



TECHNICAL NOTE

Using a deep learning framework to forecast reservoir water availability in India

Samantha Kuzma, Lucas Kruitwagen, Christopher Arderne, Sahana Goswami, Anupriya Goyal, Thomas Lees, Lisa Thalheimer, and Samrat Basak

CONTENTS

Abstract.....1
 1 Introduction 2
 2 Methodology.....3
 3 Results 9
 4 Data Services 16
 5 Limitations17
 6 Next Steps17
 Appendices 18
 Abbreviations20
 Endnotes..... 21
 References..... 21
 Acknowledgments24
 About the authors.....24

Technical notes document the research or analytical methodology underpinning a publication, interactive application, or tool.

Suggested Citation: Kuzma, S., L. Kruitwagen, C. Arderne, S. Goswami, A. Goyal, T. Lees, L. Thalheimer, and S. Basak. 2023. "Using a deep learning framework to forecast reservoir water availability in India." Technical Note. Washington, DC: World Resources Institute. Available online at: <https://doi.org/10.46830/wri.tn.21.00088>.

ABSTRACT

India’s power sector is dominated by thermal power generation, with half of all installed electricity capacity coming from coal. This makes it highly exposed to water risk as thermal power is dependent on the uninterrupted availability of freshwater for a range of operational needs. Disruptions in power supply due to droughts have compounding effects across a range of users from industry to agriculture to human health.

In areas with high water stress, having access to timely information on forecasted water availability could help decision-makers avoid the risk of acute water-driven power outages and advocate for long-term, water-prudent policies and management. We introduce a daily forecast of reservoir water volumes for the coming 90 days using a deep learning framework called the Bayesian long short-term memory sequence-to-sequence-to-sequence model. We show that it is possible to create a high-quality, timely reservoir forecast using global meteorological data. On average, the 11 pilot reservoirs in this study achieved a coefficient of determination score of 92 percent for a short-term (1–14 day) forecast, and 56 percent for the long-term (15–90 day) forecast. Our approach, which does not rely on specific reservoir operations management data, can still provide a time- and cost-efficient solution compared with traditional hydrologic models, and can serve a variety of applications, including power production, food security, urban water supply, and resilience building. The forecast can be used to flag when drought-like conditions threaten water supply, but it should not be used to monitor human interventions or as a tool to inform reservoir management operations.

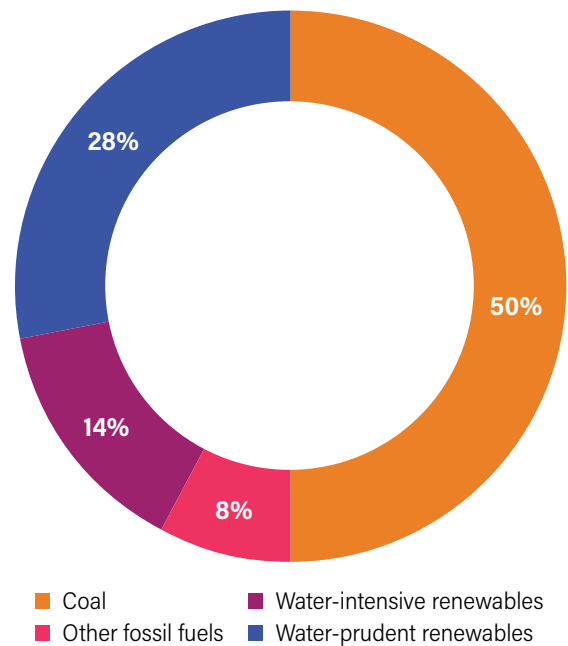
1 INTRODUCTION

By mid-2023, India is expected to have the largest population in the world (AP News 2023). At the same time, it is one of the globe’s most water-stressed countries (Hofste et al. 2019a). India is home to 18 percent of the Earth’s population, but only 4 percent of its freshwater (India-WRIS 2021). As India’s population and economy continue to grow, so will the competition over its limited water resources. India’s water stress is intensified by severe droughts, which amplify the strain on water supply. Droughts reduced India’s gross domestic product by an estimated 2–5 percent between 1998 and 2017 (UNDRR 2021). By the end of the century, flash droughts—characterized by the rapid depletion of soil moisture—are projected to increase sevenfold due to climate change (Mishra et al. 2021a).

The increasing demand for its water resources, coupled with changes in supply patterns caused by climate change, pose serious implications for the future health and well-being of India’s economy and people. One often overlooked consequence of drought is the risk of power outages due to insufficient water supply. While the energy sector accounts for only about 3 percent¹ of India’s water withdrawals (compared with at least 80 percent for agriculture) (de Oliveira Bredariol et al. 2021; FAO 2020), it is still susceptible to cut-offs when water availability is low. Water is a necessary resource for thermal energy production, from cooling equipment to cleaning ash. India’s current power mix and generation is largely comprised of thermal power—coal represents 50 percent of installed capacity as shown in Figure 1—and is highly vulnerable to water scarcity, droughts, and climate change (CEA 2022). Enormous volumes of India’s thermoelectric-heavy power production—up to 14 terawatt-hours of electricity generation during the 2016 drought, enough energy to power Sri Lanka for a year—are lost due to insufficient water to cool power plants (Luo et al. 2018). These shutdowns are costly, too, losing around \$1.4 billion in potential revenue from sales from 2013 to 2016 (Luo et al. 2018). Further, water-intensive power plants also siphon water away from essential food production, and the daily needs of cities, businesses, and households.

India has set a nationally determined contribution target of 50 percent non-fossil fuel-based electric power installed capacity by 2030 to curb greenhouse gas emissions (PIB Delhi 2022). It is imperative that this effort be paired with water-saving policies—such as retrofitting older plants with new cooling methods like dry cooling, or decommissioning water-intensive plants. One study found India’s water withdrawals for electricity could grow ninefold by 2050 from a 2010 baseline if no changes to the power mix or water usage are made (Srinivasan

Figure 1 | Installed electric capacity mix in India as of September 2022



Notes: Water-prudent power sources include wind, solar, biomass, and waste-to-energy. Water-intensive power sources include hydro and nuclear.

Source: Data from CEA 2022.

et al. 2018). This same study found that water withdrawals could also increase under a low-carbon scenario if water-intensive technologies like nuclear are prioritized (Srinivasan et al. 2018). In such a scenario, the power sector would continue to be constrained by water availability.

Having access to current and forecasted water availability in reservoirs that supply the power sector with water could help illuminate the chronic issues around India’s dependence on water for power and avoid acute instances of power outages due to insufficient cooling water supply. Trusted and open-source water risk information—especially if rooted in the government’s own data—can help drive continued on-the-ground, evidence-driven decision-making related to future energy development and operation of existing power plants in water-stressed scenarios, as well as advance the social right to water in a highly competitive water use landscape.

World Resources Institute (WRI) and its partner Vasudha Foundation present Water4Power (www.water4power.wri-india.org), an initiative designed to create decision-ready information on the water-energy nexus (WRI India 2020). Our pilot project provides open access to near-real-time alerts of potential water shortages for 11 reservoirs in India, all of

which source water for thermal coal power generation. The pilot reservoirs serve roughly 20 percent² of India's installed electric capacity from coal and include a broad range of facilities in terms of location and size. In addition, some but not all the pilot reservoirs have faced cut-offs due to water stress. For example, the Seoni thermal power plant, which receives water from the Bargi reservoir, was shut down for 32 days in the summer of 2017 due to raw water unavailability (VPIH 2023). The diverse mix of pilot reservoirs will help decision-makers weigh multiple factors as they begin to prioritize policies ranging from upgrading water-saving cooling technologies to decommissioning plants.

WRI found another research partner, H2Ox, through a hackathon designed to solicit innovative solutions to forecasting reservoir water levels (see Box 1). Together, we introduce a cutting-edge approach to forecast daily changes in reservoir water availability for up to 90 days using a Bayesian long short-term memory (BLSTM) neural network.

Box 1 | A hackathon success story

In 2021, WRI, Microsoft, and BlackRock hosted the Wave2Web Hackathon.^a We invited teams of students, professionals, and start-ups from all over the world to develop a machine learning-based model to forecast water levels for the reservoirs supplying the city of Bengaluru with drinking water. Out of the 26 teams that submitted proposals, H2Ox, coauthors on this paper, won the hackathon with their sophisticated model and dashboard (<https://www.h2ox.org/>).^b Following the hackathon, WRI and H2Ox shared a vision to bring their winning sequence-to-sequence long short-term memory (LSTM) model into production. Our work has culminated with the launch of the Water4Power dashboard, presented in this paper. We have also released a research paper that evaluates the impact of the various model structures on performance.^c In that paper, the research covers 66 reservoirs in India (including the original reservoirs from the hackathon) and serves as a companion piece to this technical note for those interested in machine learning science.

