JGR Solid Earth

RESEARCH ARTICLE

10.1029/2020JB021053

Key Points:

- Propagation of Local Undamped Motion (PLUM) detects offshore events that produce at least moderate onshore ground motions, including an event problematic for other algorithms
- PLUM's detection times are as timely, and sometimes faster, than other earthquake early warning detection times
- PLUM correctly avoids erroneous detections for all teleseismic, calibration pulses, and anomalous signals in the test suite

Supporting Information:

Supporting Information may be found in the online version of this article.

Correspondence to:

D. Kilb,
dkilb@ucsd.edu

Citation:

Kilb, D., Bunn, J. J., Saunders, J. K., Cochran, E. S., Minson, S. E., Baltay, A., et al. (2021). The PLUM earthquake early warning algorithm: A retrospective case study of West Coast, USA, data. *Journal of Geophysical Research: Solid Earth*, 126, e2020JB021053. <https://doi.org/10.1029/2020JB021053>

Received 30 OCT 2020

Accepted 4 JUN 2021

The PLUM Earthquake Early Warning Algorithm: A Retrospective Case Study of West Coast, USA, Data

D. Kilb¹ , J. J. Bunn² , J. K. Saunders³ , E. S. Cochran⁴ , S. E. Minson³ , A. Baltay³ , C. T. O'Rourke⁵ , M. Hoshiba⁶ , and Y. Kodera⁶

¹Scripps Institution of Oceanography, UCSD, San Diego, CA, USA, ²California Institute of Technology, Pasadena, CA, USA, ³U.S. Geological Survey, Moffett Field, CA, USA, ⁴U.S. Geological Survey, Pasadena, CA, USA, ⁵U.S. Geological Survey, Seattle, WA, USA, ⁶Meteorological Research Institute, JMA, Japan

Abstract The PLUM (Propagation of Local Undamped Motion) earthquake early warning (EEW) algorithm differs from typical source-based EEW algorithms as it predicts shaking directly from observed shaking without first deriving earthquake source information (e.g., magnitude and epicenter). Here, we determine optimal PLUM event detection thresholds for U.S. West Coast earthquakes using two data sets: 558 M3.5+ earthquakes (California, Oregon, Washington; 2012–2017) and the ShakeAlert test suite of historic and problematic signals (1999–2015). PLUM computes Modified Mercalli Intensity (I_{MMI}) using velocity and acceleration data, leveraging co-located sensors to avoid problematic signals. An event detection is issued when the observed I_{MMI} exceeds a given threshold(s). We find a two-station detection method using I_{MMI} trigger thresholds of 4.0 and 3.0 for the first and second stations, respectively, is optimal for detecting M4.5+ earthquakes. PLUM detected 79 events in the 2012–2017 data set, reporting (not including telemetry or alert dissemination) detection times on par, and sometimes faster than current EEW methods (mean 8 s; median 6 s). As expected, detection times were slower for the older 1999–2015 earthquakes ($N = 21$; mean 11 s; median 6 s) when station coverage was sparser. Of the 31 PLUM detected M5+ events (10 2012–2017; 21 1999–2015), theoretically 20 (~65%) could provide timely warnings. PLUM issued no false detections and avoided issuing detections for all calibration/anomalous signals, regional and teleseismic events. We conclude PLUM can successfully identify I_{MMI} 4+ shaking from local earthquakes and could complement and enhance EEW in the U.S.

Plain Language Summary Earthquake early warning (EEW) detection schemes require (1) ample seismic information to identify where large ground motions are underway; (2) determining if these ground motions are significant enough to issue a detection; and (3) detecting large ground motions in a timely fashion. Some EEW methods estimate earthquake source parameters like magnitude and location and then input those parameters into a ground-motion prediction equation, while other methods use observations of the ground motions to directly forecast shaking. We explore the latter approach using the PLUM (Propagation of Local Undamped Motion) method to detect earthquakes that produce shaking above a target value. In this work, we test PLUM's ability to detect earthquakes using two data sets: 558 earthquakes magnitude 3.5 and above from California, Oregon, and Washington (2012–2017) and a test suite of historic and problematic signals (1999–2015) curated by ShakeAlert. We find a two-station detection method is preferred over a one-station method as two-stations can greatly minimize false detections. The PLUM method is also 100% successful at avoiding non-earthquake anomalous signals and can successfully differentiate ground shaking from local and distant earthquakes. We conclude that PLUM may be a promising candidate for integration into the U.S. EEW system.

1. Introduction

Around the globe, scientists are now implementing earthquake early warning (EEW) systems, leveraging the ongoing technological improvements in data recording, archiving, and processing, coupled with the data dissemination speeds that can push scientific information to the public in near real-time (R. M. Allen et al., 2019; Cochran et al., 2019; Chung et al., 2019; Hoshiba et al., 2008; Hsiao et al., 2009; Ruhl et al., 2019; Kodera et al., 2020; Strauss & Allen, 2016). The goal of EEW is to issue alerts that provide advanced warning of potentially damaging ground shaking related to earthquakes so that mitigative actions can be taken (e.g., Allen, Brown, et al., 2009; Allen, Gasparini, et al., 2009; R. M. Allen & Kanamori, 2003; R. M. Allen &

Melgar, 2019; R. M. Allen et al., 2018; Chung et al., 2019; Cochran & Husker, 2019; Cochran et al., 2019; Given et al., 2014; Heaton, 1985; Kodera et al., 2020; Kohler et al., 2018, 2020; Meier, 2017; Minson et al., 2018; Meier et al., 2020; Wald, 2020). In 2018, the ShakeAlert® Earthquake Early Warning system (ShakeAlert System, ShakeAlert) began the rollout of public EEW alerts for the western states of the U.S. (California, Oregon, and Washington). The goal of the ShakeAlert System is to issue advance warnings to people who may experience potentially damaging shaking so they can take protective actions (Given et al., 2018; Hartog et al., 2016; Kohler et al., 2018, 2020; Thakoor et al., 2019). Currently, the operational ShakeAlert System includes two independent algorithms (EPIC and FinDer) and plans are in place to investigate if integration of additional algorithms would be beneficial (Böse et al., 2018; Given et al., 2018; Kohler et al., 2020). EEW alerts via ShakeAlert (called “ShakeAlert Messages”) are typically issued when an earthquake is estimated to be larger than M5.0, since ground motions for smaller earthquakes are generally lower amplitude and more spatially compact, making timely warning difficult (Wald, 2020).

The amount of early warning can vary, depending on the earthquake location and size, the seismic network configuration in relation to the earthquake, data telemetry latencies, the spatial extent of the alerting region, the level of shaking for which an alert is issued, alert delivery latencies, and the relative importance of necessary alerts and avoidance of false alarms (R. M. Allen & Melgar, 2019; Kong et al., 2019; Kohler et al., 2020; Minson et al., 2018, 2019, 2020). Expected warning times are not hours or minutes, as would be the case for early warnings for other natural hazards like tornadoes or flash floods, but instead are on the order of seconds (e.g., Minson et al., 2018).

In this work, we test the Propagation of Local Undamped Motion (PLUM) EEW algorithm, which was originally designed for use in Japan (Hoshiba, 2013; Kodera et al., 2016, 2020). PLUM is designed to detect when ground motions exceed a threshold, and can issue alerts directly from these observations without first deriving a magnitude or location for the earthquake. This design was to improve detection timeliness, avoid alerting errors caused by significant earthquake source mischaracterization, and function more robustly during aftershock sequences (Kodera et al., 2016). Because PLUM is not a source-based algorithm, it does not need to untangle source information from multiple concurrent events (Hoshiba et al., 2011). The simplicity of PLUM directly using observed ground motion to forward-predict nearby (10s of km) ground motions has been able to reduce uncertainties in EEW alerts (Kodera et al., 2020), and also has the potential to appropriately issue detections for high stress drop earthquakes for which traditional EEW algorithms would fail to issue warnings because the shaking is greater than expected for the earthquake magnitude (Minson et al., 2019).

Here, we use a modified version of the PLUM algorithm first tested with Southern California Seismic Network data by Cochran et al. (2019). We have improved the algorithm to select more reliable data at co-located sensors and have reevaluated triggering thresholds to reliably detect M5+ earthquakes while minimizing detections of smaller events for which timely detections are often not possible (Kodera, 2019). Our goal is to detect all M5+ earthquakes that are expected to generate Modified Mercalli Intensity (I_{MMI}) of 4 or higher, which is consistent with the targets of ShakeAlert (Cochran et al., 2018; Given et al., 2014, 2018; Kohler et al., 2018, 2020). In the U.S. Geological Survey (USGS) “Did You Feel It” program, a citizen scientist project that collects information about the public’s perception of shaking from earthquakes, the shaking level for $I_{MMI} = 4$ is categorized as “light shaking” and generally felt by everyone, however, $I_{MMI} = 4$ is described by those who experience it as “strong shaking” despite not causing damage (Dengler & Dewey, 1998; Dewey et al., 1995; Stover & Coffman, 1993; Wald et al., 1999). PLUM is configurable, so if the ShakeAlert standards change, PLUM can be adapted accordingly.

We apply different Instances of the PLUM algorithm to two retrospective data sets (i.e., not real-time) with the aim of finding optimal I_{MMI} trigger thresholds for detecting M5+ earthquakes within the West Coast states of California, Oregon, and Washington (Figure 1; Table S1 provides region boundaries) while minimizing detections of smaller magnitude events. However, smaller events ($M < 5.0$) may generate ground motions of interest ($I_{MMI} = 4+$) potentially warranting alerts to populations near the epicenter, but these smaller quakes have a smaller spatial footprint making it very difficult to provide timely warnings. Because of this, we do not prioritize detection of these smaller events. We also assess PLUM’s ability to correctly avoid detecting non-earthquake signals that have been problematic for the ShakeAlert System in the past (Cochran et al., 2018).

Figure 1. Map of WC data within the ShakeAlert Pacific Northwest (PNW) and California regions (dashed lines; polygon nodes defining these regions listed in Table S1). Of the 558 earthquakes (2012–2017), 504 have magnitudes below 4.5 (gray circles). The other events are color-coded by missed M4.5–M5.0 ($N = 23$; orange stars), missed M5+ ($N = 10$; brown stars), detected $M < 4.5$ ($N = 58$; light purple squares), detected M4.5–M5 ($N = 11$; light purple stars) and detected M5+ earthquakes ($N = 10$; larger purple stars). Two M5+ events south of the Salton Sea are at similar locations, making it appear as if there are nine large purple stars instead of 10. The 10 missed M5+ events are not theoretically detectable, and no PNW events are theoretically detectable. Notable earthquakes discussed in the text are labeled.

2. Terminology

We begin by defining key terms used in this manuscript (Table 1). The term *trigger* is used to indicate that ground motion at a given station exceeded a specified intensity threshold, resulting in a *station trigger*. Importantly, the terms *PLUM event detection* and *PLUM alert* cannot be used interchangeably. A *PLUM event detection* is issued when a predefined number of stations issue a trigger. The *PLUM detection time* is defined as the elapsed time between the origin time of the earthquake and the issuance of the PLUM event detection and corresponds to the earliest possible time PLUM could issue an EEW alert. *PLUM alert* regions are subsequently determined using the *PLUM event detection* information, and other predefined parameters (e.g., Minson et al., 2020). As the focus of this paper is on PLUM event detections, we leave the optimization of PLUM alert regions in the U.S. for future analysis. In these retrospective tests, a *match* is assigned when

Table 1

Terms Used to Define the Status of an Event or Data Block and Other Associated Terminology Used in This Work

Term	Description
Alert region	PLUM alert regions are derived from the PLUM detections. We do not investigate PLUM alert regions in this work.
Alert region method	Numerous alerting methods exist, including station-based, grid-based and spatial configurations of different sizes (Minson et al., 2020). We do not implement alert region methods in this work.
Event detection	An event detection is issued if PLUM detects ground motions above a specified triggering threshold(s). A single threshold is used for a one-station detection method; two thresholds are used for a two-station detection method.
False detection	A false detection is issued when no earthquake is associated with a PLUM detection.
Instance	An Instance number is used to track different threshold configuration tests, see Table 2.
Match	A match is assigned for an earthquake when associated with a PLUM detection.
Missed	A missed event is assigned when an earthquake has no associated PLUM event detection.
Station trigger	A station will trigger when the I_{MMI} recorded at that station exceeds a threshold value.
Abbreviation: PLUM, Propagation of Local Undamped Motion.	

the PLUM event detection can be associated with a specific earthquake (described in more detail below). We label an event detection as *false* if there is no corresponding earthquake within the Advanced National Seismic System (ANSS) earthquake catalog.

