



Improving Out of Network Earthquake Locations Using Prior Seismicity for Use in Earthquake Early Warning

Amy Williamson^{*1}, Angela Lux¹ , and Richard Allen¹ 

ABSTRACT

Timely alerts sent through earthquake early warning (EEW) programs allow those alerted to take protective actions to mitigate their risk from potentially damaging shaking. Over the past few years, ShakeAlert, the EEW program focused on the west coast of the contiguous United States, has grown, alerting communities within California, Oregon, and Washington about earthquakes where damaging shaking is expected. ShakeAlert uses a set of algorithms including the point-source algorithm, earthquake point-source integrated code (EPIC), to determine the location, magnitude, and origin time of potential earthquakes. Although EPIC produces low-latency and low error solutions for many events originating within the seismic network on land, numerous recent small earthquakes rupturing offshore of northern California have EPIC location solutions with high error (> 50 km compared to USGS locations). Because most events are occurring offshore, there is a limited number of stations that can trigger and contribute information in a timely manner for use in EEW. To better constrain location solutions in this region, we propose to include information about contemporary past seismicity into EPIC's grid-search algorithm through a Bayesian framework. This prior information layer downweights high error locations where EPIC's proposed event location coincides with an area of low prior seismicity in preference for locations with a similar level of data fit that also have higher past seismicity. This addition to EPIC lowers the mean location error offshore northern California from 58 to 14 km.

KEY POINTS

- Revised grid-search technique produces lower average error earthquake locations for edge of network earthquakes.
- The new model uses recent seismicity as a weight, reducing high error locations.
- These better location estimates also improve magnitude estimates.



Supplemental Material

INTRODUCTION

Earthquake early warning (EEW) is the rapid detection and alerting of regions that are expected to experience damaging ground motion from an earthquake (Allen and Melgar, 2019). EEW systems need to quickly detect the location and magnitude of an earthquake, estimate ground-motion intensities, and send alerts to end users. An efficient system leveraging a dense seismic network can provide seconds to tens of seconds of warning. In brief, this amount of time is enough for those alerted to take simple but effective protective measures

such as to drop, cover, and hold on. Automated systems also benefit greatly from EEW. With seconds of notice, high-speed trains can slow, reducing the risk of derailment and allowing for an assessment of the track ahead. For example, California's Bay Area Rapid Transit system has been an early adopter of EEW alerts for their transit system (Strauss and Allen, 2016). Japan's high-speed Shinkansen trains also utilize an EEW system to mitigate risk (Nakamura and Saita, 2007).

Through partnerships, the U.S. Geological Survey (USGS) and participating academic institutions have developed and implemented ShakeAlert, an EEW system spanning California, Oregon, and Washington. Since 2012, ShakeAlert has been issuing alerts to pilot users and select partnering institutions. By 2018, public alerts were made available in

1. Berkeley Seismological Laboratory, University of California, Berkeley, Berkeley, California, U.S.A.,  <https://orcid.org/0000-0002-3767-6018> (AL);  <https://orcid.org/0000-0003-4293-9772> (RA)

*Corresponding author: amy.williamson@berkeley.edu

Cite this article as Williamson, A., A. Lux, and R. Allen (2023). Improving Out of Network Earthquake Locations Using Prior Seismicity for Use in Earthquake Early Warning, *Bull. Seismol. Soc. Am.* **XX**, 1–12, doi: [10.1785/0120220159](https://doi.org/10.1785/0120220159)

© Seismological Society of America

California, and, as of 2022, ShakeAlert issues public alerts across the entire U.S. West Coast. Communities can receive phone alerts through the MyShake App for earthquakes with a magnitude greater than 4.5 and a Modified Mercalli Intensity (MMI) \geq III (Strauss *et al.*, 2020). For regions with an expected shaking of MMI \geq IV and earthquake magnitude greater than M 5, users may also receive a wireless emergency alert (WEA; Kohler *et al.*, 2020). A recent example of an event that prompted the issuance of both a WEA and MyShake App alert was the 25 October 2022 M 5.1 Alum Creek, California, earthquake, rupturing just southeast of San Jose, California.

ShakeAlert uses three operational layers to handle incoming seismic data, production, and outgoing alerts. The data layer ingests and handles ground-motion data from seismic stations participating in the ShakeAlert network. These data are then sent to the production layer where EEW specific algorithms, such as earthquake point-source integrated code (EPIC; Chung *et al.*, 2019), and FinDER (Böse *et al.*, 2015, 2018) are housed. These algorithms detect earthquakes, solve for location and magnitude, and estimate the intensity and extent of the expected shaking. The alert layer analyzes the EEW solution and sends a product if the event passes set quality checks, has a large enough magnitude, and has high enough expected ground motion to merit an alert. All EEW processing occurs on redundant servers located in Seattle, Menlo Park, Berkeley, and Pasadena. This allows for operational continuity in the event of a data outage at any site.

In this study, we propose bEPIC, an improvement to the methodology of the EEW algorithm EPIC, with the goal of limiting high location error solutions for offshore or edge of network earthquakes. We employ a Bayesian framework to modify EPIC's current grid-search location algorithm, including the contemporary seismic history over the western United States as prior information. We test bEPIC on a catalog of recent earthquakes within the ShakeAlert region. The revised earthquake locations through bEPIC are compared against historic EPIC solution, and performance improvement is calculated with respect to the final USGS solution.

The use of a Bayesian framework for earthquake characterization is not new. Yin *et al.* (2018) incorporates the recent seismic activity rate to determine the likelihood that a trigger at any individual station is related to an earthquake. This can reduce the likelihood of nonseismic triggers contaminating location algorithms, producing high-error results. The GaMMA model (Zhu *et al.*, 2022) has a component that similarly models the probability of a phase pick given the potential for multiple local earthquakes and with this information associates the *P*- and *S*-wave phases at nearby seismic stations and determines earthquake source parameters such as location and magnitude. When solving the rupture extent for large earthquakes, Minson *et al.* (2013) uses a Bayesian methodology to solve finite-fault inverse problems in synthetic environments and then use the same model to focus on the 2011 Tohoku–Oki earthquake (Minson *et al.*, 2014). Lomax *et al.* (2009) lay out a

generalized methodology for Bayesian grid-search techniques, including test examples using the locations of stations and sources from past Italian earthquakes. In this study, the introduction of bEPIC fulfills a need for an algorithm that will operate under the temporal limitations of EEW without the need for *S* wave or full waveform information in the initial grid search.

ALGORITHM BACKGROUND

This study focuses on improvements affecting EPIC's location algorithm; formally known as the ElarmS algorithm (Allen *et al.*, 2009; Chung *et al.*, 2019). Prior to locating an event, a separate triggering algorithm uses a short-term average/long-term average method to detect the arrival of *P* waves on individual station channels. Although this triggering algorithm is separate from the location algorithm, new triggers are regularly fed into EPIC when made available. As time elapses and seismic waves propagate farther from the hypocenter, the number of triggers increases, especially in regions where the seismic network is dense. With one trigger, a preliminary epicentral location is set at the location of the triggered station. With two triggers from two different stations, the preliminary location moves to a point between the first and second station. Once provided with triggers from at least four unique stations, EPIC uses a direct grid search to determine the best-fitting event location (Chung *et al.*, 2019, 2020). The grid space initially is centered on the preliminarily epicentral location determined from the sub-four triggers estimate. As better location estimates are made available through the inclusion of more triggers at unique stations, the grid space is updated and the grid center moves to the previous timestep's best-event location. The best location from this new search is the grid node that has the lowest data misfit between observed and modeled station trigger times.

As a way of limiting the computational cost and reducing the number of parameters to be solved, EPIC preassigns a depth of 8.0 km for all events. This simplification is reasonable as most target events in the ShakeAlert reporting region occur on shallow crustal faults (Wurman *et al.*, 2007; Thompson *et al.*, 2022). In tandem with the location algorithm, the magnitude is calculated using a distance and *P*-wave peak displacement (P_d) scaling relation (Wurman *et al.*, 2007; Kuyuk and Allen, 2013; Chung *et al.*, 2020). As new stations are triggered and more waveform information prior to *S*-wave arrival is available, EPIC updates its location and magnitude solution. Numerous quality checks are included to differentiate between noise, calibration pulses, and earthquake shaking. This is primarily done within the trigger algorithm before passing on to the location phase. EPIC also limits false triggers by including a “filter bank” discriminator to reject picks related to large teleseismic earthquakes (Chung *et al.*, 2020).

