

Geostatistical and machine learning approaches for high-resolution mapping of vaccination coverage

C. Edson Utazi^{a,b,c,*}, Ortis Yankey^a, Somnath Chaudhuri^a, Iyanuloluwa D. Olowe^a, M. Carolina Danovaro-Holliday^d, Attila N. Lazar^a, Andrew J. Tatem^a

^a WorldPop, School of Geography and Environmental Science, University of Southampton, SO17 1BJ, United Kingdom

^b Southampton Statistical Sciences Research Institute, University of Southampton, SO17 1BJ, United Kingdom

^c Nnamdi Azikiwe University, PMB 5025, Awka, Nigeria

^d World Health Organization, Geneva 1211, Switzerland

ARTICLE INFO

Keywords:

Bayesian geostatistics
Machine learning
Vaccination coverage
Demographic and health surveys
Health and development indicators
INLA-SPDE

ABSTRACT

Recently, there has been a growing interest in the production of high-resolution maps of vaccination coverage. These maps have been useful for uncovering geographic inequities in coverage and improving targeting of interventions to reach marginalized populations. Different methodological approaches have been developed for producing these maps using mostly geolocated household survey data and geospatial covariate information. However, it remains unclear how much the predicted coverage maps produced by the various methods differ, and which methods yield more reliable estimates. Here, we explore the predictive performance of these methods and resulting implications for spatial prioritization to fill this gap. Using Nigeria Demographic and Health Survey as a case study, we generate 1 × 1 km and district level maps of indicators of vaccination coverage using geostatistical, machine learning (ML) and hybrid methods and evaluate predictive performance via cross-validation. Our results show similar predictive performance for five of the seven methods investigated, although two geostatistical approaches are the best performing methods. The worst-performing methods are two ML approaches. We find marked differences in spatial prioritization using these methods, which could potentially result in missing important underserved populations, although broad similarities exist. Our study can help guide map production for other health and development metrics.

1. Introduction

Immunization is a fundamental component of primary healthcare, playing a critical role in reducing vaccine preventable morbidity and mortality (Shattock et al., 2024). It is also essential for achieving the Sustainable Development Goals (United Nations, 2015) and other global health policy goals, such as the Immunization Agenda 2030 (World Health Organization, 2020) and Gavi Strategy 5.0 (Gavi The Vaccine Alliance, 2020). However, immunization and other health services remain out of reach for vulnerable and marginalized populations, including those in remote rural areas, urban slums, and conflict-affected and humanitarian settings (Chopra et al., 2020; UNICEF and the Bill and Melinda Gates Foundation, 2021; Wigley et al., 2022).

To design effective strategies to reach these underserved populations, there is a need for accurate, spatially detailed maps of vaccination coverage and other health and development indicators (HDIs)

such as maternal literacy, poverty, school attendance, malaria prevalence, malnutrition and skilled birth attendance (Bosco et al., 2017; Mosser et al., 2019; Weiss et al., 2019; Kinyoki et al., 2020; Sbarra et al., 2021). Such maps enable decision makers to identify geographic and other inequities in service coverage and utilization, thereby supporting more targeted and effective interventions - a key focus of the growing field of precision public health (Dowell et al., 2016). Moreover, by providing current, robust and actionable evidence base, high-resolution maps help bridge the data gap that exist in many low- and middle-income countries where health management information systems and other administrative data sources such as vital registration are often incomplete and unreliable (Scobie et al., 2020; Mwinnyaa et al., 2021).

Data for producing maps of vaccination coverage and other HDIs often come from nationally representative, geolocated household surveys such as the Demographic and Health Surveys, Multiple Indicator

* Corresponding author.

E-mail address: c.e.utazi@oton.ac.uk (C.E. Utazi).

<https://doi.org/10.1016/j.sste.2025.100744>

Received 24 December 2024; Received in revised form 7 August 2025; Accepted 18 August 2025

Available online 19 August 2025

1877-5845/© 2025 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Cluster Surveys and national vaccination coverage surveys. Due to their high operational costs, these surveys are generally designed to provide estimates at the provincial or first administrative level. As a result, classical survey analysis methods, such as direct weighted estimators (Rao, 2005), can only generate reliable estimates at this coarse spatial scale. However, accurate and timely estimates are most valuable at lower administrative levels, e.g., the district or second administrative level at which vaccination programs and other interventions are planned and implemented. This need, along with advances in geostatistical modelling techniques and computing power, has spurred the widespread use of geostatistical and machine learning (ML) approaches to produce gridded estimates of HDIs from survey data. These approaches leverage the direct and proximate relationships between HDI outcomes measured at survey cluster locations and geospatial covariates, along with spatial and spatiotemporal dependence, to model and predict the spatial distributions of HDIs for single or multiple timepoints. By producing estimates at the grid level, typically 1 km or 5 km resolution, these outputs are not constrained by changing political or administrative boundaries and can be flexibly aggregated to operationally relevant areas of interest. Moreover, when integrated with other geospatial datasets, e.g., high-resolution population maps (Tatem, 2017) and geolocated health facility data (Lim et al., 2008; Johns et al., 2022), precise estimates of at-risk or underserved populations can be produced. Research and survey programs such as WorldPop through its VaxPop project (Utazi et al., 2018b, 2019, 2021, 2022), the Institute for Health Metrics and Evaluation (IHME) (Mosser et al., 2019; Sbarra et al., 2021) and the DHS program (Janocha et al., 2021) now routinely produce and distribute maps of HDIs.

A range of geostatistical, ML and hybrid approaches have been employed to produce high-resolution maps of vaccination coverage and other HDIs. Prominent examples include geostatistical models (GEOS) (Bosco et al., 2017; Utazi et al., 2021, 2022; Alegana et al., 2024), generalized additive models (GAMs) (Takahashi et al., 2017; Kawakatsu et al., 2024), stacked generalization (STG) (Mosser et al., 2019; Sbarra et al., 2021), boosted regression trees (BRT) (Kawakatsu et al., 2024), random forests (Browne et al., 2021), least absolute shrinkage and selection operator (LASSO) regression and deep learning/artificial neural networks (ANN) (Bosco et al., 2017). Model-based geostatistics (Diggle et al., 1998) explicitly accounts for spatial autocorrelation and the (non) linear effects of covariates, and is often implemented in a Bayesian framework using the INLA-SPDE approach or MCMC techniques, with INLA-SPDE being more popular recently due to its computational efficiency. When non-linear (or smooth) functions of covariates are incorporated into a geostatistical model, the result is a semiparametric geostatistical model (SGEOS) (Wood, 2011; Wang et al., 2018), which eliminates the need for covariate data transformation. A key advantage of the Bayesian implementation of geostatistical models is the natural framework to account for uncertainty in both model predictions and input data. ML and hybrid approaches are particularly suitable for modelling complex nonlinear relationships and interactions in the data, though this often comes at the expense of interpretability. ML approaches can automatically identify relevant covariates/features in the data, unlike geostatistical modelling which may require a separate covariate selection process. While ML approaches rely only on covariates to make predictions and would be expected to perform well when these are highly informative, geostatistical and hybrid approaches additionally exploit residual spatial (and temporal) autocorrelation to improve predictive performance. In general, ML approaches are computationally less demanding, can handle large-scale and high-dimensional data better, and are sometimes less challenging to implement (e.g., GAM, LASSO and BRT) (James et al., 2013; Berrocal et al., 2020). However, some ML approaches such as BRT, ANN and LASSO do not produce uncertainty estimates, necessitating the use of supplementary techniques for uncertainty quantification (Veronesi and Schillaci, 2019; Berrocal et al., 2020).

Currently, little is known about the comparative predictive

performance of these ML and geostatistical approaches in the context of mapping vaccination coverage. There is a lack of substantial evidence on how much the predicted maps produced by these approaches differ and which approaches yield more reliable estimates for vaccination coverage mapping. This gap may be attributed to the technical complexity involved in implementing these models and, in some cases, insufficient emphasis on methodological rigour. As maps of vaccination coverage and other HDIs become increasingly popular, it is crucial to assess the strengths and limitations of these modelling approaches. The goal of this study is, therefore, to critically evaluate widely used approaches for mapping vaccination coverage and other HDIs in terms of their predictive accuracy and associated uncertainties. Specifically, we investigate four machine learning approaches (ANN, BRT, GAM and LASSO), two geostatistical models (GEOS and SGEOS) and one hybrid approach (STG). Our evaluation is based on a case study mapping the coverage of the first dose of the diphtheria-tetanus-pertussis (DTP1) and the first dose of the measles-containing vaccine (MCV1) vaccines using the 2018 Nigeria Demographic and Health Survey (NDHS) (National Population Commission - NPC and ICF, 2019).