Sources:

a: WRI India 2021.

b: Kruitwagen et al. 2021, 2.

c: Kruitwagen et al. 2022b.

2 METHODOLOGY

2.1 Baseline model

As part of the Water4Power Initiative, we set out to develop an adaptable model that could produce daily forecasts of reservoir water volumes for the coming 90 days. We wanted to create a nonprescriptive model that could be applied to different locations and contexts, and that was more time- and cost-efficient to run than traditional hydrologic models. We call this the baseline model as it has the fundamental elements needed to create the forecast. Through the baseline model, we show that it is possible to generate high-quality, reliable, near-real-time reservoir forecasts using mainly meteorological data derived from open-source satellite products. We intentionally kept the method generalized—meaning we avoided using local datasets and limited the number of inputs—so that it can easily be adapted to regions outside of our pilot study area. Just as importantly, we also selected data that could support the near-real-time production of the forecast. As we describe in “Data services,” the baseline model can be customized with site-specific and other types of information such as local weather data.

As a result, our model does not explicitly include reservoir operations—human interventions that significantly impact the amount of water available in a reservoir. Rather, the model infers patterns among water availability, upstream and downstream reservoir flows, and climate-related indicators to create its forecast. Therefore, the baseline model can be used as a tool to alert stakeholders when drought-like conditions threaten water supply; it cannot be used to alert or inform stakeholders about upcoming reservoir operations that might impact supply.

Overview

The objective of the baseline model is to create a transparent, timely forecast of daily water volumes for the coming 90 days to alert stakeholders of potential water shortages before they happen. Our pilot focuses on 11 reservoirs in India, all of which source water for thermal coal power generation. We used a BLSTM sequence-to-sequence-to-sequence (seq2seq2seq) deep learning model with a graph convolutional layer to produce the 90-day forecast. Specifically, the model forecasts the net change in water volume from the previous day, which can be converted back to volumes in post-processing. We created a unique model for the six major river basin networks featured in our pilot. A reservoir's forecast was created by running its basin's model with reservoir-specific daily inputs.

Modeling extent

While our pilot study focuses on only 11 reservoirs, we acknowledge that these reservoirs do not exist in a vacuum. The management of upstream and downstream reservoirs will likely impact water availability. Therefore, we included reservoirs³ adjacent to the 11 pilot reservoirs as model inputs via a graph convolutional layer (called an adjacency matrix) to capture interbasin interactions—specifically inflows and outflows that impact the main pilot reservoirs. In total, 38 reservoirs were involved in our modeling extent (11 pilot reservoirs and 27 auxiliary reservoirs, as illustrated in Figure 2) mapping across six major river basins.

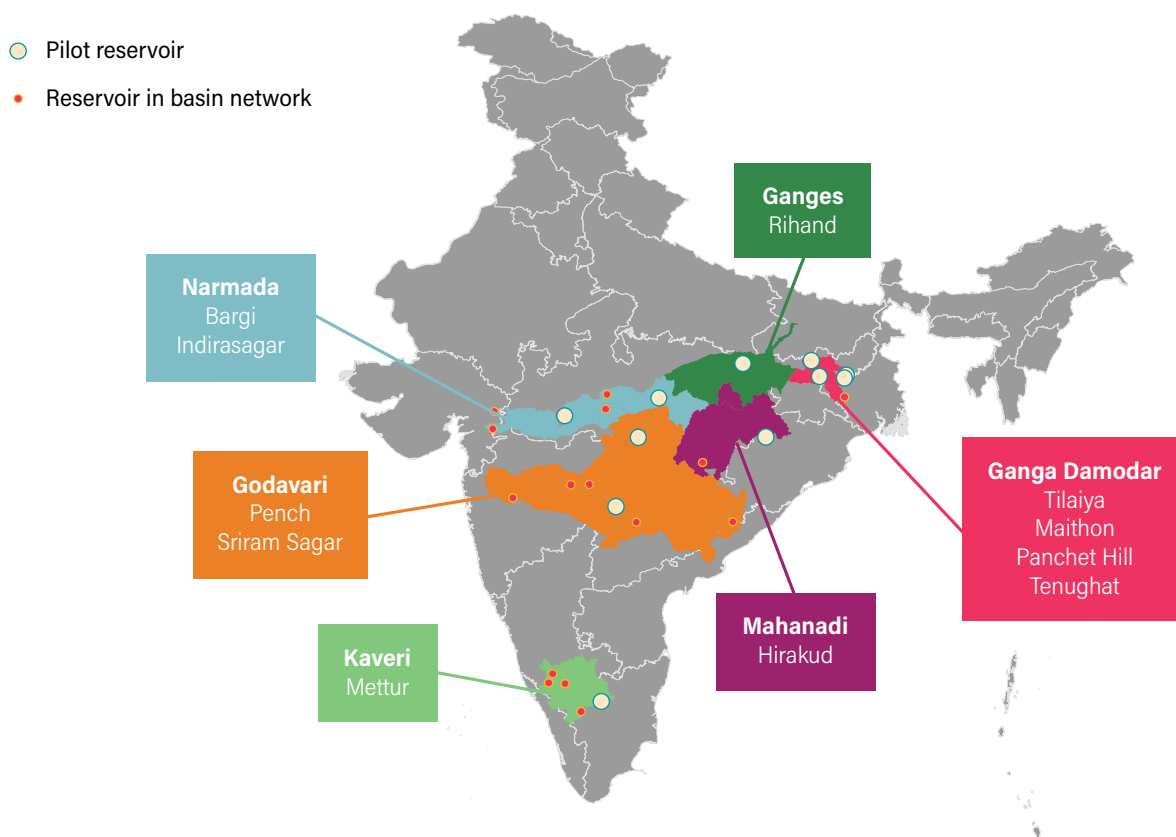
Data sources

India's governmental Central Water Commission (CWC) publishes near-real-time data on reservoir water levels and volumes via the India Water Resources Information System (India-WRIS 2008). We selected CWC's reservoir water availability data because the data are published on an open-source platform and the CWC has a minimum 10-year record

of daily data (the longer the record, the more instances of dry spells our model can learn from). CWC data have the added benefit of being an official government dataset, which can help with local acceptance and application of the research.

Precipitation and temperature were identified as strong predictors of future reservoir water levels during the hackathon. While there are many potential sources of this information, our research required data to be open source and readily accessible (via an application programming interface [API] or similar functionality), with a 10-year-record minimum, and of consistent quality throughout India. The accessibility element ruled out most local weather datasets for the baseline model. We identified two sources that met our criteria: the European Centre for Medium-Range Weather Forecasts (ECMWF) and the Climate Hazards Group InfraRed Precipitation with Station data (CHIRPS). Both provide near-real-time, globally available historic meteorological data. ECMWF has the added benefit of also providing weather predictions for 14 days.

Figure 2 | **Locations of the 11 pilot Water4Power reservoirs**



Notes: Reservoirs are grouped by their basin network, named at the top of each box. In addition, we show neighboring reservoirs within each network that fed inflows and outflows to the pilot reservoirs via an adjacency matrix.

Source: Authors. Data used: Reservoir locations (Lehner et al. 2011); HydroBASINS catchments (Lehner and Grill 2014).

Within the ECMWF family, there were two datasets of interest: the ECMWF ReAnalysis Version 5 Land (ERA5-Land) (Muñoz Sabater 2019), and THORPEX Interactive Grand Global Ensemble (TIGGE) (Swinbank et al. 2016). ERA5-Land provides high-resolution (0.5 degree, hourly time step), cloud-optimized total precipitation and two-meter surface air temperature, but with a two-to-three month lag in data delivery at the time of the research.⁴ TIGGE provides daily forecasts of total precipitation and two-meter surface air temperature for the coming 14 days, also in high resolution (0.5 degree, six-hour time step) but with no lag in delivery. Therefore, we selected TIGGE for both the ex ante meteorological data and historic temperature data (data from Day 0 were used to represent historic conditions; all other days (1–14) were used to represent forecasted conditions).

The CHIRPS dataset provides daily total precipitation values available with a 0.05-degree resolution (Funk et al. 2015). CHIRPS precipitation data are derived from satellite imagery

and bias corrected using gauge data. Studies show that CHIRPS data are comparable to gauge-based precipitation estimates in India (Prakash 2019). CHIRPS was selected as the source for historic precipitation data because of its high resolution (higher than TIGGE), its proven quality in India, and its near-real-time availability.

The final input included in the model does not have a source because it is date-based. We used a day-of-year variable to help the model track seasonality within a year. In India, there are four main hydrologic seasons: monsoon (June–September), post-monsoon (October–December), winter (January–February), and summer (March–May) (Attri and Tyagi 2010). Day-of-year are the only data used to train the model during every stage of the BLSTM. A list of all data sources can be found in Table 1.

