CAUCHY – Jurnal Matematika Murni dan Aplikasi
Volume 11 (2) (2026), Pages 2704–2719
p-ISSN: 2086-0382; e-ISSN: 2477-3344

Spatial Classification Modelling of Imbalanced Response Using Firth Penalized Class-Weighted Geographically Weighted Logistic Regression

Zheryl Zabrina Hibatullah, Suci Astutik*, and Henny Pramoedyo

Department of Statistics, Faculty of Science, Technology, and Mathematics (FSTeM), Brawijaya University, Indonesia

Abstract

Binary regional classification of food security is often complicated by class imbalance and spatially heterogeneous predictor–outcome relationships. This study applies Firth Penalized Class-Weighted Geographically Weighted Logistic Regression (Firth CW-GWLR), which combines geographically weighted estimation, class weighting, and Firth's penalty on each local likelihood to reduce small-sample bias and protect against separation. The model is applied to the 2025 food security status of the 22 regencies/cities of East Nusa Tenggara (NTT), Indonesia (18 food insecure and 4 food secure), using the open unemployment rate (X_1), the share of agricultural workers (X_2), and the number of markets (X_3). VIF values (1.12–2.03) indicate no multicollinearity. The Breusch–Pagan test ($p = 0.0375$) and Moran's I of the global-model residuals ($I = 0.2025$, $p = 0.0248$) suggest that the global model is misspecified and motivate local modelling, but they do not by themselves show that the regression coefficients vary spatially. An adaptive Gaussian kernel with 21 nearest neighbours was selected (AICc = 21.170). GWLR, CW-GWLR, Firth-GWLR, and Firth CW-GWLR were compared by leave-one-out cross-validation (LOOCV), in which the class weights, bandwidth, and coefficients were re-estimated in each training fold; because the kernel type was selected on the full data, the LOOCV results are conditional on that choice. All local estimates converged. CW-GWLR achieved the highest balanced accuracy (0.764) and F_1 (0.545), Firth-GWLR the highest specificity (0.944), and Firth CW-GWLR identified three of the four food-secure regencies (recall 0.750) but classified five food-insecure regencies as food secure. The fully local model had a slightly lower LOOCV Brier score than models with one global coefficient, which is descriptive rather than formal evidence of spatial variation. With only four food-secure regencies, these results are indicative rather than conclusive.

Keywords: Firth's Penalized Likelihood; Class Imbalance; Geographically Weighted Logistic Regression; Spatial Heterogeneity; Food Security.

Copyright © 2026 by Authors, Published by CAUCHY Group. This is an open access article under the CC BY-SA License (https://creativecommons.org/licenses/by-sa/4.0)

1. Introduction

Food security is a multidimensional condition in which every individual has physical, social, and economic access to sufficient, safe, and nutritious food [1]. In Indonesia, it is monitored at the regency/city level through the Food Security Index (*Indeks Ketahanan Pangan*, IKP) of the National Food Agency [2].

*Corresponding author. E-mail: suci_sp@ub.ac.id

Received: August 7, 2026 Revised: October 30, 2026 Accepted: November 11, 2026
DOI: https://doi.org/10.18860/cauchy.v11i2.45533

The IKP is a composite index of 12 indicators grouped into three aspects. Availability covers the ratio of normative consumption to net carbohydrate production, energy and animal-protein availability, and regional rice reserves. Affordability covers the poverty rate, food price variability, and the prevalence of undernourishment. Utilization covers women's mean years of schooling, households without access to clean water, food safety, the Desirable Dietary Pattern score, and stunting prevalence. The indicators are normalized, weighted by expert judgement, and combined into a score on a 0–100 scale, which classifies each regency/city into six priority groups: Priority 1 (< 45.59, highly vulnerable), Priority 2 (45.60–53.42, vulnerable), Priority 3 (53.43–61.47, slightly vulnerable), Priority 4 (61.48–69.52, slightly secure), Priority 5 (69.53–77.35, secure), and Priority 6 (> 77.35, highly secure). Priorities 1–3 are designated as food insecure and Priorities 4–6 as food secure [2]. Treating this official dichotomy as a spatial classification problem raises two difficulties at once: class imbalance and spatial heterogeneity.

East Nusa Tenggara (*Nusa Tenggara Timur*, NTT) is frequently identified as food insecure [1, 3] and is the province with the largest number of food-insecure regencies/cities in Indonesia; only four of its 22 regencies/cities are food secure in the 2025 IKP (imbalance ratio $4.5 : 1$) [2]. Under such imbalance, ordinary logistic regression favours the majority class, yielding high accuracy but poor minority-class sensitivity [4, 5], and the problem worsens as imbalance increases [6, 7]. Weighted Logistic Regression (WLR) compensates by re-weighting observations by class proportions [8, 9], although such corrections can distort predicted probabilities [10], and it still assumes spatially constant effects. Geographically Weighted Logistic Regression (GWLR) [11] relaxes this assumption and has been applied to regional binary outcomes such as COVID-19 status [12], underdevelopment [13], human development [14], malaria [15], illiteracy [16], stunting [17], and food security in East Java [18]. In these applications, the local likelihoods are not weighted by class, so the majority-class dominance of the global likelihood carries over to each local likelihood.

Local likelihoods estimated from small, rare-event samples are also prone to quasi-complete separation, under which the maximum likelihood estimates diverge [19]. Firth's penalized likelihood [20, 21] reduces small-sample bias and yields finite estimates under separation [22], but its behaviour when combined with both spatial and class weights in GWLR has not been examined for regional food security data.

This study therefore applies and evaluates Firth Penalized Class-Weighted GWLR (Firth CW-GWLR), which combines three components: (i) geographically weighted estimation, allowing predictor effects to vary locally; (ii) class weighting, giving the minority class greater influence in each local likelihood; and (iii) Firth penalization, reducing small-sample bias and protecting against separation. Its contributions are a formulation of the weighted penalized local likelihood and its score function that handles zero weights directly, and a LOOCV comparison of GWLR, CW-GWLR, Firth-GWLR, and Firth CW-GWLR that separates the effects of class weighting and Firth's penalty. The model is applied to the 22 regencies/cities of NTT using three predictors that are not IKP indicators: the open unemployment rate (X_1), the share of workers in the agricultural sector (X_2), and the number of markets (X_3). Section 2 presents the data and methods, Section 3 the results, and Section 4 the conclusions.

2. Methods

Section 2.1 describes the data and Section 2.2 the analytical procedure. Sections 2.3–2.5 review logistic, weighted logistic, and geographically weighted logistic regression with the associated diagnostic tests, and Sections 2.6–2.9 present Firth's penalized likelihood, the imbalance ratio, the performance measures, and leave-one-out cross-validation.

2.1. Data and Variables

This study uses secondary cross-sectional data for all 22 regencies/cities in East Nusa Tenggara (NTT), Indonesia. The response is the 2025 regency/city IKP [23], coded as $Y = 0$ for food-

insecure (groups 1–3) and $Y = 1$ for food-secure (groups 4–6) regencies/cities (Section 1). The predictors are the open unemployment rate (X_1) and the share of agricultural workers (X_2) from the August 2025 Sakernas, and the number of markets (X_3) in 2025, all from BPS (Table 1). None of them is an IKP indicator or a component of one, so there is no mechanical predictor–response circularity; because they are still conceptually related to the IKP aspects, their coefficients are interpreted as associations rather than causal effects.

Table 1: Description of Variables Used in the Study

Variable	Description	Scale
Y	Food security status (1 = food secure, 0 = food insecure)	Binary
X_1	Open unemployment rate (%)	Ratio
X_2	Share of workers in the agricultural sector (%)	Ratio
X_3	Number of markets	Ratio

Each observation $i = 1, 2, \ldots, n$ ($n = 22$) is associated with the centroid coordinates of its regency/city, recorded as WGS84 longitude–latitude and projected onto UTM Zone 51S (in metres) before any distance-based computation.

