



Household-level spatial modelling of child height-for-age in Northern Province, Rwanda: comparison of geographically weighted and neural network weighted regression models

Clarisse Kagoyire,^{1,2} Gilbert Nduwayezu,^{1,3} Rachid Oucheikh,¹ Jean Pierre Bizimana,⁴ Petter Pilesjö,¹ Ali Mansourian¹

¹Department of Earth and Environmental Sciences, Lund University, Lund, Sweden; ²Centre for GIS and Remote Sensing, College of Science and Technology, University of Rwanda, Kigali, Rwanda; ³Department of Civil, Environmental and Geomatics Engineering, College of Science and Technology, University of Rwanda, Kigali, Rwanda; ⁴Department of Spatial Planning, School of Architecture and Built Environment, College of Science and Technology, University of Rwanda, Kigali, Rwanda

Abstract

Childhood stunting remains a major public health concern in low- and middle-income countries, with persistent sub-national disparities. This study examined spatially varying associations between Height-for-Age z-scores (HAZ) and selected child, maternal, and household factors, and compared Geographically Weighted Regression (GWR), Multi-scale GWR (MGWR) and Geographically Neural Network Weighted Regression (GNNWR) in their ability to capture spatial heterogeneity and non-linearity. We analysed data from a cross-sectional survey conducted in December 2021 in the Northern Province of Rwanda. The survey covered 615 households, with analyses performed for 601 children aged 1–36 months. After imputation, multicollinearity screening, and feature selection, HAZ was modelled using Ordinary Least Squares (OLS), GWR, MGWR and GNNWR. Model performance was assessed using coefficient of determination (R^2), Root Mean Square Error (RMSE), Akaike Information Criterion/AIC Corrected (AIC/AICc), and Moran's I of residuals. Overall stunting prevalence was 27.1%. Spatially varying associations were observed for child age, sex, birth-weight, underweight status, selected childcare practices, maternal support, and household living conditions. GNNWR achieved the best training performance ($R^2 = 0.66$; RMSE=0.74) followed by MGWR ($R^2 = 0.51$; RMSE=0.88) and GWR ($R^2 = 0.43$; RMSE=0.95). Validation performance declined for all models, although GNNWR retained a higher validation ($R^2=0.28$) with negligible residual spatial autocorrelation (Moran's $I = -0.005$). GNNWR provided the strongest overall fit among the compared models and highlighted local spatial variation in associations with HAZ. However, the modest validation performance indicates that the findings should be interpreted cautiously. The workflow offers a portable framework for analysing DHS-like geocoded household datasets, but broader applicability should be assessed using spatial cross-validation, and sensitivity analyses.

Key words: child height-for-age, childhood stunting, spatial heterogeneity, geographically weighted regression, Rwanda.

Correspondence: Clarisse Kagoyire, Department of Earth and Environmental Sciences, Lund University, Sölvegatan 12, 223 62 Lund, Sweden.
Tel.: +46 769200501. E-mail: clarisse.kagoyire@mgeo.lu.se

Introduction

Childhood stunting remains a significant global public health challenge, particularly in low- and middle-income countries. In 2022, approximately 148.1 million of children under five years of age were stunted, with a global prevalence of 22.3% (Benjamin-Chung *et al.*, 2023; UNICEF *et al.*, 2023). Stunting, defined as a Height-for-Age z-score (HAZ) 2 Standard Deviations (SDs) below the World Health Organization (WHO) Child Growth Standards median, reflects impaired linear growth resulting from chronic or recurrent under-nutrition (WHO, 2006). Studies have shown that childhood stunting is closely associated with poor maternal and child health, low birth-weight, recurrent childhood illness, inadequate feeding and care practices, food insecurity and inadequate household water, sanitation and hygiene conditions (Black *et al.*, 2013; Vaivada *et al.*, 2020; Mertens *et al.*, 2023; UNICEF *et al.*, 2023). Its impacts extend beyond immediate health issues, with lasting effects on cognitive development, educational attainment,

economic productivity in adulthood and overall human capital accumulation (Black *et al.*, 2008; Victora *et al.*, 2008; Perkins *et al.*, 2017; Mertens *et al.*, 2023). Despite notable global progress in reducing stunting rates, substantial spatial disparities persist within individual countries, highlighting the need for context-specific and targeted strategies (Amegbor *et al.*, 2020; Osgood-Zimmerman *et al.*, 2018; Ssentongo *et al.*, 2021; Amir-ud-Din *et al.*, 2022; Uwiringiyimana *et al.*, 2022; Kuse & Debeke, 2023;).

These persistent, within-country disparities mean that national averages alone are insufficient for guiding effective interventions. Understanding where child linear growth faltering is concentrated and whether associated determinants vary geographically is therefore crucial for designing targeted interventions and allocating resources effectively (Osgood-Zimmerman *et al.*, 2018; Kinyoki *et al.*, 2020). District-level studies in India, for example, reveal significant clustering of childhood stunting despite national progress (Striessnig & Bora, 2020; Biswas, 2022). Similarly, geospatial analyses in Rwanda, Mozambique and Indonesia have identified

persistent sub-national hotspots of stunting, emphasising the limitations of one-size-fits-all interventions (Eryando *et al.*, 2022; Kalinda *et al.*, 2023; Tamir *et al.*, 2024). These findings highlight local spatial heterogeneity in child growth outcomes and associated factors, emphasising the need for spatially explicit analyses to inform geographically targeted public health planning.

To assess and respond to such spatial disparities, appropriate analytical methods are needed that can go beyond global average associations. Conventional global regression methods, such as Ordinary Least Squares (OLS), assume spatial stationarity and estimate a single association across the study area, potentially overlooking significant local variations in factors associated with child linear growth (Biswas, 2022; Seboka *et al.*, 2022; Yi *et al.*, 2024). Spatial econometric models, including Spatial Lag Model (SLM) and Spatial Error Model (SEM), incorporate spatial dependencies into regression analyses, addressing spatial autocorrelation and improving model accuracy (Puri *et al.*, 2020; Kuse & Debeko, 2023). Similarly, exploratory spatial statistics tools like Moran's *I* and Local Indicators of Spatial Association (LISA) are frequently used to detect spatial clustering and identify local hotspots (Khan & Mohanty, 2018; Hailu *et al.*, 2020; Puri *et al.*, 2020). Furthermore, Bayesian methods have been increasingly adapted within spatial regression frameworks, integrating prior knowledge and accommodating both spatial dependencies and non-linear relationships (Klein *et al.*, 2020; Moonga *et al.*, 2021; Muche *et al.*, 2021; Uwiringiyimana *et al.*, 2022; Ma & Huang, 2023). However, these approaches do not necessarily characterise how associations vary continuously across space, which motivates the use of geographically weighted methods in this study.

Geographically weighted models allow associations to vary across space rather than imposing a single global effect. Geographically Weighted Regression (GWR) estimates location-specific coefficients using a kernel-based weighting scheme, and Multi-Scale Geographically Weighted Regression (MGWR) extends this framework by allowing different predictors to operate at different spatial scales (Fotheringham *et al.*, 2017, 2003; Comber *et al.*, 2023). Despite these methodological advances, GWR-based models rely on predefined kernel functions, limiting their flexibility in modelling complex, non-linear relationships common in spatial health data (Luo *et al.*, 2021; Quiñones *et al.*, 2021). Moreover, many spatial health studies continue to rely on aggregated areal units, which are susceptible to the Modifiable Areal Unit Problem (MAUP) and potentially obscuring important local variations (Fotheringham & Wong, 1991; Haque *et al.*, 2025). Recently, hybrid approaches integrating spatial modelling and artificial neural networks have emerged to address these limitations (Du *et al.*, 2020). The Geographically Neural Network Weighted Regression (GNNWR) combines the spatial interpretability of GWR with the adaptive learning of neural networks, offering a flexible way to model both local spatial heterogeneity and non-linearity (Sun *et al.*, 2022; Du *et al.*, 2020; Siqi *et al.*, 2023; Zhao *et al.*, 2024). While initial applications of GNNWR have demonstrated potential across diverse fields such as environmental monitoring and urban studies (Sun *et al.*, 2022; Siqi *et al.*, 2023), its application in public health remains limited (Luo *et al.*, 2021; Quiñones *et al.*, 2021). Applying and validating this spatial deep learning method in the context of childhood stunting promises a substantial methodological contribution with global public health relevance, given the inherently complex spatial and non-linear dynamics often observed in its determinants.

