



TECHNICAL BRIEF 2

Modelling Gridded Forced Displacement Across Africa

ABSTRACT

We present a predictive modelling approach to estimate monthly forced displacement outflows at a 0.5-degree grid-cell resolution across 25 African countries using data from UNHCR's own refugee registration records. The refugee records were spatially gridded on a consistent 0.5-degree grid-scale and combined with environmental, conflict, demographic, and socio-economic covariates at the same resolution to create a comprehensive dataset of displacement drivers with climate as a specific focus. The target variable presents significant statistical challenges due to excess zeros, and infrequent, short-lived spikes with large displacement counts. These challenging characteristics required the implementation of multiple classification algorithms for prediction. Count patterns were analysed using tree-based models (lightGBM), convolutional neural networks (CNN), and ConvLSTM architectures to handle spatial-temporal dependencies in the data. Results show good predictive performance for low and medium counts, but difficulties in detecting large and acute displacement spikes.

Table of Contents

Introduction	3
Literature Overview	4
Data	6
Data Sources	6
Spatial and Temporal Coverage	6
Target Variable Characteristics	6
Exploratory Data Analysis	7
Feature variables	11
Methodology	15
Baseline Model	15
Tree-based Models	16
Neural networks	17
Ensemble model	20
Results	21
Conclusion	26
Bibliography	27

Introduction

Climate change is increasingly recognized as a significant driver of forced displacement, operating through its effects on conflict, economic shocks, and food insecurity. As climate impacts intensify, understanding and predicting climate-driven displacement becomes essential to anticipate future population displacements and enable more effective humanitarian preparedness.

Early approaches to climate-driven displacement prediction relied on simplistic assumptions about linear cause-and-effect relationships between climate hazards and population movements, often focusing on mapping out hazards and assuming complete displacement from affected areas without taking adaptation measures into account. These theoretical approaches generated broad, often catastrophic, projections of future climate driven mobility, but lacked the rigor and timeliness needed for humanitarian planning. More recent research focused on understanding how climate impacts intersect with existing mobility systems, and affect the nature, volume and direction of displacement flows. (see Schewel et al., 2024, for an overview). In this context, recognizing that climate change acts indirectly on pre-existing migration systems, and is often not the proximate cause for migration, researchers incorporate climate-related indicators into migration and population models to forecast climate-driven mobility.

However, several limitations characterizing the field of climate-related displacement forecasting remain, many of them related to data availability. In particular, a lack of data on non-climate related drivers of displacement such as social, political, and economic factors, and insufficient frequency and geographical disaggregation of displacement flows data may hamper the understanding of the localized mechanisms through which climate impacts actually operate. Environmental stressors interact with local conflict dynamics, economic conditions, and social vulnerabilities in spatially heterogeneous ways that national averages fail to capture adequately. Without detecting climate-displacement relationships at the scales where they operate, we may miss important signals of environmentally influenced displacement.

This gap limits opportunities for early intervention and anticipatory resource allocation, perpetuating reactive humanitarian responses.

We attempt to address this gap by developing a grid-based displacement forecasting framework that operates at the more granular spatial scale. Our approach combines high-resolution gridded displacement data from UNHCR's internal registration database with advanced deep learning architectures to predict forced displacement outflows at 0.5-degree resolution across 25 Sub-Saharan African countries. Building on insights from O'Loughlin et al. (2012), who demonstrate that climate-conflict relationships only becomes statistically detectable at fine spatial scales, we apply this principle to displacement forecasting. Since most forced displacement is conflict-driven and climate-conflict links seem to operate mainly at localized scales, displacement forecasting requires similar spatial granularity to capture the environmental drivers of population movements.

Our methodology integrates spatial-temporal modelling through tree-based models (lightGBM), convolutional neural networks, and ConvLSTM architectures to handle the complex spatial dependencies and temporal dynamics inherent in displacement processes. Deep learning approaches are particularly suited to this task because they can capture non-linear relationships between climate variables and displacement outcomes while accounting for spatial autocorrelation and temporal dependencies that traditional statistical methods struggle to model. This framework enables monthly displacement predictions up to six months in advance on three operationally relevant scales: small-scale movements (0-10 people), medium-scale events (11-500 people) and large-scale crises (>500 people).

The remainder of this paper is organized as follows. Section 2 gives a brief overview of the literature, and section 3 describes data sources and the spatial gridding methodology. Section 4 presents the modelling framework and deep learning architectures. Section 5 discusses the modelling results, and Section 6 concludes.

Literature Overview

This paper is related to a number of different literatures that contribute insights into the challenges of anticipating forced displacement outflows. Early quantitative work by Schmeidl (1997) used pooled time-series analysis to identify structural factors driving forced displacement from 1971-1990, establishing that measures of generalized violence and military interventions have high predictive power of refugee outflows. This study, together with the growing availability of violence indicators, was instrumental to the pioneering paper by Schmeidl and Jenkins (1998), which posited the feasibility of applying Early Warning Models (EWMs) to the prediction of humanitarian disasters such as large refugee crises. Moore and Shellman (2004, 2006) extended this foundational work by analysing forced displacement from 1952 to 1995 and developing frameworks to predict whether individuals would become refugees or internally displaced persons. These foundational studies demonstrated the feasibility of data-driven prediction approaches, while revealing the complexity of modelling forced displacement. Evidence that violence serves as the main push factor in the case of forced displacement is also established by Moore and Shellman (2007), while in a country case study, Ibáñez and Velez (2008) found that actual violent events, in addition to perceptions of violence and a generalized lack of security, were central determinants of forced population movements in Colombia. These authors also determined that modelling forced displacement critically differs from traditional migration modelling in that the threat of violence vastly increases the costs of staying.

Recent methodological advances, such as machine and deep learning approaches, have enabled more sophisticated forecasting approaches through improved computational methods and data availability. Carammia et al. (2022) developed adaptive machine learning algorithms to forecast asylum applications to EU countries up to four weeks ahead, integrating administrative statistics

with non-traditional data sources, including internet searches and conflict event databases. Their approach demonstrates that individual country-to-country displacement flows could be modelled accurately using dynamic variable selection methods. Henningsen (2025) created country-level early warning models using gradient boosting classification to predict displacement flows and sudden increases in movement, generating monthly risk indices for 176 countries with prediction horizons extending to six months, reaching an accuracy level of over 80% to predict sudden increases in displacement numbers. Suleimenova et al. (2017) proposed agent-based modelling approaches to predict refugee destinations during conflicts, synthesizing data from UNHCR, conflict databases, and mapping services to achieve over 75% accuracy in forecasting camp destinations across three major African conflicts. Finally, Hoffmann Pham and Luengo-Oroz (2022) provided a comprehensive methodological framework for prediction of IDP and refugee flows, highlighting the growing interest in machine learning applications while noting the lack of standardized approaches to structured prediction problems.

Another strand of literature this study relates to looks at the empirical relationship between climate, conflict, and displacement. In the most comprehensive study so far, Burke, Hsiang and Miguel (2015) use a hierarchical meta-analysis to estimate the mean effect and quantify the degree of variability across 55 studies on climate and conflict, and find that deviations from moderate temperatures and precipitation patterns systematically increase conflict risk. Contemporaneous temperature has the largest average impact, with each 1 standard deviation increase in temperature increasing interpersonal conflict by 2.4% and intergroup conflict by 11.3%. With respect to the climate-migration axis, the conceptual framework of Black et al. (2011), guides much of the empirical work by stating that in addition to exercising a direct influence, climate change indirectly affects migration decisions by affecting other drivers of

migration. Indeed, some macro-level studies provide support of the indirect effect of climatic factors on international migration through reduction of crop yields (Cai et al., 2016), and wage differentials between origin and destination (Beine and Parsons, 2015). In a study that uses data from 115 countries between 1960 and 2000 to analyse the effect of differential warming trends across countries on the probability of either migrating out of the country or from rural to urban areas, Cattaneo and Peri (2016) find increased migration with higher temperature for middle-income countries, whilst migration is suppressed in low-income countries. Additionally, using a gravity model accounting for endogenous selection, Abel et al. (2019) exploit bilateral data on asylum seeking applications for 157 countries over the period 2006–2015 to empirically establish the links between climate change, conflict and migration. They find that climatic conditions, by affecting drought severity and the likelihood of armed conflict, played a significant role as an explanatory factor for asylum seeking in the period 2011–2015.

Finally, this work relates to a strand of literature that assesses the phenomena at much finer level of geographical granularity. Using a conflict database that contains 16,359 individual geolocated violent events for East Africa from 1990 to 2009, and climate indicators at gridded 1° resolution (\sim 100 km), O’ Loughlin et al. (2012) find that much warmer than normal temperatures raise the risk of violence,

whereas average and cooler temperatures have no effect. Similarly, Harari and La Ferrara (2018) conduct a geographically disaggregated analysis of civil conflict in Africa between 1997 and 2011 taking as units of observation 110×110 km subnational “cells,” and estimate the incidence of conflict as a function of weather shocks and a number of other covariates in both the cell and neighbouring areas. Using a model that includes spatially and temporally autoregressive terms to account for the fact that conflict may be persistent over time and that both the covariates and the presence of conflict may be correlated across space, they find that a 1 standard deviation shock during the growing season to the Standardized Precipitation-Evapotranspiration Index (SPEI), which considers the joint effects of precipitation, potential evaporation, and temperature, is associated with a 1.3 percentage point increase in conflict likelihood in the subsequent year, relative to the cell’s historic mean.

Despite evidence from climate-conflict research that environmental effects on human behaviour require fine spatial scales for detection, displacement prediction continues to operate primarily at country or regional levels. To the best of our knowledge, the present study is the first to address this gap by developing displacement forecasting using a new gridded panel dataset with a rich set of georeferenced covariates at the cell/year level.

Data

Data Sources

Our analysis uses forced displacement data from the UNHCR's PRIMES database, which contains individual-level records of all refugees and asylum seekers registered by UNHCR. The spatial gridding methodology and data processing procedures for converting these individual-level registration records

into gridded displacement outflows counts are detailed in Wells et al. (2025). This approach enables consistent spatial aggregation of displacement events at 0.5-degree grid resolution while preserving temporal granularity at the monthly level.

Spatial and Temporal Coverage

The data set covers monthly displacement data for 25 West, Central and East African countries¹ from January 2010 to September 2025. This 15-year time series provides 189 monthly observations for each grid cell, creating a comprehensive spatial-temporal dataset for displacement modelling.

The 0.5-degree grid resolution balances computational feasibility with the spatial granularity required to detect localized displacement patterns, particularly those driven by environmental and

conflict factors operating at sub-national scales. The geographic coverage includes major displacement-affected regions across three African sub-regions, capturing diverse conflict contexts, climate zones, and socio-economic conditions. This spatial scope enables the analysis of displacement patterns across varied environmental and political contexts while maintaining sufficient observations for statistical modelling.

Target Variable Characteristics

The displacement data exhibits several challenging statistical properties that significantly influence modelling approaches. The distribution is characterized by extreme right skewness, with most cell/year observations recording zero displacement counts. Of the 6221 grid cells, only 1,742 (28.0%) have experienced at least some displacement and

of the total 1,175,769 observations in the dataset (6,221 grid cells x 189 months), 1,063,920 (90.5%) have displacement values of zero, while only 111,849 (9.5%) record one or more displaced persons. To address these statistical challenges posed by

¹ The 25 countries of focus include: Somalia, Eritrea, Ethiopia, Kenya, Tanzania, Burundi, Rwanda, Uganda, South Sudan, Sudan, Chad, Central African Republic, Democratic Republic of the Congo, Angola, Cameroon, Nigeria, Niger, Benin, Togo, Ghana, Burkina Faso, Mali, Côte d'Ivoire, Senegal, and Mauritania.

the raw count data, we implement a three-category classification scheme that aligns with operational humanitarian response frameworks:

- **Small-scale movements (0-10 people):** Captures minor displacement and population movements
- **Medium-scale events (11-500 people):** Represents significant displacement events requiring humanitarian attention

- **Large-scale crises (>500 people):** Identifies major displacement emergencies that may require immediate large-scale response

This binning approach addresses the statistical challenges of modelling highly skewed count data while creating operationally meaningful categories corresponding to different levels of humanitarian response capacities.