2.2. Analytical Procedure

The predictors were first described and checked for multicollinearity with the VIF in Eq. (15), and class imbalance was quantified by the imbalance ratio in Eq. (20). The Breusch–Pagan statistic in Eq. (16) and Moran's I in Eq. (17), with a k-nearest-neighbour weight matrix ($k = 4$), were computed from the residuals of the global models (Sections 2.5.2–2.5.3) as indicative, not confirmatory, evidence for local modelling.

Spatial weights were derived from the distances in Eq. (8) between UTM-projected centroids, and the adaptive kernels in Eqs. (9)–(11) were compared on the full data, each with k^* selected by minimum AICc in Eq. (12). With the selected kernel, GWLR, CW-GWLR, Firth-GWLR, and Firth CW-GWLR were estimated at every location, the penalized models by Fisher scoring with the score in Eq. (39), and the local coefficients were tested with the Wald statistic in Eq. (41) and the Benjamini–Hochberg adjustment.

The models were compared using the LOOCV scheme in Eq. (26) using the measures in Eqs. (21)–(24), with a held-out regency/city classified as food secure when $\hat{p}_i \geq 0.5$. The class weights, k^*, and all coefficients were re-estimated in each fold, whereas the kernel type was fixed at the full-data choice, so the LOOCV results are conditional on it (Section 2.9). Finally, the fully local Firth CW-GWLR was compared by the LOOCV Brier score in Eq. (25) with mixed models in which one coefficient is fixed at a global value.

2.3. Logistic Regression

Logistic regression models a binary response $Y_i \sim \text{Bernoulli}(\pi_i)$, with $\pi_i = P(Y_i = 1 \mid \mathbf{x}_i)$ related to the predictors through the logistic function [24, 25]

$$\pi(\mathbf{x}_i) = \frac{\exp\left(\sum_{j=0}^{p} \beta_j x_{ij}\right)}{1 + \exp\left(\sum_{j=0}^{p} \beta_j x_{ij}\right)}, \qquad x_{i0} \equiv 1. \tag{1}$$

Given n independent observations, the log-likelihood is

$$\ell(\boldsymbol{\beta}) = \sum_{i=1}^{n} \left[y_i \ln \pi(\mathbf{x}_i) + (1 - y_i) \ln(1 - \pi(\mathbf{x}_i)) \right], \tag{2}$$

and setting its derivative to zero gives the score equations

$$\frac{\partial \ell}{\partial \beta_j} = \sum_{i=1}^{n} x_{ij}(y_i - \pi_i) = 0, \qquad j = 0, 1, \ldots, p, \tag{3}$$

which are solved iteratively by Newton–Raphson with gradient $\mathbf{X}^T(\mathbf{y}-\boldsymbol{\pi})$ and Hessian $-\mathbf{X}^T\mathbf{V}\mathbf{X}$, $\mathbf{V}=\mathrm{diag}[\pi_i(1-\pi_i)]$, yielding the maximum likelihood estimate (MLE) $\hat{\boldsymbol{\beta}}$ [24].

2.4. Weighted Logistic Regression (WLR)

Under class imbalance, the ordinary MLE is biased toward the majority class [4]. WLR corrects this by assigning each observation i a class weight

$$C_i = \begin{cases} \dfrac{\tau}{\bar{y}}, & y_i = 1 \\ \dfrac{1-\tau}{1-\bar{y}}, & y_i = 0, \end{cases} \tag{4}$$

where $\bar{y}$ is the observed proportion of the minority class and $\tau = 0.5$, so that both classes contribute equally to the likelihood [4, 8]. The weighted log-likelihood is

$$\ell_C(\boldsymbol{\beta}) = \sum_{i=1}^{n} C_i \left[y_i \ln(\pi_i) + (1-y_i)\ln(1-\pi_i) \right], \tag{5}$$

whose weighted score equations $\sum_{i=1}^{n} C_i x_{ij}(y_i - \pi_i) = 0$ are solved by the IRLS update

$$\boldsymbol{\beta}^{(t+1)} = \boldsymbol{\beta}^{(t)} + \left(\mathbf{X}^T \mathbf{C} \mathbf{V}^{(t)} \mathbf{X} \right)^{-1} \mathbf{X}^T \mathbf{C} \left(\mathbf{y} - \boldsymbol{\pi}^{(t)} \right), \tag{6}$$

with $\mathbf{C} = \mathrm{diag}(C_1, \ldots, C_n)$, yielding the class-weighted estimate $\hat{\boldsymbol{\beta}}_C$.

2.5. Geographically Weighted Logistic Regression (GWLR)

GWLR relaxes the spatial-constancy assumption in Eq. (1) by giving each location i, with coordinates (u_i, v_i), its own parameter vector $\boldsymbol{\beta}(u_i, v_i)$:

$$\pi(\mathbf{x}_i) = \frac{\exp\left(\sum_{k=0}^{p} \beta_k(u_i, v_i) x_{ik}\right)}{1 + \exp\left(\sum_{k=0}^{p} \beta_k(u_i, v_i) x_{ik}\right)}. \tag{7}$$

The weight of observation j in the local fit at i depends on the Euclidean distance

$$d_{ij} = \sqrt{(u_i - u_j)^2 + (v_i - v_j)^2}, \tag{8}$$

where (u_i, v_i) are UTM Zone 51S-projected centroids, so d_{ij} in Eq. (8) is in metres. Each location uses an adaptive bandwidth b_i, the distance to its k^*-th nearest neighbour (including itself), in the Gaussian, bisquare, or tricube kernel of Eqs. (9)–(11); the Gaussian kernel gives every observation a positive weight, whereas the other two assign zero weight beyond b_i:

$$w_{ij} = \exp\left[-\tfrac{1}{2} \left(d_{ij}/b_i \right)^2 \right] \quad \text{for all } j, \tag{9}$$

$$w_{ij} = \begin{cases} \left[1 - (d_{ij}/b_i)^2 \right]^2, & d_{ij} \le b_i \\ 0, & d_{ij} > b_i, \end{cases} \tag{10}$$

$$w_{ij} = \begin{cases} \left[1 - (d_{ij}/b_i)^3 \right]^3, & d_{ij} \le b_i \\ 0, & d_{ij} > b_i. \end{cases} \tag{11}$$

The kernel and k^* are chosen by minimum AICc [11],

$$AIC = 2\nu - 2\ln(\hat{L}), \qquad AIC_c = AIC + \frac{2\nu(\nu+1)}{n-\nu-1}, \tag{12}$$

where $\hat{L}$ is the maximized weighted likelihood and ν the effective number of parameters.

With $\mathbf{W}_i = \mathrm{diag}(w_{i1}, \ldots, w_{in})$, the local log-likelihood

$$\ell(\boldsymbol{\beta}(u_i, v_i)) = \sum_{j=1}^{n} w_{ij} \left[y_j \ln(\pi_j) + (1 - y_j) \ln(1 - \pi_j) \right] \tag{13}$$

has the form of Eq. (5) with w_{ij} in place of C_i, so Eq. (13) is maximized by the local IRLS update

$$\boldsymbol{\beta}^{(m+1)}(u_i, v_i) = \boldsymbol{\beta}^{(m)}(u_i, v_i) + \left[\mathbf{X}^T \mathbf{W}_i \mathbf{V}^{(m)} \mathbf{X} \right]^{-1} \mathbf{X}^T \mathbf{W}_i \left(\mathbf{y} - \boldsymbol{\pi}^{(m)} \right), \tag{14}$$

and Eq. (14) is solved independently at each $i = 1, \ldots, n$.

2.5.1. Multicollinearity Test

Multicollinearity inflates coefficient variance [26] and was assessed via the VIF in Eq. (15),

$$VIF_k = \frac{1}{1 - R_k^2}, \tag{15}$$

with R_k^2 from regressing X_k on the remaining predictors; $VIF_k > 10$ indicates serious multicollinearity [26, 27].

