

Simple smoothed seismicity earthquake forecasts for Italy

J. Douglas Zechar^{1,2,*} and Thomas H. Jordan³

¹ ETH Zurich, Swiss Seismological Service, Zurich, Switzerland

² Lamont-Doherty Earth Observatory, Columbia University, Palisades, NY, USA

³ University of Southern California, Department of Earth Sciences, Los Angeles, USA

Article history

Received February 18, 2010; accepted July 23, 2010.

Subject classification:

Smoothed seismicity, Earthquake predictability, Forecast optimization, Area skill score.

ABSTRACT

Several earthquake forecast experiments in Italy have been initiated within the European testing center of the global Collaboratory for the Study of Earthquake Predictability. In preparation for these experiments, we developed space-rate-magnitude forecasts based on a simple model that incorporates spatial clustering of seismicity. This model, which we call the Simple Smoothed Seismicity model (TripleS), has a minimal number of free parameters and is based on very few assumptions; therefore, it can be considered as a model of least information with which others can be compared. The fundamental TripleS parameter controls the spatial extent of the smoothing, and we selected its value based on an optimization procedure that was applied to retrospective forecast experiments. In this article, we present the motivation for developing TripleS and describe the construction of forecasts for Italian seismicity. We also discuss the research questions that remain to be answered with respect to TripleS and, more generally, the smoothed seismicity approach to earthquake forecasting.

Introduction

Evaluation of earthquake forecast experiments requires careful consideration of the appropriate reference models. For example, a forecast that incorporates some form of spatial clustering is likely to perform significantly better than a reference forecast that is based on a spatially uniform seismicity model. Therefore, a reference model should capture the clustering that is observed in seismicity [Michael 1997]. One straightforward approach to incorporate this clustering is to smooth past seismicity, thereby allowing each past earthquake to provide some smoothed contribution to an estimate of the seismic rate density. Kernel smoothing is a commonly used technique for density estimation, and the objective of forecast experiments in the Collaboratory for the Study of Earthquake Predictability (CSEP) – namely, to accurately forecast the space-rate-magnitude distribution of earthquakes – is similar to the estimation of a predictive density of the seismicity rate.

Many previous studies have applied smoothed seismicity approaches that have incorporated various levels of complexity [e.g., Kagan and Jackson 1994, Frankel 1995, Jackson and Kagan 1999, Kagan and Jackson 2000, Stirling et al. 2002, Stock and Smith 2002a, Stock and Smith 2002b, Helmstetter et al. 2006, Helmstetter et al. 2007, Kagan et al. 2007, Petersen et al. 2008, Werner et al. submitted BSSA]. Several of these analyses were made as part of regional and national seismic hazard mapping projects, and the Global Earthquake Model (<http://www.globalquakemodel.org>) will probably include a smoothed component. In this study, we considered a very basic model, which we call the Simple Smoothed Seismicity model (TripleS), that applies an isotropic Gaussian smoothing to the locations of past earthquakes to forecast the density of future seismicity. TripleS is characterized by one parameter: the smoothing distance, σ , which is equivalent to the standard deviation of a one-dimensional Gaussian distribution.

Zechar and Jordan [2008] introduced a performance metric, the area skill score, which can be used for evaluating earthquake forecasts. When an area skill score test is performed, a reference model is explicitly specified in terms of an alarm function, which is a general format for ranking regions of space/time/magnitude in terms of the probability that a future target earthquake will occur in a given region. A wide range of earthquake forecasts, including all of those participating in CSEP Italy experiments, can be interpreted in terms of alarm functions, and this interpretation allows for rigorous comparative testing. Given that there is little convincing evidence that complex prediction algorithms substantially outperform very crude smoothed seismicity models, we decided to essentially invert the testing problem: rather than building increasingly complex models to forecast seismicity, can we optimize a simple smoothed seismicity model? In this study, we used retrospective forecast experiments to optimize the smoothing distance with respect to the area skill score.

Using TripleS, we constructed and submitted forecasts for the CSEP Italy five-year and ten-year experiments, and we submitted codes for inclusion in the three-month experiment [Schorlemmer et al. 2010]. Because, at the time of writing, the three-month forecasts have not yet been computed, we emphasize here the five- and ten-year forecasts, although the methods are the same for each experiment.

