Global models for short-term earthquake forecasting and predictive skill assessment

Working Paper**Author(s):**

Nandan, Shyam; Kamer, Yavor; Ouillon, Guy; Hiemer, Stefan; Sornette, Didier

Publication date:

2020-12-16

Permanent link:

<https://doi.org/10.3929/ethz-b-000464460>

Rights / license:

[Creative Commons Attribution 4.0 International](#)

Originally published in:

EarthArXiv, <https://doi.org/10.31223/X5G02H>

Global models for short-term earthquake forecasting and predictive skill assessment

This manuscript has been accepted for publication in the present form in the THE EUROPEAN PHYSICAL JOURNAL SPECIAL TOPICS with doi <https://doi.org/10.1140/epjst/e2020-000259-3>

Global models for short-term earthquake forecasting and predictive skill assessment

Shyam Nandan^{1,*}, Yavor Kamer², Guy Ouillon³, Stefan Hiemer², Didier Sornette^{4,5}

¹ Windegstrasse, 5, 8953 Dietikon, Zurich, Switzerland

² RichterX.com

³ Lithophyse, 4 rue de l'Ancien Sénat, 06300 Nice, France

⁴ ETH Zurich, Zurich, Switzerland

⁵ SUSTech, Shenzhen, China

* shyam4iiser@gmail.com

ABSTRACT

We present rigorous tests of global short-term earthquake forecasts using Epidemic Type Aftershock Sequence models with two different time kernels (one with exponentially tapered Omori kernel (ETOK) and another with linear magnitude dependent Omori kernel (MDOK)). The tests are conducted with three different magnitude cutoffs for the auxiliary catalog (M3, M4 or M5) and two different magnitude cutoffs for the primary catalog (M5 or M6), in 30 day long pseudo prospective experiments designed to forecast worldwide $M \geq 5$ and $M \geq 6$ earthquakes during the period from January 1981 to October 2019. MDOK ETAS models perform significantly better relative to ETOK ETAS models. The superiority of MDOK ETAS models adds further support to the multifractal stress activation model proposed by Ouillon and Sornette (2005). We find a significant improvement of forecasting skills by lowering the auxiliary catalog magnitude cutoff from 5 to 4. We unearth evidence for a self-similarity of the triggering process as models trained on lower magnitude events have the same forecasting skills as models trained on higher magnitude earthquakes. Expressing our forecasts in terms of the full distribution of earthquake rates at different spatial resolutions, we present tests for the consistency of our model, which is often found satisfactory but also points to a number of potential improvements, such as incorporating anisotropic spatial kernels, and accounting for spatial and depth dependant variations of the ETAS parameters. The model has been implemented as a reference model on the global earthquake prediction platform RichterX, facilitating predictive skill assessment and allowing anyone to review its prospective performance.

Keywords ETAS models · magnitude dependent Omori law · multifractal stress activation · global earthquake forecasts · consistency · RichterX

1 Introduction

Over the last 40 years, the number of people living in earthquake prone regions has almost doubled, from an estimated 1.4 to 2.7 billion [Pesaresi et al., 2017], making earthquakes one of the deadliest natural hazards. Currently there are no reliable methods to accurately predict earthquakes in a short time-space window that would allow for evacuations. Nevertheless, real-time earthquake forecasts can provide systematic assessment of earthquake occurrence probabilities that are known to vary greatly with time. These forecasts become especially important during seismic sequences where the public is faced with important decisions, such as whether to return to their houses or stay outside. Short term earthquake probabilities vary greatly from place to place depending on the local seismic history and their computation requires scientific expertise, computer infrastructure and resources . While in most developed countries, such as Japan, New Zealand, Italy, earthquake forecasts are publicly available, the vast majority of the seismically vulnerable population is residing in developing countries that do not have access to this vital product. In this sense there is a global need for a system that can deliver worldwide, publicly accessible earthquake forecasts updated in real-time. Such forecasts would not only inform the public and raise public risk awareness but they would also provide local authorities with an independent and consistent assessment of the short-term earthquake hazard.

In addition to its social utility, a real-time earthquake forecasting model with global coverage would be an essential tool for exploring new horizons in earthquake predictability research. In the last two decades, the Collaboratory for the Study of Earthquake Predictability (CSEP) has facilitated internationally coordinated efforts to develop numerous systematic tests of models forecasting future seismicity rates using observed seismicity [Jordan, 2006]. These efforts, commendable as they are, address only a very specific type of models, namely seismicity based forecasts which express their forecast as occurrence rates under the assumption of Poissonian distribution. Thus, studies investigating the predictive potential of various dynamic and intermittent non-seismic signals, (such as thermal infrared, electromagnetic waves, electric potential differences, ground water chemistry etc) are effectively left out since they cannot be adequately tested in the provided CSEP framework. Recent studies have also pointed out deficiencies that introduce biases against models that do not share the assumptions of the CSEP testing methodology. Moreover, some researchers have expressed concerns regarding the effective public communication of the numerous test results associated with each model. Some have argued that test metrics should be tailored not only to the small community of highly-sophisticated statistical seismologists, but to the intuitive understanding of the general public and civil protection agencies.

We suggest that the drawbacks and limitations of the previous testing methodologies can be addressed by establishing a real-time global earthquake forecasting model that serves as a benchmark for evaluating not only grid based seismicity rate models but also a wide variety of alarm based methods. We thus introduced RichterX: a global earthquake prediction platform where participants can query the probability of an earthquake (or a number of earthquakes) above a given magnitude to occur in a specific time-space window, and issue to-occur or not-to-occur predictions [Kamer et al., 2020]. By analyzing the prospective outcomes of the issued predictions we can establish if they exhibit significant skill compared to our global reference model, and if they do, we can successfully rank them accordingly.

This work documents the development of such a global earthquake forecasting model derived from the Epidemic Type Aftershock Sequence (ETAS) family and presents a novel set of rigorous tests tailored to the specific needs of short term earthquake forecasting and benchmarking of short term earthquake predictions. ETAS based models have been shown to be the best contenders in the horse race organised within CSEP. Moreover, they contain generic and parsimonious assumptions that provide consistent descriptions of the statistical properties of realised seismicity. Specifically, the ETAS family of models is based on the following assumptions: (i) the distinction between foreshocks, mainshocks and aftershocks is artificial and all earthquakes obey the same empirical laws describing their probability to occur and their ability to trigger future earthquakes; (ii) earthquakes have their magnitude distributed according to the Gutenberg-Richter distribution; (iii) the rate of triggered events by a given earthquake obeys the Omori-Utsu temporal law of aftershocks; (iv) the number of earthquakes triggered by a given event obeys a productivity law usually linking the average number of offsprings to the exponential of the triggering earthquake magnitude; (v) triggered events are distributed in space according to a spatially dependent power law function.

Here, we develop horse races between two ETAS models differing in their specification of their time kernels, one with an exponentially tapered Omori kernel (ETOK) and another with a magnitude dependent Omori kernel (MDOK). We define three different training settings for the auxiliary catalog's magnitude cutoff (3, 4 or 5) and two different training settings for the primary catalog's magnitude cutoff (5 or 6), in 362 pseudo prospective global experiments designed to forecast $M \geq 5$ and $M \geq 6$ earthquakes at three different spatial resolutions, spanning scales from $45km$ to $180km$. While previous works have shown the importance of accounting for spatially varying ETAS parameters [Nandan et al., 2017], here we assume the same ETAS parameters hold for the whole Earth. This assumption is made for computational simplicity and with the intention to have a uniform global reference model allowing for an easier interpretation of the participants' predictive performance. This also allows us to focus on the key question we want to investigate, namely the role of a possibly magnitude dependent Omori exponent on forecasting skills. This hypothesis has been derived from a physics-based model of triggered seismicity based on the premises that 1) there is an exponential dependence between seismic rupture and local stress and 2) the stress relaxation has a long memory [Ouillon and Sornette, 2005; Sornette and Ouillon, 2005]. These physical ingredients predict that the exponent of the Omori law for triggered events is an increasing function of the magnitude of the triggering event. This prediction has been corroborated by systematic empirical studies for California and worldwide catalogues [Ouillon and Sornette, 2005], as well as for Taiwanese [Tsai et al., 2012] and Japanese [Ouillon et al., 2009] catalogues.

Therefore we consider the addition of magnitude dependant Omori law as a potential improvement to a global implementation of the ETAS model. We propose a general pseudo-prospective testing experiment that can be applied to any future candidate model. The testing experiment is designed to address the specific needs of a global, short-term earthquake forecasting application and correct for some of the defects highlighted by our previous work [Nandan et al., 2019c,b]. In particular, we use equal sized meshes compatible with the spatial scales available on the RichterX platform (radius in the range of $30-300km$) and target duration of 30 days, which is the maximum prediction time window in RichterX.

The organisation of the paper is as follows. Section 2 presents the data used in our tests and its main properties. Section 3 starts with a description of the pseudo-prospective forecasting experiments. Then, it defines the two ETAS models that are compared. We explain how parameters inversion is performed and the details of the simulations used to construct the forecasts. Section 4 presents the results, starting with a rate map and full distributions of earthquake numbers in the different cells covering the Earth at the eleven different resolution levels. The model comparisons are performed in terms of pair-wise cumulative information gains, and we calculate the statistical significance of right tailed paired t-tests. We study in details the sensitivity of our results to the spatial resolution, number of simulations and inclusion of smaller magnitudes in the auxiliary catalogs. We also describe how the best performing model is adopted as a benchmark for the RichterX global earthquake prediction contest, which is presented in more details in a companion paper. Section 5 concludes by summarising and outlining further developments. A supplementary material document provides additional figures and descriptions for the interested readers.

2 Data

We use the global earthquake catalog obtained from the Advanced National Seismic System (ANSS) database. To maintain the same precision for all reported earthquakes in the catalog, we first bin the reported magnitudes at 0.1 units. In this study, we use all $M \geq 5$ earthquakes that occurred between January 1981 and October 2019 as our primary data source for target earthquakes.

Figure 1 shows the different features of this dataset. In Figure 1a we show the location, time and magnitudes of these earthquakes. Figure 1b shows the spatial density of $M \geq 5$ earthquakes. This spatial density is obtained by first counting the number of earthquakes in $1 \times 1 \text{ deg}^2$ pixels, normalizing the counts by the area of each pixel and then smoothing the resultant density. We also show the time series of cumulative number of $M \geq 5$ and $M \geq 6$ earthquakes and the magnitudes vs. times of $M \geq 7$ earthquakes in Figures 1c and 1d, respectively.

