Explainable artificial intelligence for short-term data-driven aftershock forecasts

FOTEINI DERVISI[1,2], MARGARITA SEGOU[1], BRIAN BAPTIE[1], PIERO POLI[3], IAN MAIN[2], ANDREW CURTIS[2]

[1] The Lyell Centre, British Geological Survey, Edinburgh, United Kingdom [2] School of GeoSciences, The University of Edinburgh, Edinburgh, United Kingdom

[3] Dipartimento di Geoscienze, Università degli Studi di Padova, Padova, Italy

British Geological Survey

THE UNIVERSITY of EDINBURGH

Università degli Studi di Padova

SPIN MONITORING A RESTLESS EARTH

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 955515 – SPIN ITN (www.spin-itn.eu).

1 INTRODUCTION & BACKGROUND

The catastrophic nature of earthquakes drives the need for understanding seismic events, as well as for providing forecasts of when these are likely to occur. In this study, we use a convolutional neural network to develop a data-driven spatiotemporal model to forecast next-day seismicity. Leveraging the predictive power of deep learning, our model uncovers complex patterns within earthquake catalogues and produces next-day expected seismicity rate and magnitude forecasts in regions of interest.

We employ SmaAt-UNet **[1]**, a convolutional neural network with convolutional block attention modules and depthwise-separable convolutions that is able to create meaningful representations of grid structures. We evaluate the performance of our forecasting model using data science and earthquake forecasting metrics and compare against a null hypothesis, the persistence model, which assumes no change between consecutive time steps. We also use an enhanced gradient-weighted class activation mapping technique (Smooth Grad-CAM++ **[2]**) to provide insights into the model's decision-making process. Finally, we use a time series forecasting foundation model, TimesFM **[3]**, to generate next-day aftershock forecasts on the same dataset and compare these results against those produced by the convolutional neural network.

2 DATA

We assemble a large dataset containing data from diverse tectonic regions using publicly available earthquake catalogues for Southern California, Northern California, New Zealand, Italy, Greece and Japan. We use earthquakes with magnitude ≥ 2 and depth ≤ 40km. We randomly split the dataset into training, validation and test sets using 80% of the data for training, 10% for validation and 10% for testing.

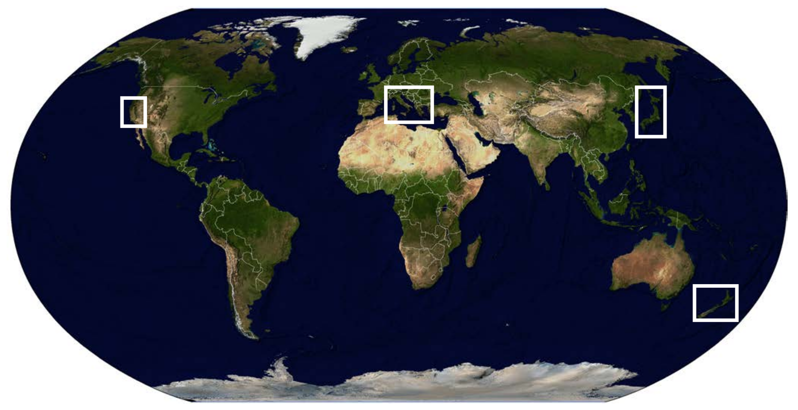

Fig 1 World map highlighting geographical regions that are included in our dataset. The map was created using https://ian.macky.net/pat/.

3 METHODOLOGY

We create spatiotemporal sequences of daily maps by splitting the spatial area into bins of 0.1 degrees of longitude and latitude. Three types of two-dimensional daily maps are used as input, containing: i) the number of events (rate) per grid cell, ii) the maximum magnitude of events in each grid cell and iii) the average depth of events in each grid cell. We identify events with magnitude ≥ 4 and for each one of these events, we aim to forecast the next day's seismicity within a spatial area of 1 longitude/latitude degree around the event. We use 7 days of rate, magnitude and depth maps as input and pass them through a deep learning model to produce next-day rate and magnitude forecasts.