As a means of limiting the total number of tested points in the location grid search, EPIC also computes the activity level at each grid-search node based on the number of nearby active (triggered) and inactive (not triggered) stations (Serdar

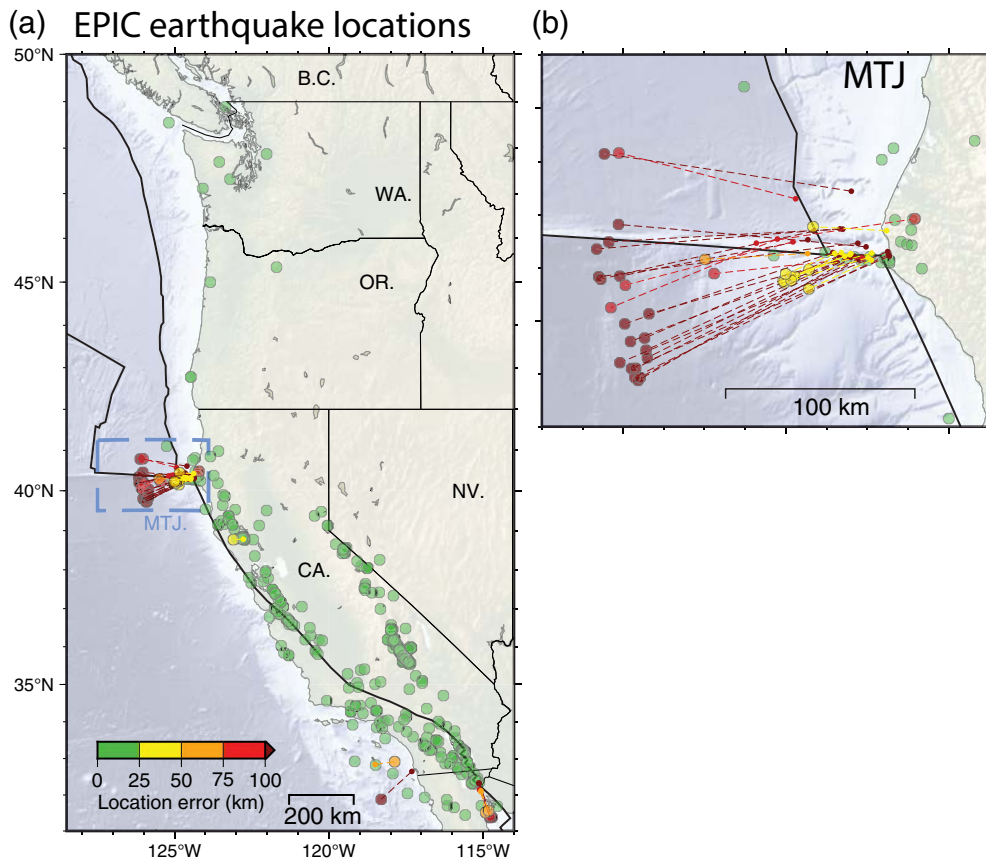


Figure 1. (a) Summary of EPIC location and magnitude performance from October 2018 to May 2022. EPIC final event locations are plotted as green, yellow, orange, red, and dark red dots when the location error is below 25 km, between 25 and 50 km, between 50 and 75 km, between 75 and 100 km, and greater than 100 km, respectively. Dots are plotting at EPIC's location. For each event, a corresponding colored dashed line connects the EPIC location with the final U.S. Geological Survey (USGS) location. The blue dashed box outlines what we define as the Mendocino Triple Junction (MTJ) region and is used to classify the offshore events. (b) Zoomed in view of the MTJ region, and location error between EPIC and USGS locations. BC, British Columbia; WA, Washington; OR, Oregon; NV, Nevada; CA, California. The color version of this figure is available only in the electronic edition.

Kuyuk *et al.*, 2014). At each timestep, EPIC tallies the number of triggered (active) and untriggered (inactive) stations that are located inside the grid-search spatial domain. At each grid node in the search, the maximum distance between the node and active stations is calculated. Then, a circle centered on the grid node with a radius equal to distance of the farthest active station is drawn. The percent of active stations inside this circle must exceed a preset threshold. Here, we set this value to be 30%. This means that at least 30% of stations near this grid node need to have triggered for this node to be considered as a viable earthquake location. The purpose of this exercise is to limit erroneous earthquake locations on land in areas where many functioning stations did not trigger. However, this activity level filter has little effect on limiting potential offshore grid node locations due to a lack of offshore stations used in the ShakeAlert network.

The latest version of EPIC excels at providing low latency location and magnitude solutions for earthquakes located on

land, within the dense seismic network. This is illustrated in Figure 1, which shows the location and magnitude estimate performance of EPIC for earthquakes from October 2018 to May 2022. One area of potential improvement is the rapid location of earthquakes rupturing offshore of northern California in a seismically active area around the Mendocino Triple Junction (MTJ). Almost all events with high error (>50 km with respect to final USGS locations) occur within the MTJ region. Here, the average location error for EPIC's final event solution is 58 km. In contrast, the average location error for events that originate outside the MTJ is only 4 km. High location errors also can contribute to large errors in the magnitude calculation. Because the magnitude is calculated with a scaling law that depends on station-epicenter distances, a large, and erroneous, increase in distance will contribute to an overestimate in magnitude; in some poorly located cases the magnitude is a full unit higher than the true solution. Although some

recent earthquakes used in our test data set were too small to merit the issuance of a public alert, larger magnitude earthquakes within the MTJ area are possible. For example, the 1992 *M* 7.1 Cape Mendocino earthquake ruptured just offshore, possibly on the subduction interface (Oppenheimer *et al.*, 1993). This was followed by the 1994 *M* 6.9 earthquake, farther west along the Mendocino fault zone, and numerous other *M* 6.0+ earthquakes along the fault zone and within the Gorda deformation zone (Rollins and Stein, 2010). Improved detection and characterization of all offshore events raises confidence in the less common, larger offshore earthquakes, limiting potential future false alarms.

DATA

To test the performance of the bEPIC algorithm, we compiled a replay catalog composed of past events from ShakeAlert. Each event in the replay catalog has an EPIC solution containing the

earthquake location, magnitude, and origin time over the entire time history of the event, that is, over a duration of tens of seconds as the seismic network collected real-time data. The catalog also contains the event ID and source information for the corresponding USGS event solution to allow for an easy comparison. We include earthquakes with either EPIC or USGS determined magnitudes greater than 3.5 that occurred between October 2018 and May 2022. For each replay event, we gather information on which stations were triggered and their trigger times. We visually inspect the triggers for each event and manually discard events for which the trigger data set is contaminated with triggers from noise, boxcar functions, the passage of teleseismic waves, and cases where the seismic waves of more than one event are visible within the seconds around the trigger (Fig. S01, available in the supplemental material to this article). Of an initial catalog with 628 events, we discarded 86 events due to poor data quality, leaving us with a replay catalog of 542 earthquakes. Although events with high noise and poor-quality data are also prone to producing high location error solutions, we wish to first find improvements to EPIC for cases where the data quality is good, but the grid-search location algorithm performs poorly. Although it is an important problem for early warning, we leave the question of how to identify and handle the ingest of poor-quality data in real time to a future study. Many of the remaining earthquakes rupture within California, but a few events also occur in British Columbia, Canada, Washington, Oregon, Nevada, and Baja California, Mexico.

For each event in the replay catalog, we create a table of all the stations that triggered, the trigger times, and the timestamp when EPIC included that data into the location algorithm. This table is queried at each timestep in the replay. In a typical solution, EPIC will recompute the earthquake location as more station triggers are made available due to the passage of body waves at farther and farther distances from the source. Although EPIC will often recompute the earthquake location at subsecond intervals, the exact timing of each recomputation is dependent on the time between new station triggers. EPIC ceases computing locations when no new stations are triggered due to the P wave's attenuation with distance to below the triggering algorithm's detection level or due to stations exceeding a set distance limit of 200 km away from the previously computed location, whichever is closer. By having our replay model mirror EPIC's inclusion of stations over time, we ensure a better solution comparison between our replay and real-time performance.