Data

In preparation for the CSEP Italy experiment, two earthquake catalogs were analyzed and summarized for the participants: the parametric catalog of Italian earthquakes (Catalogo Parametrico dei Terremoti Italiani, CPTI) and the catalog of Italian seismicity (Catalogo della Sismicità Italiana, CSI) [Schorlemmer et al. 2010]. The CPTI contains information on earthquakes that occurred in Italy between 1901 and 2006. (All date ranges in this section are inclusive.) The earthquake magnitudes are specified in terms of the moment magnitudes, and it is believed that the CPTI contains all of the earthquakes with a moment magnitude (M_w) of at least 4.8 [Schorlemmer et al. 2010]. The CSI reports on the earthquakes in Italy that occurred from 1981 to 2002. The CSI uses a local magnitude scale (M_L), and is believed to contain all events with a M_L of at least 2.8, although the level of completeness might have changed substantially as the network evolved [Schorlemmer et al. 2010]. The CPTI and CSI are not the only catalogs that report earthquakes in Italy. For example, the Earthquake Catalog of Switzerland (ECoS) [Faeh et al. 2003] provides information for events that happened within 200-300 km of the border between Italy and Switzerland; this catalog is thought to be as complete as the CSI (J. Woessner, personal communication).

The fundamental hypothesis that motivates a smoothed seismicity approach to earthquake forecasting is that the locations of past earthquakes provide information about the locations of future earthquakes – in other words, «the future will most of all resemble the past» [Marcus 2006]. However, it is not obvious that all past earthquakes provide equal predictive information, and it has been suggested that the inclusion of smaller earthquakes increases the predictive skill of earthquake forecasts [Helmstetter et al. 2006, Helmstetter et al. 2007, Werner et al. submitted BSSA]. In the CSEP Italy experiment, there is a clear trade-off between using the CSI and the CPTI: the CSI includes smaller events, but only over a relatively limited time period, and while the CPTI provides information on events over a much longer time period, it only includes the very largest events during this time. Rather than arbitrarily choosing one catalog as the most appropriate for smoothing, we generated three forecasts, each of which was based on a different catalog: 1) the CSI; 2) the CPTI; and 3) a hybrid catalog that comprises the CPTI events from 1901 to 1980 and from April 16, 2005 to 2006, plus the CSI events from 1981 to 2002, plus the ECoS events from 2003 to April 15, 2005.

We note that the ECoS does not cover the entire CSEP Italy testing region, and therefore its use might have caused some slight bias towards higher apparent seismicity rates in northern Italy. However, we only used the ECoS catalog for 27.5 months of data in a combined catalog that covers 1,272 months, and the potential bias would not affect the entire testing region. Additionally, this bias is likely to be much smaller than the Poisson forecast uncertainty imposed on each bin.

We filtered the catalogs to remove earthquakes outside the latitude, longitude, and depth range of 34.3° N to 49.2° N, 4.3° E to 20.9° E, and 0 km to 30 km, respectively, although we retained events with unconstrained depths (of which there are many in the CPTI). This region contains the CSEP Italy testing region and extends beyond it by at least $\sim 1^\circ$ in latitude/longitude to reduce edge effects when smoothing. (See Schorlemmer et al. [2010] for the CSEP Italy grid definition.) To reduce bias caused by fluctuations in data quality (in space and time), we further filtered the catalogs to remove all of the earthquakes below the above-mentioned respective completeness magnitudes.

After the initial submission of our forecasts, and based on the analysis of Werner et al. [2010], it became obvious that the M_w values in the CPTI had not been converted to the M_L values required for evaluation with respect to the CSI. This error for the forecasts based on the CPTI and the hybrid catalog was consequently corrected, and our forecasts were resubmitted; the magnitude conversion suggested by MPS Working Group [2004] was used, which is based on regression analysis between these CPTI and CSI magnitudes: $M_L = (M_w - 1.145) / 0.812$. In this article, we present the results for these corrected forecasts, rather than those in the initial submissions.

