



A computationally lightweight model for ensemble forecasting of environmental hazards: General TAMSAT-ALERT v1.2.1

Emily Black^{1,2}, John Ellis³, and Ross I. Maidment^{1,2,4}

¹Department of Meteorology, University of Reading, Reading, RG6 6BB, UK

²National Centre for Atmospheric Science, Leeds, LS2 9PH, UK

³Department of Computer Science, University of Warwick, Coventry, CV4 7AL, UK

⁴National Centre for Earth Observation, Leicester, LE4 5SP, UK

Correspondence: Emily Black (e.c.l.black@reading.ac.uk)

Received: 10 April 2024 – Discussion started: 6 May 2024

Revised: 8 August 2024 – Accepted: 19 September 2024 – Published: 25 November 2024

Abstract. Efficient methods for predicting weather-related hazards are crucial for the effective management of environmental risk. Many environmental hazards depend on the evolution of meteorological conditions over protracted periods, requiring assessments that account for evolving conditions. The TAMSAT-ALERT approach addresses this challenge by combining observational monitoring with a weighted multi-year ensemble. In this way, it enhances the utility of existing systems by enabling users to combine multiple streams of monitoring and meteorological forecasting data into holistic hazard assessments. TAMSAT-ALERT forecasts are now used in a number of regions in the Global South for soil moisture forecasting, drought early warning and agricultural decision support. The model presented here, General TAMSAT-ALERT, represents a significant scientific and functional advance on previous implementations. Notably, General TAMSAT-ALERT is applicable to any variable for which time series data are available. In addition, functionality has been introduced to account for climatological non-stationarity (for example due to climate change), large-scale modes of variability (for example El Niño) and persistence (for example of land-surface conditions). In this paper, we present a full description of the model, along with case studies of its application to the prediction of central England temperature, Pakistan vegetation conditions and African precipitation.

1 Introduction

Transparent, robust and computationally efficient methods for hazard assessment are of great value to stakeholders dealing with environmental risk (for example, Boulton et al., 2022). Weather-related hazards may depend on the evolution of multiple meteorological variables over a protracted period. For example, crop yield is affected by precipitation and temperature throughout the growing season. In-season updates therefore require monitoring of past conditions as well as forecasting. Combining observations and forecasts is, however, challenging – especially when the hazard in question is affected by more than one variable. Extending the aforementioned example, when making an in-season assessment of the risk of low crop yield, it is necessary to consider both meteorological conditions in the period since planting and the probability of adverse conditions during the remaining season. One approach would be to drive a crop model, such as AquaCrop (Steduto et al., 2009), with observations of each driving variable up until the present and with an ensemble of numerical weather prediction (NWP) model forecasts for the future. There are two difficulties with this. On a practical level, it may be problematic to obtain forecasts of all driving variables at the required time and spatial resolution. A further challenging problem is the drift in predictions that occurs as NWP models move from their initial state into equilibrium with their own physics (for example, Manzanar, 2020). If model predictions were to be spliced directly onto historical observations, the drift would cause systematic bias in yield assessments, the magnitude of which would depend on the stage of the growing season at which the meteorological

logical forecasts were initiated. This is a challenging error to correct using standard bias correction methods because the magnitude of the bias depends on the lead time of the forecast. The TAMSAT-ALERT approach addresses these issues by splicing together historical time series for the past with an ensemble comprised of time series extracted from individual historical years (hereafter referred to as a “climatology”) for the future (Asfaw et al., 2018; Boulton et al., 2020; Brown et al., 2017). The use of consistent historical data avoids the issue of model drift and accounts for complex interactions between variables, such as precipitation, temperature and net radiation (Asfaw et al., 2018). From a practical perspective, seamless integration of past and future conditions facilitates in-season updates for slowly developing hazards, such as drought (Brown et al., 2017). This feature of TAMSAT-ALERT provides a simple means of calculating ensemble statistics that take into account both historical time series and future predictions. Furthermore, the progressive incorporation of observational monitoring gradually reduces uncertainty in risk assessments as the season evolves – facilitating decision-making.

In its default state, TAMSAT-ALERT treats all members of the climatological ensemble as equally likely. When calculating ensemble statistics, such as the ensemble mean and standard deviation, it is, however, in principle possible to weight ensemble members individually if there is evidence that some are more likely to occur than others. For example, climatological ensemble members may be weighted more strongly if the El Niño phase during their associated year is close to that at the initiation of the forecast. Extending this idea, the ensemble can be weighted using meteorological forecasts, based on the similarity of the predicted weather to the weather during the year associated with each ensemble member. A strength of the methodology is that NWP output can be incorporated, even when forecasts are not available for the variable being assessed. In Kenya, for example, incorporation of skilful precipitation tercile forecast probabilities output by the European Centre for Medium-Range Weather Forecasts (ECMWF) dynamical forecasting system improves the skill of the normalised difference vegetation index (NDVI) and yield forecasts during the secondary rainy season (Young et al., 2020; Boulton et al., 2020). The use of a weighted climatological ensemble thus enables historical and forecast data to be combined into a holistic view of risk (see Appendix A and B for a more comprehensive description of the weighting methodology).

Previous work on the TAMSAT-ALERT method has described how the system can be used for agricultural and drought forecasting (Asfaw et al., 2018; Boulton et al., 2020; Black et al., 2024a) and for short-term decision support (Black et al., 2024b). Although originally designed for application in Africa, there have been applications in Guyana (David, 2023) and Pakistan (Black et al., 2024c). A similar approach is used for the FEWS NET drought outlooks (Shukla et al., 2014; Turner et al., 2022) and for precipitation

predictions within the seasonal performance probability tool (Novella and Thiaw, 2016).

In this paper, we present General TAMSAT-ALERT – a versatile implementation of the TAMSAT-ALERT framework. Unlike the previous system, which required the use of a land-surface or crop model, General TAMSAT-ALERT can be applied to any variable for which time series data are available. General TAMSAT-ALERT is thus a self-contained model rather than a modelling framework. During the development of General TAMSAT-ALERT, we took the opportunity to extend the methodology and completely revise the way that users interact with the code. A key innovation is the option to increment variables from the forecast initialisation date, enabling forecasts to account for persistence of environmental conditions in time. In addition, as well as enabling users to weight the ensemble with time series of climate indices, General TAMSAT-ALERT permits weighting by proximity of the climatological year to the target year. The latter development avoids the assumption of climatological stationarity – a weakness of the original methodology.

The paper is structured as follows. Section 2 summarises the design and implementation of General TAMSAT-ALERT and describes the novel developments in TAMSAT-ALERT. In Sect. 3, the usage of the system is illustrated through three case studies: probabilistic prediction of central England temperature statistics, NDVI forecasting in Pakistan and continually updated predictions of the standardised precipitation index (SPI) for Africa. The paper closes with some reflections on how General TAMSAT-ALERT fits into the ever-expanding ecosystem of environmental forecast models.

2 Model description

2.1 Conceptual approach and implementation

The inputs and outputs of the system are illustrated by Fig. 1, and the procedure is described more fully in the following example, which considers a metric calculated over a period of interest with the forecast initiated on a date within the period of interest.

The preparatory steps are as follows:

1. Specify the period of interest (POI) start and end dates, specify the initiation date of the forecast, and supply the time series data for the forecasts.
2. Determine the dominant periodicity at the time resolution of the data (see Appendix C). For example, for a time series with an annual dominant periodicity, the periodicity would be 365 for daily data or 12 for monthly data. To simplify the description, for the rest of this section, we will assume that the dominant periodicity is annual.

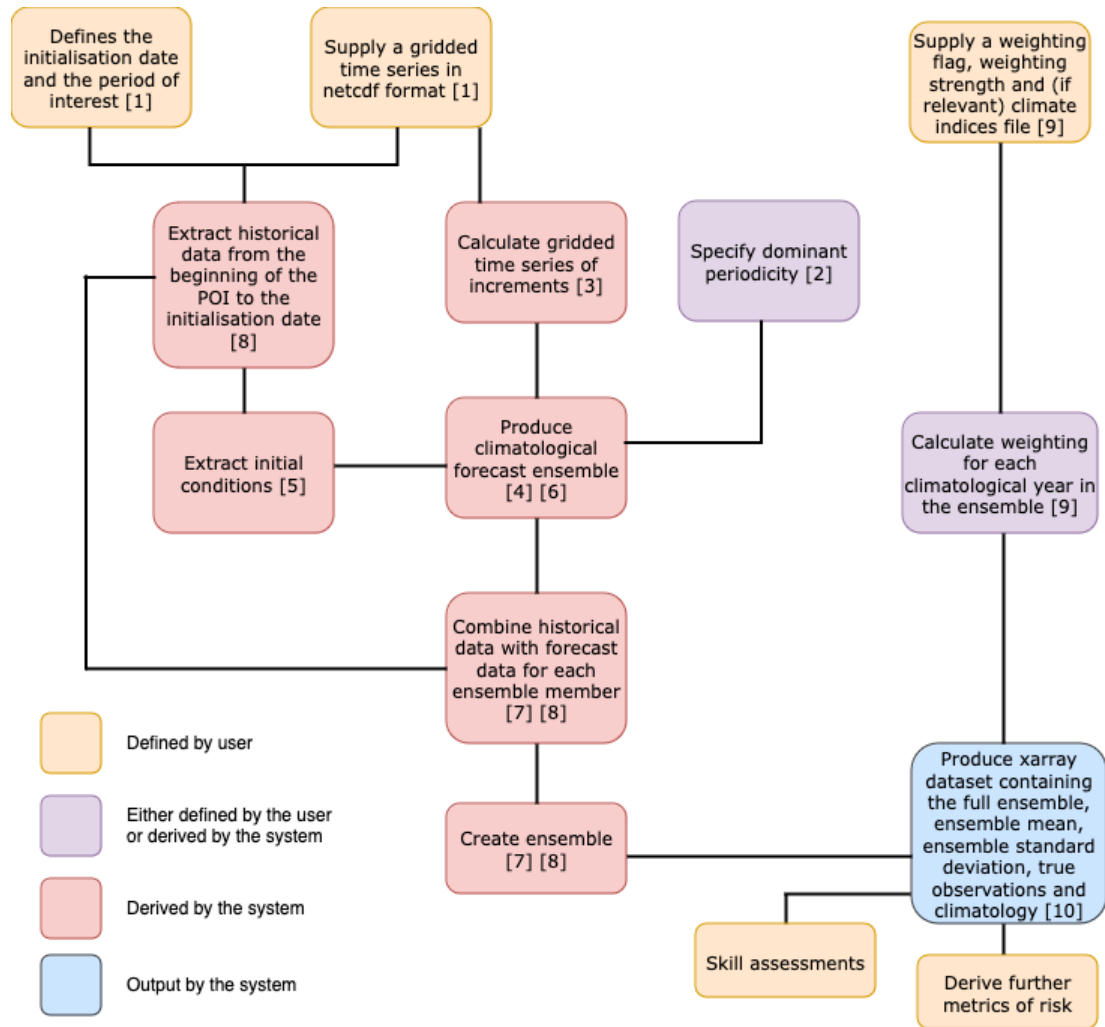


Figure 1. Flow chart showing how the General TAMSAT-ALERT implementation works. The numbers in square brackets refer to the modelling steps described in Sect. 2.1.

