A Bayesian spatio-temporal statistical analysis of Out-of-Hospital Cardiac Arrests

Stefano Peluso* Antonietta Mira† Håvard Rue‡ Nicholas John Tierney§
Claudio Benvenuti¶ Roberto Cianella‖ Maria Luce Caputo** Angelo Auricchio††

May 30, 2019

Abstract

We propose a Bayesian spatio-temporal statistical model for predicting Out-of-Hospital Cardiac Arrests (OHCA). Risk maps for Ticino, adjusted for demographic covariates, are built for explaining and forecasting the spatial distribution of OHCAs and their temporal dynamics. The occurrence intensity of the OHCA event in each area of interest, and the cardiac risk-based clustering of municipalities are efficiently estimated, through a statistical model that decomposes OHCA intensity into overall intensity, demographic fixed effects, spatially structured and unstructured random effects, time polynomial dependence and spatio-temporal random effect. In the studied geography, time evolution and dependence on demographic features are robust over different categories of OHCAs, but with variability in their spatial and spatio-temporal structure. Two main OHCA incidence-based clusters of municipalities are identified.
Keywords. Cardiac risk map, Integrated Nested Laplace Approximation, Temporal and spatial heterogeneity.

1 Introduction

The adoption of spatial statistical tools for the analysis of cardiac arrests is gaining more and more attention in the medical literature, as longer time series of geolocalized data and faster computational methods become available. Spatio-temporal datasets of Out-of-Hospital Cardiac Arrests (OHCA) are now available for Toronto [Sun et al., 2017], Taiwan [Chen et al., 2015], Brazil [Rodrigues et al., 2015], Australia [Straney et al., 2015], Japan [Onozuka and Hagihara, 2017], Korea [Han et al., 2017] and Canton Ticino [Tierney et al., 2018b].

*Corresponding author. Università Cattolica del Sacro Cuore, Department of Statistical Sciences, Largo Gemelli 1, 20123 Milan (Italy). E-mail: stefano.peluso@unicatt.it.

†Università della Svizzera italiana, Institute of Computational Science and Università dell'Insubria, Department of Science and High Technology

‡King Abdullah University of Science and Technology

§Monash University, Department of Econometrics and Statistics

¶Fondazione Ticino Cuore

‖FCTSA Federazione Cantonale Ticinese Servizi Autoambulanze

**Fondazione Cardiocentro Ticino, Division of Cardiology, University of Pavia, Department of Molecular Medicine and Fondazione IRCCS Policlinico San Matteo, Coronary Care Unit and Cardiovascular Clinical Research Center

††Fondazione Cardiocentro Ticino, Division of Cardiology and Università della Svizzera italiana, Center for Computational Medicine in Cardiology

From 2002 a systematic prospectively designed data collection started to shape the extensive registry of OHCAs occurring in Swiss Canton Ticino. Each ambulance intervention in the OHCA database is consecutively and uniformly recorded, geolocalized, and annotated with key intervention data such as year, month, day, hour and minute, characteristics of first responder time to ambulance arrival and transport time to the hospital. A first analysis on OHCA incidences in Canton Ticino is conducted in Mauri et al. [2015] and, more recently, in Caputo et al. [2017].

In the present paper we aim to estimate and predict the OHCA risk for the municipalities of Swiss Canton Ticino, through a model which accounts for temporal and spatial heterogeneity, space-time interactions and demographic features, and which can be conveniently estimated by exploiting the INLA methodology [Rue et al., 2009]. We refer the reader to Rue and Held [2005] and Gelfand et al. [2010, pp. 171-200] for an introduction to Gaussian Markov Random Fields (GMRF), and to the two review papers of Rue et al. [2017] and Bakka et al. [2018] for the large applicability of INLA to Bayesian spatial models based on latent GMRFs.

Inference through INLA has been recently conducted on a series of spatio-temporal problems with medical relevance. A Bayesian hierarchical varying-coefficient model is estimated with INLA by Osei and Stein [2017] to study spatio-temporal heterogeneities in diarrhoea morbidities. Spatio-temporal variations of substance abuse mortality in Iran with a log-Gaussian Cox point process model is proposed by Rostami et al. [2017]. Prostate cancer mortality data in 50 Spanish provinces over the period 1986-2010 is studied with INLA-related methods by Goicoa et al. [2016]. A tensor product spline model with a Markov random field prior on the coefficients of the basis functions, for hand, foot and mouth disease data in the central-north region of China, is estimated by INLA in Bauer et al. [2016]. To our knowledge, it is the first time that a spatio-temporal model is adopted to predict the OHCA risk map of a region. Furthermore, the methodology of INLA has never been applied to the analysis of spatial and temporal OHCAs. Nevertheless, INLA can provide results within few seconds, encouraging the adoption of the proposed model and estimation methodology to other OHCA datasets.

The closest approach to the one we propose is by Tierney et al. [2018a]: they conduct a spatial analysis on the same data, but there are several differences with our approach: (i) their model does not account for temporal and spatio-temporal heterogeneity, but only for the spatial dimension of the data; (ii) we introduce different timed demographic covariates that we appropriately manipulate to have coherence with a changing municipality subdivision of Ticino over time, whilst their covariates are fixed in time; (iii) our grid is irregular and coincides with municipalities, to cast the conclusions to the politically relevant level of municipality, whilst the grid used in Tierney et al. [2018a] is regular and does not follow municipalities boundaries. Another closely related work is Lin et al. [2016], in which a Bayesian spatial analysis is applied to estimate the spatial OHCA risks in Taiwan, but again they do not consider the dimension of time.

Other statistical approaches to the analysis of OHCA data can be found in Chen et al. [2015]: they apply logistic regression with kriging to identify risk factors in Taiwan OHCAs, identifying spatial heterogeneity (temporal dimension is neglected) of emergency medical resources between rural and urban areas. Rodrigues et al. [2015] analyse the space-time distribution of cardiovascular diseases in urban areas of Mato Grosso State through neighbourhood weighted means of mortality rates, with the purpose of identifying spatio-temporal clusters, but the statistical model is confined to single years. Spatially-smoothed logit and Poisson models are adopted in Straney et al. [2015] on Australian data, but separately in each spatial cell and with years arbitrarily collapsed. On Toronto data from 2007 to 2015, Sun et al. [2017] identify and rank businesses and municipal locations by spatio-temporal OHCA coverage, also studying the temporal stability of the rankings, but no prediction of OHCA risk is inferred. Poisson kriging and Hotspot analysis is used on the same data by Przybysz and Bunch [2017], but without temporal heterogeneity, finding that the influence of spatial neighbours does not extend beyond two adjacent units. Time-series Poisson regression analysis combined with a distributed lag non-linear model is implemented by Onozuka and Hagihara [2017] to assess the temporal variability in the effects of extremely high temperature on OHCA incidence in Japan. Their two-stage statistical analysis is focused on the estimation of the association between temperature and OHCA incidence, with no estimate of OHCA risk that accounts for spatial and temporal effects.

The statistical model and the related estimation methodology are detailed in Section 2, whilst the results are reported in Section 3, with emphasis on estimation and prediction of OHCAs (Section 3.1) and on municipality classification (Section 3.2). Results will be anonymized in a way that the ranking among municipalities in terms of their OHCA incidence will not be explicit, but we are available to privately release the exact ranking to interested practitioners and authorities. Finally Section 4 highlights conclusions and directions of investigation.

2 Statistical framework

2.1 Data preparation

A first look at the data clearly shows complex spatial structure in the OHCA distribution, reasonably related to demographic characteristics. From left Figure 1, we observe a concentration of cardiac events in the Southern part of Ticino, with a spatial distribution that follows the geographical configuration of valleys. A counting of cardiac events in each municipality show a striking difference between the map with the number of events (left Figure 1) and the map with event incidences over municipality populations (right Figure 1): as expected, the absolute numbers of OHCAs are strictly higher in more populated areas, whilst municipalities in the Northern and Western part of the region with less population concentration show relevant arrest incidences.

Starting from the whole dataset of 4638 OHCAs from Jan 2005 to Dec 2015, we first eliminate 2289 cases with unknown or non-cardiac aetiology (as for trauma-related events). We then remove 78 cases occurred close but beyond the border of Canton Ticino. The remaining 2271 cases represent the studied population and are represented in left Figure 1. For each OHCA case, the record of features we adopt in the statistical analysis are: date of the event, geographic coordinates, age, residential status, sex, time of 144 call, aetiology, if witnessed or not, with cardiac pulmonary resuscitation (CPR) attempted or not.

Time series from 2010 to 2016 of the total populations of the 117 Ticino municipalities are collected from the Land Register of Canton Ticino, together with sex and age composition. Demographic trends appear very stable over time, and populations at municipality level before 2010 have been extrapolated using a linear regression model. The geographical configuration of municipalities has changed over the studied years through municipality mergers and splits occurring twice per year. To have a coherent time series of demographic covariates we therefore reconstruct the time series backwards accounting for all mergers and splits, so that the data are consistent with the most recent available Ticino map in 2017.

