AGILE: GIScience Series, 7, 16, 2026. https://doi.org/10.5194/agile-giss-7-16-2026
Proceedings of the 29th AGILE Conference on Geographic Information Science, 16–19 June 2026.
Eds.: Evelyn Uuemaa, Alexander Kmoch, Holger Virro, and Simon Scheider
This contribution underwent peer review based on a full paper submission.

Spatio-Temporal Knowledge Graph from Unstructured Texts: A Multi-Scale Approach for Food Security Monitoring

Charles Abdoulaye Ngom [1,2,3], Landy Rajaonarivo [1,2], Maguelonne Teisseire [1,2], and Sarah Valentin [2,4]

[1]INRAE, F-34398 Montpellier, France
[2]TETIS, Univ. Montpellier, AgroParisTech, CIRAD, CNRS, INRAE, Montpellier 34090, France
[3]Université de Montpellier, Montpellier, France
[4]CIRAD, F-34398, Montpellier, France

Correspondence: Charles Abdoulaye Ngom (charles.ngom@inrae.fr)

Abstract.
Food security monitoring in West Africa requires timely, fine-grained, and interpretable information to support decision-making and crisis prevention. However, current approaches rely largely on structured survey data (e.g., FEWS NET), which, while valuable, do not capture the fine-grained, local-scale spatial and temporal dynamics that can be extracted from unstructured textual sources. To address these limitations, we propose an integrated framework combining natural language processing (NLP) and knowledge graph modeling to extract, structure, and analyze food security-related information from press articles. This representation enables multi-scale spatial queries and temporal pattern detection through the knowledge graph, allowing the identification of cascading or co-occurring risks. Applied to a West African country, the framework demonstrates how unstructured textual data can complement conventional food security assessments by providing both localized insights and broader spatio-temporal perspectives, highlighting the combined value of NLP-based extraction and knowledge graph reasoning for food security monitoring.

Submission Type. Full paper track, model, algorithm, case study

BoK Concepts. [GC7] Artificial intelligence, [AM7] Spatio-temporal analysis, [CV6] Thematic mapping, [DA5] Knowledge discovery

Keywords. Knowledge Graphs, Food Security, Spatial Information Extraction, Natural Language Processing

1 Introduction

The food security situation is deteriorating in many countries around the world, particularly in the Global South. This issue is one of the Sustainable Development Goals[1] (SDGs), specifically SDG 2, which calls on countries around the world to adapt their policies and legislation to eliminate hunger, ensure food security, improve nutrition, and promote sustainable agriculture. International organizations such as the United Nations (UN), Food and Agricultural Organization (FAO), United Nations High Commissioner for Refugees (UNHCR), and United Nations Children's Fund (UNICEF) are working with governments and local authorities to guide the policies and programs adopted to combat food crises. To do this, economic and geospatial data are analyzed. Questions arise when analyzing food security, such as: Where do people who are food insecure or vulnerable live? Why are they exposed to this situation? How is the situation evolving over time and space, and what are the risks threatening them? A network of more than 150 analysts around the world is seeking answers to these key questions through the World Food Programme (WFP). This program aims to provide concrete information on food security in each country where it operates. It thus helps identify populations affected by food insecurity around the world and determine the causes of this insecurity. However, assessments take time and require the collection of information on several aspects, such as food consumption, the economic and political situation, natural risks, etc., from various sources. The WFP has set up a platform[2] for classifying levels of insecurity in order to visualize, via mapping, the levels of food insecurity for regions in different countries to aid decision-making.

[1]https://sdgs.un.org/goals
[2]https://www.ipcinfo.org

Other actors are working on food crisis prevention, such as the Food Crisis Prevention Network[3] (RPCA). They have set up a platform for dialogue and policy coordination. The network provides data and evidence-based analysis on food and nutrition security for 17 countries in the Sahel and West Africa. This platform aims to improve countries' preparedness to prevent and effectively manage food crises. However, it is largely absent from regional and national political debates, partly due to the lack of available data on the reality of the situation. It is therefore essential to have reliable data on the factors and dynamics of food crises in the regions to anticipate the impacts on food security and inform food and agricultural policies.

Some approaches demonstrate the value and importance of capturing information on food security, the geographical location of crises, and the production of analytical indicators from informal data sources(Carneiro et al., 2024; Ahn et al., 2023; Deleglise et al., 2024). Real-time and local information is crucial for risk management, as it can save lives and prevent food insecurity.

In this paper, we present an approach for extracting spatio-temporal and thematic entities from news articles and exploiting knowledge graphs for food crisis monitoring. The study area of this research work is Burkina Faso, a country in West Africa. Our main contributions focus on:

- Detecting fine-grained geographic entities (street, village, municipality) in limited resources;
- Representing data using knowledge graphs to exploit the representation of semantic links between data in the food security domain;
- Exploiting knowledge graphs to generate answers to questions that stakeholders (local authorities, international organizations, ministerial institutes, etc.) may ask, or to produce cartographic and thematic visualizations to obtain a comprehensive and detailed view of the available data and perform analyses.

This paper is organized as follows: Section 2 reviews the state of the art in spatial and temporal entity extraction, event argument extraction, and knowledge graphs for food security. Section 3 presents our methodology, including data collection, the triplet formation algorithm, and knowledge graph construction. Section 4 evaluates the pipeline components and demonstrates the exploitation of the knowledge graph for food security analysis. Finally, Section 5 summarizes our contributions and discusses future directions.

[3]https://www.oecd.org/en/networks/food-crisis-prevention-network-rpca.html

2 Related Work

The construction of spatio-temporal knowledge graphs from unstructured texts for food security monitoring requires addressing three interconnected challenges: the extraction of spatial and temporal entities from texts with varying levels of geographic granularity, the association of these spatial and temporal entities with thematic information through implicit relations dispersed across documents, and the representation of this information in a structured format that enables temporal tracking and spatial analysis.

2.1 Extraction of spatial and temporal entities

Spatial entity extraction, or geoparsing, combines toponym recognition (identifying place names) and toponym resolution (linking them to geographic coordinates or gazetteer entries) (Middleton et al., 2018). It is critical for food security monitoring, as the impact of crisis triggers depends on their spatial extent and granularity. Village-cluster level models can identify 83%–99% of the most food-insecure populations, outperforming provincial-scale approaches (McBride et al., 2022). Fine-scale localization is essential where food insecurity arises from localized factors such as market disruptions, conflicts, or micro-climatic variability. Spatial typologies support targeted interventions by accounting for local heterogeneity, crucial in West Africa, where thousands of villages may require differentiated responses (Marivoet et al., 2019).

Despite progress, existing spatial extraction methods struggle to achieve the granularity needed for food security. Transformer-based models have improved NER—CamemBERT for French (Martin et al., 2020), GLiNER for task-agnostic extraction (Zaratiana et al., 2024), GeoLM and TopoBERT for geographic entities in English (Li et al., 2023; Zhou et al., 2023)—but in African contexts, training data bias toward cities under-represents villages (Adelani et al., 2021), geographic databases lack sub-provincial coverage, and spelling variations reduce accuracy.

To address these issues, we propose an approach enabling village-level spatial entity extraction, satisfying operational needs for fine-grained, location-specific monitoring (Marivoet et al., 2019). This supports targeted analysis of, for instance, villages affected by market disruptions rather than entire provinces, improving response precision. For temporal extraction, we use HeidelTime (Strötgen et al., 2013), which supports French TIMEX3 normalization with performance comparable to English and Spanish (Moriceau and Tannier, 2014). Neural alternatives such as TEI2GO (Sousa et al., 2023) and GeospaCy (Mehtab Alam et al., 2024) trade accuracy for speed or focus on English. Given limited annotated West African data, we adopt fine-tuned CamemBERT for spatial extraction and HeidelTime for temporal normalization.