3. Optionally, transform the time series into a time series of incremental change (i.e. the difference between each value and the one that came before it in the time series).

4. Divide the time series into individual years.

The forecast is performed as follows:

5. Extract the data value on the initiation date of the forecast.

6. Extract the future forecast period from every member of the climatological ensemble. The first day of the forecast period is the initiation date, and the last day is the end of the period of interest.

7. If the option is taken to increment from the initiation date, the data derived in step 5 are added to the state on the initiation date of the forecast; otherwise the ensemble consists of the raw data derived in step 5.

The metric over the whole period of interest is calculated as follows:

8. The period of interest spans a period starting prior to the initiation date as well as a period afterwards. For each ensemble member, the forecast ensemble member (derived in step 6) is spliced together with the observations from the start of the period of interest to the initiation date.

The ensemble statistics are calculated as follows:

9. Any weighting of the ensemble is applied at this stage, with weights allocated to each climatological ensemble member based on the conditions experienced during the year from which the ensemble member is derived (see Appendix A and B for further details of the methodology). The code includes three options for weighting.

- a. *No weighting applied (weighting flag set to 0)*. In this case, equal weights are assigned to all ensemble members.
- b. *Weighting by proximity of the climatological ensemble member year to the actual year (weighting flag set to 1)*. This option allows the user to account for long-term variability and trends in the variable being forecast. The strength of the weighting is assigned by the user.
- c. *Weighting using an index of climate variability, such as El Niño–Southern Oscillation (ENSO)*. This option allows the user to favour climatological ensemble members in which climate conditions are similar to those at the initialisation of the forecast. Both the data file used and the strength of the weighting are specified by the user.

It is recommended that users experiment with different weightings and strengths, ideally carrying out skill assessments for a range of different set-ups (see case study 1 for an example of this).

10. Any ensemble forecast statistic can be calculated from the ensemble derived in step 7, using any weighting options derived in step 9. These include the ensemble mean, the ensemble spread, confidence intervals and the probability that the metric being forecast will be less than a user-prescribed threshold. The inbuilt returns from General TAMSAT-ALERT are the weighted ensemble mean and weighted ensemble standard deviation.

2.2 Novel methodological developments

The original TAMSAT-ALERT implementation was designed as a framework for producing ensemble crop yield forecasts (Asfaw et al., 2018). Following the publication of the original TAMSAT-ALERT method, there was interest from the agricultural and humanitarian sectors in its wider application to drought forecasting and agricultural decision support (Boult et al., 2020; Black et al., 2024b). These subsequent applications did not require significant development of the underlying methodology for the framework, and so there was no need to develop a new model. The wider use of the TAMSAT-ALERT methodology, however, highlighted an opportunity to develop the methodological framework described in these papers into a new more general model. General TAMSAT-ALERT thus builds on the success of TAMSAT-ALERT by enabling users to derive forecasts directly from observations and reanalysis without the need for the use of land-surface/crop models or NWP forecasts. This has required several methodological and scientific extensions and a complete redesign and rewrite of the underlying code. For this release of General TAMSAT-ALERT, we have also taken the opportunity to rethink the way that users interact

with the system, with the aim of enabling a wider range of users to use the system to produce quantitative forecasts.

Scientific developments are as follows:

- In the previous implementation of TAMSAT-ALERT, there was an assumption of climatological stationarity, which is clearly over-simplistic when analysing variables with significant trends (such as surface temperature). General TAMSAT-ALERT includes options to account for trends in variables by incrementing from the last day of the historical time series and/or by weighting the climatology to favour years closer to the forecast/hindcast initiation date. These methodologies are demonstrated in case study 1 – prediction of central England temperature extremes.
- The previous implementation of TAMSAT-ALERT did not explicitly account for persistence, instead relying on land-surface models to represent persistence in soil moisture. General TAMSAT-ALERT directly accounts for temporal persistence (and implicitly long-term trends) by allowing the user to select an option for incrementing forecasts from the last day of historical observations. This functionality is illustrated by case study 2 – NDVI forecasting for Pakistan – and by case study 1 – prediction of central England temperature extremes.
- Predictability of many environmental variables is amplified by large-scale modes of variability, such as ENSO. In General TAMSAT-ALERT, a method is implemented for incorporating climate indices into ensemble predictions. This is illustrated by case study 3 – probabilistic prediction of SPI3 for Africa.

Methodological developments are as follows:

- General functions for weighting ensembles with climatological indices have been developed and incorporated into the code (allowing the user to specify the strength of the weighting). Further details are included in Appendix B.
- An extension has been included to identify the dominant periodicity in input data, and the code has been generalised to deal with data with any periodicity and time resolution. Further details are included in Appendix C.

Code developments are as follows:

- General TAMSAT-ALERT is released as a fully documented and publicly issued Python package (`general_tamsat_alert`).
- The whole procedure described in Fig. 1 and Sect. 2.1 is encompassed by a single function [`do_forecast()`], which ingests netCDF data and produces ensemble forecasts for a user-specified initiation date and output period. The function includes all

of the functionality for weighting and incrementing described above.

3 Model application

The application of General TAMSAT-ALERT is demonstrated through three case studies that are designed to demonstrate the functionality of the system beyond conventional analyses of meteorological time series data. In particular, the case studies demonstrate how General TAMSAT-ALERT improves skill by incorporating non-stationarity, persistence of environmental variables in time and wider modes of variability such as ENSO.

- *Case study 1: prediction of central England temperature (CET) extremes.* There is a strong anthropogenic positive trend in CET, superposed on strong decadal and interannual variability. Application of General TAMSAT-ALERT to the prediction of CET demonstrates how the system handles non-stationarity in a historical climate.
- *Case study 2: prediction of NDVI in Pakistan.* NDVI in semi-arid regions is highly persistent, and so this case study demonstrates how General TAMSAT-ALERT exploits persistence to make skilful predictions with lead times of up to 2 months.
- *Case study 3: probabilistic prediction of the 3-month standardised precipitation index (SPI3) for Africa.* SPI3 on a particular day depends on cumulative precipitation in the preceding 3 months. This case study demonstrates how the system combines observations and predictions into probabilistic forecasts. September–November precipitation in eastern Africa is, furthermore, strongly influenced by El Niño, and the case study demonstrates how weighting with the Oceanic Niño Index (ONI) improves forecast skill.

In order to assess the skill of the forecasts objectively, two standard skill scores were utilised in the case studies:

- r^2 (square of Pearson's moment correlation coefficient) is a measure of how closely the hindcast interannual time series of NDVI captures the observed time series. Specifically, r^2 gives the proportion of variance in the time series of observations is explained by the hindcasts. An r^2 of 1 denotes a perfect linear correlation, for which all of the observed variability is represented by the forecast; for a 40-year time series, $r^2 < 0.07$ ($r < 0.26$) denotes that the correlation is insignificant at the 95 % level.
- ROC–AUC (receiver operating characteristics–area under the curve) is a metric of how reliably events can be predicted. In this context, an event is the predictand below a user-defined threshold (for example, below

the 20th percentile). An ROC–AUC of 0.5 or less indicates no skill in distinguishing false alarms from true positives (hits). An ROC–AUC of 1 indicates a perfect forecast. For a concise guide to ROC–AUC, see <https://www.metoffice.gov.uk/research/climate/seasonal-to-decadal/gpc-outlooks/user-guide/interpret-roc> (last access: 17 November 2024).

3.1 Case study 1: prediction of central England temperature (CET) seasonal extremes

3.1.1 Case study 1: introduction

In this case study, the TAMSAT-ALERT system provides probabilistic predictions of whether CET in a given season will exceed the 90th, 95th or 99th percentiles. CET exhibits a pronounced trend, which is reflected in increased occurrence of 90th-, 95th- and 99th-percentile events in the later part of the time series (Fig. 2). The TAMSAT-ALERT system can account for non-stationarity in two ways. Firstly, for future periods, the climatological ensemble can be incremented from the conditions on the initiation date of the forecast. Any trend in the variable being predicted on the initiation date will be implicitly represented in the ensemble. Secondly, the facility to weight the climatological ensemble allows the system to favour years that are close in time to the period of interest (see Appendix A and B). The CET case study thus illustrates how accounting for non-stationarity improves prediction skill for variables with a strong trend.

3.1.2 Case study 1: data and methodology

The dataset used in this study is HadCET (available from <https://www.metoffice.gov.uk/hadobs/hadcet/>, last access: 17 November 2024). HadCET is a time series of temperature diagnostics for a region of the UK, roughly encompassed by a triangle enclosed by Lancashire, London and Bristol. Monthly CET data are provided from 1659, with daily maximum and minimum temperatures available from 1878 (see Parker et al., 1992, for a full description). This case study utilises the monthly-mean daily maximum temperatures for 1882–2021, focusing on July. It can be seen from the time series displayed in Fig. 2 that there is pronounced decadal variability throughout the time series, with a clear warming trend evident from ~ 1980 onwards.