Table 1 | List of predictors used to train the LSTM model along with their data sources

TYPE	SOURCE (CITATION)	SPATIAL RESOLUTION	TEMPORAL RESOLUTION	UPDATED DATA AVAILABILITY	BLSTM STAGE		
					Stage 1 (Day 0)	Stage 2 (Days 1–14)	Stage 3 (Days 15–90)
Reservoir water volumes	CWC (India-WRIS 2008)	Reservoir-wise	Daily	Every 1–7 days	X		
Historic precipitation	CHIRPS (Funk et al. 2015)	0.05 degrees	Daily	Every 5 days	X		
Historic two-meter surface air temperature	ECMWF TIGGE (Swinbank et al. 2016)	0.5 degrees	6-hour time steps; forecasts for up to 14 days	Daily	X		
Forecasted precipitation	ECMWF TIGGE (Swinbank et al. 2016)	0.5 degrees	6-hour time steps; forecasts for up to 14 days	Daily		X	
Forecasted two-meter surface air temperature	ECMWF TIGGE (Swinbank et al. 2016)	0.5 degrees	6-hour time steps; forecasts for up to 14 days	Daily		X	
Day-of-year	N/A	Reservoir-wise	Daily	Daily	X	X	X

Source: BLSTM = Bayesian long short-term memory; CWC = India's Central Water Commission; CHIRPS = Climate Hazards Group InfraRed Precipitation with Station data; ECMWF = European Centre for Medium-Range Weather Forecasts; TIGGE = THORPEX Interactive Grand Global Ensemble; N/A = not applicable.

Source: Authors.

Data processing

The BLSTM model required that data be in a standardized format to interpret them together; therefore, data inputs (historic water volumes, historic and forecasted precipitation and temperature, and day-of-year) were normalized to a scale from -1 to 1 and tied to a forecast date and reservoir name, as shown in Table 2. Missing data were treated as follows: null data with fewer than 15 missing consecutive days were imputed using a cubic interpolation approach; data with more than 15 missing consecutive days were deemed invalid and removed.

WATER VOLUMES

We first converted the reservoir water volumes into net change by subtracting the previous day's volume; we then z-scored the data using the reservoir-wise mean and standard deviation to normalize the data to a scale of -1 to 1 .

TEMPERATURE AND PRECIPITATION

The raw meteorological data came in a gridded format and therefore needed to be reduced to a reservoir-specific area to become an input dataset. To define this area per reservoir, we used the HydroBASINS Level 8 nested catchment network (Lehner and Grill 2014) to identify catchment

geometries at and uniquely upstream of each reservoir, as this area captures the weather events that directly impact water availability. We specify uniquely upstream because the flow from adjacent reservoirs is accounted for in the model architecture via the graph convolutional layer, and we wanted to avoid double-counting adjacent precipitation as an input. The process started at the most upstream reservoir within each basin, where we merged its upstream catchments into a single geometry. The process continued reservoir by reservoir, cascading down the basin's hydrological graph so that every catchment flowed into only one reservoir, and every reservoir was assigned its unique upstream area. The resulting areas were used to reduce the gridded daily meteorological data into a mean and standard deviation per reservoir, which were then z-scored to a scale of -1 to 1 .

DAY-OF-YEAR

The day-of-year input represents the date sequence for each day within the 90-day forecast. For example, the data sequence for a forecast made on January 1 runs from January 2 to April 1 (i.e., 90 days from January 1). To create the day-of-year variable, we converted the date sequences into day-of-year values (January 1 is 1, December 31 is 366), and applied a cosine transformation to normalize it to a scale of -1 to 1 (see Figure 3).

Table 2 | Example input data structure for the 90-day forecast

Stage	FORECAST PARAMETERS			MODEL INPUTS					
	Reservoir	Date of forecast	Step (day from forecast)	Water volume	Historic precipitation	Historic temperature	Forecasted precipitation	Forecasted temperature	Day of year
1	A	2011-01-01	0	0.0018	-0.3325	-0.5419	NaN	NaN	0.999852
2	A	2011-01-01	1	N/A	N/A	N/A	-3.722	0.1247	0.99407

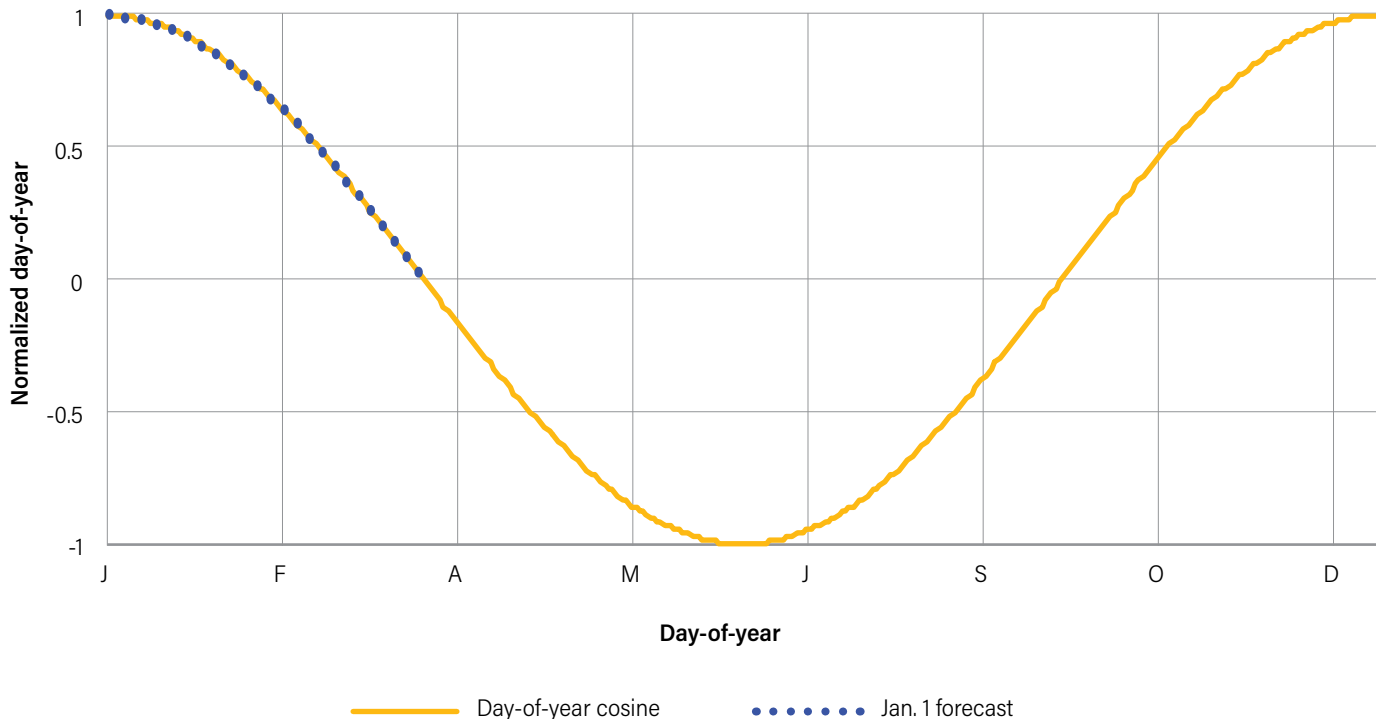
	A	2011-01-01	14	N/A	N/A	N/A	0.2493	0.0211	0.966848
3	A	2011-01-01	15	N/A	N/A	N/A	N/A	N/A	0.962309

	A	2011-01-01	90	N/A	N/A	N/A	N/A	N/A	0.004304

Source: N/A = not applicable.

Source: Authors.

Figure 3 | Normalized day-of-year input variable



Notes: An example of the day-of-year input data for a January 1 forecast is shown in blue.

Source: Authors.

