

1 Estimation of seismic moment tensors using variational 2 inference machine learning

3 **Andreas Steinberg^{1,*}, Hannes Vasyura-Bathke^{2,*}, Peter Gaebler¹, Matthias**
4 **Ohrnberger², Lars Ceranna¹**

5 ¹Federal Institute for Geosciences and Natural Resources (BGR), Hannover, Germany

6 ²Institute for Earth and Environmental Sciences, University of Potsdam, Potsdam, Germany

7 *These authors contributed equally to this work.

8 **Key Points:**

- 9 • Ensemble of moment tensors of earthquakes are determined with Bayesian Neu-
10 ral Networks trained on synthetic waveforms
- 11 • Uncertainties in centroid location and time as well as uncertainties in Earth's
12 structure are considered by variational inference
- 13 • Application to a subset of the 2019 Ridgecrest sequence and comparison to
14 independent moment-tensor estimates shows robust performance

Corresponding author: Andreas Steinberg, andreas.steinberg@bgr.de

Abstract

We present an approach for estimating in near real-time full moment tensors of earthquakes and their parameter uncertainties based on short time windows of recorded seismic waveform data by considering deep learning of Bayesian Neural Networks. The individual neural networks are trained on synthetic seismic waveform data and corresponding known earthquake moment-tensor parameters. A monitoring volume has been pre-defined to form a three-dimensional grid of locations and to train a Bayesian neural network for each grid point. Variational inference on several of these networks allows us to consider several sources of error and how they affect the estimated full moment-tensor parameters and their uncertainties. In particular, we demonstrate how estimated parameter distributions are affected by uncertainties in the earthquake centroid location in space and time as well as in the assumed Earth structure model. We apply our approach on seismic waveform recordings of aftershocks of the Ridgecrest 2019 earthquake with moment magnitudes ranging from Mw 2.7 to Mw 5.5. Overall, good agreement has been achieved between inferred parameter ensembles and independently estimated parameters using classical methods. Our developed approach is fast and robust, and therefore, suitable for operational earthquake early warning systems.

1 Introduction

Robust and fast estimation of the source mechanism of earthquakes, i.e., the seismic moment tensor (MT), is important for many near-real time hazard assessments (earthquake early warning), and provides helpful information for evaluating appropriate measures and responses. Furthermore, hazard assessments and physics based aftershock probability calculations can be improved by using the inferred full seismic MT. Routine operational monitoring frameworks such as the United States Geological Survey (USGS) and GEOFON provide automatic centroid moment tensor (CMT) point-source solutions within minutes for moderate and large earthquakes ($>Mw 4.5$), usually in telseismic distances (Ekström, Nettles, & Dziewoński, 2012; Hanka & Kind, 1994). However, the MTs for smaller regional or local earthquakes, are often only analysed after manual inspection with delay times of up to days. The estimation of the full MT of smaller earthquakes ($>Mw 3$) can be important for detailed analysis of fore- and aftershock sequences, inference of local stress redistribution and especially, for seismicity monitoring in geotechnical applications (Cesca, Şen, & Dahm, 2014), where significant non double-couple (DC) components due to volumetric changes can be expected.

CMTs are usually estimated as solutions to an inverse problem by iterative comparison of synthetic and observed waveform data until a sufficient match is achieved. Forward modelling of synthetics is typically performed by assuming a point source and by considering a range of potential source model parameters and their combinations; whereas the uncertainties of the estimated parameters are quantified by considering data errors and theory errors which are introduced by the measurement and the assumptions in the inverse problem, respectively (Vasyura-Bathke et al., 2020). Uncertainties can be obtained through probabilistic approaches (Duputel, Rivera, Fukahata, & Kanamori, 2012; Kühn, Heimann, Isken, Ruigrok, & Dost, 2020; Stähler & Sigloch, 2014, 2016; Vackář, Burjánek, Gallovič, Zahradník, & Clinton, 2017; Vasyura-Bathke et al., 2020, e.g.), but these methods are computationally expensive and the estimation of CMT parameter densities can take tens of minutes to hours. Faster estimates would greatly increase the capabilities of earthquake early warning systems. Machine learning algorithms have been shown to be helpful and fast for seismic signal detection and localisation (Kriegerowski, Petersen, Vasyura-Bathke, & Ohrnberger, 2019; Smith, Ross, Azizzadenesheli, & Muir, 2021), phase picking (Mousavi, Ellsworth, Zhu, Chuang, & Beroza, 2020; Ross, Meier, & Hauksson, 2018) as well as initial characterization of the seismic source (e.g., Käuffl, Valentine, O’Toole, & Trampert, 2014; van den Ende &

67 Ampuero, 2020). P-wave first-motion polarity can be used to determine the MT of
 68 earthquakes assuming a DC source, which has been shown to be fast and reliably to
 69 enhance MT catalogs using deep learning (Hara, Fukahata, & Iio, 2019; Ross et al.,
 70 2018; Uchide, 2020). Recently, deep learning has been used to train the so called Fo-
 71 cal Mechanism Network (FMNet) to determine pure DC MTs based on full waveform
 72 synthetics Kuang, Yuan, and Zhang (2021). The FMNet has 16 trainable layers and
 73 was applied to four 2019 Ridgecrest earthquakes with magnitude larger than Mw 5.4.
 74 The network was trained on subjectively pre-defined Gaussian distributions as labels,
 75 describing the assumed distribution of the DC parameters strike, dip and rake.

76 Here, we present a machine learning framework employing several Bayesian Neu-
 77 ronal Networks (BNN) and using variational inference. Comprehensive consideration
 78 of errors are especially important for estimates obtained from unsupervised machine
 79 learning algorithms, as these are often treated and used as black boxes. Our BNNs are
 80 trained on synthetic waveforms with the aim to estimate MT parameters in near-real
 81 time considering errors in measurement and theory. We validate our approach on a
 82 subset of earthquakes from the aftershocks of the Californian Ridgecrest 2019-2020
 83 sequence (Ross et al., 2019), as the Ridgecrest area is exceptionally well monitored
 84 with a dense station distribution, both in azimuth and distance (Fig. (1,a)). The main
 85 shock of the 2019 Ridgecrest sequence was the Mw 7.1 2019-07-06 03:19:52 earthquake,
 86 preceded by several foreshocks of which the largest was the Mw 6.4 2019-07-04 17:33:49
 87 earthquake. The following months several hundred aftershocks \geq Mw 3 were recorded
 88 (Ross et al., 2019). For the subset of earthquakes from the 2019 Ridgecrest sequence
 89 we investigate earthquakes with moment magnitudes M_W between 2.7 and 5.5. We
 90 compare our estimations with the moment tensors provided by the Southern California
 91 Earthquake Data Center (SCEDC).

92 **2 Variational inference Neural Network estimation of Moment Ten-** 93 **sors**

94 Our main goal is to infer the radiation pattern and the orientation of the earth-
 95 quake source. We train location specific neural networks for each point of a pre-defined
 96 grid of potential hypo-centers based on full sets of synthetic waveforms with associated
 97 source model parameters to be learned. We use a set of 41 broadband stations within
 98 a range up to 150 km around the center of our grid (Fig. 1,a). The grid (Fig. 1,b)
 99 extends horizontally 10.5 by 10.5 km, with a step size of 1.5 km. The vertical extent
 100 ranges from 2 km to 10 km depth, in 2 km steps.

101 As prior information our proposed framework needs a detection of an earthquake
 102 and the associated approximate source time. Furthermore, an approximate earthquake
 103 location can be considered. Nevertheless, it has already been demonstrated that detec-
 104 tion and location of earthquakes are timely deliverable by other established machine
 105 learning based algorithms (Kriegerowski et al., 2019; Mousavi et al., 2020). Our ap-
 106 proach does not estimate earthquake moment magnitudes and is indirectly limited to
 107 a range of magnitudes (e.g. between Mw 3 and 5) as the network training depends on
 108 signal processing parameters. The magnitude of earthquakes can be readily estimated
 109 in real time by other approaches (van den Ende & Ampuero, 2020).