Exploratory Data Analysis

The displacement data exhibit significant temporal variation that reflects major conflict and political events in the study region. Figure 1 shows the average displacement count across all grid cells over the 15-year study period, revealing several distinct phases. Counts escalated from 2011 onward, with average displacement counts reaching peaks of nearly 30 persons per grid cell during 2014-2016, corresponding to major conflicts in Mali, the Central African Republic, South Sudan, and the Lake Chad

Basin. The average displacement spiked again in late 2016 through 2017, reflecting increased conflict in South Sudan. There was a decrease in average displacement from 2018 to 2023, with a spike in 2023, reflecting the outbreak of the civil war in Sudan and increased conflict in the border region between Burkina Faso, Mali and Niger. The time series demonstrates substantial volatility, with sharp spikes followed by periods of relative stability, highlighting the episodic nature of forced displacement events.

Average count across grids over time

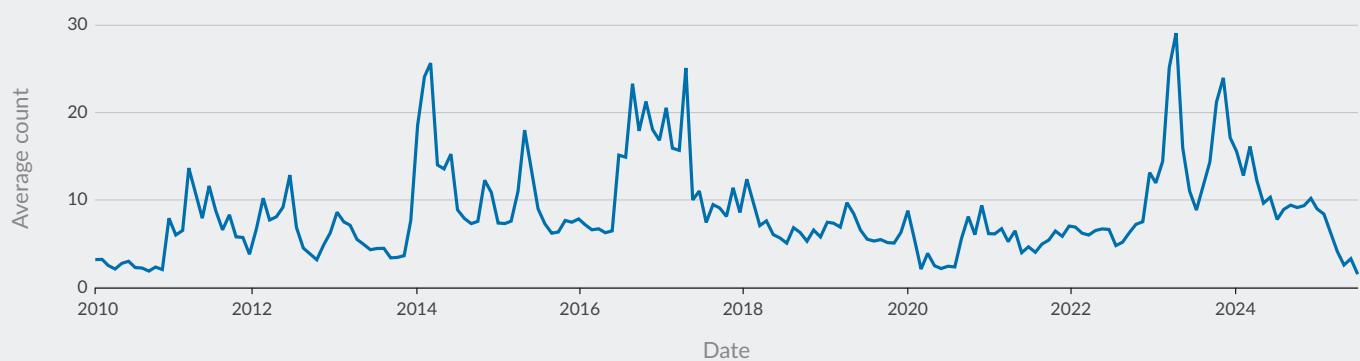


Figure 1. Average displacement per grid cell from 2000 to 2025.

The temporal dynamics of displacement events reveal additional complexity beyond simple time series patterns. Figure 2 shows the distribution of the top ten grids by displacement counts over time, coloured by country of origin. This visualization reveals the episodic nature of major displacement events, which

is characterized by many large and sudden spikes in the displacement figures. It is also noteworthy that 7 of the 10 grid cells with the most overall displacement fall within 3 countries; Sudan, Democratic Republic of the Congo, and South Sudan.

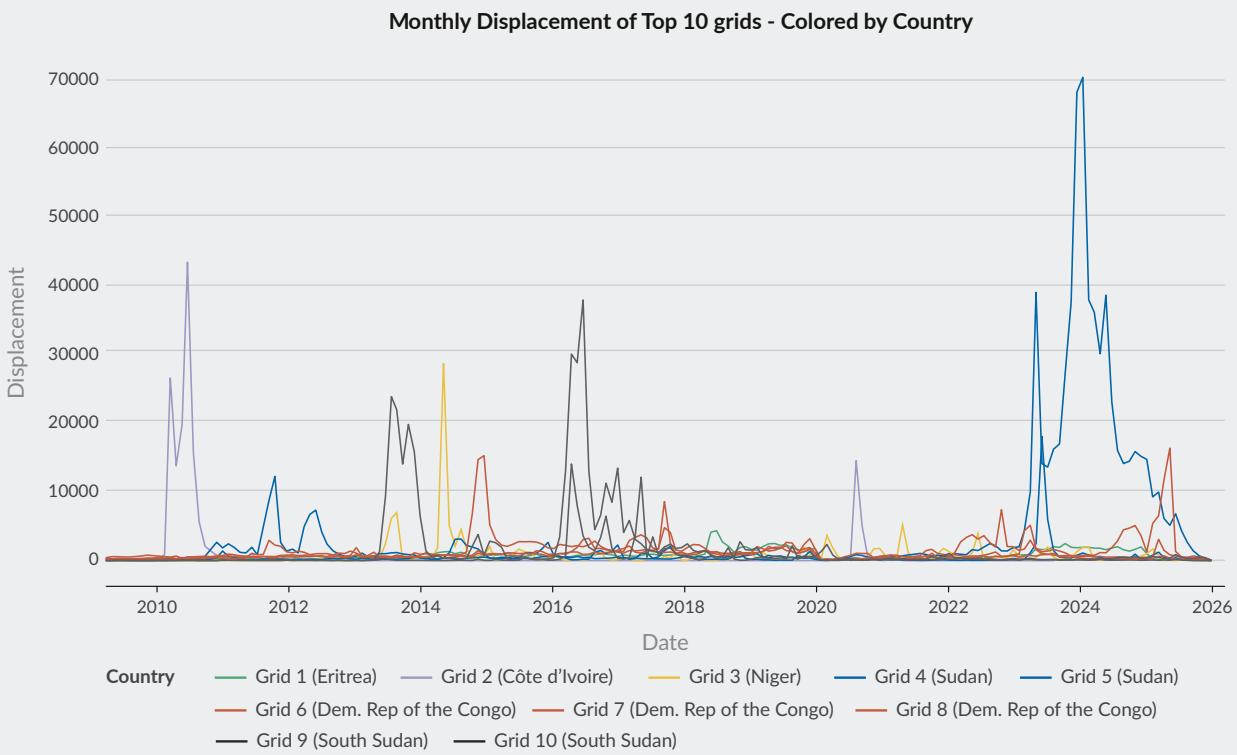


Figure 2. Top ten grid cells by maximum peak over study period.

The duration characteristics of the displacement spikes provide further insight into the temporal structure of forced displacement events. Figure 3 displays the distribution of the duration of the spikes in all observations, where a spike is defined as a continuous period of elevated displacement above the baseline (monthly average level in the previous year).

The distribution shows that the majority of displacement spikes are short-lived, with approximately 10,000 events lasting only 1-3

months. The frequency drops for longer durations, with fewer than 2,000 events lasting 4 months, and few extending beyond 6 months. This pattern suggests that most major displacement events are acute responses to specific triggers rather than sustained processes, although the small number of longer-duration events may represent protracted displacement situations requiring different analytical and operational approaches.

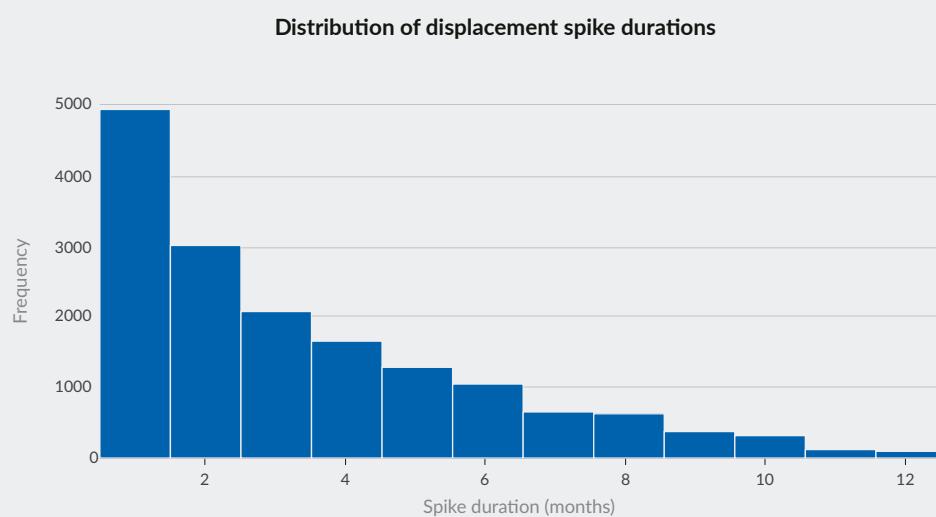


Figure 3. Distribution of spike duration over all grid cells measured as significant consecutive deviation from a grid's baseline.

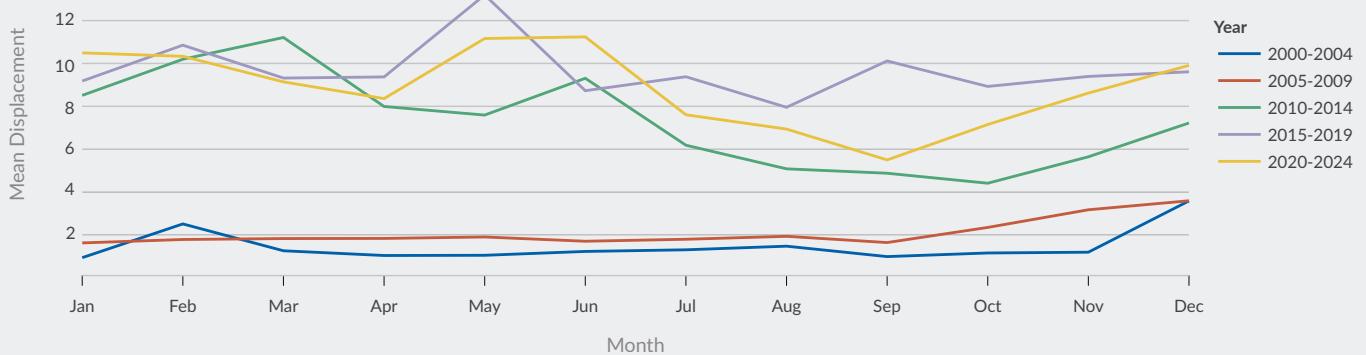


Figure 4. Seasonality pattern of gridded forced displacement relevant to predict forced displacement.

Figure 4 shows monthly displacement per grid averaged over 5-year segments. The emerging pattern shows that seasonality in displacement numbers is increasing over time, with peak displacement coinciding with the late dry season and early planting period, when food stocks are lowest and competition for resources is highest. This might indicate that environmental factors are increasingly relevant to predict forced displacement, as populations become more vulnerable to seasonal stressors.

The spatial distribution of displacement events shows clear geographical clustering that reflects conflict patterns and refugee flow corridors. Figure 5 displays the year of maximum displacement for each grid cell from 2000 to 2025 in colour and the aggregate displacement figure in the same year represented by the bubble size. Lighter bubbles show grid cells that have seen their highest year of displacement closer to 2000 and darker ones experienced the highest year of displacement more recently. Blank areas represent grid cells that have never experienced any displacement during this period. The colours and sizes of the bubbles reflect the reality of when and how many people were forced to leave certain areas, mainly due to conflict. The lightest colours in Angola reflect the displacement that occurred prior to the end of the 27-year-long civil war in 2002. Southern Somalia and Côte d'Ivoire are a green-blue. In Somalia, this reflects the formation of the extremist group Al Shabaab in the 2006 and a civil war which started in 2009. In Côte d'Ivoire, this reflects the outbreak of the Second Ivoirian Civil War in late-2010. The bubbles in eastern Mali and Central African Republic are slightly darker blue, representing the displacement that took place at the outbreak of civil

wars in these countries, each in 2012. South Sudan experienced its highest levels of displacement during the civil war of 2014 to 2020, which is represented by the darker bubbles. Finally, Burkina Faso, northwest Nigeria, and Sudan all have dark bubbles, showing that these grid cells experienced most displacement in the last couple of years. Some of these bubbles are very large, indicating the high levels of displacement from certain grid cells, particularly in southwest Burkina Faso, southern Sudan, and eastern Sudan. An additional note is that some areas show a variation in colouring. For instance, the Darfur region in western Sudan has both light and dark coloured bubbles. This reflects high levels of displacement due to the Darfur Crisis in the early 2000s and the recent conflict stemming from the Sudanese Civil War, since 2023.

