Improving Earthquake Monitoring and Early Warning Using GNSS Velocities and Machine Learning

by

S. Timothy Dittmann

B.S., Princeton University, 2005

M.S., University of Colorado, 2019

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Aerospace Engineering

2023

Committee Members:

Y. Jade Morton, Chair

R. Steven Nerem

Kathleen Hodgkinson

Yang Wang

David Mencin

Dittmann, S. Timothy (Ph.D., Aerospace Engineering Sciences)

Improving Earthquake Monitoring and Early Warning Using GNSS Velocities and Machine Learning

Thesis directed by Prof. Y. Jade Morton

Earthquake ground motions intersecting a population and seismically-triggered tsunamis have cost over 800,000 lives globally in the last 20 years. Distributed measurements of earthquake ground motions: 1) diagnose shaking intensity for rapid disaster response, 2) can alert to minimize damage ahead of the most destructive shaking and 3) are fundamental to understanding past events to inform future preparedness. Nearfield, conventional inertial seismic instruments saturate during the largest, most destructive earthquakes and may have limited regional availability. High-rate GNSS ground-station networks offer an alternative source of unsaturated ground motion measurements of medium to larger earthquakes. However, elevated noise levels of relative motion from space borne signals and minimal high-rate, larger magnitude event catalogs have limited the contribution of GNSS for current operational seismic monitoring, alerting and research.

This thesis builds upon previous GNSS seismology research to address this gap between sensitivity and current functional range through a data-driven approach. A hemispheric network noise comparison determined time differenced carrier phase velocities is the geodetic processing method most sensitive to seismic signals. This method does not require external corrections and is more computationally efficient for our signal of interest at the highest rates and potentially on the network edge. This thesis then presents a supervised random forest classifier that outperformed existing detection methods when trained and tested on a catalog of high-rate GNSS velocity seismic waveforms to discriminate between signal and noise. This classifier can be run with minimal latency at high rates for robust stand-alone seismic-event detection. Lastly, zero-baseline inertial waveforms were augmented with stochastic GNSS

noise time series to expand the GNSS seismic catalog. An expanded catalog improved generalization and will enable deeper learning. The analysis and models presented in this thesis lay a foundation for components of the next generation geodetic network.

Dedication

This thesis is dedicated with love to Levi, Cosmo, Jane and Katherine.

Acknowledgements

With gratitude for committee members Steve, Yang, Kathleen and Dave, for your commitment to critically mentor me through this journey.

With gratitude for my colleagues at UNAVCO and the CU SenSE Lab, for your fellowship, inspiration and support.

With gratitude for my parents and sisters, for instilling a responsibility to learning and "thoroughly finishing" any endeavor.

With gratitude for Cosmo, Jane and Levi, for your patience with your dad and grounding me with love.

With gratitude for Jade, for your wisdom, patient enthusiasm and devotion to your students and colleagues.

With gratitude for Katherine, for pushing our family to live large and dream bigger, and loving us when we fall short.

Contents

Chapter

Tables

Table

Figures

Figure

Chapter 1

Introduction

Earthquakes have cost the lives of over three quarters of a million people worldwide in the last two decades [NGD,], and cost the United States building stock alone $6.1B US dollars a year [(FEMA) et al., 2017]. It is estimated that over 300,000 people lost lives from the 2011 M7.0 Haiti earthquake alone [NGD, , Kam et al., 2021]. The majority of these costs of life and property result from the impacts of the most destructive coseismic ground motions from the largest events intersecting a population or infrastructure that have a range of seismic resilience. In the event of a large earthquake, faithfully observing coseismic ground motions in space and time offers information for targeting post disaster responses and irreplaceable knowledge for preparing for future hazards. Additionally, if available with limited latency, these observations enable a class of extremely rapid hazard alerting to reduce impacts known as Earthquake Early Warning (EEW) [Allen and Melgar, 2019]. Unsaturated ground motion measurements are a tool for addressing seismic hazards in the recent past, in real-time, and in the future.

Ground motion measurements historically have been made by zero-baseline inertial instruments with a range of signal bandwidths predominantly engineered for minimum detectability. The effective cost of this instrumental sensitivity combined with sensor physical limitations is that they struggle to capture large amplitude critical information that differentiates the most damaging events. The global proliferation of real-time Global Navigation Satellite System (GNSS) networks offers an additional source of these vital measurements

that has unbound upper geophysical dynamic range. However, the relatively high noise levels of GNSS measurements limit their sensitivity and detection methods adopted from seismology trigger unreasonable amounts of false alerts. This in turn has limited their contributions to global seismic monitoring. We address these limitations through data analysis and data-driven machine learning model development. The processing assessment, detection model development and data catalog assembled here will expand GNSS's contribution to rapid hazard monitoring and assessment.

1.1 Research Background

1.1.1 Seismic Ground Motions

A ground-motion sensor experiencing an earthquake encounters a series of phases: First arrivals are body waves, including p- or primary compressional waves, followed by s- or secondary shear waves (Figure 1.1). Subsequent surface waves include energy transferred along the upper crust in the form of Rayleigh or Love waves [Aki et al., 2002]. These waves have the largest amplitudes and are typically the most destructive. The timing, frequency and amplitude of these phases are a function of source parameters, such as slip distribution and fault geometry, and the crustal medium the waves propagate through.

Observing and identifying these phases is critical to earthquake characterization and used for a variety of hazard applications. Many of these applications and observations are tightly coupled, where the accuracy of derivative applications is dependent on the quantity, quality and distribution of ground motion observations available. For example, prior to an event seismic ground motion propagation models are generated empirically from existing catalogs of these peak ground motion measurements [Ancheta et al., 2014]. During and immediately after an event, teleseismic body and surface wave measurements are joined with these local-to-regional observations and fed into finite fault model inversions [Goldberg et al., 2022]. These finite fault models [Hayes, 2017] are produced to estimate slip distribution

Figure 1.1: Prototypical seismic waveform [MTU, 2022]

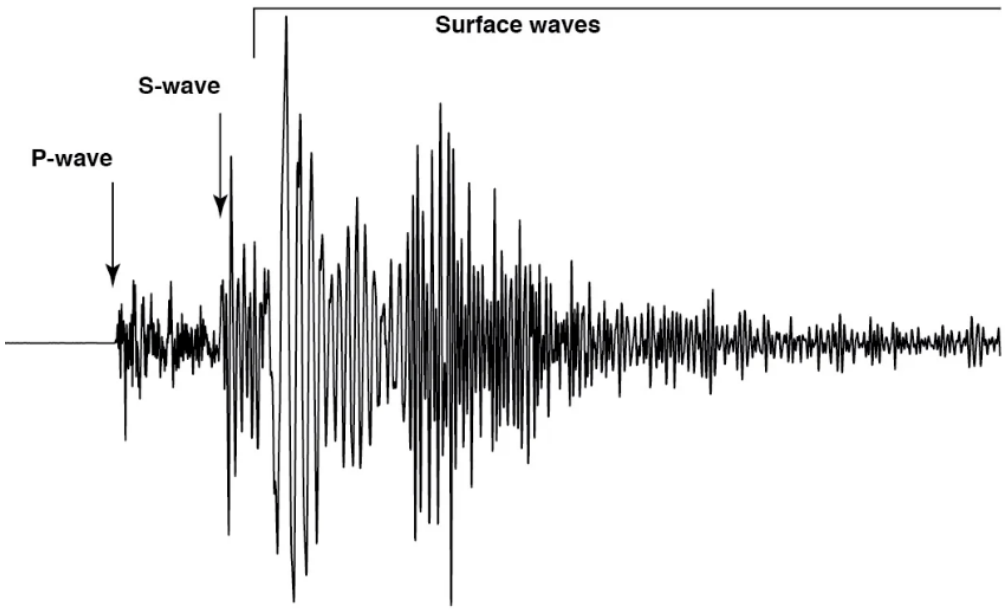

along a strained fault as a function of space and time (Figure 1.2-b). A ground motion map, such as a ShakeMap [Worden et al., 2020], applies such a ground propagation model to a finite fault model to estimate spatial distribution of peak ground motion dynamics (Figure 1.2-a). These predictions are interpolated between any additional in situ measurements observed during the event. A Shakemap overlaid on a population profile in combination with existing building codes is a critical product for diagnosing the extent of shaking and damage. In summary: additional historical measurements improve propagation models, and more current measurements improve source characterization estimates and propagation observations– the more observations, both historical and current, the more accurate the product to inform effective response and future actions.

Figure 1.2: (a) is a prototypical Shakemap and color scheme intensity legend. (credit: USGS). (b) is a prototypical finite fault model using GNSS and INSAR inputs. (credit: USGS)

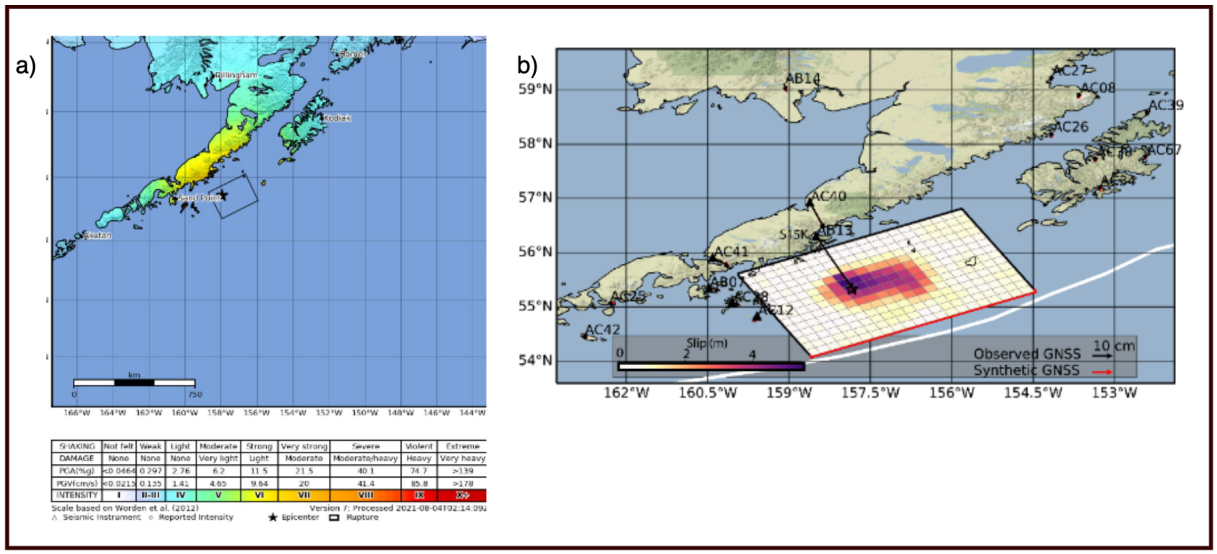

Additionally, information encoded in these earthquake phases is also required for earthquake early warning (EEW). EEW is the "the delivery of ground shaking alerts" [Allen and Melgar, 2019]; alert implying that notice is given in advance of the ground shaking of concern. The societal and economic benefits of an effective system grossly outweigh the costs

[Strauss and Allen, 2016], particularly if additional source instrumentation infrastructure is subsidized by alternative interests. Current operational approaches [Douglas D. Given, 2018] for EEW consist of rapid source characterization by way of finite fault [Böse et al., 2017, Crowell et al., 2016] or point source [Chung et al., 2019] algorithms. Alternatively, a class of alerts from propagating observed ground motions [Kodera et al., 2018] avoid source estimation altogether. Each strategy proposed has various advantages and disadvantages, but all rely on accurate empirical ground motion models and in situ low-latency unsaturated ground motion measurements. Furthermore, while the issue of determinism, or at which phase within an incoming seismic waveform final earthquake magnitude can be determined, is still debated [Meier et al., 2016, Rydelek and Horiuchi, 2006, Melgar et al., 2019], it is well understood that differentiating between the largest magnitude events requires longer period signals [Goldberg et al., 2019] that are difficult to extract from these sensitive inertial measurements due to clipping or baseline errors [Bock and Melgar, 2016]. For these events, their ground motions and tsunamigenic potential make them the most deadly and destructive. In short, all sensors are not equivalent but complementary with respect to their EEW contribution, and the most destructive events require a specific dynamic range outside of much of the inertial network infrastructure.

1.1.2 Ground Motion Instruments

Historically, ground motion observations are made exclusively by one of several types of seismic instruments amongst a broader class of inertial sensors. This mature technology has evolved from masses on a pendulum to electronic negative feedback loops capable of extended dynamics [IRIS, 2021]. For the medium to larger event hazard monitoring, these are typically either a broadband seismometer or a strong motion accelerometer. Broadband seismometers offer wide dynamic ranges and weaker signal sensitivity, but will clip in the nearfield of larger amplitude signals [Larson et al., 2003]. Strong motion accelerometers do not suffer from digital clipping, but require a single or double integration to get to velocity

or displacement, which proves problematic in recovering critical longer period signals from larger magnitude events due to instrument baseline effects [Melgar et al., 2013, Trifunac and Todorovska, 2001]. beginning in the 1990's Global Positioning System (GPS) supported civilian signals and began answering long period geophysical research questions about plate velocities and reference frame definitions. A variety of processing techniques converted raw GPS phase and range measurements into absolute or relative positions, the precision of which benefited from averaging over longer 24-hour arcs. Depending on the re-sampling interval, non-continuous measurements could measure the offset from an earthquake or the longer term velocity of plate motion. Scientific returns on investment combined with decreases in hardware costs lead to global, continuous permanent geodetic networks being constructed.

As receiver and processing technology advanced, higher rate ($\geq$1Hz) GNSS processed kinematically captured seismic waveform dynamics [Nikolaidis et al., 2001, Larson et al., 2003]. Remote communication telemetry also progressed, and real-time streaming of raw observations allowed for low-latency capture of static offsets and peak dynamics for finite fault models [Crowell et al., 2009, Allen and Ziv, 2011], magnitude estimation [Melgar et al., 2015, Ruhl et al., 2018], and inclusion in operational early warning systems [Crowell et al., 2016, Murray et al., 2018]. Expansion to multi-national positioning satellite constellations (collectively termed Global Navigation Satellite Systems, or GNSS) transmitting on multiple frequencies and signal structures have resulted in exponential observation growth. This space asset and signal proliferation has been largely driven by the needs of the position, navigation and timing infrastructure that is now part of the global critical infrastructure, from cadastral surveying, machine localization and financial institution timing to 1.5B/year smartphone GPS/GNSS chipsets. The GNSS industry and derivatives has a market revenue of €200B with 6.5B devices in use [EUSPA, 2022]. GNSS networks have proliferated globally, supported by a variety of non-seismic commercial, government and research applications. Any available high rate observational stream can become a seismometer.

Inclusion of GNSS amongst operational seismic observations offers two distinct bene-

fits. In most parts of the world, the existing seismic and geodetic infrastructure were built independently for different research, hazards or commercial applications. Certain regions of the globe (e.g. Japan, California) have extensive seismic networks that have operated for close to a century, but that is far from the norm along global seismic hazard zones that might intersect populations or are tsunamigenic. Real-time, higher rate GNSS can spatially complement or act independently in areas where seismic networks are sparse [Grapenthin et al., 2017]. Secondly, GNSS overcomes instrument and early phase saturation issues in EEW. It cannot be understated the effect this has had in recent history [Hoshiba and Iwakiri, 2011, Colombelli et al., 2013] and will continue dire consequences if not accounted for. For these reasons GNSS has been included in operational ShakeAlert, but dependent on a (accurate) seismic trigger, without a robust means of separating true signal from noise.

1.1.3 Machine Learning

Data driven, machine learned models address earth science questions that are difficult to represent in physics-based models [Bergen et al., 2019]. Earth scientists have benefited from adopting algorithms developed for ulterior motivations such as natural language processing, computer vision and forecasting to answer novel research questions, draw new inferences and optimize scientific systems [Jordan and Mitchell, 2015]. Most ML methods can be grouped into a few broad classes:

- Supervised learning involves a learned mapping from a labeled data catalog, whereas unsupervised learning applies statistical principles to unlabeled samples to draw inferences or discover patterns.

- Classification consists of categorical labels while regression models target continuous variables.

- Deep learning typically refers to layers of learned features (e.g. RNN, CNN, ANN), whereas conventional or "shallow" ML typically refers to domain-experts (humans)

Figure 1.3: Station AB13 in Chignik, Alaska, an example of a permanent continuous GNSS station used in this study (credit: UNAVCO/Ellie Boyce)

curated feature extraction for a variety of models (e.g. support vector machines, random forest, ridge regression, etc) [LeCun et al., 2015]. The depth adjective generally refers to the overall model structure, but semantic variations appear common.

[Bergen et al., 2019] outlines more detailed descriptions of models commonly used in earth sciences, and [Kong et al., 2019] and [Mousavi and Beroza, 2022] review seismological questions addressed with ML and deep learning, respectively. We evaluate ML for our research problem with two primary goals: 1) we demonstrate that a data-driven approach outperforms the current "state of the art," model driven approach and 2) we provide explainable, interpretable results. We adopt ML not as a black box but as a tool to recast our problem in a different space to draw inferences, akin to a physicist adopting a Fourier transform or a sociologist adopting logarithmic transformations. As model complexity increases, or deepens,and supervision decreases, explainability may become harder to come by despite potentially promising results. This discussion is outside the scope of our work, but certainly something for scientists and engineers to consider as artificial intelligence explodes in the scientific zeitgeist.

1.2 The Research Problem

GNSS-seismology research over the last two decades since [Nikolaidis et al., 2001] and [Larson et al., 2003] demonstrated the usefulness of GNSS for coseismic applications, but few efforts exploit the full range of valid GNSS seismic measurements. Ingesting the extent of valid GNSS seismic observations involves embracing the challenging noise signatures present in these measurements [Melgar et al., 2020] that can interfere with or obscure the signals of interest. However, most current detection approaches adopt methods developed for traditional seismic instrumentation. Many applications use position displacement processing that relies on external corrections services, such as precise satellite clock and position information. Alternative processing methods have been proposed [Colosimo and Mazzoni,

2011, Francesca Fratarcangeli, 2018, Grapenthin and Freymueller, 2018], but their relative event detection sensitivities are not well understood. In addition, existing automatic algorithms must minimize high false alerting by cross-correlation of positive events with nearby stations [Colombelli et al., 2013, Ohta et al., 2012, Psimoulis et al., 2018] or time windowing their allowable search space based upon seismic triggers and estimated wave propagation speeds [Crowell et al., 2016]. This introduces points of failure and limits effectiveness, such as when an earthquake is mis-located or in very sparse networks. Finally, most have relatively high magnitude thresholds initiated by the seismic system, limiting the magnitude range of contributions. **There is a gap between the sensitivity of using GNSS as a seismometer, and the current functional range given existing signal detection methods. We aim to minimize this gap to enable expanded adoption of high-rate GNSS reference stations as integral sources of ground motion measurements to improve global seismic hazard monitoring.** We will do so by addressing the following research questions:

(1) What is the sensitivity range of detection for GNSS as a seismometer, and what is the optimal processing method to maximize this range?

(2) Without external inputs, how well can we detect these signals with respect to minimizing false alerts and missed detections?

(3) Can we expand the limited available data catalogs to allow more sophisticated models to further improve GNSS seismology contributions?

This thesis' objectives address these questions to expand the scope of integral satellite seismology measurements. For monitoring, these GNSS seismometer measurements are particularly valuable when existing sensors are sparse or when the most destructive and costly largest magnitude events occur. For research, an improved stand alone detection module can enhance highest rate data capture for future research. For further development, an expanded

labeled catalog will promote sophisticated methods to further improve GNSS contribution. The primary limitation of this data-driven approach, which is partially addressed within the work itself, is the availability of highest quality, classifiable data of interest. The long tail of larger magnitude earthquakes mean relatively few have been observed by GPS satellites launched in the 1990's. Furthermore, while optimal sampling rates for seismic signals are $\geq$1Hz, receiver and telemetry constraints limit the availability of these higher rate archives. Various geodetic archives had the foresight to collect what was considered high-rate data (5-10Hz) prior to many applications using it; we hope this work can enable the continuation of that model by triggering much higher rate data capture. Finally, labeling samples close to the noise floor is both time consuming and error prone. We address many of these issues in Chapter 4's work using synthetics and augmentation.

1.3 Research Overview

In chapter two, we compare the ambient noise of two processing methods to their respective earthquake magnitude scaling laws. We quantify the anticipated sensitivity as a function of magnitude and hypocentral distance, and compare these results as a function of processing method to determine the optimal approach for our research.

In chapter three, we generate a catalog of processed earthquake waveforms and ambient timeseries to train a random forest classifier on identifying seismic events amongst noise. We optimize this model on a balance of missed detection and false alerting, and compare this classifier to the current state of the art methods.

In chapter four, we address the limited data challenge of GNSS seismic waveforms. We model probabilistic GNSS processing noise to generate stochastic noise timeseries to add to inertial measurement data catalogs. We validate this approach by training a similar classifier on exclusively pseudo-synthetic, augmented GNSS waveforms and compare detection on real GNSS data relative to the limited real data approach.

In chapter five, we discuss the potential impact of the work presented in this thesis

and future directions that can be extended from this work.

Chapter 2

Comparing Sensitivities of Geodetic Processing Methods for Rapid Earthquake Magnitude Estimation

This chapter is adopted from the following manuscript:
Tim Dittmann, Kathleen Hodgkinson, Jade Morton, David Mencin, Glen S. Mattioli; Comparing Sensitivities of Geodetic Processing Methods for Rapid Earthquake Magnitude Estimation. Seismological Research Letters 2022;; 93 (3): 1497–1509.

2.1 Abstract

Rapid earthquake magnitude estimation from real-time space-based geodetic observation streams provides an opportunity to mitigate the impact of large and potentially damaging earthquakes by issuing low-latency warnings prior to any significant and destructive shaking. Geodetic contributions to earthquake characterization and rapid magnitude estimation have evolved in the last 20 years, from post-processed seismic waveforms to, more recently, improved capacity of regional geodetic networks enabled real-time Global Navigation Satellite System (GNSS) seismology using precise point position (PPP) displacement estimates. In addition, empirical scaling laws relating earthquake magnitude to peak ground displacement (PGD) at a given hypocentral distance have proven effective in rapid earthquake magnitude estimation, with an emphasis on performance in earthquakes larger than M6.5 where nearfield seismometers generally saturate. While the primary geodetic contributions to date in earthquake early warning (EEW) have focused on the use of 3-D position estimates and displacements, concurrent efforts in time differenced carrier phase (TDCP) de-

rived velocity estimates also have demonstrated that this methodology has utility, including similarly derived empirical scaling relationships. This study builds upon previous efforts in quantifying the ambient noise of three-component ground displacement and ground velocity estimates. We relate these noise thresholds to expected signals based on published scaling laws. Finally, we compare the performance of PPP-derived PGD to TDCP derived peak ground velocity (PGV) given several rich event datasets. Our results indicate that TDCP-PGV is more likely than PPP-PGD to detect intermediate magnitude ($M_w5.0 - M_w6.0$) earthquakes, albeit with greater magnitude estimate uncertainty and across smaller epicentral distances. We conclude that TDCP derived PGV magnitude estimation is a computationally lightweight, independently processed complement to PPP-derived PGD magnitude estimates, which could be produced at the network edge at high rates and with increased sensitivity to ground motion than current PPP estimates.

2.2 Introduction

2.2.1 GNSS Processing

The Global Navigation Satellite System (GNSS) consists of multiple medium-earth-orbit satellite constellations broadcasting precise timing L-band multi-frequency signals towards Earth. A GNSS receiver and antenna combination amplifies and demodulates these incoming low power signals to continuously track the broadcast signal and convert it to range and carrier phase measurements. These phase measurements are highly precise but contain unknown integer biases. Precise point positioning (PPP) [Zumberge et al., 1997] with ambiguity resolution, also known as PPP-AR [Bertiger et al., 2010], is a position estimation method in which these carrier phase biases are estimated to achieve uncertainties of several centimeters in horizontal antenna positions Many approaches exist to accomplish this level of precision, but most depend on reducing uncertainties through precise satellite clock error estimates, precise satellite orbit determination estimates, and multi-frequency carrier-phase

data to remove first-order ionospheric effects. This study uses PPP estimated in real-time by the Trimble RTX [R. Leandro, 2011] software. These positions are estimated once per second from 1 Hz GNSS observations streamed from the field receiver to a central network operations center (NOC) and are typically available with 1-2 seconds latency, where latency is defined as the time difference between observations encoded at the receiver and availability from the data center [Murray et al., 2019]. The low latencies for the real-time PPP solutions from the GNSS component of the Network of the Americas (NOTA) make them useful to safety of life and hazard applications [Mattioli et al., 2020]. Once at the NOC, positions are computed in a geodetic reference frame and translated to a topocentric reference frame as relative displacements from an a priori position. The RTX software is a commercial product using proprietary positioning software and data products, including the satellite orbits and clocks, and it has proven capable of tracking medium to large (M6.0+) seismic displacements [Hodgkinson et al., 2020].

A partially independent geodetic processing method using the same "raw" GNSS observations is known as the time-differenced carrier phase method (TDCP) [Graas and Soloviev, 2003], and was first proposed for capturing seismic waveforms by [Colosimo and Mazzoni, 2011]. This processing method estimates antenna velocity by differencing the phase observations between two epochs instead of solving for the unknown integer ambiguity at each epoch. This change in phase over time is equivalent to the pseudorange rate or doppler shift. By assuming no change in carrier phase integer ambiguity nor substantial epoch-wise dynamic atmospheric change, the integer ambiguity and atmospheric effects are effectively removed. For a complete description of this method, see [Misra and Enge, 2011, Colosimo and Mazzoni, 2011, Grapenthin and Freymueller, 2018, Crowell, 2021].

The novel VADASE (Variometric Approach for Displacements Analysis Stand-alone Engine) method [Colosimo and Mazzoni, 2011] integrates these velocities over time to estimate displacements. In our study, we complete our analysis in velocity space using the open-source SNIVEL package (SNIVEL; see Data and Resources). The SNIVEL software

uses the narrow lane, GPS-only, dual-frequency phase combination and currently does not account for tides, higher-order ionospheric effects, atmospheric gradients, or cycle slips, which are partially addressed below. Without the need for sophisticated corrections or additional earth models for real-time processing, a major advantage of the TDCP processing method is that it can be pushed to the network edge in real-time with a relatively low computational and financial cost.

2.2.2 Peak geodetic signals for rapid magnitude estimation

The ability for high rate GNSS to track co-seismic and dynamic displacements is well established [Larson et al., 2003, Genrich and Bock, 2006], and advances in network design, coupled with improved real-time processing, enabled relatively precise estimates of these waveforms in real-time and furthering their utility for hazards applications [Allen and Ziv, 2011]. [Crowell and Geng, 2013, Melgar et al., 2015, Ruhl et al., 2018] demonstrated the usefulness of the empirical relationship between peak ground displacement (PGD) estimated from higher rate, unsaturated PPP displacements, earthquake magnitude and distance from the seismic source using scaling law coefficients of log-linear regression models from catalogs of recent earthquakes. In addition, these higher rate estimates enable geodetic infrastructure to act as a strong motion instrument [Larson, 2009, Crowell, 2021] and complement existing accelerometer and seismometer early warning infrastructure.

[Fang and Liu, 2020] developed a similar log-linear regression model relating earthquake magnitude, hypocentral distance, and PGV estimated from time-differenced carrier phase velocities using a catalog of recorded recent earthquakes. The GNSS TDCP-derived velocities can be sensitive to compression (P) waves [Hohensinn and Geiger, 2018] and show good agreement with USGS ShakeMap velocities [Grapenthin and Freymueller, 2018, Crowell, 2021], further enhancing the value of this processing method for hazards research and real-time assessments by complementing existing seismic infrastructure and potentially filling geographic or temporal availability gaps. Here, we present a noise model of the three-

component displacements and velocities and compare these to the expected co-seismic signals based upon the existing empirical scaling laws. We then select a threshold to quantify and compare the probability of false alerts and probabilities of detection across the processing methods. In addition, we investigate possible station-dependent sources of noise correlated with station attributes. Finally, we compare these expected signals and noise with results from two actual events of different magnitudes.

2.3 Methods

2.3.1 Dataset

This study establishes a noise assessment of PPP-PGD and TDCP-PGV using a dataset of concurrent 1 Hz Trimble RTX PPP estimates and 1 Hz GPS raw observations at 592 stations distributed across North America from 16:00-17:00 UTC on 4 July 2019. The 2019 Ridgecrest earthquake sequence initiated at 17:50 UTC on this day with a M6.4 event. Larger seismic events (>M6.0) within the NOTA footprint initiate a data response to recover complete high-rate data files [Mattioli et al., 2020]. The time interval for this study was selected because of the availability of complete higher rate data due to the earthquake-initiated data response during a period of no known external dynamic signals. There is no prior seismic activity in the time window greater than M3.0 (as published by the Advanced National Seismic System Comprehensive Catalog of Earthquake Events and Products) within the NOTA footprint. In addition, space weather indices indicate calm space weather conditions, with the daily AP index reported at 4 nT, which implies a low ionospheric gradient at the NOTA primarily mid-latitude stations. We infer that the completeness of the NOTA dataset combined with the atmospheric and tectonic stability of the temporal epoch makes it optimal for 1 Hz ambient noise characterization. This study investigated different variance estimation periods by examining 24-hours of data from station P509 (Figure 2.1) and conclude that 1 hour was a reasonable window for the characterization of our noise estimates.

Figure 2.1: Twenty-four hour evaluation of variance of 1 hr noise estimates for station P509. The black dots are the 1 Hz precise point positioning (PPP) displacement and time-differenced carrier phase (TDCP) velocity estimates derived from the three-component waveforms. The yellow bars are the hourly noise estimates based upon a 99.5% confidence level upper threshold given an NCX2 distribution. The blue bar and orange bar are the displacement and velocity estimates, respectively, used in the temporal window of this study.

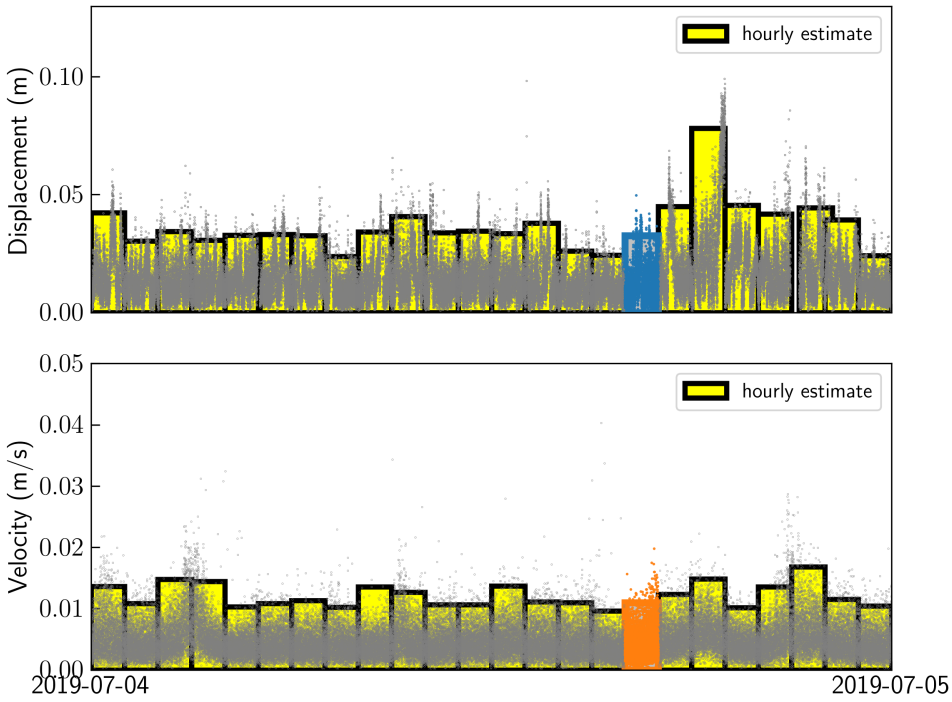

The ambient noise of the 1 Hz position displacement estimates is dependent on many factors, including the performance of the positioning algorithm, the GNSS constellations tracked and included in the processing, multipath noise, unmodeled atmospheric conditions, and local radio interference. Our PPP noise characterization assumes the same positioning algorithm and no physical displacement dynamics occur in the window of characterization, such that any relative variance in station-dependent PPP noise is assumed to be a function of raw observation quality including multipath, antenna or receiver performance, differences in signals tracked, and any unlikely but possible systematic station bias in corrections model errors. This assessment mitigates the latter impact by consistent windowing in time and space. Typical noise values are at the centimeter level. Previous studies have demonstrated the value of using only horizontal displacements in earthquake magnitude inversions to minimize the impact of the noisier vertical position estimates [Melgar et al., 2019]. For consistency, however, we use the three-dimensional scaling laws, which might prove invaluable for an event with significant vertical signals such as those expected from great earthquakes generated on the subduction interface. Ambient noise in TDCP velocity estimates is directly related to receiver carrier phase noise. Receiver manufacturers are able to minimize phase noise by using higher quality oscillators and advanced carrier phase tracking loop filtering, including radio frequency interference or signal multipath mitigation. Typical noise values are around a centimeter per second.

