

<https://doi.org/10.1038/s43247-024-01698-9>

Forecasting trends in food security with real time data

Check for updates

Joschka Herteux¹, Christoph Raeth²✉, Giulia Martini¹, Amine Baha³, Kyriacos Koupparis¹, Iliaria Lauzana¹ & Duccio Piovani¹✉

Early warning systems are an essential tool for effective humanitarian action. Advance warnings on impending disasters facilitate timely and targeted response which help save lives and livelihoods. In this work we present a quantitative methodology to forecast levels of food consumption for 60 consecutive days, at the sub-national level, in four countries: Mali, Nigeria, Syria, and Yemen. The methodology is built on publicly available data from the World Food Programme's global hunger monitoring system which collects, processes, and displays daily updates on key food security metrics, conflict, weather events, and other drivers of food insecurity. In this study we assessed the performance of various models including Autoregressive Integrated Moving Average (ARIMA), Extreme Gradient Boosting (XGBoost), Long Short Term Memory (LSTM) Network, Convolutional Neural Network (CNN), and Reservoir Computing (RC), by comparing their Root Mean Squared Error (RMSE) metrics. Our findings highlight Reservoir Computing as a particularly well-suited model in the field of food security given both its notable resistance to over-fitting on limited data samples and its efficient training capabilities. The methodology we introduce establishes the groundwork for a global, data-driven early warning system designed to anticipate and detect food insecurity.

Conflict^{1,2}, climate extremes, and soaring food, fertilizer and energy prices³ on the heels of an incomplete recovery from the COVID-19 pandemic⁴ have created a food crisis of unprecedented proportions⁵. The war in Ukraine further complicated the situation with millions of people a step away from starvation. In such an uncertain environment it is incumbent upon humanitarian agencies to deploy effective early warning systems that monitor food security conditions in the most vulnerable countries⁶. The World Food Programme (WFP) operates in emergency contexts and its operations include direct delivery of food assistance, cash based transfers, nutrition support, and are characterized by rapid response. This is done by leveraging an extensive logistic network in partnership with other UN agencies, NGOs, government counter parts and local communities, and requires swift assessments of local needs and vulnerabilities. In this context early warning systems serve as the foundation for preparedness and quick response to potential food crises, whether man-made or natural disasters. They help humanitarian organizations better target assistance to where it is needed the most, hence minimizing duplication and waste. Forecasting systems are an integral part of early warning systems, allowing them to anticipate potential hazards and issue alerts in a timely and effective manner. A robust approach to creating accurate forecasting tools is to leverage

modern time series prediction methods that are based on Machine Learning (ML)⁷. These methods can be applied to a growing number of available data streams, and have proven to be successful across a range of diverse fields. Examples of these fields include monitoring epidemics⁸⁻¹³, predicting financial markets¹⁴⁻¹⁷, tracking energy consumption¹⁸ and forecasting weather patterns and climate change impacts^{19,20}.

The scientific community has been studying and modeling food security for many decades and recent research has shown that ML can also be used in this context. By using ML, it is possible to develop complex, data-driven models with minimal feature engineering, thereby avoiding the need for an in-depth understanding of the underlying processes, provided that enough training data is available. This makes ML an attractive candidate for the study of food security, where the problem is complex and reliant on local conditions, making it challenging to generalize any claims that may lead to a knowledge-based model. An ML approach, if well designed, is able to reconstruct the underlying causal pathways without pre-established modelling choices, offering an alternative to the modelling from first principles. There exists a comprehensive body of literature on methods that leverage ML techniques, either standalone or in combination with other approaches. For instance, the Food and Agriculture Organization of the United Nations

¹World Food Programme, Research, Assessment and Monitoring Division (RAM), Via Cesare Giulio Viola 68, 00148 Rome, Italy. ²Deutsches Zentrum für Luft- und Raumfahrt (DLR), Institute for AI Safety and Security, Wilhelm-Runge-Straße 10, 89081 Ulm, Germany. ³World Food Programme Innovation Accelerator, Buttermelcherstrasse 16, 80469 Munich, Germany. ✉e-mail: christoph.raeth@dlr.de; duccio.piovani@wfp.org

(FAO) has investigated long-term forecasts at the country level for several years^{21,22}. In^{23,24} ML algorithms were used to classify households based on their caloric intake in Uganda, from data collected by the Bureau of Statistics. Studies carried out by the World Bank focus on the modeling and forecasting of Integrated Food Security Phase Classification (IPC) phases based on data from the Famine Early Warning Systems Network (FEWS.NET) across fifteen countries^{25–27}. The studies include attempts to identify the onset of food crises by building a binary crises indicator from the IPC phases and reconstruct the national share of population living in districts labeled as IPC3+²⁵. Moreover in ref. 26 forecasts sub-national IPC phases 1, 2, and 3+ employing a panel vector–auto regression model and using the Least Absolute Shrinkage and Selection Operator (LASSO) technique to pinpoint the most critical drivers of food security. Notably in ref. 28, based on the same FEWS.NET data, the authors extract significant text-based features from news articles accessed via Factiva. In their work they demonstrate that the text features greatly enhance the predictive capabilities of a Random Forest model, thus exploiting the potential of an abundant data stream that has not been thoroughly explored in the context of food insecurity.

Using data from the Living Standards Measurements Study (LSMS) in ref. 29, a LASSO regression was used to forecast food security indicators such as the Food Consumption Score (FCS), the Household Dietary Diversity Score (HDDS) and the reduced Coping Strategy Index (rCSI), in Malawi. In³⁰ an XGBoost algorithm is trained on secondary data in Ethiopia spanning a wide range of available data sets to forecast the future behaviour of the FEWS.NET IPC phases, classified as deteriorations, improvements and no change. A recent study from the French National Centre for Scientific Research (CNRS) has investigated use of Deep Learning approaches to forecast FCS and HDDS indicators as collected annually from the government in Burkina Faso. While researchers from UCLA recently used remotely sensed soil moisture data and food prices to predict changes in food security, using data provided by IPC in 7 countries²⁸. Leveraging real-time data collected by the World Food Programme, in ref. 31 authors developed a cross-country prediction model to *Nowcast* FCS and rCSI indicators across sub-national units in multiple countries, while in ref. 32 authors trained 30 distinct XGBoost models to construct a 30 day forecasting model for FCS. In this manuscript we used the same target and secondary variables used in refs. 31,32.

The World Food Programme regularly collects data on food security in countries where it operates. The methodology we present in this manuscript builds upon data sourced from the Real-Time Monitoring (RTM) system, developed by the organization over the last decade and publicly available via HungerMap^{LIVE}³³. Collected through Computer Assisted Telephone Interviews (CATI)³⁴, the data provides daily updates on food security indicators at the household level in near 20 countries. One such indicator is the Ffood Consumption Score (FCS), a composite measure that assesses the frequency of consumption of different food groups in the seven days prior to the survey, categorizing households as having poor, borderline, or acceptable food consumption. Households with insufficient food consumption refers to those with poor or borderline food consumption. The aggregation of this variable to the first administrative level in in countries is the target variable for our study. In contrast to other food security monitoring initiatives like IPC, CH, or FEWSNet, WFP's RTM program stands out due to several distinctions. Other programs release country updates at intervals of four or more months while the RTM program provides daily updates throughout the entire year. This unique feature enables real-time monitoring of the dynamic changes in food security levels, setting it apart as a more responsive and timely tool for assessing and addressing evolving challenges. Furthermore, the underlying process to generate the indicators in the RTM is purely quantitative, and does not rely on domain expert consensus or qualitative observations.

The primary aim of this research is to develop a forecasting methodology that can predict insufficient food consumption, as described in the WFP's RTM system, over a continuous 60 day period. While national development plans can depend on low-frequency data released on fixed

dates, the strategic allocation and distribution of resources in emergency scenarios require more frequent updates, such as monthly or weekly, and constant prioritization for optimal impact. Our early warning system is specifically designed for the dynamic and rapidly changing conditions of humanitarian aid operations in crisis settings. The forecasts are updated daily, enhancing their versatility and flexibility, making them continuously useful throughout the year, unlike other methodologies that are restricted by fixed report release schedules. This feature supports frequent re-targeting and optimization, which are crucial for effective humanitarian responses. Given the logistical challenges in hardship conditions where resource distribution can take up to 2 weeks, our 60 day forecasts are valuable for assessing the risk of deterioration and enabling planning based on updated priorities. Therefore, a 60 day forecasting horizon, with frequent updates, is considered operationally valuable for managing resource distribution programs effectively.

We tested out the methodology in four countries, Mali, Nigeria (north east), Syria, and Yemen, at the first sub-national level, using secondary data that represent the key-drivers of food insecurity: conflict, extreme weather events and economic shocks⁶. The selection of dataset categories for secondary variables in our study draws inspiration from the choices made in previous research^{25–32}, with some adjustments based on availability and access convenience. A detailed description is found Table 1. Our forecasting methodology innovates by leveraging RTM data as the primary target, capitalizing on the unique aspects of the monitoring and incorporating several desirable features that are particularly beneficial for informing humanitarian assistance. In order to increase the transparency and explainability of the model, an aspect that has been identified as crucial in the literature^{35–37}, we further implement a simple method for dynamic feature selection. This allows us to assess the importance of different variables in the model and their impact on the prediction.

