PROCEEDINGS, 45th Workshop on Geothermal Reservoir Engineering
Stanford University, Stanford, California, February 10-12, 2020
SGP-TR-216

Automatic Seismic Phase Picking Using Deep Learning for the EGS Collab project

Chengping Chai[1], Monica Maceira[1,2], Hector J. Santos-Villalobos[1], Singanallur V. Venkatakrishnan[1], Martin Schoenball[3], and EGS Collab Team[†]

[1]Oak Ridge National Laboratory, Oak Ridge, TN 37830

[2]Deparment of Physics and Astronomy, University of Tennessee, Knoxville, TN 37996

[3]Lawrence Berkeley National Laboratory, Berkeley, CA 94720

Email: chaic@ornl.gov, maceiram@ornl.gov, hsantos@ornl.gov, venkatakrisv@ornl.gov, schoenball@lbl.gov

Keywords: Transfer Learning, Deep Learning, seismic tomography, geothermal energy, EGS Collab

ABSTRACT

Microseismic monitoring plays an important role in many energy-related and environmental industries. The microseismic event catalog and seismic structure of the subsurface are two of the primary outputs of the microseismic monitoring system. Though rough locations of microseismic events can be estimated automatically, obtaining high-resolution microseismic event locations requires a significant amount of human labor especially on seismic phase picking. Unlike traditional automatic pickers that are usually less precise than human analysts, a few recently proposed algorithms based on deep neural networks (DNN) were able to match or surpass human performance for earthquake signals. Due to differences in the spatial scale of the study area, sensor sampling rate, and geometry of the monitoring system, it is not clear whether these deep neural network models can be used to speed up microseismic data processing. In this paper, we adapted the DNN based technique for automatic phase picking of microseismic signals. We used microseismic data recorded at the experiment 1 site of the enhanced geothermal system (EGS) Collab project and designed a workflow that we call transfer-learning aided double-difference tomography (TADT), that combines transfer learning and seismic tomography. We re-train an existing DNN with our data to obtain a new model using around 3500 seismograms and associated manual phase picks. This transfer learned model is able to reach human performance but much faster than human analysts. The transfer-learning-derived phase picks were used to improve microseismic event locations and image the subsurface. The results are similar to or slightly better than those obtained with manual phase picks.

1. INTRODUCTION

Micro-earthquake (MEQ) monitoring plays a significant role in the oil and gas industry, underground mines, carbon capture and storage, and geothermal industry due to both economic and safety concerns. Valuable information such as fracture development and elastic properties of the subsurface can be extracted from data recorded with the MEQ monitoring system. Spatial dimensions and temporal evolution of hydraulic and/or reactivated natural fractures are usually estimated with MEQ catalogs. The location and origin time of micro-seismic events are determined by arrival times of seismic phases at multiple seismic sensors. The arrival times of primary (P) and secondary (S) waves are also necessary for subsurface seismic imaging that measures elastic properties of the subsurface. Manually picking arrival times of seismic phases is a very time-consuming task especially for small-scale projects since high temporal sampling rate is required. Reliable automatic phase pickers are much needed, if not required, for these projects. Traditional automatic pickers such as Short-Term Average/Long-Term Average (STA/LTA; Allen, 1978) and Auto Regression-Akaike Information Criterion (AR-AIC;

[†]J. Ajo-Franklin, T. Baumgartner, K. Beckers, D. Blankenship, A. Bonneville, L. Boyd, S. Brown, J.A. Burghardt, C. Chai, Y. Chen, B. Chi, K. Condon, P.J. Cook, D. Crandall, P.F. Dobson, T. Doe, C.A. Doughty, D. Elsworth, J. Feldman, Z. Feng, A. Foris, L.P. Frash, Z. Frone, P. Fu, K. Gao, A. Ghassemi, Y. Guglielmi, B. Haimson, A. Hawkins, J. Heise, C. Hopp, M. Horn, R.N. Horne, J. Horner, M. Hu, H. Huang, L. Huang, K.J. Im, M. Ingraham, E. Jafarov, R.S. Jayne, S.E. Johnson, T.C. Johnson, B. Johnston, K. Kim, D.K. King, T. Kneafsey, H. Knox, J. Knox, D. Kumar, M. Lee, K. Li, Z. Li, M. Maceira, P. Mackey, N. Makedonska, E. Mattson, M.W. McClure, J. McLennan, C. Medler, R.J. Mellors, E. Metcalfe, J. Moore, C.E. Morency, J.P. Morris, T. Myers, S. Nakagawa, G. Neupane, G. Newman, A. Nieto, C.M. Oldenburg, T. Paronish, R. Pawar, P. Petrov, B. Pietzyk, R. Podgorney, Y. Polsky, J. Pope, S. Porse, J.C. Primo, C. Reimers, B.Q. Roberts, M. Robertson, W. Roggenthen, J. Rutqvist, D. Rynders, M. Schoenball, P. Schwering, V. Sesetty, C.S. Sherman, A. Singh, M.M. Smith, H. Sone, E.L. Sonnenthal, F.A. Soom, P. Sprinkle, C.E. Strickland, J. Su, D. Templeton, J.N. Thomle, V.R. Tribaldos, C. Ulrich, N. Uzunlar, A. Vachaparampil, C.A. Valladao, W. Vandermeer, G. Vandine, D. Vardiman, V.R. Vermeul, J.L. Wagoner, H.F. Wang, J. Weers, N. Welch, J. White, M.D. White, P. Winterfeld, T. Wood, S. Workman, H. Wu, Y.S. Wu, E.C. Yildirim, Y. Zhang, Y.Q. Zhang, Q. Zhou, M.D. Zoback

Sleeman and van Eck, 1999) pickers require intensive human involvement and refinement. When applied to MEQ, the accuracy of traditional automatic pickers is usually not satisfactory, especially for noisy data. Recent applications of deep-learning-based automatic seismic phase pickers (e.g. Ross et al., 2018; Zhu and Beroza, 2018; Zhu et al., 2019) have shown remarkable accuracy and processing speed for seismic signals originated from natural earthquakes. However, whether these deep learning phase pickers can be used for MEQ monitoring remains unclear.