2.3.2 Establishing a robust 3 component noise threshold

By estimating a station-dependent, ambient noise threshold, [Hodgkinson et al., 2020] established NOTA real-time GNSS sensitivity to published PGD scaling law magnitude detection for earthquakes that occur within its spatial footprint. Each contributing station's three-component magnitude time series threshold was derived from a heuristic statistic of the time series median plus three times the time-series median absolute deviation (MAD) over a specified period. This station threshold was then used to represent the noise floor

of the three-component magnitude above which the signals of interest, peak displacement waveforms, can be detected at any NOTA real-time station. Peaks below the established station-specific threshold are considered noise and thus not included in any scaling law magnitude inversion estimate. Non-seismic peak ground motions above the threshold are considered false positives. We build upon the previous noise threshold approach established by [Hodgkinson et al., 2020] by more accurately representing the sampling distribution model as follows:

We observe Trimble RTX position time series topocentric displacements, $\delta_{north}, \delta_{east}, \delta_{up}$ and TDCP velocity components, $vel_{north}, vel_{east}, vel_{up}$ are normally distributed variables with respective non-zero means, $\mu_{north}, \mu_{east}, \mu_{up}$. The ground displacement (GD) and ground velocity (GV) are related to the sum of the squares of the normally distributed components:

$$GD = \sqrt{\delta_{east}^2 + \delta_{north}^2 + \delta_{up}^2} \tag{2.1}$$

$$GV = \sqrt{vel_{east}^2 + vel_{north}^2 + vel_{up}^2} \tag{2.2}$$

A random variable representing the sum of the squares of normally distributed, independent random variables can be approximated by the noncentral chi-squared distribution (NCX2) [Press, 1966], defined by k degrees of freedom equal to three and noncentrality parameter, lambda, equal to the sum of the square of the means of each component:

$$\lambda = \sum_{n=1}^{3} \mu_i^2 \tag{2.3}$$

We observe that the distribution of the sum of the squares of the components can be approximated by a chi-squared distribution. The aforementioned k and estimated λ values seed the Scipy scientific package for fitting a maximum likelihood estimate (MLE, See Data and Resources) of the scale and location of a distribution that fits each station's hour-long time series square of components summation. This allows us to characterize each station's

ambient noise by utilizing the fitted noncentral chi-squared probability distribution function (PDF):

$$f(x,k,\lambda) = \frac{1}{2}(e^{-\frac{\lambda+x}{2}})(x/\lambda)^{\frac{k-2}{4}} I_{(k-2)/2}(\sqrt{\lambda x}) \tag{2.4}$$

Where $I_{k-2/2}$ denotes the modified Bessel function of order $(k-2)/2$ (See Data and Resources: Scipy Package). We acknowledge that PPP displacements and TDCP velocities are not independent measurements, and thus this distribution is mathematically inaccurate. Empirically we find this distribution offers a more meaningful noise model over the previous heuristic threshold metric of Hodgkinson et al. (2020). Future work into applying a more sophisticated model, such as a parameterized gamma distribution [Ferrari, 2019], could improve accuracy.

Table 2.1: Probability of False Positive (FP) Rate Observed Over 24 Hr, in Which FP is the Percentage of FPs Out of the Total Number of Epochs in the Test Window.

Station	**Method**	**Hodgkinson FP**	**99.95 NCX2 FP**	**99.99 NCX2 FP**
P509	PPP-GD	2.6%	1.4%	0.5%
	TDCP-GV	2.8%	0.5%	0.2%
CLGO*	PPP-GD	2.7%	1.4%	0.5%
	TDCP-GV	3.5%	1.8%	0.9%
SC00	PPP-GD	3.0%	1.6%	0.7%
	TDCP-GV	3.0%	1.4%	0.6%
P041	PPP-GD	2.8%	1.5%	0.5%
	TDCP-GV	3.5%	1.9%	1.3%

†PPP-GD, precise point positioning-derived ground displacement; TDCP-GV, time-differenced carrier phase-derived ground velocity.
*CLGO was tested over a 22 hr arc.

The NCX2 PDF of the three-component squared summation enables a significance quantile, α, to be used as a threshold to balance the probability of detection, PD (or true positive), with a probability of false positive (FP) or false alert (PFA), where PFA is $\approx 1-\alpha$. Observed false positive rates exceed anticipated sample significance quantiles in Table 2.1; this is possibly a result of correlated noise that is not accounted for in this distribution.

Figure 2.2 illustrates our detection sensitivity balance: increasing α shifts the vertical red dashed line to the right, which has the effect of decreasing PFA (gray area in left panels), while increasing the probability of a missed detection or false negative (gray area in the right panels). These expected signals are based upon the normally distributed empirical scaling laws, using the coefficients of [Melgar et al., 2015, Fang and Liu, 2020]. Table 2.1 quantifies these false alert rates, with α as a tunable parameter for the decision balance in such a system. We consider our selection of the quantile threshold further in the next section.

2.3.3 Determining an optimal noise threshold

We evaluate a range of probabilities of the best fit NCX2 distribution to explore the relationship between PFA and PD (Figure 2.3) given a median noise model across the 592 stations. Assessing societal false-alert tolerance in combination with human and infrastructure costs of missed earthquake early warning is outside the scope of this work. We recognize an operational early warning system, however, would require a rigorous analysis for this impactful parameter selection [Minson et al., 2019]. In this study we select the 99.5% quantile for our methodology comparison. At any quantile level chosen, individual station false alerts within an operational EEW system could be mitigated by correlating with spatially adjacent GNSS and seismic equipment. The thresholds are applied consistently across both processing methods for this relative evaluation.

Our study uses the displacement scaling coefficients of [Melgar et al., 2015] for modeling the expected displacement signal at a given hypocentral distance and magnitude. This expected signal is then related to the station-dependent estimated displacement noise, the result of which we consider the sensitivity to an earthquake given its distance and magnitude. We used a windowing method similar to that of [Hodgkinson et al., 2020] to evaluate the ambient noise: for each site, a 1-hour window is divided into 15-minute segments, in which the median displacement of the first 300 seconds is used as the reference for the entire window. This reduces the effect of longer period drift present in RT-GNSS PPP solutions

Figure 2.2: Panels (a) and (c) are the displacement and velocity NCX2 noise models (panel a: PPP-derived peak ground displacement [PGD]; panel b: TDCP-derived peak ground velocity [PGV]) for an example station (P509) with the selected threshold of a 99.5% significance interval. The region of false alert is shaded in gray (PFA). Panels (b) and (d) are the probability distribution functions of a range of different magnitude events given normally distributed expected (b) PGD and (d) PGV signals and variances of the [Melgar et al., 2015, Fang and Liu, 2020] scaling laws at a fixed hypocentral distance of 100 km. The selected thresholds are mapped to displacement and velocity spaces to capture probability of detection (PD) for these events. Panels (b) and (d) illustrate TCDP velocity has higher PD for the medium events (M_w 5.0 and M_w6.0). They also indicate the larger magnitude events (M_w 8.0) gain signal at a higher rate in displacement space, evident with the relative spacing between the threshold and the expected signals.

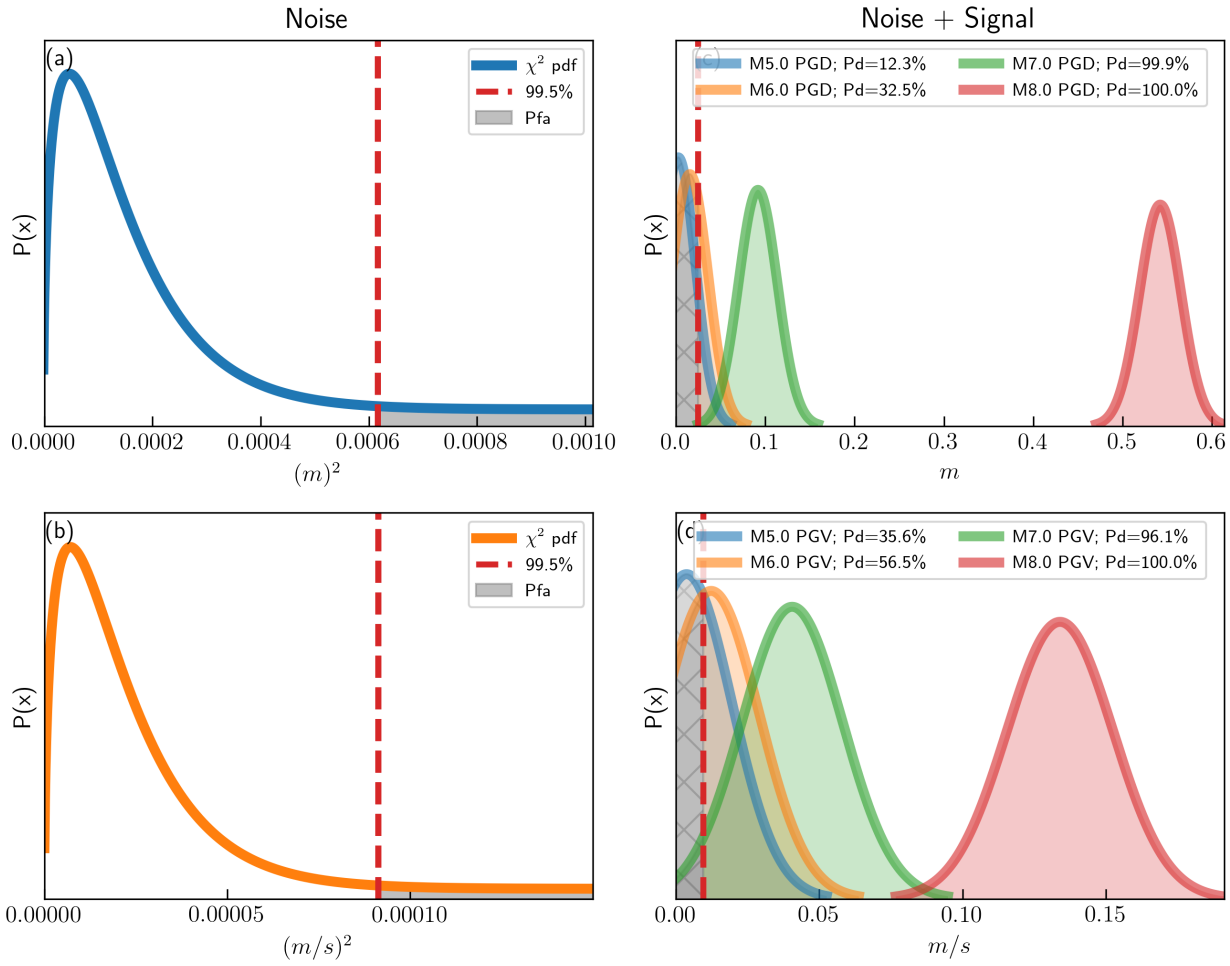

and is more comparable to a moving window process that would be implemented in real-time monitoring. This approach has an effect similar to removing the long-term average (LTA) used by traditional seismic event pickers, such as one applied to GNSS data by [Ohta et al., 2012]. We find fitting a distribution to the sample data is sporadically affected by large outliers; accordingly, we do a simple threshold mask and remove any non-geophysical single epoch offsets greater than 2 meters, which could easily be implemented in real-time. The station-dependent noise threshold is set as the 99.5% significance of the noncentral chi-squared distribution fit to the quiescent one-hour time window of each site.

For TDCP PGV derived magnitude estimates and sensitivity, this study uses the velocity scaling coefficients of [Fang and Liu, 2020] to model the expected TDCP velocity signals at a given hypocentral distance and earthquake magnitude. In our TDCP processing method, cycle slip impacts are limited to individual epoch differences, but they are not detected in this current SNIVEL implementation. We currently minimize the impact of cycle slips and other possible non-geophysical outliers using a coarse threshold set to 1 m/s, filtering presumed non-geophysical velocities in the time series that would be geophysically considered extreme or greater [Worden et al., 2020]. Future work should implement more rigorous methods to detect cycle slips in real-time to accurately mitigate their impact and a more sophisticated approach to remove non-geophysical signals while preserving the response to extreme events. This approach to detection and removal or repair could be done on the phase measurements directly (eg [de Lacy et al., 2012, Li and Melachroinos, 2018]), or potentially identified in least-squares outliers of the phase-derived velocities (eg [Qian et al., 2016, Francesca Fratarcangeli, 2018]). Using a very similar method as employed for the RT-GNSS displacements, the station-dependent velocity noise threshold is set as the 99.5% significance of the NCX2 distribution fit to the quiescent 1-hour time window for each site.

Figure 2.3: Probability of detection as a function of earthquake magnitude over a range of α values fixed hypocentral distances of 10 and 100 km, in which α is the quantile used within a best-fit NCX2 distribution. We use the 592-station median threshold at each quantile for the basis of the noise distribution, and the scaling laws of [Melgar et al., 2015, Fang and Liu, 2020] as the basis for the expected signal plus noise. It is worth noting that this illustration does not account for accuracy of detection.

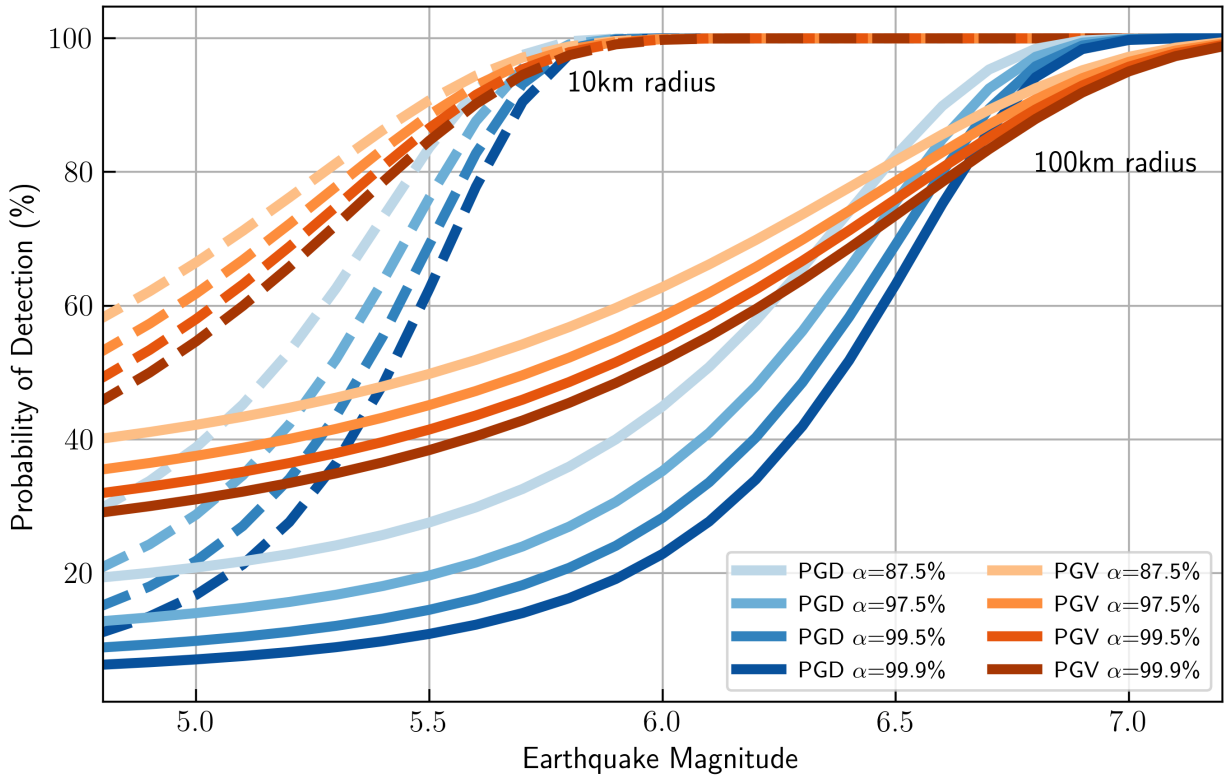

2.3.4 Evaluating events from real-time time series

For this study's evaluation of actual events, hindsight allows us to confidently identify the time series peaks to test the processing method, not the detection algorithm. This is distinctly different from rapid, real-time detection, where the peak displacement grows with the evolving, seismic waveforms. To initiate a time window in this study, the IRIS-DMC travel-time 1-D spherical model web-service is used to estimate the S and P wave arrival of a given event at a given station (see Data and Resources). In displacement space, we reference all topocentric displacements relative to the arrival of the P-wave. PGD always occurs after the S wave arrival, the amount largely dependent on distance but generally tens of seconds behind the P wave [Goldberg et al., 2018]. While not exact, the advantage of using this a priori reference versus a fixed one is it reduces the impact of longer-period RT-GNSS PPP drift of stations further from the source. We also window the detection time domain to allow for the full PGD temporal evolution to occur while also limiting the likelihood of inverting spurious noise. In addition, we also estimate the noise threshold of the 2 minutes prior to the earthquake origin time using the NCX2 approach and use this as a station- and time-dependent integrity validation: if the peak is not above this noise threshold, the peak does not contribute to the evolving event magnitude estimate. Finally, we filter any non-physical outliers that are a result of the processing (such as cycle slips), using a coarse threshold of 2 meters in displacement space and 1 m/s in velocity space. This limits the impact of potential non-geophysical signals affecting magnitude estimates outside the range of any feasible geophysical signals.

2.4 Results

We compare ambient TDCP velocity noise with ambient PPP displacement noise across the same stations over the identical time window to ensure that any potential uncorrected temporal, spatial, or receiver noise dependency is consistent across both processing meth-

ods. We bin each station's threshold for both velocity and displacement into quartiles of the entire network. In Figure 2.4, a map is presented where each station is colored by these quartile estimates. Most notable are the Q4 (red) stations, which have the highest threshold magnitudes and therefore considered the noisiest. We observe a general correlation between GD and GV, which is logical given they share many environmental noise factors. We observe regional clusters of lower-performing (red) stations, such as those in the Caribbean, Yellowstone National Park, the Pacific Northwest, and Southern California. The Southern California and Pacific Northwest areas (insets A and D) have higher station density, so this is partially a visual effect from map layering Q4 last, but could also be near-coastal effects such as unmodeled tropospheric gradients. Yellowstone (inset B) has a high density of stations that are in a sub-optimal environment in terms of obscured satellite sky-view and vegetation multipath, which likely leads to noisier high-rate position and velocity estimates. The Hispaniola/Puerto Rico sector (inset C) has most of the non-braced monuments, which are less stable and also commonly in sub-optimal radio frequency environments, such as on roofs with higher multipath or partially obscured satellite visibility.

Two other parameters we assess related to noise thresholds are GNSS receiver type and monument type at a given station. Monumentation effects do not appear to be clearly correlated to noise performance in Figure 2.5. Much research has been put into evaluating secular velocity noise deriving from monumentation type, but we infer at these higher frequencies monument stability plays less of a role in ambient noise estimates. Certainly, monumentation plays a role in dynamic station-dependent responses, such as ringing or amplification of surface waves [Hodgkinson et al., 2020] from monumentation instability. A receiver-type comparison offers an interestingly salient correlation: PPP noise estimates do not differ by receiver type, while velocity threshold estimates do. We acknowledge the potential that additional interdependence on antenna type or receiver environment might influence this result but infer this receiver correlation result is predominantly related to the respective processing methods: TDCP-GV derives directly from the receiver-independent

Figure 2.4: (a) Spatial distribution of position displacement noise estimated from PPP-derived ground displacement and (b) velocity noise estimated for TDCP-derived ground velocity. Inset A is the Pacific Northwest of the United States, inset B is Yellowstone National Park and vicinity, inset C is Hispaniola and vicinity, and inset D is Southern California. The estimated thresholds of each station are binned by quartile of the entire network distribution of thresholds by method. These quartiles are represented by the blue, Q1 markers (the lowest noise) to the red Q4 markers (the highest noise). We observe higher noise (red) regional clusters such as in Yellowstone National Park, Hispaniola, and the Pacific Coastline of the United States. For absolute values, quantitative threshold ranges are included in the legend.

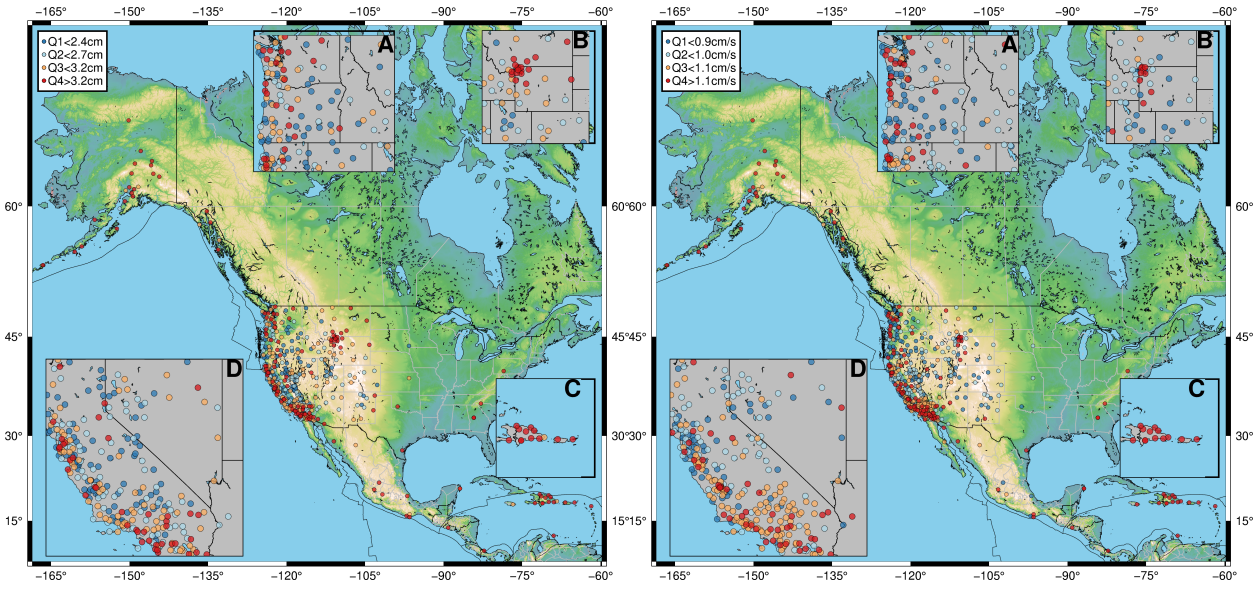

broadcast ephemeris and receiver-dependent carrier phase observations. Receiver dependent carrier phase effects such as receiver dependent multipath or radio frequency interference mitigation, and oscillator drift rate not accounted for in local oscillator drift estimates impact these TDCP velocity estimates. PPP GD is derived from the PPP estimates, which are stateful estimates of which the noise is largely driven by the modeled common-mode effects such as the commercial satellite orbits, clocks, and atmospheric products. Possible TDCP receiver-dependent dynamic response performance is not evaluated here.

We take the median of these 592-station noise estimates to represent the ambient noise level of the respective methods; a signal larger than this threshold is considered detectable. Given the scaling laws of [Melgar et al., 2015, Fang and Liu, 2020] as our signal model, the noise threshold is used to estimate the detection sensitivity of each method across a range of magnitudes and hypocentral distances. It is worth noting we take the network median to compare methods across a large sample set, flattening detection to a binary result. This time-domain station- or time-dependent thresholds, however, can also be compared for assessment or potentially real-time integrity monitoring, like the proposed approach of [Melgar et al., 2019] in the frequency domain; please refer to Figure 2.3 for a more probabilistic assessment of these sensitivities. In Figure 2.6, we observe across a range of distances from 5 to 1000 kilometers a lower magnitude sensitivity threshold for the TDCP PGV relative to PPP PGD method given the respective TDCP PGV and PPP PGD average noise estimates. We also observe a higher variance in the GV noise estimates that leads to a more distributed station-dependent noise threshold. This aligns with the reported standard errors of the scaling law magnitude uncertainties: 0.27 magnitude units for [Melgar et al., 2015] PGD method, and 0.389 magnitude units for [Fang and Liu, 2020]. Figure 2.6 presents an additional feature to the detection threshold comparison: while the PGV method outperforms with respect to the lower magnitude noise detection threshold, as events get larger in magnitude, the signal of PGD gains magnitude at a greater rate than PGV relative to the static ambient noise, as evident in the nearly converging distal arcs. Given these results, for a given magnitude and

Figure 2.5: Distributions of noise threshold estimates for the experiment set by antenna monument type (panels a,b) and by receiver type (panels c,d). Panels (a) and (c) are displacement-derived noise thresholds; panels (b) and (d) are velocity-derived noise thresholds. The drill braced monuments include deep and short drill monuments, which are generally stainless steel rods drilled 2–10 m in 3–4 azimuthal directions coupled to a vertical rod just below the antenna mount. "Other" consists of all other monument types, including masts, wellheads, roof mounts, and pillars. The receivers included in the study are Trimble NetR9s, Septentrio PolaRx5s, and Trimble NetRSs. NetRS solutions are Global Positioning System [GPS]-only; Trimble NetR9s and Septentrio PolaRx5s RTX PPP estimates use GPS and GLONASS.

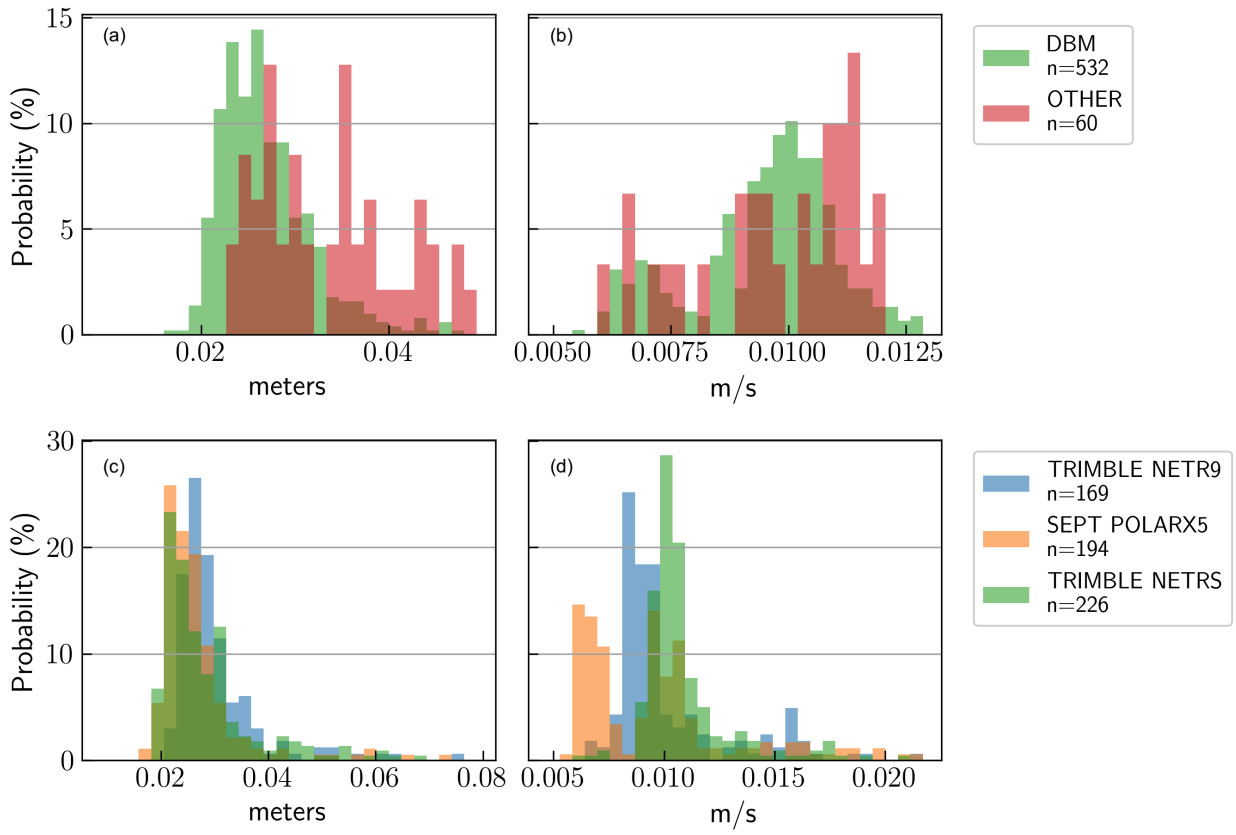

distance window, we infer this would lead to the TDCP PGV method detecting an event by more stations although with greater uncertainty.

2.4.1 Comparison of PPP-PGD vs TDCP-PGV in actual events

We use two well-recorded events as examples of the response of these methods in supporting event detection and real-time magnitude estimates. The GNSS PPP-PGD network response in the Ridgecrest 2019 M7.1 was nominal and is well documented ([Hodgkinson et al., 2020, Melgar et al., 2019, Mattioli et al., 2020, Melbourne et al., 2020]). The real-time GNSS network density surrounding this event enables a direct comparison of the PPP-PGD and TDCP-PGV methods. In this analysis, 250 sites are used within 375 kilometers to contribute to the PGV and PGD estimates (Figure 2.7). Consistent with the previously published studies cited above, PPP-PGD is within $\pm$0.2M of the final ANSS Comprehensive Earthquake Catalog (ComCat) magnitude of 7.1 in under 20 seconds. Network latency of position availability is not accounted for in these estimates, though typically these are around 1-2 seconds [Mattioli et al., 2020].

In this larger magnitude event with many nearfield stations, detection sensitivity does not play a noticeable role in rapid magnitude estimation with the majority of waveform magnitudes well above the noise thresholds on arrival at the stations. In Figure 2.7, the running median estimate by each processing method is depicted by the continuous orange or blue lines; the black dashed line represents the estimates if no thresholding was implemented, which is barely visible due to its minor variance from the thresholded estimate. The increased uncertainties of the PGV estimates though impact the accuracy of the magnitude estimate, evident in the larger spread of the individual station velocity-derived estimates. This result is consistent with the variance of the station-dependent ambient noise thresholds from Figure 2.6. Except for the initial PGV estimate, both methods initially underestimate the magnitude before converging almost exactly on the final ComCat magnitude of 7.1. This response might be explained by the limitations of existing scaling laws in the near

Figure 2.6: Detection thresholds of each method as a function of hypocentral distance from a given magnitude event. The scaling coefficients used are noted in the legend. The top panels cor- respond to the color-coded vertical bars in the lower panel. These bars are slices at different hypocentral distances that illustrate the distribution of ambient noise threshold estimates across the 592 stations in the set.