The model we identify as the most suitable for forecasting trends in food insecurity is based on Reservoir Computing (RC)^{38–40}, a type of recurrent neural network (RNN). Several features make this approach particularly appealing in this context. First, when training the RC, only the parameters of the final layer are “learned,” while all other layers are kept constant. This simplification effectively makes RC equivalent to a linear regression, eliminating the need for gradient descent and reducing the number of parameters to learn drastically, speeding up the process and making it suitable for data-constrained contexts. Despite this simplification, the untrained layers and activation functions maintain the model's ability to capture non linear and complex patterns in the data. Additionally, RC allows for the dynamic treatment of secondary variables which are iteratively estimated along with the target variable. The model learns to project these secondary variables into the future, facilitating a straightforward treatment of multidimensional time series. Furthermore, the architecture of RC enables the easy addition of exogenous variables whose future values are known and do not need to be estimated. This feature enhances the model's flexibility and accuracy by incorporating external information directly into the forecasting process. Despite its simplicity RC has proven to be a valid alternative to advanced machine learning prediction techniques particularly in data-scarce contexts as can be found in refs. 41–43 when applied to synthetic data. Specifically our methodology relies on an ensemble of RC models, exploiting the low resource requirements of training RC and stabilizing predictions. This was also recently attempted in ref. 44 where RC is shown to outperform LSTMs on forecasting food prices.

To assess the goodness of our approach, we conducted a comprehensive evaluation by comparing the Root Mean Squared Error (RMSE) of various methodologies, that span deep learning, machine learning and classical statistical model: ARIMA, CNN, LSTM and the XGBoost from a previous attempt³². Inspired by the approach in ref. 30, we also developed a simple classifier to label future behavior as Deteriorating, No Change, or Improving based on the forecasted values. Our results demonstrate that the RC model outperforms other tested approaches in both types of tasks, and in particular in anticipating dramatic deterioration. This is notable given that curves that show deterioration and severe deterioration are less common in

Table 1 | Input data

| Variable | Yemen | Syria | Mali | Nigeria | frequency | spatial resolution |
|---------------------------------------|-------|-------|------|---------|-----------|--------------------|
| Insufficient food consumption | ✓ | ✓ | ✓ | ✓ | daily | sub-national |
| Food based coping | ✓ | ✓ | ✓ | ✓ | daily | sub-national |
| Rainfall | ✓ | ✓ | ✓ | ✓ | dekad | sub-national |
| Rainfall 1 month anomaly | ✓ | ✓ | ✓ | ✓ | dekad | sub-national |
| Rainfall 3 months anomaly | ✓ | ✓ | ✓ | ✓ | dekad | sub-national |
| NDVI | ✓ | ✓ | ✓ | ✓ | dekad | sub-national |
| NDVI 1 month anomaly | ✓ | ✓ | ✓ | ✓ | dekad | sub-national |
| ALPS/PEWI | ✓ | ✓ | ✓ | ✓ | monthly | sub-national |
| Food inflation | | | ✓ | ✓ | monthly | national |
| Headline inflation | | | ✓ | ✓ | monthly | national |
| Currency exchange official | ✓ | ✓ | ✓ | ✓ | daily | national |
| Currency exchange unofficial | ✓ | | | | daily | national |
| Battle fatalities | ✓ | ✓ | ✓ | ✓ | daily | sub-national |
| Violence against civilians fatalities | ✓ | ✓ | ✓ | ✓ | daily | sub-national |
| Remote violence/explosions fatalities | ✓ | ✓ | ✓ | ✓ | daily | sub-national |
| Seasonal calendars | ✓ | ✓ | ✓ | ✓ | dekad | sub-national |
| Ramadan | ✓ | ✓ | ✓ | ✓ | daily | national |
| Day of the year | ✓ | ✓ | ✓ | ✓ | daily | national |
| Number of input variables | 266 | 160 | 114 | 360 | | |

Variables used in the model for each of the four pilot countries. Time resolution of the data is indicated under frequency. The level of spatial aggregation is shown under spatial resolution. The last row shows the total number of input variables used in each model, which is summed over all sub-national time series available for the specified features.

the dataset than the stable curves. For these reasons we believe that the findings presented in this manuscript highlight Reservoir Computing (RC) as a robust option for forecasting applications in the field of food security, where high frequency data is scarce and skewed towards stable behaviour, and for producing early warning signals of impending food crises.

Results

The target variable of our modelling is the time series of the prevalence of households with insufficient food consumption at the first sub-national level. An example is shown in the left panel of Fig. 1, together with time series examples of a sub sample of the secondary data. Our choice is inspired by the approach used in previous works^{31,32} and makes use of available datasets describing conflict, economic shocks, and weather events. More specifically these consist of rainfall and vegetation levels, spikes in market food prices, food and headline inflation, currency exchange rates and number of fatalities due to different kinds of conflict. A detailed description of the input data is provided in the Methods section. Since Islam is the main religion in all four analysis countries, the holy month of Ramadan (fasting) is expected to significantly influence insufficient food consumption (see the Methods section). Therefore, the dates for Ramadan are used as an external input in the RC model, since it is available not only for the past, but also for the prediction period.

To standardize all input variables to the same format of sub-national daily time series we performed some basic pre-processing steps. These include linear interpolation of variables which are only available on monthly basis, aggregating information to the desired sub-national level, averaging intermittent variables over a rolling window as well as de-noising and smoothing all data. Given the stochastic nature of the RC algorithm, to generate our output we repeat the prediction procedure 100 times with different seeds, with the median of all results then considered the final prediction.

A comparison between the predicted curves obtained with the different tested methodologies and the actual data is shown in the plots in Fig. 2a where the blue curves are the actual data representing the sub-national curves of insufficient food consumption, the colored curves representing the

forecasts. In Fig. 2b, we can see how the variation of the curves is distributed in the dataset, where $\Delta = fcs_{60} - fcs_0$, expressed in percentage points, represents the difference between the last and first values in the forecasting window. A positive Δ implies that the last value of the window is higher than the first, therefore indicating an increase of food insecurity. On the other hand a negative value indicates an improvement. We refer to curves with positive Δ as deteriorating curves, while those with a negative Δ as improving curves. We can see from the figure that the dataset is skewed towards small values of Δ , indicating a majority of stable curves.

The objective of this work is to test the goodness of the predictions for 60 consecutive days. To ensure a comprehensive comparison among all considered algorithms and to assess their robustness across all seasonal trends, we have evaluated their performances in 12 60 day splits. The forecasts for each split starts at the beginning of every month, covering the period from June 2022 to May 2023. To simulate the use of such methodology in a production environment we employed a walk-forward validation scheme, meaning the model was retrained for every new forecast adding the data of the previous month, therefore expanding the training moving forward in time. At this stage retroactive delays in data updates were not factored in and the models were trained on all available data at present. In a production environment one must take into consideration the update frequency of the secondary variables as well as the target, that may not happen regularly. In section 2 of the Supplementary Material we further address this issue. For every split the hyperparameters of the different algorithms were updated and selected through a grid search of the performances on previous splits. Furthermore the algorithms were allowed to choose the subset of input variables that were found to be more informative, i.e. minimized the errors similarly to what was measured in ref. 31.

The aggregate errors are shown in Figs. 3 and 4 and are the result of ca. 70,000 CPU hours of calculations accumulated over several 10,000 runs for each country and model on the German Aerospace Agency's High Performance Cluster. This extensive computation period was essential to thoroughly assess the performances achieved with a vast array of hyperparameter combinations across all tested splits and countries, ensuring a fair

