A framework for estimating and visualising excess mortality during the COVID-19 pandemic

by G. Konstantinoudis, V. Gómez-Rubio, M. Cameletti, M. Pirani, G. Baio, M. Blangiardo

Abstract COVID-19 related deaths underestimate the pandemic burden on mortality because they suffer from completeness and accuracy issues. Excess mortality is a popular alternative, as it compares observed with expected deaths based on the assumption that the pandemic did not occur. Expected deaths had the pandemic not occurred depend on population trends, temperature, and spatio-temporal patterns. In addition to this, high geographical resolution is required to examine within country trends and the effectiveness of the different public health policies. In this tutorial, we propose a framework using R to estimate and visualise excess mortality at high geographical resolution. We show a case study estimating excess deaths during 2020 in Italy. The proposed framework is fast to implement and allows combining different models and presenting the results in any age, sex, spatial and temporal aggregation desired. This makes it particularly powerful and appealing for online monitoring of the pandemic burden and timely policy making.

Introduction

Estimating the burden of the COVID-19 pandemic on mortality is an important challenge (Weinberger et al., 2020). COVID-19 related deaths are subject to testing capacity and changes in definition and reporting, raising accuracy and completeness considerations (Aburto et al., 2021; Konstantinoudis et al., 2021). In addition, COVID-19 related deaths give an incomplete picture of the burden of the COVID-19 pandemic on mortality, as they do no account for the indirect pandemic effects due to, for instance, disruption to health services (Kaczorowski and Del Grande, 2021). Alternatively, excess mortality has been extensively used to evaluate the impact of the COVID-19 pandemic on mortality (Rossen et al., 2020; Islam et al., 2021; Kontis et al., 2020; Konstantinoudis et al., 2021; Verbeeck et al., 2021).

Excess mortality is estimated by comparing the observed number of deaths in a particular time period with the expected number of deaths under the counterfactual scenario that the pandemic did not occur. Typically this counterfactual scenario is calculated using a comparison period, for instance several previous years (`https://www.euromomo.eu/`). Calculating accurately the expected deaths requires accounting for factors such as population trends, seasonality, temperature, public holidays and spatio-temporal dependencies. While accounting for the above-mentioned factors, most studies to date have estimated excess mortality at national level (e.g., Rossen et al., 2020; Weinberger et al., 2020) and a few have looked across different countries reporting important geographical discrepancies (Islam et al., 2021; Kontis et al., 2020, 2021).

While national level estimates of excess mortality give valuable insights into the total number of excess deaths, they do not allow the evaluation of within country geographical disparities. Such information is essential to understand the country's transmission patterns and effectiveness of local policies and measures to contain the pandemic (Kontopantelis et al., 2021). Temporal trends in excess can substantially differ across regions of the same country, which makes national-based comparisons even more challenging (Blangiardo et al., 2020). Therefore, understanding the burden of the COVID-19 pandemic on mortality requires higher spatial resolution and models that account for spatial, temporal and spatio-temporal dependencies.

When working at high spatio-temporal resolution, data are generally sparse, leading to excess mortality estimates that are highly variable. This is aggravated by the fact that they are subject to spatial and temporal correlation, making it essential to use statistical methods that account for these dependencies in order to provide robust and accurate estimates. The disease mapping framework, which is commonly employed in epidemiology to study the spatio-temporal variation of diseases (Waller et al., 1997; Moraga, 2018), can be exploited to estimate excess mortality at subnational and weekly level. The Bayesian approach is naturally suited in this context, as it is able to model complex dependency structures, as well as to incorporate uncertainty in the data and modelling process. In addition, while fully propagating the uncertainty, it allows us to summarise the results at any desired spatio-temporal aggregation (using for instance coarser geographical units suitable for policy implementation). This in combination with fast approximate methods to inference, such as the Integrated Laplace Approximation (INLA, Rue et al. (2009)), make this framework particularly powerful and appealing for monitoring of the pandemic burden and timely policy making.

In this tutorial, we describe how to run a study on excess mortality at high spatial and temporal resolution using a Bayesian approach and R. This tutorial provides a detailed implementation of the approach followed previously in 5 European regions (Konstantinoudis et al., 2021). Figure 1 illustrates graphically the workflow followed in this paper together with the main r-packages used. Source code for replicating the data wrangling (r-files starting with 01, 02 or 03), analysis (r-files starting with 04) and post-processing steps (r-files starting with 05 or 06 and the App folder) and data files are available from GitHub at https://github.com/gkonstantinoudis/TutorialExcess. This tutorial is structured as follows: we first describe the modelling framework and present the case study in Italy. We then show how to run and evaluate the model, and extract and visualise the results. Finally, we present an R-shiny app which makes the results effectively and easily disseminated.
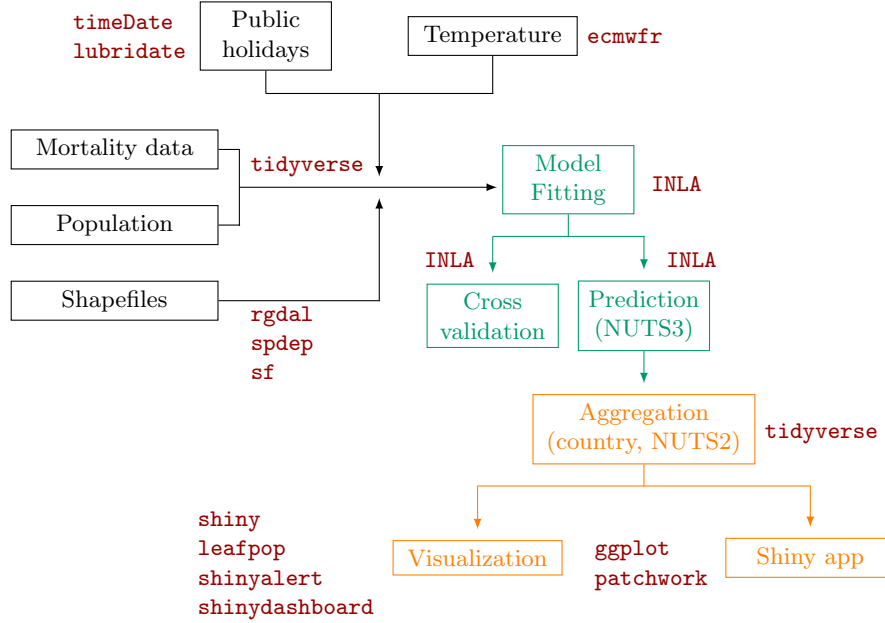

Figure 1: Diagram of the workflow: data wrangling in black, analysis in green, post-processing in orange and main packages in red. NUTS stands for Nomenclature of Territorial Units for Statistics with NUTS3 being the highest spatial resolution available and NUTS2 coarser but appropriate for policy making.

Bayesian hierarchical spatio-temporal model to estimate excess mortality

We propose a Bayesian hierarchical model to quantify spatio-temporal excess mortality under extreme events such as the COVID-19 pandemic, stratified by specific age-sex population groups. To do so, we first describe the statistical model for predicting the number of deaths from all-causes based on historical data, in the counterfactual scenario in which the pandemic did not take place. Then, we show how to estimate the magnitude of excess deaths over a specific period of time, with associated uncertainty, by comparing the predicted versus the actual number of deaths.