Model Structure

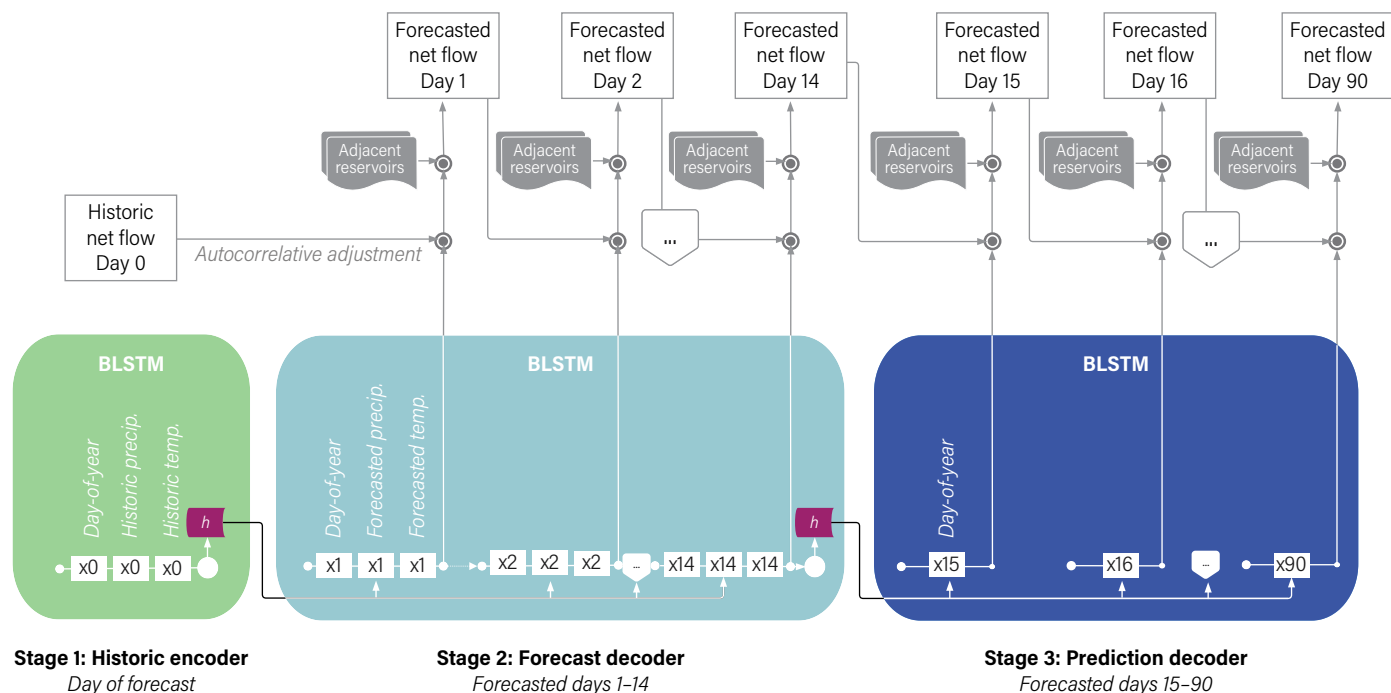
OVERVIEW OF TERMS

We used a BLSTM seq2seq2seq deep learning framework with a graphical convolutional layer to produce the daily forecast of reservoir water volumes for the coming 90 days. The term Bayesian simply refers to the probabilistic nature of the model. In short, it allowed us to quantify uncertainty (see the “Uncertainty” section). An LSTM (Hochreiter and Schmidhuber 1997) is a type of neural network adept at handling a long-term sequence like daily reservoir water volumes thanks to its use of gates to control the flow of information—only data that help the model learn patterns are passed on. It enables the model to learn from longer sequences without running into decaying error issues (Hochreiter and Schmidhuber 1997). Seq2seq2seq refers to a type of LSTM that enables a model to take in a sequence of data and return an output sequence of a different length—like our 14-day

forecasted meteorologic input data with our 90-day forecast window. Sequence-to-sequence, first published by Sutskever et al. (2014) as a language translation tool, has shown strong skill in forecasting hydrologic topics such as monsoon spell categories in India (Viswanath et al. 2019); reservoir inflows (Lee and Kim 2021); and surface water runoff (Xiang et al. 2020).

Following Xiang et al. (2020)’s seq2seq structure, we made use of encoders and decoders to pass sequences of varying length between stages. An encoder maps input data to a fixed-length vector called a hidden state, which is then passed to the next stage as an input. A decoder translates the hidden state to its original format per time step within its stage. We added an additional decoder, meaning our model contains three sequenced BLSTMs, as shown in Figure 4.

Figure 4 | Schematic of the BLSTM seq2seq2seq



Notes: Each stage represents a Bayesian long short-term memory (BLSTM) model. In Stages 1 and 2, input data are encoded into a hidden state variable (h) and passed along to the decoder in the next stage. In Stages 2 and 3, the BLSTM outputs are first adjusted by the previous day's volume and then run through a graph convolutional layer to account for the flow of water from adjacent reservoirs. Within each BLSTM, x represents the step from the forecast date (x_0). For example, x_1 is the first step (Day 1), x_{15} is the 15th step (Day 15), and so on. precip. = precipitation; temp. = temperature.

Source: Authors, based on Kruitwagen et al. 2022b.

MODEL ARCHITECTURE

In Stage 1, the historic encoder, the day-of-year variable, and the historic meteorological data were encoded into a hidden state variable, which was passed on to Stage 2 along with historic water volume (converted into net change).

Next, Stage 2, the forecast decoder, was conditioned using the day-of-year variable and the 14-day ex ante meteorological data (the forecast window of this stage is limited to the 14-day lead time of its input data). Upon each time step (day), the decoder inherited the hidden state from Stage 1 and output an estimated net change in volume, which was then appended with the previous day's value via a linear header (i.e., autocorrelative model). Then we applied the reservoir adjacency matrix⁵ via a graphical convolutional layer⁶ to adjust the final output (net inflow/outflow) based on neighboring reservoir activity.

Stage 3, the prediction decoder, was designed to extend the prediction horizon to an arbitrary length: in our case, Days 15–90. It was conditioned on the day-of-year variable and inherited the updated hidden state variable from Stage 2. Like in Stage 2, the output from Stage 3 is first appended to the previous day's volume and then run through the graphical convolutional layer before producing the final output of net change.

Once the full forecast is run, the net change values are converted back to volumes by de-z-scoring and adding to the previous day's volume. One model is trained for every basin network—six in total for the Water4Power pilot. That model is then run per pilot reservoir to generate a localized forecast.

Uncertainty

Uncertainty is a critical tool for understanding model performance. It helps us understand the level of confidence we can have in a given prediction. However, with neural networks such as an LSTM, quantifying uncertainty is a challenge due to the black-box nature of the models; there are too many possible configurations of weights to consider (Blundell et al. 2015). Therefore, we followed Blundell et al. (2015)'s Bayes by Backprop method to quantify uncertainty and reduce error. In short, each weight within our model was normalized to its mean and standard deviation. We then applied two terms within the optimization step to minimize loss in our weights: mean squared error (MSE), which measures likelihood loss; and Kullback-Leibler (KL) divergence, which measures complexity loss. MSE quantifies the magnitude of errors made by the model by measuring the difference between the model's predicted versus actual values. A larger MSE means the prediction is far from the actual value. KL divergence, then, measures the difference between the predicted distribution and the target distribution. We found that our best results came when heavily weighting the likelihood loss term and relaxing the assumptions around the complexity term.

Training

The full pool of data spanned 11 years, from 2010 through 2021. We segmented the data by year into three pools: train, validate, and test. Data reserved for training were used to create the model architecture and weights; the validation years were used to create hyperparameters for the model; the test data were withheld from the model creation process and used only to compare the model predictions to the actual results at the end.

Our model was trained on seven years of data (highlighted in dark blue in Figure 5), meaning we have around 2,500 data points per reservoir to fit to the model (though some reservoirs have less data availability). We validate and test on two years (730 days) of data (each) (medium blue and light blue in Figure 5, respectively). By having at least a decade of data, we can train the model on multiple dry episodes and test its ability to predict those events. In fact, this is why we did not reserve the final four years for validation and testing, as is typical to avoid look-ahead biases; we wanted to test performance in years with different climatic conditions.

3 RESULTS

Given that the forecast results are made available as water volumes rather than net changes in volume (i.e., the model's native output), we evaluate performance in terms of volume. Volume is a more actionable metric for our stakeholders, who need to understand the change relative to the current water supply. Results can be found in Table 3 and seen spatially in Figure 6.

Overall, performance varies based on the prediction horizon (i.e., how far out the estimate is from the date of prediction), the reservoir, and the season. The short-term forecast from Stage 2 (coming 14 days) had the least amount of variation. All pilot reservoirs had a strong short-term forecast, averaging a coefficient of determination (R^2) of 92 percent. The R^2 value measures the goodness of fit from a scale of 0 to 1. The closer to 1, the closer the forecast is to the actual results. One reason the short-term forecast performs well is because it includes the ex ante forecasts for precipitation and temperature.

Figure 5 | Years used to train, validate, and test the model

	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021
Train	Dark Blue	White	Dark Blue	White	Dark Blue	Dark Blue	Dark Blue	White	White	Dark Blue	White	Dark Blue
Validate	White	White	White	Medium Blue	White	White	White	White	Medium Blue	White	White	White
Test	White	Light Blue	White	White	White	White	White	White	White	White	Light Blue	White

Note: The year 2017 was excluded from the dataset due to corrupted data from TIGGE (THORPEX Interactive Grand Global Ensemble).