A modified version of the original PLUM method developed for EEW in Japan was first applied to southern California data in a study that explored the effectiveness of one-station and two-station detection strategies for event detection (Cochran et al., 2019). The single-station strategy issues an event detection when the ground motions on a single station exceed a specified I_{MMI} threshold. In the two-station strategy, a trigger is issued when the ground motion at any station exceeds a primary threshold and is confirmed by shaking at a neighboring station exceeding a secondary threshold. Cochran et al. (2019) tested six different implementations of PLUM (“Instances”) with varying primary and secondary triggering thresholds. Cochran et al. (2019) found that using two stations with primary and secondary triggering thresholds of I_{MMI} 4.0 and 2.5, respectively, was preferred to reduce false detections and maximize event detection for southern California data. In this work, the focus is on PLUM event detections and PLUM detection times, computation and evaluation of PLUM alert regions is not performed, as this is a secondary task left for future studies.

3. Reassessing PLUM Implementation in the U.S.

3.1. The PLUM Algorithm

PLUM continually observes ground-motion amplitudes and initiates a PLUM event detection when the algorithm detects ground motions above specified threshold(s). Here, we evaluated the same six instance configurations as in Cochran et al. (2019), which tested southern California data, to determine which configuration is optimal for the west coast of the United States.

We expect event detection to be more challenging in Oregon and Washington where station spacing is sparser than in southern California. The goal is to minimize false detections and minimize the detection of smaller magnitude events ($M < 4.5$). Detections of $M < 4.5$ earthquakes, even if they produce shaking $> I_{MMI}$ 4.0, are less likely to provide timely and actionable alerts because for these smaller earthquakes the spatial extent of high-amplitude shaking is small. Our trigger threshold selection also considers the time required to make a detection, that is, the time from when an earthquake starts to the first issuance of an event detection. We use the PLUM algorithm as implemented in Cochran et al. (2019), but with two modifications to reduce corruption by seismic noise, described below.

3.2. Modification #1

The first modification is to mitigate against bad data such as clipped waveforms and data with a high noise floor. To accomplish this, when available, we compare data from co-located sensors. We begin by computing separate intensity estimates from both the weak-motion velocity sensor and the strong-motion accelerometer sensor. The peak ground motions we use are the maximum absolute values of the acceleration and velocity over one second of data. These data have been demeaned and bandpass filtered at 0.5–25 Hz to reduce the very low frequencies commonly affected by baseline drifts (which can still be present after demeaning) and to remove high-frequency noise. To estimate peak acceleration from the velocity sensor data, we use differentiated velocity data, and to estimate peak velocity data from the acceleration sensor, we use integrated acceleration data. For a given sensor, this nets a time series of measurements taken each second for velocity and acceleration:

$$\text{MAXVEL}_{i,j} = \max \left(\sqrt{N(t)_{i,j}^2 + E(t)_{i,j}^2 + Z(t)_{i,j}^2} \right) \quad (1a)$$

$$\text{MAXACC}_{i,j} = \max \left(\sqrt{N(t)_{i,j}^2 + E(t)_{i,j}^2 + Z(t)_{i,j}^2} \right) \quad (1b)$$

where MAXVEL and MAXACC represent the maximum values of the velocity and the acceleration at a sensor, respectively, i is the time index, j is the station index and $N(t)_{i,j}$, $E(t)_{i,j}$ and $Z(t)_{i,j}$ represent a 1-second-duration seismogram snippet for the three data components (N for north-south, E for east-west and Z for vertical) for velocity (1a) and acceleration (1b). For each station with a co-located acceleration- and velocity-sensor, we use the ground-motion-to-intensity conversion equations of Worden et al. (2012; their equations 3 and 9) to compute separate I_{MMI} time series for acceleration sensor data ($I_{A:\text{MMI}}(i,j)$) and the velocity sensor data ($I_{V:\text{MMI}}(i,j)$).

We next confirm that the I_{MMI} values computed using the velocity sensor data ($I_{V:\text{MMI}}(i,j)$) are similar to the values computed using the acceleration sensor data ($I_{A:\text{MMI}}(i,j)$). To do this, we examine the ratio of the $I_{A:\text{MMI}}(i,j)$ and $I_{V:\text{MMI}}(i,j)$ values for each second of data (i) at each station (j).

$$R(i,j) = \frac{I_{A:\text{MMI}}(i,j)}{I_{V:\text{MMI}}(i,j)} \quad (2)$$

If $R(i,j)$ is within an acceptable range (0.9–1.1) the final $I_{\text{MMI}}(i,j)$ value is assigned the average of the two sensor values. This is consistent with the work of Li et al. (2019) that found comparing data from co-located sensors is an effective quality control technique for southern California data. When R is outside of the range 0.9–1.1, it could indicate a problem with one of the sensors, suggesting we should eliminate the contribution from one of the sensors.

Typically, velocity sensors have lower noise floors than accelerometers and are more reliable for low ground motions, whereas for high ground motions, velocity sensors may clip while the accelerometer data are more apt to remain on-scale. Given this, if the ratio is not within the acceptable range and $I_{V:\text{MMI}}(i,j) < 3$, we assume something is suspect with the acceleration sensor measurement and we use $I_{V:\text{MMI}}(i,j)$ as the final $I_{\text{MMI}}(i,j)$. If instead the ratio is not within the accepted range and $I_{V:\text{MMI}}(i,j) \geq 3$, we assume the velocity sensor data may have clipped and we use $I_{A:\text{MMI}}(i,j)$ as the final $I_{\text{MMI}}(i,j)$. A summary of this method is outlined in Table S2.

As an example, consider the 2014 M5.1 La Habra, California, earthquake. We processed 7 min of data (from 2 min before to 5 min after the earthquake origin time) into 1 sample-per-second intensity estimates from 552 Stations. These intensity estimates include 841 (3.4%) ratios above 1.1, 181 (0.8%) ratios below 0.9, and 23,548 (95.8%) ratios within the expected bounds (Figure 2a). Overall, we find an asymmetry in the ratio distributions, showing over ~4 times more ratios that exceed 1.1 than fall below 0.9. The prevalence of large ratios stems from sensors observing high amplitude shaking that could cause the velocity sensor to clip. We inspect these ratios in more detail highlighting the velocity sensor (Figures 2b–2d) and acceleration sensor (Figures 2e–2g) data from station CI.OLI, which is located 1.6 km from the earthquake's epicenter. When

Figure 2

the first seismic waves arrive, the ratio test flags one second of data, and the velocity sensor $I_{V:MMI}(i,j)$ value is deemed most appropriate. In the following 4 s, the ground motions are larger and the acceleration sensor $I_{A:MMI}(i,j)$ values are preferred (Figure 2h). Outside of this small portion of data, the ratios are near one and the average of the acceleration sensor $I_{A:MMI}(i,j)$ and velocity sensor $I_{V:MMI}(i,j)$ values are used. This method does not flag all clipped velocity sensor data; rather, when the ratio is within the allowed limits we assume the selected average of the two sensor values ($I_{A:MMI}(i,j)$ and $I_{V:MMI}(i,j)$) suffices.

3.3. Modification #2

The second modification is a two-station event detection logic. In this logic, a PLUM event is detected when a first station trigger is confirmed by a trigger at a second neighboring station within a timeout period (15 s, configurable). This modification reduces the chance of false detections resulting from uncorrelated, time-separated, noise on two far-away stations. This modification is only used to assure the two triggers used for event detection are identifying the same phenomena; it is not a measure of the spatial extent of the PLUM alert region. The 15 s limitation was chosen based on the distribution of distances to the nearest-neighbor station. For most stations in the ShakeAlert network, the distance to its nearest neighbor is less than 45 km (based on the mean plus two standard deviations, Figure S1). If we assume that significant shaking is carried by the S-wave, and that S-wave speeds are ~ 3 km/s (Minson et al., 2020), then it should take no longer than 15 s for the initial trigger from one station to be confirmed by a neighboring station. Although more detailed calculations could be made using the exact locations of the triggering stations, we found that was not necessary and that these approximations successfully capture suspect station pairs. As the ShakeAlert network becomes denser, and the average interstation-spacing becomes reduced, this 15 s threshold may need to be reassessed.

4. Two U.S. West Coast Data Sets

We apply the PLUM algorithm to two retrospective data sets recorded by stations within the ShakeAlert System reporting regions in California and the Pacific Northwest (PNW): (A) $M_L 3.5+$ earthquakes between 2012 and 2017 and (B) a subset of waveforms containing earthquakes and other signals recorded between 1999 and 2015 used by ShakeAlert to assess algorithm performance (Cochran et al., 2018). Of the 39 local earthquakes in the latter data set, 13 are also in the former data set and will net the same results. This duplication is included to produce complete and separate reports for both data sets. Since these data sets are archived and do not include data telemetry latencies, this analysis is retrospective and not intended to be interpreted as a measure of PLUM's real-time performance. Information about how PLUM performed in real-time for the July 2019 M6.4 and M7.1 Ridgecrest, California, earthquakes can be found in Minson et al. (2020).

4.1. West Coast (WC) Data Set

The first retrospective data set from the U.S. West Coast (WC data set) includes 558 $M_L 3.5+$ earthquakes (Data Set S1, ds01.csv) over a 6-year period (2012–2017; see Data Availability Statement). Of these 558 events, 54 are located in the PNW and 506 in California (Figure 1). The median event depth of the WC data set is 8.84 km and the deepest earthquake is 58 km. The median event magnitude is M3.82 and the largest

Figure 2. Data from the M5.1 La Habra earthquake (March 29, 2014) used to present the assessment of the quality of acceleration $I_{A:MMI}$ and velocity $I_{V:MMI}$ using the ratio method for co-located sensors outlined in Table S2. (a) Ratio of $I_{A:MMI}$ divided by $I_{V:MMI}$ as a function of $I_{V:MMI}$ (points). Each point represents the result from processing one-second of data at an individual station, and each sensor contributes 420 data points to the plot representing each second in the 7-min waveform ($420 = 7 \times 60$ s). For robust data, the ratios should have a value close to unity, and here we assume ratios between 0.9 and 1.1 are acceptable (white band), in which case the final I_{MMI} value is assigned the average of $I_{V:MMI}$ and $I_{A:MMI}$. Because this plot is for illustrative purposes only, extreme outlier data are not shown and we limit the y-axis to between 0.8 and 1.5. When the ratio is not within the acceptable limit we use the $I_{V:MMI}$ value if $I_{V:MMI} < 3$ (pink bands) and the $I_{A:MMI}$ value when $I_{V:MMI} \geq 3$ (yellow bands). Results from station CI.OLI (squares) are highlighted. (b)–(d) Velocity data from station CI.OLI for the east-west, north-south and vertical components, respectively. Color swatches encompass times when the ratio is deemed unacceptable. (e)–(g) As in (b)–(d) but for acceleration data. (h) Temporal evolution of the PLUM I_{MMI} values at CI.OLI. Flagged values when the ratio is outside the expected bounds are color-coded; velocity preferred (pink, $N = 1$) and acceleration is preferred (yellow, $N = 4$).

Figure 3. Temporal evolution of seismic stations and earthquakes over the 6-year West Coast study period (January 1, 2012–December 31, 2017). (a) Number of $M \geq 3.5$ earthquakes recorded each year. (b) Number of stations as a function of time, where the y-axis indicates the total number of stations. In 2012 there were ~ 500 operational stations and at the end of 2017 the number of stations exceeded 800. The white horizontal stripe starting in late January 2016 is caused by the deployment of five stations (NP.5430, NP.5425, NP.5080, NP.5434, and NP.5431) that were operational for only about a month.

magnitude earthquake is $M 6.8$. The annual number of earthquakes within our 6-year study varies, with an average of 93 earthquakes per year.