Methods

Perhaps the simplest method for smoothing seismicity is to discretize the region of interest and to count the number of earthquakes that have occurred in each cell, allowing each point-source epicenter to be smoothed over the cell into which it falls (Rundle et al. [2002] called this «the relative-intensity method»). However, the results of such smoothing are not stable with respect to changes of the discretization parameters, i.e., the cell size and grid alignment, and the smoothing itself is anisotropic and un-physical. For example, epicenters occurring in opposite corners of a cell are treated as if they both occurred in the center of the cell. To minimize these effects, and to relax the constraint that each epicenter contributes only to the cell in which it occurs, we smoothed the earthquakes using a continuous kernel function that allowed for a wider region of influence. For TripleS, we chose a two-dimensional isotropic Gaussian smoothing kernel governed by a single parameter. The primary computational task was to calculate the contribution of an earthquake epicenter at (x_{eqk}, y_{eqk}) to a grid cell with bounds x_1, x_2, y_1 and y_2 . In general, for a bivariate smoothing kernel

$K(x, y)$, this contribution takes the form of Equation (1):

$$K(x_{eqk}, y_{eqk}, x_1, x_2, y_1, y_2) = \int_{y_1}^{y_2} \int_{x_1}^{x_2} K(x_{eqk}, y_{eqk}) dx dy. \quad (1)$$

Gaussian kernels have been used in several previous studies of smoothed seismicity [Stock and Smith 2002a, Stock and Smith 2002b, Helmstetter et al. 2006, Helmstetter et al. 2007], and they form the basis for modeling «background» seismicity in the national seismic hazard maps of the United States of America [Petersen et al. 2008] and New Zealand [Stirling et al. 2002]. Frankel [1995] first suggested Gaussian smoothing for seismic hazard, although he formulated the kernel in terms of a correlation distance c that differs from the pure Gaussian form, so we here explicitly describe our formulation. We denote our isotropic Gaussian kernel, K_σ , as in Equation (2):

$$K_\sigma(x, y) = \frac{1}{2\pi\sigma^2} \exp\left[-\frac{x^2 + y^2}{2\sigma^2}\right] \quad (2)$$

where σ is the smoothing distance. Upon substituting Equation (2) into Equation (1) and solving the integral, we find the Gaussian contribution of an epicenter to a cell, as Equation (3):

$$K_\sigma(x_{eqk}, y_{eqk}, x_1, x_2, y_1, y_2) = \frac{1}{4} \left[\operatorname{erf}\left[\frac{x_{eqk} - x_2}{\sigma\sqrt{2}}\right] - \operatorname{erf}\left[\frac{x_{eqk} - x_1}{\sigma\sqrt{2}}\right] \right] \left[\operatorname{erf}\left[\frac{y_{eqk} - y_2}{\sigma\sqrt{2}}\right] - \operatorname{erf}\left[\frac{y_{eqk} - y_1}{\sigma\sqrt{2}}\right] \right] \quad (3)$$

Analytically, the kernel described by Equation (2) extends infinitely in both dimensions. However, within the precision provided by modern computers, $\operatorname{erf}(a) \sim \operatorname{sgn}(a)$ when $|a| > 5.92$. Therefore, any cell satisfying the conditions of Equation (4) will cause Equation (3) to equal zero – i.e., the cell will receive no contribution from the epicenter.

$$\begin{aligned} |x_{eqk} - x_2| > 5.92\sigma\sqrt{2} \quad \text{and} \quad |x_{eqk} - x_1| > 5.92\sigma\sqrt{2} \\ \text{or} \\ |y_{eqk} - y_2| > 5.92\sigma\sqrt{2} \quad \text{and} \quad |y_{eqk} - y_1| > 5.92\sigma\sqrt{2} \end{aligned} \quad (4)$$

To reduce the computation time, we determined which cells would obtain some contribution from the epicenter rather than evaluating Equation (3) everywhere in the gridded region. A cell with bounds x_1, x_2, y_1 , and y_2 will obtain some contribution from an epicenter at (x_{eqk}, y_{eqk}) only if the conditions of Equation (5) are met:

$$\begin{cases} x_{eqk} - x_1 < 5.92\sigma\sqrt{2} \\ x_{eqk} - x_2 > -5.92\sigma\sqrt{2} \\ y_{eqk} - y_1 < 5.92\sigma\sqrt{2} \\ y_{eqk} - y_2 > -5.92\sigma\sqrt{2} \end{cases} \quad (5)$$

Particularly for small σ and large catalogs, the use of these conditions provides for a substantial reduction in the

computing time.

