Draft version 2026/Mar/02

ICASX
2026
8–10 July 2026 • Kraków, Poland

AI-Augmented spatio-temporal modelling of acute malnutrition risks embedded in livelihoods systems: lessons from South Sudan

Luís Silva e Silva[1,] Hamed Dehkharghani[1], Brahima Diarra[1], Flavio Vata[1], Mary Njenga[1], Tom Oguta[1], Claudia Okley[1], Duaa Sayed[1], Matieu Henry[1], Christian Mongeau[1], Nicholas Kerandi[1], Gwenaelle Luc[1]
[1]Food and Agriculture Organization

Abstract

Acute child malnutrition in South Sudan remains persistently high and displays pronounced spatial and seasonal variability that is insufficiently captured by conventional nutrition analysis systems. This study presents a pilot proof-of-concept integrating Bayesian spatio-temporal modeling with an AI-assisted analytical interface to enhance subnational early warning systems for global acute malnutrition (GAM) prevention. A multi-scale relational database was developed linking child-level survey data (2015–2024; 150,799 children under five) with household, environmental, and conflict indicators. Acute malnutrition was modeled using a Bayesian hierarchical logistic regression with structured spatial effects and temporal dynamics. Results reveal high geographic clustering of acute malnutrition with persistent localized hotspots, with predicted incidence ranging from approximately 10% to 30% of GAM across counties. Significant associations with risks factors were identified at child level (sex and morbidity), community level (water access), and lagged environmental and conflict exposures, influencing two recurring seasonal peaks of GAM happening during the height of the dry season (pastoralism lean season) and at the end of the main rainy season (before harvesting season) systems. Model outputs were embedded within the LUMINA AI Assistant, enabling interactive exploration of datasets to detect spatial and temporal risk patterns embedded in local livelihood systems. The framework illustrates how spatio-temporal statistics and augmented intelligence can support evidence-informed nutrition early warning systems and anticipatory actions in a protracted crisis context.

1 Introduction

Acute malnutrition among children remains critically high in South Sudan, with prevalence rates often surpassing the World Health Organization's (WHO) emergency threshold of 15%, despite decades of humanitarian and development interventions [1, 2]. Relapse rates also remain significant, with up to 63% of treated children relapsing or dying within six months [3].

Simultaneously, funding for humanitarian nutrition is projected to decline substantially, with anticipated reductions of 44 percent for nutrition programmes and 49 percent for acute malnutrition treatment [4]. In this context of diminishing resources and persistent vulnerability, improving the precision, timeliness, and cost-efficiency of risk

---

[1] *The authors certify that the text, figures, and materials presented in this paper are free from any copyright violations.*

detection is essential. Failure to implement more anticipatory and locally grounded strategies may increase the risk of acute malnutrition, potentially reversing previous progress and further constraining lifesaving interventions.

Although current nutrition information systems are essential for providing snapshots of current and projected caseloads at national or large administrative levels, they remain inadequate for capturing the underlying drivers and localized dynamics of acute malnutrition. Existing risk-modelling frameworks similarly emphasize immediate and food-system-based drivers while insufficiently accounting for the structural determinants shaping food security and health outcomes. This narrow focus may lead to incomplete risk assessments and misrepresentation of vulnerability patterns, as models often rely on simplified assumptions — including linear relationships between food insecurity and nutritional outcomes. Moreover, such approaches frequently fail to capture seasonal variability and the localised clustering of risks embedded in livelihood systems and contextual factors, as emphasised in adapted conceptual frameworks for acute malnutrition [5].

Recent longitudinal evidence from African drylands further challenges conventional assumptions. Acute malnutrition does not demonstrate a straightforward spatial or seasonal alignment with food security, nor does it adhere to a single, predictable lean-season pattern. Rather, it may exhibit multiple seasonal peaks that differ across ecological zones and livelihood systems [6]. Furthermore, an expanding literature shows that acute malnutrition risks are highly clusteed at very local levels, as evidenced by the consistently high design effects observed in nutrition anthropometric surveys [7, 8]. This underscores the importance of granular, timely, and geographically targeted analysis and responses, especially in resource-constrained settings.

Addressing these complexities necessitates modelling frameworks that can simultaneously represent continuous spatial variation, temporal dependence, and lagged effects.

To address this gap, the present study introduces a Bayesian hierarchical spatio-temporal modelling framework for estimating the probability of acute malnutrition at fine spatial and temporal resolutions in South Sudan. The model integrates individual, household, and area-level predictors within a unified structure, incorporating spatial random effects to capture geographic dependence and temporal random effects to model smooth seasonal and long-term trends. By leveraging information from neighbouring administrative units and adjacent time periods, the framework enhances estimation in data-sparse and crisis-affected settings, where sampling gaps and insecurity frequently limit direct observation.