Let y_{jtsk} and P_{jtsk} be the number of all-cause deaths and the population at risk for the j-th time point (e.g., week or month) of the t-th year preceding the pandemic, the s-th spatial unit, and k-th age-sex group. Also let x_j and z_{jts} denote the adjustment covariates, respectively public holidays (i.e., $x_j = 1$ if week j contains a public holiday and 0 otherwise) and temperature. We assume that the historical observed number of deaths, conditional on the risk r_{jtsk}, follows a Poisson distribution, with a log-linear model on the risk:

$$y_{jtsk}|r_{jtsk} \sim \text{Poisson}(\mu_{jtsk} = r_{jtsk}P_{jtsk}),$$
$$\log\left(r_{jtsk}\right) = \beta_{0t} + \beta x_j + f(z_{jts}) + b_s + w_j. \tag{1}$$

Here, β_{0t} is the year specific intercept given by $\beta_{0t} = \beta_0 + \epsilon_t$, where β_0 is the global intercept and ϵ_t is an unstructured random effect representing the deviation of each year from the global intercept, which is modelled using independent and identically (iid) distributed Gaussian prior distribution with zero-mean and variance equal to τ_ϵ^{-1}. The parameter β is an unknown regression coefficient

associated to the public holidays covariate x_j. The effect of temperature is allowed to be non-linear $f(\cdot)$ by specifying a second-order random walk (RW2) model:

$$z_{jts} \mid z_{(j-1)ts}, z_{(j-2)ts}, \tau_x \sim \mathrm{Normal}\left(2z_{(j-1)ts} + z_{(j-2)ts}, \tau_z^{-1}\right),$$

where τ_z^{-1} is the variance.

Terms b_s and w_j are spatial and temporal random effects, respectively. We specify the spatial random effect term, b_s, with a convolution prior (Besag et al., 1991), and the temporal random effect term, w_j with a non-stationary in time prior. In detail, we model $\boldsymbol{b}$ using a reparameterisation of the popular Besag-York-Mollié prior, which is a convolution of an intrinsic conditional autoregressive (CAR) model and an iid Gaussian model. Let u_s be the spatially structured component defined by an intrinsic CAR (Besag, 1974) prior $u_s|u_i, i \in \partial_s \sim (\bar{u}, \tau_u^{-1}/m_s)$, where $\bar{u}$ is the local mean and ∂_s and m_s are respectively the set and the number of neighbours of area s, τ_u^{-1} the conditional variance, and v_s the unstructured component with prior $v_s \overset{iid}{\sim} \mathrm{Normal}(0, \tau_v^{-1})$. We model b_s as follows (Besag et al., 1991; Riebler et al., 2016; Konstantinoudis et al., 2020):

$$b_s = \frac{1}{\sqrt{\tau_b}}\left(\sqrt{1-\phi}v_s^\star + \sqrt{\phi}u_s^\star\right)$$

where $u_s^\star$ and $v_s^\star$ are standardised versions of u_s and v_s to have variance equal to 1 (Simpson et al., 2017). The term $0 \leq \phi \leq 1$ is a mixing parameter, which measures the proportion of the marginal variance explained by the structured effect. Finally, we assign to the temporal random effect term, w_j, a Gaussian random walk model of order 1 (RW1), specified as:

$$w_j \mid w_{j-1}, \tau_w \sim \mathrm{Normal}(w_{j-1}, \tau_w^{-1}).$$

The Bayesian representation of the above model is completed once we select priors for the fixed effects β_0 and β and the hyperparameters: $\tau_\epsilon, \tau_b, \tau_w$, and ϕ. For the fixed effects we selected minimally informative Normal distributions whereas, we specified "penalising complexity" (PC) priors (Simpson et al., 2017) for the hyperparameters. PC priors are defined by penalising deviations from a "base" model (e.g., specified in terms of a specific value) and have the effect of regularising inference, while not implying too strong information. Technically, PC priors imply an exponential distribution on a function of the Kullback–Leibler divergence between the base model and an alternative model in which the relevant parameter is unrestricted. This translates to a suitable "minimally informative", regularising prior on the natural scale of the parameter.

In order to quantify the weekly excess mortality at sub-national level for specific age-sex population groups, we need to predict the number of deaths had COVID-19 not occurred. In Bayesian analysis, this can be performed by drawing random samples from the posterior predictive distribution (that is, the distribution of unobserved values conditional on the observed values from previous years). Specifically, letting $\boldsymbol{\theta}$ be the model parameters, $\mathcal{D}$ be the observed data, and y_{jt^*sk} be the count of deaths that we want to predict, we have:

$$p(y_{jt^*sk} \mid \mathcal{D}) = \int p(y_{jt^*sk} \mid \boldsymbol{\theta})p(\boldsymbol{\theta} \mid \mathcal{D})d\boldsymbol{\theta}. \tag{2}$$

Operationally, we first generate random samples from the joint posterior marginal of the fitted linear predictor specified in equations (1) at the highest spatial resolution available (NUTS3 regions; Nomenclature of Territorial Units for Statistics 3 regions, https://ec.europa.eu/eurostat/web/nuts/background/). Successively, we use these to be in turn the mean parameter of a Poisson distribution, to obtain the predicted number of deaths, which represents the baseline number of deaths assuming the pandemic did not take place.

Finally, to estimate the magnitude of excess deaths, the predicted number of deaths is compared against the actual observed number of deaths, allowing the computation of the relative change in the mortality (i.e., relative to what we could expect if the pandemic did not occur). This is obtained by (i) subtracting the predicted number of deaths from the observed number of deaths in each time point t and spatial unit s (number of excess deaths or NED), and (ii) dividing NED by the predicted number of deaths for each sample and multiplying by 100 (% relative excess mortality or REM).

Bayesian inference for the model is computed using Integrated Nested Laplace Approximation (INLA; Rue et al. (2009), which performs approximate Bayesian inference on the class of latent Gaussian models (Rue and Held, 2005). Unlike simulation based Markov chain Monte Carlo method, INLA is a deterministic algorithm, which employs analytical approximations and efficient numerical integration schemes to provide accurate approximations of the posterior distributions in short computing times. The INLA software is provided through the R package INLA, which can be downloaded from https://www.r-inla.org/download-install.
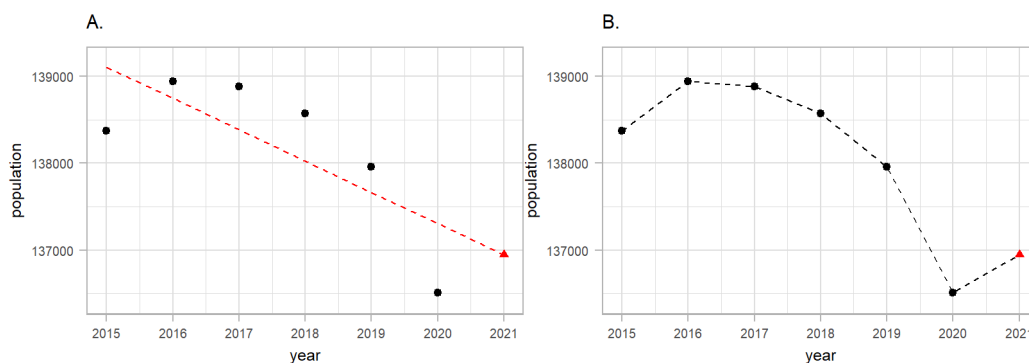
Figure 2: A schematic representation of the weekly population estimation procedure focusing on females aged 40-49 in Venice during 2015-2019 as an example. On panel A we show how we used the historical data (black points) and fit a linear regression (red dashed line) to predict 2021 (red triangle). On the panel B we show how we predicted weekly population by drawing lines between the years.