2.5.2. Spatial Heterogeneity Test

The studentized Breusch–Pagan test ($H_0 : \sigma_1^2 = \cdots = \sigma_n^2$) checks whether the residual variance is constant across observations. With e_i the residuals of the linear regression of Y on X_1–X_3, the auxiliary regression of e_i^2 on the same predictors gives

$$BP = n R_{e^2}^2, \tag{16}$$

where $R_{e^2}^2$ is its coefficient of determination; under H_0, the statistic in Eq. (16) follows χ_p^2, with $p = 3$. A significant BP indicates heteroscedastic residuals, an indirect signal of possible global-model misspecification; because the variance of a binary response depends on the predictors by construction, it does not by itself demonstrate spatially varying coefficients [11].

2.5.3. Spatial Autocorrelation Test

Moran's I in Eq. (17) measures spatial dependence in a variable e (Pearson residuals of the global logistic regression, so that autocorrelation already explained by X_1, X_2, X_3 is not conflated with unexplained spatial structure),

$$I = \frac{n}{\sum_i \sum_j w_{ij}} \cdot \frac{\sum_i \sum_j w_{ij} (e_i - \bar{e})(e_j - \bar{e})}{\sum_i (e_i - \bar{e})^2}, \tag{17}$$

with values near $+1$, -1, and 0 indicating positive, negative, and no spatial autocorrelation, respectively. The spatial weight matrix $\mathbf{W} = (w_{ij})$ is a row-standardized k-nearest-neighbour matrix with $k = 4$, and significance is assessed with a two-sided normal approximation.

2.6. Firth's Penalized Likelihood

Under small samples or severe imbalance, the MLE in Eq. (3) may fail to exist: near-perfect separation drives $\|\boldsymbol{\beta}\| \to \infty$, and even when it exists it is finite-sample biased [19, 20].

Firth [20] reduces this bias to $O(n^{-1})$ by maximizing the penalized log-likelihood in Eq. (18),

$$\ell^*(\boldsymbol{\beta}) = \ell(\boldsymbol{\beta}) + \frac{1}{2} \ln \left| \mathbf{I}(\boldsymbol{\beta}) \right|, \tag{18}$$

with $\ell(\boldsymbol{\beta})$ from Eq. (2) and $\mathbf{I}(\boldsymbol{\beta}) = \mathbf{X}^T \mathbf{V}(\boldsymbol{\beta}) \mathbf{X}$, $\mathbf{V}(\boldsymbol{\beta}) = \mathrm{diag}[\pi_i(1 - \pi_i)]$; the penalty is equivalent to a Jeffreys prior shrinking extreme estimates toward zero [20, 21].

Differentiating Eq. (18) gives the penalized score

$$\frac{\partial \ell^*}{\partial \beta_j} = \sum_{i=1}^{n} x_{ij}(y_i - \pi_i) + \frac{1}{2}\sum_{i=1}^{n} x_{ij}(1 - 2\pi_i)h_i, \tag{19}$$

where h_i is the leverage (diagonal of $\mathbf{H} = \mathbf{V}^{1/2}\mathbf{X}(\mathbf{X}^T\mathbf{V}\mathbf{X})^{-1}\mathbf{X}^T\mathbf{V}^{1/2}$), equivalent to replacing y_i with $y_i^* = y_i + h_i(\frac{1}{2} - \pi_i)$, which keeps the estimates finite under separation [20, 21]. Setting Eq. (19) to zero gives the penalized score equations, which are solved iteratively by a modified Newton–Raphson (IRLS) algorithm.

2.7. Imbalance Ratio

Class imbalance was quantified by

$$IR = \frac{N_{maj}}{N_{min}}, \tag{20}$$

where $IR \approx 1$ indicates balance and large IR indicates severe imbalance, biasing classifiers toward the majority class [9].

2.8. Confusion Matrix and Classification Performance Measures

Performance was evaluated via the confusion matrix (TP, TN, FP, FN) [28], shown in Table 2.

Table 2: Confusion Matrix for Binary Classification

	Actual Value	
Predicted Value	Positive ($Y = 1$)	Negative ($Y = 0$)
Positive ($\hat{Y} = 1$)	*True Positive* (TP)	*False Positive* (FP)
Negative ($\hat{Y} = 0$)	*False Negative* (FN)	*True Negative* (TN)

Since accuracy is misleading under imbalance, precision, recall, F_1, G-mean, and balanced accuracy were used instead:

$$\text{Precision} = \frac{TP}{TP + FP}, \qquad \text{Recall} = \frac{TP}{TP + FN}, \qquad \text{Specificity} = \frac{TN}{TN + FP}, \tag{21}$$

$$F_1 = \frac{2\,TP}{2\,TP + FP + FN}, \tag{22}$$

$$G\text{-mean} = \sqrt{\text{Recall} \cdot \text{Specificity}}, \tag{23}$$

$$\text{Balanced Accuracy} = \frac{1}{2}\left(\text{Recall} + \text{Specificity}\right). \tag{24}$$

The accuracy of the predicted probabilities $\hat{p}_i$ was evaluated by the Brier score,

$$B = \frac{1}{n}\sum_{i=1}^{n}\left(\hat{p}_i - y_i\right)^2, \tag{25}$$

where smaller values indicate more accurate probabilities; it is used to compare the fully local and mixed Firth CW-GWLR specifications (Section 3.8).

2.9. Leave-One-Out Cross-Validation (LOOCV)

Given the small sample size, all performance measures and the Brier score were computed via LOOCV [29]. For each regency/city i, the model is refitted without observation i, and the held-out predicted probability is

$$\hat{p}_i = \pi\Big(\mathbf{x}_i; \hat{\boldsymbol{\beta}}_{(-i)}(u_i, v_i)\Big), \qquad i = 1, \ldots, n, \tag{26}$$

where $\hat{\boldsymbol{\beta}}_{(-i)}(u_i, v_i)$ is the local coefficient vector at location i estimated from the remaining $n-1$ observations. The held-out regency/city is classified as food secure ($\hat{y}_i = 1$) when $\hat{p}_i \geq 0.5$, and the n held-out predictions from Eq. (26) are used to compute the measures in Eqs. (21)–(25).

Within each training fold, the class weights in Eq. (4), the bandwidth k^* in Eq. (12), the global coefficients of the mixed models (Section 3.8), and all local coefficients are estimated from the $n-1$ training observations only. The kernel type, however, was selected once on the full data and held fixed in every fold, so the held-out response entered the kernel choice. The reported LOOCV performance is therefore conditional on this full-data kernel selection, is not a fully nested estimate, and may be slightly optimistic; a fully nested scheme would repeat the kernel selection within each fold.

All analyses used R: `car` (VIF), `lmtest` (Breusch–Pagan), `spdep` (Moran's I), and `GWmodel` [30] for kernel and bandwidth selection by AICc. The four local models were estimated with a purpose-written R routine.

3. Results and Discussion

Section 3.1 derives the Firth CW-GWLR estimator, its estimation algorithm, and its local significance test. Sections 3.2–3.5 present the descriptive statistics, the multicollinearity and spatial tests, and the kernel selection for the NTT data; Section 3.6 verifies convergence, Section 3.7 compares the local coefficients, Section 3.8 assesses spatially varying effects out of sample, and Section 3.9 compares the LOOCV classification performance.