Finally, in Figure 1e, we show the empirical magnitude distribution of $M \geq 5$ and $M \geq 6$ earthquakes. To each of them, we separately fit the Gutenberg-Richter (GR) law. The maximum likelihood estimate of the parameters of the GR distribution for $M \geq 5$ and $M \geq 6$ earthquakes are indicated in the inset in figure 1e. In order to obtain the exponent of the GR distribution, we use the analytical maximum likelihood estimator for binned magnitude derived by Tinti and Mulargia [1987]. Having obtained the exponent, the prefactor of the GR distribution can be analytically estimated. The GR law exponents obtained for both magnitude thresholds are 1.05 ($M \geq 5$) and 1.02 ($M \geq 6$) and are thus nearly identical. Such consistency is often treated as an indication of the completeness of the catalog [Cao and Gao, 2002; Mignan and Woessner, 2012]. With this reasoning, the consistency of GR exponents indicates the completeness of the catalog for $M \geq 5$ in our case.

For the appropriate calibration of the ETAS model, we also use the $M \geq 3$ earthquakes between January 1975 and October 2019 as auxiliary dataset. The use of auxiliary dataset is often encouraged in ETAS literature [Schoenberg et al., 2010; Wang et al., 2010; Seif et al., 2017], as it allows to minimize the biases in the genealogy tree of earthquakes

due to the missing sources [Sornette and Werner, 2005a]. Earthquakes in the auxiliary catalogs can act only as sources during the calibration of the ETAS model, thus, their completeness is not required.

For the sake of reproducibility, since catalogs are subject to updates, the catalog used in this study is provided as a supplementary material.

Figure 1: Features of primary data used in this study; (a) Location of $M \geq 5$ earthquakes since 1981; Sizes of the earthquakes scale with their magnitude and colors show the year of occurrence; (b) Spatial density of $M \geq 5$ earthquakes since 1981 obtained by smoothing the counts of earthquakes in $1 \times 1 \text{ deg}^2$ grid with a Gaussian filter; (c) Cumulative number of $M \geq 5$ and $M \geq 6$ earthquakes, with their respective scaled axes on the left and right side of the plot, respectively; (d) Magnitude vs. Time plot of $M \geq 7$ earthquakes; (e) Empirical frequency magnitude distribution of $M \geq 5$ and $M \geq 6$ earthquakes; Lines show the best fit Gutenberg Richter Distribution to the empirical distribution of $M \geq 5$ and $M \geq 6$ earthquakes.

3 Method

In this study, our aim is to compare the performance of different models (described in Section 3.2) in forecasting future earthquakes. We do this by means of pseudo prospective experiments.

3.1 Pseudo Prospective Forecasting Experiments

Prospective forecasting experiments are a powerful tool allowing scientists to check if the improvements lead to better forecasts of future unseen observations. Truly prospective experiments are time consuming, as they require many years before enough future observations accumulate to strengthen (or weaken) the evidence in favor of a model [Kossobokov, 2013]. A practical solution is to conduct pseudo prospective experiments. In these experiments, one uses only the early part of the dataset for the calibration of the models and leaves the rest as virtually unseen future data. Subsequently, the

calibrated models are used to simulate a forecast of future observations and the left out data is used to obtain a score for each of the forecasting models. These scores can then be compared to identify the best model.

Although, (pseudo) prospective experiments have started to catch up in the field of earthquake research [Kagan and Jackson, 2010; Zechar and Jordan, 2010; Eberhard et al., 2012; Kagan and Jackson, 2011; Zechar et al., 2013; Ogata et al., 2013; Hiemer et al., 2014; Hiemer and Kamer, 2016; Schorlemmer et al., 2018; Nandan et al., 2019b,c], they still have not become the norm. In this regard, the work done by the collaboratory for the study of earthquake predictability (CSEP) and others [Schorlemmer et al., 2018; Kossobokov, 2013] has been very commendable, as they have tried to bring the prospective model validation on the center stage of earthquake forecasting research. However, the prescribed prospective validation settings, in particular by CSEP, are too primitive and sometimes biased in favour of certain model types. For instance, most of the CSEP experiments have been conducted with spatial grids which are defined as $0.1 \times 0.1 \text{ deg}^2$ cells (see for instance Zechar et al. [2013]) or $0.05 \times 0.05 \text{ deg}^2$ cells (e.g. Cattania et al. [2018]), mostly for computational convenience. However, as the area of these cells vary with latitude, becoming smaller as one moves north or south from the equator, a model gets evaluated at different spatial resolutions at different latitudes, thus, by construction, yielding different performances as a function of latitude. This areal effect on the performance then gets convoluted with the underlying spatial variation in the performance of the model. For instance, modelers might find that their models yield better performance in California, but very poor performance in Indonesia at a fixed spatial resolution. Should they then infer that their models are better suited for a strike slip regime than for a subduction setting? Due to the inherent nature of the varying spatial resolution of the grid prescribed by CSEP, the answer to this question becomes obfuscated.

Another aspect of pseudo prospective experiments that is poorly handled by CSEP is that it forces the modelers to assume Poissonian rates in a given space-time-magnitude window irrespective of their best judgement. Nandan et al. [2019b] showed that this choice puts the models that do not comply with the Poissonian assumption on weaker footing than those models which agree with this assumption. As a result, the reliability of the relative rankings obtained from the models evaluated by CSEP remains questionable to some extent.

Last but not the least, in some model evaluation categories [Werner et al., 2011; Schorlemmer and Gerstenberger, 2007], CSEP evaluates the models using only "background" earthquakes, which are identified using Reasenberg's declustering algorithm as being independent [Reasenberg, 1985]. However, as the earthquakes do not naturally come with labels such as "background", "aftershocks" and "foreshocks" that can be used for validation, this a posteriori identification remains highly subjective. In regards to CSEP's use of the Reasenberg's declustering algorithm, Nandan et al. [2019c] pointed out that the subjective nature of declustering introduces a bias towards models that are consistent with the declustering technique, rather than the observed earthquakes as a whole. This puts into questions the value of such experiments, as their results are subject to change as a function of the declustering parameters.

Since the aim of our forecasting experiment is to assess which model is more suitable to serve as a reference model for the global earthquake prediction platform RichterX, it becomes important to address the drawbacks mentioned above by

designing a testing framework tailored to multi-resolution, short-term forecasts on a global scale. Accordingly we have designed the pseudo prospective experiments in this study with the following settings:

1. **Many training and testing periods:** We start testing models beginning on January 1, 1990 and continue testing till October 31, 2019, spanning a duration of nearly 30 years. The maximum duration of an earthquake prediction that can be submitted on the RichterX platform is 30 days. Using this time window, our pseudo prospective experiments are composed of 362 non overlapping, 30 days long testing periods. To create the forecast for each of the testing periods, the models are calibrated only on data prior to the beginning of each testing period for calibration as well as simulation. The forecasts are specified on an equal area mesh with predefined spatial resolution.
2. **Equal area mesh:** To create this equal area mesh, we tile the whole globe with spherical triangles whose area is constant all over the globe. This mesh is designed in a hierarchical fashion. To create a higher resolution mesh from a lower resolution one, the triangles in the lower resolution mesh are divided into four equal area triangles. In this way, we create eleven levels of resolution: at the first level, the globe is tiled with 20 equal area triangles (corresponding to an areal resolution of $\approx 25.5 \times 10^6 \text{ km}^2$ each); at the second level 80 equal area triangles tile the globe, and so on. Finally, at level eleven $\approx 21 \times 10^6$ triangles tile the globe with an areal resolution of $\approx 24 \text{ km}^2$. In this study, we evaluate the models at level six (unless otherwise stated), which has an areal resolution equivalent to a circle with radius $\approx 90 \text{ km}$. To test the sensitivity of our results to the choice of areal resolution, we also evaluate the models at level five and level seven, which correspond to an areal resolution equivalent to circles with radii $\approx 180 \text{ km}$ and $\approx 45 \text{ km}$, respectively. In principle, the models can be evaluated at all spatial resolutions (from 1 to 11). The resolutions in this study are chosen to be in accordance with the the spatial extents used on the RichterX platform (radius of 30 to 300km).
3. **Flexibility to use parametric or non-parametric probability distributions:** The model forecasts can be specified on the equal area mesh during each testing period either as the full distribution of earthquake numbers (as in Nandan et al. [2019b]) empirically obtained from the simulations, or as a Poissonian rate or as any other analytical probability distribution function that may be in line with the model assumptions.
4. **Performance evaluation using all earthquakes with $M \geq 5$ or $M \geq 6$:** We test the models against target sets consisting of $M \geq 5$ and $M \geq 6$ events that occurred during each testing period. During a given testing period, competing models forecast a distribution of earthquake numbers ($M \geq 5$ or $M \geq 6$ depending on the choice of target magnitude threshold or M_t) in each triangular pixel. We then count the actual number of observed earthquakes in each pixel. With these two pieces of information, the log likelihood LL_A^i of a given model A during the i^{th} testing period is defined as:

$$LL_A^i = \sum_{j=1}^N \ln (Pr_j^i(n_j^i)) \quad (1)$$

where Pr_j^i is the probability density function (PDF) of earthquake numbers forecasted by model A and n_j^i is the observed number of earthquakes ($\geq M_t$) during the i^{th} testing period in the j^{th} pixel. N is the total number of pixels, which at level six is equal to 20,480. Similarly, LL_B^i for another competing model (B) can be obtained. The information gain IG_{AB}^i of model A over model B for the i^{th} testing period is equal to

$$IG_{AB}^i = LL_A^i - LL_B^i . \quad (2)$$

In order to ascertain if the information gain of model A over B is statistically significant over all testing periods, we conduct a right tailed paired t-test, in which we test the null hypothesis that the mean information gain ($MIG_{AB} = \frac{\sum_i IG_{AB}^i}{362}$) over all testing periods is equal to 0 against the alternative that it is larger than 0. We then report the p -values obtained from the test. If the p -values obtained from the tests are smaller than the "standard" statistical significance threshold of 0.05, model A is considered to be statistically significantly more informative than model B.

