

One neuron versus deep learning in aftershock prediction

Arnaud Mignan^{1,2,3*} & Marco Broccardo^{2,4*}

ARISING FROM P. M. R. DeVries et al. *Nature* <https://doi.org/10.1038/s41586-018-0438-y> (2018)

Forecasting the spatial distribution of aftershocks in the aftermath of large seismic events is of great importance for improving both our understanding of earthquake triggering and post-disaster management. Recently, DeVries et al.¹ attempted to solve this scientific problem by deep learning. Using the area under the curve (AUC) of receiver operating characteristic curves, the authors showed that a deep neural network (DNN) considerably outperformed classical Coulomb stress. Here we first clarify that similar performances had already been obtained (by the same authors, in 2017)² for various scalar stress metrics, suggesting that deep learning does not actually improve prediction. Second, we reformulate the 2017 results² using two-parameter logistic regression (that is, one neuron) and obtain the same performance as that of the 13,451-parameter DNN. We further show that a logistic regression based on the measured distance and mainshock average slip (instead of derived stresses) performs better than the DNN. This demonstrates

that so far the proposed deep learning strategy does not provide any new insight (predictive or inferential) in this domain.

Operational aftershock forecasting has been possible for decades thanks to well established empirical laws^{3–5}. Spatial patterns of aftershocks are often described as a power-law decay^{5–8}. Physical models based on the Coulomb stress paradigm only outperform statistical methods when considered in physical/statistical hybrids with high-quality mainshock rupture data⁹. Coulomb stress models can on their own be as performant as statistical methods when additionally including receiver fault heterogeneities^{10,11}. Meade et al.² performed a thorough analysis of various scalar stress metrics and showed that several of them outperform classic Coulomb failure stress. We are here concerned with their follow-up article¹, which presented similar results, but via deep learning. In this study, we aim to demonstrate that while defining larger and deeper DNNs usually does not hurt model

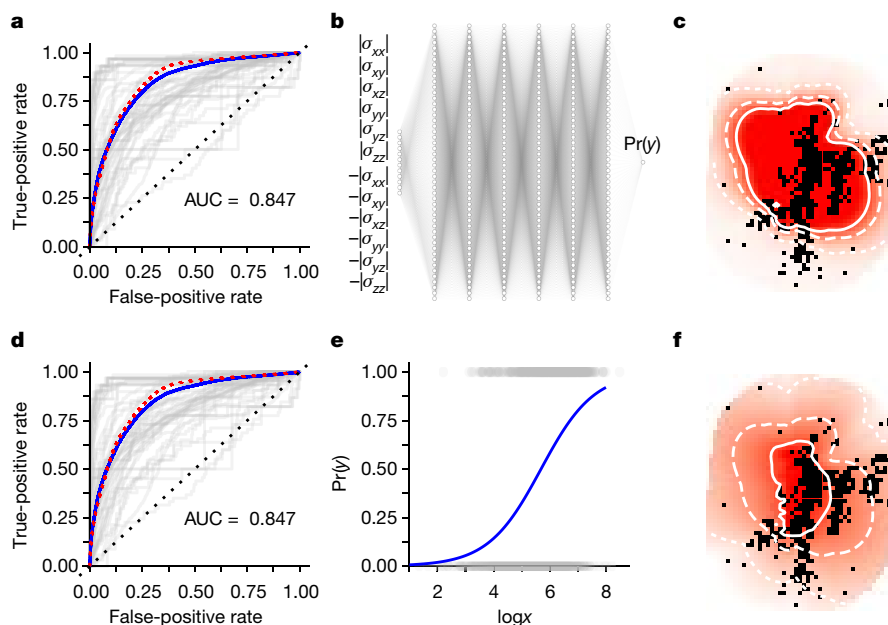


Fig. 1 | Prediction of aftershock spatial patterns based on stress features. **a**, Receiver operating characteristic (ROC) curves of test data for the DNN of ref. ¹. **b**, DNN topology; plot generated with NN-SVG (<http://alexlenail.me/NN-SVG/>). **c**, Example of DNN prediction (1999 ChiChi aftershocks). **d**, Test-data ROC curves for logistic regression with the logarithm of the sum of absolute stress components as input.

e, Logistic regression fit on training data. **f**, Example of logistic regression prediction. For both DNN and logistic regression models, 58 ROC curves are shown, for the 57 mainshocks from the test set (grey) and all combined mainshock–aftershock pairs (blue; see legend of Fig. 2 for red colour explanation). Dotted, dashed and solid curves in **c** and **f** represent $\Pr(y) = 0.3, 0.5$ and 0.7 , respectively.