3.1. Firth Penalized Class-Weighted Geographically Weighted Logistic Regression Estimation

3.1.1. Combined Spatial-and-Class Weighted Likelihood

Class-Weighted GWLR (CW-GWLR) combines the class weighting of WLR with the spatial weighting of GWLR: ordinary GWLR still lets the majority class dominate each local likelihood, whereas WLR assumes a spatially constant relationship. Each observation j therefore receives a combined weight reflecting its proximity to location i and its class,

$$\omega_{ij} = C_j \, w_{ij}(u_i, v_i), \tag{27}$$

where C_j is the class weight from Eq. (4) and $w_{ij}(u_i, v_i)$ is the spatial weight from Section 2.5. Collecting the combined weights in Eq. (27) into $\boldsymbol{\Omega}_i = \mathrm{diag}(\omega_{i1}, \ldots, \omega_{in})$, the local likelihood at location i is formed, analogous to Eq. (5), by raising the contribution of each observation j to the power of its combined weight ω_{ij},

$$L(\boldsymbol{\beta}(u_i, v_i)) = \prod_{j=1}^{n} \pi_j^{\omega_{ij} y_j} \, (1 - \pi_j)^{\omega_{ij}(1-y_j)}, \tag{28}$$

with π_j evaluated at $\boldsymbol{\beta}(u_i, v_i)$ as in Eq. (7). Taking the logarithm of Eq. (28) gives the local weighted log-likelihood,

$$\ell_i(\boldsymbol{\beta}) = \sum_{j=1}^{n} \omega_{ij} \left[y_j \ln(\pi_j) + (1 - y_j) \ln(1 - \pi_j) \right]. \tag{29}$$

Differentiating Eq. (29) with respect to β_k, exactly as in the derivation of Eq. (6), gives the ordinary local score function,

$$\frac{\partial \ell_i}{\partial \beta_k} = \sum_{j=1}^{n} \omega_{ij} x_{jk} (y_j - \pi_j), \qquad k = 0, 1, \ldots, p, \tag{30}$$

and setting Eq. (30) to zero for every k gives the local likelihood equations,

$$\sum_{j=1}^{n} \omega_{ij} x_{jk} (y_j - \pi_j) = 0, \qquad k = 0, 1, \ldots, p. \tag{31}$$

The $p+1$ equations in Eq. (31) can be written in matrix form as

$$\mathbf{q}_i = \mathbf{X}^T \mathbf{\Omega}_i (\mathbf{y} - \boldsymbol{\pi}) = \mathbf{0}, \tag{32}$$

and the Hessian of ℓ_i is

$$\nabla^2 \ell_i(\boldsymbol{\beta}) = -\mathbf{X}^T \mathbf{\Omega}_i \mathbf{V} \mathbf{X}, \qquad \mathbf{V} = \operatorname{diag}\left[\pi_j (1 - \pi_j)\right], \tag{33}$$

Combining the score in Eq. (32) with the Hessian in Eq. (33) in a Newton–Raphson step gives the ordinary IRLS update

$$\boldsymbol{\beta}^{(m+1)}(u_i, v_i) = \boldsymbol{\beta}^{(m)}(u_i, v_i) + \left[\mathbf{X}^T \mathbf{\Omega}_i \mathbf{V}^{(m)} \mathbf{X}\right]^{-1} \mathbf{X}^T \mathbf{\Omega}_i \left(\mathbf{y} - \boldsymbol{\pi}^{(m)}\right). \tag{34}$$

Eq. (34) defines the unpenalized CW-GWLR estimator, to which Firth's penalty is added below.

3.1.2. Penalized Log-Likelihood and Score Function

Because each local subsample is weighted by both distance and class, the effective number of minority-class observations, $n_{\text{eff},i} = \sum_j \omega_{ij} y_j$, can be very small. If the local predictors then (nearly) separate the weighted classes, the ordinary MLE solving Eq. (30) diverges, $\|\boldsymbol{\beta}(u_i, v_i)\| \to \infty$ [19]. Firth [20] showed that this non-existence and the finite-sample bias of the MLE are removed by adding one half the log-determinant of the Fisher information to the log-likelihood, equivalent to a Jeffreys prior on $\boldsymbol{\beta}$.

Applying this penalty to Eq. (29) gives the Firth Penalized CW-GWLR log-likelihood at location (u_i, v_i),

$$\ell_i^*(\boldsymbol{\beta}) = \ell_i(\boldsymbol{\beta}) + \frac{1}{2} \ln \left|\mathbf{I}_i(\boldsymbol{\beta})\right|, \tag{35}$$

where $\mathbf{I}_i(\boldsymbol{\beta}) = \mathbf{X}^T \mathbf{\Omega}_i \mathbf{V}(\boldsymbol{\beta}) \mathbf{X}$ is the combined spatial-and-class-weighted Fisher information matrix, whose diagonal weight $\omega_{ij} \pi_j (1 - \pi_j)$ combines the spatial weight, the class weight, and the Bernoulli variance. Differentiating Eq. (35) with respect to β_k gives

$$\frac{\partial \ell_i^*}{\partial \beta_k} = \frac{\partial \ell_i}{\partial \beta_k} + \frac{1}{2} \operatorname{tr}\left[\mathbf{I}_i(\boldsymbol{\beta})^{-1} \frac{\partial \mathbf{I}_i(\boldsymbol{\beta})}{\partial \beta_k}\right], \tag{36}$$

where the first term is the ordinary score in Eq. (30). Since $\partial \pi_j (1-\pi_j) / \partial \beta_k = x_{jk} \pi_j (1-\pi_j)(1-2\pi_j)$ for the logistic link,

$$\frac{\partial \mathbf{I}_i(\boldsymbol{\beta})}{\partial \beta_k} = \mathbf{X}^T \operatorname{diag}\left[\omega_{ij} x_{jk} \pi_j (1 - \pi_j)(1 - 2\pi_j)\right] \mathbf{X}. \tag{37}$$

Substituting Eq. (37) into Eq. (36) and writing $v_j = \pi_j (1 - \pi_j)$, the trace term becomes $\sum_j x_{jk} h_{ij} (\frac{1}{2} - \pi_j)$, where

$$h_{ij} = \omega_{ij}\, v_j\, \mathbf{x}_j^T \left(\mathbf{X}^T \mathbf{\Omega}_i \mathbf{V} \mathbf{X}\right)^{-1} \mathbf{x}_j \tag{38}$$

is the weighted leverage, i.e., the j-th diagonal element of $\mathbf{H}_i = \mathbf{\Omega}_i^{1/2} \mathbf{V}^{1/2} \mathbf{X} \left(\mathbf{X}^T \mathbf{\Omega}_i \mathbf{V} \mathbf{X}\right)^{-1} \mathbf{X}^T \mathbf{V}^{1/2} \mathbf{\Omega}_i^{1/2}$. The penalized score function is therefore

$$U_k^*(\boldsymbol{\beta}) = \sum_{j=1}^{n} x_{jk} \left[\omega_{ij} (y_j - \pi_j) + h_{ij} \left(\tfrac{1}{2} - \pi_j\right)\right], \qquad k = 0, 1, \ldots, p, \tag{39}$$

in which the combined weight enters the data term explicitly and the penalty term only through h_{ij}, so it is not applied twice. Because Eq. (39) involves no division by ω_{ij}, zero weights are handled directly: if $\omega_{ij} = 0$, then $h_{ij} = 0$ by Eq. (38), and observation j contributes nothing to the log-likelihood, the information matrix, or the penalized score, which is equivalent to excluding it from the local sample. The penalty term pulls the fitted probabilities towards $\frac{1}{2}$ in proportion to h_{ij}, which keeps the local coefficients finite when the weighted sub-sample is (quasi-)separated.