110 **2.1 Input**

111 As input we use synthetic displacement waveform data calculated for a specific
 112 earthquake source and for all considered stations in E, N and Z components. Training
 113 on synthetic data has several advantages compared to training on recorded data sets.
 114 The procedure is applicable to regions with low seismicity, and furthermore, the use of
 115 synthetic waveforms allows exploring the full range of possible CMTs. Consequently,
 116 the training is not restricted by a biased set of catalog mechanisms from available

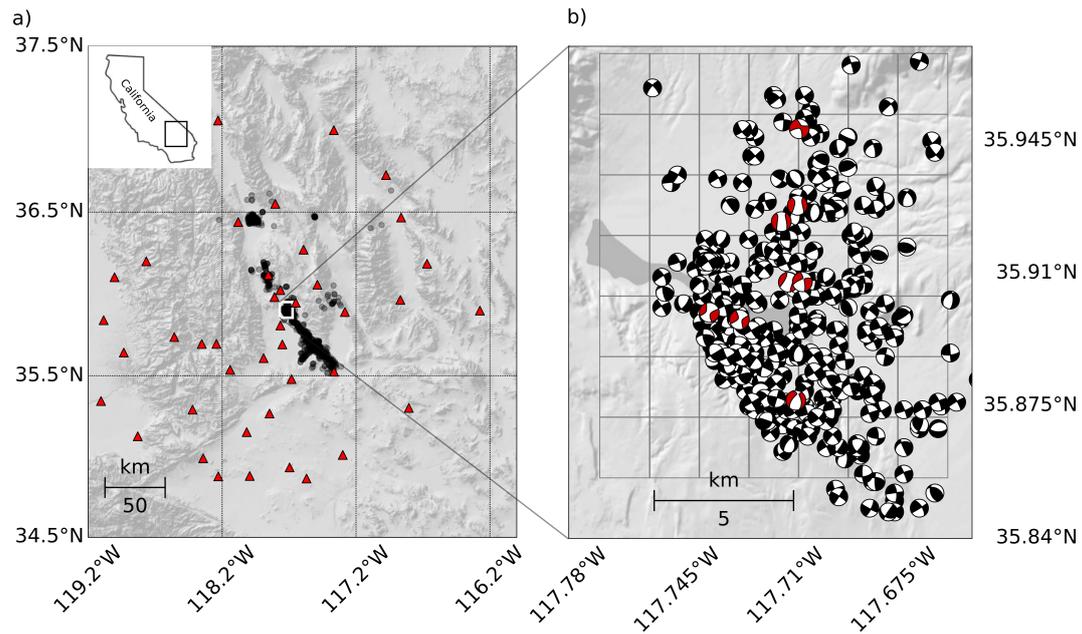


Figure 1: a) Region of interest, seismicity from 2019-07-04 to 2021-01-26 (black dots) and the station distribution (red triangles). Top-left inset shows the location of the map in California. The white rectangle shows the location of the study area. b) Zoom in to the study area. The black lines mark the grid of locations for which individual Neural Networks are trained. The focal mechanisms of earthquakes between July 2019 and December 2020 used for testing are plotted for full and double-couple CMTs in red and black, respectively. The indicated mechanisms are given as determined by SCEDC. Background in both a) and b) is a shaded relief of a digital elevation model.

117 observations, but it can be assured that the complete parameter space has been ex-
 118 plored. For fast simulation of synthetic waveforms we use pre-calculated Green’s func-
 119 tions (GF) stores from the Pyrocko software framework (Heimann et al., 2017, 2019).
 120 These GF stores are based on 1-D layered Earth structure models computed by using
 121 the reflectivity-type wavenumber integration method implemented in QSEIS (Wang,
 122 1999). We calculate three different GF stores based on 1-D velocity profiles (Supp.
 123 Fig. S1): 1) for the entire Mojave Region used by the USGS and the SCEDC, 2) for
 124 the Coso Geothermal area (Wu & Lees, 1999) and 3) for a regional shallow velocity
 125 profile based on Crust2.0 (Bassin, Laske, & Masters, 2000).

126 We train our neural networks on pure synthetic waveforms without adding noise,
 127 because the characteristics of the noise would be learned as well by the networks. We
 128 filter the waveforms with a butterworth bandpass filter of fourth order between 0.8 and
 129 2.4 Hz to avoid poor long-period response and weak long-period signals below the cor-
 130 ner frequency of Mw 3.5 earthquakes (Aki & Richards, 2002). We assume a triangular
 131 source time function of fixed duration of 0.5 seconds, representative of earthquakes in
 132 the magnitude range 3-3.5 (Aki & Richards, 2002). Therefore, our trained networks
 133 are restricted to specific frequencies. This implies that our trained networks are only
 134 valid for a pre-defined magnitude range and that for studying earthquakes of different
 135 magnitudes, additional specific networks would need to be trained. For each source
 136 grid point location and the given 1-D Earth structure model we use the expected the-
 137 oretical travel times to extract a snippet of waveform data 1 s before and 4 s after the
 138 theoretical first phase arrival. This also means that our extracted waveform snippets
 139 are relative in time and that they can be used for all possible centroid times in the
 140 training phase. To cut out real data, however, this means that the centroid time needs
 141 to be known.

142 We convert the extracted waveform snippets around the P-wave onset to form
 143 a 2D input image such that the rows represent the waveforms that are grouped first
 144 by channels (E, N, Z) and second by stations; the columns represent the samples
 145 over time. Finally, we normalize and re-scale the image by the absolute maximum
 146 amplitude of the full image such that all values fall between the closed interval of 0
 147 and 1, where 0.5 indicates zero in the original waveform amplitudes as well as missing
 148 data. Due to this normalization all synthetics can be calculated for one single (but
 149 arbitrary) magnitude. The order of stations needs to be consistent for each image
 150 and must not change. Here, we chose an alphabetical order according to the station
 151 codes as arranging by azimuth or distance would be different for each considered source
 152 location and would cause artificial patterns which in turn would make efficient training
 153 of the networks difficult.

154 2.2 Labels

155 For each set of synthetic waveforms forming an input image we know the par-
 156 ameters of the underlying source. These are the output labels that our networks
 157 predict. The common MT parameterization with six independent components (Aki
 158 & Richards, 2002; Madariaga, 2007) seems a natural choice for describing the seis-
 159 mic source. However, a uniform sampling in this parameter space does not yield a
 160 uniform unique distribution of samples in moment-tensor space (Tape & Tape, 2015).
 161 Such a non-uniform and non-unique mapping would lead to bias in learned patterns
 162 for our networks. This problem can be solved by using spherical coordinates on the
 163 unit sphere of the fundamental lune description of the moment tensor (Tape & Tape,
 164 2012b). Moreover, this parameterization allows for a uniform sampling of moment-
 165 tensors, with the advantage of only five independent parameters to describe the full
 166 spectrum of moment tensors. These five independent parameters (Tab. 1) are: κ as the
 167 strike-angle equivalent, σ as the rake-angle equivalent of the moment tensor slip angle,
 168 h as the dip-angle equivalent and the non-isotropic components v and w as the lune

Table 1: Lune parameter definitions and chosen discretization for constructing the training dataset.

Parameter	interpretation	min. value	max. value	step size
κ	strike angle	0	2π	0.1π
σ	rake angle	$-\frac{\pi}{2}$	$\frac{\pi}{2}$	0.2
h	dip angle	0	1	0.2
w	Lune latitude	$-\frac{3}{8}\pi$	$\frac{3}{8}\pi$	0.2
v	Lune co-longitude	$-\frac{1}{3}$	$\frac{1}{3}$	0.02

169 latitude and co-longitude, respectively (Tape & Tape, 2015). This parameterization
 170 of the MT clearly separates the radiation pattern from the source orientation. We
 171 choose a discretization of $0.1 \cdot \pi$ for κ , 0.2 for σ, h, w and 0.02 for v . This results to
 172 171.600 synthetic waveform datasets that we use for training for each single location
 173 grid-point.

174 2.3 Network design

175 Instead of using deterministic network layers where scalar weights and biases are
 176 learned, we use their probabilistic expression with distributions of weights and biases.
 177 Each distribution is assumed to be Gaussian with mean μ and a standard deviation $\hat{\sigma}$
 178 (e.g. Blundell, Cornebise, Kavukcuoglu, & Wierstra, 2015; Graves, 2011; Wen, Vicol,
 179 Ba, Tran, & Grosse, 2018). A neural network designed with such probabilistic layers
 180 (i.e., flipout layers) forms a Bayesian Neural Network (BNN) and can be considered
 181 as representing an ensemble of deterministic neural networks trained several times on
 182 the same input data. These BNNs allow to represent epistemic uncertainty in their
 183 inherent predictions due to limited training data and they yield a likelihood value to
 184 each drawn sample. Consequently, rather than predicting the same set of output labels
 185 given the same input data, repeated forward pass yields a distribution of output labels,
 186 i.e. uncertainties in lune parameters. This can vary for each individual BNN learned
 187 for the grid points, as the significance of specific seismic stations towards the source
 188 will vary.

189 Each single training iteration of a BNN consists of a forward pass and a back-
 190 propagation pass (Wen et al., 2018). In the forward pass a single sample is drawn
 191 from the output labels. During a backwards pass the gradients of the layer weights
 192 and bias distributions (i.e. means μ and standard-deviations $\hat{\sigma}$) are calculated with
 193 automatic differentiation and μ and $\hat{\sigma}$ are then updated to maximize an objective
 194 function depending on the input and output labels (Wen et al., 2018).