2. Methodology

2.1. Data

2.1.1. Vaccination coverage data

Data on the coverage of DTP1 and MCV1 vaccines were obtained from the 2018 NDHS (National Population Commission - NPC and ICF, 2019) for children aged 12-23 months and 9-35 months, respectively. The NDHS was conducted between August and December 2018, utilizing a stratified, two-stage sampling design to produce estimates of indicators at the national, regional and state levels, as well as for urban and rural areas. Stratification was achieved by separating each of the 36 states and the Federal Capital Territory (FCT) into urban and rural areas. Samples were drawn from within each stratum in two stages: the first stage involved the selection of survey clusters (enumeration areas) from a national sampling frame using a probability proportional to size sampling scheme, while the second stage involved selecting households randomly from household lists within the selected clusters. Detailed information on the methods employed in the survey is published elsewhere (National Population Commission - NPC and ICF, 2019). The NDHS was selected for this study because of ease of data access and having been used extensively in previous work to map coverage (Dong and Wakefield, 2021; Aheto et al., 2023; Utazi et al., 2023; Kawakatsu et al., 2024).

The survey was implemented in a total of 1,389 clusters, with 11 of the originally selected 1,400 clusters excluded due to security concerns. In Borno State, only 11 of the 27 local government areas were included in the survey due for similar reasons. For both vaccines, we used information obtained from both home-based records and maternal/care-giver recall, following DHS guidance during data extraction (Croft et al., 2023). Hence, our analysis captures crude DTP1 and MCV1 coverage estimates (World Health Organization, 2018). At the cluster level, we aggregated individual-level data to produce numbers of children surveyed, numbers vaccinated and empirical proportions of children vaccinated as shown in Fig. 1.

2.1.2. Geospatial covariate and population data

To enhance the prediction of vaccination coverage using the approaches investigated, we obtained some geospatial covariate information — see supplementary Figs. S1 and S2 and supplementary Table 1. These covariates have been successfully used in previous work (Bosco et al., 2017; Utazi et al., 2019, 2022, 2023) to model and predict vaccination coverage and other HDIs. These comprise variables measuring a range of conditions in the study country which may have direct or proximate relationships with vaccination coverage. The covariates include measures of remoteness (travel time to the nearest health facility




Fig. 1. Cluster level estimates of proportions of children aged 12-23 months and 9-35 months who had received DTP1 and MCV1, respectively, obtained using the 2018 Nigeria Demographic and Health Survey.

and distance to cultivated areas), socioeconomic status (poverty index, household wealth, maternal education), health-related factors (ownership of health or vaccination card/document, skilled birth attendance, access to media and use of mobile phone/internet) and urbanicity or development (nightlight intensity and urban/rural areas).

The externally sourced geospatial covariates (supplementary Table 1) were processed and harmonized at 1×1 km resolution, at which we planned to produce grid level coverage estimates. To extract the values of the covariates for each cluster location, we used the approach described in [Utazi et al. \(2018b\)](#) and [Perez-Haydrich et al. \(2013\)](#), which accounts for the displacement of the clusters (this displacement often occurs within districts in DHS surveys). For the DHS-derived covariates, we first calculated their values at the cluster level using detailed definitions provided in supplementary Table 1 and then used the `krig()` function in the `fields` package in R ([Nychka et al., 2017](#)) to create corresponding 1×1 km interpolated surfaces, with the optimal range parameter set to the first quartile of the distances between the clusters (other distance quartiles yielded almost the same results). The kriging interpolation was carried out using the logit-transformed cluster level data in each case, due to its underlying Gaussian assumption, after which the estimates were back-transformed to the unit interval.

We checked for multicollinearity by examining the correlations between the covariates and by fitting non-spatial binomial regression models to estimate their variance inflation factors (VIFs). Furthermore, for one of the modelling approaches (equations (1) and (2)), we examined the distributions of the covariates and their relationships with vaccination coverage (on the empirical logit scale), following which we log- or logit-transformed some skewed covariates to improve their linear relationships with vaccination coverage. The plots of the covariates and their relationships with vaccination coverage are shown in supplementary Figs. S3 and S4. All 14 covariates were retained in our study, as their VIFs were less than 5.0 for both DTP1 and MCV1. This also facilitated using the application of ML approaches, which typically benefit from a richer set of covariates.

To aggregate the coverage estimates to the district and other administrative levels, we obtained 2018 gridded estimates of numbers of children aged under 5 years from WorldPop ([Tatem, 2017](#)), which we used as a proxy population layer for the age groups included in the study.

2.2. Geostatistical and machine learning modelling approaches

We considered seven modelling approaches to predict vaccination coverage at 1×1 km resolution, as indicated previously. In all analyses, we accounted for the complex sampling design of the NDHS, specifically

urban-rural stratification, by including an urban-rural covariate and, when using geostatistical modelling approaches, between-cluster variation ([Dong and Wakefield, 2021](#); [Gascoigne et al., 2025](#)). The modelling approaches are described in detail as follows and illustrated in Fig. 2.

2.2.1. Bayesian geostatistical regression model (GEOS)

The first model we considered is a Bayesian geostatistical model with a Binomial likelihood. Let $Y(s_i)$ denote the number of children vaccinated at survey location s_i ($i = 1, \dots, n$) and $m(s_i)$ the number of children sampled at the location. The first level of the model assumes that

$$Y(s_i) | p(s_i) \sim \text{Binomial}(m(s_i), p(s_i)), \quad (1)$$

where $p(s_i)$ ($0 \leq p(s_i) \leq 1$) is the true vaccination coverage at location s_i . We model $p(s_i)$ using the logistic regression model

$$\text{logit}(p(s_i)) = \beta_0 + \sum_{j=1}^p x_j(s_i) \beta_j + \omega(s_i) + \epsilon(s_i), \quad (2)$$

where β_0 is an intercept term, $x_1(s_i), \dots, x_p(s_i)$ are covariates associated with s_i (including an urban-rural covariate), $\beta_1, \dots, \beta_p$ are the corresponding regression coefficients, $\epsilon(s_i)$ is an independent and identically distributed (iid) Gaussian random effect with variance, σ_ϵ^2 , used to model non-spatial residual variation or between-cluster variation, and $\omega(s_i)$ is a Gaussian spatial random effect used to capture residual spatial correlation in the model. That is, $\omega = (\omega(s_1), \dots, \omega(s_n))^T \sim N(0, \Sigma_\omega)$, where Σ_ω is assumed to follow the Matérn covariance function ([Matérn, 1960](#)). For identifiability reasons, we set the smoothness parameter in Σ_ω to one, see [Lindgren et al. \(2011\)](#).

To complete the Bayesian model specification, we assigned a $N(0, 10^3 I)$ prior to the regression parameter, β , and a penalized complexity (PC) ([Simpson et al., 2017](#)) prior to σ_ϵ such that $p(\sigma_\epsilon > 3) = 0.01$. Similarly, following [Fuglstad et al. \(2019\)](#), we placed a joint PC prior on the covariance parameters of the spatial random effect, ω , such that $p(r < r_0) = 0.01$ and $p(\sigma > 3) = 0.01$, with r_0 chosen to be the 5% of the extent of the country in the north-south direction.

The model was fitted using the INLA-SPDE approach implemented in the R-INLA package ([Lindgren et al., 2015](#); [R Core Team, 2021](#)). Predictions at 1×1 km resolution were obtained using the fitted model by drawing samples from the posterior predictive distributions of $p(s_i)$ at the grid locations. Throughout, predictions at the administrative level were obtained as population-weighted averages taken over all the grid cells falling within each administrative area ([Utazi et al., 2022](#)).




Fig. 2. Methodological overview. A schematic illustrating the implementation of geostatistical and machine learning approaches investigated in the study.