We observed that displacement seems to be clustered among grid cells in geographic proximity and conducted a spatial autocorrelation analysis using Moran's I statistic to test this apparent spatial pattern formally. This approach allows us to quantify the degree of spatial clustering observed in the displacement data and identify statistically significant hotspots throughout the region. Our analysis yields a global Moran's I value of 0.1885 (k -nearest neighbours with $k = 6$), indicating a weak but positive spatial autocorrelation in displacement patterns. This suggests some tendency for similar displacement values to cluster together, although the effect is not particularly strong. Despite the modest coefficient, the permutation test resulted in a p -sim value of 0.001, confirming that this spatial pattern is statistically significant and unlikely to have occurred by chance alone.

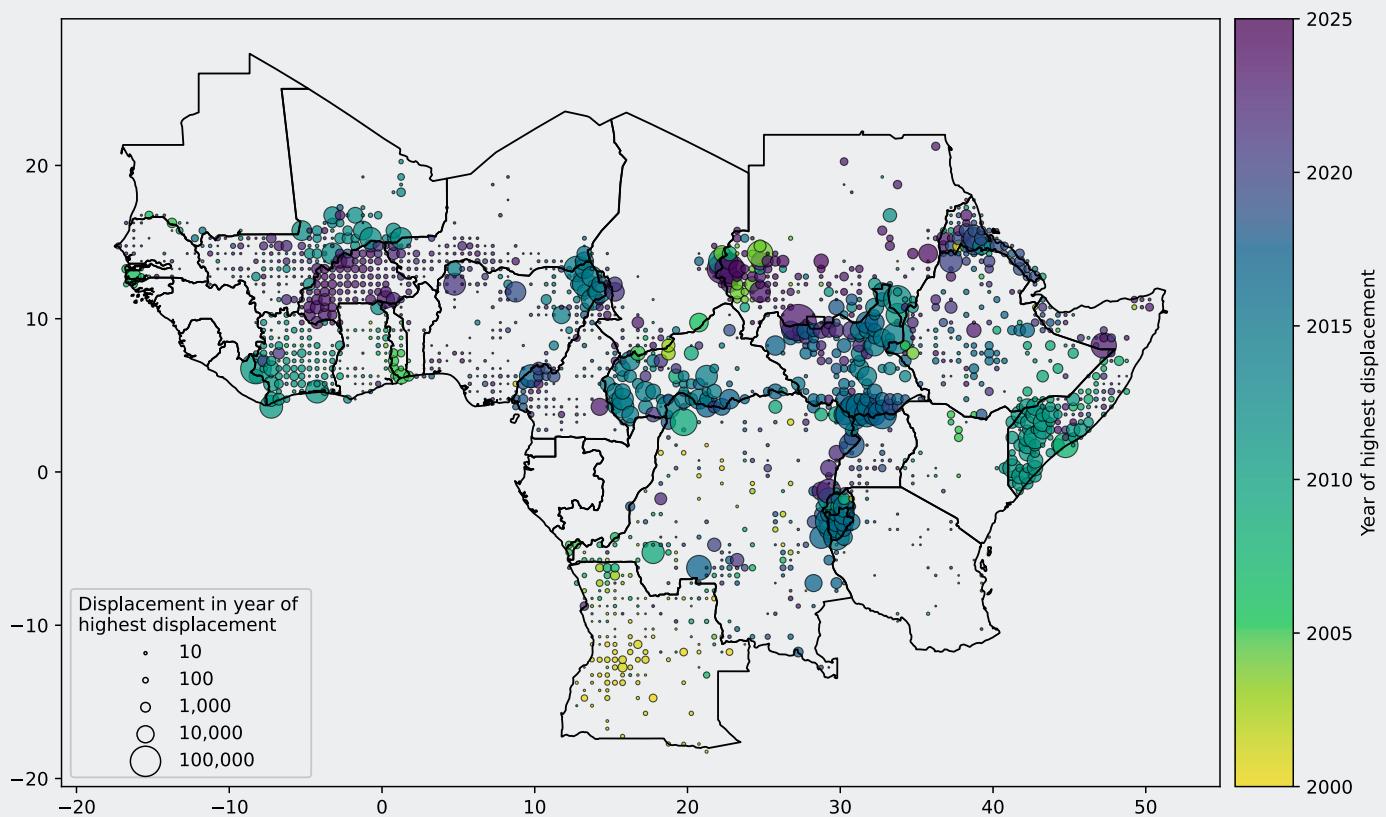


Figure 5. Year of the highest displacement count per grid cell.

The map in Figure 6 classifies each cell in the grid according to its most common value of local indicators of spatial association (LISA) over time, categorizing them into four distinct patterns: high-high, high-low, low-high, and low-low. The areas marked as high-high (red) represent displacement hotspots—locations with high displacement values surrounded by neighbouring cells that exhibit high displacement. These clusters are prominently visible across coastal West Africa (including Nigeria, Burkina Faso, Côte d'Ivoire), Ethiopia, Somalia, and parts of East Africa. In contrast, low-low clusters (light grey) indicate areas where low displacement values are spatially concentrated, visible in parts of Sudan, South Sudan, and Central Africa. The high-low (pink) and low-high (brown) designations represent spatial outliers—areas where displacement values differ markedly from their neighbours. These outliers appear more sparsely distributed and often at the borders of high-high clusters. The presence of

statistically significant spatial clustering reinforces our methodological decision to incorporate spatial effects in our predictive modelling approach.

These findings suggest that displacement is not merely a function of localized environmental or socio-political factors but is influenced by broader regional dynamics and possible spillover effects from neighbouring areas. Identifying specific hotspots provides valuable information to target preventive interventions and humanitarian responses. The clustering patterns align with known conflict zones and regions of environmental stress. The high-high clusters in the Horn of Africa correspond to areas experiencing recurrent drought conditions and long-lasting conflicts. In contrast, clusters in West Africa coincide with regions of political instability and temperature anomalies, providing some evidence displacement patterns may reflect complex interactions between environmental stressors and socio-political factors.

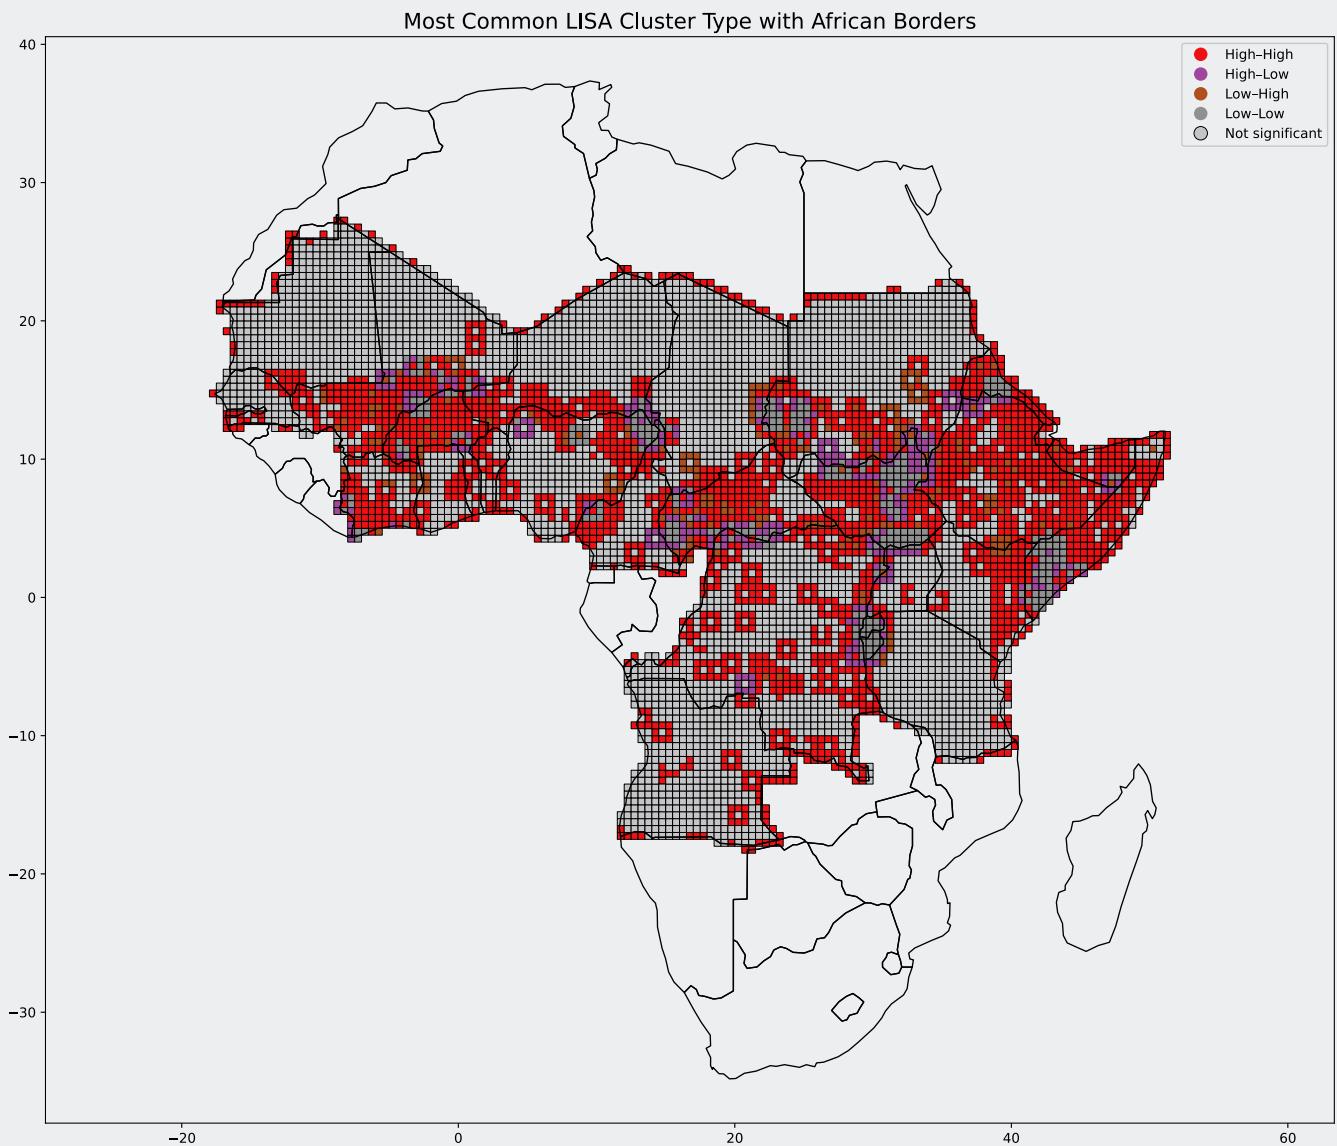


Figure 6. Most common LISA cluster type per grid cell over time.

Feature variables

Our modelling framework incorporates a comprehensive set of 180 variables organised into six thematic categories: environmental, geographic, demographic, food security, socio-economic, political, and conflict.

