

Google's AI river-flood model gives five-day warnings in ungauged basins, Nature reports

Source copy kept by wheresthe.ai for the story "Google's AI river-flood model gives five-day warnings in ungauged basins, Nature reports" (<https://wheresthe.ai/d/google-ai-flood-forecasting-nature-2024-03/>).

Source: <https://research.google/blog/using-ai-to-expand-global-access-to-reliable-flood-forecasts/>

Text saved 18 Oct 2026 (the page as published).

SHA-256 of the saved text: 854de26eb2fd8ad7d290636f7a69fc4abfecc6c8752eca0598acf7bfd3a793ca

Using AI to expand global access to reliable flood forecasts

March 20, 2024 Posted by Yossi Matias, VP Engineering & Research, and Grey Nearing, Research Scientist, Google Research

Large-scale global flood forecasting has been out of reach for a long time. In our Nature paper published today we show how breakthroughs in AI can close the gap & provide reliable flood predictions even in regions that previously lacked data.

Quick links

Paper

Share

Copy link

x

Floods are the most common natural disaster, and are responsible for roughly \$50 billion in annual financial damages worldwide. The rate of flood-related disasters has more than doubled since the year 2000 partly due to climate change. Nearly 1.5 billion people, making up 19% of the world's population, are exposed to substantial risks from severe flood events. Upgrading early warning systems to make accurate and timely information accessible to these populations can save thousands of lives per year. Driven by the potential impact of reliable flood forecasting on people's lives globally, we started our flood forecasting effort in 2017. Through this multi-year journey, we advanced research over the years hand-in-hand with building a real-time operational flood forecasting system that provides alerts on Google Search, Maps, Android notifications and through the Flood Hub. However, in order to scale globally, especially in places where accurate local data is not available, more research advances were required. In "Global prediction of extreme floods in ungauged watersheds", published in Nature, we demonstrate how machine learning (ML) technologies can significantly improve global-scale flood forecasting relative to the current state-of-the-art for countries where flood-related data is scarce. With these AI-based technologies we extended the reliability of currently-available global nowcasts, on average, from zero to five days, and improved forecasts across regions in Africa and Asia to be similar to what are currently available in Europe. The evaluation of the models was conducted in collaboration with the European Center for Medium Range Weather Forecasting (ECMWF). These technologies also enable Flood Hub to provide real-time river forecasts up to seven days in advance, covering river reaches across over 80 countries. This information can be used by people, communities, governments and international organizations to take anticipatory action to help protect vulnerable populations.

Watch the film

Link to Youtube Video

Flood forecasting at Google

The ML models that power the FloodHub tool are the product of many years of research, conducted in collaboration with several partners, including academics, governments, international organizations, and NGOs. In 2018, we launched a pilot early warning system in the Ganges-Brahmaputra river basin in India, with the hypothesis that ML could help address the challenging problem of reliable flood forecasting at scale. The pilot was further expanded the following year via the combination of an inundation model, real-time water level measurements, the creation of an elevation map and hydrologic modeling. In collaboration with academics, and, in

particular, with the JKU Institute for Machine Learning we explored ML-based hydrologic models, showing that LSTM -based models could produce more accurate simulations than traditional conceptual and physics-based hydrology models . This research led to flood forecasting improvements that enabled the expansion of our forecasting coverage to include all of India and Bangladesh. We also worked with researchers at Yale University to test technological interventions that increase the reach and impact of flood warnings. Our hydrological models predict river floods by processing publicly available weather data like precipitation and physical watershed information. Such models must be calibrated to long data records from streamflow gauging stations in individual rivers. A low percentage of global river watersheds (basins) have streamflow gauges, which are expensive but necessary to supply relevant data, and it's challenging for hydrological simulation and forecasting to provide predictions in basins that lack this infrastructure. Lower gross domestic product (GDP) is correlated with increased vulnerability to flood risks , and there is an inverse correlation between national GDP and the amount of publicly available data in a country. ML helps to address this problem by allowing a single model to be trained on all available river data and to be applied to ungauged basins where no data are available . In this way, models can be trained globally, and can make predictions for any river location.

There is an inverse (log-log) correlation between the amount of publicly available streamflow data in a country and national GDP. Streamflow data from the Global Runoff Data Center .

Our academic collaborations led to ML research that developed methods to estimate uncertainty in river forecasts and showed how ML river forecast models synthesize information from multiple data sources . They demonstrated that these models can simulate extreme events reliably , even when those events are not part of the training data. In an effort to contribute to open science, in 2023 we open-sourced a community-driven dataset for large-sample hydrology in Nature Scientific Data .

The river forecast model

Most hydrology models used by national and international agencies for flood forecasting and river modeling are state-space models, which depend only on daily inputs (e.g., precipitation, temperature, etc.) and the current state of the system (e.g., soil moisture, snowpack, etc.). LSTMs are a variant of state-space models and work by defining a neural network that represents a single time step, where input data (such as current weather conditions) are processed to produce updated state information and output values (streamflow) for that time step. LSTMs are applied sequentially to make time-series predictions, and in this sense, behave similarly to how scientists typically conceptualize hydrologic systems. Empirically, we have found that LSTMs perform well on the task of river forecasting.

A diagram of the LSTM, which is a neural network that operates sequentially in time. An accessible primer can be found [here](#) .