2.2.2. Bayesian semiparametric geostatistical regression model (SGEOS)

This model extends the GEOS model in equations (1) and (2) through using smooth functions to account for the nonlinear effects of some covariates. The model assumes that the true vaccination coverage at location s_i , $p(s_i)$, can be expressed as

$$\text{logit}(p(s_i)) = \beta_0 + \sum_{j=1}^p x_j(s_i)\beta_j + \sum_{k=1}^q f_k(z_k(s_i)) + \omega(s_i) + \epsilon(s_i) \quad (3)$$

where β_0 is an intercept term, $x_1(s_i), \dots, x_p(s_i)$ are linear covariates with regression coefficients $\beta_1, \dots, \beta_p$, and $f_1(\cdot), \dots, f_q(\cdot)$ are smooth functions used to account for the non-linear effects of the covariates $z_1(s_i), \dots, z_q(s_i)$. Other terms in the model are as defined previously in equation (2). We specified a second-order random walk prior for $f(\cdot)$ such that

$$f(u_i|u_{i-1}, u_{i-2}) \sim N(2u_{i-1} - u_{i-2}, \sigma_u^2) \quad (4)$$

which is the Bayesian equivalent of a cubic smoothing spline (Wang et al., 2018). For identifiability, a sum-to-zero constraint was imposed on each of the smooth functions since the model includes an intercept term (Wang et al., 2018). Model (3) was also fitted in a Bayesian framework using the INLA-SPDE approach. We assumed the default non-informative R-INLA log-Gamma prior on $\log(\sigma_u^2)$, i. e., $\log(\sigma_u^2) \sim \text{log} - \text{Gamma}(1, 0.00005)$.

2.2.3. Generalized additive model (GAM)

Generalized additive models also provide a mechanism to account for non-linear relationships by adding non-linear functions of all continuous covariates whilst maintaining additivity (James et al., 2013). The model is given by

$$\text{logit}(p(s_i)) = \beta_0 + \beta_1 x_1(s_i) + \sum_{k=1}^q f_k(z_k(s_i)) + g(s_i), \quad (5)$$

where $x_1(s_i)$ denotes the urban-rural covariate and $f_1(\cdot), \dots, f_q(\cdot)$ are functions used to account for the non-linear effects of other covariates. For our analyses, we chose cubic smoothing splines for $f(\cdot)$, noting that other choices are also possible (James et al., 2013). The function $g(\cdot)$ is used to account for the effect of space in the model, for which we specified a two-dimensional smoother - an isotropic smooth of latitude and longitude on the sphere with a second-order penalty and number of basis functions set equal to 100 (Wahba, 1981). The model was fitted in a frequentist framework and implemented in R using the mgcv package (Wood and Wood, 2015). We note that by including non-linear functions of all continuous covariates, our implementation of model (5) differs from the SGEOS model where smooth functions are only applied to

non-linear relationships determined beforehand.

2.2.4. Boosted Regression Model/Trees (BRT)

Boosting is a tree-based ensemble method that models complex, non-linear relationships between an outcome variable and multiple predictor variables (James et al., 2013). The method is based on the generation of a collection of sequentially fitted regression trees that optimize the predictive value of the response variable based on local predictor values. The boosting algorithm proceeds by fitting a regression tree to the data using the outcome variable as the response in the first iteration. The fitted tree is then scaled by a shrinkage parameter and added to the fitted function (this is set equal to zero in the first iteration) to update the residuals. In subsequent iterations of the algorithm, the regression trees are fitted using the residuals as the response. The process continues until a desired number of iterations or trees have been fitted. The output from the boosted model for location s_i can be expressed as

$$\hat{g}(\tilde{p}(s_i)) = \sum_{b=1}^B \lambda \hat{g}^b(\tilde{p}(s_i)), \quad (6)$$

where, $\hat{g}(\cdot)$ denotes the final prediction from the model, $\hat{g}^b(\cdot)$ is the prediction from the b th component regression tree, λ is a shrinkage parameter and B is the number of trees/iterations. λ controls the rate at which the boosting learns and is usually chosen to be small. For our application, we set $\lambda = 0.01$ as recommended in James et al. (2013) and chose $B = 10,000$. Another important tuning parameter when fitting a boosting model is the number of splits in each tree or the interaction depth, which controls the complexity of the boosted ensemble. This is often set equal to the default value of 1. The BRT model was implemented in our study using the gbm package in R (Ridgeway and Ridgeway, 2004). Due to the unavailability of the binomial distribution in the gbm package, we elected to model the logit-transformed cluster level vaccination coverage $\tilde{p}(s_i)$ using a Gaussian distribution and then back-transformed all the predictions post model-fitting. We note that as in model (5), the set of covariates used in fitting the model included the longitude and latitude coordinates to account for spatial variation.

2.2.5. Least absolute shrinkage and selection operator (LASSO) regression

Lasso regression performs both variable selection and regularization and is particularly suitable for modelling contexts where a large or considerable number of covariates are available. The method implements automatic covariate selection through a penalty term (the L_1 penalty) included in its objective function, which uses a tuning or regularization parameter to control the amount of regularization, i.e., how much the regression coefficients are shrunk towards zero. The method

finds regression coefficients $\hat{\beta}_\alpha^L$ that minimize the objective function

$$\ln L(\beta|\text{data}) + \alpha \sum_{j=1}^p |\beta_j|, \quad (7)$$

where α is the regularization parameter and all other terms are as defined previously. The first term in (7) is the log-likelihood function which can be obtained from the binomial regression model in equations (1) and (2) when the spatial and non-spatial random effects are excluded. Sufficiently large values of α will force some regression coefficients to be equal to zero. In practice, α is chosen via a grid search using cross-validation techniques. As in the GAM approach, the covariate data considered in the analysis using (7) included the longitude and latitude coordinates of the data locations. The LASSO regression model was implemented in our work using the glmnet package in R (Friedman et al., 2021).

2.2.6. Stacked generalization using a Bayesian geostatistical model (STG)

In statistical learning, stacked generalisation or stacked regression is an ensemble method for combining predictions from multiple models, often referred to as child models. In the hybrid variant implemented in our work, the child models were different ML approaches, predictions from which were combined using a geostatistical model (Bhatt et al., 2017; Mosser et al., 2019; Sbarra et al., 2021). Through these child models, the STG approach accounts for complex, nonlinear relationships between the covariates and the outcome. The geostatistical modelling framework provides a mechanism to account for residual spatial autocorrelation. The STG approach was proposed/utilized in Bhatt et al. (2017) and has been used to model vaccination coverage and various HDIs (Mayala et al., 2019; Mosser et al., 2019; Sbarra et al., 2021).

Following Sbarra et al. (2021), we considered the following child models: GAM, BRT and LASSO regression. These child models were implemented as described previously but excluding the geographical coordinates of the data locations in the covariate data. To obtain final predictions for the outcome, the predictions from these child models were included as covariates in the geostatistical model:

$$\text{logit}(p(s_i)) = \beta_0 + \beta_1 x^{\text{GAM}}(s_i) + \beta_2 x^{\text{BRT}}(s_i) + \beta_3 x^{\text{LASSO}}(s_i) + \omega(s_i) + \epsilon(s_i), \quad (8)$$

where $\beta_0, \beta_1, \beta_2$ and β_3 are regression coefficients and other terms are as described previously in equation (2). As in Sbarra et al. (2021), a sum-to-one constraint was imposed on the regression coefficients corresponding to the child models, such that $\sum_{j=1}^3 \beta_j = 1$. This constraint helps to mitigate the effect of extreme predictions in the child models included in (8) (Bhatt et al., 2017). As is usually the case in stacked generalization, Bhatt et al. (2017) recommended the use of K-fold cross-validation predictions from the child models to calibrate the model (i.e., estimate the parameters) in (8), and then refitting the child models using the full data and using the predictions from these in (8) without refitting the model. We noted that using the cross-validation predictions from the child models in (8) compared to the full data predictions did not necessarily yield improvements in predictive performance in our analyses. The STG approach was implemented in our work using the INLA-SPDE approach and the inlabru package in R (Lindgren et al., 2024).