Environmental

We use data based on four categories of environmental variables in our model. These are data on temperature, precipitation, vegetation, and drought. The temperature data come from three different sources: Climate Hazards InfraRed Temperature with Stations (CHIRTS), which offers daily temperature highs at $0.05^\circ \times 0.05^\circ$ from 01 January 1980²; Copernicus ERA-5-land post-processed daily-statistics, which offers daily temperature highs at 0.10°

$\times 0.10^\circ$ from 01 January 1950³; and Berkeley Earth, which offers monthly temperature averages at $0.25^\circ \times 0.25^\circ$ from January 1850⁴. The precipitation data come from two sources: Climate Hazards InfraRed Precipitation with Stations (CHIRPS), which offers hourly precipitation total at $0.05^\circ \times 0.05^\circ$ from 01 January 1981⁵; and Copernicus ERA-5, which offers hourly precipitation measurements at $0.10^\circ \times 0.10^\circ$ from 01 January 1950. Vegetation is measured through the Normalized Difference Vegetation Index (NDVI), which approximates the level of vegetation at a point in time through satellite data. We use NDVI data from NASA, which offers monthly averages at $0.05^\circ \times 0.05^\circ$ from January 2000⁶. Drought levels are measured through the Standardized Precipitation Evapotranspiration Index (SPEI), which is an index used to approximate the level of dryness and drought at a point in time. Data are extracted from the SPEI Global Drought Monitor, which offers monthly SPEI averages at $1.00^\circ \times 1.00^\circ$ from January 1950⁷.

Geographic

We use four different types of geographic or resource variables, which include landcover, agro-ecological zone, elevation, river presence, road presence, and market access. The landcover data are extracted at $0.05^\circ \times 0.05^\circ$ resolution point locations from Copernicus. The data are categorized into 8 general types: tree cover, bare area, shrubland, cropland, grassland, urban, cover flooded and water⁸. Additionally, urban areas are categorized through data from Natural Earth⁹. Identifying grid cells classified as urban, we developed a variable that measures the distance from each grid cell to the nearest urban grid

cell. Agro-ecological zone (AEZ) data are extracted from the International Food Policy Research Institute, which are coordinate polygons of the various zones. AEZ classifications for Africa have three dimensions: major climate zones (tropics or subtropics), moisture zones (water availability) and highland/lowland (warm or cool based on elevation)¹⁰. Elevation data come from the HarvestChoice CELL5M Database. The data give the elevation, in meters, at coordinate points at a resolution of 5 arc minutes (approximately 9.30 km)¹¹. The geographic presence of a river is based on data from Natural Earth, which offers the coordinate geometries of rivers¹². Road presence is based on data from Humanitarian OpenStreetMap. Each grid cell is labelled based on the presence of at least some of the grid cell area encompassing a paved road or unpaved road, if neither of these are present, then the grid cell is labelled as no-road¹³.

Demographic

Demographic data used for the models are predicted population and ethnic composition of each grid cell. Predicted population is based on LandScan annual population data¹⁴ along with monthly DMSP¹⁵ and VIIRS¹⁶ nightlight data. The annual population and nightlight data are used to predict monthly population within each grid cell. Ethnic group location data are extracted from ETH Zurich, which are coordinate polygons of the boundaries between ethnic groups¹⁷. In order to account for Somali clan boundaries, SWALIM data¹⁸ on Somali clan locations are used. Based on the ethnic classification of each grid cell, we created a variable to measures the distance to the nearest grid cell of a different ethnic classification.

3 <https://cds.climate.copernicus.eu/datasets/derived-era5-land-daily-statistics?tab=overview>

4 <https://berkeleyearth.org/data/>

5 <https://www.chc.ucsb.edu/data/chirps>

6 <https://www.earthdata.nasa.gov/topics/land-surface/normalized-difference-vegetation-index-ndvi>

7 <https://spei.csic.es/map/maps.html#months=1#month=8#year=2025>

8 <https://land.copernicus.eu/en/products/global-dynamic-land-cover/land-cover-2020-raster-10-m-global-annual>

9 https://github.com/nvkelso/natural-earth-vector/blob/master/50m_cultural/ne_50m_urban_areas.shx

10 <https://dataVERSE.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/HJYYTI>

11 <https://dataVERSE.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/G4TBLF>

12 https://github.com/nvkelso/natural-earth-vector/blob/master/10m_physical/ne_10m_rivers_lake_centerlines.shp

13 https://data.humdata.org/dataset/hotosm_ago_roads

14 <https://landscan.ornl.gov/>

15 https://eogdata.mines.edu/wwwdata/dmsp/monthly_composites/

16 https://eogdata.mines.edu/nighttime_light/monthly_nottile/

17 <https://icr.ethz.ch/data/epr/core/>

18 <https://spatial.faoswalim.org/>

Food security

The food security classification of a grid cell on a monthly basis is based on FEWS NET data, which is released every three months with the data on the food security situation in subnational areas for the current month, predicted over the next three months, and four to six months into the future¹⁹. Certain countries are not included in the FEWS NET data, grid cells in these countries are imputed with a food security rating based on nearest neighbour imputation.

Socio-economic

Child health variables can offer an indication of the overall wellbeing of the population within a grid cell. These variables, infant mortality rate and child malnutrition rate, are extracted from the PRIO-GRID dataset. These variables are gridded to 0.5° resolution and are time-invariant, reflecting the values of the respective variables for the year 2000²⁰. In order to approximate market access, data from the International Food Policy Research Institute are used to label each grid cell with the average time in hours to the nearest market within towns of various size: 20,000, 50,000, 100,000, 250,000 and 500,000. To capture overall accessibility to markets of varying sizes, we constructed a composite indicator called the Market Accessibility Index (MAI). Each travel-time variable was first weighted inversely by the corresponding population size, giving greater emphasis to accessibility to markets in larger urban centres. Specifically, MAI was calculated as the sum of each travel time divided by the respective population scale. The resulting index was log-transformed to reduce skewness and then normalized using min–max scaling to a 0–1 range. Finally, the scale was inverted so that higher MAI values represent greater accessibility (i.e., shorter travel times to larger towns and cities), while lower values indicate poorer accessibility.

To account for inequality within grid cells, a time-invariant Gini coefficient was calculated by gridding WorldPop data²¹ from 2020 and VIIRS nightlight radiance data, averaged over the year 2020. These variables were placed inside 0.02° grid cell. The mean population within

these 0.02° grid cells was used along with the 2020 nightlight intensity data to estimate nightlight per capita at the 0.02° grid cell-level. The nightlight per capita of each 0.02° grid cell is used to estimate the inequality (Gini index) within each grid cell.

Political

To account for the political system and quality of governance that people experience, we use data from the Fragile States Index²². These variables include group grievance, which quantifies the divisions and schisms between different groups; economic inequality; public services; human rights; security apparatus; and fractionalized elites, which quantifies the fragmentation of state institutions along ethnic, class, clan, racial, or religious lines. The fragility variable is a composite of the previous variables. These indicators are annual and at the national level, so all grid cells in a country and in the same year carry the same value for each individual indicator.

Conflict

Conflict data used for this project is extracted from the Armed Conflict Location & Event Data (ACLED), which offers daily data on conflict events with geo-coordinated point locations of each conflict event as well as a description of each event²³. The descriptions of each event include information on the actors involved, number of fatalities, whether civilians were targeted, a categorization of each actor involved (as state forces, rebel group, political militia, identity militia, civilians, etc.), and the type of event (battle, protest, riot, strategic development, etc.). Based on the date and geo-point location of conflict events, data are aggregated into grid cells on the monthly basis. We focus on the number of events and conflict fatalities in each grid cell in each month. We also count the number of conflict events and fatalities in neighbouring grid cells to account for nearby violence that might drive people to become displaced.

Table 1 provides a detailed breakdown of these variables, including their temporal and spatial resolutions, as well as data sources.

19 <https://fews.net/data/acute-food-insecurity>

20 <https://grid.prio.org/#/download>

21 <https://www.worldpop.org/>

22 <https://fragilestatesindex.org/global-data/>

23 <https://acleddata.com/>

Variable type	Unit	Source
Climate variables		
Temperature	Daily, 0.05° Daily, 0.10° Monthly, 0.25°	CHIRTS Copernicus ERA-5 Berkeley Earth
Precipitation	Daily, 0.05° Daily, 0.10°	CHIRPS Copernicus ERA-5
Normalized Difference Vegetation Index	Monthly, 0.05°	NASA
Standardized Precipitation Evapotranspiration Index	Monthly, 1.00°	SPEI Global Drought Monitor
Resource and geographic variables		
Landcover	Constant, 0.05°	Copernicus Natural Earth
Agro-ecological zone	Constant, defined regions	International Food Policy Research Institute
Elevation	Constant, 0.10°	HarvestChoice CELL5M
River	Constant, defined regions	Natural Earth
Road	Constant, defined regions	Humanitarian OpenStreetMap
Market access	Constant, 0.10°	International Food Policy Research Institute
Demographic variables		
Population density	Monthly, 0.10°	LandScan, DMSP, VIIRS
Ethnicity	Constant, 0.10°	ETH Zurich SWALIM
Food security variables		
Food security	Monthly, 0.10°	FEWS NET
Socio-economic and wellbeing variables		
Child health	Monthly, 0.50°	PRIO
Market access	Constant, 0.10°	International Food Policy Research Institute
Gini	Monthly, 0.10°	WorldPop, VIIRS
Political variables		
Fragility	Annual, national	Fragile States Index
Conflict variables		
Conflict events and fatalities	Daily, geo-point locations	ACLED

Methodology

We develop a classification model for the three categories of forced displacement defined in section 3 (low, medium, and high) across three predictive horizons of one, three, and six months. We evaluate five different types of classification models in our

framework; a baseline model, a tree-based method, three different types of neural network architectures, and combine them in an Ensemble model. To maximize performance, we train a separate model for each predictive horizon.

Baseline Model

The baseline model is a simple naïve model, where the last available observation is used as the prediction for all three horizons. The naïve model

tends to perform surprisingly well in many real-world scenarios, particularly in slowly changing or persisting situations over short-time horizons.

Table 1. Variables used by model type.

Variable	CNN/ConvLSTM/2-stage convLSTM (3 models)	LightGBM
Core conflict variables		
conflict_events	✓	✓
conflict_fatalities	✓	✓
new_event	✓	✓
new_event_decay		✓
Geographic variables		
dist_urban	✓	✓
border_dist	✓	✓
Climate variables		
high_temperature		✓
heavy_precipitation		✓
high_temperature_accumulated	✓	✓
heavy_precipitation_accumulated	✓	✓
drought_accumulated	✓	✓
Socioeconomic variables		
Population	✓	✓
Fragility	✓	✓
Food security	✓	✓
Temporal features		
count (lag1-12)		✓
count rolling average		✓
min rolling average		✓
max rolling average		✓
Other variables		
All base variables		✓

Tree-based Models

We use LightGBM, which has been selected for its computational efficiency and low memory usage, which is particularly important given the size of our dataset (Ke et al., 2017). LightGBM has shown strong performance in time series forecasting tasks (Makridakis et al., 2022) and is also well suited to handle high-dimensional data, due to its built-in feature selection via split gain to keep the most relevant variables. To account for the temporal aspect of our data, we included the lagged response variable of up to 12 months along with engineered features

such as exponentially weighted moving averages. After testing different subsets of variables, the final LightGBM model was selected to use a subset of the main 113 feature variables mentioned above, combined with temporal features. The model's inherent feature selection was particularly important in identifying key climate and environmental predictors. Extensive testing showed that including environmental variables improved the model's accuracy, showing the importance of incorporating climate data into our framework.

