

# Towards a Harmonized Operational Earthquake Forecasting Model for Europe: Reply on RC2

Marta Han<sup>1</sup>, Leila Mizrahi<sup>1</sup>, and Stefan Wiemer<sup>1</sup>

<sup>1</sup>Swiss Seismological Service (SED), ETH Zurich

*First, my apologies for the delay of this review!*

*This paper presents preliminary candidates for operational earthquake forecasting in Europe, based on variants of the well-established Epidemic Type Aftershock Sequence (ETAS) model. The authors draw on available harmonized datasets and a background model from the European Seismic Hazard model ESHM2020 to generate the pan-European models, making good use of prior painstaking work to collate the myriad of different national catalogs and methods. The study demonstrates convincingly that the ETAS models forecast some of the observed space-time clustering and may therefore be useful for real-time deployment. The detailed model evaluations include in-sample tests, which reveal some issues in some of the models as well as some issues with the applied tests, and they include a more severe out-of-sample pseudo-prospective test, and some sober and honest analysis of the relative performance of the models. They conclude that indeed 1-2 of the proposed ETAS models perform best and could serve as initial candidates (to be improved in the future).*

*Please see my (many) comments below. It's a well written and structured paper – and I have only a couple of major comments, everything else is largely around presentation and suggestions for improvements. I have one request to assess a forecast/simulation choice in more detail, and another to pursue some more analysis in the evaluation.*

[We are grateful to the reviewer for the time invested in reviewing the study in a detailed and constructive manner and for providing insightful suggestions to improve the manuscript. We acknowledge that steps outlining both the major and minor changes to the manuscript are needed to address the current shortcomings.](#)

*What is the influence of the chosen minimum probability level in histograms where no simulations have filled bins? How many number bins of the full space-time (and magnitude?) bins are set to this level? How frequently are they*

“observed”? How frequently do these empty number bins sit next to filled bins, suggesting perhaps that interpolation would be a better approach?

Thank you for this comment. Many bins are indeed empty – the idea behind Fig. 5 in the original manuscript was to show that total information gain is dominated by cells where events do occur rather than ones where the event count is zero. We perform an additional analysis of how the choice of how to deal with empty bins affects the results, and add the relevant details to the revised manuscript.

*The authors rightly identify some undesirable features in the retrospective hypothesis tests applied, which is a useful contribution. But that means they should pursue other/additional methods, e.g. comparing the cumulative number trajectories of observed and synthetic catalogs (see specific comment below), comparing (at least visually) the spatial distribution of the observed and synthetic catalogs, etc. Similarly, Figure 5 of the pseudo-prospective information gains don't seem very insightful – this is a first step to identify what's gone wrong/right, but would really benefit from some analysis, e.g. are the aftershocks not well captured in space, time or both? Why are the ETAS models sometimes performing worse than the ESHM2020 model? Even qualitative visual analysis would be useful.*

In the revised manuscript, we address this point in two steps: adding a plot comparing the cumulative number of events in Sect. 4 (Results), and performing more analyses to get a better understanding of the models' performance comparison in space and time, the methodology of which we add to Sect. 3 and the output to Sect. 4.

*Otherwise, please consider the below comments as hopefully constructive suggestions.*

## **Abstract**

*L2 – is the purpose an aftershock forecast model? It sounds like it based on the first sentence. Please clarify.*

*The abstract would benefit from a sentence on the data the model is fit to, and any issues with “the” European catalog. Ie – what's harmonized?*

*It would help to clarify that these are currently 2D seismicity forecasts, not extended ruptures with depth (which might be fruitful areas for the future).*

*A clearer statement of the different models explored and their rankings would be of interest.*

*L10 How do the findings highlight these promising areas (last sentence)? Can you express it concisely or are these speculations or a wish list?*

The goal is to develop a ‘full’ earthquake forecasting model, i.e., one that produces forecasts continuously. With this sentence, we meant to express the fact that the main component of clustering is due to aftershocks, but we change

this sentence to avoid confusion. We amend the abstract as allowed by the word limit (that we are already close to) to address the listed points.