Hindcasts were generated for the whole data period, and each hindcast uses data from all years to generate the climatological ensemble (i.e. from the past as well as from the present). The metrics derived from the system were probabilistic forecasts that the 90th, 95th and 99th percentiles of CET would be exceeded. The percentiles were calculated using data from the whole time series based on the ensemble mean and standard deviation (i.e. assuming Gaussian behaviour). As was described above, General TAMSAT-ALERT can optionally be run with each ensemble member




Figure 2. Time series of monthly-mean daily maximum July central England temperature (CET), highlighting the 90th, 95th and 99th percentiles.

incremented from the initiation dates of the forecasts and/or weighted according to some measure of how similar the climatological year is to the year in question. The intention of the incrementing process is to generate ensembles of increments of the historical time series, relative to the previous time step. Therefore, let us say for a given climatological year that July had an average temperature of 18°C and August had an average temperature of 17.5°C . In this case the incremental value for August would be -0.5°C . The climatological ensemble member time series thus comprises a time series of increments, which are then applied from the initialisation date. Practically speaking, the increment time series are generated by subtracting the initial state for the climatological year and then adding the result to the current year.

In the case study, the system was run for four set-ups. All set-ups are run for a 1-month lead forecast (i.e. initialised by the June monthly mean). The set-ups are

- no incrementing or weighting (i.e. the ensemble of values for historical data);
- ensemble members weighted by the proximity of the ensemble climatological year to the hindcast year but with no incrementing, where the exponential weighting factor is arbitrarily set to 1 (see Appendix B);
- no weighting but ensemble members incremented from the temperature at the time of initialisation;
- ensemble members weighted by the proximity of the ensemble climatological year to the hindcast year and ensemble members incremented from the temperature on the day of initialisation.

3.1.3 Case study 1: results

Figure 3 shows forecasts for each of the set-ups described above. It can be seen that monthly-mean daily maximum temperature for the example month chosen (July 2021) fell just below the 90th percentile. When no weighting or incrementing is included, the probabilities of exceeding the 90th, 95th and 99th percentiles are close to the climatological expectation of 0.1, 0.05 and 0.01, respectively. When the weighting and/or incrementing is applied, the probabilities of exceeding the threshold increase markedly. This is, in part, because both weighting and incrementing implicitly account for the strong positive observed trend.

A more formal analysis is shown in Fig. 4, which displays a time series of hindcast probabilities of exceeding the 90th, 95th and 99th percentiles. When neither incrementing nor weighting is applied (Fig. 4a), the probabilities are close to climatological expectation and do not change from year to year. Weighting the distribution by proximity of the climatological ensemble year to the observed year (Fig. 4b) has the expected effect of reducing the probabilities in the early part of the time series and increasing them in the later part – reflecting the trend. The incrementing (Fig. 4c) has a similar effect, with the probabilities of exceeding each threshold being consistently elevated in the later part of the time series. Interestingly, when the incrementing is included, the probabilities are consistently higher than the climatological expectation – suggesting that there is some degree of non-Gaussian behaviour that is picked up when the forecasts are initialised from observations. In this case, the elevated probabilities suggest that the climatological percentiles (as estimated by




Figure 3. Example forecasts with a 1-month lead time for July 2021. (a) No weighting or incrementing of the time series; (b) weighting of the climatology by year proximity but no incrementing of the time series; (c) incrementing of the time series from the initiation but no weighting; (d) weighting by year proximity and incrementing of the time series. In all of the plots, observations are blue; the 90th, 95th and 99th percentiles are pale yellow, pale orange and pale pink; and individual ensemble members are grey. The vertical green line is the date of forecast initiation, and the red line is the target date. In (b) and (d), the weighting is denoted by the darkness of the lines.

assuming Gaussian behaviour) are too cold. When both the incrementing and the weighting are applied (Fig. 4d), elements of both types of behaviour are evident, with strong increases in hindcast probabilities in the later part of the time series and greater-than-expected climatological probabilities throughout. Figure 4 displays the results of a formal skill assessment (ROC–AUC). When neither incrementing nor weighting is applied, the ROC–AUC is close to the climatological expectation of 0.5. When weighting and/or incrementing is applied, the skill improves, with ROC–AUC in the region of 0.7–0.8. For the most extreme cases (95th and 99th percentiles), slightly better scores are achieved when both the incrementing and weighting are applied, but the differences are small.

3.2 Case study 2: prediction of NDVI for Pakistan

The implementation of General TAMSAT-ALERT for NDVI forecasting is illustrated through a case study of Pakistan and the surrounding region. Pakistan was chosen for the case study for several reasons. Firstly, the climate and topography of Pakistan vary considerably (see Fig. 5), enabling us to test the NDVI forecasting method in a range of environmental settings. The topographic zones include the highlands of the north, large river plains in Punjab and Sindh, and the Balochistan Plateau. The climate also varies considerably. For example, precipitation ranges from $< 100 \text{ mm yr}^{-1}$ in the deserts of Balochistan to $> 1000 \text{ mm yr}^{-1}$ in northern

regions affected by the southwest monsoon. Secondly, vegetation in Pakistan is partly rainfed and partly fed by river overflow and glacial melt. This means that there is a disconnect between variability in precipitation and vegetation conditions. Direct observation of vegetation may thus be the most appropriate method of monitoring drought and crop conditions, and NDVI forecasts thus have practical value for the region. The final reason for choosing Pakistan was pragmatic. In 2020, the Start Network commissioned the TAMSAT group to develop a new drought monitoring service for Punjab and Sindh, which, unlike pre-existing services, targets the secondary winter growing season. The development of the drought monitoring system necessitated extending the TAMSAT-ALERT method from soil moisture in Africa to NDVI in Pakistan (Black et al., 2024c).

3.2.1 Case study 2: data and methodology

The dataset used in this study is the Blended Vegetation Health Product (Blended-VHP). Blended-VHP is a multiple-product dataset issued by NOAA's Center for Satellite Applications and Research (STAR) that produces global vegetation health products. The Blended-VHP products include data from different satellite sensors (VIIRS for 2013–present and AVHRR for 1981–2012) that have been reprocessed into a single homogenous time series (Yang et al., 2021). The vegetation-health-related variables provided include NDVI, brightness temperature, the vegetation con-

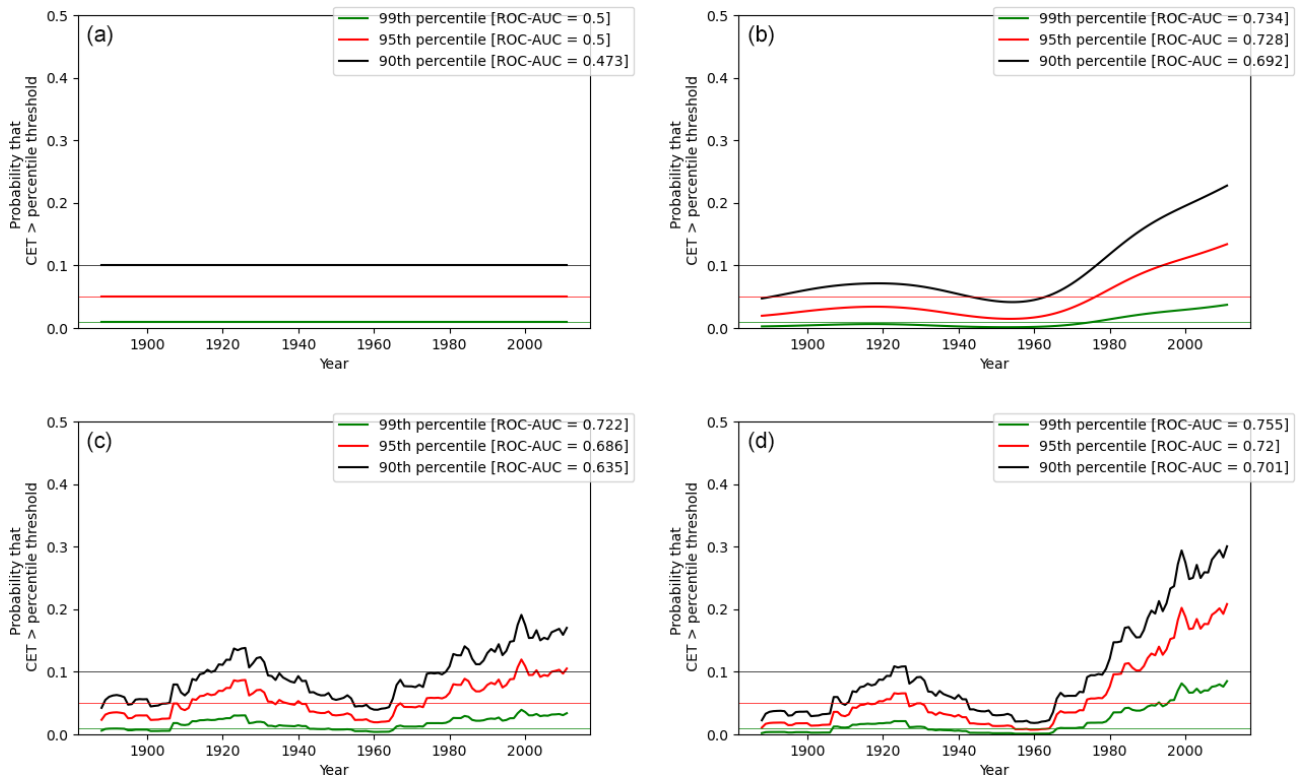


Figure 4. Time series of predicted probabilities that the July central England temperature (CET) will exceed the 1850–2020 90th, 95th and 99th percentiles for a 1-month lead time – i.e. forecasts initiated in June. (a) No weighting or incrementing of the time series; (b) weighting of the climatology by year proximity but no incrementing of the time series; (c) incrementing of the time series but no weighting; (d) both weighting and incrementing applied. The bold lines show modelled probabilities of exceedance of the 90th-, 95th- and 99th-percentile thresholds, and the fine lines show the climatological probabilities.