Rather than arbitrarily choosing a fixed value of σ for TripleS forecasts, we used retrospective experiments and the optimization procedure outlined by Zechar [2008]. For each experiment, we considered the CSEP Italy testing region with $0.1^\circ \times 0.1^\circ$ latitude by longitude cell sizes, and we explored the smoothing that resulted from each of the smoothing distance values given in Equation (6):

$$\Sigma = \{5, 10, 15, 20, 25, 50, 75, 100, 200\} \text{ km} \quad (6)$$

To build the five-year forecasts, we designated the most recent five years as the target period, and the earthquakes in this period with magnitudes of at least 4.95 were the target events. For the ten-year forecasts, we used the earthquakes with magnitudes of at least 4.95 in the most recent ten years as the target earthquakes. We smoothed the locations of all events that occurred before the target period, and we repeated this process for each candidate value of σ in Σ . This smoothing process resulted in several estimates of seismic density, each of which was an alarm function. These alarm functions were used as forecasts of the target events, and for each candidate value of σ , the average misfit of the area skill score was computed. These were determined by selecting one of the alarm functions as the reference model and computing the area skill score for all of the alarm functions with respect to this reference. Because no reference model was preferred *a priori*, we used a round-robin procedure, and therefore each candidate value of σ had its own average misfit. For additional details of the average misfit calculation, we refer readers to Zechar [2008, p. 94].

We considered the smoothing distance that yields the minimum average misfit, σ^* , to be optimal. To construct the prospective forecasts for the CSEP Italy experiment, we smoothed the locations of all of the events in the complete catalog and thereby obtained a spatial distribution. Because these CSEP Italy experiments required space-rate-magnitude forecasts, the spatial distribution was scaled such that the overall forecast rate matched the average rate of target events in the catalog, and we used an untapered Gutenberg-Richter distribution with a b -value of unity to form the magnitude distribution.

Results

Figure 1 shows the results of the TripleS retrospective optimization procedure in terms of the average misfit as a function of the smoothing distance. Note that in all of the plots the average misfit values form a convex set (i.e., the local minimum is also the absolute minimum). We interpret this phenomenon as representing a transition from too little smoothing at small smoothing distances to too much smoothing at large smoothing distances. It is not clear whether the optimal smoothing distance can be related to any