Against this background, Northern Rwanda provides a relevant setting for examining sub-national disparities in child linear growth, given the persistently high burden of childhood stunting despite national progress (National Institute of Statistics of Rwanda, Rwanda, 2021; Uwiringiyimana *et al.*, 2022; Kalinda *et*

al., 2023). Using household-level geocoded survey data, this study had two objectives. Firstly, it examined spatially varying associations between child HAZ and selected maternal, child and household factors, with particular attention to childcare practices and maternal social support, which have received comparatively less attention in stunting analyses, particularly spatial analyses (Miconi *et al.*, 2018; Ruiz-Casares *et al.*, 2018; Utumatwishima *et al.*, 2025), while household living conditions were included as well-established factors associated with child linear growth (Black *et al.*, 2013; Mertens *et al.*, 2023). Secondly, it compared GWR, MGWR and GNNWR to assess how well these models represent local spatial heterogeneity and non-linearity in these associations. Within this broader set of factors, the primary analytic focus was on proximal childcare, maternal support and household-condition factors, while acknowledging environmental influences as important distal determinants beyond the primary analytic scope. The resulting analytical framework may also inform analyses of comparable DHS-like geocoded household datasets in other settings.

Materials and Methods

Study area

The study was conducted in the Northern Province of Rwanda, one of the country's four administrative provinces and the City of Kigali. Rwanda has an estimated population of 13.2 million, approximately 12.9% of whom are children aged 0 to 4 years (NISR, 2021). Over the past two decades, Rwanda has made significant progress in child health outcomes, reducing under-five child mortality and decreasing childhood stunting by approximately 15% (Amoroso *et al.*, 2018; NISR, 2021; Ekholuenetale *et al.*, 2023). Despite this national progress in child health, the Northern Province was specifically chosen for this study due to its persistently high prevalence of childhood stunting, estimated at approximately 41%, which is substantially above the national average of 33% (NISR, 2021). This province is administratively divided into five districts: Musanze, Burera, Gakenke, Rulindo and Gicumbi, which are further subdivided into 89 sectors, 414 cells and 2,744 villages. A population-based cross-sectional household survey was conducted across the five districts in December 2021 (Figure 1).

Design and sampling technique

The target population consisted of children aged from 1 to 36 months and their mothers or guardians, who were at least 18 years old and available to participate in the interview. A two-stage cluster random sampling approach was used to ensure both representativeness and spatial distribution across the province. In the first stage, villages were selected as primary sampling clusters using probability proportional to size; 137 villages corresponding to 5% of all villages in the province, were sampled. This ensured that villages with larger populations had proportionally greater representation. In the second stage, households were randomly selected within each village cluster. The community health workers provided village-specific lists of households with eligible children under three years of age, which served as the sampling frame at this stage. Households were then selected randomly within each village, proportionate to village population density. The minimum required sample size was calculated based on an expected stunting prevalence of 40% (NISR, 2021), with a 95% Confidence Interval (CI), a significance level of 0.05 (indicative of a 5% Type I error rate), and adjusted for a 10% non-response rate. A total of 615 households were visited, with structured interviews successfully conducted with the mothers or guardians of each eligible child.

Geographic coordinates were recorded for the sampled households and used in the spatial analyses. After exclusion of 14 observations with substantial missing data during data pre-processing, 601 records remained for analysis.

Variables

The outcome variable

The primary outcome variable was the Height-for-Age Z-score (HAZ), a standardised measure of linear growth relative to age- and sex-specific references from WHO growth standards (WHO,

2006). Lower, more negative HAZ values indicated poorer linear growth and greater growth faltering. Stunting, defined as a HAZ value below two SDs from the reference median (Benjamin-Chung *et al.*, 2023; WHO, 2006), was used for descriptive summaries and prevalence reporting, but the regression models used continuous HAZ as the dependent variable. Anthropometric measurements, including height/length, weight and age were collected by trained enumerators according to standard procedures, and HAZ scores were subsequently computed. Figure 2 presents the overall and district-level distribution of HAZ within the study population.




Figure 1. Map of the study area, the Northern Province of Rwanda.




Figure 2. District-level and overall distributions of height-for-age z-scores (HAZ). (A) Boxplots showing the distribution of children's HAZ by district (Rulindo, Gicumbi, Gakenke, Burera, and Musanze). (B) Histogram of the overall HAZ distribution across all districts. The vertical red dashed line marks the stunting threshold (HAZ = -2), while the grey dotted line indicates the overall mean HAZ.

Explanatory variables

The explanatory variables included in this study were selected based on conceptual relevance and empirical evidence on factors associated with child stunting (Black *et al.*, 2013; Benjamin-Chung *et al.*, 2023; Mertens *et al.*, 2023). Child health factors such as age in months, sex, birth weight, underweight status, multiple births and current breastfeeding practices were included due to their established influence on early childhood nutrition (Dusingizimana *et al.*, 2021; Quamme & Iversen, 2022; Mertens *et al.*, 2023). Childcare-related indicators included the number of days the child was left unattended for more than 1 hour within the past two weeks and the frequency of growth monitoring during the preceding six months. Leaving a child unattended was included as a proxy for potential care-giving constraints, which may affect feeding regularity and child supervision (Ruiz-Casares *et al.*, 2018; Daelmans *et al.*, 2021). Although childcare-related deprivation is recognised in maternal and child nutrition frameworks (WHO *et al.*, 2018; Daelmans *et al.*, 2021), direct indicators of inadequate supervision remain relatively limited in stunting analyses. The variable was therefore included to assess whether this underexplored aspect of care-giving was associated with HAZ in the present setting. Maternal social support was represented by mothers' access to reliable food assistance, and practical help in resolving personal problems. Household-level factors included the availability of hand-washing facilities as a proxy for hygiene standards within the household, access to electricity reflecting the household ability to access to basic amenities, the size of owned farmland and the

ownership of a home-garden (or kitchen-garden) which are indicative of the household's capacity for food production, and availability of milk for consumption in the household. Although environmental factors such as rainfall and elevation are commonly incorporated into spatial analyses of nutritional outcomes, they were not included in the primary analysis, consistent with the study's focus on proximal childcare, behavioural, maternal support and household-condition factors. The descriptions of the study variables are provided in *Supplementary Table S1*. Descriptive statistics for these variables, stratified by children's stunting status, are presented in *Supplementary Table S2*.

Analytical approach

Figure 3 outlines the methodological framework adopted in this study. The workflow begins with study design, spatial sampling and data pre-processing, including cleaning, multicollinearity screening by using the Variance Inflation Factor (VIF), imputation of missing values and feature selection. We then conduct exploratory spatial data analysis and estimate a global OLS model, with spatial autocorrelation assessed using Moran's *I*. Subsequently, three local spatial models: GWR, MGWR and the proposed GNNWR are fitted and evaluated. Model performance is compared using the coefficient of determination (R^2), adjusted R^2 , Akaike Information Criterion (AIC), corrected AIC (AICc), root mean square error (RMSE) and residual spatial autocorrelation diagnostics. The final outputs include maps of locally estimated coefficients, intended to support geographically targeted policy and interventions.




Figure 3. Methodological framework and analytical workflow.

Data pre-processing

The raw dataset included 615 observations and a wide range of potential explanatory variables covering child health, nutritional status, childcare practices, maternal health and support, maternal experiences of violence, household socioeconomic conditions as well as sanitation and agricultural factors. Rigorous data pre-processing ensured data quality, completeness, and analytical robustness. Observations with significant missing data (14 households, ~2%) were excluded, reducing the dataset to 601 observations. Explanatory variables with more than 3% missing data were removed, while remaining missing values were imputed using the MiceForest algorithm, which uses Light Gradient-Boosting Machines (LightGBM) and predictive mean matching for robust imputation (Azur *et al.*, 2011; Samuel, 2023). Several variables related to agriculture and livestock practices were excluded due to a high non-response rate of approximately 75%. For variables with more than two classes (k categories), k-1 dummy variables were created, with one category retained as the reference level to avoid multicollinearity. A multicollinearity assessment using VIF diagnostics led to the exclusion of variables with values greater than 10 (James *et al.*, 2023).

Feature selection was performed using three complementary methods: random forest regressor-based importance scores (Breiman, 2001; Archer & Kimes, 2008), ridge regression with L2 regularization (Wu, 2021) and forward selection regression. Variables that consistently ranked as important across these methods, or those strongly supported by theoretical frameworks, were included in the final analysis. The final set of explanatory variables was structured into coherent thematic categories, including child health and care factors, maternal characteristics and household socioeconomic factors; which provided a foundation for subsequent statistical modelling and interpretation.