2.2 Model formulation

We can measure Y_{it}, the number of OHCAs occurring in a given spatio-temporal cell, that is the number of cardiac events in year t and municipality i, for $t = t_1, \dots, T$ and $i = 1, \dots, 119 =: n$, where $t_1 = 2005$ is the first year of observations, and $T = 2015$ is the last one. Response Y_{it} sums up cardiac events occurring at a constant rate, specific of the year and of the municipality. The expected number of events in the specific cell is $O_{it}\mu_{it}$: O_{it} is the known or plug-in estimate of the baseline population for which events can occur (the resident population of municipality i in year t) and μ_{it} is the cardiac event incidence per unit of population. Inference on μ_{it} is crucial to highlight regions in Ticino with high/low exposure to the factors underlying cardiac risk, whilst inference on $O_{it}\mu_{it}$ is crucial for intervention purposes.

The model has to be complex enough to include over-dispersion on the Poisson realizations potentially caused by spatial and temporal effects, spatio-temporal interaction, and covariates effects. The OHCA incidence is composed of six parts: (i) the overall cardiac risk rate μ (in log scale), common to all municipalities and years; (ii) a fixed effect induced by covariates $\boldsymbol{X}_{it}$ and corresponding regression parameters $\boldsymbol{\beta}$, which in our problem correspond to demographic features: male/female population proportions, age proportions distinguishing the age classes of $[0, 40)$, $[40, 50)$, $[50, 60)$, $[60, 70)$, $[70, 80)$ and $[80, \infty)$, and interactions between sex and age features; (iii) the spatially related component u_i and its parameter ϕ that captures the proportion of variability in Y_{it} explained by the geographical heterogeneity; (iv) the temporal effect through the direct inclusion into the model of a polynomial of degree 3 in the demeaned year, with related parameters γ_1, γ_2 and γ_3 that should highlight possible changes over time in the distributions of the Y_{it}; (v) the spatio-temporal interaction component w_{it}, with related precision in the parameter ξ; (vi) an unstructured random effect v_i not related to the previous components. Therefore we can write, for $i = 1, \dots, n$ municipalities and years $t = t_1, \dots, T$:

$$
\begin{aligned}
Y_{it} &\sim \frac{e^p}{1+e^p}\mathbb{1}_{Y_{it}=0} + \frac{1}{1+e^p} Poisson(O_{it}\mu_{it}), && (1)\\
\ln \mu_{it} &= \mu + \boldsymbol{X}_{it}\boldsymbol{\beta} + \frac{1}{\sqrt{\tau}}\left(\sqrt{1-\phi}v_i + \sqrt{\phi}u_i\right) + \\
& \quad + \sum_{k=1}^{3} \gamma_k (t^k - \bar{t}^k) + \frac{1}{\sqrt{\xi}} w_{it}, && (2)
\end{aligned}
$$

where the additional parameters τ^{-1} and ξ^{-1} have, respectively, the role of the variance of both the spatially structured and unstructured component, and of the variance of the unstructured spatio-temporal component [Riebler et al., 2016], whilst $\bar{t}^k$ is the sample mean of the years powered k.

2.3 Spatial components

For the spatially structured and unstructured part of the statistical model, we adopt a reparameterized version of the BYM (acronym from the authors' names of Besag et al. 1991) model, as introduced by Riebler et al. [2016]. We prefer the BYM model to the similar and simpler Conditional Autoregressive (CAR) Model of Besag [1974] since the former is able to manage the extreme case of no spatial variability, without giving misleading parameter estimates [Breslow et al., 1998]. The reparameterization follows the suggestions of Simpson et al. [2017] and Riebler et al. [2016], and it is justified by a better interpretation of hyperprior specifications and by a reduced sensitivity of spatial smoothness to the hyperpriors. To be more detailed, our specification builds on the model of Besag [1974], which relies on the reasonable assumption that areas close in space behave similarly. But the main limitation of Besag [1974] is that spatial variability is only driven by a structural random effect

$$\boldsymbol{u} = (u_1, \dots, u_n) \sim \exp\left(-\frac{\tau}{2}\boldsymbol{u}'\boldsymbol{Q}\boldsymbol{u}\right),$$

where the $n \times n$ matrix $\boldsymbol{Q}$ reflects the graphical structure of the map under study:

$$\boldsymbol{Q} = \begin{cases} n_i, & i = j \\ -1, & i \sim j \\ 0, & \text{otherwise,} \end{cases}$$

where $i \sim j$ denotes the set of municipalities $j \neq i$ neighbours of municipality i, and n_i is the cardinality of this set. Therefore this model cannot handle situations where there is no spatially structured effect without introducing some distortions in the estimation [Breslow et al., 1998]. This prior choice for $\boldsymbol{u}$ is equivalent to assign a prior CAR distribution with Markovian structure to the spatial effects u_i: u_i conditionally on all other values of u_j, $j \neq i$, is Gaussian and centred on the mean value of all geographical neighbours u_j, with a variance that depends inversely on the size of this neighbourhood:

$$u_i | \{u_j, j \neq i\}, \tau \quad \sim \quad N\left(\frac{1}{n_i}\sum_{i \sim j} u_j, \frac{1}{n_i \tau}\right). \tag{3}$$

The BYM model of Besag et al. [1991] then introduced an unstructured component $\boldsymbol{v} = (v_1, \dots, v_n) \sim N(\boldsymbol{0}, \tau^{-1}\boldsymbol{I})$ of the spatial effect, with $\boldsymbol{0}$ and $\boldsymbol{I}$ being, respectively, the null vector and the identity matrix of appropriate dimensions, with the purpose of dealing with cases where there is no structural component. But the BYM model has two important limitations: (a) the structured and unstructered components are not identifiable and (b) the marginal variance of the structural component is not scaled. Successive models of Leroux et al. [2000] and Dean et al. [2001] solve the former limitation, but are still affected by the latter. A structured component that is not scaled implies that the interpretation of the related hyperparameters can change from application to application. Both limitations above are overcome in the reparameterized version of the BYM model proposed in Simpson et al. [2017] and Riebler et al. [2016] and that we adopt in our setting. Then the unstructered spatial component in Equation (2) is $\boldsymbol{v} \sim N(\boldsymbol{0}, \boldsymbol{I})$, whilst the structured component is $\boldsymbol{u} \sim N(\boldsymbol{0}, \boldsymbol{Q}/\sigma^2_{\boldsymbol{u}})$, where $\sigma^2_{\boldsymbol{u}}$ is known as generalized variance, computed as

$$\sigma^2_{\boldsymbol{u}} = \exp\left(\frac{1}{n}\sum_{i=1}^{n} \ln\left(\boldsymbol{Q}^{-1}\right)_{ii}\right).$$

With this formulation, the proportions of spatial variability explained the structured and unstructured components are, respectively, ϕ and $1 - \phi$, and the model reduces to pure overdispersion for $\phi = 0$ and to the model of Besag [1974] for $\phi = 1$.

2.4 Other components and Penalized Complexity priors

Highly non-informative priors have been chosen for μ, $\boldsymbol{\beta}$ and γ_k, $k = 1, 2, 3$:

$$
\begin{aligned}
\mu &\sim N(0, 1000) && (4)\\
\boldsymbol{\beta} &\sim N(\mathbf{0}, 1000\boldsymbol{I}) && (5)\\
\gamma_k &\sim N(0, 1000). && (6)
\end{aligned}
$$

The unstructured and the spatio-temporal effects are set a priori to follow a standard Gaussian distribution:

$$
\begin{aligned}
v_i &\sim N(0, 1), && (7)\\
w_{it} &\sim N(0, 1). && (8)
\end{aligned}
$$

whilst the zero-inflated parameter p is assigned a priori a Gaussian distribution $N(-1, 0.2)$ to its logistic transformation $\ln(p/(1-p))$. Variance decomposition results below will attribute very little weight to the spatio-temporal component, and this behaviour, together with needs of model parsimony, led us to consider a model with unstructured spatio-temporal random effect a good fit for our data. Finally, for the parameters τ, ξ and ϕ we fix Penalised-Complexity (PC) priors, shown in Simpson et al. [2017] to be principle-based and with nice desiderata. More precisely, if we measure model complexity in terms of Kullback-Leibner distance KL between the model and some reference base model, PC priors favour simpler models by putting on this distance a constantly decaying prior mass. This means a prior exponential distribution on $\sqrt{2KL}$, in turn implying some prior on the parameter of interest from which the distance depends. It is shown in Simpson et al. [2017, Appendix A.1] that the PC priors for $1/\sqrt{\tau}$ and for $1/\sqrt{\xi}$ both correspond to an exponential distribution with rate 4.60517. For the PC-prior on ϕ we refer the reader to Simpson et al. [2017, Appendix A.3] for the exact form of the distance KL: in this case the implied prior on ϕ is not available in closed form and it is derived from the prior on KL and the numerical approximation of the Jacobian transformation from KL to ϕ.

2.5 Estimation methodology

To infer the OHCA incidences in Ticino municipalities and all the remaining parameters in the model introduced aobe, we implement the Integrated Nested Laplace Approximation (INLA). We refer the reader to the original article of Rue et al. [2009] for the details of how the methodology works: here we focus on its assumptions and we highlight its suitability to the current problem setting.