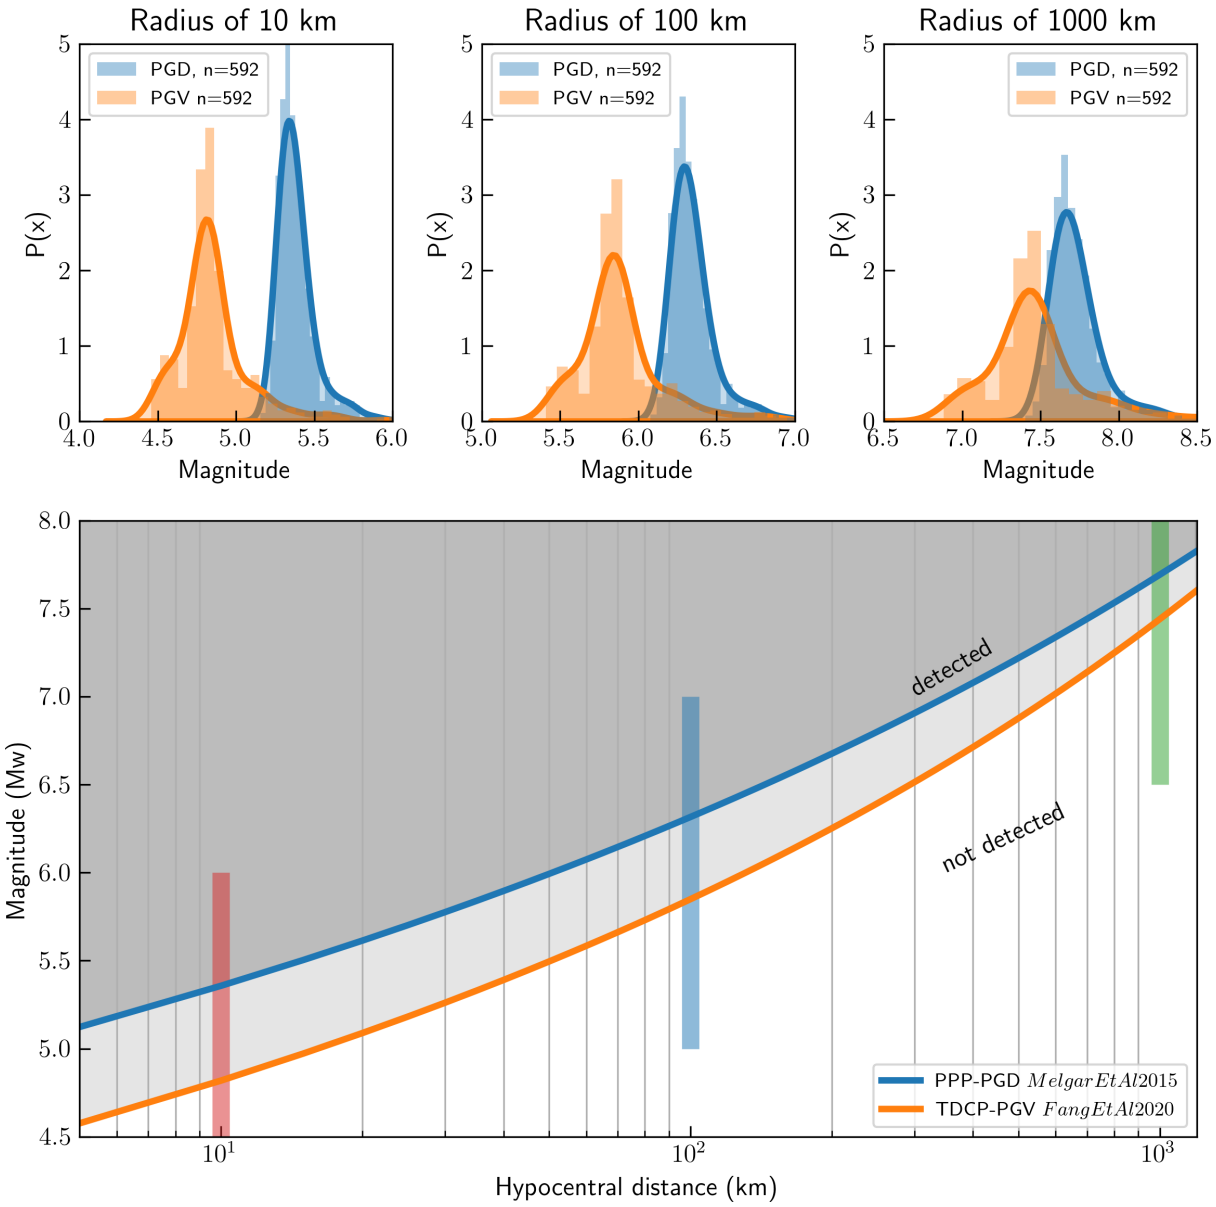

Figure 2.7: Magnitude estimation evolution over time for the 2019 Mw 7.1 Ridgecrest sequence given individual stations TDCP-PGV and PPP-PGD magnitude estimates. The points represent the individual station magnitude estimates, and the solid lines are the running median event estimate given the available estimates at the time. The stations that are above the noise estimate using the running previous 2 min window are included in this median. The shading represents one standard deviation of the estimates contributing to the magnitude at each epoch. The barely visible black dashed lines are the estimates if no thresholding takes place; this is added to demonstrate that thresholding has little apparent effect in the relative near field of larger events. In this figure, all station's contributions to the overall magnitude estimate are evenly weighted over time. The U.S. Geological Survey Comprehensive Earthquake Catalog final magnitude is the dashed red line.

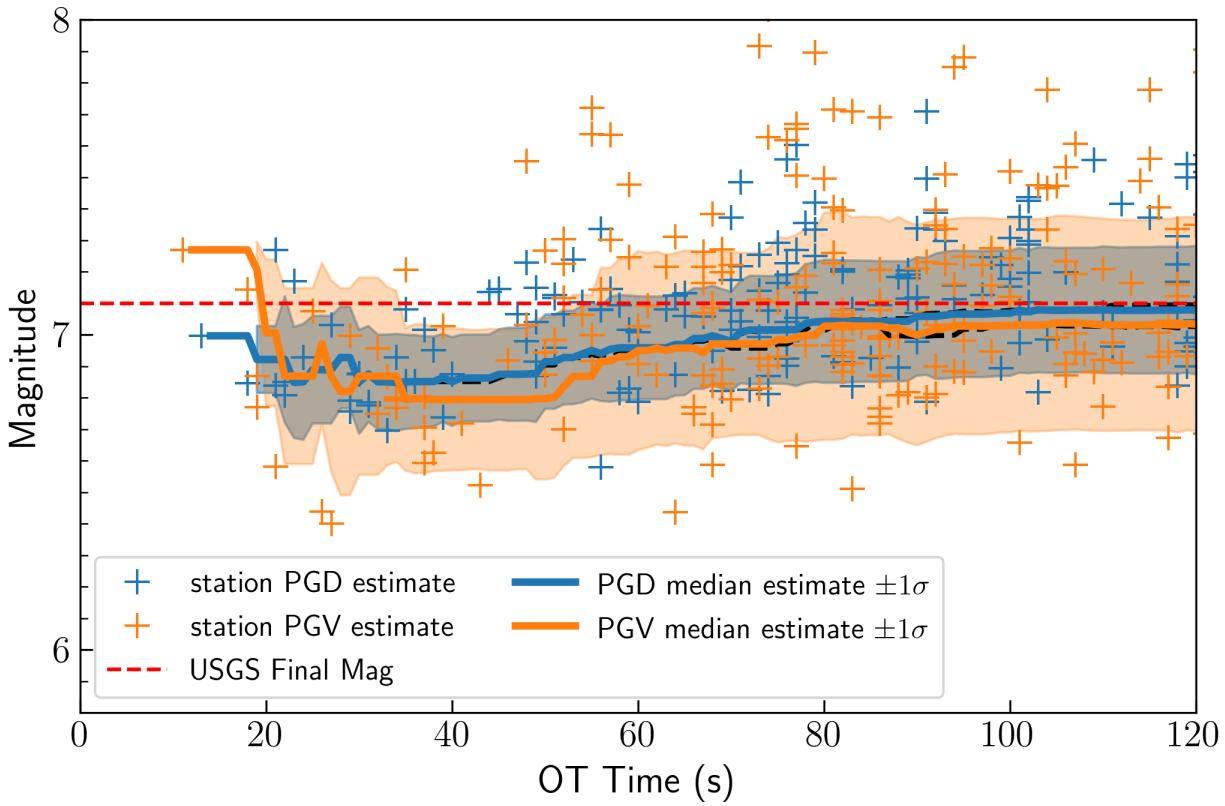

field treating the events as a simple point source [**?**]. Within 25 seconds, both methods processed independently, reach an estimate within ±0.2 magnitude units of the ComCat value to provide useful information in a hazard alert system. Combining or considering both independent estimates would not address the variance in observed peak deformation from the current point source, isotropic attenuation empirical model, but would add redundancy to the respective estimation processes.

For the M5.3 event, 2 sites are presented within 12 km hypocentral distance (Figure 2.8). We present this smaller set as evidence of the characteristics of the ambient noise model assessed previously: In the right two panels of Figure 2.8, the two nearest (≤12km) receivers detect the M5.3 event in TDCP velocity space given the aforementioned magnitude threshold. Only the nearest station detects the event above the estimated threshold in PPP displacements. This event is close to the boundary threshold of detection for PPP 3D displacement method given these individual station noise thresholds. Peak velocity and displacement signals estimate the magnitude within ±0.2M of the final ComCat magnitude 5 seconds after the origin time. Further visual inspection of the left panels of Figure 2.8's time-series of directional components indicates a clear signal in the velocities and a more subtle signal in the displacements. The largely horizontal nature of this event makes the peak magnitude partially obscured by the vertical noise: this is especially true in the 2nd station, P495, velocity time-series, where the peak velocity is above the threshold but so are spurious outliers that would lead to greater uncertainty in real-time peak identification. This is also evident in the PPP time series of P507, where a rapid southward motion is evident but ultimately obscured in accumulated noise of the 3-component magnitude. The only stations capable of detection are very near the epicenter, and thus likely more responsive to near-field source variations, as suggested by [Goldberg et al., 2021], and yet at this lower end of detectable seismic deformation signals, the scaling velocity relationship passes for these two receivers and for one in displacement space. That the P495 displacement signal fails to go above the threshold yet still estimates the correct magnitude raises some concern about

Figure 2.8: Near-field, two station response in displacement and velocity space for the 2021 Mw 5.3 Calipatria event. Panels (a)–(d) are P507, approximately 9 km hypocentral radius, and panels (e)–(h) are from P495, approximately 12 km hypocentral radius. Each station has four panels: panels (a), (b), (e), and (f) are the time series of vertically offset displacement components (panels a, e) and magnitude (panels b,f), and panels (c), (d), (g), and (h) are vertically offset velocity components (panels c,g) and velocity magnitude (panels d,h), with the running 2 min windowed 99.5% noise estimate prior to the event represented by the horizontal, dashed red line. A green vertical line highlights the peak signal detected by the method in displacement or velocity space that is above the running noise threshold; a gray-vertical line highlights the peak signal detected is below the running noise threshold and thus would not be included in a magnitude estimate. The vertical, dashed black, blue, and red lines are the earthquake origin time, estimated time of arrival of the P wave, and estimated arrival of the S wave at the station location, respectively. These two stations clearly illustrate the strength of signal relative to the estimated noise in the velocity estimates. This aligns with the modeled signals using the scaling laws from Figure 2.

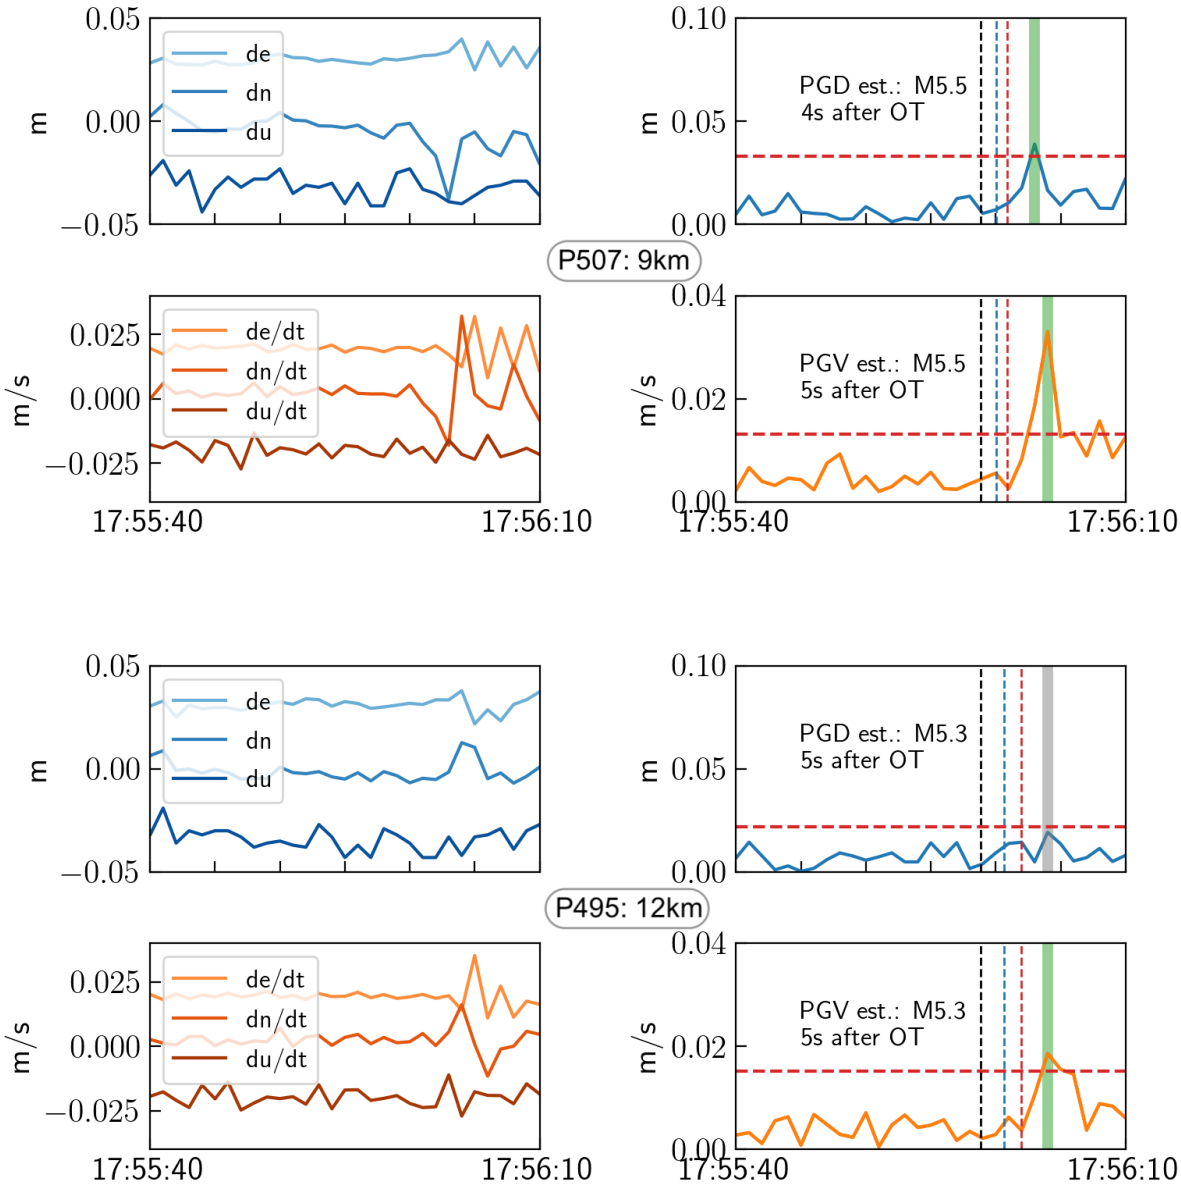

the threshold level if this is indeed a false negative. Future work should further evaluate this balance of false alerts with sensitivity across many events and stations close to the boundary. In summary, these TDCP velocity estimates would offer additional information to detect and characterize co-seismic antenna motion in real-time that the PPP-derived displacements might not.

2.5 Conclusions

For nearfield relatively smaller events (M5-6.5) and/or larger events at greater hypocentral distances, TDCP velocities are more sensitive to detecting seismic surface waves than PPP displacements. PGD grows larger with magnitude relative to PGV, however, thus this advantage of sensitivity falls off with increased magnitude in the relative nearfield. It is important to note while the seismic infrastructure signal saturation is not an issue at lower magnitude earthquakes, the TDCP method provides an independent estimate of event detection at that magnitude range. [Crowell, 2021] emphasizes the role these independent velocity estimates could play in contributing to USGS ShakeMaps of Modified Mercalli Intensity, a valuable metric for evaluating spatially distributed impacts from co-seismic shaking. Future efforts to evaluate lowering this threshold to increase sensitivity without increasing the probability of false alerts could prove useful for incorporating these methods into smaller magnitude detection, complementing existing seismic infrastructure especially in more sparsely instrumented regions. Station-specific TDCP GV stochastic noise estimates vary more than station-specific PPP GD noise, which leads to increased variability in estimated sensitivity. This ambient noise difference aligns with historical dynamic signal regressions: current PGD-derived scaling laws have less variance than PGV scaling laws in their log-linear fit and thus are more robust in their magnitude estimate. [Shu et al., 2018] demonstrate that aliasing is present in 1 Hz velocity estimates, which would most likely impact the near-field, higher frequency signals. The increased variance in TDCP PGV is likely related to estimation in velocity space: velocity signal scattering results in a more complex structure ([Wu and Zhao,

2006]).

TDCP velocities can be computed in real-time at the network edge and at higher rates to complement existing PPP displacements and/or traditional seismic methods, including magnitude estimates and Shakemap velocity estimates. This processing method is open source and free of commercial corrections and proprietary algorithms. Estimating TDCP at higher rates and at the network edge would reduce potential spectral aliasing in lower rate (¡=1 Hz) estimates with substantially less telemetry bandwidth cost compared to centralized processing of higher rate full fidelity observational streams. The lower cost and increased sensitivity of TDCP velocities, particularly in the near field, makes this method well suited for a potential low cost, dense array experiment to capture and potentially alert for ground motion. The representative noise model of 3-component ground motion (velocity or displacements) magnitudes recommended here can be implemented for a real-time/near real-time temporal and station dependent characterization. It is important to note that in this analysis we compare two separate processing methods, PPP and TDCP, in two separate detection domains, displacement and velocity. We did this intentionally to assess the respective methods "as is," however further work could decouple the processing method from the detection domain.

Chapter 3

Supervised Learning of High Rate GNSS Velocities for Earthquake Strong Motion Signals.

This chapter is adopted from the following manuscript:
Dittmann, T., Liu, Y., Morton, Y., and Mencin, D. (2022). Supervised machine learning of high rate GNSS velocities for Earthquake strong motion signals. Journal of Geophysical Research: Solid Earth, 127, e2022JB024854. https://doi.org/10.1029/2022JB024854

3.1 Abstract

High rate Global Navigation Satellite System (GNSS) processed time series capture a broad spectrum of earthquake strong motion signals, but experience regular sporadic noise that can be difficult to distinguish from true seismic signals. The range of possible seismic signal frequencies amidst a high, location-varying noise floor makes filtering difficult to generalize. Existing methods for automatic detection rely on external inputs to mitigate false alerts, which limit their usefulness. For these reasons, geodetic seismic signal detection makes for a compelling candidate for data-driven machine learning classification.

In this study we generated high rate GNSS time differenced carrier phase (TDCP) velocity time series concurrent in space and time with expected signals from 77 earthquakes occurring over nearly 20 years. TDCP velocity processing has increased sensitivity relative to traditional geodetic displacement processing without requiring sophisticated corrections. We trained, validated and tested a random forest classifier to differentiate seismic events from noise. We find our supervised random forest classifier outperforms the existing detec-

tion methods in stand-alone mode by combining frequency and time domain features into decision criteria. The classifier achieves a 90% true positive rate of seismic event detection within the dataset of events ranging from M_W4.8-8.2, with typical detection latencies seconds behind S-wave arrivals. We conclude the performance of this model provides sufficient confidence to enable these valuable ground motion measurements to run in stand-alone mode for development of edge processing, geodetic infrastructure monitoring and inclusion in operational ground motion observations and models.

3.2 Introduction

Real-time, low-latency observations of medium to great earthquake ground motions are vital to rapid hazard assessment and earthquake early warning (EEW) systems. These measurements have historically been recorded by inertial seismometers. Higher rate ($\geq$ 1Hz) continuous GNSS measurements capture stronger dynamic motions and permanent displacements of propagating seismic waveforms [Nikolaidis et al., 2001, Larson et al., 2003]. These geodetic strong motion measurements [Larson, 2009] will rarely clip nor require double integration that leads to magnitude saturation in the near-field of larger, destructive earthquakes common to inertial velocity sensors [Bock et al., 2004, Crowell and Geng, 2013, Colombelli et al., 2013]. Furthermore, additional material low-latency observations densify existing seismic ground motion measurements. These observations are particularly valuable when damaging seismic events occur in sparsely instrumented regions [Grapenthin et al., 2017] or when networks or infrastructure fails.

However, geodetic deformation timeseries are noisier than traditional inertial sensors [Melgar et al., 2020]. GNSS seismic waveforms are observed from spaceborne radio signals traveling over 20,000 kilometers through often convoluted atmospheric signal paths. The radio signals are tracked by a variety of receivers using antenna situated in a range of radio frequency environments. As a result, these measurements have relatively high and complex noise signatures, making separating signal from noise challenging. Signal amplitudes from

the largest, most costly events can be difficult to distinguish from non geophysical events, such as filter reconvergence or signal loss of lock (Fig: 3.1-a). Medium magnitude events, often difficult to detect above the geodetic noise floor, can be destructive or tsunamigenic. The ability to make accurate, low-latency distinction between true signals and noise in stand-alone mode, without external sensors or information, minimizes points of failure and decision latency and maximizes the value of these integral network decision inputs and potential edge processing capabilities.

Current approaches to detect motion use variations of time domain thresholds to flatten the decision to a function of signal amplitude. Several existing approaches make use of low-pass filters similar to traditional STA/LTA seismological phase picking [Allen and Ziv, 2011, Ohta et al., 2012, Minson et al., 2014, Kawamoto et al., 2016, Goldberg and Bock, 2017] that extract static offsets for finite fault inversion but filter valuable dynamics information. Recent interest in peak geodetic dynamic signals [Melgar et al., 2015, **?**, Fang and Liu, 2020, Crowell, 2021] prompted use of unfiltered timeseries to capture peak signals for magnitude scaling laws and ground motion intensity measurements (Fig: 3.1-b,c). These epoch-wise threshold detection methods [Crowell et al., 2009, Psimoulis et al., 2018, Hohensinn and Geiger, 2018, Hodgkinson et al., 2020, Dittmann et al., 2022a] use instantaneous measurements to estimate motion onset and therefore are a step-forward for inclusion of GNSS-seismology waveforms, but have limited "real-world" testing and most importantly mitigate false alerts for operational systems by correlating detections with proximal stations within networks or windowing in time from seismic triggers. These processes reduce the usefulness of these measurements for rapid, stand-alone decision criteria. The evolution of these detection methods has been vital for the vanguard of GNSS-based seismology, but fall short for real-time operational hazard systems to ingest the full temporal and frequency range of these valuable measurements into models with minimal stand-alone false alerting.

In this work, we evaluate whether existing GNSS hardware can: more reliably detect motion signals that are 1) constellating near the ambient temporal noise floor 2) with min-

Figure 3.1: An example of the difficulties of differentiating a relatively weak seismic GNSS signal event and a GNSS noise disturbance using existing detection methods. The signal depicted is the east component of station P507 observing a M_W5.41 (USGS event ID: ci15200401) at approx. 23 km; velocities are presented in seconds relative to the event origin time (OT). The proximal noise disturbance depicted is a non-geophysical processing artifact or signal propagation effect that might result from sources such as cycle slips, ephemeris, multipath, or other signal path effects. Panel (a) is the 5Hz timeseries, in addition to a low-pass filtered (corner frequency of 0.5Hz) timeseries to emphasize the signal and noise for the reader. Gray shading represent areas within the noise estimate for each respective method. Vertical dashed lines are estimated (iasp91 model) P- and S-wave arrival times. Panel (b) illustrates a static threshold taken from [Hodgkinson et al., 2020]. This approach is sensitive to the weak signal, but equally sensitive to noise. The threshold has limited memory and rapidly alerts to the onset of the noise disturbance, and also issues several additional false alerts around 105 seconds OT. Panel (c) is a variation on an STA/LTA approach implemented from [Psimoulis et al., 2018] called RT-Shake with a moving threshold of 3 times the moving standard deviation. This approach detects the signal event later in the waveform with little information regarding the event duration. The noise disturbance adds an initial false alert, after which the noise region expands to minimize additional false alerts. However this memory would result in missed detection should such a noise disturbance occur immediately prior or during a seismic event.

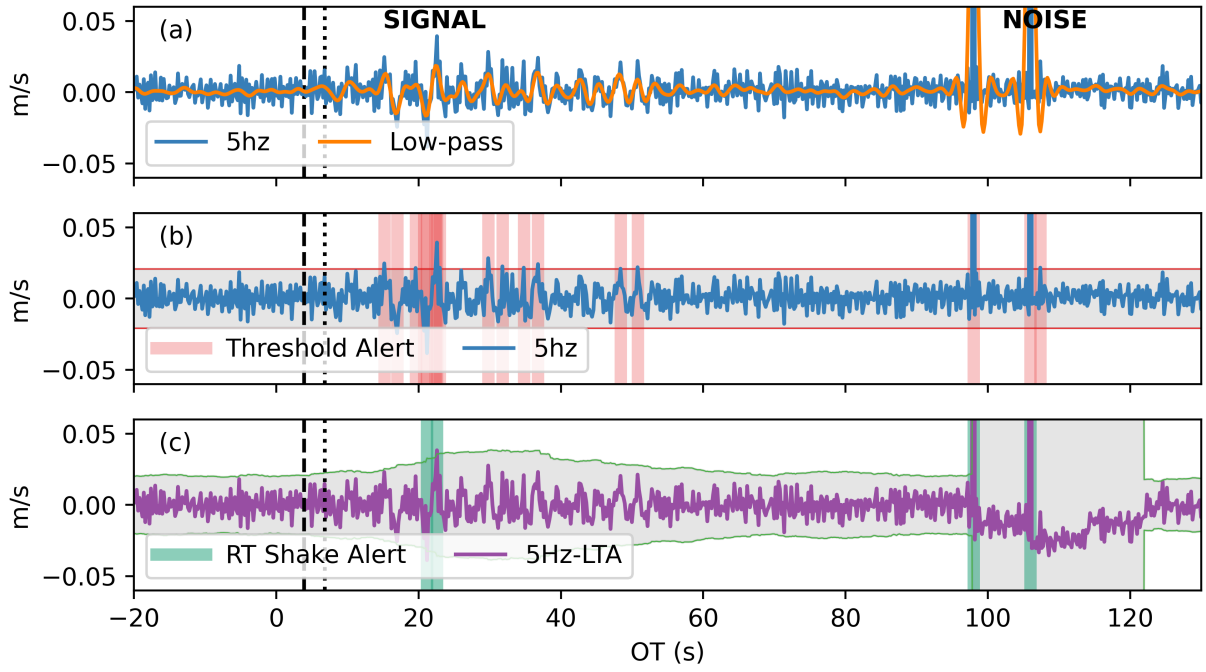

imal false alerting 3) in a low-latency, stand-alone mode and 4) with no specific fault or network geometry. We trained a machine learning classifier on a supervised dataset of GNSS velocity time series concurrent in space and time with known seismic source signals. We assembled, processed and labeled a dataset of 1701 earthquake-station high rate (5Hz) time series pairs. We also include a substantial seismic event-free noise dataset to improve model generalization. We optimized the classifier on these combined datasets with applied domain knowledge to feature selection and feature engineering that is able to combine time and frequency domain information. We present the superior performance of this classifier relative to existing methods within this motivational context. We offer advantages and implications of deploying this processing and trained model at scale for network wide monitoring, with particular emphasis on the improved sensitivity and integrity of stand-alone GNSS seismic event detection without external inputs.

3.3 Methods

3.3.1 Signals of Interest

We define our detection domain as a binary **event** or **no event** state classification. A critical component of developing a robust classification model is a substantial dataset from which to train, validate and test the model. For optimal results, this dataset requires broad spectrum noise and signal samples such that the model can "learn" and generalize our classification and distinguish signal from noise. We assembled a catalog of 1701 station-event pairs from 77 events by cross referencing available 5Hz GNSS observational data in the UNAVCO geodetic archive with Advanced National Seismic System Comprehensive Earthquake Catalog (COMCAT) of earthquakes greater than M_W4.5. While 1Hz data is more readily available, 1Hz observable decimation undersamples certain event velocity spectra [Joyner, 1984]. We observe this effect in reduced velocity amplitudes from 1Hz data when compared to 5Hz observables in several nearfield TDCP velocity timeseries, such as the M_W6.2 2021

Petrolia event. For larger magnitude events it's likely that sampling closer to 10 Hz is necessary to avoid aliasing [Shu et al., 2018], but we balance this design parameter with the need for sufficiently large available datasets for training. We assigned a conservative radius of detection for each event using ambient noise estimation from [Dittmann et al., 2022a]. For each station-event pair within this spatial footprint, a time series window began 2 minutes prior to earthquake origin time (OT), and extends out in time as a function of radius (Figure 3.2). We conservatively buffered the radius and time window to mitigate limiting this result from the existing model. We also processed the available 5Hz observables for a 30 minute window in the hour prior to event times of the event catalog from 2017-2021. This noise catalog consisted of 1507 unfiltered station-noise timeseries from 904 unique stations across a range of receiver types, geographic locations, antenna environments and atmospheric conditions, among other potential TDCP noise variance sources. Inclusion of this extended, real-world noise dataset in training and validating will improve the model's generalization, or performance on unseen data.

Current use of GNSS-derived seismic ground motion for operational EEW [Murray et al., 2018] use precise point positioning (PPP) derived topocentric coordinates to capture dynamic waveforms or static offsets relative to a stations *a priori* position. Instead, we align synchronous carrier phase epoch-wise changes, predicted satellite orbital velocity and line-of-sight geometry to accumulate coherent energy with respect to the shared receiver clock drift rate and directional velocities in a local reference frame. Variations of this geodetic processing method, known as time differenced carrier phase (TDCP) [Graas and Soloviev, 2003] or variometric velocities, can record co-seismic velocity waveforms [Grapenthin and Freymueller, 2018, Hohensinn and Geiger, 2018, Crowell, 2021] as well as integrated over time into seismic displacement waveforms [Colosimo and Mazzoni, 2011, Branzanti et al., 2013, Francesca Fratarcangeli, 2018]. We processed these 5hz measurements with the open-source SNIVEL package [Crowell, 2021] using broadcast ephemeris and narrow lane phase combinations. We chose TDCP over PPP because it is more sensitive to motion [Fang

Figure 3.2: Map of seismic focal mechanisms used in this work and distribution of 77 event magnitudes. The number of stations used in each event is a function of the ground station network density and the magnitude-dependent sensitivity radii.

