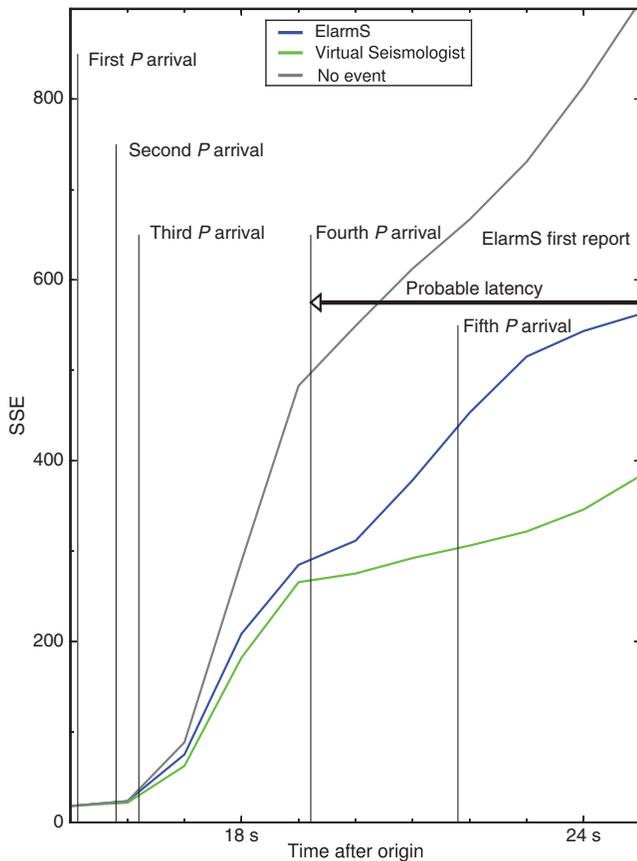


Figure A1. Sum of squared errors (SSE) misfits between earthquake early warning (EEW) algorithm predictions and observations as a function of differing data availability. The predicted envelopes of acceleration from the first report from each EEW algorithm (Earthquake Alarm Systems [ElarmS], Virtual Seismologist) for the 2014 M_w 6.8 off Cape Mendocino earthquake are compared with the observed accelerations over a 5 min window, ending at the time given on the x axis. The P -wave arrival times at the five stations nearest to the epicenter and the time of the first EEW report (issued by ElarmS) are marked.

(EEW) algorithms to analyze. Further, the CDM design proposed here will always be able to confirm or reject EEW algorithm reports because it will always be able to compare the predictions from the EEW reports to whatever data those algorithms used to issue their reports.

Perhaps the most important questions are what duration of observations is required for the CDM to make a robust ground-motion prediction, and is this duration greater than the duration of observed ground motion required for an EEW algorithm to issue an alert? If the CDM analysis can function with the same (or fewer) data as the original EEW algorithms, then we may expect that the CDM can produce a robust ground-motion prediction as soon as an EEW algorithm issues an alert, without requiring any additional observations. To explore these questions, we computed the misfit between the predicted and observed envelopes of acceleration from the first report from each EEW algorithm (Earthquake Alarm Systems [ElarmS], Virtual Seismologist [VS]) for

the 2014 M_w 6.8 off Cape Mendocino earthquake (Fig. A1). We computed this misfit for 5 min windows ending at different times, thus simulating the effects of observing different durations of data or, equivalently, the effects of different data latencies. The first ElarmS solution was issued ~ 25 s after origin time (OT). Using data up to 25 s after OT is more than sufficient for the sum of squared errors (SSE) misfit to clearly show that the predictions from the EEW algorithms are more probable than the predictions for No Event. Thus, similar to the three events in the Examples section, the CDM can issue a robust ground-motion prediction as soon as the first EEW algorithm issues its alert, without any additional observational latency required. However, this comparison may be overly optimistic because real seismic networks have data latency, and thus not all 25 s of data would be available at 25 s after OT. ElarmS requires data from at least four stations to issue an alarm. This means that the actual data latency for the 2014 off Cape Mendocino earthquake could not have been longer than the difference between when the first ElarmS report was issued and when the P wave reached the fourth station from the epicenter. Even if we only used data from the time that the P wave reached the nearest four stations and earlier, the SSE misfit for No Event is about $2\times$ larger than for either the ElarmS or VS solutions. Further, Figure A1 shows that the data favor both the ElarmS and VS solutions over No Event long before the fourth P wave is recorded, indicating that our CDM analysis can distinguish between real events and false alarms with fewer data than the EEW algorithms require to issue an alert. Thus, our proposed methodology has the ability to distinguish real earthquakes from false alarms in real time without any additional delay beyond that needed for the EEW algorithms to issue their first reports.