3.1.3. Estimation Algorithm and Local Significance Testing

Setting $\mathbf{U}^*(\boldsymbol{\beta}) = \mathbf{0}$ and solving by Fisher scoring, using the combined-weighted information matrix as the working Hessian, gives the update

$$\boldsymbol{\beta}^{(t+1)}(u_i, v_i) = \boldsymbol{\beta}^{(t)}(u_i, v_i) + \lambda_t \left[\mathbf{X}^T \boldsymbol{\Omega}_i \mathbf{V}^{(t)} \mathbf{X}\right]^{-1} \mathbf{X}^T \left\{\boldsymbol{\Omega}_i \left(\mathbf{y} - \boldsymbol{\pi}^{(t)}\right) + \mathbf{h}_i^{(t)} \odot \left(\tfrac{1}{2}\mathbf{1} - \boldsymbol{\pi}^{(t)}\right)\right\}, \quad (40)$$

where $\odot$ is the element-wise product, $\mathbf{h}_i^{(t)}$ is the vector of weighted leverages h_{ij} evaluated at $\boldsymbol{\beta}^{(t)}$, and the step length $\lambda_t \in \{1, \frac{1}{2}, \frac{1}{4}, \dots\}$ is halved until the penalized log-likelihood does not decrease. Because Eq. (40) uses h_{ij} directly, no division by ω_{ij} is required. Eq. (40) is solved independently at every location, and an estimate is accepted only when the maximum absolute penalized score satisfies $\max_k |U_k^*| < 10^{-6}$.

Local significance is assessed using the Wald statistic in Eq. (41),

$$Z_k(u_i, v_i) = \frac{\hat{\beta}_k(u_i, v_i)}{SE\left(\hat{\beta}_k(u_i, v_i)\right)}, \quad (41)$$

where $SE\left(\hat{\beta}_k(u_i, v_i)\right)$ is obtained from the square root of the diagonal of $\left[\mathbf{X}^T \boldsymbol{\Omega}_i \mathbf{V}(\hat{\boldsymbol{\beta}}) \mathbf{X}\right]^{-1}$ at convergence, and $H_0 : \beta_k(u_i, v_i) = 0$ is rejected at level α when the statistic in Eq. (41) satisfies $|Z_k(u_i, v_i)| > Z_{\alpha/2}$. For each coefficient, the 22 local p-values are adjusted using the Benjamini–Hochberg false discovery rate procedure [31].

3.2. Descriptive Statistics

The dataset covers 22 regencies/cities in NTT, of which only 4 (Sikka, Ende, Ngada, and Manggarai) are classified as food secure ($Y = 1$) and the remaining 18 as food insecure ($Y = 0$). Figure 1 shows the resulting class distribution.

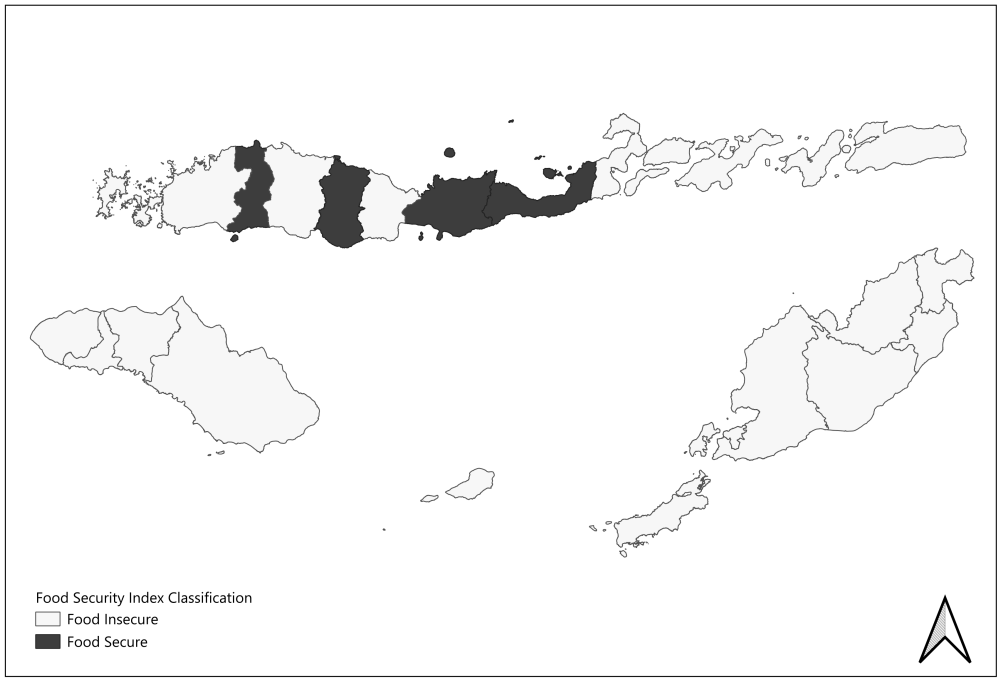


Fig. 1: Spatial Distribution of Observed Food Security Status Across NTT Regencies/Cities

Substituting $N_{maj} = 18$ and $N_{min} = 4$ into Eq. (20) gives Eq. (42),

$$IR = \frac{N_{maj}}{N_{min}} = \frac{18}{4} = 4.5, \quad (42)$$

which confirms severe imbalance, with food-insecure regencies/cities outnumbering food-secure ones by a factor of 4.5. Table 3 summarizes the predictors.

Table 3: Descriptive Statistics of the Predictors ($n = 22$)

Variable	Min.	Q1	Median	Mean	Q3	Max.	SD
X_1	1.24	2.07	2.61	3.05	3.35	8.61	1.75
X_2	6.72	38.36	43.25	44.19	52.82	72.64	12.85
X_3	51	112.75	159.00	154.73	192.25	276	66.17

Kota Kupang is an outlier on two predictors, with the highest unemployment rate (8.61%) and the lowest agricultural share (6.72%), whereas Manggarai Timur has the lowest unemployment rate (1.24%) and the highest agricultural share (72.64%). The food-secure regencies have, on average, more markets than the food-insecure ones (mean 221.5 vs. 139.9) and slightly lower unemployment (2.43% vs. 3.18%) and agricultural shares (41.6% vs. 44.8%).

3.3. Multicollinearity Test

The VIF for each predictor is reported in Table 4.

Table 4: Variance Inflation Factors

	X_1	X_2	X_3
VIF	2.0274	1.1169	1.8909

All VIF values are well below the conventional threshold of 10 [26], indicating no serious multicollinearity, so all three predictors are retained.

3.4. Spatial Heterogeneity and Spatial Autocorrelation Tests

Table 5 reports the studentized Breusch–Pagan test on the residuals of the linear regression of Y on X_1–X_3 and Moran's I of the Pearson residuals of the global logistic regression (GLR) with a k-nearest-neighbour weight matrix ($k = 4$) (Sections 2.5.2–2.5.3).

Table 5: Spatial Heterogeneity and Spatial Autocorrelation Tests

Test	Statistic	p-value
Breusch–Pagan	$BP = 8.4555$	0.0375
Moran's I (KNN, $k = 4$)	$I = 0.2025$ $Z = 2.2442$	0.0248

Table 5 shows that the Breusch–Pagan test rejects homoscedasticity ($BP = 8.4555$, $p = 0.0375$), although for a binary response non-constant variance is expected even when the coefficients are constant, so this result is only supporting evidence. The GLR residuals are positively autocorrelated ($I = 0.2025$; $p = 0.0248$): neighbouring regencies tend to have residuals of the same sign, in line with the Flores cluster of food-secure regencies in Figure 1. Taken together, these tests motivate examining local models; they do not establish that the regression coefficients themselves vary significantly across regencies.

3.5. Kernel and Bandwidth Selection

Three adaptive kernels were compared using AICc on the UTM-projected coordinates; results are shown in Table 6.

Table 6: Comparison of Adaptive Kernels for GWLR

Kernel	Bandwidth k^*	AICc
Adaptive Gaussian	**21**	**21.1697**
Adaptive Tricube	21	21.2927
Adaptive Bisquare	21	22.6561

The adaptive Gaussian kernel with $k = 21$ had the lowest AICc (Table 6) and was used for all local models. A bandwidth of $k = 21$ means that each location uses its 21 nearest neighbours, including itself; because the Gaussian kernel is untruncated, all 22 observations still receive a positive weight that decreases with distance. In each LOOCV fold ($n = 21$), the bandwidth was reselected by AICc from the training data, giving $k = 20$, whereas the kernel type was kept fixed at the adaptive Gaussian kernel selected on the full data (Section 2.9). The adaptive kernel is compatible with Firth CW-GWLR because the estimator only requires a non-negative weight $\omega_{ij} = C_j w_{ij}$ for each observation.