Preliminary spatial data analysis

A preliminary spatial data analysis was performed to explore spatial dependency (autocorrelation) and identify initial relationships between HAZ and explanatory variables before doing a detailed spatial modelling. Initially, we generated descriptive statistics and mapped the spatial distribution of selected variables (Figure 4). For visualisation and to facilitate the interpretation of spatial patterns, both the distribution of explanatory variables and model-derived spatial outputs (such as local coefficient estimates and residuals) were aggregated within a uniform 1x1 km grid using a zonal statistics approach. The zonal statistics and the final maps were created using ArcGIS Pro 3.2. This level of spatial aggregation, performed specifically to enhance visual clarity, assumes relative homogeneity of variables within each grid cell, thereby preserving the spatial trends identified in the original point-based analyses. All regression models (GWR, MGWR, GNNWR) were fitted using unaggregated household-level point data recorded for sampled households; the grid-based aggregation was used exclusively for visualisation and interpretation of spatial patterns.

Spatial dependence was quantitatively assessed using the global Moran's *I* statistic, a common measure of spatial autocorrelation indicating whether the distribution of a variable is clustered, dispersed, or random across geographical locations (Anselin, 1995; Getis, 2007). The univariate Moran's *I* tests were applied separately to selected explanatory variables to identify their underlying spatial patterns. Variables with significant values indicate spatial clustering, while non-significant results suggest a random spatial pattern. OLS was used as initial global analytical step to estimate average associations between HAZ and the selected explanatory variables without considering spatial variations explicitly (Anselin *et al.*, 1996). The OLS model was expressed by:

$$y_i = \beta_0 + \sum_j^k \beta_j X_{ij} + \varepsilon_i \quad \text{Eq. 1}$$

where y_i is the outcome variable (HAZ) for observation i ; X_{ij} the j^{th} explanatory variable (predictor) for the i^{th} observation; β_0 the intercept; β_j the regression coefficient; and ε the error term associated with observation i

Negative coefficients indicate associations with lower HAZ and poorer linear growth, whereas positive coefficients indicate associations with higher HAZ and better linear growth. Following OLS estimation, we assessed the model residuals for spatial dependency using Moran's *I*. Positive autocorrelation indicates violations of OLS independence assumptions, motivating the use of spatially explicit methods. Lagrange multiplier tests further confirmed the presence and type (lag or error) of spatial autocorrelation, informing subsequent advanced spatial analyses. A SLM was subsequently estimated as an additional global benchmark before fitting the local spatial models. The dataset was randomly partitioned into training (70%) and validation (30) subsets for comparative model evaluation.

Local spatial heterogeneity models

To examine local variations in the associations between HAZ and the selected explanatory variables, geographically weighted regression techniques were used. Unlike global regression methods, geographically weighted approaches explicitly allow relationships to vary spatially, rather than imposing a single average association for the entire study area. All local models were fitted to unaggregated household-level point data. In all local models, negative coefficients indicate associations with lower HAZ and poorer linear growth, whereas positive coefficients indicate associations with higher HAZ and better linear growth.

Geographically weighted regression

GWR enables modelling of spatially varying relationships by estimating unique regression parameters at each observation location based on surrounding data points. Mathematically, the GWR model can be formulated as follows (Fotheringham *et al.*, 2003):

$$y_i = \beta_0(u_i v_i) + \sum_j^p \beta_j(u_i v_i) X_{ij} + \varepsilon_i \quad \text{Eq. 2}$$

where y_i is the outcome variable at observation location i ; $(u_i v_i)$ the geographical coordinates; $\beta_0(u_i v_i)$ the local intercept; $\beta_j(u_i v_i)$ the local coefficient for explanatory variable X_{ij} ; and ε_i the error term at location i .

Local coefficients are estimated through a spatial weighting matrix defined by a kernel function that assigns greater weights to observations closer to the focal location and lower weights to distant observations. In this study, both fixed and adaptive bi-square kernel specifications were evaluated. For each specification, the bandwidth was selected by minimising the AICc, allowing the neighbourhood size to be determined empirically from the data (Leung *et al.*, 2000). The better-fitting GWR specification was then retained for substantive interpretation in the results.

Multi-scale geographically weighted regression

While GWR uses a single bandwidth for all predictors or explanatory variables, MGWR extends this approach by allowing each predictor variable to have a distinct bandwidth (Fotheringham *et al.*, 2017). MGWR model is specified as:

$$y_i = \beta_0(u_i v_i) + \sum_j^p \beta_{bw_j}(u_i v_i) X_{ij} + \varepsilon_i \quad \text{Eq. 3}$$




Figure 4. Spatial distribution of selected explanatory variables in the Northern province of Rwanda. Geocoded household point data were aggregated using 1×1 km grid. For categorical variables, mapped values represent within-area proportions; for continuous variables, they represent area-level means. All regression models were fitted to household-level point data.

where bw_{mj} is the bandwidth explicitly optimised for each explanatory variable X_{ij} . In this study, MGWR optimally determines individual bandwidths for each predictor through an iterative back-fitting procedure guided by the AICc. This approach allows the model to distinguish between associations that vary over relatively broad spatial ranges and those that vary more locally, thereby providing a more flexible representation of multi-scale spatial heterogeneity in the outcome variable (HAZ).

Geographically neural network weighted regression

While GWR and MGWR effectively capture spatial heterogeneity through location-specific coefficients, their dependence on predefined kernel functions limits their ability to model complex, non-linear spatial relationships between predictors and the outcome variable. To overcome this limitation, we used GNNWR, a method that leverages the flexibility of artificial neural networks to learn spatially adaptive weights directly from the data, thus capturing complex, non-linear spatial variations (Du *et al.*, 2020). GNNWR improves upon conventional geographically weighted approaches by using a Spatially Weighted Neural Network (SWNN) to compute spatially varying regression coefficients. Particularly, the local regression coefficients in GNNWR are obtained by multiplying globally estimated OLS coefficients by neural-network-generated non-stationary spatial weights, as follows:

$$\hat{\beta}_{ij} = \mathbf{w}_{ij} \times \beta_j \quad \text{Eq. 4}$$

where $\hat{\beta}_{ij}$ is the spatially varying coefficient for the j_{th} explanatory variable at observation location i ; β_j the global coefficient estimate from the initial OLS regression (as shown in Eq.1); and $\mathbf{w}_{ij}$ the non-stationary spatial weight computed or learned adaptively by the neural network, varying according to each observation's spatial context. In contrast to GWR and MGWR, where spatial weights are explicitly defined through kernel functions, spatial weights $\mathbf{w}_{ij}$ in the GNNWR model are implicitly learned from the data using a neural network. The neural network captures spatial relationships based on spatial distances d_{ij} between observation locations. These spatial distances, representing the distance from observation location i to all other sample points, serves as inputs to the SWNN, processes these inputs through multiple hidden layers to produce output weights tailored to each explanatory variable at each observation location, as follows:

$$\mathbf{w}_i = \mathbf{W}(\mathbf{u}_i, \mathbf{v}_i) = \text{SWNN}([d_{i1}, d_{i2}, \dots, d_{in}]^T) \quad \text{Eq. 5}$$

where $[d_{i1}, d_{i2}, \dots, d_{in}]^T$ represent the spatial distances from the observations point and other sample observations. The GNNWR model is defined as:

$$\mathbf{y}_i = \mathbf{w}_0(\mathbf{u}_i, \mathbf{v}_i) \times \beta_{0OLS} + \sum_j \mathbf{w}_j(\mathbf{u}_i, \mathbf{v}_i) \times \beta_{jOLS} \mathbf{X}_{ij} + \epsilon_i \quad \text{Eq. 6}$$

where $\mathbf{w}_j(\mathbf{u}_i, \mathbf{v}_i)$ represents the non-stationary spatial weights of the OLS coefficients β_{jOLS} .

The SWNN architecture implemented in this study comprises four hidden layers, each using batch normalization, a parametric rectified linear unit (PReLU) activation, and drop-out techniques (He *et al.*, 2015; Ioffe & Szegedy, 2015). To further minimise overfitting, model hyperparameters (such as learning rate, drop-out rate, and network depth) were selected based on performance on a held-out validation set. Early stopping was considered to avoid excessive training (Prechelt, 1998). The complexity of the network systema-

tically decreases across hidden layers from an initial 256 nodes down to 15 nodes in the final output layer, matching the number of predictors. Training was conducted using the Stochastic Gradient Descent (SGD) optimiser with a learning rate of 0.01, drop-out rate of 0.2, and weight decay of 0.001 to ensure convergence and generalization to unseen data. The GNNWR approach, implemented here using open-source neural network library (Du *et al.*, 2020; Wu *et al.*, 2021), is adaptable to geocoded health datasets, allowing for the integration of country-specific or global health factors and supporting methodological transfer across contexts. Performance of the GWR, MGWR and GNNWR models was evaluated and compared using R^2 , adjusted R^2 , AIC, AICc and RMSE. Higher R^2 values indicate better explanatory capability, while lower AIC/AICc and RMSE values signify improved model fit. The spatial autocorrelation in the residuals from each model was further examined using global Moran's I statistics, with values close to zero showing minimal residual autocorrelation, which indicates minimal remaining spatial dependence after modelling. By comparing these geographically weighted models, the analysis aimed to characterise locally varying associations between the selected variables and child HAZ, with the broader goal of informing targeted public health interventions. The GWR and MGWR models were implemented using the *mgwr* Python package (Oshan *et al.*, 2019), while the GNNWR model was developed using the *gnnwr* package, a PyTorch-based implementation for spatially weighted neural network regression (Du *et al.*, 2020; Wu *et al.*, 2021).