2.2 Event Argument Extraction

While identifying where and when food security events occur provides essential spatial and temporal anchors, this information becomes actionable only when coupled with understanding what specific indicators characterize each crisis. Although a press article mentions a risk to food security (e.g., market disruptions, displacement, reduced harvests), the challenge lies in identifying which geographical areas are affected among all those mentioned in the article. This challenge relates to Event Argument Extraction (EAE), which identifies the elements (including locations and times) that fill specific roles in detected events (Li et al., 2021). Recent approaches to EAE at the document level employ various strategies: generative methods that formulate extraction as conditional text generation (Li et al., 2021), classification-based approaches modeling relationships between entities and events (Wang et al., 2022), prompt-based methods using selective scope mechanisms (Ma et al., 2022), and approaches leveraging large language models (Zhang et al., 2024; Shuang et al., 2024).

However, these methods present fundamental limitations for operational food security monitoring in resource-constrained contexts: they require significant memory and processing time, and cannot detect precise geographical information such as coordinates or administrative hierarchies essential for food security analysis. Humanitarian organizations require systems capable of determining whether a food shortage is limited to a village, spans neighboring communities, or warrants interventions at provincial or village-cluster levels. This entails identifying place names while preserving normalized coordinates, administrative hierarchies, and spatial relationships. Existing EAE methods often ignore such structure, struggle with implicit spatial references (e.g., "northern villages"), and fail to capture relations requiring geographical knowledge such as boundaries, proximity, or regional typologies. They also face challenges with transitivity, long-distance dependencies, and hallucination, reducing reliability and risking misallocation of humanitarian resources (Zhang et al., 2024).

Existing approaches depend on large annotated datasets linking events to their spatial arguments, unavailable for food security in the French-speaking West African press, and creating such data would require costly expert annotation. To address these gaps, rather than learning generic syntactic relations, we construct a knowledge graph via a projection-based approach integrating three specialized components.

2.3 Knowledge Graph for food security

Knowledge graphs (KGs) offer a unified and semantically structured framework for integrating heterogeneous food security data, including climatic, agricultural, economic, and sociopolitical information. By explicitly modeling entities (e.g., locations, hazards, indicators) and their relations, KGs enable the identification of complex patterns such as cascading shocks and multi-factorial drivers of food insecurity. The importance of textual data in anticipating crises has been demonstrated at scale: Balashankar et al. (2023) analyzed 11.2 million news articles across 21 food-insecure countries, showing that text-based signals improve crisis prediction up to 12 months ahead. However, these approaches rely on aggregated embeddings and temporal correlations, without explicit event structures, spatial hierarchies, or interpretable causal chains. Similarly, *SemEval-2025* Task 9 focused on hazard classification without addressing spatio-temporal or event-based modeling (Randl et al., 2025).

Recent work highlights the added value of structured representations for humanitarian and food security analysis. Large-scale humanitarian KGs integrate multi-source crisis data through explicit entity–relation modeling, supporting interoperability, traceability, and downstream analytics (Consoli et al., 2025). In the food security domain, geospatial KGs connect environmental indicators, administrative units, and socio-economic variables, enabling spatial queries and cross-scale analyses (Liu and Kuo, 2024). Structured representations also support interpretable composite indicators, as shown by Carneiro et al. (2025), who combine remote sensing, market, and conflict data to capture both short-term shocks and long-term vulnerabilities. Despite these advances, most existing approaches remain limited by static relations, aggregated indicators, and weak modeling of temporal ordering and event recurrence.

Building on these foundations, food security indicators can be automatically derived once a KG is constructed: encoding observations and events as graph nodes and relations makes composite metrics queryable and explainable, such as recurrent droughts inferred from weather-event sequences or price-spike frequencies from market relations. Our framework introduces specialized relation types to explicitly model temporal order and periodically recurring events.

3 Proposed approach

This section describes our methodology for building spatio-temporal knowledge graphs from French news articles to monitor food security in Burkina Faso. The approach addresses three key challenges: (i) detecting fine-grained geographic entities down to the village level in a context where standard NLP resources poorly cover

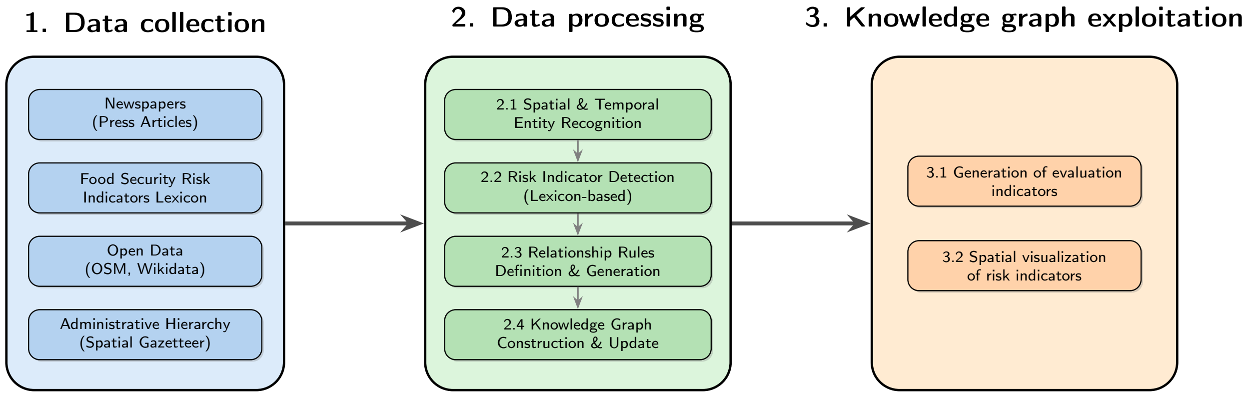

Figure 1. Global architecture of the proposed approach.

African toponyms, (ii) correctly anchoring risk indicators in space and time despite the frequent use of implicit references in journalistic texts, and (iii) organizing the extracted information into a queryable knowledge graph that supports both historical analysis and operational decision-making.

Our approach comprises three main modules (Figure 1): (1) data collection, (2) data processing, and (3) knowledge graph exploitation. As our primary contributions lie in data extraction and representation, we focus on the data processing stage, particularly tasks 2.1 (*Spatial and Temporal Entity Recognition*), 2.3 (*Relationship Rules*), and 2.4 (*Knowledge Graph Construction and Update*).

3.1 Data Collection

This study relies on multiple data sources. The corpus comprises 15,000 unannotated French news articles published in 2009, collected via web scraping from major West African news outlets (e.g., LeFaso.net, Burkina24). A domain-specific lexicon of food security risk indicators was developed by domain experts who are co-authors (particulary Maître d'Hôtel, E., Bégué, A.) : researchers/experts in food security of this paper, building on the IPC/Cadre Harmonisé framework. The lexicon is publicly available Deléglise et al. (2021) and comprises 433 unique entries organised into 92 concepts and eight thematic domains, with each entry iplinked to a single concept and its corresponding risk phase.

The administrative hierarchy used as our spatial reference is derived from the 2022 Statistical Yearbook of Territorial Administration published by the Burkina Faso Ministry of Territorial Administration[4]. It provides the complete administrative structure: 13 regions, 45 provinces, 351 departments, and 7,936 villages. Despite spanning 2009–2022, Burkina Faso's administrative hierarchy remained stable at regional, provincial, and departmental levels.

[4]http://cns.bf/IMG/pdf/matds_annuaire_at_2022.pdf

3.2 Data Processing

Data processing consists of spatio-temporal and thematic entity extraction, relationship definition, and the construction and updating of knowledge graphs.