¹Swiss Federal Institute of Technology, Institute of Geophysics, Zurich, Switzerland. ²Swiss Seismological Service, Zurich, Switzerland. ³Institute of Risk Analysis, Prediction and Management, Academy for Advanced Interdisciplinary Studies, Southern University of Science and Technology, Shenzhen, China. ⁴Swiss Federal Institute of Technology, Institute of Structural Engineering, Zurich, Switzerland. *e-mail: arnaud.mignan@sed.ethz.ch; bromarco@ethz.ch

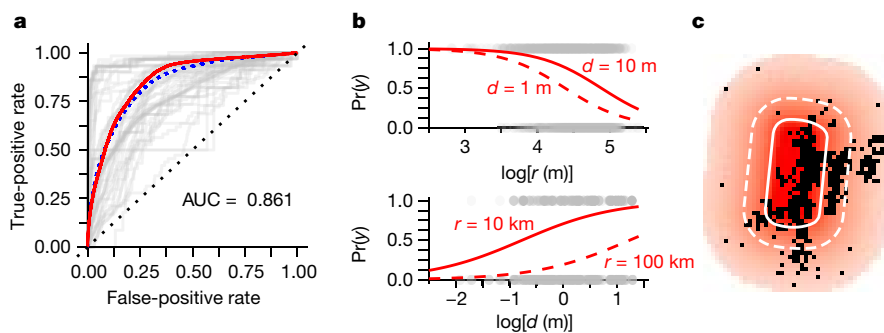


Fig. 2 | Prediction of aftershock spatial patterns using the distance, r , and the slip, d . **a**, Test-data ROC curves for the logistic regression. **b**, Logistic regression fit on training data. **c**, Example of logistic regression prediction. 58 ROC curves are shown, for the 57 mainshocks of the test set

performance, it decreases model interpretability for physical inference. Specifically, we show that a single neuron provides similar results as the DNN of DeVries et al.¹ with AUC = 0.85, and that the primary strain inputs of ref.¹ (used for stress computation) yield a better performance than their stress-based features with AUC = 0.86.

First, we reproduce the results of DeVries et al.¹, as illustrated in Fig. 1a–c (see the methods section of ref.¹). Their model contains 13,451 free parameters (weights and biases) because their DNN is made of six layers, each composed of 50 nodes. Their inputs are 12 engineered stress components: the absolute values of the tensor's six independent components, $|\sigma_{ij}|$, and their six opposites, $-|\sigma_{ij}|$. Their output is the probability that a spatial cell is in binary class y (1 if one or more aftershocks are present, 0 otherwise). We retrieve AUC = 0.85 and a precision of 5.4% at a threshold of $P = 0.5$, as in the original publication (we note that the slightly different precision is due to the random subsampling of the balanced training dataset). However, a similar AUC can be obtained when using only one input (scalar stress metric) and one node (two free parameters with one weight, β_1 , and one bias, β_0). This is illustrated in Fig. 1d–f, where we introduce a logistic regression for direct comparison with the DNN of DeVries et al. The classifier is defined as $\text{Pr}(y) = 1/\{1 + \exp[-(\beta_0 + \beta_1 \log x)]\}$, where y is the spatial cell class and x the chosen scalar stress metric. We obtain AUC = 0.85 and a precision of 5.4% at a threshold of $P = 0.5$ for the sum of absolute stress components (Fig. 1d–f), maximum shear stress and von Mises yield criterion.