## Introduction

*L60 – missing a word after  $m_c$  – maybe “more”?*

*The introduction could probably be shortened – e.g. the background to  $m_c$  and the GR law etc seems more relevant for a thesis than a paper – unless the authors are specifically trying to reach a broader audience.*

*L75: As I understand it, the Mizrahi et al. (2021) model accounts for temporal – not spatial – variations in  $m_c$ . Is the model being extended? If so, mention it in the abstract and here.*

*L81: ESHM20 contains many models, some idea of how that’s incorporated into ETAS in both the abstract as well as more discussion here would be appropriate.*

We address these points in the Introduction of the revised manuscript. The Mizrahi et al. (2021) model is based on the approximation that  $m_c$  is locally constant and that observed seismicity at any point is representative of the unobserved seismicity below  $m_c$ . Although in their study, they only explore temporally varying  $m_c$ , no extension of the model is needed to apply it with space-time varying  $m_c$ .

*L112: These statistics of  $m_c$  presumably come from Danciu et al. (2021)? Please clarify.*

Yes. We are adding the clarification in the revised version of the manuscript.

*Figure 1: the red dots look black because of the marker edge colour. It’s impossible to differentiate red from green dots in this figure. Try a bigger size or different/no edge colour.*

Thank you for pointing this out, we will amend the figure.

*L137: what do you mean here by “such differences”? Do you mean a difference in completeness magnitude, or a difference in the spatial locations? The second is indeed expected, the former perhaps less so. Please clarify.*

We were referring to differences both in composition methods and content of new catalogs. We recognize that changes in completeness would normally not be an expected occurrence, especially that it increases with time and rephrase this statement. However, given that catalog is composed of ‘smaller’ catalogs issued by local agencies, it is not unreasonable to expect at least one of the subregions will differ in properties compared to the training period at any given point in the testing period (especially in ‘real-time’ deployment when homogenized data may not be available).

*L140: Good to clarify. Does the binning require the discrete version of the  $b$ -value estimation procedure? Which one was used in Fig 1d? What are the uncertainties? Why show  $b=0.99$ ? Perhaps this will be discussed below, then guide the reader.*

Yes, and we use this estimation procedure introduced in Tinti and Mulargia (1987). This is mentioned later, but we agree it should also be included when introducing the binning and the figure. The two  $b$ -values shown in the plot are one used by ETAS variants later, which we also explain later, but now also add at this point.

*L142: I would appreciate a figure of the long-term rates in these two models.*

The figure is added in the revised manuscript, either in the main manuscript or as supplementary material.

*L146-153: this sounds like Methods, rather than data. I recommend more discussion of the two long-term models (ie a figure of the two models and a bit more background on both models). Then move the discussion of how you map these into the background rate of ETAS into a subsection in Methods, including some discussion of how this long-term seismicity might or might not be compatible with the background (independent) rate in an ETAS model would be appropriate (or else just extend this section).*

While this data requires some preprocessing, the procedure itself is not a method we develop, but a product of ESHM20 (Danciu et al., 2021) used here as input. We agree that the discussion would benefit from more details and we address this point in the revised manuscript.

*L187: again, if the ESHM20 considers long-term rate as the declustered rate, there is some consistency (even that's debatable), but if the ESHM20 rate is the long-term, i.e. average rate, then setting this to the background rate may lead to an overestimation of the total rate (instead the average model rate should equal the ESHM20 rate, see e.g. Field et al. (2017)). Or are you solely considering the relative spatial information, rather than the absolute rate? I think clarification of this point earlier in the paper would help. (You clarify this in L242, and I agree this is a good approach, but I'd mention it earlier.)*

We are indeed considering only the relative spatial information, exactly because it would lead to the overestimated background rate otherwise. As suggested, we clarify this earlier in the revised version.

*L193: ETAS<sub>USGS</sub>: A few more details would help – isn't this an aftershock-only model? Are these truly one-size fits all for the entire globe or are they regionalized?*

The true USGS software uses parameters specific to tectonic regimes, which were obtained by fitting global data from the respective tectonic regimes, but they also have a set of parameters representing the global average, which we employ for aftershock modelling, while using ‘our’ background seismicity inferred by ETAS<sub>0</sub>. We add this additional explanation to the revised manuscript.

*L203: burn period → burn-in period*

Thank you for noticing - corrected.