Motivating example: Italy

Outcome data

We retrieved all-cause mortality data during 2015-2020 in Italy from the Italian National Institute of Statistics (https://www.istat.it/). Data were available weekly (ISO week), by age (5-year age groups), sex and NUTS3 regions. As the COVID-19 mortality rates increase with age, we aggregated mortality counts based on the following age groups: <40, 40-59, 60-69, 70-79 and 80 years and older (Davies et al., 2021).

Population data

Population data in Italy during 2015-2020 were retrieved from the Italian National Institute of Statistics. The data represent the population in Italy on January 1st of every year and are available by age (5-year age groups), sex and NUTS3 regions. To retrieve weekly population, we performed linear interpolation by the selected age groups (<40, 40-59, 60-69, 70-79 and 80+), sex and NUTS3 regions using populations at January 1st of the current and the next year. Population counts on January 1st 2021, which takes COVID-19 deaths in 2020 into consideration, were available at the time of analysis. Our goal was, however, to predict mortality for 2020 had the pandemic not occurred. Thus we performed an additional linear interpolation by age, sex and NUTS3 regions to predict the population at January 1st 2021, using the years 2015-2020 (Figure 2, panel A). Object pop is a `tibble` containing the NUTS3 region ID (`NUTS318CD`), age group (`ageg`), sex (`sex`), year (`year`) and population counts (`population`):

```
pop
# A tibble: 6,420 x 5
# Groups:   ageg, sex [10]
   NUTS318CD ageg    sex       year population
   <chr>     <fct>   <chr>    <dbl>      <dbl>
 1 TO        less40  female    2015     435758
 2 TO        less40  female    2016     427702
 3 TO        less40  female    2017     420498
 4 TO        less40  female    2018     413141
 5 TO        less40  female    2019     406937
 6 TO        less40  female    2020     402768
 7 TO        less40  male      2015     449605
 8 TO        less40  male      2016     443941
 9 TO        less40  male      2017     439522
10 TO        less40  male      2018     433365
# ... with 6,410 more rows
```

We can use the following code (based on `tidyverse` and "piping" principles) to calculate the population of the 1st of January 2021 by NUTS3 regions, sex and age:
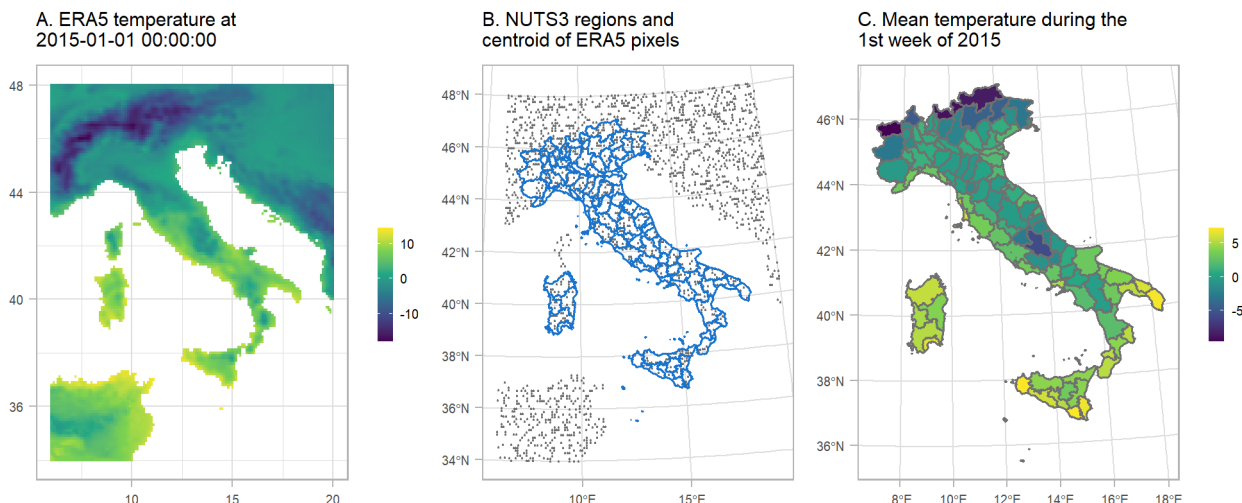
Figure 3: Schematic representation of the temperature misalignment procedure. A. The temperature obtained by ERA5 at 2015-01-01 00:00:00. B. NUTS3 regions in blue and a sample of the centroids of the pixels from the ERA5 raster. C. Mean temperature per NUTS3 region during the 1st week of 2015.

```
pop %>% group_by(NUTS3, sex, age) %>%
summarise(pop = as.vector(coef(lm(pop ~ year)) %*% c(1, 2021))) %>%
mutate(year = 2021) -> pop2021
```

Once we obtained the year 2021 we performed an additional linear interpolation to calculate weekly number of population as shown on Figure 2, panel B.

Covariates data

We used covariates related with ambient temperature and national holidays to help the model predictions. Data on air-temperature during 2015-2020 in Italy at 2m above the surface of land were retrieved from the ERA5 reanalysis data set of the Copernicus climate change program (Hersbach et al., 2020). The geographical resolution of the ERA5 estimates is $0.25° \times 0.25°$ (panel A of Figure 3). We calculated the weekly mean by the centroids of the $0.25° \times 0.25°$ grid (panel B of Figure 3) and then averaged the weekly temperature over the ERA5 centroids that overlay with the NUTS3 regions (panels B and C of Figure 3).

Fitting the model

The modelling process in `INLA` consists of three main steps: (1) the selection of priors; (2) definition of the model "formula" (which sets out the expression for the generalised linear predictor); and (3) the call to the main function `inla`, which computes the estimates.

In particular, we constructed the PC priors for $\sigma_\epsilon = \sqrt{1/\tau_\epsilon}, \sigma_z = \sqrt{1/\tau_z}, \sigma_b = \sqrt{1/\tau_b}$ and $\sigma_w = \sqrt{1/\tau_w}$ based on statement that it is unlikely to have a relative risks higher than $\exp(2)$ based solely on spatial, yearly and seasonal variation, Figure 4, panel A. For the mixing parameter ϕ, we set $\Pr(\phi < 0.5) = 0.5$ reflecting our lack of knowledge about whether overdispersion or strong spatial autocorrelation should dominate the field b, Figure 4, Figure 4, panel B.