We use MEQ data from experiment 1 of the enhanced geothermal system (EGS) Collab project to test whether one of the deep-learning-based automatic phase picker, PhaseNet (Zhu and Beroza, 2018) is applicable to mesoscale monitoring systems. The experiment was conducted at the 4850-foot level of the Sanford Underground Research Facility (SURF) located in Lead, South Dakota (Kneafsey et al., 2019). The testbed consists of one injection, one production, and six 60m-long monitoring boreholes. The MEQ monitoring system was equipped with multiple types of geophysical instruments including 24 hydrophones and 12 accelerometers. An 8-core workstation with an automated processing flow was deployed at the experiment site. The processing scripts are capable of detecting MEQ events (triggered), finding initial P-wave phase picks, and inverting for initial MEQ event locations and origin times. MEQ event locations and origin times were then improved with human reviewed and refined phase picks. The original MEQ catalog was processed using a uniform seismic velocity model (Schoenball et al., 2019b). Several hydraulic stimulations were performed since May 2018. We focused on seismic signals associated with stimulations between May 22 and December 21of 2018.

We applied the PhaseNet model and retrained it with a subset of MEQ data from experiment 1 of the EGS Collab project. The performance of the resulting transfer learned (TL) model was compared with a traditional automatic picker, the original PhaseNet model, and human analysts. The TL model was then applied to all the seismograms from the triggered MEQ events. The resulting TL-derived phase picks and a double-difference tomography package (tomoDD; Zhang and Thurber, 2006) were used to constrain subsurface seismic velocities and update MEQ event locations. The results were compared with those using manual picks.

2. DATA

Our data consist of seismograms from triggered micro-seismic events between May 2018 and December 2018 at the Experiment 1 site of the EGS Collab project, manually picked P-wave and S-wave arrival times, and the original microearthquake (MEQ) catalog from Schoenball et al. (2019a). We used 35 seismic sensors (one hydrophone is defective) with a 100 kHz sampling rate that were deployed in six 60m-long monitoring wells (see Figure 1). Twelve of these sensors are three-component accelerometers. The rest are single-component hydrophones. The micro-seismic events were detected and located using a standard STA/LTA routine, the PhasePAPy package (Chen and Holland, 2016), and a modified version of Hypoinverse (Klein, 2002). Triggered seismograms were cut to 0.11 s long segments around the P-wave arrival times and filtered with a bandpass filter between 3 kHz and 20 kHz. We used a total of 69444 waveform segments. P-wave arrival times were automatically measured with the PhasePAPy package. The MEQ monitoring system can automatically detect MEQ events and estimate MEQ locations. All S-wave arrival times were picked manually. The P-wave arrival times were then manually reviewed and refined to obtain MEQ locations with high precision. Additional details about the monitoring system and data preprocessing procedures of the original MEQ catalog can be found in Schoenball et al. (2019b).

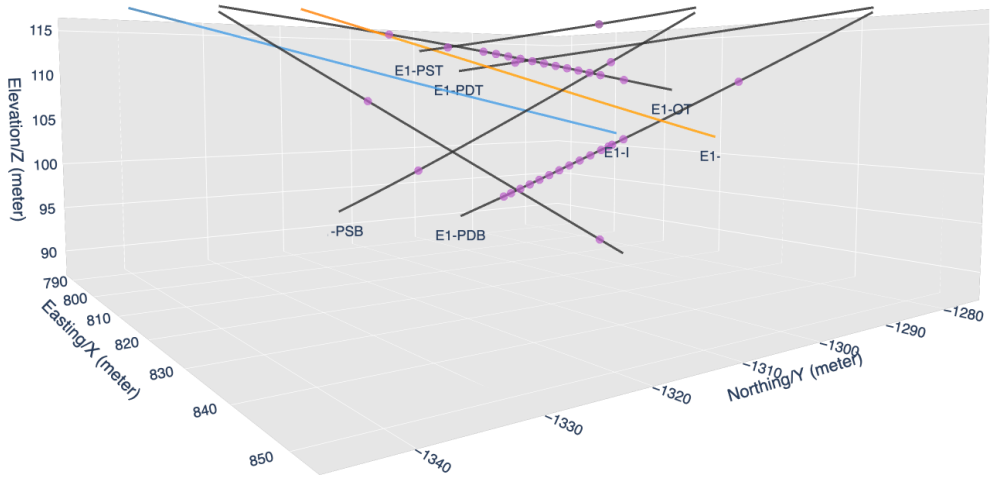

Figure 1: Seismic sensors (purple dots) and boreholes (lines). The blue line is the injection well, while the yellow line is the production well.

3. METHOD

We designed a workflow (Figure 2), transfer-learning aided double-difference tomography (TADT), that takes advantage of two existing technologies, deep neural networks (DNN) and seismic double-difference tomography. We started with the pre-trained DNN model - PhaseNet (Zhu and Beroza, 2018) - that is trained with over 0.8 million seismic recordings recorded in and around northern California for natural earthquakes. The PhaseNet model was trained using 90 seconds long seismograms sampled at 100 Hz. The earthquake-station distance for these data is on the order of kilometers. The monitoring system for our data samples at 100 kHz and the

source-sensor distance is on the order of meters. Despite the three orders of magnitude difference in both sampling rate and source-sensor distance between the PhaseNet data and our data, we found that the PhaseNet produced acceptable results when we applied to our data. To further improve the performance, the PhaseNet model was updated with a subset of seismic data (transfer learning) that meet the training data requirements (three-component seismograms with both P- and S-wave picks) for the PhaseNet. The resulting TL model was then applied to all the triggered seismograms (30 milliseconds long). We obtained transfer-learning derived P- and S-wave phase picks afterward. The tomoDD package (Zhang and Thurber, 2006) was used to update the MEQ catalog and image the subsurface simultaneously. The TADT workflow allows us to reduce the human cost significantly.

3.1 Transfer Learning

During the transfer learning process, we use the same network architecture as the PhaseNet and initialize the weights with the PhaseNet model. A total of 3821 three-component seismograms meet the training data requirements of PhaseNet. We visually inspected the seismograms and excluded 343 (9%) incorrect phase picks. The remaining 3478 seismograms were randomly divided into training, validation, and test data sets with 2443, 345, and 690 waveforms, respectively (Figure 3). The training set was used to retrain the DNN model. The validation set was used to select the optimal model from different training runs. The test set was used to evaluate performance. Similar to Zhu and Beroza (2018), we used a Gaussian distribution with a standard derivation of 0.1 milliseconds centered on the manual picks to represent manual pick uncertainty. The entire neural network was allowed to change during the transfer learning. Our tests indicate using data filtered in the frequency domain leads to better performance than raw data. A bandpass filter with corner frequencies of 3 kHz and 20 kHz was applied to the seismograms before feeding them into the neural network for training, which is different from Zhu and Beroza (2018). For a fair comparison, filtered (frequency domain) data were used throughout this study.