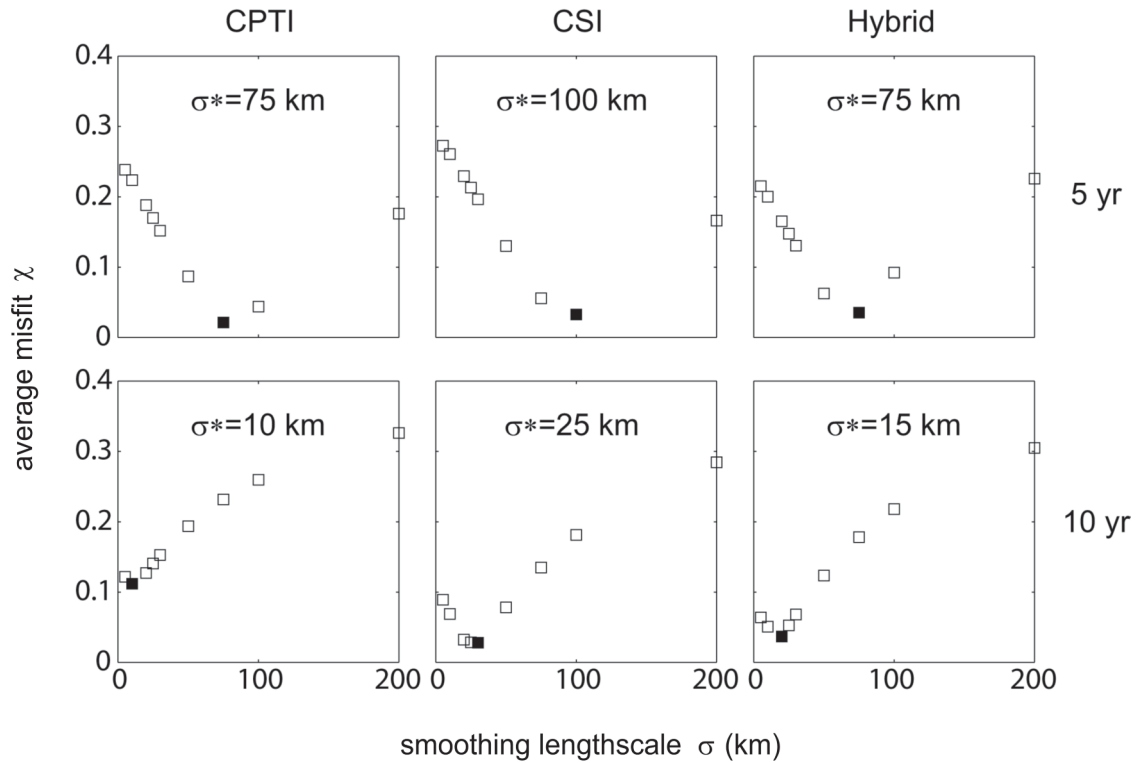


Figure 1. Plots of average misfit, χ , as a function of kernel bandwidth, σ , for the retrospective five-year and ten-year tests using the CPTI, CSI, and hybrid catalogs, as indicated. The filled square in each plot shows the minimum average misfit, with the corresponding optimal smoothing bandwidth, σ^* , also given.

particular physical or geological characteristic; it probably depends upon the minimum input magnitude, the minimum target magnitude, the particular target period, and the configuration of the seismic network, among other characteristics [Zechar 2008]. Certainly, if the spatial seismicity distribution were stationary and the catalog sampled the distribution fully, σ^* would be infinitesimal, but it is unlikely that both of those conditions were met in this study.

In constructing TripleS space-rate-magnitude forecasts, we computed the historical average rate of target earthquakes in the testing region for each catalog: 1.74 events per year in the CPTI, 1.32 events per year in the CSI, and 1.69 events per year in the hybrid catalog. Note that these rates are for the rectangular prism region described in the Data section; therefore, the rates in the CSEP Italy testing region – and thus our overall annual forecast rates – are slightly lower.

Figure 2 shows a map view of the corrected, submitted TripleS five-year and ten-year forecasts for the CSEP Italy experiment. Each five-year forecast is much smoother than its ten-year counterpart because of the larger values of σ^* in the retrospective five-year experiments. The large discrepancies of the σ^* values are directly related to the samples upon which each forecast was optimized. In the electronic supplement (<http://www.annalsofgeophysics.eu/index.php/annals/rt/suppFiles/4845/0>), we include map-view animations of these optimizations; these animations show that the target earthquakes of the most recent five years were substantially less clustered near regions of past high seismicity than the target earthquakes of the most recent ten years. Therefore,

larger smoothing distances were optimal for the five-year experiments. To check that these discrepancies were not caused by the area skill score metric itself, the optimization experiments were repeated using the spatial likelihood of Zechar et al. [2010], and similar discrepancies were found; this provides further support for our claim that the differences are due to the particular samples.

Overall, the forecasts are similar, and all of the forecasts suggest high rates throughout central Italy and, to a lesser extent, in Sicily. Nevertheless, we see large differences between the ten-year forecasts in northeast Sicily, and the forecasts based on the CSI tend to lack some of the high seismicity rates in the northeast region of the Italian mainland. The initial three-month forecasts cannot be shown because they had not been computed at the time of writing; the automated TripleS codes will generate the three-month forecasts at the beginning of the first experiment, and the forecasts and associated evaluations will be archived at <http://www.cseptest.org>. For qualitative and quantitative comparisons of TripleS forecasts with other forecasts in the CSEP Italy experiments, the reader is referred to the retrospective analyses by Werner et al. [2010], bearing in mind that the uncorrected TripleS forecasts were considered in that study.

Discussion and conclusion

During the CSEP Italy experiments, the TripleS forecasts described in this article will be compared with several other forecasts, many of which derive from more complex models that incorporate multiple hypotheses regarding earthquake

occurrence [Schorlemmer et al. 2010, and references therein]. Because TripleS is relatively simple and includes fairly standard assumptions, only a few explanations are possible if another forecast demonstrates greater predictive skill than TripleS in the prospective experiments: 1) the superior forecast more accurately represents spatial clustering of seismicity; 2) the superior forecast ignores spatial clustering but incorporates information that is more predictive than spatial clustering; or 3) the superior forecast incorporates predictive information in addition to spatial clustering. Certainly, these are generalized explanations and the details of the specific experiments, forecasts and observed target earthquake distributions will require careful analysis. But the simplicity of TripleS is advantageous because the success or failure of the resulting forecasts can be understood clearly, without simultaneous testing of multiple dependent hypotheses obscuring the interpretation.

