



Moderation effects and elasticities in compositional regression with a total. Application to bayesian spatiotemporal modelling of all-cause mortality from environmental stressors

Germà Coenders^{1,2} · Javier Palarea-Albaladejo³ · Marc Saez^{1,2} · Maria A. Barceló^{1,2}

Received: 5 June 2025 / Accepted: 23 January 2026 / Published online: 28 February 2026
© The Author(s) 2026

Abstract

Compositional regression models with a real-valued response variable can generally be specified as log-contrast models subject to a zero-sum constraint on the model coefficients. This formulation emphasises the relative information conveyed in the composition, while the overall total is regarded irrelevant. In this work, such a setting is extended to account not only for total effects, formally defined in a so-called $\mathcal{T}$ -space, but also for moderation or interaction effects. This is applied in the context of complex spatiotemporal data modelling, through an adaptation of the integrated nested Laplace approximation (INLA) method within a Bayesian estimation framework. Particular emphasis is placed on the interpretation of model coefficients and results, both on the original scale of the response variable and in terms of elasticities. The methodology is demonstrated through a detailed case study investigating the relationship between all-cause mortality and the interaction between extreme temperatures, air pollution composition, and total air pollution in Catalonia, Spain, during the summer of 2022. The results indicate that extreme temperatures are associated with an increased risk of mortality four days after exposure. Additionally, exposure to total air pollution, especially to NO_2 , is linked to elevated mortality risk regardless of temperature. In contrast, particulate matter is associated to increased mortality only when exposure occurs on days of extreme heat.

Keywords Compositional data · Log-contrast model · Integrated nested Laplace approximation (INLA) · Spatiotemporal model · Moderation · T-space · Compositional regression · Extreme temperatures · Air pollutants · All-cause mortality

Mathematics subject classification 62F15 · 62J99 · 62M30 · 62P12

✉ Marc Saez
marc.saez@udg.edu

Germà Coenders
germa.coenders@udg.edu

Javier Palarea-Albaladejo
javier.palarea@udg.edu

Maria A. Barceló
antonia.barcelo@udg.edu

¹ Research Group on Statistics, Econometrics and Health (GRECS), University of Girona, C. Universitat de Girona 10, 17003 Girona, Spain

² Centro de Investigación Biomédica en Red de Epidemiología y Salud Pública, Instituto de Salud Carlos III, Av. Monforte de Lemos 5-7, 28029 Madrid, Spain

³ Dept. of Computer Science, Applied Mathematics and Statistics, University of Girona, C. Maria Aurèlia Capmany 61, 17003 Girona, Spain

1 Introduction

Compositional data arise in various scientific disciplines where multivariate non-negative observations referring to parts of a whole are collected. In practice, they might be originally expressed in any relative unit of measurement (e.g. parts per million, ppm; weight percentages, wt%; milligrammes per litre, mg/L; and so on), although they are often rescaled (*closed*) by dividing each part by their total sum. Thus, they are equivalently re-expressed in proportions or percentages. This is commonly the case when referring to, for instance, nutritional compositions, chemical concentrations, unemployment rates across regions, and the like. Note that, in doing so, the researcher implicitly assumes that the total is irrelevant for the scientific question and, formally, the data are projected onto a simplex as sample space

(Aitchison 1982; Pawlowsky-Glahn and Egozcue 2001). Importantly, the relevant relative information contained in the data remains the same regardless of the scale and, hence, the results of any statistical analysis should be compatible. In air pollution studies, the compositional approach enables taking the full pollution mix and its intrinsic interdependencies into account, rather than focusing on a specific pollutant (AL-Dhurafi et al. 2018; Sánchez-Balseca and Pérez-Foguet 2020). A recent overview of the state of the art in compositional data analysis since Aitchison's early work can be found in Coenders et al. (2023).

It has been thoroughly shown that applying conventional statistical methods directly to compositional data can produce misleading or spurious results [see e.g. Aitchison, 1986, Pawlowsky-Glahn et al. 2015, Filzmoser et al. 2018, 93 as general references]. Instead, the mainstream approach involves mapping them onto the real space through the definition of sensible logratios between the parts, which focuses the data analysis solely on the relative information. By using logratio transformations or coordinates, statistical analysis can, with some caveats, be equivalently conducted through standard methods for real-valued data.

The so-called additive logratio (alr) and orthonormal/isometric logratio (olr/ilr) coordinates have been commonly used in regression analysis with compositional covariates (Aitchison and Bacon-Shone 1984; Coenders and Pawlowsky-Glahn 2020; van den Boogaart et al. 2021). Working with logratios guarantees desirable properties, such as the irrelevance of both the scale and the number of parts forming the composition. It also prevents technical issues with closed compositional data, such as singularity of covariance matrices and perfect multicollinearity in regression analysis. Note that these logratio representations are particular cases of a log-contrast, i.e. a linear combination of the log-transformed parts of the composition where the coefficients are constrained to add up to zero, so that the expression remains scale invariant. As detailed later on, log-contrasts will be used as the basis to formulate our modelling approach.

However, there are situations where the practitioner is actually interested in the potential effect of the total value of the composition in their modelling, so long as this is not constant (i.e. the data are not closed). That is, not only the relative, but also the absolute information is relevant for the scientific question therein. This is clearly the case in air pollution studies, where total pollution levels should not be ignored (Mota Bertran et al. 2022). Combining both relative and absolute information coherently within a statistical model draws on the theory of $\mathcal{T}$ -spaces (Pawlowsky-Glahn et al. 2015). In regression analysis, this allows for the separation of the effects of relative and absolute changes in the composition. It is important to note that simply applying a

logarithm transformation of the compositional variables, as often done in environmental sciences (e.g. when analysing pollutant concentrations), fails to achieve such separation, since log-transformed values inherently mix both types of information (Coenders et al. 2017).

In this work, compositional regression analysis with a total is extended to encompass potential moderation or interaction effects. Briefly, moderation analysis allows exploring how a certain variable influences the strength or direction of the relationship between explanatory and response variables, which in regression implies considering an interaction term between them (Jaccard and Turrissi 2003). With a focus on interpretability, a novel parameterisation is formulated which enables the discussion of results in terms of elasticities. This can, in fact, be readily integrated into wider families of models and accommodate both frequentist and Bayesian inference approaches. In particular, the required zero-sum constraint on the model coefficients to account for compositionality is embedded here for the first time into the integrated nested Laplace approximation (INLA) method for Bayesian inference.

In the following, Sect. 2 presents the compositional regression model with compositional covariates and a total in terms of log-contrasts and some issues when fitting such models with Bayesian inference. Section 3 elaborates on its extension to consider moderation effects and the interpretation of the associated model coefficients. To illustrate the use, in Sect. 4 the method is embedded in Bayesian spatio-temporal modelling, using a novel adaptation of the INLA method to deal with the compositional aspects. This model is applied to investigate the relationship between all-cause mortality and the interaction between extreme temperatures, air pollution composition, and total air pollution in the region of Catalonia, Spain, during the summer of 2022. Finally, Sect. 5 concludes with some final remarks.

2 Compositional linear regression model with real-valued response and compositional covariates

Let us consider a non-closed composition $\mathbf{x}$ comprising D strictly positive parts, that is

$$\mathbf{x} = (x_1, x_2, \dots, x_D) \in \mathbb{R}^{D+}, \text{ with } x_j > 0. \quad (1)$$

The purpose is to model the association between $\mathbf{x}$, in the role of explanatory or predictor variable, and a real-valued variable y in the role of dependent or response variable. The most fundamental form of a linear regression model in this setting is the so-called log-contrast model, as originally formulated in Aitchison and Bacon-Shone (1984).