Input Data Examples

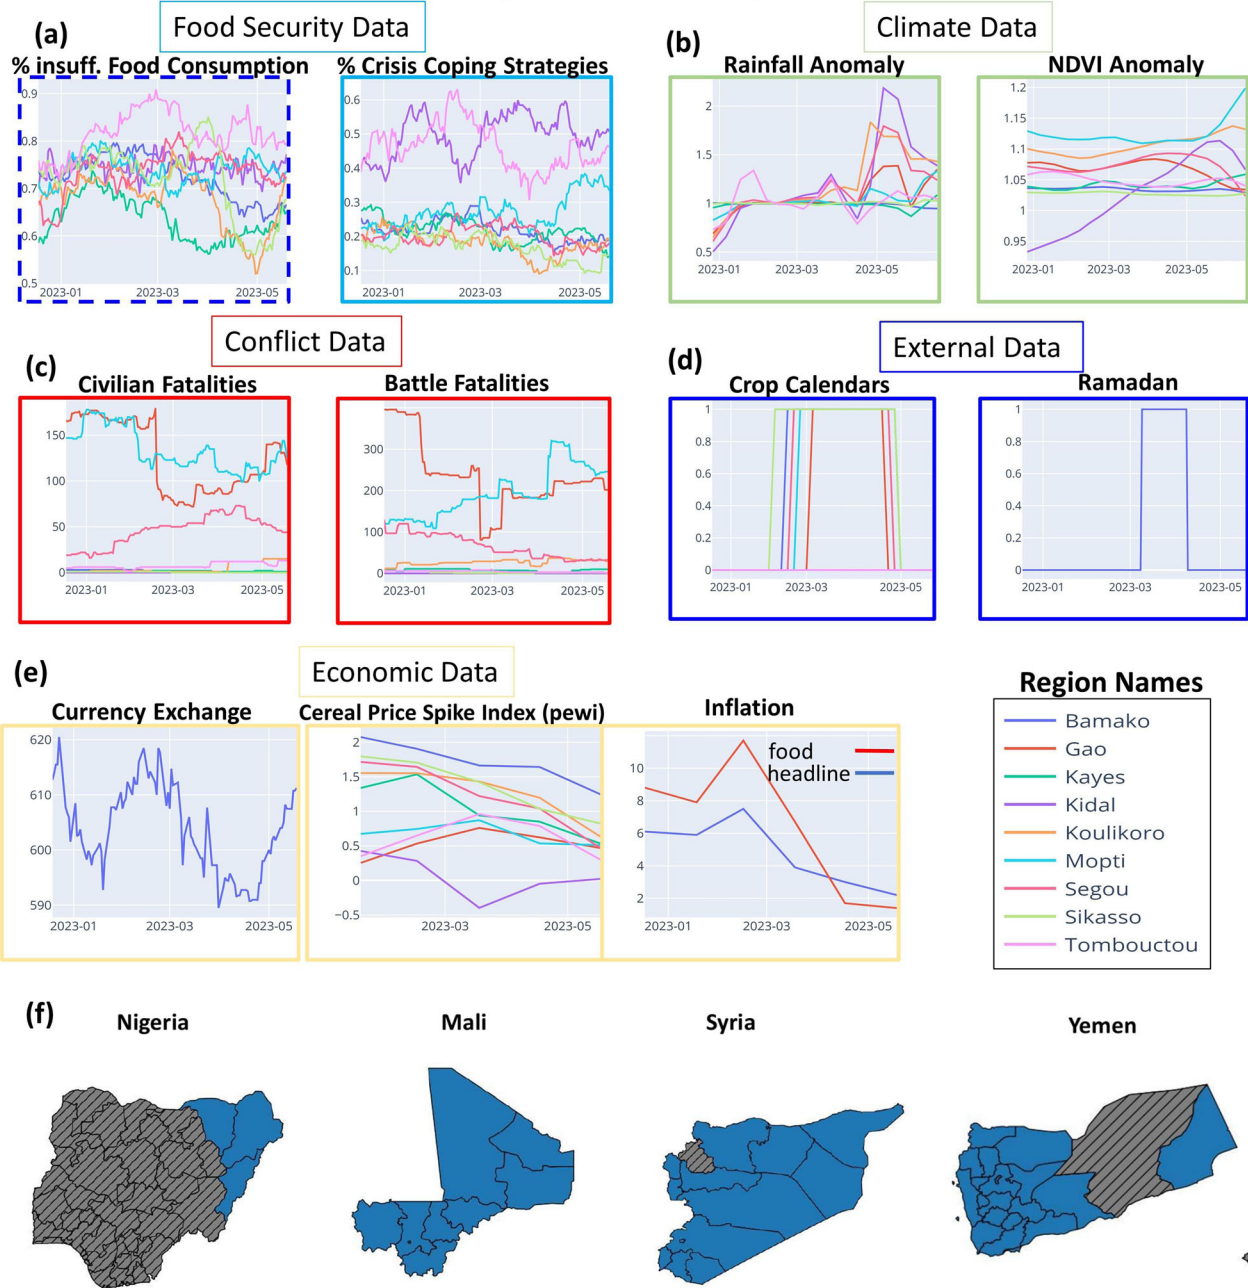


Fig. 1 | Input data. The figure shows the time series of the data used in constructing the forecasting methodology. The target variable, highlighted by the blue dashed curve in **a**, is the regional prevalence of insufficient food consumption extracted from the Food Consumption Score (FCS). In addition to the historical values of the target, the methodology incorporates predictors coming from another food security indicator, **b** climate, **c** conflict, and **e** economic data. External datasets with known future values, including crop calendars and Ramadan days, are also considered (**d**). While

the figures are based on data from Mali, the framework remains consistent across all countries. A detailed breakdown of the data for each country can be found in Table 1, with comprehensive information on data sources available in the methods section. **f** The map displays the first administrative level boundaries in the four tested countries, where grey dashed polygons indicate regions that were excluded due to data unavailability.

comparison between methodologies. It's noteworthy that the actual training times, on a local standard machine, for an individual hyper-parameter combination typically fall under 1 min for all methodologies (See the Supplementary Material section 5 for more details). As illustrated in Fig. 3, our analysis of the performances includes a comprehensive study from multiple angles, emphasizing systematic differences in the outputs.

In Fig. 3a we start by exploring the temporal evolution of the errors, showing the aggregate error for each forecasting step. We can see how the RC model's error rate increases at a slower pace compared to all other

considered algorithms. While the RC model is initially outpaced up to the 15th forecasting step, it distinctly becomes the preferable option beyond this point. Figure 3b presents the second type of aggregation for the errors: the degree of change in the data. In the figure we observe that the difference between the performances of the RC and the other tested model widens as the target variable undergoes more substantial change, with RC model outperforming all benchmarks by a clear margin under severely deteriorating conditions. We interpret the general increase across all models in errors for larger data variations as due to the infrequency of such extreme

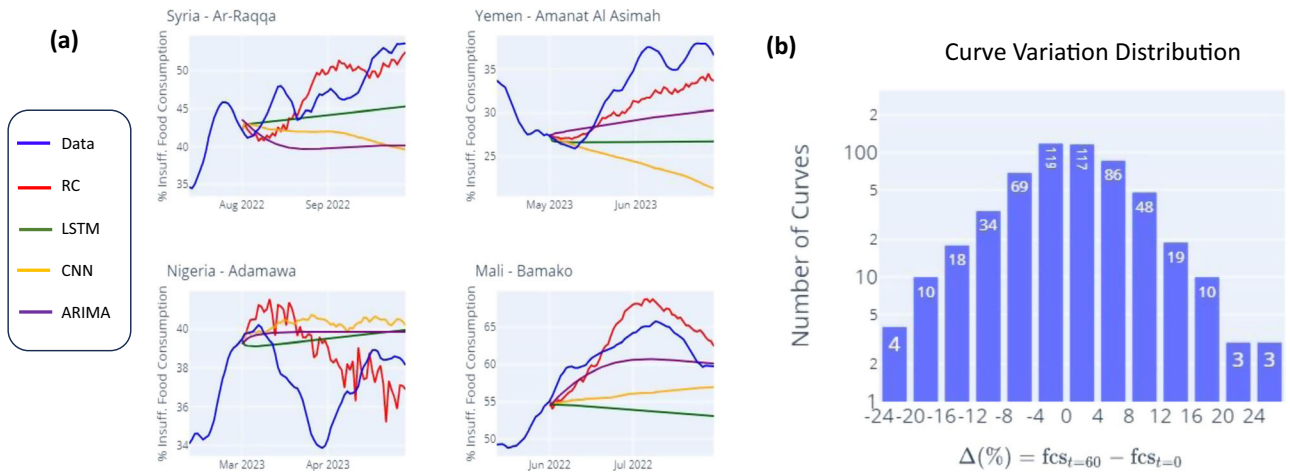


Fig. 2 | Forecasts and data. a 60 day forecasts examples generated using the RC, CNN, LSTM, and ARIMA models for four specific sub-national regions: Yemen, Syria, Mali, and Nigeria. In the visual representation, the blue curve represents the actual data, while each of the other curves depicts the prediction of one of the models.

b The observed distribution of the variation of the prevalence of insufficient food consumption on the 60 day windows used to train and test the algorithms. We can appreciate how the dataset is biased towards curves of small variation.




Fig. 3 | Forecasting aggregated performances. Performances of the LSTM, CNN and ARIMA models measured in median RMSE, aggregated in three different ways: Forecasting time step (a), variation of the target variable (b), per country (c). Figure (a) shows how the RC model outperforms the other methodologies after the 15th forecasted step with an aggregated RMSE at the end of the 60 day window of 4.9 percentage points. In figure (b) we see RC tends to be consistently the best performer when aggregating on curves according to the variation of their target variable

$\Delta = fcs_{60} - fcs_0$. This is more evident for high values of Δ , that indicate a sharp increase of levels of insufficient food consumption. The barplot in figure (c) aggregates the error per country. Despite the fact the RC is among the top performers for every country there is not a clear preferred methodology. This is due to the fact that regional time series belonging to the same country can have very different behaviours and have no clear national characteristic.

changes, as highlighted in Fig. 2b. Enhancing our approach to better capture these less common types of curves will be the aim of our future work. Figure 3c, shows the results of aggregating the methodologies on different countries, to assess the algorithms' robustness to different sets of available input data sets as well as diverse underlying relationships with the target variable. A complete reference of the available input data sets per country is found in Table 1. Figure 3c reveals that both the RC and the ARIMA models emerge as performing better than other methodologies overall, but no clear preferred choice emerges from this aggregation. We believe that this is due to the fact that regional time series of insufficient food consumption do not show any evident national property or characteristics and can assume very different behaviour, even within successive splits. The good performance shown by the simple ARIMA model, frequently among the top choices in all aggregation, aligns with findings from several other studies⁴⁴⁻⁴⁸, where the ARIMA model often surpasses or matches the effectiveness of more sophisticated algorithms. As a further evaluation of our proposed approach we compared the performance on 30 day forecasts with those shown in ref. 32 on the same dataset. The comparison is found in section 6 of

the Supplementary Material where we can see how RC once again stands out the best choice.

Inspired by ref. 30, we classified curves into three categories based on their variation: "No Change" for curves with a variation $|\Delta| < 0.04$, "Improvement" for a $\Delta < -0.04$, and "Deterioration" for $\Delta > 0.04$. Once again $\Delta = fcs_{60} - fcs_0$ is the measure of change seen in the curves. We also applied the same classification criterion to the forecasted curves, turning the regression problem into a classification. From Fig. 4 we can see that the RC model shows superior performance across all metrics with the exception of the accuracy on stable curves. Other methodologies tend to label the majority of curves as being stable therefore falling in the accuracy paradox, where high performances depend on the class unbalance of the dataset rather than on correct prediction. When looking at the minority class *Deterioration*, RC clearly stands out as the best choice in accuracy, recall and precision. It is worth underlying that all models, including the RC, show a conservative behaviour and tend to forecast a stable outcome. We can see this from the very low levels of the recall on the Deterioration class. This is in line with what was seen in ref. 25 where the False positive rates shown by the