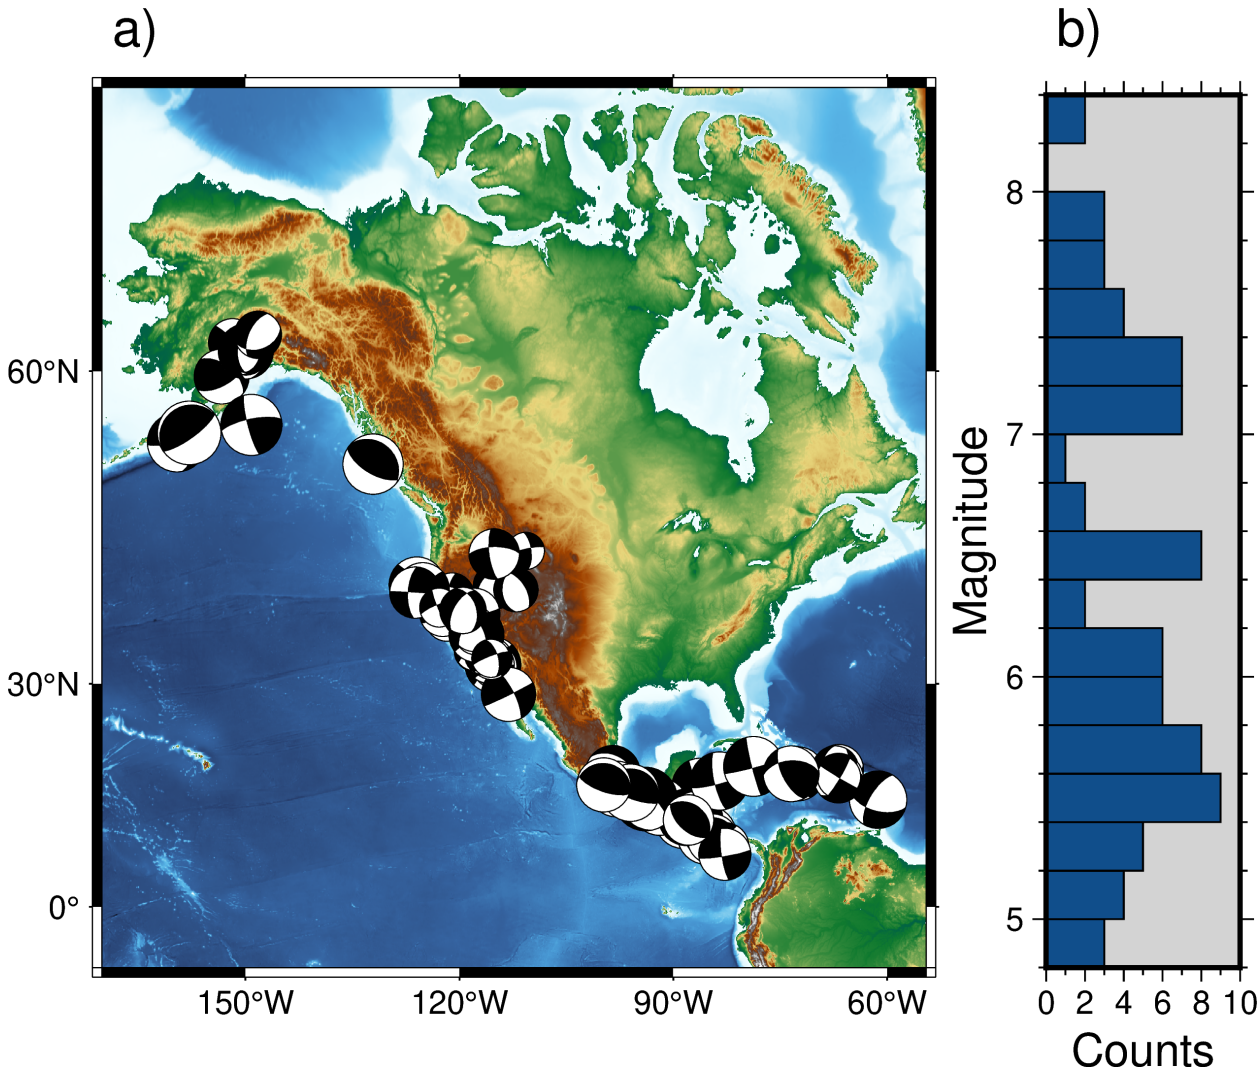

and Liu, 2020, Dittmann et al., 2022a], and it is "lightweight" in that it does not require sophisticated corrections and is computationally inexpensive. From a machine learning perspective, this could be considered a first step in our feature engineering, or applying domain knowledge to extracting features that are correlated with motion in observed carrier phase measurements.

3.3.2 Feature Engineering Pipeline

Data-driven supervised machine learning models are widely used in computer vision and natural language processing due to their superior accuracy for challenging classification, regression and clustering problems. Earth scientists have adopted many of these models for geoscience research [Kong et al., 2019]. Recent catalogs of historic seismic data training sets (eg. Stanford Earthquake Data Set [Mousavi et al., 2019], INSTANCE [Michelini et al., 2021]) have contributed to benchmarking improvements of earthquake detection, phase picking, localization, and magnitude estimation (eg. [Meier et al., 2019, Mousavi et al., 2020, Kong et al., 2019]. These extensive labeled data sets enable sophisticated data-driven classifiers and deep learning models using inertial seismic data. Several geodetic applications of machine learning algorithms have demonstrated promising results with respect to seismic processes. [Crocetti et al., 2021] used a random forest classifier for antenna offset detection, including due to earthquake offsets, from low-rate, 24-hour position solutions. [Habboub et al., 2020] applied a neural network to coordinate time series anomaly detection applicable to specific regional datasets well above the noise floor. [Dybing et al., 2021] used neural networks for earthquake detection and [Lin et al., 2021] employed deep learning used for rapid event magnitude estimation; both of these studies used extensive synthetic displacement waveforms derived from real-world fault geometries and real-world PPP noise models.

In our study, we used a random forest algorithm for our classifier [Breiman, 2001] of GNSS velocities. Random forest is an ensemble of decision trees; a single decision tree is a classifier where input features are split along thresholds to separate source, or root, data

from end node classifications, or leaves. An ensemble or forest of trees each vote on the feature decision criteria to select the optimal decisions towards minimizing correlated noise. Due to the infrequent nature of larger magnitude earthquakes, the event classes are naturally imbalanced but by pre-selecting specific time series of events, we have reduced this imbalance for training (Table 3.1) and testing. Random forest hyperparameters were selected using a grid search over the number of decision trees used, the maximum decision splits within a tree, and imbalance classification weighting strategies.

Figure 3.3: Schematic of our classification workflow: Inputs were 5Hz GPS phase measurements and broadcast ephemeris, which are processed using narrow lane combinations using SNIVEL. Target labeling combined with feature extraction were used for training a supervised random forest classification model to predict event classification on testing subsets.

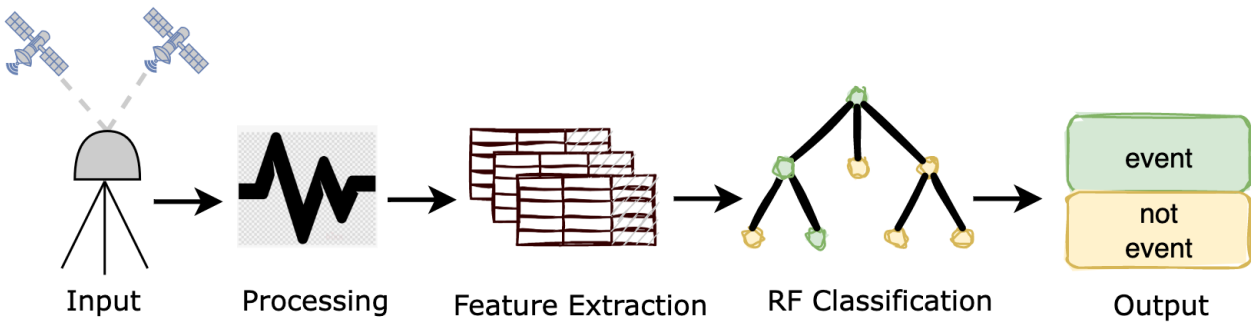

SNIVEL TDCP processing generates 5 Hz time series of the three topocentric velocity components and the clock drift rates. From these event-station pair time series of velocities, we generated feature sets to label for our supervised classification (Figure 3.3). Our feature samples consisted of three directional components of 30 second windows overlapping every 10 seconds; within these windows we included the four maximum component norm window values, window median, window median absolute deviation and window power spectral densities from the lowest frequencies bins containing periods 1-30s as features. These features and windowing allowed our model to incorporate signal and noise amplitude in the time domain, akin to the traditional threshold approach, as well as power spectra in the frequency

domain. In our binary classification, an **event** is seismic ground motion in an individual component. Labels were assigned through visual inspection as **not event** or 0, **event** or 1, and **maybe** for windows that we are not able to distinguish and excluded from testing and training. Each directional component was labeled independently. This resulted in 140,334 labels for the on-average 30 samples for 1701 station event pairs of three component velocity time series. The event-free, noise dataset included an additional 266,739 labels of 3 component, non-overlapping velocity time series.

We evaluated two feature extraction models. Feature set #1 was a combined array of all 3 directional components with a single label at each window. The horizontally concatenated components resulted in $3 \times m$ features and n samples, where m is the number of features per component ($m = 36$ in our pipeline) and n is the number of window samples. If any component was labeled "1" for event, the feature set #1 sample label was "1" for event. If a **maybe** label was present without **yes** events on the other concurrent components, the window was excluded from training/testing. Feature set 2 included a target vector for each component but excluded the noisier vertical signals. These vertically concatenated components resulted in m features and $2 \times n$ samples. In this extraction case any **maybe** labels were excluded from training and testing.

We employed a nested cross validation approach for unbiased testing of our dataset. We initialized 10 different folds of randomly splitting the 77 events and noise catalog samples into 90% training and 10% testing. By splitting on events we avoided "leakage" of information from our training into our testing, including correlation of seismic waveforms from any given event observed across a network. By cross validating over 10 folds we minimized biasing our result by the relatively small testing subsets of events, and can quantify the ability of our classification model to generalize for future events. Each event was observed by a different number of stations depending on network density and sensing radius, and each station-event pair had differing number of time samples; consequently the feature vectors of training and testing were not precisely 90/10 split in samples. In each fold, we held the test set aside as

Table 3.1: Distribution of classification sample labels used in training/testing datasets by component and label. For more information regarding the distribution of peak values, see fig. 3.5-c

	East (n=135,671)	North (n=135,671)	Up (n=135,671)
Non-Event	94%	94%	99%
Event	5%	5%	1%
Maybe*	1%	1%	<1%

*Maybe's excluded from training/testing

"unseen", and tuned our model using K-fold cross validation [Bishop and Nasrabadi, 2007] on the remaining training set. We implemented 5 inner folds in our K-fold cross validation to find the best hyperparameters. This cross validation approach allowed us to minimize overfitting the training dataset and evaluate the performance of our model on unseen data as though it were running such a classifier on yet-to-occur events.

The traditional "accuracy" metric, or the ratio of the correctly classified labels relative to the total number of labels, of our classification will be less sensitive regardless of optimization choices due to the infrequent events of our imbalanced classification. Instead, we optimized on metrics that reflect accurately classifying the infrequent events. Precision, or positive predictive value, is equal to the number of true positives (TP) over the sum of TP and false positives (FP).

$$Precision = \frac{TP}{TP + FP} \tag{3.1}$$

Recall, or sensitivity, is the number of TP over the sum of TP and false negtives (FN).

$$Recall = \frac{TP}{TP + FN} \tag{3.2}$$

F1 is the harmonic mean of precision and recall:

$$F_1 = 2 \times \frac{precision \times recall}{precision + recall} \tag{3.3}$$

Here, positive denotes motion and vice versa.

Precision and recall are approximately inversely related and each is a function of our random forest decision threshold (Figure 3.4). Quantifying missed detections and false alert

Table 3.2: 10 fold nested cross validation results comparing Feature Set 1 is where all 3 components are combined for each window, and Feature Set 2 is where each horizontal component is tested independently.

		Feature Set #1	Feature Set #2
Precision	mean	0.72	0.64
	stdev	0.19	0.19
Recall	mean	0.70	0.65
	stdev	0.10	0.15
F1	mean	**0.70**	0.63
	stdev	0.12	0.14

rates is imperative for the effectiveness of any EEW system [Minson et al., 2019]. We optimized hyperparameters on F1 scores, a balance of precision and recall, but this parameter is a knob available to tune depending on societal missed detection of false alerting tolerances of a future operational system.

3.4 Results and Discussion

We evaluate the two optimal feature selection strategies and a range of random forest hyperparameters using a grid search. Given the F1 scores of our 10 fold nested cross validation approach (Table 2), our optimal model used feature set #1, with all available spatial components with a single target label to accumulate as much signal as possible towards our binary classification. Each train/test fold selected different optimal hyperparameter combinations for optimizing F1 scores, but the majority selected 100-200 decision trees, 100 decision splits and no class weighting with a decision threshold of 0.4 (Figure 3.4). This decision threshold was selected inside the cross validation of each split and applied to testing sets along with the other hyperparameters selected. Our mean and one standard deviation nested cross validation F1 score of 0.70 ± 0.12 indicates our ability to successfully train a model using random forest. The variance in our results as a justifies our nested cross validation approach to quantify the variability in results as a function of the testing set; presumably some variability will resolve with expanded target catalogs.

Figure 3.4: Mean precision, recall and F1 as a function of decision thresholds for the 10 fold nested cross validation evaluation. The shaded regions are the standard deviations across the 10 folds as a function of threshold. The dashed vertical lines are the maximum F1 decision threshold, with the dashed horizontal lines being the corresponding maximum F1 score.

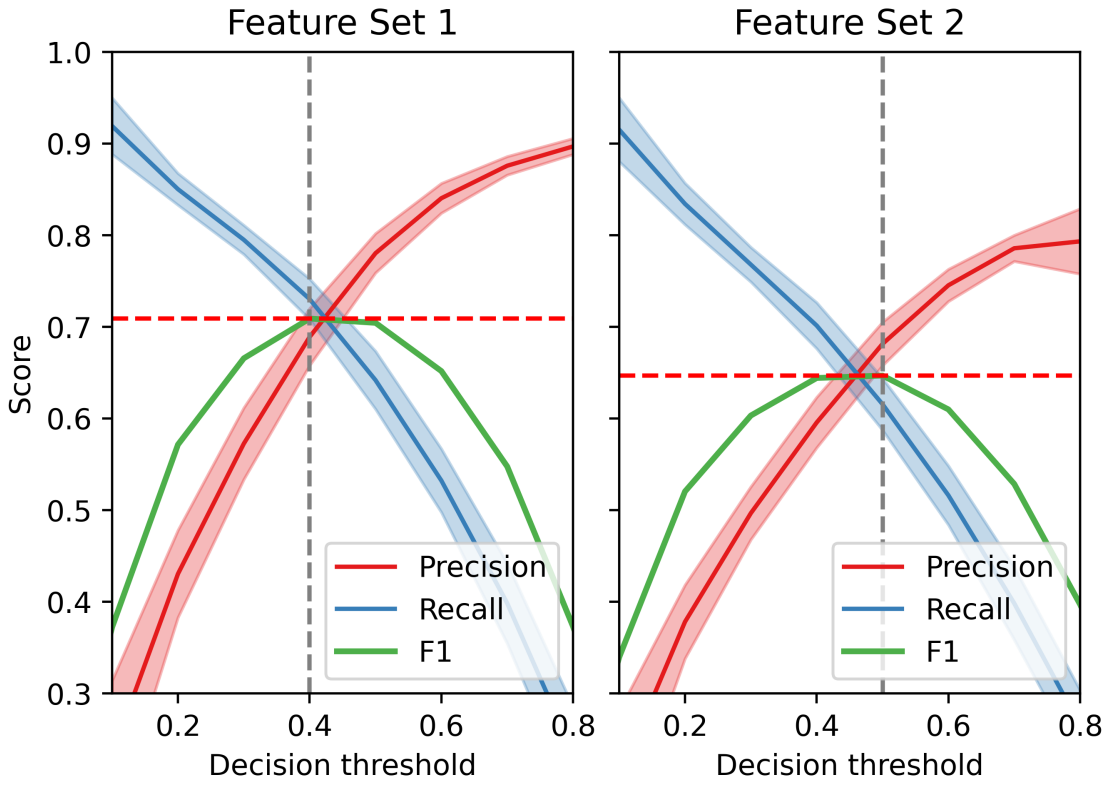

3.4.1 Feature Importance

A benefit of random forest is that individual feature importance is readily extracted from the trained model. When evaluating feature set 1, we find several aspects of the feature importances that align with our domain knowledge and therefore contribute to the explainability of our trained model. The horizontal velocity components dominate the contribution to the model (Figure 3.5a). GNSS ambient noise on the vertical component is much higher than that of the horizontal components and vertical seismic signal amplitudes are diminished relative to horizontal motion along horizontal strike-slip fault mechanics that are common in the spatial region of this study. These less frequent signals amidst a higher relative noise floor were harder to detect and thus contributed less to the empirical classification model. Within a horizontal component, the lower frequency spectral features had the most influence (Figure 3.5b). The most important frequency bins were between 6-15 second periods, aligned with the dominant frequencies of seismic surface waves. Our 5Hz sampling, as compared to lower rates, boosted the detectability around the noise floor, and avoided corner frequency aliasing of certain magnitudes.

The time domain features contributed to the model, albeit much less than the lower frequency spectral content and with a more complex relationship. Figure 3.5(c) shows increasing F1 score with increasing peak velocity up until approximately a peak velocity of 25 cm/s in the east, followed by diminishing performance. We infer this to be the result of readily visibly identifiable signal events experiencing strong to very-strong shaking around 5-20cm/s [Worden et al., 2020], well above the median noise floor. Infrequent, highest peak velocities ($\geq$ 25 cm/s) might either be the result of the largest events or noise disturbances; the latter are likely degrading the performance within these peak velocity bins. Figure 3.5(d) presents a more straightforward feature relationship in the frequency features, where the greater the accumulated power in the frequency bands of greatest importance (b), the higher the performance metrics (F1, recall, precision). After an initial evaluation, we re-

moved the highest frequency power spectral densities from our features; these are logically "noise" in our classification and not contributing. Altogether, these feature importances illustrate a key attribute of such a machine learning approach: combining features in an explainable way into an effective decision process.

3.4.2 Comparison with Existing Methods

A critical performance indicator is evaluating how our classification model performs over a range of test events relative to existing threshold approaches. Logic was applied to map existing continuous epoch-wise time domain threshold detection to our 30 second overlapping window target labels. For a threshold method comparison similar to the approach of [Hodgkinson et al., 2020] and [Dittmann et al., 2022a], we estimated the noise threshold in the 2 minute window prior to seismic origin time. [Hodgkinson et al., 2020] characterized the stand-alone sensitivity of detection using ambient noise antecedent to an event as a Gaussian heuristic threshold. [Dittmann et al., 2022a] approximated the 2 minute window of ground velocities as a non central chi-squared (NCX2) distribution with 3 degrees of freedom, and then set the 0.995 confidence level value of this distribution as a noise floor approximation. Any three dimensional GV magnitude above this noise threshold after this window is considered an event, and evaluated on whether it falls within a window labeled motion or not. RT-Shake [Psimoulis et al., 2018] evolved the previous geodetic STA/LTA algorithms [Allen and Ziv, 2011, Ohta et al., 2012] by differencing instantaneous measurements from 80 epoch moving averages and then related these values to a moving window noise threshold estimate set to three times the standard error of the previous 80 epochs. This method was run on each component independently, with a single boolean for the presence of motion on any component, and each sample window assigned a boolean based on the presence of any motion. The [Dittmann et al., 2022a] implementation of the threshold window in time was based upon S-Wave speeds [Crowell and Geng, 2013], and [Psimoulis et al., 2018] modified STA/LTA correlated with surrounding stations to minimize false alerts; we did not add this

Figure 3.5: Panel (a) are the distribution of the feature importances across the horizontally concatenated, three spatial components feature set #1 testing. Panel (b) is a close up of the east component, with the features labeled across the x axis for closer inspection. From the left, the first 6 of each component are time domain features (max, min, mad) within the 30 second windows; the next 15 are the power from given frequency bins of the periodogram of the 30 second 5Hz data, increasing in frequency from left to right, with the periods indicated. Panel (c) are precision, recall and F1 scores binned by peak velocity of each sample's east component (denoted in (b) with **). The gray shading are the counts of samples falling within the bins. Panel (d) is a similar performance measure to (c) but binned by accumulated power in the lowest frequency bins that had the highest feature importance in the model, 5-30s period (denoted in (b) with *).

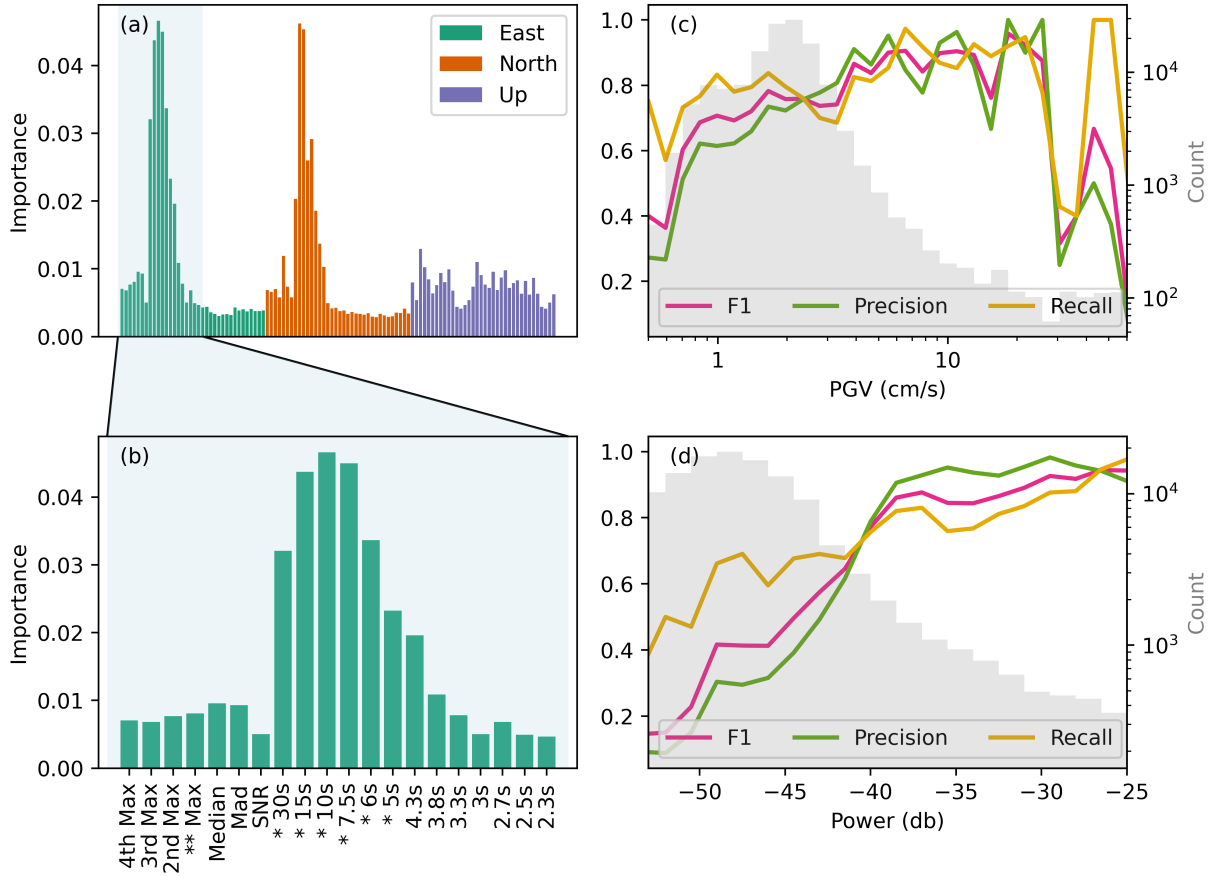

logic so that we could simulate running as a stand-alone instrument.

The mean precision, F1 and accuracy from our 10 fold test of our random forest classifier outperforms the existing threshold approaches (Figure 3.6). In the threshold approach, recall is higher than the random forest classifier; given the large number of false positives that this method triggers, we believe this value is boosted by chance noise triggers occurring in windows of true motion triggering the motion boolean. This further demonstrates the value of optimizing on F1 as a balance of precision and recall to reduce biasing one decision criteria. Precision is low for both the threshold method and the STA/LTA, but for different reasons; while the precision values (Equation 4.1) are nearly identical, the threshold method suffers from a relatively high amount of false positives, whereas the STA/LTA method low score is due to a lower amount of true positives. This discrepancy is evident in the accuracy scores, where the STA/LTA outperforms the threshold approach. False positives would be decreased if using additional external information as their authors' suggest, such as stricter time window approaches and correlating in space within networks. Such an approach would also likely improve the random forest classifiers performance but limit the utility of a stand-alone detection node. Spatio-temporal information could be incorporated into future network decision criteria.

3.4.3 Edge Sensitivity Detection

Detecting the largest amplitude velocity waveforms relative to ambient noise does not present a significant challenge outside of mitigating false alerting from sporadic outliers (Figure 3.7-a), with a 98% true positive rate of events greater than M_W6.0 and less than 100km radius. The random forest classifier's balance of improved false alerting relative to thresholds and improved sensitivity relative to the STA/LTA is evident for these highest seismic risks. To further investigate the random forest model performance we evaluate detecting signals closer to the noise floor. For simplicity, we bin seismic motion edge case detection into two distinct classes in what is a continuous distribution: large magnitude

Figure 3.6: Performance metrics for 3 methods in stand-alone mode without external triggers or correlation. Threshold is the NCX2-995 approach used by [Dittmann et al., 2022a] that thresholds the noise based upon the 0.995 significance of a non-central chi-squared distribution of the ambient noise. STA/LTA is based on [Psimoulis et al., 2018] GNSS motion detection modified STA/LTA algorithm. RF-ML is the method presented in the work here. Optimizing on F1 in this study allows us to balance missed detections (recall) with false alerts (precision); given the amount of false alerts of the Threshold and STA/LTA, the higher recall score could be a result of regular noise triggering events.

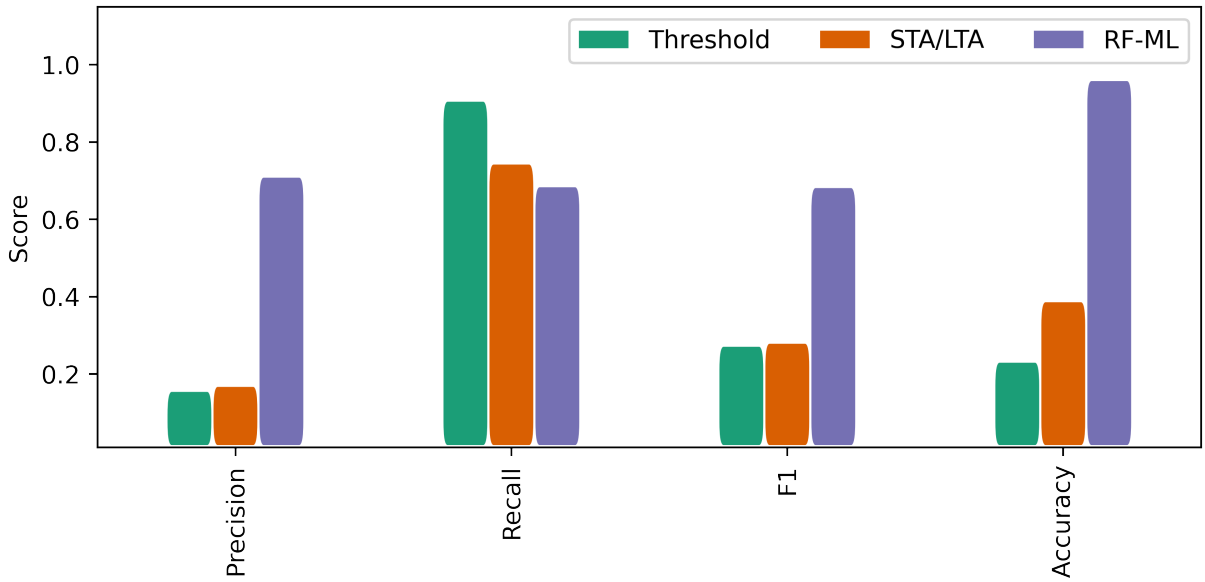

Figure 3.7: Performance of Random forest model developed in the work here across the entire event catalog. We reduce detection of events to a single binary for the figure. In this, each event is evaluated in a "test" split during the nested validation pipeline. This approach ensures each result depicted was evaluated as "unseen" relative to the best fit model from the training subset, and therefore representative of our model's future performance.

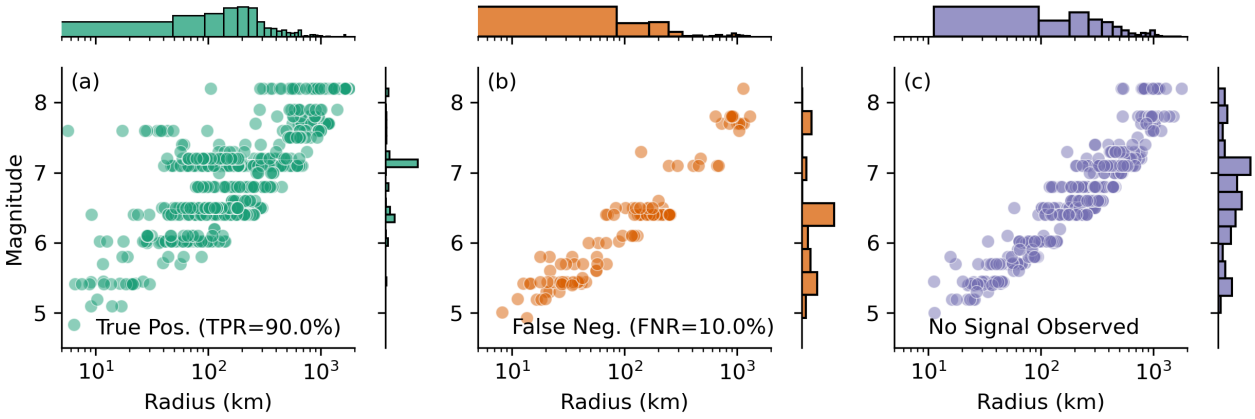

event seismic motion detection in the far field, and smaller magnitude events detected in the nearfield.

In the relative nearfield, much of the seismic energy passes through a station in shorter duration, varied frequency signals. Earthquake focal depth and fault slip distribution in time and space can significantly vary these waveforms as observed. Critically, the waveform signatures can appear similar to those of non geophysical processing outliers which we wish to ignore for this classification. Most existing STA/LTA methods filter these noise signals but also these valuable higher frequency dynamics. In the previous threshold methods, detection of these edge cases was a function of the ambient noise level, with low precision resulting (Figure 3.6) as a result of a high false positive rate. Our classifier has far less false alerts than the threshold approach in these signals, but nevertheless still presents the hardest detection domain for our classifier, evident in the missed detections of Figure 3.7-b of events less than M_W6.0. The left panels of figure 3.8 is an example of a smaller magnitude event (M_W5.4) in the relative nearfield (21km). The top 4 panels ([a:d]-0) on the left of figure 3.8) demonstrate that accurately detecting such an event using the threshold or modified STA/LTA approach is difficult; not only does the true signal barely exceed the noise floor, but there are numerous false alerts using both methods. The random forest classifier captures the majority of labeled motion window in addition to "ignoring" the spurious disturbance around 100s OT that triggers all other methods evaluated 3.8(e-0).