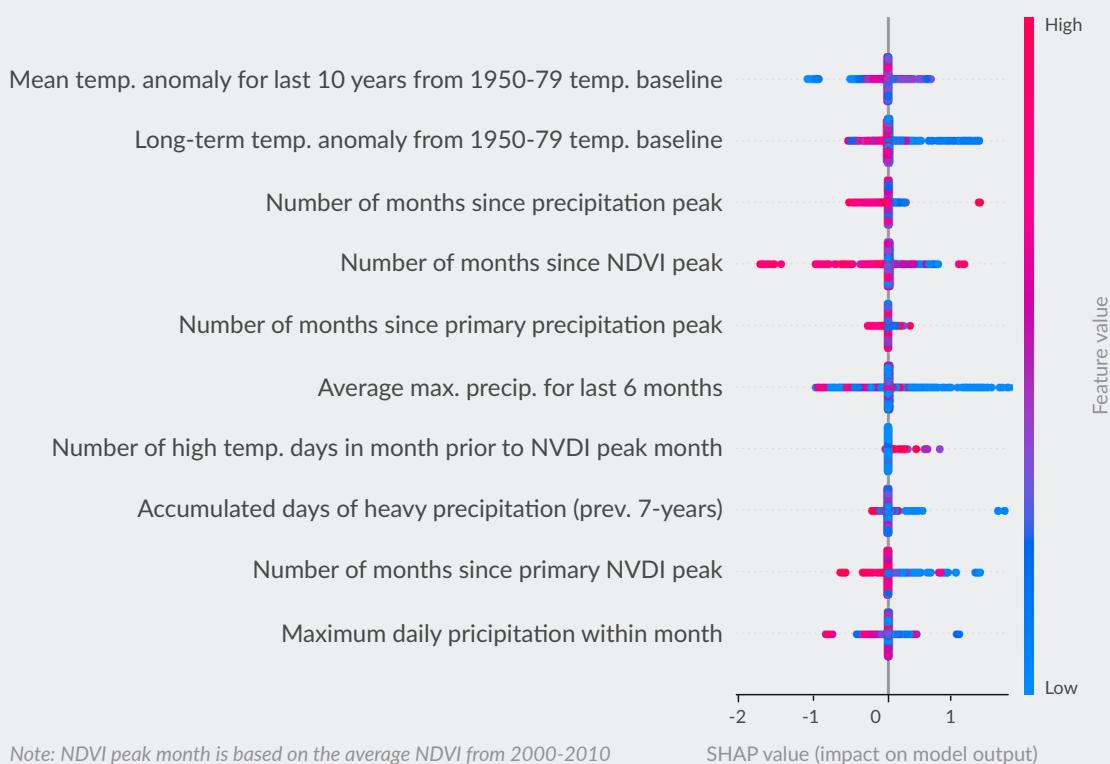


Figure 7. SHAP importance of top 10 climate variables for class 2 prediction (6-month horizon LightGBM model). Dots show individual data point predictions; colour represents feature value.

Figure 7 presents the top 10 climate variables ranked by the highest average impact on the model predictions for the LightGBM model predicting 6 months ahead. The x-axis represents the SHAP value – a measure that quantifies a feature's contribution to a specific prediction. The y-axis displays the features, with each dot signifying the SHAP value of a particular feature for a given data point. The colour

of the dot represents the associated feature value from low (blue) to high (red), allowing to gauge the extent of each feature's contribution to the prediction. For example, the number of high temperature days in the month prior to the NDVI peak month shows that higher values are associated with positive SHAP values increasing the likelihood of class 2 predictions.

Neural networks

The three neural networks that we employ are a simple convolutional neural network (CNN), a convolutional LSTM with a simple architecture, and a hierarchical convolutional LSTM. Due to dimensionality issues in spatial deep learning (Tomaso Poggio, 2018) the CNN and ConvLSTM models employ a focused set of 11 core variables which were selected using SHAP values from the final LightGBM model which include the most important identified climate, conflict, and socioeconomic variables as displayed in Table 1.

Convolutional Neural Network (CNN)

The 3D CNN processes spatiotemporal sequences using stacked 3D convolutional layers that jointly capture temporal and spatial dependencies. The idea is to have changing kernel sizes for our temporal fields to capture short- and long-term temporal patterns.

CNNs have been widely used in the context of spatiotemporal modelling, especially in the context of flood predictions, where they have demonstrated the ability to capture large- and small-scale spatial patterns well (Yan et al., 2022). Furthermore, another advantage of 3D CNNs is that they process space and time simultaneously rather than sequentially, which helps to preserve spatio-temporal correlations throughout the network.

The network processes the information through 5-dimensional tensors that encode spatiotemporal information in the standard format: [batch size, sequence length, channels, height, width] (Chollet & Watson, 2025). This format is typically used for video sequences. In our case, this can be seen as a sequence of multi-channel spatial frames over time, where each frame represents the geospatial data at a specific point in time. Each tensor contains 80 consecutive time points as input, followed by the target displacement category at the specified forecast horizon. The height and width dimensions of the tensor refer to the spatial information, with the original 0.5° grids mapped to cells within the tensor grid. Adjacent cells in the original geographic grid are assigned adjacent positions in the height and width dimensions of the tensor, preserving

spatial relationships and allowing the models to learn from spatial dependencies. The channel dimension refers to the number of feature variables, such as geographic and climate variables, plus the response variable. The batch size determines how many complete sequences can be processed simultaneously during training.

The model employs a rolling-origin cross validation approach for training and optimization across different predictive horizons. This approach is essential for time series forecasting as it mimics real-world scenarios by respecting the temporal nature of the data. The tensors are created according to this sliding-window approach, where each tensor is built using the specified consecutive time points as input and the target displacement category at the specified forecast horizon. The origin then rolls forward through time by one step and new tensors are created. The number of training samples depends on the length of the time series and the step size.

The architecture of the CNN consists of three convolutional layers followed by a projection layer. The first 3D convolutional layer uses a kernel size of (12,3,3), over the temporal and spatial dimensions with 16 output channels and applied padding to preserve resolution. After convolution, we apply batch normalization to stabilize training, ReLU activation to introduce nonlinearity, and a 3D dropout rate of 20% to prevent overfitting. The second 3D convolutional layer increases the temporal kernel size to (24, 3, 3) to enable the model to capture longer-range temporal patterns and expand the feature space to 32 channels. Again, we apply batch normalization, ReLU and dropout at the end. As a third convolutional layer, we have a feature interaction layer which uses a smaller kernel size of (3,3,3). This has been introduced as a layer to learn more complex interactions between the features it has already extracted. After the convolutional blocks, the model extracts the last time slice from the feature map, and a 1x1x1 projection layer produces the final spatial prediction for our three classes.

Training employs a weighted Cross Entropy loss function where the weight for the minority class 2 is treated as a hyperparameter. The model undergoes

hyperparameter optimization over the weight of the minority class, learning rates, number of epochs, batch size using a simple grid search.

Convolutional Long Short Term Memory (ConvLSTM)

We use a simple convolutional LSTM, which was inspired by Shi et al. (2015) to perform a sequence-to-frame prediction which has been shown to be stable and powerful in capturing the temporal and spatial information encoded in a data set.

It is implemented through a ConvLSTM cell by integrating 2D convolutional layers with the memory mechanism of LSTMs (input, forget, cell, and output gates) to preserve spatial information while enabling temporal memory. Unlike 3D CNNs which process fixed temporal windows, ConvLSTM uses gating mechanisms, which allow us to selectively determine which past information to remember and forget. The sequential processing approach allows variable-length sequences without the computational overhead of convolutions, facilitating the use of long temporal sequences.

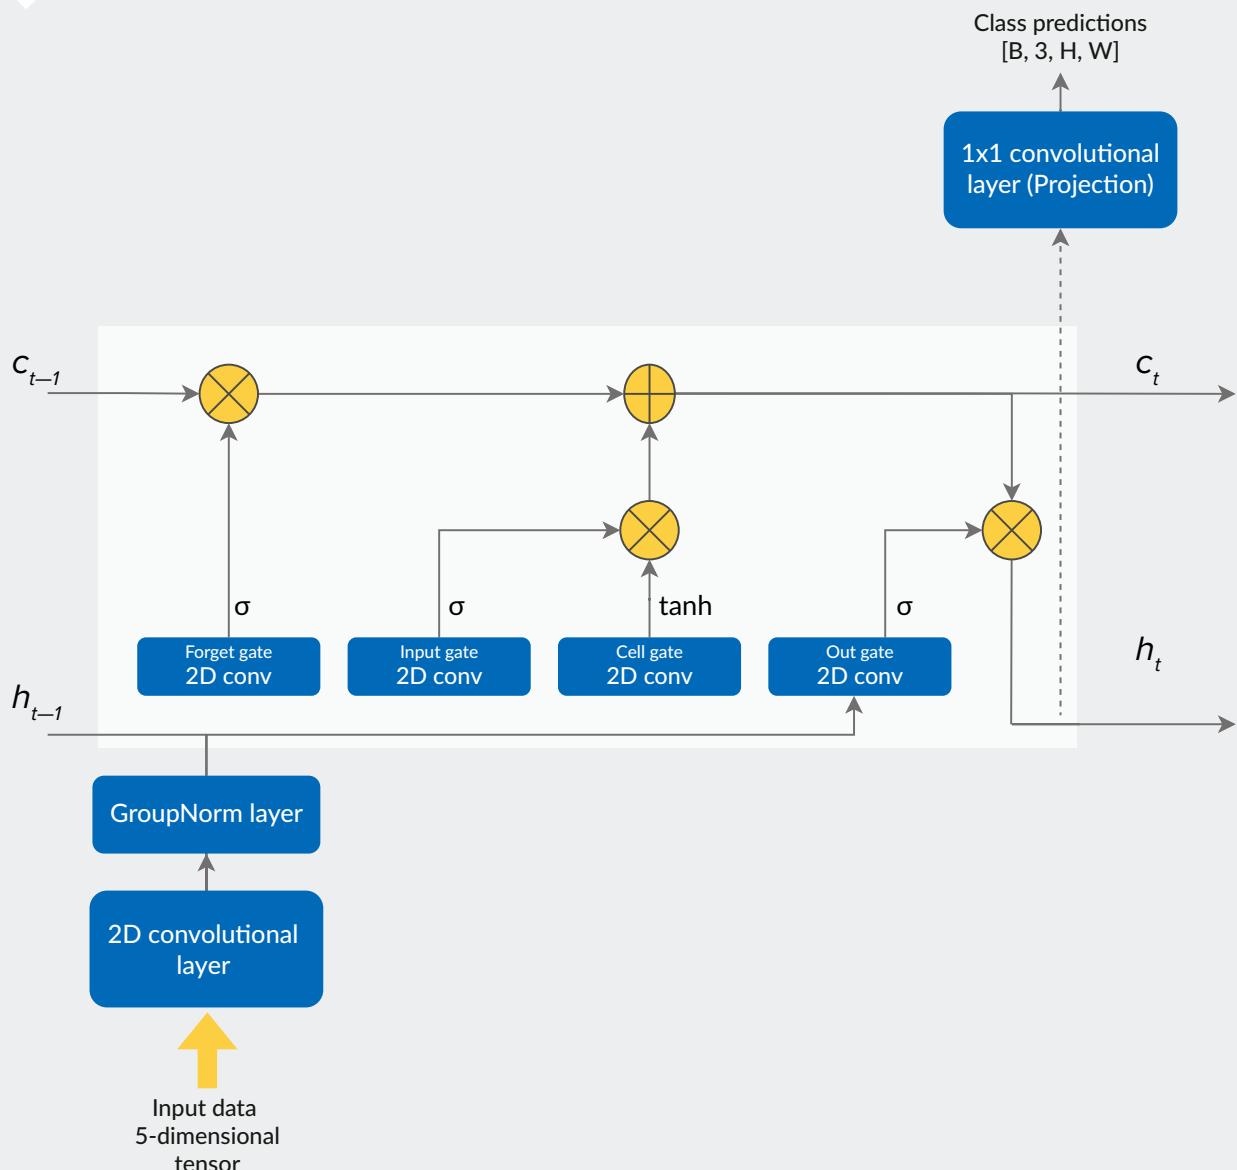


Figure 8. Architecture of the convLSTM model. σ refers to sigmoid activation and \tanh refers to hyperbolic tangent. The circles refer to element-wise multiplication or addition.