2.2.7. Artificial neural networks (ANN)

An artificial neural network (ANN) is a ML technique that mimics the functioning of the animal brain. An ANN model is particularly useful in modelling contexts where data are large and complex, with potential nonlinearities and interactions between the covariates. The network consists of layers of connected neurons that serve as data processing units, where each neuron applies a linear transformation to its inputs, followed by a non-linear activation function. For our work, we used a multilayer perceptron network (Park and Lek, 2016), which consists of

an input layer, multiple hidden layers and an output layer. The input layer receives the features from the data, processes and transmits these to the hidden layers which process the information further through interconnected neurons, while the output layer produces the final predictions. For a spatial location s with covariate vector $x(s) = (x_1(s), x_2(s), \dots, x_p(s))^T$, the predicted value from an ANN with a single hidden layer can be expressed as:

$$\begin{aligned} \text{Outputlayer : } \widehat{p}(s) &= b^3 + \sum_{j=1}^{L_2} w_j^3 z_j^2(s), \\ \text{Hiddenlayer : } z_l^2(s) &= f\left(b_l^2 + \sum_{j=1}^{L_1} w_{jl}^2 z_j^1(s)\right), \quad l = 1, \dots, L_2, \\ \text{Inputlayer : } z_l^1(s) &= f\left(b_l^1 + \sum_{j=1}^p w_{jl}^1 x_j(s)\right), \quad l = 1, \dots, L_1, \end{aligned} \quad (9)$$

where L_1 and L_2 are the numbers of neurons in the input and hidden layers, respectively, $f(\cdot)$ is the activation function, b_l^1, b_l^2, b^3 and w_{jl}^1, w_j^2, w_j^3 are bias and weight parameters estimated to minimize mean squared error in the training data. Furthermore, $z_l^1(s), z_l^2(s)$ and $\widehat{p}(s)$ are outputs from the layers as shown in equation (9).

Fitting an ANN requires tuning the number of hidden layers, the number of neurons in each layer, and choosing the activation function. Other parameters such as the number of epochs (the number of times the entire data is passed through the network during training), stopping metric, stopping tolerance and stopping rounds are also tuned during model fitting. These early stopping criteria help to avoid overfitting in the model. A common choice for the activation function $f(\cdot)$ is the rectified linear unit (relu), defined as $f(x) = \max\{0, x\}$. The model was fitted using the `h2o.deeplearning()` function in the H2O package in R (Fryda et al., 2024). Since the H2O package does not support the binomial distribution, we elected to model the logit-transformed cluster-level vaccination coverage, denoted by $\widehat{p}(s)$ in equation (9) using a Gaussian distribution and then back-transformed the predictions post model fitting. Based on a hold-out cross-validation exercise with an 80% training and 20% testing split, the final selected model had two hidden layers with 100 neurons each, with the number of epochs set to 100. The chosen stopping metric was the root mean square error (RMSE) while the stopping tolerance and rounds were set equal to 0.001 and 5, respectively. We checked the sensitivity of these choices by running several cases with different justifiable parameter values but obtained the same results each time.

2.3. Uncertainty estimation using delete-a-block jackknife cross-validation

To estimate the uncertainties associated with the ML approaches: BRT, LASSO and ANN, we employed a delete-a-block jackknife technique. This is a variant of the delete-1 jackknife (Wang and Yu, 2021) in which a block of observations is deleted at a time. The spatial blocks were formed by drawing observations at random from the observed data. These can also be formed using spatially contiguous observations, but this approach is more likely to affect the underlying spatial structure in the data and can potentially introduce some artificial patterns in the uncertainty estimates, depending on the sizes of the blocks. The choice of the block size was guided by the need to have as many iterations as computationally logical (relative to the number of observations in the data) whilst preserving the underlying spatial correlation in the data. Having many iterations ensures stability in the results (i.e., the uncertainty estimates) and also reduces the numbers of observations deleted at each iteration. We noted during test runs that block sizes of up to $b = 40$ observations produced variogram estimates that were very similar to those of the full data in our applications (supplementary Figs. S5 and S6). We also noted that there were no material differences in the

estimates obtained for numbers of replicates $r \geq 100$. We, therefore, set $r = 100$ in our work, corresponding to block sizes of n/r , where n is the number of observations or spatial locations in the data as defined previously in (1). For all three ML approaches, we obtained the jackknife estimates of the uncertainties (i.e., the standard deviations) associated with the grid level predictions as $\sqrt{\frac{r-1}{r} \sum_{j=1}^r (\hat{p}_j(s) - \hat{p}_{jack}(s))^2}$, where $\hat{p}_j(s)$ is the j th coverage estimate for grid location s and $\hat{p}_{jack}(s)$ is the jackknife estimate of the mean across all the replicates.

2.4. Model validation using k -fold cross-validation and variogram analysis

To evaluate the out-of-sample predictive performance of the modelling approaches, we conducted a k -fold cross-validation exercise, setting $k = 10$. For each indicator-method combination, we created the cross-validation folds in two ways: random folds and spatially stratified folds. For the random folds, the survey locations were assigned to each of the k folds in a random manner; whereas with the spatially stratified method, each fold comprised neighbouring cluster locations. We calculated the following measures of predictive performance: the correlation between observed (p_i) and predicted ($\hat{p}_i$) values, root mean square error ($RMSE = \sqrt{\sum_i (\hat{p}_i - p_i)^2 / m}$), mean absolute error ($MAE = \frac{1}{m} \sum_{i=1}^m |\hat{p}_i - p_i|$), average bias ($AVG_BIAS = \frac{1}{m} \sum_{i=1}^m (\hat{p}_i - p_i)$) and the continuous ranked probability score ($CRPS(F_i, p_i) = E_{F_i} |X_i - p_i| - \frac{1}{2} E_{F_i} |X_i - X_i^*|$) (Gneiting and Raftery, 2007), where $F_i(\cdot)$ is the cumulative distribution function corresponding to the predictive distribution of the i th cluster level estimate, and X_i and X_i^* are two independent random variables distributed according to $F_i(\cdot)$. With r posterior samples, the CRPS can be estimated as $CRPS(F_i, p_i) = \frac{1}{r} \sum_{j=1}^r |\hat{p}_i^{(j)} - p_i| - \frac{1}{2r^2} \sum_{j=1}^r \sum_{l=1}^r |\hat{p}_i^{(j)} - \hat{p}_i^{(l)}|$, which is then averaged over all the locations within each fold and over all the k folds. While the other metrics (also averaged over all the k folds) measure the accuracy of

the point predictions produced by the approaches, the CRPS measures the accuracy of both the point and uncertainty estimates as it utilizes the entire posterior predictive distribution to determine the discrepancies between the observations and the predictions. Also, the CRPS was only computed for the three Bayesian approaches (GEOS, SGEOS and STG) in our work as it requires the posterior distributions of the estimates. The closer the AVG_BIAS, MAE and RMSE estimates are to zero and the smaller the CRPS, the better the predictions. Correlation values closer to one indicate better predictive ability.

Additionally, to further examine the fits of the different methods, we checked their (standardized) in-sample residuals for spatial autocorrelation using variograms and the associated variogram envelopes, which were obtained by permutation, using the geoR package in R (Ribeiro Jr et al., 2024).

3. Results

3.1. In- and out-of-sample predictive performance using cross-validation and variogram analysis

With respect to the metrics used to evaluate the accuracy of the point estimates produced by the methods at the cluster level (correlation, RMSE, MAE and AVG_BIAS), we found that GEOS, SGEOS and, to a great extent, LASSO had the best out-of-sample predictive performance in most cases (Fig. 3 and supplementary Table 2). The values of these metrics for GAM and STG were also very close to those of the three best approaches, indicating only slightly worse predictive performance. In contrast, BRT and ANN generally had the worst predictive performance, which can be clearly seen when considering the AVG_BIAS and RMSE estimates in Fig. 3.

Among the three Bayesian approaches for which we computed the CRPS metric, we found that GEOS and SGEOS outperformed the STG method based on this metric, which is also consistent with the results obtained using the other metrics. All the methods had fairly similar predictive performance under the two types of cross-validation folds investigated (i.e., random and spatially stratified folds) according to all




Fig. 3. Model validation. Out-of-sample predictive performance of the geostatistical and machine learning approaches investigated based on a k -fold cross-validation exercise using cluster level data (see supplementary Table 2).

the metrics except the correlations which showed that nearly all the methods had better predictive performance under the random folds. These results indicate that the methods can reasonably predict not only random but also spatial blocks of missing values in unsampled areas. There was no evidence of improved predictive performance for MCV1 despite having relatively larger cluster level sample sizes than DTP1 (supplementary Fig. S7). This is likely due to the cluster level sample sizes for MCV1 not being large enough to induce noticeable improvements in predictive performance.

Furthermore, when examining the out-of-sample predictions in low coverage areas (i. e., areas with cluster level proportions $p(s_i) \leq 0.4$ - supplementary Figs. S8 and S9), we found that the prediction errors (RMSE, random folds) for ANN and BRT were consistently larger ($RMSE \geq 0.33$) than those of the other approaches ($0.24 \leq RMSE \leq 0.3$), although there was evidence of overestimation in all the approaches. For DTP1, the lowest prediction errors were obtained for the GEOS and SGEOS methods, whereas for MCV1, these were obtained for SGEOS,

GEOS and STG.