*L213: which events’ locations exactly – the new/simulated background events, or do you not include all background events in the burn-in period?*

The locations of simulated background events in the period we are simulating (the burn-in period is only used to simulate aftershocks of the events from that period that occur in the period we are simulating). Clarification added in the sentence.

*L258: It would help to a bit more precise here, and to consider the role of the maximum magnitude or a tapered GR law. “Exploding” aftershock behavior may occur with some finite probability when the branching ratio  $r$  is  $> 1$ , and sequences are finite with probability 1 when  $r = 1$ . The calculation of  $r$  involves the GR law, and thus depends on your choice of pure GR vs tapered/cut GR, and thus in the latter case on  $M_{\text{corner}}/M_{\text{max}}$ . First, please include a short discussion/justification on your choice earlier when you introduce the magnitude distribution (including the potential for spatial variations). Second, in the context of constraining parameters for subcritical branching, these modifications matter and thus the parameter constraints change. The branching ratio  $r$  may well be  $< 1$  if the GR is tapered/cut while it may not with a pure GR law. Sornette and Werner (2005) derived  $r$  equations for a truncated GR. Finally, clarify what ‘ $e$ ’ is (another variable or Euler’s number?).*

We take into account the maximum magnitude when calculating the branching ratio in cases when  $\alpha - \rho\gamma > \beta$ . Using more detailed assessments of maximum magnitudes per region and spatial variations in magnitude distributions in general provided by ESHM20 might help with the explosive behaviour, however, we assign most of this effect to the tapered exponential temporal kernel. In ETAS<sub>USGS</sub>, we set  $\alpha = \beta$  with no taper in temporal kernel, and simulations not only converge, but underestimate the number of events in retrospective CSEP tests, due to a significant portion of aftershocks occurring after the observed period, resulting in a lower ‘effective’ branching ratio lower than the total one higher than 1. We clarify this point further in the revised version of the manuscript. We address other points as well and will add details about this to the revised manuscript.

*L264: “standard” in the USGS software?*

The formulation we are referring to is indeed one used by USGS, by ‘standard’ formulation we mean one that uses  $(t+c)^{-p}$  rather than  $(t+c)^{-(1+\omega)}$  (and similar for spatial and productivity) that we use, which is a bit less common in literature. We rephrase for clarity.

*L286: Does this mean that the explicitly spatially-variable models forecast spatially uniform background rates in each spatial ESHM zone? I don’t believe so, but please clarify – if I understand correctly, the  $ETAS_{bg}$  models use the same procedure to simulate background events as the  $ETAS_0$  model, but both the probability of being a background rate has changed because of different parameters  $\mathcal{E}$  the additional constraint that the rates in different zones are relatively constrained. Is that correct?*

The procedure is different, the models with spatially-varying ESHM-informed background rate do simulate locations uniformly within each grid cell ( $0.1^\circ\text{lat} \times 0.1^\circ\text{lon}$ ), the total number of background events is distributed according to the ETAS inversion output, but their relative locations are enforced to follow the ESHM input. We clarify this in the revised manuscript.

*L292: Ah, so it is a truncated GR law – does this influence the branching ratio constraint  $\alpha = e$ ? See earlier comment.*

We mention the truncated GR law earlier in the article when the previous question arose.

*L295: “true”  $\rightarrow$  observed? Here and below.*

Yes, clarified in the revised manuscript.

*L320: Why only focus on retrospective consistency testing? 7 years of daily out-of-sample testing seems like a great start, even if not 25 years. I recommend doing these or similar tests, or some of them – the information gain assess a different aspect of the models, but tell us less about how close to the data the models are (and while the non-Poisson cell-wise LL scores are a great improvement, they still neglect known correlations between spatial cells).*

Thank you for this suggestion. In response to other comments, also of the other reviewer, we perform additional analyses both for the retrospective and pseudo-prospective tests and add the relevant findings to the revised manuscript.

*L338: What is the influence of this implementation and the particular choice of a minimum waterlevel on the overall scores?: how many zero-bins are there in a 100k simulation set per day? How does the waterlevel compare with the background rate (and its Poisson process implied probabilities)? How “rough” is the zero-bin distribution, ie could an interpolation procedure approximate the ETAS forecasts better than this water-level?*

We add this analysis in the revised manuscript, see also our response to the previous (major) comment.