Theoretical Foundation for the Method

Problem Definition and System Outline

Currently, ShakeAlert includes three algorithms (Onsite, ElarmS, and VS) that provide optimal estimates of the earthquake source parameters with inconsistent uncertainty quantifications. In the future, different kinds of algorithms may also be included in the ShakeAlert system, such as algorithms that directly provide ground-motion estimates (e.g., Hoshiya, 2013; Hoshiya and Aoki, 2015) or algorithms that provide uncertainty quantification under the Bayesian framework (e.g., Minson *et al.*, 2014). To prepare a platform for consistently combining predictions from all kinds of algorithms, we develop a system based on the Bayesian model selection framework.

We begin with the details of our input data. At a given time t , our system observes waveform envelope data $D_{1:t}^j = \{z_1^j, \dots, z_t^j\}$ from n seismic stations, that is, $j = 1, \dots, n$. (We denote the full data set as $\mathcal{D}_{1:t} = \{D_{1:t}^1, \dots, D_{1:t}^n\}$.) Also, the system receives information about the potential earthquake inferred by different algorithms A_i , in which $i = 1, \dots, N$. Most algorithms will provide estimates of the standard earthquake

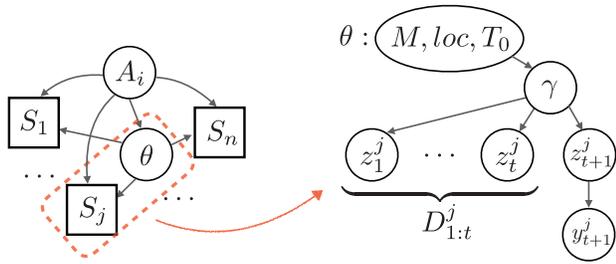


Figure A2. Graphical representation of the stochastic model class. S_j denotes the j th station. (Left) The algorithm A_i may predict the ground motion at the j th station directly or through a set of source parameters θ . For the latter case of a specific station S_j (right), θ is the input for evaluating parameters γ in a ground-motion prediction model (e.g., Cua, 2005) that can, in turn, be used to predict the ground-motion data $D_{1:t}^j$ and future ground motion z_{t+1}^j . Then that information can be used for prediction of any quantity of interest y_{t+1}^j (e.g., peak ground acceleration [PGA] or shaking intensity).

parameters $\theta = \{\text{magnitude, location, OT}\}$ directly. However, some algorithms may infer only part of θ and/or other earthquake information such as the spatial extent of the rupture. Some may even directly provide waveform predictions. As discussed in the main article, the ultimate goal of EEW is to provide local shaking information to the users. Therefore, we set the waveform predictions at each location as the common ground to combine all the predictions. In this study, our waveform data are the envelopes of the observed waveforms and, for algorithms that only provide predictions of the standard earthquake parameters, we impose an envelope-based forward prediction model (Cua, 2005) to extend the source parameter predictions to waveform predictions.

Our goal is to establish a unified framework for predicting waveform values z_{t+1}^j and any other quantity of interest y_{t+1}^j (e.g., peak ground acceleration [PGA], peak ground velocity [PGV], shaking intensity [SI]) at a future time step $t + 1$, using information from all the algorithms. More precisely, there are two stages in our system output:

1. Real-time accuracy assessment of the algorithms: we relate the accuracy of an algorithm A_i to the probability of A_i , conditional on the data and our stochastic model class, that is, $P(A_i|\mathcal{D}_{1:t})$, which is a measure of the plausibility of A_i , conditional on the information in $\mathcal{D}_{1:t}$.
2. Hyper-robust posterior prediction for the quantity of interest: this prediction corresponds to finding $p(y_{t+1}^1, \dots, y_{t+1}^n | \mathcal{D}_{1:t})$, in which $\mathcal{D}_{1:t} = \{D_{1:t}^1, \dots, D_{t+1}^n\}$, and the parameter uncertainties from all algorithms are marginalized out so that the posterior uncertainty arising from all of the algorithms is explicitly treated in making the predictions.

We summarize the information dependency between different stochastic variables using a graphical representation shown in Figure A2.