195 Our goal is to use a simple neuronal network architecture to avoid over-fitting
 196 and to allow for straightforward interpretation of the individual training steps. The
 197 network design (Fig. 2,a) is similar in rationale to Kriegerowski et al. (2019). We use
 198 three 2-D convolutional flipout hidden layers. The first two hidden layers are sensitive
 199 to the information over time only (Fig. 2,a). The first hidden layer has 8 filters
 200 and a 1 by 2 kernel and the second layer has 10 filters and a 1 by 30 kernel. The
 201 last 2-D hidden layer collects information over the station components with 12 filters
 202 and a 3 by 1 kernel. We use a dropout of 0.2 between convolutional flipout layers
 203 to robustly handle data errors and missing waveform data at particular stations. We
 204 downsample the output data of the convolutional flipout layers with a 2-D max pooling
 205 layer with a 3 by 4 kernel (example activations see Fig. 2,c) followed by flattening
 206 the data into a vector and feeding them into a fully connected dense flipout layer.
 207 The relatively simple network design allows for visual inspection of the activations
 208 in each layer (Supp. Fig S2). All convolutional flipout layers are activated using a

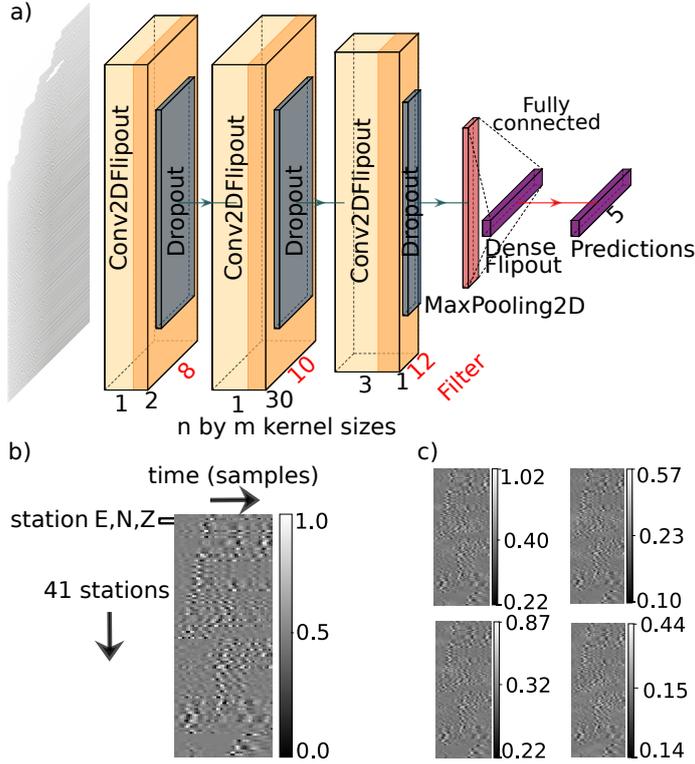


Figure 2: a) Design of an individual Bayesian Neural Network. b) example 2d normalized array input from synthetic waveforms. Blue arrows indicate hidden layers with RELU activation. c) exemplary activation of the pooling layer given the input of b)

209 Rectified Linear Unit function (RELU) (Glorot, Bordes, & Bengio, 2011). Finally, a
 210 non activated lambda distribution layer is used to hold the resulting distributions of
 211 predicted source parameters. As objective function (loss function) we use the negative
 212 log-likelihood and as optimizer the Adam algorithm (Kingma & Ba, 2014).

213 2.4 Variational inference from multiple BNN

214 The probabilistic output of the BNNs allows to combine inferences at several
 215 likely locations and centroid times of the earthquake's source. Each evaluation of
 216 a network with inputs yields a single prediction of the 5 MT parameters and the
 217 associated negative log-likelihood. The inferences from all these individual evaluations
 218 of networks can be combined and the source's errors in both centroid location and time
 219 can be propagated to uncertainties in MT parameter marginals through variational
 220 inference, yielding an ensemble of possible source mechanisms.

221 We consider an error in location within an ellipse around the assumed centroid
 222 location and evaluate the respective BNNs with the given input (Fig.3). Note that
 223 waveform snippets are extracted differently from the waveform input according to
 224 theoretical arrival times at each receiver location (sec. 2.1).

225 Errors in centroid time result in shifts of the predicted theoretical arrival times
 226 and the extracted waveform snippets. We assume uniform distributed errors in timing
 227 and draw random samples within the timing errors and therefore, all BNNs at the
 228 considered grid points are evaluated several times. Consequently, we get different
 229 likelihoods to the differently extracted waveform snippets. In classical approaches in

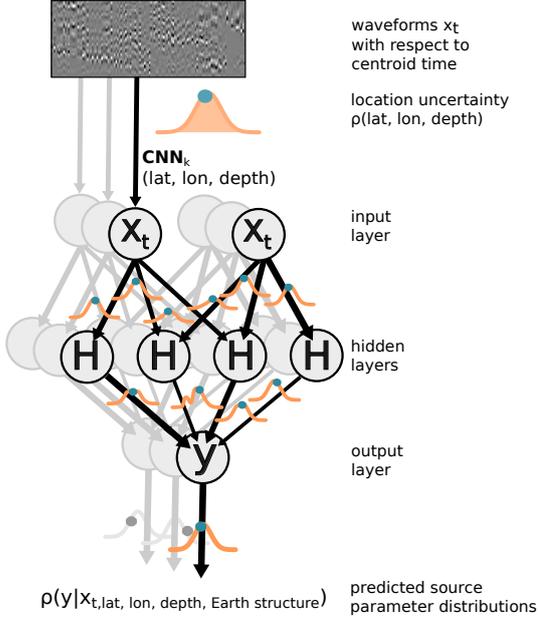


Figure 3: Training scheme, the normalized waveforms are timeshifted n -times within a timing error. The location uncertainty of an event determines the number of k BNN's that are used for prediction and contribute to the posterior probability density $\rho(\mathbf{m}|x)$.

230 seismology this corresponds roughly to shifting the waveforms to find the maximum
 231 correlation (e. g. Kühn et al., 2020).

232 Finally, errors in the Earth structure model can to be taken into account for ro-
 233 bust inference on the estimated source mechanism (Vasyura-Bathke, Dettmer, Dutta,
 234 Mai, & Jónsson, 2021). The theory error from the choice of the 1-D Earth structure
 235 model can be included in our framework by training and evaluating BNNs on each
 236 grid-point for each Earth structure. This requires calculation of the full set of syn-
 237 thetic waveforms for different Earth structures; in our case, three structures (Supp.
 238 Fig S1). This results in total to over 100 Million waveform datasets on which the 588
 239 BNNs (196 grid points times three Earth structures) are trained. The calculation of
 240 synthetic waveforms and the network training was done in parallel on several machines
 241 with a total of 128 CPUs over a period of three months. By using GPUs this time
 242 could be drastically reduced to a few days.

243 2.5 Moment tensor ensemble similarity

244 To asses the similarity between the predicted ensemble of MTs and a reference
 245 solution, e.g. from a catalog, we use the omega angle measure (Tape & Tape, 2012a).
 246 The omega angle has the advantage that focal mechanisms with opposite polarities
 247 are considered most dissimilar in contrast to other measures, e.g., the Kagan angle
 248 (Cesca et al., 2014; Tape & Tape, 2012a). The normalized omega angle distance d
 249 (Cesca et al., 2014; Tape & Tape, 2012a) between two moment tensors U_1 and U_2 with
 250 components I and J is calculated by:

$$d_\omega = \frac{1}{2} \left[1 - \frac{U_1 \cdot U_2}{\|U_1\| \|U_2\|} \right] = \frac{1}{2} \left[1 - \frac{\sum_{i,j=1}^{I,J} U_{1ij} \cdot U_{2ij}}{(\sum_{i,j=1}^{I,J} U_{1ij}^2)^{\frac{1}{2}} (\sum_{i,j=1}^{I,J} U_{2ij}^2)^{\frac{1}{2}}} \right] \quad (1)$$

251 It is defined between 0 and 1, for identical and opposite seismic radiation patterns
 252 between the two compared moment tensors, respectively. Note, that in order to calcu-
 253 late d_ω we need to convert our predicted MT ensemble from the Lune parameterization
 254 to the North-East-Down coordinate system (Aki & Richards, 2002).

255 3 Application to the Ridgecrest 2019 earthquake aftershock sequence

256 We train our networks for an area South of the Coso geothermal field (Fig. 1),
 257 which is known to host both induced and tectonic earthquakes (Monastero et al., 2005;
 258 Schoenball, Davatzes, & Glen, 2015). Significant non-DC components can be expected
 259 for earthquakes in this region (Ichinose, Anderson, Smith, & Zeng, 2003), potentially
 260 also for tectonic earthquakes, due to the influence of the geothermal reservoir. To
 261 test the performance of our framework we use recorded waveform data of the after-
 262 shocks that occurred between July 2019 and December 2020 to the Mw 7.1 Ridgecrest
 263 earthquake.

264 For these aftershocks, 8 full moment tensor solutions and 198 pure DC MT
 265 solutions (Fig. 1,b) are calculated (Hauksson & Unruh, 2007; Jordan & Maechling,
 266 2003) and made publicly available by the SCEDC ((SCEDC), 2013). We compare the
 267 MT estimates of our approach to the moment tensors as determined independently by
 268 USGS and SCEDC (Hutton, Woessner, & Hauksson, 2010).

269 We download the waveform data for all events and for the 41 stations from the
 270 Southern California Seismic Network (California Institute Of Technology And United
 271 States Geological Survey Pasadena, 1926). Missing waveform data for any station and
 272 time period are replaced by zero values in the waveform data, which are then mapped
 273 to 0.5 values in the normalized input images. Measured waveform data are treated in
 274 the same way as our synthetic waveforms (sec. 2.1), i.e. data is restituted to ground
 275 displacement and down-sampled to match the Green’s function sampling rate of 14 Hz.