These assumptions can be encoded using the following code:

```
# Defines the priors
hyper.bym <- list(
theta1 = list('PCprior', c(1, 0.01)),
theta2 = list('PCprior', c(0.5, 0.5))
)
hyper.iid <- list(theta = list(prior="pc.prec", param=c(1, 0.01)))

# Defines the model "formula"
```
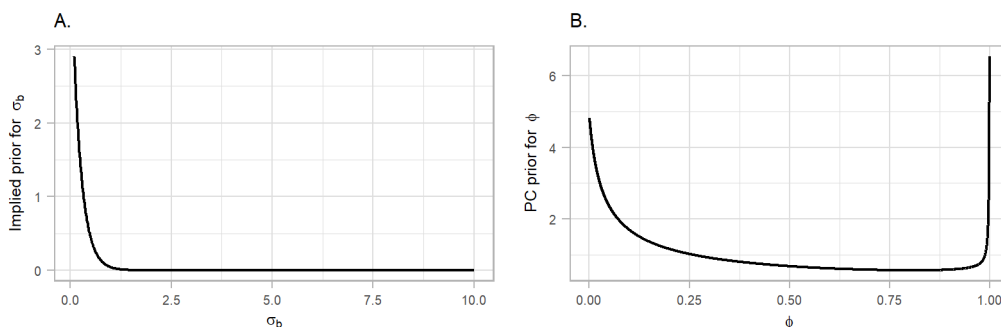
Figure 4: Penalised complexity (PC) priors for the hyperparameters of the spatial field. A. The implied PC prior for the standard deviation (as the original scale of the prior is the precision). B. The PC prior for the mixing parameter ϕ.

```
formula =
deaths ~ 1 + offset(log(population)) + hol +
f(id.tmp, model='rw2', hyper=hyper.iid, constr = TRUE, scale.model = TRUE) +
f(id.year, model='iid', hyper=hyper.iid, constr = TRUE) +
f(id.time, model='rw1', hyper=hyper.iid, constr = TRUE, scale.model = TRUE,
cyclic = TRUE) +
f(id.space, model='bym2', graph="W.adj", scale.model = TRUE, constr = RUE,
hyper = hyper.bym)

control.family=inla.set.control.family.default()

# Calls INLA to fit the model
inla.mod = inla(formula,
data=dat,
family="Poisson",
verbose = TRUE,
control.family=control.family,
control.compute=list(config = TRUE),
control.mode=list(restart=T),
num.threads = round(parallel::detectCores()*.8),
control.predictor = list(link = 1))
```

After fitting the model, we take 1000 samples from the (approximated) posterior distribution of the linear predictor and we use each drawn sample as the mean of a Poisson distribution to retrieve the predicted mortality counts:

```
post.samples <- inla.posterior.sample(n = 1000, result = inla.mod)
predlist <- do.call(cbind, lapply(post.samples, function(X)
exp(X$latent[startsWith(rownames(X$latent), "Pred")])))

pois.samples <- apply(predlist, 2, function(Z) rpois(n = length(Z), lambda = Z))
```

This allows us to estimate the entire predictive posterior distribution of the mortality counts incorporating both the sampling and the linear predictor uncertainty.

Model validation

To examine model validity, we performed a cross validation leaving out one historical year at a time and predicting the weekly number of deaths by NUTS3 regions for the year left out. For each stratum, we calculated the correlation between observed and fitted and a coverage probability, i.e. the probability that the observed death fall into the 95% credible interval (95% CrI) of the predicted.

Results

Cross validation

We overall found good predictive ability of our model. The correlation estimates varies from 0.38 (95% CrI: 0.36, 0.40) in females 40< to 0.96 (95% CrI: 0.94, 0.96) in females 80> and coverage probability from 0.91 in females 40< to 0.95 in females and males 40-59, 60-69, 70-79:

```
res
#             Correlation Coverage
# less40F 0.38 (0.36, 0.40)    0.91
# 40-59F  0.78 (0.77, 0.78)    0.95
# 60-69F  0.83 (0.82, 0.84)    0.95
# 70-79F  0.91 (0.90, 0.92)    0.95
# 80plusF 0.96 (0.94, 0.96)    0.93
# less40M 0.51 (0.49, 0.52)    0.93
# 40-59M  0.84 (0.83, 0.84)    0.95
# 60-69M  0.87 (0.85, 0.88)    0.95
# 70-79M  0.92 (0.91, 0.93)    0.95
# 80plusM 0.95 (0.93, 0.95)    0.94
```

Expected number of deaths

The object `pois.samples.list` contains 1000 samples from the posterior predictive distribution 2, i.e. 1000 samples of the expected number of deaths by age, sex, NUTS3 regions and week, had the pandemic not occurred. We can access the different age-sex groups as follows:

```
names(pois.samples.list)
# "F_less40" "F_40_59" "F_60_69" "F_70_79" "F_80plus"
# "M_less40" "M_40_59" "M_60_69" "M_70_79" "M_80plus"
```

where F stands for females and M for males across the different age groups. To get an idea about the structure of the data, we can use the `head()` function for females 60-69 and check the first 10 samples of excess deaths for the first 6 weeks of 2020 in the 001 region

```
pois.samples.list$F_60_69 %>%
select(paste0("V", 1:10), EURO_LABEL, ID_space, year) %>%
head()
#     V1 V2 V3 V4 V5 V6 V7 V8 V9 V10 EURO_LABEL ID_space year
# 262 13 20 20 17 19 20  9 20 27  23   2020-W01      001 2020
# 263 17 21 13 30 20 18 21 19 12  24   2020-W02      001 2020
# 264 17 18 18 22 13 18 30 22 24  16   2020-W03      001 2020
# 265 21 15 23 12 21 14 18 23 18  12   2020-W04      001 2020
# 266 18 16 15 17 14 16 20 25 16  19   2020-W05      001 2020
# 267 15 18 11 22 21 13 18 13 17  19   2020-W06      001 2020
```

We can also calculate median and 95% CrI expected number of deaths for this specific age-sex group over the year by NUTS3 region:

```
pois.samples.list$F_60_69 %>%
select(starts_with("V"), "ID_space") %>%
        group_by(ID_space) %>%
        summarise_all(sum) %>%
        rowwise(ID_space) %>%
        mutate(median = median(c_across(V1:V1000)),
LL = quantile(c_across(V1:V1000), probs= 0.025),
UL = quantile(c_across(V1:V1000), probs= 0.975)) %>%
        select(ID_space, median, LL, UL) %>%
        head()

# A tibble: 107 x 4
# Rowwise:  ID_space
#   ID_space median    LL    UL
#   <chr>     <dbl> <dbl> <dbl>
# 1 001         866  757.  985.
# 2 002        76.5    57    99
```

```
#  3 003        148    118    180
#  4 004        226    189.   265.
#  5 005         92     70    117
#  6 006        188    156.   228
#  7 007         49     36     67
#  8 008         92     72    117
#  9 009        120.    95    146
# 10 010        363    310.   425
# ... with 97 more rows
```

Excess mortality

The above results can be combined in different ways using the functions get2020data() and get2020weeklydata() to calculate excess mortality. The function get2020data() aggregates over the entire country, NUTS2 regions, sex, age and time resulting in the object d, whereas the function get2020weeklydata() over the entire country, NUTS2 regions, sex and age but not over time resulting in the object d_week.

```
names(d)
# "province" "region"    "country"
names(d$province)
# "none"    "age"      "sex"      "agesex"
names(d$province$age)
# "40<"     "40-59" "60-69" "70-79" "80+"
names(d$province$sex)
# "F" "M"
names(d$province$agesex)
# "F40<" "F40-59" "F60-69" "F70-79" "F80+"    "M40<"    "M40-59" "M60-69" "M70-79" "M80+"
```

