

Secondary Publication



Frink, Nicolas; Schmid, Timo

Small area prediction of counts under machine learning-type mixed models

Date of secondary publication: 28.08.2025

Version of Record (Published Version), Article

Persistent identifier: urn:nbn:de:bvb:473-irb-109919x

Primary publication

Frink, Nicolas; Schmid, Timo (2025): Small area prediction of counts under machine learning-type mixed models, in: Computational Statistics & Data Analysis, Amsterdam: Elsevier Science, Vol. 211, Nr. 108218, pp. 1–16, doi: 10.1016/j.csda.2025.108218.

Legal Notice

This work is protected by copyright and/or the indication of a licence. You are free to use this work in any way permitted by the copyright and/or the licence that applies to your usage. For other uses, you must obtain permission from the rights-holders.

This document is made available under a Creative Commons license.



The license information is available online:

<https://creativecommons.org/licenses/by/4.0/legalcode>



Small area prediction of counts under machine learning-type mixed models

Nicolas Frink, Timo Schmid ^{*,*}

Institute of Statistics, Otto-Friedrich-Universität Bamberg, Bamberg, Germany

ARTICLE INFO

Keywords:

Bootstrap
Generalized linear mixed models
Overdispersed count data
Random forest
Small area estimation

ABSTRACT

Small area estimation methods are proposed that use generalized tree-based machine learning techniques to improve the estimation of disaggregated means in small areas using discrete survey data. Specifically, two existing approaches based on random forests - the Generalized Mixed Effects Random Forest (GMERF) and a Mixed Effects Random Forest (MERF) - are extended to accommodate count outcomes, addressing key challenges such as overdispersion. Additionally, three bootstrap methodologies designed to assess the reliability of point estimators for area-level means are evaluated. The numerical analysis shows that the MERF, which does not assume a Poisson distribution to model the mean behavior of count data, excels in scenarios of severe overdispersion. Conversely, the GMERF performs best under conditions where Poisson distribution assumptions are moderately met. In a case study using real-world data from the state of Guerrero, Mexico, the proposed methods effectively estimate area-level means while capturing the uncertainty inherent in overdispersed count data. These findings highlight their practical applicability for small area estimation.

1. Introduction

Linear mixed models (LMM) are a staple for analyzing unit-level survey data and estimating area-level means. These models account for the hierarchical structure of observations through random effects. An example of such a model is the nested error regression model (Battese et al., 1988), which requires the availability of unit-level survey data and administrative auxiliary information, such as data from a register or census. However, it is crucial to recognize the limitations imposed by the distributional and structural assumptions of these models, which may not always hold in small area estimation (SAE) applications (Rao and Molina, 2015; Tzavidis et al., 2018). For example, when working with LMMs, it is necessary to assume a linear relationship between the covariates and the outcome variable. This assumption may not always align with empirical evidence. Thus, it is critical to ensure the validity of model assumptions for optimal results and predictive performance of model-based SAE. If assumptions are not met, parameter estimates may be biased, and the reliability of mean squared error (MSE) estimates may be compromised (Jiang and Rao, 2020).

To circumvent the parametric assumptions inherent in LMMs, employing machine learning techniques is a viable methodological alternative. These techniques are capable of extracting predictive relationships from data, encompassing complex interactions among

* Corresponding author. Institute of Statistics, Otto-Friedrich-Universität Bamberg, Feldkirchenstraße 21, 96045 Bamberg, Germany.
E-mail address: timo.schmid@uni-bamberg.de (T. Schmid).

<https://doi.org/10.1016/j.csda.2025.108218>

Received 3 July 2024; Received in revised form 22 May 2025; Accepted 22 May 2025

Available online 30 May 2025

0167-9473/© 2025 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

covariates, without relying on explicit model assumptions (Hastie et al., 2009; Varian, 2014). Krennmair and Schmid (2022) provide a framework utilizing tree-based machine learning methods for SAE. This framework presents a non-linear, data-driven, and semi-parametric approach for continuous variables to estimate area-level means using Mixed Effects Random Forests (MERF). Random forests (Breiman, 2001) are renowned for their robust predictive performance, robustness with respect to model-misspecification, and they inherently address model-selection issues (Biau and Scornet, 2016). MERFs incorporate these benefits while also modeling hierarchical dependencies. Originally developed for continuous target variables similar to LMMs, Frink and Schmid (2024) have expanded the application of MERFs to estimate poverty indicators using binary response variables within the context of SAE. However, their focus is limited to binary variables, and the behavior of MERFs with count data remains unexplored.

Various methods for SAE with count data have been investigated. Parametric models are discussed by Ghosh et al. (1998), Jiang and Lahiri (2006), Dreassi et al. (2014), Chen et al. (2015), Boubeta et al. (2016), Hobza and Morales (2016), Boubeta et al. (2017) and Chandra et al. (2017). In contrast, Chambers et al. (2014) and Tzavidis et al. (2015) have examined semi-parametric unit-level models utilizing M-quantiles, which only partially rely on the distributional assumptions of the Poisson and Negative Binomial families. However, all of these methodologies for count data assume a linear relationship within the systematic component of the underlying model.

Building on the work of Krennmair and Schmid (2022) and Frink and Schmid (2024), who originally developed the (generalized) MERF approaches for continuous and binary data, this paper introduces semi-parametric small area estimation methods for count data using random forests. These extensions allow for a more flexible specification of the relationship between the outcome variable and the covariates, enabling their application to scenarios involving count outcomes.

The generalized MERF (GMERF) is presented as a flexible and data-driven approach that employs the Poisson distribution to model the mean behavior of count data. Furthermore, a parametric MSE bootstrap scheme is proposed to evaluate the uncertainty associated with area-level estimates. While commonly used for modeling count data, the Poisson distribution often underestimates variability in cases of overdispersion (Ver Hoef and Boveng, 2007). To address this challenge, we introduce a semi-parametric MSE bootstrap procedure. This bootstrap scheme for the GMERF operates with reduced reliance on distributional assumptions, providing a safeguard against mild deviations from the Poisson distribution. The numerical findings from our simulation studies and real-world application indicate that GMERF performs best when the assumptions of the Poisson distribution are moderately met, particularly in scenarios with minimal to low overdispersion and non-linear interactions between covariates.

Although the original MERF approach is designed for continuous outcome variables, we extend it to accommodate count data. Our numerical analysis demonstrates that the MERF provides a valid alternative for modeling count data, effectively capturing complex relationships and handling severe overdispersion without relying on explicit distributional assumptions like Poisson or Negative Binomial. Furthermore, we evaluate its numerical performance in comparison to GMERF through simulation studies. Additionally, we introduce a modified non-parametric bootstrap method for MERF, inspired by Krennmair and Schmid (2022), to account for the additional uncertainty that arises when treating a count variable as continuous.

The paper is organized as follows: Section 2 introduces GMERFs as a methodology that merges random forests with hierarchical modeling to capture dependencies among unit-level (count) observations. Section 2.2 details the process for constructing area-level estimates. To accurately evaluate the MSE of these estimates, Section 3 describes three bootstrap schemes: parametric and semi-parametric approaches for GMERFs, and an adjusted non-parametric method for the MERF. Section 4 assesses and compares the effectiveness of the proposed GMERF and MERF methods against the empirical best plug-in predictor (EBPP, Jiang (2003)) under the Poisson generalized linear mixed model through model-based simulation studies. Section 5 applies these methods to estimate area-level means of educational attainment among women in the Mexican state of Guerrero, also addressing the uncertainty of these estimates. Section 6 presents a design-based simulation study, based on data from Guerrero, to assess the proposed estimators in a close-to-reality environment. The study concludes in Section 7 with recommendations for further research in the field of SAE.

2. Methodology

Krennmair and Schmid (2022) introduced the MERF for continuous data in the context of SAE. This section introduces an extension of the MERF, the GMERF, a semi-parametric unit-level model designed to handle Quasi-Poisson-distributed count outcomes. The model employs a regression forest approach to estimate area-level means, utilizing both unit-level survey and unit-level administrative data. This technique effectively addresses the challenges of modeling count data by combining the flexibility of machine learning methods, which can account for non-linear relationships, with parametric Poisson assumptions.