INLA is a methodology proposed in Rue et al. [2009] for fast statistical analysis of a large series of complex problems, also accounting for spatial and temporal heterogeneity. More in details, it is a Bayesian estimation approach useful for Latent Gaussian Models [Fahrmeir and Tutz, 2013], that is for models that can be written as

$$
\begin{aligned}
\boldsymbol{y}|\boldsymbol{\theta}, \boldsymbol{\psi} &\sim \prod_{i=1}^{n} p(y_i|\boldsymbol{\theta}, \boldsymbol{\psi})\\
\boldsymbol{\theta}|\boldsymbol{\psi} &\sim N(\mathbf{0}, \boldsymbol{\Sigma}_{\boldsymbol{\psi}})\\
\boldsymbol{\psi} &\sim p(\boldsymbol{\psi}),
\end{aligned}
$$

for some generic data $\boldsymbol{y}$, parameter vector $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_p)$, covariance matrix $\boldsymbol{\Sigma}_{\boldsymbol{\psi}}$ and parameter vector $\boldsymbol{\psi}$. Furthermore, INLA assumes that (i) $\boldsymbol{\theta}$, often of large dimension, admits a conditional independence structure $\boldsymbol{\theta}_I \perp\!\!\!\perp \boldsymbol{\theta}_J | \boldsymbol{\theta}_{-I\cup J}$, for some index sets I and J of $\mathbb{N}_p := \{1, 2, \ldots, p\}$ and with $-I \cup J$ defined as $\mathbb{N}_p \setminus (I \cup J)$. Hence, the latent field is a sparse Gaussian Markov random field (GMRF); (ii) the dimension of the parameter vector $\boldsymbol{\psi}$ is small, typically less or equal to 6.

All the properties required by INLA suits particularly well to our problem: rewriting

$$\begin{aligned}
\boldsymbol{y} &= (Y_{11}, Y_{12}, \ldots, Y_{nT}),\\
\boldsymbol{\theta} &= (\mu, \boldsymbol{\beta}', \gamma_1, \gamma_2, \gamma_3, u_1, \ldots, u_n, v_1, \ldots, v_n, w_{11}, \ldots, w_{nT}),\\
\boldsymbol{\psi} &= (p, \tau, \phi, \xi)
\end{aligned}$$

we can cast our model as a Latent Gaussian Model, and $\boldsymbol{\theta}$ admits a sparse conditional independence structure: μ, all v_is, w_{it}s and all components of $\boldsymbol{\beta}$ are independent from the rest, and u_is are independent from u_js not neighbours. Furthermore, the parameter space is large, being the dimension of $\boldsymbol{\theta}$ equal to $21 + 2n + nT = 1776$, and this makes INLA more convenient than other MCMC or deterministic methods. Note that we are able to estimate a model with more parameters than observations: the number of parameters is not the right measure for the complexity of the model, as strong dependence between the parameters will reduce the "effective number of parameters". It is the latter that should be compared with the number of observations (following a frequentist point of view) for a model to be identifiable or not. In the Bayesian framework we adopt, this is very different, as proper priors would give proper posteriors under quite general assumptions, even with no observations.

3 Empirical investigation

3.1 Inferential and prediction results

We split the available sample in two parts: 2005-2013 for estimating the models, and 2014-2015 for out-of-sample prediction purposes. We implement the INLA methodology to different categories of OHCAs: to all OHCAs, to witnessed and not (at the moment of collapse) OHCAs, with CPR attempted and not, of resident and not resident person, and to OHCAs occurred in the daytime window $[7, 19]$ or $[19, 7]$. Estimates and posterior credible intervals at 95% probability reported in Table 1 reveal important facts. First of all there is a time effect for which the overall event risk is decreasing over time, with all degrees of the time polynomial being relevant. Also, the age factor seems to be more important than sex when both are taken into account. Specifically, we estimate some negative relationship between the OHCA rate and proportion of people below 50 years old. This result is perfectly coherent with clinical intuition. Furthermore, interaction terms reveal that sex differences are evident within younger age classes: there is a posteriori a negative dependence between the event rate and males below 40 years old and females below 60 years old, again in line with clinical observations. Results for subcategories of OHCAs are not reported for brevity. With the only exception of the age class $[50, 60)$ for females, we estimate exactly the same time decrease in OHCA rate and the same effects of demographic features, when we limit the analysis to witnessed OHCAs, not witnessed, CPR-attempted, not CPR-attempted, OHCAs in window $[7, 19]$, to resident and not resident OHCAs. Therefore our estimation results are very robust among diverse OHCA categories. For the subset of OHCAs occuring in $[19, 7]$ we also observe some negative dependence with proportion of females older than 70 years old.

Once the effect of demographic variables is accounted for, $\phi\xi/(\tau+\xi)$ is the proportion of the left variability in the occurrence of OHCAs that can be explained by the structurally spatial component, whilst $\tau/(\tau+\xi)$ is due to the spatio-temporal variability. In the model with all OHCAs, we estimate a proportion of 22.07% due to spatial heterogeneity, 1.97% of spatio-temporal variability, and the remaining is attributed to unstructured spatial overdispersion. The same variance decomposition is reported in Table 2 for all models considered. Not surprisingly, the importance of the structured spatial component has its minimum for not resident OHCAs (7.22%) and its maximum for all OHCAs (22.07%), with large variability among the models. For the category of not resident events, we also observe the minimum spatio-temporal relevance, whilst maximum importance to spatio-temporal heterogeneity is shown for OHCAs with attempted CPR. Cardiac events with no CPR, that are witnessed, with 144 call after 7pm, or not resident, are the classes of cardiac events better linked to unknown features, other than demographic ones, related to the territory: this is clear from low DIC, associated with an important unstructured spatial variability.

The combined space and space-time effect is minimum for not resident data, but overall the DIC measure for this model is low, meaning that other spatial features could play a proper role in explaining the not resident OHCAs distribution in space and time. Then our model performs well in explaining the variability of cardiac events for people resident outside Ticino, but as expected this result is less related to $\boldsymbol{\beta}$, since demographic features are measured on the resident population. Still, non-negligible random effects suggest that cardiac events of foreigners could be related to other unknown factors, linked to the territory and not currently accounted in our model, as pollution levels, meteorological covariates, educational level, socio-economic indicators, closeness to the borders, tourist and work attractiveness.

Furthermore, we note that the relative quality of the model measured by DIC is worst for all OHCAs and for resident OHCAs, suggesting that the formulation of statistical models based on a more detailed classification of OHCAs is indeed relevant. In this direction, we can capture the improvement driven by more precisely classified OHCAs by comparing the respective DICs: whilst a model on all OHCAs shows a DIC of 2605.58, the models that distinguish between CPR attempted and not OHCAs have on average a DIC of 1923.49. The other classifications in witnessed and not, $[7-19]$ and $[19-7]$, resident and not, also brings a relevant decrease of roughly the same amount.

To detail the prediction results, we measure the prediction error between the true and predicted number of events per municipality. We focus on the OHCAs categories with CPR attempted and not, and to forecast years 2014, 2015 and 2016, years not included for the estimation of the model parameters. Similar results can be obtained for all categories. Note that in 2014 and 2015 both the true data and the demographic covariates are available, whilst in 2016 true events are not at disposal. We plot in Figure 2 the barplots for mistakes in the prediction of CPR attempted (first row) and not (second row) events, in 2014 (first column) and in 2015 (second column). For the great majority of municipalities there is no error or an error of one in the prediction of OHCAs, and the error is roughly always below 3, with notable exceptions that occurred in the most populated municipalities. There seems to be no worsening in the prediction from 2014 to 2015, despice predictions in 2015 have been conducted on a model estimated on events up to two years before. In Figure 3 we show the predicted map for OHCAs with CPR attempted and not in 2016, and similar maps can be provided for other kinds of OHCAs, years or subsets of municipalities.

Overall, by linearly regressing yearly true OHCAs on predicted ones, our model with demographic characteristics of the territory, spatial, temporal and spatio-temporal effects, is able to explain more than 90% of the variation of OHCAs in Ticino, as shown in Figure 4. We report the

R^2 for the model on all OHCAs, highlighting validation years. In the last two years we do not observe, relative to the part of sample dedicated to estimate the model, any notable decrease in the performance, suggesting against hypotheses of model overfitting. Furthermore, subcategories of OHCAs reveal that events with no CPR attempted, witnessed and non-resident are more difficult to predict. These classes of OHCAs are observed more rarely (respectively 20%, 25% and 10% of all events) and therefore the fewer counts of OHCAs affect negatively the prediction ability of the models. Coherently with the analysis above, particularly worse are the predictions of OHCAs for non-resident patients, for which beyond the few counts of OHCAs there is also a detachment from the demographic features that concern the resident population.

3.2 Municipality classification by incidence

One of the main objectives of the spatio-temporal statistical analysis of OHCAs in Ticino is to identify regions with anomalous OHCA incidence. This is important for creating a risk map of the territory that can serve for interventions and optimal allocation of defibrillators in areas of higher cardiac risk [Tierney et al., 2018b]. Following the methodology in Section 2.5, we provide point estimates $\hat{\mu}_{it}$ and 95% posterior credible intervals $CI_{\hat{\mu}_{it}} = (L_{\hat{\mu}_{it}}, U_{\hat{\mu}_{it}})$ of incidences of single municipalities in each year of the sample, for $i = 1, \dots, n$ and $t = t_1, \dots, T$. We can therefore detect regions of relevant differences (according to the estimated posterior probabilities) through comparison of municipalities: for instance the anonymized Municipality A and B show the temporal pattern in the left panel of Figure 5. Solid lines are the point estimates of OHCA incidences, with corresponding dashed lines denoting the credible intervals. We consider two municipalities to have a different incidence if their credible intervals do not overlap. It is clear that from 2009 there is a noteworthy difference in posterior distribution of the two incidences, with Municipality B higher than A, and that this is due to a roughly stable incidence in Municipality A, against worsening conditions in B. Further results not shown for brevity reveal that the difference is limited to OHCAs of residents, in the time range from 7am to 7pm, not witnessed and with CPR attempted.