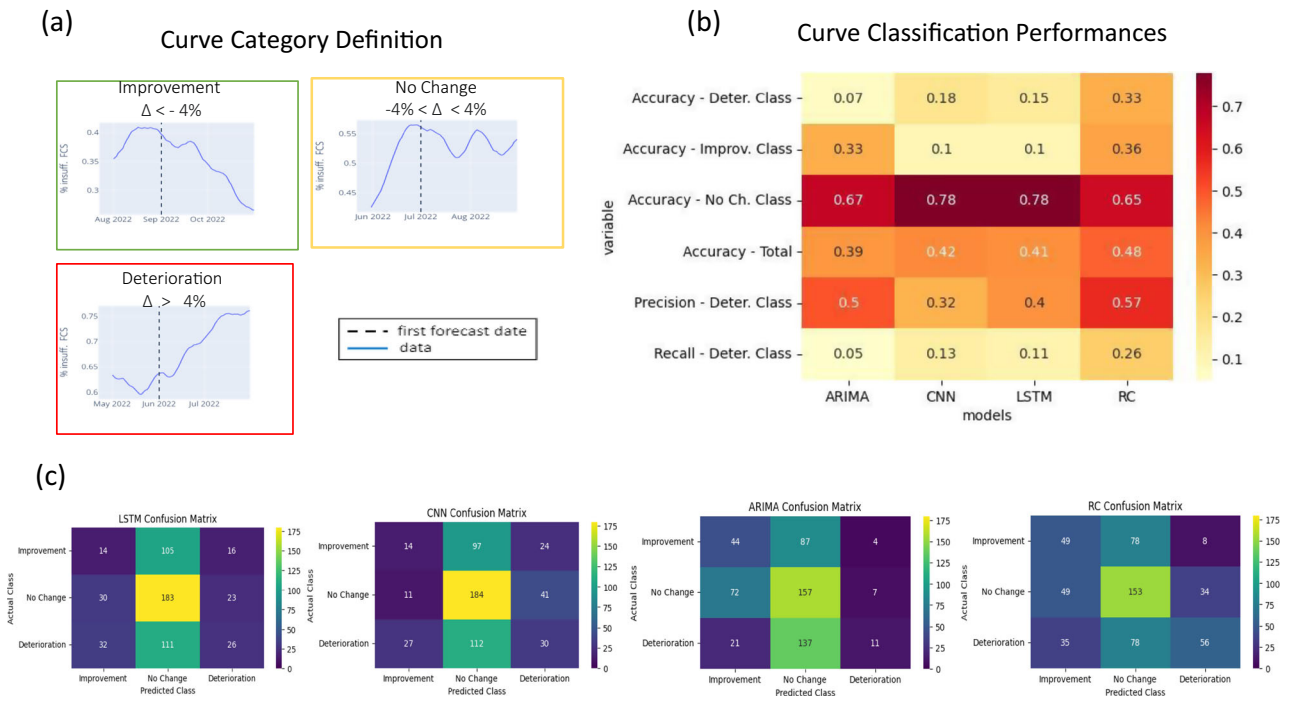


Fig. 4 | Classification performance. This figure compares the performances of the ARIMA, CNN, LSTM and RC models on their ability to distinguish different classes of behaviour: No Change, Improvement and Deterioration as defined in (a). The same criterion is applied to actual curves and predicted curves to label the behaviour.

b Performances metrics on the classification task. Accuracy on single classes are computed treating the problem as a binary classification. **c** Confusion Matrices of all tested methodologies.

methodology are one order of magnitude lower than the False negative rates, and is a limitation of the current methodology that will be addressed in future work.

Finally, we have studied what subset of secondary variables produced the lowest error for each split that was tested. In Fig. 5 the subsets of input variables are reported along with the frequency of their selection for every methodology. It is clear that most of the predictive power for all tested methodologies comes from the auto-regressive element, the rCSI indicator and the seasonal external variables. Understanding how to optimise the triangulation of the data sources will be the object of future work. We further analyse this important point in the Training Procedure sub-section of the Methods section.

Discussion

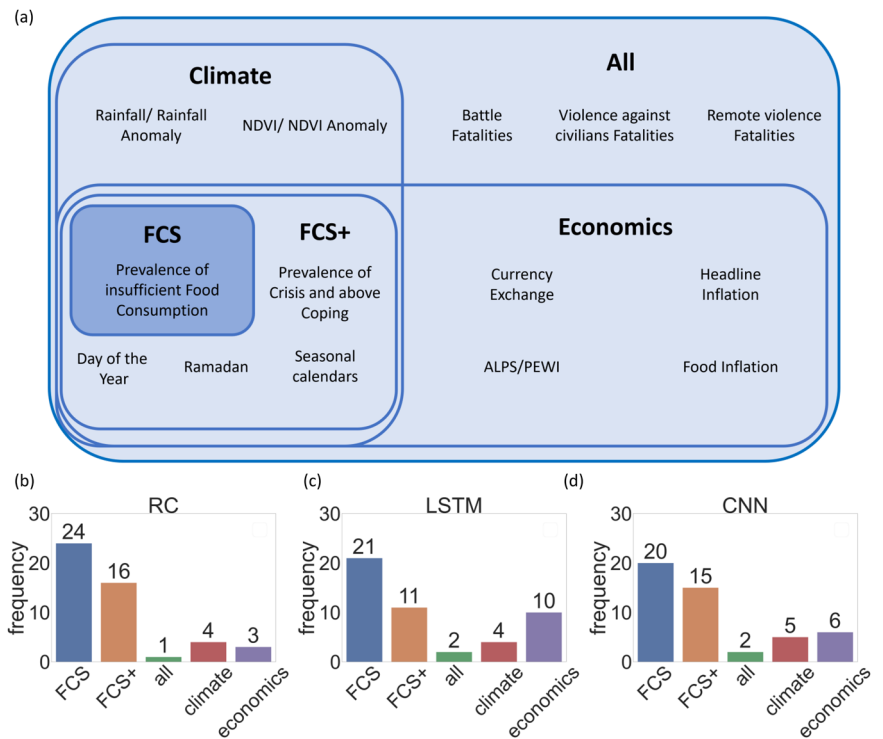
This manuscript applies several purely quantitative time series prediction methodologies to the problem of forecasting insufficient food consumption of a horizon of 60 days, based on the data collected by the World Food Programme’s real-time monitoring system. In direct comparison with LSTM, CNN and ARIMA we find Reservoir Computing to be the most suitable for this task in terms of performance, resource requirements, robustness and ease of use. We further proved this approach to surpass a previous attempt using XGBoost³² (See supplementary material section 6).

The Reservoir Computing algorithm has demonstrated to be a solid framework for the construction of a purely quantitative early warning system, owing to several advantageous features. Its ability to process multidimensional time series and learn complex nonlinear relations in target and features gives it the advantage over simple statistical models like ARIMA. The iterative prediction scheme allows the use of exogenous variables like seasonal and religious calendars where future values are already known as well as make implicit forecasts for those where this is not the case. The relative simplicity of the approach mitigates the danger of overfitting when compared to the more sophisticated Deep Learning methodologies CNN and LSTM, in line with what was shown in ref. 44. The limited resources needed for the training process allow an ensemble approach, the size of which can be finely tuned and adjusted based on the

available computational resources and accuracy requirements of the field and context of application. Furthermore, we saw how the RC methodology was able to accommodate a degree of class imbalance in the training data. Indeed the above benchmark performance shown by the RC model does not depend on the majority of stable curves that form the training set, but is actually more evident in the subset of curves that represent dramatic increases in the levels of insufficient food consumption. These findings suggest the potential advantages of RC over models of this kind in data-scarce contexts, even when confronted with high-dimensional, noisy, and authentic datasets. The ensemble architecture appears to perform well without requiring specific modifications for large input dimensions, as seen in approaches like Local States⁴⁹. Although the comparisons with methodologies based on different target variables is not straightforward, its worth noting that the performances measured with RMSE are in line or better of what presented in other works. In ref. 26 with IPC3+ target the RMSE is of 13%, while in ref. 50, the best model performs between 5% and 13% RMSE, although on different time horizons. In ref. 25 the performances vary considerably for different countries up to a maximum of 14.7% RMSE for near term forecasts and to 17.5% on medium term. Other works used different metrics to measure the errors, such as R^2 , making it more difficult, to compare with what obtained in this work.

The results shown in this manuscript indicate that a forecasting methodology based on the RC algorithm is able to predict 60 days into the future the levels of insufficient food consumption with an average RMSE of 5 percentage points. This increases to around 10 percentage points when only focusing on the subset of curves that show a severe deterioration. From the classification exercise we see that the RC correctly forecasts the type of behaviour once every two forecasts (Total accuracy 0.48), and the number goes down to once every three forecasts when only considering classifying future behaviour as deteriorating or not (accuracy 0.33), or improving or not (accuracy 0.36), according to our definitions. We believe that this is a solid starting point for the development of an early warning system based on the RC forecasting algorithm. More work will be done both to improve the general performances and to tailor the methodology to on deteriorating curves. To achieve this goal we will concentrate our future work on four

Fig. 5 | Feature selection. Feature groupings that were considered in the grid search for the RC, CNN and LSTM models to implement a dynamic feature selection. **a** Schema of how the features were classified into 5 different feature groupings. **b–d** frequency with which each feature grouping was selected for in (a) the RC model, (b) the LSTM model and (c) the CNN model.



main pillars: extending the forecasting horizon, tailoring the methodology to detect severe deterioration’s of food security levels, developing methods to explain the role of the input variables and implementing the methodology in a production environment. Initial results obtained on 90 day forecasts confirm the RC algorithm as the preferred choice with respect to other benchmarks, despite a general increase of the errors (see Supplementary Material section 1). Aggregating the underlying data to weekly or monthly time granularity, reducing the number of data points that need to be forecasted and smoothing uninformative fluctuations, seems like a promising possibility. This would also eliminate the need for interpolation and other pre-processing steps on secondary variables, thereby facilitating the deployment and maintenance of the approach. Introducing a custom weighting system in the training process and duplicating data splits which show severe deterioration will be attempted to focus the methodology on deteriorating curves, that are the most important to anticipate. Also testing other and more modern architectures of Reservoir Computing like the Next Generation RC⁵¹ or the Trend-Seasonality decomposition RC⁴⁴ variation might represent a valid choice to improve overall performances and extend the horizon. In section 2 of the supplementary material we have described a procedure to compute the confidence intervals of the forecasted curves. The confidence intervals may be added considering ± 2 -standard errors confidence bands. This methodology can be applied to all algorithms including RC. Finally, allowing the algorithms to autonomously select the most informative subset of features for each forecast, as done in this manuscript, marks an initial effort to comprehend the underlying drivers influencing the predictions. A comprehensive analysis of the results and the continued development of this approach will be crucial for enhancing the explanatory power of the methodology.