We use the ObsPy Mass Downloader (Beyreuther et al., 2010; Krischer et al., 2015) to retrieve waveform data (MSEED format) and station XML files for our $M_L 3.5+$ target events, which includes requesting all available data from 15 networks (Table S3). For each event, we request 7 min of data at each available station beginning 2 min prior to the catalog time of the event. Prior to processing, if the 7-min time windows of two or more events in the $M_L 3.5+$ catalog overlap, the windows are combined into a single, longer time window. Assuming PLUM might identify events below magnitude 3.5 that may have produced seismic signals within these time windows, the WC ground-truth catalog used for matching PLUM event detections with earthquakes also includes $16,984 \leq M_L < 3.5$ earthquakes.

The number of operational stations increased during the WC study period from ~ 500 stations in 2014 to over ~ 800 stations by the end of 2017 (Figure 3). As of January 2020, the number of stations used by the ShakeAlert network has increased to 1,111. The test suite of stations used in PLUM testing (which included some stations from networks not on the official ShakeAlert contributor list) can be found in Data Set S2, ds02.csv. So, these retrospective tests are a conservative estimate of the detection ability of a present-day real-time PLUM implementation.

4.2. West Coast Testing and Certification (WCTC) Data Set

The second data set is a modification of the West Coast Testing and Certification (WCTC) data suite designed by ShakeAlert to determine how EEW algorithms perform with large historical events as well as

Table 2
We Test Six Different Configurations of PLUM (Instances)

Instance index	I_{MMI} primary threshold (station #1)	I_{MMI} secondary threshold (station #2)	Threshold differential ΔI_{MMI}
1	4.00	--	--
2	4.00	4.00	0.00
3	4.00	3.75	0.25
4	4.00	3.50	0.50
5*	4.00	3.00	1.00
6	4.00	2.50	1.50

Note. For the data in this study, we favor Instance #5 (denoted with an asterisk).

Abbreviation: PLUM, Propagation of Local Undamped Motion.

problematic data (such as large magnitude earthquakes, deep teleseismic events and calibration pulses) that have caused the system to issue false detections in the past (Cochran et al., 2018; see Data Availability Statement). These data include (1) Historical events from southern California ($N = 17$; Data Set TC1), northern California ($N = 12$; Data Set TC2) and the PNW within Oregon and Washington ($N = 10$; Data Set TC3); (2) Calibration and re-centering tests in southern California ($N = 3$; Data Set TC4), northern California ($N = 23$; Data Set TC5), and known anomalous signals in the PNW ($N = 10$; Data Set TC6). Some of these data were specifically selected because they generated false alerts in the real-time system; (3) Known problematic teleseismic and regional events ($N = 9$; Data Set TC7); and (4) Large magnitude teleseismic events (2000–2015) within 10,000 km of southern California both shallow ($N = 10$; Data Set TC8) and deep ($N = 10$; Data Set TC9).

The historical local WCTC data (TC1, TC2, TC3, and TC7) contain waveforms that are 7 min in duration, windowed 2 min prior to and 5 min post the theoretical P-wave arrival time at the closest ShakeAlert station at the time of the event (Cochran et al., 2018). For the calibration test signals,

recentering and anomalous signal data have waveform durations ranging from 7 min up to 6 h. The waveforms for the regional and teleseismic events were selected based on a theoretical P-wave arrival, where the waveform snippets contain 2 min of waveform prior to the expected P-wave arrival and 5 min after, for a total duration of 7 min (Cochran et al., 2018). The longer 6-h data sets were processed using the provided data-subsets, typically each 30 min in duration.

A limiting factor of PLUM is that it is not designed to ingest data recorded in trigger-mode where, for a given station, data from different channels have different lengths, which makes it impossible to correctly apply Equations 1a and 1b presented above. For a given event, it is not a problem if only a small portion of the data have mismatched lengths, but when a large percentage of the data are mismatched the data are not viable for PLUM processing. Of the 105 WCTC events and signals, data from three earthquakes were deemed non-viable because of significant mismatching lengths and were removed from consideration: the 2014 M6.5 earthquake near Vancouver Island, Canada; the 2001 M6.8 Nisqually earthquake near Puget Sound, Washington; and the 2000 M7.6 Volcano Island, Japan, teleseismic earthquake. Removal of these three events brings the total number of signals in the WCTC data set from 105 to 102. Additionally, seven stations were identified to have incorrect gains (Deborah Smith, written communication, 2020; NC.CSL, NC.CVP; NC.PCH, NC.PCHB, NC.PCHC, NC.PHOB, and NC.PMM), and waveforms from these stations were removed.

5. Testing Implementation and Performance Assessment

We follow the method of Cochran et al. (2019) and explore six different triggering threshold configurations, which we call Instances. One Instance issues a detection when a single station exceeds a given I_{MMI} threshold. The other five Instances require a second neighboring station to exceed a secondary I_{MMI} threshold. These Instances explore threshold differentials (i.e., the difference in I_{MMI} thresholds for the first and second triggering stations) spanning 0.0–1.5 I_{MMI} units (see Table 2). For both the one-station and two-station methods we found a primary I_{MMI} threshold of 4.0 returned a larger number of correct event detections than the other tested primary thresholds (Figure 4). Moving forward, we only conduct tests using a primary I_{MMI} threshold of 4.0, which is the I_{MMI} level that ShakeAlert aims to detect (Kohler et al., 2020).

For each event, our modified PLUM detection algorithm is applied to consecutive one-second waveform data blocks from all available stations. For the single-station Instance, when a station's observed I_{MMI} exceeds the primary trigger threshold, an event is declared, the associated information is placed in the Active Event Repository, and PLUM issues an event detection. For the two-station Instances, a second, neighboring, station must exceed the secondary triggering threshold within 15 s of the first station's trigger. The information from the first station trigger is thus placed in the Pending Event Repository. Then the first

Figure 4. Exploration of the distribution of detections by different primary I_{MMI} thresholds for the West Coast data. We find that for both the one-station and two-station methods, a primary I_{MMI} threshold of 4.0 nets the most correct event detections (matched, green). (a) One station method results. (b) Example two-station method results using a secondary I_{MMI} threshold 1.0 unit below the primary I_{MMI} threshold. Note the false detections (brown) are completely removed when using a two-station approach.

station's six nearest neighboring stations are examined for amplitudes at or above the secondary trigger threshold every second for 15 s following the first station trigger until one or more of the neighboring stations observes an amplitude above the secondary threshold (Cochran et al., 2019). If so, the event is moved from the Pending Event Repository to the Active Event Repository and a PLUM event detection is issued. Otherwise, the event is removed from the Pending Event Repository and no detection is issued. For detections, our convention is to assign the packet end time to the PLUM detection time, consequently this introduces an uncertainty of one second to these values. For an event in progress, the I_{MMI} observed at each station within the event is monitored, and once all stations observe less than I_{MMI} 2.0 for 10 s, the shaking from the earthquake is assumed to be complete and the event is removed from the Active Event Repository.

5.1. Matching PLUM Event Detections With Earthquakes

We evaluate PLUM's ability to detect earthquakes that cause I_{MMI} 4+ shaking within the ShakeAlert alerting boundary and properly ignore teleseismic events, calibration pulses, and anomalous signals. This evaluation is event-based, that is, we match PLUM event detections to catalog earthquakes. Because PLUM does not estimate earthquake source parameters, we have to build a methodology for matching PLUM event detections to cataloged earthquakes. Our matching criteria is twofold. Following Cochran et al. (2019), the first criteria is based on timing, where we require the first station trigger to occur after the arrival of the P-wave (with assumed speed of 6.5 km/s) and no later than the passing of surface waves (assumed speed 2 km/s). To account for variations in seismic velocities throughout our study region, and the one-second ambiguity in the time of the trigger, we assume the P-wave might arrive up to one second earlier than expected and that the surface wave might arrive up to 5 s later than expected. Thus, the temporal restrictions for identifying when

a PLUM event detection can be classified as a match with a cataloged earthquake requires the hypocentral distance to the first triggered station to be such that the PLUM detection time falls within these bounds:

$$T_{\min} = \left(\frac{(\text{hypo_dist})}{(6.5 \text{ km/s})} \right) - 1 \text{ s} \quad (3a)$$

$$T_{\max} = \left(\frac{(\text{hypo_dist})}{(2.0 \text{ km/s})} \right) + 5 \text{ s} \quad (3b)$$

where hypo_dist is the hypocentral distance in km between the earthquake source and the first triggering station. If a PLUM event detection time falls within this temporal range corresponding to an earthquake in the ANSS comprehensive catalog (ComCat) catalog, the two are paired as a potential match.

Our second criteria requires that the observed maximum I_{MMI} at the first triggered station is not implausibly large given the magnitude of the earthquake and hypocentral distance. To confirm this, we check that the hypocentral distance is smaller than the maximum distance at which the observed I_{MMI} would be expected based on the intensity prediction equation (IPE) of Allen et al. (2012, their equation 2). This plausibility test requires:

$$\text{Hypo_dist} < 10^{(c1 + c2 * M + c3 * \max\{I_{\text{MMI}}\})} \text{ km} \quad (4)$$

where Hypo_dist is the hypocentral distance between the station and the catalog earthquake, M is the catalog magnitude of the earthquake, $\max\{I_{\text{MMI}}\}$ is the maximum I_{MMI} observed at the first triggering station, and we use constants from the Allen et al. (2012) IPE that allow for three standard deviations in ground motion amplitude ($c1 = 1.5604$; $c2 = 0.3939$; and $c3 = -0.2718$). If both criteria are met, the PLUM event detection is deemed a match to that catalog earthquake.

5.2. Identification of Detectable Events

It is important to differentiate between an earthquake that PLUM did not detect because it produced low I_{MMI} values within the ShakeAlert alerting boundaries and an earthquake that PLUM missed for other reasons such as sparse station coverage. The former is a desirable outcome: the earthquake, regardless of its magnitude, produced ground motions too small to be of interest and no event detection (and subsequent EEW alert) should be issued. However, it is problematic if an event is missed for any other reason. Of particular interest is identifying regions where the station coverage is so sparse that it is impossible for PLUM to detect M5+ earthquakes that do produce $I_{\text{MMI}} > 4$ shaking.

To this aim, we first determine which events in our two data sets are theoretically detectable by PLUM (subsequently referred to as theoretically detectable). To do this we first calculate the median-expected spatial extent of I_{MMI} energy for each earthquake using the particular Instance's I_{MMI} threshold(s). These extents are then compared with the station configuration to determine if a given PLUM Instance could theoretically detect the earthquake. In this evaluation, it is essential to use the station network configuration that was operational at the time of the earthquake as the network changes considerably over our 6-year study period (Figure 3). Notably, we do not use the Allen et al. (2012) IPE for this task as used above for event matching. This is because the Allen et al. (2012) IPE is a global IPE known to generate larger-than-generally-observed I_{MMI} extents in the west coast of the U.S. This makes the Allen et al. (2012) IPE a good choice when associating PLUM detections with earthquakes because it allows leeway to account for high-stress drop earthquakes, but it will be too generous for estimating the theoretical detectability of earthquakes.

To determine which events are theoretically detectable, we use the ground-motion models used by the ShakeAlert System to estimate I_{MMI} extents (Kohler et al., 2020). In this procedure, different ground-motion prediction equations (GMPs) are implemented for different regions to estimate the median-expected Peak ground velocity (PGV) and Peak ground acceleration (PGA) with distance. For events in California,

the Boore and Atkinson (2008) GMPE is used for $M \geq 5.7$ earthquakes and the Atkinson and Boore (2011) GMPE is used for $M < 5.7$ earthquakes. For events in the PNW, the Chiou and Youngs (2008) GMPE is used. The ground motion intensity conversion equations of Worden et al. (2012; their equations 6 and 9) is then used to compute I_{MMI} values from the median-expected PGV and PGA values with distance estimated by the GMPEs for a given earthquake magnitude. Lookup tables presenting the spatial extent of I_{MMI} 2.5 to 6.0 with magnitude are included as Data Sets 3 and 4 (ds03.csv and ds04.csv). Note, because the GMPE regressions used epicentral distances, these GMPEs rely on epicentral distance, not hypocentral distance. When evaluating if an earthquake is theoretically detectable, we use the station network configuration that was operational at the time of the earthquake; this is essential as the network changes considerably over our 6-year study period (Figure 3).