This study contributes by developing a multi-scale relational geospatial database that links child anthropometric outcomes with demographic, health, livelihoods systems, environmental, and conflict indicators. These indicators are derived from survey data and Earth observation sources. The integration enables systematic examination of lagged environmental signals, such as rainfall and vegetation anomalies, and their relationship with subsequent nutritional outcomes, thereby strengthening the statistical framework for anticipatory analysis.

Building upon this modelling architecture, the study introduces an AI-augmented analytical interface (LUMINA) designed to improve the interpretability and operational usability of complex statistical outputs. Unlike autonomous predictive systems, the AI component functions within an augmented intelligence paradigm: all quantitative inferences are produced exclusively by the Bayesian model, while the AI layer synthesizes model outputs with curated contextual evidence to support structured and transparent interpretation. This design connects advanced statistical modelling with decision-making requirements while maintaining methodological traceability and statistical integrity.

By integrating hierarchical spatio-temporal modelling with AI-assisted interpretation, this study advances agricultural and nutrition statistics in three key areas: *(i)* enhancing the precision of localized risk estimation within heterogeneous livelihood systems; *(ii)* formalizing the integration of multi-source geospatial data into a coherent statistical framework; and *(iii)* improving the accessibility of model outputs for early warning and anticipatory planning. This approach offers a scalable template for crisis-affected contexts where malnutrition risk is influenced by interacting livelihood systems, formal and informal institutions and, environmental dynamics across space and time.

The structure of the paper is as follows. Section 2 details the data sources and integration process. Section 3 outlines the spatio-temporal modeling framework and the AI-assisted architecture. Section 4 presents the empirical results and the interface implementation. Section 5 discusses the implications and limitations.

2 Data

To construct the feature space for the spatio-temporal Bayesian hierarchical model and the AI-assisted interpretation layer, we developed a multi-scale relational geospatial database integrating environmental, socioeconomic, and human security indicators with food security and nutrition survey data on child anthropometry and household characteristics.

The consolidated dataset spans 2015–2024 and includes 150,799 children under five years of age. By merging survey data with proxy indicators corresponding to each component of the conceptual framework, the system harmonizes fragmented data sources into a coherent structure, enabling fine-scale, multi-level analysis of acute malnutrition risk.

Child Acute Malnutrition outcomes were derived from nationally coordinated surveys conducted between 2015 and 2024, primarily the Food Security and Nutrition Monitoring System (FSNMS) and SMART (Standardized Monitoring and Assessment of Relief and Transitions) surveys. Both systems employ two-stage cluster sampling designs and collect standardized anthropometric measurements among children under five years of age, including weight, height (or length), and mid-upper arm circumference (MUAC). MUAC is defined as the circumference of the upper arm measured at the midpoint between the acromion and olecranon and is widely used as a field indicator of acute malnutrition. Acute malnutrition status was determined using the World Health Organization (WHO) Child Growth Standards. Acute malnutrition outcomes in this study were based on weight-for-height z-scores, bilateral pitting oedema and MUAC thresholds.

Predictors were structured according to the UNICEF conceptual framework. **Immediate drivers** included child dietary diversity, breastfeeding practices, and recent morbidity (symptoms in the past 14 days). **Underlying drivers** comprised household food security (Household Dietary Diversity Score, Reduced Coping Strategy Index, and

Household Hunger Score), the social and care environment (women's empowerment, women's dietary diversity, and women's mid-upper arm circumference), and access to health and sanitation services (health centres, hygiene practices, safe water sources, and latrine use). **Basic drivers** captured livelihood systems (production systems, agricultural and livestock assets, cultivation practices), formal and informal institutions (conflict, markets, land access, demographic factors, and household displacement status), and environmental seasonality (climate and water availability).

The FSNMS integrates modules on nutrition, food security, livelihoods, water, sanitation and hygiene (WASH), and health, providing household-level socioeconomic context alongside anthropometric outcomes. Although both survey systems are cross-sectional, repeated survey rounds over time enable reconstruction of temporal patterns when embedded within a hierarchical spatio-temporal modelling framework.

2.1 Multiscale spatial structure and geospatial data

All survey and geospatial data were organized within a multiscale spatial framework designed to preserve fine-grained variability while supporting policy-relevant aggregation. Three complementary spatial representations were used: *(i)* official administrative units (Admin 1–State, Admin 2–County, Admin 3–Payam) aligned with national reporting structures; *(ii)* a uniform 25 km² analytical grid to support scale-invariant modelling and reduce dependence on boundary changes; and *(iii)* household and community GPS coordinates linking individual anthropometric observations to local environmental and socioeconomic conditions. This configuration enables modelling at multiple aggregation levels while maintaining spatial precision at the point level.