However, following Müller et al. (2018) and Coenders and Pawlowsky-Glahn (2020), here we specify the log-contrast element of the model in terms of logarithms in base 2 to simplify the interpretation of its coefficients:

$$y = \beta_0 + \beta_1 \log_2(x_1) + \beta_2 \log_2(x_2) + \cdots + \beta_D \log_2(x_D) + \varepsilon, \quad (2)$$

with $\sum_{j=1}^D \beta_j = 0$ and $\varepsilon \sim N(0, \sigma^2)$.

The constraint $\sum_{j=1}^D \beta_j = 0$, apart from guaranteeing scale invariance of the log-contrast as noted above, ensures in practice that not all parts can simultaneously increase in relative terms. The ε corresponds to the normally distributed error term of mean zero and constant variance σ^2 , as commonly included in linear regression models.

Recall that compositionality implies that a part can increase if and only if at least one other part decreases. An increase in parts with positive coefficients β_j at the expense of a decrease in parts with negative coefficients β_j will be associated with a higher expected value of the response y . In more precise terms, due to the use of logarithms in base 2, a coefficient β_j is interpreted as the expected increase in y when the ratio of a part x_j to any other part doubles (Coenders and Pawlowsky-Glahn 2020).

2.1 Model fitting

In an ordinary frequentist estimation setting, the above model (2) can be directly fitted by the constrained least squares method, subject to $\sum_{j=1}^D \beta_j = 0$, or by ordinary least squares after the alr coordinate transform is applied to the D -part composition $\mathbf{x}$ (Aitchison and Bacon-Shone 1984). That is, using e.g. the last part x_D as reference in the denominator of the logratios without loss of generality, the $(D-1)$ -dimensional alr-vector is given by $(x_1/x_D, x_2/x_D, \dots, x_{D-1}/x_D)$, and the model is then specified as

$$y = \beta_0 + \beta_1 \log_2\left(\frac{x_1}{x_D}\right) + \beta_2 \log_2\left(\frac{x_2}{x_D}\right) + \cdots + \beta_{D-1} \log_2\left(\frac{x_{D-1}}{x_D}\right) + \varepsilon, \quad (3)$$

with $\beta_D = -\beta_1 - \beta_2 - \cdots - \beta_{D-1}$ and $\varepsilon \sim N(0, \sigma^2)$.

Note that using the least squares method, it does not matter which part is chosen as denominator for the alr representation. Thus, the same results would be obtained using, say, the part x_1 as reference instead of x_D , leading to the model specification

$$y = \beta_0 + \beta_2 \log_2\left(\frac{x_2}{x_1}\right) + \beta_3 \log_2\left(\frac{x_3}{x_1}\right) + \cdots + \beta_D \log_2\left(\frac{x_D}{x_1}\right) + \varepsilon, \quad (4)$$

with $\beta_1 = -\beta_2 - \beta_3 - \cdots - \beta_D$ and $\varepsilon \sim N(0, \sigma^2)$.

However, it has been shown that this will not necessarily hold for other estimation methods. This is the case of

Bayesian inference using INLA (Gómez-Rubio 2020; Sánchez-Balseca and Pérez-Foguet 2021) or Markov chain Monte Carlo (MCMC) estimation (Le et al. 2025; Sánchez-Balseca and Pérez-Foguet 2020). Here, the prior distribution interferes with the logratio representation, making the model specifications (3) and (4) no longer equivalent as it should be.

Within a Bayesian framework, the fitting of model (2) requires incorporating the compositional constraint $\sum_{j=1}^D \beta_j = 0$ into the specification of the prior distributions. Following Zhang et al. (2025, 2024a) and Zhang et al. (2024b), this can be achieved using MCMC methods by implementing a *soft* constraint of the form

$$\sum_{j=1}^D \beta_j \sim N(0, 0.001D). \quad (5)$$

Or, alternatively, it can be done by assuming the *hard* constraint $\sum_{j=1}^D \beta_j = 0$ as in Scott et al. (2023). The latter implies that the multivariate priors for the vector of model coefficients $(\beta_1, \dots, \beta_D)$ will have a singular covariance matrix.

Importantly, including the constraint, either in its soft or hard variant, ensures the reproducibility of results and, simultaneously, avoids having to choose an arbitrary reference part to obtain an alr coordinate representation of the compositional predictor. This is also compatible with the principle of permutation invariance in compositional data analysis (Aitchison 1986), by which the statistical results should be invariant to permutations of the parts of the composition. Although it is theoretically possible, none of these variants of the constraint had been so far implemented in the INLA method. Note that the constraint is linked to the model specification and not to any particular Bayesian computation method.

This is the approach we adhere to in the current work, particularly through the novel implementation of the soft constraint in the INLA method. This is done on the R system for statistical computing (R Core Team 2025) through a new function called `A.local` in the package R-INLA (R INLA project 2025). The `A.local` function allows users to impose local constraints on the linear predictor. Previously, if a linear predictor contained more than one element of the same model component, it required a global redefinition through complicated design matrices, which also involved irrelevant parts of the model. With `A.local`, this is avoided in most cases by efficiently applying constraints directly to specific subsets of the model predictors.

2.2 Considering the total effect

As previously discussed, there are cases where, in addition to the potential effect of relative changes among parts of the composition, their actual absolute values may also be contributing to the variation in the response variable (Coenders et al. 2017). To account for this, a term capturing the total magnitude of the composition can be included into the regression model. A coherent way to do this, while respecting the geometric properties of the sample space of compositions, is provided by the theory of $\mathcal{T}$ -spaces (Pawlowsky-Glahn et al. 2015). A few different formulations have been discussed to summarise such total information, with the so-called multiplicative total being preferred according to Pawlowsky-Glahn et al. (2015). This is defined as

$$\sqrt{D} \log(\sqrt[D]{x_1 x_2 \cdots x_D}) = \frac{1}{\sqrt{D}} (\log(x_1) + \log(x_2) + \cdots + \log(x_D)). \quad (6)$$

However, for estimation and interpretation purposes within a linear modelling framework using INLA, Mota Bertran et al. (2024) simplified it to

$$t = D \log_2(\sqrt[D]{x_1 x_2 \cdots x_D}) = \log_2(x_1 x_2 \cdots x_D) = (\log_2(x_1) + \log_2(x_2) + \cdots + \log_2(x_D)). \quad (7)$$

Hence, a log-contrast model (2) extended to include the information about the total (9) is written as

$$y = \beta_0 + \beta_1 \log_2(x_1) + \beta_2 \log_2(x_2) + \cdots + \beta_D \log_2(x_D) + \beta^{(t)} t + \varepsilon, \quad (8)$$

with $\sum_{j=1}^D \beta_j = 0$ and $\varepsilon \sim N(0, \sigma^2)$.

The new coefficient $\beta^{(t)}$ can be interpreted as $1/D$ times the effect on the expected value of the response y when the absolute values of all parts are simultaneously doubled, while preserving the relative structure encoded in their logratios. Moreover, a coefficient β_j , for $j = 1, \dots, D$, is interpreted as the expected increase in the response y when the ratio of the part x_j to any other part doubles, while the total of the composition is kept constant. This means that the parts $x_1, x_2, \dots, x_{j-1}, x_{j+1}, \dots, x_D$ decrease by a common factor to compensate for the increase in x_j , thereby keeping the total t , as defined in (7), constant (Coenders et al. 2017).