We gather earthquake magnitude and location information from the Advanced National Seismic System (ANSS) Comprehensive Earthquake Catalog (ComCat) for use in our prior seismicity function. The base catalog contains all events from 1 January 2000 to May 2022 of magnitude 3.0 or greater in a region enveloping the ShakeAlert reporting zone. As new earthquakes occur within this zone, the catalog

can be updated at regular temporal intervals to reflect the new information. All events, regardless of magnitude and age, are weighed equally in the seismicity catalog. We test our algorithm with the assumption that bEPIC will have access to the seismicity catalog, which would be updated over time to include new events. When testing bEPIC with our replay catalog, we only include in seismicity up to but not including the time of our test event. This removes the potential that the test event's appearance in the ComCat catalog, or any potentially related aftershocks, biases our solutions. When not testing past replay events, the full catalog up to present is used.

METHODS

For each step in the processing of each earthquake in our replay catalog, we calculate the earthquake location (latitude and longitude) and magnitude. We draw our preferred earthquake location from the most probable grid node using a Bayesian framework. In the following section, we describe how we formulate the bEPIC algorithm, including the construction of the likelihood and prior seismicity functions.

The posterior probability density function (pdf) of the model parameters, $P(\mathbf{m}|\mathbf{d})$, can be obtained through Bayes' theorem:

$$P(\mathbf{m}|\mathbf{d}) = \kappa\rho(\mathbf{m})\alpha(\mathbf{m})P(\mathbf{d}|\mathbf{m}),$$

in which \mathbf{d} is a vector containing the P -wave arrival times at N stations provided by EPIC's triggering algorithm: $\mathbf{d} = (t_1^{\text{obs}}, \dots, t_N^{\text{obs}})$. The components of the model parameter vector are the x and y coordinates of all potential hypocenter locations, $\mathbf{m} = (x, y)$, that are possible within the grid-search domain.

The prior information is expressed by $\rho(\mathbf{m})$, a seismicity probability for all grid nodes in model vector \mathbf{m} . The likelihood $P(\mathbf{d}|\mathbf{m})$ indicates how well the data fit the model at each point on the grid. The constant κ is a normalization constant that ensures that the integral of the posterior probability density function is equal to 1. An additional parameter that is currently implemented in EPIC $\alpha(\mathbf{m})$ represents the station activity level at each grid node. This binary vector of length \mathbf{m} limits the posterior solution to parts of the grid-search domain in which the total percent of nearby stations that triggered from the event are above a preset value. In practice, this mask removes grid nodes (by setting their value to zero) that are close to untriggered stations using the logic that if the earthquake was close to those stations, they would have triggered. This mask is effective for events located onland inside the seismic network. It has limited use for offshore events because most if not all potential earthquake locations are located away from stations.

The likelihood function is a nonnormalized measure of how well each grid point explains the observed data while also considering uncertainties inherent in the assumptions included in

both the data and algorithm. It provides an estimate of the best-event location over the spatial domain. The conditional probability from the likelihood function is

$$P(\mathbf{d}|\mathbf{m}) \propto L(\mathbf{m}) = \exp\left(-1/2 \sum_{i=1} [t_i - t_i^{\text{calc}}(\mathbf{m})]^2 / \sigma_i^2\right),$$

in which t_i is the observed P -wave arrival time at each station, i and t_i^{calc} is the calculated P -wave arrival time. The station calculated P -wave arrival times are derived from the station and grid node distances, the expected travel time using a velocity model, and a calculated origin time specific to each grid node.

The arrival-time uncertainty at each station is denoted by σ , a vector of length N . The value used for the uncertainty in arrival times stems from incomplete knowledge of the velocity structure along the ray path, errors in the triggering algorithm, and uncertainties stemming from the coarseness of the grid-search spatial domain. A grid node with a likelihood value that approaches the value one has a very low misfit, thus high likelihood of being the location of the earthquake. A low likelihood has a high misfit and is less likely to be the true location of the earthquake.

The prior function is a representation of our best knowledge about where earthquakes likely occur and is based on past local seismicity. The data included in the prior are pulled from the ANSS ComCat catalog. In replay mode, only the earthquakes with rupture dates prior to the test event are included in the data set.

We construct our prior pdf $\rho(\mathbf{m})$ using a 2D kernel density estimate (kde). The kde is a way to characterize the probability of a random variable, in this case, the location of past earthquakes. The epicenter of each past earthquake that occurred inside the grid-search domain is characterized as a Gaussian kernel K and takes the form

$$K(\mathbf{m}|\boldsymbol{\mu},\Sigma) = \frac{1}{2\pi\sqrt{|\Sigma|}} \exp\left(-\frac{1}{2}(\mathbf{m} - \boldsymbol{\mu})^T \Sigma^{-1}(\mathbf{m} - \boldsymbol{\mu})\right),$$

in which the bandwidth Σ is a 2×2 covariance matrix controlling the shape and extent of smoothing of each kernel. We assume the location in both the x and y direction of the past seismicity are independent variables and assign a bandwidth value following Scott's rule (Scott, 2015): A smaller bandwidth value will limit the amount of smoothing; a larger value will extend the smoothing of each Gaussian kernel. The prior $\rho(\mathbf{m})$ is the summation and normalization of all kernels:

$$\rho(\mathbf{m}) = \frac{1}{n} \sum K(\mathbf{m}|\boldsymbol{\mu},\Sigma).$$

In application, the prior kde is applied only over the model space, limiting the number of events from the prior seismicity catalog to those that occur inside the grid-search domain. If

there is a situation for which no past earthquakes occurred inside the model domain, then the value of the prior will be constant and, therefore, have no impact on the final posterior solution.

The expression of the prior is shown across the western United States in Figure 2. Here, we show all past seismicity incorporated into our reference catalog. In this example, all past data are incorporated into the kernel density estimate to illustrate areas of low and high seismicity. Inset figures show the distribution of seismicity within the MTJ (Fig. 2b) and, for reference, in a section of California (Fig. 2c). In execution, the prior is only computed over the same spatial domain as the location grid search. In addition, when replaying past earthquakes in our testing catalog, we take care to only include earthquakes that occurred at times prior to our replay event. Although the effect of the addition of one or a handful of earthquakes to the past seismicity catalog is small, we want to limit potentially biasing effects in our results.

EPIC's triggering algorithm provides a list of stations that triggered, which is then used as inputs into the location algorithm. However, stations that are within our grid-search domain but did not trigger can be used as a data set to define our activity mask. The activity mask, $\alpha(\mathbf{m})$, contains either a 1 or a 0 depending on whether the grid node is within or outside of the network. The network is defined here as any grid node where at least 30% of the nearby stations are triggered. To determine what is inside or outside this network, we calculate the distance between the grid node and all stations inside the grid-search domain. From this, we draw a circle with a radius equal to the distance of the farthest triggered station and assess how many stations inside this circle triggered and how many did not. If more than 30% of stations are triggered inside the circle, then this grid node is considered inside the network. The value for alpha at the index of this grid node is then assigned to be 1. Because new stations have been added to the ShakeAlert network over the duration of our test catalog, we only include stations that were installed prior to the time of the replay event.

Finally, the posterior pdf is the product of the likelihood, prior, and activity mask functions, normalized by a constant κ . The best solution is drawn from the maximum value of the posterior function. This best solution is then used as the new initial location for creating a grid search in subsequent iteration of solving the event location, if more information from new seismic stations is made available. This best location is also used when calculating the earthquake magnitude.