In Step 1 of our theoretical tests, we identify the closest station to the earthquake (i.e., station with the smallest hypocentral distance) that we assume to be the first triggering station. We next determine the six closest neighboring stations to this first station, and of these, the station with the smallest hypocentral distance to the earthquake is deemed the second station to theoretically trigger on large ground motions. We assume these are the two triggering stations, although operationally that might not be the case because of any number of factors that contribute to ground motion variability such as directivity or local amplification effects. We then use our lookup tables (with the earthquake magnitude and epicentral distances to the first and second stations) to determine if shaking at these two closest stations is theoretically strong enough to trigger a PLUM detection using the PLUM Instance triggering thresholds.

Next, we determine if the configuration of the two closest stations can meet our requirement that the second detection be issued within 15 s of the first detection. Although for large magnitude earthquakes PLUM has been observed to trigger on the P-wave in the near field (Minson et al., 2020), here we take a conservative approach and assume PLUM triggers on S-wave signals, traveling at $V_s = 3.0$ km/s. We compute the source-to-station travel times to both stations and compute the time difference ($\Delta t_{\text{sta1_to_sta2}}$) between the expected triggers:

$$\Delta t_{\text{sta1_to_sta2}} = \left(\frac{\text{hypo_dist}_2 - \text{hypo_dist}_1}{V_s} \right) \quad (5)$$

where hypo_dist_1 and hypo_dist_2 are the hypocentral distances to the closest station and second closest to the catalog location, respectively. We can then classify which earthquakes meet the 15 s restriction by PLUM and which are not as:

$$\text{Meets 15 s requirement: } \Delta t_{\text{sta1_to_sta2}} \leq 15 \text{ s} \quad (6a)$$

$$\text{Does not met the 15 s requirement: } \Delta t_{\text{sta1_to_sta2}} > 15 \text{ s} \quad (6b)$$

We use median-expected ground motions to compute these binary theoretical limits (i.e., detectable or not detectable). However, because of ground motion variability and variability in source processes, it is possible that PLUM can detect an event that is not theoretically detectable, and conversely, that PLUM does not detect an event that is theoretically detectable. Indeed, this is what we find.

6. Results

6.1. WC Data: M3.5+ Events (2012–2017)

We first investigate the WC data (M3.5+ earthquakes between 2012 and 2017), applying the six different Instance configurations (see Table 2; Figure 5a). For all six Instances, the mean detection times (i.e., elapsed time between the earthquake origin time and its detection by PLUM) for the M4.5+ earthquakes range between ~10 and 13 s and the mean detection times for the M3.5+ earthquakes range between ~6 and 9 s (Figure 5b), showing consistently that detection times of the smaller earthquakes are lower than those of the larger earthquakes. We expected lower detection times for large earthquakes since it is more likely the triggering threshold for large events can be exceeded in the P-wave (Minson et al., 2018, 2020). However,

Figure 5. Summary of results from the six configuration Instances for a primary I_{MMI} threshold of 4; only successfully detected events are included. (a) Stacked histograms of detected events for each Propagation of Local Undamped Motion (PLUM) Instance, color-coded by magnitude. (b) Average detection time (i.e., time between event origin time and its detection by PLUM) for each Instance, color-coded by magnitude.

this turns out not to be the case for two reasons. First, larger earthquakes produce signals that can be detected at larger distances, so $M5+$ events can be detected even in areas with sparse station spacing, albeit with larger detection times. Conversely, $M < 4.5$ earthquakes generally have a small footprint of $I_{\text{MMI}} 4+$ shaking, so $M < 4.5$ earthquakes that are detected are those with stations located near the epicenter and thus have inherently lower detection times. Second, most ($>68\%$ for the one-station test) of the detected $M4.5+$ earthquakes occurred prior to 2015, when the network was sparser (see Figure 3), equating to larger source/station distances and longer detection times.

Initially, Instance #1 (single-station event triggering) looks the most favorable because it detects the most earthquakes with the lowest detection times (Figure 5), however, Instance #1 issues 51 false detections using an $I_{\text{MMI}} 4$ trigger threshold (Figure 4a). Although these false detections were generated at only eight stations (Table S4), we find no obvious connections between them (i.e., similar instrumentation, geographical region, or location characteristics). These types of false detections can help establish schemes to identify errant signals and unanticipated noise patterns in the future. As shown below, the single-station Instance #1 (single station) is also prone to issuing false detections for the other data sets. Given this, Instance #1 is eliminated from contention.

False detections are completely eliminated when using a two-station method (Figure 4b). For all two-station methods tested (Instances #2–6; see Table 2), PLUM issued no false detections for the $M3.5+$ retrospective data. This confirms that the best way to safeguard PLUM from making false detections on spurious noise signals is to implement a two-station triggering method. We next eliminate Instances #2 and #3 because the number of events they detect is substantially less than the other two-station Instances, including missing a number of $M5+$ earthquakes. Instance #4 is also removed because of the larger detection times, especially for

Table 3

List of 23 Earthquakes That PLUM can Theoretically Issue Timely Warnings (Detection Issued Before the S-wave Reaches the Theoretical Expected Extent of $I_{MMI} = 4$ Shaking) Using Instance #5

Date time	Earthquake magnitude	Potential maximum warning time (second)	Data Set
16-Oct-1999 09:46:44	M7.1	28	WCTC
22-Feb-2003 12:19:10	M4.99	1	WCTC
22-Dec-2003 19:15:56	M6.5	16	WCTC
28-Sep-2004 17:15:24	M5.97	13	WCTC
06-Jan-2005 14:35:27	M4.42	<1	WCTC
12-Jun-2005 15:41:46	M5.2	4	WCTC
31-Oct-2007 03:04:54	M5.45	9	WCTC
29-Jul-2008 18:42:15	M5.44	7	WCTC
10-Jan-2010 00:27:39	M6.5	22	WCTC
04-Apr-2010 22:40:42	M7.2	39	WCTC
07-Jul-2010 23:53:33	M5.42	5	WCTC
13-Feb-2012 21:07:02	M5.6	2	WC
06-Aug-2012 07:35:49	M4.45	<1	WC
26-Aug-2012 20:57:58	M5.41	<1	WC
21-Oct-2012 06:55:09	M5.29	1	WCTC
11-Mar-2013 16:56:06	M4.7	<1	WC/WCTC
10-Mar-2014 05:18:13	M6.8	14	WC/WCTC
17-Mar-2014 13:25:36	M4.4	1	WC
29-Mar-2014 04:09:42	M5.1	4	WC/WCTC
29-Mar-2014 21:32:45	M4.14	<1	WC
24-Aug-2014 10:20:44	M6.02	19	WC/WCTC
21-May-2015 03:15:29	M4.1	<1	WC
10-Jun-2016 08:04:38	M5.19	3	WC

Note. Four events are in both the WC and WCTC data sets and the maximum potential warning times agree to within the 1-s uncertainty, thus, for simplicity, only the WC value is listed. Warning time estimates do not include telemetry and public alert dissemination.

Abbreviations: PLUM, Propagation of Local Undamped Motion; WC, West Coast; WCTC, West Coast Testing and Certification.

the larger magnitude events (Figure 5b). This leaves Instances #5 and #6 as our main candidates for the preferred PLUM event triggering threshold strategy.

Looking at the detection times in more detail, PLUM Instance #5 detects 79 earthquakes and has detection times ranging from 1.5 to 37 s with mean and median values of 8.0 and 5.9 s, respectively. PLUM Instance #6 detects 92 earthquakes, and has the same detection time range, mean and median as Instance #5. Two events (2014 M6.8 off-shore and 2014 M4.89 south of the California-Mexico border) have detection times that exceed 30 s, which is why the mean detection time is higher than the median.

We next explore the expected timeliness of the PLUM detections to determine if one Instance provides timelier detections than the others. A detection is timely if an alert can be issued when at least one location has yet to experience shaking at or above $I_{MMI} = 4$. Again, we assume this level of shaking occurs with the arrival of the S-wave traveling 3 km/s. Although near-source detection speeds can be highly variable, this assumed 3 km/s speed is consistent with the analysis of Minson et al. (2018) as well as the $I_{MMI} 4+$ propagation speed derived by Minson et al. (2020) for the majority of the $I_{MMI} 4+$ areas in the 2019 M6.4 and M7.1 Ridgecrest, California, earthquakes. To determine which event detections are timely and which are not we use the following steps. (1) Use the earthquake magnitude and lookup tables to determine the epicentral distance corresponding to $I_{MMI} = 4$; using this epicentral distance and the earthquake depth, a hypocentral distance is computed. (2) Using the 3 km/s velocity estimate and the hypocentral distance, a maximum theoretical elapsed time can be calculated. (3) Compare the theoretical elapsed time with the PLUM detection time. If the PLUM event detection is issued after the theoretical limit (i.e., the PLUM detection time is greater than the theoretical elapsed time), the detection is deemed too late (i.e., untimely). If the PLUM event detection is issued before the theoretical limit, the detection is deemed timely. The derived theoretical limits are similar to the USGS ShakeMaps $I_{MMI} = 4$ extents for the five on-shore WC M5+ earthquakes (excludes the 2016 M5 Geyser event that did not have an Atlas or NEIC (US) ShakeMap), suggesting the theoretical timeliness analysis is viable (Figure S2).

As an example of how the theoretical values are computed, we evaluate the 2014 M6.02 South Napa earthquake (see Figure 1). Based on the California lookup tables, an earthquake of this size will produce $I_{MMI} = 4$ energy out to an epicentral distance of 68.6 km. For this earthquake's 11.1 km depth, this corresponds to a hypocentral distance of 69.5 km. The elapsed time for an S-wave traveling at 3 km/s to travel 69.5 km is 23 s, demarking the theoretical limit for timely detections. Both PLUM Instance #5 and

Instance #6 detect the South Napa earthquake in 4 s. Thus, theoretically there will be a maximum-available potential warning time of 19 s. For simplicity, we use a binary metric (i.e., timely or not timely), assuming all potential warning times exceeding zero are timely warnings (Table 3). This is because warning times can vary significantly depending on the final alerting strategy employed (Minson et al., 2020).

The potential warning times for the stricter Instance #5 are not appreciably slower than those from Instance #6. In fact, for the 11 events that are detected by both Instances the potential warning times agree to within the 1 s uncertainties. These findings suggest both Instances #5 and #6 would perform well, the primary difference being that Instance #6 is less conservative, issuing detections for an additional 14 $M < 4.3$ earthquakes that Instance #5 would ignore, but whose detections would not be timely. Our goal is detecting M5+ earthquakes, where timely detections are more likely to be possible, so we favor Instance #5 that detects the

Table 4
Summary of WC Results

Description		Number (All events, $N = 558$)	Number M4+ ($N = 189$)	Number M4.5+ ($N = 54$)	Number M5+ ($N = 20$)
PLUM detected		79	50	21	10
Theoretically detectable		85	38	15	10
PLUM Detected? [Y/N]	Theoretically Detectable? [Y/N]				
Y	Y	51	32	15	10
Y	N	28	18	6	0
N	Y	34	6	0	0
N	N	445	133	33	10

Note. Comparison of the number earthquakes that are detected by PLUM (Instance #5: I_{MMI} primary and secondary thresholds 4.0 and 3.0, respectively) and number theoretically detectable. Given ground-motion variabilities, the 79 events detected by PLUM Instance #5 include some earthquakes that were not within the subset of the 85 theoretically detectable earthquakes. Non-zero values listed for “PLUM detected (yes) and theoretically detectable (no)” do not indicate false detections, instead, they signal that the ground motions were larger than expected.