Geospatial datasets were harmonized into four predictor domains: *(i)* environmental state indicators (topography, soils, land cover, surface water) capturing structural constraints; *(ii)* environmental stress and variability indicators (precipitation, NDVI, temperature anomalies, floods) representing seasonal and interannual dynamics; *(iii)* socioeconomic and access indicators (e.g., population density, health facilities); and *(iv)* conflict and human security indicators (conflict events, displacement).

Together, these predictors capture both structural vulnerability and short-term shocks that may influence child nutritional outcomes, enabling explicit modelling of spatial dependence and temporally lagged exposures.

2.2 Integration of the data sources

Following the comprehensive processing of geospatial and survey datasets, all derived indicators were systematically integrated into a unified relational database designed to function as a complete, analysis-ready feature matrix. This database constitutes the analytical backbone of the study, explicitly linking individual-level malnutrition outcomes with household characteristics and multi-scale geospatial predictors across both space and time. Within a relational database architecture, the schema is anchored by individual survey records, which serve as the core observational units. These records are relationally linked to household-level attributes—such as livelihood strategies, livestock ownership, and residential status—and further extended through spatial joins to grid- and administrative-level spatiotemporal indicators. By consolidating previously fragmented datasets into a coherent, normalized structure, the database enables robust, scalable machine-learning workflows while preserving the hierarchical and contextual relationships inherent in the data.

The integrated dataset has a hierarchical and multi-scale structure, connecting individual anthropometric outcomes with household, community and systems level attributes and covariates that are matched in space and time. Because the data show spatial dependence, repeated survey rounds, and lagged environmental exposures, we need a statistical approach that can handle geographic autocorrelation, seasonal changes, and the hierarchical nature of the data simultaneously. For this reason, we use a Bayesian hierarchical logistic regression model that includes spatial and temporal random effects and their interaction, as explained below.

3 AI-Augmented spatio-temporal modelling

This section explains the methods used to study acute malnutrition in South Sudan and how we interpreted the findings with interactive tools. First, we describe the spatio-temporal statistical model used to estimate associations between individual, household, environmental, and contextual factors and Acute malnutrition. This model accommodates both spatial and temporal patterns and includes delayed effects of environmental and shock variables. Next, we explain the Augmented Intelligence approach in the LUMINA AI Assistant, which combines model results with explanatory and contextual details to help users explore and interpret complex findings.

3.1 Spatial-Temporal Statistical Modelling

This study uses a Bayesian hierarchical logistic regression framework to estimate the drivers and spatio-temporal dynamics of acute malnutrition (Acute malnutrition) in South Sudan. This modelling strategy allows simultaneous estimation of fixed effects for observed covariates and latent effects capturing structured spatial and temporal variation, while accounting for correlations induced by geographic proximity and evolving temporal patterns. The main objective of this model is to generate interpretable, robust estimates of associations that can be communicated transparently through the LUMINA-AI dashboard.

3.1.1 Model Formulation

The probability of Acute malnutrition among children is modeled using a hierarchical spatial logistic regression. Let Y_{it} denote the binary Acute malnutrition outcome for child i observed at time t, taking value 1 if the child wasted and 0 otherwise. The conditional probability of Acute malnutrition is modelled on the logit scale as a function of fixed covariates, spatial random effects, and temporal random effects:

$$\operatorname{logit}(\Pr(Y_{it} = 1)) = \boldsymbol{X}_{it}\beta + f_{space}(ADM2_i) + f_{time}(t)$$

where $\boldsymbol{X}_{it}$ represents the design matrix of observed predictors for child i at time t, including individual, household, environmental, and shock-related covariates; β denotes the vector of fixed effect coefficients associated with these predictors; $f_{space}(ADM2_i)$ is a structured spatial random effect defined at the administrative level 2 (ADM2); $f_{time}(t)$ is a temporal random effect capturing smooth time (monthly) variation.