Scan to see abstract

Sharing is encouraged

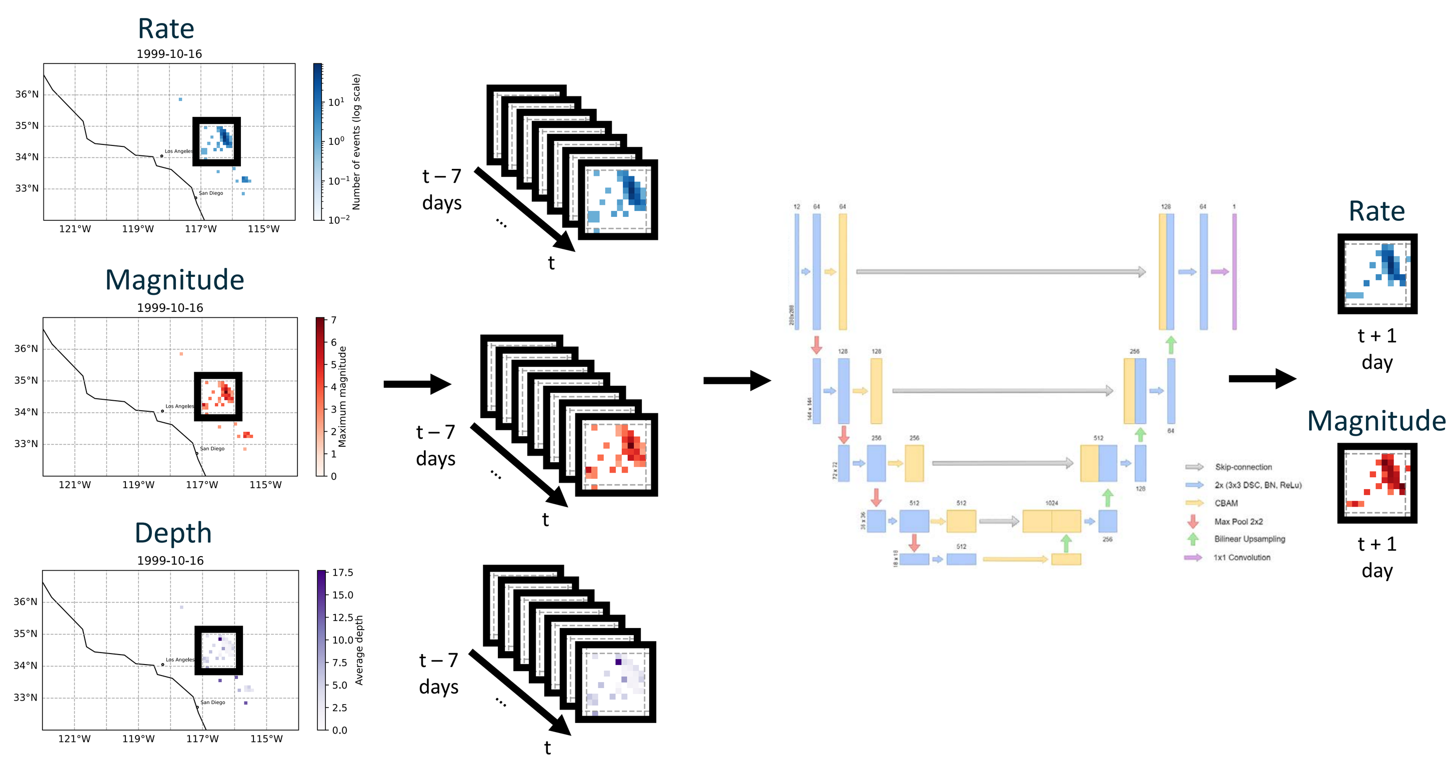

Fig 2 We identify events with magnitude 4 and above, create a spatial grid around them and produce deep learning-based next-day rate and magnitude forecasts using the rate, magnitude and depth maps of the previous 7 days as input. The neural network visualisation is from **[1]**.

4 QUANTITATIVE EVALUATION

			Classification								Regression	
			$\frac{TP+TN}{TP+FP+TN+FN}$	$\frac{TP}{TP+FP}$	$\frac{TP}{TP+FN}$	$\frac{2TP}{2TP+FP+FN}$	$\frac{TP}{TP+FP+FN}$	$\frac{FP}{TP+FP}$			$\frac{1}{N}\sum_{i=1}^{N}\lvert x_i - y_i\rvert$	$\sqrt{\frac{1}{N}\sum_{i=1}^{N}(x_i - y_i)^2}$
Forecasting Model	**Number of Parameters**	**Training Time (min)**	**Accuracy↑**	**Precision↑**	**Recall↑**	**F1 Score↑**	**Critical Success Index↑**	**False Alarm Ratio↓**	**Area Under ROC Curve↑**	**Area Under Precision-Recall Curve↑**	**Mean Absolute Error↓ (μ, σ)**	**Root Mean Squared Error↓ (μ, σ)**
SmaAt U-Net	4032205	62.8	0.997	0.944	0.821	0.878	0.783	0.056	0.910	0.884	0.058, 0.132	0.194, 0.457
Persistence (previous day)	-	-	0.975	0.511	0.594	0.550	0.379	0.489	0.789	0.558	0.204, 0.506	0.605, 1.610
Persistence (avg of prev. 7 days)	-	-	0.976	0.361	0.721	0.481	0.317	0.639	0.687	0.517	0.223, 0.546	0.652, 1.723

