



MARTY MELVILLE/AFP/GETTY

Figure 1 | Damage caused by the 2011 Christchurch earthquake. Large earthquakes can be followed by thousands of smaller ones, called aftershocks. In February 2011, an aftershock struck the city of Christchurch, New Zealand, and was more destructive than the earthquake it followed. The image shows the smoking ruins of the six-storey Canterbury Television building, which collapsed and caught fire in the aftershock, killing more than 100 people.

GEOPHYSICS

Aftershock forecasts turn to AI

Understanding how earthquakes interact is key to reliable earthquake forecasting. A machine-learning study reveals how the stress change induced by earthquakes at geological faults affects these interactions. [SEE LETTER P.632](#)

GREGORY C. BEROZA

All major earthquakes are followed by smaller ones, called aftershocks, which can themselves be hazardous. The forecasting of aftershocks is an area of long-standing seismological interest¹. It is receiving renewed attention² because of earthquake sequences in Italy, New Zealand and Japan over the past decade, in which the first earthquake in the sequence was not the most destructive (Fig. 1). On page 632, DeVries *et al.*³ use machine-learning tools to take a

fresh look at how changes in geological stress generated by earthquakes influence the spatial distribution of aftershocks. The authors' work provides more-accurate forecasts of aftershock locations than does the standard approach⁴.

Deterministic earthquake prediction remains an elusive goal, but seismologists are working intently to make quantitative probabilistic forecasts of future earthquake occurrences. Prominent among the factors thought to affect earthquake probabilities is the change in stress induced by one earthquake at the potential initiation site of another.

Probabilistic forecasting depends on well-established statistical properties of seismicity — the spatial and temporal distribution of earthquakes. Although earthquakes cluster in space and time, large ones are rare, which makes documenting the interactions between these earthquakes intrinsically challenging.

Large earthquakes, however, can be followed by thousands of aftershocks, which are indistinguishable from other earthquakes. Aftershocks occur by the same mechanism, on the same geological faults and under the same conditions as for other earthquakes. It is therefore

reasonable to assume that understanding the interactions between the largest earthquake in a sequence (the mainshock) and its aftershocks will enhance general understanding of earthquake interactions.

DeVries and colleagues studied these mainshock–aftershock interactions using a database of published distributions of mainshock-induced slip — the relative movement of geological features on opposite sides of a fault. From these distributions, the authors calculated the stress changes induced by the mainshocks. They fed this information into an artificial-intelligence system known as an artificial neural network, which was trained to determine the likelihood that aftershocks would occur in a particular location on a spatial grid.

The authors withheld a randomly selected 25% of the mainshock–aftershock sequences from the training data, and used this subset to validate the predictive power of their machine-learning method. They report that the trained network can predict the locations of aftershocks more accurately than can the standard forecasting approach, which considers only one aspect of the induced shift in stress, known as the Coulomb failure stress change⁴. The authors find that other characteristics of the stress change play a crucial part in triggering aftershocks. The paper therefore demonstrates how machine learning could aid research in seismology.

However, for several reasons, it might be premature to infer that DeVries and colleagues' work has led to an improved physical understanding of aftershock triggering. One reason is that the current study — and earlier studies of aftershock triggering — focused on static stress changes that occur and persist long after the passage of seismic waves⁵. But dynamic stress changes caused by seismic waves can also trigger earthquakes⁶. The combination of static and dynamic stress changes leads to a spatial distribution of aftershocks that differs from the pattern caused by static stress changes alone⁷.

Another reason for caution is that the authors' analysis relies on factors that are fraught with uncertainty. Uncertainties in earthquake locations are probably small, but uncertainties in slip distributions, on which the stress-change analysis depends, are large and potentially problematic. It is well documented that estimates of slip made by different investigators are subject to substantial differences⁸. The inferred stress change, which is input to the authors' machine-learning algorithm, depends on the rate of change of these slip distributions with respect to position, such that slip uncertainty is amplified. This issue is more problematic close to the fault than it is farther away from it, but most aftershocks occur close to the fault.

The situation is compounded by the fact that slip estimates invariably assume that slip occurs on faults that are planar or composed of multiple planes. However, fault geometry is known to be complex at all scales⁹.

This complexity leads to strong, local stress concentrations that can trigger aftershocks¹⁰, but that will not be included in slip models that assume planar faults. This could explain why the authors see no evidence of a lack of aftershocks near faults — caused by an overall decrease in stress — despite the fact that this feature is readily apparent in situations in which data and circumstances allow it to be clearly observed¹¹. These issues concerning uncertainty are not particular to the authors' study, but they counsel some temperance in calling for new physical models to explain the current results.