Following Mota Bertran et al. (2024), and with the aim of making the decomposition between compositional (zero-sum) and total (constant) effects more evident, model (8) can be reformulated as

$$y = \beta_0 + (\beta_1 + \beta^{(t)}) \log_2(x_1) + (\beta_2 + \beta^{(t)}) \log_2(x_2) + \cdots + (\beta_D + \beta^{(t)}) \log_2(x_D) + \varepsilon, \quad (9)$$

with $\sum_{j=1}^D \beta_j = 0$ and $\varepsilon \sim N(0, \sigma^2)$.

In this case, a coefficient $(\beta_j + \beta^{(t)})$ is interpreted as the effect of multiplying x_j by 2 while keeping the absolute values of the parts $x_1, x_2, \dots, x_{j-1}, x_{j+1}, \dots, x_D$ constant. Furthermore, if y is log-transformed, the model becomes

$$\log_2(y) = \beta_0 + (\beta_1 + \beta^{(t)}) \log_2(x_1) + (\beta_2 + \beta^{(t)}) \log_2(x_2) + \cdots + (\beta_D + \beta^{(t)}) \log_2(x_D) + \varepsilon, \quad (10)$$

with $\sum_{j=1}^D \beta_j = 0$ and $\varepsilon \sim N(0, \sigma^2)$,

and the coefficients are readily interpreted as elasticities.

Alternative formulations of compositional regression models have been proposed to enable interpretation in terms of elasticities (Dargel and Thomas-Agnan 2024; Morais and Thomas-Agnan 2021). However, these approaches involve complex functions of the model parameters. Contrarily, in our proposed model (10), the sum of coefficients $(\beta_j + \beta^{(t)})$ is directly interpretable as an elasticity. Thus, if a part x_j has a small relative increase (say by 1%), while keeping the remaining parts unaltered, the response variable y will increase by $(\beta_j + \beta^{(t)})\%$.

This elasticity can be further decomposed into its compositional and total contributions. On the one hand, if the ratio of a part x_j to any other part increases by 1%, while keeping the total constant, then y will increase by $\beta_j\%$. On the other hand, if all parts increase simultaneously by $1/D\%$, then y will increase by $\beta^{(t)}\%$.

Finally, note that the base of the logarithm used in (10) does not affect the interpretation, as long as it is applied consistently to both the response and explanatory sides of the model.

3 Adding a moderation effect

3.1 Formulation in the original scale of the response variable

Building on model (8), we introduce the extension to incorporate a moderation or interaction effect. This involves the composition, the total, and a new moderating variable z . The variable z may be defined either as a dummy categorical variable, coded $\{0, 1\}$, or as a numeric variable. In the latter case, we assume z is mean-centred to aid interpretation. Similarly, $\log_2(x_1), \dots, \log_2(x_D)$ and t are also centred around their respective means. The full model can thus be written as

$$y = \beta_0 + \beta_1 \log_2(x_1) + \beta_2 \log_2(x_2) + \cdots + \beta_D \log_2(x_D) + \beta^{(z)}z + \beta_1^{(i)}z \log_2(x_1) + \beta_2^{(i)}z \log_2(x_2) + \cdots + \beta_D^{(i)}z \log_2(x_D) + \beta^{(t)}t + \beta^{(t,i)}tz + \varepsilon, \quad (11)$$

with $\sum_{j=1}^D \beta_j = 0$, $\sum_{j=1}^D \beta_j^{(i)} = 0$ and $\varepsilon \sim N(0, \sigma^2)$.

The following points summarise the interpretation of the model coefficients:

- The coefficients β_j , $\beta^{(t)}$, and $\beta^{(z)}$ represent the main effects, each corresponding to the scenario where the other variables are held at zero. Specifically, the coefficient β_j quantifies the expected change in y when the ratio of x_j to any other part doubles, assuming that the total remains constant and $z = 0$ (i.e. the variable z is held at its mean value if numeric or at the reference level 0 if binary). The coefficient $\beta^{(t)}$ captures the effect (scaled by $1/D$) of simultaneously doubling the absolute values of all parts of the composition, while $z = 0$ is assumed. Thanks to the mean-centring, the coefficient $\beta^{(z)}$ represents the expected change in y associated with a one-unit increase in z (or a change from level 0 to 1 in the binary case), when the terms $\log_2(x_1), \dots, \log_2(x_D)$ are held at zero (i.e. at their respective mean values).
- The coefficients $\beta^{(i)}$ account for the moderating or interaction effects, i.e. the change in the compositional and total effects when $z = 1$ (meaning that the variable z is one unit above its mean value if numeric or at level 1 if binary). In other words, when $z = 0$, the log-contrast coefficients are $\beta_1, \beta_2, \dots, \beta_D$; whereas when $z = 1$, the log-contrast coefficients become $(\beta_1 + \beta_1^{(i)}), (\beta_2 + \beta_2^{(i)}), \dots, (\beta_D + \beta_D^{(i)})$.
- The coefficient $\beta^{(t,i)}$ also refers to an interaction effect, in this case indicating changes in the total effect when z either raises one unit over the mean (numeric case) or shifts from level 0 to 1 (binary case). That is, when $z = 0$ the total effect is just $\beta^{(t)}$, but for $z = 1$ it includes the term of the interaction with z and is given by $(\beta^{(t)} + \beta^{(t,i)})$.

Note that these interaction effects are often of substantial practical interest. For instance, the impact of a relative increase in a specific part x_j on the response y may vary depending on the level of z . This is governed by the sign of $\beta_j^{(i)}$: a positive value indicates a gain in influence with increasing values of z , while a negative value indicates a loss. Similarly, the impact of an absolute increase in all parts of the composition on y with higher values of z will be conditioned by the sign of $\beta^{(t,i)}$, being strengthened or weakened in the case of positive or negative signs respectively.

3.2 Formulation in terms of elasticities

As discussed in Sect. 2, expressing the response variable y in logarithmic scale in the model allows a moderation effect to be interpreted in terms of changes in elasticities. Thus, specifying the model in the form

$$\log_2(y) = \beta_0 + \beta_1 \log_2(x_1) + \beta_2 \log_2(x_2) + \cdots + \beta_D \log_2(x_D) + \beta^{(z)}z + \beta_1^{(i)}z \log_2(x_1) + \beta_2^{(i)}z \log_2(x_2) + \cdots + \beta_D^{(i)}z \log_2(x_D) + \beta^{(t)}t + \beta^{(t,i)}tz + \varepsilon, \quad (12)$$

with $\sum_{j=1}^D \beta_j = 0$, $\sum_{j=1}^D \beta_j^{(i)} = 0$ and $\varepsilon \sim N(0, \sigma^2)$,

leads to the following interpretations:

- Regarding main effects, the sum $(\beta_j + \beta^{(t)})$ is directly interpretable as an elasticity. Thus, a small relative increase (e.g. by 1%) in x_j , while keeping the remaining parts unaltered, leads to an increase of $(\beta_j + \beta^{(t)})\%$ in the response y . Note that this is so only when $z = 0$ (i.e. with z at its mean value if numeric or at the reference level 0 if binary). Furthermore, these elasticities can be decomposed into total and compositional effects as follows. When all parts of the composition increase simultaneously by $1/D\%$, then the response y increases by $\beta^{(t)}\%$ while holding $z = 0$. When the ratio of x_j to any other part increases by 1%, while keeping the total constant, the response y increases by $\beta_j\%$, with $z = 0$.
- The coefficient $\beta^{(z)}$ represents a semi-elasticity. When z increases by one unit (above its mean value in the numerical case or from level 0 to 1 in the binary case), the response y increases by a factor of $2^{\beta^{(z)}}$, for $\log_2(x_1), \dots, \log_2(x_D)$ held at value zero (i.e. at the mean composition).
- As to the moderation or interaction effects, the coefficients $\beta_j^{(i)}$ and $\beta^{(t,i)}$ represent the change in the compositional and total effects for the case of $z = 1$ (i.e. when z is one unit above its mean value if numeric or at level 1 if binary). If a certain part x_j increases by 1%, while the other parts remain constant, the response y increases by $(\beta_j + \beta^{(t)} + \beta_j^{(i)} + \beta^{(t,i)})\%$. As noted before, these elasticities can be further decomposed into total and compositional effects. Thus, if all the parts increase simultaneously by $1/D\%$, then the response y increases by $(\beta^{(t)} + \beta^{(t,i)})\%$. If the ratio of x_j to any other part increases by 1%, the response y increases by $(\beta_j + \beta_j^{(i)})\%$.