The number of municipalities with higher and lower OHCA incidence gives a simple score for a given municipality in a year, from which we can create a ranking of municipalities evolving over time. More formally, the score of municipality i in year t is

$$\mathcal{S}_{it} = \sum_{j=1}^{n} \left(\mathbb{1}_{[0, L_{\hat{\mu}_{jt}})}(U_{\hat{\mu}_{it}}) - \mathbb{1}_{(U_{\hat{\mu}_{jt}}, \infty)}(L_{\hat{\mu}_{it}}) \right),$$

where $\sum_j \mathbb{1}_{[0, L_{\hat{\mu}_{jt}})}(U_{\hat{\mu}_{it}})$, the number of municipalities with higher incidence, contributes positively to the score, whilst the number of municipalities with lower incidence, $\sum_j \mathbb{1}_{(U_{\hat{\mu}_{jt}}, \infty)}(L_{\hat{\mu}_{it}})$, negatively. A municipality with a higher score is *better* and has a lower rank. Rank comparison between Municipality A and B in the right panel of Figure 5 provides further information: Municipality B with its increasing OHCA incidence loses 7 ranking positions from the beginning to the end of the sample, whilst Municipality A maintains a fairly stable incidence and at the same time gains 4 positions in the ranking. For better visual inspection, we report in the top-left panel of Figure 6 the 10 most virtuous (in blue) and vicious (in red) municipalities all over the years, and it is striking the geographical closeness of the municipalities with highest OHCA incidence.

The main limitations of pairwise comparisons of OHCA incidences are the impossibility to capture the relationships among multiple municipalities. For instance, in the top-right and bottom panels of Figure 6, we report the time evolution over the periods 2005-2008, 2009-2012 and 2013-2016 of the heatmaps for six anonymized municipalities, included those in Figure 5. From yellow to red we mark higher and higher distances among municipalities, where the distance is interpreted in terms of number of years in the period for which the credible intervals of the involved municipalities do not overlap. In other words, the number highlighted in each cell is the number of years in the period at which the municipalities at the margins of the cell have a posteriori different OHCA incidences. It is apparent the consolidating behaviours over time of two clusters from the years 2005-2008 when it is not even clear how many clusters there are, to the final years in which the distances between municipalities from different clusters are almost always at the maximum allowed.

A *complete linkage* hierarchical clustering method is implemented, based on the dissimilarity matrix given by the distances among municipalities (Everitt et al. [2011], Legendre and Legendre [2012]). The number of clusters has been optimally chosen to be three according to the silhouette principle of Rousseeuw [1987]. As mentioned above, the distances are the number of years in the period for which any two municipalities have OHCA incidences different a posteriori. Municipalities with less than 100 inhabitants have been excluded from the analysis, to not artificially alter the results with null incidences due to very low municipality populations. In Figure 7, overall in Ticino, we observe two main clusters in the two sub-periods covering respectively the training and test part of the data: a big cluster of roughly 80 municipalities spread out all over the region, mainly following the valleys, presenting on average an OHCA rate lower than 7/10,000 and a second one, characterized by a more limited number of municipalities (around 25), in a more geographically limited area, with a posteriori higher OHCA incidence. In the validation period, the main cluster remains mainly unaltered: two municipalities move to the second cluster and two others enter from the second cluster. Also, our model predicts the creation of a third cluster made of four municipalities with higher OHCAs incidence. These municipalities were part of the second cluster in the training period, and their leaving causes a slight improvement in the OHCA incidence of the second cluster during the test period.

We finally report in Figure 8 maps of standardized morbidity ratio (SMR), a common measure of interest in such analyses, for comparing the observed events to expected events. It helps identifying regions that have unusually higher risk of morbidity, beneficial to anyone looking to use these results to identify areas of higher risk than expected. The SMR for municipality i and time t is estimated as $\mu_{it}/\bar{\mu}_t$, where $\bar{\mu}_t = \sum_{i=1}^{n} O_{it}\mu_{it}/\sum_{i=1}^{n} O_{it}$ is the overall OHCA indicence, weighted by municipality populations, for all municipalities with the exceptions of those with less than 100 people. A SMR close to one denotes a municipality in line with the overall incidence of the year, whilst a SMR above (below) one indicates a municipality with incidence estimated above (below) the overall incidence. For instance, a SMR close to 1.25 means that the municipality is estimated to have a OHCA incidence 25% higher the overall rate. As expected, SMR maps closely follow the clustering structure found above.

4 Conclusions and further directions

We provide a spatio-temporal statistical model for Out-of-Hospital Cardiac Arrests (OHCAs), with the purpose of constructing complex model-based cardiac risk maps for Ticino. The inferential methodology adopted is the Integrated Nested Laplace Approximation (INLA), an efficient numerical method that estimates, in each area of interest, the occurrence intensity of the OHCA event, and how it evolves over time. Relative to the past literature we provide, to our knowledge for the first time, a statistical model for the prediction of spatially and temporally heterogeneous OHCAs. Our algorithm is very efficient, providing valid answers within few seconds, against huge computational efforts of more traditional techniques. Furthermore, our statistical method is Bayesian, and therefore can provide quantification of uncertainty in the intensity estimation.

The statistical analysis is validated on a portion of the sample, and repeated for subcategories of OHCAs, also accounting for demographic features of the municipalities. We provide a decomposition of OHCA variability into spatial and spatio-temporal components, and prediction of the number of events and OHCA incidences per resident population. Finally, two main clusters of municipalities consolidating over time and coherent with the standard morbidity ratios are identified, according to the predicted OHCA risk, together with one additional minor cluster with worse incidence appearing in the validation set.

The developed framework can be extended in various directions. First, the realistic map of OHCA risk serves to reliably predict incidence of OHCAs over the upcoming years in a given municipality or living area. The model and the related methodology can be validated outside the Swiss Canton Ticino, in different geographies. Second, the statistical model can be enriched by meteorological, climatic, environmental and socio-economic factors. Third, once a realistic risk map is generated, optimization of the territorial deployment of defibrillators can take place. The spatial and temporal variations of the OHCA events can help the aetiology of the phenomenon and identify the risk factors to support the decision of defibrillators locations near the population more at risk, as well as to optimize the number and location of ambulances and rescuer teams, with the purpose of remarkably reducing time to their arrival. Finally, INLA can be extended to models with natural barriers that affect the spatial dependence of OHCAs: two municipalities of Ticino separated by mountains affect each other to a lesser extent than two municipalities of Ticino not separated by natural barriers, other things being equal.

Acknowledgements

Financial support from Fondazione Fratelli Agostino Enrico Rocca is acknowledged. Antonietta Mira was partially supported by SNF grant 105218_166504. The authors thank the Swiss Cardiology Foundation.

References

Haakon Bakka, Håvard Rue, Geir-Arne Fuglstad, Andrea Riebler, David Bolin, Janine Illian, Elias Krainski, Daniel Simpson, and Finn Lindgren. Spatial modeling with r-inla: A review. *Wiley Interdisciplinary Reviews: Computational Statistics*, 10(6):e1443, 2018.

Cici Bauer, Jon Wakefield, Håvard Rue, Steve Self, Zijian Feng, and Yu Wang. Bayesian penalized spline models for the analysis of spatio-temporal count data. *Statistics in medicine*, 35(11):1848–1865, 2016.

Julian Besag. Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 192–236, 1974.

Julian Besag, Jeremy York, and Annie Molli. Bayesian image restoration, with two applications in spatial statistics. *Annals of the Institute of Statistical Mathematics*, 43(1):1–20, 1991.

Norman Breslow, Brian Leroux, and Robert Platt. Approximate hierarchical modelling of discrete data in epidemiology. *Statistical Methods in Medical Research*, 7(1):49–62, 1998.

Maria Luce Caputo, Sandro Muschietti, Roman Burkart, Claudio Benvenuti, Giulio Conte, François Regoli, Romano Mauri, Catherine Klersy, Tiziano Moccetti, and Angelo Auricchio. Lay persons alerted by mobile application system initiate earlier cardio-pulmonary resuscitation: A comparison with sms-based system notification. *Resuscitation*, 114:73–78, 2017.

Chien-Chou Chen, Chao-Wen Chen, Chi-Kung Ho, I-Chuan Liu, Bo-Cheng Lin, and Ta-Chien Chan. Spatial variation and resuscitation process affecting survival after out-of-hospital cardiac arrests (ohca). *PloS one*, 10(12):e0144882, 2015.

CB Dean, MD Ugarte, and AF Militino. Detecting interaction between random region and fixed age effects in disease mapping. *Biometrics*, 57(1):197–202, 2001.