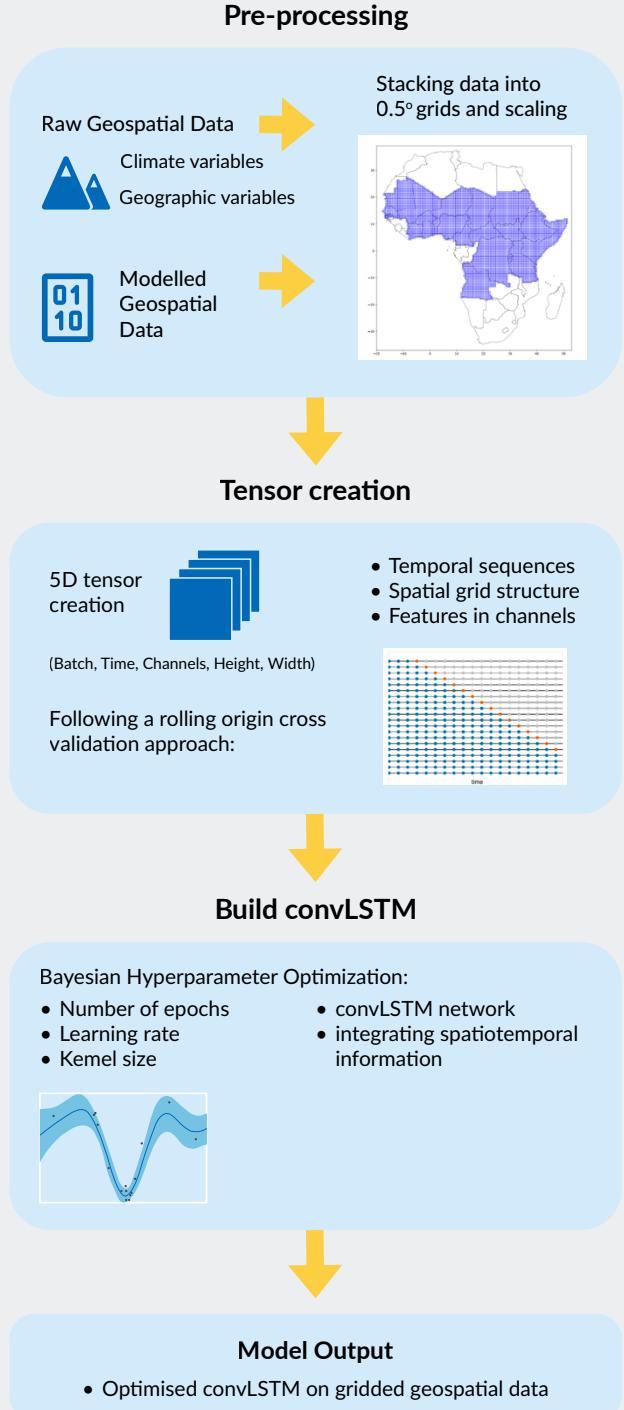


Figure 9. Framework of neural network methodology.

Again, we used up to 80 time points to capture long-range temporal dependencies. As with the CNN, the data are represented and processed as 5-dimensional tensors (batch, time, channels, height, width) which encode the spatiotemporal information and follow a cross-validation technique. The architecture consists of a single 2D convolutional layer and a GroupNorm layer that processes the concatenated input and hidden state. The input data set is processed sequentially through the convLSTM cell. At each timestep, the current input is concatenated with the previous hidden state and fed through the convolutional layer followed by the normalisation layer, producing features that are split into four sets of channels. The outputs are then split and fed into the four LSTM gates (input, forget, cell, and output gates) with sigmoid activations for the input, forget, and output gates, and tanh activation for the cell gate. The gates selectively decide which spatiotemporal information to retain in the cell state and which to discard. Unlike traditional LSTMs, this network is constructed so that the spatial structure is preserved through convolutions rather than fully connected layers, maintaining spatial resolution throughout the temporal sequence processing. After processing all timesteps, a 1x1 convolutional projection layer maps the learned spatio-temporal features to the final class-specific spatial prediction map for sequence-to-frame tasks.

Hierarchical convLSTM

The simple convolutional LSTM is then further extended through a hierarchical encoder-decoder architecture that captures local and global spatiotemporal patterns, following the approach of Shi et al. (2015). The encoder consists of two sequential ConvLSTM cells. The first cell extracts local spatiotemporal features maintaining the base number of hidden channels. These local features are then passed to the second cell processes, which learn more global spatiotemporal patterns, which is done by doubling the hidden channel size to allow for increased complexity. Each stage processes the full temporal sequence, with stage 2 taking the hidden state outputs from stage 1 as its inputs, allowing the network to build hierarchical spatiotemporal representations.

A final 1x1 convolutional projection layer maps the decoder's output to class-specific spatial predictions. The architecture implements several design choices to improve performance and flexibility. Each ConvLSTM cell includes adaptive GroupNorm layers that dynamically adjust the number of groups (ranging from 1 to 32) based on the channel count to ensure numerical stability. The progressive channel expansion from the base hidden channels in Stage 1 to double in Stage 2, followed by contraction back to the base size in the decoder, creates a bottleneck architecture that forces the network to learn compressed global representations. In particular, each ConvLSTM cell generates class predictions at each time step through internal projection layers of 1x1, enabling possible intermediate supervision,

although only the final output of the decoder is used for the final prediction. The decoder's efficiency is enhanced by processing only a single timestep rather than the full sequence, reducing computational cost while maintaining the benefits of hierarchical feature integration.

The decoder then combines the two scales of information through a fusion mechanism. This hierarchical approach basically enables specialized learning of local and global spatial patterns, which are then integrated to produce more accurate predictions. The architecture should provide superior spatiotemporal modelling compared to single-stage approaches, with the trade-off of increased computational cost.

Ensemble model

To exploit the strengths of different model architectures, we implement a weighted ensemble approach over all classification models. The ensemble weights are optimized on the validation set. For each forecast horizon and current displacement level, we determine the optimal linear combination of model predictions that minimize the average F1 score. To ensure that forecasts remain accurate and reliable, models are re-trained every three months incorporating the latest data and trends, and the ensemble temporally adjusts weights according to the models' predictive performance. This regular updating allows the system to adapt to changing conditions on the ground and maintain its usefulness in humanitarian planning.

Results

Each method is used to develop predictive models for the three crises classes defined above. We use the F1, and PR AUC scores for all classes to evaluate model performance. The F1 score is a harmonic average of the precision and recall metrics, calculated separately for each class:

$$F_1 = 2 * \frac{\text{precision} * \text{recall}}{\text{precision} + \text{recall}}$$

The harmonic average ensures that both metrics must be reasonably high for a good F1 score. We complement this with the PR AUC score for a probability-based metric which is also particularly suitable for class imbalance. The PR AUC summarizes the precision-recall trade off and ranges from 0 to 1, with higher values indicating better performance.

Our analysis shows that Ensemble modelling approaches show high predictive performance for displacement levels across all forecast horizons, with particularly interesting results for high displacement events. The LightGBM model achieves highest overall accuracy with a PR AUC of at least 0.64 for the 6 months horizon, while the Ensemble approach is better at capturing the most critical class 2 events. Even at 6 months into the future, the Ensemble model achieves a PR AUC of 0.50 for class 2, representing large-scale displacement events that account for less than 0.3% of the overall dataset. Table 2 below presents the detailed evaluation approach including results across all models, forecast horizons, and countries.

Table 2. Precision-Recall AUC per model and predictive horizon.

Horizon (months)	Model	PR AUC	F1
1	Baseline	0.87188	0.989252
	LightGBM	0.999654	0.990196
	Simple convLSTM	0.999227	0.982280
	CNN	0.998685	0.825868
	Hierarch. convLSTM	0.998400	0.979131
	Ensemble	0.999571	0.988363
3	Baseline	0.587037	0.756006
	LightGBM	0.848264	0.780281
	Simple convLSTM	0.773426	0.694041
	CNN	0.714724	0.655764
	Hierarch. convLSTM	0.755635	0.728660
	Ensemble	0.838713	0.757251
6	Baseline	0.435189	0.652512
	LightGBM	0.644033	0.598180
	Simple convLSTM	0.622529	0.589392
	CNN	0.609971	0.560585
	Hierarch. convLSTM	0.604661	0.642479
	Ensemble	0.699143	0.666627

Overall, model performance displayed in Table 2 shows that the LightGBM, followed by the Ensemble model, achieves the highest accuracy across all predictive horizons. Even at horizon 6, the PR AUC and F1 score remain relatively high. However, these aggregated metrics mask class-specific patterns, particularly for class 2, which represents less than 0.3% of observations but constitutes the most important prediction target.

Class-specific metrics in Table 3 show distinct modelling strengths. All models demonstrate consistently strong performance in predicting the majority class (class 0) with PR AUC scores exceeding 0.99 and negligible differences between modelling approaches. This reflects the prevalence of this class and the ability of the model to identify stable regions. Class 1, which represents moderate and ongoing displacement, is best predicted by the LightGBM model across all horizons, using the main environmental, geographic, conflict, and socioeconomic variables to capture non-linear dynamics.

Class 2 represents the most complex and important displacement dynamics, characterized by rare and volatile events which makes them difficult to forecast. The ensemble model performs best at capturing this

class for all horizons, showing its strength in leveraging different models. As expected, the metrics for class 2 are lower than the ones for class 1 but are still substantially higher compared to the Baseline model. Performance degradation is also low across longer predictive horizons for classes 0 and 1, while it is slightly more pronounced for class 2. This pattern is expected as class 0 represents most of the observations in the dataset and stable situations and class 2 represents volatile events which are more difficult to forecast over extended periods.

Testing different variable subsets also showed the importance of including climate variables into our framework. Including them improved the PR AUC score for class 2 by up to 5% mainly in East and West Africa, showing how critical they are for the most important class. Feature importance analysis also showed that they were among the top predictors across horizons 3 and 6 for class 2 and class 1. This highlights that climate variables are not only supplementary features but important components for understanding displacement risk, particularly for the most severe humanitarian crises.

Based on these results, our final prediction model is a hybrid approach where we use LightGBM for classes 0 and 1, and the Ensemble model for class 2 as our final model.

Table 3. Under Curve (PR-AUC) by displacement class and forecast horizons. Best values are in bold, second-best are underlined.