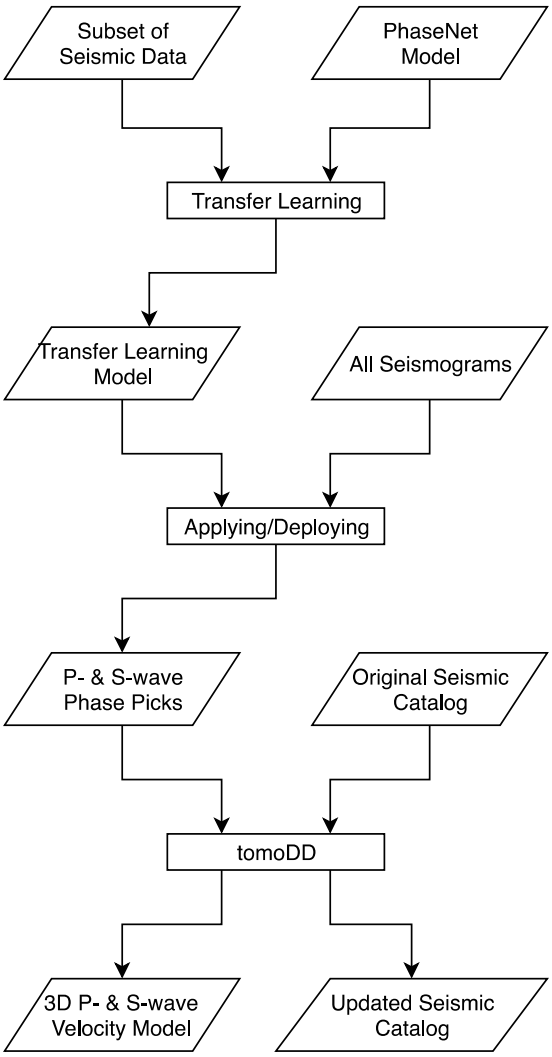

Figure 2: A flowchart of transfer-learning-aided seismic tomography using PhaseNet and tomoDD.

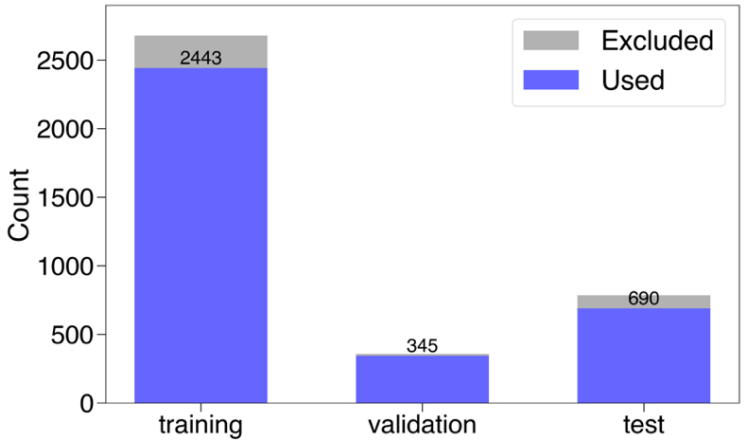

Figure 3: Number of waveform sets (three waveforms per set) used for transfer learning.