Source: Authors.

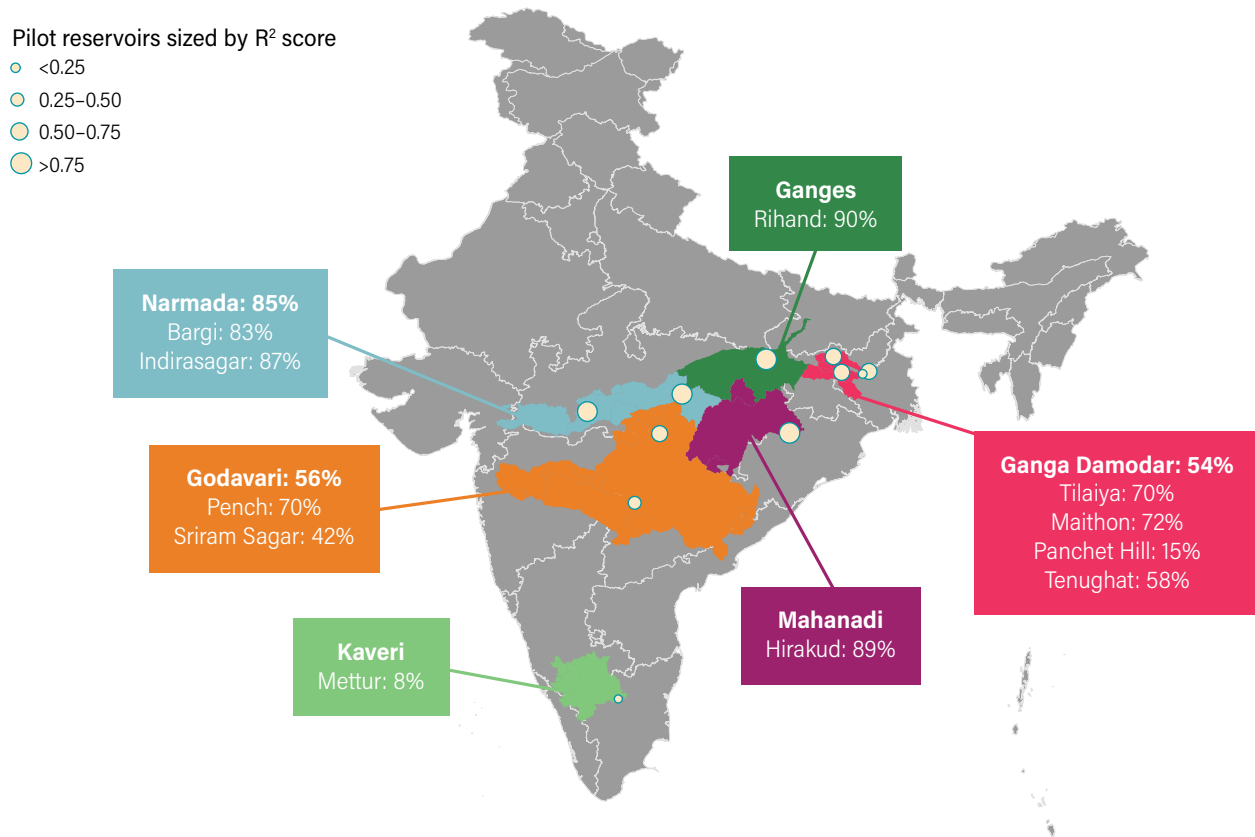
Table 3 | Coefficient of determination of predicted water volume for test data

BASIN	PILOT RESERVOIR	FULL STORAGE (MCM)	PREDICTION HORIZON		
			STAGE 2 (DAYS 1-14)	STAGE 3 (DAYS 15-90)	OVERALL (DAYS 1-90)
Ganges	Rihand	10,600	0.99	0.88	0.90
Ganga Damodar	Maithon	1,094	0.93	0.67	0.72
	Panchet	1,497	0.85	0.00	0.15
	Tenughat	1,021	0.76	0.54	0.58
	Tilaiya	395	0.96	0.60	0.70
Godavari	Pench	1,241	0.97	0.65	0.70
	Sriram Sagar	3,172	0.94	0.32	0.42
Kaveri	Mettur	2,707	0.82	-0.06	0.08
Mahanadi	Hirakud	5,896	0.97	0.88	0.89
Narmada	Bargi	3,920	0.98	0.81	0.83
	Indira Sagar	12,220	0.97	0.85	0.87
Average			0.92	0.56	0.62
Standard deviation			0.07	0.32	0.28

Note: The performance for each pilot reservoir is shown for three prediction horizons with the model structure: Stage 2 (forecast decoder), which predicted days 1-14; Stage 3 (prediction decoder), which predicted days 15-90; and overall (combination of Stages 2 and 3). The R^2 value was calculated by comparing the predicted water volumes to the actual values for the days within each prediction horizon from the Test subset. The table also includes the full (gross) storage volume in million cubic meters (MCM).

Source: Authors.

Figure 6 | Overall coefficient of determination of predicted water volume for pilot reservoirs



Notes: The mean overall performance for each pilot reservoir. Reservoirs are grouped into their basin networks. R² = coefficient of determination.

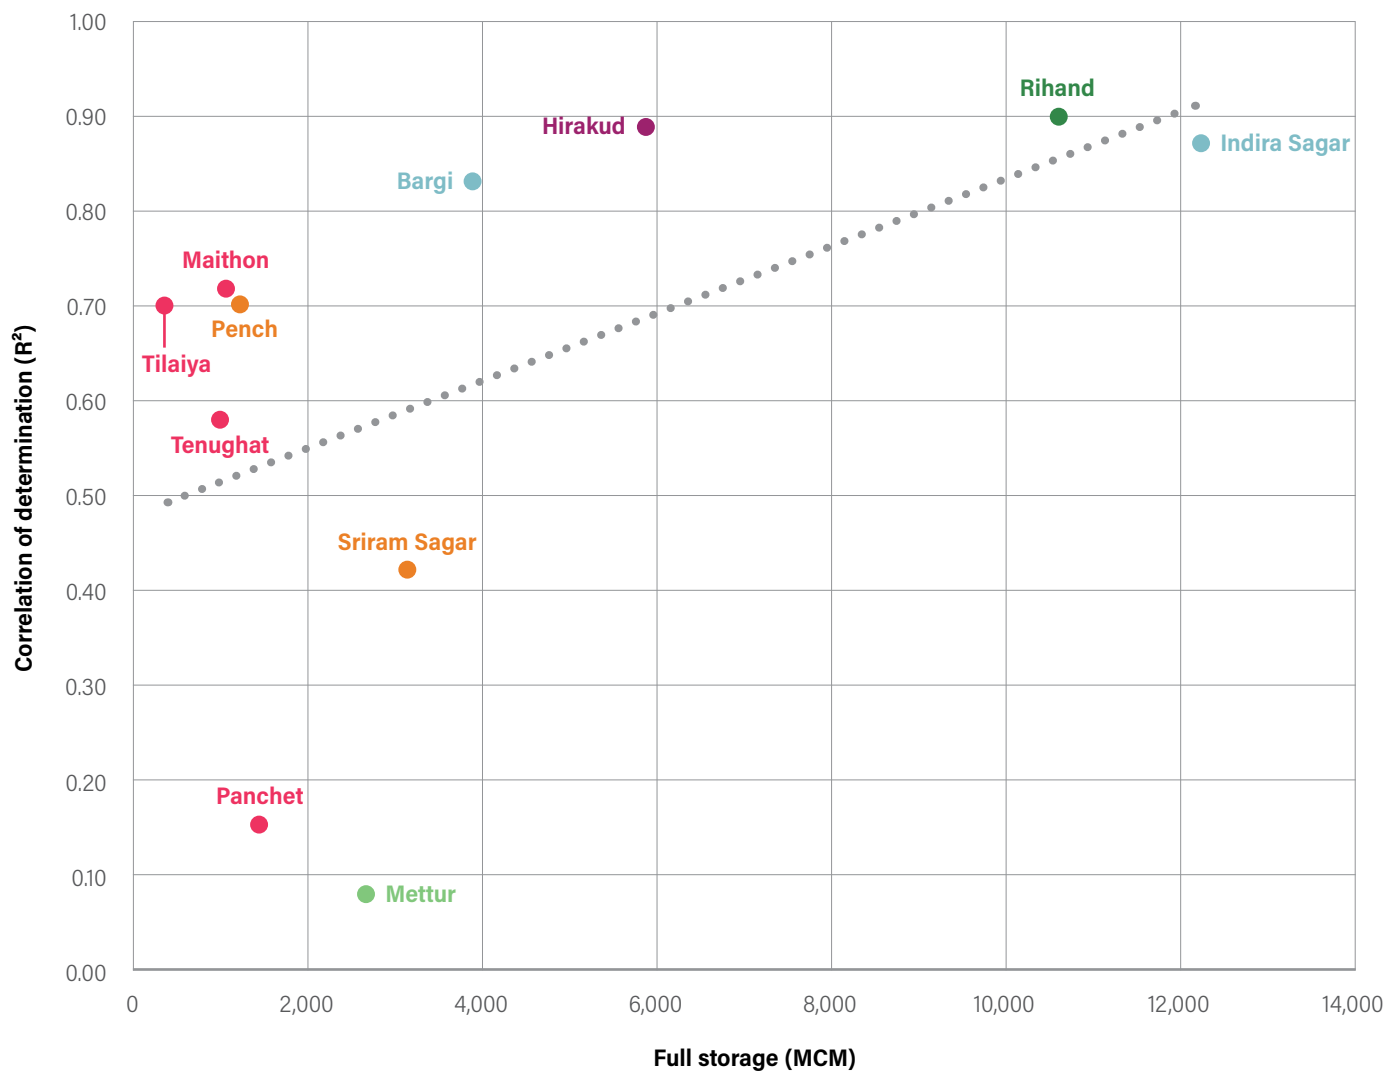
Source: Data used: Reservoir locations (Lehner et al. 2011); HydroBASINS catchments (Lehner and Grill 2014).