The spatial structure in the model was represented using a Conditional Autoregressive (CAR) specification for the random effects associated with administrative units at the second level (ADM2). The CAR formulation [9, 10] incorporates an adjacency matrix $\boldsymbol{W}$ that encodes neighbourhood relationships among spatial units, and a precision parameter τ_s that governs the degree of smoothing across adjacent areas. Neighbourhood structure is defined under a first-order queen contiguity criterion, whereby two spatial units (polygons) are considered adjacent if they share at least one boundary. By borrowing strength from neighbouring units, this spatial term captures residual geographic dependence not explained by the fixed covariates, while still allowing for local variation and helping to distinguish systematic spatial patterns from random noise. Temporal variation was modelled using a random walk of order 1 (RW1) on the discrete time index, with a penalised complexity prior [11] placed on the associated precision parameter τ_s. The RW1 specification encourages smooth evolution of temporal effects over the study period, accommodating trends and potential seasonal patterns that are characteristic of acute malnutrition dynamics.

The hierarchical spatial–temporal models are fitted using the Integrated Nested Laplace Approximation (INLA) framework [12]. INLA offers a computationally efficient alternative to Markov Chain Monte Carlo methods for approximate Bayesian inference in latent Gaussian models, including generalised linear mixed models and spatially structured models [13]. Its use of sparse precision matrices enables scalable estimation in the presence of structured and unstructured random effects, making it well-suited for large datasets with complex dependency structures such as children nested within households and households nested within spatial units. This capability is essential for accurately modelling the multifactorial determinants of child Acute malnutrition in South Sudan.

3.2 LUMINA AI Assistant

Unlike traditional Artificial Intelligence systems that automate analysis or decision-making, LUMINA adopts an **Augmented Intelligence** approach [14]. This approach is designed to enhance human understanding, support reasoning, and reinforce expert judgment, ensuring humans retain control over the analytical process.

The LUMINA AI **Assistant** integrates conversational and visual analytics to bridge advanced statistical modelling with practical decision-making. It features an interactive chatbot and dynamic visualisations, leveraging an AI engine that synthesises outputs from the spatio-temporal malnutrition model with relevant scientific and operational research.

By unifying components of the LUMINA project, the Assistant facilitates user engagement and contextualises complex analyses for policy application. Rather than generating new predictions, it summarises evidence, elucidates model results, and aids data exploration.

The LUMINA AI Assistant addresses three primary challenges in nutrition analytics: the misalignment between complex statistical outputs and decision-maker needs, reliance on a limited expert pool which impedes scalability, and the demand for timely, context-specific insights in humanitarian contexts. By translating technical results into accessible, narrative-driven explanations, the Assistant enhances the transparency and utility of information for practical decision-making.

3.2.1 Formulation

The AI Assistant is built upon a Retrieval-Augmented Generation (RAG) architecture, extended with a system of context-aware functions and analytical tools that are dynamically executed based on the user query and interaction context [15].

In its traditional form, RAG retrieves relevant text chunks from a knowledge base and provides them to a large language model (LLM) to ground responses in external evidence. This mechanism is directly implemented within LUMINA to retrieve contextual information from curated collections of scientific publications, technical reports, and operational guidance on acute malnutrition, livelihoods, and environmental drivers.

However, a key limitation of standard RAG approaches is that not all relevant information is natively textual. Outputs from spatio-temporal statistical models, such as posterior distributions, uncertainty intervals, spatial effects, and temporal trends are often numerical, probabilistic, or geospatial in nature, and therefore not always suitable for direct conversion into static text fragments. To address this limitation, the LUMINA AI Assistant extends the RAG paradigm through the integration of analytical functions and visualization tools that can be invoked on demand. These functions are designed to: *(i)* Query model outputs programmatically rather than retrieve pre-written text; *(ii)* Compute derived indicators (e.g. contrasts over time, spatial differentials, contribution of covariates); *(iii)* Generate contextual summaries conditioned on location, period, and livelihood system; *(iv)* Render interactive visualizations that complement narrative explanations.

In this extended formulation, the LLM acts as an orchestration and reasoning layer, determining when to retrieve textual evidence, when to call analytical functions, and how to synthesize the resulting outputs into coherent, interpretable responses. The Assistant therefore operates as an augmented analytical interface, rather than a static question-answering system.

Crucially, the AI Assistant does not perform independent inference or prediction. All quantitative results are derived exclusively from the underlying statistical model, ensuring consistency, traceability, and methodological integrity. The role of the LLM is limited to interpretation, synthesis, and explanation, guided by explicitly defined tools and controlled inputs.

4 Results

As this study represents a pilot (proof-of-concept) implementation, variable selection was constrained by data availability and harmonization feasibility. Not all covariates were included in the tested PoC model across space and time. The model therefore, includes a subset of individual-, household-, environmental-, and shock-related predictors that met coverage and quality criteria for integration into the relational database.

4.1 Statistical Model Results

Table 1 posterior mean estimates and 95% credible intervals (CrI) for fixed-effect covariates from the Bayesian spatio-temporal binomial logit model. Coefficients represent the change in the log-odds of Acute malnutrition associated with a one-unit increase in the predictor, conditional on spatial (BYM2) and temporal (RW1) random effects.