276 For each aftershock we evaluate the BNNs for a total of 6000 samples. However,
 277 the number of activated BNNs depends on the uncertainties in centroid location and
 278 time as provided by the SCEDC catalog. The location uncertainty in horizontal and
 279 vertical position as given by the SCEDC is increased 10 times, as reported uncertainties
 280 are in the order of few hundreds of meters. The total ensemble of samples is then
 281 obtained by evaluating the activated BNNs equally.

282 3.1 Inferences for full moment tensors

283 We focus primarily on 8 aftershocks for, which a full moment tensor solution
 284 is available in the SCEDC catalog. We refer to these solutions as ”reference” in the
 285 following.

286 We first evaluate only the waveform input with the BNN’s trained using synthet-
 287 ics based on the *Mojave* Earth structure, which is the same as used by the SCEDC
 288 to determine their focal mechanisms (Supp. Fig. S1). Consequently, the reference and
 289 predicted MTs should be consistent in their epistemic uncertainty as the same Earth
 290 structure model and (mostly the same) dataset is used. We use the SCEDC catalog
 291 values for source position and centroid time. For the comparison we only consider un-
 292 certainty in centroid location. We find very good agreement of our predicted ensembles
 293 to most of the 8 reference moment tensors, with most of the omega angle distances d_ω
 294 being below 0.1 (Fig 4,a-h). Histograms of d_ω show their maximum mostly within the
 295 first few bins. Only, two ensembles of predicted moment tensors show small system-
 296 atic errors (Fig 4,f and g). For those also the histograms of d_ω show their maxima at
 297 distances above zero.

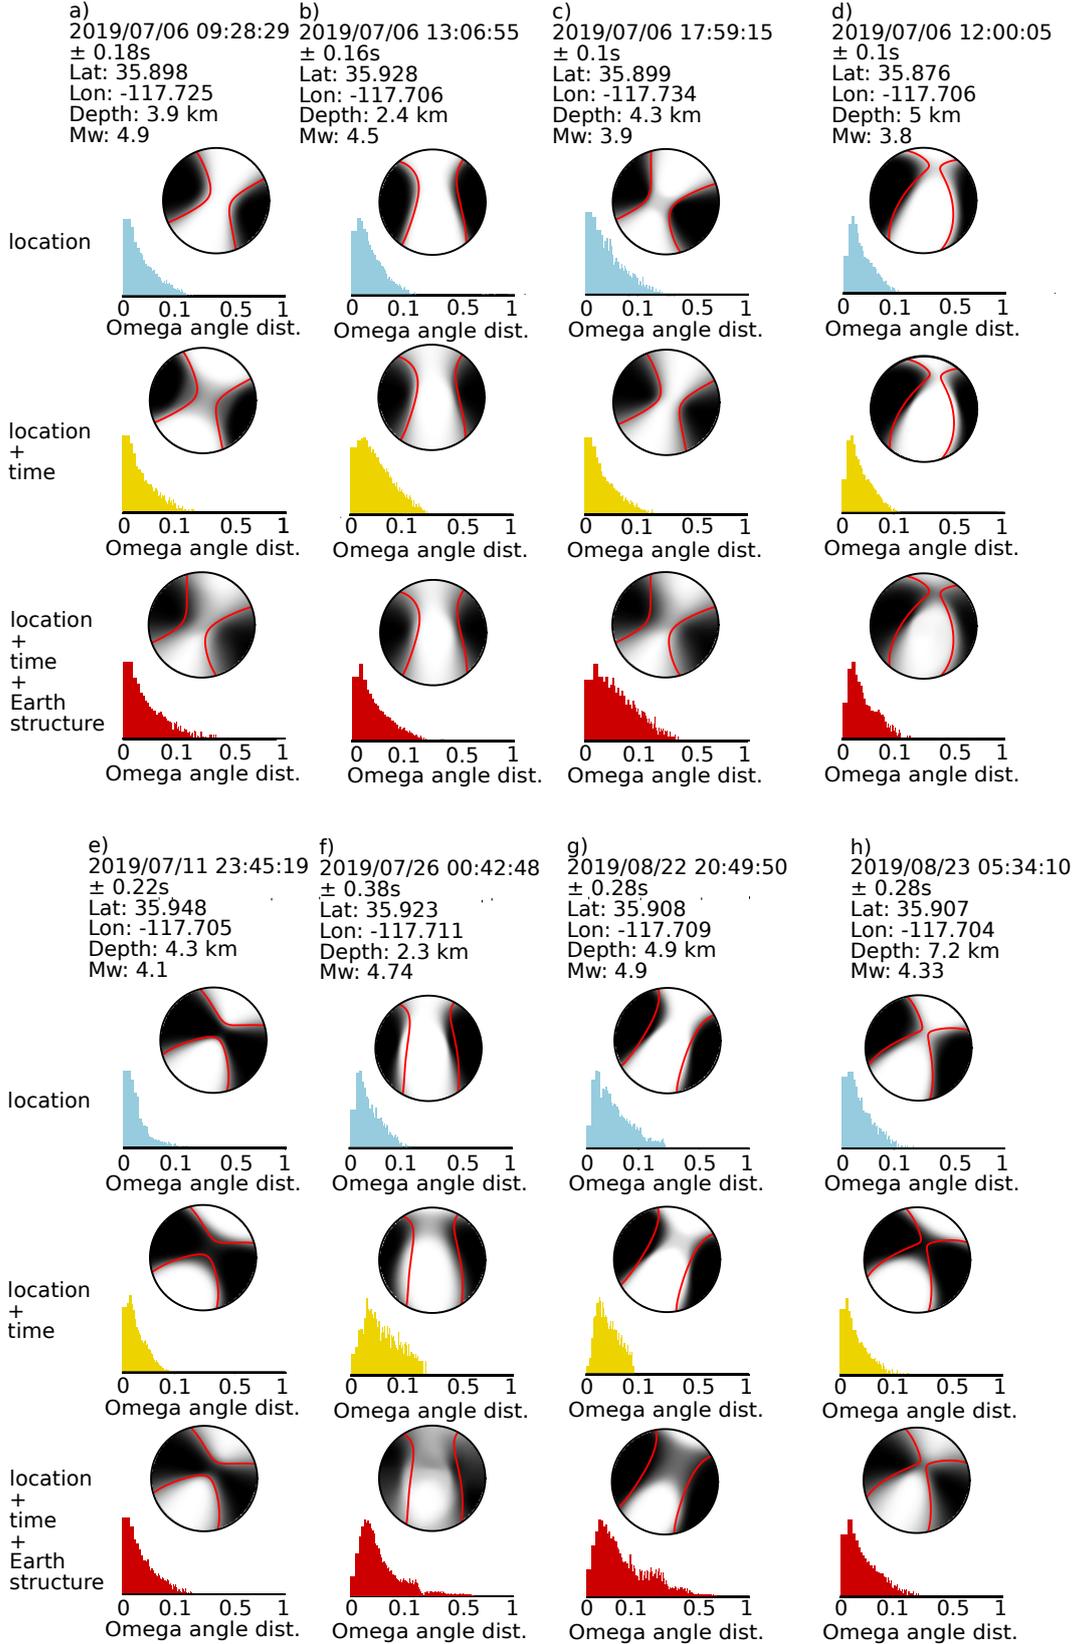


Figure 4: Inferred ensembles of full MTs considering different uncertainties. a) to h) show three fuzzy beachballs (BB), each based on 6000 MT predictions, where the reference MT (from SCEDC) is marked by red lines. The top BB is based on the predictions from BNNs trained on the *Mojave* Earth structure model and considering uncertainty in centroid location. The middle BB additionally includes uncertainty in the centroid time. The bottom BB shows the inferences of BNNs additionally considering inferences from all three Earth structure models. The normalized histogram of the omega angle distances d_{ω} between the reference MT and the ensemble of MT predictions is shown below each BB. Note that the x-axis for all subfigures scales quadratic.

298 In addition to uncertainty in centroid location we consider in the following un-
 299 certainties in the centroid time, which are also provided by the SCEDC catalog for
 300 each event. These uncertainties differ from earthquake to earthquake but they do
 301 not exceed 0.4s for the considered aftershocks. When uncertainties in centroid times
 302 are considered the widths of some of the d_ω histograms increase for some ensembles
 303 of MT predictions (Fig 4,b,f) confirming the quality of the absolute centroid times of
 304 these aftershocks determined by SCEDC. However, it is worth mentioning that the
 305 widths of some of the d_ω histograms also decrease for some ensembles of MT predic-
 306 tions (Fig 4,g,h,c) suggesting biased absolute centroid times for those aftershocks in
 307 the catalogue.

308 Finally, in addition to uncertainties in centroid location and time we consider
 309 uncertainties in Earth structure. We evaluate the BNNs that have been trained on the
 310 synthetics from three different Earth structures (Supp. Fig. S1). The expected arrival
 311 times and thus extracted waveform snippets will be systematically different for each
 312 Earth structure. For some of the inferred MT ensembles the spread in d_ω histograms
 313 increases and some show values of up to 0.5 (Fig 4,f-g). For those events the *Mojave*
 314 structure model seems to be the most appropriate one and therefore uncertainties in
 315 Earth structure are overestimated. For other MT ensembles (Fig 4,a,d,e,h) the spread
 316 in d_ω histograms decreases or stays similar, meaning that the uncertainties in Earth
 317 structure are less crucial for those events. Nevertheless, the resulting ensembles of
 318 predicted MTs also comprise the solutions of considering only location uncertainty
 319 (Fig. 6) and, the maximum a-posterior (MAP) solution still shows good agreement
 320 between extracted waveform data snippets and synthetic waveforms calculated from
 321 the predicted source parameters (Fig. 5).