The variograms of the in-sample residuals for DTP1 and MCV1 shown in supplementary Figs. S10 and S11 indicate that of all seven approaches investigated, there was strong evidence of residual spatial autocorrelation in the ANN and BRT methods. The variograms for both methods closely resembled those of the outcome variables (i. e., the cluster level proportions of vaccinated children - supplementary Figs. S5 and S6). Also, the lack of evidence of spatial autocorrelation in the residuals is strongest for the geostatistical approaches - GEOS, SGEOS and STG.

3.2. 1×1 km estimates of vaccination coverage and associated uncertainties

The rationale for the differences observed in the out-of-sample predictive performance of the approaches is apparent when investigating the 1×1 km predicted maps of vaccination coverage and associated




Fig. 4. Predicted DTP1 coverage estimates. 1×1 km modelled estimates of DTP1 coverage produced through using different geostatistical and machine learning approaches shown using (top) maps and (bottom) density plots.

uncertainties produced through using these approaches. Fig. 4 (a) shows strong similarities between the predicted surfaces produced by GAM, LASSO, GEOS, SGEOS and STG. Broadly similar patterns demonstrating a north-south divide in coverage can also be seen in the predicted maps produced using ANN and BRT, but their estimates are closer to the extremes of the unit interval and smoother in the lower and higher coverage areas than those of the other approaches.

The over-smoothing of the coverage estimates by ANN and BRT relative to the other approaches is evident in the distributions of the grid level DTP1 predictions shown in Fig. 4 (b). All the methods produced bimodal distributions reflecting the characteristic spatial distribution of vaccination coverage in Nigeria (Utazi et al., 2018b; Dong and Wakefield, 2021; Utazi et al., 2022,2023). However, the grid level estimates produced by ANN and BRT are more peaked near zero and one than those produced by the other approaches, suggesting overestimation in

high coverage areas and underestimation in low coverage areas by both approaches. This also explains the higher AVG.BIAS and RMSE values for both approaches relative to other approaches. For MCV1, supplementary Figs. S12 (a-b) show similar patterns in the grid level estimates produced by all the approaches, with strong evidence of over-smoothing in low and high coverage areas by ANN and BRT relative to the other approaches.

The uncertainties associated with the predictions have broadly similar spatial patterns across the methods, with lower uncertainties in areas where coverage estimates are either close to the endpoints (an artefact of the binomial distribution) of the unit interval or where data locations are dense, and higher uncertainties in areas where the estimates are closer to 0.5 or where data locations are sparse (Fig. 5 (a) and (b)).

However, due to the relative over-smoothing by ANN and BRT, the




Fig. 5. Uncertainty estimates for DTP1. Estimates of the uncertainties associated with 1×1 km estimates of DTP1 coverage produced using different geostatistical and machine learning approaches shown using (a) maps and (b) density plots.

uncertainties associated with both approaches are much smaller than those of other approaches (Fig. 5 (b)) in areas of lower and higher coverage, even in comparison with LASSO for which we used the same jackknife approach to produce its uncertainty estimates. In areas with mid-level coverage estimates, the uncertainties associated with the estimates produced by BRT are noisier and relatively much higher than other approaches. For MCV1 (supplementary Figs. S13 (a-b)), similar patterns can be observed, with the uncertainties associated with both ANN and BRT being much higher in many areas relative to the other approaches.

At the national level, the estimates produced through using these approaches revealed that ANN (and BRT to some extent; and GEOS – MCV1 only) overestimated coverage for both DTP1 and MCV1 relative to the direct survey estimate that is often considered to be the gold standard (Fig. 6a-b). On the other hand, whilst there are strong correlations between the grid level estimates produced by these approaches (Fig. 6c-d), it is evident that ANN and BRT are most dissimilar to other approaches, particularly for DTP1.

3.3. Exploring spatial prioritization using district level coverage estimates

To further investigate the utility of the coverage estimates produced

by the methods for spatial prioritization, we computed district level coverage estimates using their respective 1×1 km predicted maps and then ranked the districts based on these estimates. We note that the comparisons undertaken here using rankings obtained from the district-level coverage estimates are purely for illustration since estimates of numbers of unvaccinated children can characterise disease risk more accurately and are better suited for this purpose.

Fig. 7 (a-c) demonstrate that although there are broad similarities between the rankings of the district level DTP1 coverage estimates produced by the different methods, remarkable differences exist, both when examining groups of ranks (Fig. 7a) and, more evidently, the individual ranks (Fig. 7b). The differences between the rankings generally appear relatively smaller in areas of lower coverage in the northern parts of the country and much larger in higher coverage areas (Fig. 6c). Also, these differences appear more pronounced when considering smaller numbers of areas (e.g., the 80th to 100th lowest coverage areas) than larger numbers of areas (e.g., the 100 lowest coverage areas) (Fig. 7b). The median of the ranges of the ranks per district (Fig. 7c) at the national level is 112.5 (interquartile range (IQR) = 100, maximum value = 428), indicating marked differences among the methods. Among the five methods with similar predictive performance (i.e., GAM, LASSO, GEOS, SGEOS and STG), the median of the ranges of the ranks per district




Fig. 6. National and grid level coverage estimates. Comparisons between national (a-b) and grid level (c-d) coverage estimates produced through using geostatistical and machine learning approaches. In panels (a-b), the national estimates are compared with the direct survey estimates (dotted red lines) while panels (c-d) show the correlations between the grid level estimates.




Fig. 7. Spatial prioritization. (a-b) Ranking of districts from lowest (rank 1) to highest (rank 774) based on DTP1 coverage estimates produced using geostatistical and machine learning approaches, and (c) the range of the ranks per district. In panel (b), the districts have been ordered using the ranking produced by the GEOS method.

reduces to 83 (IQR = 89, maximum value = 337), which still indicates considerable differences. However, when examining pairs of methods with more similar predictive accuracy, there are large reductions in the differences between the rankings. For example, for the GEOS and SGEOS methods, the median of the ranges of the ranks per district is 17 (IQR=32).

Similar patterns were observed for MCV1 (supplementary Fig. S14), with the median of the ranges of the ranks per district estimated to be 141 (IQR=114, maximum value = 499) for all the methods at the national level, 87 (IQR=82, maximum value = 336) for GAM, LASSO, GEOS, SGEOS and STG, and 26 (IQR=50) for the GEOS and SGEOS methods. These differences in the rankings produced by the methods are also apparent in the bivariate plots of the ranks shown in supplementary Figs. S15 and S16.

4. Discussion

This study systematically evaluated the performance of seven geostatistical and ML approaches for producing high-resolution estimates of

vaccination coverage. All the methods, with the exception of SGEOS, were implemented using standard desktop computers, each requiring less than three hours (some ML methods completed much faster) to produce predictions at 1×1 km resolution. The SGEOS method, due to its computational demands, was run on a high-memory computer with a total runtime of ≈ 2.5 hours.

Our results revealed similar out-of-sample predictive performance at the cluster level for five of the methods - GEOS, SGEOS, LASSO, GAM and STG, although stronger predictive performance was observed for GEOS, SGEOS and LASSO methods. Among all seven approaches investigated, ANN and BRT had the poorest predictive performance. The relative over-smoothing observed in both approaches is likely due to the use of a different outcome distribution (Gaussian, instead of binomial), or how their algorithms learn from data. We further explored the impact of likelihood choice by evaluating the predictive performance of the remaining methods under a Gaussian likelihood (Supplementary Table 5). While ANN and BRT continued to underperform (see Supplementary Tables 2 and 5), the predictive performance of the other methods deteriorated under the Gaussian likelihood, reinforcing the

suitability of the binomial distribution for our application. Moreover, the poor performance of ANN and BRT was consistent regardless of the set of covariates used (e.g., kriged DHS covariates vs. other geospatial covariates), as over-smoothing remained evident in both cases (supplementary Fig. S17). Interestingly, among the Bayesian approaches, GEOS and SGEOS generally outperformed the hybrid STG method, despite the latter's widespread use in mapping HDIs (Bhatt et al., 2017; Mosser et al., 2019; Sbarra et al., 2021). These findings were further supported by in-sample assessments of residual spatial autocorrelation and comparisons of out-of-sample predictions in low coverage areas. We did not find evidence of better predictive performance for MCV1 due to larger cluster level sample sizes relative to DTP1. This may be because the increase in sample size was insufficient to yield measurable improvements in predictive accuracy. A more detailed examination of the effect of cluster-level sample size on predictive accuracy can be found in Utazi et al. (2022).