Table 1 . Posterior Mean Estimates for Fixed Effects from the Bayesian Spatio-Temporal Model. Estimates represent the change in the log-odds of a child being acutely malnourished (wasted) associated with the predictor, conditional on all other covariates, the spatial effect (Besag), and the temporal trend (RW1) remaining constant. 95% CI is the 95% Bayesian Credible Interval.

Predictor	Mean (Log-Odds)	95% CI	Interpretation
Intercept	-1.66	(-1.707, -1.613)	Represents the baseline log-odds of Acute malnutrition when all continuous covariates are zero and categorical covariates are at their reference level.
Child age (months)	-0.01	(-0.011, -0.009)	Protective Effect: For every one-month increase in age, the odds of Acute malnutrition decrease by approximately 1%.
Male child	0.19	(0.162, 0.218)	Vulnerability: Male children have approximately 21% higher odds of Acute malnutrition compared to female children (the reference group)
Diarrhea	0.251	(0.205, 0.297)	Strong Risk Factor: Recent experience of diarrhea significantly increases the odds of Acute malnutrition by about 28%.
Fever	0.084	(0.047, 0.122)	Risk Factor: Recent fever is also associated with a 9% increase in the odds of acute malnutrition
Land access (Yes)	-0.058	(-0.097, -0.019)	Protective Effect: Households with access to land show a 6% decrease in the odds of Acute malnutrition.
Precipitation	-0.094	(-0.195, 0.007)	Protective (Borderline): The current month's precipitation has a weakly protective effect. The CI just barely overlaps zero, suggesting that current high rainfall slightly reduces risk, possibly by increasing immediate water availability.
Precipitation (lag 1)	0.126	(0.026, 0.227)	Precipitation one month prior significantly increases the odds of Acute malnutrition by 13%. This likely captures the delayed negative effects of heavy rainfall, such as waterlogging, destruction of sanitation infrastructure, and increased pathogen load that spreads disease (e.g., diarrhea).
Conflict (lag 3)	0.015	(0.001, 0.030)	Delayed Risk Factor: Conflict incidence three months prior significantly increases the odds of Acute malnutrition by about 1.5%. This effect, while small, is robust and highlights the delayed impact of insecurity on livelihoods, market access, and displacement.

The results indicate that child-level morbidity (diarrhea, fever) and male sex are positively associated with Acute malnutrition, while age shows a modest protective association. At the household level, piped water access is associated with lower odds of Acute malnutrition, whereas reliance on protected wells/springs is positively associated with risk, suggesting spatially clustered vulnerability. Environmental and shock variables exhibit lagged effects: precipitation lagged one month and conflict lagged three months are positively associated with Acute malnutrition, indicating delayed pathways through environmental stress and insecurity.

Together, these findings illustrate the joint contribution of individual health, household conditions, and lagged environmental and conflict exposures to acute malnutrition risk.

4.1.1 Predicted Acute malnutrition Probability and Temporal Dynamics

The temporal random effect is modeled using a first-order random walk (RW1). This component captures the smooth evolution of Acute malnutrition risk over time after controlling for observed covariates and spatial effects. In practical terms, it represents the underlying monthly trend that is not explained by individual, household, environmental, or spatial factors.

The estimated values are expressed on the log-odds scale and should be interpreted relative to the overall study period average: Positive values indicate that the risk of Acute malnutrition during that month was higher than the overall temporal average, conditional on all other model components. Negative values indicate that the risk was lower than the study-period average. Values close to zero indicate that the estimated risk during that month was similar to the overall average temporal risk.

Because the RW1 structure imposes smoothness across consecutive months, adjacent time points are expected to show gradual changes rather than abrupt shifts. This allows the model to capture seasonal patterns and medium-term fluctuations while reducing short-term noise. As shown in Fig. **1**, the temporal random effect exhibits gradual seasonal fluctuations, with elevated log-odds of Acute malnutrition in April and August and lower values toward October, consistent with a smooth seasonal pattern.