3.2.1 Spatial and temporal entity recognition

Recognizing fine-grained geographic entities in West African news articles presents unique challenges, including the underrepresentation of African toponyms in standard NER training corpora and significant orthographic variability in place names (e.g., "Ouagadougou" vs. "Ouaga"). To address these limitations, we fine-tuned a CamemBERT model (Martin et al., 2020) on the complete administrative hierarchy of Burkina Faso. The fine-tuning process uses the official gazetteer to automatically generate training data through distant supervision, i.e., automatically generating labeled training data by matching entity mentions in text against entries in the official gazetteer, followed by a manual correction of the automatically generated annotations. The training corpus was split into 70% for training, 15% for validation, and 15% for testing. Crucially, the test set includes previously unseen entities: 19.9% of villages, 1.1% of departments, and 2.9% of provinces in the test set never appeared during training, enabling evaluation of the model's generalization to new toponyms. During fine-tuning, the first 10 encoder layers were frozen to preserve CamemBERT's pre-trained French representations, updating only the last 2 encoder layers and the classification head. Training used the AdamW optimizer with a learning rate of 5e-5, a batch size of 8, and early stopping (patience = 2) over 5 epochs.

Human validation. To ensure annotation quality, the automatically generated spatial annotations underwent systematic human review. Domain experts manually validated detected entities, correcting recognition errors and resolving ambiguous cases where place names could refer to multiple locations (e.g., homonymous village names across different provinces). This human-in-the-loop validation was essential given the orthographic variations common in West African toponyms and the presence of place names that coincidentally match common French words.

For temporal expressions, we adopt HeidelTime, a rule-based system supporting French that normalizes expressions according to the TIMEX3 standard. HeidelTime requires a document creation time (DCT) to resolve relative temporal expressions such as "yesterday," "last week," or "next month." We configure HeidelTime to use the publication date of each article as the DCT, ensuring accurate normalization of the relative temporal expressions that predominate in news articles.

3.2.2 Risk indicator detection.

From the domain-specific lexicon of food security risk indicators (Section 3.1), this step uses lexicon-based matching to identify food security risk indicators within the text. To account for the lexical variations commonly found in press writing, we apply a rule-based morphological expansion to each lexicon entry. Specifically, for each term t in the lexicon, we systematically generate three suffix-based variations: the plural form (t+*s*), the feminine/extended plural form (t+*es*), and the elided form (t+'), in addition to retaining the base form. Each detected term is mapped back to its base lexicon entry and associated with its corresponding concept, theme, and crisis phase as defined in the expert-validated lexicon.

3.2.3 Triplet formation algorithm

With risk indicators and spatio-temporal entities detected, we employ a projection-based algorithm to form complete triplets (Risk Indicator, Location, Date). This algorithm addresses two key challenges: filtering lexically ambiguous terms that generate false positives, and resolving implicit spatial and temporal references common in journalistic writing, where locations and dates are established early and subsequently implied rather than repeated.

Vague term filtering. Certain terms in the food security lexicon exhibit high ambiguity when used without a qualifying context. We identified eight vague terms: *production*, *campaign / agricultural season*, *price*, *cost*, *crisis*, *stock / inventory*, *rain*, and *rainfall* that frequently appear in non-relevant contexts (e.g., discussing general economic trends rather than specific food security events). These eight terms were identified through an iterative process: after initial triplet extraction, domain experts reviewed the most frequent false positives and flagged terms that systematically appeared in non-food-security contexts. To mitigate false positives, we apply conditional filtering: a vague term is retained only if it satisfies one of two conditions:

(i) **Co-occurrence with non-vague term**: The vague term co-occurs with at least one non-vague term from the lexicon within the same sentence. For example, "price" alone would be filtered, but "cereal prices" (where "cereal" is a non-vague food security term) would be retained.

(ii) **Presence of variation trigger**: The vague term is associated with a linguistic marker expressing a directional change, indicating either an increase or a decrease. Variation triggers include augmentation-related terms (e.g., *increase*, *rise*, *growth*, *improvement*, *upward trend*, *elevation*, *intensification*, *increase*, *progression*) and diminution-related terms (e.g., *decrease*, *drop*, *reduction*, *fall*, *decline*, *setback*, *regression*, *degradation*, *deterioration*, *weakening*), including their conjugated forms.

This filtering ensures that ambiguous terms like "price" are included only when contextually anchored to a specific variation event (e.g., "price increase"), reducing noise while preserving semantically meaningful indicators.

Spatial projection (SP). When a sentence contains a relevant risk indicator, the algorithm associates it with the most specific spatial entity available, following a hierarchical preference order: $L_v \prec L_d \prec L_p \prec L_r \prec L_c$, where L_v, L_d, L_p, L_r, and L_c denote village, department, province, region, and country respectively. The projection operates as follows:

- SP1: If one or more locations are detected in the current sentence, select the one with the finest granularity (lowest administrative level).
- SP2: If no location is found in the current sentence, propagate the most recent location from the preceding context. The algorithm searches up to two sentences forward from the last known location.
- SP3: If no location can be resolved within the local context, default to the country level (Burkina Faso).

This projection mechanism addresses the common journalistic pattern where locations are introduced at the beginning of an article or paragraph and subsequently implied rather than explicitly repeated. **Temporal projection (TP).** Temporal resolution follows an analogous inheritance strategy:

- TP1: If temporal expressions are detected in the current sentence, select the date closest to (but not exceeding) the article's publication date. This constraint ensures temporal plausibility by excluding future dates that may appear in articles discussing planned events because our framework focuses on reported events rather than planned or anticipated ones.

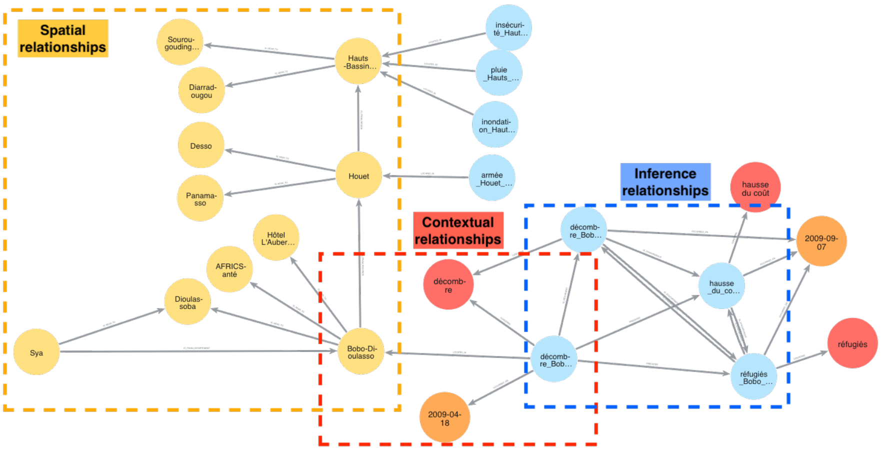

Figure 2. Extract of knowledge graph showing three types of relationships: (1) spatial relationships (yellow box), (2) contextual relationships (red box), and temporal inference relationships (blue box).

- TP2: If no temporal expression is found in the current sentence, inherit the most recent date from preceding sentences.
- TP3: If no explicit temporal expression can be resolved, use the article's publication date as the default temporal anchor.

Context enrichment. To preserve interpretability and support downstream validation, each extracted triplet is augmented with a context window comprising the source sentence plus one adjacent sentence before and after. This enriched context provides analysts with sufficient surrounding text to assess the relevance and meaning of detected indicators without requiring access to the full article.

Geographic metadata enrichment. Each extracted location is enriched with metadata retrieved from Wikidata, including geographic coordinates (latitude, longitude) and the official Wikidata identifier (Q-code).

3.2.4 Annotation campaign for triplet evaluation

To evaluate the triplet formation algorithm, we conducted an expert annotation campaign on a stratified sample of extracted triplets. Annotators assessed three dimensions:

(i) **Term relevance**: Correct identification of risk indicators as food-security relevant, excluding hypothetical, preventive, or unrelated mentions.