The 1×1 km predicted maps of the indicators revealed that GAM, LASSO, GEOS, SGEOS and STG produced very similar results, whereas ANN and BRT produced relatively over-smoothed estimates, with values clustering toward the extremes of the coverage scale. Although the uncertainty estimates produced by these approaches had very similar spatial distributions, the uncertainties from ANN and BRT were either relatively smaller (for DTP1) – an artefact of over-smoothing – or appeared relatively noisier and higher (for MCV1) in certain areas. Correlations between the grid level estimates produced by the different approaches were generally high (≥ 0.78), but these also indicated relatively lower correlations between ANN and BRT and other approaches, particularly for DTP1. Further comparisons with direct survey estimates at the national level revealed that ANN consistently over-estimated coverage, with some evidence of similar overestimation by BRT. These discrepancies suggest that ANN and BRT are also likely to yield other subnational (e.g., provincial-level) estimates that are not well aligned with direct survey estimates.

Considering the importance of district level estimates of vaccination coverage and corresponding estimates of numbers of zero-dose and under-vaccinated children for program planning and implementation, we further investigated the utility of the coverage estimates produced by the different approaches for spatial prioritization. We found remarkable differences in their rankings of the districts, although there were broad similarities especially when considering larger numbers of areas. The differences were most pronounced in areas of higher coverage and more modest in lower coverage areas, which might have been affected by the spatial distribution of vaccination coverage in the study country. We further observed a reduction in differences in rankings among methods with similar predictive performance, as expected, and even substantial reductions between pairs of methods with similar predictive performance. These results hold significant implications for vaccination programming, especially in resource-constrained settings where only a limited number of areas can be targeted per time, since inaccurate identification of priority areas for interventions could result in missing important vulnerable populations, suboptimal resource allocation, reduced impact and persistence of disease circulation or outbreaks. The predictive accuracy of these approaches should therefore guide their use for map production and operationalization.

Although our study is the first to systematically compare geostatistical, ML and hybrid approaches for vaccination coverage estimation, similar studies have been conducted in other application domains. For example, Berrocal et al. (2020) and Veronesi and Schillaci (2019) evaluated geostatistical and ML approaches for mapping air pollution and soil organic carbon, respectively, and found that geostatistical models outperformed ML methods – findings that align closely with our results. Zhu et al. (2024), in a related study on air pollution mapping, reported that random forests, a ML method, outperformed geostatistical models. Similarly, other studies utilizing kriging-based geostatistical approaches (e.g., Chen et al (2019) and Molla et al. (2023)) found that ML approaches performed better. In the context of mapping HDIs, our

results are somewhat different from those of Bosco et al. (2017), who compared ANN with a Bayesian geostatistical model across multiple countries and various HDIs. While they found similar predictive performance between the two approaches, they preferred the geostatistical model due to its ease of implementation (i.e., not requiring many tuning parameters). It is important to note that direct comparisons between our results and those from the literature are challenging due to substantial differences in study designs, geographic settings and applications, including choice of validation methods, amount of available data and parameter tuning. These contextual factors likely contribute to the mixed evidence regarding the relative performance of ML and geostatistical approaches, suggesting that model performance is, to some extent, context specific. We further assessed the generalizability of our findings to other settings by conducting an additional case study using data from the 2021 Cote d'Ivoire DHS (Institut National de la Statistique and ICF, 2023) - see supplementary materials for details. The results we obtained also showed that the other approaches investigated generally had better out-of-sample predictive performance than ANN and BRT. Also, among these other approaches, SGEOS and GEOS were the best performing methods, further corroborating our findings using the 2018 NDHS and reinforcing the robustness of our conclusions across different settings.

Our study has some limitations that should be acknowledged. As noted earlier, our implementation of the ANN and BRT approaches in the R programming language did not permit the use of a binomial likelihood for the outcome variables. Implementing these approaches using other programming languages such as Python may facilitate this. Although we found consistent results in our case study using the 2018 Nigeria DHS and the additional analyses using the 2021 Cote d'Ivoire DHS (supplementary materials), it will be useful to also explore how these methods perform in other geographical settings with additional antigens and potentially different sampling designs, degrees of spatial autocorrelation in vaccination coverage and numbers and types of covariates and their relationships with vaccination coverage (Bosco et al., 2017). It may be the case that the underperformance observed in ANN and BRT is dependent on some of these factors, but this also reveals a lack of robustness of both approaches to some modelling contexts or limitations in some of their current software implementations. For geostatistical models, these attributes have been investigated in detail in previous work using simulation studies (e.g., Utazi et al. (2018a)). However, we note that a simulation study would not be ideal when comparing geostatistical and ML approaches, as this would require simulating data from a geostatistical model or a sampling design based on geostatistical techniques, which would confer an undue advantage on these models over ML techniques. Furthermore, other approaches for estimating the uncertainties associated with the ML approaches are also possible. For example, a spatial bootstrap algorithm (this did not perform well in our study during initial trials) or an approach that involves interpolating spatial cross-validation residuals to create an uncertainty map, similar to Blanco et al (2018), could be used.

Whilst the use of geostatistical and ML approaches to produce high-resolution maps of HDIs has grown in popularity, other small area estimation methods for producing maps of HDIs exist (Tzavidis et al., 2018; Utazi et al., 2021; Paige et al., 2022), but these assume a discrete spatial domain, meaning that estimates can only be produced for a given administrative level at a time. Some of these methods are well explored in Utazi et al. (2021). Furthermore, in the ML arena, there are other hybrid approaches aiming to overcome the limitation of ML approaches not explicitly accounting for spatial autocorrelation in the data through (i) creating features that imitate the spatial autocorrelation in the outcome and using these as additional covariates in conventional ML methods (Sekulić et al., 2020; Fouedjio and Arya, 2024), (ii) combining ML predictions with the kriging of the prediction residuals (Kaya et al., 2022) and (iii) locally calibrated ML algorithms (Hagenauer and Helbich, 2022; Fouedjio and Arya, 2024). Future work in mapping vaccination coverage and other HDIs may involve the exploration of these

hybrid approaches. In geostatistical models, spatially varying coefficient models (Gelfand et al., 2003) could also be used to account for the spatial non-stationarity in the regression relationship between vaccination coverage and geospatial covariate information.

In conclusion, our results provide valuable guidance to practitioners regarding the utility of these modelling approaches for producing maps of vaccination coverage and other HDIs. While most of the approaches we investigated had good predictive accuracy and produced similar results, some approaches were relatively better, with significant implications for spatial prioritization. Effort should be made to either identify the best modelling framework for each analytical context or to use approaches that have been shown to be more robust and reliable in a similar setting.

Funding

CEU, ANL and AJT were supported by funding from Gavi, the Vaccine Alliance. CEU, SC and ANL were supported by funding from UNICEF (Grant number: 43387656).

Data and code availability

All the data used in the study are available from the sources referenced in the manuscript. The authors do not have the rights to redistribute these data. All R code used in the analyses is available via GitHub https://github.com/EdsonUtazi/GEOS_ML_paper.

Ethical approval

The study utilized anonymized secondary data. Ethics approval was provided by the University Ethics Committee (Application ID: 48522.A1), University of Southampton, UK.

CRediT authorship contribution statement

C. Edson Utazi: Writing – review & editing, Writing – original draft, Visualization, Supervision, Software, Methodology, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Ortis Yankey:** Writing – review & editing, Software, Methodology, Formal analysis, Data curation. **Somnath Chaudhuri:** Writing – review & editing, Writing – original draft, Visualization, Software, Methodology, Formal analysis, Data curation. **Iyanuloluwa D. Olowe:** Writing – review & editing, Data curation. **M. Carolina Danovaro-Holliday:** Writing – review & editing, Conceptualization. **Attila N. Lazar:** Writing – review & editing, Supervision, Funding acquisition. **Andrew J. Tatem:** Writing – review & editing, Supervision, Funding acquisition.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

C. Edson Utazi reports financial support was provided by UNICEF. Attila N. Lazar reports financial support was provided by UNICEF. M. Carolina Danovaro-Holliday reports a relationship with World Health Organization that includes: employment. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The authors would like to thank the VaxPop Project team for their support during the study.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at [doi:10.1016/j.sste.2025.100744](https://doi.org/10.1016/j.sste.2025.100744).

Data availability

The authors do not have permission to share data.