322 3.2 Inferences for double-couple moment-tensors

323 The SCEDC catalog also contains 198 pure double-couple focal mechanisms for
 324 events that occurred in the area of interest, which we refer to as reference in the fol-
 325 lowing. Without visual inspection we let for the waveform data of each of those events
 326 our BNNs infer ensembles of 6000 MT solutions considering centroid location and time
 327 uncertainty. We compare the 198 reference focal mechanisms with our ensembles of
 328 MT parameter predictions from our framework by setting the predicted v and w values
 329 to zero, representing a pure double-couple source (Fig. 7,a). We also show d_ω between
 330 the reference mechanism and the predicted full seismic MT ensembles (Fig. 7,b). The
 331 additional degree of freedom of full MT solutions versus DC constrained solutions
 332 results in broadening and a slight shift of the histogram towards higher d_ω (Fig. 7,b).

333 With decreasing earthquake magnitude the spread of d_ω of the trained networks is
 334 increasing comparing the predicted ensembles of MT for the 198 earthquakes (Fig. 7,c).
 335 This spread is expected as the signal-to-noise ratio decreases with lower magnitude and
 336 larger d_ω values are expressions of an increase in uncertainty of the MT ensembles.
 337 However, the bulk part of d_ω shows distances below 0.1 and the predicted ensembles
 338 are in good agreement with the reference solutions across different magnitudes 2.7-
 339 4.5 (Fig. 7,d-g). We also notice a slight increase in the omega angle distances between
 340 reference and predicted source mechanisms for the largest of the 198 earthquakes. This
 341 might indicate a need for incorporating non-DC components in the source mechanism;
 342 whereas these components are missing in the catalogue descriptions.

343 4 Discussion

344 In general, we find a good agreement between the ensemble of predicted MTs and
 345 the independently determined and unseen moment tensor solutions from the SCEDC.
 346 Only a few predicted moment tensor ensembles show systematic differences (Fig 4,f and
 347 g), which could be due to several reasons, e. g. differences in the station configurations.

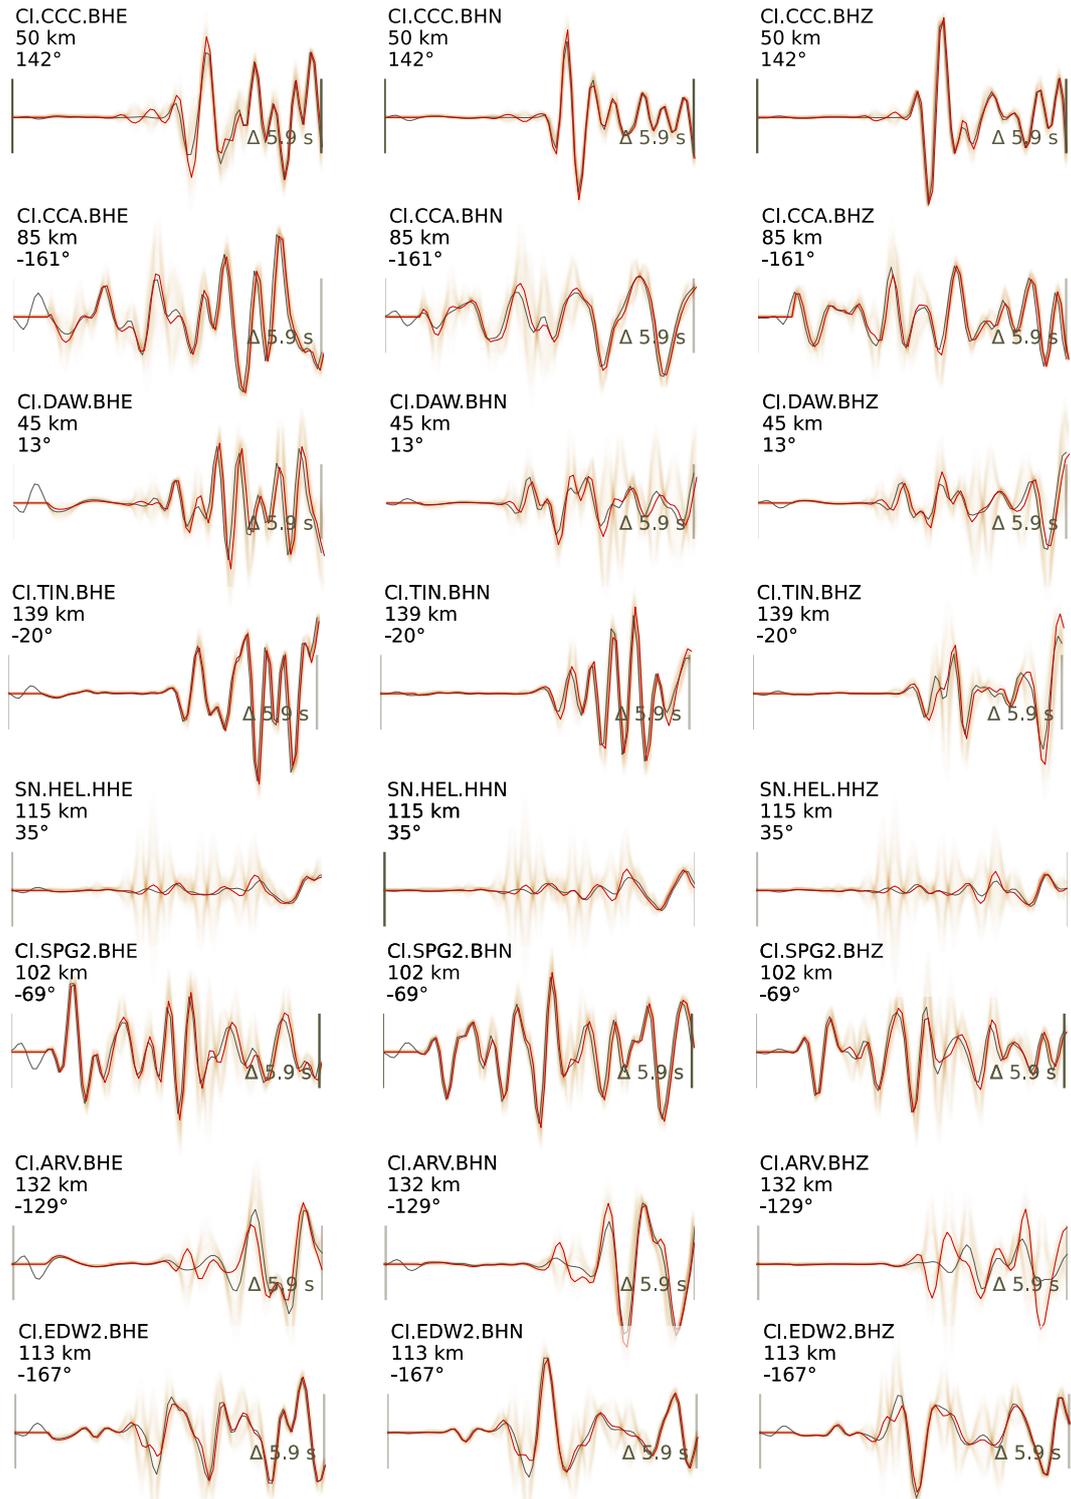


Figure 5: Exemplary waveform fits between observed waveforms (black) and synthetic waveforms based on the ensemble of estimated MT parameters (brown) with the MAP in red, for the Mw 4.1 earthquake on 2019/07/11 23:45:19. Note that the waveforms are displacements and normalized as described in section 2.1.