Looking at the longer-term forecast from Stage 3 (coming 15–90 days), we see that the performance degrades, especially in the Ganga Damodar, Godavari, and Kaveri basins. This is not surprising given that day-of-year is the only input data available for this stage (aside from the hidden state).

The performance among reservoirs may vary for a few reasons, such as size and location. Figure 7 highlights that reservoirs smaller than 4,000 billion cubic meters perform worse than the larger reservoirs, with an average R² value of 48 percent versus 87 percent. Figure 7 shows that performance is also clustered by basin. For example, the reservoirs in the Ganga Damodar have some of the weakest prediction power. The

Ganga Damodar reservoirs are operated as a network (Saha et al. 2017) to coordinate water releases together to manage floods and meet water demands for irrigation, water supply, and electrical power generation (Sen 2021). While it is common for reservoirs to be operated in a network, the close spatial proximity of these reservoirs may mean that water exchanges may happen more frequently and/or be felt more rapidly than farther-spaced reservoirs. In other words, human intervention may be a main driver of water volumes in these nested reservoirs, which is something our meteorological predictors cannot fully learn.

Figure 7 | Forecast performance organized by reservoir size and major basin



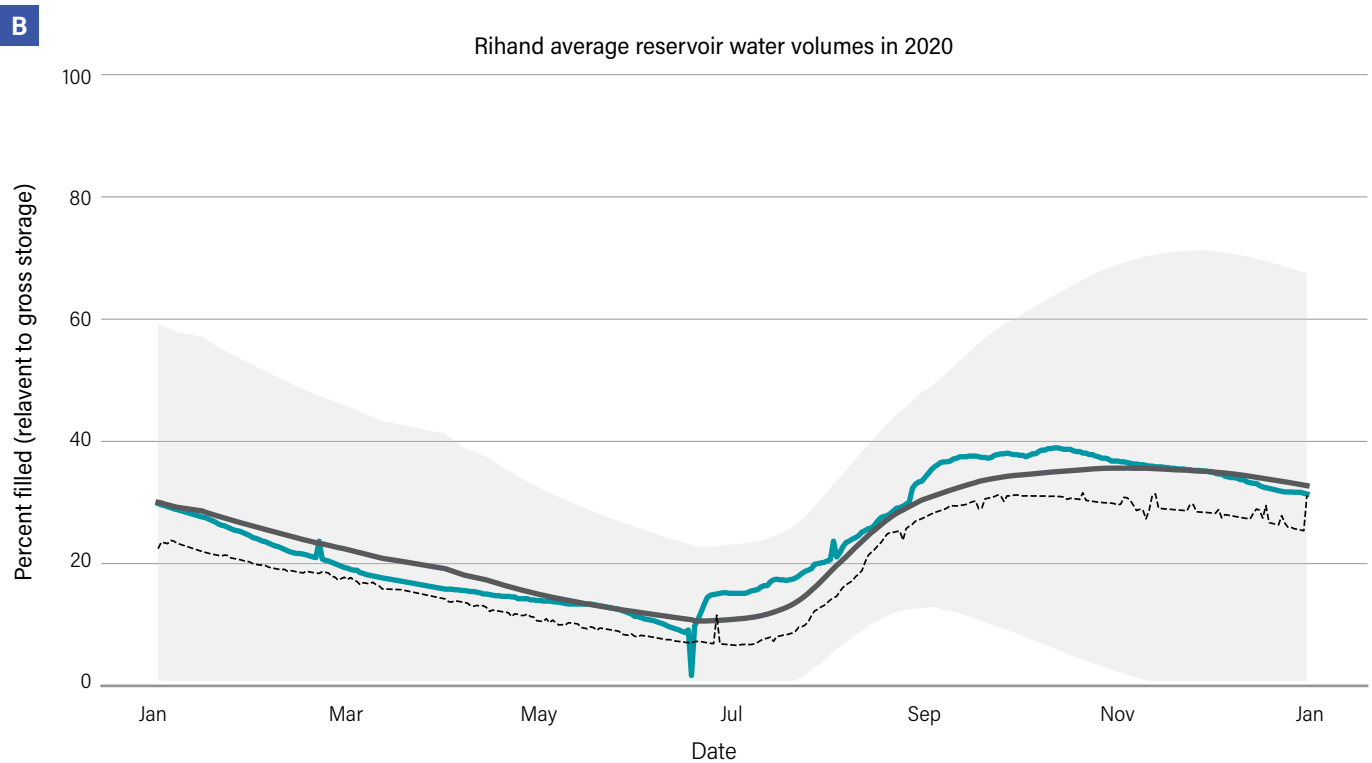
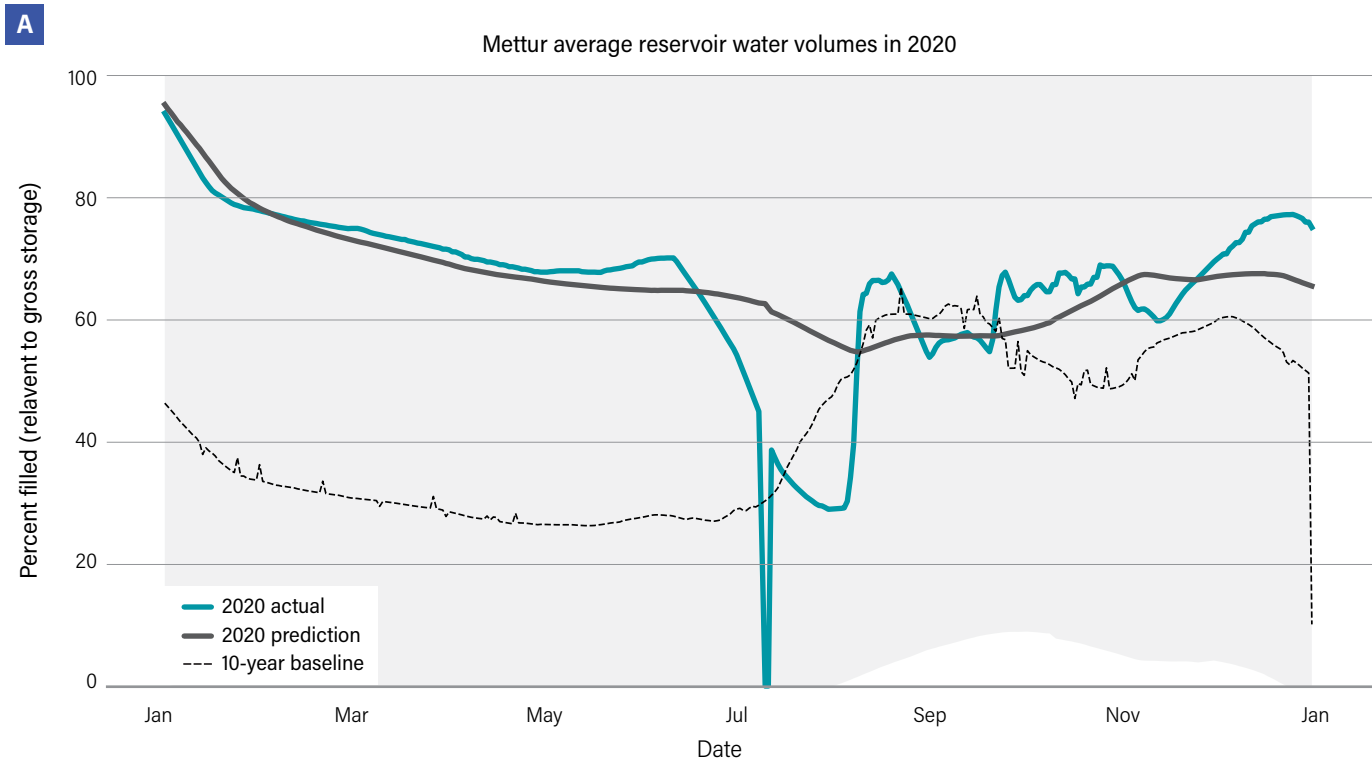
Note: Colors correspond to basin groupings. R^2 = coefficient of determination; MCM = million cubic meters.