3.2 Competing Models

3.2.1 Preliminaries on ETAS models

In this study, we conduct a contest between different variants of ETAS models only. The reasons for this are multifold:

1. Our target use case for RichterX requires providing near-real time short time earthquake forecasts on a global scale. ETAS models are suitable for such applications as they rely only on a timely stream of earthquake locations and magnitudes, whereas models based on stress transfer require additional data, such as fault plane orientations, rupture extent, slip distributions etc. which are often available only after a few days.
2. Due to the abundance of target events (the world-wide average number of $M \geq 5$ and $M \geq 6$ earthquakes per month, since 1990, is ≈ 137 and ≈ 13 , respectively), global models providing short term (here, monthly) earthquake forecasts can be tested with a greater statistical significance (even at high magnitude threshold, such as $M \geq 5$ and $M \geq 6$, of the testing catalog), compared to their regional counterparts.
3. On the global scale, there exist some "long term" models [Bird et al., 2015; Kagan and Jackson, 2011, 2010]. However, there is no model (to the best of our knowledge) that provides short term forecasts. We intend to fill this gap with the best performing ETAS model of this study.
4. On the regional scale, ETAS models [Nandan et al., 2019c] have been shown to be much more effective than standard smoothed seismicity models [Werner et al., 2011; Helmstetter et al., 2006], which provide forecasts of future earthquakes by smoothing the location of past background earthquakes. However, their forecasting effectiveness on the global scale remains to be assessed.

In this study, our goal is not to provide a comprehensive test between various types of state-of-the-art forecasting approaches [Bach and Hainzl, 2012; Helmstetter and Werner, 2014; Gerstenberger et al., 2005; Steacy et al., 2014;

Reverso et al., 2018; Cattania et al., 2014, 2018], but rather to design an experiment in which short term global earthquake forecasting models can be developed, compared and enhanced. Furthermore, we only conduct the horse race between the simplest ETAS models. This means that we exclude space and time variation of its parameters, which have been actively reported by numerous authors, at regional and global scales [Page et al., 2016; Zhang et al., 2020; Chu et al., 2011; Nandan et al., 2017; Ogata, 2011], to lead to enhanced performance. Furthermore, the ETAS models considered in this study do not make use of any other datasets such as fault networks, global strain rates, source models, focal mechanisms and so on. Some authors [Cattania et al., 2018; Guo et al., 2015; Bach and Hainzl, 2012] have shown in case studies that these additional datasets enhance the forecasting potential of ETAS type models. We disregard these complexities, not due to any underlying belief that they are not informative, but because their limited availability would hinder a real-time implementation on the RichterX platform. We also maintain that the initial model should be simple enough. Then, adding complexities should follow sequentially, only if they can be justified by their forecasting gains over simpler models. With these points in mind, in the following, we describe the different variants of ETAS models that we have compared in this study.

3.2.2 ETAS model with exponentially tapered Omori kernel (ETOK)

Model Description In this model, the seismicity rate $\lambda(t, x, y | \mathcal{H}_t)$ at any time t and location (x, y) depends on the history of the seismicity $\mathcal{H}_t$ up to t in the following way:

$$\lambda(t, x, y | \mathcal{H}_t) = \mu + \sum_{i: t_i < t} g(t - t_i, x - x_i, y - y_i, M_i) \quad (3)$$

where μ is the time-independent background intensity function, g is the triggering function and (t_i, x_i, y_i, M_i) represents the time, x-coordinate, y-coordinate and magnitude of the i^{th} earthquake in the catalog, respectively.

The memory function g in Equation 3 is formulated as:

$$\begin{aligned}
 g(t - t_i, x - x_i, y - y_i, M_i) = & K e^{a(M_i - M_c)} \times \frac{e^{-\frac{t - t_i}{\tau}}}{(t - t_i + c)^{1+\omega}} \\
 & \times \left[(x - x_i)^2 + (y - y_i)^2 + d e^{\gamma(M_i - M_c)} \right]^{-1-\rho}
 \end{aligned} \quad (4)$$

which is the product of three kernels:

1. *The productivity kernel* $K e^{a(M_i - M_c)}$ quantifies the expected number of aftershocks triggered by an earthquake with magnitude M_i above the magnitude of completeness, M_c , where K and a are the productivity constant and exponent respectively.
2. *The exponentially tapered Omori kernel* $\frac{e^{-\frac{t - t_i}{\tau}}}{(t - t_i + c)^{1+\omega}}$ quantifies the time distribution of the direct aftershocks of an earthquake that occurred at t_i . The exponential taper term $e^{-\frac{t - t_i}{\tau}}$ ensures that the parameter ω can attain even negative values during the calibration of the model, which is not possible for a pure power law distribution, as it becomes unnormalizable for exponents smaller than 1.

3. *The isotropic power-law spatial kernel* $[(x - x_i)^2 + (y - y_i)^2 + de^{\gamma(M_i - M_c)}]^{-1-\rho}$ quantifies the spatial distribution of the aftershocks of an earthquake with magnitude M_i that occurred at (x_i, y_i) .

Note that the model defined in Equations 3 and 4 implicitly assumes that the magnitudes of both the background and the triggered earthquakes follow the Gutenberg Richter (GR) distribution, which is described by the following probability density function (PDF):

$$f(M) = \beta e^{-\beta(M - M_c)} \quad (5)$$

Note that the exponent β in expression (5) is related to the b -value reported above in the inset of Figure 1e via $\beta = b \ln 10 \approx 2.3b$. Thus a value of β in the range 2.3-2.4 as shown in Figure S1 in the supplementary materials corresponds to a b -value in the range 1.0-1.04.

Due to its commonality in both the background and the triggering function, GR law is usually factored out of the explicit formulation of the ETAS model. However, one could imagine other formulations of ETAS models in which such factoring out is not possible. For instance, simply allowing the background earthquakes and aftershocks to follow GR distribution with different exponents (β_1 and β_2) makes the factoring impossible and the exponents β_1 and β_2 then have to be jointly inverted with the other ETAS parameters. In this context, Nandan et al. [2019a] showed that, using the Californian earthquake catalog, not only the exponents corresponding to the background earthquakes are distinct from those of aftershocks, but also that the magnitude distribution of direct aftershocks is scaled by the magnitude of their mainshock. Despite these findings, we make the simplifying assumption that both background earthquakes and aftershocks follow the same GR distribution, and factor it out from the explicit formulation of the ETAS model. Nevertheless, the GR law plays an explicit role when the ETAS model is used to forecast the magnitude of the future earthquakes and its parameter β has thus to be inverted from the training period.

Simulation We follow the standard algorithms for the simulation of synthetic earthquake catalogs for the ETAS model [Zhuang et al., 2004, 2002; Nandan et al., 2019b]. For completeness, a detailed description of the simulation is also provided in the Supplementary Text S1.

Parameter Inversion And Modelling Choices As described in Supplementary Text S1, the set of parameters $\{\mu, K, a, c, \omega, \tau, d, \rho, \gamma, \beta\}$ are necessary for the simulation of future earthquake catalogs. The values of these parameters are not known in practice and they have to be inverted from the training data. The parameters $\{\mu, K, a, c, \omega, \tau, d, \rho, \gamma\}$ can be inverted by calibrating the model (Equation 3) on the real data by means of the Expectation Maximization (EM) algorithm proposed by Veen and Schoenberg [2008]. To obtain the parameter β , we first bin the magnitudes of the earthquakes in the ANSS catalog in 0.1 units and then use the analytical maximum likelihood estimator derived by Tinti and Mulargia [1987] for binned magnitudes.

An important consideration before calibrating the ETAS model is the choice of the primary and auxiliary catalogs. The main difference between these two catalogs is that earthquakes in the primary catalog can act both as targets and sources during the calibration of the ETAS model, while the earthquakes in the auxiliary catalogs can act only as

sources. In ETAS literature [Schoenberg et al., 2010; Wang et al., 2010; Seif et al., 2017], the use of auxiliary catalogs is encouraged during inversion of ETAS parameters, so as to minimize the biases in the genealogy tree of earthquakes due to the missing sources [Sornette and Werner, 2005a; Saichev and Sornette, 2006b]. In this study, we calibrate the ETAS model (described in Equations 3 and 4) using primary catalogs with two different magnitude thresholds: $M_{pri} = 5$ and 6. Both these primary catalogs start in year 1981 and include earthquakes from all over the globe. For the auxiliary catalogs, which start in 1975 and are also composed of earthquakes from all over the globe, we use three different magnitude thresholds: $M_{aux} = 3, 4$ and 5 during calibration. We use different magnitude thresholds for the primary catalogs to test the hypothesis that better forecasting potential can be achieved for higher magnitude thresholds if we specifically train our models for them. We use different magnitude thresholds for the auxiliary catalog to test the hypothesis that smaller earthquakes play an important role in triggering and can improve the forecasting potential of the ETAS models.

Note that, even though the available ANSS catalog extends down to magnitude 0, we do not use such a low magnitude threshold for the auxiliary and the primary catalogs because: (1) in the formulation of the ETAS model, the primary catalog should follow a GR law and be complete above the considered threshold magnitude. These two criteria can not be fulfilled for the global ANSS catalog at magnitude thresholds lower than 5, and extending back to year 1981; (2) lowering the magnitude threshold of both the primary and auxiliary catalog increases enormously the computational burden for both the inversion and simulations.

In Figure S2, we show the time evolution of the estimates of the parameters for the ETAS model with exponentially tapered kernel for $M_{aux} = 3, 4$ and 5 and $M_{pri} = 5$. The time evolution for $M_{pri} = 6$, for the same model and the three auxiliary magnitude settings, is shown in Figure S3. Beside the "usual" ETAS parameters, Figure S2 shows the time series of the branching ratio. This parameter quantifies the average number of triggered earthquakes of first generation per triggering event, as well as the fraction of triggered earthquakes in the training catalog [Helmstetter and Sornette, 2003]. For a branching ratio < 1 , the system is in the sub-critical regime. For a branching ratio > 1 , the system is in the super-critical regime [Helmstetter and Sornette, 2002]. In addition, in Figure S1, we show the time evolution of the parameter β for two M_{pri} settings. Since the parameter β is only estimated from the primary catalog, only two time series are obtained and not six (one for each of the two M_{pri} settings) as in the case of other ETAS parameters. The time series of all the parameters is composed of 362 points, each corresponding to one of the training catalogs preceding the 362 testing periods. We notice that parameters show conspicuous variation with time, with a tendency to stabilise after about 2011, perhaps reflecting a better global catalogue completeness. We cannot exclude a genuine trend resulting from the shortness of the time series, which are strongly impacted by the two great earthquakes of magnitude larger than 9 that occurred in 2004 (great Indian ocean earthquake) and 2011 (Tohoku, Japan). Furthermore, some of the parameter pairs (μ , n or branching ratio), (c , ω) and so on exhibit cross correlations. In addition, the parameters also seem to be systematically dependent on the choices of M_{aux} and M_{pri} . Investigating the sources of these time variations, cross correlations and dependencies on auxiliary and primary magnitude thresholds is beyond the scope of this paper. In this article, we focus on evaluating the importance of these hyper-parameter choices

(M_{aux} and M_{pri}) in terms of forecasting performance. Nevertheless, we report the time evolution of these parameter estimates as it would aid the readers in reproducing the results presented in later sections.