We emphasise once again that the base of the logarithm does not affect the interpretation of the compositional and total effects, as long as the base is applied consistently to both the response and explanatory sides of the

model. Only the interpretation of the main effect of z will change, referring to an increase in y by a factor $b^{\beta(z)}$, where b stands for the base of the logarithm.

Finally, note that all the above, elaborated for the ordinary linear regression model, can be readily extended to generalised linear models (GLMs) with non-normal response variables (Coenders et al. 2017). The case study developed in the following section provides an illustration of such extension using a zero-inflated negative binomial regression model.

4 Bayesian spatiotemporal modelling of all-cause mortality from environmental stressors

4.1 Background

In environmental health studies, the compositional data methodology has been used to identify patterns of pollution which can be associated with health risks (Sánchez-Balseca and Pérez-Foguet 2020, 2022). In this context, jointly considering relative and absolute information, through their representation in a $\mathcal{T}$ -space, makes it possible to study the relative importance of pollutant concentrations along with the impact of the overall pollution level (Mota Bertran et al. 2022, 2024). Furthermore, it is well known that exposure to certain air pollutants can modify the impact of extreme heat on mortality. In particular, it has been found that ozone exposure influences the effects of heatwaves on cardiovascular, respiratory, and overall mortality (Alari et al. 2023; Analitis et al. 2014; Dear et al. 2005; Du et al. 2024; Qi et al. 2023). Exposure to PM_{10} and $PM_{2.5}$ have also been found to play some moderating role on mortality (Analitis et al. 2014; Xu et al. 2023).

When working within a spatiotemporal framework, the estimation of a second-order stationary Gaussian field (GF) is known to often suffer from the so-called big n problem. This refers to the high computational burden associated to model fitting with increasing size and variability along both space and time dimensions (Lindgren and Rue 2015). Computational burden can be alleviated by representing a GF as a Gaussian Markov random field (GMRF) (Rue et al. 2009). This is defined by a precision matrix with a sparse structure, allowing inference to be performed in a computationally efficient way. The GF and GMRF are linked through the stochastic partial differential equations approach (Lindgren et al. 2011). This makes it possible to find a GMRF with local neighbourhood and sparse precision matrix (instead of the spatiotemporal covariance function and dense covariance matrix of a GF) which best represents a Matérn field (Cameletti et al. 2013; Lindgren et al. 2011). This often

renders a frequentist approach impractical, favouring the use of Bayesian alternatives instead. In this setting, the INLA approach, as used in this work, has been shown to be much faster and computationally efficient than traditional Bayesian methods based on MCMC, producing accurate approximations to posterior distributions, even for very complex models (Lindgren and Rue 2015). In this vein, the model specification detailed in Saez and Barceló (2022) offers a fast yet sufficiently accurate solution. Such specification is used as reference for the model presented here. Our application concerns an ecological small-area study conducted at the level of the 379 Basic Health Areas (ABS) in which the territory of Catalonia is divided. The dataset comprises daily observations from June to September 2022. In the summer of 2022, Spain as a whole experienced an unprecedented series of 41 consecutive days of extreme temperatures due to three successive heatwaves. This broke previous records in 2015 and 2004 by 12 and 21 days respectively. The heatwave in July 2022 hit 44 out of the 52 Spanish provinces, surpassing any previous event in terms of number of provinces affected and registering a temperature anomaly of 3.7 °C above the monthly average. In addition, the first heatwave (12–18 June) was the second-earliest on record, whereas the second heatwave (30 July to 15 August) was the most extensive and intense observed to date (Agencia Española de Meteorología -AEMET- 2024). In Catalonia specifically, the summer of 2022 was the hottest on record, with over two weeks of sustained above-average temperatures in July, ranking among the most persistent heatwaves experienced so far in the region (METEOCAT 2022).

4.2 Model formulation

The use of the methodology introduced above is adapted and illustrated here in the context of Bayesian modelling of environmental spatiotemporal data using the INLA method for parameter estimation. The purpose is to investigate the association between all-cause mortality as response variable and extreme temperatures, air pollution and the interaction between them as explanatory variables.

The mortality data used were obtained from the Spanish National Statistics Institute (INE) and consist of daily all-cause mortality counts for the summer months (June to September) of 2022. They are aggregated at the ecological level for each ABS and are available for the total population.

The meteorological data were provided by the XEMA network (Network of Automatic Meteorological Stations; open data available at https://analisi.transparenciacatalunya.cat/ca/Medi-Ambient/Dades-meteorol-giques-de-la-XEMA/nzvn-apec/about_data), which comprises 189 automatic weather stations distributed across Catalonia. This excluded 12 stations located at altitudes of 1500 metres

or higher (Departament de Territori, Habitatge i Transició Ecològica 2025a). The key variable here is the maximum daily temperature.

Lastly, the pollution data are daily averages of hourly levels of the concentrations of five pollutants: particulate matter (PM₁₀), nitrogen dioxide (NO₂), ozone (O₃), carbon monoxide (CO) and sulphur dioxide (SO₂). These were measured for the 95 automatic monitoring stations of the XVPCA network (Catalan Network for Pollution Control and Prevention; open data available at https://analisi.transparenciacatalunya.cat/Medi-Ambient/Qualitat-de-l-aire-als-punts-de-mesurament-autom-t/tasf-thgu/about_data) (Departament de Territori, Habitatge i Transició Ecològica 2025b). Note that for ABSs lacking a nearby weather or pollution monitoring station in their territory, the corresponding meteorological or pollutant concentrations were estimated using a spatiotemporal predictive model following the approach in (Mota Bertran et al. 2022; Saez and Barceló 2022).

Building on the set-up of model (12), we adapted it to account for the fact that the response variable, all-cause mortality, arises from a count process and that random effects are included. Thus, model selection was confined to the family of generalised linear mixed models (GLMMs) within the Bayesian framework provided by the R-INLA package (Rue et al. 2009, 2017) used in experimental mode (Van Niekerk et al. 2023). These included Poisson and negative binomial distribution specifications for the response variable, with the latter accounting for data overdispersion. Moreover, since a number of ABSs did not report any deaths over several days, we also considered their zero-inflated variants. Based on the Watanabe-Akaike Information Criterion (WAIC) (Watanabe 2010), the best fit to the data (lowest WAIC) was reached with a zero-inflated negative binomial GLMM formulation, using a natural logarithm link function.