We further show that a logistic regression based directly on the mainshock average slip, d , and the minimum distance, r , between space cells and mainshock rupture (that is, the simplest of the possible models, with orthogonal features) provides comparable or better accuracy than the stress-based models mentioned above. Both d and r (in metres) were obtained from the SRCMOD database—the same primary data as those used by DeVries et al.¹ to compute their stress metrics. The results are shown in Fig. 2. The performance is improved to AUC = 0.86 and the same precision 5.4% compared to the ones obtained with stress features^{2,3}. The distance–slip probabilistic model is described by

$$\text{Pr}(y) = \frac{1}{1 + e^{-(\beta_0 + \beta_1 \log r + \beta_2 \log d)}} \quad (1)$$

where $\beta_0 = 10.18 \pm 0.07$, $\beta_1 = -2.32 \pm 0.02$ and $\beta_2 = 1.16 \pm 0.01$. Equation (1) provides a transparent and interpretable model to forecast aftershock patterns from geometric and kinematic data, retrievable almost in real time after a mainshock. This approach is in line with both the literature on operational earthquake forecasting^{3–5} and statistical seismology^{6–8}.

This communication shows that, given the same datasets and same accuracy assessments proposed by DeVries et al.¹, deep learning does not offer new insights or better accuracy in predicting

(grey) and all combined mainshock–aftershock pairs (red; see legend of Fig. 1 for blue colour explanation). Dashed and solid curves in c represent $\text{Pr}(y) = 0.5$ and 0.7, respectively.

aftershock patterns. However, we strongly believe that deep learning is revolutionizing data analytics in many domains^{12,13}, including statistical seismology¹⁴. Therefore, the objective of our study is not to restrain its use in this field, but to stimulate a further research effort¹⁵.

Data availability

The data that support the findings of this study are available from the SRCMOD fault rupture catalogue (<http://equake-rc.info/SRCMOD>), the International Seismological Centre earthquake catalogue (<http://www.isc.ac.uk/iscgem>) and from DeVries et al.¹ at <https://github.com/phoebemrdevries/Learning-aftershock-location-patterns>.

Code availability

Original codes by DeVries et al.¹ are available at <https://github.com/phoebemrdevries/Learning-aftershock-location-patterns>. An R code including the distance–slip feature definition and logistic regression training/testing is available from the corresponding authors on request.

Online content

Any methods, additional references, Nature Research reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41586-019-1582-8>.

Received: 3 April 2019; Accepted: 30 July 2019;

Published online 2 October 2019.

- DeVries, P. M. H., Viégas, F., Wattenberg, M. & Meade, B. J. Deep learning of aftershock patterns following large earthquakes. *Nature* **560**, 632–634 (2018).
- Meade, B. J., DeVries, P. M. R., Faller, J., Viégas, F. & Wattenberg, M. What is better than Coulomb failure stress? A ranking of scalar static stress triggering mechanisms from 10^5 mainshock–aftershock pairs. *Geophys. Res. Lett.* **44**, 11,409–11,416 (2017).
- Reasenber, P. A. & Jones, L. M. Earthquake hazard after a mainshock in California. *Science* **243**, 1173–1176 (1989).
- Reasenber, P. A. & Jones, L. M. Earthquake aftershocks: update. *Science* **265**, 1251–1252 (1994).
- Gerstenberger, M. C., Wiemer, S., Jones, L. M. & Reasenber, P. A. Real-time forecast of tomorrow's earthquakes in California. *Nature* **435**, 328–331 (2005).
- Felzer, K. R. & Brodsky, E. E. Decay of aftershock density with distance indicates triggering by dynamic stress. *Nature* **441**, 735–738 (2006).
- Richards-Dinger, K., Stein, R. S. & Toda, S. Decay of aftershock density with distance does not indicate triggering by dynamic stress. *Nature* **467**, 583–586 (2010).
- Mignan, A. Utsu aftershock productivity law explained from geometric operations on the permanent static stress field of mainshocks. *Nonlinear Process. Geophys.* **25**, 241–250 (2018).
- Stacy, S., Gerstenberger, M., Williams, C., Rhoades, D. & Christophersen, A. A new hybrid Coulomb/statistical model for forecasting aftershock rates. *Geophys. J. Int.* **196**, 918–923 (2014).

10. Cattania, C., Hainzl, S., Wang, L., Roth, F. & Enescu, B. Propagation of Coulomb stress uncertainties in physics-based aftershock models. *J. Geophys. Res. Solid Earth* **119**, 7846–7864 (2014).
11. Cattania, C. et al. The forecasting skill of physics-based seismicity models during the 2010–2012 Canterbury, New Zealand, earthquake sequence. *Seismol. Res. Lett.* **89**, 1238–1250 (2018).
12. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
13. Jordan, M. I. & Mitchell, T. M. Machine learning: trends, perspectives, and prospects. *Science* **349**, 255–260 (2015).
14. Kong, Q. et al. Machine learning in seismology: turning data into insights. *Seismol. Res. Lett.* **90**, 3–14 (2019).
15. Beroza, G. C. Aftershock forecasts turn to AI. *Nature* **560**, 556–557 (2018).