In conclusion, this research lays the groundwork for a dedicated quantitative tool for modeling and forecasting food insecurity. Our approach complements existing methodologies with its year-round versatility by offering real time updates on future levels of food insecurity. By leveraging the advantageous features of real-time monitoring, the methodology proves particularly valuable for foreseeing and quantifying the impact of events that occur between the releases of other monitoring systems. Our work presents the potential to furnish humanitarian

organizations with the capability to rapidly institute early warning systems to preemptively help vulnerable communities. The utilization of extant and expanding datasets in conjunction with the low requirement of computational power renders the proposed approach simple and economical to implement. As such, it could offer much-needed technical support to humanitarian actors in their efforts to anticipate, quantify and promptly address impending food crises.

Methods

In this section, we provide an in-depth exploration of the forecasting methodologies employed in our study, along with a comprehensive overview of the input datasets utilized to generate the results discussed in the preceding sections. Particular emphasis is given to the selected Reservoir Computing framework, because of it being less utilised and its central role in this work. In addition, a detailed overview of the training procedures and the hyperparameter selection process for each methodology is shown.

Data

As mentioned in the Introduction and Results sections, the models were trained on data representing the prevalence of food consumption and coping strategies as well as the global key-drivers of food insecurity. The following paragraphs introduce each of the input data sources and the features extracted from them. In Table 1 we can see the availability of the data for each country.

Prevalence of people with insufficient food consumption. This measure is built on the Food Consumption Score (FCS), a commonly used food security indicator, and is obtained estimating the percentage of households in a region with an FCS below a country-dependent threshold indicating poor or borderline food consumption. This way a single value between 0 and 100 per region is gained. The Food Consumption Score is collected by WFP both by face to face surveys, which happen approximately once a year and through CATI (Computer Assisted Telephone Interviews) on a daily basis. The score evaluates the amount as well as the nutritional value of a household’s food consumption. For some regions

not enough data is available and they were thus excluded from the study. This is the target variable of our modelling.

Prevalence of people using crisis or above crisis food based coping. Similar to insufficient food consumption, the prevalence of people using crisis or above crisis food based coping is based on a food security indicator, the reduced Coping Strategy Index (rCSI), which represents a different dimension of the concept. The rCSI is calculated from the results of the same surveys as the FCS and therefore the same methodology is used. Again a prevalence or percentage is defined by counting the number of people in a region scoring below a certain threshold. The rCSI is in particular a measure of the coping strategies adopted by the victims of food insecurity. Examples include limiting the sizes of food portions or the amount of meals per day, borrowing food or relying on help from friends or relatives, and choosing less preferred or less expensive food. Since both indicators stem from the same source, the availability of the data is the same.

Rainfall and vegetation. Five rainfall and vegetation features were constructed, sourced from WFP’s Climate Explorer⁵² and are originally provided by CHIRPS⁵³ (rainfall) and MODIS⁵⁴ (vegetation) satellite imagery. These include the amount of rainfall in mm and the NDVI (Normalized Difference Vegetation Index) value, which is a measure for the amount of vegetation in a region. In addition the anomalies of both indicators are included. The anomaly is calculated by dividing the average value of a variable during a period (e.g. 1 month) by the historical average for the same period in all previous years. To inform the model about the comparison of current and historical seasonal behavior we include the 1- and 3 month anomaly of rainfall and the 1 month anomaly of NDVI. From these WFP derives regional seasonal calendars which mark agriculturally relevant seasons. They can be considered a binary time series of either 0 or 1. All of these variables are given in 10 day intervals (“dekads”) for every region.

Conflict. To represent conflict as a driver of food insecurity we use data from the Armed Conflict Location & Event Data Project (ACLED)⁵⁵ which is a publicly available repository providing real-time and historical data on political violence and protest events in nearly 100 countries⁵⁶. We build three features from the fatalities reported for events in an administrative region of the categories Battle, Violence against Civilians and Explosions/Remote Violence by summing each of them in a rolling window over the last 90 days in order to create daily time series for each subnational region.

PEWI index. Food prices are an important factor for food security as they affect how much food a household can purchase. The forecasting power of this kind of data has been shown empirically⁵⁷. The ALPS indicator (Alert for Price Spikes)⁵⁸ measures relative spikes in monthly prices by comparing them to estimated values and is calculated monthly as follows

$$ALPS = (Price_t - \hat{Price}_t) / \sigma_\epsilon \quad (1)$$

where t gives the month and σ_ϵ describes the historic standard deviation of the residuals ($Price_t - \hat{Price}_t$ for previous months). WFP monitors commodity prices in local markets monthly and makes them publicly available through its Economic Explorer⁵⁹. For our model, we only use commodities in the “cereals and tubers” category. The data is averaged over all commodities within this category and across all markets in a given region. This aggregates the data to single time series per first-level administrative unit. The ALPS indicator is sometimes also referred to as PEWI (Price Early Warning Indicator).

Macro-economics. To gain a better understanding of the overall situation in a country we also look at several economic indicators, which are available at the country level. We include food and headline inflation

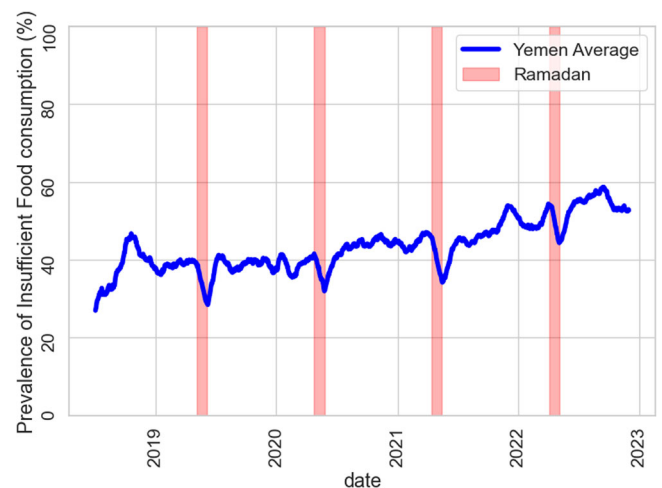


Fig. 6 | Ramadan effect. Prevalence of people with insufficient food consumption averaged over all available regions of Yemen in blue. Time periods of Ramadan shown by areas shaded red.

to complement our monthly data on food commodity prices. We also use daily effective exchange rates for local currency. All of these indicators can be found in WFP’s Economic Explorer⁵⁹.

Ramadan. Ramadan is the month of fasting in the Islamic faith, which is the predominant religion in all four countries considered in this study. The traditional practice of fasting for the whole month can clearly be expected to have a strong influence on food consumption and in the way it is self-reported. The exact relationship is likely complex, but the impact can empirically be seen in the data. In Fig. 6 we can see how the target data averaged over sub-national units is affected. The target time series consistently starts to go down during the time of Ramadan and recovers to roughly the same level afterwards in a similar time frame. These dips are among the most noticeable features of the time series. To accommodate this, we include the occurrence of Ramadan as a binary variable. If a date lies during Ramadan, this feature has a value of 1, otherwise it is 0. Since we know even for future dates if they occur in Ramadan, this feature is processed separately as external data. It is available during prediction and does not need to be inferred.

Day of the year. To give the models the ability to learn seasonal effects as well as possible we include a feature to represent the day of the year as a number ranging from 1 to 365. Similarly to Ramadan this feature is of course known for the future and can be used as an external variable.

Algorithms

For this study we have investigated a number of forecasting methodologies ranging from the simple statistical model ARIMA (Autoregressive Integrated Moving Average) to the Deep Learning techniques CNN (Convolutional Neural Network) and LSTM (Long Short-Term Memory) where Reservoir Computing (RC) can be understood as a hybrid between the two extremes.

Ensemble reservoir computing model. Reservoir Computing is a simplified Recurrent Neural Network(RNN) algorithm known to be easy to train and to be well-suited for problems with limited amounts of data, where it outperforms other state-of-the-art prediction algorithms⁴¹⁻⁴³. It is able to process multidimensional inputs as well as past values through its untrained hidden layer called the reservoir and use them to iteratively predict multidimensional time series. As an RNN it is well suited for temporal, sequential data. Because only the last readout layer is explicitly trained through a simple linear regression, it is still considered less complex than typical Deep Learning methodologies like LSTM. To utilize

the low resource requirements of RC and stabilize any random fluctuations caused by the initialization of the untrained network weights we implemented a simple ensemble scheme for the model by training 100 separate models with the same hyperparameters similar to what has been applied in ref. 44. For each single model only the seed for the random creation of the reservoir network is changed. The final output of the ensemble model is derived by calculating the median of the compilation of single outputs.