Abbreviations: PLUM, Propagation of Local Undamped Motion; WC, West Coast.

larger magnitude events while reducing the detection of smaller events. This is the main driver in selecting Instance #5 over Instance #6 and explains why, counter-intuitively, when examining data from the full West Coast that has lower station density than the prior southern California study (Cochran et al., 2019; Instance #6 preferred) that we prefer an increase, rather than a decrease, of the second triggering threshold.

Additionally, as discussed below, using the WCTC data set, Instance #5 successfully avoids issuing a detection for a box-car calibration test whereas Instance #6 does not, which also supports selecting Instance #5. We expect Instance #5 to be less likely to issue triggers for noise signals when implemented in a real-time system given the slightly higher second station trigger threshold. Moving forward, we only present results from Instance #5 ($I_{\text{MMI}} = 4.0$ primary threshold; $I_{\text{MMI}} = 3.0$ secondary threshold). A general summary of Instance #5 performance using the WC data is presented in Table 4 and a full summary of the results is presented in Data Set S1 (ds01.csv).

6.1.1. Large Events (M4.5+; WC Data Set)

Of the 20 M5+ earthquakes, 10 are theoretically detectable and PLUM successfully detected all 10. These included two offshore events (a M6.8 in 2014 and a M5.72 in 2015; Table S5) of the 54 M4.5+ earthquakes, 15 are theoretically detectable and again PLUM detected them all (Figure 6; Figure S3). Of the 79 detected events, only 11 of the theoretical warning times would be potentially timely, of which 6 are M5+ (Table 3).

For the largest earthquake in our WC data set, a M6.8 earthquake that occurred offshore of Ferndale, CA, in 2014 (see Figure 1), PLUM reported $I_{\text{MMI}} 4+$ values at only station NC.KMPB located 97 km from the earthquake epicenter. Despite the fact that the M6.8 event was ~75 km offshore, it was sufficiently large to generate moderate onshore ground motions and provide a timely detection (see Table 3). For the 2015 offshore M5.72 event, two stations exceeded $I_{\text{MMI}} 4+$ (NC.KCT and NC.KMR). However, despite the near-shore location, the PLUM event detection was not timely.

6.1.2. Theoretically Detectable Earthquakes That PLUM Missed (WC Data Set)

Of the 85 theoretically detectable earthquakes in the WC data set (Table 4), PLUM detected 51 and ignored 34 whose magnitudes were <M4.5 (Data Set S1; ds01.csv). None of the ignored events would have had timely detections in regions of $I_{\text{MMI}} \geq 4$.

6.1.3. Theoretically Undetectable Events

A large percentage of events in the WC data set are theoretically undetectable because of sparse station density at the time of the event (473/558; 85%); in fact, no earthquakes in the PNW for the WC data set are theoretically detectable (Figure 1). Of the 473 theoretically undetectable events, the majority (434) are below magnitude M4.5 and not a concern. In the WC data set, 39 M4.5+ earthquakes are theoretically

Figure 6. Results from the West Coast 6-year (2012–2017) retrospective study using Instance #5, showing the theoretical epicentral distance to the first triggered station as a function of earthquake magnitude. Included are theoretical curves (ShakeAlert ground-motion prediction equation for CA and Pacific Northwest [PNW]; Worden et al., 2012 ground motion intensity conversion equations) showing estimates for California ($\text{lat} \leq 42^\circ$; solid line) and PNW ($\text{lat} > 42^\circ$; dotted line). Propagation of Local Undamped Motion detected all 15 of the theoretically detectable M4.5+ earthquakes. For clarity, we have truncated the plot at an epicentral distance of 150 km, which removes a M6.5 earthquake, and three $M < 5$ earthquakes; all truncated events are theoretically undetectable. The 2015 high stress drop event is labeled. A version of this figure with a logarithmic x-axis is shown in Figure S3.

undetectable because of minimal station density near the source. Of these, PLUM detected 6, indicating PLUM is detecting more than theoretically expected. This is reasonable given that we use median-expected ground motions to determine if an event is theoretically detectable, and because of the inherent, and significant, variability in an earthquake's ground motions (Minson et al., 2021).

6.2. WCTC: ShakeAlert Testing and Certification Data (1999–2015)

We next present PLUM's performance for the WCTC data sets. These data contain 102 signals and include local historical events, problematic regional and teleseismic events, calibration pulses, and other anomalous signals. The results are presented below according to category and a summary of the results are listed in Tables 5 and 6. Here, we take a two-fold approach: first, we assess how PLUM performs on the specified data; and second, we determine if, by chance, any additional earthquakes are present within these data and if so, whether they were properly detected. As with the first data set, for these data we apply PLUM Instance #5 using a primary I_{MMI} threshold of 4.0 and a secondary threshold of 3.0.

6.2.1. M5+ Events in Local Historical Catalogs TC1, TC2, and TC3

We first examine the 19 M5+ earthquakes in an aggregated catalog containing earthquakes in the local historical catalogs TC1, TC2, and TC3 (six in TC1, 10 in TC2; three in TC3; see Figure 7; Table 5 and Data Set 5, ds05.csv). PLUM detected 14 of these 19 events, where the five earthquakes PLUM missed are theoretically

Table 5

Summary of PLUM's Ability to Correctly Detect Local Earthquakes in the WCTC Catalogs TC1, TC2 and TC3 Using PLUM Instance #5

Data Set	Signal type	#Earthquakes	% PLUM correctly detected		#Earthquakes detected by PLUM versus #theoretically detectable
			$M < 5$	M5+	
TC1	Historic Local (So. Cal)	17	64% (7/11)	100% (6/6)	93% (13/14)
TC2	Historic Local (No. Cal)	12	0% (0/2)	70% (7/10)	100% (7/7)
TC3	Historic Local (PNW)	8*	0% (0/5)	33% (1/3)	100% (1/1)
Total TC1-TC3	--	37	39% (7/18)	74% (14/19)	95% (21/22)

Note. The TC3 data sets contained two non-viable data sets not included in this study, bringing the total number of events in TC3 from 10 to 8 (denoted with an asterisk). An expanded listing of these results can be found in Data Set S5. It is possible for PLUM to detect a theoretically undetectable event, which is why summation of the “% PLUM Correctly Detected” denominators do not always equal denominators in the last column.

Abbreviations: PLUM, Propagation of Local Undamped Motion; PNW, Pacific Northwest; WCTC, West Coast Testing and Certification.

not detectable (Figure 8; Figure S4). Of these five missed events, three were offshore (2005 M7.2 offshore California, 2010 M5.2 offshore Oregon and 2013 M5.5 near Vancouver Island, Canada) and two were on-shore in areas of sparse station spacing.

Many of the historical local earthquakes in the WCTC TC1, TC2, and TC3 catalogs occurred over a decade ago, at a time when the network was relatively sparse. When using the newer January 2020 ShakeAlert network configuration ($N = 1,111$) to recompute the theoretically detectable events we net an additional three detectable earthquakes: (1) 2013 M5.69 Greenville, California, earthquake (listed as the Canyon Dam earthquake in the original Cochran et al., 2018, WCTC data set); (2) 2005 M7.2 Offshore Northern California, earthquake; and (3) 2006 M4.5 Cowlitz Chimneys, WA earthquake. If the 555 additionally planned ShakeAlert stations were also operational the 2004 M5.4 Adobe Hills swarm event would be theoretically detectable. The only WCTC M5+ earthquakes that would not be theoretically detectable, even with the addition of the expanded and planned ShakeAlert network, are the 2013 M5.5 ($z = 10$ km) offshore earthquake near Vancouver Island, Canada, and the 2010 M5.2 ($z = 10$ km) offshore Oregon earthquake. Both of these earthquakes did not produce large enough ground motions on land to trigger PLUM. However, large enough offshore events that produce $I_{MMI} 4+$ ground motions on land are expected to be detectable by PLUM, as is the case for the 2014 M6.8 earthquake discussed above that PLUM detected.

Within the TC1 waveform data, in addition to the targeted TC1 earthquakes, there was also a M5.2 Baja California earthquake on April 4, 2010 at 22:43:01 that was an aftershock of the 2010 M7.2 El Mayor-Cucapah earthquake on that same day. PLUM Instances #5 and #6 do not detect this event, consistent with this event not being theoretically detectable and with the results reported in Cochran et al. (2019).

Table 6

Summary of PLUM's Ability to Correctly Avoid Anomalous and Teleseismic Signals Using PLUM Instance #5

Data Set	Signal type	# Signals	% Correctly ignored
TC4	Calibration and re-centering (So. California)	3	100% (3/3)
TC5	Calibration and re-centering (No. California)	23	100% (23/23)
TC6	Known anomalous signals (PNW)	10	100% (10/10)
TC7	Problematic teleseismic/regional	10	100% (10/10)
TC8	Teleseismic (2000–2015; Large M)	10	100% (10/10)
TC9	Teleseismic (2000–2015; deep)	9*	100% (9/9)
Total TC4–TC9	--	65	100% (65/65)

Note. The TC9 data sets contained 1 non-viable data set (Volcano Islands, Japan, 2000; M7.6) not included in this study, bringing the total number of events in TC9 from 10 to 9 (denoted with an asterisk).

Abbreviations: PLUM, Propagation of Local Undamped Motion; PNW, Pacific Northwest.

Figure 7. Results from West Coast Testing and Certification test suite data sets TC1, TC2, TC3, and TC7. Of the 39 local events within the ShakeAlert boundary (shown with dashed lines) two events do not have the necessary data for testing (gray stars), 21 were correctly detected by Propagation of Local Undamped Motion (PLUM; purple symbols), 5 were missed (brown stars labeled) and 11 were properly ignored (orange stars and dots). PLUM correctly detected 14 of the M5+ events (purple stars). For the regional events (TC7), that all locate outside of the ShakeAlert boundary, all nine were successfully not detected (orange crosses outside of the ShakeAlert boundary; Alaska earthquake not shown). An indication if the event should be theoretically detectable (black outline) or not theoretically detectable (no black outline) is also included. In this way events that are outlined (theoretically detectable) and orange (not detected) can be identified, however, this only occurs for $M < 5$ events (orange symbols with black outlines), highlighting that all missed M5+ earthquakes (labeled brown stars) are not theoretically detectable. For a map showing which detected earthquakes were timely see Figure S5.

6.2.2. Timely Detection of Earthquakes in the Local Historical Catalogs TC1, TC2, and TC3

Collectively, of the 39 earthquakes in the TC1, TC2, and TC3 local earthquake catalogs PLUM detected 21, with mean and median detection times of 11 and 6 s, respectively. As in the WC data, the mean detection time is elevated compared with the median because of anomalously high detection times (>25 s) for three offshore earthquakes (2004 M4.98; 2011 M6.4; 2014 M6.8). Applying the timeliness metric, as outlined above, of the 21 detections we find 16 are theoretically timely as the detections were issued prior to the theoretical arrival of the $I_{\text{MMI}} = 4$ energy (see Table 3 and Figures S5). All of the timely detected earthquakes locate south of latitude 42° (i.e., not in the PNW region). Of the 16 detected events deemed timely, 14 are M5+.

6.2.3. Small Events ($M < 5$) in the Local Historical Catalogs TC1, TC2, and TC3

We next report our findings for the $M < 5$ earthquakes in the TC1, TC2, and TC3 catalogs, of which PLUM detected seven of eighteen (see Table 5). Five of the 11 earthquakes that were missed have magnitudes below the M4.5 target and will not be investigated further. Of the remaining six events, four are theoretically

Figure 8. As in Figure 6 but constructed using West Coast Testing and Certification local historical earthquake catalogs TC1, TC2, and TC3 ($N = 37$ viable). Unlike Figure 6 the x-axis is not truncated; all data are shown. Of the 19 detected events 14 are $M5+$. Propagation of Local Undamped Motion (PLUM) did not detect 5 $M5+$ earthquakes (labeled year and magnitudes). The 11 missed $M < 5$ events are smaller than the ShakeAlert $M5+$ target and not a concern. The labeled high stress drop 2004 $M4.98$ event offshore San Clemente Island in southern California, that PLUM detected is considered in the discussion section. A version of this figure with a logarithmic x-axis is shown in Figure S4.

not detectable. This leaves two $M4.5$ – 5.0 theoretically detectable events that were not detected by PLUM: the southern California 2002 Yorba Linda $M4.8$ earthquake, and the 2005 $M4.5$ Obsidian Butte earthquake. Both of these events occurred over a decade ago, on land and close to the theoretically detectable limit based on the network configurations at the time. Undoubtedly, with the densification of ShakeAlert stations in these regions events like these would be detected today.