Schematically, the fitted model has the following structure. Being y_{id} the all-cause mortality count in the i -th ABS on day d , f_{NB} the probability function of the associated negative binomial model NB(μ_{id} , θ), with μ_{id} and θ denoting the corresponding mean and dispersion parameters respectively, and π_{id} the probability of an excess zero; the data generation process is given by

$$y_{id} \sim \begin{cases} 0 & \text{with probability } \pi_{id} + (1 - \pi_{id}) \cdot f_{NB}(0 \mid \mu_{id}, \theta) \\ \text{NB}(\mu_{id}, \theta) & \text{with probability } (1 - \pi_{id}) \cdot f_{NB}(y_{id} \mid y_{id} > 0; \mu_{id}, \theta). \end{cases} \quad (13)$$

The relationship with the explanatory terms is established through the (natural) logarithm link function as the linear predictor

$$\begin{aligned} \log(\mu_{id}) = & \beta_0 + \beta_1 \log(x_{1id-4}) + \beta_2 \log(x_{2id-4}) \\ & + \dots + \beta_5 \log(x_{5id-4}) + \beta^{(z)} z_{id-4} + \\ & + \beta_1^{(i)} z_{id-4} \log(x_{1id-4}) + \beta_2^{(i)} z_{id-4} \log(x_{2id-4}) \\ & + \dots + \beta_5^{(i)} z_{id-4} \log(x_{5id-4}) + \\ & + \beta^{(t)} t_{id-4} + \beta^{(t,i)} t_{id-4} z_{id-4} + \beta^{(p)} \log(p_i) \end{aligned} \quad (14)$$

$$+ \eta_i + S(ABS_i) + \tau s_{id},$$

with $\sum_{j=1}^5 \beta_j = 0$ and $\sum_{j=1}^5 \beta_j^{(i)} = 0$.

Regarding the fixed effects, the variables x_{1id-4} , x_{2id-4} , x_{3id-4} , x_{4id-4} and x_{5id-4} refer, respectively, to the daily average concentrations of the following air pollutants in the i -th ABS on day $d - 4$: particulate matter (PM₁₀), nitrogen dioxide (NO₂), ozone (O₃), carbon monoxide (CO) and sulphur dioxide (SO₂). These compositional variables were log-transformed using the natural logarithm, which allows the model parameters to be interpreted as elasticities, as discussed in previous sections. Moreover, the variable z_{id-4} is a dummy variable indicating the daily extreme maximum temperature. It takes value 1, when the maximum temperature (or its estimated value if it was originally absent) exceeds the trigger threshold in the i -th ABS on day $d - 4$, and value 0 otherwise. Such a trigger temperature is defined as the 95-th percentile of the predicted maximum temperature in the ABS during the Summer months of 2022 (June to September). Out of the 122 days under study, 15.7% days (average across all ABS) exhibited extreme temperatures. Note that all these variables lagged by 4 days relative to the response variable. Following the same strategy as Barceló and Saez (2025), the selection of lag structures was based on a systematic comparison of alternative specifications. Namely, we estimated models including all combinations of the contemporaneous value of extreme maximum temperature and its lags up to 7 days, together with the contemporaneous value and lags of air pollutants up to 7 days. The selected specification corresponds to the combination that yielded the highest WAIC, i.e. model (14), which considers a lag of 4 days for extreme maximum temperature. Finally, the p_i is an offset term accounting for the total population in the i -th ABS.

As to the random effects, the η_i term represents an unstructured random effect accounting for heterogeneity in the ABS, e.g. due to unobserved confounders specific to the ABS and assumed invariant in time. The $S()$ term represents a structured random effect accounting for the spatial dependence among neighbouring ABSs. It is assumed to be normally distributed with zero mean and a Matérn covariance function, where the covariance only depends on the distance between the ABSs involved (Lindgren et al. 2011; Saez and Barceló 2022). The τs_{id} term contains structured random effects accounting for seasonality over time. Within this framework, the random effects are used as smoothers

to model non-linear dependency on covariates in the linear predictor (14).

4.3 Results

The model was implemented and fitted using the R-INLA package. As noted above, the soft constraint (5) was applied through the novel `A.local` function. To estimate the model parameters, we relied on a class of prior distributions that penalises complexity, called PC priors (Simpson et al. 2017). A reproducible R script implementing the analysis can be found in the accompanying Appendix A.

Table 1 summarises the model coefficient estimates in terms of posterior means and posterior probabilities of being positive or negative. Note that we considered those effects with an associated posterior probability higher than 90% for interpretation.

In the following, we provide detailed practical interpretations of these results:

- The coefficient β_2 has a posterior probability of being positive equal to 0.914. Then, if the ratio of the concentration of NO_2 to any other pollutant increases by 1%, while keeping the total pollution constant, the expected number of deaths four days later will increase by 0.010%. Note that this may equally apply to days of extremely high temperature or days of normal temperature, given that the corresponding interaction coefficient $\beta_2^{(i)}$ has a low posterior probability.

Table 1 Results from the INLA estimation of the specified zero-inflated negative binomial generalised mixed model

	Mean ¹	Prob. ²
β_1	− 0.002	0.623
β_2	0.010	0.914
β_3	− 0.001	0.558
β_4	− 0.006	0.856
β_5	− 0.001	0.561
$\beta^{(t)}$	0.006	0.983
$\beta^{(z)}$	0.071	> 0.99
$\beta_1^{(i)}$	0.025	0.937
$\beta_2^{(i)}$	0.005	0.613
$\beta_3^{(i)}$	− 0.011	0.766
$\beta_4^{(i)}$	0.003	0.580
$\beta_5^{(i)}$	− 0.022	0.922
$\beta^{(t,i)}$	0.002	0.629

¹Posterior means of the relevant model coefficients

²Posterior probabilities of being positive or negative

- The coefficient $\beta^{(t)}$ has a posterior probability equal to 0.983 of being positive. Then, if the total pollution increases by $1/D = 1/5\%$, while the relative composition remains unaltered, the expected number of deaths four days later will increase by 0.006%. If we scale the results for an easier interpretation, when the total pollution increases by 1%, the expected number of deaths four days later will increase by $0.006 \times 5 = 0.030\%$. Note that this may equally be applied for days of extremely high temperature and days of normal temperature, since the corresponding interaction coefficient $\beta^{(t,i)}$ has a low posterior probability of 0.629.
- The coefficient $\beta^{(z)}$ is estimated to have a posterior probability greater than 0.999 of being positive. This indicates that the expected number of deaths will increase by 7.4% ($e^{0.071} = 1.074$) four days after the occurrence of an extreme temperature event when pollutant concentrations are on their average.
- The coefficient $\beta_1^{(i)}$ has a posterior probability of being positive equal to 0.937. Then, if the ratio of the concentration of PM_{10} to any other pollutant increases by 1%, while keeping the total pollution constant, the expected number of deaths will increase by 0.025%. Note that this may only hold during days of extremely high temperature, since the corresponding main effect coefficient β_1 has a low posterior probability of 0.623. Moreover, the negative sign and high absolute value of $\beta_5^{(i)}$ (posterior probability of being negative equal to 0.922) indicates that increased mortality is expected mainly when PM_{10} is traded for SO_2 .

The potential existence of temporal autocorrelation in the model residuals was evaluated using within-ABS autocorrelation functions (ACF). No significant autocorrelation was observed at any lag, with all estimated correlations lying within the 95% limits (Fig. 1). This indicates that the temporal components of the model sufficiently captured the underlying temporal dynamics.

5 Final remarks

Bayesian compositional regression modelling had already been successfully addressed through INLA inference for the case of compositions in a response variable role and without considering the total (Adin et al. 2024; Figueira et al. 2025; Martínez-Minaya et al. 2023; Martínez-Minaya and Rue 2024). The novelty of our approach lies in handling the case of a composition in an explanatory role, for which only the less computationally efficient MCMC estimation method had been used so far, and also in considering the information provided by the total. This has implied the




Fig. 1 Autocorrelation function (ACF) of the model residuals within each Basic Health Area (ABS)

novel implementation of the soft constraint used in MCMC into the INLA framework. We have also extended the modelling to include moderation effects, using a parametrisation which enhances interpretability. Thus, when the response variable is log-transformed, or a logarithmic link function is used in GLM(M)s, such interpretation can be made in terms of elasticities. Notably, the proposed approach does not rely on any particular logratio coordinates. This, hence, avoids discussions about the convenience of using one or another representation or the lack of equivalence between them.