Real-Time Accuracy Assessment of Algorithms

We would like to have a consistent framework for quantifying the accuracy of an algorithm through the probability value $P(A_i|\mathcal{D}_{1:t})$. This is essentially the model selection/assessment problem in the Bayesian framework (e.g., Beck and Yuen, 2004; Beck, 2010). Using Bayes' theorem and assuming, before observing any data, that all algorithms are equally plausible as predictors, we can derive that:

$$P(A_i|\mathcal{D}_{1:t}) \propto p(\mathcal{D}_{1:t}|A_i)P(A_i) \propto p(\mathcal{D}_{1:t}|A_i). \quad (\text{A1})$$

For algorithms that are based on a Bayesian framework, $p(\mathcal{D}_{1:t}|A_i)$ is the evidence (or marginal likelihood) for the model class and can be obtained in principle from

$$p(\mathcal{D}_{1:t}|A_i) = \int p(\mathcal{D}_{1:t}|\theta_i)p(\theta_i|A_i)d\theta_i, \quad (\text{A2})$$

in which $p(\theta_i|A_i)$ is the chosen prior for the model parameters θ_i used in algorithm A_i , and

$$p(\mathcal{D}_{1:t}|\theta_i) = \prod_{j=1}^n p(D_{1:t}^j|\theta_i), \quad (\text{A3})$$

in which $p(D_{1:t}^j|\theta_i)$ can be evaluated by a probabilistic ground-motion prediction equation (GMPE) that is appropriate for source parameters θ_i . For example, the current generation of EEW algorithms, which only provide location and magnitude information, may rely on a GMPE that only depends on location and magnitude (e.g., Cua and Heaton, 2007), whereas more detailed source models that include information about the fault geometry and style of faulting may use GMPEs that include that information as well (e.g., Campbell and Bozorgnia, 2008).

Current algorithms in the ShakeAlert system do not provide consistent uncertainty quantification. To implement our system, we can either impose consistent probabilistic models upon all algorithms or infer reasonable values for the probabilities, $P(A_i|\mathcal{D}_{1:t})$, using a regression model based on the observed waveform data. Although it is theoretically possible to evaluate equation (A2) by choosing an appropriate prior $p(\theta_i|A_i)$, the probabilistic version of the envelope-based forward model we use for waveform predictions does not give an integral that can be evaluated analytically, and the evaluation of equation (A2) numerically is very time-consuming. Further, although we can substitute for equation (A2) the approximations given by Akaike's information criterion (Akaike, 1974) or the Bayesian information criterion (Akaike, 1976; Schwarz, 1978), these approximations of the evidence $p(\mathcal{D}_{1:t}|A_i)$ can be quite poor (Beck and Yuen, 2004; Muto and Beck, 2008; Oh et al., 2008). Further, these approximations have limited applicability to parameter-free EEW algorithms (e.g., Hoshiya, 2013; Hoshiya and Aoki, 2015).

Therefore, we estimate $w_i = P(A_i|\mathcal{D}_{1:t})$ by solving the following optimization problem:

$$w_i = \operatorname{argmin}_{w_i} \left\{ \sum_{j=1}^n \|D_{1:t}^j - \sum_{i=1}^N w_i \hat{D}_{1:t,i}^j\|_2^2 \right\},$$

$$\text{such that } \sum_{i=1}^N w_i = 1 \quad \text{and} \quad 0 \leq w_i \leq 1, \quad (\text{A4})$$

in which $\hat{D}_{1:t,i}^j$ is the waveform envelope for the j th station predicted by algorithm A_i . Specifically, $\hat{D}_{1:t,i}^j$ represents the observations predicted by the maximum *a posteriori* (MAP) estimate of the model parameters θ_i of algorithm A_i . The MAP estimate $\hat{\theta}_i$ is defined as $\hat{\theta}_i = \operatorname{argmax}_{\theta_i} p(\theta_i | \mathcal{D}_{1:t}, A_i)$. (Note that equation A4 is a convex constrained optimization and so has a unique minimum.)

Although the approach of choosing $P(A_i | \mathcal{D}_{1:t})$ by a least-squares (LSQ) fit is somewhat *ad hoc*, it is also completely justifiable. What the optimization in equation (A4) is doing is adjusting the posterior hyper-robust predictive (PHRP) probability density function (PDF) so that its mean matches the actual observed data. The proof is as follows:

Consider the mean of the PHRP PDF for $y = \tilde{D}_{1:t}^j$:

$$E[\tilde{D}_{1:t}^j | \mathcal{D}] = \sum_{i=1}^N E[\tilde{D}_{1:t}^j | \mathcal{D}, A_i] P(A_i | \mathcal{D}), \quad (\text{A5})$$

in which $\tilde{D}_{1:t}^j$ is the predicted data at station j (and is thus a stochastic variable), as opposed to $D_{1:t}^j$, which is the observed data at station j .