Province stands for the NUTS3 regions (the resolution we used to fit the models), region for the NUTS2 (coarser than NUTS3, appropriate for policy making) and country for the nationwide aggregation. Within these aggregations users can select the option "none" being the total aggregation by age and sex, "age" by sex, "sex" by age and "agesex" refers to no age and sex aggregation. The objects d and d_week have similar structure and contain summary statistics for REM and NED and posterior probabilities of a positive REM or NED:

```
head(d$province$none)
# Simple feature collection with 6 features and 24 fields
# Geometry type: POLYGON
# Dimension:     XY
# Bounding box:  xmin: 6.626865 ymin: 44.06028 xmax: 9.21355 ymax: 45.95041
# CRS:           +proj=longlat +datum=WGS84
#   ID_space SIGLA    DEN_UTS observed population mean.REM median.REM    sd.REM
# 1      001    TO     Torino    32478 2226317.2 15.42716   15.40348 2.951326
# 2      002    VC   Vercelli     3216  168761.7 25.58699   25.57595 4.227239
# 3      003    NO     Novara     5274  364529.4 17.45157   17.51337 3.563529
# 4      004    CN      Cuneo     8716  585479.3 14.28525   14.25575 3.137300
# 5      005    AT       Asti     3745  211437.1 20.21055   20.14758 3.873325
# 6      006    AL Alessandria    7916  416090.6 21.99010   21.93469 3.454054
#      LL.REM   UL.REM exceedance.REM median.REM.cat exceedance.REM.cat median.pred
# 1  9.666010 21.52144             1       [15%, 20%)          (0.8, 1]       28144
# 2 17.327303 33.61169             1            20%>          (0.8, 1]        2561
# 3 10.423095 24.18324             1       [15%, 20%)          (0.8, 1]        4488
# 4  8.272284 20.53825             1       [10%, 15%)          (0.8, 1]        7629
# 5 12.628196 28.38643             1            20%>          (0.8, 1]        3117
# 6 15.107077 28.99179             1            20%>          (0.8, 1]        6492
#   LL.pred UL.pred mean.NED median.NED     sd.NED    LL.NED   UL.NED exceedance.NED
# 1   26727   29630 4322.414     4335.0 719.41805 2862.625 5751.850              1
# 2    2407    2743  652.324      655.0  86.38414  474.950  809.025              1
# 3    4247    4783  779.497      786.0 136.79073  497.825 1027.050              1
# 4    7231    8053 1083.727     1087.5 209.56953  665.925 1485.100              1
# 5    2917    3329  626.401      628.0 100.52396  419.900  828.025              1
# 6    6137    6880 1421.752     1424.0 183.90238 1038.925 1779.175              1
#   median.NED.cat exceedance.NED.cat                              geometry
# 1          1000>           (0.8, 1] POLYGON ((7.859044 45.59758...
```

```
# 2   [500, 1000)      (0.8, 1] POLYGON ((8.204465 45.93567...
# 3   [500, 1000)      (0.8, 1] POLYGON ((8.496878 45.83934...
# 4        1000>       (0.8, 1] POLYGON ((7.990897 44.82381...
# 5   [500, 1000)      (0.8, 1] POLYGON ((8.046805 45.12815...
# 6        1000>       (0.8, 1] POLYGON ((8.405489 45.20148...
```

Notice that the object d$province$none is a simple feature collection, making mapping it straightforward.

In Figure 5 (plots 1A, 2A and 3A) we show the median posterior of REM for total age and sex at the national, NUTS2 and NUTS3 regional level. For these plots we used the object d with the selection "none" and plot the median REM (median.REM), for example for 3A:

```
prov <- "Foggia"
d$province$none %>% filter(DEN_UTS == prov) %>%
select(geometry) %>%
ggplot() +
geom_sf(data = d$province$none, aes(fill = median.REM.cat)) +
geom_sf(fill = NA, col = col.highlight, size = .8) +
scale_fill_manual(values=colors, name = "", drop=FALSE) + theme_light() +
ggtitle(paste0("3A. NUTS3 regions: ", prov))
```

Overall, the REM in Italy during 2020 was between 5-10%, Figure 5, panel 1A. When the higher geographical resolution is assessed, it is revealed that north and in particular Lombardia is the region hit the worst, with the REM exceeding 20%, Figure 5, panels 1B and 1C. Figure 6 shows a measure of uncertainty of the REM, now for the different age groups and both sexes (selection "age"). The probability of a positive excess (exceedance.REM) in older people and in the north of Italy is larger than 0.80, Figure 6.

Panels 1B and 1C of Figure 5 show the median temporal nationwide excess together with 95% CrI by sex after aggregating the different age groups (using d_week and the "sex" selection). We observe a clear first pandemic wave during March and May and a second one during mid October and December in 2020. During the first pandemic wave, there were weeks when the median relative excess death reached almost 100% in males, Figure 5.

Panels 2B, 2C, 3B and 3C of Figure 5 show the median spatio-temporal excess together with 95% CrI by sex after aggregating over the different age groups. Panels 2B and 2C highlight the region of Puglia, where during the first wave of the pandemic experienced a positive, similar with the nationwide excess. When we increase the spatial resolution in panels 3B and 3C, we highlight the region of Foggia, where there was insufficient evidence of a positive excess during the first wave, but strong during the second.

Shiny Web-Application

To be able to effectively examine and communicate the different aggregation levels of the output of our modelling framework, we have also developed a Shiny Web-Application (WebApp), Figure 7. The WebApp provides spatial, temporal and spatio-temporal analysis tabs, and within each tab there are plots and summary statistics for the level of aggregation selected from the drop-down menu. Users can select across different variables (relative excess mortality or number of excess deaths), statistics (median or posterior probability), sex (males, females or both), age group (40<, 40-59, 60-69, 70-79, 80> and all) and different geographical level (national, NUTS2 or NUTS3 regions). Summary statistics for each area are available and they are displayed in a pop-up window, which is activated by clicking on the area of interest. In addition, graphical pop-ups are provided to show the weekly estimates for each area with the leafpop R-package (Appelhans and Detsch, 2021), in the spatio-temporal analysis tab.

The WebApp that we have developed is hosted at http://atlasmortalidad.uclm.es/italyexcess.