References

- Aheto, J.M.K., Olowe, I.D., Chan, H.M., Ekeh, A., Dieng, B., Fafunmi, B., Setayesh, H., Atuhaire, B., Crawford, J., Tatem, A.J., Utazi, C.E., 2023. Geospatial analyses of recent household surveys to assess changes in the distribution of zero-dose children and their associated factors before and during the covid-19 pandemic in Nigeria. *Vaccines* 11.
- Alegana, V.A., Ticha, J.M., Mwenda, J.M., Katsande, R., Gacic-Dobo, M., Danovaro-Holliday, M.C., Shey, C.W., Akpaka, K.A., Kazembe, L.N., Impouma, B., 2024. Modelling the spatial variability and uncertainty for under-vaccination and zero-dose children in fragile settings. *Sci. Rep.* 14, 24405.
- Berocal, V.J., Guan, Y., Muyskens, A., Wang, H., Reich, B.J., Mulholland, J.A., Chang, H. H., 2020. A comparison of statistical and machine learning methods for creating national daily maps of ambient pm2.5 concentration. *Atmos. Environ.* 222, 117130.
- Bhatt, S., Cameron, E., Flaxman, S.R., Weiss, D.J., Smith, D.L., Gething, P.W., 2017. Improved prediction accuracy for disease risk mapping using gaussian process stacked generalization. *J. R. Soc. Interface* 14, 20170520.
- Bosco, C., Alegana, V., Bird, T., Pezzulo, C., Bengtsson, L., Soricchetta, A., Steele, J., Hornby, G., Ruktanonchai, C., Ruktanonchai, N., Wetter, E., Tatem, A.J., 2017. Exploring the high-resolution mapping of gender-disaggregated development indicators. *J. R. Soc. Interface* 14, 20160825.
- Browne, C., Matteson, D.S., McBride, L., Hu, L., Liu, Y., Sun, Y., Wen, J., Barrett, C.B., 2021. Multivariate random forest prediction of poverty and malnutrition prevalence. *PLoS ONE* 16, e0255519.
- Chen, L., Ren, C., Li, L., Wang, Y., Zhang, B., Wang, Z., Li, L., 2019. A comparative assessment of geostatistical, machine learning, and hybrid approaches for mapping topsoil organic carbon content. *ISPRS Int. J. GeoInf.* 8, 174.
- Chopra, M., Bhutta, Z., Chang Blanc, D., Checchi, F., Gupta, A., et al., 2020. Addressing the persistent inequities in immunization coverage. *Bull. World Health Organ.* 98, 146–148.
- Croft, T.N., Allen, C.K., Zachary, B.W., 2023. Guide to DHS Statistics. ICF, Rockville, Maryland, USA.
- Diggle, P.J., Tawn, J.A., Moyeed, R.A., 1998. Model-based geostatistics. *J. R. Stat. Soc. Ser. C* 47, 299–350.
- Dong, T.Q., Wakefield, J., 2021. Modeling and presentation of vaccination coverage estimates using data from household surveys. *Vaccine* 39, 2584–2594.
- Dowell, S.F., Blazes, D., Desmond-Hellmann, S., 2016. Four steps to precision public health. *Nature* 540, 189–191.
- Fouedjio, F., Arya, E., 2024. Locally varying geostatistical machine learning for spatial prediction. *Artif. Intell. Geosci.* 5, 100081.
- Friedman, J., Hastie, T., Tibshirani, R., Narasimhan, B., Tay, K., Simon, N., Qian, J., 2021. Package ‘glmnet’. *CRAN R. Reposit.* 595.
- Fryda, T., Ledell, E., Gill, N., Aiello, S., Fu, A., Candel, A., Click, C., Kraljevic, T. & Nykodym, T. 2024. R package ‘h2o’: R interface for the ‘h2o’ scalable machine learning platform.
- Fuglstad, G.-A., Simpson, D., Lindgren, F., Rue, H., 2019. Constructing priors that penalize the complexity of gaussian random fields. *J. Am. Stat. Assoc.* 114, 445–452.
- Gascoigne, C., Smith, T., Paige, J., Wakefield, J., 2025. Estimating subnational under-five mortality rates using a spatio-temporal age-period-cohort model. *Spat. Spatiotemporal Epidemiol.* 52, 100708.
- Gavi the Vaccine Alliance, 2020. Gavi Strategy 5.0, 2021–2025. Available. <https://www.gavi.org/our-alliance/strategy/phase-5-2021-2025> (Accessed 25 June 2021).
- Gelfand, A.E., Kim, H.-J., Sirmans, C.F., Banerjee, S., 2003. Spatial modeling with spatially varying coefficient processes. *J. Am. Stat. Assoc.* 98, 387–396.
- Gneiting, T., Raftery, A.E., 2007. Strictly proper scoring rules, prediction, and estimation. *J. Am. Stat. Assoc.* 102, 359–378.
- Guio Blanco, C.M., Brito Gomez, V.M., Crespo, P., Ließ, M., 2018. Spatial prediction of soil water retention in a páramo landscape: methodological insight into machine learning using random forest. *Geoderma* 316, 100–114.
- Hagenauer, J., Helbich, M., 2022. A geographically weighted artificial neural network. *Int. J. Geograph. Inf. Sci.* 36, 215–235.
- Institut National De La Statistique & Icf, 2023. Côte d’Ivoire Enquête Démographique et de Santé 2021 Rapport Final. INS et ICF, Rockville, Maryland, USA et la Côte d’Ivoire.
- James, G., Witten, D., Hastie, T., Tibshirani, R., 2013. An Introduction to Statistical Learning: With Applications in R. Springer.
- Janocha, B., Donohue, R.E., Fish, T.D., Mayala, B.K., Croft, T.N., 2021. Guidance and Recommendations for the Use of Indicator Estimates at Subnational Administrative Level 2. ICF, Rockville, Maryland, USA. DHS Spatial Analysis Report 20.
- Johns, N.E., Hosseinpour, A.R., Chisema, M., Danovaro-Holliday, M.C., Kirkby, K., Schlottheuber, A., Shibeshi, M., Sodha, S.V., Zimba, B., 2022. Association between childhood immunisation coverage and proximity to health facilities in rural settings: a cross-sectional analysis of service provision assessment 2013–2014 facility data