Source: Authors.

To understand the difference in performance further, it is useful to examine the volumes themselves. Figure 8A shows the 2020 volume data for the worst-performing reservoir, Mettur, which has a gross storage capacity of 2,700 million cubic meters (MCM). The time series is noisy—its daily volumes jump around—making it difficult to predict. This may reflect human intervention in its water management, such as transboundary water transfers from the upstream state of

Karnataka (Ghosh and Bandyopadhyay 2009). Regardless, the impact is a lower-quality forecast—which can be seen in both the R^2 value as well as the massive range of uncertainty shown in gray. Comparatively, Figure 8B shows the Rihand reservoir, with a gross capacity of just above 10,000 MCM, in the Ganges basin. This is larger than Mettur and has a smoother—more predictable—time series on average. It also has much less uncertainty surrounding its forecast.

Figure 8A-B | Average water volumes in 2020 for low-performing and high-performing reservoirs



Note: Water volumes for Mettur, with one of the lowest overall R^2 values of 8 percent, compared with Rihand, with one of the highest at 90 percent. Mettur is about a quarter of the size of Rihand. The actual and predicted volumes represent the 2020 average for that day of year (a single day may be included in up to 90 forecasts). The ± 95 percent confidence intervals of the forecast are shown in gray. The 10-year baseline represents the average actual value for that day of year using data from the 2010–2020 average. R^2 = coefficient of determination.

Source: Authors.

Seasonality is the final factor to consider. On average, about 80 percent of India’s annual rainfall occurs during the monsoon season (June-September) (Mishra 2020). We do not include monsoon-specific forecasts, which could help capture the varying timing and intensity of a particular year’s cycle, and, as a result, it is the second-worst-performing season with an R² value of 37 percent. Ironically, the post-monsoon season, which typically has the most stagnate volumes, performs the worst. In fact, as Table 4 shows, three of the reservoirs had R² values below –1, meaning a horizontal line would have matched the actual volumes better than our forecast. In fact, that is exactly the problem: Our model is unable to learn how to “store” water during the post-monsoon (see Figure 9),

and is heavily penalized by the R² value as a result. Winter is the best-performing season (average R² value of 66 percent), with many reservoirs experiencing steady drawdowns to meet supply due to the lack of precipitation. That said, some reservoirs like Panchet and Tilaiya still store water without many releases; in those cases, the winter forecast fails. The model performs comparatively well during the summer season (average R² value of 47 percent), which is surprising given the number of operational decisions that must happen during this time of the year. As water planners wait for the monsoon, they must ration what’s left in the reservoir, while making sure there is enough room to store the upcoming monsoon’s rainfall.

Table 4 | Coefficient of determination by season

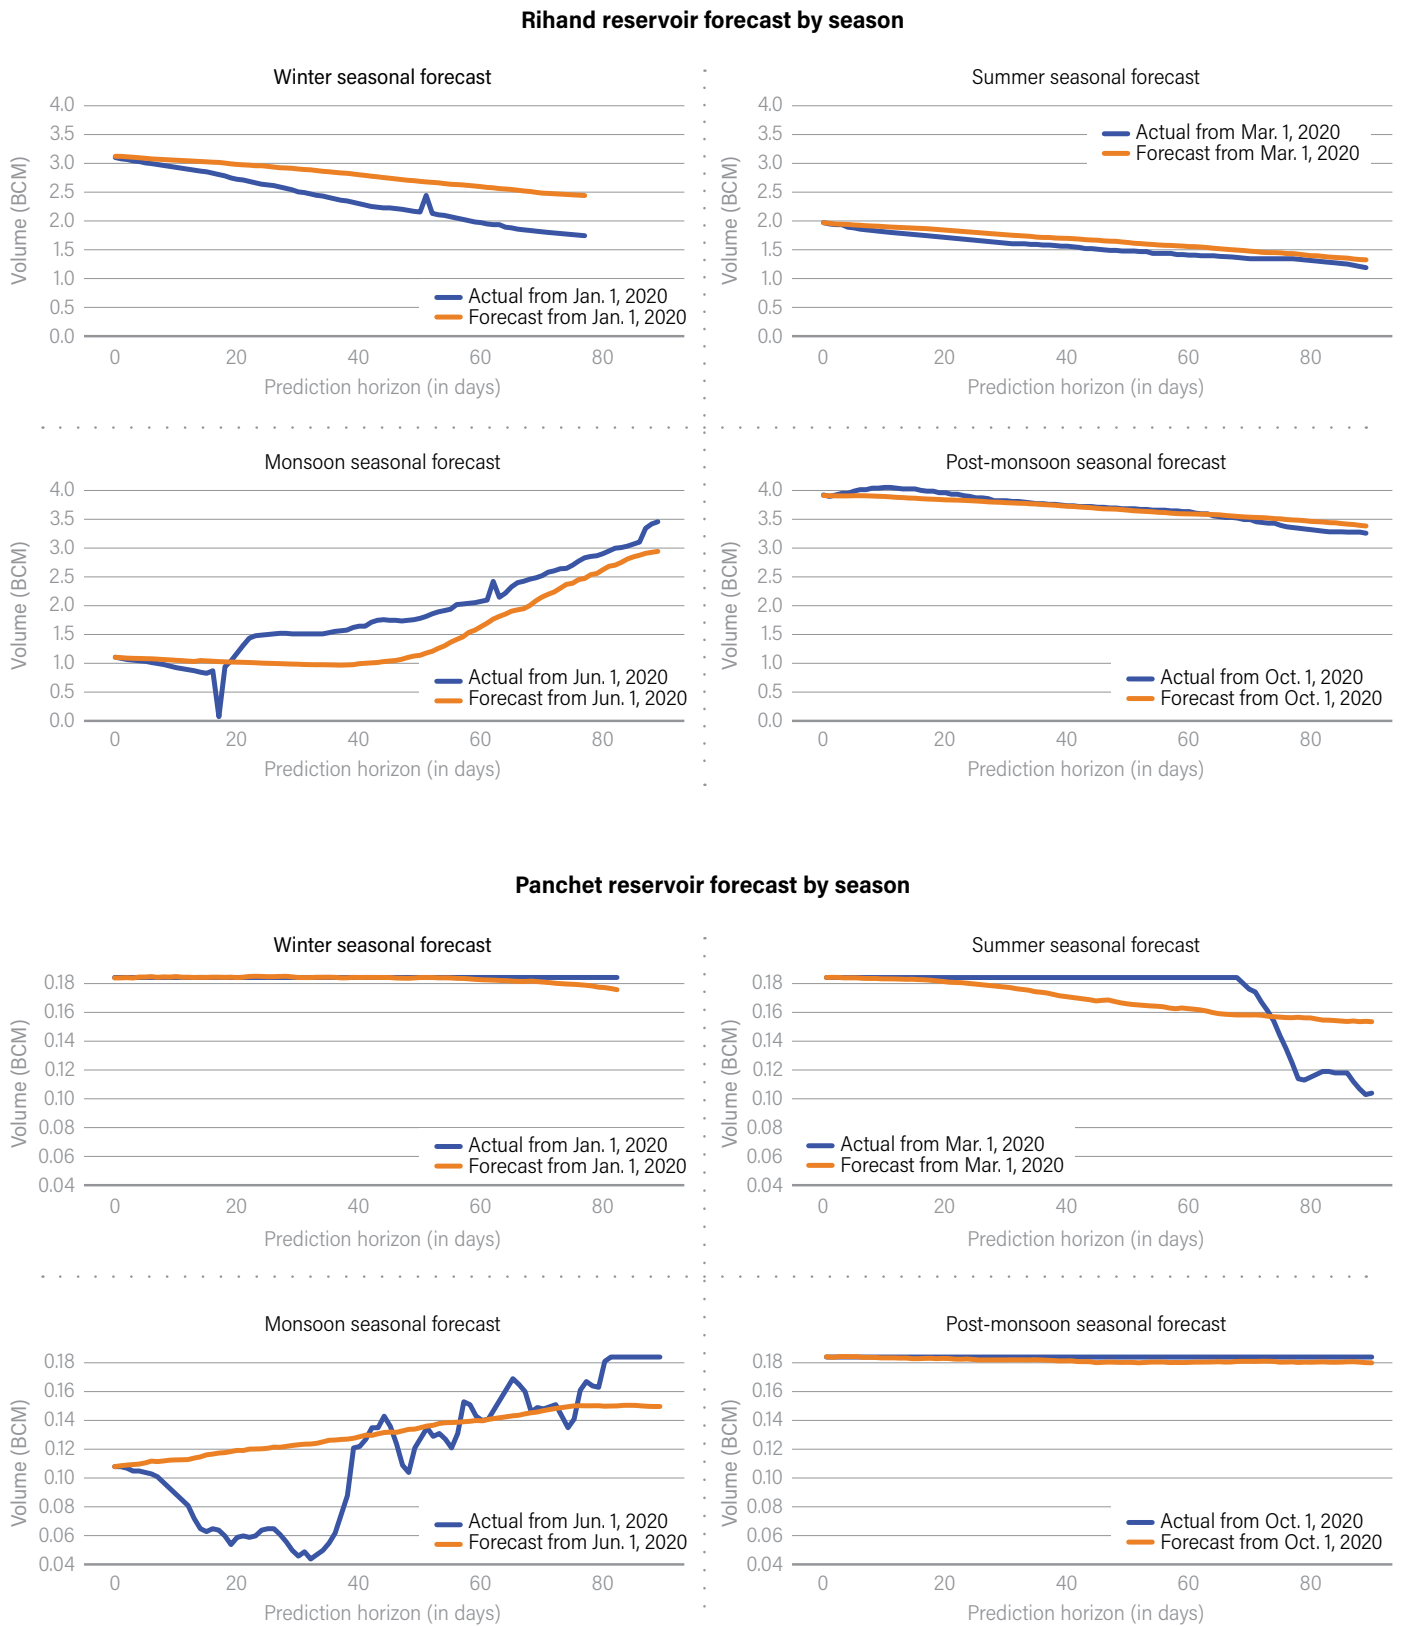
BASIN	PILOT RESERVOIR	SEASON			
		MONSOON (JUNE-SEPTEMBER)	POST-MONSOON (OCTOBER-DECEMBER)	WINTER (JANUARY-FEBRUARY)	SUMMER (MARCH-MAY)
Ganges	Rihand	0.78	-0.41	0.91	0.85
Ganga Damodar	Maithon	0.29	< -1*	0.88	0.79
	Panchet	-0.36	< -1*	< -1*	0.24
	Tenughat	-0.05	-0.17	0.95	0.76
	Tilaiya	0.83	-0.22	< -1*	-0.9
Godavari	Pench	0.67	-0.67	0.59	0.32
	Sriram Sagar	0.1	< -1*	0.48	< -1*
Kaveri	Mettur	-0.6	0	0.16	0.3
Mahanadi	Hirakud	0.83	-0.32	0.61	0.81
Narmada	Bargi	0.75	0.37	0.74	0.66
	Indira Sagar	0.83	0.53	0.66	0.82
Seasonal average		0.37	-0.11	0.66	0.47
Minimum		-0.6	-0.67	0.16	-0.9
Maximum		0.83	0.53	0.95	0.85