The variant of RC we implement for each single model is based on the original version of RC known as Echo State Network (ESN)³⁹. An ESN is a Recurrent Artificial Neural Network consisting of three layers: an input layer, a hidden reservoir layer and a read-out. While the input layer is simply responsible for distributing the incoming information in the network and the function of the readout is to shape the final output, the reservoir's purpose is firstly the projection of the data to a high-dimensional, nonlinear space and secondly to serve as the memory of the model. The dynamics of the ESN center around the reservoir state $\mathbf{r}_t \in \mathbb{R}^N$, which develops according to the update equation

$$\mathbf{r}_{t+1} = \tanh(\mathbf{A}\mathbf{r}_t + \mathbf{W}_{in}\mathbf{x}_t) \quad (2)$$

where the adjacency matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ represents the network, $\mathbf{x}_t \in \mathbb{R}^{d_x}$ is the input fed into the reservoir and $\mathbf{W}_{in} \in \mathbb{R}^{d_x \times N}$ is the input matrix. The hyperbolic tangent is a standard choice for the activation function. Since the state of the reservoir is a function of all the past inputs it received, it retains some information about the past if they are provided in sequential order. This is what makes ESN and RC in general well suited for time series analysis and prediction. We create \mathbf{A} as a sparse random network since low connectivity has been found to be advantageous⁶⁰. The weights of the network are then drawn uniformly from $[-1, 1]$ and afterwards rescaled to fix the spectral radius ρ to some fixed value. ρ is a free hyper-parameter. We chose \mathbf{W}_{in} to be also sparse, in the sense that every row has only one nonzero element. This means every reservoir node is only connected to one degree of freedom of the input⁶¹. The nonzero elements are drawn uniformly from the interval $[-1, 1]$ and then rescaled with a factor s_{input} which is another free hyperparameter called input strength.

From this we can then compute the output $\mathbf{y}_t \in \mathbb{R}^{d_y}$. For our predictive model we train the ESN to approximate $\mathbf{y}_t \approx \tilde{\mathbf{x}}_{t+1}$. Where $\tilde{\mathbf{x}}$ is not just the target data, but also all of the secondary data that is not known for the future. This way the model also learns to estimate the secondary variables and project them into the future. The readout is characterized by

$$\mathbf{y}_t = \mathbf{W}_{out}\tilde{\mathbf{r}}_t \quad (3)$$

where $\tilde{\mathbf{r}} = \{r_1, r_2, \dots, r_N, r_1^2, r_2^2, \dots, r_N^2\}$. This is a commonly used⁶² nonlinear transformation of \mathbf{r}_t which additionally serves to break the anti-symmetry the equations would otherwise have⁶³. The readout matrix $\mathbf{W}_{out} \in \mathbb{R}^{d_y \times N}$ is the only part of the ESN that is trained. This is done via simple Ridge Regression⁶⁴ controlled by the regularization parameter β . The resulting model is as a first step capable of forecasting the target variable as well as the required secondary variables one day into the future. To do so the time series of past input data is fed into the reservoir sequentially. During this process the reservoir stores information about past and present values of these variables. The readout can then be used to map from the reservoir to the prediction. For the purpose of predicting time series of any length the model is used as a "closed loop". The output is fed back into the model allowing the sequence of input data to continue. This way the forecast can be extended to any length, although the error tends to grow with every iteration, thus effectively limiting the reasonable time window. Secondary variables have to be projected into the future in the same way for the procedure to work. Since this creates an additional source of error, we developed a way to include future values of the data in case they are known. For some features like Ramadan or growth seasonality, we do not train the model to forecast the values. Instead, the known future values are used as input during the respective prediction steps. An avenue for future research

and an improved feature selection could be the inclusion of independent forecasts for secondary variables like rainfall and conflict, since the RC model is not optimized to predict their behavior.

ARIMA. The Autoregressive Integrated Moving Average (ARIMA) is a simple statistical model designed to predict future behavior of a time series solely based on its past values, without incorporating additional variables. In this model, each time step is presumed to have a linear dependence on the preceding p values of the time series, along with q error terms representing white noise. To address non-stationary time series, the data undergoes differencing d times. While these assumptions limit the model's ability to capture complex nonlinear relationships, ARIMA is recognized for its robustness, efficiency, and ease of training. Consequently, it remains widely applied in various domains such as financial time series⁴⁷ and epidemiology⁶⁵.

Convolutional Neural Networks (CNN). a widely recognized category of Artificial Neural Network (ANN) particularly renowned for image recognition tasks. 1-dimensional CNNs have been successfully applied in time series prediction tasks even though they are not primarily designed for sequential data. These networks can incorporate multiple hidden layers and are able to learn temporal dependencies between the target data, its past values and secondary data, which they achieve via an advanced backpropagation-based training scheme. Our model is based on the keras implementation of CNN and consists of a variable number of 1D Convolution layers with ReLU activation functions each followed by a 1D Max Pooling layer. The number of filters, the kernel size and the pool size are among the free hyper-parameters of the architecture which are selected through the grid search. The readout is comprised of two Dense layers, the first employing ReLU activation functions and a variable number of units followed by a linear layer that maps to the output. In contrast to the iterative scheme utilized by RC, the model takes a variable number of past values of the employed features as input, and directly outputs a vector of predicted values. During training, the model parameters are trained by minimizing the Mean Squared Error (MSE) using the Adam optimizer for 200 epochs with a variable learning rate. Early Stopping was employed, using 20% of the training data for validation, to halt the training once the validation error fails to improve for five consecutive epochs.

Long Short-Term Memory (LSTM). known as the most widely adopted variant of Recurrent Neural Network, it is part of the larger category of ANNs where it stands out by its sophisticated mechanism to deal with long-range temporal dependencies through selectively retaining past information with the help of its forget gate⁶⁶. The parameters of this model are trained via backpropagation through time (BPTT). While this is the most sophisticated algorithm we tested in this study, it is also the hardest to train and to fine-tune and carries the highest risk of overfitting. Similarly to our approach with CNNs we built our LSTM model using keras. The architecture consists of two LSTM layers connected by a Repeat Vector Layer. The activation functions used are of the type ReLU while the number of units in both layers is a free hyperparameter subject to tuning in the grid search. The same is true for the dropout rate which we apply to both layers on the input as well as the recurrent connections. The readout of the model consists of a Time Distributed Dense layer mapping to a vector of predicted values. We only apply the model to differenced data as opposed to making this a part of the hyperparameter tuning as in the RC and CNN model due to the fact that we saw a complete inability to learn the original data without differencing with the LSTM in preliminary trials. The training of the LSTM model mirrors that of the CNN model using an Adam optimizer and Early Stopping.

Training procedure

To ensure a fair comparison of all models and determine their optimal configurations, we conducted an extensive grid search on a high-

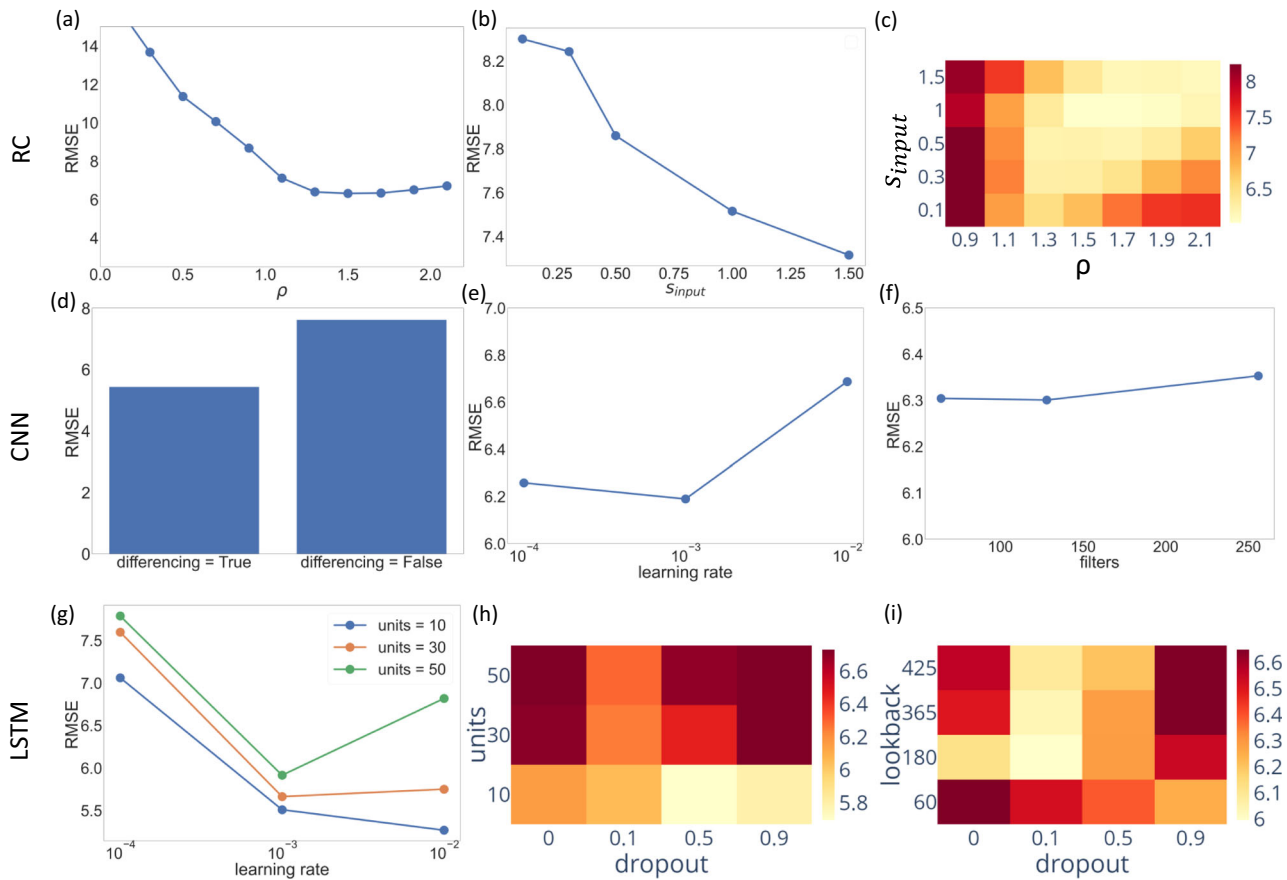


Fig. 7 | Grid search results. Additional results on the influence of different hyperparameters derived from the grid search on Yemen. The value shown is always the median RMSE over all splits and unspecified hyperparameters. **a–c** RC results for spectral radius and input strength, which are often considered the most important

hyperparameters in RC. **d–f** CNN results on differencing, learning rate and number of filters. **g–i** LSTM results on learning rate, number of units, dropout rate and length of lookback.

performance cluster for each model and country. In order to simulate real-life scenarios, we implemented a walk-forward validation scheme across the 12 splits under consideration covering a full year. As a preprocessing step all input data was resampled to a daily resolution and to eliminate isolated missing values linear interpolation was used. A trailing 10 day moving average was applied to smooth the data and reduce the amount of noise. For every prediction, we record the RMSE and the utilized hyperparameters. The result is a rich dataset that enables an in-depth analysis of our models' behavior and the significance of various hyperparameters.