Summary

This tutorial provides a detailed implementation of the framework followed previously to calculate excess mortality during the COVID-19 pandemic in 5 European regions (Konstantinoudis et al., 2021). We have proposed a Bayesian framework for estimating excess mortality and shown how to use **R** and **INLA** to retrieve fast and accurate estimates of the excess mortality. The proposed framework also allows combining different models and presenting the results in any age, sex, spatial and temporal aggregation desired. We have given a practical example of how to use the proposed framework
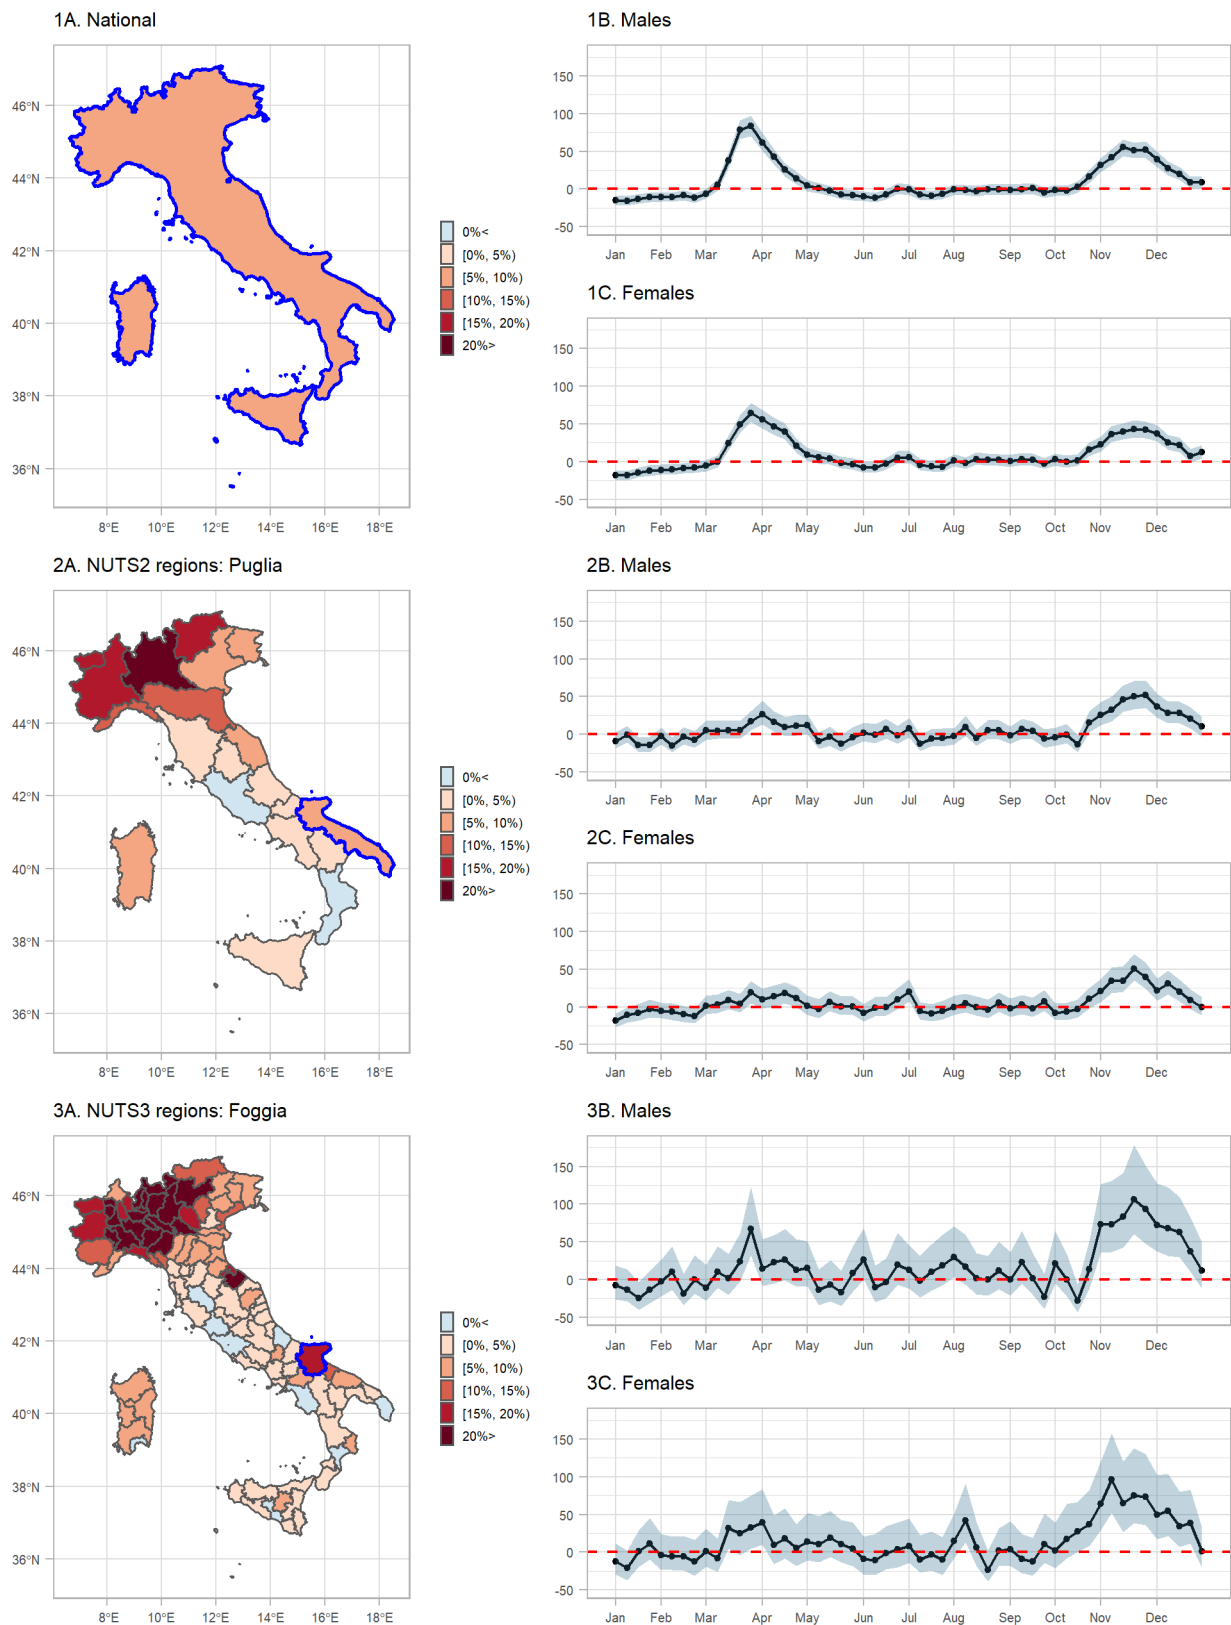
Figure 5: Median relative excess mortality by spatial region during 2020 at: 1A. nationwide, 2A. NUTS2 and 3A. NUTS3 level, and weekly median relative excess mortality by sex during 2020 in: 1B. males nationwide, 1C. females nationwide, 2B. males in Puglia, 2C. females in Puglia, 3B. males in Foggia and 3C. females in Foggia.
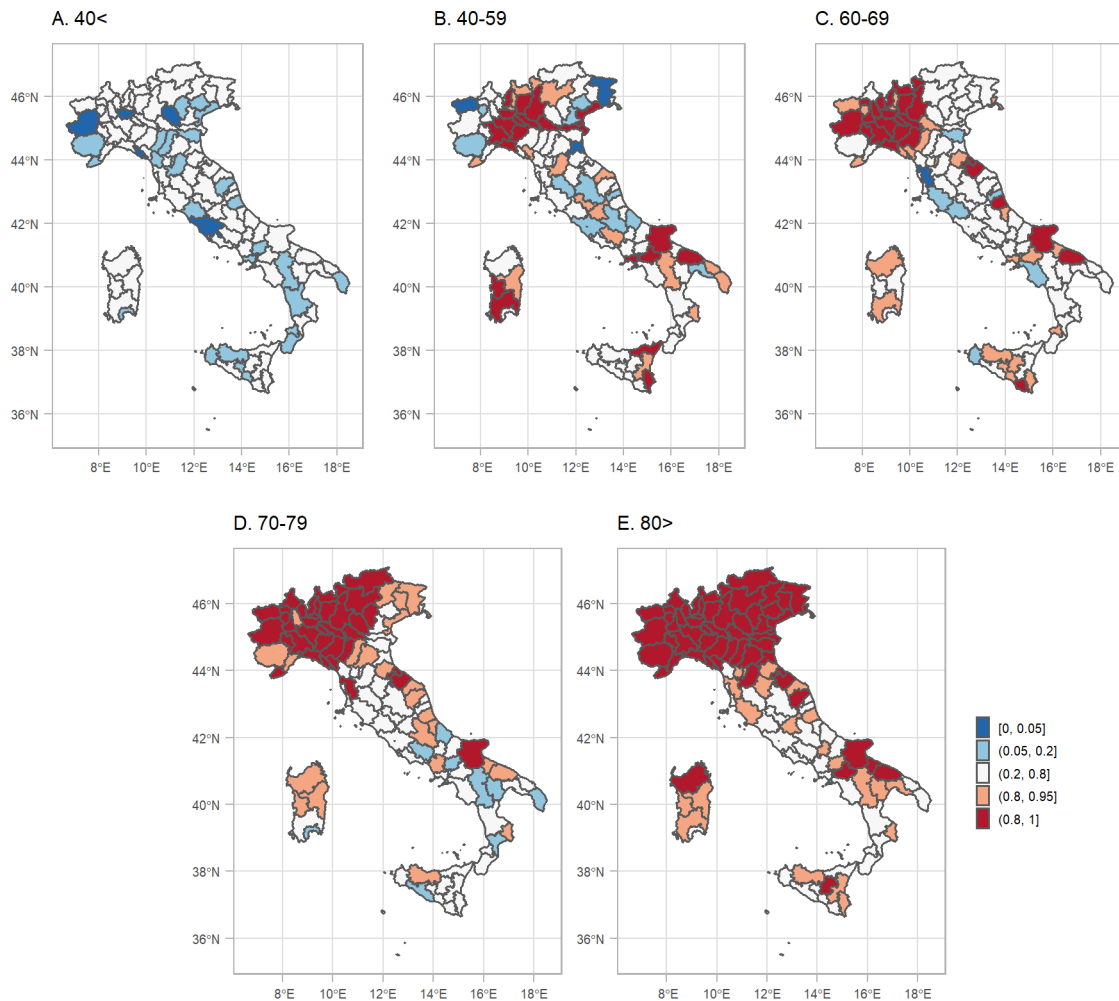
Figure 6: Posterior probability that the relative excess mortality is positive for both sexes during 2020 by age group and NUTS3 region.