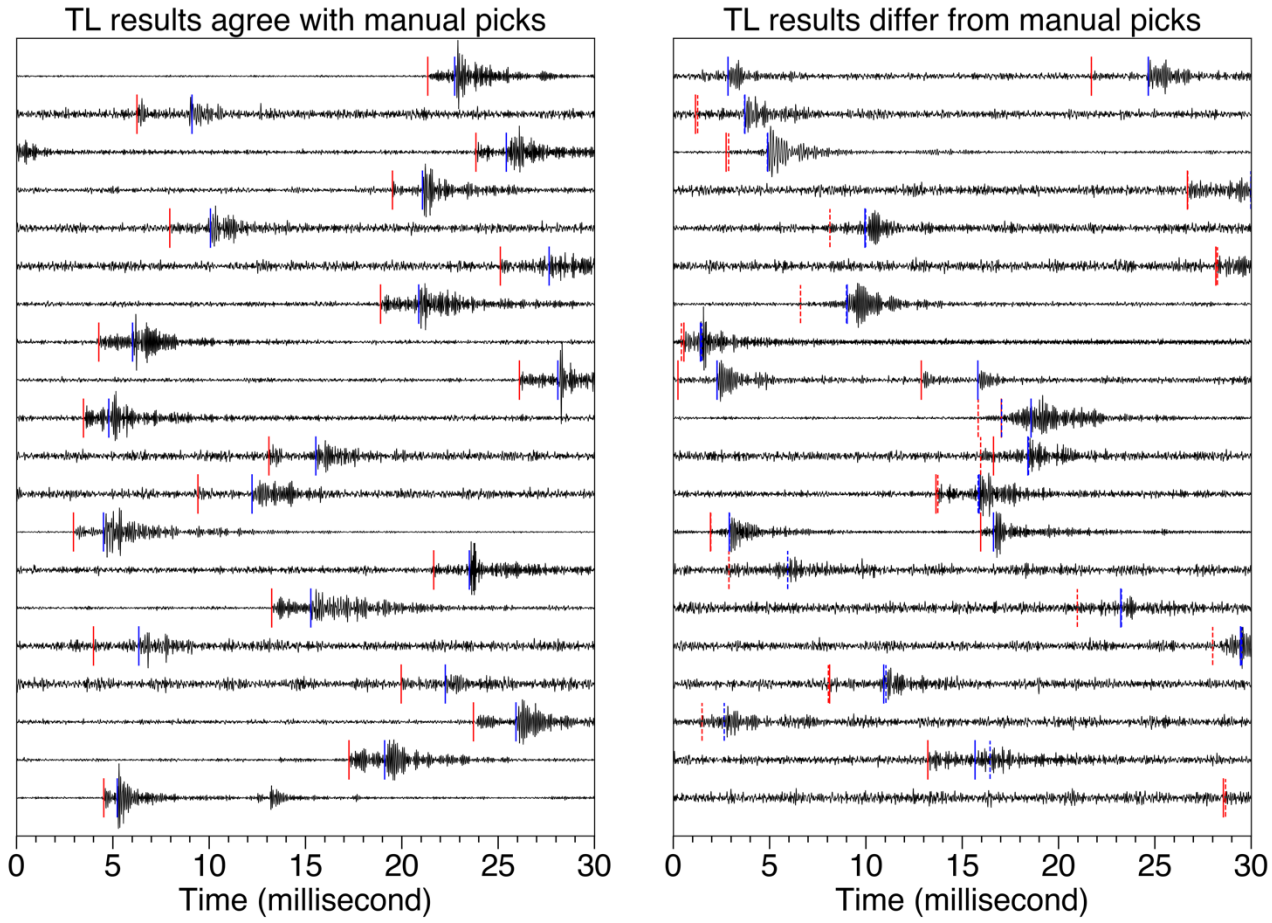

Figure 4: Randomly selected example waveforms and transfer-learning-derived (TL) phase picks that agree with (left) and differ from (right) the manual picks. Only one component seismogram is shown here, but three component seismograms were used for phase picking. Red vertical lines represent P-wave picks. Blue vertical lines are used for S-waves. Solid vertical lines are transfer-learning-derived picks. Dashed lines are manual picks. When the temporal difference between TL results and manual picks are larger than 0.1 millisecond, we consider the results to be different.

3.2 Double Difference Tomography

The original MEQ catalog was processed with a uniform seismic velocity model (Schoenball et al., 2019b). We used the double-difference tomography package tomoDD (Zhang and Thurber, 2006) to reduce the MEQ location uncertainty due to spatial seismic velocity variations and constrain the 3D subsurface P-wave and S-wave velocity models for the seismically active region. Since the tomoDD package was originally designed for kilometer-scale problems, we made some modifications (e.g. input and output format, coordinate system) specifically for meter-scale projects. Both P- and S-wave seismic velocity models are constrained. We relocated around 1800 MEQ events. A 3D volume of 77 m (easting), 83 m (northing), and 40 m (vertical) was discretized with 1 m^3 cubes. The tomography started with a uniform seismic model that has a P-wave speed of 5.9 km/s and an S-wave speed of 3.5 km/s. These two velocities were obtained from curve fitting of travel-time observations (travel-time versus distance). The final velocity models and updated MEQ catalog were obtained after eight iterations. Numerous previous studies (e.g. Syracuse et al., 2016; Chai et al., 2019) have shown that appropriate inversion parameters are required for a well-constrained seismic velocity model. We used an L-curve analysis (Hansen, 1992) to find the optimal set of inversion parameters. An optimal weight of 10 was used for smoothing and 200 for damping.

4. RESULTS

Our results consist of transfer-learning-derived phase picks, updated MEQ locations, and 3D seismic velocity models. Hyper Text Markup Language (HTML) based visualizations (similar to Chai et al., 2018) were used to inspect MEQ locations and seismic velocity models.

4.1 Phase Picks

The TL model we obtained is able to find phase picks from seismograms with high accuracy. Some randomly selected waveforms and associated phase picks from the test dataset are shown in Figure 4. We can see that transfer learning results agree with manual picks even when the background noise level is high. Inspecting the data when transfer learning results differ from manual picks, we noticed that the TL model is able to correct some human errors or skip difficult-to-pick signals (more often for P-waves than for S-waves). The difference between transfer learning results and manual picks is just slightly larger than the threshold (0.1 milliseconds), for many cases in Figure 4. The transfer learning model is more prone to error when signals are very complex or when the signal to noise ratio (SNR) is small.

We compared the transfer learning results with those using the Obspy (Beyreuther et al., 2010) implementation of AR picker (Akazawa, 2004), the original PhaseNet, and human analysts (Figure 5 for P-waves and Figure 6 for S-waves). We use precision and recall to quantify and compare the performance. Precision is the ratio of the number of correctly predicted phase picks over the total of predicted phase picks. Recall is the ratio of the number of correctly predicted phase picks over the sum of the number of correctly predict phase picks and the number of manual picks missed. The performance of human analysts was estimated by having three analysts manually pick the phase arrival times from the same 100 three-component seismograms. For each seismogram, the median of the three manual picks is considered the ground truth. The human accuracy for each analyst is measured by comparing results from each analyst against the ground truth. The human accuracy in Figure 5 and Figure 6 was the average of the three analysts. Since we did not have enough seismograms with no P-waves and S-waves, the recall was not computed for human analysts. The original PhaseNet produced much better results than the AR picker for both P- and S-waves. The TL model outperformed the original PhaseNet with an improvement of roughly 0.1 in precision and 0.3 in recall, highlighting the importance of re-training the DNN with our data. The TL model is the only one among the three automatic pickers that has a performance comparable to human analysts. The TL model performed slightly better on S-waves than human analysts, which could be due to larger SNRs comparing to P-waves. The TL model is able to achieve human performance in a fraction of the time.

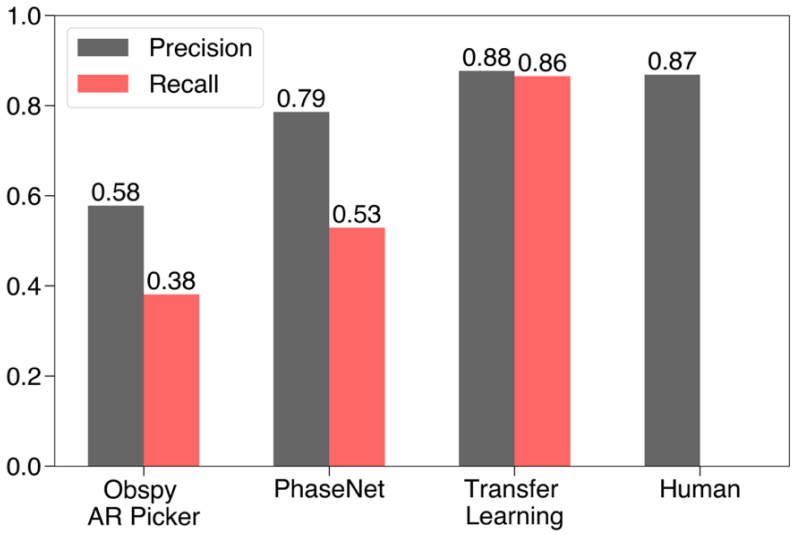

Figure 5: A comparison of accuracy between human (three analysts), Obspy AR picker (www.obspy.org), directly applying PhaseNet, and the transfer learning derived model for P-waves.