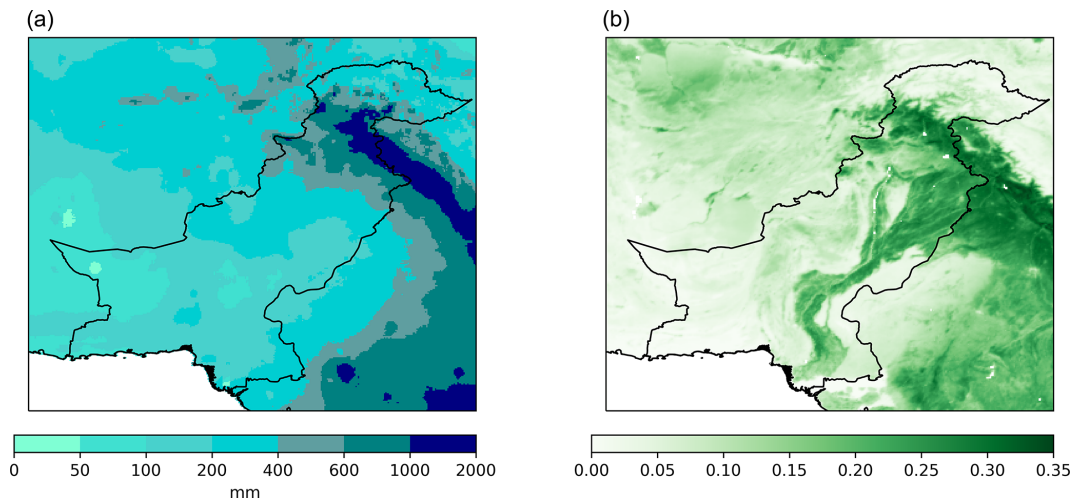


Figure 5. Maps of precipitation and NDVI annual means. (a) Annual mean precipitation in the case study region; (b) annual mean maximum 15 d NDVI in the case study region. The country outline shows the border of Pakistan.

dition index, the thermal condition index and the vegetation health index. The raw Blended-VHP data are released as weekly files at 4 km resolution. For this study, the data were re-gridded to 0.05° and interpolated to twice-monthly time steps (15th of the month, last day of the month). The forecasts were produced using incrementing but no weighting, with the TAMSAT-ALERT forecasts generated independently for each grid point.

3.2.2 Case study 2: results

An example of an NDVI forecast for 1 August 2018, with a lead time of up to 60 d, is shown in Fig. 6, and Fig. 7 shows r^2 , together with ROC–AUC, for forecasts of NDVI that are below the 20th percentile on 1 August. The skill scores are calculated for each point for the gridded forecasts. Even 2 months in advance, there is some skill compared to climatology, and by 1 July (1-month lead time), the forecasts are highly skilful over most of the region. Although there is some spatial variation, in all regions there is skill at a 60 d lead time, with excellent skill at a 30 d lead time (ROC–AUC > 0.85; $r^2 > 0.75$). Since no weighting is applied and the forecasts are for an NDVI snapshot rather than a seasonal mean, the only source of skill in the forecasts is the time persistence of NDVI. As was discussed for CET, a potential limitation of the TAMSAT-ALERT method is the implicit assumption that we can predict the future using historical climatologies. The difficulty of dealing with data with a strong trend is partly addressed by deriving ensembles using time step changes in NDVI rather than using the absolute values of the historical observations. It should be noted, however, that although starting the forecasts from the NDVI on the date of initiation accounts for the trend in NDVI on the date of initiation, it does not account for trends in the magnitude of the NDVI time increments.

3.3 Case study 3: 3-month standardised precipitation index (SPI3) prediction for Africa

Lower-than-usual precipitation has severe consequences in Africa because of the region's dependence on rainfed agriculture. Precipitation in Africa varies strongly in both space and time, with most regions having pronounced dry and wet seasons (see Fig. 8). A commonly used metric of the precipitation anomaly is the standardised precipitation index (SPI) (McKee et al., 1993). For a given time, SPI is essentially the normalised and standardised cumulative precipitation anomaly derived for a user-defined number of preceding months. The choice of the number of months depends on how the SPI data will be utilised. For assessment of agricultural drought, a 3-month period is typically used; for hydrological drought, longer periods (6 or 9 months) may be more appropriate.

In this application, the TAMSAT-ALERT system is used to predict SPI3 for Africa with lead times of up to 3 months.

As was described above, SPI3 is based on cumulative precipitation during the 3 months prior to the target date. At a 3-month lead time, therefore, the TAMSAT-ALERT forecast is based entirely on an ensemble of future precipitation. As the season progresses, the ensemble progressively incorporates observations. For example, at a 1-month lead time, each ensemble member will include 2 months of observations and 1 month of forecasts. The primary aim of this case study is thus to demonstrate how TAMSAT-ALERT can be used to combine historical observations and climatological information into probabilistic predictions.

3.3.1 Case study 3: data and methodology

The dataset chosen was the GPCP monthly gridded precipitation (a gridded dataset based on rain gauge observations). For this study, the 1° resolution “Full Data Monthly Product” was used (Schneider et al., 2016). The weighting was carried out using the Oceanic Niño Index (ONI) provided by NOAA (<https://psl.noaa.gov/data/climateindices/list/>, last access: 17 November 2024). SPI was calculated for each ensemble member precipitation time series using the procedure outlined in Keyantash et al. (2023). The ROC–AUC scores are for the detection of SPI < -0.75 (mild to moderate drought). The probability that SPI breaches the drought threshold is derived from the ensemble mean and standard deviation, assuming Gaussian behaviour. This choice is justified because SPI is constructed in such a way as to make it Gaussian, and an empirical approach tends to produce noisy results, especially for extremes. For the weighted ensembles, the weightings were derived from the ONI time series, with weights based on the difference between the ONI at the forecast initiation and the ONI for the climatological ensemble.

3.3.2 Case study 3: results

Figure 9 shows example SPI forecasts initiated at the beginning of August, September and October 2015 for November SPI3. In 2015, there was a large El Niño event, resulting in high precipitation in eastern Africa and low precipitation in southern Africa. It can be seen that in August, when the forecast is based entirely on the climatological ensemble, as would be expected, the unweighted ensemble predicts the climatology – i.e. zero anomalies. As additional data are included, the unweighted forecasts approach the observations, with most observed features evident by October. For this case (a large El Niño event), comparison between the weighted and unweighted forecasts shows that the ONI weighting has a pronounced effect (Fig. 9). Even in August, slight positive anomalies are predicted in eastern Africa and slight negative anomalies are predicted in southern Africa, in line with the expected El Niño teleconnection (Reason, 2017; Black et al., 2003). As the season progresses, the anomalies approach observations, with significantly positive anomalies evident in southern and eastern Africa at a 2-month lead time.

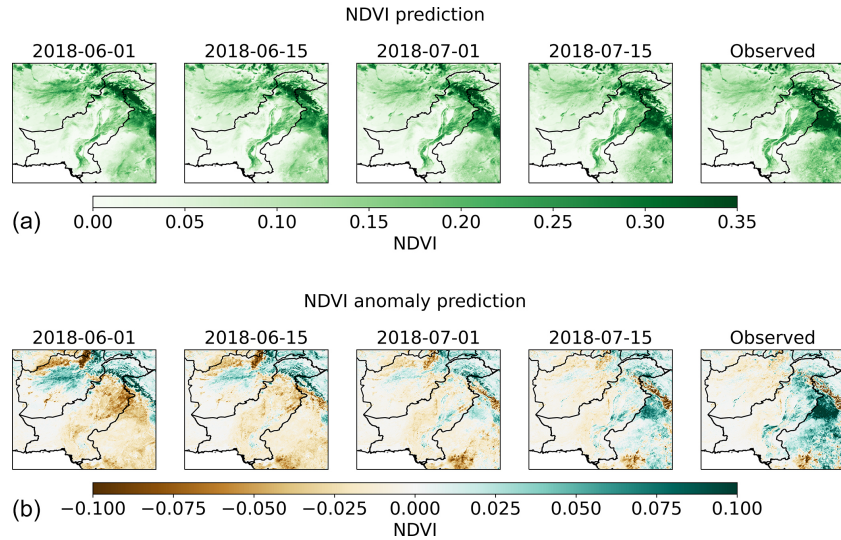


Figure 6. Predicted and observed NDVI in an example month (August 2018). The dates in the sub-panel titles are the initialisation dates for the forecasts (date format: year-month-day). **(a)** Predicted and observed NDVI; **(b)** predicted and observed NDVI anomaly.




Figure 7. Skill scores for NDVI predicted for 1 August. The dates in the sub-panel titles are the initialisation dates for the forecasts (date format: month-day). **(a)** ROC–AUC for detecting a 20th-percentile event; **(b)** r^2 (i.e. the square of the Pearson correlation coefficient).

The improvement in skill as the lead time reduces, suggested by the 2015 example, is consistent with the formal skill assessment shown in Fig. 10, which shows the ROC–AUC for detection of $SPI < -0.75$ (mild/moderate drought). For the unweighted November forecasts, at a 3-month lead time, the ROC–AUC is ~ 0.5 for all of Africa, indicating that the skill is no better than the climatology. As observations are incorporated, the skill increases until it is > 0.7 in most regions at a lead time of 1 month. The weighting has the greatest effect at the longer lead times, with the greatest impact on eastern Africa. This is consistent with the well-known El Niño teleconnection with the eastern Africa October–December rainy season (Black et al., 2003).