2.1. Generalized semi-parametric unit-level model

We examine a finite population P divided into D disjunct areas P_i each with a sub-population size of N_i , where $i = 1, \dots, D$ specifies the areas and $N = \sum_{i=1}^D N_i$ defines the population size of all areas. The sample s consists of area-specific sub-samples s_i with a total size $n = \sum_{i=1}^D n_i$. In contrast, non-sampled observations are denoted as r_i with a size of $N_i - n_i$. We denote individual units within each area as $j \in s_i$ for sampled and $j \in r_i$ for non-sampled observations. The discrete unit-level outcome variable is represented by y_{ij} , with information available for n observations within the sample s . The vector $\mathbf{x}_{ij} = [x_1, x_2, \dots, x_p]^T$ encompasses p auxiliary covariates, and these auxiliary variables are known for all N units in our population P . We assume that y_{ij} follows a generalized semi-parametric unit-level model:

$$\eta_{ij} = f(\mathbf{x}_{ij}) + \nu_i, \quad (1)$$

$$v_i \sim N(0, \sigma_v^2),$$

$$E(y_{ij} | v_i) = \mu_{ij} = g(\eta_{ij}),$$

where the fixed part, $f(\cdot)$, is defined by a random forest and models the conditional mean of the linear predictor (η_{ij}) given the covariates. The term v_i denotes a domain-specific random intercept with its variance σ_v^2 , capturing small-area variations in the conditional distribution of the linear predictor given \mathbf{x}_{ij} . $g(\cdot)$ is a monotonic, differentiable link function that specifies the function of the mean (μ_{ij}) equated to the systematic component. In this work, we assume that the individual y_{ij} values in domain i are independent Poisson random variables with $Var(y_{ij} | v_i) = E(y_{ij} | v_i)$. Consequently, the mean of the response variable is connected to the linear predictor via the logarithmic link function: $\mu_{ij} = \exp(\eta_{ij})$. If y_{ij} were a continuous variable following a Gaussian distribution, the link function would simplify to the identity function. Thus, in model (1), the GMERF reduces to the MERF, as proposed by Krennmair and Schmid (2022). Furthermore, model (1) can be seen as a generalized linear mixed model (GLMM) by setting $f(\mathbf{x}_{ij}) = \mathbf{x}_{ij}^T \beta$, where $\beta = [\beta_1, \dots, \beta_p]^T$ represents the regression parameters.

Given that inference using GLMMs presents computational hurdles due to the presence of high-dimensional integrals in the likelihood, which are not amenable to analytical evaluation, we employ an estimation strategy for model (1) akin to Frink and Schmid (2024). The key distinction in our approach lies in the use of a different link function, tailored to the count data in our model, while the other elements of the estimation strategy align closely with those described by Frink and Schmid (2024). Specifically, this approach utilizes the penalized quasi-likelihood (PQL) method (Breslow and Clayton, 1993; Stroup, 2012) in conjunction with the EM-algorithm (Moon, 1996) to address the computational challenges inherent in GLMM estimation. More concretely, a weighted MERF pseudo-model is devised, leveraging a linearized target variable $y_{L,ij} = g(\mu_{ij}) + (y_{ij} - \mu_{ij})g'(\mu_{ij})$ and weights $w_{ij} = (v_{ij}g'(\mu_{ij})^2)^{-1}$, with v_{ij} being a known variance function. The algorithm follows a doubly iterative process, with micro iterations nested within macro iterations. During each macro-iteration, the linearized response variable and weights are updated. These revised values then serve as the response variable and weights for the subsequent micro-iterations. To fit model (1) on survey data, the GMERF algorithm uses initial estimates for μ_{ij} , weights w_{ij} and $y_{L,ij}$ from a generalized linear model (GLM) during the first macro-iteration. In the subsequent micro-iterations, the algorithm (i) estimates the forest function, assuming the random effects to be correct, and (ii) estimates the random effects part using the weighted pseudo-model, assuming the Out-of-Bag (OOB) predictions from the regression forest to be valid. Convergence is monitored by the relative change in the log-likelihood of model (1) between two micro-iterations. It is considered achieved when its relative change is less than a specified number. Once convergence is achieved, the initial values for the next macro-iteration are updated. For the specific case of y_{ij} being Poisson distributed, the initial values are updated as follows:

$$\begin{aligned} \eta_{ij} &= \hat{f}(\mathbf{x}_{ij})^{OOB} + \hat{v}_i, \\ \mu_{ij} &= \exp(\eta_{ij}), \\ y_{L,ij} &= \log(\mu_{ij}) + \frac{(y_{ij} - \mu_{ij})}{\mu_{ij}}, \\ w_{ij} &= \frac{1}{v_{ij}(\mu_{ij})}, \end{aligned} \tag{2}$$

where $\hat{f}(\mathbf{x}_{ij})^{OOB}$ and \hat{v}_i reflect their estimated values from the micro-level convergence achieved in the preceding macro-iteration. The predictive performance of our method is significantly influenced by two key tuning parameters: the number of split candidates at each node, which controls the degree of decorrelation, and the number of trees, which together enhance the model's stability and accuracy. For further methodological details, we refer to Frink and Schmid (2024).

The Poisson distribution is very common for modeling the mean behavior of count outcomes. However, it may underestimate variability in cases of overdispersion, where the observed variance of the response variable exceeds what the Poisson distribution predicts. Overdispersion is a common occurrence in count data analysis and can significantly affect result interpretation. Failure to account for overdispersion in the model can result in overly narrow standard errors and confidence intervals, as well as excessively lenient significance tests, leading to the detection of effects that do not truly exist (Ver Hoef and Boveng, 2007). One widely used approach to address overdispersion is to employ a Quasi-Poisson model (Gourieroux et al., 1984). In Quasi-Poisson models, $E(y_{ij}^{qp} | v_i) = E(y_{ij} | v_i) = \mu_{ij}$ and $Var(y_{ij}^{qp} | v_i) = \theta \mu_{ij}$, where θ is a dispersion parameter. The algorithm mentioned above can be directly expanded to incorporate Quasi-Poisson models to address overdispersion concerns. Nevertheless, given our primary focus on predicting means within the machine learning framework and our lesser concern with standard error inference, the Quasi-Poisson model does not yield any discernible advantage within the generalized mixed effect regression forest model. Results for the Quasi-Poisson distribution in terms of GMERFs are available and can be requested from the authors.

2.2. Domain-level estimator for counts

The proposed estimator for the area-level mean is given by:

$$\begin{aligned} \hat{\eta}_i &= \hat{f}(\mathbf{x}_i) + \hat{v}_i, \\ \hat{\mu}_i &= \frac{1}{N_i} \sum_{j=1}^{N_i} \exp(\hat{\eta}_{ij}) \text{ for } i = 1, \dots, D. \end{aligned} \tag{3}$$

For non-sampled areas, the proposed estimator for the area-level mean simplifies to the fixed component obtained from the random forest:

$$\hat{\eta}_{ij} = \hat{f}(\mathbf{x}_{ij}),$$

$$\hat{\mu}_i = \frac{1}{N_i} \sum_{j=1}^{N_i} \exp(\hat{\eta}_{ij}).$$

Our proposed estimator (3) offers an advancement over existing SAE methods for count data by addressing issues commonly associated with model selection. This improvement arises from employing the random forest technique, which inherently optimizes model selection to capture higher-order effects and non-linear interactions. This feature is practically demonstrated in our application to real-world data from the state of Guerrero in Section 5. Furthermore, the random forest is well-suited for analyzing high-dimensional covariate data, effectively managing scenarios where the number of covariates exceeds the sample size (Hastie et al., 2009).

However, the GMERF estimator relies on the assumptions of the Poisson distribution due to the linearized target variable and the weights used in the algorithm. In contrast, the MERF, originally designed for continuous outcomes, offers a viable alternative for modeling count data. This is achieved by having the MERF iteratively estimate a random forest, assuming the random effects term to be correct, and then estimating the random effects components (and variance components) by decomposing the random forest residuals using a linear mixed model. Moreover, the MERF has the advantage of not depending on the assumptions of the Poisson or Negative Binomial distributions, making it an alternative option when dealing with severe overdispersion.