Results

Initial findings from the exploratory data analysis

Among 601 children aged 136 months surveyed in the Northern Province of Rwanda, the overall prevalence of stunting (HAZ <2 SD) was 27.1%. Stunting rates varied by district, Musanze had the highest proportion of stunted children (31 of 96 children; 32%) and the lowest median HAZ (−1.63), whereas Rulindo had the lowest proportion of stunted children (22 of 101 children; 22%) and the highest median HAZ (−1.25). *Supplementary Table S2* shows clear unadjusted differences between stunted and non-stunted children. Compared with non-stunted children, stunted children were older on average (22 months compared with 16 months), had lower mean birth-weight (3,045 g compared with 3,247 g), and were more often underweight (19.6% compared with 2.1%). They also were more likely to live in households without electricity (74.8% compared with 57.8%) and without a hand-washing facility (95.7% compared with 85.8%). Maternal food support and help with personal problems were also less favourable among stunted children, although these differences were smaller. By contrast, days left unattended (>1 hour) and farmland size differed little between the groups, whereas growth monitoring was slightly more frequent among stunted children, possibly reflecting greater service contact among children already at higher nutritional risk. Considering that *Supplementary Table S2* presents unadjusted comparisons by binary stunting status, these findings should be interpreted as exploratory rather than as direct equivalents of the adjusted regression coefficients for continuous HAZ. Figure 4 shows the spatial distribution of selected explanatory variables, and *Supplementary Table S3* provides district-level summaries.

Global regression analyses

The OLS model explained approximately 34% of the variance in HAZ ($R^2 = 0.34$; F-statistic: 14.6, $p < 0.001$). Table 1 presents the




Figure 5. Aggregated local coefficient maps from the GWR model for selected predictors, Northern Province of Rwanda. Red tones indicate more negative coefficients, blue tones more positive coefficients, and pale tones coefficients closer to zero.

estimated OLS coefficients, standard errors and *p*-values considering that the outcome variable, HAZ, was analysed as a continuous measure of child linear growth rather than as a binary stunting indicator. Negative coefficients indicate associations with lower HAZ and poorer linear growth, whereas positive coefficients indicate associations with higher HAZ and better linear growth. Accordingly, lower HAZ reflects greater growth faltering and a higher likelihood of stunting. Variables negatively associated with HAZ included child age ($\beta = -0.152$, $p < 0.001$), being underweight ($\beta = -0.376$, $p < 0.001$), twin or triplet birth ($\beta = -0.103$, $p < 0.018$), lack of maternal emotional support ($\beta = -0.142$, $p < 0.004$), absence of electricity ($\beta = -0.125$, $p < 0.004$), inadequate sanitation facilities ($\beta = -0.113$, $p < 0.011$), absence of kitchen gardens ($\beta = -0.107$, $p < 0.017$), and ownership of farmland larger than one hectare ($\beta = -0.114$, $p < 0.010$). Positive associations with HAZ were observed for birth weight ($\beta = 0.175$, $p < 0.001$), no longer being breastfed ($\beta = 0.145$, $p < 0.013$), and leaving a child unattended ($\beta = 0.096$, $p < 0.037$). Growth monitoring, no food support, and no milk consumption were not statistically significant in the OLS model. Spatial autocorrelation in OLS residuals was detected, as indicated by significant Moran's *I* statistics, prompting the use of spatial regression models. SLM was also fitted, yielding a modest improvement in pseudo R^2 (0.35) compared to the OLS model.

Local spatial heterogeneity analysis

Geographically weighted regression (GWR) and multi-scale GWR (MGWR)

The GWR model used an adaptive and fixed bandwidth kernel functions, with the adaptive kernel achieving better training model fit ($R^2 = 0.43$; $RMSE = 0.95$), whereas the fixed-kernel GWR showed slightly better internal validation performance ($R^2 = 0.24$ compared to $R^2 = 0.19$). The local coefficient summaries reported in Table 2 correspond to the adaptive-bandwidth GWR and the MGWR models. MGWR improved on adaptive-bandwidth GWR in training performance ($R^2 = 0.51$; $RMSE = 0.88$) and identified predictor-specific spatial scales. In both models, mean local coefficients for age of the child, male sex, being underweight, twin or triplet birth, lack of maternal personal support, growth monitoring, no hand-washing facility, no access to electricity, absence of a kit-

chen garden, farmland size greater than 1 ha, and no milk consumption were negative, indicating associations with lower HAZ. By contrast, birth-weight, no longer being breastfed, and leaving a child unattended showed positive mean coefficients in both models. MGWR model identified predictor-specific spatial scales, with the greatest local variability observed for no longer being breastfed ($SD = 0.124$) and farmland size greater than 1 ha ($SD = 0.130$). Figure 5 and Figure 6 show the spatial distributions of GWR and MGWR local coefficients, respectively. In both figures, local coefficients are displayed using a diverging colour scale, with red-orange tones indicating more negative coefficients, blue tones indicating more positive coefficients, and pale tones indicating coefficients closer to zero.

Geographically neural network weighted regression (GNNWR)

Among the evaluated models, GNNWR demonstrated the highest explanatory performance on the training dataset, achieving an $R^2 = 0.66$, $AIC = 981$, $AICc = 1001$, and $RMSE = 0.74$. When evaluated on the validation dataset, the GNNWR model yielded an $R^2 = 0.28$ and $RMSE = 1.07$, indicating a modest out-of-sample performance than in the training data. Table 3 presents the summary statistics of the estimated GNNWR local coefficients for each predictor. For instance, the age of the child and being underweight were consistently negatively associated with HAZ across the study area (age: mean coefficient -0.585 , range -0.838 to -0.294 ; being underweight: mean coefficient -0.378 , range -0.533 to -0.291). Most other predictors showed spatially mixed associations, with local coefficients crossing zero. Mean coefficients remained positive for birth-weight (0.165), no longer being breastfed (0.122), child left unattended (0.121), and no food support (0.026), and negative for male sex (-0.078), twin or triplet birth (-0.083), growth monitoring (-0.091), lack of support with personal problems (-0.166), no hand-washing facility (-0.099), no access to electricity (-0.088), absence of a kitchen garden (-0.089), farmland size greater than 1 ha (-0.085), and no milk consumption (-0.053). Figure 7 presents maps of the spatial distribution of GNNWR coefficients for the selected predictors, where darker colours indicate areas with higher magnitude coefficients, revealing clusters of stronger positive associations.

Table 1. Results for Ordinary Least Squares (OLS) regression.

PREDICTOR	Coefficient	Standard error	p
Intercept	-1.378	0.043	0.000*
Age of the child	-0.549	0.060	0.000*
Child sex (male)	-0.152	0.044	0.000*
Birth-weight	0.175	0.045	0.000*
Underweight (yes)	-0.376	0.044	0.000*
Twin or triplet	-0.103	0.044	0.018*
No longer breastfed	0.145	0.058	0.013*
Child left alone more than 1 hour	0.096	0.046	0.037*
Number of times of growth monitoring	-0.064	0.045	0.153
No food support	0.071	0.049	0.141
No support with personal problems	-0.142	0.049	0.004*
No hand-washing facility	-0.113	0.044	0.011*
No access to electricity	-0.125	0.044	0.004*
Without a kitchen garden	-0.107	0.044	0.017*
Farmland > than 1Ha	-0.114	0.044	0.010*
No milk consumption	-0.085	0.046	0.063

* Statistically significant at the 0.05 level.




Figure 6. Aggregated local coefficient maps from the MGWR model for selected predictors, Northern Province of Rwanda. Red tones indicate more negative coefficients, blue tones more positive coefficients, and pale tones coefficients closer to zero.

Figure 8 shows the actual spatial distribution of aggregated observed HAZ, the predicted HAZ by the GNNWR model and the residuals of the GNNWR model across the study area. The residual spatial autocorrelation, as measured by global Moran's I , was minimal in the GNNWR model (Moran's $I = -0.0051$, $p = 0.89$), indicating that the spatial structure of the data was adequately accounted for by the model.