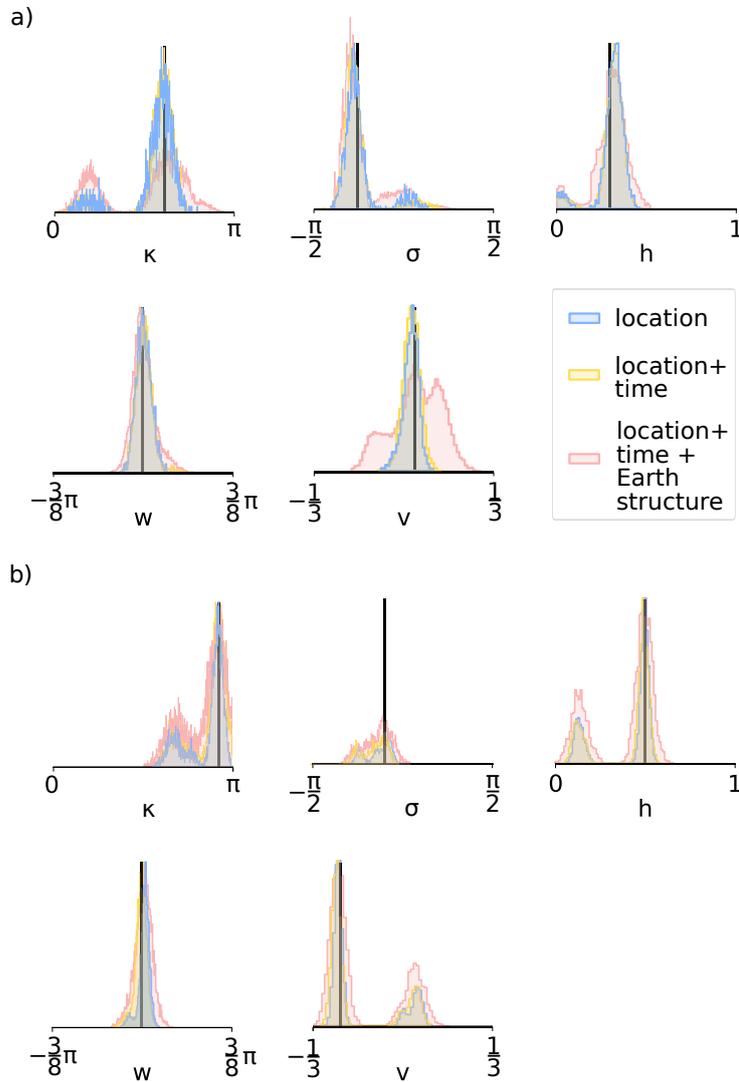


Figure 6: Normalized histograms of the ensemble of 6000 source parameter predictions for a) the Mw 3.8 earthquake on 2019/07/06 12:00:05 and b) the Mw 4.74 earthquake on 2019/07/26 00:42:48. Blue colors indicate the ensemble of predictions when using only the *Mojave* structure model and only considering error in location, yellow colors when using the *Mojave* structure model and considering error in location and timing and red colors the ensemble from all three considered Earth structure models and also considering errors in centroid location and timing.

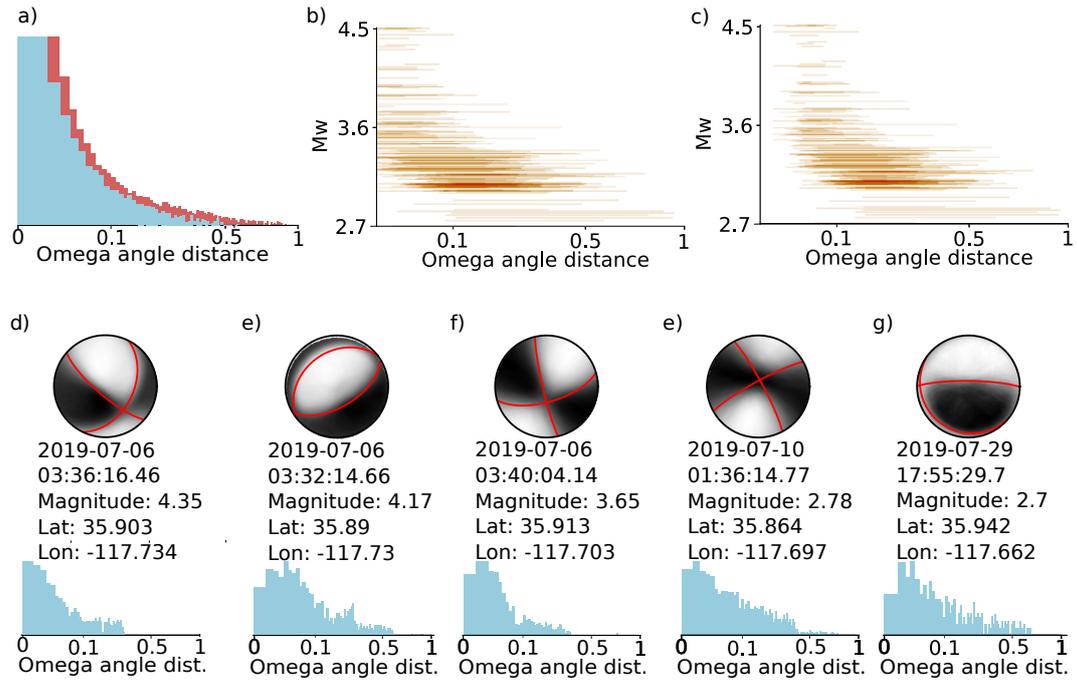


Figure 7: MT ensemble predictions for 198 pure double-couples focal mechanisms in the SCEDC catalog. a) omega angle distance as normalized histograms for v and w values set to zero in blue and in red with v and w left open, b) Density plot of lines drawn by the minimum and maximum omega angle distance d_{ω} for each earthquake magnitude as provided by the SCEDC catalog, with v and w set to zero, c) the same as in b) but with v and w left open, d) to g) show fuzzy beachballs, each based on 6000 moment tensor predictions, where the reference moment tensor (from SCEDC) is marked by red lines. Below: normalized histograms of the omega angle distances d_{ω} between the reference MT and the ensemble of MT predictions. Note that the x-axis for all subfigures scales quadratic.

348 Some of those systematic differences partly vanished by including also uncertainty in
349 centroid time into the variational inference scheme (Fig 4, g). As we estimate the full
350 seismic moment tensor the distribution and density of the non-DC components from
351 the predicted ensemble can be inferred (Supp. Fig. S3 and S4). The main regions of
352 high probability of solutions is consistent considering different sources of theory error.
353 However, larger uncertainties for both the CLVD as well as the isotropic components,
354 i.e. the lunge v and w parameters, can be observed when additionally considering
355 errors in Earth structure models (Supp. Fig. S3). It has been shown that an error in
356 Earth's structure is often compensated by increased CLVD and isotropic components
357 (Vasyura-Bathke et al., 2020).

358 We note that we evaluate the prediction accuracy of our framework by comparison
359 with SCEDC cataloged moment tensors. These solutions, however, could potentially
360 also be biased, deviating from the unknown "true" earthquake source. Variance reduction
361 could be used to estimate the precision with respect to the real waveform data.
362 We observe larger ω angle distances between the predicted MT ensemble and reference
363 MTs when considering the inferences from several Earth structure models (Fig4,a-h).
364 This is not unexpected, because the reference solutions are estimated with only one
365 of the Earth structure models. However, it is also possible that the "true" unknown
366 solution is better represented by our ensemble of predictions considering other Earth
367 structure models. In regions with well known structure this approach likely overesti-
368 mates the parameter uncertainties, but in regions with poorly known structure it might
369 provide a more realistic representations of parameter uncertainties (Vasyura-Bathke
370 et al., 2021).

371 The observation of a relation between spread of inferred parameter uncertainties
372 with magnitude is a result of parameter selections before learning, such as filter and
373 time window length, as well as decreasing signal-to-noise ratios for lower magnitudes.
374 Our considered filter frequencies are optimal for earthquakes with magnitudes Mw 3
375 to 4, of which hundreds occurred during the 2019 Ridgecrest sequence (Ross et al.,
376 2019). The station distribution around the Ridgecrest area and the good quality of the
377 waveform data due to mostly remote station locations is exceptional and together with
378 the statistically significant number of earthquakes this study area is bench-marking
379 showcase to demonstrate the robustness and performance of our approach. It remains
380 to be evaluated whether our approach performs equally well in areas with a sparse
381 station network under worse noise conditions.

382 The novelty of our proposed framework lies in the estimation of ensembles of
383 the full seismic MTs yielding uncertainties in parameter estimates based on seismic
384 waveforms. A shortcoming in our approach is the current limited transferability of
385 the trained BNNs to other study areas, unlike P-wave first motion polarity based
386 approaches (e.g. Ross et al., 2018). We assume that under operational conditions
387 on live incoming waveform data the prediction of the ensemble of full seismic MTs
388 using the presented framework can be done a few tens of seconds after the earthquake,
389 being almost near-realtime. Main factors that influence this response time are: 1) Our
390 algorithm considers a waveform window of 5 s. 2) The safe restitution of the waveform
391 data into displacement to avoid filter effects requires that at least several seconds of
392 data are available (around 2 s for the chosen frequencies in the case study). 3) In its
393 current form our approach requires the detection and location of an earthquake, which
394 can be used to infer a centroid time and optionally, its uncertainty as prior knowledge.
395 However, these can be delivered fast by other deep learning methods (Kriegerowski et
396 al., 2019). 4) Finally, the evaluation of the waveform data by a single trained BNN
397 takes a few hundred milliseconds and can be done in parallel for several BNNs at the
398 same time. Hence, approaches based on P-wave first motion polarity only (Hara et al.,
399 2019; Ross et al., 2018; Uchide, 2020) will likely outperform our proposed framework
400 in terms of response time. Nevertheless, these time factors are not of importance

401 for already cataloged data in a database, which can be searched fast by keeping the
402 recorded waveform data in memory.

403 In principle, the presented method can be made independent of the particular
404 station configuration at the expense of computational cost. This could be accom-
405 plished by calculating the synthetic waveforms for a distance-depth grid of locations
406 and shifting the source and receiver relatively or by assuming a location grid of ab-
407 stract receivers (van den Ende & Ampuero, 2020). The actual station locations can
408 then be mapped to such an abstract receiver grid by interpolation or nearest neigh-
409 bour. However, we do not expect that the framework could be made transferable to
410 other regions, because of the characteristics of the assumed Earth structure models
411 that are learned by the BNNs.