4 Discussion and reflections

There are many excellent forecasting systems that contribute to the early warning of climate-related hazards. Such systems are based on cutting-edge numerical weather prediction models and on observations. These range from full dynamical systems run by meteorological agencies, such as the Met Office, ECMWF and the Bureau of Meteorology, to simple statistical regression models (e.g. Diro et al., 2008; Gissila et al., 2004). In recent years, there has been a proliferation of machine-learning-based algorithms capable of emulating numerical forecasting systems or of inferring future conditions using historical data (see Ren et al., 2021, for a review). In parallel with forecast improvements, there have




Figure 8. Seasonal cycle in precipitation over Africa: precipitation accumulations for DJF (a), MAM (b), JJA (c) and SON (d).

been significant developments in observation, and now that we are 40 years into the meteorological satellite era, traditional ground-based observational networks are routinely supplemented by satellite estimates. Long-standing satellite datasets include the CHIRPS and TAMSAT precipitation data (Funk et al., 2015; Maidment et al., 2017), both of which provide data back to the early 1980s. Although generally less accurate than satellite and ground-based observations (for example, Lavers et al., 2022), meteorological reanalyses offer long and consistent time series of a wide range of variables – some of which are not easily observed by more direct methods. Widely used examples include the NCEP and the ERA5 reanalyses (Kalnay et al., 2018; Hersbach et al., 2020).

So – taking these examples into account – what is the point of yet another ensemble forecasting system? From a technical perspective, General TAMSAT-ALERT is complementary to existing systems, which by and large focus either on observations or on forecasts. It is challenging to provide even qualitative assessments that take into account both the past and the future. For example, one can envisage a situation in which regional precipitation is forecast to be high following a poor start to the rainy season. Does this mean that a seasonal meteorological drought is likely? Or is the forecast high precipitation sufficient to outweigh the current dry conditions? General TAMSAT-ALERT provides a straightforward way to combine observations and forecasts into quantitative fore-

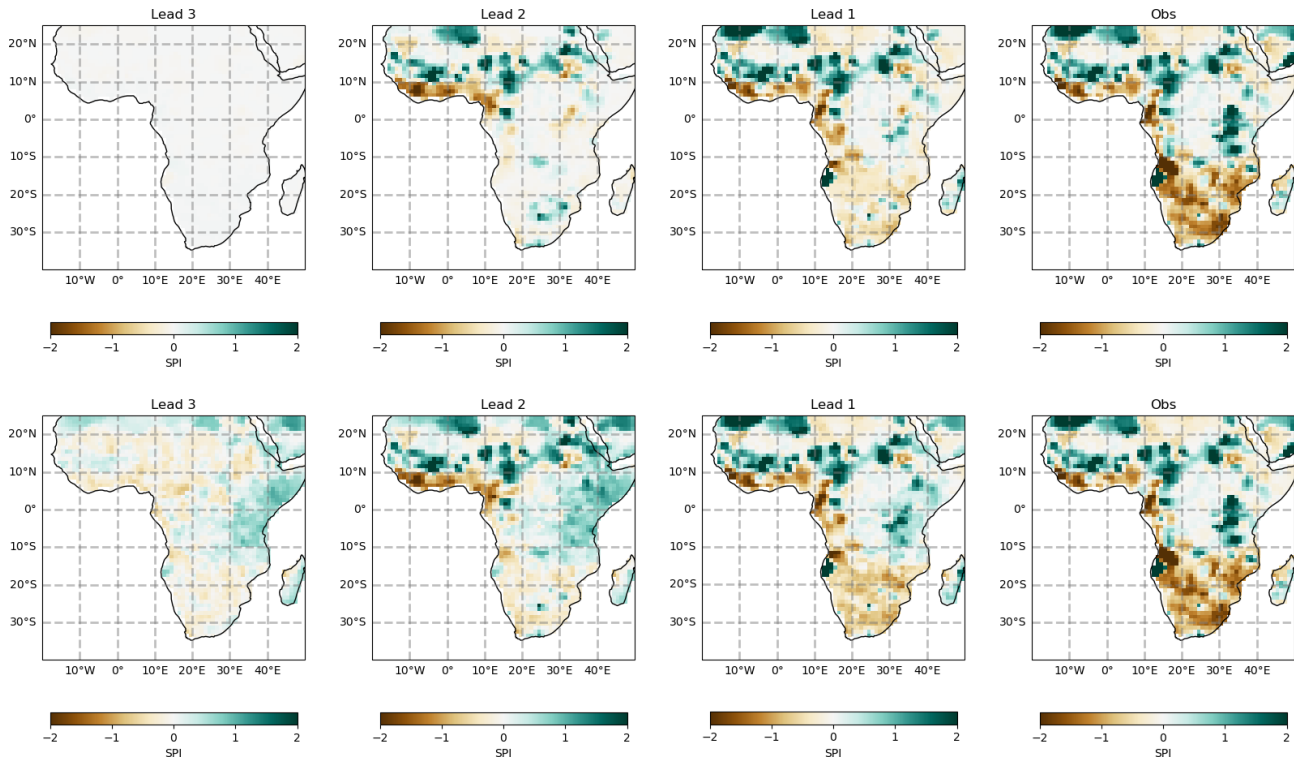


Figure 9. Example forecast for 3-month SPI3 on 15 November 2015. From left to right: 3-month lead time (forecast initiated on 15 August), 2-month lead time (forecast initiated on 15 September), 1-month lead time (forecast initiated on 15 October) and observed SPI3. The top row displays forecasts with no weighting applied to the ensemble. The bottom row displays forecasts with the ensemble weighted using the Oceanic Niño Index (ONI).

casts (Boult et al., 2020). Extending the example above, in some cases, meteorological forecasts may not be available, or they may have poor skill. In this situation, the method presented here enables decision-makers to judge the probability of seasonal drought based purely on observations of the season so far (as with case study 3).

An additional issue addressed by General TAMSAT-ALERT is that of climate non-stationarity. The assumption of stationarity is demonstrably false in a changing climate and has been shown to have a detrimental effect on the skill of empirical statistical predictions (Salvi et al., 2016; Adeyemi and Akerele, 2017). An advantage of the approach introduced in General TAMSAT-ALERT is that the weighting methodology accounts for non-stationarity without making assumptions about the structure of long-term trends or variability. This means that the forecasts account for both natural decadal variability and anthropogenic climate change.

A further challenge for stakeholders is that the major forecasting organisations issue predictions and observations for a limited number of variables. Even ECMWF, which provides both forecasts and an extensive and continually updated meteorological reanalysis, only considers meteorological and land-surface variables that can be output directly from NWP models. Metrics of vegetation condition, such as NDVI, are not included and could not be handled by the orig-

inal TAMSAT-ALERT modelling framework. Lack of forecasts for particular variables may make it difficult for stakeholders to accommodate forecasts into their operations – especially if their current monitoring protocols utilise variables, such as NDVI, that are not routinely predicted. The general approach to time series forecasting presented here provides a solution for such users.

The climatological approach central to the TAMSAT-ALERT methodology enables users to exploit publicly available meteorological forecasts and observations to improve skill (case studies 2 and 3 and Boult et al., 2020). Although incorporation of skilful meteorological forecasts into TAMSAT-ALERT is undoubtedly useful (Boult et al., 2020), weighting with observed data is also potentially of value to decision-makers. For example, El Niño can exacerbate the risks of extreme weather in many regions (for example Goddard and Gershunov, 2020; Kay et al., 2022), and existing systems and warnings may report possible sectorial impacts of El Niño (for example Nobre et al., 2019; Kim et al., 2021). However, in practice, such reports tend to be semi-quantitative and reliant on expert judgement. The option to weight climatological ensembles with observations of climate indices (weighting flag 2 in General TAMSAT-ALERT) supplements this information by providing quantitative assessments of how environmental variables are mod-