Class	Model	1 Month	3 Months	6 Months
0	Baseline	0.991458	0.991339	0.991797
	LightGBM	0.999728	0.999664	0.999600
	Simple convLSTM	0.999388	0.999090	0.998210
	CNN	0.999415	0.998260	0.999195
	Hierarch. convLSTM	0.998337	0.998489	0.998799
	Ensemble	0.999729	0.999537	0.999496
1	Baseline	0.613626	0.549013	0.496395
	LightGBM	0.882590	0.822296	0.770439
	Simple convLSTM	0.804354	0.656441	0.620376
	CNN	0.819927	0.747606	0.655166
	Hierarch. convLSTM	0.789561	0.710296	0.661075
	Ensemble	0.871661	0.788185	0.754425
2	Baseline	0.525052	0.356618	0.222105
	LightGBM	0.781365	0.568037	0.447263
	Simple convLSTM	0.704381	0.534503	0.351497
	CNN	0.678212	0.523443	0.424356
	Hierarch. convLSTM	0.660205	0.502952	0.406210
	Ensemble	0.874733	0.607959	0.490638

Precision - recall curve for class 2

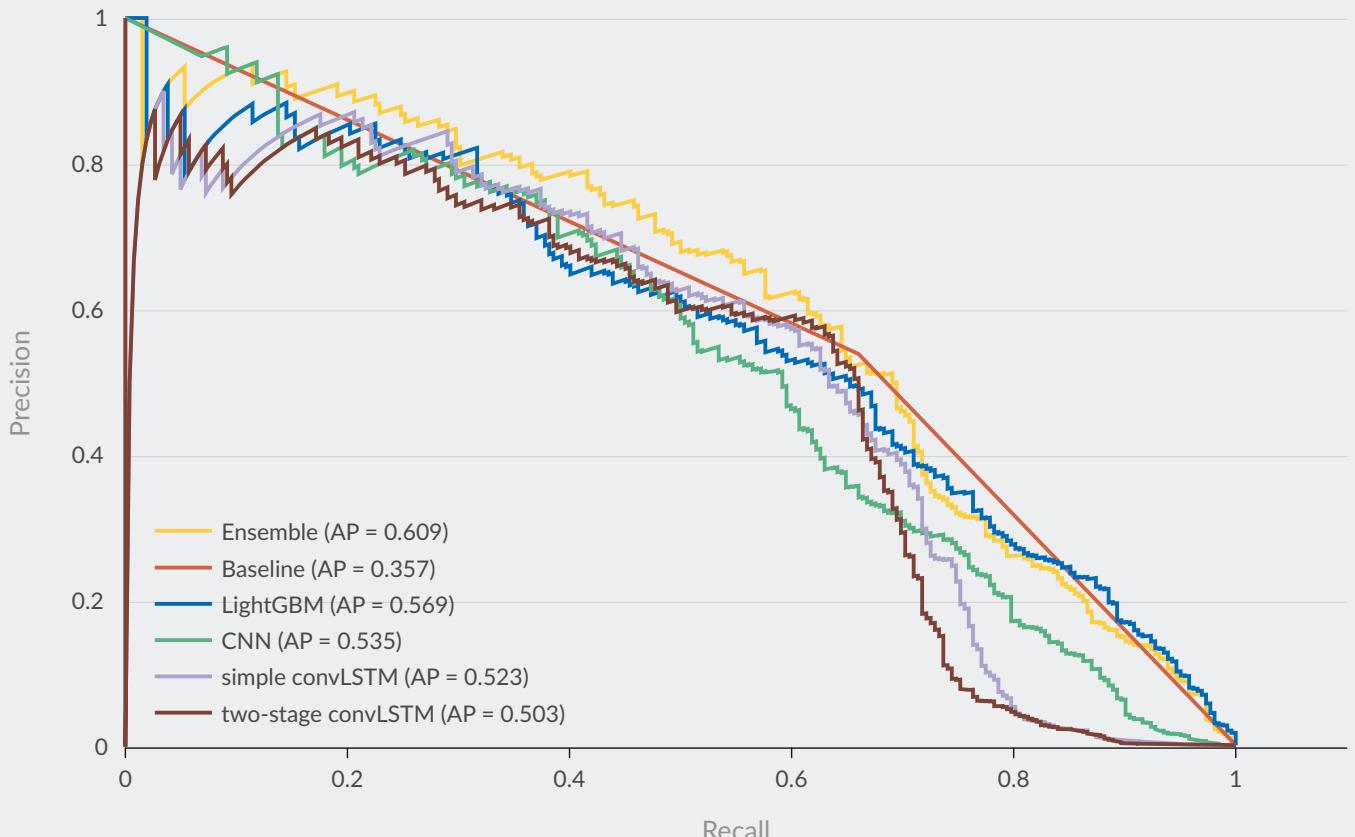


Figure 10. PR AUC for horizon 3 and class 2 over the testing period (2024-2025).

Figure 10 shows the PR AUC for the 3 months horizon and class 2 for all models, with the Ensemble model showing the highest PR AUC.

Figures 11 to 13 summarize the geographic accuracy of the final model. They represent the weighted PR AUC for each grid over the testing period and for the 1, 3, and 6 months horizons. The testing period has been selected to include the years 2024 to 2025. In the maps, the hue of the colour refers to different PR AUC levels, with the green colour indicating higher predictive accuracy, the orange colour medium accuracy, and red colour low accuracy. Different shades of the colours indicate displacement levels in the grid, with lighter colours corresponding to grids with lower displacement, darker colours to higher levels of displacement, and grids with dark colour and thick border representing areas of very high displacement, which are the most critical for humanitarian planning.

Overall, we can see that the model works very well at correctly capturing areas of low risk of forced displacement and produces few false positives where no displacement occurs. The more challenging regions to predict are the ones with isolated higher levels of forced displacement, which is unsurprising.

The highest levels of displacement over our testing period are observed in Burundi, parts of Sudan, South Sudan, Somalia, the Lake Chad region, and Mali. The model achieves particularly high accuracy in Burundi, DRC, and South Sudan where displacement tends to cluster spatially, and more historical data is available for training. In contrast, accuracy is lower in the Lake Chad basin parts of West Africa due to them being spatially isolated and having less historical data coverage for class 2 compared to Eastern Africa. As it is to be expected, prediction accuracy slightly decreases with longer time horizons.

Horizon 1

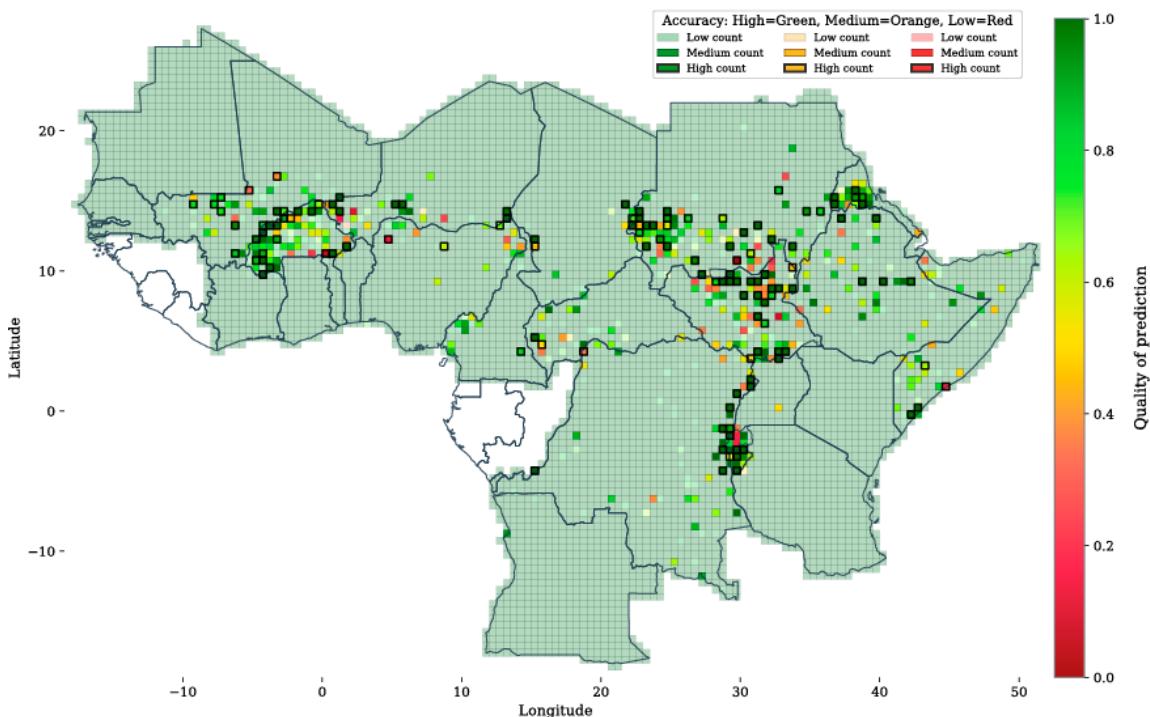


Figure 11 . Weighted PR AUC by grid over the testing period (2024-2025) for our selected model. Hue represents the metric, whereas colour depth refers to the level of forced displacement. Greener areas correspond to more accurate predictions and darker shades represent higher levels of displacement.

Horizon 3

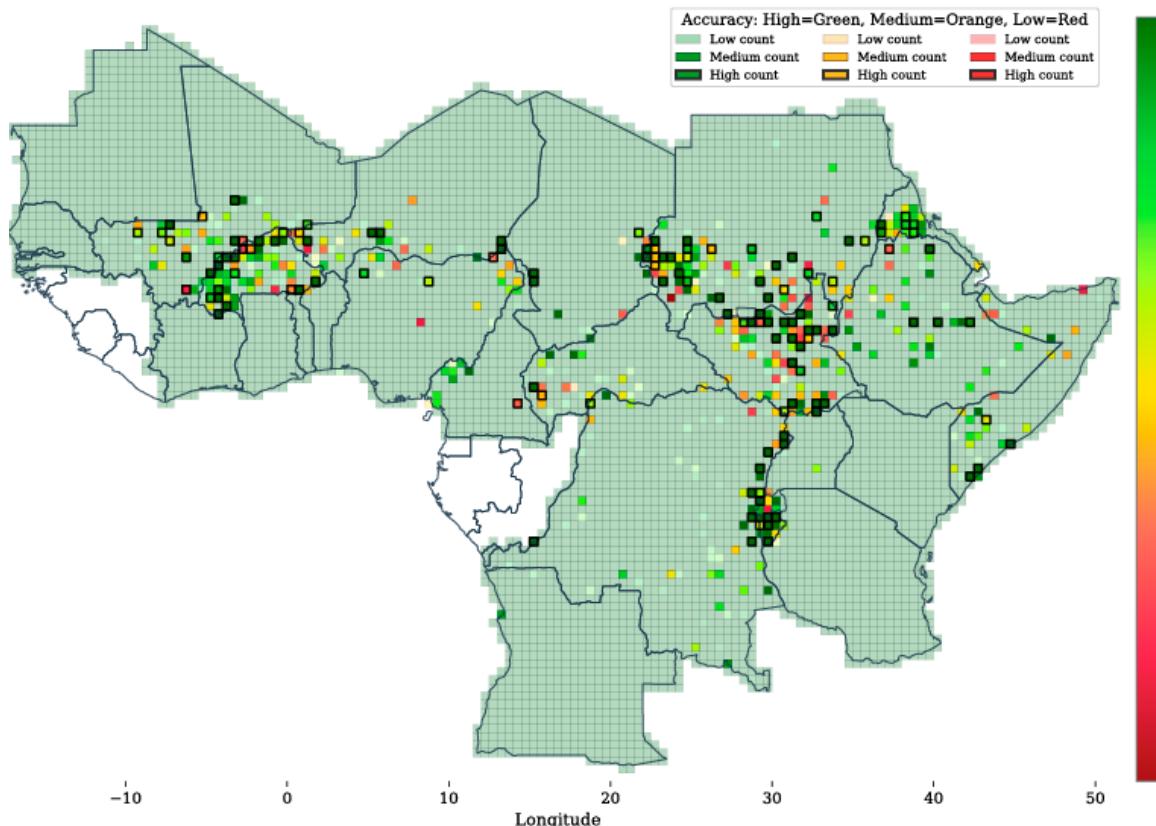


Figure 12 . Weighted PR AUC by grid over the testing period (2024-2025) for our selected model. Hue represents the metric, whereas colour depth refers to the level of forced displacement. Greener areas correspond to more accurate predictions and darker shades represent higher levels of displacement.