The sensitivity of GNSS to longer period surface waves are apparent at relatively great radii in the 5 hz TDCP velocity time series (Figure 3.7). The model detects teleseismic surface waves in unfiltered GNSS velocities from a M_W8.2 (USGS event ID: ak0219neiszm) at 1780km epicentral radius in real-time with no external corrections; the right-hand panels of figure 3.8 provides an example of this detection. Future analysis could investigate the range of geodetic teleseism detection with respect to larger magnitude event directivity, attenuation and observational networks. In Figure 3.8d-1, the amplitude of the ground velocity magnitude of these long period signals is insufficient to cross the traditional noise thresh-

Figure 3.8: Velocity and detection time series of two stations observing different events. The left column is from P507 observing a M5.41 at 21km epicentral radius; the right is from AB18 observing a M_W7.9 from ∼1400km epicentral radius. In the top velocity component panels (a-c), we include a downsampled running mean so that the reader may readily visualize the lower frequency surface waves passing through. The teal vertical lines are alerts from the STA/LTA classifier [Psimoulis et al., 2018] on each component. Panel (d) green timeseries is the 3 component ground velocity; the red horizontal line is the sensitivity threshold of a 0.995 non central chi squared (ncx2) noise model [Dittmann et al., 2022a], with orange vertical lines indicating a potential alert where GV greater than the threshold. (e) panels are a comparison of the labeled feature set 1 for these event-station pairs in purple, and the results of the model prediction in red. Shading is used to distinguish overlapping windows. This event-station pair prediction is extracted from the test or unseen event collection.

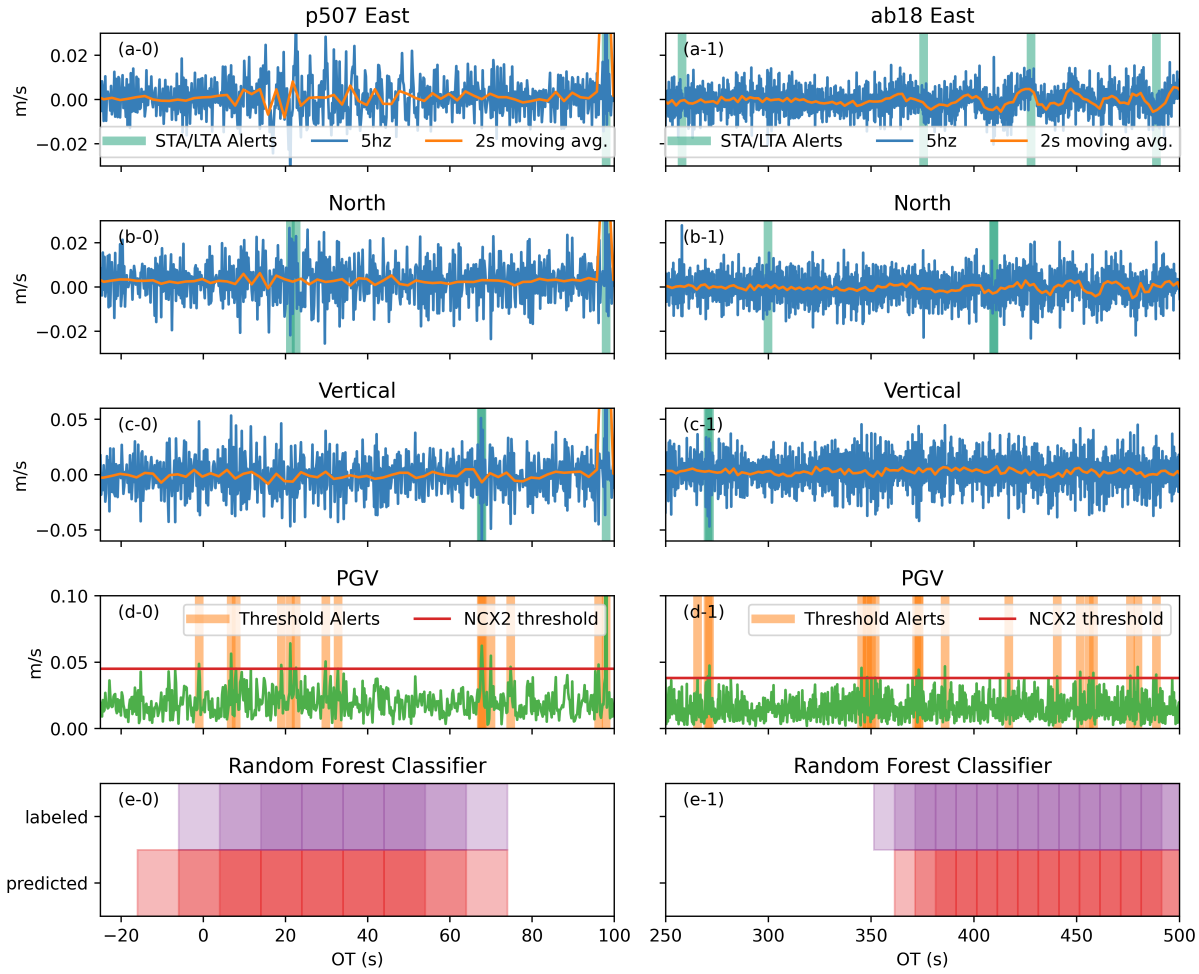

old with consistency, and there are many antecedent false alerts. The modified STA/LTA RT-Shake approach does not identify the majority of the long period waves either (Figure 3.8a/c-1), while the random forest classifier in the bottom panel only misses the first window (Figure 3.8e-1).

3.4.4 Decision Latency

Delay in alerting is critical to EEW. While our model is trained, tested, and validated on overlapping windows every 10 seconds, we evaluate running the model at once per second, the current US EEW [Murray et al., 2018] geodetic input rate (Figure 3.9). On testing data not used in model training, we find a delay relative to the estimated P-wave, ∼3-5 seconds under 15km) exists in the current approach. Coarse P- and S-wave arrivals are estimated using the iasp91 model [Kennett and Engdahl, 1991]; future work more accurately quantifying these phase arrivals such as the approach of [Goldberg et al., 2018] would not only more accurately represent timing performance but also useful for training more sophisticated ground motion models. GNSS velocities using this current approach cannot reliably be used for earliest phase picking, but can rapidly contribute to ground motion models or peak motion scaling laws [Fang and Liu, 2020]. Given the feature importances of the classifier (Fig 3.5-a), we interpret delays to be the result of the classifier trained on the relatively longer period signals visible within the noise. Depending on source magnitude and travel path, these could be P-, S-, surface-waves or some convolution of the energy traveling through the GNSS location. Variance in delays in the near field are likely due to inherent limitations of modeling rupture as a point source at proximal locations [Goldberg et al., 2021] and possibly related to errors of the iasp91 travel time model. Future work will address the possible limitations or delays introduced by our visual classification labeling. It is worth repeating that this assessment uses no external input or seismic triggering.

Figure 3.9: Time of first detection of all individual event-station pairs within 70km radius relative to estimated P-wave arrival time (iasp91 travel time model) as a function of radius. Green dots are the estimated S-wave arrivals at the event-station pairs used in this study shown for reference. Purple circles are centered on the time of first detection after the OT, where the diameter is scaled to the event magnitude. These results are from the classifier run at 1Hz on unseen testing sets to simulate a real-time operational mode.

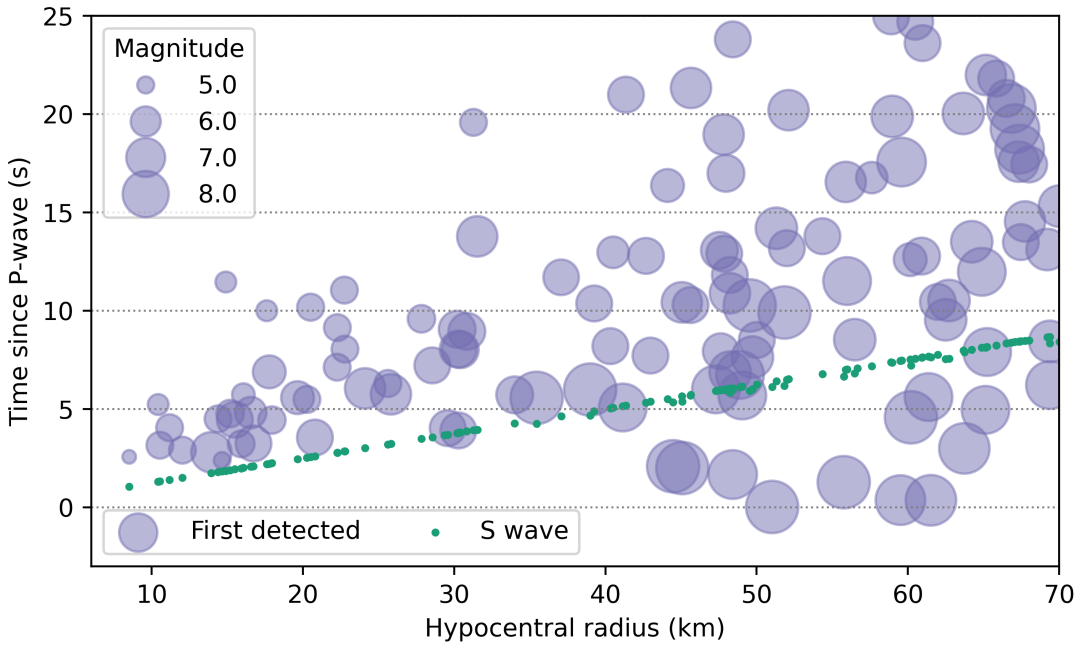

3.4.5 Ambient Noise Dataset

In addition to evaluating the performance within the bespoke event and noise data sets, we also evaluated the performance of the method during periods of quiescence to further quantify relative false alert rates. Our unseen testing set consisted of 1321 30-minute velocity timeseries from 2019-2021, not included with the original nested cross-validation data. We ran 5-fold cross validation on the entire event and noise labeled dataset from the nested cross-validation pipeline (Sect. 3.3.1) to select hyperparameters for training a complete model on all available labeled data for future "unseen" events. Such "unseen" events include this set-aside noise testing set. We confirmed there were no concurrent events greater than M_W4.0 in the USGS COMCAT catalog within the relevant spatial footprint and all other sources of noise or disturbances (signal multipath, oscillators, atmospheric anomalies, etc) remained in the test set. We assigned labels of **non event** to all target vectors associated with feature extraction. This allowed us to quantify ambient noise performance, or false alarm rate (Figure 3.10-a) using the detection methods previously described in section 3.4.2 from 860 unique stations from Alaska to the Caribbean across a range of potential TDCP noise or disturbance sources.

The random forest classifier was less susceptible to false alerts over the window tested than the threshold and STA/LTA approaches. The two threshold models have the highest rates to false alerting, an anticipated result based upon the precision metric reported in Figure 3.6. Station variations present in the random forest approach (Figure 3.10-b) suggest the current random forest model has some station or time noise dependence not correlated with the variations of other detection methods. Future inclusion of more extensive noise training datasets into our detection classifier and possibly data augmentation techniques would likely be beneficial towards training on the widest variety of noise scenarios and optimizing feature engineering for these complex noise environments.

Figure 3.10: Panel (a) is mean false positive rates (FPR) from 1321 spatially distributed, 30 minute duration of TDCP 5Hz velocities from windows prior to events in 2019-2021. Methods include: median plus 3 times the median absolute deviation threshold of [Hodgkinson et al., 2020], non-central chi square of [Dittmann et al., 2022a] NCX2 using alpha value of 0.995, the modified STA/LTA implemented by [Psimoulis et al., 2018] and the random forest machine learning classifier developed in this work (RF-ML). Panel (b) is a distribution of each method of a randomly chosen subset of stations to illustrate some of the station variability to the reader.

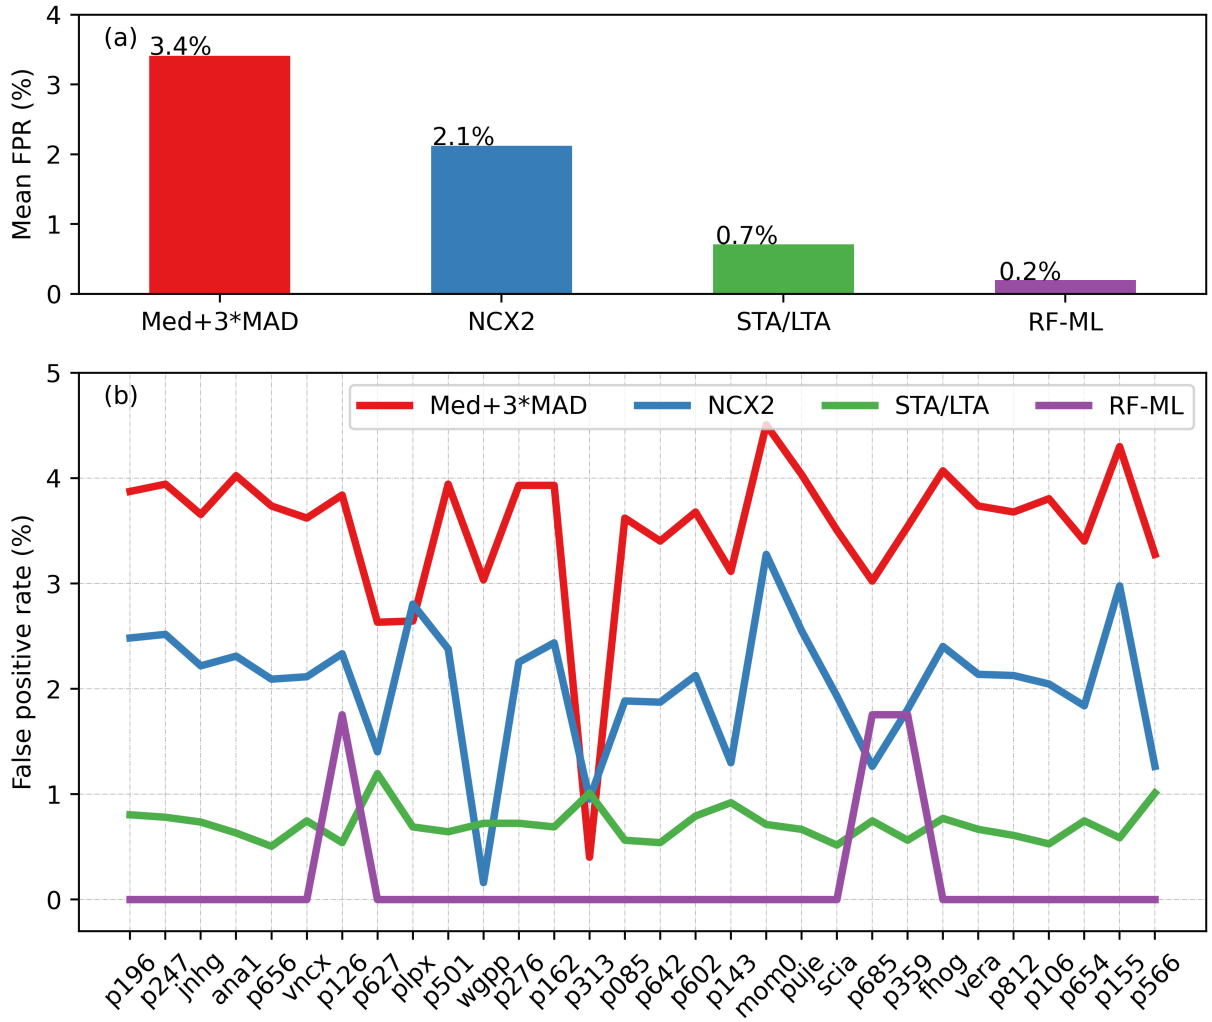

3.5 Conclusion

We applied an existing machine learning algorithm and sample splitting pipeline techniques to training, validating and testing a seismic motion detection classifier from 5Hz TDCP GNSS velocities. We leveraged nearly 20 years of 5Hz GNSS data archives for training a classification model that outperforms existing threshold approaches for detecting motion in stand-alone mode. The classifier combines time domain and frequency domain features to match the sensitivity of the threshold method without the false alerts, and matches the minimal false alerting of the STA/LTA with improved sensitivity. Given the agreement that GNSS velocities have with existing ground motion models [B. Crowell and Ghent, 2022] and the increased confidence in separating signal from noise demonstrated here, these GNSS velocities can operationally contribute to ground motion measurements. The alert latency of this current model does not match the sensitivity of existing inertial infrastructure. A complementary approach using the information available at the time, including lowest latency p-wave characterization from inertial sensors and unsaturated velocity estimation from GNSS provides an optimal solution for existing dense multi-sensor networks. For less dense networks of either sensor type, it is more critical to establish a decision criteria for balancing timing, noise and accuracy of these independent observation systems. Further investigation of integrating the processing and classifying approach of this manuscript with the sensitivity of co-located MEMS sensors [Goldberg and Bock, 2017] would advantageously overlap seismic and geodetic traditional boundaries.

Current 5 Hz GNSS observation data streams are too verbose for many bandwidth limited remote hardware; this presents an exciting opportunity for edge processing at potentially much higher rates [Shu et al., 2018], or experimental lean 5 Hz carrier phase data streams. Our method presented here does not use a sophisticated machine learning model, yet has improved detection relative to existing approaches; much improvement remains, especially with expanded datasets across global geodetic networks and/or synthetics or data augmentation

for training, validation and testing of neural networks and deep learning models.

With an expanding availability and access to real-time GNSS streaming networks, the seismological community stands to benefit from this signal of opportunity for rapid ground motion detection for earthquake and tsunami source characterization. Furthermore, the vast industry of GNSS position, navigation and timing users catalyzing the expansion of these GNSS real-time networks will benefit from improved automated alerting of reference station motion onset. Future work will include integrating this classifier amongst existing and future automated GNSS carrier phase disturbance characterization methods, including space weather disturbances [Jiao et al., 2017], oscillator anomalies [Liu and Morton, 2022], radio frequency interference and signal multipath.

3.6 Open Research

The 5Hz GNSS data used for TDCP processing in the study are available from the Geodetic Facility for the Advancement of Geoscience (GAGE) Global Navigation Satellite Systems (GNSS) archives as maintained by UNAVCO, Inc. The data are available in RINEX (v.2.11) format at `https://data.unavco.org/archive/gnss/highrate/5-Hz/rinex/`. Earthquake depths, locations, and magnitudes came from the Advanced National Seismic System (ANSS) Comprehensive Catalog of Earthquake Events and Products (`https://earthquake.usgs.gov/data/comcat/`). Arrival times are calculated using the iasp91 velocity model as implemented by Incorporated Research Institutions for Seismology (IRIS) Web Services (`http://service.iris.edu/irisws/traveltime/`). SNIVEL code used for TDCP velocity processing is developed openly at `https://github.com/crowellbw/SNIVEL` (Accessed December 2021)[Crowell, 2021]. SNIVEL 5Hz velocity timeseries used in this study are preserved at `https://doi.org/10.5281/zenodo.6588601`. Version 1.0.1 of the scikit-learn software used for random forest classification is preserved at `https://doi.org/10.5281/zenodo.5596244` and developed openly at `https://github.com/scikit-learn/scikit-learn` [Pedregosa et al., 2011]. Version v0.5.0 of PyGMT used

for generating the map is preserved at https://doi.org/10.5281/zenodo.5607255 and developed openly at https://github.com/GenericMappingTools/pygmt[Wessel et al., 2019]. Figures were made with Matplotlib version 3.5.1 [Caswell et al., 2021], available under the Matplotlib license at https://matplotlib.org/.

This material is based on services provided by the GAGE Facility, operated by UNAVCO, Inc., with support from the National Science Foundation, the National Aeronautics and Space Administration, and the U.S. Geological Survey under NSF Cooperative Agreement EAR-1724794. High-rate processing and and machine learning for geoscience and hazards research is supported by NSF OAC-1835791.

Chapter 4

GNSS Velocity Noise Characterization for Augmented Nearfield Strong Motion Learning

4.1 Abstract

Data-driven approaches to identify geophysical signals have proven beneficial in high dimensional environments where physics based model-driven methods fall short. GNSS offers a source of unsaturated ground motion observations; these signals are the data currency of ground motion forecasting and rapid seismic hazard assessment and alerting. But, these signals are superposed onto hardware-, location- and time-dependent noise signatures influenced by the earth's atmosphere, low-cost and or spacebourne oscillators and complex radio frequency environments. Eschewing physics-based models for a data-driven approach in this context is a step forward in autonomous signal discrimination. However, the performance of a data-driven approach depends upon substantial representative samples with accurate classifications; more complex algorithm architectures for deeper scientific insights compound this need. The existing catalogs of high-rate ($\geq$1Hz) GNSS ground motions are relatively limited. In this work, we model the probabilistic noise of GNSS velocity measurements over a hemispheric network. We generate stochastic noise time series to augment transferred low-noise strong motion signals from an existing inertial catalog. We leverage known signal and noise information to assess feature extraction strategies and quantify augmentation benefits. When validated against an existing real-GNSS velocity dataset, this pseudo-synthetic catalog offers a framework for future enhanced data driven approaches.

4.2 Introduction

Distributed observations of coseismic ground motions are the backbone of accurate ground motion models, finite fault modeling and early warning. If available in real-time, GNSS-derived high rate time differenced carrier phase (TDCP) velocities [Graas and Soloviev, 2003] applied to seismology [Colosimo and Mazzoni, 2011] are an additional source of these intrinsic measurements that traditionally are produced by dedicated inertial sensor networks. If available for post processing, GNSS velocities can contribute to catalogs of ground motion measurements used for empirical regional and local ground motion models. GNSS spatially complements or replaces existing inertial ground motion observations [B. Crowell and Ghent, 2022], especially valuable in sparse networks [Grapenthin et al., 2017]. Furthermore, GNSS expands the dynamic range of most inertial measurements, and contributes to magnitude estimation [Melgar et al., 2015, Ruhl et al., 2018, Fang and Liu, 2020]when inertial sensors saturate [Melgar et al., 2013].

However, ambient GNSS velocity noise remains well above the noise floor of inertial sensors, largely due to sources of error of space-based ranging of weak radio frequency signals. Analysis of high rate positioning noise [Bock et al., 2004], carrier phase noise [Wang et al., 2020], and TDCP velocities [Shu et al., 2018, B. Crowell and Ghent, 2022] has shed valuable insight into the factors that influence the ambient noise floor of these GNSS velocities. To date, the GNSS velocity noise frequency spectrum has not been evaluated across sufficiently large temporal and spatial scales to statistically report on the ambient noise across a network. Ambient noise characterization methods developed in the seismic community offer a rigorous approach to characterize ambient noise for sensor network monitoring and calibration. Understanding the probabilistic spectrum of GNSS velocity noise offers insight into the limit of seismic signal detection in GNSS. This spectrum can also be preserved in representative synthetic noise generation.

Improved separation of this noise from true signals will broaden the range in which

these GNSS seismic waveforms will contribute material observations. Various methods for overcoming this signal to noise (SNR) challenge exist: variations on a short term average over long term average (STA/LTA) detection adopted from inertial seismic sensors resolve static offsets [Allen and Ziv, 2011, Colombelli et al., 2013] but filter valuable dynamics encoded in the waveforms; threshold based detection methods [Crowell et al., 2009, Hodgkinson et al., 2020, Dittmann et al., 2022a] capture dynamics but struggle to balance sensitivity with false alerting, and must mitigate false alerts with external dependencies such as spatially correlating or temporally windowing from seismic triggers; machine learning models combine a range of feature inputs to improve the decision confidence in separating signal from noise [Dittmann et al., 2022b] in stand-alone mode. However, the generalization performance of any such classifier or deeper learning model will ultimately be limited by the model selection, the extent of the labeled catalog for training and the quality of the labels themselves.

Previous GNSS seismic catalogs illustrate how limited the GNSS observed long-tail larger magnitude seismic events datasets are [Ruhl et al., 2018]. For example, the UNAVCO continuous geodetic archive began in 1993 and 5Hz high rate data retrieval in 2006. Decreased hardware costs coupled with commercial and scientific demand only relatively recently allowed for global high-rate network proliferation. Additional geodetic networks (eg INGV, GEONET, NZ) can supplement UNAVCO's high rate catalog worthy of inclusion, but on the order of doubling, not the order(s) of magnitude needed for deeper learning to answer more sophisticated questions. One solution to this small data challenge is synthesizing waveforms, such as FakeQakes [Melgar et al., 2016, Williamson et al., 2020, Lin et al., 2021]. This promising approach presents great potential but is expensive for generalizing across global rupture scenarios. Furthermore, great care must be taken to not bias results with **"unknown unknowns"** of fault models and ground motion propagation that might under perform in future events; cutting edge research is ongoing in this space. An intermediate real-world-data driven alternative is to transfer samples from a separate source of our signals of interest [Hoffmann et al., 2019]. Inertial sensors have existed at more locations for

far longer than the first positioning satellite was launched. Event catalogs of zero-baseline inertial measurements offer low-noise ground motion velocities to be transferred as our truth waveforms of accurately labeled samples. The GNSS noise probabilistic power spectral density (PPSD) characterization offers the necessary information to superpose stochastic noise for training over a range of noise conditions. The final component to improved generalization are the learning training decisions, including model selection and feature engineering. With appropriately applied domain knowledge to increasingly larger data volumes, the revolution of transferable classification and regression model algorithm development is readily adaptable to earth science questions [Bergen et al., 2019, Kong et al., 2019] To improve our understanding of GNSS velocity sensitivity relative to ambient noise, expand the quantity of available labeled training data and improve detection classification performance in a highly variable noise environment, we characterized the GNSS velocity noise frequency spectrum from which we augmented inertial velocity waveforms observed over 80 years with synthetic GNSS velocity time series. This manuscript presents a framework for expanding the available, accurately labeled GNSS velocity waveforms and evaluates the improved signal detection gained from learning on such a catalog. We evaluate additional high-return features with an eye towards a goal of combined, larger GNSS velocity catalog with improved feature selection allowing evolving machine learning models, within an immediate objective to minimize overfitting of small datasets by expanding the available catalog through transferred signals and data augmentation.

4.3 Materials and Methods

4.3.1 TDCP Processing

A GNSS receiver generates precise relative phase estimations of the signal carrier wave through its phase lock loop. In order to achieve absolute positioning using carrier phase measurements, a suite of measurement error source models must be estimated to account

for thermal noise, satellite and receiver oscillators, multipath reflections, atmospheric and ionospheric effects from a 20,000 kilometer signal propagation path, and unknown carrier cycle integer offsets [Teunissen, 2020]. These correction models inject a cost, both computationally, potentially monetarily, and in performance for resolving carrier phase ambiguities to estimate absolute position. In past and current implementations of using geodetic measurements for capturing earthquakes, absolute positions are differenced from an a priori position to extract relative topocentric motion, the signal of interest. Time-differenced carrier phase or variometric processing [Graas and Soloviev, 2003] differences these precise carrier phase measurements in consecutive epochs to remove temporally correlated error sources and consistent integer ambiguities. TDCP uses the precision of these measurements to its advantage, by foregoing absolute positioning in exchange for precise relative velocity measurements while still benefiting from measurement observability across a visible satellite constellation. TDCP does not require ambiguity resolution convergence, the lack of complex error models minimizes measurement noise, and the lack of computational requirements combined with the simplicity of the method makes it ideal for seismic ground deformation applications [Colosimo and Mazzoni, 2011, Benedetti et al., 2014, Hohensinn and Geiger, 2018, Grapenthin and Freymueller, 2018] at higher rates and potentially on the network edge. We use the SNIVEL processing method [Crowell, 2021] for estimating 5Hz GPS TDCP. This method uses the narrow lane GPS-only L1/L2 phase combination, the Klobuchar ionospheric correction, the Niell tropospheric correction and broadcast satellite ephemeris. Observations are weighted as a function of satellite elevation angle with a seven degree elevation mask. While development accommodating precise orbits [Shu et al., 2020], multi-GNSS, cycle slip detection/mitigation [Francesca Fratarcangeli, 2018] and higher order noise source mitigation is ongoing and warranted, the current method is capable of capturing ground motions of nearfield M4.9 and larger, and larger sources at teleseismic distances [Crowell, 2021, B. Crowell and Ghent, 2022, Dittmann et al., 2022b].

4.3.2 TDCP Noise Model and synthetic time series

The UNAVCO geodetic archive captures 5Hz data of stations recording concurrent with larger magnitude earthquakes. This includes at least 1 hour of "ambient" 5Hz data antecedent to the hour in which the event takes place. We process with SNIVEL all available 5Hz pre-event hour long windows for our ambient GNSS velocity dataset. This dataset consists of 1507 hours from 904 stations since 2007 distributed from the Caribbean to Alaska. We use this sample space to be representative of GNSS velocity distributions both spatially and temporally.

We evaluated the spectrum of GNSS TDCP noise over this sample set by adopting a seismic ambient noise characterization method of [McNamara and Buland, 2004] modified for GNSS displacements by [Melgar et al., 2020]. In this approach, further modified for 5Hz GNSS velocities, we calculated the power spectral density of 10 minute 5Hz single component velocity windows. We evaluated power spectral densities (PSD) at periods from 205s down to 0.4s in 512 bins. PSDs were smoothed in octave intervals and then stacked across 73 aligned frequency bins over all available PSD segments. The result is a probabilistic power spectral density (PPSD), or distribution of power spectral densities over the samples included. These PPSDs have been adopted for seismic network monitoring [Casey et al., 2018] and offer valuable insight for anticipated signal sensitivity. We combined horizontal topocentric components into a single PPSD and then estimated an independent vertical PPSD, given GNSS position noise is approximately 3-5 times larger.

We stored 19 distribution slices (every 5th percentile from 5% to 95%) of the real-world noise quantiles from which to generate synthetic stochastic noise time series. We adopted the approach of [Melgar et al., 2020], first proposed by [Boore, 1983] and further developed by [Graves and Pitarka, 2010]. In this approach, we were able to maintain the frequency content of the noise at respective reference levels while randomizing the phase for generating multiple unique time series. We accommodated amplitude loss in the domain

transformations with linear scaling. To see a similar approach used with synthetic GNSS displacement noise/waveforms, see [Lin et al., 2021].

4.3.3 Strong Motion Observations and Augmentation

Our signals of interest are velocity waveforms from medium to larger earthquakes (¿M5.0) of which GNSS velocities are sensitive to [Dittmann et al., 2022a]. The Next Generation Attenuation for Western United States 2.0 (NGAW2) project [Ancheta et al., 2014] is a database of global strong motion measurements and response spectral ordinates spanning over 75 years including 21,339 three component records from 599 events ranging M3.0 to M7.9. Global seismic networks contribute strong motion accelerograms or broadband velocity measurements that are processed by the NGAW2 project into acceleration, velocity, and displacement waveforms. The processing consists of an acausal Butterworth filter to reduce high- and low-frequency noise and an instrument response correction; further information regarding processing can be found at [Ancheta et al., 2014]. The records are visually inspected for corner frequency determination, quality and completeness, making the catalog an ideal source of low-noise larger ground motion measurements. The primary applications of such a catalog is for ground motion prediction research to inform seismic engineering; we use the processed velocity waveforms as our noise-free signal. It's worth noting the seismic community has extensive labeled datasets to extend their potential model selection. This has facilitated rapid growth in neural networks and deep learning models pushing the limits of detection, signal denoising and earthquake source characterization and ground motion propagation. We considered the several curated seismic ML catalogs [Woollam et al., 2022], but found these catalogs emphasized smaller magnitude signals; logical given the signal-to-noise challenges from weaker inertial measurements looking to ML for use in seismology, but insufficient amplitudes for synthesizing our GNSS strong motion geodetic measurements. We focus our effort on the portion of the database containing nearfield ($\leq$70 km radius) observations of M5.0 to M7.9 within expected sensitivity radius of 1cm/s given the scaling

Figure 4.1: Example of three noise levels of GNSS TDCP synthetic noise generated from PPSD slices.