(ii) **Spatial association**: Correct linkage of the term to the actual location of the risk event, not contextual or comparative locations.

(iii) **Temporal association**: Correct linkage of the term to the event date, excluding historical or future references.

This evaluation allowed separate measurement of lexicon-based detection, spatial and temporal projections, and helped identify systematic errors and edge cases. Results guided refinements of the vague term list, trigger vocabulary, and projection windows.

3.2.5 Knowledge Graph construction and update

This step involves constructing and updating the KG using the formed triplets. The resulting knowledge graph comprises four node types and three categories of relationships designed to capture both explicit and inferred spatio-temporal patterns, as illustrated in Figure 2.

This figure provides an overview of the complete knowledge graph structure centered on Bobo-Dioulasso, with the three relationship categories highlighted in distinct boxes: spatial relationships (yellow dashed frame), contextual relationships (red dashed frame), and temporal inference relationships (blue dashed frame). The color coding distinguishes entity types throughout the graph: yellow for locations, orange for temporal entities, blue for events, and red for thematic risks.

Node types. The knowledge graph comprises four node types. *Risk nodes* represent thematic indicators and store the theme, concept, and crisis phase. *Location nodes* encode geographic entities with their administrative level (country, region, province, department, or village), coordinates, and Wikidata identifier. *Time nodes* represent normalized temporal entities with year, month, and day attributes. *Event nodes* model food security situations and are defined as structured triplets (*Risk, Location, Time*), capturing what happened, where it occurred, and when it took place. Each event corresponds to a unique combination of these three nodes and includes additional attributes such as article frequency, publication dates, and supporting contextual sentences.

Relationship categories. The nodes are connected through three categories of relationships: spatial relationships, contextual relationships, and temporal inference relationships (Figure 3).

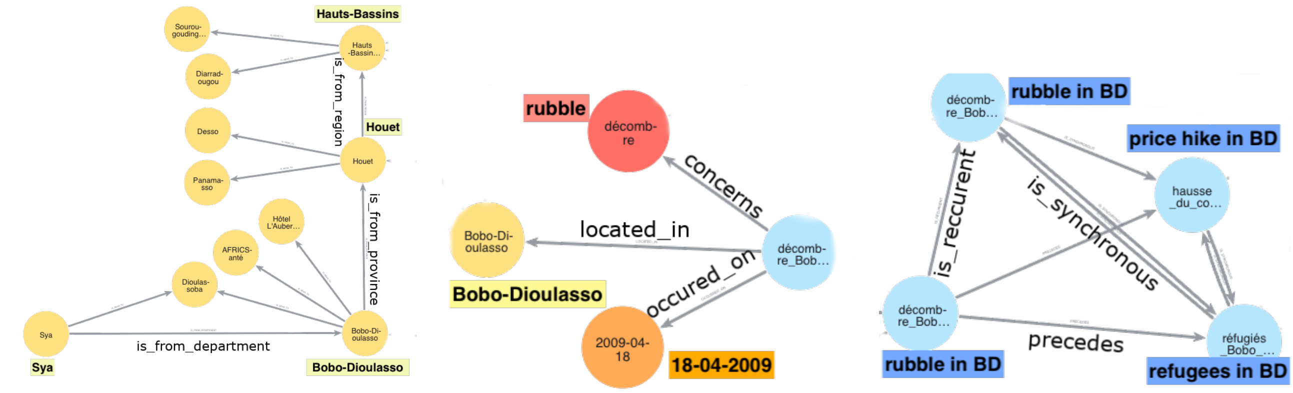

(a) Spatial relationships between location-type nodes

(b) Contextual relationships

(c) Temporal inference relationships between situation-type nodes (BD: Bobo-Dioulasso)

Figure 3. Illustration of the extracted knowledge graph relationships: spatial, contextual, and temporal inference views.

Spatial relationships encode the administrative hierarchy of Burkina Faso, linking locations to their parent administrative units through three edge types: `IS_FROM_DEPARTMENT` (village → department), `IS_FROM_PROVINCE` (department → province), and `IS_FROM_REGION` (province → region). Figure 3a illustrates these hierarchical connections: the village of Sya is linked to its parent department Bobo-Dioulasso via `IS_FROM_DEPARTMENT`, Bobo-Dioulasso is connected to its province Houet via `IS_FROM_PROVINCE`, and Houet is linked to the Hauts-Bassins region via `IS_FROM_REGION`. This hierarchical structure enables multi-scale spatial queries, allowing analysts to aggregate events from village level up to regional level or to drill down from regional patterns to specific localities.

Contextual relationships link each event node to its three constituent entities, materializing the triplet structure within the graph. As shown in Figure 3b, the event node "décombre_Bobo-Dioulasso" (rubble in Bobo-Dioulasso) is connected to: (i) its associated risk type ("rubble/décombre") through the `CONCERNS` relationship, (ii) its spatial anchor (Bobo-Dioulasso) through the `LOCATED_IN` relationship, and (iii) its temporal anchor (2009-04-18) through the `OCCURRED_ON` relationship. These three relationships constitute the core semantic structure that answers the fundamental questions of food security monitoring: what risk occurred, where, and when.

Temporal inference relationships capture event dynamics and enable the discovery of patterns across time. Figure 3c illustrates the three types of inter-event relationships:

- `IS_RECURRENT` connects events sharing the same risk type and location but occurring at different times, with a *duration_days* attribute indicating the interval between occurrences. In the figure, the two "rubble in BD" events from April and September 2009 are linked by this relationship, indicating a recurring phenomenon.
- `PRECEDES` links temporally ordered events at the same location involving different risk types, enabling the identification of potential cascading effects. For example, the "rubble in BD" event precedes the "refugees in BD" event, suggesting a possible causal sequence where infrastructure damage leads to population displacement.
- `IS_SYNCHRONOUS` connects events occurring simultaneously at the same location but involving different risk types, capturing co-occurring phenomena that may indicate compound crises. This bidirectional relationship identifies situations where multiple risks affect a location at the same time, potentially amplifying their combined impact.

3.3 Data and Software Availability

To ensure reproducibility and support future research, the resources associated with this work are publicly available. The complete pipeline implementation, including spatial entity recognition, triplet formation, and knowledge graph construction, is released on GitHub.[5] The fine-tuned CamemBERT model for spatial entity recognition in Burkina Faso is available on Hugging Face,[6] along with the associated training and evaluation datasets.[7] These resources include the domain-specific food security lexicon, the Burkina Faso administrative gazetteer, the extracted knowledge graph, and the expert-validated triplet evaluation dataset. Due to copyright restrictions on the press articles used in this study, the full-text corpus cannot be made publicly available. Researchers wishing

[5]https://github.com/CharlemagneBrain/STKG-FS

[6]https://huggingface.co/CharlesAbdoulaye/BF_NER

[7]https://huggingface.co/datasets/CharlesAbdoulaye/BF_NER_datasets

to reproduce the experiments on the complete corpus are invited to contact the authors to arrange access.

4 Results and Discussion

This section presents the evaluation results for each component of our pipeline: spatial entity recognition, relevant indicator classification, and knowledge graph exploitation.

4.1 Spatial entity recognition

To evaluate spatial entity recognition, we first compared two baseline approaches before fine-tuning: CamemBERT and GLiNER, a generalist zero-shot NER model.

Before presenting the fine-tuning results, we briefly define the evaluation metrics. *Precision* measures the proportion of correctly identified entities among all entities detected by the model; high precision indicates few false alarms. *Recall* measures the proportion of actual entities that the model successfully detected; high recall indicates few missed entities. The *F1-score* is the harmonic mean of precision and recall, providing a single balanced metric ranging from 0 to 1.