Comparative assessment of the local spatial models

Among the three spatial modelling approaches, the GNNWR achieved the best performance on the training set, explaining approximately 66% of the variability in HAZ values with the smallest AIC value ($R^2 = 0.66$, RMSE = 0.74, AIC = 981). In compari-

son, MGWR attained moderate accuracy ($R^2 = 0.51$, RMSE = 0.88, AIC = 1268), while GWR performed least effectively ($R^2 = 0.43$, RMSE = 0.95, AIC = 1288). This indicates that GNNWR captured more effectively non-linear and location-specific variations in the relationships between HAZ and selected predictors than the kernel-based models. Despite a relatively low validation with $R^2 = 0.28$, the GNNWR model demonstrated the strongest generalisation performance among the evaluated models. These results suggest that, while caution is warranted in interpreting the absolute level of explained variance, GNNWR represents a valid alternative to conventional spatial models, particularly in applications where modelling both spatial heterogeneity and non-linearity is critical. Table 4 summarises the accuracy metrics for all models. Notably,

Table 2. Summary statistics (mean and standard deviation) for parameter estimates from the investigated models, using the bi-square kernel with an adaptive weighting function and respective optimal bandwidths for the predictors.

PREDICTOR	Geographically Weighted regression (GWR)			Multi-scale Geographically Weighted regression (MGWR)		
	Mean coefficient	SD	Bandwidth	Mean coefficient	SD	Bandwidth
Intercept	-1.396	0.058	298	-1.389	0.153	95
Age of the child	-0.604	0.098	298	-0.618	0.078	298
Child sex (male)	-0.169	0.086	298	-0.167	0.094	254
Birth-weight	0.191	0.078	298	0.183	0.071	298
Underweight (yes)	-0.387	0.031	298	-0.395	0.034	298
Twin or triplet	-0.092	0.040	298	-0.093	0.039	298
No longer breastfed	0.195	0.099	298	0.213	0.124	129
Child left alone for more than 1 hour	0.104	0.044	298	0.106	0.049	298
Number of times of growth monitoring	-0.049	0.084	298	-0.068	0.069	298
No food support	0.062	0.048	298	0.054	0.035	298
No support with personal problems	-0.119	0.079	298	-0.115	0.089	231
No hand-washing facility	-0.074	0.050	298	-0.065	0.049	298
No access to electricity	-0.140	0.050	298	-0.151	0.055	298
Without a kitchen garden	-0.092	0.061	298	-0.088	0.057	298
Farmland > than 1Ha	-0.093	0.064	298	-0.089	0.130	141
No milk consumption	-0.075	0.077	298	-0.082	0.099	235

Table 3. Summary statistics for local parameter estimates from geographically neural network weighted regression (GNNWR) model.

PREDICTOR	Mean coefficient	Minimum	Maximum	SD
Age of the child	0.585	0.838	0.294	0.107
Child sex (male)	0.078	0.450	0.223	0.103
Birthweight	0.165	0.233	0.394	0.113
Underweight (yes)	0.378	0.533	0.291	0.048
Twin or triplet	0.083	0.272	0.104	0.052
No longer breastfed	0.122	0.028	0.255	0.067
Child left alone more than 1 hour	0.121	0.150	0.328	0.083
Number of times of growth monitoring	0.091	0.256	0.152	0.058
No food support	0.026	0.125	0.232	0.054
No support with personal problems	0.166	0.426	0.048	0.101
No hand-washing facility	0.099	0.277	0.128	0.084
No access to electricity	0.088	0.275	0.058	0.049
Without a kitchen garden	0.089	0.246	0.036	0.060
Farmland > 1Ha	0.085	0.363	0.148	0.086
No milk consumption	0.053	0.190	0.093	0.059

SD, standard deviation.




Figure 7. Aggregated local coefficient maps from the GNNWR model for selected predictors, Northern province of Rwanda. Red tones indicate more negative coefficients, blue tones more positive coefficients, and pale tones coefficients closer to zero.

MGWR improved upon standard GWR by assigning predictor-specific bandwidths, thereby better accounting for multi-scale spatial heterogeneity. Conversely, the OLS model, which lacks spatial weighting, exhibited the lowest accuracy (training $R^2 = 0.34$; validation $R^2 = 0.25$).

The residual spatial autocorrelation diagnostics (Global Moran's I statistics) indicated minimal residual spatial autocorrelation across all models, particularly for GNNWR, confirming effective handling of spatial dependencies (Table 5).

Figure 9 further illustrates observed and predicted HAZ values for each model (GWR, MGWR, GNNWR), highlighting both accuracy and limitations. Additional residual distribution maps and predicted HAZ maps for GWR and MGWR are available as supplementary material (Figure S2).

Discussion

This study examined spatially varying associations between child HAZ and selected maternal, child, household, and childcare factors in Northern Province of Rwanda, and compared GWR, MGWR and GNNWR in their ability to represent local spatial heterogeneity and non-linearity. Across the evaluated models, GNNWR provided the strongest fit on the training data and the smallest residual spatial autocorrelation, while MGWR also improved on standard GWR by allowing predictor-specific spatial scales. This pattern is methodologically valid because, unlike GWR and MGWR, GNNWR learns spatial weights through a neural-network framework rather than relying solely on predefined kernels, which may better accommodate local non-linearity alongside spatial heterogeneity (Du *et al.*, 2020; Quiñones *et al.*, 2021; S. Wu *et al.*, 2021). At the same time, the validation performance of GNNWR remained modest, indicating that the model should be interpreted as the strongest-performing approach within this comparison rather than as evidence of strong generalisability. Child growth faltering is a multi-factorial process, and associations with child HAZ may vary across places and at different levels of a predictor. Models that assume only linear relationships may therefore miss thresholds, plateaus, or stronger effects in particular local contexts (Du *et al.*, 2020; Mertens *et al.*, 2023). Our findings indicate that lower HAZ was consistently associated with older child age and underweight across

the study area, while male sex, twin or triplet birth, limited maternal personal support, and poorer household living conditions also tended, on average, to be associated with lower HAZ. In particular, lack of electricity, absence of a hand-washing facility, absence of a kitchen garden, farmland size greater than 1 ha, and no milk consumption were generally associated with poorer linear growth, although the magnitude of these associations varied spatially. These patterns are broadly consistent with previous spatial and multilevel analyses in Ethiopia and Rwanda that have linked child growth faltering to biological vulnerability, household deprivation, and inadequate living conditions (Hailu *et al.*, 2020; Uwiringiyimana *et al.*, 2022; Kuse & Debeke, 2023). The positive association between birth-weight and HAZ was likewise expected and is consistent with broader evidence that prenatal and early-life conditions shape subsequent linear growth trajectories (Benjamin-Chung *et al.*, 2023; Mertens *et al.*, 2023). A number of childcare factors showed more complex and spatially mixed associations. In particular, no longer being breastfed, frequency of growth monitoring, and days a child was left unattended more than 1 hour should not be interpreted as direct protective or risk effects on child HAZ, as these observed associations do not support causal inference. The positive average association for no longer being breastfed is likely influenced by child age and reverse causality, since prolonged breastfeeding may reflect an adaptive response to smaller size or poor growth rather than its cause (Atsbeha *et al.*, 2015; Kramer *et al.*, 2011; Marquis *et al.*, 1997). The association with being left unattended may likewise capture local childcare arrangements, age-related supervision patterns, or broader care-giving constraints rather than a simple deprivation gradient. This interpretation is consistent with the limited and still fragmented literature on child supervision in low- and middle-income settings (Miconi *et al.*, 2018; Ruiz-Casares *et al.*, 2018). The mixed local coefficients observed for several childcare and maternal support variables underscore the importance of interpreting behavioural indicators in context. A predictor may show different associations with HAZ in different localities because it is embedded in different social, service, or household environments. This is precisely where local spatial models are useful as they allow the analysis to move beyond average associations and to identify where a factor appears more strongly or more weakly associated with child growth. At the same time, such local patterns require careful interpretation, especially when reverse causality or residual

Table 4. Model performance diagnostics for the different models on training and validation datasets.

Models	Training					Validation	
	R^2	Adjusted R^2	AIC	AICc	RMSE	R^2	RMSE
OLS	0.34	0.32	1276	1279	1.02	0.25	1.08
GWR*	0.39	0.34	1281	1286	0.98	0.24	1.09
GWR**	0.43	0.36	1288	1303	0.95	0.19	1.13
MGWR*	0.51	0.41	1268	1299	0.88	-	-
GNNWR	0.66	0.64	981	1001	0.74	0.28	1.07

R^2 coefficient of determination; AIC, Akaike information criterion; AICc, corrected Akaike information criterion; RMSE, root mean square error; *Fixed bandwidth kernel;

**adaptive bandwidth kernel.

Table 5. Global Moran's I statistics for residuals of the investigated models indicating minimal residual spatial autocorrelation.