Our river forecast model uses two LSTMs applied sequentially: (1) a “hindcast” LSTM ingests historical weather data (dynamic hindcast features) up to the present time (or rather, the issue time of a forecast), and (2) a “forecast” LSTM ingests states from the hindcast LSTM along with forecasted weather data (dynamic forecast features) to make future predictions. One year of historical weather data are input into the hindcast LSTM, and seven days of forecasted weather data are input into the forecast LSTM. Static features include geographical and geophysical characteristics of watersheds that are input into both the hindcast and forecast LSTMs and allow the model to learn different hydrological behaviors and responses in various types of watersheds. Output from the forecast LSTM is fed into a “head” layer that uses mixture density networks to produce a probabilistic forecast (i.e., predicted parameters of a probability distribution over streamflow). Specifically, the model predicts the parameters of a mixture of heavy-tailed probability density functions, called asymmetric Laplacian distributions , at each forecast time step. The result is a mixture density function, called a Countable Mixture of Asymmetric Laplacians (CMAL) distribution, which represents a probabilistic prediction of the volumetric flow rate in a particular river at a particular time.

LSTM-based river forecast model architecture. Two LSTMs are applied in sequence, one ingesting historical weather data and one ingesting forecasted weather data. The model outputs are the parameters of a probability distribution over streamflow at each forecasted timestep.

Input and training data

The model uses three types of publicly available data inputs, mostly from governmental sources: Static watershed attributes representing geographical and geophysical variables: From the HydroATLAS project , including data like

long-term climate indexes (precipitation, temperature, snow fractions), land cover, and anthropogenic attributes (e.g., a nighttime lights index as a proxy for human development). Historical meteorological time-series data : Used to spin up the model for one year prior to the issue time of a forecast. The data comes from NASA IMERG , NOAA CPC Global Unified Gauge-Based Analysis of Daily Precipitation , and the ECMWF ERA5-land reanalysis . Variables include daily total precipitation, air temperature, solar and thermal radiation, snowfall, and surface pressure. Forecasted meteorological time series over a seven-day forecast horizon : Used as input for the forecast LSTM. These data are the same meteorological variables listed above, and come from the ECMWF HRES atmospheric model . Training data are daily streamflow values from the Global Runoff Data Center over the time period 1980 - 2023. A single streamflow forecast model is trained using data from 5,680 diverse watershed streamflow gauges (shown below) to improve accuracy .

Location of 5,680 streamflow gauges that supply training data for the river forecast model from the Global Runoff Data Center .

Improving on the current state-of-the-art

We compared our river forecast model with GloFAS version 4 , the current state-of-the-art global flood forecasting system. These experiments showed that ML can provide accurate warnings earlier and over larger and more impactful events. The figure below shows the distribution of F1 scores when predicting different severity events at river locations around the world, with plus or minus 1 day accuracy. F1 scores are an average of precision and recall and event severity is measured by return period . For example, a 2-year return period event is a volume of streamflow that is expected to be exceeded on average once every two years. Our model achieves reliability scores at up to 4-day or 5-day lead times that are similar to or better, on average, than the reliability of GloFAS nowcasts (0-day lead time).

Distributions of F1 scores over 2-year return period events in 2,092 watersheds globally during the time period 2014-2023 from GloFAS (blue) and our model (orange) at different lead times. On average, our model is statistically as accurate as GloFAS nowcasts (0-day lead time) up to 5 days in advance over 2-year (shown) and 1-year, 5-year, and 10-year events (not shown).

Additionally (not shown), our model achieves accuracies over larger and rarer extreme events, with precision and recall scores over 5-year return period events that are similar to or better than GloFAS accuracies over 1-year return period events. See the paper for more information.

Looking into the future

The flood forecasting initiative is part of our Adaptation and Resilience efforts and reflects Google's commitment to address climate change while helping global communities become more resilient. We believe that AI and ML will continue to play a critical role in helping advance science and research towards climate action. We actively collaborate with several international aid organizations (e.g., the Centre for Humanitarian Data and the Red Cross), to provide actionable flood forecasts. Additionally, in an ongoing collaboration with the World Meteorological Organization (WMO) to support early warning systems for climate hazards, we are conducting a study to help understand how AI can help address real-world challenges faced by national flood forecasting agencies. While the work presented here demonstrates a significant step forward in flood forecasting, future work is needed to further expand flood forecasting coverage to more locations globally and other types of flood-related events and disasters, including flash floods and urban floods. We are looking forward to continuing collaborations with our partners in the academic and expert communities, local governments and the industry to reach these goals.

Labels:

Climate & Sustainability

Machine Intelligence

Quick links

Paper

Share

Copy link

×

Other posts of interest

October 14, 2026

Analyzing technical documents with Deep Review

General Science

.

Machine Intelligence

.

Natural Language Processing

October 14, 2026

Putting Earth AI to work: Geospatial agents for our planet's biggest challenges

Earth AI

.

General Science

.

Machine Intelligence

September 29, 2026

How Diffusion Controller unifies and simplifies AI image generation

Algorithms & Theory

.

Machine Intelligence

× < >

Location of 5,680 streamflow gauges that supply training data for the river forecast model from the Global Runoff Data Center .

Distributions of F1 scores over 2-year return period events in 2,092 watersheds globally during the time period 2014-2023 from GloFAS (blue) and our model (orange) at different lead times. On average, our model is statistically as accurate as GloFAS nowcasts (0-day lead time) up to 5 days in advance over 2-year (shown) and 1-year, 5-year, and 10-year events (not shown).

A diagram of the LSTM, which is a neural network that operates sequentially in time. An accessible primer can be found here .

LSTM-based river forecast model architecture. Two LSTMs are applied in sequence, one ingesting historical weather data and one ingesting forecasted weather data. The model outputs are the parameters of a probability distribution over streamflow at each forecasted timestep.

There is an inverse (log-log) correlation between the amount of publicly available streamflow data in a country and national GDP. Streamflow data from the Global Runoff Data Center .