Two key advantages of our proposal in relation to more conventional approaches used in environmental epidemiology are worth highlighting. Firstly, if the entire composition of pollution concentrations, commonly expressed as percentages adding up to 100 or similar, is included in an ordinary model, it will suffer from perfect multicollinearity and, hence, routine parameter estimation would not be feasible. Secondly, ordinary modelling does not allow estimating the effect of the total on the response variable, since combining the composition, regardless of whether it is closed or not, with the total in the same model leads again to perfect collinearity. Moreover, our method allows us to interpret the estimated coefficients associated to individual pollutants as exchanges. If total pollution is included as a covariate in the model, then keeping it constant makes the interpretation of any change in the concentration of a pollutant to be balanced by an equivalent and opposite change in at least one other pollutant. This generally makes sense in air pollution studies, as certain known geographical and climatological

factors tend to increase some pollutants at the expense of reducing others [see 12, Mota Bertran et al. 2022, and references therein].

The method has been illustrated through an investigation of the moderation effect of air pollution and extreme temperatures on all-cause mortality. The proposed modelling brings in simple interpretations based on elasticities. Such elasticities refer to percentage increases in death counts in response to percentage increases in overall pollution or to percentage increases in the ratio of concentrations of one pollutant (NO_2) to the rest. Additionally, these elasticities may be modified by the interacting variable. Some pollutant(s) (PM_{10}) may affect mortality only, or mainly, when temperatures are extreme.

Although the main objective of this study was not to investigate the biological mechanisms underlying the observed associations, it is nevertheless important to briefly discuss the plausibility of the selected lag structure. The effects of extreme maximum temperatures on health are widely recognised to be non-contemporaneous, as adverse outcomes often arise after sustained exposure rather than immediately. In line with the definition adopted by the Spanish Ministry of Health (Ministerio de Sanidad. Gobierno de España 2025), a heatwave is identified when extreme maximum temperatures persist for three or more consecutive days across a substantial proportion of monitoring stations. Within this framework, it is biologically plausible that the health impacts of extreme heat only become apparent after several days of accumulated thermal stress, making a lag of 4 days a reasonable and interpretable choice.

Lastly, note that the proposed approach remains valid and applicable even when the total is not considered. In such cases, the t and tz terms can simply be omitted from the model equation. While total levels are often of primary interest in contexts such as air, soil, and water pollution, this is not always the case in other fields. For example, in physical activity and time-use epidemiology, the total time spent in various movement behaviours each day is fixed at 24 h (or 1440 min) for all individuals, rendering the total inherently irrelevant (Dumuid et al. 2020; Kuzik et al. 2025). Moreover, the approach presented here includes only moderation effects. Related topics such as non-linear (e.g. quadratic) effects (Aitchison and Bacon-Shone 1984) or mediation effects (Sohn and Li 2019) are possible avenues for further research, particularly within a Bayesian framework.

Appendix A Reproducible R script implementing the model fitting pipeline

```
# Load the data, the necessary packages
and options
```

```

library(INLA)
inla.setOption(inla.mode="experimental")
library(haven)
load("file.RData")
attach(file)
options(contrasts=c("contr.poly", "contr.treatment"))
# Define the mesh for the spatial covariance structure
mesh=inla.mesh.2d(cbind(UTM_x/1000, UTM_y/1000),
max.edge=c(15,40),
offset=c(15,40), cutoff=0.5)
spde=inla.spde2.pcmatern(mesh,
constr=TRUE,
prior.range=c(0.01,0.01), prior.sigma=c(100,0.1))
field=meshloc
# Soft constraint,
# main effects and centering
lag_total_4=lag_PM10_4*lag_NO2_4*lag_O3_4*lag_CO_4*lag_SO2_4
log_lag_total_4=log(lag_total_4)-mean(log(lag_total_4),na.rm=T)
X=as.matrix(cbind(
log(lag_PM10_4)-mean(log(lag_PM10_4),na.rm=T),
log(lag_NO2_4)-mean(log(lag_NO2_4),na.rm=T),
log(lag_O3_4)-mean(log(lag_O3_4),na.rm=T),
log(lag_CO_4)-mean(log(lag_CO_4),na.rm=T),
log(lag_SO2_4)-mean(log(lag_SO2_4),na.rm=T)))
# Interaction and centering
X_interaction=as.matrix(cbind(
(log(lag_PM10_4)-mean(log(lag_PM10_4),na.rm=T))*lag_extreme_heat_4,
(log(lag_NO2_4)-mean(log(lag_NO2_4),na.rm=T))*lag_extreme_heat_4,
(log(lag_O3_4)-mean(log(lag_O3_4),na.rm=T))*lag_extreme_heat_4,
(log(lag_CO_4)-mean(log(lag_CO_4),na.rm=T))*lag_extreme_heat_4,
(log(lag_SO2_4)-mean(log(lag_SO2_4),na.rm=T))*lag_extreme_heat_4))
idx=rep(NA, dim(X)[1])
idx_interaction=idx
# constr=TRUE sets the sum of coefficients in a submatrix
# of predictors to zero.

```

```

formula=dead_count ~ 1 +
f(idx, model = "iid",
hyper=list(prec=list(prior="gaussian",param=c(0,10000))),
constr=TRUE,
A.local=X,
values=1:dim(X)[2]) +
f(idx_interaction, model="iid",
hyper=list(prec=list(prior="gaussian",param=c(0,10000))),
constr=TRUE,
A.local=X_interaction,
values=1:dim(X_interaction)[2]) +
log_lag_total_4*lag_extreme_heat_4 +
f(field, model=spde) +
f(month, model="rw1", scale.model=T,
hyper=list(theta = list(prior="pc.prec", param=c(0.5,0.01))))
+ offset(log(Poblacion))
ptm<- proc.time()
result=inla(formula, data=file, family="zeroinflatednbinomial1",
control.compute=list(dic=T,waic=T),
control.predictor=list(link=1,compute=TRUE),
control.fixed=list(expand.factor.strategy='inla', prec=0.01,
prec.intercept = 0.01))
proc.time() - ptm
summary(result)
# Posterior probabilities
matriz=matrix(NA,
length(names(result$marginals.fixed)),1)
for(i in 1:length(names(result$marginals.fixed))){
matriz[i,]=ifelse(result$summary.fixed[i,1]>0,1-
inla.pmarginal(0,result$marginals.fixed[[i]]),
inla.pmarginal(0,result$marginals.fixed[[i]]))
prob.coef.dif.zero=matriz[,1]
prob.coef.dif.zero
# Goodness of fit
resultdic
resultwaic

```

Acknowledgements The authors wish to thank Håvard Rue for his advice in implementing and programming the new feature `A.local` that allowed the soft constraint to be realized in R INLA. We are also grateful to the anonymous reviewers for their constructive comments on an earlier version of this manuscript, which have notably improved it. This study was carried out within the ‘Atlas of Social and Environmental Determinants of Health’ subprogram of CIBER of

Epidemiology and Public Health (CIBERESP). The usual disclaimer applies.

Author contributions Conceptualization GC, MS; Data curation MS, MAB; Formal analysis MS, MAB; Funding acquisition MAB; Methodology GC, JPA, MS; Project administration MAB; Supervision MS, JPA; Writing – original draft all authors; Writing – review & editing all authors.