Fig. **2** presents the posterior mean predicted probability of Acute malnutrition aggregated at the Administrative Level 2 (ADM2). The spatial distribution indicates substantial geographic heterogeneity in model-based risk, with the highest estimated probabilities concentrated in Upper Nile, Jonglei, and parts of Unity State. The full model identifies Maiwut (Upper Nile) and Fangak (Jonglei) as the highest-risk counties, each with an estimated probability of 30%, followed by Uror (27%), and Ayod, Longochuk, and Nyirol (25%). Additional high-risk areas include Melut (24%), Luakpiny/Nasir (23%), Mayom (23%), and Twic East (22%). These counties also represent substantial sample sizes, ranging from 547 to 4,023 children, supporting the stability of the estimates. Overall, predicted probabilities range from approximately 10 percent to 30 percent, highlighting persistent subnational disparities. As these estimates are derived from the complete spatio-temporal model, which incorporates fixed effects as well as structured BYM2 spatial and RW1 temporal components, they reflect both observed risk factors and residual spatial dependence. This approach highlights areas where vulnerability remains elevated after accounting for measured covariates.

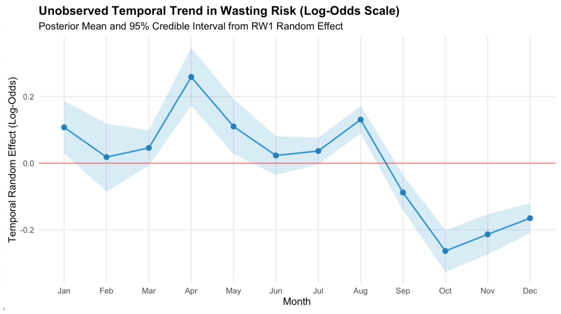

Fig. 1. Unobserved temporal trend in Acute malnutrition risk (log-odds scale). The figure displays the posterior mean (solid blue line) and 95% credible interval (shaded area) of the temporal random effect modeled using a first-order random walk (RW1).

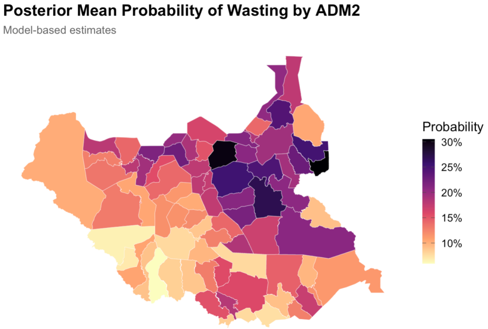

Fig. 2. Posterior Mean Probability of Acute Malnutrition (Acute malnutrition) by ADM2. The map displays the aggregated mean predicted probability of Acute malnutrition, across Administrative Level 2 (ADM2) units. This probability is derived from the complete Spatio-Temporal model.

Overall, the spatio-temporal modelling results demonstrate that acute malnutrition risk in South Sudan is shaped by the combined influence of individual health conditions, household characteristics, environmental variability, and conflict exposure, operating across space and time. The integration of structured spatial (BYM2) and temporal (RW1) components allows the model to disentangle localized clustering from broader seasonal dynamics, producing geographically differentiated and temporally coherent risk estimates. As a pilot proof-of-concept, these findings illustrate the feasibility of integrating multi-source data within a unified Bayesian framework to generate policy-relevant, subnational evidence. The results provide a quantitative foundation for geographically targeted and seasonally informed interventions, while also highlighting areas where additional data and model refinement could further strengthen inference.

4.2 AI Assistant – User-interface

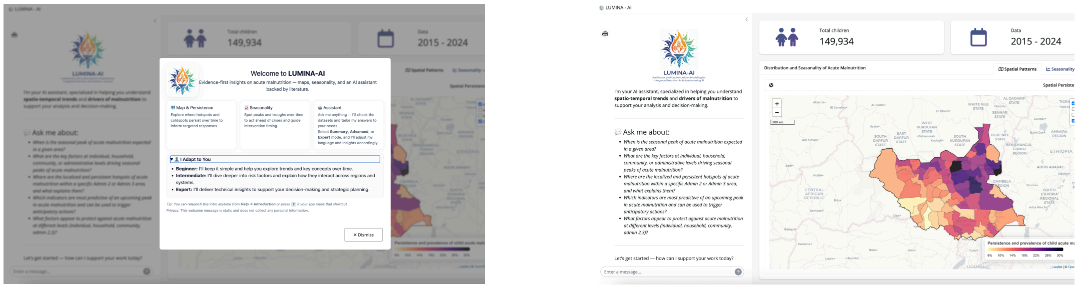

Welcome to **LUMINA-AI**

Evidence-first insights on acute malnutrition — maps, seasonality, and an AI assistant backed by literature.

Map & Persistence
Explore where hotspots and coldspots persist over time to inform targeted responses.

Seasonality
Spot peaks and troughs over time to act ahead of crises and guide intervention timing.

Assistant
Ask me anything — I'll check the datasets and tailor my answers to your needs.
Select **Summary**, **Advanced**, or **Expert** mode, and I'll adjust my language and insights accordingly.