Brian S Everitt, Sabine Landau, Morven Leese, and Daniel Stahl. Hierarchical clustering. *Cluster analysis*, 5, 2011.

Ludwig Fahrmeir and Gerhard Tutz. *Multivariate statistical modelling based on generalized linear models*. Springer Science & Business Media, 2013.

Alan E Gelfand, Peter Diggle, Peter Guttorp, and Montserrat Fuentes. *Handbook of spatial statistics*. CRC press, 2010.

T Goicoa, MD Ugarte, J Etxeberria, and AF Militino. Age–space–time car models in bayesian disease mapping. *Statistics in medicine*, 35(14):2391–2405, 2016.

Sang Hun Han, Kyoung Ok Kim, Eun Jong Cha, Kyung Ah Kim, and Ho Sun Shon. System framework for cardiovascular disease prediction based on big data technology. *Symmetry*, 9(12):293, 2017.

Pierre Legendre and Louis Legendre. Numerical ecology. 3rd english ed. *Developments in environmental modelling*, 24, 2012.

Brian G Leroux, Xingye Lei, and Norman Breslow. Estimation of disease rates in small areas: a new mixed model for spatial dependence. In *Statistical models in epidemiology, the environment, and clinical trials*, pages 179–191. Springer, 2000.

Bo-Cheng Lin, Chao-Wen Chen, Chien-Chou Chen, Chiao-Ling Kuo, I-chun Fan, Chi-Kung Ho, I-Chuan Liu, and Ta-Chien Chan. Spatial decision on allocating automated external defibrillators (aed) in communities by multi-criterion two-step floating catchment area (mc2sfca). *International journal of health geographics*, 15(1):17, 2016.

Romano Mauri, Roman Burkart, Claudio Benvenuti, Maria Luce Caputo, Tiziano Moccetti, Alessandro Del Bufalo, Augusto Gallino, Carlo Casso, Luciano Anselmi, Tiziano Cassina, Catherin Klersy, and Angelo Auricchio. Better management of out-of-hospital cardiac arrest increases survival rate and improves neurological outcome in the swiss canton ticino. *Europace*, 18(3):398–404, 2015.

Daisuke Onozuka and Akihito Hagihara. Spatiotemporal variation in heat-related out-of-hospital cardiac arrest during the summer in japan. *Science of the Total Environment*, 583:401–407, 2017.

FB Osei and A Stein. Diarrhea morbidities in small areas: Accounting for non-stationarity in sociodemographic impacts using bayesian spatially varying coefficient modelling. *Scientific reports*, 7(1):9908, 2017.

Raymond Przybysz and Martin Bunch. Exploring spatial patterns of sudden cardiac arrests in the city of toronto using poisson kriging and hot spot analyses. *PloS one*, 12(7):e0180721, 2017.

Andrea Riebler, Sigrunn H Sørbye, Daniel Simpson, and Håvard Rue. An intuitive bayesian spatial model for disease mapping that accounts for scaling. *Statistical methods in medical research*, 25(4):1145–1165, 2016.

Poliany CO Rodrigues, Emerson S Santos, Eliane Ignotti, and Sandra S Hacon. Space-time analysis to identify areas at risk of mortality from cardiovascular disease. *BioMed research international*, 2015, 2015.

Mehran Rostami, Younes Mohammadi, Abdollah Jalilian, and Bashir Nazparvar. Modeling spatio-temporal variations of substance abuse mortality in iran using a log-gaussian cox point process. *Spatial and spatio-temporal epidemiology*, 22:15–25, 2017.

Peter J Rousseeuw. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20:53–65, 1987.

Havard Rue and Leonhard Held. *Gaussian Markov random fields: theory and applications*. Chapman and Hall/CRC, 2005.

Håvard Rue, Sara Martino, and Nicolas Chopin. Approximate bayesian inference for latent gaussian models by using integrated nested laplace approximations. *Journal of the royal statistical society: Series b (statistical methodology)*, 71(2):319–392, 2009.

Håvard Rue, Andrea Riebler, Sigrunn H Sørbye, Janine B Illian, Daniel P Simpson, and Finn K Lindgren. Bayesian computing with inla: a review. *Annual Review of Statistics and Its Application*, 4:395–421, 2017.

Daniel Simpson, Håvard Rue, Andrea Riebler, Thiago G Martins, Sigrunn H Sørbye, et al. Penalising model component complexity: A principled, practical approach to constructing priors. *Statistical Science*, 32(1): 1–28, 2017.

Lahn D Straney, Janet E Bray, Ben Beck, Judith Finn, Stephen Bernard, Kylie Dyson, Marijana Lijovic, and Karen Smith. Regions of high out-of-hospital cardiac arrest incidence and low bystander cpr rates in victoria, australia. *PloS one*, 10(10):e0139776, 2015.

Christopher LF Sun, Steven C Brooks, Laurie J Morrison, and Timothy CY Chan. Ranking businesses and municipal locations by spatiotemporal cardiac arrest risk to guide public defibrillator placementclinical perspective. *Circulation*, 135(12):1104–1119, 2017.

N John Tierney, Antonietta Mira, H Jost Reinhold, Giuseppe Arbia, Stefano Peluso, S. Clifford, Tiziano Moccetti, Angelo Auricchio, and Kerrie Mengersen. Evaluating health facility access using bayesian spatial models and location analysis methods. 2018a.

N John Tierney, H Jost Reinhold, Antonietta Mira, Martin Weiser, Roman Burkart, Claudio Benvenuti, and Angelo Auricchio. Novel relocation methods for automatic external defibrillator improve out-of-hospital cardiac arrest coverage under limited resources. *Resuscitation*, 125:83–89, 2018b.

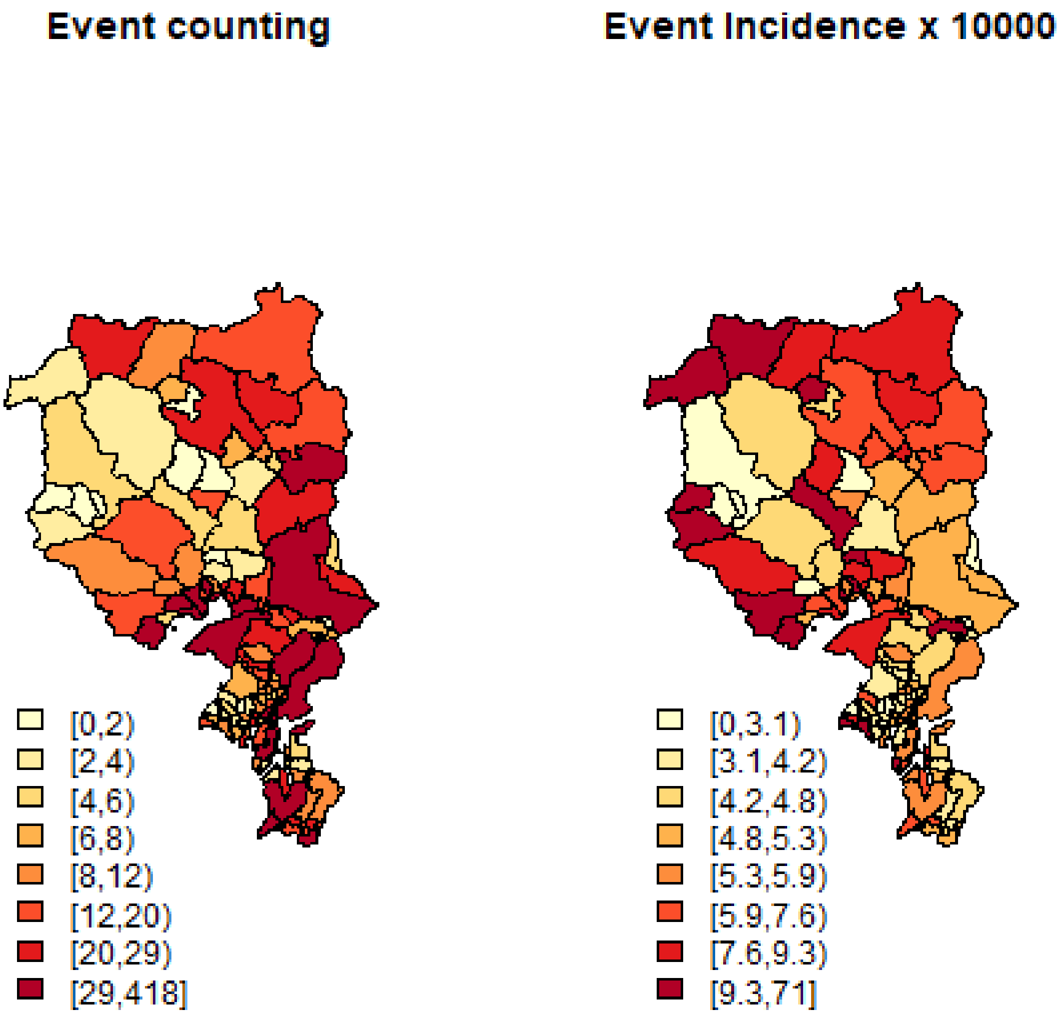

Figure 1: All OHCAs in Ticino from 2005 to 2015. Left: Number of OHCAs in each municipality. Right: number of OHCAs per unit of municipality population.