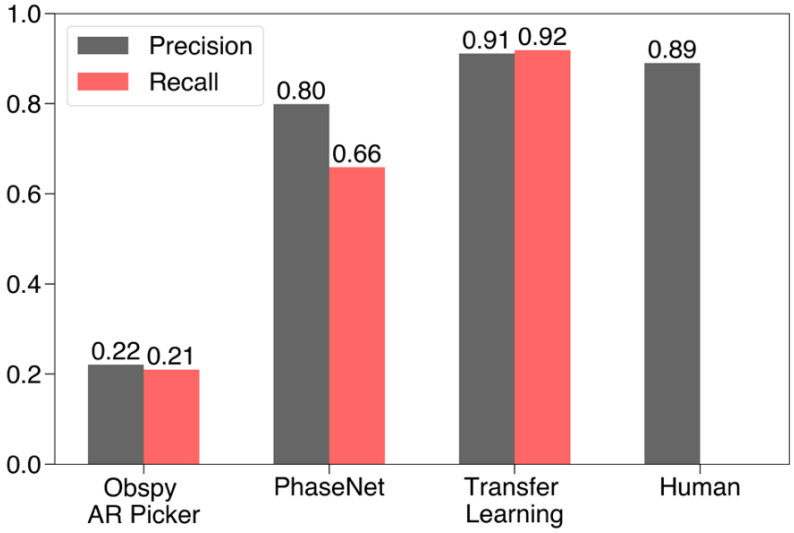

Figure 6: A comparison of accuracy between human (three analysts), Obspy AR picker (www.obspy.org), directly applying PhaseNet, and the transfer learning derived model for S-waves.

When we apply the TL model to all the triggered seismograms, the TL model was able to find more S-wave picks than the human expert. We performed 3D double-difference tomography using manual picks and transfer-learning derived picks with the same inversion parameters. Though fewer P-wave picks were obtained by the TL model comparing to the expert, the updated MEQ locations using transfer-learning-derived picks show tighter patterns compared to that using manual picks (see the next section for details). Specifically, we found 18543 acceptable P-wave picks and 8935 S-wave picks from the human expert using a total of 69444 seismograms. The transfer-learning model was able to identify 12050 acceptable P-wave picks (20% of which are included in the training dataset) and 13297 S-wave picks (18% of which were included in training dataset).

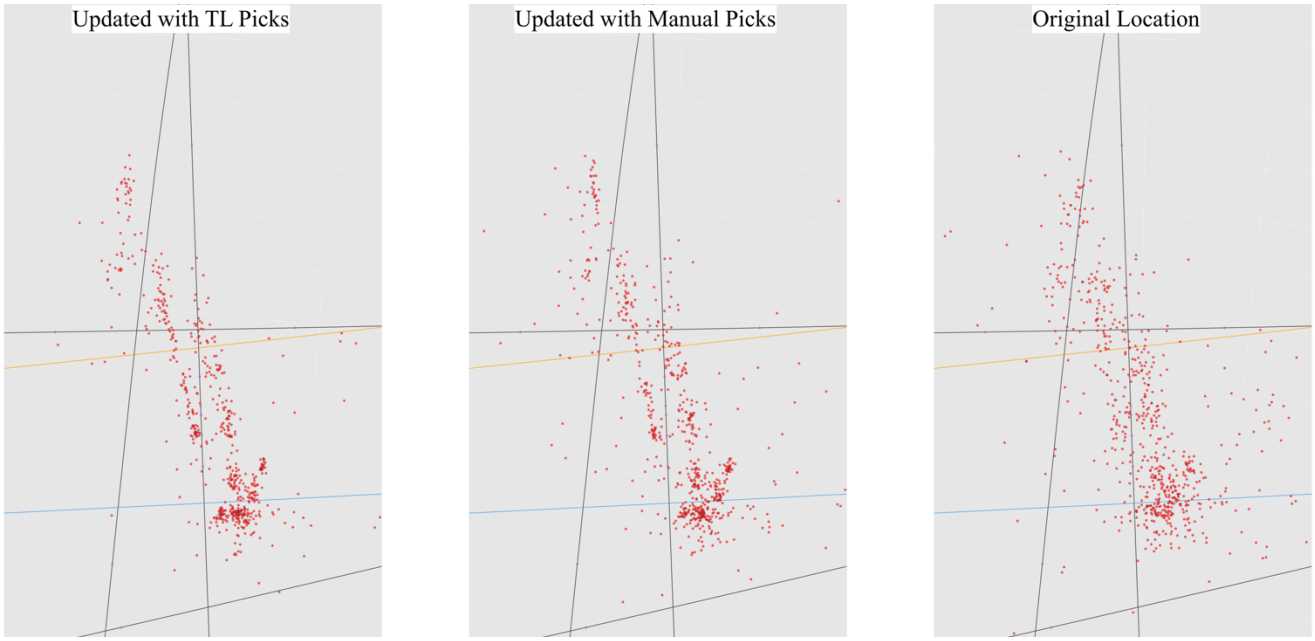

Figure 7: A comparison of original (right, using a uniform seismic velocity model) and updated (left and middle, using data-constrained 3D seismic velocity models) microseismic event locations associated with the May 2018 stimulations. Seismic events in the left panel show locations updated with transfer-learning-derived phase picks (using tomoDD). The middle panel shows locations updated with manual picks (using tomoDD).

4.2 Updated MEQ Location

We examine MEQ locations associated with the May 2018 stimulations, June 2018 stimulation, and December 2018 stimulations in Figure 7, Figure 8, and Figure 9, respectively. Comparing the original MEQ locations, the updated locations from the double-difference tomography using either manual picks or the transfer-learning-derived picks show more detailed geometry of the fractures. For the May 2018 stimulations (Figure 7), the update MEQ locations show two parallel fractures that are not clear in the original locations. Since these two fractures ruptured into one monitoring borehole, we were able to confirm these fractures with independent temperature data recorded in the borehole. Using the transfer-learning-derived picks leads to tighter clustered MEQ locations. As for the June 2018 stimulation (Figure 8), both the original and updated MEQ locations show two fractures. The updated locations show a slightly tighter pattern compared to the original. As for the December 2018 stimulations (Figure 9), the original locations show two intersecting fractures, but the geometry of these fractures was not well constrained especially near the two ends of the fractures. When we updated the MEQ locations with manual picks, these two fractures showed a tighter pattern. When transfer-learning-derived picks were used to update the MEQ locations, these two fractures were constrained even tightly. We can see the two ends of the fractures more clearly. So, the transfer-learning-derived picks are equivalent to or better than the manual picks as indicated by the updated MEQ locations.