*L419: It would be illustrative to give the range of values in days or years after which the Omori law is significantly tapered in this formulation (ie give the range of tau in days). (I see the range in L436, but give units in years).*

The ranges are given in more interpretable units in the revised version of the manuscript.

*Parameters: Can you connect your discussion of the ETAS parameter estimation with the literature more directly (e.g. do you see similar patterns to Seif et al. (2017), and others...)?*

We extend and clarify the parameter inversion output discussion in the revised manuscript.

*Do you have LL scores for your various ETAS model fits? Can you compare the in-sample fit using LL and AIC? It's a useful indicator of fit, though not necessarily about predictive skill.*

We calculate these values and will add them to the manuscript if it turns out to add valuable information to the study.

*Figure 3: Please clarify these are simulations over the entire training period. Also, L279 states 100k catalogs were simulated, but here it states only 10k. I would clarify the timeframe of the retrospective testing also in the relevant Methods section around L279. These are not next-day forecasts, as in the pseudo-prospective testing section.*

Thank you for spotting the inconsistency - the suggested changes and corrections are done in the revised version.

*Figure 3a: I would recommend looking at the cumulative counts in each of the catalogs and compare with the observed count. You could look at the cumulative 95% model range and compare with the observed number. But you also get important insights into how well the model is qualitatively and quantitatively reproducing the features you are hoping it will, namely aftershock clustering and background rate.*

*Fig3: As part of these formal tests, I would also look at visual fits, e.g. Fig 1d seems to show a comparison of observed data and the magnitude forecasts from the different b-value models. I know the formal tests do this comparison quantitatively, but I would discuss the visual fit (which quite clearly favours the higher b-value).*

We are adding the figure showing cumulative event counts comparison in the revised manuscript. We also acknowledge the need to compare the magnitude distributions visually and agree on the assessment that the higher  $b$ -value estimate looks more correct.

*Supercritical branching when alpha is fixed: these branching ratios are indeed very high, but supercritical has been found before, e.g. Seif et al. (2017) and some of the other references you used have found similar issues. There's clearly model misspecification that's driving this bias, so it would be interesting to discuss possibilities: changing background rates, anisotropic spatial aftershock triggering, spatial variation in incompleteness?*

Thank you for the suggestion - we are extending this discussion.

*L476: "suspiciously well" in spatial terms means that the models are too smooth, ie that the events occur too close to likelihood peaks without the scatter expected if the model were the data generator. That should be spelt out.*

We agree and amend the phrasing.

*L479: refer to Fig 1d in this discussion of magnitude distribution fit to data.*

We agree this is an additional argument for the  $b$ -positive ETAS variants and add it to the discussion.

*L481: is it true that "ETAS<sub>bg</sub><sup>b+</sup> uses the same spatial distribution for placing the background events" as ETAS<sub>0</sub><sup>b+</sup>? As commented above, I understood two differences: the relative rates between zones are constrained, and the absolute rate is different, and the parameters are different, so the probability of being a background event is also affected. Could you clarify (here and above where commented)?*

*Irrespectively, and again, I would visually compare the forecast and observed magnitude distributions and use this to support your case that the difference in magnitude forecast performance of the two models is indeed a surprising result and indicates an issue with the M-test (and the S-test, although I'm not sure I understand your detailed argument here, see above).*

They indeed don't use the same spatial distribution of background events. This paragraph is corrected in the revision. We add the magnitude distribution comparisons to the discussion as mentioned in earlier comments, too.

*L482: "known issue": as you know, Francesco Serafini and others identified a correlation between the M-test and the N-test, which is indeed problematic and being fixed. You might cite Francesco Serafini (personal communication, 2024), or perhaps refer to the manuscript in prep that might be submitted by the time this is published.*



*Given the (reasonable) caution you advise in interpreting these M/S/PL results, you might use visual checks, including of the spatial distributions, appropriately averaged (or not) over the many simulations, to show the extent to which models reproduce the features you expect them to.*

We agree that visualizations would help assess the spatial and magnitude distributions. As described in our responses to previous comments, we will perform these additional analyses and include relevant findings either in the main manuscript or as supplementary material. We also agree that the study being conducted by Serafini et al. is highly relevant to this discussion and cite it in the revised manuscript.