parameter	mean	sd	2.5%	97.5%
μ	**-0.4605**	0.2014	-0.8559	-0.0654
$\gamma_1 \times 10^4$	**-2.5240**	1.0185	-4.5238	-0.5260
$\gamma_2 \times 10^7$	**-1.2482**	0.5066	-2.2428	-0.2544
$\gamma_3 \times 10^{11}$	**-6.1718**	2.5195	-11.1184	-1.2294
male	-0.6802	0.4052	-1.4758	0.1146
age $[0, 40)$	**-1.9487**	0.4592	-2.8500	-1.0477
age $[40, 50)$	**-2.6450**	1.0456	-4.6982	-0.5941
age $[50, 60)$	-1.7337	1.2395	-4.1674	0.6975
age $[60, 70)$	0.2233	1.3294	-2.3875	2.8305
age $[70, 80)$	-1.3302	1.8026	-4.8692	2.2063
male $\times[0, 40)$	**-3.2942**	0.9173	-5.0952	-1.4947
male $\times[40, 50)$	-3.8549	2.0057	-7.7974	0.0756
male $\times[50, 60)$	-1.8188	2.3769	-6.4896	2.8400
male $\times[60, 70)$	2.5200	2.5355	-2.4613	7.4911
male $\times[70, 80)$	0.2366	3.3197	-6.2818	6.7488
male $\times[80, \infty)$	-3.5095	3.5476	-10.4800	3.4447
female $\times[0, 40)$	**-4.4316**	0.8672	-6.1324	-2.7288
female $\times[40, 50)$	**-6.5037**	1.9997	-10.4309	-2.5822
female $\times[50, 60)$	**-4.9819**	2.4046	-9.7024	-0.2643
female $\times[60, 70)$	-1.8030	2.5382	-6.7875	3.1748
female $\times[70, 80)$	-5.8306	3.4537	-12.6092	0.9467

Table 1: Parameter posterior expectation, standard deviation and 95% credible interval, for the model estimated on all OHCAs. In bold those parameters different from 0 with posterior probability higher than 95%.

Model	Spatial	Spatio-temporal	Unstructured	DIC	# events
all OHCAs	22.07%	1.97%	75.97%	2605.58	2271
CPR	19.52%	2.74%	77.74%	2392.69	1828
not CPR	8.45%	0.08%	91.47%	1454.28	443
witnessed	9.05%	0.15%	90.8%	1576.69	599
not witnessed	18.98%	0.78%	80.24%	2316.36	1672
$[7-19]$	15.11%	0.82%	84.06%	2222.28	1493
$[19-7]$	10.98%	0.21%	88.81%	1680.68	778
resident	17.11%	1.74%	81.16%	2475.41	2032
not resident	7.22%	0.02%	92.76%	1283.21	239

Table 2: Variance decomposition for OHCAs, after accounting for demographic effects, Deviance Information Criterion of overall model fit, number of events for each OHCA category.

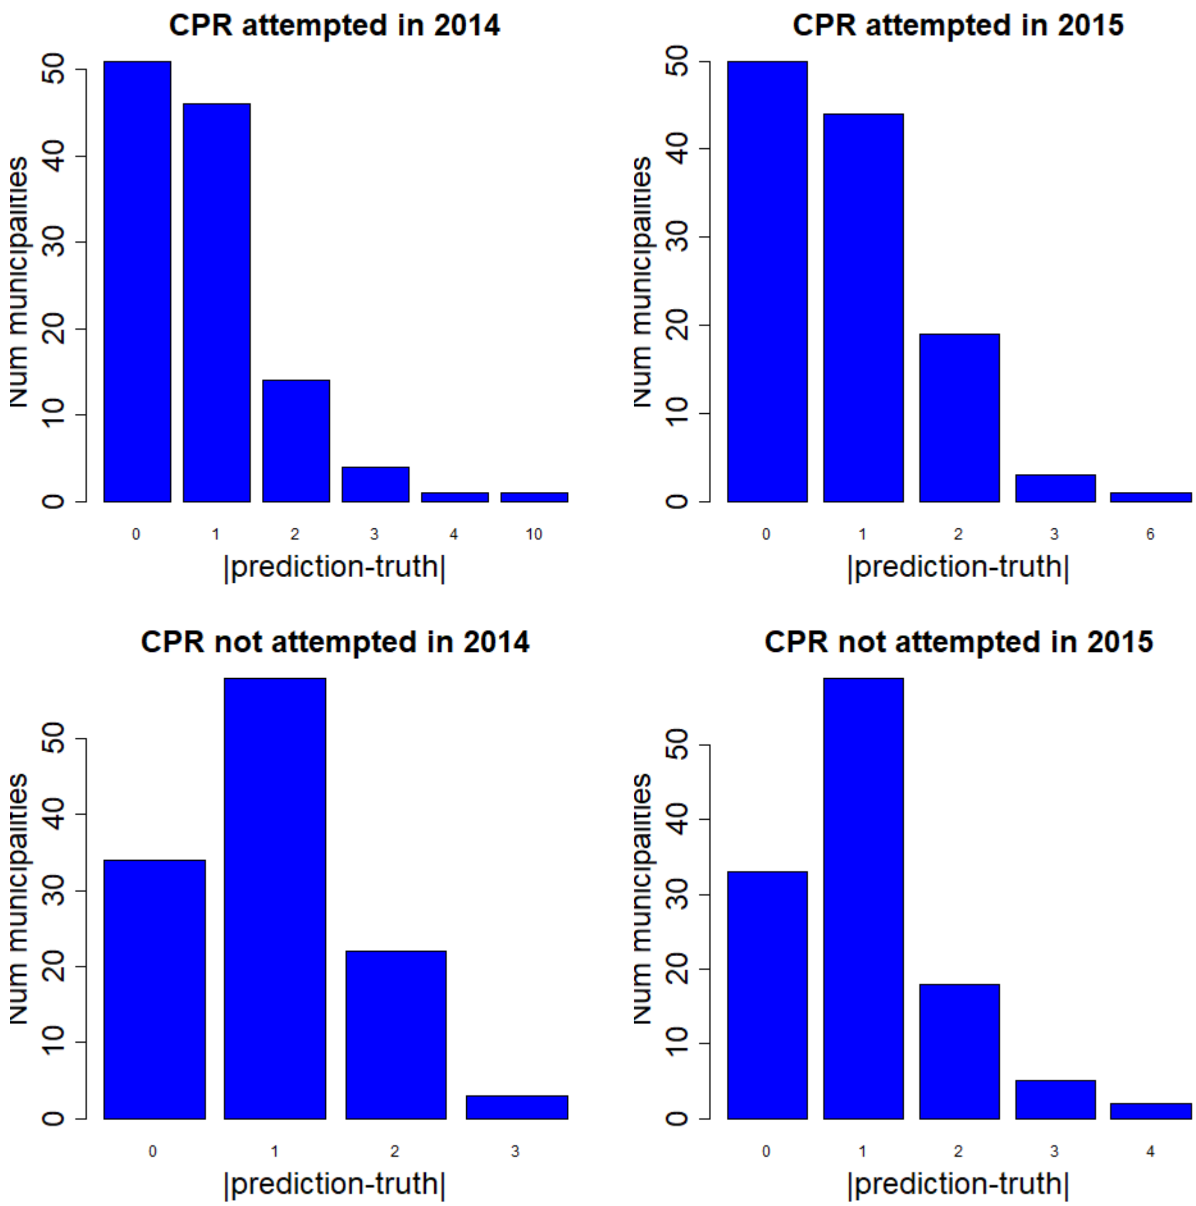

Figure 2: Prediction error barplot for CPR attempted (first row) and CPR not attempted (second row) OHCAs in 2014 (first column) and 2015 (second column)

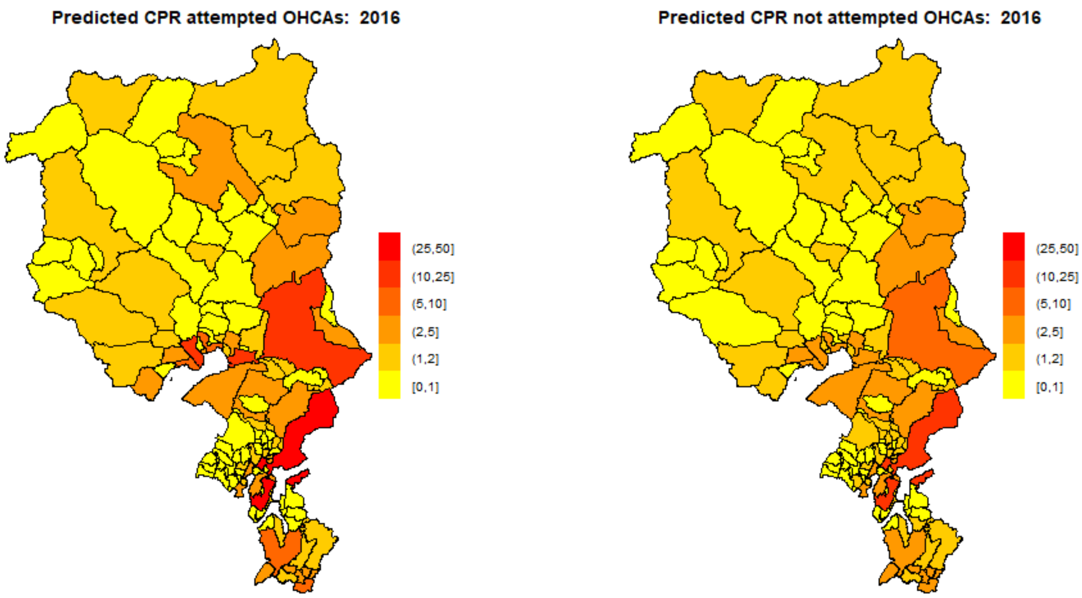

Figure 3: Prediction map for CPR attempted (left) and not (right) OHCAs in 2016.