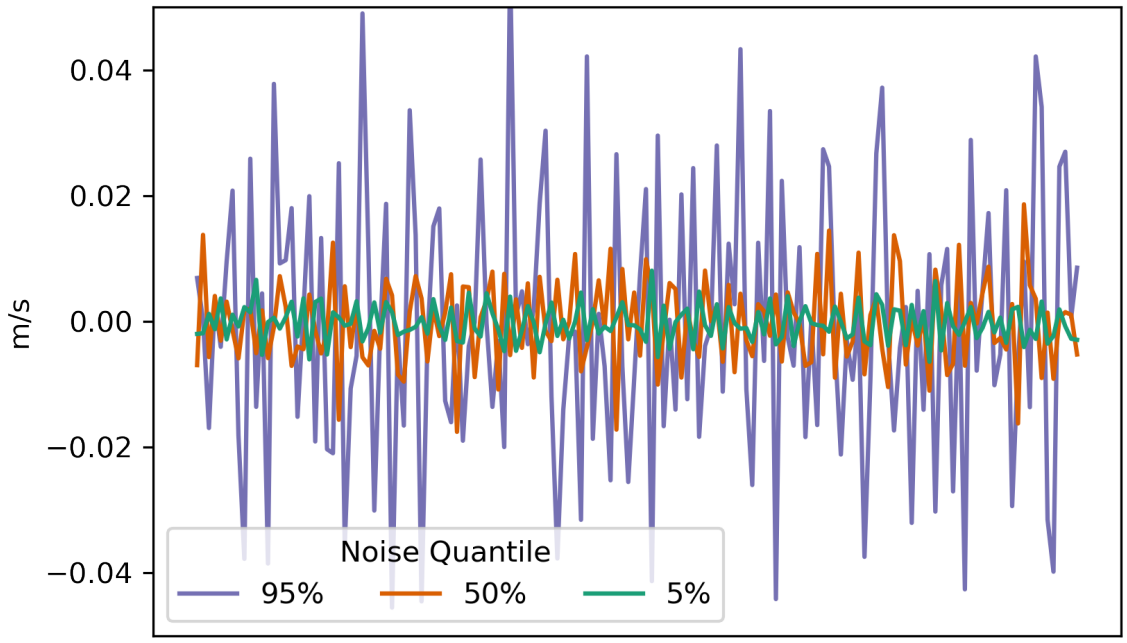

laws of [Fang and Liu, 2020] for rapid hazard applications; future work is extensible to the limits of detection above the noise floor (>1000km). We collected 1846 waveforms from 173 events. The processed velocity time series are offered at either 100 or 200 Hz sampling rate; we low pass filtered with a filter corner frequency of 2.5Hz and then downsampled to 5Hz. We adopted a recursive short-term average over long-term average (STA/LTA) detection algorithm to label ground motion on each individual component. We found this is a sufficient detector given its performance [Withers et al., 1998] and factoring in the subsequent noise injected into our system. We used a 5 second short-term window and 10 second long-term window with a detection threshold ratio of 1.5; through trial and error we found this a reasonable balance of sensitivity for our larger strong motion signals of interest [Trnkoczy, 2012]

Figure 4.2: (a) is a histogram comparing the UNAVCO 5Hz catalog ("GNSS') with the NGA West-2 database ("NGAW2") for events observed by stations within 70 kilometers and sensitivity radii. The scatter plot in (b) are the individual event magnitudes as a function of time, and the secondary axis line plot is the cumulative station count over time observing the events.

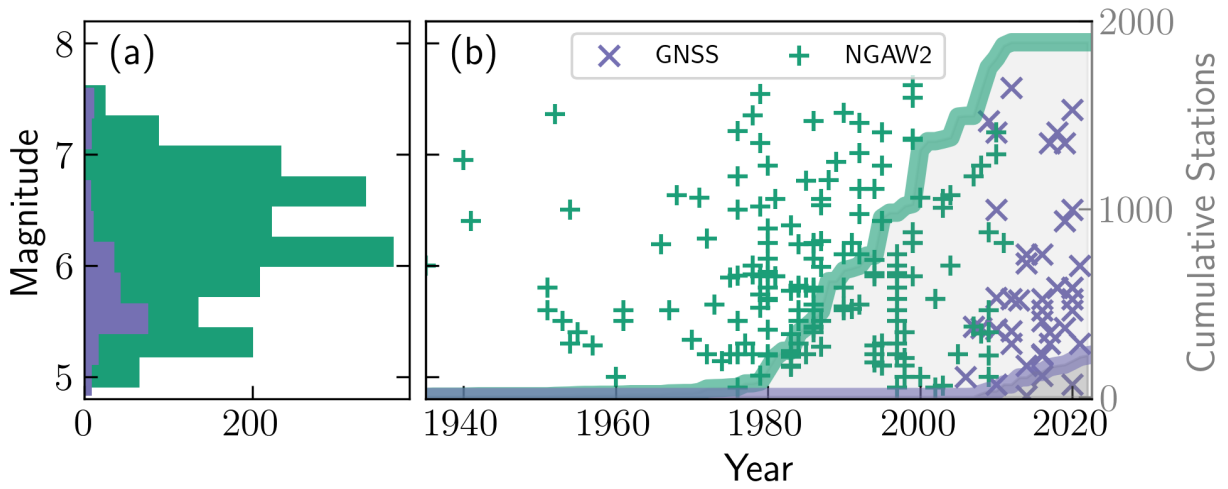

We exploited our "noise-free" signal waveforms and realistic stochastic noise generation by adopting data augmentation of transferred signals. Data augmentation is a form of

regularization in which the size of a data catalog is artificially increased by creating augmented copies of our original waveforms [Zhu et al., 2020]. Augmentation not only expands extents of a data catalog, valuable in relatively limited event datasets such as ours, but also improves generalization [Bishop, 1995]. Successful augmentation trains a classifier to learn features or patterns in the presence of a larger range of authentic noise factors [Iwana and Uchida, 2021]. In our application, we injected a synthetic noise time series derived from a single reference level of noise spectrum with unique phase values. We did this at seven noise reference levels from the 5th to 95th percentile to augment each strong motion waveform, a form of magnitude augmentation or jitter. We also buffered each augmented waveform with a random number of samples to misalign the samples in time relative to each waveform replica. This resulted in seven different pseudo-synthetic observational waveforms for each station-event pair. This approach minimized overfitting in our models by training on a range of noise for a given signal at different offsets in each feature window, and expanded our catalog seven-fold from 1,846 strong motion waveforms to 12,922 pseudo synthetic GNSS velocity waveforms.

Additionally, we included the ambient catalog used in creation of the PPSDs to ensure the classifier is both trained on and tested against real-world GNSS velocity noise. This strategy was particularly important for potential disturbances not captured by the ambient synthetic noise generation process, such as the most infrequent events that might get statistically removed from the stochastic power spectrum but could prove detrimental to false alerts if their signature is not learned. For example, the lowest frequency offsets from processing artifacts are infrequent enough to barely impact the probabilistic spectrum but if not trained on could present as a synchronized event. We validated the performance of training a classifier on this synthetic catalog against the previously labeled UNAVCO 5Hz GNSS velocities (see Open Research). For description of this dataset, please refer to [Dittmann et al., 2022b]. This curated catalog of GNSS velocity waveforms was processed identically as the noise catalog of this work; but one fundamental difference is it is labeled through visual

Figure 4.3: Example of three component waveforms from a single event NGAW2 waveform from Chi-Chi, Taiwan (2003, M6.2 50Km radius) with three levels of synthetic noise added.

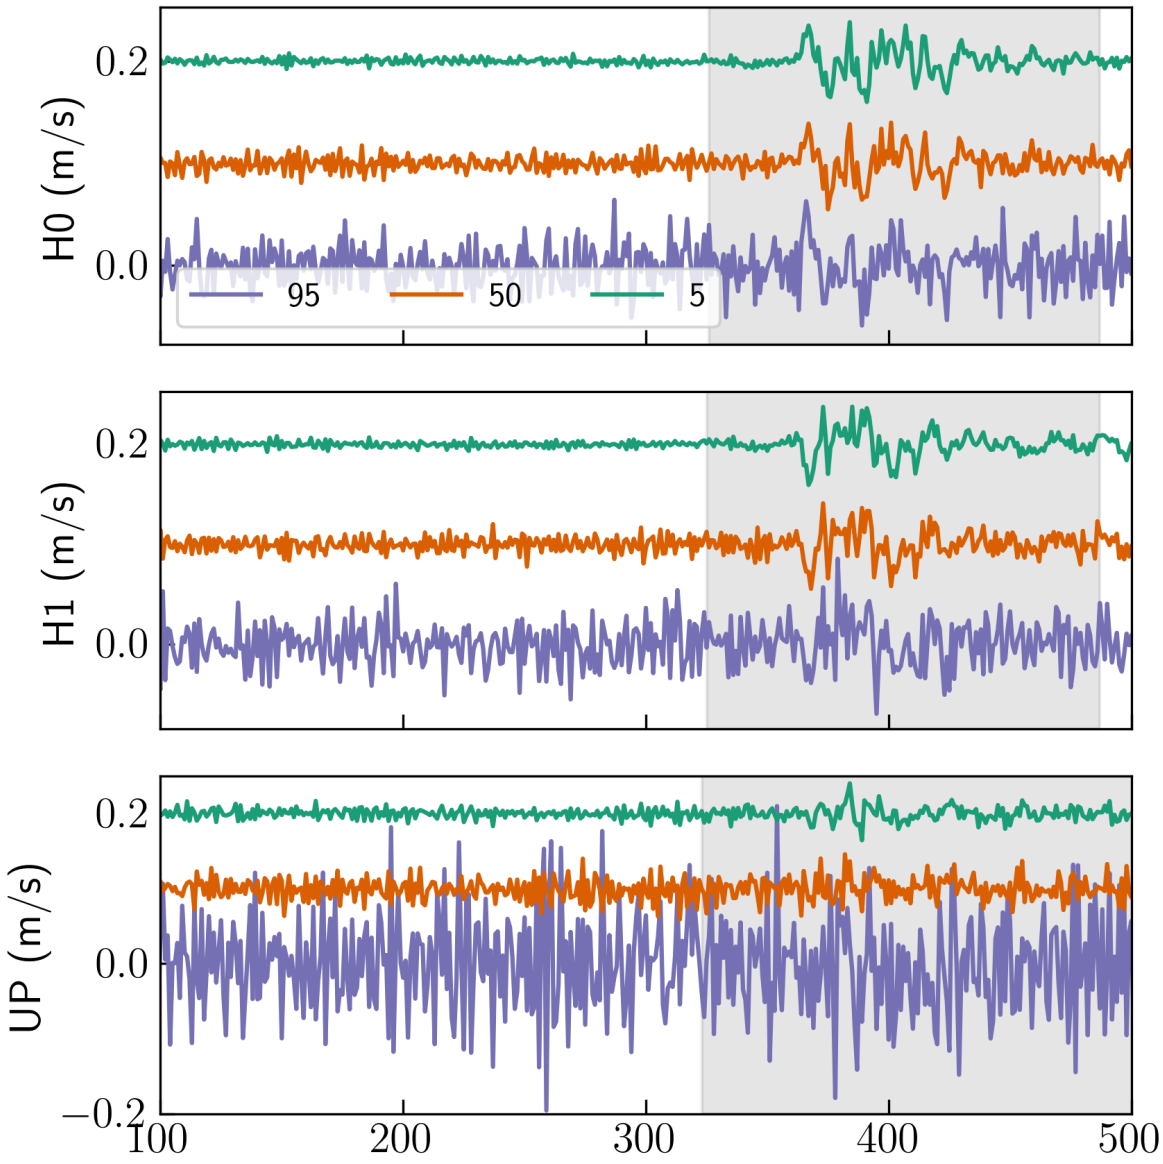

inspection instead of a known "truth" of our lowest noise inertial waveforms.

Table 4.1: Extent and strategy of catalogs used in this research of noise and M5+ events within expected detectability and 70km radius.

	Number of Station-Event Waveforms	**Number of Samples**	**Labeling Strategy**
GNSS Event Catalog (<70km)(Dittmann 2022b)	247	5,187	visual inspection
Ambient Noise Training	1,507	88,893	assumed event-free
Ambient Noise Testing	1,507	85,806	assumed event-free
NGAW2	1,846	44,548	zero-noise truth labels
NGAW2 with Augmentation	12,922	311,836	zero-noise truth labels

4.3.4 Model selection, Feature Engineering and Training

First we validated the performance of a classifier trained on our strong motion waveforms relative to our previous GNSS velocity catalog approach. We used a random forest classifier [Breiman, 2001] for our detection model. Random forest is an ensemble method of decision trees. A decision tree is an algorithm that splits inputs along features to classify samples. A single decision tree can be biased by the initial features selected to seed the splitting; random forest overcomes this by running an ensemble of decision trees and having each cast a vote. We set up a binary classification that is demonstrated to have high accuracy and balance of sensitivity and false alerting in GNSS velocities. By keeping our model consistent with our previous work, we validated the newly formed catalog.

For validation comparison, we preserved our strategy from [Dittmann et al., 2022b] of 30s overlapping windows; future work will further optimize this sampling strategy with respect to sensitivity and real-time performance. From each window sample, we extracted

a series of features to test their performance for our signal detection classification. In the time domain, we extract metrics akin to the traditional thresholding methods, including the four largest amplitudes, the median and the median absolute deviation. In the frequency domain we included the entire PSD range over the 5Hz sampling of 30s windows, which includes periods from 0.4 seconds to 30 seconds. Variations on both of these time and frequency metrics were evaluated in our previous work, with the lower frequency (3s-15s period) horizontal PSD the most influential for the classifier model. However, while the overall performance over the entire catalog was excellent, variability in the false positive rate of the ambient dataset combined with missed detections of nearfield smaller magnitude events encouraged investigation of expanded features.

The strength of the previous classification model is attributed to the ability to combine information from the time and frequency domain. The continuous wavelet transform is an appealing compromise for detection that provides temporal information while still accumulating energy in scale (approximately equivalent to frequency) bands for our signals of interest [Gabor, 1946]. Wavelets consist of a host of compact mathematical functions that are convolved with an array of data to estimate correlation energy as a function of time and scale; scale can be related to frequency. These have been successfully applied to a variety of non-stationary geophysical processes [Kumar and Foufoula-Georgiou, 1997], including seismic signal detection [Simons et al., 2006], GPS deformation monitoring [Min et al., 2006] and GPS de-noising [Kaloop and Kim, 2016, Wu et al., 2011]. In our particular 30s window classification, we did not preserve the valuable temporal information gained from this approach, and instead aimed to maximize the non-stationary deterministic features present in the resulting wavelet coefficients to boost early or weak detection in challenging noise environments [Bakshi, 1998]. We used a Morlet wavelet [Morlet et al., 1982], native to obspy and pywt software [Krischer et al., 2015, Lee et al., 2019]; we chose this wavelet for its past performance in seismological applications [Kristeková et al., 2009], although future wavelet selection analysis is warranted.

Each sample consisted of one or a combination of these features for 30 second windows for all three components (Figure 4.4). STA/LTA labels were reduced to a single positive or negative outcome from 450 samples (150 samples per window x 3 components). Given our knowledge of signal relative to noise in this synthetic dataset, we also assigned a SNR metric for each sample, which was the peak single difference between signal power and noise power across all frequency bins. We employed a similar nested cross validation approach to our previous work for comparison and validation. Because the number of discrete events is still relatively small, we wished to minimize the potential bias from random validation and testing set selection.

In nested cross validation [Bishop and Nasrabadi, 2007], we ran 10 different testing scenarios, where each scenario keeps aside a different subset of one tenth of the events. Within each fold, we also ran an inner loop of 5 fold cross validation across a grid search of hyperparameters. This technique further minimized overfitting hyperparameters by cross validating across a range of sample subsets. Our hyperparameters included the depth of nodes, or the number of decision splits, the number of estimators or decision trees, class weighting, a strategy that can assist with imbalanced datasets such as ours, and finally a SNR training threshold. This last hyperparameter was uniquely available to this pseudo-synthetic dataset; we generated the noise added to the signal, and so with this information we can accurately quantify the relative detectability. Using this as a hyperparameter allowed us to optimize training sets to include the largest extent of low signal-to-noise samples that benefit the model, while avoiding degrading model performance with undetectable low SNR. We optimized our model on F1 scores, a balance of precision and recall.

Precision, or positive predictive value, is equal to the number of true positives (TP) over the sum of TP and false positives (FP).

$$Precision = \frac{TP}{TP + FP} \tag{4.1}$$

Recall, or sensitivity, is the number of TP over the sum of TP and false negtives (FN).

$$Recall = \frac{TP}{TP + FN} \tag{4.2}$$

F1 is the harmonic mean of precision and recall:

$$F_1 = 2 \times \frac{precision \times recall}{precision + recall} \tag{4.3}$$

Here, positive denotes motion and vice versa.

4.4 Results and Discussion

4.4.1 Noise Characteristics

Understanding GNSS noise is imperative to applying GNSS observations to answer complex geophysical questions. Such investigations range from low frequency estimation of secular plate velocities [Williams et al., 2004] to higher frequency (>1Hz) signals, including structural monitoring [SHEN et al., 2019, Hohensinn et al., 2020], space weather [Yang et al., 2017] and deformation monitoring [Geng et al., 2018, Avallone et al., 2011]. Previous studies inform that GNSS position noise is a combination of white and colored or power-law noise [Langbein and Bock, 2004]. Starting from lowest frequencies, the "dam profile" of exponentially decaying noise with increased frequency is inferred to be a result of correlated signal path and processing contributions including multipath, ephemerides, clocks and atmospheric effects. GNSS highest frequency position noise is attributed to receiver thermal noise and often presented as a white spectrum [Genrich and Bock, 2006]. Receiver thermal noise is parameterized as a function of incoming signal strength and carrier phase tracking filter design, including filter bandwidth and sample integration time. These baseband signal tracking loop design choices balance dynamic stress response with thermal noise mitigation [Yang et al., 2017], and are reflected in this highest frequency noise profile [Moschas and Stiros, 2013, Häberling et al., 2015]. As an aside, for these reasons a calibrated high frequency instrument response, similar to what has become the defacto standard in digital

Figure 4.4: Demonstration of waveforms, noise and feature selection. The green timeseries in (a) is a downsampled NGAW waveform of a relatively weak signal for our application (a M5.5 at 12.5 km); the orange is a randomly generated noise time series using the 50th percentile noise spectrum; the gray shading is the region of detection triggered by the recursive STA/LTA. The sum of these time series (b) is then used as our observation; in the time domain the features selected include the 4 largest amplitudes, the median and the median absolute deviations, all indicated for this waveform in magenta. We compute a continuous wavelet transform for each 30s window, shown in (c.) and extract the peak values for each scale/frequency, examples of which are shown in magenta. Finally, we also compute the power spectral density using a periodogram (in purple) and extract the power at each frequency bin. The original signal and noise powers are shown as well, for reference, though they are not included in the feature extraction.

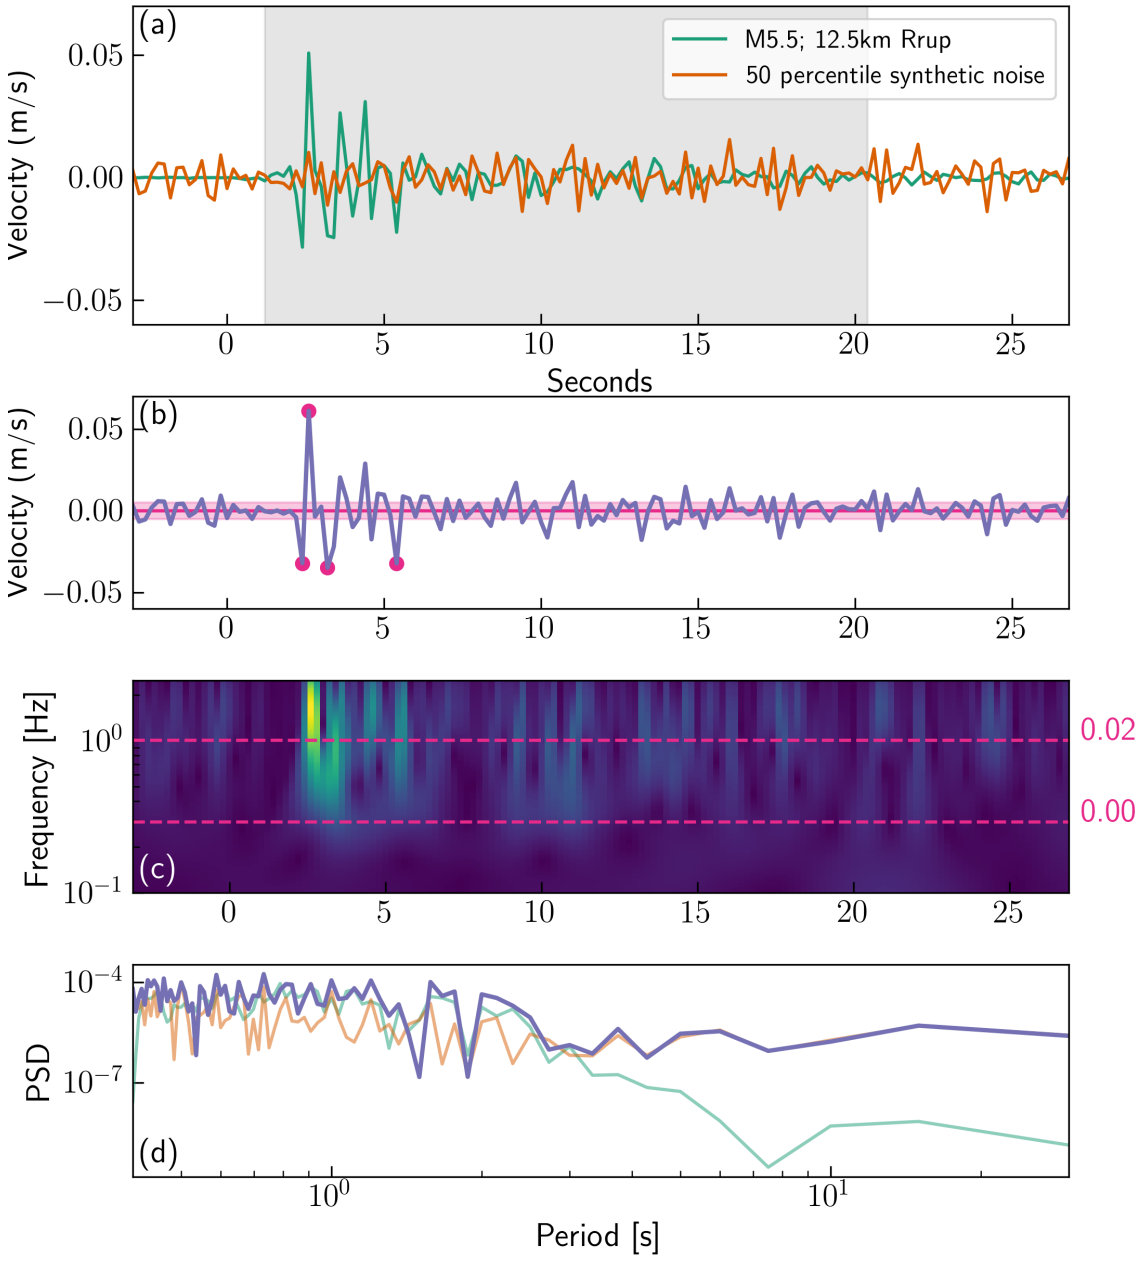

inertial instruments, has been proposed [Ebinuma and Kato, 2012]. We note this as worthy of further investigation for future efforts integrating TDCP velocity noise into monitoring but have not yet observed an instrument bias with respect to capturing seismic strong motion signals in 5Hz velocities.

In TDCP velocity noise, we observe a V-shaped noise spectral profile in the PPSD (Figure 4.5). Periods longer than 6s follow a power law profile, likely reflecting correlated errors such as multipath and atmospheric effects not completely removed in the time differencing. This result is aligned with [Melgar et al., 2020], which identified 1Hz PPP displacement noise as a red noise with a dam profile down to their Nyquist frequency. They infer that multipath and troposphere are the primary sources of the PPP "random walk" correlated noise signature (5s-200s period), and anticipate a spectral flattening to white noise around their maximum resolvable frequency (0.5 Hz). The Melgar, et al. (2020) 1Hz PPP PPSD had a corner around 3 seconds, while in TDCP the lower frequency power law corner is at 6 seconds period. Another notable difference with TDCP processing reflected in this profile is the absence of absolute atmospheric models; in TDCP, the single slant path phase differences with first order corrections remove all but higher order gradients. Unfiltered time-differenced velocities will not accumulate error from potentially biased corrections models, a challenge of PPP. [Shu et al., 2020] noted that inclusion of precise satellite clocks and orbits can significantly reduce longer period drifts existing in displacements derived from GNSS variometric velocities that otherwise must be detrended.

At approximately 4-6s period the noise spectrum inflects and begins increasing at a mirrored power law exponential to the lower frequencies. In TDCP at higher rates (>1Hz), [B. Crowell and Ghent, 2022] observes in multiple sample rates from a single receiver that TDCP velocities have increased noise in the time domain, roughly a factor of 7 of standard deviation from 1 Hz to 10 Hz velocities. In the frequency domain these velocities present as a reverse power law of increasing noise as frequency increases, flattening at a corner around 0.2s period (5Hz). We observe a similar spectral shape in our PPSDs. Furthermore, [Shu

et al., 2018] processed up to 50Hz and identified a spectral "knee" around 3.5Hz; the highest frequencies observed in our study terminated at this "knee". We infer this highest frequency (>1Hz) correlated noise to be predominantly influenced by receiver thermal noise, and likely receiver baseband design dependent [Moschas and Stiros, 2013]. [B. Crowell and Ghent, 2022] also identifies the lowest noise power in the frequency domain exists in the 1-10s periods of the highest sample rate observations (20Hz in their study), notable given this intersects the spectral region of the seismic ground motion waveforms of interest. Given the spectrum at higher sampling rates, there is likely potential for improved screening of TDCP velocities for our signals of interest to reduce temporal aliasing [Hohensinn et al., 2020, B. Crowell and Ghent, 2022].

A future PPSD product from continuous single station measurements would enable quantitative comparisons of the ambient noise levels from one station to another. These noise levels, presented in a domain familiar scheme, are a meaningful proxy for the relative sensitivity to observe ground motions. Routine outliers can be observed and correlated to disturbances or events, a potentially valuable tool for network monitoring. In our study, without continuous 5Hz observations, it is not possible to assess time or spatially related variability outside the semi-arbitrary windows we have available, but note it offers a promising framework for inclusion in network monitoring and relative TDCP performance analysis.

4.4.2 Pseudo Synthetic Model Performance

We evaluated four different feature selection strategies by deploying four independent scenarios of random forest model fitting on the same training and testing splits. Our feature sets were time, frequency, wavelet and a combined set. First we tested each feature set independently, and then we also ran a test where we concatenated all available features into a single feature vector. Each run selected its own hyperparameters through k-fold grid search optimization. Overall, we found the highest performance from the largest feature vector of all available features, though we note that an operational pipeline would factor computational

Figure 4.5: GNSS velocity PPSDs. Panel (a) is the combined horizontal components, panel (b) is the vertical component. The black lines represent the 7 distribution quantile slices used for stochastic noise generation.

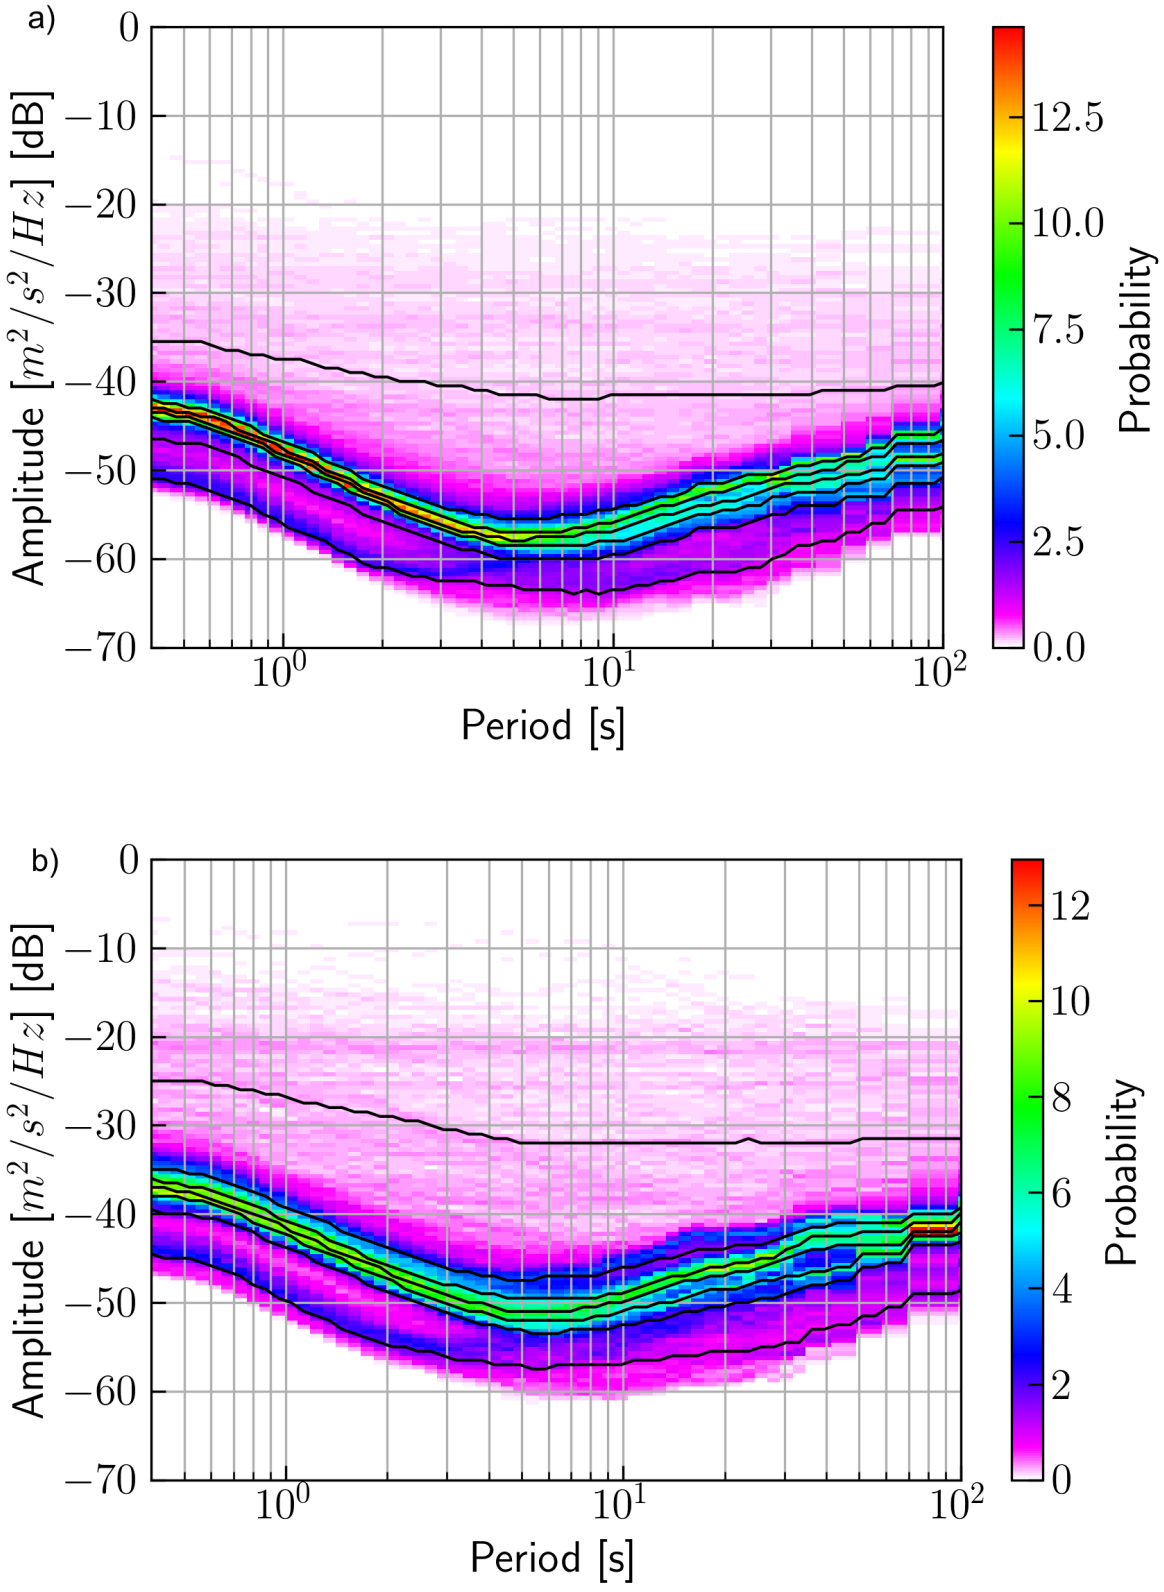

costs as well. We found the PSD-only performance similar to the "all" combined feature vectors, which aligns with our feature importances from [Dittmann et al., 2022b].