Funding Open Access funding provided thanks to the CRUE-CSIC agreement with Springer Nature. This research was supported by the Spanish Ministry of Science and Innovation and ERDF-a way of making Europe [Grant Number PID2021-123833OB-I00], the Spanish Ministry of Health [Grant Number CIBERCB06/02/1002], the Department of Research and Universities of Generalitat de Catalunya [Grant Numbers 2021SGR01197 and 2023-CLIMA-00037] and AGAUR and the Department of Climate Action, Food and Rural Agenda of Generalitat de Catalunya [Grant Number 2023-CLIMA-00037]. The funders had no role in the research process from study design to submission.

Data Availability In accordance with Article 2 of Regulation 223/2009 of the Council and the European Parliament on European Statistics; Articles 13, 17.3, and 17.4 of the Spanish Law on the Public Statistical Function; Article 25.1 of Spanish Organic Law 3/2018, of December 5, on the Protection of Personal Data and Guarantee of Digital Rights; and Regulation 2016/679 of the European Union, the mortality databases of the National Institute of Statistics (INE) are subject to Statistical Confidentiality; therefore, there are restrictions on their transfer to third parties, and data are not publicly available. The datasets (properly anonymized) used and analyzed during the current study are available from the corresponding author upon reasonable request. For the rest of variables we used freely available open data. • Meteorological variables: <https://analisi.transparenciacatalunya.cat/ca/Medi-Ambient/Dades-meteorol-giques-de-la-XEMA/nzvn-apee/about-data> • Air pollutants: <https://analisi.transparenciacatalunya.cat/Medi-Ambient/Qualitat-de-l-aire-als-punts-de-mesurament-autom-t/tasf-thgu/about-data> • Population data. Continuous Register Statistics: <https://www.ine.es/dyngs/INEbase/en/operacion.htm?c=EstadisticaC&cid=1254736177012&menu=resultados&secc=1254736195461&idp=1254734710990#tabs-1254736195557>

Declarations

Conflict of interest The authors declare no Conflict of interest.

Declaration of generative AI and AI-assisted technologies During the writing of the article the authors have not used any type of AI and AI-assisted technologies.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Aitchison J (1982) The statistical analysis of compositional data (with discussion). *J R Stat Soc Ser B Stat Methodol* 44(2):139–177
- Pawlowsky-Glahn V, Egozcue JJ (2001) Geometric approach to statistical analysis on the simplex. *Stoch Env Res Risk Assess* 15(384–398):10. <https://doi.org/10.1007/s004770100077>
- Al-dhurafi NA, Masseran N, Zamzuri ZH (2018) Compositional time series analysis for air pollution index data. *Stoch Env Res Risk Assess* 32:2903–2911. <https://doi.org/10.1007/s00477-018-1542-0>
- Sánchez-Balseca J, Pérez-Foguet A (2020) Spatio-temporal air pollution modelling using a compositional approach. *Heliyon* 6(9):04794. <https://doi.org/10.1016/j.heliyon.2020.e04794>
- Coenders G, Egozcue JJ, Fačevićová K, Navarro-López C, Palarea-Albaladejo J, Pawlowsky-Glahn V, Tolosana-Delgado R (2023) 40 years after Aitchison's article "the statistical analysis of compositional data": where we are and where we are heading. *SORT Stat Op Res Trans* 47(2):207–228. <https://doi.org/10.57645/20.8080.02.6>
- Aitchison J (1986) *The Statistical Analysis of Compositional Data*. Monographs on Statistics and Applied Probability. Chapman and Hall, London, UK
- Pawlowsky-Glahn V, Egozcue J, Tolosana-Delgado R (2015) *Modeling and analysis of compositional data*. Wiley, Chichester, UK
- Filzmoser P, Hron K, Templ M (2018) *Applied compositional data analysis*. Springer, New York, NY
- Aitchison J, Bacon-Shone J (1984) Log contrast models for experiments with mixtures. *Biometrika* 71(2):323–330. <https://doi.org/10.2307/2336249>
- Coenders G, Pawlowsky-Glahn V (2020) On interpretations of tests and effect sizes in regression models with a compositional predictor. *SORT Stat Op Res Trans* 44(1):201–220. <https://doi.org/10.2436/20.8080.02.100>
- Boogaart KG, Filzmoser P, Hron K, Templ M, Tolosana-Delgado R (2021) Classical and robust regression analysis with compositional data. *Math Geosci* 53(5):823–858. <https://doi.org/10.1007/s11004-020-09895-w>
- Mota Bertran A, Saez M, Coenders G (2022) Compositional and Bayesian inference analysis of the concentrations of air pollutants in Catalonia. *Spain Environ Res* 204(112388):10. <https://doi.org/10.1016/j.envres.2021.112388>
- Pawlowsky-Glahn V, Egozcue JJ, Lovell D (2015) Tools for compositional data with a total. *Stat Model* 15(2):175–190. <https://doi.org/10.1177/1471082X14535526>
- Coenders G, Martín-Fernández JA, Ferrer-Rosell B (2017) When relative and absolute information matter: compositional predictor with a total in generalized linear models. *Stat Model* 17(6):494–512. <https://doi.org/10.1177/1471082X17110398>
- Jaccard J, Turrissi R (2003) *Interaction effects in multiple regression, vol 72. Quantitative Applications in the Social Sciences*, Sage, Newbury Park, CA
- Müller I, Hron K, Fišerová E, Šmahaj J, Cakirpaloglu P, Vančáková J (2018) Interpretation of compositional regression with application to time budget analysis. *Austrian J Stat* 47(2):3–19. <https://doi.org/10.17713/ajs.v47i2.652>
- Gómez-Rubio V (2020) *Bayesian inference with INLA*. Chapman and Hall/CRC Press, Boca Raton, FL
- Sánchez-Balseca J, Pérez-Foguet A (2021) Compositional spatio-temporal PM2.5 modelling in wildfires. *Atmosphere* 12(10):1309. <https://doi.org/10.3390/atmos12101309>
- Le F, Stanford TE, Dumuid D, Wiley JF (2025) Bayesian multilevel compositional data analysis: Introduction, evaluation, and application. *Psychol Methods* <https://doi.org/10.48550/arXiv.2405.03985>