To ensure evaluation focused specifically on Burkina Faso locations, we applied a post-processing filter using OpenStreetMap to retain only entities within the country's administrative boundaries. On a test corpus of 1,000 press articles, baseline CamemBERT achieves high recall (0.81) but low precision (0.41; $F_1 = 0.55$), producing many false positives. GLiNER offers a better balance (precision 0.66, recall 0.63; $F_1 = 0.65$) but struggled to converge on the same data volume, limiting its adaptation to our domain-specific gazetteer. As missing a location may hide an emerging crisis signal, recall remains critical, yet improved precision is necessary to reduce noise. This motivated fine-tuning CamemBERT on the full administrative hierarchy of Burkina Faso to enhance overall performance. Table 1 presents the performance of fine-tuned CamemBERT, evaluated on a held-out test set containing approximately 20% of unseen entities at each hierarchical level. To maximize coverage of underrepresented locations, particularly villages, the fine-tuning was performed on 15,000 press articles. The model achieves near-perfect performance for higher administrative levels (country, region, province, department), with F1-scores of 0.99 across all four categories. Performance for village-level entities remains strong (F1 = 0.94), though slightly lower due to the greater diversity of village names (over 7,900 distinct entries) and the presence of previously unseen villages in the test set (19.9% of test villages did not appear in training data). These results demonstrate that domain-specific fine-tuning effectively addresses the challenge of recognizing African toponyms underrepresented in general-purpose NER models. At the village level, precision (0.94) and recall (0.93) are almost equal, with the occasional confusion arising from common words matching village names, and a small number of entities being missed among previously unseen toponyms. This is an acceptable trade-off for early warning systems, where both false alarms and missed detections incur operational costs.

Table 1. Performance of fine-tuned CamemBERT for spatial entity recognition across administrative levels in Burkina Faso.

Entity Type	Precision	Recall	F1-score
Country	0.99	0.99	0.99
Region	1.00	0.99	0.99
Province	0.99	0.98	0.99
Department	0.99	0.99	0.99
Village	0.94	0.93	0.94

4.2 Evaluation of the triplet formation algorithm

Table 2. Evaluation of the triplet formation algorithm on 1,133 annotated instances. Each triplet was assessed for correct spatial and temporal association.

Category	Count	Percentage
Location and Date correct	928	81.91%
Location corrected only	166	14.65%
Date corrected only	27	2.38%
Both corrected	12	1.06%
Total	1,133	100%

To evaluate the projection-based triplet formation algorithm, we annotated 1,133 extracted triplets. Experts assessed each triplet on two dimensions: (i) correct spatial association of the risk indicator with its location, and (ii) correct temporal association with its date. Table 2 summarizes the results. Overall, 81.91% of triplets were fully correct. Temporal projection performed particularly well, with 96.56% accuracy (1,094 correct associations), benefiting from HeidelTime's robust normalization of temporal expressions in French news and the publication date as a default anchor. Spatial projection reached 84.29% accuracy (955 correct associations), highlighting the main area for improvement. Analysis of 178 spatial errors reveals two patterns:

- **Over-aggregation to capital or country level**: The algorithm often defaulted to "Ouagadougou" or "Burkina" when more specific locations appeared outside the two-sentence projection window, reflecting the journalistic tendency to present national context before local details.
- **Ambiguous location references**: Errors occurred when multiple locations of similar granularity were mentioned, leading the algorithm to select the wrong one based on proximity rather than semantic relevance.

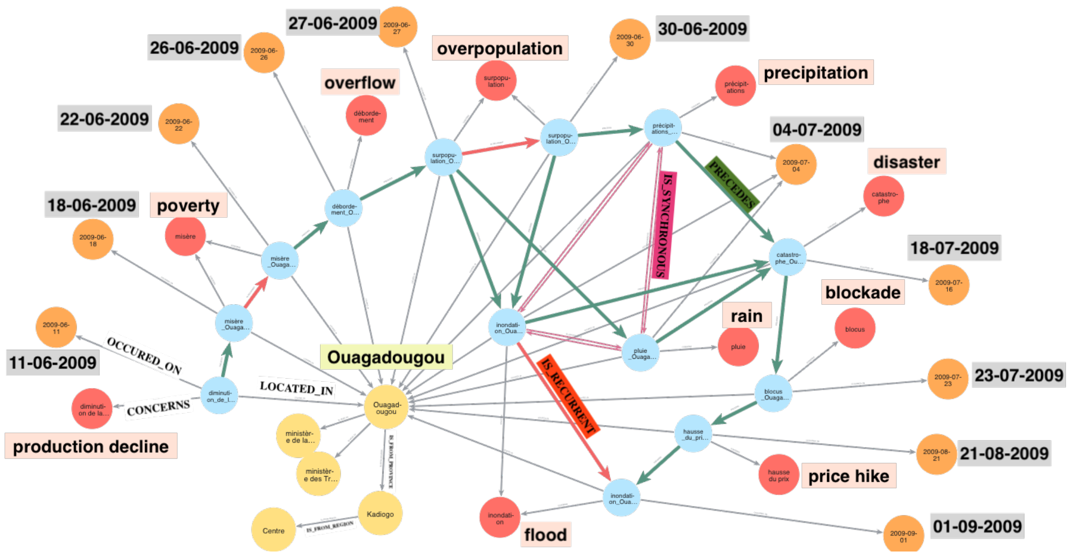

Figure 4. Knowledge graph extract: temporal monitoring of events in the Ouagadougou department (the color of edges between situation-type nodes indicates their type of inference: green for the relation PRECEDES, red for IS_RECURRENT, and double-direction pink edge for IS_SYNCHRONOUS)

Multi-location associations. During annotation, we observed that some risk terms were semantically associated with multiple spatial entities, rather than only with the finest-grained location assumed by the algorithm. To account for this, a post-processing step expanded the triplet set: when annotators indicated that an event affected multiple locations simultaneously, additional triplets were generated with the same risk term and date but different spatial anchors. This process added 21 triplets, yielding a total of 1,154 annotated triplets.

Relevance filtering. Among the 1,154 triplets, expert annotators assessed whether each detected risk indicator was genuinely relevant to food security in context. The evaluation identified 541 triplets (46.9%) as relevant and 613 as non-relevant. The non-relevant triplet rate reflects the inherent challenge of lexicon-based detection in journalistic texts, where food security terminology frequently appears in policy discussions, prevention campaigns, or historical reviews rather than in actual crisis reports. For extension to other press corpora, the domain-specific lexicon provides a transferable starting point for term detection, but integrating an automatic relevance classifier such as a transformer-based binary model fine-tuned on the 1,154 expert-annotated triplets with supervised contrastive learning Gunel et al. (2021) would be necessary to filter contextual mentions from genuine crisis reports.

Event aggregation. Before knowledge graph construction, we merged triplets referring to the same real-world event. Two or more triplets are fused into a single event node if they share the same risk term, normalized location, and normalized date. During fusion, article identifiers, enriched contexts (up to 10 per event), mention frequencies, and publication dates are aggregated. After aggregation of the 541 relevant triplets, the knowledge graph contains 376 unique events, 77 distinct risk types, 71 distinct locations, 147 distinct dates, and 200 source articles.

Corpus statistics. The distribution of location types shows that department-level entities (44.2%) and country-level references (34.7%) dominate, followed by villages (12.8%), regions (6.7%), and provinces (1.6%). The most frequent locations are Burkina (402 occurrences), Ouagadougou (319), and Bobo-Dioulasso (41), reflecting press coverage bias toward the capital. Thematically, sociopolitical events (29.9%) and environmental events (23.8%) are most prevalent, followed by agricultural (16.3%), general crisis (11.4%), and economic (10.9%). The top risk indicators include *inondation* (*flood*, 28.1%), *catastrophe* (*disaster*, 13.1%), *insécurité* (*insecurity*, 10.9%), *corruption* (*corruption*, 10.9%), and *incendie* (*fire*, 10.9%). These results confirm the effectiveness of the projection-based triplet formation approach. By combining spatial hierarchy preferences, temporal inheritance, and vague term filtering, the proposed heuristics achieve competitive performance without relying on annotated training data for event argument extraction. The method attains an accuracy of 81.91% for spatial and temporal associations, supporting reliable knowledge graph construction. Furthermore, the aggregation step consolidates 541 relevant triplets into 376 distinct events, reducing redundancy while preserving multi-source evidence for each food security situation.