▾ I Adapt to You

- **Beginner:** I'll keep it simple and help you explore trends and key concepts over time.
- **Intermediate:** I'll dive deeper into risk factors and explain how they interact across regions and systems.
- **Expert:** I'll deliver technical insights to support your decision-making and strategic planning.

Tip: You can relaunch this intro anytime from ***Help → Introduction*** *or press ? if your app maps that shortcut.*

Privacy: This welcome message is static and does not collect any personal information.

LUMINA-AI
Livelihoods and Undernutrition Modeling for Integrated Nutrition Anticipation using AI

I'm your AI assistant, specialized in helping you understand **spatio-temporal trends** and **drivers of malnutrition** to support your analysis and decision-making.

Ask me about:

- *When is the seasonal peak of acute malnutrition expected in a given area?*
- *What are the key factors at individual, household, community, or administrative levels driving seasonal peaks of acute malnutrition?*
- *Where are the localized and persistent hotspots of acute malnutrition within a specific Admin 2 or Admin 3 area, and what explains them?*
- *Which indicators are most predictive of an upcoming peak in acute malnutrition and can be used to trigger anticipatory actions?*
- *What factors appear to protect against acute malnutrition at different levels (individual, household, community, admin 2,3)?*

Let's get started — how can I support your work today?

Enter a message...

(a) Application welcome and navigation overlay:

(b) Spatial distribution of acute malnutrition (map view)

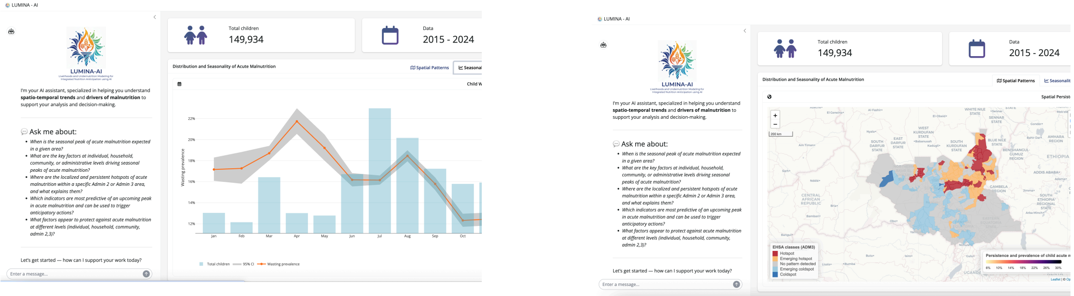

(c) Temporal evolution of acute malnutrition (time-series view):

(d) Vulnerability and contextual risk mapping

Fig. 3 LUMINA-AI application interface and analytical views.

The LUMINA interface operationalizes the spatio-temporal model outputs within an interactive analytical environment. Users can identify high-risk geographic areas and examine associated drivers through a structured dialogue with the AI assistant, enabling a combined visual and narrative exploration of results.

The assistant adapts its explanations according to three predefined user profiles—Beginner, Intermediate, and Expert—ensuring accessibility while preserving technical detail where required.

Fig**. 3** presents the core components of the interface: (a) onboarding and navigation; (b) an ADM2-level choropleth map of predicted Acute malnutrition probability; (c) a seasonality view with credible intervals; and (d) an Emerging Hotspot Analysis at ADM3 level. Together, these elements translate model outputs into an interpretable and decision-oriented analytical workflow.

Overall, the LUMINA AI Assistant demonstrates how advanced spatio-temporal modeling can be translated into an interactive and policy-oriented analytical tool. By integrating visualization, structured explanation, and adaptive user profiles within a single environment, the interface facilitates systematic exploration of spatial and temporal risk patterns. The results illustrate the feasibility of embedding model-based evidence into a practical decision-support workflow, bridging statistical analysis and operational application.

5 Conclusion

This study demonstrates the feasibility of integrating the acute malnutrition conceptual framework with Bayesian spatio-temporal modelling and AI-powered analytics to support subnational monitoring of acute malnutrition risk.The findings reveal significant differences across locations and seasons, which highlights the need for analysis that considers both geography and time. As a pilot, the LUMINA framework demonstrates how data from different sources can be brought together into a single statistical system and turned into an interactive decision-making tool.

There are also some limitations to note. The reliability of AI-assisted interpretations depends on the quality and completeness of the data and evidence used. Human oversight is needed to interpret the model results, since understanding uncertainty and causes still requires local knowledge. Ethical and governance issues, especially transparency, accountability, and data protection, are very important when using AI tools in humanitarian and policy work. The AI Assistant is not a replacement for statistical modelling or expert judgment; instead, it brings together and explains results from the spatio-temporal framework.