412 The choice of training a BNN for each considered grid point instead of training
413 a single large neural network with waveforms from all possible locations, such as in
414 Kuang et al. (2021), is a key point in our approach which allows us for estimating
415 MT parameter uncertainties considering uncertainty in centroid time and location as
416 well as uncertainty in Earth structure. Training a single large neural network with
417 waveforms from all potential source locations would require to estimate additionally
418 three location parameters (latitude, longitude and depth) as labels. This significantly
419 increases the non-linearity of the problem and, consequently increases the required
420 complexity of the neural network architecture, i.e. the number of trained filter weights
421 and biases. In our view, a simple network architecture with few trainable parameters
422 is favorable (Mignan & Broccardo, 2019) and, therefore, we chose to train multiple,
423 but individually rather simple networks.

424 As a by-product of our approach it turns out that our BNNs also learned to be
425 sensitive to the centroid location. Assuming that an earthquake occurred somewhere in
426 the grid of BNNs, each BNN can be queried to return the log-likelihoods for the input
427 data. The highest log-likelihoods should stem from BNNs learned for grid locations
428 close to the true centroid location. We test this assumption for a Mw 3.9 earthquake
429 included in the SCEDC catalog and indeed find a correlation of the log-likelihood
430 values with distance to the centroid location (Fig. 8). As prior information only the
431 centroid time and optionally its uncertainty is needed.

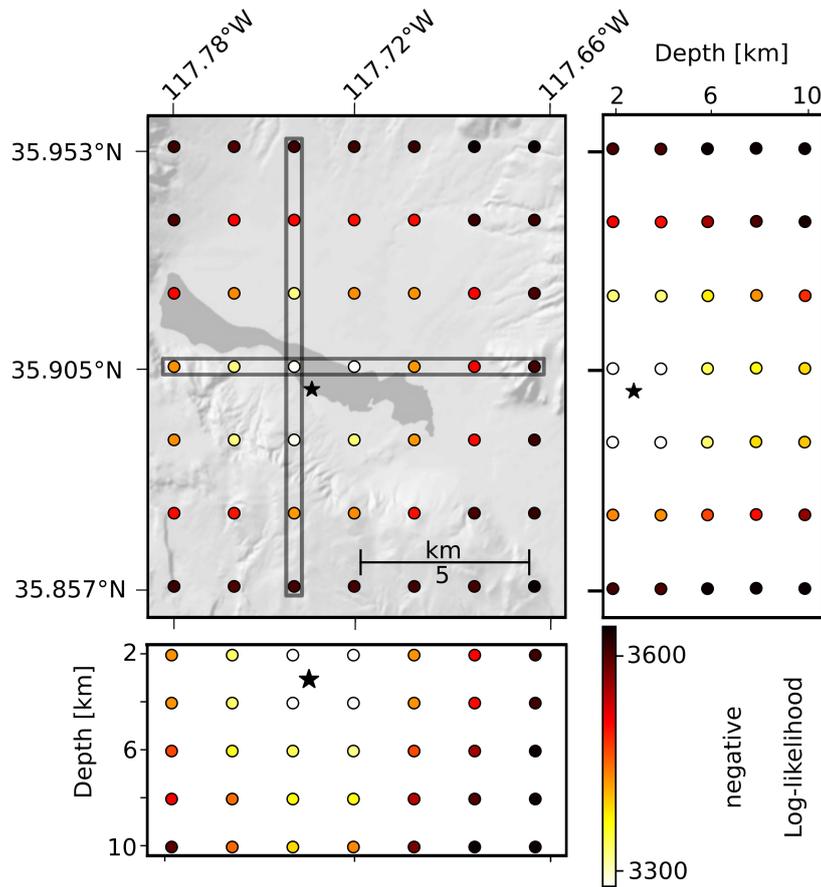


Figure 8: Earthquake centroid location inference. The grid points are colored by the negative log-likelihood values as inferred from evaluation of the BNNs for the real waveform data of the Mw 3.9 at 2019-07-06 17:59:15. The map view shows grid points at 4 km depth, whereas side views left and bottom show the grid-points at depth versus latitude and longitude along the profiles outlined with grey rectangles in the map view, respectively. The black star marks the centroid location as given by SCEDC for this earthquake.

5 Conclusions

We demonstrated that variational inference based on deep learning of Bayesian Neural Networks shows the capability to not only reproduce optimum parameter estimates of classical full moment tensor inversion, but it also yields uncertainties of the inferred MT parameters in near-real time. Our presented approach is flexible enough to optionally account for various cases of theory error that are well known to affect MT parameter estimates, i.e. errors in centroid location and time as well as errors in the assumed Earth structure.

The presented method has been successfully applied on local scale using field data of a subset of the Ridgecrest 2019 aftershock sequence, comprising 206 earthquakes with magnitudes Mw 2.7 to 5.5. The inferred ensembles of MT parameters have been compared to independently determined source mechanisms by the SCEDC.

One limitation of the presented approach is the non-transferable nature of the trained networks as they are trained for a specific Earth structure model, station setups, frequency filters and phase arrival time windows.

Our approach demonstrates the capabilities and the potential of machine learning for near-real time earthquake source mechanism estimation of small earthquakes with associated uncertainties. These are important information for hazard assessments and for providing other products to policy makers and public which are based on earthquake source analysis, e.g. shakemaps. The presented framework has the potential to be expanded upon and to be used in standardized automatic operational procedures.

Data availability statement

Data from regional seismometers are available via FDSN services from GEOFON and IRIS. The Caltech/USGS Southern California Seismic Network (SCSN) earthquake catalog, along with metadata and other ancillary data, such as moment tensors and focal mechanisms as been used and are available at <http://www.data.scec.org/index.html>. The Green's function stores used here are uploaded on Zenodo under DOI: 10.5281/zenodo.4643478 We make the code available and only use open-access waveform data for testing. We make the software available as jupyter notebook in the supplement and with pre-calculated example data as-well on Zenodo under DOI:10.5281/zenodo.4646666.

Acknowledgments

A.S was funded by the German Federal Ministry for Economic Affairs and Energy (BMWi) and was supervised by Project Management Jülich (PtJ) (grand number 03EE4003A). H.V-B was partially supported by Geo.X, the Research Network for Geosciences in Berlin and Potsdam under the project number SO_087_GeoX.

We thank Nima Nooshiri for helpful discussions. We thank the colleagues at SCEDC for maintaining open data policies and open access catalogs. We use Keras (Chollet & Others, 2015) and TensorFlow probability (Abadi et al., 2016; Dillon et al., 2017) to build our network architecture.

References

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., ... others (2016). Tensorflow: A system for large-scale machine learning. In *12th {USENIX} symposium on operating systems design and implementation ({OSDI} 16)* (pp. 265–283).
- Aki, K., & Richards, P. G. (2002). *Quantitative seismology*.