$$E[\tilde{D}_{1:t}^j | \mathcal{D}, A_i] = \int \tilde{D}_{1:t}^j p(\tilde{D}_{1:t}^j | \mathcal{D}, A_i) d\tilde{D}_{1:t}^j, \quad (\text{A6})$$

in which $p(\tilde{D}_{1:t}^j | \mathcal{D}, A_i) = \int p(\tilde{D}_{1:t}^j | \mathcal{D}, \theta_i, A_i) p(\theta_i | \mathcal{D}, A_i) d\theta_i$.

Let $\hat{\theta}_i$ be the MAP estimate of θ . Then the Laplace asymptotic approximation yields

$$p(\tilde{D}_{1:t}^j | \mathcal{D}, A_i) \approx p(\tilde{D}_{1:t}^j | \mathcal{D}, \hat{\theta}_i, A_i) = p(\tilde{D}_{1:t}^j | \hat{\theta}_i, A_i) \quad (\text{A7})$$

(Beck and Katafygiotis, 1998), if \mathcal{D} is irrelevant for prediction by algorithm A_i when θ_i is given.

Thus:

$$E[\tilde{D}_{1:t}^j | \mathcal{D}, A_i] \approx \int \tilde{D}_{1:t}^j p(\tilde{D}_{1:t}^j | \hat{\theta}_i, A_i) d\tilde{D}_{1:t}^j = \hat{D}_{1:t,i}^j, \quad (\text{A8})$$

$\hat{D}_{1:t,i}^j$ is the posterior mean prediction given by A_i for MAP estimate $\hat{\theta}_i$.

Using $w_i = P(A_i | \mathcal{D})$, equation (A5) becomes

$$E[\tilde{D}_{1:t}^j | \mathcal{D}] = \sum_{i=1}^N w_i \hat{D}_{1:t,i}^j. \quad (\text{A9})$$

Thus, when we choose the values of w_i by an LSQ match of the predicted waveform envelopes to the actual data at each station j (equation A4), we are actually choosing w_i , such that the mean ground-motion prediction from the CDM

(the mean of the PHRP PDF) will match the mean of the observed data.

As an aside, if we set $\hat{D}_{1:t,i}^j = E[\tilde{D}_{1:t}^j | \mathcal{D}, A_i]$, then the Laplace asymptotic approximation is not needed in the theory but only in the computation.

Robust Posterior Prediction of y

For any quantity of interest y , according to the total probability theorem, the posterior predictive distribution is the linear combination of the predictive distribution from each algorithm, weighted with the posterior probability of each algorithm after observing waveform data $\mathcal{D}_{1:t}$ from the stations:

$$p(y | \mathcal{D}_{1:t}) = \sum_{i=1}^N p(y | \mathcal{D}_{1:t}, A_i) P(A_i | \mathcal{D}_{1:t}). \quad (\text{A10})$$

In real time, we observe waveform data z_1^j, \dots, z_t^j . Rather than predict future waveform data, we instead want to use z_1^j, \dots, z_t^j to predict other future quantities of interest y_{t+1}^j , such as PGA, PGV, or SI. In this section, we demonstrate the validity of this approach.

In our stochastic model class, y_{t+1}^j is chosen as stochastically independent of all other variables given z_{t+1}^j (this can be easily generalized to z_1^j, \dots, z_{t+1}^j):

$$p(y_{t+1}^1, \dots, y_{t+1}^n | \mathcal{D}_{1:t}) = \int \left(\prod_{j=1}^n p(y_{t+1}^j | z_{t+1}^j) \right) p(z_{t+1}^1, \dots, z_{t+1}^n | \mathcal{D}_{1:t}) dz_{t+1}^1 \dots dz_{t+1}^n. \quad (\text{A11})$$

In practice, $p(y_{t+1}^1, \dots, y_{t+1}^n | \mathcal{D}_{1:t})$ contains much information, and it may be difficult for the user to receive and make the best use of it. Hence, we will focus on the marginal distribution $p(y_{t+1}^j | \mathcal{D}_{1:t})$. Here, we consider two ways to evaluate $p(y_{t+1}^j | \mathcal{D}_{1:t})$ and show that they are mathematically the same.