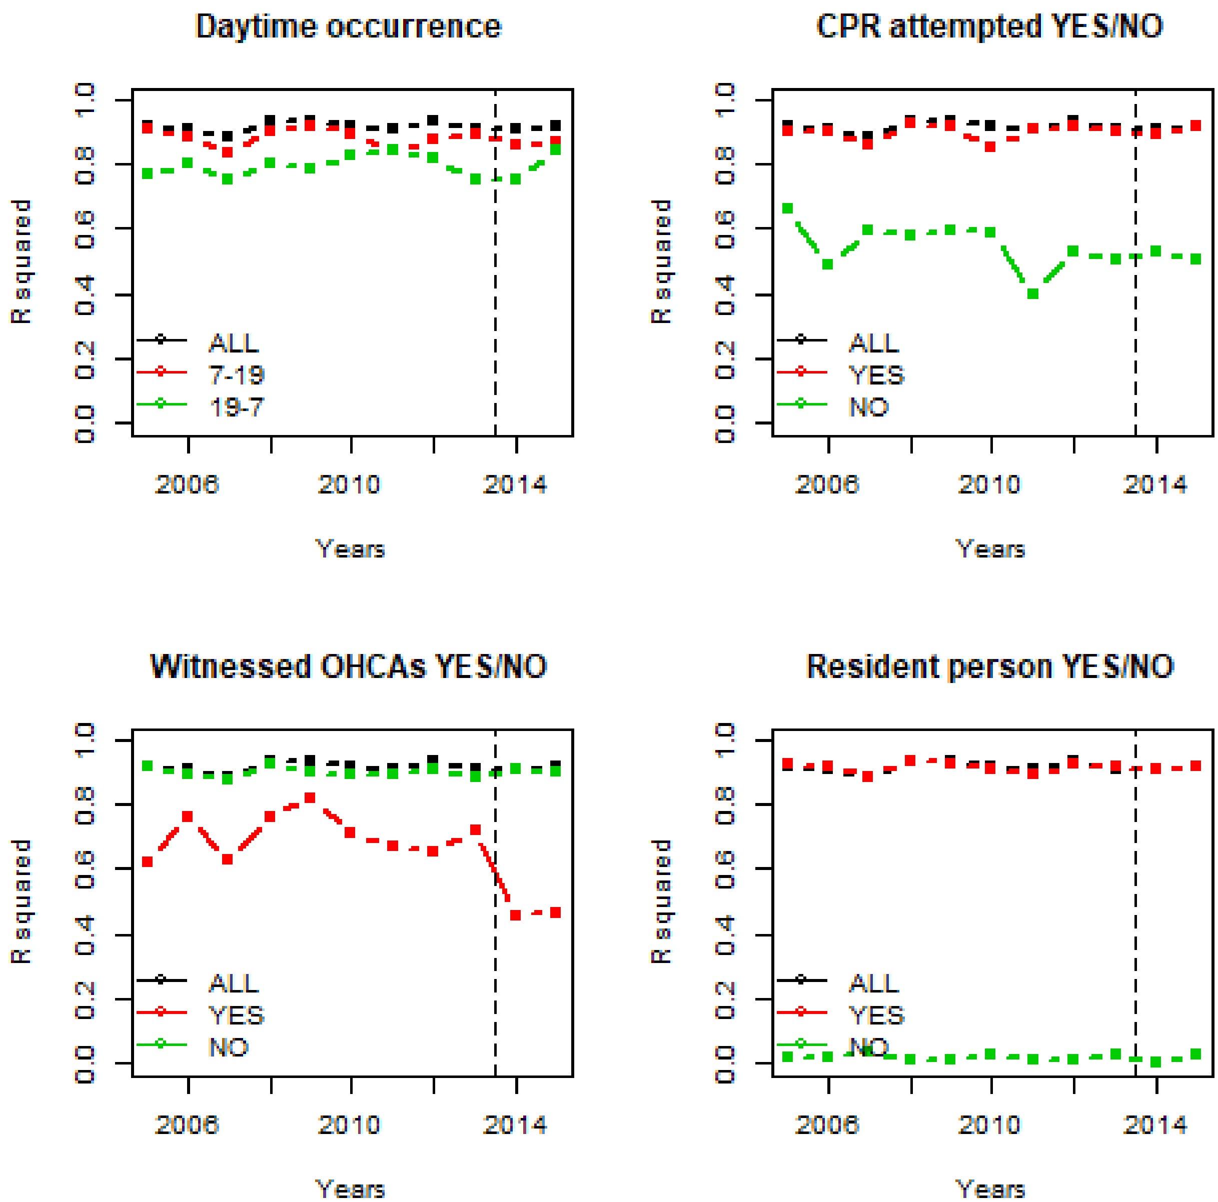

Figure 4: Time evolution of R^2 of true OHCAs on predicted ones, for all subcategories and highlighting the two years of 2014 and 2015 in the validation set.

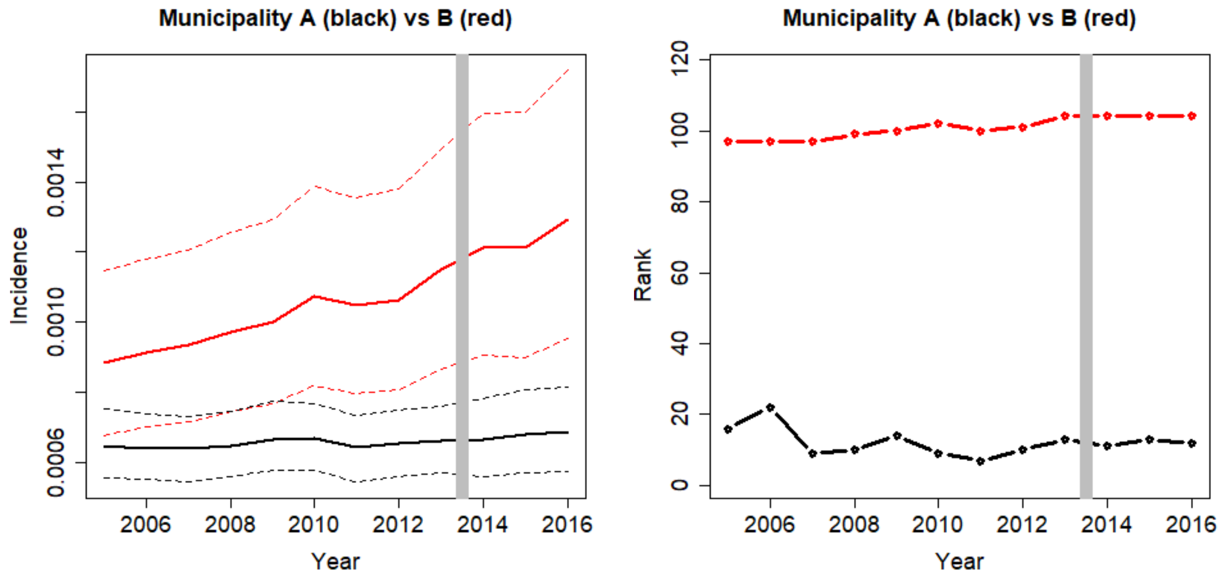

Figure 5: Left: Point estimates (solid lines) and corresponding 95% credible intervals (dashed lines) for OHCA incidences in Municipality A (black lines) and B (red lines). Right: Time evolution of $\mathcal{S}_{it}$-based rank for Municipality A (black) and B (red). Values for 2014, 2015 and 2016 are predicted.