Models	Moran's I	z-score	p
GWR	0.0342	1.2121	0.23
MGWR	-0.0120	-0.3497	0.73
GNNWR	-0.0051	-0.1432	0.89

GWR, geographically weighted regression model; MGWR, multi-scale geographically weighted regression model; GNNWR, geographically neural network weighted regression.

confounding are plausible. For this reason, the coefficient maps should be viewed as tools for identifying context-specific priorities and hypotheses rather than as stand-alone causal evidence.

Implications for spatial health modelling and policy planning

The findings highlight the need to interpret behavioural indicators within their local context and policy environment. The association of “unattended status” with child HAZ requires further investigation, particularly to isolate potential contextual factors that may confound interpretation. For instance, Rwanda’s Early Childhood Development (ECD) policy has expanded access to ECD centres since its implementation, though coverage remains uneven; spatial models could help identify areas where combinations of nutritional, care-giving, and environmental disadvantages are most concentrated (Binagwaho *et al.*, 2016; McMahan & Sprague, 2024). Future research should explore how access to ECD services interacts with household childcare practices to shape nutritional outcomes, especially in areas where poor childcare shows stronger associations with HAZ. Moreover, the demonstrated performance of GNNWR highlights the potential of hybrid

spatial models for understanding child growth outcomes and guide targeted resource allocation. However, this study does not imply that the mapped coefficients can be translated directly into policy action. Rather, they provide a spatially informed starting point for programme targeting, local validation, and deeper investigation of contextual pathways.

Transferability and reproducibility

A practical contribution of this study is the demonstration of an open-source workflow that can be adapted to other geocoded household surveys, including DHS- or MICS-like datasets. The workflow combines data quality check and imputation; exploratory spatial data analysis, global baselines (OLS/SLM), kernel-weighted local models (GWR/MGWR), neural network weighted local modelling, and comparative model diagnostics using available Python tools. Portability rests on three conditions: i) availability of geocoded outcomes and covariates (explanatory variables) with sufficient spatial coverage, ii) minimal multicollinearity after screening, and careful harmonisation of variable coding and interpretation. Given modest validation performance here, we recommend spatial cross-validation and sensitivity analyses (bandwidths/kernels/architecture) and possibly multi-country replication.




Figure 8. Spatial distributions of aggregated observed HAZ, GNNWR-predicted HAZ and GNNWR residuals across the study area. Mapped as 1×1 km grid-based area-level summaries for visualisation, whereas the GNNWR model was fitted to household-level point data.




Figure 9. Comparison of observed versus predicted HAZ values for GWR, MGWR and GNNWR models.

Strengths and limitations of the study

A key strength of this study lies in its application of GNNWR, a novel hybrid spatial modelling approach that effectively captures both spatial heterogeneity and non-linear relationships, applied particularly to household-level health data. This model demonstrated improved predictive performance over conventional geographically weighted techniques with minimal residual autocorrelation; offering valuable insights into localised drivers of child growth outcomes through fine-scale local coefficients maps. The side-by-side comparison (OLS, SLM, GWR, MGWR, GNNWR) with residual spatial diagnostics strengthens inference that learned spatial weights improve fit and reduce ‘unmodeled’ spatial structure. Moreover, a globally portable workflow is proposed, using data elements mirrored to DHS/MICS-style household surveys, and implementation relies on open-source Python packages (mgwr, gnnwr), enabling researcher and practitioners in other localities to replicate and adapt the approach with comparable variables.

However, the study has also limitations. First, the analysis is based on cross-sectional data and therefore cannot establish causality. Second, although GNNWR showed the best overall fit on training data, the validation performance remained modest in the held-out subset. This likely reflects the inherently multidimensional nature of child growth faltering, which arises from interacting prenatal, biological, care-giving, household, and broader contextual determinants that cannot be fully captured in a single household survey or study; similar limits to explained variance have been reported in other child-growth studies. Third, the deliberate focus on proximal childcare, maternal support, and household-condition factors meant that broader environmental and service-accessibility influences were not included, even though they may shape child growth indirectly. Finally, although the sample size was adequate for the planned household survey, it remains moderate for highly flexible local models. This reflects a practical trade-off in geocoded survey research, where larger samples would strengthen local estimation and independent validation, but household survey design must also balance coverage, cost-effectiveness, and fieldwork feasibility. The modest validation performance observed here should therefore be interpreted partly as a consequence of applying a flexible local model to a complex child-growth outcome under realistic survey constraints. At the same time, the more flexible models still added value by revealing spatially varying and potentially non-linear associations that simpler global models may smooth over.

Future directions

Future research should address the limitations of the GNNWR model and expand its applicability in spatial health modelling, covering multi-country while leveraging DHS data. Enhancing predictive accuracy may require incorporating additional predictors, including environmental and service accessibility indicators. Exploring alternative neural architectures, such as graph neural networks or convolution neural networks, could improve the model’s ability to capture complex spatial dependencies while enhancing interpretability. To strengthen the model validation, adopting spatial cross-validation would better assess generalisation. Further investigation is also needed to clarify the role of childcare practices in shaping nutrition outcomes. Specifically, understanding how access to ECD services interacts with household factors could provide critical insights for policy and intervention program design. This is particularly relevant in areas where poor childcare practices showed stronger associations with HAZ, and where reverse causality may obscure the interpretation of behavioural indicators.

Conclusions

Among the compared models, GNNWR provided the strongest fit to the training data and the smallest residual spatial autocorrelation, indicating that accounting for both spatial heterogeneity and non-linearity can improve modelling of child linear growth. However, the modest validation performance and computational demands underscore the need for further refinement and validation across diverse settings (with larger and more diverse datasets). The results indicated that lower HAZ was most consistently associated with older child age and underweight status, while several household and support-related factors also showed predominantly adverse associations that varied across space. These spatial patterns may help inform geographically targeted nutrition, maternal-support, and household-environment interventions, but behavioural indicators such as breastfeeding status, growth monitoring, and child supervision require cautious interpretation because reverse causality and contextual confounding are plausible. More broadly, the workflow may offer a useful framework for analysing comparable DHS-like geocoded household datasets, provided that its broader applicability is further evaluated through spatial cross-validation, sensitivity analyses, and replication in additional settings.

References

- Amegbor PM, Zhang Z, Dalgaard R, Sabel CE, 2020. Multilevel and spatial analyses of childhood malnutrition in Uganda: Examining individual and contextual factors. *Sci Rep* 10:20019.
- Amir-ud-Din R, Fawad S, Naz L, Zafar S, Kumar R, Pongpanich S, 2022. Nutritional inequalities among under-five children: A geospatial analysis of hotspots and cold spots in 73 low- and middle-income countries. *Int J Equity Health* 21:135.
- Amoroso CL, Nisingizwe MP, Rouleau D, Thomson DR, Kagabo DM, Bucyana T, Drobac P, Ngabo F, 2018. Next wave of interventions to reduce under-five mortality in Rwanda: A cross-sectional analysis of demographic and health survey data. *BMC Pediatr* 18:27.
- Anselin L, 1995. Local Indicators of Spatial Association—LISA. *Geograph Anal* 27:93–115.
- Anselin L, Bera AK, Florax R, Yoon MJ, 1996. Simple diagnostic tests for spatial dependence. *Regional Sci Urban Econ* 26:77–104.
- Archer KJ, Kimes RV, 2008. Empirical characterization of random forest variable importance measures. *Computat Stat Data Analysis* 52:2249–60.
- Atsbeha DM, Nayga RM, Rickertsen K, 2015. Can prolonged breastfeeding duration impair child growth? Evidence from rural Ethiopia. *Food Policy* 53:46–53.
- Azur MJ, Stuart EA, Frangakis C, Leaf PJ, 2011. Multiple imputation by chained equations: What is it and how does it work? *Int J Methods Psychiatr Res* 20:40–9.
- Benjamin-Chung J, Mertens A, Colford JM Jr, Hubbard AE, van der Laan MJ, Coyle J, Sofrygin O, Cai W, Nguyen A, Pokpongkiat NN, Djajadi S, Seth A, Jilek W, Jung E, Chung EO, Rosete S, Hejazi N, Malenica I, Li H, Hafen R, Subramoney V, Häggström J, Norman T, Brown KH, Christian P, Arnold BF; Ki Child Growth Consortium, 2023. Early-childhood linear growth faltering in low- and middle-income countries. *Nature* 621:550–7.
- Binagwaho A, Scott KW, Harward SH, 2016. Early childhood development in Rwanda: A policy analysis of the human rights