Weighting over y

Here, we first predict y_{t+1}^j for each algorithm and then combine them:

$$p(y_{t+1}^j | \mathcal{D}_{1:t}) = \sum_{i=1}^N p(y_{t+1}^j | \mathcal{D}_{1:t}, A_i) P(A_i | \mathcal{D}_{1:t}) \quad (\text{A12})$$

$$p(y_{t+1}^j | \mathcal{D}_{1:t}, A_i) = \int p(y_{t+1}^j | z_{t+1}^j) p(z_{t+1}^j | \mathcal{D}_{1:t}, A_i) dz_{t+1}^j$$

$$= \int p(y_{t+1}^j | z_{t+1}^j) p(z_{t+1}^j | \mathcal{D}_{1:t}, \hat{\theta}_i) dz_{t+1}^j. \quad (\text{A13})$$

Note that by the Laplace asymptotic approximation (Beck, 2010):

$$\begin{aligned}
 p(z_{t+1}^j | \mathcal{D}_{1:t}, A_i) &= \int p(z_{t+1}^j | \mathcal{D}_{1:t}^j, \theta_i) p(\theta_i | \mathcal{D}_{1:t}, A_i) d\theta_i \\
 &= p(z_{t+1}^j | \mathcal{D}_{1:t}^j, \hat{\theta}_i). \tag{A14}
 \end{aligned}$$

We can easily extend this to multiple y for multiple stations:

$$\begin{aligned}
 &p(y_{t+1}^1, \dots, y_{t+1}^n | \mathcal{D}_{1:t}) \\
 &= \sum_{i=1}^N p(y_{t+1}^1, \dots, y_{t+1}^n | \mathcal{D}_{1:t}, A_i) P(A_i | \mathcal{D}_{1:t}) \tag{A15}
 \end{aligned}$$

We can easily see that the two methods are the same by re-ordering the summation and the integration. Depending on the information needed from the distribution (e.g., the expected value and variance), either of the above formulations can be used for practical convenience.

Earthquake Science Center
 U.S. Geological Survey
 345 Middlefield Road, MS 977
 Menlo Park, California 94025
 sminson@usgs.gov
 (S.E.M.)

$$\begin{aligned}
 p(y_{t+1}^1, \dots, y_{t+1}^n | \mathcal{D}_{1:t}, A_i) &= \int p(y_{t+1}^1, \dots, y_{t+1}^n | z_{t+1}^1, \dots, z_{t+1}^n) p(z_{t+1}^1, \dots, z_{t+1}^n | \mathcal{D}_{1:t}, A_i) dz_{t+1}^1, \dots, dz_{t+1}^n \\
 &= \int p(y_{t+1}^1, \dots, y_{t+1}^n | z_{t+1}^1, \dots, z_{t+1}^n) p(z_{t+1}^1, \dots, z_{t+1}^n | \mathcal{D}_{1:t}, \hat{\theta}_i) dz_{t+1}^1, \dots, dz_{t+1}^n. \tag{A16}
 \end{aligned}$$

In short, the posterior-robust prediction of y equals the sum of the posterior-robust predictions from all algorithms weighted by $P(A_i | \mathcal{D}_{1:t})$.

Weighting over z

Here, we first combine the predictions of z_{t+1}^j from all algorithms and then propagate the uncertainty to the prediction of y_{t+1}^j :

$$p(y_{t+1}^j | \mathcal{D}_{1:t}) = \int p(y_{t+1}^j | z_{t+1}^j) p(z_{t+1}^j | \mathcal{D}_{1:t}) dz_{t+1}^j \tag{A17}$$

$$\begin{aligned}
 p(z_{t+1}^j | \mathcal{D}_{1:t}) &= \sum_{i=1}^N p(z_{t+1}^j | \mathcal{D}_{1:t}, A_i) P(A_i | \mathcal{D}_{1:t}) \\
 &= \sum_{i=1}^N p(z_{t+1}^j | \mathcal{D}_{1:t}, \hat{\theta}_i) P(A_i | \mathcal{D}_{1:t}). \tag{A18}
 \end{aligned}$$

Research and Development Center for Data Assimilation
 Institute of Statistical Mathematics
 10-3 Midori-cho
 Tachikawa-shi, Tokyo, Japan
 (S.W.)

Division of Engineering and Applied Science
 California Institute of Technology
 1200 East California Boulevard, MC 9-94
 Pasadena, California 91125
 (J.L.B.)

Seismological Laboratory
 Division of Geological and Planetary Sciences
 California Institute of Technology
 1200 East California Boulevard, MC 252-21
 Pasadena, California 91125
 (T.H.H.)

Manuscript received 31 October 2016;
 Published Online 13 June 2017