6.2.4. Calibration, Re-Centering Tests and Anomalous Signals (TC4, TC5, and TC6)

The TC4 data from southern California contains three test data sets. The first two tests are box-car calibration signals recorded on stations within ~ 20 km of each other. Using one-station Instance #1 or two-station Instance #6, PLUM would issue a false detection for the 2014 box-car calibration test, but the remaining two-station Instances (#2, #3, #4, and #5) would not issue any detections. This is because the amplitude of the box-car function at the second station exceeds the triggering threshold for Instance #6, but does not exceed the triggering threshold for the other two-station methods. The third TC4 data set contains recentering pulses. These data have inter-station spacings that exceed ~ 75 km. These large distances between stations preclude PLUM from issuing detections, given PLUM's restriction that a second station must issue a detection within 15 s of the first station detection (i.e., the two stations must be within 45 km of each other). In summary, our preferred Instance #5 method successfully avoids issuing detections for all signals in the TC4 data set.

For the northern California calibration and re-centering tests (TC5 data set), PLUM successfully did not issue detections for any of the 23 calibration and re-centering tests. We conclude that non-seismic sources, such as re-centering tests and calibration pulses, are not problematic for PLUM when using our preferred detection triggering criteria (PLUM Instance #5). Cochran et al. (2019) similarly found that, in southern California, PLUM did not issue false detections for recentering tests and calibration pulses.

For the 10 anomalous signals in the PNW (TC6 data set), when using a one-station method, PLUM issues detections on seven of the 10 anomalous signals. However, when using any two-station method, PLUM correctly does not issue any detections for these 10 signals. These results again highlight the importance of implementing a two-station method to avoid issuing detections on anomalous signals.

6.2.5. Problematic Teleseismic and Regional Events (TC7)

PLUM Instance #5 successfully issued no detections for the nine problematic teleseismic and regional events in TC7 data set (Table S6). PLUM Instance #6, however, did issue a detection for the 03 August 2009 M6.9 Canal de Ballenas, Mexico, earthquake (also referred to as the Sonora, Mexico, earthquake). Issuing a detection for this event is not necessarily in error as this event is known for anomalously high amplitude dynamic strain (Agnew & Wyatt, 2014), producing shaking amplitudes six-times larger than expected at southern Californian station AZ.L-VA2 (Cochran et al., 2019). We favor the Instance #5 results that do not issue a detection for this event because the spatial extent of the large amplitudes is very localized and it is unlikely that the detection would be timely.

6.2.6. Teleseismic Events (2000–2015; TC8 and TC9)

Erroneously identifying teleseisms as one or more local events has been a recurring problem with EEW systems, including ShakeAlert. Robust EEW algorithms should not issue detections for non-local events such as teleseisms. Of the 20 teleseismic events in the catalogs TC8 and TC9, the 2000 M7.6 Volcano Island, Japan, data were non-viable (i.e., lacked the required continuous waveforms). Of the remaining 19 earthquakes, PLUM issued no detections for the 10 shallow events in the TC8 catalog and no detections for the nine deep events in the TC9 catalog. If using the Instance #1 one-station approach, a detection would have been issued for the M7.8 2005 Tarapaca, Chile, earthquake. However, no two-station approach issues a detection for any of these events. These findings indicate that PLUM can successfully avoid (100% success rate) issuing detections on teleseismic signals when using the preferred two-station approach.

7. Discussion

We find PLUM is 100% successful at not issuing detections for teleseismic events and problematic signals such as calibration pulses in the WCTC data sets (TC5, TC6, and TC7). These types of signals are known to be problematic for other ShakeAlert algorithms (Kohler et al., 2020). Taken in tandem, PLUM's 100% success rate avoiding these non-local earthquake signals, paired with ShakeAlert's erroneous history with them, suggests PLUM would be a beneficial addition to ShakeAlert.

In this work, we do not investigate alerts and alerting strategies, we only assess PLUM's event detection capabilities. This makes it impossible to precisely compare PLUM with ShakeAlert, but we can make some estimates. During ShakeAlert's initial testing phase, a ShakeAlert Message (alert) was issued (including telemetry) within 5.3 s of the start of the March 29, 2014, M5.1, La Habra earthquake to internal test users (Kong et al., 2016). ShakeAlert also issued an alert (including telemetry) 5.1 s after the start of the August 24, 2014 M6.0 South Napa earthquake (Grapenthin et al., 2014). Both of these events are in our data set and were successfully detected by PLUM within 3 s (La Habra) and 4 s (Napa) relative to the origin time (maximum values that might be elevated by 1 s; does not include telemetry or public alert dissemination). To be on par with the ShakeAlert results would require a few seconds or less for telemetry, which is within reason (Ogwen et al., 2019). PLUM also performed comparably (including telemetry), if not better, than the ShakeAlert algorithms in issuing alerts to areas of I_{MMI} 4+ ground motion during the 2019 M6.4 and M7.1 Ridgecrest earthquakes (Chung et al., 2020; Minson et al., 2020).

If the January 2020 ShakeAlert monitoring network were in place at the time of the two missed onshore M5+ earthquakes in the WC and WCTC data sets (2004 M5.4 Adobe Hills earthquake in the WCTC data set; 2013 M5.69 Greenville earthquake in both the WC & WCTC data sets) both would be theoretically

detectable. The increase in the number of stations and densification of the network will also increase the speed of the detections as the source to station distances will be reduced. However, even with the 2020 ShakeAlert network, the 2016 M5.09 Upper Lake, California, earthquake in the WC data would remain at the edge of theoretical detectability. The epicentral distance to the nearest station NC.GHGB would be 22 km, a distance just beyond the 19.6 km estimated limit for the I_{MMI} 4+ energy. However, with the addition of planned station BK.NS0062 the Upper Lake earthquake would be theoretically detectable.

In general, offshore earthquakes are problematic for EEW algorithms because the one-sided station azimuthal coverage can cause inaccurate epicenter estimates (Hsu et al., 2018). PLUM successfully avoids these types of location errors by using direct observations of the ground motions to issue detections. For PLUM, the offshore earthquakes are only difficult to detect if they do not produce substantial regions of large I_{MMI} at the onshore stations. Of course, if an offshore earthquake does not produce I_{MMI} 4 on land, then an EEW detection for that earthquake is not necessary. For example, three of the five missed M5+ local earthquakes in the WCTC catalog (data sets TC1, TC2, and TC3) were offshore. Theoretically, only the largest of these three (2005, M7.2) would be detectable if the January 2020 ShakeAlert network had been in place at that time. The other two missed events are simply too small (2010 M5.2 and 2013 M5.5) and are too far offshore (i.e., they both locate in the western offshore region of the ShakeAlert boundary) to produce I_{MMI} 4.0+ shaking on land. We do not consider this lack of event detection a problem for these offshore earthquakes because of the absence of I_{MMI} 4+ shaking on land.

We can make strides, however, to improve PLUM's offshore earthquake detection capabilities. By definition, offshore earthquakes will have relatively large azimuthal gaps (i.e., contiguous azimuthal range surrounding an earthquake that lacks a seismic recording station). Thus, the two stations detecting an offshore earthquake will typically have similar back-azimuths. Therefore, in general, to improve PLUM's performance for detecting offshore events, coastal stations should have higher station densities than inland station configurations. This is especially true for stations on Vancouver Island, Canada, where the build-out of EEW stations is just getting underway (Bent et al., 2020). Similarly, a denser station configuration near the ShakeAlert boundary would also improve PLUM's detection capabilities.

For the WC data set, PLUM detected two offshore Cascadia M5+ earthquakes: a M5.7 (January 28, 2015; depth = 16.18 km) and a M6.8 (March 10, 2014, depth = 16.44 km). While both of these earthquakes were also detected by ShakeAlert, the alerts issued contained large location errors (Kohler et al., 2020). The 2015 M5.72 earthquake was ~30 km from the coastline and the epicenter estimate by ShakeAlert was mislocated by 20 km in the initial alert and increased to 90 km error in the final alert update (Angela Chung, written communication 2020). For the 2014 M6.8 earthquake, ShakeAlert misreported this earthquake as a series of two smaller magnitude earthquakes (Kohler et al., 2020). PLUM's design is to report large shaking, not estimate source parameters (i.e., earthquake location and magnitude), so these offshore events are not as troublesome for PLUM. Additionally, untangling consecutive small earthquakes, or handling many near-simultaneous earthquakes as can occur in a robust aftershock sequence, is not a significant issue for PLUM (Kodera et al., 2018).

Detection of earthquakes that were deemed theoretically undetectable (Figures 6 and 8) allows for the identification of high stress drop events that produce ground motions above the median-expected values. Two events in particular exhibit this behavior. The first is in the WC data set, a M4.89 earthquake ($z = 12.5$ km) on February 9, 2015 (labeled in Figure 6) located south of the US/Mexico border near Maneadero, B.C., Mexico. For this event, the Atlas ShakeMap (García et al., 2012; Worden & Wald, 2016) reports an unusually large maximum $I_{\text{MMI}} = 7$ near the source. The second is in the WCTC data set, a M4.98 earthquake ($z = 9.8$ km) on June 15, 2004 (labeled in Figure 8) located offshore San Clemente Island, in southern California. For this second event, the southern California ShakeMap reports a relatively large maximum of $I_{\text{MMI}} = 5$ at station CE.03121 at the Ocean Beach Fire Station in San Diego, 78 km from the epicenter, which we expect was caused by local amplification effects. Thus, detecting these events is probably desirable since they truly produced significant shaking for which the public should be warned (given ShakeAlert's current target warning level of MMI 4+ shaking).

Using Instance #5, PLUM issued timely detections (i.e., detected the earthquake prior to the S-wave passing the theoretically expected extent of the I_{MMI} 4+ energy, not including telemetry or public alert dissemination)

for only 11 out of the 79 detected earthquakes in the WC data set (Table 3). For the WCTC local data (S1, S2, and S3) all detected earthquakes are located below latitude 42° (i.e., none in the PNW region). Of these 16, 14 were M5+ (Table 3; Figure S5). This reflects a known difficulty in issuing timely alerts for small earthquakes (e.g., Wald, 2020). To determine the optimal tuning of the PLUM method, or any EEW algorithm, we will need to better understand the trade-offs of different EEW strategies, including societal tolerance of late alerts and false alarms, as well as the ability of the public to quickly take protective action upon receipt of an EEW alert (McBride et al., 2019, 2020; Minson et al., 2020; Sutton et al., 2020; Thompson et al., 2020).

All data in our retrospective tests target specific types of signals, whether local, regional, or teleseismic earthquake or noise sources that have been problematic for the existing ShakeAlert System. Given this, our tests do not necessarily account for the full range of signals that might be present in continuous, real-time data. We also use precompiled data that avoids issues such as station reporting delays, data gaps, or incomplete subsets of channels reported at a station. As we move to test the PLUM algorithm on continuous real-time data, further modifications may be required to appropriately handle non-earthquake signals within the real-time data feed. As the ShakeAlert network densifies and extends its spatial footprint, this could reduce detection times at individual stations and thus reduce PLUM's time-to-alert using the two-station triggering method. Given this, reassessment of our parameters may be needed as we test new data using new network configurations.