Figure 8: Same to Figure 7 but for microseismic events associated with the June 2018 stimulation.

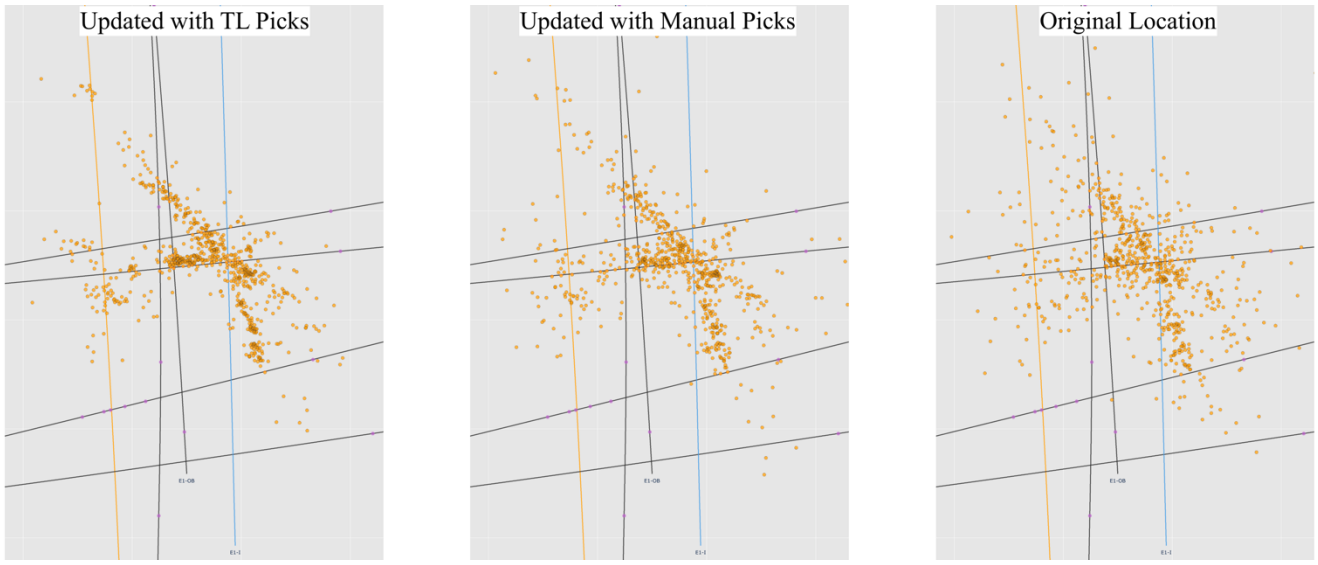

Figure 9: Same to Figure 7 but for microseismic events associated with the December 2018 stimulations.

4.3 3D Seismic Velocity Model

We constrained both P-wave and S-wave seismic velocities since both P- and S-wave picks were used in the double-difference tomography. The P-wave velocity model shows significant spatial heterogeneity when either manual picks or transfer-learning-derived picks were used in the tomography. A slice of the 3D P-wave and S-wave velocity models, obtained using transfer-learning-derived picks, is shown in Figure 10 and Figure 11, respectively. The P-wave velocity model contains a couple of small-scale high-velocity anomalies. To the first order, the P-wave velocity is slower at a lower elevation compared to a higher elevation. The S-wave velocity model shows a clearer pattern with a low-velocity zone imaged at an elevation below 105 meters.

To identify the volume that we can reliably image, we performed checkerboard tests using data simulated according to both manual picks and transfer-learning-derived picks. For the checkerboard tests, we started with an artificial model with alternating high and low velocities. Synthetic P-wave and S-wave phase picks were computed using the artificial model when we have observations between MEQ event and seismic sensor pairs. The synthetic phase picks were then used in tomography with a uniform starting model. The recover/inverted model is compared to the true model to identify the volume that is well-constrained by data. To measure the volume, we first compute the absolute difference between the recovered model and the true model at each grid. A grid cell is considered well-constrained when the recovered seismic velocities are less than 0.1 km/s away for P waves or 0.06 km/s for S waves from the ground truth. The well-constrained volume is smoothed by applying a spatial Gaussian filter to all of the well-constrained grid. Slices of the recovered S-wave velocity models using manual picks and transfer-learning-derived picks are shown in Figure 12 and Figure 13, respectively. For the P-wave velocity model, the well-constrained volume is 2678 m^3 for manual picks and 2465 m^3 (8% decrease) for

transfer-learning-derived picks. For the S-wave velocity model, the well-constrained volume is 815 m^3 for manual picks and 1895 m^3 (133% increase) for transfer-learning-derived picks.

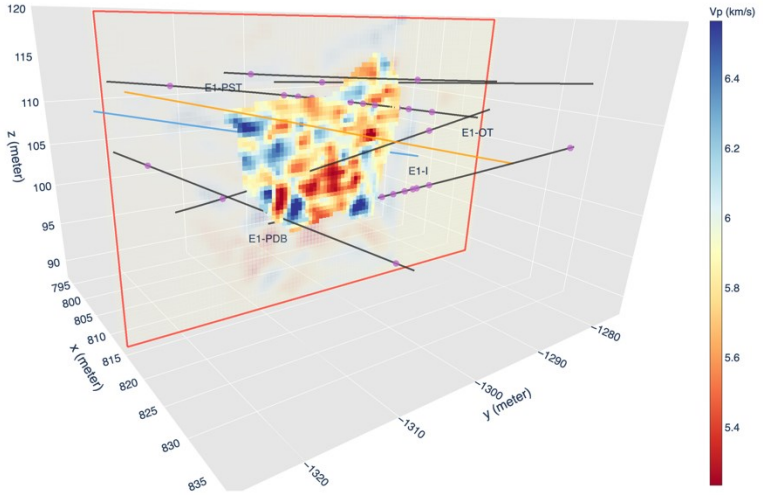

Figure 10: A slice of the 3D P-wave velocity model using transfer-learning-derived picks. The highlighted regions are well constrained as indicated by synthetic tests.