- and demographic and health survey 2015–2016 individual data in Malawi. *BMJ Open* 12, e061346.
- Kawakatsu, Y., Mosser, J.F., Adolph, C., Baffoe, P., Cheshi, F., Aiga, H., Watkins, D.A., Sherr, K.H., 2024. High-resolution mapping of essential maternal and child health service coverage in Nigeria: a machine learning approach. *BMJ Open* 14, e080135.
- Kaya, F., Keshavarzi, A., Francaviglia, R., Kaplan, G., Başayigit, L., Dedeoğlu, M., 2022. Assessing machine learning-based prediction under different agricultural practices for digital mapping of soil organic carbon and available phosphorus. *Agriculture* 12, 1062.
- Kinyoki, D.K., Osgood-Zimmerman, A.E., Pickering, B.V., Schaeffer, L.E., Marczak, L.B., et al., 2020. Mapping child growth failure across low- and middle-income countries. *Nature* 577, 231–234.
- Lim, S.S., Stein, D.B., Charrow, A., Murray, C.J.L., 2008. Tracking progress towards universal childhood immunisation and the impact of global initiatives: a systematic analysis of three-dose diphtheria, tetanus, and pertussis immunisation coverage. *Lancet* 372, 2031–2046.
- Lindgren, F., Bachl, F., Illian, J., Suen, M. H., Rue, H. & Seaton, A. E. 2024. Inlabru: Software for fitting latent gaussian models with non-linear predictors. *arXiv preprint arXiv:2407.00791*.
- 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* 73, 423–498.
- Lindgren, F., Rue, H., Lindström, J., 2015. Bayesian spatial modelling with r-Inla. *J. Stat. Softw.* 63, 25.
- Matérn, B., 1960. Spatial variation. Springer-Verlag, Berlin, Germany.
- Mayala, B., Dontamsetti, T., Fish, T., Crof, T., 2019. Interpolation of DHS Survey Data at Subnational Administrative Level 2. ICF, Rockville. *Dhs Spatial Analysis Reports* no. 17.
- Molla, A., Zhang, W., Zuo, S., Ren, Y., Han, J., 2023. A machine learning and geostatistical hybrid method to improve spatial prediction accuracy of soil potentially toxic elements. *Stoch. Environ. Res. Risk Assess.* 37, 681–696.
- Mosser, J.F., Gagne-Maynard, W., Rao, P.C., Osgood-Zimmerman, A., Fullman, N., et al., 2019. Mapping diphtheria-pertussis-tetanus vaccine coverage in africa, 2000 - 2016: a spatial and temporal modelling study. *Lancet* 393, 1843–1855.
- Mwinnyaa, G., Hazel, E., Maiga, A., Amouzou, A., 2021. Estimating population-based coverage of reproductive, maternal, newborn, and child health (rmnch) interventions from health management information systems: a comprehensive review. *BMC Health Serv. Res.* 21, 1083.
- National Population Commission - Npc & Icf, 2019. Nigeria Demographic and Health Survey 2018 - Final Report. NPC and ICF, Abuja, Nigeria.
- Nychka, D., Furrer, R., Paige, J., Sain, S., 2017. Fields: Tools for Spatial Data. R package version, 9, D6W957CT.
- Paige, J., Fuglstad, G.-A., Riebler, A., Wakefield, J., 2022. Design- and model-based approaches to small-area estimation in a low- and middle-income country context: comparisons and recommendations. *J. Surv. Stat. Methodol.* 10, 50–80.
- Park, Y.S., Lek, S., 2016. Chapter 7 - artificial neural networks: multilayer perceptron for ecological modeling. In: Jørgensen, S.E. (Ed.), *Developments in Environmental Modelling*. Elsevier.
- Perez-Haydrich, C., Warren, J.L., Burgert, C.R., Emch, M.E., 2013. Guidelines on the Use of DHS GPS Data. ICF International, Calverton, Maryland, USA. *DHS Spatial Analysis Reports* No. 8.
- R Core Team, 2021. A Language and Environment for Statistical Computing. Vienna, Austria.
- Rao, J.N., 2005. Small Area Estimation. John Wiley & Sons.
- Jr, Ribeiro, J. P., Diggle, P.J., Christensen, O., Schlather, M., Bivand, R., Ripley, B., 2024. The GEOR Package: Analysis of Geostatistical Data.
- Ridgeway, G., Ridgeway, M.G., 2004. The GBM package. R Foundation for Statistical Computing, Vienna, Austria, p. 5.
- Sbarra, A.N., Rolfe, S., Nguyen, J.Q., Earl, L., Galles, N.C., et al., 2021. Mapping routine measles vaccination in low- and middle-income countries. *Nature* 589, 415–419.
- Scobie, H.M., Edelstein, M., Nicol, E., Morice, A., Rahimi, N., Macdonald, N.E., Carolina Danovaro-Holliday, M., Jawad, J., 2020. Improving the quality and use of immunization and surveillance data: summary report of the working group of the strategic advisory group of experts on immunization. *Vaccine* 38, 7183–7197.
- Sekulić, A., Kilibarda, M., Heuvelink, G.B., Nikolić, M., Bajat, B., 2020. Random forest spatial interpolation. *Remote Sens.* 12, 1687.
- Shattock, A.J., Johnson, H.C., Sim, S.Y., Carter, A., Lambach, P., et al., 2024. Contribution of vaccination to improved survival and health: modelling 50 years of the expanded programme on immunization. *Lancet* 403, 2307–2316.
- Simpson, D., Rue, H., Riebler, A., Martins, T.G., Sørbye, S.H., 2017. Penalising model component complexity: a principled, practical approach to constructing priors. *Stat. Sci.* 32, 1–28.
- Takahashi, S., Metcalf, C.J.E., Ferrari, M.J., Tatem, A.J., Lessler, J., 2017. The geography of measles vaccination in the African great Lakes region. *Nat. Commun.* 8, 15585.
- Tatem, A.J., 2017. Worldpop, open data for spatial demography. *Sci. Data* 4, 170004.
- 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* 181, 927–979.
- Unicef and the Bill and Melinda Gates Foundation, 2021. Equity Reference Group for Immunization Advocacy Brief. Available. https://drive.google.com/file/d/1VpuVX85RWd_vq6FJ4lcmCnPOYJp1AhuM/view (Accessed 5 May 2021).
- United Nations, 2015. Transforming Our World: The 2030 Agenda for Sustainable Development. Available. http://www.un.org/ga/search/view_doc.asp?symbol=A/RES/70/1&Lang=E (Accessed 20 June 2017).
- Utazi, C.E., Aheto, J.M.K., Chan, H.M.T., Tatem, A.J., Sahu, S.K., 2022. Conditional probability and ratio-based approaches for mapping the coverage of multi-dose vaccines. *Stat. Med.* 41, 5662–5678.
- Utazi, C.E., Aheto, J.M.K., Wigley, A., Tejedor-Garavito, N., Bonnie, A., Nnanatu, C.C., Wagai, J., Williams, C., Setayesh, H., Tatem, A.J., Cutts, F.T., 2023. Mapping the distribution of zero-dose children to assess the performance of vaccine delivery strategies and their relationships with measles incidence in Nigeria. *Vaccine* 41, 170–181.
- Utazi, C.E., Nilsen, K., Pannell, O., Dotse-Gborgborsi, W., Tatem, A.J., 2021. District-level estimation of vaccination coverage: discrete vs continuous spatial models. *Stat. Med.* 40, 2197–2211.
- Utazi, C.E., Thorley, J., Alegana, V.A., Ferrari, M.J., Nilsen, K., Takahashi, S., Metcalf, C.J.E., Lessler, J., Tatem, A.J., 2018a. A spatial regression model for the disaggregation of areal unit based data to high-resolution grids with application to vaccination coverage mapping. *Stat. Methods Med. Res.* 28, 3226–3241.
- Utazi, C.E., Thorley, J., Alegana, V.A., Ferrari, M.J., Takahashi, S., Metcalf, C.J.E., Lessler, J., Cutts, F.T., Tatem, A.J., 2019. Mapping vaccination coverage to explore the effects of delivery mechanisms and inform vaccination strategies. *Nat. Commun.* 10, 1633.
- Utazi, C.E., Thorley, J., Alegana, V.A., Ferrari, M.J., Takahashi, S., Metcalf, C.J.E., Lessler, J., Tatem, A.J., 2018b. High resolution age-structured mapping of childhood vaccination coverage in low and middle income countries. *Vaccine* 36, 1583–1591.
- Veronesi, F., Schillaci, C., 2019. Comparison between geostatistical and machine learning models as predictors of topsoil organic carbon with a focus on local uncertainty estimation. *Ecol. Indic.* 101, 1032–1044.
- Wahba, G., 1981. Spline interpolation and smoothing on the sphere. *SIAM J. Sci. Stat. Comput.* 2, 5–16.
- Wang, L., Yu, F., 2021. Jackknife resample method for precision estimation of weighted total least squares. *Commun. Stat. Simul. Comput.* 50, 1272–1289.
- Wang, X., Yue, Y.R., Faraway, J.J., 2018. Bayesian Regression Modeling with Inla. Chapman and Hall/CRC.
- Weiss, D.J., Lucas, T.C.D., Nguyen, M., Nandi, A.K., Bisanzio, D., et al., 2019. Mapping the global prevalence, incidence, and mortality of plasmodium falciparum, 2000–17: a spatial and temporal modelling study. *Lancet* 394, 322–331.
- Wigley, A., Lorin, J., Hogan, D., Utazi, C.E., Hagedorn, B., Dansereau, E., Tatem, A.J., Tejedor-Garavito, N., 2022. Estimates of the number and distribution of zero-dose and under-immunised children across remote-rural, urban, and conflict-affected settings in low and middle-income countries. *PLOS Glob. Public Health* 2, e0001126.
- Wood, S., Wood, M.S., 2015. Package 'mgcv', p. 729. R Package Version, 1.
- Wood, S.N., 2011. Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *J. R. Stat. Soc. Ser. B* 73, 3–36.
- World Health Organization, 2018. World Health Organization Vaccination Coverage Cluster Surveys: Reference Manual. Available. <https://apps.who.int/iris/handle/10665/272820>.
- World Health Organization, 2020. Immunization Agenda 2030: A Global Strategy to Leave No One Behind. Available. https://www.who.int/immunization/immunization_agenda_2030/en/ (Accessed 25 June 2020).
- Zhu, Q., Lee, D., Stoner, O., 2024. A comparison of statistical and machine learning models for spatio-temporal prediction of ambient air pollutant concentrations in scotland. *Environ. Ecol. Stat.* 31, 1085–1108.