modelling the excess mortality during the 2020 COVID-19 pandemic in Italy at small area level. We also developed a Webapp to effectively communicate the results. This framework can be extended to monitor mortality from other extreme events for instance natural hazards such as hurricanes (Acosta and Irizarry, 2020). All the above make the proposed framework particularly powerful, generalisable and appealing for online monitoring of the pandemic burden and timely policy making. Potential extensions include different ways of modelling the younger age groups to increase the predictive ability, for instance using a zero-inflated Poisson.
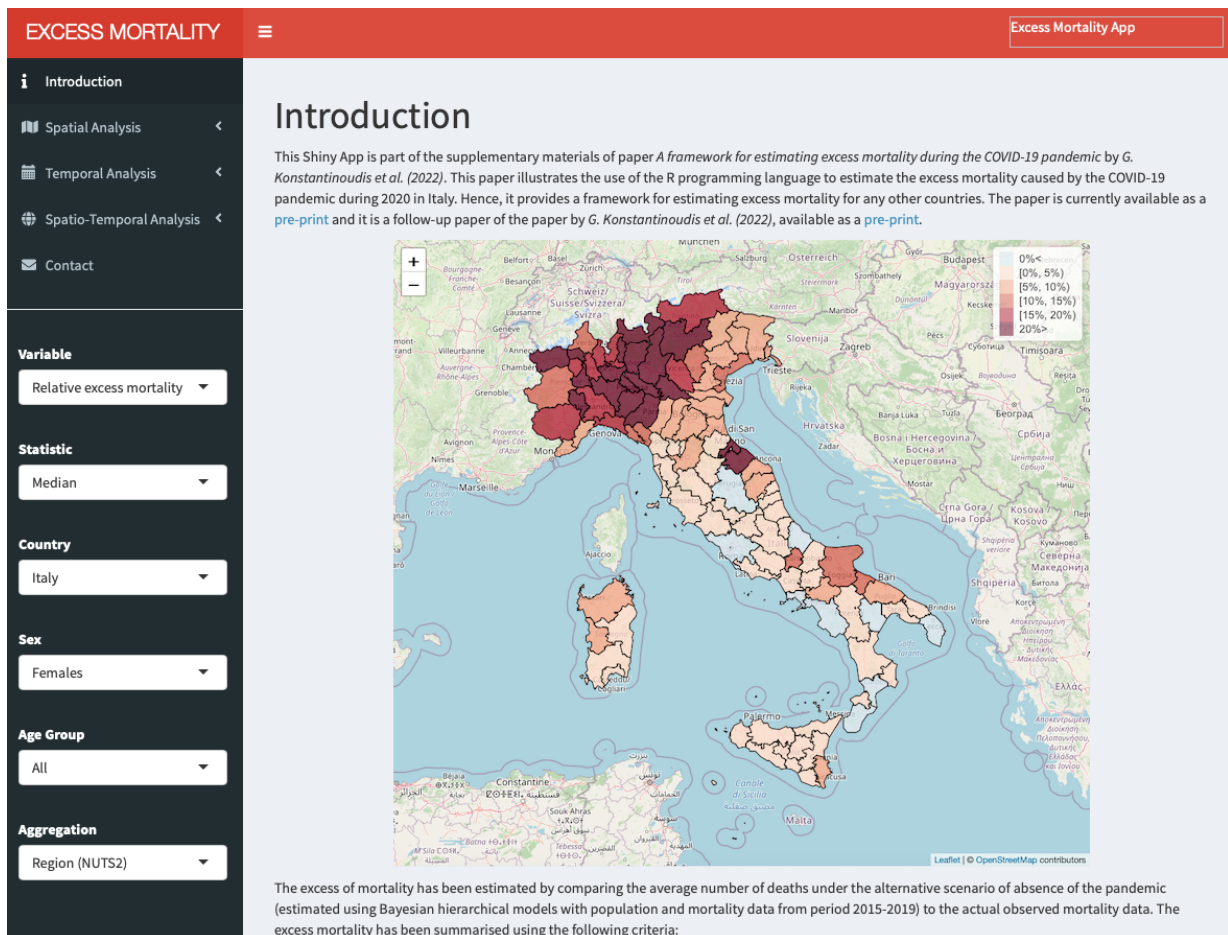
Acknowledgements

Figure 7: Shiny App developed to navigate through the excess mortality estimates during 2020 in Italy across the different aggregations.

Author contributions

V.G.R. conceived the study. M.B. supervised the study. G.K. developed the initial study protocol and discussed it with M.B., M.C., M.P. and G.B.. G.K. developed the statistical model, prepared the population and covariate data and led the acquisition of mortality data. M.C. and V.G.R. validated and modified accordingly the code. G.K. ran the analysis. G.K., V.G.R., M.B., M.C. and M.P wrote the initial draft and all the authors contributed in modifying the paper and critically interpreting the results. V.G.R. developed the Shiny app. All authors read and approved the final version for publication.

Competing interests

The authors declare no competing interests.

Ethics

The study is about secondary, aggregate anonymised data so no additional ethical permission is required.