- legal framework. *BMC Int Health Human Rights* 16:1.
- Biswas M, 2022. Identifying geographical heterogeneity in associations between under-five child nutritional status and its correlates across Indian districts. *Spatial Demography* 10:143–87.
- Black RE, Allen LH, Bhutta ZA, Caulfield LE, De Onis M, Ezzati M, Mathers C, Rivera J, 2008. Maternal and child undernutrition: Global and regional exposures and health consequences. *Lancet*, 371:243–60.
- Black RE, Victora CG, Walker SP, Bhutta ZA, Christian P, De Onis M, Ezzati M, Grantham-McGregor S, Katz J, Martorell R, Uauy R, 2013. Maternal and child undernutrition and overweight in low-income and middle-income countries. *Lancet* 382:427–51.
- Breiman L, 2001. Random Forests. *Machine Learning* 45:5–32.
- Comber A, Brunsdon C, Charlton M, Dong G, Harris R, Lu B, Lü Y, Murakami D, Nakaya T, Wang Y, Harris P, 2023. A route map for successful applications of geographically weighted regression. *Geograph Anal* 55:155–78.
- Daelmans B, Manji SA, Raina N, 2021. Nurturing care for early childhood development: global perspective and guidance. *Indian Pediatr* 58:11–15.
- Du Z, Wang Z, Wu S, Zhang F, Liu R, 2020. Geographically neural network weighted regression for the accurate estimation of spatial non-stationarity. *Int J Geograph Informat Sci* 34:1353–77.
- Dusingizimana T, Weber JL, Ramilan T, Iversen PO, Brough L, The role of care practices related to child feeding and health. *Br J Nutrition* 126:1203–14.
- Ekholuenetale M, Okonji OC, Nzopotam CI, Edet CK, Wegbom AI, Arora A, 2023. Socioeconomic disparities in Rwanda's under-5 population's growth tracking and nutrition promotion: Findings from the 2019–2020 demographic and health survey. *BMC Pediatrics* 23:467.
- Eryando T, Sipahutar T, Budhiharsana MP, Siregar KN, Nur Aidi M, Minarto D, Utari DM, Rahmaniati M, Hendarwan H, 2022. Spatial analysis of stunting determinants in 514 Indonesian districts/cities: Implications for intervention and setting of priority. *Geospat Health* 17:1055.
- Fotheringham AS, Brunsdon C, Charlton M, 2003. Geographically weighted regression: The analysis of spatially varying relationships. John Wiley & Sons.
- Fotheringham AS, Wong DWS, 1991. The modifiable areal unit problem in multivariate statistical analysis. *Environ Planning A: Economy Space* 23:1025–44.
- Fotheringham AS, Yang W, Kang W, 2017. Multiscale Geographically Weighted Regression (MGWR). *Ann Am Assoc Geographers* 107:1247–65.
- Getis A, 2007. Reflections on spatial autocorrelation. *Regional Science and Urban Economics* at 35. A Retrospective/Prospective Special Issue, 37:491–6.
- Hailu BA, Bogale GG, Beyene J, 2020. Spatial heterogeneity and factors influencing stunting and severe stunting among under-5 children in Ethiopia: Spatial and multilevel analysis. *Sci Rep* 10:16427.
- Haque S, Price A, Mengersen K, Hu W, 2025. Evaluating the impact of the Modifiable Areal Unit Problem on ecological model inference: A case study of COVID-19 data in Queensland, Australia. *Infect Dis Modelling* 10:1002–19.
- He K, Zhang X, Ren S, Sun J, 2015. delving deep into rectifiers: surpassing human-level performance on ImageNet Classification. 2015 IEEE International Conference on Computer Vision (ICCV) 1026–1034.
- Ioffe S, Szegedy C, 2015. Batch normalization: accelerating deep network training by reducing internal covariate shift. In F. Bach & D. Blei (Eds), *Proceedings of the 32nd International Conference on Machine Learning* 37:448–56.
- James G, Witten D, Hastie T, Tibshirani R, Taylor J, 2023. Linear regression. In *An introduction to statistical learning: With applications in python* (pp. 69–134). Springer.
- Kalinda C, Phiri M, Simona SJ, Banda A, Wong R, Qambayot MA, Ishimwe SMC, Amberbir A, Abebe B, Gebremariam A, Nyerere JO, 2023. Understanding factors associated with rural-urban disparities of stunting among under-five children in Rwanda: A decomposition analysis approach. *Mat Child Nutr* 19:e13511.
- Khan J, Mohanty SK, 2018. Spatial heterogeneity and correlates of child malnutrition in districts of India. *BMC Public Health* 18:1027.
- Kinyoki DK, Osgood-Zimmerman AE, Pickering BV, Schaeffer LE, Marczak LB, Lazzar-Atwood A, Collison ML, Henry NJ, Abebe Z, Adamu AA, Adekanmbi V, Ahmadi K, Ajumobi O, Al-Eyadhy A, Al-Raddadi RM, Alahdab F, Alijanzadeh M, Alipour V, Altirkawi K, ... Local Burden of Disease Child Growth Failure Collaborators. (2020). Mapping child growth failure across low- and middle-income countries. *Nature* 577:231–4.
- Klein N, Kneib T, Marra G, Radice R, 2020. Chapter 5—Bayesian mixed binary-continuous copula regression with an application to childhood undernutrition. In Y. Fan, D. Nott, M. S. Smith, & J.-L. Dortet-Bernadet (Eds), *Flexible Bayesian Regression Modelling Academic Press*. pp. 121–152.
- Kramer MS, Moodie EEM, Dahhou M, Platt RW, 2011. Breastfeeding and infant size: evidence of reverse causality. *Am J Epidemiol* 173(9), 978–983.
- Kuse KA, Debeke DD, 2023. Spatial distribution and determinants of stunting, wasting and underweight in children under-five in Ethiopia. *BMC Public Health* 23:641.
- Leung Y, Mei C-L, Zhang W-X, 2000. Statistical tests for spatial nonstationarity based on the geographically weighted regression Model. *Environ Planning A* 32:9–32.
- Luo Y, Yan J, McClure S, 2021. Distribution of the environmental and socioeconomic risk factors on COVID-19 death rate across continental USA: a spatial nonlinear analysis. *Environ Sci Pollut Res* 28:6587–99.
- Ma Z, Huang Z. 2023. A Bayesian implementation of the multiscale geographically weighted regression model with INLA. *Ann Am Assoc Geographers* 113:1501–15.
- Marquis GS, Habicht JP, Lanata CF, Black RE, Rasmussen KM, 1997. Association of breastfeeding and stunting in Peruvian toddlers: An example of reverse causality. *Int J Epidemiol* 26:349–56.
- McMahan LD, Sprague C, 2024. The varied perspectives of organisational effectiveness: What's at stake for early childhood development programmes in Rwanda? *Global Public Health* 19:237280.
- Mertens A, Benjamin-Chung J, Colford JM Jr, Coyle J, van der Laan MJ, Hubbard AE, Rosete S, Malenica I, Hejazi N, Sofrygin O, Cai W, Li H, Nguyen A, Pokpongkiat NN, Djajadi S, Seth A, Jung E, Chung EO, Jilek W, Subramoney V, Hafen R, Häggström J, Norman T, Brown KH, Christian P, Arnold BF; Ki Child Growth Consortium, 2023. Causes and consequences of child growth faltering in low-resource settings. *Nature* 621:568–76.
- Miconi D, Beeman I, Robert E, Beatson J, Ruiz-Casares M, 2018. Child supervision in low- and middle-income countries: A scoping review. *Children Youth Serv Rev* 89:226–42.
- Moonga G, Böse-O'Reilly S, Berger U, Harttgen K, Michelo C, Nowak D, Siebert U, Yabe J, Seiler J, 2021. Modelling chronic malnutrition in Zambia: A Bayesian distributional regression approach. *PLOS ONE* 16:e0255073.