Future efforts will aim to include more data, improve validation steps, and strengthen human oversight to make the system more reliable and responsible. Overall, the LUMINA approach offers a clear way to turn advanced statistical modelling into evidence that is useful for policy, while keeping methods sound and results trustworthy.

The framework illustrates how spatio-temporal statistics and augmented intelligence can support evidence-informed nutrition early warning systems and anticipatory actions in a protracted crisis context.

References

1] FAO and WFP, *Hunger Hotspots. FAO–WFP early warnings on acute food insecurity,* Rome: FAO ; WFP, 2025, p. 84.

2] H. Young and A. Marshak, "Persistent global acute malnutrition," *Feinstein International Center, Tufts University,* vol. 55, 2017.

3] S. a. M. A. a. D.-G. L. a. Y. E. a. C. S. M. a. S. S. a. B. M. a. D. S. a. M. F. A. a. O. M. S. a. o. King, "Rates and risk factors for relapse among children recovered from severe acute malnutrition in Mali, South Sudan, and Somalia: a prospective cohort study," *The Lancet Global Health,* vol. 13, no. 1, pp. e98--e111, 2025.

4] S. Osendarp, M. Ruel, E. Udomkesmalee, M. Tessema and L. Haddad, "The full lethal impact of massive cuts to international food aid," *Nature,* vol. 640, no. 8057, pp. 35--37, 2025.

H. Young, A. Osman, A. Radday, A. Marshak, E. Olesambu and N. Jenkins, 5] "Improving the way we address acute malnutrition in Africa's drylands," *Field Exchange,* vol. 65, p. 14, May 2021.

6] A. Venkat, A. Marshak, H. Young and E. N. Naumova, "Seasonality of Acute Malnutrition in African Drylands: Evidence From 15 Years of SMART Surveys," *Food and Nutrition Bulletin,* vol. 44, no. 2_suppl, pp. S94--S108, 2023.

7] G. Luc, M. Keita, F. Houssoube, E. Wabyona, A. Constant, A. Bori, K. Sadik, A. Marshak and A. M. Osman, "Community clustering of food insecurity and malnutrition associated with systemic drivers in Chad," *Food and Nutrition Bulletin,* vol. 44, no. 2 suppl, pp. S69--S82, 2023.

8] G. Luc, M. Keita, B. Bazongo, B. Diarra, A. Virchenko, C. Okley, T. Seyvet, Y. Ngongang Mbunang, D. Molla, D. Dujanovic, A. Nur, A. Marshak, A. Rahim and A. M. Osman, "Special focus: Wasting patterns in Somalia," *Field Exchange Magazine 72,* 2024.

9] A. Freni-Sterrantino, M. Ventrucci and H. Rue, "A note on intrinsic Conditional Autoregressive models for disconnected graphs," 13 May 2017. [Online]. Available: https://arxiv.org/abs/1705.04854. [Accessed January 2026].

10] S. Banerjee, B. P. Carlin and A. E. Gelfand, Hierarchical Modeling and Analysis for Spatial Data, Boca Raton, FL: Chapman & Hall/CRC, 2014.

11] D. Simpson, H. Rue, A. Riebler, T. G. Martins and S. H. Sørbye, "Penalising model component complexity: A principled, practical approach to constructing priors," *Statistical Science,* vol. 32, no. 1, pp. 1 -- 28, 2017.

12] R. Håvard, S. Martino and N. Chopin, "Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations," *Journal of the Royal Statistical Society Series B: Statistical Methodology,* vol. 71, no. 2, pp. 319--392, 2009.

13] M. Blangiardo and M. Cameletti, Spatial and Spatio-temporal Bayesian Models with R-INLA, John Wiley & Sons, 2015.

14] K.-L. A. Yau, H. J. Lee, Y.-W. Chong, M. H. Ling, A. R. Syed, C. Wu and H. G. Goh, "Augmented Intelligence: Surveys of Literature and Expert Opinion to Understand Relations Between Human Intelligence and Artificial Intelligence," *IEEE Access,* vol. 9, pp. 136744-136761, 2021.

15] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel, S. Riedel and D. Kiela, "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks," *Advances in Neural Information Processing Systems,* pp. 9459--9474, 2020.

16] A. Marshak, A. Venkat, H. Young and E. N. Naumova, "How Seasonality of Malnutrition Is Measured and Analyzed," *Int J Environ Res Public Health,* vol. 18, no. 4, p. 1828, Feb 2021.

17] C. Maitra, A review of studies examining the link between food insecurity and malnutrition, FAO, 2018.