Horizon 6

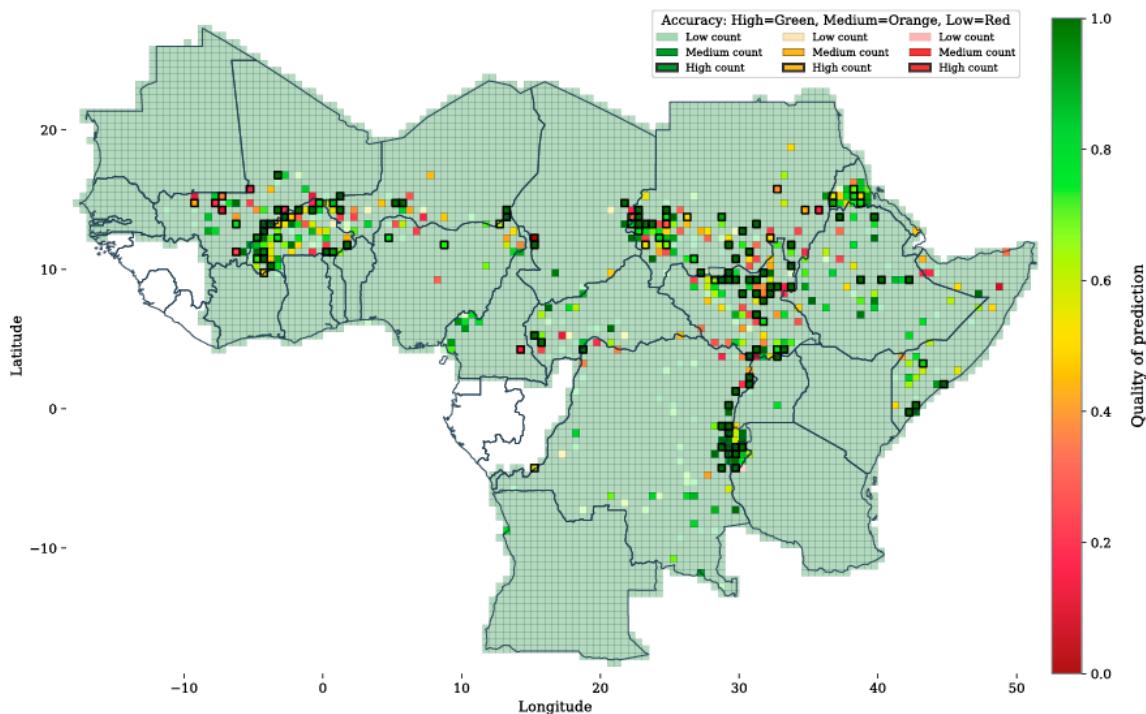


Figure 13. Weighted PR AUC by grid over the testing period (2024-2025) for our selected model. Hue represents the metric, whereas colour depth refers to the level of forced displacement. Greener areas correspond to more accurate predictions and darker shades represent higher levels of displacement.

The superior performance of LightGBM for classes 0 and 1 reflects its efficient handling of tabular features and ability to capture non-linear interactions between conflict indicators, climate variables, and socioeconomic factors without requiring extensive spatial modelling. The Ensemble model's advantage for class 2 stems from its integration of spatial and temporal patterns through convolutional and LSTM components. Large displacement movements often exhibit spatial patterns and temporal acceleration that tree-based models cannot easily capture. By combining LightGBM's feature-based predictions with deep learning's spatiotemporal modelling, the Ensemble leverages complementary strengths. This architectural choice proves particularly valuable for the rare but critical events that are important for humanitarian resource allocation.

Taken together, the results show the value of integrating tree-based and deep-learning approaches to achieve high accuracy in regions with severe class imbalance and complex spatiotemporal interactions between climate and conflict variables. LightGBM has been selected for regions which predominantly have displacement levels in classes 0 and 1, while the Ensemble model has been deployed in areas showing class 2 dynamics, where the increased complexity of our Ensemble model improves predictive performance while taking into account computational constraints.

Conclusion

This study represents a significant methodological advance in displacement forecasting by addressing a critical gap in the field: the mismatch between the spatial scales at which climate-displacement relationships operate and the aggregated levels at which they are typically modelled. By developing a gridded forecasting framework at 0.5-degree resolution across 25 African countries, we demonstrate that machine learning approaches can provide actionable early warning for forced displacement at the spatial granularity required to capture localized climate-conflict-displacement dynamics.

This work makes several methodological contributions to displacement forecasting. We introduce a comprehensive spatial gridding approach that preserves both temporal granularity (monthly observations) and spatial precision (0.5-degree resolution) while integrating diverse data sources—from satellite imagery to conflict event databases to socioeconomic indicators. The 180-variable feature set spanning geographic, demographic, environmental, conflict, political, and economic dimensions represents one of the most comprehensive covariate frameworks applied to displacement prediction at sub-national scales.

Our ensemble architecture, which combines the complementary strengths of tree-based feature learning and deep learning's spatiotemporal pattern recognition, demonstrates how hybrid approaches can address the multiple scales and temporal dynamics inherent in displacement processes. The hierarchical ConvLSTM architecture, in particular, shows promise for capturing both local displacement triggers and regional displacement systems, though at increased computational cost that may limit real-time deployment in resource-constrained settings.

However, several important limitations warrant acknowledgement. First, despite achieving operationally meaningful accuracy, our models struggle with the large displacement situations—

precisely the events of greatest humanitarian concern. The extreme rarity and inherent unpredictability of these events, combined with short spike durations (most lasting only 2-3 months), pose fundamental forecasting challenges that may require complementary approaches such as scenario analysis or qualitative expert assessment alongside quantitative prediction.

Second, our classification scheme, while operationally motivated, involves discretizing a continuous phenomenon in ways that may obscure important dynamics. The binning approach addresses statistical challenges posed by extreme skewness but sacrifices information about displacement magnitudes within categories. Future work should explore whether direct count regression, potentially with zero-inflated models or other approaches designed for rare events, can complement classification results.

Third, while our models capture correlations between environmental variables and displacement, they do not establish causal mechanisms. The “black box” nature of some deep learning components limits interpretability regarding how specific climate variables influence predictions. Future research should integrate causal inference frameworks and mechanistic understanding to improve our understanding of climate-displacement pathways.

As we approach an era where climate-influenced displacement may affect hundreds of millions of people, developing robust early warning systems represents both a moral imperative and a practical necessity. This study provides one methodological pathway toward that goal, demonstrating that with appropriate data, methods, and spatial resolution, we can begin to anticipate rather than simply react to forced displacement crises. The humanitarian community must now translate these technical capabilities into operational systems that can save lives, reduce suffering, and support more dignified responses to displacement in an era of climate change.

Bibliography

Abel, B. C. (2019). Climate, conflict and forced migration. *Global Environmental Change*, 239-249.

Beine, M., & Parsons, C. (2015). *Climatic factors as determinants of international migration. The Scandinavian Journal of Economics*, 723-767.

Black, R., Adger, W. N., Arnell, N. W., Dercon, S., Geddes, A., & Thomas, D. (2011). The effect of environmental change on human migration. *Global Environmental Change*, S3-S11.

Burke, M., Hsiang, S. M., & Miguel, E. (2015). Global non-linear effect of temperature on economic production. *Nature*, 235-239.

Cai, F. O. (2016). Climate variability and international migration: The importance of the agricultural linkage. *Journal of Environmental Economics and Management*, 135-151.

Carammia, M., Iacus, S. M., & Wilkin, T. (2022). Forecasting asylum-related migration flows with machine learning and data at scale. *Scientific Reports*, 12(1), 1457. doi:<https://doi.org/10.1038/s41598-022-05241-8>

Cattaneo, C., & Peri, G. (2016). The migration response to increasing temperatures. *Journal of Development Economics*, 127-146.

Chollet, F., & Watson, M. (2025). *Deep Learning with Python* (3rd ed.). Manning Publications Co.

Harari, Mariaflavia, & La Ferrara, E. (2018). Conflict, Climate, and Cells: A Disaggregated Analysis. *The Review of Economics and Statistics*, 594-608.

Henningsen, G. (2025). An Early Warning Model for Forced Displacement. Retrieved from <https://arxiv.org/abs/2505.06249>

Hoffmann Pham, K., & Luengo-Oroz, M. (2022). Predictive modelling of movements of refugees and internally displaced people: towards a computational framework. *Journal of Ethnic and Migration Studies*, 49(2), 589-613. doi:<https://doi.org/10.1080/1369183X.2022.2100546>

Ibáñez, A. M., & Vélez, C. E. (2008). *Civil Conflict and Forced Migration: The Micro Determinants and Welfare Losses of Displacement in Colombia*. *World Development*, 659-676.

Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., . . . Liu, T.-Y. (2017). LightGBM: A Highly Efficient Gradient Boosting. 31st Conference on Neural Information Processing System. Long Beach, CA, USA.

Lebakula, V. G. (2024). *LandScan*. Oak Ridge National Laboratory. Retrieved from <https://doi.org/10.48690/1532445>

Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2022). M5 accuracy competition: Results, findings, and conclusions. *International Journal of Forecasting*, 38(4), 1346-1364. doi:<https://doi.org/10.1016/j.ijforecast.2021.11.013>

Moore, W. H., & Shellman, S. M. (2004). Fear of persecution: Forced migration, 1952-1995. *Journal of Conflict Resolution*, 48(5), 723-740. doi:10.1177/0022002704267767

Moore, W. H., & Shellman, S. M. (2006). Refugee or internally displaced person?: To where should one flee? *Comparative Political Studies*, 39(5), 599-622. doi:10.1177/0010414005276457

Moore, W. H., & Shellman, S. M. (2007). Whither will they go? A global study of refugees' destinations, 1965-1995. *International Studies Quarterly*, 811-834.

O'Loughlin, J., Witmer, F. D., Linke, A. M., Laing, A., Gettelman, A., & Dudhia, J. (2012). Climate variability and conflict risk in East Africa, 1990-2009. *Proceedings of the National Academy of Sciences*, 109(45), 18344-18349. doi:<https://doi.org/10.1073/pnas.1205130109>

Schewel, K., Dickerson, S., Madson, B., & Nagle Alverio, G. (2024). How well can we predict climate migration? A review of forecasting models. *Frontiers in Climate*.

Schmeidl, S. (1997). Exploring the causes of forced migration: A pooled time-series analysis, 1971–1990. *Social Science Quarterly*, 78(2), 284–308.

Schmeidl, S., & Jenkins, J. C. (1998). The Early Warning of Humanitarian Disasters: Problems in Building an Early Warning System. *The International Migration Review*, 471-486.

Sebastian, K. (2009). Agro-Ecological Zones in Africa. Harvard Dataverse. doi: <https://doi.org/doi:10.7910/DVN/HJYYTI>

Shi, X., Chen, Z., Wang, H., Yeung, D.-Y., Wong, W.-K., & Woo, W.-C. (2015). Convolutional LSTM Network: A Machine Learning Approach for Precipitation Nowcasting. Retrieved from <https://arxiv.org/abs/1506.04214>

Suleimenova, D., Bell, D., & Groen, D. (2017). A generalized simulation development approach for predicting refugee destinations. *Scientific Reports*, 7(1), 13377. doi:<https://doi.org/10.1038/s41598-017-13828-9>

Tollefson, A., Strand, H., & Buhaug, H. (2012). PRIO-GRID: A Unified Spatial Data Structure. *Journal of Peace Research*, 49(2), 363-374. doi:[doi:10.1177/0022343311431287](https://doi.org/10.1177/0022343311431287)

Tomaso Poggio, Q. L. (2018). Theory I: Deep networks and the curse of dimensionality. *Bulletin of the Polish Academy of Sciences, Technical Science* 66, 761-773.

Vogt, M., Bormann, N.-C., Rüegger, S., Cederman, L.-E., Hunziker, P., & Girardin, L. (2015). Integrating Data on Ethnicity, Geography, and Conflict: The Ethnic Power Relations Data Set Family. *Journal of Conflict Resolution*, 59(7), 1327-1342.

Wells, A., Henningsen, G., & Kengne, B. B. (2025). Gridding Forced Displacement using Semi-Supervised Learning. Retrieved from <https://arxiv.org/abs/2506.08019>

Yan, X., Mohammadian, A., Ao, R., & Liu, J. (2022). Two-dimensional Convolutional Neural Network Outperforms Other Machine Learning Architectures for Water Depth Surrogate Modeling. *Journal of Hydrology*, 616(1-2), 128812. doi:[10.1016/j.jhydrol.2022.128812](https://doi.org/10.1016/j.jhydrol.2022.128812)



**COMPLEX
RISK
ANALYTICS**
Fund