3. Uncertainty estimation

Estimating the MSE of small area estimates poses a substantial challenge (Rao and Molina, 2015). In this section, we present an adaption of two existing bootstrap schemes, proposed by González-Manteiga et al. (2007) and Chambers and Chandra (2013), to estimate the MSE of the small area estimator introduced in equation (3). The introduced bootstrap methods account for the additional uncertainty associated with modeling count outcomes. The primary distinction between the two bootstrap procedures lies in the mechanism used to generate the bootstrap population. Following González-Manteiga et al. (2007) and González-Manteiga et al. (2008), the first bootstrap scheme generates bootstrap realizations of the random effects parametrically and explicitly employs the Poisson distribution to simulate y_{ij} . This approach is widely used in SAE applications, such as in Boubeta et al. (2016) for Poisson mixed models. It is particularly effective in scenarios where the target variable follows a Poisson distribution, as it directly leverages the properties of this distribution. However, count data often exhibit overdispersion, which violates the assumptions underlying the Poisson distribution. As a result, the parametric bootstrap may become unsuitable in such cases. To address this, we propose adopting a semi-parametric bootstrap approach for the GMERF - at least as a supplementary MSE estimation method - drawing on the work of Chambers and Chandra (2013), Flores-Agreda and Cantoni (2019) and Krennmair and Schmid (2022). Since this approach is less reliant on the assumptions of the Poisson distribution, we expect the semi-parametric bootstrap to serve as a safeguard against deviations from the Poisson distribution, particularly in the presence of overdispersion.

Furthermore, we introduce an adjusted non-parametric MSE bootstrap scheme for the MERF with count outcomes. Building on the work of Chambers and Chandra (2013) and Krennmair (2022), this bootstrap procedure is adjusted to account for the uncertainty of the estimation, enabling the generation of a discrete count variable in the bootstrap population. Detailed information on this bootstrap method can be found in Appendix A.1.

3.1. Parametric bootstrap

The procedural steps outlined for the proposed parametric bootstrap of the GMERF closely resemble those detailed in Frink and Schmid (2024) and can be summarized as follows:

1. For $b = 1, \dots, B$:
 - (a) Create bootstrap random effects for each of the D areas by $v_i^{(b)} \sim N(0, \hat{\sigma}_v^2)$.
 - (b) Create a bootstrap population of size N , by generating values $y_{ij}^{(b)}$ from a Poisson distribution with

$$\mu_{ij}^{(b)} = \exp\left(\hat{f}(\mathbf{x}_{ij}) + v_i^{(b)}\right).$$

- (c) Determine the true bootstrap population area means $\mu_i^{(b)}$ as $\frac{1}{N_i} \sum_{j=1}^{N_i} y_{ij}^{(b)}$ for all $i = 1, \dots, D$.
 - (d) For each bootstrap population, select a bootstrap sample that matches the original sample size n_i and estimate $\hat{f}^{(b)}()$ and $\hat{v}_i^{(b)}$. Calculate area-level means $\hat{\mu}_i^{(b)}$.
2. Utilizing the B bootstrap samples, the MSE estimator is derived as follows:

$$\widehat{MSE}_i = \frac{1}{B} \sum_{b=1}^B \left(\mu_i^{(b)} - \hat{\mu}_i^{(b)} \right)^2.$$

3.2. Semi-parametric bootstrap

The parametric bootstrap may be unsuitable in the presence of overdispersion. To address this, we discuss a semi-parametric bootstrap approach that reduces reliance on strict Poisson distribution assumptions. This bootstrap scheme for the GMERF ensures that a discrete count target variable is generated in the bootstrap population without explicitly relying on the Poisson distribution. The proposed bootstrap scheme is described below:

1. For given $\hat{f}()$, calculate the marginal pearson residuals $z_{ij} = \frac{y_{ij} - \exp(\hat{f}(\mathbf{x}_{ij}))}{\sqrt{\exp(\hat{f}(\mathbf{x}_{ij}))}}$.
2. Employing the marginal pearson residuals z_{ij} , calculate level-2 residuals for each area by

$$\bar{z}_i = \frac{1}{n_i} \sum_{j=1}^{n_i} z_{ij} \quad \text{for } i = 1, \dots, D.$$

3. Compute the vector of level-1 residuals by $\hat{z}_{ij} = z_{ij} - \bar{z}_i$. Similar to Krennmair and Schmid (2022), we adjust the residuals \hat{z}_{ij} to the bias-corrected variance and center them, indicated by \hat{z}_{ij}^c . Our approach differs in that we use the marginal Pearson residuals of the GMERF, as described in step 1 above, instead of the marginal raw residuals of the MERF for the adjustment process. Level-2 residuals \bar{z}_i are additionally adjusted for the estimated variance $\hat{\sigma}_v^2$ and centered, indicated by $\bar{z}^c = (\bar{z}_1^c, \dots, \bar{z}_D^c)$.
4. For $b = 1, \dots, B$:

(a) Independently draw samples with replacement from the scaled and centered level-1 and level-2 residuals:

$$z_{ij}^{(b)} = \text{srswr}(\hat{z}_{ij}^c, N) \quad \text{and} \quad \bar{z}_i^{(b)} = \text{srswr}(\bar{z}^c, D).$$

(b) Calculate $\mu_{ij}^{(b)} = \exp(\hat{f}(\mathbf{x}_{ij}) + \bar{z}_i^{(b)})$ and the corresponding $\tilde{y}_{ij}^{(b)} = \mu_{ij}^{(b)} + \sqrt{\mu_{ij}^{(b)}} \times z_{ij}^{(b)}$ for $j = 1, \dots, N$. Note that $\tilde{y}_{ij}^{(b)}$ is defined on a continuous scale. Following Flores-Agreda and Cantoni (2019), the inclusion of the residuals $z_{ij}^{(b)}$ ensures that the bootstrap data capture the variability observed in the original data, thereby preventing the underestimation of uncertainty. Scaling these residuals by the estimated standard deviation $\sqrt{\mu_{ij}^{(b)}}$ helps to maintain consistency in variability between the bootstrap and original data.

(c) To obtain a count target variable, we match $\tilde{y}_{ij}^{(b)}$ with the set of estimated unit-level predictors from the sample $\{\hat{\mu}_t \mid \hat{\mu}_t = \exp(\hat{f}(\mathbf{x}_{it}) + \hat{v}_t), \text{ for } t = 1, \dots, n\}$ by finding the corresponding index \tilde{t} solving

$$\min_t \left| \tilde{y}_{ij}^{(b)} - \hat{\mu}_t \right|.$$

(d) Define the bootstrap population by $y_{ij}^{(b)} = y_{\tilde{t}}$ for $j = 1, \dots, N$ and calculate the true bootstrap population mean $\mu_i^{(b)}$ for $i = 1, \dots, D$.

(e) For each bootstrap population, select a bootstrap sample that matches the original sample size n_i and estimate $\hat{f}^{(b)}()$ and $\hat{v}_i^{(b)}$. Obtain estimates for the mean $\hat{\mu}_i^{(b)}$.

5. Using the B bootstrap samples, the MSE estimator is computed as follows:

$$\widehat{MSE}_i = \frac{1}{B} \sum_{b=1}^B \left(\mu_i^{(b)} - \hat{\mu}_i^{(b)} \right)^2.$$

Both bootstrap schemes we described in this section are empirically evaluated in Section 4.

4. Model-based simulation study

The use of model-based simulations enables a controlled empirical assessment of our proposed methods for point and uncertainty estimates. We compare the point and uncertainty estimates for domain-level means derived from the GMERF model (1) with those obtained from two competing models. In particular, we examine the performance of GMERFs in comparison to the EBPP, which is based on a Poisson GLMM:

$$\hat{\mu}_i^{EBPP} = \frac{1}{N_i} \left(\sum_{j \in s_i} y_{ij} + \sum_{j \in r_i} \hat{\mu}_{ij} \right),$$

where $\hat{\mu}_{ij} = \exp(\mathbf{x}_{ij}^T \hat{\beta} + \hat{v}_i)$ (Jiang, 2003). By contrasting the performance of this linear-based competitor (on the linear predictor sale) with our more flexible approach, which incorporates semi-parametric and non-linear modeling, we aim to showcase the advantages of the GMERF methodology. Additionally, we compare the GMERF method with the MERF (for originally continuous outcomes):

$$\hat{\mu}_i^{MERF} = \frac{1}{N_i} \sum_{j=1}^{N_i} \hat{f}(\mathbf{x}_{ij}) + \hat{v}_i.$$

Table 1
Model-based simulation scenarios.