- 478 Bassin, C., Laske, G., & Masters, G. (2000). The current limits of resolution for sur-
479 face wave tomography in North America. *EOS Trans. AGU*, *81*(F897).
- 480 Blundell, C., Cornebise, J., Kavukcuoglu, K., & Wierstra, D. (2015). Weight Uncer-
481 tainty in Neural Networks. In *Proceedings of the 32nd international conference*
482 *on machine learning (icml)* (Vol. 37, pp. 1613–1622).
- 483 California Institute Of Technology And United States Geological Survey Pasadena.
484 (1926). *Southern california seismic network*. International Federation of
485 Digital Seismograph Networks. Retrieved from [http://www.fdsn.org/doi/](http://www.fdsn.org/doi/10.7914/SN/CI)
486 [10.7914/SN/CI](http://www.fdsn.org/doi/10.7914/SN/CI) doi: 10.7914/SN/CI
- 487 Cesca, S., Şen, A. T., & Dahm, T. (2014). Seismicity monitoring by cluster analysis
488 of moment tensors. *Geophysical Journal International*, *196*(3), 1813–1826.
- 489 Chollet, F., & Others. (2015). *Keras*. [\url{https://keras.io}](https://keras.io).
- 490 Dillon, J. V., Langmore, I., Tran, D., Brevdo, E., Vasudevan, S., Moore, D., ...
491 Saurous, R. A. (2017). Tensorflow distributions. *CoRR*, *abs/1711.10604*.
492 Retrieved from <http://arxiv.org/abs/1711.10604>
- 493 Duputel, Z., Rivera, L., Fukahata, Y., & Kanamori, H. (2012). Uncertainty estima-
494 tions for seismic source inversions. *Geophys. J. Int.*, *190*(2), 1243–1256. doi: 10
495 .1111/j.1365-246X.2012.05554.x
- 496 Ekström, G., Nettles, M., & Dziewoński, A. M. (2012). The global {CMT} project
497 2004–2010: {C}entroid-moment tensors for 13,017 earthquakes. *Physics of the*
498 *Earth and Planetary Interiors*, *200*, 1–9.
- 499 Glorot, X., Bordes, A., & Bengio, Y. (2011). Deep sparse rectifier neural networks.
500 In *Proceedings of the fourteenth international conference on artificial intelli-*
501 *gence and statistics* (pp. 315–323).
- 502 Graves, A. (2011). Practical variational inference for neural networks. *Advances in*
503 *Neural Information Processing Systems (NIPS)*, 2348–2356.
- 504 Hanka, W., & Kind, R. (1994). The {GEOFON} program. *Annals of Geophysics*,
505 *37*(5).
- 506 Hara, S., Fukahata, Y., & Iio, Y. (2019). P-wave first-motion polarity determi-
507 nation of waveform data in western japan using deep learning. *Earth, Planets*
508 *and Space*, *71*(1), 1–11.
- 509 Hauksson, E., & Unruh, J. (2007). Regional tectonics of the Coso geothermal
510 area along the intracontinental plate boundary in central eastern California:
511 Three-dimensional Vp and Vp/Vs models, spatial-temporal seismicity pat-
512 terns, and seismogenic deformation. *Journal of Geophysical Research: Solid*
513 *Earth*, *112*(B6).
- 514 Heimann, S., Kriegerowski, M., Isken, M., Cesca, S., Daout, S., Grigoli, F., ...
515 Dahm, T. (2017). Pyrocko - An open-source seismology toolbox and li-
516 brary. *GFZ Data Services*, *v. 0.3*. Retrieved from www.pyrocko.org doi:
517 <http://doi.org/10.5880/GFZ.2.1.2017.001>
- 518 Heimann, S., Vasyura-Bathke, H., Sudhaus, H., Isken, M. P., Kriegerowski, M.,
519 Steinberg, A., & Dahm, T. (2019). A Python framework for efficient use of
520 pre-computed Green’s functions in seismological and other physical forward
521 and inverse source problems. *Solid Earth*, *10*(6), 1921–1935.
- 522 Hutton, K., Woessner, J., & Hauksson, E. (2010). Earthquake monitoring in south-
523 ern California for seventy-seven years (1932–2008). *Bulletin of the Seismologi-*
524 *cal Society of America*, *100*(2), 423–446.
- 525 Ichinose, G. A., Anderson, J. G., Smith, K. D., & Zeng, Y. (2003). Source pa-
526 rameters of eastern california and western nevada earthquakes from regional
527 moment tensor inversion. *Bulletin of the Seismological Society of America*,
528 *93*(1), 61–84.
- 529 Jordan, T. H., & Maechling, P. J. (2003). The SCEC community modeling environ-
530 ment: An information infrastructure for system-level earthquake science. *Seis-*
531 *mological Research Letters*, *74*(3), 324–328.
- 532 Käüfl, P., Valentine, A. P., O’Toole, T. B., & Trampert, J. (2014). A framework for

- 533 fast probabilistic centroid-moment-tensor determination—inversion of regional
 534 static displacement measurements. *Geophysical Journal International*, 196(3),
 535 1676–1693.
- 536 Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv*
 537 *preprint arXiv:1412.6980*.
- 538 Kriegerowski, M., Petersen, G. M., Vasyura-Bathke, H., & Ohrnberger, M. (2019).
 539 A deep convolutional neural network for localization of clustered earthquakes
 540 based on multistation full waveforms. *Seismological Research Letters*, 90(2A),
 541 510–516.
- 542 Kuang, W., Yuan, C., & Zhang, J. (2021). Real-time determination of earthquake
 543 focal mechanism via deep learning. *Nature Communications*, 12(1), 1432. Re-
 544 trieved from <https://doi.org/10.1038/s41467-021-21670-x> doi: 10.1038/
 545 s41467-021-21670-x
- 546 Kühn, D., Heimann, S., Isken, M. P., Ruigrok, E., & Dost, B. (2020). Probabilistic
 547 moment tensor inversion for hydrocarbon-induced seismicity in the groningen
 548 gas field, the netherlands, part 1: Testing. *Bulletin of the Seismological Society*
 549 *of America*, 110(5), 2095–2111. doi: 10.1785/0120200099
- 550 Madariaga, R. (2007). Seismic source theory.
- 551 Mignan, A., & Broccardo, M. (2019). One neuron versus deep learning in aftershock
 552 prediction. *Nature*, 574(7776), E1–E3.
- 553 Monastero, F. C., Katzenstein, A. M., Miller, J. S., Unruh, J. R., Adams, M. C.,
 554 & Richards-Dinger, K. (2005). The Coso geothermal field: A nascent meta-
 555 morphic core complex. *Geological Society of America Bulletin*, 117(11-12),
 556 1534–1553.
- 557 Mousavi, S. M., Ellsworth, W. L., Zhu, W., Chuang, L. Y., & Beroza, G. C. (2020).
 558 Earthquake transformer—an attentive deep-learning model for simultaneous
 559 earthquake detection and phase picking. *Nature communications*, 11(1), 1–12.
- 560 Ross, Z. E., Idini, B., Jia, Z., Stephenson, O. L., Zhong, M., Wang, X., ... Others
 561 (2019). Hierarchical interlocked orthogonal faulting in the 2019 Ridgecrest
 562 earthquake sequence. *Science*, 366(6463), 346–351.
- 563 Ross, Z. E., Meier, M.-A., & Hauksson, E. (2018). P wave arrival picking and first-
 564 motion polarity determination with deep learning. *Journal of Geophysical Re-*
 565 *search: Solid Earth*, 123(6), 5120–5129.
- 566 (SCEDC), S. C. E. C. (2013). Southern california earthquake center. *Caltech.*
 567 *Dataset*.
- 568 Schoenball, M., Davatzes, N. C., & Glen, J. M. G. (2015). Differentiating induced
 569 and natural seismicity using space-time-magnitude statistics applied to the
 570 Coso Geothermal field. *Geophysical Research Letters*, 42(15), 6221–6228.
- 571 Smith, J. D., Ross, Z. E., Azizzadenesheli, K., & Muir, J. B. (2021). Hyposvi:
 572 Hypocenter inversion with stein variational inference and physics informed
 573 neural networks. *arXiv*.
- 574 Stähler, S. C., & Sigloch, K. (2014). Fully probabilistic seismic source inversion –
 575 Part 1 : Efficient parameterisation. *Solid Earth*, 5, 1055–1069. doi: 10.5194/
 576 se-5-1055-2014
- 577 Stähler, S. C., & Sigloch, K. (2016). Fully probabilistic seismic source inversion –
 578 Part 2 : Modelling errors and station covariances. *Solid Earth*(7), 1521–1536.
 579 doi: 10.5194/se-7-1521-2016
- 580 Tape, W., & Tape, C. (2012a). Angle between principal axis triples. *Geophysical*
 581 *Journal International*, 191(2), 813–831.
- 582 Tape, W., & Tape, C. (2012b). A geometric setting for moment tensors. *Geophysical*
 583 *Journal International*, 190(1), 476–498.
- 584 Tape, W., & Tape, C. (2015). A uniform parametrization of moment tensors. *Geo-*
 585 *physical Journal International*, 202(3), 2074–2081.
- 586 Uchide, T. (2020). Focal mechanisms of small earthquakes beneath the japanese is-
 587 lands based on first-motion polarities picked using deep learning. *Geophysical*

- 588 *Journal International*, 223(3), 1658–1671.
- 589 Vackář, J., Burjánek, J., Gallovič, F., Zahradník, J., & Clinton, J. (2017). Bayesian
 590 ISOLA: New tool for automated centroid moment tensor inversion. *Geophysical*
 591 *Journal International*, 210(2), 693–705. doi: 10.1093/gji/ggx158
- 592 van den Ende, M. P. A., & Ampuero, J.-P. (2020). Automated Seismic Source
 593 Characterization Using Deep Graph Neural Networks. *Geophysical Research*
 594 *Letters*, 47(17), e2020GL088690.
- 595 Vasyura-Bathke, H., Dettmer, J., Dutta, R., Mai, P. M., & Jónsson, S. (2021,
 596 jan). Accounting for theory errors with empirical Bayesian noise models in
 597 nonlinear centroid moment tensor estimation. *Geophysical Journal Inter-*
 598 *national*. Retrieved from <https://doi.org/10.1093/gji/ggab034> doi:
 599 10.1093/gji/ggab034
- 600 Vasyura-Bathke, H., Dettmer, J., Steinberg, A., Heimann, S., Isken, M. P., Zielke,
 601 O., . . . Jónsson, S. (2020, jan). The Bayesian Earthquake Analysis Tool.
 602 *Seismol. Res. Lett.*, 91(2A), 1003–1018. Retrieved from [https://doi.org/](https://doi.org/10.1785/0220190075)
 603 [10.1785/0220190075](https://doi.org/10.1785/0220190075) doi: 10.1785/0220190075
- 604 Wang, R. (1999). A simple orthonormalization method for stable and efficient com-
 605 putation of Green’s functions. *Bulletin of the Seismological Society of America*,
 606 89(3), 733–741.
- 607 Wen, Y., Vicol, P., Ba, J., Tran, D., & Grosse, R. (2018). Flipout: Efficient
 608 pseudo-independent weight perturbations on mini-batches. *arXiv preprint*
 609 *arXiv:1803.04386*.
- 610 Wu, H., & Lees, J. M. (1999). Three-dimensional P and S wave velocity structures
 611 of the Coso geothermal area, California, from microseismic travel time data.
 612 *Journal of Geophysical Research: Solid Earth*, 104(B6), 13217–13233.

6 Supplement