3.2.3 ETAS Model With Magnitude Dependent Omori Kernel (MDOK)

Model Description While the primary equation of the seismicity rate for this ETAS model remains the same (Equation 3), the triggering kernel is modified to account for a possible magnitude dependence of Omori-Utsu parameters c and ω . The triggering kernel for this model is redefined as:

$$g(t - t_i, x - x_i, y - y_i, M_i) = K e^{a(M_i - M_c)} \times \frac{e^{-\frac{t - t_i}{\tau}}}{[t - t_i + c(M_i)]^{1 + \omega(M_i)}} \times \left[(x - x_i)^2 + (y - y_i)^2 + d e^{\gamma(M_i - M_c)} \right]^{-1 - \rho} \quad (6)$$

where $c(M_i) = 10^{c_0 + c_1 M_i}$ and $\omega = \omega_0 + \omega_1 M_i$.

The functional form for $c(M)$ is inspired from the works of Shcherbakov et al. [2004], Davidsen et al. [2015] and Hainzl [2016a]. All three authors found the c -value to increase exponentially with the mainshock magnitude. While the first two authors interpreted the c -value dependence on the mainshock magnitude as a part of a self-similar earthquake generation process (i.e. a physical process), Hainzl [2016a] attributed this dependence to the rate dependent aftershock incompleteness (i.e. a data sampling issue). The latter would require to replace the missing events in some way, as they play a role in triggering of future events. Yet no such procedure has ever been proposed. Note that several other authors [Scholz, 1968; Dieterich, 1994; Narteau et al., 2005] have also argued for the magnitude-dependence of the onset of the power-law decay based on ideas such as stress corrosion and rate and state dependent friction. However, these authors suggest that the c -value would correlate negatively with mainshock magnitude, as their model predicts that the larger the stress perturbation, the shorter would be the duration between the mainshock and the onset of the power-law decay. Regardless of the underlying mechanism for the dependence of the c -value on mainshock magnitude, the evidence for such an exponential dependence is rather clear, and thus warrants an explicit formulation within the ETAS model.

The linear dependence of the Omori exponent ω on the mainshock magnitude is based on the work of Ouillon and Sornette [2005] and Sornette and Ouillon [2005], who reported strong empirical evidence together with a physics-based theory for such a dependence for mainshocks in Californian and worldwide catalogs. Tsai et al. [2012] confirmed this observation for the Taiwanese catalog and Ouillon et al. [2009] for the Japanese catalog. These authors used a wealth of different techniques, such as various space-time windowing methods, binned aftershock time-series, wavelet analysis and time evolution of aftershocks maximum magnitude, in order to ascertain the robustness of the results and that the observed magnitude dependence of ω would not be due to some bias induced by a specific method. Ouillon and Sornette [2005] and Sornette and Ouillon [2005] proposed a theoretical statistical physics framework in which the seismic rate results from an exponential Arrhenius like activation with an energy barrier influenced by the total stress fields induced by past earthquakes and far-field tectonic loading. These authors showed that the combination of the exponential activation rate together with the long memory kernel of stress relaxation leads to temporal multifractality

expressed empirically as a magnitude-dependent Omori exponent ω . They coined this model the multifractal stress activation (MSA) model. More precisely, the MSA model can be rationalized as follows:

1. the stress at any location is the sum of the far-field contribution due to tectonic loading and the stress fluctuations due to past events;
2. each earthquake ruptures a complex set of patches whose number increases exponentially with the magnitude of the event;
3. each failing patch redistributes stress in its surrounding according to the laws of linear elasticity, so that positive or negative stress contributions add up as patches fail and consecutive earthquakes occur. The stress transferred by a failed patch at any target location can be treated as a random variable distributed according to a Cauchy law, i.e., decaying as a power law with exponent $(1 + \nu)/2$ [Kagan, 1992]. The effect of the earthquake rupture at the target location is thus the sum of the corresponding random variables. The exponent ν thus encompasses all the geometrical complexity of the problem: the (fractal) nature of the fault system, the Gutenberg–Richter law (i.e., the size of the source events), the distribution of focal mechanisms, the (possibly self-affine) morphology of slip along the rupture plane, and the spatial decay of the stress Green’s function;
4. the memory of local past stress fluctuations decays as a power-law of time, due to rock (nonlinear) viscosity, with exponent $1 + \theta$. This function encapsulates all brittle and ductile relaxation phenomena such as dislocations motion, pressure-dissolution, slow earthquakes or even those too small to be detected. In that sense, θ characterizes the whole complexity of stress relaxation in the crust.
5. at any location, the seismicity rate depends exponentially on the local shear stress, in agreement with many known underlying failure processes;

The model then predicts that the seismicity rate consists in a time invariant base rate due to the tectonic loading, nonlinearly modulated by a time varying term depending on past seismicity. This term can increase the rate if past (and/or most recent) stress fluctuations are positive, but may also decrease if they are negative. When solved self-consistently by considering all (statistical) mechanical interactions between events, the model predicts that the Omori exponent of the triggered sequence following an event of magnitude M decays with time with an exponent p increasing linearly with M . This peculiar feature is indeed predicted to hold exactly when the condition $\nu(1 + \theta) = 1$ is fulfilled, which can be viewed as the consequence of the space-time self-organization of fault networks in the brittle crust. Reviewing the possible values of parameters ν and θ for the Earth’s crust, Ouillon et al. [2009] showed that their estimations allowed them to bracket this criterion, thus evidencing another analogy with second-order phase transitions where critical exponents are linked by such relationships close to a critical point.

In this forecasting experiment, we aim to systematically test the idea that explicitly taking account of magnitude dependence in these two Omori parameters would lead to an improvement in the forecasting ability of the modified ETAS models relative to the ones in which these dependencies are ignored.

Simulation Given the set of parameters $\{\mu, K, a, c_0, c_1, \omega_0, \omega_1, \tau, d, \rho, \gamma, \beta\}$, the simulation of the time, location and magnitude of the future earthquakes proceed in the same way as for a standard ETAS model (see Supplementary Text S1), except for one difference. In this case, the times of the direct aftershocks of an earthquake with magnitude M_i are simulated using the time kernel whose parameters depend on M_i in the way described in Equation 6. This means that, despite the fact that the MSA model is by construction nonlinear, we here consider a linear approximation for the purpose of tractability. Indeed, in the MSA model, the exponential nonlinearity occurs in the stress space, a variable that is not computed within the ETAS formulation which focuses only on rates. A full MSA approach would require to compute the stress transfer (and its time dependence) due to all past events, taking account of their individual rupture complexity, and assessing all their uncertainties. As this remains challenging in the present state of seismological research, we bypass this obstacle and provide a simplified approach by introducing a magnitude-dependent Omori kernel.

Parameter Inversion And Modelling Choices Again in this case, we adapt the EM algorithm proposed by Veen and Schoenberg [2008] to invert the parameters of the model (Equation 6). For the sake of completeness, we also calibrate these models with six primary and auxiliary catalog settings as described in Section 3.2.2. Again, without going into the possible underlying causes of the time variation of the estimated parameters and their dependence on the choice of M_{aux} and M_{pri} hyper-parameters, we report the time evolution of the estimated parameters for ETAS model with magnitude dependent Omori kernel in figures S4 and S5.

3.3 Summary of competing models and experiment settings

In summary, we have twelve competing models: six models belong to the ETOK class and six belong to the MDOK class. In each of these two classes, three models are calibrated with a primary catalog magnitude threshold $M_{pri} = 5$ and three others are calibrated with a threshold $M_{pri} = 6$. These three models can be distinguished based on the different magnitude thresholds for the auxiliary catalog, $M_{aux} = 3$, $M_{aux} = 4$ and $M_{aux} = 5$, used during calibration and simulations.

Each of these twelve models are shown in Table 1 and are individually calibrated on the 362 training period periods. We then compare their forecasting performance using the $M \geq 5$ and $M \geq 6$ earthquakes under the validation settings prescribed in Section 3.1. Only models that have been been calibrated with $M_{pri} = 5$ are used to forecast $M \geq 5$ earthquakes to avoid "extrapolated" forecasts of models trained with only $M \geq 6$ earthquakes. All the models are used to forecast the $M \geq 6$ earthquakes as targets during each 30 days long validation period. In summary, six models are validated and scored using $M \geq 5$ earthquakes and all the twelve models are validated and scored using $M \geq 6$ earthquakes.

Figure 2: (a) Total Rate of $M \geq 5$ earthquakes ($km^{-2} month^{-1}$) between 12/03/2011 and 11/04/2011 forecasted using ETAS model with MDOK, $M_{aux} = 4$ and $M_{pri} = 5$; The Mw 9.1 Tohoku earthquake occurred on 11/03/2011; (b-d) Rates of three types of earthquakes (see Section S1) that are superposed to create the final forecast shown in panel a; (b) background events; (c) aftershocks of type 1; (d) aftershocks of type 2; the average number of each earthquake type in the final forecast is indicated in panels b-d; (e-l) Probability density functions (PDF) of earthquake numbers that are forecasted by the model in the circular geographic region of radius = 300 km around “earthquake prone” cities of the world.

	Model Numbers											
	1	2	3	4	5	6	7	8	9	10	11	12
Omori Kernel	ET	ET	ET	ET	ET	ET	MD	MD	MD	MD	MD	MD
M_{pri}	5	5	5	6	6	6	5	5	5	6	6	6
M_{aux}	3	4	5	3	4	5	3	4	5	3	4	5

Table 1: All twelve models resulting from different calibration choices; ET and MD stand for ETAS models with exponentially tapered Omori kernel and magnitude dependent Omori kernels, respectively.

4 Results and Discussion

4.1 Forecasted Rate Map And Full Distribution Of Earthquake Numbers

In this section, we illustrate how the forecasts of different models are constructed. We do this only for a selected model and a particular testing period, as the procedure for all other testing periods and models is the same.