Scenario	Linear predictor	Mean	y	x_1	x_2	v
<i>Normal-Poisson</i>	$\eta = 2 + x_1 + x_2 + v$	$\mu = \exp(\eta)$	$Pois(\mu)$	$U(-1, 1)$	$N(-1, 1)$	$N(0, 0.3^2)$
<i>Interaction-Poisson</i>	$\eta = 2 + 2x_1x_2 + x_2^2 + v$	$\mu = \exp(\eta)$	$Pois(\mu)$	$U(-1, 1)$	$N(-1, 1)$	$N(0, 0.3^2)$
<i>NB3</i>	$\eta = 2 + 2x_1x_2 + x_2^2 + v$	$\mu = \exp(\eta)$	$NB(\mu, 3)$	$U(-1, 1)$	$N(-1, 1)$	$N(0, 0.3^2)$
<i>NB1</i>	$\eta = 2 + 2x_1x_2 + x_2^2 + v$	$\mu = \exp(\eta)$	$NB(\mu, 1)$	$U(-1, 1)$	$N(-1, 1)$	$N(0, 0.3^2)$

The aim of comparing both machine learning methods is to highlight the advantages of GMERF when the Poisson distribution is applicable or only moderately violated, and to demonstrate the suitability of MERF for count data in situations of severe overdispersion.

The simulation setup involves a finite population U with a total size of $N = 50,000$, divided into $D = 50$ disjoint areas U_1, \dots, U_D each having $N_i = 1,000$ units. We generate samples using stratified random sampling, where the 50 small areas are treated as strata. This approach yields a total sample size of $n = \sum_{i=1}^D n_i = 921$. The number of sampled units per area varies from 8 to 29, with a median of 18. These sample sizes are consistent with the area-level sample sizes found in the application discussed in Section 5.

We consider four scenarios denoted as *Normal-Poisson*, *Interaction-Poisson*, *NB3*, *NB1* and repeat each scenario independently $M = 500$ times. The comparison of competing model-estimates under these four scenarios allows us to examine the performance under two major dimensions: Firstly, the presence of overdispersed data delineated by the Negative Binomial distribution and secondly, the presence of unknown non-linear interactions between covariates on the linear predictor scale.

We initially establish a baseline scenario, *Normal-Poisson*, for Poisson GLMMs generating Poisson-distributed outcomes without interactions or quadratic terms. Since the assumption of linearity in the model is satisfied, our goal is to demonstrate that GMERFs perform similar to their linear counterparts in the reference scenario. In contrast, the *Interaction-Poisson* scenario introduces a more intricate model, incorporating quadratic terms and interactions at the linear predictor level. This scenario aims to underscore the benefits of non-linear modeling approaches. Overdispersion frequently complicates the analysis of count data. To more realistically mirror these situations, we apply a Negative Binomial distribution (with scale parameter 3 and 1) in scenarios *NB3* and *NB1*. These scenarios utilize complex models with interaction effects but vary according to the scale parameter s , which influences the degree of overdispersion. A smaller s value indicates increased overdispersion. Further information on the data generation process for each scenario is detailed in Table 1.

We evaluate the point estimates for the area-level means using two quality measures: bias and root mean squared error (RMSE). To assess the proposed MSE estimators, we examine the relative bias of the root mean squared error (RB-RMSE) and the relative root mean squared error of the RMSE:

$$BIAS_i = \frac{1}{M} \sum_{m=1}^M \left(\hat{\mu}_i^{(m)} - \mu_i^{(m)} \right),$$

$$RMSE_i = \sqrt{\frac{1}{M} \sum_{m=1}^M \left(\hat{\mu}_i^{(m)} - \mu_i^{(m)} \right)^2},$$

$$RB-RMSE_i = \frac{\sqrt{\frac{1}{M} \sum_{m=1}^M MSE_{est_i}^{(m)}} - RMSE_i}{RMSE_i},$$

$$RRMSE-RMSE_i = \frac{\sqrt{\frac{1}{M} \sum_{m=1}^M \left(\sqrt{MSE_{est_i}^{(m)}} - RMSE_i \right)^2}}{RMSE_i},$$

where $\hat{\mu}_i^{(m)}$ denotes the estimated mean for area i derived from any of the aforementioned methods, and $\mu_i^{(m)}$ represents the true mean for area i in simulation round m . The estimation of $MSE_{est_i}^{(m)}$ is carried out using the proposed bootstrap methods detailed in Section 3 and Appendix A.1.

To conduct the model-based simulation, we utilize R (R Core Team, 2023). The EBPP estimates are generated using the *lme4* package (Bates et al., 2015). For the proposed GMERF and MERF methods, we use the *ranger* package (Wright and Ziegler, 2017) alongside *lme4* (Bates et al., 2015). To ensure the algorithms converge, we apply a precision of $1e^{-5}$ for the relative change in the log-likelihood criterion (applicable to the MERF algorithm as well) and a precision of 0.001 for the relative change in $\hat{\eta}$.

Table 2 reports the empirical bias and the RMSE of each method across the four scenarios. In the *Normal-Poisson* scenario the EBPP, MERF, and GMERF estimators show varying degrees of positive bias, with the median bias magnitude increasing in that order. For the *Interaction - Poisson* data-generating process, biases not only increase but also change sign for all methods, indicating more complex error dynamics. In the Negative Binomial scenarios, designed to examine model performance under conditions of overdispersion, the MERF demonstrates a lower bias in comparison to the GMERF. The GMERF shows the highest bias across these scenarios, likely due to its reliance on the approximation method in its algorithm.

Fig. 1, along with Table 2, evaluates the RMSE for each estimation method across various scenarios. In the *Normal-Poisson* scenario, EBPP outperforms both MERF and GMERF. This suggests that EBPP's adherence to the model's fixed effects leads to more accurate estimates when the assumptions of the Poisson distribution are met. In the more complex *Interaction-Poisson* scenario, MERF shows

Table 2
Mean and median of bias and RMSE across areas for point estimates.

	<i>Normal-Poisson</i>		<i>Interaction-Poisson</i>		<i>NB3</i>		<i>NB1</i>	
	Median	Mean	Median	Mean	Median	Mean	Median	Mean
Bias								
EBPP	0.0054	0.0123	-0.0895	-0.0887	-0.0609	-0.0489	-0.0894	-0.1013
GMERF	0.0774	0.0834	-0.1376	-0.1587	0.4432	0.4981	3.1600	3.1510
MERF	0.0075	0.0099	-0.0628	-0.0628	0.0830	0.0592	0.0675	0.0750
RMSE								
EBPP	1.0910	1.1525	1.5450	1.6360	4.3740	4.5650	7.3670	7.5960
GMERF	1.2449	1.2979	1.4300	1.4860	4.2140	4.3490	8.2460	8.8140
MERF	1.3610	1.4370	1.9300	2.0120	4.1360	4.2170	5.6160	5.6930