8. Summary and Conclusions

This work examines PLUM's earthquake detection capabilities as applied to two different data sets: a historical catalog of 558 M3.5+ earthquakes that occurred in the U.S. West Coast between 2012 and 2017 (WC); and the ShakeAlert Testing and Certification Platform of historical earthquakes and known problematic signals between 1999 and 2015 (WCTC). In this analysis, we consider whether the earthquakes in our catalogs are theoretically detectable based on median-expected shaking distributions and the station distribution at the time of the earthquake. We explore six different triggering configurations (Instances; see Table 2) and found that for both the one- and two-station configurations a primary triggering threshold of $I_{MMI} = 4.0$ is preferred. We also found that a two-station triggering approach successfully avoids false detections from problematic signals in our data sets. Specifically, Instance #5, with $I_{MMI} = 4.0$ and $I_{MMI} = 3.0$ thresholds for the first and second stations, respectively, is our preferred two-station method because it maximizes the detection of M5.0+ events while minimizing detection of $M < 4.5$ earthquakes.

For the 6-year WC data set (558 M3.5+), PLUM Instance #5 detects 79 M3.5+ events, all located south of latitude 42°; no earthquakes were detected within the PNW ShakeAlert boundary. The detection times of these 79 detections range between 1.5 and 37 s (mean 8 s; median 6 s). Of the 20 M5+ events, only 10 are theoretically detectable, of which PLUM Instance #5 detected all 10. PLUM also detected 11 M4.5-M5.0 events, of which six were theoretically undetectable.

Of the 79 detected WC events 68 produce untimely detections (i.e., detections issued after the S-wave surpasses the expected $I_{MMI} = 4$ extent). Of these untimely detections, 14 were issued for M4.5+ earthquakes where station densification might assist in detecting these types of earthquakes in the future, whereas 54 were issued for $M < 4.5$ quakes where the footprint of the $I_{MMI} = 4$ energy is so small that timely detections are unlikely. Of the 11 potentially timely detections (Table 3), only two have the potential to issue warning times exceeding 10 s.

Of the 19 M5+ local earthquakes in the WCTC data sets TC1, TC2, and TC3, PLUM detected 14 and ignored 5. Three of the missed M5+ events locate offshore and did not produce $I_{MMI} 4+$ shaking on land, and therefore not issuing a detection is the correct response. The other two missed M5+ events were in regions of sparse station coverage at the time (Greenville, previously called Adobe Hills, 2004; Canyon Dam, 2013) and theoretically, both would be detectable if the planned ShakeAlert network were operational at the time of these quakes. PLUM also detects 39% of $M < 5$ events.

Of the 21 PLUM detected local earthquakes in the WCTC data set the detection times ranged between 3 and 30 s. This large range stems from these data spanning 15 years, a period of time when the seismic network density was drastically changing. For these 21 PLUM detections, the detection time mean and median are 11

and 6 s, respectively. The elevated mean, in comparison with the median, results from three off-shore earthquakes reporting detection times exceeding 25 s. Of the 21 detections, 16 are potentially timely, of which 7 and 3 could provide warnings exceeding 10 and 20 s, respectively. Of the regional earthquakes, calibration/anomalous signals and teleseismic signals in the WCTC data sets (TC4–TC9), PLUM Instance #5 correctly does not issue event detections (100% success rate; Table 6).

The results presented in our study are consistent with past analyses that use a two-station version of PLUM to study U.S. data (Cochran et al., 2019; Minson et al., 2020). However, in this work we use a more conservative version of PLUM, applying a higher secondary triggering threshold of 3.0 instead of the lower 2.5 value in past work (Cochran et al., 2019; Minson et al., 2020). With the higher secondary threshold, we successfully maximize the detection of $M \geq 5$ events while at the same time minimizing the detection of $M < 5$ events, many of which would be so small that timely detections would not be possible (i.e., the I_{MMI} 4+ energy would have passed prior to the event detection). Additional tests using real-time data will be needed, but overall, these findings suggest PLUM is a viable EEW method for the West Coast of the U.S.

Data Availability Statement

The WC data set (magnitude, location, and depth) were obtained from the ComCat catalog (data last accessed July 2019). The miniSEED waveform data used in the WC data set were obtained from <https://scedc.caltech.edu/research-tools/eewtesting.html>, which included associated metadata (data last accessed November 2019). Data from the following networks were used in this study: Anza regional network (AZ; <https://doi.org/10.7914/SN/AZ>), Red Sísmica del Noroeste de México (BC; <https://doi.org/10.7914/SN/BC>), Berkeley Digital Seismograph Network (BK; <https://doi.org/10.9732/BDSN>), Cascade Chain Volcano Monitoring (CC; <https://doi.org/10.7914/SN/CC>), California Strong Motion Instrumentation Program (CE: no doi available), Southern California Seismic Network (CI; <https://doi.org/10.7914/SN/CI>), Canadian National Seismograph Network (CN; <https://doi.org/10.7914/SN/CN>), Global Seismograph Network – IRIS/USGS (IU; <https://doi.org/10.7914/SN/IU>), USGS Northern California Network (NC; <https://doi.org/10.7914/SN/NC>), United States National Strong-Motion Network (NP; <https://doi.org/10.7914/SN/NP>), NEPTUNE Canada (NV; no doi available), University of Oregon and Pacific Northwest Seismic Network (UO; <https://doi.org/10.7914/SN/UO>), Ocean Observatories Initiative (OO; <https://doi.org/10.7914/SN/OO>), United States National Seismic Network (US; <https://doi.org/10.7914/SN/US>), Pacific Northwest Seismic Network (UW; <https://doi.org/10.7914/SN/UW>). Data were also obtained from the Southern California Earthquake Center (<https://doi.org/10.7909/C3WD3xH1>) and IRIS (<http://ds.iris.edu/ds/nodes/dmc/data/>). Maps (Figures 1 and 7) were created with Matlab. Waveform gain and instrument response information used to produce Figure 2 were obtained via the IRIS-WS Library Java API (IRIS-WS v. 2.0.18), which allows direct access to data and information at the IRIS Data Management Center (DMC; data last accessed April 2020). The WCTC data, metadata, and waveform data (miniseed) were downloaded from the ShakeAlert Testing Event Archive ftp site housed at the Southern California Earthquake Data Center (<http://service.scedc.caltech.edu/ftp/TestingEventArchive>; last accessed October 2019). Data with incorrect gains (NC.CSL, NC.CVP; NC.PCH, NC.PCHB, NC.PCHC, NC.PHOB, and NC.PMM) that were identified in early 2020 were removed (Deborah Smith, written communication 2020). ShakeAlert station network information for January 2020 was obtained from Valerie Thomas (written communication, 2020).

Acknowledgments

This work benefited from helpful discussions with those from the extended ShakeAlert team and ShakeAlert working groups. The authors thank Jessica Murray and Andy Barbour for detailed feedback on an initial draft of this manuscript. The authors also thank Editor Rachel Abercrombie and an anonymous associate editor and two anonymous reviewers for their thoughtful suggestions that greatly improved the manuscript. Additional feedback from Jeff McGuire, Shane Detweiler and Mike Diggles also improved the manuscript. Any use of trade, firm, or product names is for descriptive purposes only and does not imply endorsement by the U.S. Government.

References

- Agnew, D. C., & Wyatt, F. K. (2014). Dynamic strains at regional and teleseismic distances. *Bulletin of the Seismological Society of America*, 104, 1846–1859. <https://doi.org/10.1785/0120140007>
- Allen, R. M., Brown, H., Hellweg, M., Khainovski, O., Lombard, P., & Neuhauser, D. (2009). Real-time earthquake detection and hazard assessment by ElarmS across California. *Geophysical Research Letters*, 36, L00B08. <https://doi.org/10.1029/2008GL036766>
- Allen, R. M., Cochran, E. S., Huggins, T., Miles, S., & Otegui, D. (2018). Lessons in public alerting from Mexico's earthquake early warning system. *Eos*, 99. <https://doi.org/10.1029/2018EO105095>
- Allen, R. M., Gasparini, P., Kamiguchi, O., & Bose, M. (2009). The status of earthquake early warning around the world: An introductory overview. *Seismological Research Letters*, 80, 682–693. <https://doi.org/10.1785/gssrl.80.5.682>
- Allen, R. M., & Kanamori, H. (2003). The potential for earthquake early warning in southern California. *Science*, 300, 786–789. <https://doi.org/10.1126/science.1080912>
- Allen, R. M., Kong, Q., & Martin-Short, R. (2019). The MyShake Platform: A Global Vision for Earthquake Early Warning. *Pure and Applied Geophysics*, 177, 1–1712. <https://doi.org/10.1007/s00024-019-02337-7>