Once the earthquake location is known, the magnitude at each station is computed using a distance and P -wave peak displacement scaling relation (Kuyuk and Allen, 2013):

$$M = a \times \log_{10}(P_d) + b \times \log_{10}(R) + c,$$

in which a , b , and c are empirical constants and set as 1.23, 1.39, and 5.39, respectively. R is the distance from the station

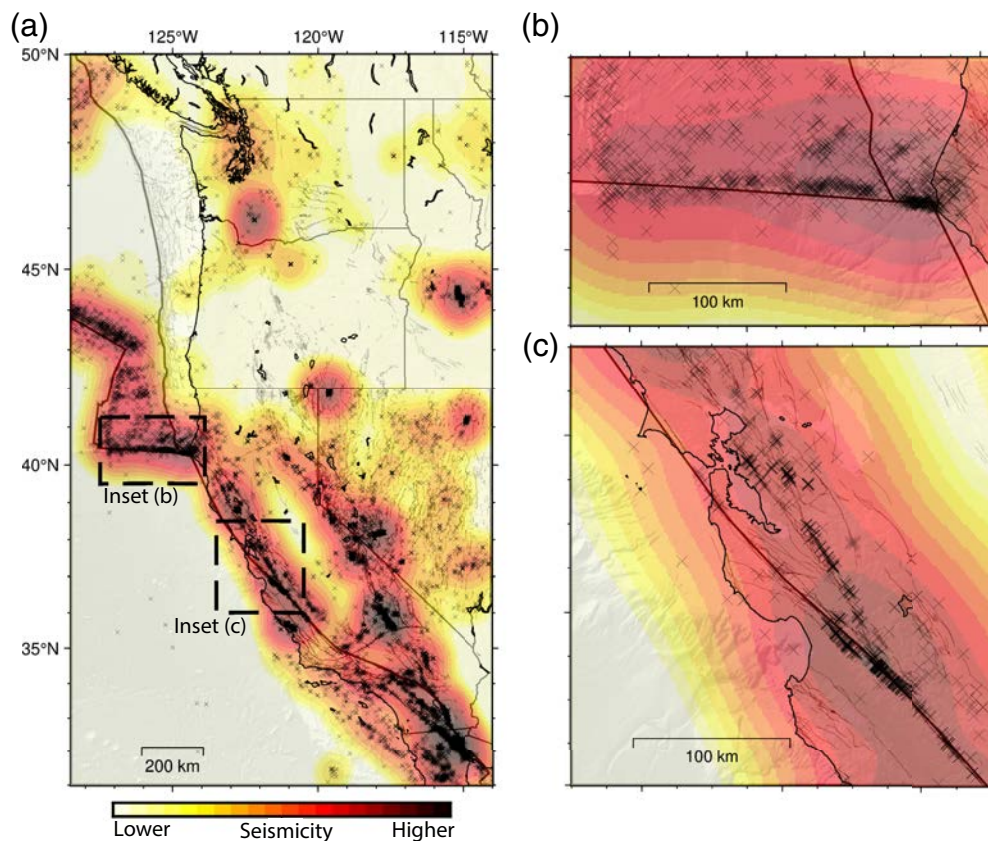


Figure 2. (a) Contemporary seismicity of M 3.0 and larger events across the western United States, marked with black x icons. Black dashed lines indicate the inset panels (b) and (c). Panel (b) shows the extent of the MTJ region. Panel (c) shows a subsection of California. The entire region is shaded based on levels of seismicity where lighter colors have lower seismicity, and darker shades have higher seismicity. The color version of this figure is available only in the electronic edition.

to the hypocenter, in kilometers. P_d is the peak displacement observed on the sensor in up to the first 4 s following the P -wave arrival time. The final earthquake magnitude is an average of the individual station magnitudes. As the location is updated and more of the initial waveform is made available at triggered stations, the magnitude is updated.

RESULTS

The inclusion of the prior seismicity layer to EPIC's location algorithm greatly improves accuracy for earthquakes nucleating in the MTJ. Figure 3 shows the improvement in location for events across the ShakeAlert reporting region and specifically within the MTJ (Fig. 3b). Figure 4 shows the improvement in location accuracy when using bEPIC for the entire reporting area (Fig. 4a) and for events within the MTJ (Fig. 4b). In the original EPIC location solutions, the mean and median location error for events inside the MTJ is 58 km and 34 km, respectively. The solutions from bEPIC reduce the mean error to 14 km and the median error to 7 km. On average, bEPIC improves the location accuracy of earthquakes in

the MTJ by 44 km. The inclusion of prior seismicity also does not negatively affect the location quality for events that occurred onland and outside of the MTJ. With EPIC, the events that occurred outside the MTJ had a mean and median location error of 4 and 2 km. bEPIC produces results with mean and median location errors of 3 and 2 km for this same data set. Two events outside of the MTJ had EPIC location errors that exceeded 100 km. Both occurred near the southern California border with Mexico. bEPIC improves the location estimate for both of these cases as well.

EPIC and bEPIC use the same scaling laws to compute the earthquake magnitude. Because this scaling law relies in part on the distance between the station and the earthquake, an improvement in earthquake location also has the potential to improve the magnitude estimate. Figure 4c,d shows an improvement in the accuracy of the magnitude calculation

when using bEPIC. The average magnitude error for earthquakes within the MTJ improves from 0.4 units with EPIC to 0.06 magnitude units using bEPIC. For events outside of the MTJ, the average magnitude error for EPIC and bEPIC events is similar at 0.06 and 0.07 magnitude units, respectively. Overall, the magnitude error is strongly related to the location error, for which an increase in location error also increases the magnitude error (Fig. S02). A reduction in the former will likely cause an improvement in the latter.

DISCUSSION

A common feature in EPIC solutions for earthquakes located within the MTJ is a location overshoot. This is when there is not enough station data in the right places to constrain the grid-search location, creating a large area of reasonably high likelihood extending far offshore. Because of the similar high likelihood, it is common for EPIC to draw a location far offshore and away from the true location. Then, as more data are made available, EPIC redraws its grid-search domain based on this overshoot location, which can perpetually move the EPIC

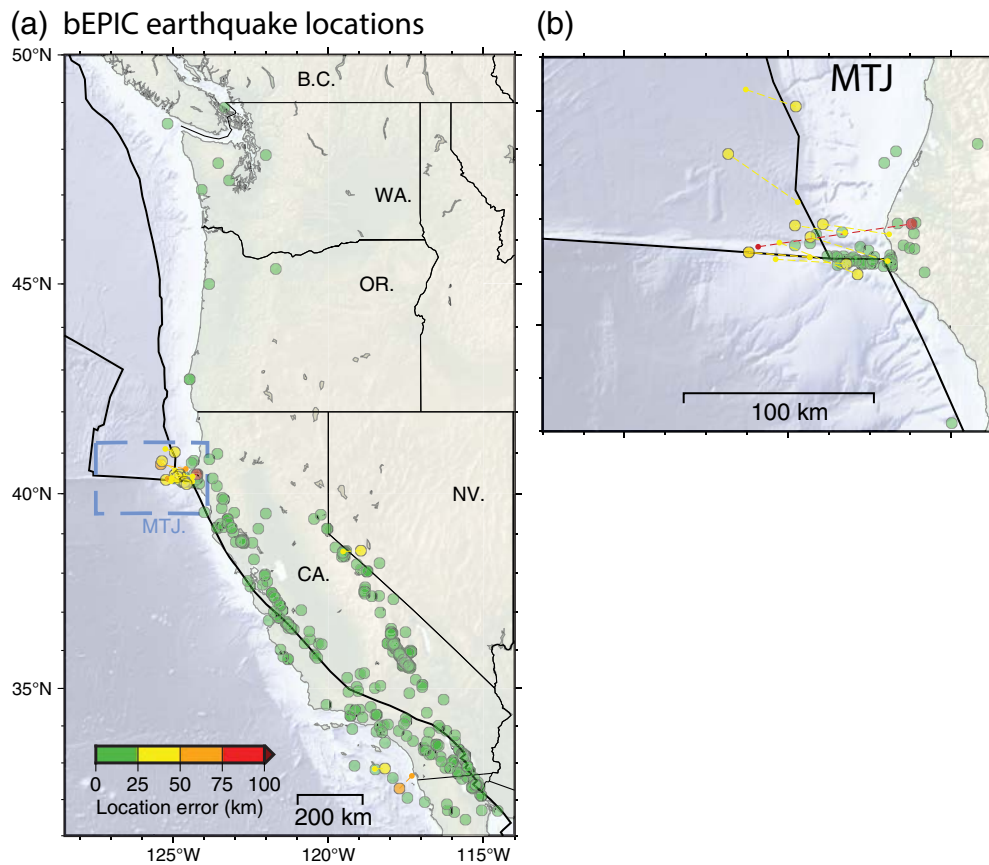


Figure 3. (a) Earthquake location accuracy across the ShakeAlert network when using the bEPIC modification for the same events as in Figure 1 across the western United States. Blue dashed line outlines our defined MTJ zone. (b) Zoomed in view of the MTJ region (blue dashed line region in panel a), and location error between bEPIC and USGS locations. The color version of this figure is available only in the electronic edition.

location farther and farther offshore. Once a location is more than 200 km from any station, EPIC discards that event. This can lead to missed events, for which potentially an alert may have been warranted, but is never issued. The use of prior seismicity in bEPIC limits this location-overshoot phenomenon.