Given the additional information we have about SNR in our synthetic catalog, we evaluated relative performance of the feature selections with respect to various regions of SNR. All 4 feature engineering strategies performed well in the highest SNR samples; the precision and recall differences largely occured in our middle range of SNR, from approximately 0 to 80 dB SNR differences. In this region, the time and wavelet features have relatively higher precision but lower recall. We optimize on F1 in order to balance these two metrics, and the net effect is lower F1 score for the time method. The performance difference between the PSD and the wavelet is an interesting result; further investigation into the additional potential of the wavelet transform is warranted, including de-noising ML strategies for true signal extraction or peak dynamics [Yin et al., 2022].

Figure 4.6 suggests that a successful classifier will primarily use the frequency domain features, given the relative performance of the individual feature performances. A benefit of our random forest model is readily extracted feature importance information (Figure 4.7). When our random forest model was presented with the time, frequency and wavelet transform information, the model learned distributed feature importances across spatial components and features. The horizontal components contributed more than the vertical, consistent with previous findings aligned with increased vertical GNSS noise relative to signals. Lower frequency power spectra (¿2s period) contributed the most within the unshaded frequency domain of figure 4.7, but certain features of the time domain (pink shaded in Figure 4.7) and wavelet features (gray shading) contributed more. From a model explainability perspective, we interpret that this reflects the strength of the randomized decision tree algorithm to distribute its decisions across all features with encoded information to optimize performance. From a domain interpretability perspective, frequency domain features contribute valuable information for an optimized classifier the majority of the time, but to increase marginal performance, likely on the rarer edge cases of both signal and noise, unique information from

Figure 4.6: Testing feature extraction strategies across the NGAW2 synthetic dataset. Panel (a) shows precision, recall and F1 scores as a function of feature extraction strategies across the entire catalog in 10 fold nested cross validation. "PSD" are the frequency domain features; "time" are the time domain features, and "wavelet" are features extracted from wavelet transforms. See section for further details of the features themselves. "All" is a concatenated array of all 3 feature strategies. The second row of panels are absolute precision (b-1), recall (b-2) and F1 (b-3) as a function of known peak SNR. This SNR metric is the peak difference between signal and noise power in the frequency domain for a given sample across all components. The third row presents the same measurements from the 2nd row, only now presented relative to the "All" scores, to emphasize the relative differences across detection regions. The gray barplot on the right is a distribution of the samples at each SNR bin.

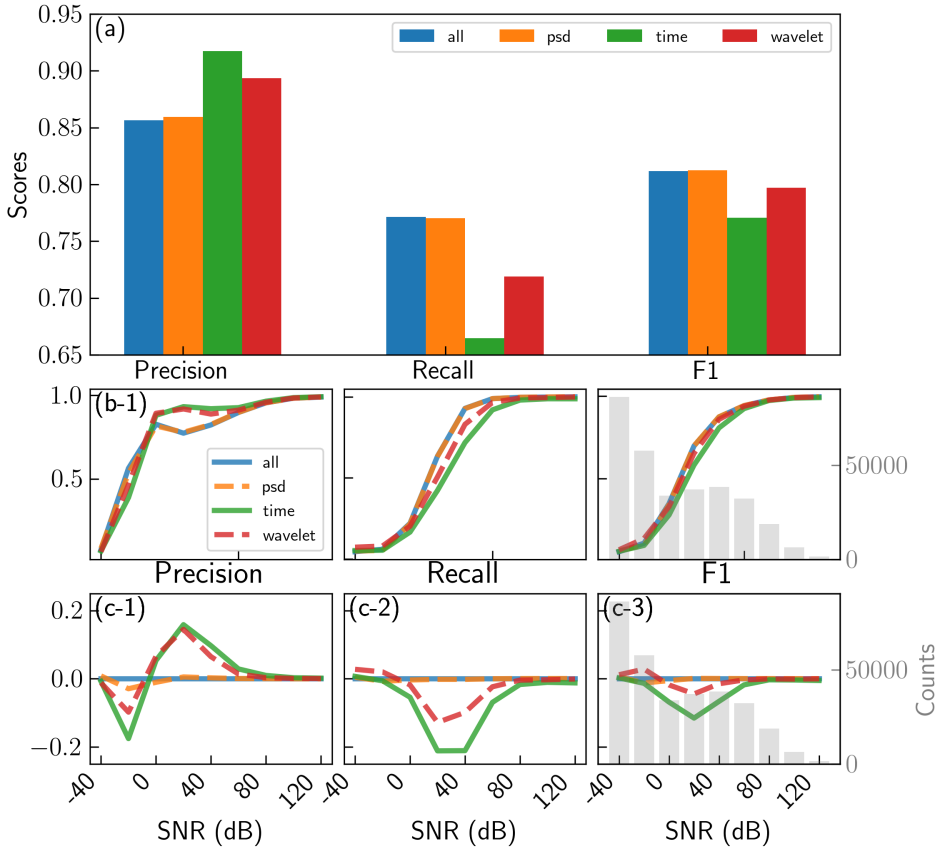

Figure 4.7: Feature importances for our random forest classification model cross validated and trained on the entire NGA West 2 synthetic GNSS dataset and the ambient noise dataset. Panel (a) shows the concatenated importances for all features across all components when a model is trained on all the features at once; the pink shading represents time domain features, the unshaded section are the frequency domain features and the grey shading are the wavelet transform features. The second panel (b) is a close up of the North component features, with the same background shading schema.

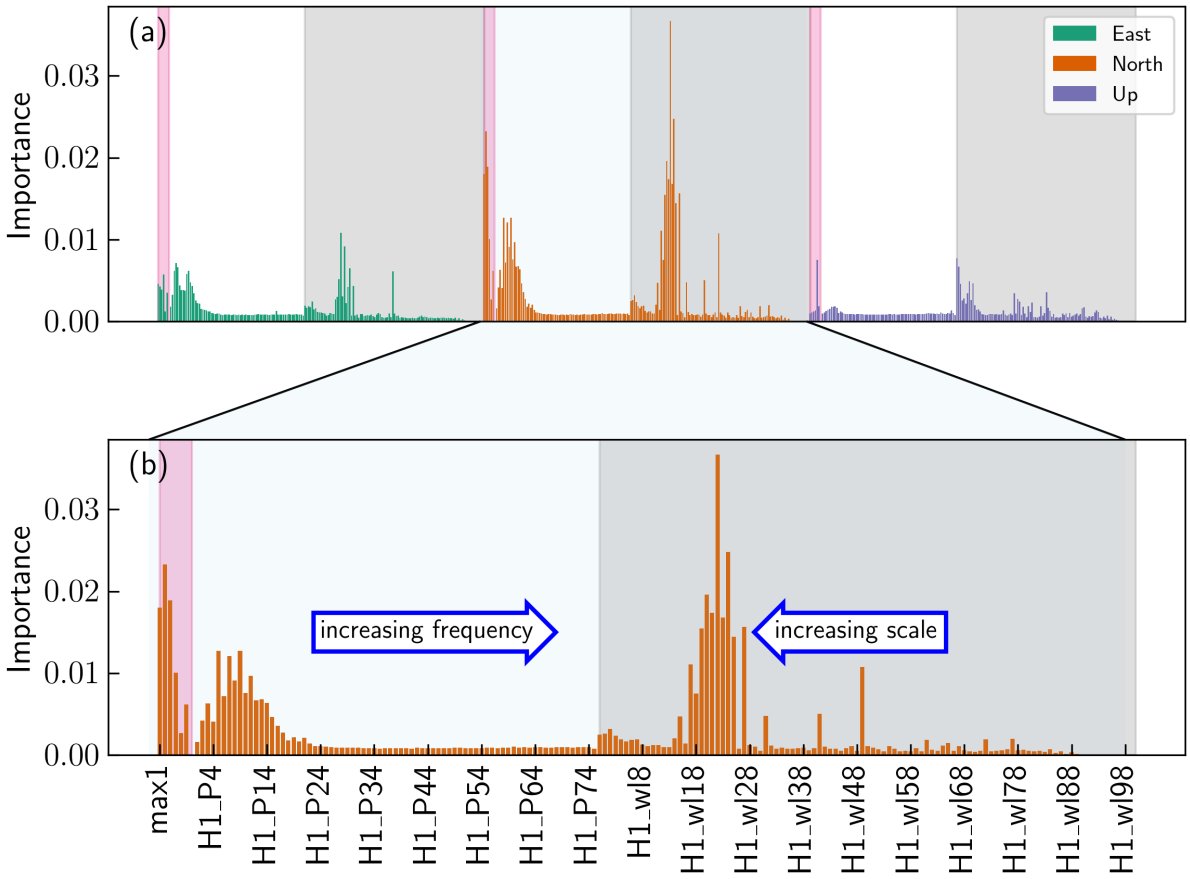

other features, such as responding to non-stationary signals, are required. An unexpected result is that the lower scale (∼higher frequency) vertical wavelet transform features played a significant role. This result is worthy of future investigation.

4.4.3 Quantifying Augmentation

Figure 4.2 and table 4.1 make evident that transferred signals with data augmentation significantly expanded the GNSS velocity catalog with respect to the number of unique waveforms. Additionally, data augmentation is an opportunity to expand sample feature space by leveraging our knowledge of the signals relative to the noise to train high quality labels with elevated noise environments [Zhu et al., 2020]. We quantified the performance impact of augmentation by comparing models trained with and without augmentation. We ran identical complete nested cross validation testing scenarios using two different training tactics. In the first we allowed the model to train on all 7 replicas of each waveform; in the second we only provided the lowest noise waveform in training. Panels (a-b) of figure 4.8 are from the first training scenario with augmentation. We tested on all replicas of the testing set waveforms, but for visualization purposes the left panel (a) is the performance of the 50th percentile median noise waveforms, and the right panel (b) is the performance of the 95th percentile high noise waveforms. As previously mentioned, SNR values derived from the known noise time series and known signal periodograms. With data augmentation, we observed decreasing SNR for the same catalog while testing against increased noise levels, from 92% true positive rate to 84%. When we compared the 50% noise levels with and without data augmentation (panels a,c), we notice a similar drop in performance without augmentation. Finally, when we looked at the highest noise samples without augmentation, we see a dramatic decrease in performance despite testing on the identical waveforms with the same SNR.

Figure 4.8: Comparing event detection with training on augmented noise samples across noise levels. Each panel includes the peak SNR of the waveform for each event as a function of radius from the event. This SNR metric is the peak of signal power to noise power for any frequency bin of periodograms calculated for all samples for all components for a given station-event waveform. The plot marker radius is determined by the event magnitude. The top panels (a-b) are testing the 50th noise model and 95th noise model of each station-event waveform using a classifier trained on all augmented samples. The markers are colored by a binary detected/not detected. The bottom panels are testing the 50th/95th noise model waveforms with no data augmentation. This illustrates the value of augmentation for detection in noise, in addition to the approximate threshold of detection given our knowledge of signal and noise in this pseudo synthetic dataset.

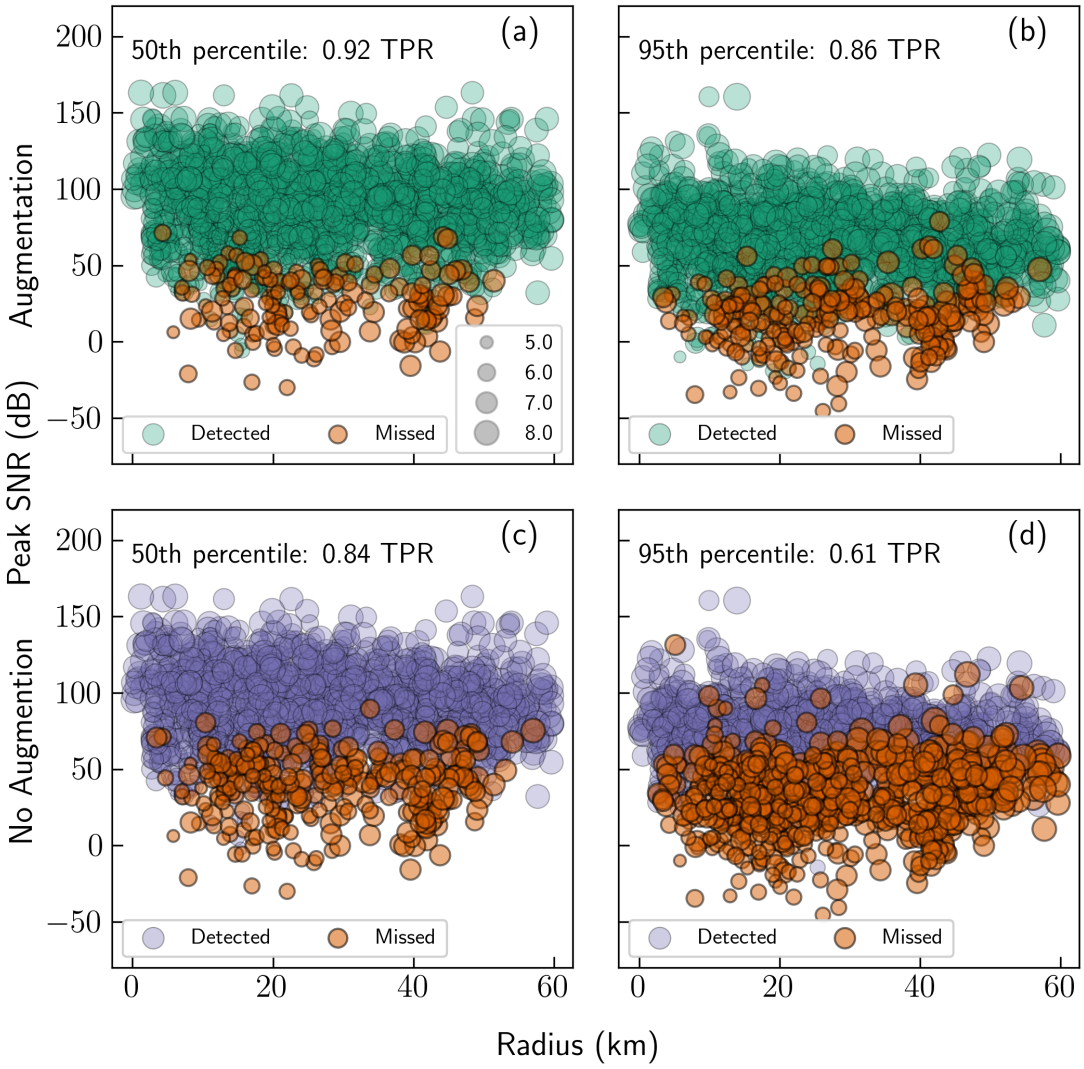

4.4.4 Validation with Real GNSS TDCP Velocities

Finally, to validate our synthesis of GNSS velocity waveforms against real-world GNSS velocities, we reran a nested cross validation experiment with the entire real-world GNSS velocity catalog of [Dittmann et al., 2022b] as a reference to compare the synthetically generated model. Similar in testing design to the previous comparison of data augmentation, we evaluated the performance of two classification models against the same semi-random testing subsets in the nested cross validation loops and report on the mean performance. In this scenario, one model was fit on the remaining data using hyperparameters extracted from k-fold cross validation for each training set, the other model was fit on the entire synthetic GNSS velocities catalog. All other feature engineering strategies were held consistent and both models were evaluated against the same testing sets. The synthetic GNSS trained model yielded better performance metrics, including increased precision, recall and F1. This performance can best be explained by the extent of training sets: the synthetic model was trained on 1846 waveforms, where the "real" model was trained on ~200, depending on the nested cross validation run testing slice. The added extent and density of information in the transferred and augmented training data improved generalization for this smaller data subset.

We ran an additional test where we take the best fit model from each dataset and applied it to a yet unseen ambient noise dataset (for dataset description, see table 4.1). We found the GNSS velocity trained model had a nearly identical false positive rate, where false positive rate is one minus the true negative rate. This further validates that our noise training and augmentation strategy was effective in improving performance in difficult noise conditions, as our performance improvement in the event catalog did not come at the expense of ambient performance.

From these improved classification results we infer that transferred, augmented "synthetic" waveforms are a valid substitute for high-rate GNSS measurements to partially over-

Figure 4.9: Testing performance of real GNSS velocity events as a function of training catalog used. These are the mean scores across the 10 testing folds of the nested cross validation. The purple results are from a model generated using cross validation of the remaining real gnss dataset; the green results are from the bulk model fit to the entire NGAW2 synthetic dataset. Each uses the "all" feature extraction method (time, frequency and wavelet features). The ambient true negative rate (TNR) is estimated using a separate dataset of unseen ambient data. TNR is [true negatives / (true negatives + false positives)], or equivalent to one minus the false positive rate. The annotated text is the difference between the two approaches.

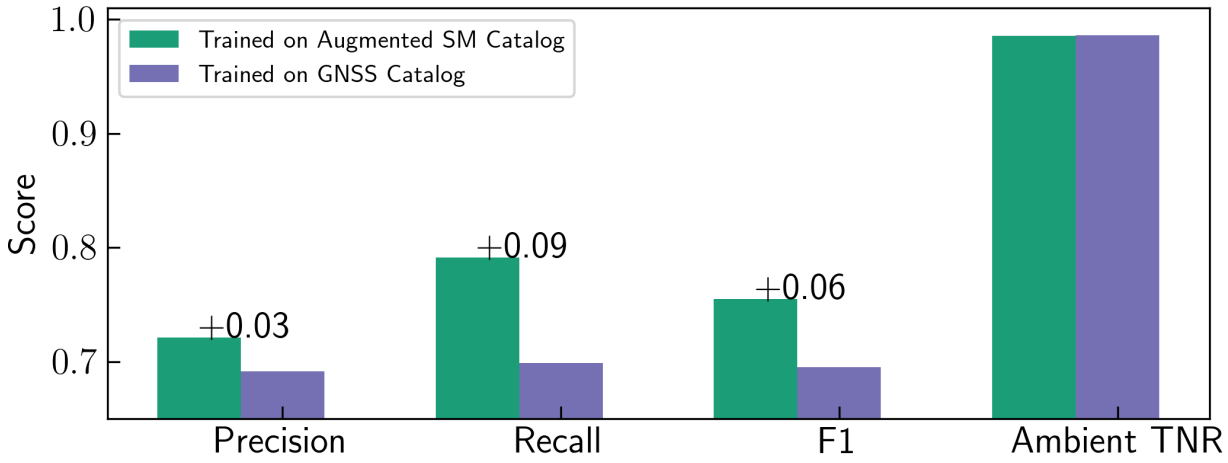

come modern, smaller GNSS seismic datasets. A future deployed classifier will be trained on the combination of data catalogs to achieve the best generalization performance for yet-to-occur events. This real-world versus pseudo-synthetic comparison and validation result also suggests that evolved transfer learning across measurement domains, including exploration of fine-tuning of more mature seismic deep learning models with GNSS velocities, could further advance GNSS seismology challenges.

4.5 Conclusions

We find the ambient GNSS velocity noise over periods from 0.4s to 100s has a v-shape power spectrum centered around 6 seconds. Lower frequency power law noise is attributed to error sources such as multipath and orbits, and higher frequency (>1Hz) noise is attributed to receiver thermal noise present in carrier phase observables. These findings are consistent with previous high-rate GNSS positioning noise analysis and useful for future network monitoring and realistic synthetic noise generation. We find that frequency, time and wavelet feature extraction strategies vary slightly under different SNR regions, with the frequency domain having the highest overall F1 scores of 0.81 over the entire catalog when run in a 10-fold nested cross validation. Data augmentation boosts overall performance from a 0.83 true postive rate to 0.92, and the majority of this improvement is from enhanced performance under higher noise settings.

Finally, we found that a model trained on these pseudo-synthetic waveforms, with the full suite of augmentation, outperformed the model trained on strictly GNSS velocity waveforms, including an increase in overall F1 score from 0.7 to 0.81 with equivalent ambient false alarm rate. The immediate benefit is an improved classification model from an expanded catalog that can be retrained on the combined pseudo-synthetic and real catalog for unseen events. The subsequent benefit is the expanded catalog, and framework for continued expansion, presents and opportunity for deeper learning models that are “data hungry” [Mousavi and Beroza, 2022]. This includes expanding functional learning outputs, such as denoising,

regression for magnitude inversion and forecasting. A "loose coupling" of inertial waveforms and these expanded geodetic seismic waveforms enables fine-tuning [Yosinski et al., 2014] or transfer of existing seismic detection models, such as [Mousavi et al., 2020, Seydoux et al., 2020]. "Tighter coupling" will include geodetic sources directly in catalogs and operational systems. This further blurs distinctions between inertial and geodetic sensors, moving from representations of different fields of earth sciences towards observational inputs with respective responses and noise models for optimizing immediate safety-of-life hazards monitoring contributions.

4.6 Open Research

The inertial seismic records are available from the Pacific Earthquake Engineering Research Center (PEER) Next Generation Attenuation for Western United States 2.0 [Ancheta et al., 2014]. `https://ngawest2.berkeley.edu/` The 5Hz GNSS data used for TDCP processing in the study are available from the Geodetic Facility for the Advancement of Geoscience (GAGE) Global Navigation Satellite Systems (GNSS) archives as maintained by UNAVCO, Inc. The data are available in RINEX (v.2.11) format at `https://data.unavco.org/archive/gnss/highrate/5-Hz/rinex/`. Earthquake depths, locations, and magnitudes came from the Advanced National Seismic System (ANSS) Comprehensive Catalog of Earthquake Events and Products (`https://earthquake.usgs.gov/data/comcat/`). SNIVEL code used for TDCP velocity processing is developed openly at `https://github.com/crowellbw/SNIVEL` (Accessed December 2021)[Crowell, 2021]. SNIVEL 5Hz velocity timeseries used in this study are preserved at `https://doi.org/10.5281/zenodo.6588601`. Version 1.0.1 of the scikit-learn software used for random forest classification is preserved at `https://doi.org/10.5281/zenodo.5596244` and developed openly at `https://github.com/scikit-learn/scikit-learn` [Pedregosa et al., 2011]. Version v0.5.0 of PyGMT used for generating the map is preserved at `https://doi.org/10.5281/zenodo.5607255` and developed openly at `https://github.com/GenericMappingTools/pygmt`[Wessel et al., 2019].

Figures were made with Matplotlib version 3.5.1 [Caswell et al., 2021], available under the Matplotlib license at `https://matplotlib.org/`. Version 1.2.1 of the obspy software used for seismic data handling and PPSD generation is preserved at `https://doi.org/10.5281/zenodo.3706479` and developed openly at `https://github.com/obspy/obspy` [Krischer et al., 2015]. Version 1.3.0 of pywavelets software used for estimating wavelet coefficients is preserved at `https://doi.org/10.5281/zenodo.6347505` and developed openly at `https://github.com/PyWavelets/pywt` [Lee et al., 2019].

This material is based on services provided by the GAGE Facility, operated by UNAVCO, Inc., with support from the National Science Foundation, the National Aeronautics and Space Administration, and the U.S. Geological Survey under NSF Cooperative Agreement EAR-1724794. High-rate processing and and machine learning for geoscience and hazards research is supported by NSF OAC-1835791.

Chapter 5

Conclusion

5.1 Research Summary

This research addresses a gap between GNSS detection sensitivity to seismic ground motions and the current functional use of these valuable measurements. We build on the GNSS seismology research over the last 20 years to offer actionable components of improved operations. We compared the detection sensitivity to peak seismic ground motions of two different GNSS processing strategies, PPP and TDCP. We model both processing strategies' ambient peak ground motion noise as a chi-squared distribution, and then use existing empirical scaling laws to estimate probability of detection of an event given magnitude, radius and false alarm tolerance. We conclude that TDCP derived PGV magnitude estimation is a computationally lightweight, independently processed complement to PPP-derived PGD magnitude estimates, which could be produced at the network edge at high rates and with increased sensitivity to ground motion than current PPP estimates

A supervised random forest classifier outperforms existing detection methods in standalone mode to enable these valuable ground motion measurements inclusion in development of edge processing, geodetic infrastructure monitoring and inclusion in operational ground motion observations and models. Our training, validation and testing pipeline used our assembled 5Hz TDCP visually labeled waveform dataset from an event catalog of 80 different earthquakes (M4.9-8.2) across a hemispheric scale over nearly 20 years.

We addressed the limited data challenge of GNSS seismic waveforms by expanding the

available GNSS seismic waveform catalog. A classifier trained on transferred strong motion signals augmented with authentic TDCP noise has improved F1 on real GNSS velocity waveforms than one trained on exclusively GNSS waveforms. This expanded catalog provides a framework for continued expansion to enable deeper learning.

5.2 Considerations Moving Forward

This study provides a quantitative analysis of GNSS measurement noise and sensitivity. This lays the foundation for future integration with inertial sensors to allow the complementary natures of the two types of sensors to be best utilized.

The classifier developed here is readily deployable for next generation geodetic network operations, including post-event research and real-time operations, dependent on a real-time high rate GNSS velocity client. This classifier and client could be deployed for centralized processing or distributed edge processing. GNSS velocity edge processing with a detection module will support higher rate intelligent signal capture in limited bandwidth environments as well as support triggering even higher rate data buffers for future research. We will address the noise frequency dependence on sample rate, as well as design and hopefully execute dynamic tests to investigate potential instrument responses.

A combined synthetic/real catalog can be further expanded, both with additional augmented inertial waveforms as well as additional GNSS velocity waveforms. From a ground motion observation perspective, an expanded catalog can further refine a similar classifier to enable improved inclusion in ground motion historical catalogs as well as operational hazard systems, such as ShakeMaps and EEW algorithms. In addition, expanded catalogs will support more sophisticated learning outputs, such as denoising and spatio-temporal forecasting that could be transferred or fine-tuned from existing seismic models. Finally, further design and analysis of a combined local/semi-local tight integration of all available sensors offers a solution that maximizes each sensor's contributions.

Bibliography

[NGD,] National Geophysical Data Center / World Data Service (NGDC/WDS): NCEI/WDS Global Significant Earthquake Database.

[MTU, 2022] (2022). Seismology.

[Aki et al., 2002] Aki, Richards, Keiiti and, and G., P. (2002). Quantitative Seismology. University Science Books, 2 edition.

[Allen and Melgar, 2019] Allen, R. M. and Melgar, D. (2019). Earthquake Early Warning: Advances, Scientific Challenges, and Societal Needs. Annual Review of Earth and Planetary Sciences, 47(1):1–28.

[Allen and Ziv, 2011] Allen, R. M. and Ziv, A. (2011). Application of real-time gps to earthquake early warning. Geophys. Res. Lett., 38(L16310).

[Ancheta et al., 2014] Ancheta, T. D., Darragh, R. B., Stewart, J. P., Seyhan, E., Silva, W. J., Chiou, B. S.-J., Wooddell, K. E., Graves, R. W., Kottke, A. R., Boore, D. M., Kishida, T., and Donahue, J. L. (2014). NGA-West2 Database. Earthquake Spectra, 30(3):989–1005.

[Avallone et al., 2011] Avallone, A., Marzario, M., Cirella, A., Piatanesi, A., Rovelli, A., Alessandro, C. D., D'Anastasio, E., D'Agostino, N., Giuliani, R., and Mattone, M. (2011). Very high rate (10 Hz) GPS seismology for moderate-magnitude earthquakes: The case of the Mw 6.3 L'Aquila (central Italy) event. Journal of Geophysical Research: Solid Earth, 116(B2).

[B. Crowell and Ghent, 2022] B. Crowell, J. DeGrande, T. D. and Ghent, J. (2022). Validation of Peak Ground Velocities Recorded on Very-high rate GNSS Against NGA-West2 Ground Motion Models. Seismica (under review).

[Bakshi, 1998] Bakshi, B. R. (1998). Multiscale PCA with application to multivariate statistical process monitoring. AIChE Journal, 44(7):1596–1610. https://doi.org/10.1002/aic.690440712.

[Benedetti et al., 2014] Benedetti, E., Branzanti, M., Biagi, L., Colosimo, G., Mazzoni, A., and Crespi, M. (2014). Global Navigation Satellite Systems Seismology for the 2012 Mw 6.1 Emilia Earthquake: Exploiting the VADASE Algorithm. Seismological Research Letters, 85(3):649–656.

[Bergen et al., 2019] Bergen, K. J., Johnson, P. A., Hoop, M. V. d., and Beroza, G. C. (2019). Machine learning for data-driven discovery in solid Earth geoscience. Science, 363(6433).

[Bertiger et al., 2010] Bertiger, W., Desai, S. D., Haines, B., Harvey, N., Moore, A. W., Owen, S., and Weiss, J. P. (2010). Single receiver phase ambiguity resolution with gps data. Journal of Geodesy, 84(5):327–337.

[Bishop, 1995] Bishop, C. M. (1995). Training with Noise is Equivalent to Tikhonov Regularization. Neural Computation, 7(1):108–116.

[Bishop and Nasrabadi, 2007] Bishop, C. M. and Nasrabadi, N. M. (2007). Pattern Recognition and Machine Learning. J. Electronic Imaging, 16:049901.

[Bock and Melgar, 2016] Bock, Y. and Melgar, D. (2016). Physical applications of GPS geodesy: a review. Reports on Progress in Physics, 79(10):106801.

[Bock et al., 2004] Bock, Y., Prawirodirdjo, L., and Melbourne, T. I. (2004). Detection of arbitrarily large dynamic ground motions with a dense high-rate GPS network. Geophysical Research Letters, 31(6):n/a–n/a.

[Boore, 1983] Boore, D. M. (1983). Stochastic simulation of high-frequency ground motions based on seismological models of the radiated spectra. Bulletin of the Seismological Society of America.