Figure 2(a) shows the net forecasted rate of earthquakes (per km^2 per month) in the time period immediately following the Tohoku earthquake (between 12/03/2011 and 11/04/2011) for the ETAS model with magnitude dependent Omori kernel (MDOK) and auxiliary magnitude setting of $M_{aux} = 4$ and primary magnitude setting of $M_{pri} = 5$. Figures 2(b-d) show the contributions of the three type of earthquakes to the net forecasted rate. The first contribution comes from the background earthquakes that are expected to occur during the testing period (Figure 2b). The second contribution is from the cascade of aftershocks (Aftershock Type 1) that are expected to be triggered by the earthquakes in the training period (Figure 2c). The third and the final contribution comes from the cascade of aftershocks (Aftershock Type 2) that are expected to be triggered by the background earthquakes occurring during the testing period (Figure 2d). In this particular testing period, the Type 1 aftershocks have the highest contribution, with ≈ 264 earthquakes on average, while the contributions of the background earthquakes and Type 2 aftershocks are relatively minuscule. The occurrence of the Tohoku earthquake just before the testing period is the main cause of this dominance. However, it is important to note that the relative importance of these three components depends on the time scale of the testing period. For longer testing periods, such as on the order of a decade to a century, the contribution of background earthquakes and especially of Type 2 aftershocks becomes significant, if not dominating compared to the Type 1 aftershocks.

It is important to mention here that these average rate maps are just used for the sake of illustration of the generated forecasts, as they provide a convenient representation. In reality, each pixel on the globe is associated with a full distribution of forecasted earthquake numbers. To illustrate this, we show in Figure 2 (e-l) the probability density function (PDF) of earthquake numbers that is forecasted by the model in circular geographic regions (with 300 km radii) around some of the “earthquake prone” cities of the world. These PDFs are obtained by counting the number of simulations in which a certain number of earthquakes were observed and then by dividing those by the total number of simulations that were performed. In this study, we perform 100,000 simulations for all the models and for all testing periods. Notice that the PDF of the forecasted number of earthquakes varies significantly from one city to another, despite the fact that none of the competing models feature spatial variation of the ETAS parameters. This variation can be attributed to variation in the local history of seismicity from one place to another. Other factors that control the shape

of these distributions include the time duration of the testing period and the size of the region of interest (see Figure 1 in Nandan et al. [2019b]). It is also evident that the forecasted distributions of earthquake numbers around these selected cities display thick tails and cannot be approximated by a Poisson distribution. In fact, Nandan et al. [2019b] showed that, if a Poissonian assumption is imposed, the ETAS model yields a worse forecast relative to the case in which it was allowed to use the full distribution. Therefore we use the full distribution approach proposed by Nandan et al. [2019b] to evaluate the forecasting performance of the models in the following section.

4.2 Model Comparison

Figure 3: (a) Time evolution of cumulative information gain (CIG) of the six magnitude dependent Omori kernel (MDOK) models when forecasting $M \geq 5$ earthquakes during the 362 testing periods over the base model; base model is calibrated with exponentially tapered Omori kernel (ETOK), $M_{aux} = 5$ and $M_{pri} = 5$; (b) Same as panel (a) except that the twelve competing models are used to forecast $M \geq 6$ earthquakes during the testing periods; the solid (resp. dashed) lines track the CIG evolution for the models with $M_{pri} = 5$ (resp. $M_{pri} = 6$).

Cumulative Information Gain (CIG) In Figure 3, we show the time series of cumulative information gain of all competing models over the base ETAS model in the two experiments designed to forecast $M \geq 5$ (Figure 3a) and $M \geq 6$ (Figure 3b) earthquakes during the 362 testing periods. The base model has been calibrated with the exponentially tapered Omori kernel (ETOK), $M_{pri} = 5$ and $M_{aux} = 5$. The six models shown in Figure 3a have been trained with either magnitude dependent Omori kernel (MDOK) or ETOK, auxiliary magnitude threshold (M_{aux}) of 3, 4 or 5 and primary magnitude threshold (M_{pri}) of 5. In Figure 3b, we show the cumulative information gain of these six models along with those variants that have been trained specifically with $M_{pri} = 6$. The performance of these models has been tracked with dashed lines. The configurations of all the twelve models is indicated in Figure 3b.

From both panels in Figure 3, we can make the following observations:

1. All other model settings being the same, the ETAS models with MDOK achieves higher CIG over the base model than the ETAS models with ETOK. This observation is independent of the M_{aux} and M_{pri} settings in both experiments, i.e. when forecasting $M \geq 5$ earthquakes as well as when forecasting $M \geq 6$ earthquakes.
2. There is a slight deterioration (possibly a saturation) in the model performance with the decreasing magnitude threshold of the auxiliary catalog. For instance, when forecasting $M \geq 5$ earthquakes, the performance of ETAS model with MDOK increases substantially when decreasing M_{aux} from 5 to 4 but then slightly diminishes when decreasing M_{aux} further from 4 to 3. Similarly, using ETOK, the model performance first substantially increases and then only shows a marginal increase when decreasing M_{aux} from 5 to 4 and then from 4 to 3 respectively. Similar observations can be made in Figure 3b.
3. Except in one case (MDOK, $M_{aux} = 5$), the models that have been specifically trained with $M_{pri} = 6$ show either no improvement or only marginal improvement over the models that have been trained with $M_{pri} = 5$.
4. Model performance increases either by changing M_{aux} from 5 to 4 or by switching the time kernel from ETOK to MDOK, or both, leading to the model with MDOK, $M_{aux} = 4$ and $M_{pri} = 5$ being the best performing model (albeit marginally) in both experiments.

Figure 4: (a) Pairwise mean information gain (MIG, per testing period) matrix of the six models used to forecast $M \geq 5$ earthquakes; (i, j) element indicates the MIG of the i^{th} model over the j^{th} model; (b) MIG matrix of twelve models in the experiments dealing with forecasting $M \geq 6$ earthquakes; Black and grey labels correspond to models trained with $M_{pri} = 5$ and $M_{pri} = 6$, respectively; (c) p-value matrix obtained from right tailed paired t-test, testing the null hypothesis that the MIG of the i^{th} model over the j^{th} model, when forecasting $M \geq 5$ earthquakes, is significantly larger than 0 against the alternative that it is not; (d) same as panel (c) but when forecasting $M \geq 6$ earthquakes.

Mean Information Gain (MIG) and Statistical Significance

So far, we have compared all models to a common null model and then compared their cumulative information gain over this null model to each other. In order to assess whether one model performs significantly better than others, we also compare the models pairwise. In Figure 4a, we show the pairwise mean information gain (MIG) per testing period corresponding to the six models that are used to forecast $M \geq 5$ earthquakes. In this matrix, (i, j) element indicates the MIG of the i^{th} model over the j^{th} model. The MIG_{ij} terms are computed by averaging the information gain of the i^{th} model over the j^{th} model in the 362 testing periods. Note that this matrix is antisymmetric. In Figure 4(b), we show the MIG matrix for the twelve models in the experiments dealing with forecasting $M \geq 6$ earthquakes. The models that have been trained with $M_{pri} = 6$ are labelled in grey while the ones trained with $M_{pri} = 5$ are labelled in black.

In order to find if the MIG of one model over the other is statistically significant, we perform right tailed paired t-test. In this test, we test the hypothesis that the MIG of the i^{th} model over the j^{th} model is significantly larger than 0 against the alternative that it is not. Figure 4c and d shows the matrix of $\log_{10}(\text{p-value})$ obtained from the t-test corresponding to the MIGs shown in panel a and b, respectively. From these MIG and p-value matrices, we can make the following observations:

1. MIG matrices echo the observations made from Figure 3.
2. All other configurations being the same, the models with MDOK almost always perform statistically significantly (at a standard significance level of 0.05) better than the models with ETOK when forecasting both $M \geq 5$ and $M \geq 6$ earthquakes.
3. We also find that, when decreasing M_{aux} from 5 to 4, the models tend to always perform statistically significantly better (all other settings being the same), not only when forecasting $M \geq 5$ earthquakes but also nearly always when forecasting $M \geq 6$ earthquakes. In the latter case, there is just one exception, i.e. when the MDOK kernels are used with $M_{pri} = 6$ setting. However, the same trend does not hold when decreasing the M_{aux} from 4 to 3.
4. We also find that the models that have been trained specifically with $M_{pri} = 6$ almost never significantly outperform the models trained with $M_{pri} = 5$, with one exception being the model with MDOK and $M_{aux} = 5$.

Sensitivity to the Spatial Resolution

To investigate if the observations from Figures 3 and 4 exhibit sensitivity to the spatial resolution, Figures S6 show the time evolution of CIG for two different spatial resolutions (level 5 and level 7). For these two resolutions, we also present the table of pairwise MIG and p-value in Figures S7 and S8, respectively. We find that the observations made earlier from Figures 3 and 4 are robust with respect to the choice of spatial resolution.

Sensitivity to the Number of Simulations

An important point to consider when evaluating and comparing the models is the number of simulations to perform. As the models are evaluated based on the empirical distribution of earthquake numbers that they provide in a given space-time-magnitude bin, performing too few simulations would

introduce random fluctuations in the log likelihood (Equation 1), thus making the model comparisons unreliable. This is due to the fat-tailed distribution of seismic rates [Saichev et al., 2005; Saichev and Sornette, 2006a, 2007], which implies strong sample to sample fluctuations and a slow convergence of statistical properties [Sornette, 2006].

On the other hand, more simulations come at higher computational costs. As a result, it is important to optimize this trade-off. Figure S9 shows the net log-likelihood (summed over all testing periods) that a model obtains as a function of the number of simulations. The default number of simulations (100,000) considered in this study to obtain all the results is indicated with a shaded vertical bar. At 100,000 simulations, all the models show a slow convergence towards their "true" log likelihood score. Furthermore, the relative ranking of the models seem to be stable for more than 100,000 simulations at all spatial resolutions, further justifying this choice.