4.3 Knowledge graph exploitation

We adopted the hypotheses that the KG would allow us to: (H1) discover information about historical trends in food security risks and (H2) provide answers to the questions that stakeholders in the field may have about

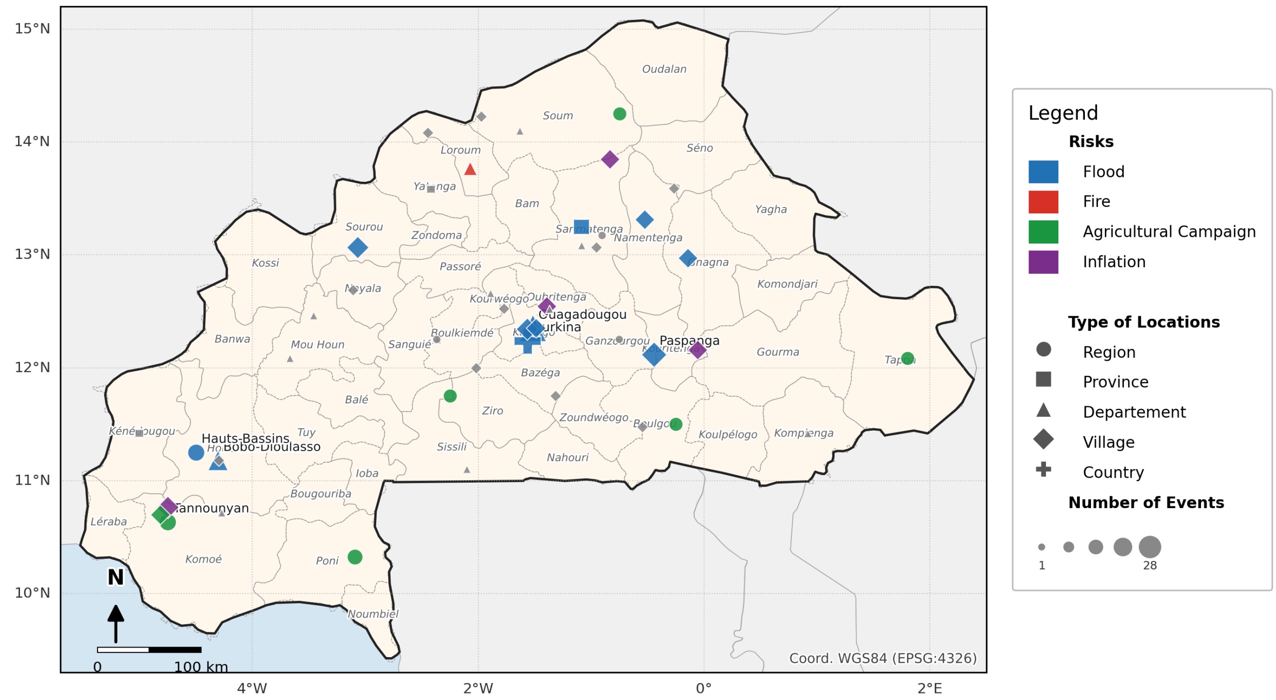

Figure 5. Vulnerable areas across the country that were affected by frequent crises in 2009: environmental (floods in blue and fires in red), agricultural (production decline in green) and economic (inflation in purple). A marker icon indicates a village, a rectangular icon indicates a department, and a circular icon indicates a region.

risk indicators. Some questions have been raised to help us assess the hypotheses. We will then request the KGs to gather information relating to these questions. These questions are categorized according to their respective hypothesis.

With reference to the hypothesis *H1*, the following questions arise: (Q1.1) What historical events occurred in vulnerable areas during a specific period? (Q1.2) What is the geographical distribution of risks across the country? Regarding hypothesis *H2*, the questions are (Q2.1): Which indicators do stakeholders use to identify and assess food security risks? and (Q2.2): What kind of indicators can be produced from KGs to complement the information provided by the standard analytical method?

To answer question Q1.1, we have chosen to analyze the Ouagadougou department, as it is the most vulnerable area in the country. Figure 4 illustrates the historical risk events that occurred in this department in 2009, as reported in newspapers published that year. The city was exposed to various types of risk (economic, demographic, environmental, and socio-political) for just four months, between June and September. Some crises are recurrent, such as poverty, overpopulation and flooding (represented by red edges). The synchronisation of precipitation, rainfall and flooding indicates the severity of the disaster that occurred on 4 July 2009 (represented by the double-direction pink edges). Looking at Figure 4, we can see that the blockade event in July occurred after the severe disaster, followed by a price hike in August 2009. This knowledge graph extract enables us to explore historical events that took place in the Ouagadougou department during a particular time of period.

In an attempt to answer question Q1.2, we have identified the three most common types of crisis (namely, environmental, agricultural and economic) and located the places affected by them.

Figure 5 illustrates the locations of the vulnerable areas that were affected by environmental crises (represented by blue for floods and red for fires), agricultural crises (represented by red for production decline), and economic crises (represented by purple for inflation) across the country. The location of all detected flood events is given in great detail, i.e. at village level (diamond markers). Although the environmental risk level at *Bobo-Dioulasso* and *Cascades* is low, they are more affected by agricultural and economic crises.

To answer questions Q2.1 and Q2.2, we referred to literature indicating that food security risk indicators encompass a range of dimensions, including agriculture, environment, economics, health, diet, education, and standard of living, among others (Dasgupta and Robinson, 2022; Lee et al., 2024).

Table 3 illustrates the number of risks of each type (environmental, sociopolitical, agricultural, economic, dietary, general crisis and health) in each region in 2009. The number of provinces, departments and villages in which events took place is indicated in columns 2, 3 and 4, respectively. As we can see, the *Centre* region was

Table 3. Regional-level aggregated distribution of thematic events in Burkina Faso. All events from provinces, departments, and villages are aggregated to their respective regions. n_p: provinces, n_d: departments, n_v: villages. ENV: Environment, SOC: Sociopolitical, AGR: Agriculture, ECO: Economic, DIE: Dietary, CRI: General Crisis, HEA: Health.

Region	n_p	n_d	n_v	**ENV**	**SOC**	**AGR**	**ECO**	**DIE**	**CRI**	**HEA**	**Tot.**
National	–	–	–	16	24	19	24	8	13	4	108
Centre	1	1	2	34	26	16	10	7	14	0	107
Hauts-Bassins	1	1	2	9	12	11	2	2	2	1	39
Boucle du Mouhoun	6	3	9	4	8	5	1	1	5	0	24
Plateau Central	3	3	5	12	5	1	2	0	2	1	23
Centre-Nord	3	2	2	7	0	7	0	1	1	0	16
Cascades	1	2	2	0	3	8	3	0	1	0	15
Centre-Ouest	2	3	0	2	1	2	2	0	2	1	10
Centre-Est	3	1	5	2	0	5	1	0	1	0	9
Sahel	2	2	0	2	0	3	0	1	2	1	9
Nord	1	1	1	3	1	1	0	0	1	0	6
Sud-Ouest	2	0	2	0	2	3	0	0	0	0	5
Est	2	1	1	1	0	2	0	0	0	0	3
Centre-Sud	1	0	1	0	0	2	0	0	0	0	2
Total	–	–	–	92	82	85	45	20	44	8	376