Bibliography

J. M. Aburto, R. Kashyap, J. Schöley, C. Angus, J. Ermisch, M. C. Mills, and J. B. Dowd. Estimating the burden of the covid-19 pandemic on mortality, life expectancy and lifespan inequality in england and wales: a population-level analysis. *J Epidemiol Community Health*, 2021. [p1]

R. J. Acosta and R. A. Irizarry. A flexible statistical framework for estimating excess mortality. *medRxiv*, pages 2020–06, 2020. [p11]

T. Appelhans and F. Detsch. *leafpop: Include Tables, Images and Graphs in Leaflet Pop-Ups*, 2021. URL https://github.com/r-spatial/leafpop. R package version 0.1.0. [p9]

J. Besag. Spatial interactions and the statistical analysis of lattice systems (with discussion). *Journal of the Royal Statistical Society: Series B*, 36:192–236, 1974. [p3]

J. Besag, J. York, and A. Mollié. Bayesian image restoration, with two applications in spatial statistics. *Annals of the Institute of Statistical Mathematics*, 43(1):1–20, 1991. [p3]

M. Blangiardo, M. Cameletti, M. Pirani, G. Corsetti, M. Battaglini, and G. Baio. Estimating weekly excess mortality at sub-national level in Italy during the COVID-19 pandemic. *PLOS ONE*, 15(10): 1–15, 10 2020. doi: 10.1371/journal.pone.0240286. URL https://doi.org/10.1371/journal.pone.0240286. [p1]

B. Davies, B. L. Parkes, J. Bennett, D. Fecht, M. Blangiardo, M. Ezzati, and P. Elliott. Community factors and excess mortality in first wave of the COVID-19 pandemic in England. *Nature Communications*, 12(1):1–9, 2021. [p4]

H. Hersbach, B. Bell, P. Berrisford, S. Hirahara, A. Horányi, J. Muñoz-Sabater, J. Nicolas, C. Peubey, R. Radu, D. Schepers, et al. The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730):1999–2049, 2020. [p5]

N. Islam, V. M. Shkolnikov, R. J. Acosta, I. Klimkin, I. Kawachi, R. A. Irizarry, G. Alicandro, K. Khunti, T. Yates, D. A. Jdanov, et al. Excess deaths associated with COVID-19 pandemic in 2020: age and sex disaggregated time series analysis in 29 high income countries. *BMJ*, 373, 2021. [p1]

J. Kaczorowski and C. Del Grande. Beyond the tip of the iceberg: direct and indirect effects of COVID-19. *The Lancet Digital Health*, 3(4):e205–e206, 2021. [p1]

G. Konstantinoudis, D. Schuhmacher, H. Rue, and B. D. Spycher. Discrete versus continuous domain models for disease mapping. *Spatial and Spatio-temporal Epidemiology*, 32:100319, 2020. ISSN 1877-5845. doi: https://doi.org/10.1016/j.sste.2019.100319. URL https://www.sciencedirect.com/science/article/pii/S1877584519301297. [p3]

G. Konstantinoudis, M. Cameletti, V. Gómez-Rubio, I. L. Gómez, M. Pirani, G. Baio, A. Larrauri, J. Riou, M. Egger, P. Vineis, et al. Regional excess mortality during the 2020 COVID-19 pandemic: a study of five european countries. *medRxiv*, 2021. [p1, 2, 9]

V. Kontis, J. E. Bennett, T. Rashid, R. M. Parks, J. Pearson-Stuttard, M. Guillot, P. Asaria, B. Zhou, M. Battaglini, G. Corsetti, et al. Magnitude, demographics and dynamics of the effect of the first wave of the COVID-19 pandemic on all-cause mortality in 21 industrialized countries. *Nature Medicine*, pages 1–10, 2020. [p1]

V. Kontis, J. Bennett, R. Parks, T. Rashid, J. Pearson-Stuttard, P. Asaria, B. Zhou, M. Guillot, C. Mathers, Y. Khang, M. McKee, and M. Ezzati. Lessons learned and lessons missed: impact of the coronavirus disease 2019 (covid-19) pandemic on all-cause mortality in 40 industrialised countries prior to mass vaccination [version 1; peer review: 2 approved with reservations]. *Wellcome Open Research*, 6(279), 2021. doi: 10.12688/wellcomeopenres.17253.1. [p1]

E. Kontopantelis, M. A. Mamas, J. Deanfield, M. Asaria, and T. Doran. Excess mortality in england and wales during the first wave of the covid-19 pandemic. *J Epidemiol Community Health*, 75(3):213–223, 2021. [p1]

P. Moraga. Small Area Disease Risk Estimation and Visualization Using R. *The R Journal*, 10(1):495–506, 2018. doi: 10.32614/RJ-2018-036. URL https://doi.org/10.32614/RJ-2018-036. [p1]

A. Riebler, S. H. Sørbye, D. Simpson, and H. Rue. An intuitive Bayesian spatial model for disease mapping that accounts for scaling. *Statistical Methods in Medical Research*, 25(4):1145–1165, 2016. [p3]

L. M. Rossen, A. M. Branum, F. B. Ahmad, P. Sutton, and R. N. Anderson. Excess deaths associated with COVID-19, by age and race and ethnicity—United States, january 26–october 3, 2020. *Morbidity and Mortality Weekly Report*, 69(42):1522, 2020. [p1]

H. Rue and L. Held. *Gaussian Markov random fields: theory and applications*. CRC press, 2005. [p3]

H. Rue, S. Martino, and N. 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. [p1, 3]

D. Simpson, H. Rue, A. Riebler, T. G. Martins, S. H. Sørbye, et al. Penalising model component complexity: A principled, practical approach to constructing priors. *Statistical Science*, 32(1):1–28, 2017. [p3]

J. Verbeeck, C. Faes, T. Neyens, N. Hens, G. Verbeke, P. Deboosere, and G. Molenberghs. A linear mixed model to estimate covid-19-induced excess mortality. *Biometrics*, pages 1–9, 2021. doi: https://doi.org/10.1111/biom.13578. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/biom.13578. [p1]

L. A. Waller, B. P. Carlin, H. Xia, and A. E. Gelfand. Hierarchical spatio-temporal mapping of disease rates. *Journal of the American Statistical association*, 92(438):607–617, 1997. [p1]

D. M. Weinberger, J. Chen, T. Cohen, F. W. Crawford, F. Mostashari, D. Olson, V. E. Pitzer, N. G. Reich, M. Russi, L. Simonsen, et al. Estimation of excess deaths associated with the COVID-19 pandemic in the United States, March to May 2020. *JAMA Internal Medicine*, 180(10):1336–1344, 2020. [p1]

Garyfallos Konstantinoudis
MRC Centre for Environment and Health,
Imperial College London,
St Mary's Campus, Praed St,
W2 1NY, London
United Kingdom
g.konstantinoudis@imperial.ac.uk

Virgilio Gómez-Rubio
Departamento de Matemáticas, Escuela Técnica Superior de Ingenieros Industriales-Albacete,
Universidad de Castilla-La Mancha,
Albacete, Spain
virgilio.gomez@uclm.es

Michela Cameletti
Department of Economics, University of Bergamo,
Bergamo, Italy
michela.cameletti@unibg.it

Monica Pirani
MRC Centre for Environment and Health,
Imperial College London,
St Mary's Campus, Praed St,
W2 1NY, London
United Kingdom
monica.pirani@imperial.ac.uk

Gianluca Baio
Department of Statistical Sciences,
University College London,
United Kingdom
g.baio@ucl.ac.uk

Marta Blangiardo
MRC Centre for Environment and Health,
Imperial College London,
St Mary's Campus, Praed St,
W2 1NY, London
United Kingdom
m.blangiardo@imperial.ac.uk