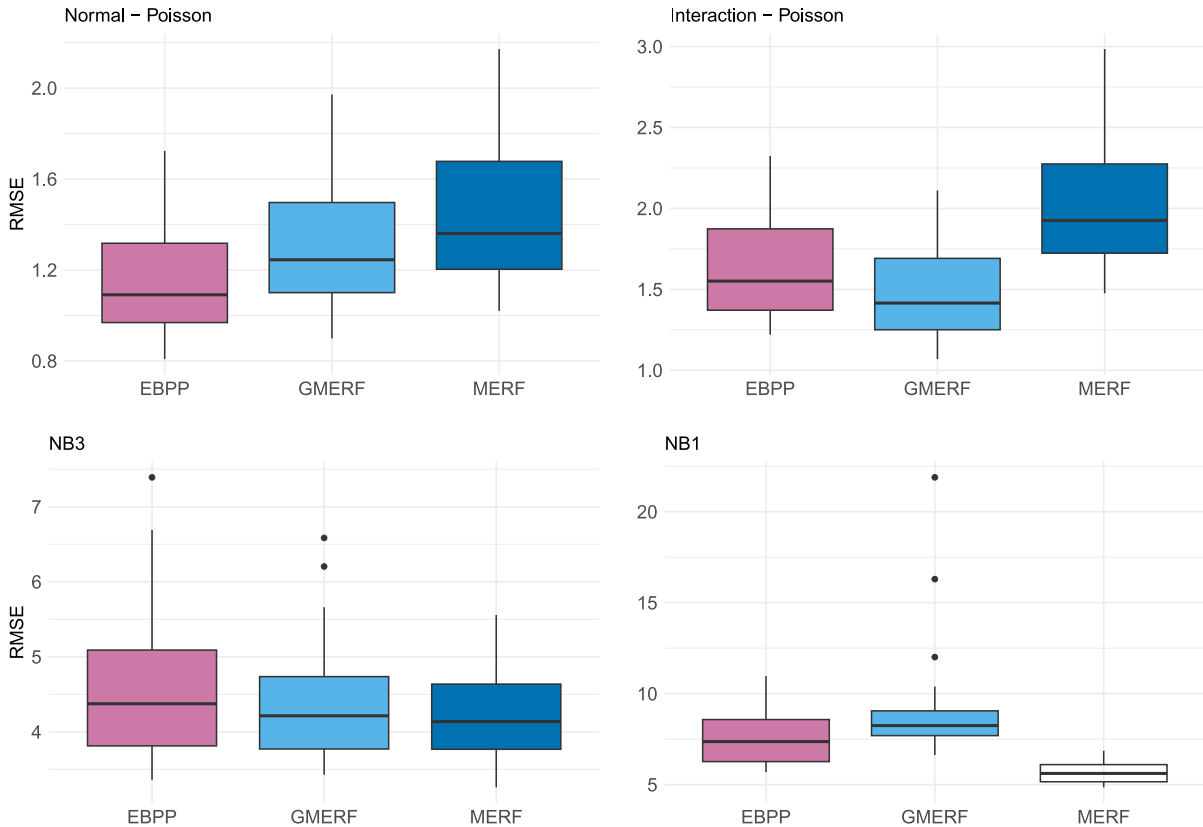


Fig. 1. Performance of area-specific point estimates. Comparison of empirical RMSEs for area-level averages under four scenarios.

a higher RMSE compared to EBPP, indicating lower accuracy. MERF’s performance suggests a greater sensitivity to deviations of distribution assumptions compared to violations of linearity assumptions, as EBPP remains more efficient under the Poisson assumption with non-linear relationships in the fixed effects part of the model. Notably, GMERF’s point estimates surpass those of EBPP, highlighting its effectiveness in complex models that involve interactions and non-linear relationships. In scenarios incorporating a Negative Binomial distribution, MERF’s performance tends to improve as the degree of overdispersion increases. This superiority (especially in the *NB1* scenario) highlights MERF’s ability to effectively handle overdispersed count data, benefiting from its flexibility in modeling both continuous and discrete targets without strict reliance on Poisson model assumptions. Overall, the insights from Table 2 and Fig. 1 indicate that GMERF delivers competitive performance when the Poisson distribution assumptions are moderately met, whereas MERF excels in scenarios where these assumptions are severely violated. Additional results using unbalanced population sizes, available in the online supplementary material, yield consistent findings and further support the stability of our results.

We now assess the performance of the MSE estimators presented in Section 3 and Appendix A.1 using $B = 200$ bootstrap replications. Table 3 displays the performance of the three bootstrap procedures, evaluated using RB-RMSE and RMSE-RMSE. Specifically, the proposed MSE estimators show reasonably low to moderate relative bias in terms of both mean and median values across all scenarios. The parametric bootstrap estimator (GMERF P) performs similarly in simpler scenarios, such as *Normal-Poisson*, and in more complex scenarios like *Interaction-Poisson*. The results for Negative Binomial settings using this parametric bootstrap are affected by

Table 3
Performance of bootstrap MSE estimators in model-based simulation: mean and median of RB-RMSE and RRMSE-RMSE across areas.

	Normal-Poisson		Interaction-Poisson		NB3		NB1	
	Median	Mean	Median	Mean	Median	Mean	Median	Mean
RB-RMSE [%]								
GMERF P	3.57	3.45	-0.36	5.33	-26.94	-23.17	-29.83	-4.90
GMERF SP	8.23	8.21	-2.24	-1.29	5.81	6.71	-8.54	2.21
MERF NPC	1.97	2.03	0.28	0.48	3.95	4.21	-0.83	-1.21
RRMSE-RMSE [%]								
GMERF P	10.16	10.56	47.91	57.39	56.30	59.87	61.85	87.57
GMERF SP	23.01	26.26	24.42	24.77	42.79	47.21	60.75	74.41
MERF NPC	7.90	8.36	12.86	13.29	15.27	15.67	23.97	24.35

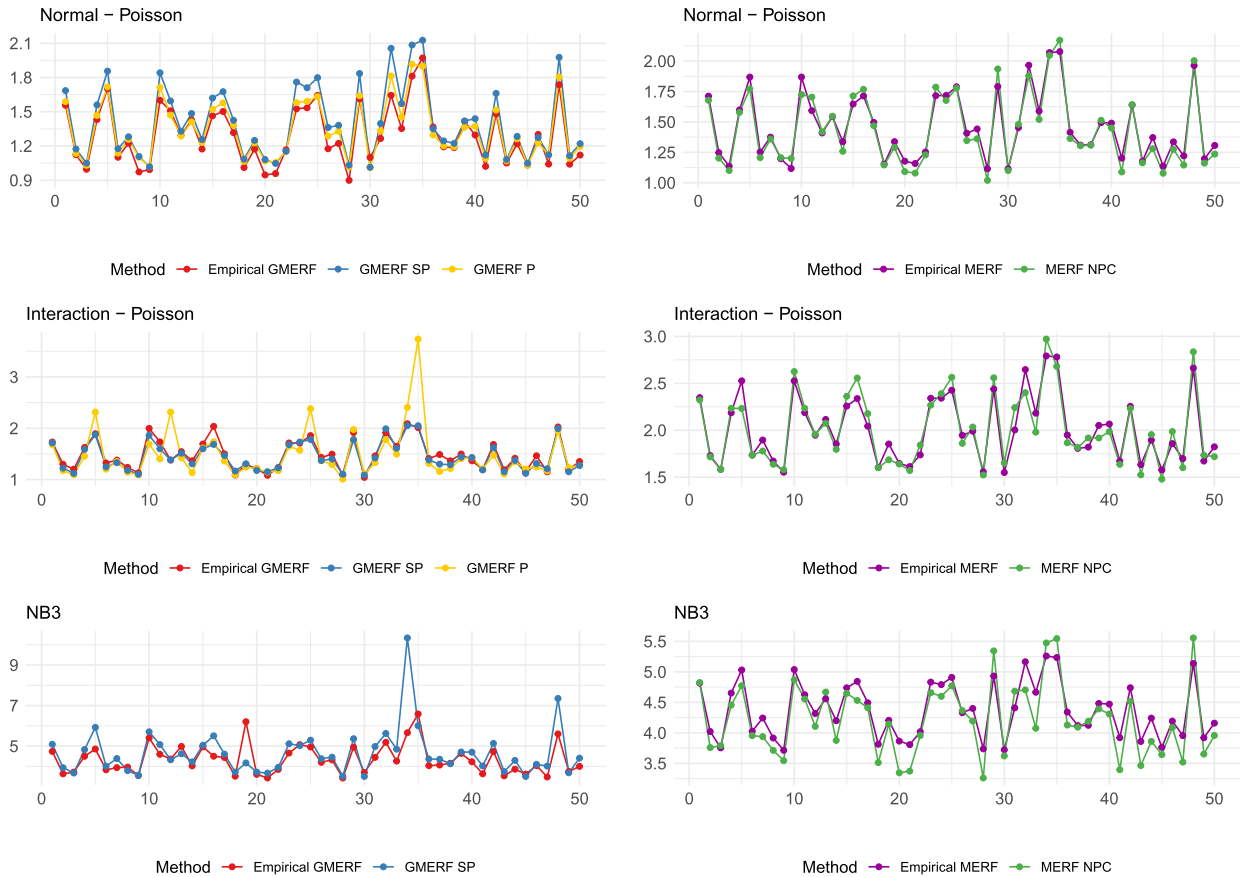


Fig. 2. Empirical and bootstrapped area-level RMSEs for three scenarios.