Table 4. Hierarchical geographic distribution of thematic events in the Centre and Cascades regions of Burkina Faso. The table shows the complete administrative hierarchy (provinces, departments, villages) where events occurred. n_p: number of provinces with events, n_d: number of departments with events, n_v: number of villages with events. ENV: Environment, SOC: Sociopolitical, AGR: Agriculture, ECO: Economic, DIE: Dietary, CRI: General Crisis, HEA: Health

Region	**Province**	**Department**	**Village**	**ENV**	**SOC**	**AGR**	**ECO**	**DIE**	**CRI**	**HEA**	**Tot.**
Centre (n=107, n_p=1, n_d=1, n_v=2)											
Centre	–	–	–	0	1	0	0	0	0	0	1
	Kadiogo	Ouagadougou	–	31	25	16	10	7	14	0	103
	Kadiogo	Ouagadougou	Pissy	2	0	0	0	0	0	0	2
	Kadiogo	Ouagadougou	Bogodogo	1	0	0	0	0	0	0	1
Cascades (n=15, n_p=1, n_d=3, n_v=2)											
Cascades	–	–	–	0	1	4	0	0	0	0	5
	Comoé	Banfora	–	0	2	0	0	0	1	0	3
	Comoé	Sidéradougou	–	0	0	2	0	0	0	0	2
	Comoé	Bérégadougou	Bérégadougou	0	0	1	3	0	0	0	4
	Comoé	Banfora	Karfiguela	0	0	1	0	0	0	0	1

affected by almost all crises and ranked first among all regions. The vulnerability of these regions is explained by the frequency of various risks throughout the year. This information supplements the findings of the 2020 edition of Burkina Faso's Environmental Statistical Yearbook, which reported that the *Centre* region had the largest increase in the number of people affected by crises. The crisis events in the *Centre* region are confined to one province and department, whereas those in the *Boucle du Mouhoun* and *Plateau Central* regions occurred in multiple provinces, departments and villages.

Table 4 compares the *Centre* region and *Cascades* region in terms of crises. According to the KG extracted from our new corpus, the *Centre* region is mainly affected by environmental and sociopolitical crises, whereas the *Cascades* region is mainly affected by agricultural and economic crises and is not affected by environmental crises. The crises in the Cascades region affect four departments, whereas in the Centre region, it affects only one. This is due to the type of crisis. We have also discovered that environmental crises tend to be associated with specific locations, such as villages, whereas agricultural and economic crises are associated with larger areas, such as departments or provinces (Figure 5). The experiment that identifies food security events requires the detection of spatial entities at a fine-grained level, and that represents data using knowledge graphs, reveals historical event trends in terms of both time and space. Furthermore, some indicators used in traditional analytical methods to assess food security can be derived from knowledge graphs. Information from knowledge graphs supplements standard analyses by providing an overview of crises (e.g. distribution by region and category of risks) and details on specific locations (e.g. affected villages and the occurrence, sequence and specifics of events). The limitations of our approach lie in the data collected, which may be influenced by media coverage. For example, the capital city, Ouagadougou, receives a lot of media attention compared to other districts, which can affect how the data is distributed. Focusing on the local media in specific districts may provide more information about them.

5 Conclusion

This paper presents an approach for constructing spatio-temporal knowledge graphs from French press articles to monitor food security in Burkina Faso. Our contributions include: a fine-tuned CamemBERT model achieving high accuracy for village-level spatial entity recognition, addressing the underrepresentation of African toponyms; a projection-based triplet formation algorithm linking risk indicators to spatial and temporal anchors without annotated training data; and a knowledge graph structure with temporal inference relations enabling detection of historical trends, cascading effects, and compound crises. Experiments show that press-derived information complements standard food security analyses by providing both regional overviews and village-level details. Several limitations remain. Heuristic projection rules produce many irrelevant triplets requiring manual filtering, and media coverage is biased toward highly populated locations (especially the capital city and its department Kadiogo), potentially skewing the geographic representation of food security events. The spatial projection also fails when location references exceed the contextual window, causing incorrect associations or over-aggregation to higher administrative levels. Future work will develop an automatic triplet classifier, expand training data for village-level recognition, and extend the approach across West Africa over longer time periods. A planned journal extension will leverage systematic Cypher queries to detect cascading crises, co-occurring risks, and spatially spreading phenomena.

Declaration of Generative AI in writing

The authors declare that Generative AI tools were used solely for proofreading. All scientific content and conclusions are original work conducted without AI assistance.

References

Adelani, D. I., Abbott, J., Neubig, G., D'souza, D., Kreutzer, J., Lignos, C., Palen-Michel, C., Buzaaba, H., Rijhwani, S., Ruder, S., Mayhew, S., Azime, I. A., Muhammad, S. H., Emezue, C. C., Nakatumba-Nabende, J., Ogayo, P., Anuoluwapo, A., Gitau, C., Mbaye, D., Alabi, J., Yimam, S. M., Gwadabe, T. R., Ezeani, I., Niyongabo, R. A., Mukiibi, J., Otiende, V., Orife, I., David, D., Ngom, S., Adewumi, T., Rayson, P., Adeyemi, M., Muriuki, G., Anebi, E., Chukwuneke, C., Odu, N., Wairagala, E. P., Oyerinde, S., Siro, C., Bateesa, T. S., Oloyede, T., Wambui, Y., Akinode, V., Nabagereka, D., Katusiime, M., Awokoya, A., MBOUP, M., Gebreyohannes, D., Tilaye, H., Nwaike, K., Wolde, D., Faye, A., Sibanda, B., Ahia, O., Dossou, B. F. P., Ogueji, K., DIOP, T. I., Diallo, A., Akinfaderin, A., Marengereke, T., and Osei, S.: MasakhaNER: Named Entity Recognition for African Languages, Transactions of the Association for Computational Linguistics, 9, 1116–1131, https://doi.org/10.1162/tacl_a_00416, 2021.

Ahn, Y., Yan, M., Lin, Y.-R., and Wang, Z.: Hungergist: An interpretable predictive model for food insecurity, in: 2023 IEEE International Conference on Big Data (BigData), pp. 1591–1600, IEEE, 2023.

Balashankar, A., Subramanian, L., and Fraiberger, S. P.: Predicting food crises using news streams, Science Advances, 9, eabm3449, https://doi.org/10.1126/sciadv.abm3449, 2023.

Carneiro, B., Resce, G., Ruscica, G., Eilerts, G., Caravaggio, N., Santangelo, A., Tucci, G., Pacillo, G., Läderach, P., and Coffey, K.: Determinants of food insecurity: a text mining approach, 2024.

Carneiro, B., Resce, G., Caravaggio, N., Santangelo, A. E., Ruscica, G., Tucci, G., Pacillo, G., Eilerts, G., Läderach, P., and Coffey, K.: Text mining and machine learning reveal global determinants of food insecurity, Scientific Reports, 15, 36 709, https://doi.org/10.1038/s41598-025-20670-x, publisher: Nature Publishing Group, 2025.

Consoli, S., Coletti, P., Markov, P. V., Orfei, L., Biazzo, I., Schuh, L., Stefanovitch, N., Bertolini, L., Ceresa, M., and Stilianakis, N. I.: An epidemiological knowledge graph extracted from the World Health Organization's Disease Outbreak News, Scientific Data, 12, 970, https://doi.org/10.1038/s41597-025-05276-2, publisher: Nature Publishing Group, 2025.

Dasgupta, S. and Robinson, E. J.: Attributing changes in food insecurity to a changing climate, Scientific Reports, 12, 4709, 2022.

Deleglise, H., Bégué, A., Interdonato, R., Maître d'Hôtel, E., Roche, M., and Teisseire, M.: How can text mining improve the explainability of Food security situations?, Journal of Intelligent Information Systems, 62, 971–994, 2024.