- Muche A, Gezie LD, Baraki AG, Amsalu ET, 2021. Predictors of stunting among children age 6–59 months in Ethiopia using Bayesian multi-level analysis. *Sci Rep* 11:3759.
- NISR [Rwanda], Ministry of Health (MOH) [Rwanda], & ICF. (2021). Rwanda Demographic and Health Survey 2019-20 Final Report. Available from: <https://dhsprogram.com/pubs/pdf/FR370/FR370.pdf>
- Osgood-Zimmerman A, Millear AI, Stubbs RW, Shields C, Pickering BV, Earl L, Graetz N, Kinyoki DK, Ray SE, Bhatt S, Browne AJ, Burstein R, Cameron E, Casey DC, Deshpande A, Fullman N, Gething PW, Gibson H, Henry NJ, ... Hay SI, 2018. Mapping child growth failure in Africa between 2000 and 2015. *Nature* 555:41–47.
- Oshan TM, Li Z, Kang W, Wolf LJ, Fotheringham AS, 2019. mgwr: A Python Implementation of Multiscale Geographically Weighted Regression for Investigating Process Spatial Heterogeneity and Scale. *ISPRS Int J Geo-Information* 8:269.
- Perkins JM, Kim R, Krishna A, McGovern M, Aguayo VM, Subramanian SV, 2017. Understanding the association between stunting and child development in low- and middle-income countries: Next steps for research and intervention. *Soc Sci Med* 193:101–9.
- Prechelt L, 1998. Automatic early stopping using cross validation: Quantifying the criteria. *Neural Networks*, 11:761–7.
- Puri P, Khan J, Shil A, Ali M, 2020. A cross-sectional study on selected child health outcomes in India: Quantifying the spatial variations and identification of the parental risk factors. *Sci Rep* 10:6645.
- Quamme SH, Iversen PO, 2022. Prevalence of child stunting in Sub-Saharan Africa and its risk factors. *Clinical Nutrition Open Sci* 42:49–61.
- Quinones S, Goyal A, Ahmed ZU, 2021. Geographically weighted machine learning model for untangling spatial heterogeneity of type 2 diabetes mellitus (T2D) prevalence in the USA. *Sci Rep* 11:6955.
- Ruiz-Casares M, Nazif-Muñoz JI, Iwo R, Oulhote Y, 2018. Nonadult supervision of children in low- and middle-income countries: results from 61 national population-based surveys. *Int J Environ Res Public Health* 15:1564.
- Samuel VW, 2023. miceforest: Fast Imputation with Random Forests in Python. Available from: <https://github.com/AnotherSamWilson/miceforest/>
- Seboka BT, Hailegebreal S, Mamo TT, Yehualashet DE, Gilano G, Kabthamer RH, Ewune HA, Kassa R, Debisa MA, Yawo MN, Endashaw H, Demeke AD, Tesfa GA, 2022. Spatial trends and projections of chronic malnutrition among children under 5 years of age in Ethiopia from 2011 to 2019: A geographically weighted regression analysis. *J Health Population Nutr* 41:28.
- Siqi J, Yuhong W, Ling C, Xiaowen B, 2023. A novel approach to estimating urban land surface temperature by the combination of geographically weighted regression and deep neural network models. *Urban Climate* 47:101390.
- Ssentongo P, Ssentongo AE, Ba DM, Ericson JE, Na M, Gao X, Fronterre C, Chinchilli VM, Schiff SJ, 2021. Global, regional and national epidemiology and prevalence of child stunting, wasting and underweight in low-and middle-income countries, 2006–2018. *Sci Rep* 11:1–12.
- Striessnig E, Bora JK, 2020. Under-five child growth and nutrition status: spatial clustering of Indian districts. *Spatial Demography* 8:63–84.
- Sun Y, Jia W, Zhu W, Zhang X, Saidahemaiti S, Hu T, Guo H, 2022. Local neural-network-weighted models for occurrence and number of down wood in natural forest ecosystem. *Sci Rep* 12:6375.
- Tamir TT, Tekeba B, Mekonen EG, Zegeye AF, Gebrehana DA, 2024. Spatial heterogeneity and predictors of stunting among under five children in Mozambique: A geographically weighted regression. *Front Public Health* 12:1502018.
- UNICEF, WHO, & World Bank. 2023. Levels and trends in child malnutrition: Key findings of the 2023 Edition of the Joint Child Malnutrition Estimates. New York: UNICEF and WHO.
- Utumatwishima JN, Mogren I, Elfving K, Umubyeyi A, Krantz G, 2025. Association between poor mental health in mothers and child stunting: A population-based cross-sectional study in Rwanda. *BMJ Open* 15:e101117.
- Uwiringiyimana V, Osei F, Amer S, Veldkamp A, 2022. Bayesian geostatistical modelling of stunting in Rwanda: Risk factors and spatially explicit residual stunting burden. *BMC Public Health* 22:159.
- Vaivada T, Akseer N, Akseer S, Somaskandan A, Stefopoulos M, Bhutta ZA, 2020. Stunting in childhood: An overview of global burden, trends, determinants, and drivers of decline. *Am J Clin Nutrition* 112:777S–91S.
- Victora CG, Adair L, Fall C, Hallal PC, Martorell R, Richter L, Sachdev HS, 2008. Maternal and child undernutrition: Consequences for adult health and human capital. *Lancet* 371:340–57.
- WHO, 2006. WHO child growth standards: Length/height-for-age, weight-for-age, weight-for-length, weight-for-height and body mass index-for-age: Methods and development.
- WHO, UNICEF, & World Bank Group. (2018). Nurturing care for early childhood development: A framework for helping children survive and thrive to transform health and human potential.
- Wu S, Wang Z, Du Z, Huang B, Zhang F, Liu R, 2021. Geographically and temporally neural network weighted regression for modeling spatiotemporal non-stationary relationships. *Int J Geograph Informat Sci* 35:582–608.
- Wu Y, 2021. Can't ridge regression perform variable selection? *Technometrics* 63:263–71.
- Yi H, Li M, Dong Y, Gan Z, He L, Li X, Tao Y, Xia Z, Xia Z, Xue Y, Zhai Z, 2024. Nonlinear associations between the ratio of family income to poverty and all-cause mortality among adults in NHANES study. *Sci Rep* 14:12018.
- Zhao Z, Xu Z, Hu C, Wang K, Ding X, 2024. Geographically weighted neural network considering spatial heterogeneity for landslide susceptibility mapping: A case study of Yichang City, China. *CATENA*, 234:107590.

Online supplementary materials

a. Selection of sampling units (villages) and descriptive statistics of study variables

The map below illustrates the villages (clusters) selected in the districts of the Northern Province of Rwanda based on population density. The number of households visited and interviewed varied according to village size: 3–4 households for small villages, and 5–9 households for medium-sized and large villages.

Figure S1. Map of Selected Villages (Clusters) in Northern Province

Table S1. Description of outcome and explanatory variables used in the study

Table S2. Descriptive summary of selected explanatory variables.

Table S3. Descriptive summary statistics of study variables by district

b. Maps of predicted values and residual spatial patterns (from GWR and MGWR models)

The figure below provides additional spatial context to the main manuscript results (Section 3.4), illustrating how predictions and residual patterns differ between the GWR and MGWR models.

Figure S2. Spatial distributions of predicted HAZ scores and residuals for the GWR and MGWR models. (a) Predicted HAZ by GWR; (b) GWR residuals; (c) Predicted HAZ by MGWR; (d) MGWR residuals. Darker shades indicate higher residual values, highlighting spatial variability in model performance across the study area.

Received: 27 January 2026; Accepted: 29 March 2026.

Contributions: Clarisse Kagoyire and Ali Mansourian conceptualized and designed the study. Clarisse Kagoyire carried out the literature review, developed the methodology, formal analysis, writing and visualisation, under Ali Mansourian supervision. Ali Mansourian critically revised and edited the manuscript. Gilbert Nduwayezu and Rachid Oucheikh contributed to the methodology. Jean Pierre Bizimana and Petter Pilesjö critically revised the manuscript. All authors revised, read and approved the final manuscript.

Conflict of interest: the authors declare that they have no competing interests.

Ethics approval and consent to participate: ethical clearance for this study was obtained from the University of Rwanda College of Medicine and Health Sciences Institutional Review Board (reference No. 181/CHMS IRB/2021). All required permissions were also received from the National Institute of Statistics of Rwanda (reference No 0971/2020/10/NISR and No 0295/2021/10/NISR), the Ministry of Local Governance (reference No 0779/07.01), and the Ministry of Health (reference No 20/4184/DPMEHF/2021 and No NHRC/2020/PROT/047). All participants were fully informed about the study's aims, procedures, and measures to ensure confidentiality and privacy, in accordance with the Declaration of Helsinki. Written informed consent was obtained from all participants prior to their inclusion in the study. Interviews were conducted privately to safeguard participant confidentiality, and a rigorous data anonymization protocol was applied to ensure data could not be traced back to participants during analysis or in published results.

Availability of data: the datasets used and/or analysed during the current study are not publicly available due to the sensitive nature of the information they contain, which includes geo-located data on surveyed households, as well as detailed records on gender-based violence and mental health. These datasets contain personal and other highly sensitive information, necessitating restrictions to protect the privacy and confidentiality of the participants. Access to the data will be granted upon reasonable request to the corresponding author, subject to compliance with appropriate ethical standards and data sharing agreements.

Consent for publication: not applicable.

Funding: this study was funded by the Swedish International Development Cooperation Agency (SIDA): SIDA contribution No. 11277.

Acknowledgements: the authors sincerely acknowledge the invaluable contributions of all participating mothers and children who generously shared their time and discussed their personal and family conditions with us. We are also thankful to the enumerators for their essential support, as well as the community health workers and village leaders who guided us in reaching the sampled households during our data collection.

Publisher's note: all claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article or claim that may be made by its manufacturer is not guaranteed or endorsed by the publisher.

This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0).