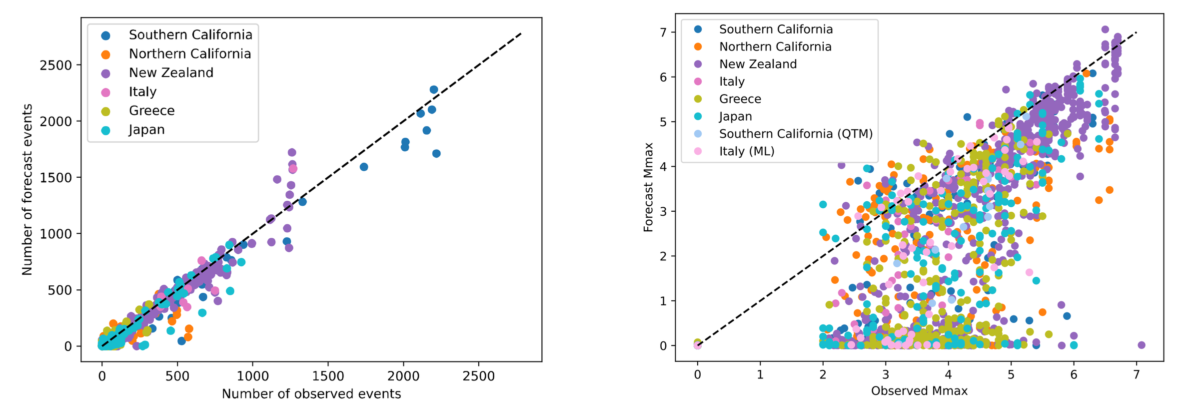

Fig 3 Observed versus forecast number of events (left) and maximum magnitude (right). Points on or close to the $y = x$ line represent forecasts that are consistent with the observations.

5 NUMBER OF FORECAST EVENTS: CNN VS FOUNDATION MODEL

CSEP N-test **[5] [6]**: $\delta_1 = 1 - F((N_{obs} - 1)|N_{fore})$, $\delta_2 = F(N_{obs}|N_{fore})$

where N_{obs} is the sum of the observed number of events over all spatial bins, N_{fore} is the sum of the forecast number of events over all spatial bins and $F(x|\mu)$ is a Poisson cumulative distribution with $F(x|\mu) < 0$ for $x < 0$. For an intended significance level $\alpha = 5\%$, a forecast is consistent if $\delta_1 > 0.025$ and $\delta_2 > 0.025$.

Forecasting Model	Rejected δ_1	Rejected δ_2
SmaAt-UNet	23.67%	4.67%
TimesFM-2.0	42.31%	3.99%