Regardless of the physical interpretation, the performance of DeVries and colleagues' artificial neural network is motivating. Until a few years ago, most statistical forecasts of aftershocks were more accurate than were physics-based forecasts, such as that of the authors. But there are now cases in which physics-based forecasting performs as well as purely statistical approaches^{12,13}. The time would seem ripe for methods based on artificial intelligence to enter the fray, and the work of DeVries *et al.* has established this beachhead.

Artificial-intelligence methods have much to offer seismology, and solid-Earth science more broadly. There are societally important phenomena to understand that are informed

by data sets growing rapidly in scale and scope, and by computational simulations growing rapidly in sophistication and realism. The application of machine-learning methods has the potential to extract meaning from these large and complex sources of information, but we are still in the early stages of this process. ■

Gregory C. Beroza is in the Department of Geophysics, Stanford University, Stanford, California 94305-2215, USA.
e-mail: berozag@stanford.edu

1. Reasenber, P. A. & Jones, L. M. *Science* **243**, 1173–1176 (1989).
2. Marzocchi, W., Taroni, M. & Falcone, G. *Sci. Adv.* **3**, e1701239 (2017).
3. DeVries, P. M. R., Viégas, F., Wattenberg, M. & Meade, B. J. *Nature* **560**, 632–634 (2018).
4. King, G. C., Stein, R. S. & Lin, J. *Bull. Seismol. Soc. Am.* **84**, 935–953 (1994).
5. Stein, R. S. *Nature* **402**, 605–609 (1999).
6. Hill, D. P. *et al. Science* **260**, 1617–1623 (1993).
7. Meng, X. & Peng, Z. *Geophys. J. Int.* **197**, 1750–1762 (2014).
8. Mai, P. M. *et al. Seismol. Res. Lett.* **87**, 690–708 (2016).
9. Tchalenko, J. S. *GSA Bull.* **81**, 1625–1640 (1970).
10. Smith, D. E. & Dieterich, J. H. *Pure Appl. Geophys.* **167**, 1067–1085 (2010).
11. Toda, S., Stein, R. S., Beroza, G. C. & Marsan, D. *Nature Geosci.* **5**, 410–413 (2012).
12. Segou, M. & Parsons, T. *Seismol. Res. Lett.* **87**, 816–825 (2016).
13. Cattania, C. *et al. Seismol. Res. Lett.* **89**, 1238–1250 (2018).

PRECISION MEDICINE

Sequence of events in prostate cancer

Whole-genome sequencing reveals the duplication of a regulatory region, called an enhancer, of the AR gene in treatment-resistant human prostate cancers. The finding shows the importance of analysing non-protein-coding regions of DNA.

KELLIE A. COTTER & MARK A. RUBIN

The publication^{1,2} of the human genome sequence in 2001 was accompanied by optimism that a rise in the availability of genomic data might improve clinical treatments. It was hoped that such data might one day enable an approach termed 'precision medicine', in which therapies are tailored to target the abnormalities specific to a particular cancer. Since then, technological advances in DNA-sequencing techniques, combined with substantially lower costs, have led to a boom in the sequencing of cancer samples. Given this progress, one might assume that the key genetic alterations that drive common cancers are already well known. However, writing in *Cell*, Takeda *et al.*³, Viswanathan *et al.*⁴ and Quigley *et al.*⁵ detail a previously unidentified type of genetic alteration that frequently occurs in late-stage human prostate cancer.

The Cancer Genome Atlas (TCGA) and the International Cancer Genome Consortium (ICGC) have undertaken some of the largest-scale projects reported so far to sequence the DNA of human cancers. These efforts have identified many DNA alterations that drive cancer growth, including mutations and genomic rearrangements. TCGA has sequenced the protein-coding regions of approximately 11,000 individual genomes and 33 types of cancer (<https://portal.gdc.cancer.gov>), whereas the ICGC has sequenced the protein-coding regions from more than 20,000 individual genomes and 22 kinds of cancer (<https://dcc.icgc.org>). Both projects have focused mainly on sequencing the protein-coding regions of genes, which represent less than 2% of the entire genome. In the Pan Cancer Analysis of Whole Genomes (PCAWG) project, the ICGC and TCGA systematically analysed whole-genome