its reliance on Poisson distribution assumptions, which can lead to severe underestimation when these assumptions are violated. The semi-parametric bootstrap estimator for GMERF (GMERF SP) exhibits slight variability in relative bias across different scenarios. The adjusted non-parametric bootstrap estimator for MERF (MERF NPC) consistently shows low bias and maintains precision across various scenarios. Notably, GMERF SP performs comparably to MERF NPC in scenarios involving a Negative Binomial distribution in terms of RB-RMSE. While Table 3 does not directly convey the area tracking properties of the estimated RMSE versus the empirical RMSE, Fig. 2 provides further insights. Based on the tracking properties, we conclude that using the parametric, semi-parametric and non-parametric bootstraps for estimating the MSE appear to have appealing properties regarding bias and stability.

5. Application to education data from Guerrero-Mexico

In this section, we describe the application of our SAE methods to estimate the average number of school years completed by women in the municipalities of Guerrero, Mexico, based on data from 2010.

Table 4
Summary statistics on in- and out-of-sample areas: area-specific sample size of census and survey data.

Municipalities	Total		In-sample		Out-of-sample	
	81		40		41	
	Min.	Q1	Median	Mean	Q3	Max.
Survey area sizes	8.00	15.00	21.00	36.12	32.00	458.00
Census area sizes	361.00	619.00	833.00	1,386.00	1,656.00	6,297.00

We combine two primary data sources: the Mexican household income and expenditure survey ENIGH (Encuesta Nacional de Ingreso y Gastos de los Hogares) and census microdata provided by the National Institute of Statistics and Geography (Instituto Nacional de Estadística y Geografía). The ENIGH survey encompasses data from 1,445 households spread across 40 municipalities in Guerrero, while the census dataset includes information on 112,265 households from all 81 municipalities in the state. The survey samples are notably varied, with individual municipality samples ranging from a minimum of 8 households to a maximum of 458, and a median of 21 households per municipality. This disparity in sample sizes results in 41 municipalities being left out of the survey sample, highlighting a significant challenge in applying SAE methods effectively. Table 4 summarizes these domain-specific data characteristics, providing an overview of the coverage and distribution of data points across the municipalities.

The implementation of SAE methodologies involves the development of a working model using survey data, which is subsequently refined with census or administrative data. In the context of Poisson GLMMs, such as the EBPP approach, selecting the appropriate variables is crucial. We apply the Akaike information criterion to identify the optimal model for predicting the average number of school years completed by women. This process resulted in the selection of 21 out of 39 potential covariates for our final Poisson GLMM model.

In contrast to the explicit variable selection required for Poisson GLMMs, random forests implement an implicit model selection process (Breiman, 2001). Fig. 3 illustrates the insights gained from this method through partial dependence plots (PDP) and variable importance plots (VIP) (Greenwell, 2017; Greenwell et al., 2020). The PDP quantifies the marginal effect of specific predictors on the target variable, highlighting the non-linear dynamics between predictors and outcomes. The VIP ranks the significance of each predictor based on the mean decrease in impurity (variance), which is aggregated from the number of splits across all trees that involve the predictor. Notably, key predictors such as years of working experience (*jexp*), household income (*inglabpc*) and the level of education in relation to average (*escol_rel_hog*) emerge as critical for modeling the educational attainment of women. Overall, the analysis in Fig. 3 underscores the complex interplay between predictors and the target variable, affirming the value of advanced modeling techniques in uncovering these relationships in the context of SAE.

Fig. 4 presents an analysis of Pearson residuals from the Poisson GLMM applied to the estimation of school years for women in Guerrero municipalities. The histogram displays a positive skew in the residuals and the presence of some large values. This pattern is further evidenced by the plot showing the distribution of residuals across municipalities, with several exhibiting numerous positive residuals. The red dashed lines in the right plot of Fig. 4 indicate the values -2 and 2, beyond which the Pearson residuals may suggest a poor model fit. Several residuals can be seen exceeding these values. These observations suggest a potential issue with overdispersion - a critical aspect considering the equidispersion assumption of Poisson models. To verify this, we conducted overdispersion tests comparing Poisson and Negative Binomial models. Both the Likelihood Ratio Test (LRT) and Dean's PB test indicate significant overdispersion, with LRT yielding a statistic of 67.471 and $p < 2.2e^{-16}$, and Dean's PB test showing 6.840 with $p = 3.953e^{-12}$.

Additionally, Fig. 5 illustrates a plot of raw residuals against fitted values, revealing a distinct pattern across the range of fitted values. In particular, for predicted school years between 0 and 10, the residuals exhibit higher variability and a discernible concentration of negative values, indicating potential difficulties in accurately predicting outcomes for individuals with fewer years of education. This pattern may indicate the presence of unaccounted variability or non-linear effects within this range, which could be indicative of potential misspecification in the Poisson GLMM. Given these findings, we argue that the model's limitations may warrant the exploration of alternative, more flexible modeling approaches. Consequently, we estimate both the GMERF and the MERF.

Fig. 6 showcases the results from three different estimation methods applied. These methods include direct estimation, which is feasible for 40 of the 81 municipalities, as well as the model-based approaches provided by GMERF and MERF. The direct estimates serve as a baseline, feasible only for municipalities with sufficient data. Conversely, GMERF and MERF extend our insights into regions where direct data are lacking, thereby enhancing our understanding of regional disparities in educational attainment. All three methods illustrate clear regional differences in educational outcomes among the municipalities. However, the point estimates from GMERF and MERF are similar, suggesting that both model-based methods provide consistent estimates despite their methodological differences. Furthermore, the maps provide useful information on the geographical distribution of the average number of years of schooling for women.

Fig. 7 highlights the variability in the estimation accuracy for the average number of school years for women, as measured by the coefficients of variation (CV). We use the calibrated bootstrap method from the R package *emdi* (Kreutzmann et al., 2019) for direct estimates, and the semi- and non-parametric bootstrap procedures outlined in Section 3 and Appendix A.1 for GMERF and MERF, each involving $B = 200$ bootstrap replications. Given the established presence of overdispersion, we opted to forgo the use of the parametric bootstrap in this application. For both in-sample and out-of-sample domains, the comparison between GMERF and MERF reveals no notable differences in precision, suggesting comparable performance across these model-based methods.




Fig. 3. Visual diagnostics on predictive relations between dependent variable *schoolFem* and predictors. Partial dependency plots (left-hand side); variable importance plot (right-hand side).

The observed similarities in CVs across different machine learning models and bootstrap methods could be attributed to the underlying strength of overdispersion in the data, as evidenced by a dispersion ratio of 1.50 in the estimated Poisson GLMM. This corresponds closely with the scale parameter of 3.071 in the Negative Binomial distribution used in the analysis, which aligns well with the model-based simulation results for the *NB3* setting (scale parameter of 3) focusing solely on in-sample domains. It seems that there are no substantial differences between GMERF and MERF in a data scenario that shows only slight to moderate overdispersion.

Further validation of these findings is provided in Section 6, where a design-based simulation study enhances the reliability of the point and MSE estimators discussed here, allowing for a more detailed exploration of their performance.

6. Design-based simulation

This section assesses the efficacy of the proposed small area estimation methods by employing a design-based simulation, using the same dataset and structure as the real-data application discussed in Section 5. The simulation was conducted with 500 independent pseudo-survey samples drawn from the fixed population census dataset of Guerrero. Each pseudo-sample mirrors the characteristics of the original survey, maintaining the same number of in-sample municipalities. This ensures that each of the 500 pseudo-survey samples has consistent structure and overall sample size. The true values are defined as the area-level means derived from the original census data. The estimation methods evaluated - EBPP, GMERF, and MERF - are applied as described in Section 5, using the same fixed working model for EBPP throughout the simulation.