3.6. Convergence and Numerical Verification

All four models converged at all 22 locations on the full data, requiring 7 iterations for GWLR and CW-GWLR, 15–18 for Firth-GWLR, and 27–167 for Firth CW-GWLR, and all produced valid held-out predictions in the 22 LOOCV folds without any fallback. For Firth CW-GWLR, the maximum absolute penalized score over the 22 local fits was 9.89×10^{-7}, below the tolerance of 10^{-6}. Because convergence alone does not verify that an estimator is implemented correctly, the estimates were also reproduced by an independent implementation of Eq. (40), which gave the same local coefficients (maximum absolute difference 1.5×10^{-6}) and the same LOOCV predicted probabilities (maximum absolute difference 1.3×10^{-8}). In the fold that holds out Ngada, the unpenalized GWLR and CW-GWLR met the convergence criterion, but their local coefficients were very large ($|\hat{\beta}| \approx 10^3$), indicating quasi-complete separation, whereas the Firth-penalized models remained finite in every fold ($|\hat{\beta}| \leq 4.0$).

3.7. Comparison of Local Coefficients Across Models

To examine how class weighting and Firth penalization affect the local estimates, the local coefficients of GWLR, CW-GWLR, Firth-GWLR, and Firth CW-GWLR were summarized across the 22 locations using their mean, minimum, maximum, and standard deviation (Table 7).

Table 7: Summary of Local Coefficients Across the Four Models (Full Data, 22 Locations)

Model	Term	Mean	Min.	Max.	SD
GWLR	Intercept	−0.5080	−0.8329	0.0174	0.2260
	X_1	−1.5666	−1.6779	−1.4313	0.0842
	X_2	−0.0940	−0.0994	−0.0882	0.0042
	X_3	0.0392	0.0372	0.0411	0.0016
CW-GWLR	Intercept	−1.7025	−2.4082	−0.7374	0.6065
	X_1	−1.4742	−1.5596	−1.3880	0.0737
	X_2	−0.0945	−0.1045	−0.0828	0.0095
	X_3	0.0529	0.0442	0.0611	0.0073
Firth-GWLR	Intercept	0.8542	0.6982	0.9746	0.1166
	X_1	−0.3564	−0.3914	−0.3185	0.0304
	X_2	−0.0873	−0.0909	−0.0835	0.0032
	X_3	0.0167	0.0151	0.0183	0.0014
Firth CW-GWLR	Intercept	1.2472	0.6686	1.7316	0.3495
	X_1	−0.5837	−0.7241	−0.4802	0.0994
	X_2	−0.0866	−0.0946	−0.0807	0.0038
	X_3	0.0235	0.0200	0.0280	0.0033

Table 7 compares the four estimators. The SDs of the local coefficients (0.0994, 0.0038, and 0.0033 for X_1–X_3 under Firth CW-GWLR) depend on the units of the predictors and should not be compared directly across predictors. Firth penalization shrinks the slopes towards zero relative to the corresponding unpenalized models (for example, the mean X_1 coefficient is -0.356 under Firth-GWLR compared with -1.567 under GWLR), as expected of a Jeffreys-prior penalty with only four events [20, 22]. Smaller coefficients are not necessarily more accurate or less biased; the benefit of the penalty has to be judged by out-of-sample performance (Section 3.9), not by coefficient magnitude.

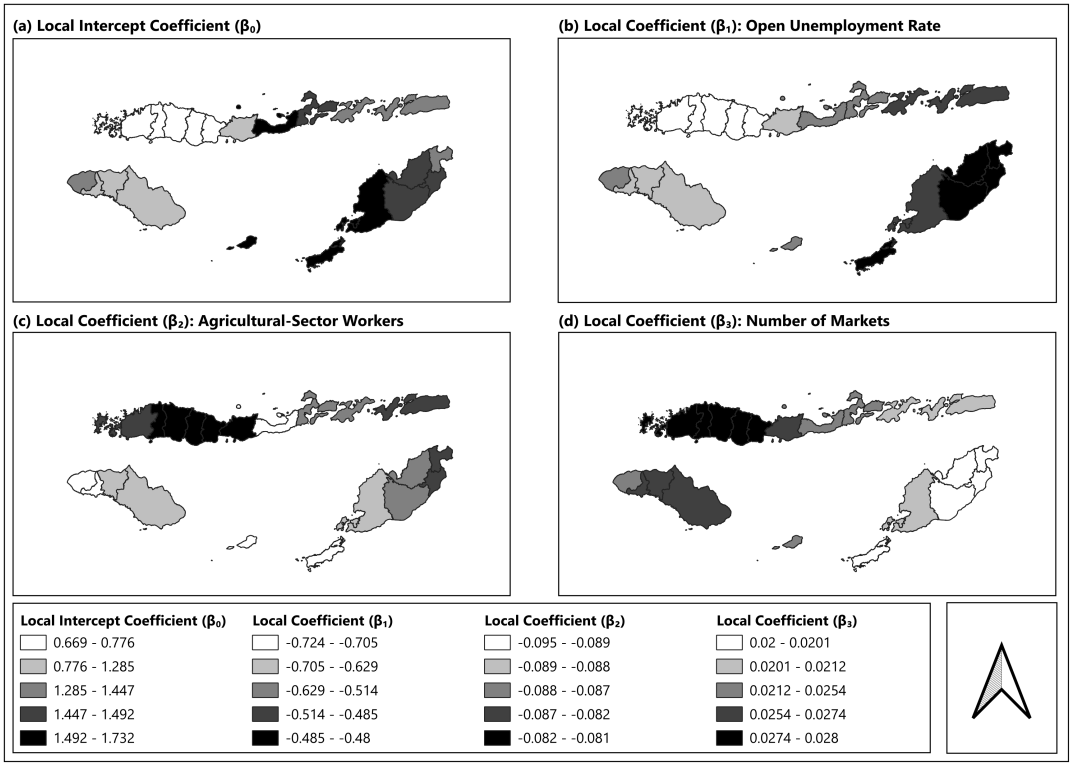


Fig. 2: Spatial Distribution of the Local Coefficients of Firth CW-GWLR: (a) Intercept, (b) Open Unemployment Rate (X_1), (c) Agricultural-Sector Workers (X_2), and (d) Number of Markets (X_3)

The local coefficients of Firth CW-GWLR, the proposed model (Figure 2), keep the same sign at all 22 locations but vary smoothly from west to east. The X_1 coefficient is negative everywhere (-0.7241 to -0.4802); it is most negative in western and central Flores (Ngada, Manggarai Timur, Nagekeo, and Manggarai) and least negative on Timor and Rote (Malaka, Rote Ndao, and Belu). The X_2 coefficient is also negative everywhere (-0.0946 to -0.0807), most strongly in Sabu Raijua, Sikka, and Sumba Barat Daya and least strongly in western Flores. The X_3 coefficient is positive everywhere (0.0200 to 0.0280), largest in western Flores and smallest on Timor and Rote. The intercept ranges from 0.6686 (Manggarai Timur) to 1.7316 (Sabu Raijua). Thus, higher unemployment and a larger share of agricultural workers are associated with lower odds of food security, whereas more markets are associated with higher odds, with the strongest associations for X_1 and X_3 in western Flores. However, none of these local coefficients remains significant after the Benjamini–Hochberg adjustment (Section 3.1.3), so the patterns are descriptive.