On the superiority of ETAS with the MDOK Kernel In summary, the results point to the significant superiority of the MDOK kernel over the ETOK when forecasting the rate of future earthquakes using the ETAS models. However, as the ETAS model with magnitude dependent Omori kernel (Equation 6) features a magnitude dependence of both $c(M)$ and $\omega(M)$, we cannot distinguish just from the model comparisons presented thus far if this model's superiority results from the magnitude dependence of $c(M)$ or of $\omega(M)$. To investigate this question, we define two variants of this model: one that features an $\omega(M)$ dependence with a magnitude independent c , and another one that features a $c(M)$ dependence with a magnitude independent ω . We then calibrate these models on all the 362 training periods. The time evolution of estimated parameters is reported in Figures S10 and S11, respectively. We use the estimated parameters to simulate 100,000 catalogs for the corresponding testing periods. To limit the needed computational resources, we only calibrate these models with the $M_{aux} = 5$ and $M_{pri} = 5$ setting and then use these two models to create and evaluate the forecasts for $M \geq 5$ earthquakes. We then compare (Figure S12) the performance of these two models to the one obtained from the ETAS model which features both $c(M)$ and $\omega(M)$ dependence and has been calibrated with the same M_{aux} and M_{pri} setting.

We find that, while the model with only $\omega(M)$ dependence outperforms the base model, the model featuring only $c(M)$ dependence systematically underperforms at all spatial resolutions. These results indicate that the superiority of MDOK models over ETOK models (Figure 3) results from the $\omega(M)$ dependence rather than a $c(M)$ dependence. In fact, the latter dependence inhibits it from realizing its true potential in forecasting (Figure S12). In other words, accounting for the $\omega(M)$ dependence is the crucial improvement for forecasting, while including a $c(M)$ dependence is detrimental.

It is thus natural to ask why do our calibrations of the MDOK model yield a positive correlation between $c(M)$ and mainshock magnitude (Figure S4)? The answer to this question potentially lies in the strong correlation between the two parameters ω and c , as seen from Figures S2 and S3. Assuming that a positive correlation between ω and mainshock magnitude exists, as proposed by Ouillon and Sornette [2005] and Sornette and Ouillon [2005] and also apparent in Figure S10, then the strong positive correlation between ω and c would artificially introduce a positive correlation between c and mainshock magnitude, masking the true underlying correlation of c and mainshock magnitude, which may indeed be negative as revealed in the model featuring only a $c(M)$ dependence (Figure S11).

In Figure S12c, we assess whether accounting for a negative correlation between c and mainshock magnitude could lead to any information gain over the model with the ETOK kernel. We find that the $c(M)$ model does not provide any systematic information gain. One possible reason for the poor performance of the $c(M)$ models in forecasting could lie in the short-term aftershock incompleteness [Hainzl, 2016b], which is present in both the training and testing catalogs. This rate-dependent incompleteness would not only dampen the negative correlation between c and mainshock magnitude, but also lead to very low information gain, as the events that would have led the $c(M)$ type model to be more informative are missing from the testing catalog in the first place.

On the Importance of Small Earthquakes in Forecasting Our results also indicate that including smaller earthquakes (to an extent) in the auxiliary catalog leads to a significant improvement in the forecast. This significant improvement can be attributed to the improved coverage and resolution of the global seismogenic zones as well as to the improved estimates of the parameters during calibration. However, the improvement starts to saturate (and sometimes even deteriorates) when even smaller earthquakes are included in the auxiliary catalog. This could potentially be due to the existence of a minimum triggering magnitude, M_0 , below which earthquakes do not trigger any aftershocks [Sornette and Werner, 2005b]. If we assume that the global average value of M_0 is somewhere in between 3 and 4, it naturally follows that we would also observe a saturation in model performance when reducing M_{aux} from 4 to 3, as the newly added earthquakes do not contribute to the triggering process. The inclusion of earthquakes smaller than the actual M_0 may even lead to deterioration in performance, as the calibration process implicitly assumes that all earthquakes have the potential to trigger aftershocks, and thus leads to biased parameter estimates. Moreover, if such a magnitude threshold exists, it could vary spatially, complicating the analysis and interpretation.

Another possible way to explain the saturation in performance improvement is by noting that, with a decrease of M_{aux} from 5 to 4, there is a nearly 5.5 fold increase in the number of earthquakes in the catalog ($M \geq M_{aux}$, between January 1975 and October 2019), while when M_{aux} is decreased from 4 to 3 the increase is only 1.5 fold, indicating a significant number of missing events in the global earthquake catalogue at these small magnitudes. This saturation in earthquakes numbers, i.e., the catalogue incompleteness, could also explain the saturation in the performance of the models, because the calibration of the ETAS models becomes intrinsically biased [Sornette and Werner, 2005a].

On the Possible Self-Similarity of the Triggering Process The insignificant difference in the performance of the models that have been trained with $M_{pri} = 6$ and $M_{pri} = 5$ suggests the existence of self similarity in triggering processes. More concretely, the models do not need to be trained specifically on $M_{pri} = 6$ to perform best in forecasting $M \geq 6$ earthquakes, as even the models trained on $M_{pri} = 5$ can do an equally good job. This observation could potentially be generalized to even higher magnitude thresholds, although we have not tested it in this work.

On the Exclusivity of the Two Model Improvements Finally, the cumulative improvements obtained by changing the time kernel from ETOK to MDOK and M_{aux} from 5 to 4, indicate that these two modifications capture, to some extent, mutually exclusive aspect of the triggering process. Furthermore, these two modifications seem to be equally

important, as they separately lead to similar information gains over the base model (see the solid orange and red curve in Figure 3).

4.3 Consistency Test

In an earthquake forecasting experiment, consistency tests play an important part, as they allow for the direct comparison of model's expectations with the observations, thus serving as necessary sanity checks. One such important sanity check is the "N-test" in which the overall number of earthquakes forecasted by a model is compared against the actual number of earthquakes observed during the testing period. Indeed, this test, along with other consistency tests such as L, M and S tests (see Rhoades et al. [2011] for details), have been used by CSEP to measure the consistency of the models relative to the data. It is important to note that these tests are not used to rank the models.

Not surprisingly, one of the hard-coded assumption in these tests, thus far in CSEP, has been that the distribution of the overall number of earthquakes forecasted by the models is Poissonian. Thus, when the numbers of earthquakes forecasted by the models are compared against the observed numbers (especially when aftershocks were deliberately not removed), most often the models are found to be inconsistent (see for e.g. Figure 9 in Werner et al. [2011]). For instance, Werner et al. [2011] have showed that, with a retrospective assumption of negative binomial distribution, the smoothed seismicity models developed in their study "passed" the N-tests for all testing periods.

Indeed, it is prohibitively reductive to enforce the same assumption on all models regardless of their formulation. Furthermore, the assumptions of the models should not be modified retrospectively. Last but not least, the assumptions in a model should be self consistent at all scales. For instance, if a model assumes that the rate of future earthquakes is Poissonian, it cannot then be evaluated using a negative binomial assumption for the N-test and a Poissonian assumption for estimating the information gain. In summary, the consistency tests should be modified to allow for simultaneous testing of models with diverse assumptions. One possible way to do this for the "N-tests" is to build an empirical PDF of earthquake numbers forecasted by the models from the numerous simulations as per Nandan et al. [2019b] and as done here.

Figure 5 shows the consistency between the PDF of forecasted number of earthquakes and the actual number of earthquakes observed during the 362 testing periods, for all the six competing models used to forecast $M \geq 5$ earthquakes. In these figures, the model type can be inferred by combining the row and the column names. Colors used in these figures show the probability of observing a certain number of earthquakes during a given testing period. The 95%ile of the PDF for each testing period is traced using the dashed red lines in the figure and the white crosses show the actual number of earthquakes ($M \geq 5$) during each testing period. Finally, the solid black line shows the mean number of earthquakes observed during all the testing periods. Figure S13 shows the same information as in Figure 5, but for the twelve models used to forecast $M \geq 6$ earthquakes. Recall that the six extra models in this case comes from the distinction introduced by the minimum threshold of the primary catalog ($M_{pri} = 5$ or 6) used to train the models.

Figure 5: (a-f) Consistency between the PDF of forecasted numbers of earthquakes and the actual number of earthquakes observed during the 362 testing periods, for all the six competing models used to forecast $M \geq 5$ earthquakes; model specifications (type of time kernel used and M_{aux} values) are indicated as row and column headings; colors show the $\log_{10}(\text{probability})$ of observing a certain number of earthquakes during a given testing period; 95%iles of the PDF for each testing period are traced using the dashed red lines in the figure; white crosses show the actual number of earthquakes ($M \geq 5$) during each testing period; solid black lines show the mean number of earthquakes observed during all the testing periods.

We can observe from both Figure 5 and S13 that the number of earthquakes forecasted by all the models seem consistent with the average number of earthquakes observed during all testing periods. However, when looking at individual testing periods, a lot of inconsistencies can be found. For instance, in testing periods immediately following very large earthquakes such as the Tohoku earthquake (11 March 2011, Mw 9.1) or the Sumatra earthquake (26 December 2004, Mw 9.3), the forecasted number is much lower than the observed number of earthquakes and not even the best model (Figure 5e and Figure S13k) is able to account for this inconsistency. This inconsistency can be primarily attributed to the isotropic assumption of the spatial kernel leading to the underestimation of the productivity exponent a (see Equation 4 and 6). Note that this effect of underestimating the productivity exponent due to the isotropic assumptions has been documented by several researchers [Hainzl et al., 2008; Helmstetter et al., 2006; Guo et al., 2015; Bach and Hainzl, 2012], who have also proposed solutions to account for anisotropy in specific case studies. In the future, we aim to generalize those solutions for real-time applications on the global scale. Moreover, other simplifications in the models, such as ignoring the spatial variation and depth dependence of parameters, may also be at the origin of some of these inconsistencies. The quantification of the extent to which each of these different factors contribute to inconsistencies will be undertaken in future studies. We also observe from Figure 5, that there are extended periods (such as between 1997-2005) in which the observed numbers of earthquakes are systematically smaller than the forecasted numbers,

possibly pointing towards a time variation of the triggering parameters and (or) background rate. Such inconsistencies are less evident for $M \geq 6$ earthquakes (Figure S13), possibly because of their sparse numbers during a given testing period, making it easier for models to pass the N-tests.

4.4 Real Time Application for Short-term Forecasts and Predictive Skill Assessment

The design of the forecasting experiment has been tailored to a global application for short term (up to 30 days) and regional (up to 300km) earthquake forecasts. Accordingly, we have operationalized the best performing ETAS model (with MDOK, $M_{aux} = 4$ and $M_{pri} = 5$) developed in this study via the RichterX platform available at www.richterX.com [Kamer et al., 2020]. On this website, the public can query the real-time model probabilities for earthquakes with $M \geq 5$ anywhere on the globe. The forecasts are provided in real-time in the sense that (1) global simulations are updated every hour as new earthquakes ($M \geq 4$) are entered in the ANSS catalog and (2) the probabilities depend on the actual time at which the user is requesting the forecast. A forecast request is performed by centering a circle at any location on the globe. The user then has the option to adjust the circle radius (30 to 300 km), time duration (1 to 30 days), minimum magnitude (M5+ to M9+) and the minimum number of earthquakes. These parameters are then used to query the database of real-time prospective simulations. The number of simulations that feature events satisfying the forecast criteria are used to construct an empirical PDF that defines the reported probability.