Despite our best efforts to construct simple forecasts, we were forced to make a number of modeling decisions that can be reconsidered in future studies. Chief among these was the choice to avoid declustering, a choice that was motivated by the desire not to convolve the effects of a given declustering algorithm and the effects of smoothing. Others have suggested that it is appropriate to decluster the input catalog and smooth the remaining events, which are thought of as «background earthquakes» [e.g., Helmstetter et al.

2007, Werner et al. submitted BSSA]. Nevertheless, the details of declustering, and particularly the preferred model parameter values, remain somewhat nebulous [van Stiphout et al. submitted BSSA].

We also decided to disregard epicenter uncertainty, based on our reasoning that the discretization of the testing region makes location errors negligible. Moreover, we chose the set of candidate smoothing distances arbitrarily; while the minimum distance of interest is related to the spatial discretization of the forecast format, the set of candidate values can be supplemented and explored iteratively. Along the same lines, rather than using the area skill score and average misfit statistic, the optimization procedure could be based on the spatial joint log-likelihood test of Zechar et al. [2010]. In constructing our hybrid catalog and the corresponding forecasts, data of varying quality was mixed (at least in terms of completeness), and in so doing we might have introduced a bias that would lend more weight to recent smaller events; this should also be reconsidered in future experiments.

The outstanding issues of TripleS are exemplary of the more general notion that many issues regarding smoothed seismicity approaches to earthquake forecasting and seismic hazard assessment are unresolved. In this study, we concentrated on a single value as an optimal spatial smoothing distance, but broader questions remain, such as:

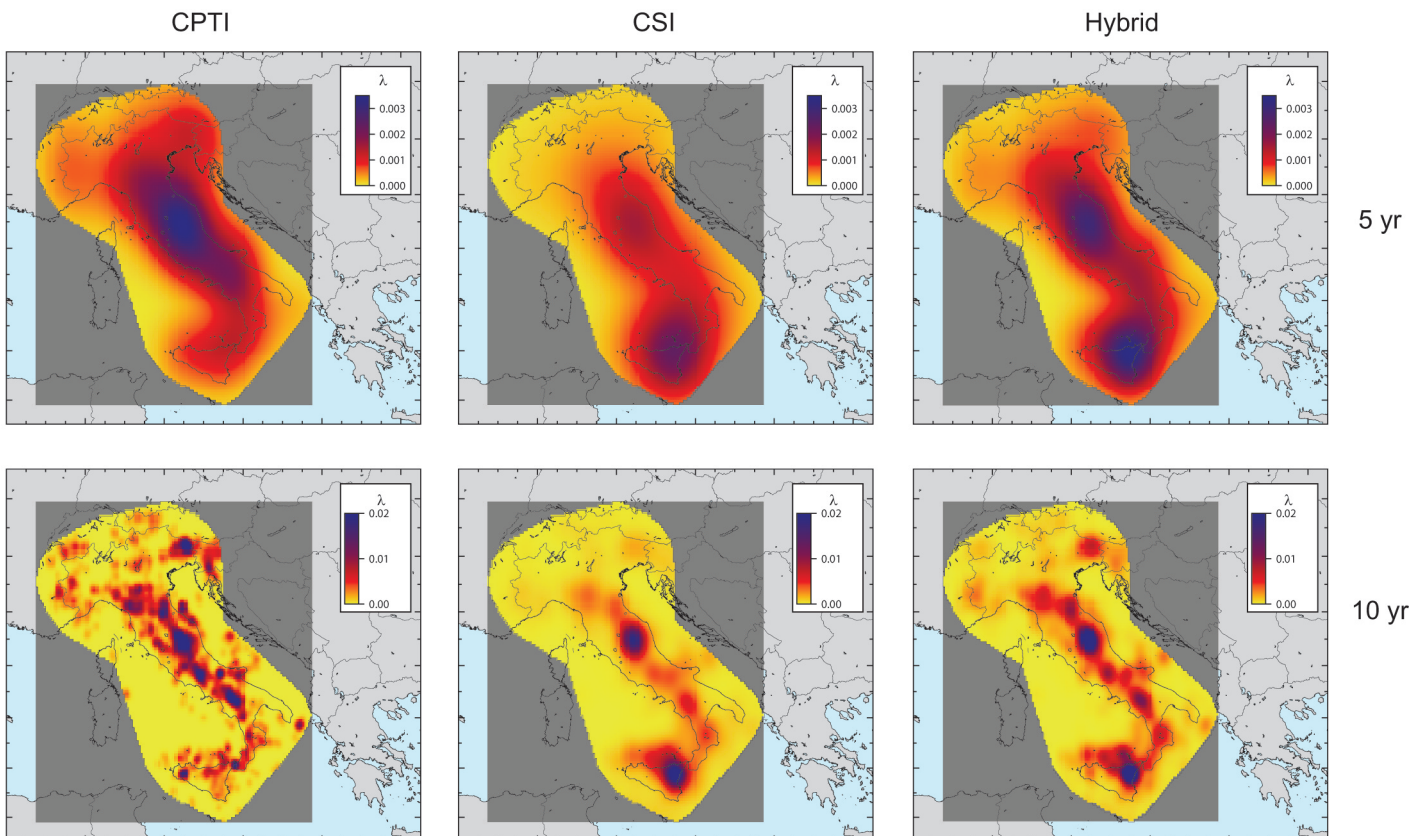


Figure 2. Map views of the corrected submitted five-year and ten-year TripleS forecasts using the CPTI, CSI, and hybrid catalogs, as indicated. These are the space-rate representations of the forecasts, as they have been summed over the magnitude bins, and the unit of measure is the expected number of earthquakes.

Should all earthquakes be smoothed? Should the contribution of an earthquake depend on its size or its origin time? What is the effect of a Gaussian kernel, as opposed to some other functional form? Should large earthquakes be smoothed using an anisotropic kernel, as Kagan and Jackson [1994] suggested? How is such anisotropy best characterized? Because seismicity smoothing is used in so many applications – including wide-ranging and long-term hazard assessments that affect building codes and insurance rates – these issues are highly relevant.

Data and sharing resources

Most of the catalog data used in this study are freely available from the CSEP Italy website (<http://cseptest.org/regions/italy>). The ECoS data used in this study are available upon request from the authors.

An electronic supplement to this article is available at <http://www.annalsofgeophysics.eu/index.php/annals/rt/suppFiles/4845/0>.

Acknowledgements. The authors thank Timothy R. Bentley and Gene Pantry for many delicious breakfasts and inspiring the name of the model, Fabian Euchner for assistance in preparing the figures, Max Werner for providing results in advance of publication, and Yan Kagan, Danijel Schorlemmer and an anonymous reviewer for carefully reading an early draft of the manuscript and providing constructive comments. The maps were created with the Generic Mapping Tools software [Wessel and Smith 1998].