Looking at our test catalog as a whole, the inclusion of prior seismicity greatly reduces the location error for earthquakes rupturing within the MTJ and performs as well as the non-Bayesian EPIC for earthquakes rupturing on land and within dense seismic networks. The test catalog includes earthquakes with varying arrangements of station locations and prior seismicity levels. Here, we draw attention to three earthquakes from the test catalog to use as points of discussion. First, an earthquake originating within the MTJ, on which our study is primarily focused. Second, an earthquake inland in southern California, where station coverage is excellent. Third, an earthquake on the coast of northern California that occurred outside the region of contemporary seismic activity. The spatial distribution of the likelihood, prior seismicity, and posteriors are mapped for each of the three examples in Figures 5–7.

tions creates a mapped posterior that has a lowered probability of the event being located far offshore and a slightly increased probability of the event occurring in the locations where both the likelihood and the prior seismicity have larger values (Fig. 5c). The bEPIC solution at the same final timestep is 8 km northeast of the USGS solution with a magnitude of 3.48. Many earthquakes from the test catalog that occurred in the MTJ have a very similar likelihood function and final posterior probability. This is in part because of the narrow azimuth of the coastal stations with respect to the source and limited availability of stations within a time frame that would allow for a timely alert to be issued. Although additional stations farther north and south of the rupture zone would allow for better azimuthal coverage, and thus a better overall control on the event location, waiting for these stations negates the use of the location for EEW. Instead, including in the prior information acts as a useful and automatic check on the location algorithm.

The second earthquake is an M 5.3 in southern California on 5 June 2021 (Fig. 6). The USGS location for this event is 33.14° N 115.635° W, which is just south of the Salton Sea.

The MTJ example earthquake is an M 3.5 earthquake that occurred on 23 March 2020 (Fig. 5). The USGS location for this event is 40.309° N 124.672° W, which is 33 km west of Petrolia, California. The final EPIC location for this event was 39.93° N 125.87° W, which is 110 km to the southwest of the USGS solution. The tendency of EPIC to locate events in this region to the southwest of their true locations is also noted for many other recent MTJ events and visible in Figure 1. This large error in location contributed to an overestimate of the magnitude by 0.7 magnitude units (M 4.2). The mapped-likelihood function (Fig. 5a) shows a clear band of near equally high likelihood extending far offshore whereas the prior seismicity (Fig. 5b) is generally concentrated along the MTJ, Mendocino fracture zone, and within more diffuse seismicity on land and within the Gorda plate. Combining the prior and the likelihood func-

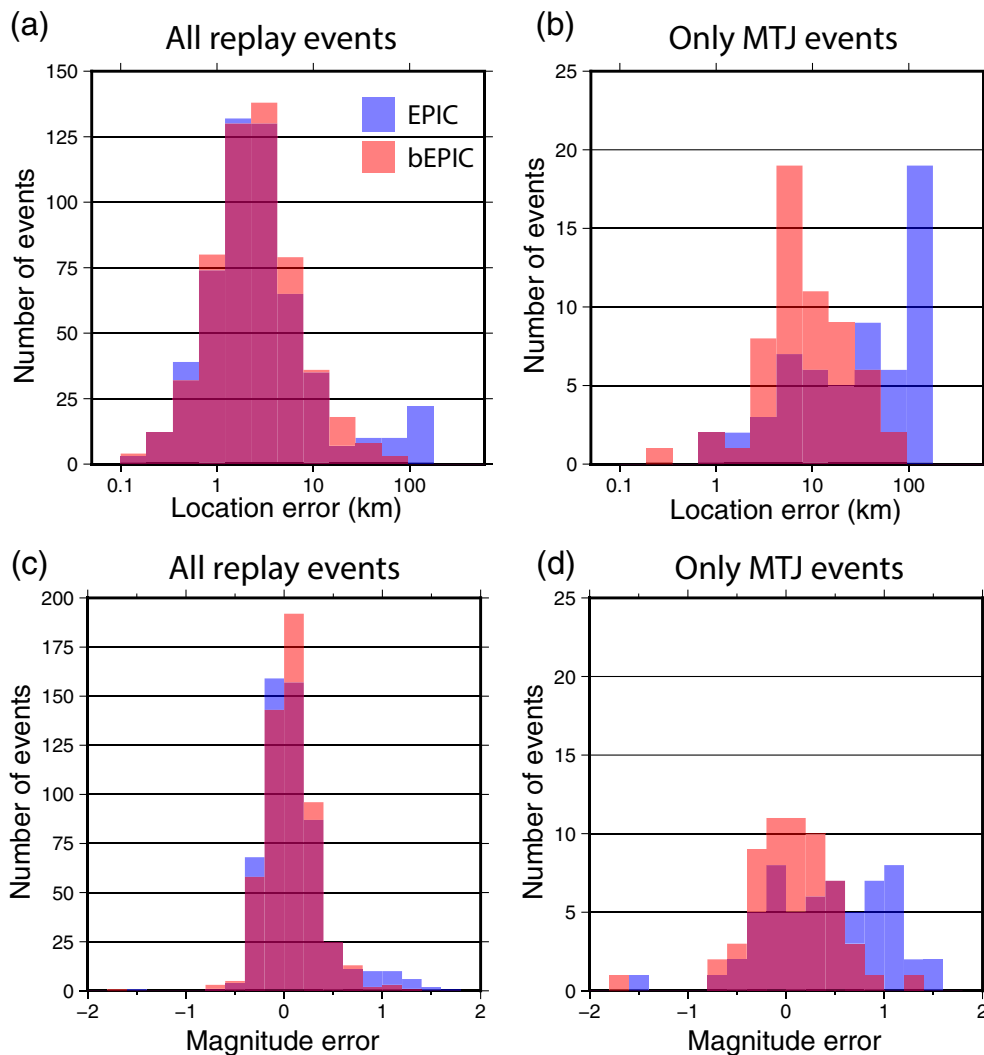


Figure 4. (a) Location error, in log scale, for EPIC (blue bars) and bEPIC (red bars) for all events in the replay catalog. (b) Location error, with same scaling as panel (a), for only earthquakes located within the MTJ region. (c) Magnitude error, linear scale, for all earthquakes in the replay catalog. (d) Magnitude error for earthquakes located within the MTJ region. The color version of this figure is available only in the electronic edition.

The earthquake occurred in an area with high seismic activity and is part of the Brawley seismic zone (Hauksson *et al.*, 2022). The region is also surrounded by a dense seismic network. Because this earthquake is surrounded by local stations, the location is well constrained with just the likelihood function alone. As shown in Figure 6a, the mapped likelihood creates a bullseye pattern with a single clear maximum point indicating where the earthquake should be located. For this event, EPIC's final solution was only 3.2 km from the final USGS solution. This is a typical average location error for earthquakes rupturing on land. The prior seismicity (Fig. 6b) is high over this region, particularly on the cross faults associated with the San Andreas and Imperial faults. Because there is only one small area with high likelihood associated with this grid search, the prior seismicity does not greatly affect the final posterior

solution (Fig. 6c). The final posterior solution (1.7 km from the USGS solution) is nearly identical to the solution provided by EPIC. Although EPIC sans the prior component produced a satisfactory solution on its own, we advocate for the inclusion of the prior anyway. Creating a system where EPIC decides when and whether to use a prior seismicity can potentially generate its own errors, particularly if a solution may appear well constrained with limited data, but then move or require a prior as more station data are incorporated into the solution with time.

The third test catalog example examines how the location algorithm will perform for regions where there is little recent documented seismicity. The example event is an *M* 3.5 earthquake that occurred on 1 October 2019, rupturing offshore and on the San Andreas fault, but south of the MTJ in a seismically quiet part of the transform system (Fig. 7). Although occurring offshore, the earthquake still had adequate station coverage to produce a satisfactory solution, as shown by the mapped likelihood in Figure 7a.