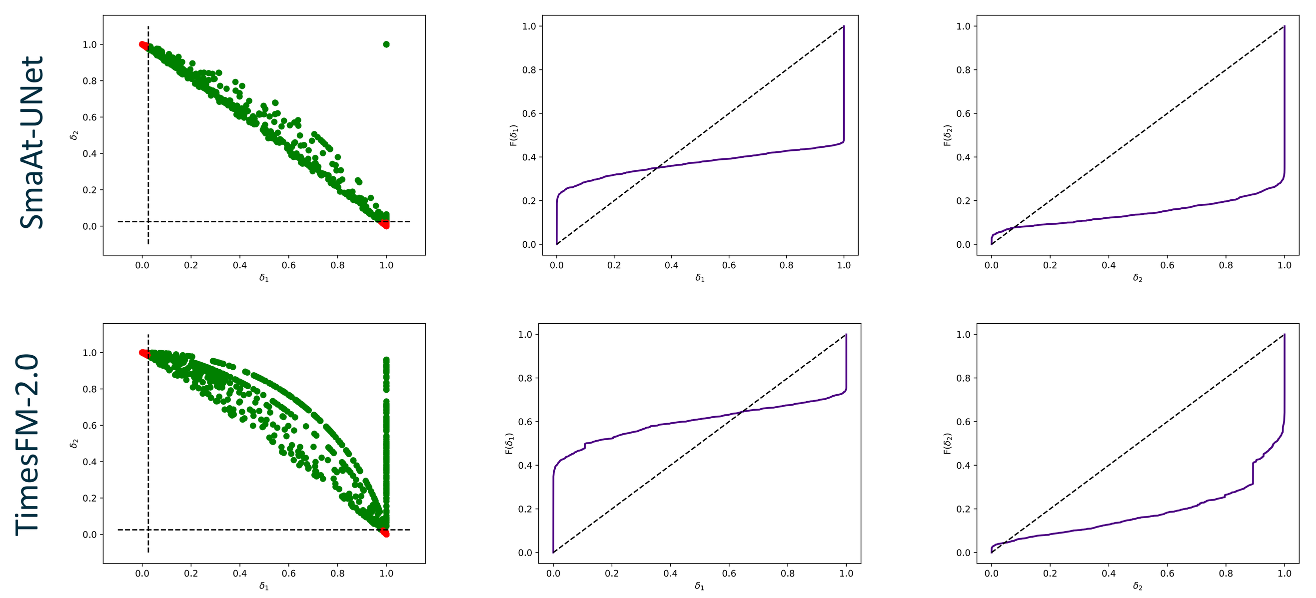

Fig 4 N-test results: Test set δ_1 and δ_2 values scatter plot and cumulative distributions. Forecasts with $\delta_1 > 0.025$ and $\delta_2 > 0.025$ are consistent with the observations.