References

- Faeh, D., D. Giardini, F. Bay, F. Bernardi, J. Braunmiller, N. Deichmann, M. Furrer, L. Gantner, M. Gisler, D. Isenegger, M. J. Jimenez, P. Kaestli, R. Koglin, V. Masciadri, M. Rutz, C. Scheidegger, R. Schibler, D. Schorlemmer, G. Schwarz-Zanetti, S. Steimen, S. Sellami, S. Wiemer and J. Woessner (2003). Earthquake catalog of Switzerland and the related macroseismic database, *Ecol. Geol. Helv.*, 96, 219-236.
- Frankel, A.D. (1995). Mapping seismic hazard in the central and eastern United States, *Seismol. Res. Lett.*, 66, 8-21.
- Helmstetter, A., Y.Y. Kagan and D.D. Jackson (2006). Comparison of short-term and time-independent earthquake forecast models for southern California, *Bull. Seismol. Soc. Am.*, 96, 90-106.
- Helmstetter, A., Y.Y. Kagan and D.D. Jackson (2007). High-resolution time-independent forecast for $M > 5$ earthquakes in California, *Seismol. Res. Lett.*, 78, 78-86.
- Jackson, D.D. and Y.Y. Kagan (1999). Testable earthquake forecasts for 1999, *Seismol. Res. Lett.*, 70, 393-403.
- Kagan, Y.Y. and D.D. Jackson (1994). Long-term probabilistic forecasting of earthquakes, *J. Geophys. Res.*, 99, 13685-13700.
- Kagan, Y.Y. and D.D. Jackson (2000). Probabilistic forecasting of earthquakes, *Geophys. J. Int.*, 143, 438-453.
- Kagan, Y.Y., D.D. Jackson, and Y. Rong (2007). A testable five-year forecast of moderate and large earthquakes in southern California based on smoothed seismicity, *Seismol. Res. Lett.*, 78, 94-98.
- Marcus, G. (2006). *The shape of things to come: prophecy and the American voice*, New York, 324 pp.
- Michael, A.J. (1997). Testing prediction methods: earthquake clustering versus the Poisson model, *Geophys. Res. Lett.*, 24, 1891-1894.
- MPS Working Group (2004). Redazione della mappa di pericolosità sismica prevista dall'Ordinanza PCM 3274 del 20 marzo 2003, Rapporto conclusivo per il Dipartimento della Protezione Civile, INGV, Milano-Roma, April 2004 (MPS04), 65 pp. + 5 appendices; <http://zonesismiche.mi.ingv.it>.
- Petersen, M.D., A.D. Frankel, S.C. Harmsen, C.S. Mueller, K.M. Haller, R.L. Wheeler, R.L. Wesson, Y. Zeng, O.S. Boyd, D.M. Perkins, N. Luco, E.H. Field, C.J. Wills and K.S. Rukstales (2008). Documentation for the 2008 update of the United States national seismic hazard maps, U.S. Geol. Surv. Open-File Rept 2008-1128, 61 pp.
- Rundle, J.B., K.F. Tiampo, W. Klein and J.S. Sá Martins (2002). Self-organization in leaky threshold systems: the influence of near-mean field dynamics and its implications for earthquakes, neurobiology, and forecasting, *Proc. Natl. Aca. Sci. USA*, 99, 2514-2521.
- Schorlemmer, D., A. Christophersen, A. Rovida, F. Mele, M. Stucchi and W. Marzocchi (2010). Setting up an earthquake forecast experiment in Italy, *Annals of Geophysics*, 53, 3 (present issue).
- Stirling, M.W., G.H. Mc Verry and K.R. Berryman (2002). A new seismic hazard model for New Zealand, *Bull. Seismol. Soc. Am.*, 92, 1878-1903.
- Stock, C. and E.G.C. Smith (2002a). Adaptive kernel estimation and continuous probability representation of historical earthquake catalogs, *Bull. Seismol. Soc. Am.*, 92, 904-912.
- Stock, C. and E.G.C. Smith (2002b). Comparison of seismicity models generated by different kernel estimations, *Bull. Seismol. Soc. Am.*, 92, 913-922.
- Van Stiphout, T., D. Schorlemmer and S. Wiemer (submitted BSSA). Uncertainties in background seismicity rate estimation, submitted to *Bull. Seism. Soc. Am.*
- Werner, M.J., A. Helmstetter, D.D. Jackson and Y.Y. Kagan (submitted BSSA). High resolution long- and short-term earthquake forecasts for California, submitted to *Bull. Seism. Soc. Am.*; draft available at <http://arxiv.org/abs/0910.4981>.
- Werner, M.J., J.D. Zechar, W. Marzocchi and S. Wiemer (2010). Retrospective evaluation of the five-year and ten-year CSEP-Italy earthquake forecasts, *Annals of Geophysics*, 53, 3 (present issue).
- Wessel, P. and W. Smith (1998). New, improved version of Generic Mapping Tools released, *Eos Trans. AGU*, 79, 579.
- Zechar, J.D. (2008). *Methods for evaluating earthquake predictions*, Los Angeles, University of Southern California Ph.D. dissertation, 134 pp.
- Zechar, J.D. and T.H. Jordan (2008): Testing alarm-based

earthquake predictions, *Geophys. J. Int.*, 172, 715-724; doi: 10.1111/j.1365-246X.2007.03676.x.

Zechar, J.D. and T.H. Jordan (2010). The area skill score statistic for evaluating earthquake predictability experiments, *Pure Appl. Geophys.*, 167 (8-9), 893-906; doi: 10.1007/s00024-010-0086-0.

Zechar, J.D., M.C. Gerstenberger and D.A. Rhoades (2010). Likelihood-based tests for evaluating space-rate-magnitude earthquake forecasts, *Bull. Seism. Soc. Am.*, 100 (3), 1184-1195; doi: 10.1785/0120090192.

*Corresponding author: J. Douglas Zechar,
ETH Zurich, Swiss Seismological Service, Zurich, Switzerland;
email: jeremy.zechar@sed.ethz.ch

Electronic supplement available at:
<http://www.annalsofgeophysics.eu/index.php/annals/rt/suppFiles/4845/0>

© 2010 by the Istituto Nazionale di Geofisica e Vulcanologia. All rights reserved.