Author contributions A.M. and M.B. contributed equally to the design and analysis of this study.

Competing interests The authors declare no competing interests.

Additional information

Reprints and permissions information is available at <http://www.nature.com/reprints>.

Correspondence and requests for materials should be addressed to A.M. or M.B.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

© The Author(s), under exclusive licence to Springer Nature Limited 2019

Reply to: One neuron versus deep learning in aftershock prediction

Brendan J. Meade^{1*}

REPLYING TO Mignan, A. & Broccardo, M. *Nature* <https://doi.org/10.1038/s41586-019-1582-8> (2019)

Before commenting on the interesting philosophical issues raised by Mignan and Broccardo¹, I note that the authors were able to reproduce the results presented in our paper² (available at <https://github.com/phoebemrdevries/Learning-aftershock-location-patterns>). In the accompanying Comment¹, the authors make two conceptual points: (1) a model with fewer parameters can explain the same feature to label mapping just as well as the neural network presented in our paper² and (2) quantities such as the average fault slip and the distance from the mainshock are precise and useful predictors that can be used instead of the elastic stresses that we used.

The fact that a neural network result can be closely approximated by a simpler model (the first point of Mignan and Broccardo) is a core result of our paper and one that we described in detail. In fact, it is the way in which our findings are interpretable. For example, the entire last paragraph of our paper is dedicated to this result, stating explicitly: “In other words, without any assumptions about receiver plane orientation or geometry, the neural network identified an aftershock location forecast that is strongly correlated with a small number of physical quantities”, including the maximum shear stress and the von Mises yield criterion. Similarly, we presented a graphical representation of this central result in figure 2d of our paper², where we presented the neural network prediction, and in figure 2b, c, where we showed how the maximum shear stress and the von Mises yield criterion provide close approximations. In the first part of their Comment¹, Mignan and Broccardo recapitulate this central result of our paper by constructing a single-node network that yields an AUC score approximately equal to that obtained by the neural network in our study² or the von Mises yield criterion alone (which, when filtered through a sigmoid, is a single-node network). The perspective presented in our paper is that it was interesting to discover that a neural network learned a simple, non-exotic combination of stresses that provided considerably improved precision.

Whereas point (1) is a restatement of a central result from our paper, point (2) represents a philosophical departure. In particular, they show that two parameters, the average mainshock fault slip and the fault–aftershock distance, are also precise and interpretable predictors of aftershock locations, serving as a parsimonious phenomenological model. We used an alternative, physics-focused approach to focus on physical parameters that appear in the equations for frictional fault failure (normal and shear stresses) and fracture (difference in principle stresses). Our approach ensures that the labels are consistent with the conservation of both mass and linear momentum. In other words, we included prior information about the physics of solid Earth by developing labels that were consistent with elastic-stress transfer. Because stresses decay with distance from a mainshock, the direct use of distance as a proxy for locally resolved stresses may be an effective approximation for operational aftershock forecasting of the type considered by Mignan and Broccardo. I thank the authors for replicating our results and for their insightful Comment.

1. Mignan, M. & Broccardo, M. One neuron versus deep learning in aftershock prediction. *Nature* <https://doi.org/10.1038/s41586-019-1582-8> (2019).
2. DeVries, P. M. H., Viégas, F., Wattenberg, M. & Meade, B. J. Deep learning of aftershock patterns following large earthquakes. *Nature* **560**, 632–634 (2018).

Author contributions B.J.M. conceived the idea and wrote the manuscript. The other authors of the original study² were not involved in the preparation of this Reply.

Competing interests The author declares no competing interests.

Additional information

Reprints and permissions information is available at <http://www.nature.com/reprints>.

Correspondence and requests for materials should be addressed to B.J.M.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

© The Author(s), under exclusive licence to Springer Nature Limited 2019

¹Department of Earth and Planetary Sciences, Harvard University, Cambridge, MA, USA. *e-mail: meade@fas.harvard.edu