3.8. Out-of-Sample Assessment of Spatially Varying Effects

Because the classical F-test of Leung et al. [32] has no established extension to a class-weighted, Firth-penalized local likelihood, an out-of-sample comparison is used instead. For each term k, a mixed model is fitted in which β_k is fixed at a global value $\hat{\gamma}_k$ (entered as an offset) while the other coefficients remain local. In each LOOCV fold, $\hat{\gamma}_k$ is estimated from the training data only by a global Firth class-weighted logistic regression, and the mixed model is compared with the

fully local Firth CW-GWLR on the LOOCV Brier score in Eq. (25), with the kernel type fixed at the full-data choice as in Section 2.9. A positive difference (mixed minus fully local) indicates that allowing β_k to vary improves out-of-sample prediction.

Table 8: Out-of-Sample Comparison of Mixed and Full Local Firth CW-GWLR (LOOCV Brier Score)

Variable	$\bar{\gamma}_k$	Brier (Mixed)	Brier (Full Local)	Difference
Intercept	0.2480	0.1881	0.1820	+0.0061
X_1	−0.6807	0.1858	0.1820	+0.0038
X_2	−0.0789	0.1872	0.1820	+0.0052
X_3	0.0279	0.1897	0.1820	+0.0077

The fully local model has a lower Brier score than all four mixed specifications, so allowing each coefficient to vary gives slightly better out-of-sample predictions than fixing it at a global value. The differences are small (0.0038–0.0077), and $\bar{\gamma}_k$ is the average of the 22 training-fold estimates rather than a pooled full-data coefficient. With $n = 22$ and only four food-secure regencies, this comparison is descriptive evidence consistent with some spatial variation, not a formal test that the coefficients vary significantly across regencies.

3.9. Classification Performance

To separate the contributions of class weighting and Firth's penalty, GWLR, CW-GWLR, Firth-GWLR, and Firth CW-GWLR were compared by LOOCV. In every fold, the bandwidth, class weights, and coefficients were recomputed from the training data only, and a held-out regency/city was classified as food secure when $\hat{p}_i \geq 0.5$. Because the kernel type was selected once on the full data, the results in Table 9 are conditional on that choice (Section 2.9). Table 9 reports the results.

Table 9: LOOCV Performance Comparison of the Four Models

Model	TP	TN	FP	FN	Precision	Recall	Specificity	F_1	G-mean	Bal. Acc.
GWLR	2	16	2	2	0.5000	0.5000	0.8889	0.5000	0.6667	0.6944
CW-GWLR	3	14	4	1	0.4286	0.7500	0.7778	0.5455	0.7638	0.7639
Firth-GWLR	1	17	1	3	0.5000	0.2500	0.9444	0.3333	0.4859	0.5972
Firth CW-GWLR	3	13	5	1	0.3750	0.7500	0.7222	0.5000	0.7360	0.7361

No model is best on all criteria. Class weighting raises recall from 0.50 (GWLR) to 0.75 (CW-GWLR) and from 0.25 (Firth-GWLR) to 0.75 (Firth CW-GWLR), but lowers specificity; CW-GWLR achieves the highest balanced accuracy (0.7639), F_1 (0.5455), and G-mean (0.7638). Firth-GWLR has the highest specificity (0.9444) but the lowest recall (0.25). Firth CW-GWLR identifies three of the four food-secure regencies (recall 0.75), but it misclassifies five food-insecure regencies (Timor Tengah Selatan, Timor Tengah Utara, Alor, Lembata, and Flores Timur) as food secure. Since $Y = 1$ denotes food security, these false positives are the more consequential error for policy, because food-insecure regencies could be omitted from intervention. With only four food-secure regencies, a single misclassification changes recall by 0.25, so these differences are indicative rather than conclusive.

4. Conclusion

This study applied Firth Penalized CW-GWLR, which combines geographically weighted estimation, class weighting, and Firth penalization, to classify the food security status of the 22 regencies/cities of NTT based on the 2025 IKP. All four local models converged, and the Firth penalty kept the local estimates finite in a LOOCV fold where the unpenalized models showed quasi-complete separation. Under LOOCV, which is conditional on the full-data kernel choice, no model was best on all criteria: CW-GWLR achieved the highest balanced accuracy (0.7639) and

F_1 (0.5455), Firth-GWLR the highest specificity (0.9444), and Firth CW-GWLR identified three of the four food-secure regencies (recall 0.7500) but misclassified five food-insecure regencies as food secure. The fully local Firth CW-GWLR had a slightly lower Brier score than each mixed specification (0.1820 against 0.1858–0.1897), which is descriptive rather than formal evidence of spatial variation. These findings are limited by the small sample ($n = 22$, with only four food-secure regencies), the binary collapse of the six-category IKP, the full-data kernel selection, which may make the LOOCV results slightly optimistic, and the absence of a formal test of spatial nonstationarity for penalized weighted local likelihoods. Future work should extend the approach to an ordinal response that preserves the six IKP categories, apply it to larger samples or multiple years, perform the kernel selection within each cross-validation fold, and develop a formal test of nonstationarity for this class of models.

CRediT Authorship Contribution Statement

Zheryl Zabrina Hibatullah: Conceptualization, methodology, formal analysis, writing—original draft. **Suci Astutik**: Supervision, validation, review. **Henny Pramoedyo**: Supervision, validation, review.

Declaration of Generative AI and AI-assisted Technologies

During the preparation of this manuscript, the authors used AI-assisted tools for language refinement, formula formatting, and manuscript structuring. All data analysis, model derivation, and model estimation were performed and verified by the authors, who take full responsibility for the content of this manuscript.

Declaration of Competing Interest

The authors declare no competing interests.

Funding and Acknowledgments

This research received no external funding. The authors thank the National Food Agency (*Badan Pangan Nasional*) and the Central Bureau of Statistics (BPS) for making the underlying 2025 data publicly available.

Data and Code Availability

The 2025 regency/city Food Security Index data were obtained from the National Food Agency [23], and the open unemployment rate, the share of workers in the agricultural sector, and the number of markets for NTT were obtained from BPS. Kernel and bandwidth selection used `GWmodel` [30], and spatial weights and coordinate projection used `sf` and `spdep`. GWLR, CW-GWLR, Firth-GWLR, and Firth CW-GWLR were estimated with a purpose-written R routine that maximizes the weighted (penalized) local log-likelihood directly. Code and data are available from the corresponding author upon reasonable request.

References

[1] Andi Rachman Salasa. “Paradigma dan dimensi strategi ketahanan pangan Indonesia”. In: *Jejaring Administrasi Publik* 13.1 (2021), pp. 35–48. DOI: 10.20473/jap.v13i1.29357.

[2] Badan Pangan Nasional. *Peta Ketahanan dan Kerentanan Pangan (Food Security and Vulnerability Atlas) Tahun 2025*. Tech. rep. Jakarta: Deputi Bidang Kerawanan Pangan dan Gizi, Badan Pangan Nasional, 2025. https://badanpangan.go.id/.

[3] Anna Ellenora Nainupu and Kadir. "Determinan Status Kerawanan Pangan Rumah Tangga di Provinsi Nusa Tenggara Timur: Analisis Regresi Logistik". In: *Jurnal Statistika Terapan* 3.1 (2023), pp. 37–49. https://jstar.id/ojs/index.php/JSTAR/article/view/38.

[4] Gary King and Langche Zeng. "Logistic regression in rare events data". In: *Political Analysis* 9.2 (2001), pp. 137–163. DOI: 10.1093/oxfordjournals.pan.a004868.

[5] Wei Chen, Kun Yang, Zhiyong Yu, Yong Shi, and C. L. Philip Chen. "A survey on imbalanced learning: latest research, applications and future directions". In: *Artificial Intelligence Review* 57.6 (2024), p. 137. DOI: 10.1007/s10462-024-10759-6.