6 QUALITATIVE EVALUATION

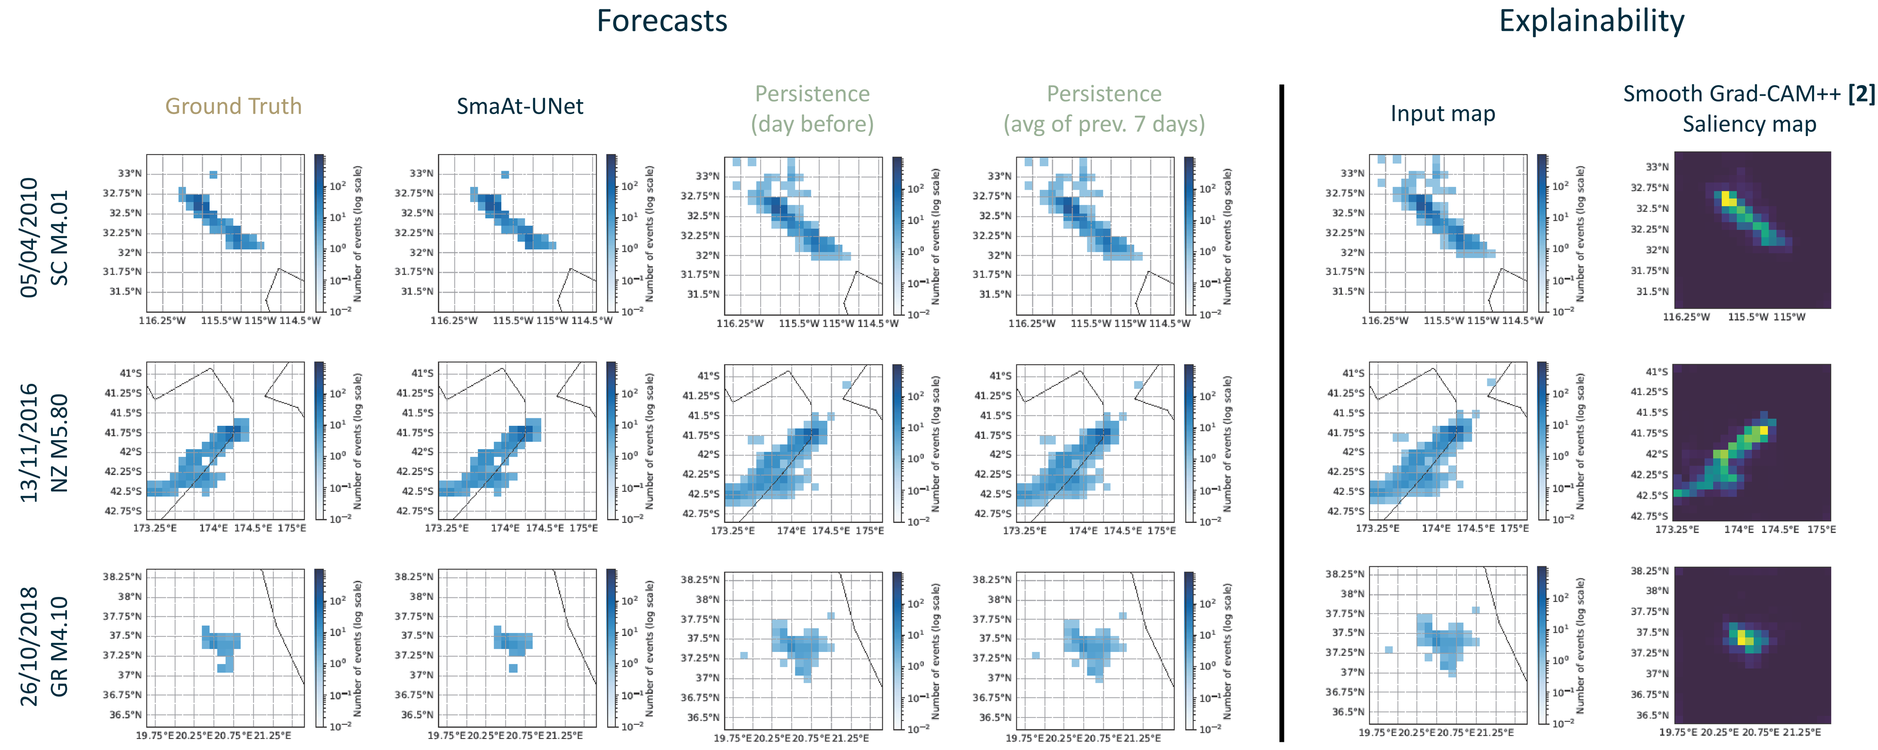

Fig 5 Examples of deep learning-based rate forecasts generated using the trained SmaAt-UNet model and comparison with the ground truth maps and the persistence baseline model. We also show the corresponding input maps and the saliency maps that have been generated with the use of Smooth Grad-CAM++ **[2] [4]**, which highlight the parts of the input maps that have the most influence on the forecasts.

7 CONCLUSION

- Persistence, the baseline model that assumes no change between consecutive time steps, is hard to beat!
- The forecasts are consistent with the observations in terms of spatial distribution of events, but we tend to underpredict the number of events and the maximum magnitude.
- Both SmaAt-UNet and TimesFM tend to underpredict the number of events, but the use of TimesFM leads to a considerably larger percentage of rate forecasts that are rejected based on the δ_1 statistic.

8 REFERENCES

[1] Trebing, K., Staǹczyk, T. and Mehrkanoon, S. (2021). "SmaAt-UNet: Precipitation Nowcasting Using a Small Attention-UNet Architecture." *Pattern Recognition Letters* 145:178–86. doi: 10.1016/j.patrec.2021.01.036, **[2]** Omeiza, D., Speakman, S., Cintas, C., & Weldermariam, K. (2019). "Smooth Grad-CAM++: An enhanced inference level visualization technique for deep convolutional neural network models." *arXiv preprint arXiv:1908.01224*., **[3]** Das, A., Kong, W., Sen, R., & Zhou, Y. (2024, July). A decoder-only foundation model for time-series forecasting. In *Forty-first International Conference on Machine Learning*. **[4]** Fernandez, F.-G. (2020). "TorchCAM: class activation explorer.", GitHub, https://github.com/frgfm/torch-cam **[5]** Zechar, J. D., Gerstenberger, M.C., & Rhoades, D. A. (2010). "Likelihood-Based Tests for Evaluating Space–Rate–Magnitude Earthquake Forecasts." *Bulletin of the Seismological Society of America*, 100(3):1184–1195. doi: 10.1785/0120090192., **[6]** Zechar, J. D., Schorlemmer, D., Liukis, M., Yu, J., Euchner, F., Maechling, P. J., & Jordan, T. H. (2010). "The Collaboratory for the Study of Earthquake Predictability perspective on computational earthquake science." *Concurrency and Computation: Practice and Experience* 22(12):1836–1847. doi: 10.1002/cpe.1519.