Results directly derived from the grid search are shown in Fig. 7. The importance of the spectral radius and the input strength in RC is confirmed in agreement with the literature. For CNN we can see that differencing improves the performance, while other hyperparameters like e.g. the number of filters only have a small impact. For LSTM we identify the learning rate, the number of units, the dropout rate and the length of lookback as having a strong impact. We can see that it is generally impossible to define a single hyperparameter as good or bad since they are usually interdependent. Nevertheless, the fact that high dropout rates are favored in some cases in LSTM indicates a tendency towards overfitting. Overall, it is observed that RC stands out as the model for which it is easiest to identify generally reasonable hyperparameter configurations that yield satisfactory performance across all splits. This implies that satisfactory forecasts could even be made with less frequent retuning of the model's hyperparameters.

Within this framework we implement a walk-forward optimization procedure. This technique is commonly used in time series analysis and forecasting to assess the performance of a predictive model. The idea is to train and test the model sequentially over time, moving forward in a step-by-step manner. This process allows the model to be evaluated on new data that

becomes available as time progresses. The process can be separated into the following steps, which are carried out for each model on each country:

- Selecting a split:** A point in the time series is chosen as the cutoff between training and evaluation. The data before this point is considered to be the training data. The data after this point is considered test data.
- Training the model:** The model is trained on the training data.
- Evaluating the model:** A forecast is created with the trained model on the time period after the cutoff and compared to the test data where the RMSE is calculated between forecast and test data.
- Walk forward:** A new split with a new cutoff point further ahead in time is selected and the process is repeated.

This strategy can be used to evaluate the model with the best hyperparameters in the following way: the process is carried out for each combination of hyperparameters considered. To evaluate the performance and simulate an real-life application with regular retraining of the model hyperparameters for each split are chosen based on their aggregate performance across all preceding splits, where the prediction period ended before the beginning of the current one. This prevents any mixing of training and test set. The metric used to evaluate this aggregate performance is the median RMSE over these splits. The choice of the median RMSE over the mean RMSE was motivated by the fact that we found usage of the latter to be more likely to lead to mean-reverting, flat predictions. Conversely, using only a smaller amount of previous splits to select the hyperparameters proved to be too unstable given the heterogenous and noisy nature of the data where the best performance in a single split can often be an outlier. With the optimal knowable configuration for a given split we can properly

Table 2 | RC grid search parameters

| Parameter | Range |
|--------------|--|
| ρ | 0.3, 0.5, 0.7, 0.9, 1.1, 1.3, 1.5, 1.7, 1.9, 2.1 |
| β | 10^{-5} , 10^{-3} , 10^{-1} , 10, 100 |
| S_{input} | 0.1, 0.3, 0.5, 1.0, 1.5 |
| features | FCS, FCS+, economics, climate, all |
| differencing | True, False |

ranges of RC hyperparameters to be searched during tuning procedure.

Table 3 | LSTM grid search parameters

| Parameter | Range |
|---------------|------------------------------------|
| learning rate | 0.0001, 0.001, 0.01 |
| lookback | 60, 180, 365, 425 |
| units | 10, 30, 50 |
| dropout rate | 0., 0.1, 0.5, 0.9 |
| features | FCS, FCS+, economics, climate, all |

ranges of LSTM hyperparameters to be searched during tuning procedure.

Table 4 | CNN grid search parameters

| Parameter | Range |
|---------------|------------------------------------|
| learning rate | 0.0001, 0.001, 0.01 |
| lookback | 60, 180, 365, 425 |
| kernel size | 2, 5, 10, 15 |
| kernel size | 64, 128, 256 |
| pool size | 2, 4, 6 |
| layers | 1, 2 |
| features | FCS, FCS+, economics, climate, all |
| differencing | True, False |

ranges of CNN hyperparameters to be searched during tuning procedure.

Table 5 | Arima grid search parameters

| Parameter | Range |
|-----------|---------------|
| d | 0, 1 |
| p | 1, 2, 3, 4 |
| q | 1, 3, 5, 7, 9 |

ranges of ARIMA hyperparameters to be searched during hyperparameter tuning.

evaluate by looking at the performance of the respective version of the model on the respective time period. Using the differences or returns of the prevalence of insufficient food consumption as target instead of the raw values proved to be particularly important for the LSTM and CNN models which were otherwise mostly unable to successfully forecast. For this reason we added this as an option for RC as well while for ARIMA differencing is a already a part of the regular setup.

Beyond the configuration of the algorithms themselves we also investigated a noteworthy modification of the input data. Because the importance of each feature as well as their information content and noise-levels can vary drastically from country to country and even between splits, we implemented a method to select the most useful features for each

prediction in a way that could be used in application. In order to achieve this we introduced a hyperparameter for different groups of features in the grid search which makes it possible to optimize it together with the model configuration. The five groups we define are called “FCS”, “FCS+”, “climate”, “economics” and all ranging from including only the target data in “FCS” to the full set of features in “all”. Exactly which features are part of the different groups is visualized in Fig. 5a. In Fig. 5b–d we present the frequency with which each of the groupings is selected by the three algorithms RC, LSTM and CNN. The smaller groups of “FCS” and “FCS+” are selected with the highest frequency indicating that the autoregressive component is the most generally predictive one. However, the larger groups especially “climate” and “economics” do make their appearance showing that they are indeed informative at specific instances. This method was not applied for ARIMA, which does not use secondary data.

For all models a large preliminary grid search was carried out on the example of Yemen to identify the most relevant hyperparameters. Based on the results the ranges and fixed parameters were selected as shown in Tables 2, 3, 4 and 5. While the hyperparameters and their ranges generally differ for each model, our selection focused on those with the highest impact on performances while keeping the overall resource requirements of the grid search similar for RC, CNN and LSTM.

Data availability

The data used in this study is available in public repository <https://github.com/ducciopiovani/FamPredAI>.

Code availability

The code used in this study is available in public repository <https://github.com/ducciopiovani/FamPredAI>.

Received: 8 February 2024; Accepted: 13 September 2024; Published online: 21 October 2024

References

- Jagtap, S. et al. The russia-ukraine conflict: Its implications for the global food supply chains. *Foods* **11**, 2098 (2022).
- Ben Hassen, T. & El Bilali, H. Impacts of the russia-ukraine war on global food security: Towards more sustainable and resilient food systems? *Foods* **11**, 2301 (2022).
- Alexander, P. et al. High energy and fertilizer prices are more damaging than food export curtailment from ukraine and russia for food prices, health and the environment. *Nat. Food* **4**, 1284–95 (2022).
- Picchioni, F., Goulao, L. F. & Roberfroid, D. The impact of covid-19 on diet quality, food security and nutrition in low and middle income countries: a systematic review of the evidence. *Clin. Nutrition* **41**, 2955–2964 (2022).
- WFP. *War in Ukraine Drives Global Food Crisis*. <https://www.wfp.org/publications/war-ukraine-drives-global-food-crisis> (2022).
- WFP. *Global Report on Food Crises - 2022*. <https://www.wfp.org/publications/global-report-food-crises-2022> (2022).
- Lim, B. & Zohren, S. Time-series forecasting with deep learning: a survey. *Philos. Trans. R Soc. A* **379**, 20200209 (2021).
- Polyvianna, Y., Chumachenko, D. & Chumachenko, T. Computer aided system of time series analysis methods for forecasting the epidemics outbreaks. In *2019 IEEE 15th International Conference on the Experience of Designing and Application of CAD Systems (CADSM)*, 1–4 (IEEE, 2019).
- Kraemer, M. U. et al. Past and future spread of the arbovirus vectors aedes aegypti and aedes albopictus. *Nat. Microbiol.* **4**, 854–863 (2019).
- Tizzoni, M. et al. Real-time numerical forecast of global epidemic spreading: case study of 2009 a/h1n1pdm. *BMC Med.* **10**, 1–31 (2012).
- Scarpino, S. V. & Petri, G. On the predictability of infectious disease outbreaks. *Nat. Commun.* **10**, 1–8 (2019).