Michael et al. [2020] showed that a statement regarding the probability of N or more earthquakes within a specific space-time-magnitude window helps the media to accurately report probabilistic earthquake forecast. Therefore, we see the RichterX platform as an important step in improving public earthquake awareness and preparedness. It is important to note that the RichterX platform does not distinguish between an aftershock or a mainshock to assess the future probability of an earthquake. Furthermore, it allows the users to interact with the probabilities, by adjusting the forecast parameters online, facilitating an intuitive understanding of the underlying hazard. In these two regards, the RichterX platform differs from the efforts of the USGS, which started to publicly release aftershock forecasts for all events $M \geq 5$ throughout the United States in September 2019 as a table of the probability of one or more earthquakes for the next day, week, month, and year for $M \geq 3$, ≥ 5 , ≥ 6 , and ≥ 7 , respectively [Michael et al., 2020].

The availability of such a publicly accessible, real-time global earthquake forecasting model allows for new testing applications. Namely, it can be used as a reference benchmark to evaluate other short-term forecasts or deterministic predictions. Since the model probabilities are based on synthetic event sets, the forecasts are independent of prescribed grids and are not hindered by assumptions about distributions. Building on this feature, we introduce the RichterX platform as a global earthquake prediction contest, where participants can challenge the reference model by issuing deterministic to-occur or not-to-occur predictions anywhere on the globe. In the accompanying paper, we introduce the platform and demonstrate metrics that allow for consistent ranking of competing models [Kamer et al., 2020]. In this way, we aim to address the deficiencies found in the current CSEP testing methodologies, allowing for the inclusion of model types that were previously deemed incompatible and encourage a broader participation. We do not intend to keep the platform limited to "seismological" experts, but rather make it accessible to experts from other fields as well as

"amateur" scientists. In fact, anyone with an idea, intuition or a model is invited to challenge the forecast developed in this study by submitting testable predictions.

5 Conclusion and Outlook

Upon rigorous testing of the two ETAS models with two different time kernels (one with exponentially tapered Omori kernel and another with magnitude dependent Omori kernel), with three different training settings for the auxiliary catalog's magnitude cutoff (3, 4 or 5) and two different training settings for the primary catalog's magnitude cutoff (5 or 6), in 362 pseudo prospective global experiments designed to forecast $M \geq 5$ and $M \geq 6$ earthquakes, we can derive the following conclusions:

1. ETAS models with Omori kernels whose parameters explicitly depend on the magnitude of the mainshock perform significantly better relative to the ETAS models that ignores such dependencies. The superiority of ETAS models with magnitude dependent Omori kernel only results from the incorporation of the magnitude dependence of the Omori exponent, thus adding further support to the multifractal stress activation model proposed by Ouillon and Sornette [2005] and Sornette and Ouillon [2005].
2. While inclusion of more data in the auxiliary catalog by lowering the minimum magnitude cutoff from 5 to 4 leads to significant improvement in the forecasting performance, the performance saturates (and even deteriorates) when even smaller magnitudes ($M \geq 3$) are included in the auxiliary catalog. This counter-intuitive observation could have its origin in biases resulting from the incompleteness of the catalogue at these small magnitudes. Alternatively or together, this may also provide an observational evidence for the theoretical concept of a minimum magnitude of earthquakes that can trigger aftershocks [Sornette and Werner, 2005b].
3. ETAS models do not need to be trained specifically with $M \geq 6$ earthquakes in the primary catalog to have outstanding forecasting performance above this magnitude threshold. Models trained using a lower magnitude threshold ($M \geq 5$) can do an equally good job. This observation could be generalized to even higher magnitude thresholds possibly pointing to the self-similarity of the triggering process.
4. The number of earthquakes forecasted by the models is not always consistent with the observed number of earthquakes during the testing period. This is especially true in experiments designed for forecasting $M \geq 5$ earthquakes. These inconsistencies possibly arise from the simplifications, such as using an isotropic spatial kernel, as well as spatially homogeneous, depth independent and time invariant ETAS parameters, hardwired in the models presented in this study.

In order to obtain a fair and reliable comparison of the model performance, we have corrected some of the obvious defects of the past model testing experiments. These corrections include:

1. using equal sized mesh to ensure homogeneity of testing scores over the globe.
2. allowing the models the flexibility to specify the forecasts in accordance with their assumptions.

3. no declustering of the testing catalogs.

The models developed and tested in this work constitute a first imperfect attempt at developing global models that are capable of making short-term operational forecasts. Several simplifications have been made, especially in terms of diversity of the models developed and tested. Some of the obvious simplifications include (a) considering only ETAS type models, (b) assuming the parameters of the ETAS models to be spatially homogeneous and time invariant, (c) ignoring the depth dependence of parameters, (d) ignoring errors in the data, (e) assuming isotropic spatial kernels and so on. Nevertheless, by introducing fair and reliable testing schemes, in which modellers have the flexibility to adhere to their best judgement consistently, this study can serve as a framework for further model developments. Indeed, by operationalizing the best performing model as a benchmark for the RichterX prediction contest, we enable fellow modellers to use our results as a stepping stone for improving their models. This also constitutes a continuing process of peer-review, whereby anyone who finds the forecast probabilities too low or high can issue a to-occur or a not-to-occur prediction, providing us with important prospective feedback to improve our model.

On more general grounds, forecasting models can be split into two broad categories, namely statistical models (such as ETAS) and physical ones (using quantities such as static or dynamic stress transfer). The latter require the knowledge of many additional parameters, including the spatial extent and orientation of each rupture, as well as a detailed description of the slip over the failure planes. Nandan [2017] showed that our ability to forecast aftershock sequences using a stress-transfer approach increased if one took into account the triggering probabilities provided by an independent ETAS declustering process (the stress-based forecast being logically more appropriate for direct aftershocks). This, in return, suggests that a better knowledge of the space-time variations of the stress field may help to improve the forecasts of ETAS-like models. Nevertheless, the difficulty of such a forecasting framework is that the details of rupture must be known in real time for all past events, and forecasted as well for all future events. As this is clearly out of scope given our very limited knowledge of the deterministic structure of fault networks in the Earth crust, the MSA model thus offers the best opportunity to encode some of the universal properties of the mechanics of brittle media within a purely statistical framework. That certainly explains the superiority of this model for forecasting purposes, even in its simplified, linearized form presented in this paper.

References

- Bach, C. and Hainzl, S. (2012). Improving empirical aftershock modeling based on additional source information. *Journal of Geophysical Research: Solid Earth*, 117(B4).
- Bird, P., Jackson, D. D., Kagan, Y. Y., Kreemer, C., and Stein, R. S. (2015). GEAR1: A Global Earthquake Activity Rate Model Constructed from Geodetic Strain Rates and Smoothed Seismicity. *Bulletin of the Seismological Society of America*, 105(5):2538–2554.
- Cao, A. and Gao, S. S. (2002). Temporal variation of seismic b-values beneath northeastern japan island arc. *Geophysical research letters*, 29(9):48–1.

- Cattania, C., Hainzl, S., Wang, L., Roth, F., and Enescu, B. (2014). Propagation of coulomb stress uncertainties in physics-based aftershock models. *Journal of Geophysical Research: Solid Earth*, 119(10):7846–7864.
- Cattania, C., Werner, M. J., Marzocchi, W., Hainzl, S., Rhoades, D., Gerstenberger, M., Liukis, M., Savran, W., Christophersen, A., Helmstetter, A., et al. (2018). The forecasting skill of physics-based seismicity models during the 2010–2012 canterbury, new zealand, earthquake sequence. *Seismological Research Letters*, 89(4):1238–1250.
- Chu, A., Schoenberg, F. P., Bird, P., Jackson, D. D., and Kagan, Y. Y. (2011). Comparison of etas parameter estimates across different global tectonic zones. *Bulletin of the Seismological Society of America*, 101(5):2323–2339.
- Davidson, J., Gu, C., and Baiesi, M. (2015). Generalized Omori–Utsu law for aftershock sequences in southern california. *Geophysical Journal International*, 201(2):965–978.
- Dieterich, J. (1994). A constitutive law for rate of earthquake production and its application to earthquake clustering. *Journal of Geophysical Research: Solid Earth*, 99(B2):2601–2618.
- Eberhard, D. A., Zechar, J. D., and Wiemer, S. (2012). A prospective earthquake forecast experiment in the western pacific. *Geophysical Journal International*, 190(3):1579–1592.
- Gerstenberger, M. C., Wiemer, S., Jones, L. M., and Reasenberg, P. A. (2005). Real-time forecasts of tomorrow’s earthquakes in california. *Nature*, 435(7040):328–331.
- Guo, Y., Zhuang, J., and Zhou, S. (2015). An improved space-time etas model for inverting the rupture geometry from seismicity triggering. *Journal of Geophysical Research: Solid Earth*, 120(5):3309–3323.
- Hainzl, S. (2016a). Apparent triggering function of aftershocks resulting from rate-dependent incompleteness of earthquake catalogs. *Journal of Geophysical Research: Solid Earth*, 121(9):6499–6509.
- Hainzl, S. (2016b). Rate-dependent incompleteness of earthquake catalogs. *Seismological Research Letters*, 87(2A):337–344.
- Hainzl, S., Christophersen, A., and Enescu, B. (2008). Impact of earthquake rupture extensions on parameter estimations of point-process models. *Bulletin of the Seismological Society of America*, 98(4):2066–2072.
- Helmstetter, A., Kagan, Y. Y., and Jackson, D. D. (2006). Comparison of short-term and time-independent earthquake forecast models for southern california. *Bulletin of the Seismological Society of America*, 96(1):90–106.
- Helmstetter, A. and Sornette, D. (2002). Subcritical and supercritical regimes in epidemic models of earthquake aftershocks. *Journal of Geophysical Research: Solid Earth*, 107(B10):ESE–10.
- Helmstetter, A. and Sornette, D. (2003). Importance of direct and indirect triggered seismicity in the etas model of seismicity. *Geophys. Res. Lett.*, 30(11):doi:10.1029/2003GL017670.
- Helmstetter, A. and Werner, M. J. (2014). Adaptive smoothing of seismicity in time, space, and magnitude for time-dependent earthquake forecasts for california. *Bulletin of the Seismological Society of America*, 104(2):809–822.
- Hiemer, S. and Kamer, Y. (2016). Improved seismicity forecast with spatially varying magnitude distribution. *Seismological Research Letters*, 87(2A):327–336.