However, this section of the San Andreas has had only two other earthquakes of magnitude 3.0 or greater in the past 20 yr: 12 February 2012 an *M* 3.0 and 24 November 2002 an *M* 3.5. In contrast, there is a higher level of recent seismicity immediately to the east and on land within California's coastal range. Given the lack of seismicity near this earthquake and much larger values on the prior seismicity near the recent on land seismicity (Fig. 7b), could lead one to expect the final solution to be pulled inland and into the zone of high seismicity. However, this is not the case here. Figure 7c maps the posterior probability, which remains highest within the area that also has a high likelihood.

The reason that the bEPIC solution does not get pulled into the zone of high seismicity is related to the difference in scale between the nonnormalized values mapped in the likelihood

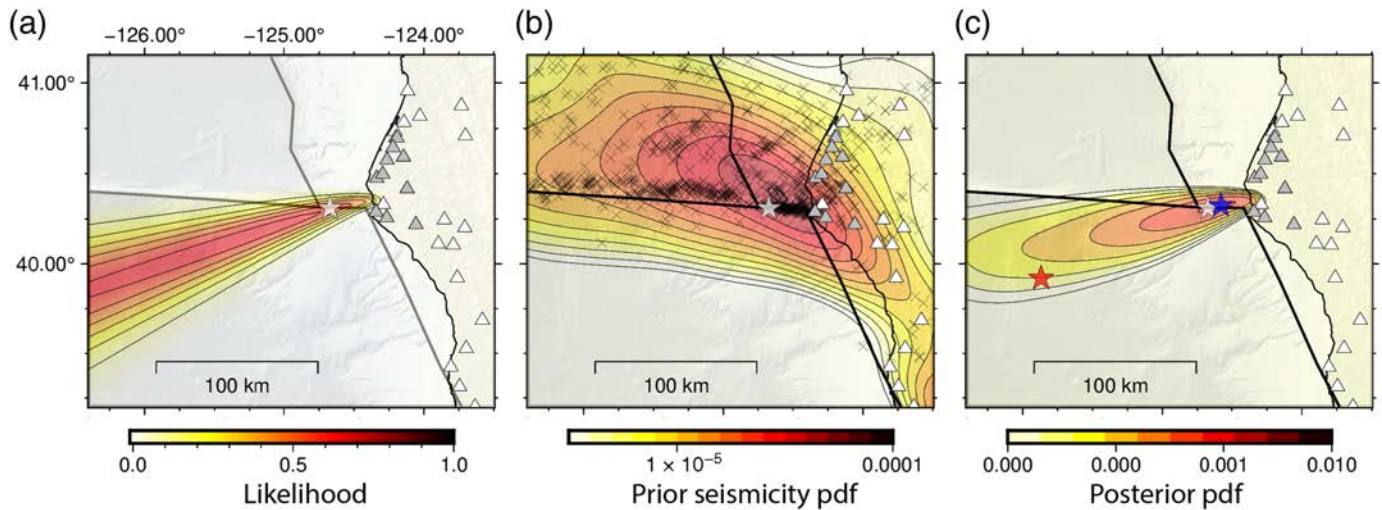


Figure 5. (a) Likelihood, (b) prior seismicity, and (c) posterior solutions for an offshore event example from the test catalog. Note the difference in color scales between the likelihood (a nonnormalized value) and the prior and posterior probability density functions (pdfs). Gray triangles mark stations in the ShakeAlert network that did trigger. White triangles mark stations that

did not trigger. The white star marks the USGS earthquake location for each event. The red star marks the EPIC event location and blue star marks the bEPIC event location. The color version of this figure is available only in the electronic edition.

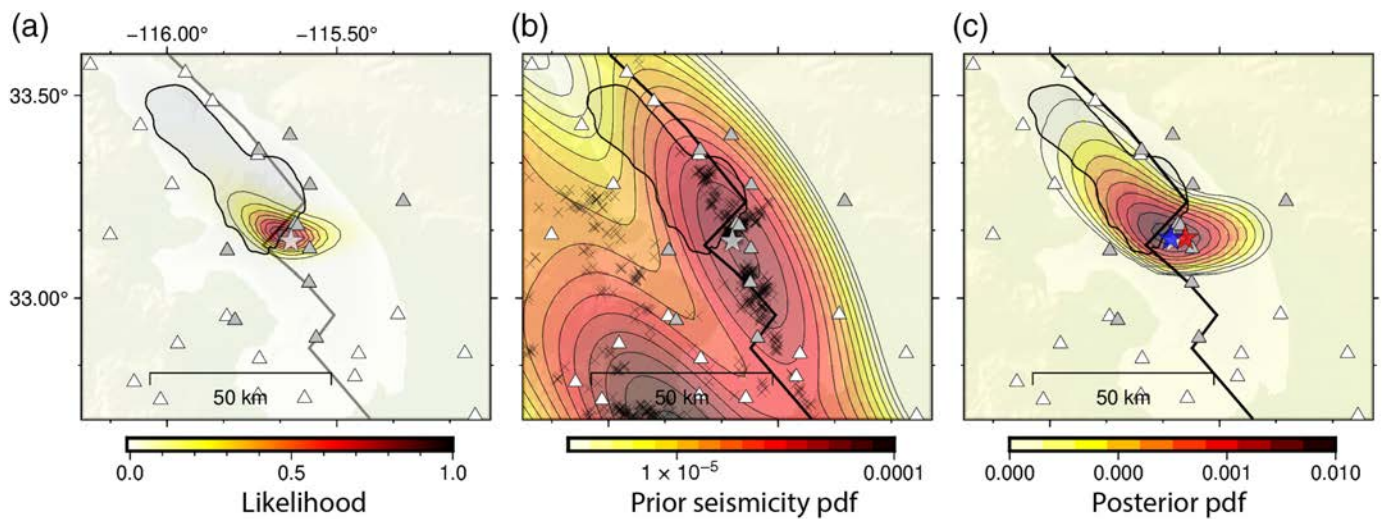


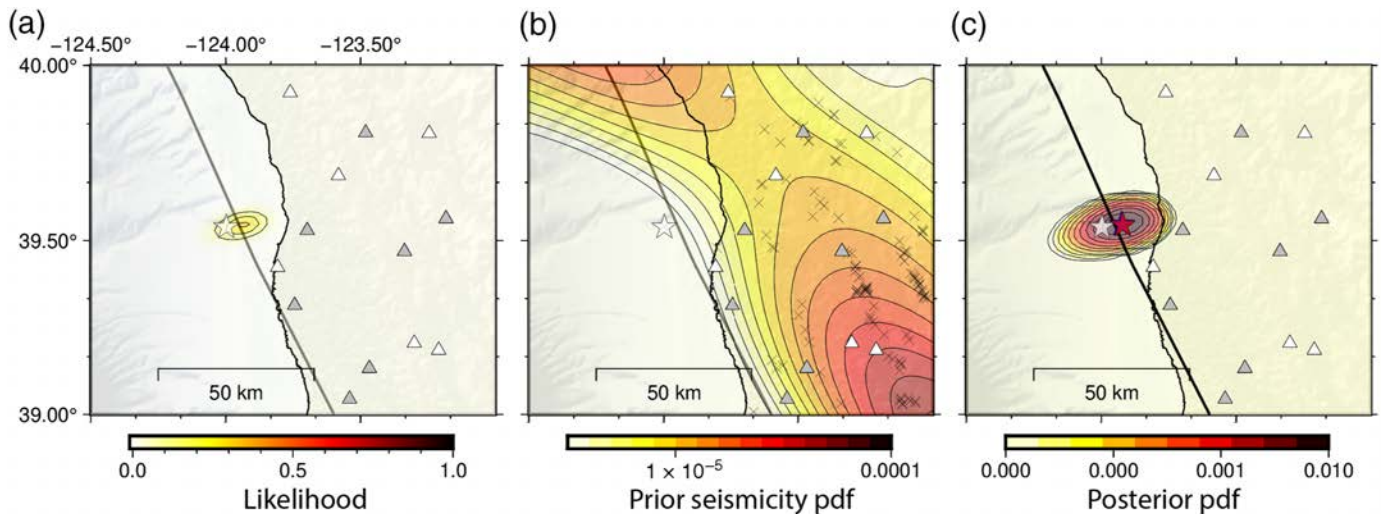
Figure 6. (a) Likelihood, (b) prior seismicity, and (c) posterior solutions for an in-network and onland event from the test catalog. Note the difference in color scales between the likelihood (a nonnormalized value) and the prior and posterior pdfs. Gray triangles mark stations in the ShakeAlert network

that did trigger. White triangles mark stations that did not trigger. The white star marks the USGS earthquake location for each event. The red star marks the EPIC event location and blue star marks the bEPIC event location. The color version of this figure is available only in the electronic edition.

function and the normalized values in the prior. The effect that this has is that although some areas with past seismicity have higher relative values than other areas with little to no seismicity, the overall magnitude of the prior is small. What this means is that when applied to the likelihood, the greatest modification only occurs among grid nodes with high and similar values where the small weight of the prior can have an effect. In the example shown here, the relatively high seismicity in the

prior at locations on land has extremely little effect as it is being applied to an area with extremely low (close to zero) likelihood.