12. Ardabili, S. F. et al. Covid-19 outbreak prediction with machine learning. *Algorithms* **13**, 249 (2020).
13. Prakash, K. B., Imambi, S. S., Ismail, M., Kumar, T. P. & Pawan, Y. Analysis, prediction and evaluation of covid-19 datasets using machine learning algorithms. *Int. J.* **8**, 2199–2204 (2020).
14. Timmermann, A. Forecasting methods in finance. *Ann. Rev. Financial Econ.* **10**, 449–479 (2018).
15. Elliott, G. & Timmermann, A. Forecasting in economics and finance. *Ann. Rev. Econ.* **8**, 81–110 (2016).
16. Sezer, O. B., Gudelek, M. U. & Ozbayoglu, A. M. Financial time series forecasting with deep learning: a systematic literature review: 2005–2019. *Appl. Soft Comput.* **90**, 106181 (2020).
17. Sirignano, J. & Cont, R. Universal features of price formation in financial markets: perspectives from deep learning. *Quant. Finance* **19**, 1449–1459 (2019).
18. Deb, C., Zhang, F., Yang, J., Lee, S. E. & Shah, K. W. A review on time series forecasting techniques for building energy consumption. *Renew. Sustain. Energy Rev.* **74**, 902–924 (2017).
19. Schultz, M. G. et al. Can deep learning beat numerical weather prediction? *Philos. Trans. R Soc. A* **379**, 20200097 (2021).
20. Mudelsee, M. Trend analysis of climate time series: a review of methods. *Earth Sci. Rev.* **190**, 310–322 (2019).
21. Organization, W. H. et al. *The State of Food Security and Nutrition in the World 2020: Transforming Food Systems for Affordable Healthy Diets*, Vol. 322 (Food & Agriculture Org., 2020).
22. Wanner, N., Cafiero, C., Troubat, N. & Conforti, P. Refinements to the fao methodology for estimating the prevalence of undernourishment indicator. *Documento. De. Trabajo*. **2014**, 14–05 (2014).
23. Mwebaze, E., Okori, W. & Quinn, J. A. Causal structure learning for famine prediction. In *2010 AAAI Spring Symposium Series* (2010).
24. Okori, W. & Obua, J. Machine learning classification technique for famine prediction. In *Proc. World Congress on Engineering*. **2**, 4–9 (2011).
25. Andree, B. P. J., Chamorro, A., Kraay, A., Spencer, P. & Wang, D. Predicting food crises. *The World Bank*. <https://hdl.handle.net/10986/34510> (2020).
26. Wang, D., Andrée, B. P. J., Chamorro, A. F. & Spencer, P. G. Transitions into and out of food insecurity: a probabilistic approach with panel data evidence from 15 countries. *World Dev.* **159**, 106035 (2022).
27. Wang, D., Andree, B. P. J., Chamorro, A. F. & Girouard Spencer, P. Stochastic modeling of food insecurity. *The World Bank*. <https://hdl.handle.net/10986/34510> (2020).
28. Krishnamurthy R, P. K., Fisher, J. B., Choularton, R. J. & Kareiva, P. M. Anticipating drought-related food security changes. *Nat. Sustain.* **5**, 956–964 (2022).
29. Lentz, E. C., Michelson, H., Baylis, K. & Zhou, Y. A data-driven approach improves food insecurity crisis prediction. *World Dev.* **122**, 399–409 (2019).
30. Westerveld, J. J. et al. Forecasting transitions in the state of food security with machine learning using transferable features. *Sci. Environ.* **786**, 147366 (2021).
31. Martini, G. et al. Machine learning can guide food security efforts when primary data are not available. *Nat. Food* **3**, 716–728 (2022).
32. Foini, P., Tizzoni, M., Martini, G., Paolotti, D. & Omodei, E. On the forecastability of food insecurity. *Sci. Rep.* **13**, 2793 (2023).
33. WFP. *Hungermap Live*. <https://hungermap.wfp.org> (2022).
34. WFP. *The World Food Programme's Real-Time Monitoring Systems: Approaches and Methodologies*. <https://docs.wfp.org/api/documents/WFP-0000135070/download/> (2021).
35. McBride, L. et al. Predicting poverty and malnutrition for targeting, mapping, monitoring, and early warning. *Appl. Econ. Persp. Policy* **44**, 879–892 (2022).
36. Baylis, K., Heckeley, T. & Storm, H. Machine learning in agricultural economics. In *Handbook of Agricultural Economics*, **5**, 4551–4612 (Elsevier, 2021).
37. Zhou, Y., Lentz, E., Michelson, H., Kim, C. & Baylis, K. Machine learning for food security: principles for transparency and usability. *Appl. Econ. Persp. Policy* **44**, 893–910 (2022).
38. Maass, W., Natschlaeger, T. & Markram, H. Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Comput.* **14**, 2531–2560 (2002).
39. Jaeger, H. The “echo state” approach to analysing and training recurrent neural networks—with an erratum note. *Ger. Inform. Technol. Rep.* **148**, 13 (2001).
40. Lukoševičius, M. & Jaeger, H. Reservoir computing approaches to recurrent neural network training. *Comput. Sci. Rev.* **3**, 127–149 (2009).
41. Chattopadhyay, A., Hassanzadeh, P. & Subramanian, D. Data-driven predictions of a multiscale lorenz 96 chaotic system using machine-learning methods: reservoir computing, artificial neural network, and long short-term memory network. *Nonlinear Proc. Geophys.* **27**, 373–389 (2020).
42. Vlachas, P. R. et al. Backpropagation algorithms and reservoir computing in recurrent neural networks for the forecasting of complex spatiotemporal dynamics. *Neural Netw.* **126**, 191–217 (2019).
43. Shahi, S., Fenton, F. H. & Cherry, E. M. Prediction of chaotic time series using recurrent neural networks and reservoir computing techniques: a comparative study. *Mach. Learn. Appl.* **8**, 100300 (2022).
44. Domingo, L., Grande, M., Borondo, F. & Borondo, J. Anticipating food price crises by reservoir computing. *Chaos. Solitons. Fractals* **174**, 113854 (2023).
45. Han, J., Lin, H. & Qin, Z. Prediction and comparison of in-vehicle co2 concentration based on arima and lstm models. *Appl. Sci.* **13**, 10858 (2023).
46. Menculini, L. et al. Comparing prophet and deep learning to arima in forecasting wholesale food prices. *Forecasting* **3**, 644–662 (2021).
47. Kobiela, D., Krefka, D., Król, W. & Weichbroth, P. Arima vs lstm on nasdaq stock exchange data. *Procedia Comput. Sci.* **207**, 3836–3845 (2022).
48. Frausto-Solís, J., Hernández-González, L. J., González-Barbosa, J. J., Sánchez-Hernández, J. P. & Román-Rangel, E. Convolutional neural network–component transformation (cnn–ct) for confirmed covid-19 cases. *Math. Comput. Appl.* **26**, 29 (2021).
49. Baur, S. & Rāth, C. Predicting high-dimensional heterogeneous time series employing generalized local states. *Phys. Rev. Res.* **3**, 023215 (2021).
50. Balashankar, A., Subramanian, L. & Fraiberger, S. P. Predicting food crises using news streams. *Sci. Adv.* **9**, eabm3449 (2023).
51. Gauthier, D. J., Bollt, E., Griffith, A. & Barbosa, W. A. Next generation reservoir computing. *Nat. Commun.* **12**, 1–8 (2021).
52. WFP. *Climate Explorer*. <https://dataviz.vam.wfp.org/climate-explorer> (2024).
53. (CHC), C. H. C. *Chirps: Rainfall Estimates From Rain Gage and Satellite Observations*. <https://www.chc.ucsb.edu/data/chirps> (2022).
54. NASA. *Moderate Resolution Imaging Spectroradiometer (Modis)*. <https://modis.gsfc.nasa.gov/> (2022).
55. ACLED. *The Armed Conflict Location & Event Data Project*. <https://acleddata.com> (2022).
56. Raleigh, C., Linke, A., Hegre, H. & Karlsen, J. Introducing acled: an armed conflict location and event dataset: special data feature. *J. Peace Res.* **47**, 651–660 (2010).
57. Araujo, C., Araujo-Bonjean, C. & Brunelin, S. Alert at maradi: preventing food crises by using price signals. *World Dev.* **40**, 1882–1894 (2012).
58. WFP. *Calculation and Use of the Alert for Price Spikes (alps) Indicator, Technical Guidance Note, Analysis and Nutrition Service*. (2014).

- https://documents.wfp.org/stellent/groups/public/documents/manual_guide_proced/wfp264186.pdf.
59. WFP. *Economic Explorer*. <https://dataviz.vam.wfp.org/economic/prices> (2022).
 60. Griffith, A., Pomerance, A. & Gauthier, D. J. Forecasting chaotic systems with very low connectivity reservoir computers. *J. Nonlinear Sci.* **29**, 123108 (2019).
 61. Lu, Z., Hunt, B. R. & Ott, E. Attractor reconstruction by machine learning. *J. Nonlinear Sci.* **28**, 061104 (2018).
 62. Lu, Z. et al. Reservoir observers: Model-free inference of unmeasured variables in chaotic systems. *J. Nonlinear Sci.* **27**, 041102 (2017).
 63. Herteux, J. & R  th, C. Breaking symmetries of the reservoir equations in echo state networks. *J. Nonlinear Sci.* **30**, 123142 (2020).
 64. Hoerl, A. E. & Kennard, R. W. Ridge regression: biased estimation for nonorthogonal problems. *Technometrics* **12**, 55–67 (1970).
 65. Benvenuto, D., Giovanetti, M., Vassallo, L., Angeletti, S. & Ciccozzi, M. Application of the arima model on the covid-2019 epidemic dataset. *Data Brief* **29**, 105340 (2020).
 66. Hochreiter, S. & Schmidhuber, J. Long short-term memory. *Neural Comput.* **9**, 1735–1780 (1997).

Acknowledgements

We would like to thank Arif Husain for the continuous support and guidance. We are also grateful to Elisa Omodei and Pietro Foini for sharing their data and their knowledge on the XGBoost approach. We also want to express our gratitude to Will Fein from Microsoft AI for Good for his suggestions that improved our training procedure. This work was supported by the World Food Programme Innovation Accelerator.

Author contributions

C.R. conceived and supervised the research, D.P. and J.H. designed the study. J.H., I.L. and D.P. conducted the calculations. J.H., C.R, G.M, A.B, K.K, I.L and D.P. interpreted and evaluated the findings. J.H., C.R, G.M, A.B, K.K, I.L and D.P. wrote and reviewed the manuscript.

Funding

Open Access funding enabled and organized by Projekt DEAL.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s43247-024-01698-9>.

Correspondence and requests for materials should be addressed to Christoph Raeth or Duccio Piovani.

Peer review information *Communications Earth & Environment* thanks Hugo Deleglise and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Primary Handling Editors: Martina Grecequet. A peer review file is available.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

  The Author(s) 2024