5. DISCUSSION AND CONCLUSIONS

We present a workflow that integrates transfer learning and seismic double-difference tomography. As demonstrated with the EGS Collab data, the workflow can produce better MEQ locations, improve subsurface imaging capabilities, and reduce the overall time cost compared to the original labor-intensive workflow. Our results also show that the TL model obtained by retraining the PhaseNet deep neural network leads to human-level performance despite the significant differences in the study area, sensor geometry, and sampling rate between the PhaseNet data and our data. Other types of geophysical observations (e.g., Chai et al., 2015; Syracuse, et al., 2016) can also be included in this workflow.

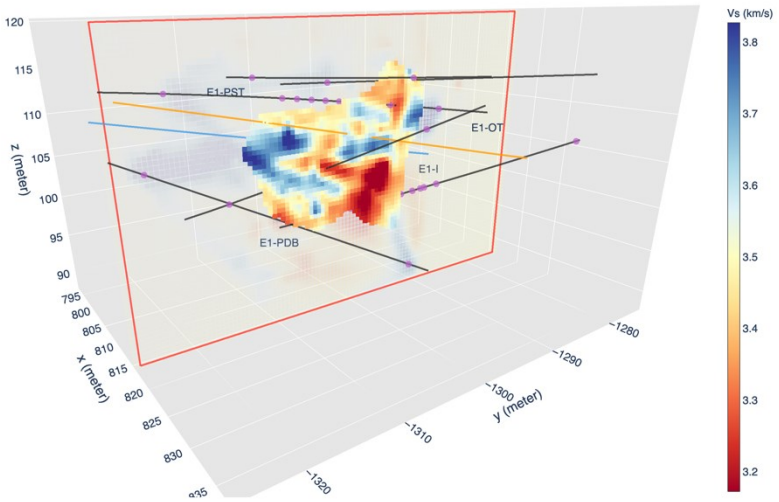

Figure 11: A slice of the 3D S-wave velocity model using transfer-learning-derived picks. The highlighted regions are well constrained as indicated by synthetic tests.

Since phase picks are the bases for both locating MEQ events and imaging the subsurface, it is valuable to find seismic phase picks quickly and reliably. The PhaseNet model leads to better picks than the Obspy implementation of AR picker. A TL model initialized with the PhaseNet model and retrained with only around 2400 three-component seismograms and associated manual picks outperform the original PhaseNet model by around 10% in terms of precision. The TL model performs equally or slightly better than a human expert. The TL model found fewer (32%) P-wave picks but more (48%) S-wave picks than the human expert. Since the double-difference tomography results that used these transfer-learning-derived phase picks show better MEQ locations comparing to those using manual picks, it is likely the TL model removed low-quality P-wave picks and added high-quality S-wave picks. The speed of the TL model (or PhaseNet) is about 1900 times (excluded training time) faster than the human expert.

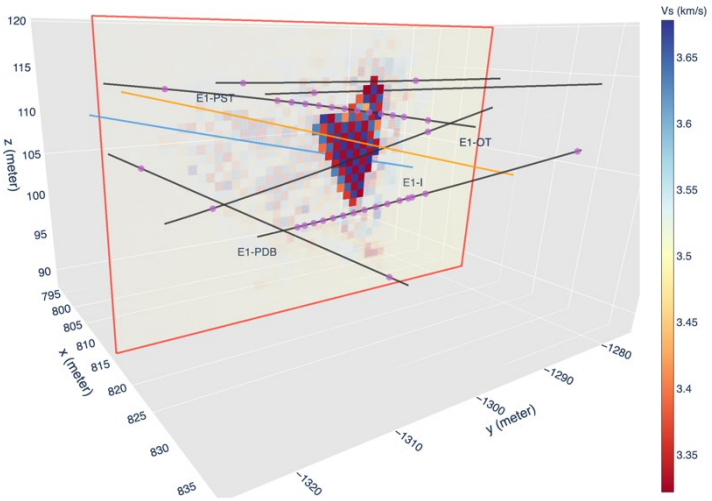

Figure 12: Checkerboard tests using manual phase picks for S-wave velocities.

Double-difference tomography tests using manual picks and transfer-learning derived picks show that transfer-learning derived picks lead to better MEQ locations and a larger (133% increase) well-constrained volume for the S-wave velocity model. Even though we obtained fewer P-wave picks with the transfer learning model compared to the human expert, the well-constrained volume for the P-wave velocity model only decreased slightly. Given the MEQ locations are better when transfer-learning derived picks are used, the transfer-learning model is likely excluded many low-quality P-wave phase picks. The improved MEQ locations allow us to see detailed structures of the fracture planes, which in turn will help us better constrain the fracture geometry. Two parallel fracture planes were confirmed with independent borehole observations.

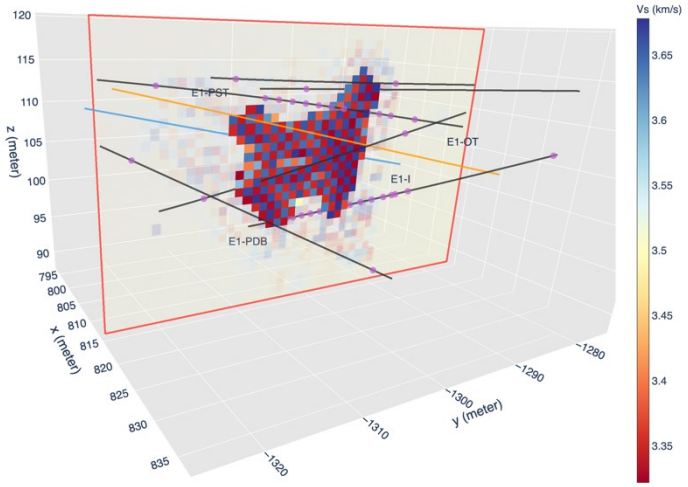

Figure 13: Checkerboard tests using transfer-learning-derived phase picks for S-wave velocities. The results show the TL model is able to pick more S-wave arrival times than a human expert.