The second reason that the bEPIC solution does not get pulled inland to the high seismicity zone in Figure 7 is that the prior seismicity has a floor in place where every point in the prior must be greater than zero. This stops the prior from having too much power in limiting new earthquakes in low seismic zones. Although there is no ceiling value set



for the prior, the requirement that the prior be normalized sets a limitation on how large prior values can be.

How EPIC and bEPIC would behave for potential future events is harder to capture without real data, but still worth examining. All recent seismicity with EPIC solutions that appear in the test catalog within the MTJ are concentrated within about 50 km of the coastline. However, there have been large past earthquakes that ruptured farther offshore on the Mendocino fracture zone but still inside the ShakeAlert reporting zone. This includes larger earthquakes such as the 1994 *M* 7.1 Ferndale earthquake as well as a handful of *M* 6.0+ earthquakes rupturing between roughly 127°–125° W. Although this region does have past seismicity that is incorporated into local priors (Fig. 2), there are clearly more recent earthquakes closer to the coastline in the database than events farther west. Would the larger density of events here have a relative pulling effect on far offshore earthquakes, moving their solutions farther east? We tested a preliminary set of synthetic events located along the Mendocino fracture zone. Early in the location search, when only 3–4 stations are available and the mapped likelihood is high over a broad region, there is an initial pulling effect toward the eastern higher seismicity region in the prior (Fig. S03). However, as even a small amount of additional station data is made available, the bEPIC solution migrates westward closer to the “true” location. In contrast, the solutions using a non-Bayesian EPIC struggle with the initial high likelihood, leading to locations that get placed hundreds of kilometers offshore. The prior seismicity, while creating an initial pulling effect, limits this location-overshoot common in EPIC. When repeating our synthetic offshore simulations, bEPIC produced on average lower location errors than EPIC.

An improvement in earthquake location affects the rest of the ShakeAlert system. Because the station-epicenter distance is used in the magnitude calculation, a more accurate location leads to a more accurate magnitude. In addition, the magnitude and location together are used to determine the extent and intensity of ground shaking. With large enough location

Figure 7. (a) Likelihood, (b) prior seismicity, and (c) posterior solutions for an offshore, non-MTJ event in an area of low past seismicity from the test catalog. Note the difference in color scales between the likelihood (a nonnormalized value) and the prior and posterior pdfs. Gray triangles mark stations in the ShakeAlert network that did trigger. White triangles mark stations that did not trigger. The white star marks the USGS earthquake location for each event. The red star marks the EPIC event location and blue star marks the bEPIC event location. Note that in this example, the red and blue stars are collocated. The color version of this figure is available only in the electronic edition.

errors, and in kind, an overestimate of the magnitude, false alerts could be created. Too many false alerts have the potential to lower community confidence in alerts received for future events. Furthermore, accurately locating and characterizing offshore earthquakes is important because of the additional hazards associated with offshore earthquakes, such as tsunamis. All cases handled in our test data set are too small to generate a tsunami; however, the MTJ has the potential to generate larger earthquakes and tsunamis. Such was the case with the 1992 Mendocino earthquake, an interslab thrust event on the subducting Gorda microplate. This event generated a small tsunami that was identified at regionally placed tide gauges as well as at a tide gauge in Hawaii (Oppenheimer *et al.*, 1993; González *et al.*, 1995). Accurately identifying the location and size of events like this, as well as limiting false alarms, is highly important.

An area of needed improvement within EPIC that can affect bEPIC solutions is the inclusion of poor trigger picks from the associator algorithm. Both EPIC and bEPIC receive station trigger data from a separate algorithm that both collects triggered stations and associates nearby triggers. Although many checks are in place to limit triggers from teleseismic events, which have in the past generated numerous false alarms (Chung *et al.*, 2019) as well as quality control metrics to weed out triggers from calibration pulses, boxcar shaped features, and errant noise, some poor data still can be sent to the

location algorithms. When this happens, neither EPIC nor bEPIC have the current functionality to identify and throw out poor picks. This means that a single station pick with a trigger time error of multiple seconds can cause an overall high misfit across all grid-search locations, creating a near-zero likelihood everywhere. This in turn means that the added prior information does little to correct the final location solution. Although uncommon, this poor data problem will generate poor location solutions. Future additions to EPIC should include additional metrics in the location algorithm to allow for the ability to reject poor station trigger information.

An interesting future direction for EEW is the incorporation of real-time offshore data into location algorithms by means of fiber optic distributed acoustic sensing (DAS) arrays (Farghal *et al.*, 2022). The incorporation of offshore data can, if positioned close to the source, provide initial trigger information much sooner than systems that wait for onland stations. Using both offshore DAS array data and traditional onshore seismometers in EEW also has the benefit of greatly reducing the level of initial uncertainty when solving for the locations of offshore earthquakes, as discussed at length here. If made available, data from currently existing transoceanic fiber optic cables would allow for greater sensitivity in detecting smaller offshore events as well as other geohazards such as sending the propagation of tsunami waves in the open ocean (Salaree *et al.*, 2022).

CONCLUSION

EEW through programs like ShakeAlert build the tools communities need to take proactive steps to mitigate risk in the seconds prior to damaging shaking. Although ShakeAlert can provide timely alerts for many of the earthquakes it encounters, the lack of seismic network coverage extending offshore makes accurate event locations for earthquakes nucleating near the MTJ challenging. In this study, we modified ShakeAlert's point-source earthquake characterization code, EPIC, to include prior information about the seismicity of a region as an additional component affecting the grid-search algorithm. Although we include this layer to specifically target poorly located events offshore, we also apply and assess the performance of the modified point-source algorithm for events across the entire ShakeAlert reporting area. Our improved code, bEPIC, produced solutions with average location errors in the MTJ of 14 km. This is an average reduction in error for events in this region of 44 km. The inclusion of the prior seismicity layer does not negatively affect network solutions in which EPIC already excelled. In these cases, the prior information has little impact on the overall posterior result, and no new sources of error are introduced. With better location estimates for offshore events, the bEPIC calculated magnitude also improves, limiting the potential of a false alert being issued for a large, offshore earthquake.

The improvement to the solved earthquake location by bEPIC limits the likelihood of false alerts for communities

along the northern California coast. This is a region that is particularly sensitive to false alerts, as large offshore earthquakes in this region also carry the possibility of other seismically related hazards such as tsunami.

DATA AND RESOURCES

bEPIC replay information including station location and trigger times for earthquakes used in this study, as well as catalog information for all test events, are available at doi: [10.5281/zenodo.6929789](https://doi.org/10.5281/zenodo.6929789). The prior seismicity information was created by querying the Advanced National Seismic System (ANSS) Comprehensive Earthquake Catalog (ComCat): <https://earthquake.usgs.gov/data/comcat/> (last accessed July 2022). Some plots were made using PyGMT (<https://zenodo.org/record/6702566>, last accessed December, 2022; Wessel *et al.* 2019; Uieda *et al.*, 2022). The source code for bEPIC is available through GitHub: <https://github.com/amy-l-williamson/bEPIC> (last accessed December 2022). The supplemental material includes three supplemental figures. Figure S1 shows an example of a poor *P*-wave trigger pick. Figure S2 shows the relationship between earthquake location error and magnitude error for the replay catalog. Figure S3 shows the spread in EPIC and bEPIC locations a potential synthetic far offshore earthquake as referenced in the Discussion section.