*Figure 4a would benefit from a panel underneath with the magnitude-time plot of the seismicity, which will help identify clustering and quieter periods and visually link them to the likelihood gains (or losses).*

*Figure 4b/c: these tables are hard to digest, despite the nice colors. Could you please plot (instead or in addition) a figure instead showing mean information gain over the ESHM20 model with 95% range (which I believe you get from the paired t-test?) to indicate significance?*

We recognize these suggestions could improve the readability of the figure and will try out alternative visualisations for the updated version of the manuscript.

*L506: “achieve a significance level below 0.05” is I believe not the right interpretation of these tests – you might just state that their p-values are below the critical value.*

True, we rephrase this statement.

*L515: “could be expected” – yes, of course, but the purpose is to develop a model that does indeed use this knowledge and puts this into practice. It might be a basic statement for some of the community, but not for others and potential users. I would emphasise the success in finding one or several models that do show substantial improvement in predictive skill, and that this skill shows up during periods of clustering (presumably, although it would be good to illustrate/visualize this explicitly, as suggested eg via the modified figure 4a that links cumulative LL scores trajectories to seismicity magnitude-time plots to show where the clustering occurs). And I would be careful with absolute statements like “poor performance” – this is relative to the other ETAS models as measured by your metric, which is (a) an approximation of the correlation of seismicity between spatial bins and (b) requires many model simulations, which may not be sufficient to fill all “bins” with simulations (see waterlevel comment). So I’d be a bit more careful and emphasise this first milestone – a European ETAS model that does what it says on the tin: forecast aftershocks and capture some of the time dependence of seismicity, clearly much better than a time-independent*

model. And yes, there are some unexpected results here, which would be good to understand in more depth in order to improve on this first attempt.

We appreciate this insightful summary of the output of the study and rephrase the conclusion according to the given suggestions.

With respect to the spatial LL comparison in Fig 5, I suspect there are patterns, and some more analysis would be required to understand these patterns. It's surprising to see such strongly negative gains between an ETAS model and the ESHM20 model, for instance. Are these individual events? Clusters gone awry? And what's going on at the mid-oceanic ridge? This section is a bit thin, so at least some more discussion would help. What explains the major differences in some cells between the two ETAS versions? Or at least what happens there?

We will look into this in more detail and provide the relevant outcomes in this section in the revised manuscript.

## References

- Danciu, L., Nandan, S., Reyes, C. G., Basili, R., Weatherill, G., Beauval, C., Rovida, A., Vilanova, S., Sesetyan, K., & Bard, P.-Y. (2021). The 2020 update of the European Seismic Hazard Model-ESHM20: Model Overview [Publisher: ETH Zurich]. *EFEHR Technical Report, 1*.
- Field, E. H., Milner, K. R., Hardebeck, J. L., Page, M. T., van der Elst, N., Jordan, T. H., Michael, A. J., Shaw, B. E., & Werner, M. J. (2017). A spatiotemporal clustering model for the third Uniform California Earthquake Rupture Forecast (UCERF3-ETAS): Toward an operational earthquake forecast. *Bulletin of the Seismological Society of America, 107*(3), 1049–1081.
- Mizrahi, L., Nandan, S., & Wiemer, S. (2021). Embracing data incompleteness for better earthquake forecasting [Publisher: Wiley Online Library]. *Journal of Geophysical Research: Solid Earth, 126*(12), e2021JB022379.
- Seif, S., Mignan, A., Zechar, J. D., Werner, M. J., & Wiemer, S. (2017). Estimating ETAS: The effects of truncation, missing data, and model assumptions. *Journal of Geophysical Research: Solid Earth, 122*(1), 449–469. <https://doi.org/10.1002/2016JB012809>
- Sornette, D., & Werner, M. J. (2005). Apparent clustering and apparent background earthquakes biased by undetected seismicity. *Journal of Geophysical Research: Solid Earth, 110*(B9), 2005JB003621. <https://doi.org/10.1029/2005JB003621>
- Tinti, S., & Mulargia, F. (1987). Confidence intervals of b values for grouped magnitudes. *Bulletin of the Seismological Society of America, 77*(6), 2125–2134. <https://doi.org/10.1785/BSSA0770062125>