- Allen, R. M., & Melgar, D. (2019). Earthquake early warning: Advances, scientific challenges, and societal needs. *Annual Review of Earth and Planetary Sciences*, 47, 361–388. <https://doi.org/10.1146/annurev-earth-053018-060457>
- Allen, T. I., Wald, D. J., & Worden, C. B. (2012). Intensity attenuation for active crustal regions. *Journal of Seismology*, 16, 409–433. <https://doi.org/10.1007/s10950-012-9278-7>
- Atkinson, G. M., & Boore, D. M. (2011). Modifications to existing ground-motion prediction equations in light of new data. *Bulletin of the Seismological Society of America*, 101, 1121–1135. <https://doi.org/10.1785/0120100270>
- Bent, A. L., Côté, T. J., Seywerd, H. C., McCormack, D. A., & Coyle, K. A. (2020). The Canadian National Seismograph Network: Upgrade and status. *Seismological Research Letters*, 91(2A), 585–592. <https://doi.org/10.1785/0220190202>
- Beyreuther, M., Barsch, R., Krischer, L., Megies, T., Behr, Y., & Wassermann, J. (2010). ObsPy: A Python toolbox for seismology. *Seismological Research Letters*, 81, 530–533. <https://doi.org/10.1785/gssrl.81.3.530>
- Boore, D. M., & Atkinson, G. M. (2008). Ground-motion prediction equations for the average horizontal component of PGA, PGV, and 5%-damped PSA at spectral periods between 0.01 s and 10.0 s. *Earthquake Spectra*, 24, 99–138. <https://doi.org/10.1193/1.2830434>
- Böse, M., Smith, D. E., Felizardo, C., Meier, M.-A., Heaton, T. H., & Clinton, J. F. (2018). FinDer v2: Improved real-time ground-motion predictions for M2–M9 with seismic finite-source characterization. *Geophysical Journal International*, 212, 725–742. <https://doi.org/10.1093/gji/ggx430>
- Chiou, B. J., & Youngs, R. R. (2008). An NGA model for the average horizontal component of peak ground motion and response spectra. *Earthquake Spectra*, 24(1), 173–215. <https://doi.org/10.1193/1.2894832>
- Chung, A. I., Henson, I., & Allen, R. M. (2019). Optimizing earthquake early warning performance: ElarmS-3. *Seismological Research Letters*, 90, 727–743. <https://doi.org/10.1785/0220180192>
- Chung, A. I., Meier, M. A., Andrews, J., Böse, M., Crowell, B. W., McGuire, J. J., & Smith, D. E. (2020). ShakeAlert earthquake early warning system performance during the 2019 Ridgecrest earthquake sequence. *Bulletin of the Seismological Society of America*, 110, 1904–1923. <https://doi.org/10.1785/0120200032>
- Cochran, E. S., Bunn, J., Minson, S. E., Baltay, A. S., Kilb, D. L., Kodera, Y., & Hoshiba, M. (2019). Event Detection Performance of the PLUM Earthquake Early Warning Algorithm in Southern California. *Bulletin of the Seismological Society of America*, 109, 1524–1541. <https://doi.org/10.1785/0120180326>
- Cochran, E. S., & Husker, A. L. (2019). How low should we go when warning for earthquakes? *Science*, 366, 957–958. <https://doi.org/10.1126/science.aaz6601>
- Cochran, E. S., Kohler, M. D., Given, D. D., Guiwits, S., Andrews, J., Meier, M. A., et al. (2018). Earthquake early warning ShakeAlert system: Testing and certification platform. *Seismological Research Letters*, 89, 108–117. <https://doi.org/10.1785/022017014010.1785/0220170138>
- Dengler, L. A., & Dewey, J. W. (1998). An intensity survey of households affected by the Northridge, California, earthquake of 17 January 1994. *Bulletin of the Seismological Society of America*, 88, 441–462.
- Dewey, J. W., Reagor, B. G., Dengler, L., & Moley, K. (1995). *Intensity distribution and isoseismal maps for the Northridge, California, earthquake of January 17, 1994*. Open-File Report. U.S. Geological Survey. (Vol. 95(92), p. 35).
- García, D., Mah, R. T., Johnson, K. L., Hearne, M. G., Marano, K. D., Lin, K. W., et al. (2012). ShakeMap Atlas 2.0: An improved suite of recent historical earthquake ShakeMaps for global hazard analyses and loss model calibration. In *World conference on earthquake engineering*.
- Given, D. D., Allen, R. M., Baltay, A. S., Bodin, P., Cochran, E. S., Creager, K., et al. (2018). *Revised technical implementation plan for the ShakeAlert system—An earthquake early warning system for the West Coast of the United States*. Open-File Report 2018-1155. U.S. Geological Survey.
- Given, D. D., Cochran, E. S., Heaton, T., Hauksson, E., Allen, R., Hellweg, P., et al. (2014). *Technical implementation plan for the ShakeAlert production system: An earthquake early warning system for the west coast of the United States*. Open-File Rept. 2014-1097. U.S. Geological Survey.
- Grapenthin, R., Johanson, I., & Allen, R. M. (2014). The 2014 M_w 6.0 Napa earthquake, California: Observations from real-time GPS-enhanced earthquake early warning. *Geophysical Research Letters*, 41, 8269–8276. <https://doi.org/10.1002/2014GL061923>
- Hartog, R. J., Kress, V. C., Malone, S. D., Bodin, P., Vidale, J. E., & Crowell, B. W. (2016). Earthquake early warning: ShakeAlert in the Pacific Northwest. *Bulletin of the Seismological Society of America*, 106, 1875–1886. <https://doi.org/10.1785/0120150261>
- Heaton, T. H. (1985). A model for a seismic computerized alert network. *Science*, 228, 987–990. <https://doi.org/10.1126/science.228.4702.987>
- Hoshiba, M. (2013). Real-time prediction of ground motion by Kirchhoff-Fresnel boundary integral equation method: Extended front detection method for earthquake early warning. *Journal of Geophysical Research: Solid Earth*, 118(3), 1038–1050. <https://doi.org/10.1002/jgrb.50119>
- Hoshiba, M., Iwakiri, K., Hayashimoto, N., Shimoyama, T., Hirano, K., Yamada, Y., et al. (2011). Outline of the 2011 off the Pacific coast of Tohoku Earthquake (M_w 9.0)—Earthquake Early Warning and observed seismic intensity. *Earth Planets and Space*, 63, 7–551. <https://doi.org/10.5047/eps.2011.05.031>
- Hoshiba, M., Kamigaichi, O., Saito, M., Tsukada, S., & Hamada, N. (2008). Earthquake early warning starts nationwide in Japan. *Eos*, 89, 73–75. <https://doi.org/10.1029/2008EO080001>
- Hsiao, N. C., Wu, Y. M., Shin, T. C., Zhao, L., & Teng, T. L. (2009). Development of earthquake early warning system in Taiwan. *Geophysical Research Letters*, 36. <https://doi.org/10.1029/2008GL036596>
- Hsu, H. C., Chen, D. Y., Tseng, T. L., Wu, Y. M., Lin, T. L., & Pu, H. C. (2018). Improving Location of Offshore Earthquakes in Earthquake Early Warning System. *Seismological Research Letters*, 89, 1101–1107. <https://doi.org/10.1785/0220170212>
- Kodera, Y. (2019). An earthquake early warning method based on Huygens principle: Robust ground motion prediction using various localized distance-attenuation models. *Journal of Geophysical Research: Solid Earth*, 124, 12981–12996. <https://doi.org/10.1029/2019JB017862>
- Kodera, Y., Hayashimoto, N., Moriawaki, K., Noguchi, K., Saito, J., Akutagawa, J., et al. (2020). First-year performance of a nationwide earthquake early warning system using a wavefield-based ground-motion prediction algorithm in Japan. *Seismological Research Letters*, 91, 826–834. <https://doi.org/10.1785/0220190263>
- Kodera, Y., Saitou, J., Hayashimoto, N., Adachi, S., Morimoto, M., Nishimae, Y., & Hoshiba, M. (2016). Earthquake early warning for the 2016 Kumamoto earthquake: Performance evaluation of the current system and the next-generation methods of the Japan Meteorological Agency. *Earth Planets and Space*, 68. <https://doi.org/10.1186/s40623-016-0567-1>
- Kodera, Y., Yamada, Y., Hirao, K., Tamaribuchi, K., Adachi, S., Hayashimoto, N., et al. (2018). The Propagation of Local Undamped Motion (PLUM) method: A simple and robust seismic wavefield estimation approach for earthquake early warning. *Bulletin of the Seismological Society of America*, 108(2), 983–1003. <https://doi.org/10.1785/0120170085>

- Kohler, M. D., Cochran, E. S., Given, D., Guiwits, S., Neuhauser, D., Henson, I., et al. (2018). Earthquake early warning ShakeAlert system: West coast wide production prototype. *Seismological Research Letters*, 89, 99–107. <https://doi.org/10.1785/0220170140>
- Kohler, M. D., Smith, D. E., Andrews, J., Chung, A. I., Hartog, R., Henson, I., et al. (2020). Earthquake early warning ShakeAlert 2.0: Public rollout. *Seismological Research Letters*, 91, 1763–1775. <https://doi.org/10.1785/0220190245>
- Kong, Q., Allen, R. M., Schreier, L., & Kwon, Y. W. (2016). MyShake: A smartphone seismic network for earthquake early warning and beyond. *Science Advances*, 2, e1501055. <https://doi.org/10.1126/sciadv.1501055>
- Kong, Q., Inbal, A., Allen, R. M., Lv, Q., & Puder, A. (2019). Machine learning aspects of the MyShake global smartphone seismic network. *Seismological Research Letters*, 90, 546–552. <https://doi.org/10.1785/0220180309>
- Krischer, L., Megies, T., Barsch, R., Beyreuther, M., Lecocq, T., Caudron, C., & Wassermann, J. (2015). ObsPy: A bridge for seismology into the scientific Python ecosystem. *Computational Science & Discovery*, 8, 014003. <https://doi.org/10.1088/1749-4699/8/1/014003>
- Li, Z., Hauksson, E., & Andrews, J. (2019). Methods for amplitude calibration and orientation discrepancy measurement: Comparing co-located sensors of different types in the southern California seismic network. *Bulletin of the Seismological Society of America*, 109, 1563–1570. <https://doi.org/10.1785/0120190019>
- McBride, S. K., Becker, J. S., & Johnston, D. M. (2019). Exploring the barriers for people taking protective actions during the 2012 and 2015 New Zealand ShakeOut drills. *International Journal of Disaster Risk Reduction*, 37, 101150. <https://doi.org/10.1016/j.ijdrr.2019.101150>
- McBride, S. K., Bostrom, A., Sutton, J., de Groot, R. M., Baltay, A. S., Terbush, B., et al. (2020). Developing post-alert messaging for ShakeAlert, the earthquake early warning system for the west coast of the United States of America. *International Journal of Disaster Risk Reduction*, 50, 101713. <https://doi.org/10.1016/j.ijdrr.2020.101713>
- Meier, M. A. (2017). How “good” are real-time ground motion predictions from earthquake early warning systems? *Journal of Geophysical Research: Solid Earth*, 122, 5561–5577. <https://doi.org/10.1002/2017jb014025>
- Meier, M. A., Kodera, Y., Böse, M., Chung, A., Hoshiba, M., Cochran, E., et al. (2020). How often can earthquake early warning systems alert sites with high-intensity ground motion? *Journal of Geophysical Research: Solid Earth*, 125. <https://doi.org/10.1029/2019JB017718>
- Minson, S. E., Baltay, A. S., Cochran, E. S., Hanks, T. C., Page, M. T., McBride, S. K., et al. (2019). The limits of earthquake early warning accuracy and best alerting strategy. *Scientific Reports*, 9(1), 1–13. <https://doi.org/10.1038/s41598-019-39384-y>
- Minson, S. E., Baltay, A. S., Cochran, E. S., McBride, S. K., & Milner, K. R. (2021). Shaking is almost always a Surprise: The earthquakes that produce significant ground motion. *Seismological Society of America*, 92(1), 460–468. <https://doi.org/10.1785/0220200165>
- Minson, S. E., Meier, M.-A., Baltay, A. S., Hanks, T. C., & Cochran, E. S. (2018). The limits of earthquake early warning: Timeliness of ground motion estimates. *Science Advances*, 4, eaag0504. <https://doi.org/10.1126/sciadv.aag0504>
- Minson, S. E., Saunders, J. K., Bunn, J. J., Cochran, E. S., Baltay, A. S., Kilb, D. L., et al. (2020). Real-time performance of the PLUM earthquake early warning method during the 2019 Ridgecrest earthquake sequence. *Bulletin of the Seismological Society of America*, 110, 1887–1903. <https://doi.org/10.1785/0120200021>
- Ogwen, L. P., Withers, M. M., & Cramer, C. H. (2019). Earthquake early warning feasibility study for the New Madrid seismic zone. *Seismological Research Letters*, 90, 1377–1392. <https://doi.org/10.1785/0220180163>
- Ruhl, C. J., Melgar, D., Chung, A. I., Grapenthin, R., & Allen, R. M. (2019). Quantifying the value of real-time geodetic constraints on earthquake early warning using a global seismic and geodetic dataset. *Journal of Geophysical Research: Solid Earth*, 124, 3819–3837. <https://doi.org/10.1029/2018JB016935>
- Stover, C. W., & Coffman, J. L. (1993). *Seismicity of the United States, 1568-1989 (revised)*. US Government Printing Office.
- Strauss, J. A., & Allen, R. M. (2016). Benefits and costs of earthquake early warning. *Seismological Research Letters*, 87, 765–772. <https://doi.org/10.1785/0220150149>
- Sutton, J., Fischer, L., James, L. E., & Sheff, S. E. (2020). Earthquake early warning message testing: Visual attention, behavioral responses, and message perceptions. *International Journal of Disaster Risk Reduction*, 49, 101664. <https://doi.org/10.1016/j.ijdrr.2020.101664>
- Thakoor, K., Andrews, J., Hauksson, E., & Heaton, T. (2019). From earthquake source parameters to ground-motion warnings near you: The ShakeAlert earthquake information to ground-motion (eqInfo2GM) method. *Seismological Research Letters*, 90, 1243–1257. <https://doi.org/10.1785/0220180245>
- Thompson, E. M., McBride, S. K., Hayes, G. P., Allstadt, K. E., Wald, L. A., Wald, D. J., et al. (2020). USGS near-real-time products—And their use—For the 2018 Anchorage earthquake. *Seismological Research Letters*, 91(1), 94–113. <https://doi.org/10.1785/0220190207>
- Wald, D. J. (2020). Practical limitations of earthquake early warning. *Earthquake Spectra*, 36, 1447, 8755293020911388. <https://doi.org/10.1177/8755293020911388>
- Wald, D. J., Quitoriano, V., Dengler, L. A., & Dewey, J. W. (1999). Utilization of the Internet for rapid community intensity maps. *Seismological Research Letters*, 70(6), 680–697. <https://doi.org/10.1785/gssrl.70.6.680>
- Worden, C. B., Gerstenberger, M. C., Rhoades, D. A., & Wald, D. J. (2012). Probabilistic relationships between ground-motion parameters and Modified Mercalli Intensity in California. *Bulletin of the Seismological Society of America*, 102, 204–221. <https://doi.org/10.1785/0120110156>
- Worden, C. B., & Wald, D. J. (2016). *ShakeMap documentation*. U.S. Geological Survey. Available at <https://usgs.github.io/shakemap/>