- Hiemer, S., Woessner, J., Basili, R., Danciu, L., Giardini, D., and Wiemer, S. (2014). A smoothed stochastic earthquake rate model considering seismicity and fault moment release for europe. *Geophysical Journal International*, 198(2):1159–1172.
- Jordan, T. H. (2006). Earthquake Predictability, Brick by Brick. *Seismological Research Letters*, 77(1):3–6.
- Kagan, Y. and Jackson, D. (2010). Earthquake forecasting in diverse tectonic zones of the globe. *Pure and applied geophysics*, 167(6-7):709–719.
- Kagan, Y. Y. (1992). Correlations of earthquake focal mechanisms. *Geophysical Journal International*, 110(2):305–320.
- Kagan, Y. Y. and Jackson, D. D. (2011). Global earthquake forecasts. *Geophysical Journal International*, 184(2):759–776.
- Kamer, Y., Nandan, S., Ouillon, G., Hiemer, S., and Sornette, D. (2020). Democratizing earthquake predictability research introducing the RichterX platform. *TBD*.
- Kossobokov, V. G. (2013). Earthquake prediction: 20 years of global experiment. *Natural Hazards*, 69(2):1155–1177.
- Michael, A. J., McBride, S. K., Hardebeck, J. L., Barall, M., Martinez, E., Page, M. T., van der Elst, N., Field, E. H., Milner, K. R., and Wein, A. M. (2020). Statistical seismology and communication of the USGS operational aftershock forecasts for the 30 november 2018 m w 7.1 anchorage, alaska, earthquake. *Seismological Research Letters*, 91(1):153–173.
- Mignan, A. and Woessner, J. (2012). Estimating the magnitude of completeness for earthquake catalogs. *Community Online Resource for Statistical Seismicity Analysis*, pages 1–45.
- Nandan, S. (2017). *Towards a Physics Based Epidemic Type Aftershock Sequence Model*. PhD thesis, ETH Zurich.
- Nandan, S., Ouillon, G., and Sornette, D. (2019a). Magnitude of earthquakes controls the size distribution of their triggered events. *Journal of Geophysical Research: Solid Earth*, 124(3):2762–2780.
- Nandan, S., Ouillon, G., Sornette, D., and Wiemer, S. (2019b). Forecasting the full distribution of earthquake numbers is fair, robust, and better. *Seismological Research Letters*, 90(4):1650–1659.
- Nandan, S., Ouillon, G., Sornette, D., and Wiemer, S. (2019c). Forecasting the rates of future aftershocks of all generations is essential to develop better earthquake forecast models. *Journal of Geophysical Research: Solid Earth*, 124(8):8404–8425.
- Nandan, S., Ouillon, G., Wiemer, S., and Sornette, D. (2017). Objective estimation of spatially variable parameters of epidemic type aftershock sequence model: Application to california. *Journal of Geophysical Research: Solid Earth*, 122(7):5118–5143.
- Narteau, C., Shebalin, P., and Holschneider, M. (2005). Onset of power law aftershock decay rates in southern california. *Geophysical research letters*, 32(22).
- Ogata, Y. (2011). Significant improvements of the space-time etas model for forecasting of accurate baseline seismicity. *Earth, planets and space*, 63(3):6.

- Ogata, Y., Katsura, K., Falcone, G., Nanjo, K., and Zhuang, J. (2013). Comprehensive and topical evaluations of earthquake forecasts in terms of number, time, space, and magnitude. *Bulletin of the Seismological Society of America*, 103(3):1692–1708.
- Ouillon, G. and Sornette, D. (2005). Magnitude-dependent Omori law: Theory and empirical study. *Journal of Geophysical Research: Solid Earth*, 110(B4).
- Ouillon, G., Sornette, D., and Ribeiro, E. (2009). Multifractal Omori law for earthquake triggering: new tests on the California, Japan and worldwide catalogues. *Geophysical Journal International*, 178(1):215–243.
- Page, M. T., Van Der Elst, N., Hardebeck, J., Felzer, K., and Michael, A. J. (2016). Three ingredients for improved global aftershock forecasts: Tectonic region, time-dependent catalog incompleteness, and intersequence variability. *Bulletin of the Seismological Society of America*, 106(5):2290–2301.
- Pesaresi, M., Ehrlich, D., Kemper, T., Siragusa, A., Florczyk, A. J., Freire, S., and Corbane, C. (2017). Atlas of the human planet 2017. *Global exposure to natural hazards*.
- Reasenber, P. (1985). Second-order moment of central california seismicity, 1969–1982. *Journal of Geophysical Research: Solid Earth*, 90(B7):5479–5495.
- Reverso, T., Steacy, S., and Marsan, D. (2018). A hybrid etas-coulomb approach to forecast spatiotemporal aftershock rates. *Journal of Geophysical Research: Solid Earth*, 123(11):9750–9763.
- Rhoades, D. A., Schorlemmer, D., Gerstenberger, M. C., Christophersen, A., Zechar, J. D., and Imoto, M. (2011). Efficient testing of earthquake forecasting models. *Acta Geophysica*, 59(4):728–747.
- Saichev, A., Helmstetter, A., and Sornette, D. (2005). Power-law distributions of offsprings and generation numbers in branching models of earthquake triggering. *Pure and Applied Geophysics*, 162:1113–1134.
- Saichev, A. and Sornette, D. (2006a). Power law distribution of seismic rates: theory and data. *Eur. Phys. J. B*, 49:377–401.
- Saichev, A. and Sornette, D. (2006b). Renormalization of branching models of triggered seismicity from total to observable seismicity. *Eur. Phys. J. B*, 51(3):443–459.
- Saichev, A. and Sornette, D. (2007). Power law distribution of seismic rates. *Tectonophysics*, 431:7–13.
- Schoenberg, F. P., Chu, A., and Veen, A. (2010). On the relationship between lower magnitude thresholds and bias in epidemic-type aftershock sequence parameter estimates. *Journal of Geophysical Research: Solid Earth*, 115(B4).
- Scholz, C. (1968). Microfractures, aftershocks, and seismicity. *Bulletin of the Seismological Society of America*, 58(3):1117–1130.
- Schorlemmer, D. and Gerstenberger, M. (2007). Relm testing center. *Seismological Research Letters*, 78(1):30–36.
- Schorlemmer, D., Werner, M. J., Marzocchi, W., Jordan, T. H., Ogata, Y., Jackson, D. D., Mak, S., Rhoades, D. A., Gerstenberger, M. C., Hirata, N., et al. (2018). The collaboratory for the study of earthquake predictability: achievements and priorities. *Seismological Research Letters*, 89(4):1305–1313.

- Seif, S., Mignan, A., Zechar, J. D., Werner, M. J., and Wiemer, S. (2017). Estimating η tas: The effects of truncation, missing data, and model assumptions. *Journal of Geophysical Research: Solid Earth*, 122(1):449–469.
- Shcherbakov, R., Turcotte, D. L., and Rundle, J. B. (2004). A generalized Omori’s law for earthquake aftershock decay. *Geophysical research letters*, 31(11).
- Sornette, D. (2006). Critical phenomena in natural sciences (chaos, fractals, self-organization and disorder: Concepts and tools). *Springer Series in Synergetics, Heidelberg*.
- Sornette, D. and Ouillon, G. (2005). Multifractal scaling of thermally activated rupture processes. *Physical review letters*, 94(3):038501.
- Sornette, D. and Werner, M. J. (2005a). Apparent clustering and apparent background earthquakes biased by undetected seismicity. *Journal of Geophysical Research: Solid Earth*, 110(B9, B09303, 10.1029/2005JB003621).
- Sornette, D. and Werner, M. J. (2005b). Constraints on the size of the smallest triggering earthquake from the epidemic-type aftershock sequence model, bâth’s law, and observed aftershock sequences. *Journal of Geophysical Research: Solid Earth*, 110(B8).
- Steacy, S., Gerstenberger, M., Williams, C., Rhoades, D., and Christophersen, A. (2014). A new hybrid coulomb/statistical model for forecasting aftershock rates. *Geophysical Journal International*, 196(2):918–923.
- Tinti, S. and Mulargia, F. (1987). Confidence intervals of b values for grouped magnitudes. *Bulletin of the Seismological Society of America*, 77(6):2125–2134.
- Tsai, C.-Y., Ouillon, G., and Sornette, D. (2012). New empirical tests of the multifractal Omori law for Taiwan. *Bulletin of the Seismological Society of America*, 102(5):2128–2138.
- Veen, A. and Schoenberg, F. P. (2008). Estimation of space–time branching process models in seismology using an em–type algorithm. *Journal of the American Statistical Association*, 103(482):614–624.
- Wang, Q., Jackson, D. D., and Zhuang, J. (2010). Missing links in earthquake clustering models. *Geophysical Research Letters*, 37(21).
- Werner, M. J., Helmstetter, A., Jackson, D. D., and Kagan, Y. Y. (2011). High-resolution long-term and short-term earthquake forecasts for california. *Bulletin of the Seismological Society of America*, 101(4):1630–1648.
- Zechar, J. D. and Jordan, T. H. (2010). Simple smoothed seismicity earthquake forecasts for italy. *Annals of Geophysics*, 53(3):99–105.
- Zechar, J. D., Schorlemmer, D., Werner, M. J., Gerstenberger, M. C., Rhoades, D. A., and Jordan, T. H. (2013). Regional earthquake likelihood models i: First-order results. *Bulletin of the Seismological Society of America*, 103(2A):787–798.
- Zhang, L., Werner, M. J., and Goda, K. (2020). Variability of η tas parameters in global subduction zones and applications to mainshock–aftershock hazard assessment. *Bulletin of the Seismological Society of America*.

- 862 Zhuang, J., Ogata, Y., and Vere-Jones, D. (2002). Stochastic declustering of space-time earthquake occurrences. *Journal*
 863 *of the American Statistical Association*, 97(458):369–380.
- 864 Zhuang, J., Ogata, Y., and Vere-Jones, D. (2004). Analyzing earthquake clustering features by using stochastic
 865 reconstruction. *Journal of Geophysical Research: Solid Earth*, 109(B5).