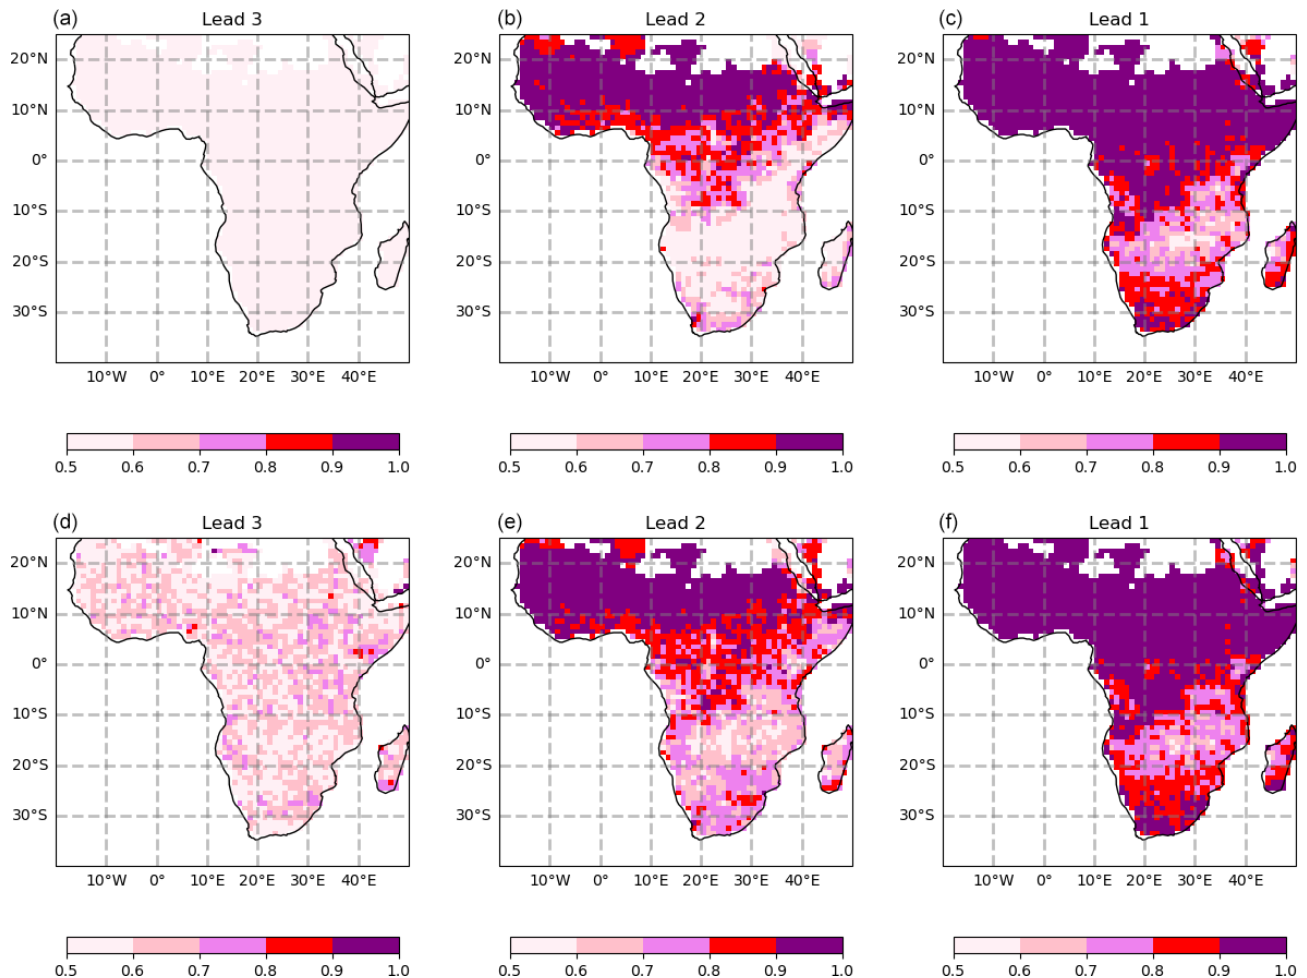


Figure 10. ROC–AUC scores for 3-month SPI3 on 15 November. From left to right: 3-month lead time (forecast initiated on 15 August; **a**, **d**), 2-month lead time (forecast initiated on 15 September; **b**, **e**) and 1-month lead time (forecast initiated on 15 October; **c**, **f**). (**a–c**) Forecasts with no weighting applied to the ensemble. (**d–f**) Forecasts with the ensemble weighted using the Oceanic Niño Index (ONI).

ulated by El Niño. The methodology implicitly accounts not only for varying teleconnection strength but also for the strength of the link between meteorological variables affected by El Niño and the metric being forecast. Figure 11 combines case study 2 (Pakistan NDVI) with case study 3 (African precipitation), comparing the effect of weighting with ONI during the strongest El Niño event in recent history. The results show that weighting with ONI data has a strong effect on African precipitation forecasts but that the predicted anomalies for Pakistan are uniformly weak – reflecting a weak link between El Niño and NDVI (although the predicted anomalies are stronger elsewhere in the region). The negative Pakistan result is of practical importance – especially given the anecdotal weight given to El Niño in the Pakistan meteorological and agricultural sectors (for example, Siyal et al., 2019). In the humanitarian sector there is often discussion on taking action when an El Niño is imminent or ongoing, and yet in some cases the connection between El Niño and impact-relevant metrics, such as NDVI, is weak.

General TAMSAT-ALERT allows users to make their own judgements about the relevance of El Niño for their particular application. Allowing users to incorporate data on modes of variability is a significant extension of the original framework.

The reliance of decision-makers on qualitative reports of El Niño highlights the limited uptake of forecasts for early action. This is in part because of the issue described above – namely that centrally issued meteorological forecasts may not target variables of interest. Further issues are that forecasts may not be available on the timetables specified in early action plans and that forecast metrics may not be presented in a way that facilitates decision-making. These issues can be addressed by some combination of collaboration, co-production and direct production of forecasts by decision-makers (for example, Dasgupta et al., 2023). Our computationally lightweight approach and the public release of General TAMSAT-ALERT through the standard Python package indexing system facilitate all three solutions.




Figure 11. Forecasts and observations of Pakistan (and surrounding region) seasonal mean NDVI anomaly (a–c) and Africa cumulative precipitation anomaly (d–f). Comparison between unweighted forecasts (a, d), weighted forecasts (b, e) and observations (c, f). Forecasts are initiated on 1 August 1997 and target the September–November 1997 season. The weighting is based on the Oceanic Niño Index.

Connected to the discussion above is the somewhat nebulous notion of “ownership” of the forecasting process. It is inevitable that forecasts are sometimes wrong. Loss of trust in forecasts is, however, not inevitable if decision-makers and end users have a detailed understanding from the outset of the underlying principles and skill and subsequent limitations (Hirons et al., 2021; Gudoshava et al., 2022). Using a transparent method to generate ensemble forecasts and derive relevant metrics is an excellent way of building such understanding. Furthermore, end users are more likely to persist with a system that they have built themselves – working with it over a number of years to design a decision-making process that accounts for error and uncertainty (Hirons et al., 2023).

5 Conclusions

In this paper, we have presented a computationally lightweight and general method for ensemble forecasting. The key scientific innovations of the method are that it accounts for non-stationarity in time series and that it exploits the predictability arising from persistence in some environ-

mental variables without the need to use an initialised numerical model. From a practical perspective, in comparison to numerical and machine learning forecasting methods, General TAMSAT-ALERT is easy to use and computationally lightweight.

General TAMSAT-ALERT is designed to ingest multiple streams of meteorological forecast and environmental observational data. Accordingly, the intention is not to replace existing forecasts and observations but rather to provide a platform for transforming such data into actionable assessments. Over the coming years, our aim is to demonstrate the use of the TAMSAT-ALERT methodology in a range of sectors and to build the capacity of decision-makers to use our code to produce bespoke hazard assessments.

Appendix A: A brief description of the TAMSAT-ALERT method

A1 Overview

TAMSAT-ALERT is a method of generating probabilistic forecasts of environmental variables, using observed time series and (optionally) meteorological forecasts or indices of large-scale modes of climatic variability. Accordingly, it enables users to combine multiple sources of data into probabilistic forecasts of environmental hazards. The way that TAMSAT-ALERT works is illustrated by Fig. A1.

It can be seen from Fig. A1 that TAMSAT-ALERT forecasts use historical time series to provide the following components.

- Monitoring information on the part of the period of interest prior to the initiation date is provided. In operational systems, forecasts are initiated as close to “today” as possible, and the observational time series is used to monitor the progress of the season/period of interest so far.
- Initial conditions for the forecasts are provided.
- A multi-year (climatological) ensemble is provided. The climatological ensemble can be thought of as multiple realisations of the future, with each possible future based on a series of events that have happened in the past.

Using observational data from multiple years from the past, spliced with observational data from the period of interest so far, TAMSAT-ALERT thus generates an ensemble of time series for the whole period of interest.

A2 Weighting the ensemble

In its default set-up, TAMSAT-ALERT treats each ensemble member as equally likely. If additional information is available about the probability of individual ensemble members, this information can be used to weight the ensemble. In practice, weighting the ensemble members is straightforward because each ensemble member is derived from a historical year. Thus, if we can infer how likely conditions during a historical year are to occur during the period of interest, we can use this information to weight the ensemble. In this paper and in previous studies, three methods of weighting the ensemble have been demonstrated.

- *Weighting method 1: by the proximity of the period of interest to the climatological ensemble year.*

Because many environmental variables show trends over time, ensemble members from years closer in time to the TAMSAT-ALERT forecast are more likely to accurately represent conditions in the future segment of the period of interest. In this weighting method, each

ensemble member is therefore weighted based on how close its associated year is to the period of interest.

- *Weighting method 2: by the similarity of climatic modes of variability during the period of interest to modes of variability during the climatological ensemble year.*

In many parts of the world, the climate is greatly affected by patterns of variability, such as El Niño–Southern Oscillation (ENSO) and the North Atlantic Oscillation. This means that weather conditions during a given period are more likely to be similar to those in past years when these patterns were in a similar phase rather than when they were in a different phase. For example, in eastern Africa, rainfall is influenced by ENSO. If the period of interest occurs during an El Niño, the rainfall will be more like the rainfall during past El Niño periods than during La Niña periods. Weighting method 2 takes this into account by giving more importance to ensemble members for which a user-provided index of a variability pattern for the relevant climatological year is similar to the index on the forecast’s start day. If the user-supplied index is poorly correlated with the variable being forecast, the effect of the weighting on the ensemble statistics will be minimal, while if it is strongly correlated, the ensemble statistics will be significantly perturbed.

- *Weighting method 3: using meteorological forecasts.*

Meteorological forecasts provide information on the likelihood that a specific weather variable, such as precipitation, will fall into one of three categories: below average (tercile 1), average (tercile 2) or above average (tercile 3).¹ This information can be used to weight ensemble members. Each ensemble member is assigned to a tercile based on how the weather conditions during the associated ensemble member year compare to the historical average. The weighting is then based on the forecasted probability for that tercile. For example, in eastern Africa, rainfall in November 1997 was above average (tercile 3). If we are using a November precipitation forecast for eastern Africa to weight the ensemble, the 1997 ensemble member will be assigned to tercile 3 and weighted according to the present-day probability of rainfall in that tercile. This approach allows precipitation forecasts to weight TAMSAT-ALERT predictions for various variables (like NDVI). Continuing the example, similarly to in weighting method 2, if the connection between rainfall and NDVI is weak, then using a precipitation forecast to weight the TAMSAT-ALERT prediction will have little impact on the ensemble statistics.

¹Other thresholds for categorising meteorological variables, such as quintiles and deciles, are equally applicable.




Figure A1. Schematic of how the TAMSAT-ALERT method can be used to translate historical time series into probabilistic predictions.

A3 Calculating the risk of an adverse event

The risk of an adverse event can be calculated similarly to other ensemble predictions. For instance, the weighted ensemble mean and standard deviation can determine the probability that NDVI or precipitation will breach a user-defined threshold. Another method involves using ensemble members to drive a model, such as a crop or land-surface model. In this approach, each ensemble member is used to run the model, the model outputs are combined into an ensemble and the risk assessment is based on the weighted ensemble statistics.