[Branzanti et al., 2013] Branzanti, M., Colosimo, G., Crespi, M., and Mazzoni, A. (2013). Gps near-real-time coseismic displacements for the great tohoku-oki earthquake. IEEE Geoscience and Remote Sensing Letters, 10(2):372–376.

[Breiman, 2001] Breiman, L. (2001). Random Forests. Machine Learning, 45(1):5–32.

[Böse et al., 2017] Böse, M., Smith, D. E., Felizardo, C., Meier, M.-A., Heaton, T. H., and Clinton, J. F. (2017). FinDer v.2: Improved real-time ground-motion predictions for M2–M9 with seismic finite-source characterization. Geophysical Journal International, 212(1):725–742.

[Casey et al., 2018] Casey, R., Templeton, M. E., Sharer, G., Keyson, L., Weertman, B. R., and Ahern, T. (2018). Assuring the Quality of IRIS Data with MUSTANG. Seismological Research Letters, 89(2A):630–639.

[Caswell et al., 2021] Caswell, T. A., Droettboom, M., Lee, A., de Andrade, E. S., Hoffmann, T., Hunter, J., Klymak, J., Firing, E., Stansby, D., Varoquaux, N., Nielsen, J. H., Root, B., May, R., Elson, P., Seppänen, J. K., Dale, D., Lee, J.-J., McDougall, D., Straw, A.,

Hobson, P., hannah, Gohlke, C., Vincent, A. F., Yu, T. S., Ma, E., Silvester, S., Moad, C., Kniazev, N., Ernest, E., and Ivanov, P. (2021). matplotlib/matplotlib: Rel: v3.5.1.

[Chung et al., 2019] Chung, A. I., Henson, I., and Allen, R. M. (2019). Optimizing Earthquake Early Warning Performance: ElarmS-3. Seismological Research Letters, 90(2A):727–743.

[Colombelli et al., 2013] Colombelli, S., Allen, R. M., and Zollo, A. (2013). Application of real-time GPS to earthquake early warning in subduction and strike-slip environments. Journal of Geophysical Research: Solid Earth, 118(7):3448–3461.

[Colosimo and Mazzoni, 2011] Colosimo, G., M. C. and Mazzoni, A. (2011). Real-time gps seismology with a stand-alone receiver: A preliminary feasibility demonstration. Journal of Geophysical Research, 116(B11302).

[Crocetti et al., 2021] Crocetti, L., Schartner, M., and Soja, B. (2021). Discontinuity Detection in GNSS Station Coordinate Time Series Using Machine Learning. Remote Sensing, 13(19):3906.

[Crowell, 2021] Crowell, B. W. (2021). Near-field strong ground motions from gps-derived velocities for 2020 intermountain western united states earthquakes. Seismol. Res. Lett, XX(1-9).

[Crowell et al., 2009] Crowell, B. W., Bock, Y., and Squibb, M. B. (2009). Demonstration of Earthquake Early Warning Using Total Displacement Waveforms from Real-time GPS Networks. Seismological Research Letters, 80(5):772–782.

[Crowell et al., 2016] Crowell, B. W., Schmidt, D. A., Bodin, P., Vidale, J. E., Gomberg, J., Hartog, J. R., Kress, V. C., Melbourne, T. I., Santillan, M., Minson, S. E., and Jamison, D. G. (2016). Demonstration of the Cascadia G-FAST Geodetic Earthquake Early Warning System for the Nisqually, Washington, Earthquake. Seismological Research Letters, 87(4):930–943.

[Crowell and Geng, 2013] Crowell, B. W., D. M. Y. B. J. S. H. and Geng, J. (2013). Earthquake magnitude scaling using seismogeodetic data. Geophys. Res. Lett, 40(6089–6094).

[de Lacy et al., 2012] de Lacy, M. C., Reguzzoni, M., and Sansò, F. (2012). Real-time cycle slip detection in triple-frequency gnss. GPS Solutions, 16(3):353–362.

[Dittmann et al., 2022a] Dittmann, T., Hodgkinson, K., Morton, J., Mencin, D., and Mattioli, G. S. (2022a). Comparing Sensitivities of Geodetic Processing Methods for Rapid Earthquake Magnitude Estimation. Seismological Research Letters, 93(3):1497–1509.

[Dittmann et al., 2022b] Dittmann, T., Liu, Y., Morton, Y., and Mencin, D. (2022b). Supervised Machine Learning of High Rate GNSS Velocities for Earthquake Strong Motion Signals. Journal of Geophysical Research: Solid Earth, 127(11).

[Douglas D. Given, 2018] Douglas D. Given, Richard M. Allen, A. S. B. P. B. E. S. C. K. C. R. M. d. G. L. S. G. E. H. T. H. H. M. H. J. R. M. V. T. D. T. T. S. Y. (2018). Revised technical implementation plan for the ShakeAlert system—An earthquake early warning system for the West Coast of the United States. Report 2018-1155, Reston, VA.

[Dybing et al., 2021] Dybing, S., Melgar, D., Thomas, A., Hodgkinson, K., and Mencin, D. (2021). Detecting earthquakes in noisy real-time gnss data with machine learning. American Geophysical Union Fall Meeting 2021.

[Ebinuma and Kato, 2012] Ebinuma, T. and Kato, T. (2012). Dynamic characteristics of very-high-rate GPS observations for seismology. Earth, Planets and Space, 64(5):369–377.

[EUSPA, 2022] EUSPA (2022). EO & GNSS Market Report.

[Fang and Liu, 2020] Fang, R., J. Z. J. G. Y. S. C. S. and Liu, J. (2020). Earthquake magnitude scaling using peak ground velocity derived from high-rate gnss observations. Seismol. Res. Lett, XX(1-11).

[(FEMA) et al., 2017] (FEMA), F. E. M. A., (USGS), U. S. G. S., and (PDC), and the Pacific Disaster Center (2017). Hazus® Estimated Annualized Earthquake Losses for the United States. Technical report.

[Ferrari, 2019] Ferrari, A. (2019). A note on sum and difference of correlated chi-squared variables.

[Francesca Fratarcangeli, 2018] Francesca Fratarcangeli, Michela Ravanelli, A. M. G. C. E. B. M. B. G. S. O. V. A. K. M. C. (2018). The variometric approach to real-time high-frequency geodesy. Rendiconti Lincei. Scienze Fisiche e Naturali, 29.

[Gabor, 1946] Gabor, D. (1946). Theory of Communication. Institution of Electrical Engineering.

[Geng et al., 2018] Geng, J., Pan, Y., Li, X., Guo, J., Liu, J., Chen, X., and Zhang, Y. (2018). Noise Characteristics of High-Rate Multi-GNSS for Subdaily Crustal Deformation Monitoring. Journal of Geophysical Research: Solid Earth, 123(2):1987–2002.

[Genrich and Bock, 2006] Genrich, J. F. and Bock, Y. (2006). Instantaneous geodetic positioning with 10–50 hz gps measurements: Noise characteristics and implications for monitoring networks. Journal of Geophysical Research: Solid Earth, 111(B3).

[Goldberg and Bock, 2017] Goldberg, D. E. and Bock, Y. (2017). Self-contained local broadband seismogeodetic early warning system: Detection and location. Journal of Geophysical Research: Solid Earth, 122(4):3197–3220.

[Goldberg et al., 2022] Goldberg, D. E., Koch, P., Melgar, D., Riquelme, S., and Yeck, W. L. (2022). Beyond the Teleseism: Introducing Regional Seismic and Geodetic Data into Routine USGS Finite-Fault Modeling. Seismological Research Letters, 93(6):3308–3323.

[Goldberg et al., 2019] Goldberg, D. E., Melgar, D., and Bock, Y. (2019). Seismogeodetic P-wave Amplitude: No Evidence for Strong Determinism. Geophysical Research Letters, 46(20):11118–11126.

[Goldberg et al., 2018] Goldberg, D. E., Melgar, D., Bock, Y., and Allen, R. M. (2018). Geodetic observations of weak determinism in rupture evolution of large earthquakes. Journal of Geophysical Research: Solid Earth, 123(11):9950–9962.

[Goldberg et al., 2021] Goldberg, D. E., Melgar, D., Hayes, G. P., Crowell, B. W., and Sahakian, V. J. (2021). A Ground-Motion Model for GNSS Peak Ground Displacement. Bulletin of the Seismological Society of America, 111(5):2393–2407.

[Graas and Soloviev, 2003] Graas, F. V. and Soloviev, A. (2003). Precise velocity estimation using a stand-alone gps receiver. In Proceedings of the ION NTM 2003, page 352–360. Institute of Navigation.

[Grapenthin and Freymueller, 2018] Grapenthin, R., M. W. C. T. M. G. and Freymueller, J. (2018). Single-frequency instantaneous gnss velocities resolve dynamic ground motion of the 2016 mw 7.1 iniskin, alaska, earthquake. Seismol. Res. Lett, 89(1040-1048).

[Grapenthin et al., 2017] Grapenthin, R., West, M., and Freymueller, J. (2017). The Utility of GNSS for Earthquake Early Warning in Regions with Sparse Seismic Networks. Bulletin of the Seismological Society of America, 107(4):1883–1890.

[Graves and Pitarka, 2010] Graves, R. W. and Pitarka, A. (2010). Broadband Ground-Motion Simulation Using a Hybrid ApproachBroadband Ground-Motion Simulation Using a Hybrid Approach. Bulletin of the Seismological Society of America, 100(5A):2095–2123.

[Habboub et al., 2020] Habboub, M., Psimoulis, P. A., Bingley, R., and Rothacher, M. (2020). A Multiple Algorithm Approach to the Analysis of GNSS Coordinate Time Series for Detecting Geohazards and Anomalies. Journal of Geophysical Research: Solid Earth, 125(2).

[Hayes, 2017] Hayes, G. P. (2017). The finite, kinematic rupture properties of great-sized earthquakes since 1990. Earth and Planetary Science Letters, 468:94–100.

[Hodgkinson et al., 2020] Hodgkinson, K. M., Mencin, D. J., Feaux, K., Sievers, C., and Mattioli, G. S. (2020). Evaluation of Earthquake Magnitude Estimation and Event Detection Thresholds for Real-Time GNSS Networks: Examples from Recent Events Captured by the Network of the Americas. Seismological Research Letters, 91(3):1628–1645.

[Hoffmann et al., 2019] Hoffmann, J., Bar-Sinai, Y., Lee, L. M., Andrejevic, J., Mishra, S., Rubinstein, S. M., and Rycroft, C. H. (2019). Machine learning in a data-limited regime: Augmenting experiments with synthetic data uncovers order in crumpled sheets. Science Advances, 5(4):eaau6792.

[Hohensinn and Geiger, 2018] Hohensinn, R. and Geiger, A. (2018). Stand-alone gnss sensors as velocity seismometers: Real-time monitoring and earthquake detection. Sensors, 18(11)(3712).

[Hohensinn et al., 2020] Hohensinn, R., Häberling, S., and Geiger, A. (2020). Dynamic displacements from high-rate GNSS: Error modeling and vibration detection. Measurement, 157:107655.

[Hoshiba and Iwakiri, 2011] Hoshiba, M. and Iwakiri, K. (2011). Initial 30 seconds of the 2011 off the Pacific coast of Tohoku Earthquake (Mw 9.0)—amplitude and c for magnitude estimation for Earthquake Early Warning—. Earth, Planets and Space, 63(7):8.

[Häberling et al., 2015] Häberling, S., Rothacher, M., Zhang, Y., Clinton, J. F., and Geiger, A. (2015). Assessment of high-rate GPS using a single-axis shake table. Journal of Geodesy, 89(7):697–709.

[IRIS, 2021] IRIS (April 2021). A brief history of the very first broadband seismic station.

[Iwana and Uchida, 2021] Iwana, B. K. and Uchida, S. (2021). An empirical survey of data augmentation for time series classification with neural networks. PLoS ONE, 16(7):e0254841.

[Jiao et al., 2017] Jiao, Y., Hall, J. J., and Morton, Y. T. (2017). Automatic equatorial gps amplitude scintillation detection using a machine learning algorithm. IEEE Transactions on Aerospace and Electronic Systems, 53(1):405–418.

[Jordan and Mitchell, 2015] Jordan, M. I. and Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. Science, 349(6245):255–260. doi: 10.1126/science.aaa8415.

[Joyner, 1984] Joyner, W. (August 1984). A SCALING LAW FOR THE SPECTRA OF LARGE EARTHQUAKES. Bulletin of the Seismological Society of America, 74(4):1167–1188.

[Kaloop and Kim, 2016] Kaloop, M. R. and Kim, D. (2016). De-noising of GPS structural monitoring observation error using wavelet analysis. Geomatics, Natural Hazards and Risk, 7(2):804–825.

[Kam et al., 2021] Kam, J., Park, J., Shao, W., Song, J., Kim, J., Gizzi, F. T., Porrini, D., and Suh, Y.-J. (2021). Data-driven modeling reveals the Western dominance of global public interest in earthquakes. Humanities and Social Sciences Communications, 8(1):242.

[Kawamoto et al., 2016] Kawamoto, S., Hiyama, Y., Ohta, Y., and Nishimura, T. (2016). First result from the GEONET real-time analysis system (REGARD): the case of the 2016 Kumamoto earthquakes. Earth, Planets and Space, 68(1):190.

[Kennett and Engdahl, 1991] Kennett, B. L. N. and Engdahl, E. R. (1991). Traveltimes for global earthquake location and phase identification. Geophysical Journal International, 105(2):429–465.

[Kodera et al., 2018] Kodera, Y., Yamada, Y., Hirano, K., Tamaribuchi, K., Adachi, S., Hayashimoto, N., Morimoto, M., Nakamura, M., and Hoshiba, M. (2018). The Propagation of Local Undamped Motion (PLUM) Method: A Simple and Robust Seismic Wavefield Estimation Approach for Earthquake Early WarningThe Propagation of Local Undamped Motion (PLUM) Method. Bulletin of the Seismological Society of America, 108(2):983–1003.

[Kong et al., 2019] Kong, Q., Trugman, D. T., Ross, Z. E., Bianco, M. J., Meade, B. J., and Gerstoft, P. (2019). Machine Learning in Seismology: Turning Data into Insights. Seismological Research Letters, 90(1):3–14.

[Krischer et al., 2015] Krischer, L., Megies, T., Barsch, R., Beyreuther, M., Lecocq, T., Caudron, C., and Wassermann, J. (2015). ObsPy: a bridge for seismology into the scientific Python ecosystem. Computational Science & Discovery, 8(1):014003.

[Kristeková et al., 2009] Kristeková, M., Kristek, J., and Moczo, P. (2009). Time-frequency misfit and goodness-of-fit criteria for quantitative comparison of time signals. Geophysical Journal International, 178(2):813–825.

[Kumar and Foufoula-Georgiou, 1997] Kumar, P. and Foufoula-Georgiou, E. (1997). Wavelet analysis for geophysical applications. Reviews of Geophysics, 35(4):385–412.

[Langbein and Bock, 2004] Langbein, J. and Bock, Y. (2004). High-rate real-time GPS network at Parkfield: Utility for detecting fault slip and seismic displacements. Geophysical Research Letters, 31(15).

[Larson, 2009] Larson, K. (2009). Gps seismology. Journal of Geodesy, 83:227–233.

[Larson et al., 2003] Larson, K. M., Bodin, P., and Gomberg, J. (2003). Using 1-hz gps data to measure deformations caused by the denali fault earthquake. Science, 300(5624):1421–1424.

[LeCun et al., 2015] LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. Nature, 521(7553):436–444.

[Lee et al., 2019] Lee, G., Gommers, R., Waselewski, F., Wohlfahrt, K., and O'Leary, A. (2019). PyWavelets: A Python package for wavelet analysis. Journal of Open Source Software, 4(36):1237.

[Li and Melachroinos, 2018] Li, T. and Melachroinos, S. (2018). An enhanced cycle slip repair algorithm for real-time multi-gnss, multi-frequency data processing. GPS Solutions, 23:1–11.

[Lin et al., 2021] Lin, J., Melgar, D., Thomas, A. M., and Searcy, J. (2021). Early Warning for Great Earthquakes From Characterization of Crustal Deformation Patterns With Deep Learning. Journal of Geophysical Research: Solid Earth, 126(10).

[Liu and Morton, 2022] Liu, Y. and Morton, Y. J. (2022). Improved Automatic Detection of GPS Satellite Oscillator Anomaly using a Machine Learning Algorithm. NAVIGATION: Journal of the Institute of Navigation, 69(1):navi.500.

[Mattioli et al., 2020] Mattioli, G. S., Phillips, D. A., Hodgkinson, K. M., Walls, C., Mencin, D. J., Bartel, B. A., Charlevoix, D. J., Crosby, C., Gottlieb, M. J., Henderson, B., Johnson, W., Maggert, D., Mann, D., Meertens, C. M., Normandeau, J., Pettit, J., Puskas, C. M., Rowan, L., Sievers, C., and Zaino, A. (2020). The GAGE Data and Field Response to the 2019 Ridgecrest Earthquake Sequence. Seismological Research Letters, 91(4):2075–2086.

[McNamara and Buland, 2004] McNamara, D. E. and Buland, R. P. (2004). Ambient Noise Levels in the Continental United States. Bulletin of the Seismological Society of America, 94(4):1517–1527.

[Meier et al., 2016] Meier, M., Heaton, T., and Clinton, J. (2016). Evidence for universal earthquake rupture initiation behavior. Geophysical Research Letters, 43(15):7991–7996.

[Meier et al., 2019] Meier, M., Ross, Z. E., Ramachandran, A., Balakrishna, A., Nair, S., Kundzicz, P., Li, Z., Andrews, J., Hauksson, E., and Yue, Y. (2019). Reliable Real-Time Seismic Signal/Noise Discrimination With Machine Learning. Journal of Geophysical Research: Solid Earth, 124(1):788–800.

[Melbourne et al., 2020] Melbourne, T. I., Szeliga, W. M., Marcelo Santillan, V., and Scrivner, C. W. (2020). 25-Second Determination of 2019 Mw 7.1 Ridgecrest Earthquake Coseismic Deformation. Bulletin of the Seismological Society of America, 110(4):1680–1687.

[Melgar et al., 2013] Melgar, D., Bock, Y., Sanchez, D., and Crowell, B. W. (2013). On robust and reliable automated baseline corrections for strong motion seismology. Journal of Geophysical Research: Solid Earth, 118(3):1177–1187.

[Melgar et al., 2015] Melgar, D., Crowell, B. W., Geng, J., Allen, R. M., Bock, Y., Riquelme, S., Hill, E. M., Protti, M., and Ganas, A. (2015). Earthquake magnitude calculation without saturation from the scaling of peak ground displacement. Geophysical Research Letters, 42(13):5197–5205.

[Melgar et al., 2020] Melgar, D., Crowell, B. W., Melbourne, T. I., Szeliga, W., Santillan, M., and Scrivner, C. (2020). Noise Characteristics of Operational Real-Time High-Rate GNSS Positions in a Large Aperture Network. Journal of Geophysical Research: Solid Earth, 125(7).

[Melgar et al., 2016] Melgar, D., LeVeque, R. J., Dreger, D. S., and Allen, R. M. (2016). Kinematic rupture scenarios and synthetic displacement data: An example application to the Cascadia subduction zone. Journal of Geophysical Research: Solid Earth, 121(9):6658–6674.

[Melgar et al., 2019] Melgar, D., Melbourne, T. I., Crowell, B. W., Geng, J., Szeliga, W., Scrivner, C., Santillan, M., and Goldberg, D. E. (2019). Real-Time High-Rate GNSS Displacements: Performance Demonstration during the 2019 Ridgecrest, California, Earthquakes. Seismological Research Letters, 91(4):1943–1951.

[Michelini et al., 2021] Michelini, A., Cianetti, S., Gaviano, S., Giunchi, C., Jozinović, D., and Lauciani, V. (2021). INSTANCE – the Italian seismic dataset for machine learning. Earth System Science Data, 13(12):5509–5544.

[Min et al., 2006] Min, Y., Hang, G., and Chengwu, Z. (2006). Application of Wavelet Analysis to GPS Deformation Monitoring. 2006 IEEE/ION Position, Location, And Navigation Symposium, pages 670–676.

[Minson et al., 2019] Minson, S. E., Baltay, A., Cochran, E., Hanks, T., Page, M. T., Mcbride, S. K., Milner, K., and Meier, M.-A. (2019). The limits of earthquake early warning accuracy and best alerting strategy. Scientific Reports, 9.

[Minson et al., 2014] Minson, S. E., Murray, J. R., Langbein, J. O., and Gomberg, J. S. (2014). Real-time inversions for finite fault slip models and rupture geometry based on high-rate GPS data. Journal of Geophysical Research: Solid Earth, 119(4):3201–3231.

[Misra and Enge, 2011] Misra, P. and Enge, P. (2011). Global Positioning System: Signals, Measurements and Performance. Ganga-Jamuna Press.

[Morlet et al., 1982] Morlet, J., Arens, G., Fourgeau, E., and Glard, D. (1982). Wave propagation and sampling theory—Part I: Complex signal and scattering in multilayered media. GEOPHYSICS, 47(2):203–221. doi: 10.1190/1.1441328.

[Moschas and Stiros, 2013] Moschas, F. and Stiros, S. (2013). Noise characteristics of high-frequency, short-duration GPS records from analysis of identical, collocated instruments. Measurement, 46(4):1488–1506.

[Mousavi and Beroza, 2022] Mousavi, S. M. and Beroza, G. C. (2022). Deep-learning seismology. Science, 377(6607):eabm4470. doi: 10.1126/science.abm4470.

[Mousavi et al., 2020] Mousavi, S. M., Ellsworth, W. L., Zhu, W., Chuang, L. Y., and Beroza, G. C. (2020). Earthquake transformer—an attentive deep-learning model for simultaneous earthquake detection and phase picking. Nature Communications, 11(1):3952.

[Mousavi et al., 2019] Mousavi, S. M., Sheng, Y., Zhu, W., and Beroza, G. C. (2019). STanford EArthquake Dataset (STEAD): A Global Data Set of Seismic Signals for AI. IEEE Access, 7:179464–179476.

[Murray et al., 2019] Murray, J. R., Bartlow, N., Bock, Y., Brooks, B. A., Foster, J., Freymueller, J., Hammond, W. C., Hodgkinson, K., Johanson, I., López-Venegas, A., Mann, D., Mattioli, G. S., Melbourne, T., Mencin, D., Montgomery-Brown, E., Murray, M. H., Smalley, R., and Thomas, V. (2019). Regional Global Navigation Satellite System Networks for Crustal Deformation Monitoring. Seismological Research Letters, 91(2A):552–572.

[Murray et al., 2018] Murray, J. R., Crowell, B. W., Grapenthin, R., Hodgkinson, K., Langbein, J. O., Melbourne, T., Melgar, D., Minson, S. E., and Schmidt, D. A. (2018). Development of a Geodetic Component for the U.S. West Coast Earthquake Early Warning System. Seismological Research Letters, 89(6):2322–2336.

[Nikolaidis et al., 2001] Nikolaidis, R. M., Bock, Y., Jonge, P. J. d., Shearer, P., Agnew, D. C., and Domselaar, M. V. (2001). Seismic wave observations with the Global Positioning System. Journal of Geophysical Research: Solid Earth, 106(B10):21897–21916.

[Ohta et al., 2012] Ohta, Y., Kobayashi, T., Tsushima, H., Miura, S., Hino, R., Takasu, T., Fujimoto, H., Iinuma, T., Tachibana, K., Demachi, T., Sato, T., Ohzono, M., and Umino, N. (2012). Quasi real-time fault model estimation for near-field tsunami forecasting based on rtk-gps analysis: Application to the 2011 tohoku-oki earthquake (mw 9.0). Journal of Geophysical Research: Solid Earth, 117(B2).

[Pedregosa et al., 2011] Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-Learn: Machine Learning in Python. J. Mach. Learn. Res., 12(null):2825–2830.

[Press, 1966] Press, S. J. (1966). Linear Combinations of Non-Central Chi-Square Variates. The Annals of Mathematical Statistics, 37(2):480–487.

[Psimoulis et al., 2018] Psimoulis, P. A., Houlié, N., Habboub, M., Michel, C., and Rothacher, M. (2018). Detection of ground motions using high-rate GPS time-series. Geophysical Journal International, 214(2):1237–1251.

[Qian et al., 2016] Qian, C., Liu, H., Zhang, M., Shu, B., Xu, L., and Zhang, R. (2016). A geometry-based cycle slip detection and repair method with time-differenced carrier phase (tdcp) for a single frequency global position system (gps) + beidou navigation satellite system (bds) receiver. Sensors, 16(12).

[R. Leandro, 2011] R. Leandro, H. Landau, M. N. M. G. S. S. X. C. A. D. M. B. F. Z. K. F. (2011). Rtx positioning: The next generation of cm-accurate real-time gnss positioning. Proc. of the 24th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS 2011), Portland, Oregon, page 1460–1475.

[Ruhl et al., 2018] Ruhl, C. J., Melgar, D., Geng, J., Goldberg, D. E., Crowell, B. W., Allen, R. M., Bock, Y., Barrientos, S., Riquelme, S., Baez, J. C., Cabral-Cano, E., Pérez-Campos, X., Hill, E. M., Protti, M., Ganas, A., Ruiz, M., Mothes, P., Jarrín, P., Nocquet, J., Avouac, J., and D'Anastasio, E. (2018). A Global Database of Strong-Motion Displacement GNSS Recordings and an Example Application to PGD Scaling. Seismological Research Letters, 90(1):271–279.

[Rydelek and Horiuchi, 2006] Rydelek, P. and Horiuchi, S. (2006). Is earthquake rupture deterministic? Nature, 442(7100):E5–E6.

[Seydoux et al., 2020] Seydoux, L., Balestriero, R., Poli, P., Hoop, M. d., Campillo, M., and Baraniuk, R. (2020). Clustering earthquake signals and background noises in continuous seismic data with unsupervised deep learning. Nature Communications, 11(1):3972.

[SHEN et al., 2019] SHEN, N., CHEN, L., LIU, J., WANG, L., TAO, T., WU, D., and CHEN, R. (2019). A Review of Global Navigation Satellite System (GNSS)-based Dynamic Monitoring Technologies for Structural Health Monitoring. Remote Sensing, 11(9):1001.

[Shu et al., 2018] Shu, Y., Fang, R., Li, M., Shi, C., and nan Liu, J. (2018). Very high-rate gps for measuring dynamic seismic displacements without aliasing: performance evaluation of the variometric approach. GPS Solutions, 22:1–13.

[Shu et al., 2020] Shu, Y., Fang, R., Liu, Y., Ding, D., Qiao, L., Li, G., and Liu, J. (2020). Precise coseismic displacements from the GPS variometric approach using different precise products: Application to the 2008 MW 7.9 Wenchuan earthquake. Advances in Space Research, 65(10):2360–2371.

[Simons et al., 2006] Simons, F. J., Dando, B. D., and Allen, R. M. (2006). Automatic detection and rapid determination of earthquake magnitude by wavelet multiscale analysis of the primary arrival. Earth and Planetary Science Letters, 250(1-2):214–223.

[Strauss and Allen, 2016] Strauss, J. A. and Allen, R. M. (2016). Benefits and Costs of Earthquake Early Warning. Seismological Research Letters, 87(3):765–772.

[Teunissen, 2020] Teunissen, P. J. (2020). GNSS Precise Point Positioning. pages 503–528.

[Trifunac and Todorovska, 2001] Trifunac, M. and Todorovska, M. (2001). A note on the useable dynamic range of accelerographs recording translation. Soil Dynamics and Earthquake Engineering, 21(4):275–286.

[Trnkoczy, 2012] Trnkoczy, A. (2012). Topic Understanding and parameter setting of STA/LTA trigger algorithm 1 Introduction. New Manual of Seismological Observatory Practice 2 (NMSOP-2).

[Wang et al., 2020] Wang, Y., Breitsch, B., and Morton, Y. T. J. (2020). A state-based method to simultaneously reduce cycle slips and noise in coherent gnss-r phase measurements from open-loop tracking. IEEE Transactions on Geoscience and Remote Sensing, pages 1–12.

[Wessel et al., 2019] Wessel, P., Luis, J. F., Uieda, L., Scharroo, R., Wobbe, F., Smith, W. H. F., and Tian, D. (2019). The generic mapping tools version 6. Geochemistry, Geophysics, Geosystems, 20(11):5556–5564.

[Williams et al., 2004] Williams, S. D. P., Bock, Y., Fang, P., Jamason, P., Nikolaidis, R. M., Prawirodirdjo, L., Miller, M., and Johnson, D. J. (2004). Error analysis of continuous GPS position time series. Journal of Geophysical Research: Solid Earth, 109(B3).

[Williamson et al., 2020] Williamson, A. L., Melgar, D., Crowell, B. W., Arcas, D., Melbourne, T. I., Wei, Y., and Kwong, K. (2020). Toward Near-Field Tsunami Forecasting Along the Cascadia Subduction Zone Using Rapid GNSS Source Models. Journal of Geophysical Research: Solid Earth, 125(8).

[Withers et al., 1998] Withers, M., Aster, R., Young, C., Beiriger, J., Harris, M., Moore, S., and Trujillo, J. (1998). A comparison of select trigger algorithms for automated global seismic phase and event detection. Bulletin of the Seismological Society of America, 88(1):95–106.

[Woollam et al., 2022] Woollam, J., Münchmeyer, J., Tilmann, F., Rietbrock, A., Lange, D., Bornstein, T., Diehl, T., Giunchi, C., Haslinger, F., Jozinović, D., Michelini, A., Saul, J., and Soto, H. (2022). SeisBench—A Toolbox for Machine Learning in Seismology. Seismological Research Letters, 93(3):1695–1709.

[Worden et al., 2020] Worden, C. B., Thompson, E. M., Hearne, M., and Wald, D. J. (2020). ShakeMap Manual Online: Technical manual, user's guide, and software guides. USGS.

[Wu et al., 2011] Wu, L., Ma, H., Ding, W., Hu, Q., Zhang, G., and Lu, D. (2011). Study of GPS Data De-noising Method Based on Wavelet and Kalman filtering. 2011 Third Pacific-Asia Conference on Circuits, Communications and System (PACCS), 1:1–3.

[Wu and Zhao, 2006] Wu, Y.-M. and Zhao, L. (2006). Magnitude estimation using the first three seconds p-wave amplitude in earthquake early warning. Geophysical Research Letters, 33(16).

[Yang et al., 2017] Yang, R., Ling, K.-V., Poh, E.-K., and Morton, Y. (2017). Generalized GNSS Signal Carrier Tracking—Part II: Optimization and Implementation. IEEE Transactions on Aerospace and Electronic Systems, 53(4):1798–1811.

[Yin et al., 2022] Yin, J., Denolle, M. A., and He, B. (2022). A multitask encoder–decoder to separate earthquake and ambient noise signal in seismograms. Geophysical Journal International, 231(3):1806–1822.

[Yosinski et al., 2014] Yosinski, J., Clune, J., Bengio, Y., and Lipson, H. (2014). How transferable are features in deep neural networks? arXiv.

[Zhu et al., 2020] Zhu, W., Mousavi, S. M., and Beroza, G. C. (2020). Seismic signal augmentation to improve generalization of deep neural networks. Advances in Geophysics.

[Zumberge et al., 1997] Zumberge, J. F., Heflin, M. B., Jefferson, D. C., Watkins, M. M., and Webb, F. H. (1997). Precise point positioning for the efficient and robust analysis of gps data from large networks. Journal of Geophysical Research: Solid Earth, 102(B3):5005–5017.