Figs. 8 and 9, as well as Table 5, present the results in terms of bias and RMSE for the estimated area-level means in Guerrero. A direct comparison between GMERF and MERF reveals no substantial difference in RMSE, indicating similar performance levels for both methods across different sample domains. Notably, both GMERF and MERF demonstrate superior performance in terms of RMSE when compared to EBPP, particularly highlighting the strength of non-parametric approaches in this simulation context.

Table 6 evaluates the performance of the proposed MSE bootstrap procedures in terms of relative bias of RMSE and relative RMSE of RMSE. In terms of RRMSE-RMSE, the adjusted non-parametric bootstrap method for the MERF appears to be the most efficient method in this design-based simulation. Regardless of the bootstrap and estimation method used, the RB-RMSE for the in-sample areas indicates an acceptable bias regarding the median and mean RB-RMSE. For out-of-sample areas, we encounter underestimation in terms of the median value and overestimation according to mean values for all three methods.




Fig. 4. Model fit diagnostics for the Poisson GLMM: Histogram of Pearson residuals (left) and box-plots of Pearson residuals by municipalities (right). (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

Table 5
Mean and median for Bias and RMSE over total, in- and out-of-sample areas for point estimates.

Municipalities	Total		In-sample		Out-of-sample	
	Median	Mean	Median	Mean	Median	Mean
Bias						
EBPP	0.1257	0.0067	0.1549	0.1356	-0.1168	-0.1192
GMERF	-0.0096	0.0110	0.0367	0.0311	-0.0693	-0.0085
MERF	0.0349	0.0631	0.0686	0.0750	-0.0190	0.0515
RMSE						
EBPP	0.4174	0.4422	0.4563	0.4711	0.3510	0.4139
GMERF	0.2209	0.2747	0.1933	0.2462	0.2666	0.3918
MERF	0.2152	0.2670	0.2051	0.2411	0.2460	0.2923

Table 6
Mean and median for relative Bias and relative RMSE of the estimated RMSE over total, in- and out-of-sample areas for MSE estimates.

Municipalities	Total		In-sample		Out-of-sample	
	Median	Mean	Median	Mean	Median	Mean
RB-RMSE [%]						
GMERF P	-5.85	2.17	-2.95	0.27	-16.50	4.03
GMERF SP	-4.16	8.14	3.21	7.19	-12.10	9.07
MERF NPC	6.44	17.59	8.08	13.17	-4.38	21.90
RRMSE-RMSE [%]						
GMERF P	53.63	57.97	45.14	48.23	58.93	67.47
GMERF SP	51.07	58.38	46.89	50.61	55.79	65.95
MERF NPC	42.68	56.11	40.84	46.07	45.09	65.90




Fig. 5. Model fit diagnostics for the Poisson GLMM: Raw residuals vs. predicted values.




Fig. 6. Estimated school years of women for the state of Guerrero based on direct estimates, GMERF and MERF.

7. Conclusion

In this paper, we introduce generalized tree-based machine learning methods for estimating disaggregated, small area means for count data. Additionally, we explore the impact of overdispersion on the performance of these estimation methods. Our findings reveal that the MERF, which operates independently of distributional assumptions with respect to the count data, performs superior under conditions of severe overdispersion compared to the GMERF. Conversely, the GMERF shows better performance when the assumptions of the Poisson distribution are met or only moderately violated.

In Section 2, we extend the GMERF procedure (Frink and Schmid, 2024) to accommodate count data and discuss generalized semi-parametric mixed models at the unit level, treating Poisson GLMM-based SAE methods such as EBPP as special cases. Furthermore, the MERF approach (Krennmair and Schmid, 2022), initially developed for continuous data, is adapted for discrete outcome variables, exploring potential advantages and limitations, especially in relation to overdispersion. We present also three bootstrap methods for evaluating point estimators for area-level means: a parametric and a semi-parametric MSE bootstrap procedure for the GMERF, and an adjusted non-parametric MSE bootstrap procedure for the MERF (originally for continuous outcomes). The efficacy of point and MSE estimates is assessed through model-based simulations. Additionally, Section 5 applies these methods to data from the Mexican state of Guerrero.




Fig. 7. Area-specific CVs for Means for in- and out-of-sample areas.




Fig. 8. Bias of area-specific point estimates including details on in- and out-of-sample areas. Comparison of empirical bias from the design-based simulation for target variable *schoolfem*.




Fig. 9. Performance of area-specific point estimates including details on in- and out-of-sample areas. Comparison of empirical RMSEs from the design-based simulation for target variable *schoolfem*.

The simulations in Section 4 demonstrate that the GMERF and MERF point estimates surpass traditional methods when dealing with non-linear interactions between covariates at the linear predictor level. Notably, in scenarios of overdispersion, the MERF remains strong performance without dependence on Poisson distribution assumptions, outperforming generalized models. The MSE-bootstrap schemes proposed are validated as reliable based on their performance in our simulations. The practical application in Guerrero and further design-based simulations also support the effectiveness of these approaches.

From a methodological standpoint, the scope for further research into the application of GMERF in SAE is broad and promising. One potential direction is to adapt GMERF to model non-linear indicators, such as quantiles. This could involve applying the smearing technique introduced by Chambers and Dunstan (1986) and further developed for machine-learning methods by Krennmair et al. (2025). Furthermore, extending the GMERF to accommodate the Negative Binomial distribution by incorporating an additional dispersion parameter into the model could address issues of overdispersion more effectively than models based solely on the Poisson distribution. This modification would potentially allow GMERF to provide more accurate estimates for count data characterized by high variability, making it a more versatile tool in the field of SAE. These expansions not only enhance the flexibility and applicability of GMERF but also open up new possibilities for addressing complex problems in statistical estimation. Finally, further research is needed to explore the theoretical properties of the discussed bootstrap schemes for area-level means. Existing theoretical results for random forests, such as the work of Wager et al. (2014) and Wager and Athey (2018) on uncertainty quantification and the consistency of individual (unit-level) predictions, could provide a foundation for developing theoretical results for MERF- and GMERF-related bootstrap schemes for area-level means.

Acknowledgements

The authors are grateful to CONEVAL for providing the data used in empirical work. The views set out in this paper are those of the authors and do not reflect the official opinion of CONEVAL. The numerical results are not official estimates and are only produced for illustrating the methods. The authors are grateful for the computation time provided by the HPC service of the Freie Universität Berlin. Finally, the authors are indebted to the Co-Editor, Associate Editor and three referees for comments that significantly improved the paper.

Appendix A

A.1. Non-parametric bootstrap for MERF with count data

1. For given $\hat{f}()$, compute the marginal residuals $z_{ij} = y_{ij} - \hat{f}(\mathbf{x}_{ij})$.
2. Utilizing the marginal residuals z_{ij} , determine level-2 residuals for each area by

$$\bar{z}_i = \frac{1}{n_i} \sum_{j=1}^{n_i} z_{ij} \quad \text{for } i = 1, \dots, D.$$

3. Derive the vector of level-1 residuals using $\hat{z}_{ij} = z_{ij} - \bar{z}_i$. Following Krennmaier and Schmid (2022), the residuals are adjusted to the bias-corrected variance and centered, indicated by \hat{z}_{ij}^c . Level-2 residuals \bar{z}_i are additionally adjusted for the estimated variance $\hat{\sigma}_v^2$ and centered, indicated by \bar{z}^c .
4. For $b = 1, \dots, B$:
 - (a) Independently draw samples with replacement from the scaled and centered level-1 and level-2 residuals:

$$z_{ij}^{(b)} = \text{srswr}(\hat{z}_{ij}^c, N) \quad \text{and} \quad \bar{z}_i^{(b)} = \text{srswr}(\bar{z}^c, D).$$

- (b) Compute $\eta_{ij}^{(b)} = \hat{f}(\mathbf{x}_{ij}) + \bar{z}_i^{(b)} + z_{ij}^{(b)}$.