DECLARATION OF COMPETING INTERESTS

The authors declare that there are no conflicts of interest recorded.

ACKNOWLEDGMENTS

The authors would like to thank Anthony Lomax and one anonymous reviewer for their helpful comments. The authors would also like to thank Tom Heaton for his constructive comments. This research was supported by the U.S. Geological Survey (USGS) under Grant Agreement Number G21AC10525 awarded to University of California, Berkeley.

REFERENCES

- Allen, R. M., and D. Melgar (2019). Earthquake early warning: Advances, scientific challenges, and societal needs, *Annu. Rev. Earth Planet. Sci.* **47**, 361–388.
- Allen, R. M., H. Brown, M. Hellweg, O. Khainovski, P. Lombard, and D. Neuhauser (2009). Real-time earthquake detection and hazard assessment by ElarmS across California, *Geophys. Res. Lett.* **36**, no. 5, doi: [10.1029/2008GL036766](https://doi.org/10.1029/2008GL036766).
- Böse, M., C. Felizardo, and T. H. 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., D. E. Smith, C. Felizardo, M. A. Meier, T. H. Heaton, and J. F. Clinton (2018). FinDer v. 2: Improved real-time ground-motion predictions for M2–M9 with seismic finite-source characterization, *Geophys. J. Int.* **212**, no. 1, 725–742.
- Chung, A. I., I. Henson, and R. M. Allen (2019). Optimizing earthquake early warning performance: ElarmS-3, *Seismol. Res. Lett.* **90**, no. 2A, 727–743.
- Chung, A. I., M. A. Meier, J. Andrews, M. Böse, B. W. Crowell, J. J. McGuire, and D. E. Smith (2020). ShakeAlert earthquake early warning system performance during the 2019 Ridgecrest earthquake sequence, *Bull. Seismol. Soc. Am.* **110**, no. 4, 1904–1923.

- Farghal, N. S., J. K. Saunders, and G. A. Parker (2022). The potential of using fiber optic distributed acoustic sensing (DAS) in earthquake early warning applications, *Bull. Seismol. Soc. Am.* **112**, no. 3, 1416–1435.
- González, F. I., K. Satake, E. F. Boss, and H. O. Mofjeld (1995). Edge wave and non-trapped modes of the 25 April 1992 Cape Mendocino tsunami, in *Tsunamis: 1992–1994*, Birkhäuser, Basel, Switzerland, 409–426.
- Hauksson, E., J. M. Stock, and A. L. Husker (2022). Seismicity in a weak crust: The transtensional tectonics of the Brawley seismic zone section of the Pacific-North America plate boundary in Southern California, USA, *Geophys. J. Int.* **231**, no. 1, 717–735.
- Kohler, M. D., D. E. Smith, J. Andrews, A. I. Chung, R. Hartog, I. Henson, D. D. Given, R. de Groot, and S. Guiwits (2020). Earthquake early warning ShakeAlert 2.0: Public rollout, *Seismol. Res. Lett.* **91**, no. 3, 1763–1775.
- Kuyuk, H. S., and R. M. Allen (2013). A global approach to provide magnitude estimates for earthquake early warning alerts, *Geophys. Res. Lett.* **40**, no. 24, 6329–6333.
- Lomax, A., A. Michelini, A. Curtis, and R. A. Meyers (2009). Earthquake location, direct, global-search methods, *Encycl. Complex. Syst. Sci.* **5**, 2449–2473.
- Minson, S. E., M. Simons, and J. L. Beck (2013). Bayesian inversion for finite fault earthquake source models I—Theory and algorithm, *Geophys. J. Int.* **194**, no. 3, 1701–1726.
- Minson, S. E., M. Simons, J. L. Beck, F. Ortega, J. Jiang, S. E. Owen, A. W. Moore, A. Inbal, and A. Sladen (2014). Bayesian inversion for finite fault earthquake source models—II: The 2011 great Tohoku-oki, Japan earthquake, *Geophys. J. Int.* **198**, no. 2, 922–940.
- Nakamura, Y., and J. Saita (2007). UrEDAS, the earthquake warning system: Today and tomorrow, in *Earthquake Early Warning Systems*, Springer, Berlin, Heidelberg, 249–281.
- Oppenheimer, D., J. Eaton, A. Jayko, M. Lisowski, G. Marshall, M. Murray, R. Simpson, R. Stein, G. Beroza, M. Magee, *et al.* (1993). The Cape Mendocino, California, earthquakes of April 1992: Subduction at the triple junction, *Science* **261**, no. 5120, 433–438.
- Rollins, J. C., and R. S. Stein (2010). Coulomb stress interactions among $M \geq 5.9$ earthquakes in the Gorda deformation zone and on the Mendocino Fault Zone, Cascadia subduction zone, and northern San Andreas fault, *J. Geophys. Res.* **115**, no. B12, doi: [10.1029/2009JB007117](https://doi.org/10.1029/2009JB007117).
- Salaree, A., B. M. Howe, Y. Huang, S. A. Weinstein, and A. E. Sakya (2022). A numerical study of SMART Cables potential in marine hazard early warning for the Sumatra and Java regions, *Pure Appl. Geophys.* doi: [10.1007/s00024-022-03004-0](https://doi.org/10.1007/s00024-022-03004-0).
- Scott, D. W. (2015). *Multivariate Density Estimation: Theory, Practice, and Visualization*, John Wiley & Sons, Hoboken, New Jersey.
- Serdar Kuyuk, H., R. M. Allen, H. Brown, M. Hellweg, I. Henson, and D. Neuhauser (2014). Designing a network-based earthquake early warning algorithm for California: ElarmS-2, *Bull. Seismol. Soc. Am.* **104**, no. 1, 162–173.
- Strauss, J. A., and R. M. Allen (2016). Benefits and costs of earthquake early warning, *Seismol. Res. Lett.* **87**, no. 3, 765–772.
- Strauss, J. A., Q. Kong, S. Pothan, S. Thompson, R. F. Mejia, S. Allen, S. Patel, and R. M. Allen (2020). MyShake Citizen Seismologists help launch dual-use seismic network in California, *Front. Commun.* doi: [10.3389/fcomm.2020.00032](https://doi.org/10.3389/fcomm.2020.00032).
- Thompson, M., J. R. Hartog, and E. A. Wirth (2022). Effect of fixing earthquake depth in ShakeAlert algorithms on performance for intraslab earthquakes, *Seismol. Soc. Am.* **93**, no. 1, 277–287.
- Uieda, L., D. Tian, W. J. Leong, M. Jones, W. Schlitzer, M. Grund, and L. Toney (2022). PyGMT: A Python interface for the Generic Mapping Tools (v0. 7.0), Computer software, *Zenodo* doi: [10.5281/zenodo.6349217](https://doi.org/10.5281/zenodo.6349217).
- Wessel, P., J. F. Luis, L. Uieda, R. Scharroo, F. Wobbe, W. H. Smith, and D. Tian (2019). The generic mapping tools version 6, *Geochem. Geophys. Geosys.* **20**, no. 11, 5556–5564.
- Wurman, G., R. M. Allen, and P. Lombard (2007). Toward earthquake early warning in northern California, *J. Geophys. Res.* **112**, no. B8, doi: [10.1029/2006JB004830](https://doi.org/10.1029/2006JB004830).
- Yin, L., J. Andrews, and T. Heaton (2018). Rapid earthquake discrimination for earthquake early warning: A Bayesian probabilistic approach using three-component single-station waveforms and seismicity forecast, *Bull. Seismol. Soc. Am.* **108**, no. 4, 2054–2067.
- Zhu, W., I. W. McBrearty, S. M. Mousavi, W. L. Ellsworth, and G. C. Beroza (2022). Earthquake phase association using a Bayesian Gaussian mixture model, *J. Geophys. Res.* **127**, no. 5, e2021JB023249, doi: [10.1029/2021JB023249](https://doi.org/10.1029/2021JB023249).

Manuscript received 28 July 2022
Published online 27 January 2023