- Zhang L, Zhang X, Leach JM, Rahman A, Howell CR, Yi N (2025) Bayesian compositional generalized linear mixed models for disease prediction using microbiome data. *BMC Bioinform* 26(98):10. <https://doi.org/10.1186/s12859-025-06114-3>
- Zhang L, Zhang X, Leach JM, Rahman AF, Yi N (2024) Bayesian compositional models for ordinal response. *Stat Methods Med Res* 33(6):1043–1054. <https://doi.org/10.1177/09622802241247730>
- Zhang L, Zhang X, Yi N (2024) Bayesian compositional generalized linear models for analyzing microbiome data. *Stat Med* 43(1):141–155. <https://doi.org/10.1002/sim.9946>
- Scott DA, Benavente E, Libiseller-Egger J, Fedorov D, Phelan J, Ilina E, Tikhonova P, Kudryavstev A, Galeeva J, Clark T, Lewin A (2023) Bayesian compositional regression with microbiome features via variational inference. *BMC Bioinform* 24(1):210. <https://doi.org/10.1186/s12859-023-05219-x>
- R Core Team: R: a language and environment for statistical computing. R foundation for statistical computing, Vienna, AT (2025). <http://www.R-project.org/> Accessed 24 April 2025
- R INLA project: R INLA project (2025). <http://www.r-inla.org/home> Accessed 24 April 2025
- Mota Bertran A, Coenders G, Plaja P, Saez M, Barceló MA (2024) Air pollution and children's mental health in rural areas. *Compos Spatio-temporal Model Sci Rep* 14(19363):10. <https://doi.org/10.1038/s41598-024-70024-2>
- Dargel L, Thomas-Agnan C (2024) The link between multiplicative competitive interaction models and compositional data regression with a total. *J Appl Stat* 51(14):2929–2960. <https://doi.org/10.1080/02664763.2024.2329923>
- Morais J, Thomas-Agnan C (2021) Impact of covariates in compositional models and simplicial derivatives. *Austrian J Stat* 50(2), 1–15 <https://doi.org/10.17713/ajs.v50i2.1069>
- Sánchez-Balseca J, Pérez-Foguet A (2022) Spatially-structured human mortality modelling using air pollutants with a compositional approach. *Sci Total Environ* 813(152486):10. <https://doi.org/10.1016/j.scitotenv.2021.152486>
- Alari A, Chen C, Schwarz L, Hdansen K, Chaix B, Benmarhnia T (2023) The role of ozone as a mediator of the relationship between heat waves and mortality in 15 French urban areas. *Am J Epidemiol* 192(6):949–962. <https://doi.org/10.1093/aje/kwad032>
- Analitis A, Michelozzi P, D'Ippoliti D, De'Donato F, Menne B, Matthies F, Atkinson RW, Iñiguez C, Basagaña X, Schneider A, Lefranc A, Paldy A, Bisanti L, Katsouyanni K (2014) Effects of heat waves on mortality: effect modification and confounding by air pollutants. *Epidemiology* 25(1):15–22. <https://doi.org/10.1097/EDE.0b013e31828ac01b>
- Dear K, Ranmuthugala G, Kjellström T, Skinner C, Hanigan I (2005) Effects of temperature and ozone on daily mortality during the august 2003 heat wave in France. *Arch Environ Occup Health* 60(4):205–212. <https://doi.org/10.3200/AEOH.60.4.205-212>
- Du H, Yan M, Liu X, Zhong Y, Ban J, Lu K, Li T (2024) Exposure to concurrent heatwaves and ozone pollution and associations with mortality risk: a nationwide study in China. *Environ Health Perspect* 132(4):047012. <https://doi.org/10.1289/EHP13790>
- Qi J, Wang Y, Wang L, Cao R, Huang J, Li G, Yin P (2023) The modification effect of ozone pollution on the associations between heat wave and cardiovascular mortality. *Innov Med* 1(10):59717 <https://doi.org/10.59717/j.xinn-med.2023.100043>
- Xu R, Huang S, Shi C, Wang R, Liu T, Li Y, Zheng Y, Lv Z, Wei J, Sun H, Liu Y (2023) Extreme temperature events, fine particulate matter, and myocardial infarction mortality. *Circulation* 148(4):312–323. <https://doi.org/10.1161/CIRCULATIONAHA.122.063504>
- Lindgren F, Rue H (2015) Bayesian spatial modelling with R-INLA. *J Stat Softw* 63:1–25
- Rue H, Martino S, Chopin N (2009) Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *J R Stat Soc Ser B Stat Methodol* 71(2):319–392
- Lindgren F, Rue H, Lindström J (2011) An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *J R Stat Soc Ser B Stat Methodol* 73(4):423–498. <https://doi.org/10.1111/j.1467-9868.2011.00777.x>
- Cameletti M, Lindgren F, Simpson D, Rue H (2013) Spatio-temporal modeling of particulate matter concentration through the SPDE approach. *ASTA Adv Stat Anal* 97(109–131):10. <https://doi.org/10.1007/s10182-012-0196-3>
- Saez M, Barceló MA (2022) Spatial prediction of air pollution levels using a hierarchical Bayesian spatiotemporal model in Catalonia. *Spain Environ Model Softw* 151(105369):10. <https://doi.org/10.1016/j.envsoft.2022.105369>
- Agencia Española de Meteorología -AEMET-: Olas de calor en España desde 1975 (2024)
- METEOCAT: Balanç d'una de les onades de calor més persistents mesurades a Catalunya, juliol 2022 (2022)
- Departament de Territori, Habitatge i Transició Ecològica: Dades meteorològiques de la XEMA (2025)
- Departament de Territori, Habitatge i Transició Ecològica: Dades dels punts de mesurament automàtics de la Xarxa de Vigilància i Previsió de la Contaminació Atmosfèrica (2025)
- Rue H, Riebler A, Sørbye SH, Illian JB, Simpson DP, Lindgren FK (2017) Bayesian computing with INLA: a review. *Ann Rev Stat Appl* 4(1):395–421
- Van Niekerk J, Krainski E, Rustand D, Rue H (2023) A new avenue for Bayesian inference with INLA. *Comput Stat Data Anal* 181(107692):10. <https://doi.org/10.1016/j.csda.2023.107692>
- Watanabe S (2010) Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory. *J Mach Learn Res* 11(12):3571–3594
- Barceló MA, Saez M (2025) Assessing excess mortality and heat-attributable risk during the summer of 2022 in Catalonia, Spain: a Bayesian spatiotemporal analysis. *J Geogr Syst* 1–27:10. <https://doi.org/10.1007/s10109-025-00475-2>
- Simpson D, Rue H, Riebler A, Martins TG, Sørbye SH (2017) Penalising model component complexity: a principled, practical approach to constructing priors. *Stat Sci* 32(1):1–46. <https://doi.org/10.1214/16-STS576>
- Adin A, Krainski ET, Lenzi A, Liu Z, Martínez-Minaya J, Rue H (2024) Automatic cross-validation in structured models: Is it time to leave out leave-one-out? *Spatial Stat* 62(100843):10. <https://doi.org/10.1016/j.spasta.2024.100843>
- Figueira M, Guarner C, Conesa D, López-Quílez A, Krisztin T (2025) Unveiling land use dynamics: insights from a hierarchical Bayesian spatio-temporal modelling of compositional data. *J Agric Biol Environ Stat*. <https://doi.org/10.1007/s13253-025-00678-6>
- Martínez-Minaya J, Lindgren F, López-Quílez A, Simpson D, Conesa D (2023) The integrated nested Laplace approximation for fitting Dirichlet regression models. *J Comput Graph Stat* 32(3):805–823. <https://doi.org/10.1080/10618600.2022.2144330>
- Martínez-Minaya J, Rue H (2024) A flexible Bayesian tool for CoDa mixed models: logistic-normal distribution with Dirichlet covariance. *Stat Comput* 34(3):116. <https://doi.org/10.1007/s11222-024-10427-3>
- Ministerio de Sanidad. Gobierno de España: Plan nacional de actuaciones preventivas de los efectos del exceso de temperaturas sobre la salud 2023 (2025). https://www.sanidad.gob.es/ciudadanos/saludAmbLaboral/planAltasTemp/2023/docs/Plan_Excesos_Temperatura_2023.pdf
- Dumuid D, Pedišić Ž, Palarea-Albaladejo J, Martín-Fernández JA, Hron K, Olds T (2020) Compositional data analysis in time-use

- epidemiology: what, why, how. *Int J Environ Res Public Health* 17(7):2220. <https://doi.org/10.3390/ijerph17072220>
- Kuzik N, Duncan MJ, Beshara N, MacDonald M, Silva DAS, Tremblay MS (2025) A systematic review and meta-analysis of the first decade of compositional data analyses of 24-hour movement behaviours, health, and well-being in school-aged children. *J Act Sedentary Sleep Behav* 4(1):4. <https://doi.org/10.1186/s44167-025-00076-w>
- Sohn MB, Li H (2019) Compositional mediation analysis for microbiome studies. *Ann Appl Stat* 13(1):661–681. <https://doi.org/10.1214/18-AOAS1210>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.