Appendix B: Additional information on the inbuilt method for weighting the climatological ensemble

Within General TAMSAT-ALERT, there are two inbuilt options for weighting the climatological ensemble. These are defined in the wrapper function by flags:

- *Weighting flag 0.* No weighting is applied.
- *Weighting flag 1.* Weight the ensemble based on the proximity of the climatological ensemble year to the year in which the forecast is initiated.

- *Weighting flag 2.* Weight the ensemble based on the similarity of some index at the point of forecast initiation to the index during the climatological year.

The method of weighting is similar for methods 1 and 2. For weighting flag 1, the ensemble is weighted as follows:

$$w_i = e^{-0.001(S\Delta I \frac{T}{24})^2},$$

where w_i is the weighting for ensemble member i , S is the weighting strength, ΔI is the difference between the initiation time index and the climatological ensemble time index, and T is the dominant period of the data.

For weighting flag 2, the ensemble is weighted as follows:

$$w_i = e^{-(S\Delta V)^2},$$

where w_i is the weighting for ensemble member i , S is the weighting strength, ΔV is the difference between the value of the climatic index at initiation and the climatic index for ensemble member i , and T is the dominant period of the data.

The weighting strength for individual ensemble members for both methods is illustrated by Fig. B1. The index used in the figure is the Oceanic Niño Index (ONI), and the forecast was initiated in November 1997 (a strong El Niño).

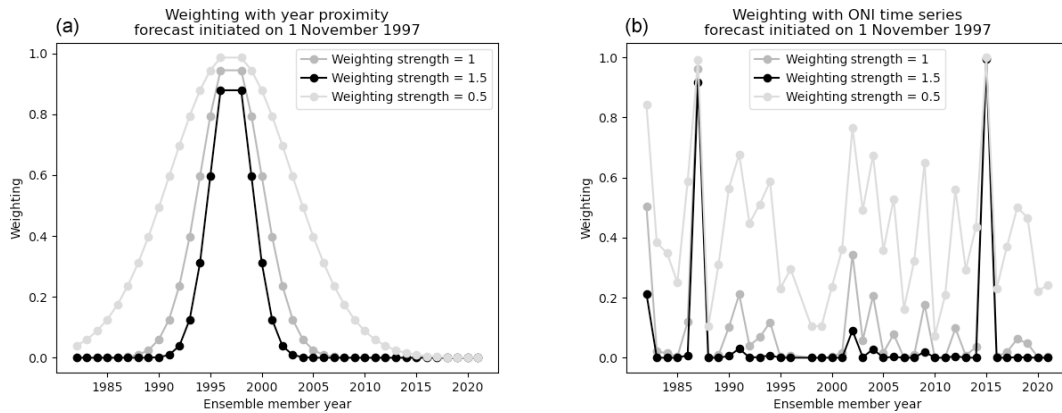


Figure B1. An example of weights assigned to climatological ensemble members for a forecast initiated on 1 November 1997. **(a)** Weighting by year proximity; **(b)** weighting by the Oceanic Niño Index. Note that the year of the forecast (in this case 1997) is left out of the climatological ensemble.

Appendix C: Additional information on the inbuilt method for determining the periodicity of the time series data

For meteorological time series data, the dominant periodicity is almost always 1 year. It is therefore recommended that periodicity is specified by the user – accounting for the time resolution of the data. For example, for the monthly precipitation and temperature data in case studies 1 and 3, the user-specified periodicity was 12, while for the twice-monthly NDVI data in case study 2, the user-specified periodicity was 24.

However, in order to make General TAMSAT-ALERT applicable to data without a known periodicity, a method for deriving periodicity is supplied within the Python package. The methodology is illustrated in Fig. C1 and summarised below:

1. The input data are first transformed into a cleaner state by subtracting the mean square offset error. This method works by quantifying how different the signal is from itself offset by a given amount, effectively correcting for small variations in phase and removing noise from the data.
2. The data are linearly detrended.
3. A Fourier transform is applied to the processed data.
4. The maximum peak in the resulting spectrum is identified as the dominant periodicity.

The algorithm described above is computationally intensive, so it is applied to a subset of the gridded input – selected by regular sampling of the grid at user-specified intervals. It should be noted that the dominant periodicity is assumed to be constant in space. If this is not the case, users are advised to use a subset of the data regionally before applying the TAMSAT-ALERT method.




Figure C1. An example of the dominant periodicity calculation method. **(a)** Example time series of processed and raw data (steps 1 and 2 above). **(b)** Fourier transform of the raw and processed data (step 3).

Code availability. All model code is open-source and publicly available at <https://pypi.org/project/general-tamsat-alert/> (Ellis and Black, 2024b) and https://github.com/brightlego/General-TAMSAT_ALERT (last access: 17 November 2024). The version of the model evaluated in this paper is persistently archived at the following DOI: <https://doi.org/10.5281/zenodo.10955490> (Ellis and Black, 2024a). A user guide to the code is included in the Supplement. The ROC–AUC score calculation code is available at <https://pypi.org/project/fastroc/> (Ellis and Black, 2024c) and <https://github.com/brightlego/fastroc> (last access: 17 November 2024).

Data availability. All datasets used in this study are publicly available, via the sources given in Sect. 2. Convenience copies of the netCDF format files used to produce the plots are available at https://gws-access.jasmin.ac.uk/public/tamsat/tamsat_alert/gmd_paper/datasets.zip (Maidment and Black, 2024).

Supplement. The supplement related to this article is available online at: <https://doi.org/10.5194/gmd-17-8353-2024-supplement>.

Author contributions. EB led the writing of the manuscript and conducted the case studies. JE wrote the General TAMSAT-ALERT and fastroc code, devised the weighting and periodicity analysis methods, and wrote Appendix B and C. RM processed the NDVI data and leads the operational applications of TAMSAT-ALERT. All authors commented on the manuscript draft.

Competing interests. The contact author has declared that none of the authors has any competing interests.

Disclaimer. Publisher’s note: Copernicus Publications remains neutral with regard to jurisdictional claims made in the text, published maps, institutional affiliations, or any other geographical representation in this paper. While Copernicus Publications makes every effort to include appropriate place names, the final responsibility lies with the authors. Regarding the maps used in this paper, please note that Figs. 5, 6, 7, and 11 contain disputed territories.

Financial support. This work was supported by the National Centre for Atmospheric Science through the NERC National Capability International Programmes Award (NC/X006263/1).

Review statement. This paper was edited by Jinkyu Hong and reviewed by Patrick Laux and two anonymous referees.

References

- Adeyemi, B. and Akerele, M. I.: The effect of seasonality and stationarity on specific humidity forecast over Nigeria, *J. Atmos. Sol.-Terr. Phy.*, 156, 57–71, <https://doi.org/10.1016/j.jastp.2017.03.001>, 2017.
- Asfaw, D., Black, E., Brown, M., Nicklin, K. J., Otu-Larbi, F., Pinnington, E., Challinor, A., Maidment, R., and Quaiye, T.: TAMSAT-ALERT v1: a new framework for agricultural decision support, *Geosci. Model Dev.*, 11, 2353–2371, <https://doi.org/10.5194/gmd-11-2353-2018>, 2018.
- Black, E., Slingo, J., and Sperber, K. R.: An observational study of the relationship between excessively strong short rains in coastal East Africa and Indian Ocean SST, *Mon. Weather Rev.*, 131, 74–94, [https://doi.org/10.1175/1520-0493\(2003\)131<0074:AOSOTR>2.0.CO;2](https://doi.org/10.1175/1520-0493(2003)131<0074:AOSOTR>2.0.CO;2), 2003.
- Black, E., Boulton, V. L., Hirons, L., and Woolnough, S. J.: The potential value of seasonal drought forecasts in the context of climate change: A case study of the African elephant conservation sector, *Meteorol. Appl.*, <https://doi.org/10.1002/met.2190>, online first, 2024a.