[6] Lili Zhang, Trent Geisler, Herman Ray, and Ying Xie. "Improving logistic regression on the imbalanced data by a novel penalized log-likelihood function". In: *Journal of Applied Statistics* 49.13 (2022), pp. 3257–3277. DOI: 10.1080/02664763.2021.1939662.

[7] Jyun-You Chiang, Yuhlong Lio, Chih-Ying Hsu, Chia-Ling Ho, and Tzong-Ru Tsai. "Binary classification with imbalanced data". In: *Entropy* 26.1 (2024), p. 15. DOI: 10.3390/e26010015.

[8] Dian Eka Apriana Sulasih, Santi Wulan Purnami, and Santi Puteri Rahayu. "The theoretical study of rare event weighted logistic regression for classification of imbalanced data". In: *Proceeding of the International Summit on Science, Technology, and Humanity (ISETH)*. 2015, pp. 159–169. DOI: 10.23917/iseth.2376.

[9] Rui Zhu, Yiwen Guo, and Jing-Hao Xue. "Adjusting the imbalance ratio by the dimensionality of imbalanced data". In: *Pattern Recognition Letters* 133 (2020), pp. 217–223. DOI: 10.1016/j.patrec.2020.03.004.

[10] Ruben Van den Goorbergh, Maarten Van Smeden, Dirk Timmerman, and Ben Van Calster. "The harm of class imbalance corrections for risk prediction models: illustration and simulation using logistic regression". In: *Journal of the American Medical Informatics Association* 29.9 (2022), pp. 1525–1534. DOI: 10.1093/jamia/ocac093.

[11] A. Stewart Fotheringham, Chris Brunsdon, and Martin Charlton. *Geographically Weighted Regression: The Analysis of Spatially Varying Relationships*. Wiley, 2002.

[12] P. Novianti and D. Rosadi. "Geographically Weighted Logistic Regression Model on Binomial Data to Explore Weather Spatial Non-Stationarity in Covid-19 Cases". In: *Engineering Letters* 31.3 (2023). https://www.engineeringletters.com/issues_v31/issue_3/EL_31_3_07.pdf.

[13] G. A. N. B. Syafii, R. D. Hanifah, R. Arisanti, and N. H. Pusponegoro. "Pemodelan Status Ketertinggalan Perekonomian Regional Menggunakan Geographically Weighted Logistic Regression (GWLR)". In: *Seminar Nasional Official Statistics*. Vol. 2024. 1. 2024, pp. 1085–1096. DOI: 10.34123/semnasoffstat.v2024i1.2208.

[14] Toha Saifudin, Leni Sartika Panjaitan, Sabrina Falasifah, and Yan Dwi. "Application of Geographically Weighted Logistic Regression in Modeling the Human Development Index in East Java". In: *Jurnal Aplikasi Statistika & Komputasi Statistik* 16.1 (2024), pp. 43–57. DOI: 10.34123/jurnalasks.v16i1.587.

[15] Ernest Konadu Appiah, Eric Nimako Aidoo, Richard Kodzo Avuglah, and Simon Kojo Appiah. "Geographically weighted logistic regression model for identifying risk factors for malaria infection among under-5 children in Ghana". In: *Scientific African* 26 (2024), e02398. DOI: 10.1016/j.sciaf.2024.e02398.

[16] Nurul Era Natasyah Beddu Solo, Muhammad Nusrang, and Zakiyah Mar'ah. "Pemodelan Faktor-Faktor Yang Berpengaruh Terhadap Angka Buta Huruf di Provinsi Sulawesi Selatan dengan Geographically Weighted Logistic Regression (GWLR)". In: *VARIANSI: Journal of Statistics and Its Application on Teaching and Research* 6.1 (2024), pp. 44–55. https://jurnalvariansi.unm.ac.id/index.php/variansi/article/view/141.

[17] Siti Hariati Hastuti, Alissa Chintyana, and Hanipar Mahyulis Sastriana. "Evaluating Kernel Weighting Functions in Geographically Weighted Logistic Regression for Spatial Modelling of Stunting in East Lombok". In: *Euler: Jurnal Ilmiah Matematika, Sains dan Teknologi* 13.3 (2025), pp. 338–345. DOI: 10.37905/euler.v13i3.33951.

[18] Vida Faiza Rochmah and Vita Ratnasari. "Pemodelan Ketahanan Pangan di Jawa Timur Menggunakan Metode Geographically Weighted Ordinal Logistic Regression (GWOLR)". In: *Jurnal Sains dan Seni ITS* 8.2 (2019), pp. D487–D492. DOI: 10.12962/j23373520.v8i2.47021.

[19] Mohammad Ali Mansournia, Angelika Geroldinger, Sander Greenland, and Georg Heinze. "Separation in logistic regression: causes, consequences, and control". In: *American Journal of Epidemiology* 187.4 (2018), pp. 864–870. DOI: 10.1093/aje/kwx299.

[20] David Firth. "Bias reduction of maximum likelihood estimates". In: *Biometrika* 80.1 (1993), pp. 27–38. DOI: 10.1093/biomet/80.1.27.

[21] Georg Heinze and Michael Schemper. "A solution to the problem of separation in logistic regression". In: *Statistics in Medicine* 21.16 (2002), pp. 2409–2419. DOI: 10.1002/sim.1047.

[22] Ioannis Kosmidis and David Firth. "Jeffreys-prior penalty, finiteness and shrinkage in binomial-response generalized linear models". In: *Biometrika* 108.1 (2021), pp. 71–82. DOI: 10.1093/biomet/asaa052.

[23] Badan Pangan Nasional. *Indeks Ketahanan Pangan (IKP) Tingkat Kabupaten Kota Tahun 2025 (12 Indikator)*. Open Data Badan Pangan Nasional. 2025. https://data.badanpangan.go.id/datasetpublications/uny/ikp-kabupaten-kota-2025.

[24] David W Hosmer Jr, Stanley Lemeshow, and Rodney X Sturdivant. *Applied logistic regression*. 3rd. John Wiley & Sons, 2013. DOI: 10.1002/9781118548387.

[25] Alan Agresti. *Categorical data analysis*. 3rd. John Wiley & Sons, 2013.

[26] Habshah Midi, Saroje Kumar Sarkar, and Sohel Rana. "Collinearity diagnostics of binary logistic regression model". In: *Journal of Interdisciplinary Mathematics* 13.3 (2010), pp. 253–267. DOI: 10.1080/09720502.2010.10700699.

[27] D. S. Young. *Handbook of Regression Methods*. CRC Press, 2018. DOI: 10.1201/9781315154701.

[28] S Sathyanarayanan and B Roopashri Tantri. "Confusion matrix-based performance evaluation metrics". In: *African Journal of Biomedical Research* 27.4S (2024), pp. 4023–4031. DOI: 10.53555/AJBR.v27i4S.4345.

[29] Sylvain Arlot and Alain Celisse. "A survey of cross-validation procedures for model selection". In: *Statistics Surveys* 4 (2010), pp. 40–79. DOI: 10.1214/09-SS054.

[30] Isabella Gollini, Binbin Lu, Martin Charlton, Christopher Brunsdon, and Paul Harris. "GWmodel: An R Package for Exploring Spatial Heterogeneity Using Geographically Weighted Models". In: *Journal of Statistical Software* 63.17 (2015), pp. 1–50. DOI: 10.18637/jss.v063.i17.

[31] Yoav Benjamini and Yosef Hochberg. "Controlling the false discovery rate: a practical and powerful approach to multiple testing". In: *Journal of the Royal Statistical Society: Series B (Methodological)* 57.1 (1995), pp. 289–300. DOI: 10.1111/j.2517-6161.1995.tb02031.x.

[32] Yee Leung, Chang-Lin Mei, and Wen-Xiu Zhang. "Statistical tests for spatial nonstationarity based on the geographically weighted regression model". In: *Environment and Planning A* 32.1 (2000), pp. 9–32. DOI: 10.1068/a3162.