- (c) To obtain a count target variable, we match $\eta_{ij}^{(b)}$ with the set of estimated unit-level predictors from the sample $\{\hat{\eta}_t \mid \hat{\eta}_t = \hat{f}(\mathbf{x}_{it}) + \hat{v}_i, \text{ for } t = 1, \dots, n\}$ by finding the corresponding index \tilde{t} solving

$$\min_t \left| \eta_{ij}^{(b)} - \hat{\eta}_t \right|.$$

- (d) Define the bootstrap population by $y_{ij}^{(b)} = y_{\tilde{t}}$ for $j = 1, \dots, N$ and calculate the true bootstrap population mean $\mu_i^{(b)}$ for $i = 1, \dots, D$.
 - (e) For each bootstrap population, select a bootstrap sample that matches the original sample size n_i and estimate $\hat{f}^{(b)}()$ and $\hat{v}_i^{(b)}$. Obtain estimates for the mean $\hat{\mu}_i^{(b)}$.
5. Utilizing the B bootstrap samples, the MSE estimator is computed as follows:

$$\widehat{MSE}_i = \frac{1}{B} \sum_{b=1}^B \left(\mu_i^{(b)} - \hat{\mu}_i^{(b)} \right)^2.$$

Appendix B. Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.csda.2025.108218>.

References

Bates, D., Mächler, M., Bolker, B., Walker, S., 2015. Fitting linear mixed-effects models using lme4. *J. Stat. Softw.* 67 (1), 1–48.

Battese, G., Harter, R.M., Fuller, W.A., 1988. An error-components model for prediction of county crop areas using survey and satellite data. *J. Am. Stat. Assoc.* 83 (401), 28–36.

Biau, G., Scornet, E., 2016. A random forest guided tour. *Test* 25 (2), 197–227.

Boubeta, M., Lombardía, M.J., Morales, D., 2016. Empirical best prediction under area-level Poisson mixed models. *Test* 25 (3), 548–569.

Boubeta, M., Lombardía, M.J., Morales, D., 2017. Poisson mixed models for studying the poverty in small areas. *Comput. Stat. Data Anal.* 107, 32–47.

Breiman, L., 2001. Random forests. *Mach. Learn.* 45 (1), 5–32.

Breslow, N.E., Clayton, D.G., 1993. Approximate inference in generalized linear mixed models. *J. Am. Stat. Assoc.* 88 (421), 9–25.

Chambers, R., Chandra, H., 2013. A random effect block bootstrap for clustered data. *J. Comput. Graph. Stat.* 22 (2), 452–470.

Chambers, R., Dreassi, E., Salvati, N., 2014. Disease mapping via negative binomial regression M-quantiles. *Stat. Med.* 33 (27), 4805–4824.

Chambers, R., Dunstan, R., 1986. Estimating distribution functions from survey data. *Biometrika* 73 (3), 597–604.

Chandra, H., Salvati, N., Chambers, R., 2017. Small area prediction of counts under a non-stationary spatial model. *Spat. Stat.* 20, 30–56.

Chen, S., Jiang, J., Nguyen, T., 2015. Observed best prediction for small area counts. *J. Surv. Stat. Methodol.* 3 (2), 136–161.

Dreassi, E., Ranalli, M.G., Salvati, N., 2014. Semiparametric M-quantile regression for count data. *Stat. Methods Med. Res.* 23 (6), 591–610.

Flores-Agreda, D., Cantoni, E., 2019. Bootstrap estimation of uncertainty in prediction for generalized linear mixed models. *Comput. Stat. Data Anal.* 130, 1–17.

Frink, N., Schmid, T., 2024. Small area estimation with generalized random forests: estimating poverty rates in Mexico. Retrieved from: <https://arxiv.org/abs/2406.03861>.

Ghosh, M., Natarajan, K., Stroud, T.W.F., Carlin, B.P., 1998. Generalized linear models for small-area estimation. *J. Am. Stat. Assoc.* 93 (441), 273–282.

González-Manteiga, W., Lombardía, M., Molina, I., Morales, D., Santamaría, L., 2007. Estimation of the mean squared error of predictors of small area linear parameters under a logistic mixed model. *Comput. Stat. Data Anal.* 51 (5), 2720–2733.

González-Manteiga, W., Lombardía, M., Molina, I., Morales, D., Santamaría, L., 2008. Analytic and bootstrap approximations of prediction errors under a multivariate Fay–Herriot model. *Comput. Stat. Data Anal.* 52 (12), 5242–5252.

Gourieroux, C., Monfort, A., Trognon, A., 1984. Pseudo maximum likelihood methods: applications to Poisson models. *Econometrica* 52 (3), 701–720.

Greenwell, B.M., 2017. pdp: an R package for constructing partial dependence plots. *R J.* 9 (1), 421–436.

Greenwell, B.M., Boehmke, B., Gray, B., 2020. Variable importance plots: an introduction to the vip package. *R J.* 12 (1), 343–366.

Hastie, T., Tibshirani, R., Friedman, J., 2009. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer Science & Business Media.

Hobza, T., Morales, D., 2016. Empirical best prediction under unit-level logit mixed models. *J. Off. Stat.* 32 (3), 661–692.

- Jiang, J., 2003. Empirical best prediction for small-area inference based on generalized linear mixed models. *J. Stat. Plan. Inference* 111 (1), 117–127.
- Jiang, J., Lahiri, P., 2006. Mixed model prediction and small area estimation. *Test* 15 (1), 1–96.
- Jiang, J., Rao, J.S., 2020. Robust small area estimation: an overview. *Annu. Rev. Stat. Appl.* 7 (1), 337–360.
- Krennmair, P., 2022. The R package SAEforest. CRAN.
- Krennmair, P., Schmid, T., 2022. Flexible domain prediction using mixed effects random forests. *J. R. Stat. Soc., Ser. C, Appl. Stat.* 71 (5), 1865–1894.
- Krennmair, P., Würz, N., Schmid, T., Tzavidis, N., 2025. Random forests and mixed effects random forests for small area estimation of general parameters: a poverty mapping case study in Mozambique. Working paper.
- Kreutzmann, A.-K., Pannier, S., Rojas-Perilla, N., Schmid, T., Templ, M., Tzavidis, N., 2019. The R package emdi for estimating and mapping regionally disaggregated indicators. *J. Stat. Softw.* 91 (7), 1–33.
- Moon, T.K., 1996. The expectation-maximization algorithm. *IEEE Signal Process. Mag.* 13 (6), 47–60.
- R Core Team 2023. R: a language and environment for statistical computing. (Computer software manual).
- Rao, J.N.K., Molina, I., 2015. Small Area Estimation. Wiley, Hoboken, New Jersey.
- Stroup, W.W., 2012. Generalized Linear Mixed Models: Modern Concepts, Methods and Applications. CRC Press.
- Tzavidis, N., Ranalli, M.G., Salvati, N., Dreassi, E., Chambers, R., 2015. Robust small area prediction for counts. *Stat. Methods Med. Res.* 24 (3), 373–395.
- Tzavidis, N., Zhang, L.-C., Luna, A., Schmid, T., Rojas-Perilla, N., 2018. From start to finish: a framework for the production of small area official statistics. *J. R. Stat. Soc., Ser. A, Stat. Soc.* 181 (4), 927–979.
- Varian, H.R., 2014. Big data: new tricks for econometrics. *J. Econ. Perspect.* 28 (2), 3–28.
- Ver Hoef, J.M., Boveng, P.L., 2007. Quasi-Poisson vs. negative binomial regression: how should we model overdispersed count data? *Ecology* 88 (11), 2766–2772.
- Wager, S., Athey, S., 2018. Estimation and inference of heterogeneous treatment effects using random forests. *J. Am. Stat. Assoc.* 113 (523), 1228–1242.
- Wager, S., Hastie, T., Efron, B., 2014. Confidence intervals for random forests: the jackknife and the infinitesimal jackknife. *J. Mach. Learn. Res.* 15 (1), 1625–1651.
- Wright, M.N., Ziegler, A., 2017. ranger: a fast implementation of random forests for highdimensional data in C++ and R. *J. Stat. Softw.* 77 (1), 1–17.