Deléglise, H., Schaeffer, C., Maître d'Hôtel, E., and Bégué, A.: Lexiques en français sur la sécurité alimentaire et les crises, https://doi.org/10.18167/DVN1/C5PU01, 2021.

Gunel, B., Du, J., Conneau, A., and Stoyanov, V.: Supervised Contrastive Learning for Pre-trained Language Model Fine-tuning, https://doi.org/10.48550/arXiv.2011.01403, arXiv:2011.01403 [cs], 2021.

Lee, C.-C., Zeng, M., and Luo, K.: How does climate change affect food security? Evidence from China, Environmental Impact Assessment Review, 104, 107 324, 2024.

Li, S., Ji, H., and Han, J.: Document-Level Event Argument Extraction by Conditional Generation, in: Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, edited by Toutanova, K., Rumshisky, A., Zettlemoyer, L., Hakkani-Tur, D., Beltagy, I., Bethard, S., Cotterell, R., Chakraborty, T., and Zhou, Y., pp. 894–908, Association for Computational Linguistics, Online, https://doi.org/10.18653/v1/2021.naacl-main.69, 2021.

Li, Z., Zhou, W., Chiang, Y.-Y., and Chen, M.: GeoLM: Empowering Language Models for Geospatially Grounded Language Understanding, in: Proceedings of the 2023 Conference on Empirical Methods in Natural Language

Processing, edited by Bouamor, H., Pino, J., and Bali, K., pp. 5227–5240, Association for Computational Linguistics, Singapore, https://doi.org/10.18653/v1/2023.emnlp-main.317, 2023.

Liu, Y.-C. and Kuo, C.-L.: Constructing Spatio-temporal Disaster Knowledge Graph from Social Media, AGILE: GIScience Series, 5, 1–8, https://doi.org/10.5194/agile-giss-5-37-2024, publisher: Copernicus GmbH, 2024.

Ma, Y., Wang, Z., Cao, Y., Li, M., Chen, M., Wang, K., and Shao, J.: Prompt for Extraction? PAIE: Prompting Argument Interaction for Event Argument Extraction, in: Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), edited by Muresan, S., Nakov, P., and Villavicencio, A., pp. 6759–6774, Association for Computational Linguistics, Dublin, Ireland, https://doi.org/10.18653/v1/2022.acl-long.466, 2022.

Marivoet, W., Ulimwengu, J., and Sedano, F.: Spatial typology for targeted food and nutrition security interventions, World Development, 120, 62–75, https://doi.org/https://doi.org/10.1016/j.worlddev.2019.04.003, 2019.

Martin, L., Muller, B., Ortiz Suárez, P. J., Dupont, Y., Romary, L., de la Clergerie, É., Seddah, D., and Sagot, B.: CamemBERT: a Tasty French Language Model, in: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, edited by Jurafsky, D., Chai, J., Schluter, N., and Tetreault, J., pp. 7203–7219, Association for Computational Linguistics, Online, https://doi.org/10.18653/v1/2020.acl-main.645, 2020.

McBride, L., Barrett, C. B., Browne, C., Hu, L., Liu, Y., Matteson, D. S., Sun, Y., and Wen, J.: Predicting poverty and malnutrition for targeting, mapping, monitoring, and early warning, Applied Economic Perspectives and Policy, 44, 879–892, https://doi.org/10.1002/aepp.13175, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/aepp.13175, 2022.

Mehtab Alam, S., Arsevska, E., Roche, M., and Teisseire, M.: GeospaCy: A tool for extraction and geographical referencing of spatial expressions in textual data, in: Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations, edited by Aletras, N. and De Clercq, O., pp. 115–126, Association for Computational Linguistics, St. Julians, Malta, https://doi.org/10.18653/v1/2024.eacl-demo.13, 2024.

Middleton, S. E., Kordopatis-Zilos, G., Papadopoulos, S., and Kompatsiaris, Y.: Location Extraction from Social Media: Geoparsing, Location Disambiguation, and Geotagging, ACM Trans. Inf. Syst., 36, https://doi.org/10.1145/3202662, 2018.

Moriceau, V. and Tannier, X.: French Resources for Extraction and Normalization of Temporal Expressions with HeidelTime, in: Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14), edited by Calzolari, N., Choukri, K., Declerck, T., Loftsson, H., Maegaard, B., Mariani, J., Moreno, A., Odijk, J., and Piperidis, S., pp. 3239–3243, European Language Resources Association (ELRA), Reykjavik, Iceland, https://aclanthology.org/L14-1382/, 2014.

Randl, K., Pavlopoulos, J., Henriksson, A., Lindgren, T., and Bakagianni, J.: SemEval-2025 Task 9: The Food Hazard Detection Challenge, in: Proceedings of the 19th International Workshop on Semantic Evaluation (SemEval-2025), edited by Rosenthal, S., Rosá, A., Ghosh, D., and Zampieri, M., pp. 2523–2534, Association for Computational Linguistics, Vienna, Austria, https://aclanthology.org/2025.semeval-1.325/, 2025.

Shuang, K., Zhou, J., Wang, Q., and Guo, J.: Thinking about how to extract: Energizing LLMs' emergence capabilities for document-level event argument extraction, in: Findings of the Association for Computational Linguistics: ACL 2024, edited by Ku, L.-W., Martins, A., and Srikumar, V., pp. 5520–5532, Association for Computational Linguistics, Bangkok, Thailand, https://doi.org/10.18653/v1/2024.findings-acl.328, 2024.

Sousa, H., Campos, R., and Jorge, A.: TEI2GO: A Multilingual Approach for Fast Temporal Expression Identification, in: Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, CIKM '23, p. 5401–5406, Association for Computing Machinery, New York, NY, USA, https://doi.org/10.1145/3583780.3615130, 2023.

Strötgen, J., Zell, J., and Gertz, M.: HeidelTime: Tuning English and Developing Spanish Resources for TempEval-3, in: Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013), edited by Manandhar, S. and Yuret, D., pp. 15–19, Association for Computational Linguistics, Atlanta, Georgia, USA, https://aclanthology.org/S13-2003/, 2013.

Wang, X., Chen, Y., Ding, N., Peng, H., Wang, Z., Lin, Y., Han, X., Hou, L., Li, J., Liu, Z., Li, P., and Zhou, J.: MAVEN-ERE: A Unified Large-scale Dataset for Event Coreference, Temporal, Causal, and Subevent Relation Extraction, in: Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, edited by Goldberg, Y., Kozareva, Z., and Zhang, Y., pp. 926–941, Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, https://doi.org/10.18653/v1/2022.emnlp-main.60, 2022.

Zaratiana, U., Tomeh, N., Holat, P., and Charnois, T.: GLiNER: Generalist Model for Named Entity Recognition using Bidirectional Transformer, in: Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers), edited by Duh, K., Gomez, H., and Bethard, S., pp. 5364–5376, Association for Computational Linguistics, Mexico City, Mexico, https://doi.org/10.18653/v1/2024.naacl-long.300, 2024.

Zhang, X. F., Blum, C., Choji, T., Shah, S., and Vempala, A.: ULTRA: Unleash LLMs' Potential for Event Argument Extraction through Hierarchical Modeling and Pair-wise Self-Refinement, in: Findings of the Association for Computational Linguistics: ACL 2024, edited by Ku, L.-W., Martins, A., and Srikumar, V., pp. 8172–8185, Association for Computational Linguistics, Bangkok, Thailand, https://doi.org/10.18653/v1/2024.findings-acl.487, 2024.

Zhou, B., Zou, L., Hu, Y., Qiang, Y., and Goldberg, D.: TopoBERT: a plug and play toponym recognition module harnessing fine-tuned BERT, International Journal of Digital Earth, 16, 3045–3064, https://doi.org/10.1080/17538947.2023.2239794, 2023.