Our results show that we can reduce the time cost significantly by adding transfer learning into the proposed workflow. The seismic phase picking is labor intensive and thus expensive. It took on the order of several days to find all the seismic phase picks from the 69444 seismograms recorded. For the presented workflow, the analyst would only need to manually pick around 3500 high-quality seismograms. Retraining the PhaseNet model takes ~5 hours (hardware dependent). Processing all the seismograms with the transfer learning model takes only 9 minutes on a laptop computer (with six 2.9 GHz Intel i9 cores). Even including the retraining time, the presented workflow takes much less time than human labor. The speed can be increased with larger computational power. Moreover, the TL model can be directly used on future seismic data from the same recording system without retraining. The proposed workflow is an economical way to monitor subsurface fracture evolvement and image subsurface seismic structure with high resolution. The workflow is applicable to new study areas.

ACKNOWLEDGEMENTS

This material was based upon work supported by the U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy (EERE), Office of Technology Development, Geothermal Technologies Office, under Award Number DE-AC05-00OR22725. The United States Government retains, and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes. The research supporting this work took place in whole or in part at the Sanford Underground Research Facility in Lead, South Dakota. The assistance of the Sanford Underground Research Facility and its personnel in providing physical access and general logistical and technical support is acknowledged. We acknowledge Weiqiang Zhu and Gregory C. Beroza for sharing PhaseNet. We also thank Haijiang Zhang and Clifford Thurber for sharing tomoDD.

REFERENCES

Allen, R. V., Automatic earthquake recognition and timing from single traces, Bulletin of the Seismological Society of America, **68**, no. 5, (1978), 1521–1532.

Akazawa, T., A technique for automatic detection of onset time of P-and S-Phases in strong motion records, 13th World Conference on Earthquake Engineering, 786, (2004).

Beyreuther, M., Barsch, R., Krischer, L., Megies, T., Behr, Y., and Wassermann, J., ObsPy: A Python Toolbox for Seismology. Seismological Research Letters, **81**, no. 3, (2010), 530–533. https://doi.org/10.1785/gssrl.81.3.530

Chai, C., Ammon, C.J., Maceira, M., and Herrmann, R.B., Inverting interpolated receiver functions with surface wave dispersion and gravity: Application to the western U.S. and adjacent Canada and Mexico, Geophys. Res. Lett., **42**, no. 11, (2015), 4359–4366, doi: 10.1002/2015GL063733.

Chai, C., Maceira, M., Santos-Villalobos, H.J., and EGS Collab Team, Subsurface seismic structure around the Sanford Underground Research Facility, in PROCEEDINGS: 44th Workshop on Geothermal Reservoir Engineering, Stanford University, Stanford, California, CA, (2019), 1365-1376.

Chai, C., Ammon, C.J., Maceira, M., and Herrmann, R.B., Interactive Visualization of Complex Seismic Data and Models Using Bokeh, Seismol. Res. Lett., **89**, no. 2A, (2018), 668–676, doi: 10.1785/0220170132.

Chen, C., and Holland, A. A. PhasePApy: A robust pure Python package for automatic identification of seismic phases, Seismol. Res. Lett., **87,** 6, (2016), 1384-1396, doi: 10.1785/0220160019.

Hansen, P.C., Analysis of Discrete Ill-Posed Problems by Means of the L-Curve, SIAM Rev., **34**, no. 4, (1992), 561–580, doi: 10.1137/1034115.

Klein, F. W., User's Guide to HYPOINVERSE-2000, a Fortran Program to Solve for Earthquake Locations and Magnitudes, U.S. Geol. Surv. Open File Rep. 02-171, (2002), doi: http://geopubs.wr.usgs.gov/open-file/of02-171/.

Kneafsey, T. J., Blankenship, D., Knox, H. A., Johnson, T. C., Ajo-Franklin, J., Schwering, P. C., et al., EGS Collab Project: Status and Progress, in Proceedings 44th Workshop on Geothermal Reservoir Engineering, Stanford University, Stanford, CA, (2019).

Ross, Z. E., Meier, M.-A., and Hauksson, E. P Wave Arrival Picking and First-Motion Polarity Determination With Deep Learning. Journal of Geophysical Research: Solid Earth, **123**, no. 6, (2018), 5120–5129. https://doi.org/10.1029/2017JB015251

Schoenball, M., Ajo-Franklin, J., Robertson, M., Wood, T., Blankenship, D., Cook, P., Dobson, P., Guglielmi, Y., Fu, P., Kneafsey, T., Knox, H., Petrov, P., Schwering, P., Rempleton, D., and Ulrich, C.. EGS Collab Experiment 1: Microseismic Monitoring. United States: N. p., (2019a), Web. doi:10.15121/1557417.

Schoenball, M., Ajo-Franklin, J., Blankenship, D., Cook, P., Dobson, P., Guglielmi, Y., Fu, P., et al. Microseismic monitoring of meso-scale stimulations for the DOE EGS Collab project at the Sanford Underground Research Facility, in Proceedings: 44th Workshop on Geothermal Reservoir Engineering, Stanford University, Stanford, CA, (2019b).

Sleeman, R., & van Eck, T. (1999). Robust automatic P-phase picking: an on-line implementation in the analysis of broadband seismogram recordings. Physics of the Earth and Planetary Interiors, 113(1–4), 265–275. https://doi.org/10.1016/S0031-9201(99)00007-2

Syracuse, E. M., Maceira, M., Prieto, G.A., Zhang, H., and Ammon, C.J., Multiple plates subducting beneath Colombia, as illuminated by seismicity and velocity from the joint inversion of seismic and gravity data, Earth Planet. Sci. Lett., **444**, (2016), 139–149, doi: 10.1016/j.epsl.2016.03.050.

Zhang, H., and Thurber, C., Development and Applications of Double-difference Seismic Tomography, Pure Appl. Geophys., **163**, no. 2–3, 373–403, (2006), doi: 10.1007/s00024-005-0021-y.

Zhu, W., and Beroza, G.C., PhaseNet: A Deep-Neural-Network-Based Seismic Arrival Time Picking Method, Geophys. J. Int., **216**, no. 1, 261–273, (2018), doi: 10.1093/gji/ggy423.

Zhu, L., Peng, Z., McClellan, J., Li, C., Yao, D., Li, Z., and Fang, L. Deep learning for seismic phase detection and picking in the aftershock zone of 2008 M7.9 Wenchuan Earthquake. Physics of the Earth and Planetary Interiors, **293**, May 2018, (2019),106261. https://doi.org/10.1016/j.pepi.2019.05.004