- Black, E., Asfaw, D. T., Sananka, A., Aston, S., Boulton, V. L., and Maidment, R. I.: Application of TAMSAT-ALERT soil moisture forecasts for planting date decision support in Africa, *Front. Clim.*, 4, 993511, <https://doi.org/10.3389/fclim.2022.993511>, 2024b.
- Black, E., Maidment, R. I., Rees, E., and Nderitu, E.: A new drought model for disaster risk management in the Punjab, Sindh and Baluchistan provinces of Pakistan, *Front. Climate*, 6, 1332233, <https://doi.org/10.3389/fclim.2024.1332233>, 2024c.
- Boulton, V. L., Asfaw, D. T., Young, M., Maidment, R., Mwangi, E., Ambani, M., Waruru, S., Otieno, G., Todd, M. C., and Black, E.: Evaluation and validation of TAMSAT-ALERT soil moisture and WRSI for use in drought anticipatory action, *Meteorol. Appl.*, 27, e1959, <https://doi.org/10.1002/met.1959>, 2020.
- Boulton, V. L., Black, E., Abdillahi, H. S., Bailey, M., Harris, C., Kilavi, M., Kniveton, D., MacLeod, D., Mwangi, E., Otieno, G., and Rees, E.: Towards drought impact-based forecasting in a multi-hazard context, *Clim. Risk Manage.*, 35, 100402, <https://doi.org/10.1016/j.crm.2022.100402>, 2022.
- Brown, M., Black, E., Asfaw, D., and Otu-Larbi, F.: Monitoring drought in Ghana using TAMSAT-ALERT: A new decision support system, *Weather*, 72, 201–205, <https://doi.org/10.1002/wea.3033>, 2017.
- Dasgupta, A., Arnal, L., Emerton, R., Harrigan, S., Matthews, G., Muhammad, A., O'Regan, K., Pérez-Ciria, T., Valdez, E., van Osnabrugge, B., and Werner, M.: Connecting hydrological modelling and forecasting from global to local scales: Perspectives from an international joint virtual workshop, *J. Flood Risk Manage.*, e12880, <https://doi.org/10.1111/jfr3.12880>, 2023.
- David, D. L.: Detecting agricultural drought risks: A case study of the rice crop (*Oryza sativa*) and the TAMSAT-ALERT system in Guyana, *Meteorol. Appl.*, 30, e2149, <https://doi.org/10.1002/met.2149>, 2023.
- Diro, G. T., Black, E., and Grimes, D. I. F.: Seasonal forecasting of Ethiopian spring rains, *Meteorol. Appl.*, 15, 73–83, <https://doi.org/10.1002/met.63>, 2008.
- Ellis, J. and Black, E.: General TAMSAT-ALERT v1.2.1, Zenodo [code], <https://doi.org/10.5281/zenodo.10955490>, 2024a.
- Ellis, J. and Black, E.: general-tamsat-alert 1.2.2, <https://pypi.org/project/general-tamsat-alert/>, last access: 17 November 2024b.
- Ellis, J. and Black, E.: fastroc, <https://pypi.org/project/fastroc/>, last access: 21 November 2024c.
- Funk, C., Peterson, P., Landsfeld, M., Pedreros, D., Verdin, J., Shukla, S., Husak, G., Rowland, J., Harrison, L., Hoell, A., and Michaelsen, J.: The climate hazards infrared precipitation with stations—a new environmental record for monitoring extremes, *Sci. Data*, 2, 1–21, <https://doi.org/10.1038/sdata.2015.66>, 2015.
- Gissila, T., Black, E., Grimes, D. I. F., and Slingo, J. M.: Seasonal forecasting of the Ethiopian summer rains, *Int. J. Climatol.*, 24, 1345–1358, <https://doi.org/10.1002/joc.1078>, 2004.
- Goddard, L. and Gershunov, A.: Impact of El Niño on weather and climate extremes, in: *El Niño Southern Oscillation in a Changing Climate*, edited by: McPhaden, M. J., Santoso, A., and Cai, W., American Geophysical Union (New Jersey, USA) and John Wiley & Sons, Inc. (Washington DC, USA), <https://doi.org/10.1002/9781119548164.ch16>, 361–375, 2020.
- Gudoshava, M., Wanzala, M., Thompson, E., Mwesigwa, J., Endris, H. S., Segele, Z., Hirons, L., Kipkogei, O., Mumbua, C., Njoka, W., and Baraibar, M.: Application of real-time S2S forecasts over Eastern Africa in the co-production of climate services, *Clim. Serv.*, 27, 100319, <https://doi.org/10.1016/j.cliser.2022.100319>, 2022.
- Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., Nicolas, J., Peubey, C., Radu, R., Schepers, D., and Simmons, A.: The ERA5 global reanalysis, *Q. J. Roy. Meteor. Soc.*, 146, 1999–2049, <https://doi.org/10.1002/qj.3803>, 2020.
- Hirons, L., Thompson, E., Dione, C., Indasi, V. S., Kilavi, M., Nkikaka, E., Talib, J., Visman, E., Adefisan, E. A., de Andrade, F., and Ashong, J.: Using co-production to improve the appropriate use of sub-seasonal forecasts in Africa, *Clim. Serv.*, 23, 100246, <https://doi.org/10.1016/j.cliser.2021.100246>, 2021.
- Hirons, L., Wainwright, C. M., Nying'uro, P., Quaye, D., Ashong, J., Kiptum, C., Opoku, N. K., Thompson, E. M., and Lamptey, B.: Experiences of co-producing sub-seasonal forecast products for agricultural application in Kenya and Ghana, *Weather*, 78, 148–153, <https://doi.org/10.1002/wea.4324>, 2023.
- Kalnay, E., Kanamitsu, M., Kistler, R., Collins, W., Deaven, D., Gandin, L., Iredell, M., Saha, S., White, G., Woollen, J., and Zhu, Y.: The NCEP/NCAR 40-year reanalysis project, in: *Renewable energy*, Routledge, London, <https://doi.org/10.4324/9781315793245>, 2018.
- Kay, G., Dunstone, N. J., Smith, D. M., Betts, R. A., Cunningham, C., and Scaife, A. A.: Assessing the chance of unprecedented dry conditions over North Brazil during El Niño events, *Environ. Res. Lett.*, 17, 064016, <https://doi.org/10.1088/1748-9326/ac6cd3>, 2022.
- Keyantash, J. and National Center for Atmospheric Research Staff (Eds.): *The Climate Data Guide: Standardized Precipitation Index (SPI)*, last modified 19 August 2023, NCAR, <https://climatedataguide.ucar.edu/climate-data/standardized-precipitation-index-spi>, last access: 6 March 2024.
- Kim, K., Chowdhury, R., Pant, P., Yamashita, E., and Ghimire, J.: Assessment of ENSO risks to support transportation resilience, *Prog. Disaster Sci.*, 12, 100196, <https://doi.org/10.1016/j.pdisas.2021.100196>, 2021.
- Lavers, D. A., Simmons, A., Vamborg, F., and Rodwell, M. J.: An evaluation of ERA5 precipitation for climate monitoring, *Q. J. Roy. Meteor. Soc.*, 148, 3152–3165, <https://doi.org/10.1002/qj.4363>, 2022.
- Maidment, R. and Black, E.: General TAMSAT-ALERT example datasets, https://gws-access.jasmin.ac.uk/public/tamsat/tamsat_alert/gmd_paper/datasets.zip, last access: 21 November 2024.
- Maidment, R. I., Grimes, D., Black, E., Tarnavsky, E., Young, M., Greatrex, H., Allan, R. P., Stein, T., Nkonde, E., Senkunda, S., and Alcántara, E. M. U.: A new, long-term daily satellite-based rainfall dataset for operational monitoring in Africa, *Sci. Data*, 4, 170063, <https://doi.org/10.1038/sdata.2017.63>, 2017.
- Manzanas, R.: Assessment of model drifts in seasonal forecasting: Sensitivity to ensemble size and implications for bias correction, *J. Adv. Model. Earth Sy.*, 12, e2019MS001751, <https://doi.org/10.1029/2019MS001751>, 2020.
- McKee, T. B., Doesken, N. J., and Kleist, J.: The relationship of drought frequency and duration to time scales, in: *Proceedings of the 8th Conference on Applied Climatology*, 17–22 January 1993, Anaheim, California, USA, pp. 179–183, 1993.

- Nobre, G. G., Muis, S., Veldkamp, T. I., and Ward, P. J.: Achieving the reduction of disaster risk by better predicting impacts of El Niño and La Niña, *Prog. Disaster Sci.*, 2, 100022, <https://doi.org/10.1016/j.pdisas.2019.100022>, 2019.
- Novella, N. S. and Thiaw, W. M.: A seasonal rainfall performance probability tool for famine early warning systems, *J. Appl. Meteorol. Clim.*, 55, 2575–2586, <https://doi.org/10.1175/JAMC-D-16-0102.1>, 2016.
- Parker, D. E., Legg, T. P., and Folland, C. K.: A new daily central England temperature series, 1772–1991, *Int. J. Climatol.*, 12, 317–342, <https://doi.org/10.1002/joc.3370120402>, 1992.
- Reason, C.: Climate of Southern Africa. Oxford Research Encyclopedia of Climate Science, <https://oxfordre.com/climatescience/view/10.1093/acrefore/9780190228620.001.0001/acrefore-9780190228620-e-513> (last access: 17 November 2024), 2017.
- Ren, X., Li, X., Ren, K., Song, J., Xu, Z., Deng, K., and Wang, X.: Deep learning-based weather prediction: a survey, *Big Data Res.*, 23, 100178, <https://doi.org/10.1016/j.bdr.2020.100178>, 2021.
- Salvi, K., Ghosh, S., and Ganguly, A. R.: Credibility of statistical downscaling under nonstationary climate, *Clim. Dynam.*, 46, 1991–2023, <https://doi.org/10.1007/s00382-015-2688-9>, 2016.
- Schneider, U., Becker, A., Finger, P., Meyer-Christoffer, A., Rudolf, B., and Ziese, M.: GPCC full data reanalysis version 7.0: Monthly land-surface precipitation from rain gauges built on GTS based and historic data, Deutscher Wetterdienst (DWD) [data set], Germany, https://doi.org/10.5676/DWD_GPCC/FD_M_V7_050, 2016.
- Shukla, S., McNally, A., Husak, G., and Funk, C.: A seasonal agricultural drought forecast system for food-insecure regions of East Africa, *Hydrol. Earth Syst. Sci.*, 18, 3907–3921, <https://doi.org/10.5194/hess-18-3907-2014>, 2014.
- Siyal, G. E. A., Kazmi, S. M. A., and Zahara, M.: Understanding El-Niño: Tracing its impacts in Sindh, Pakistan, in: 70 Years of Development: the Way Forward, edited by Aneel, S., Haroon, U. T., and Niazi, I., extracted from report issued by the Sustainable Policy Institute, <https://www.jstor.org/stable/resrep24393>. 14 (last access: 26 March 2024), 2019.
- Steduto, P., Hsiao, T. C., Raes, D., and Fereres, E.: AquaCrop – The FAO crop model to simulate yield response to water: I. Concepts and underlying principles, *Agron. J.*, 101, 426–437, <https://doi.org/10.2134/agronj2008.0139s>, 2009.
- Turner, W. A., Husak, G., Funk, C., Roberts, D. A., and Jones, C.: An improved climatological forecast method for projecting end-of-season Water Requirement Satisfaction Index, *J. Hydrometeorol.*, 23, 1281–1295, <https://doi.org/10.1175/JHM-D-21-0232.1>, 2022.
- Yang, W., Kogan, F., Guo, W., and Chen, Y.: A novel re-compositing approach to create continuous and consistent cross-sensor/cross-production global NDVI datasets, *Int. J. Remote Sens.*, 42, 6023–6047, <https://doi.org/10.1080/01431161.2021.1934597>, 2021.
- Young, M., Heinrich, V., Black, E., and Asfaw, D.: Optimal spatial scales for seasonal forecasts over Africa, *Environ. Res. Lett.*, 15, 094023, <https://doi.org/10.1088/1748-9326/ab9d38>, 2020.