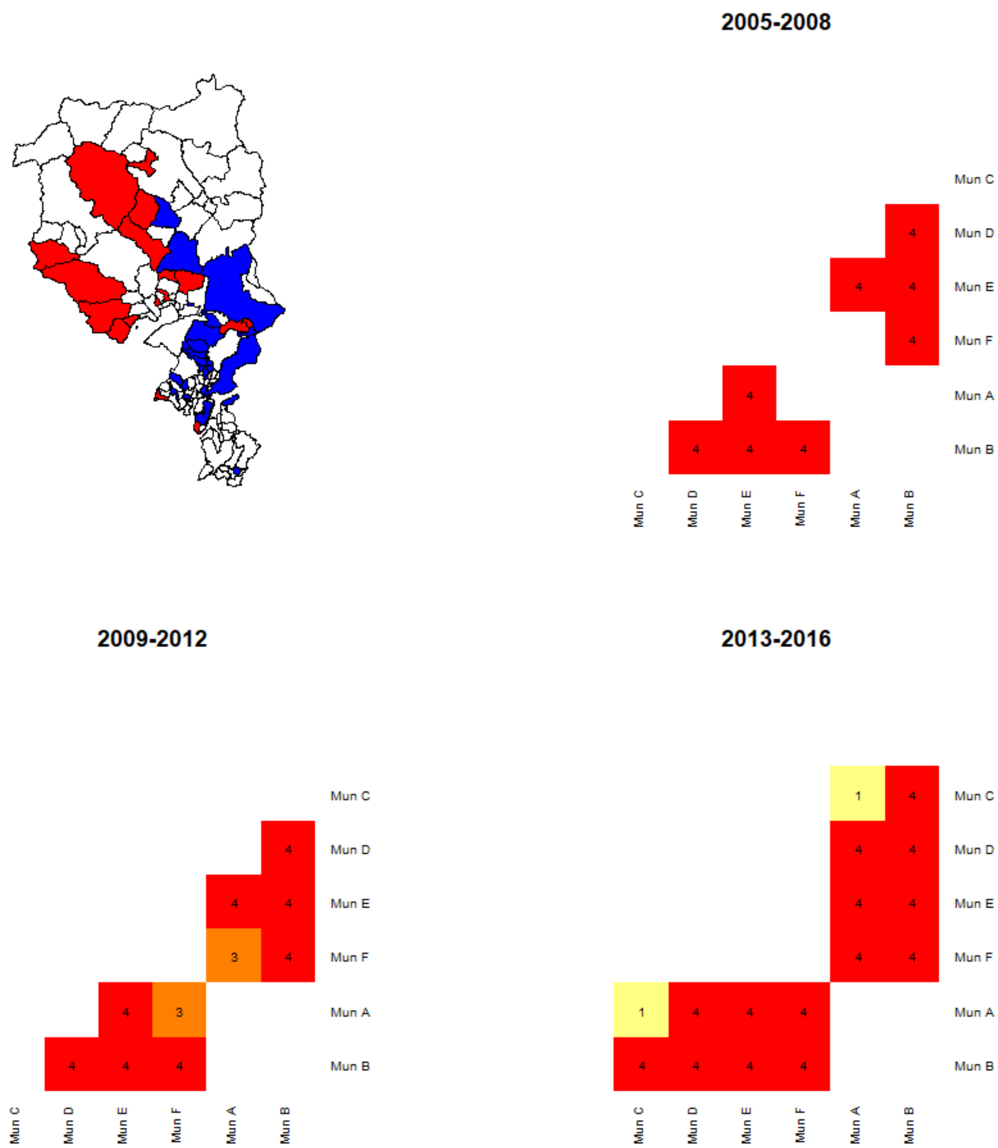

Figure 6: Top-left: Ticino map with municipalities having ten highest (in blue) and lowest (in red) scores. Top-right and bottom: heatmaps for six municipalities in the years 2005-2008, 2009-2012 and 2013-2016. From yellow to red to indicate higher distance between municipalities, with cell number denoting number of years in the period at which incidences are different a posteriori.

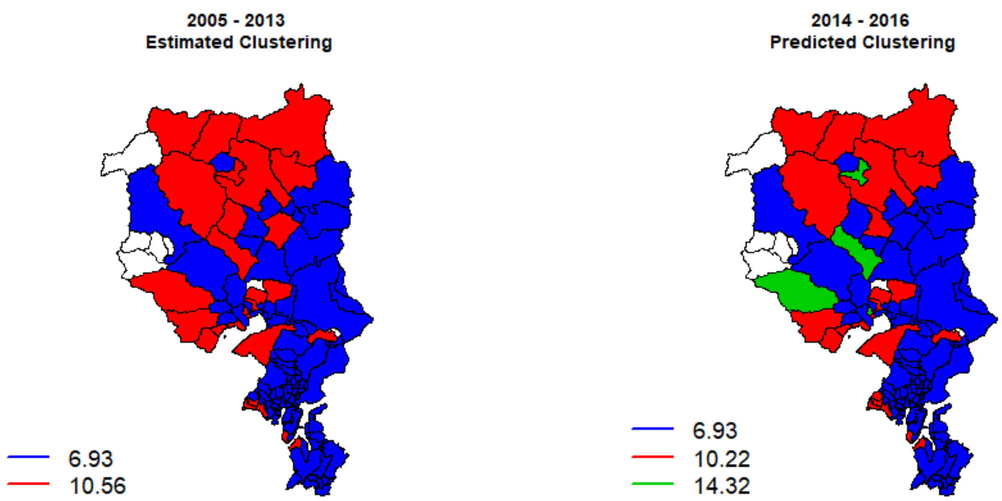

Figure 7: OHCA incidence hierarchical clustering in the training (2005-2013) and test (2014-2016) sub-periods, with respective incidence rates $\times$ 10,000.

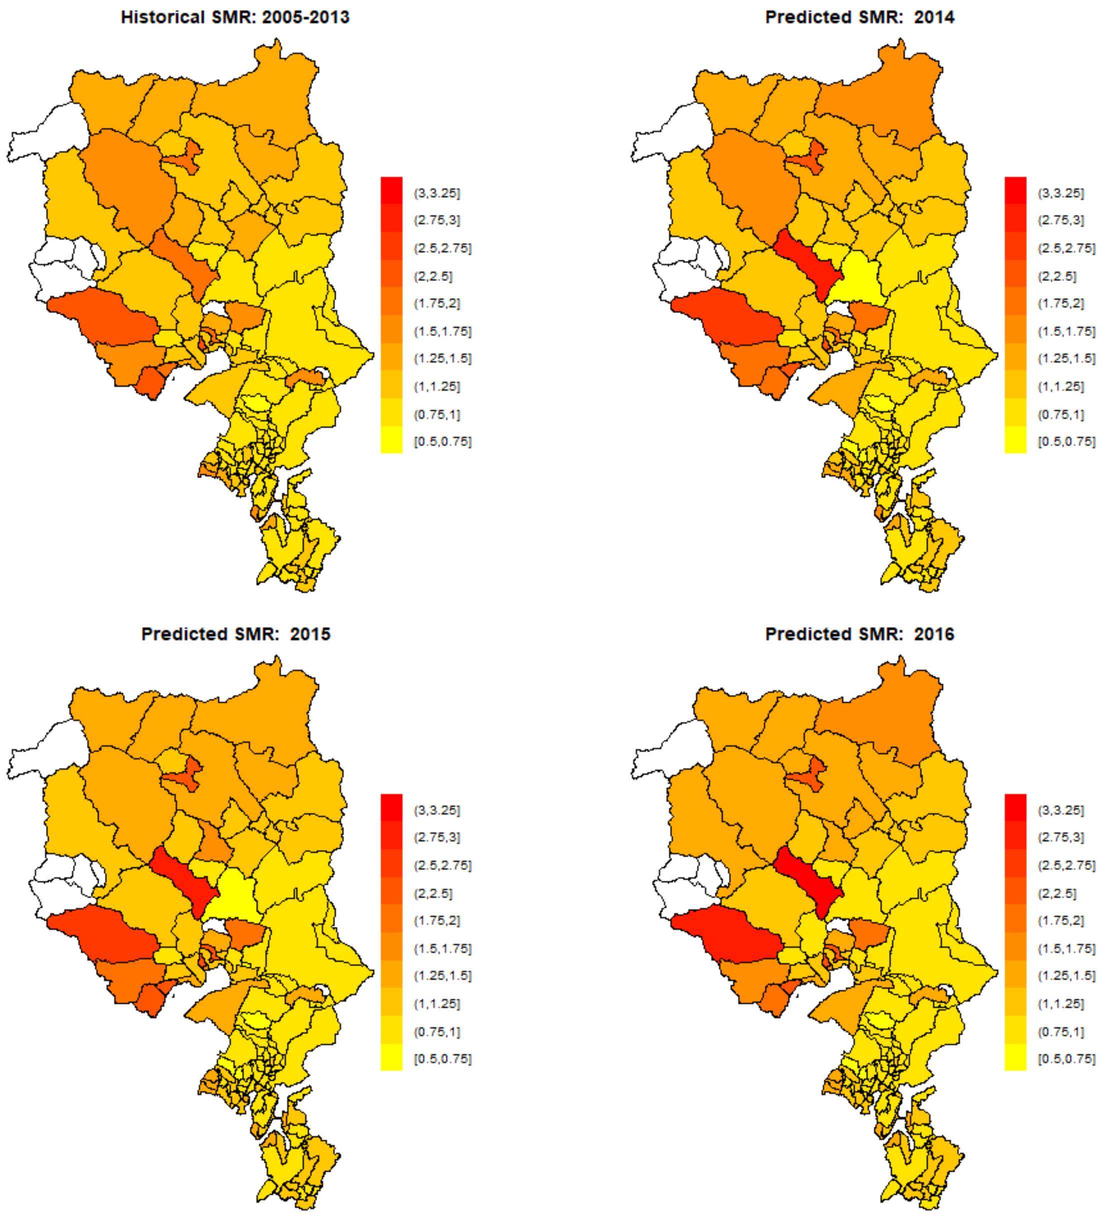

Figure 8: Standardized Morbidity Ratio maps, for all OHCAs, estimated over the period 2005-2013 and predicted for 2014, 2015 and 2016.