Note: The seasonal results are based on the valid date.^a A < -1* value means a horizontal line would have matched the actual volumes better. These reservoirs were dropped from the seasonal statistics at the bottom of the table.

a. Within our 90-day forecast, the valid date refers to the date of a forecasted value, not the date the forecast was made. For example, if a forecast was made on January 1, the valid date of the first forecasted day (i.e., first step) would be January 2; the valid date for the second step would be January 3, and so on.

Source: Authors.

Figure 9 | Seasonal forecasts for Rihand and Panchet



Note: BCM = billion cubic meters.

Source: Authors.

4 DATA SERVICES

4.1 Automated results

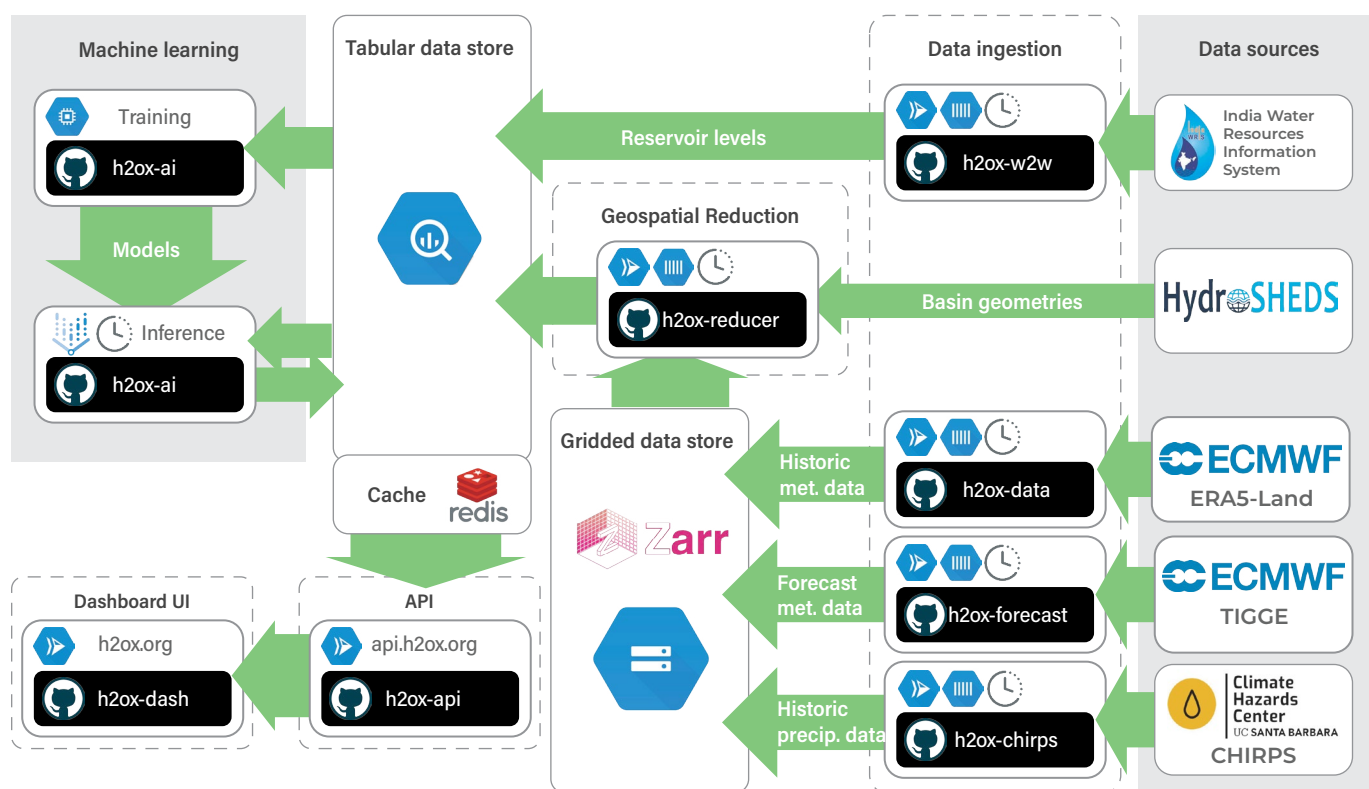
To make our baseline model actionable, its results need to be openly accessible and timely. Therefore, we have created a near-real-time computer system that automatically acquires new data, runs the model, and serves the results via an API so that the forecast can be easily integrated into existing dashboards. This data service is illustrated in Figure 10.

Essentially, the data service runs using cloud infrastructure. It searches for and downloads new versions of the input data (Table 2) as they become available (usually once every five to

seven days, when the new reservoir data are published). It then ingests, reduces, and normalizes the data according to our data processing methodology. The models are run for each reservoir to produce the forecasted change in water availability, which is then transformed back to water volume and stored. The data storage is served via an API—called the H2Ox API. This includes the 90-day forecast and the historic time series of water volumes, precipitation, and temperature.

Using this data service, we make our forecast available for viewing on the Water4Power dashboard, found on the Vasudha Power Info Hub. See Appendix A to learn more.

Figure 10 | Data service schematic



Note: Our system is deployed on the Google Cloud Platform. UI = user interface; API = application programming interface; met. = meteorological; precip. = precipitation; ECMWF = European Centre for Medium-Range Weather Forecasts; ERA5-Land = ECMWF ReAnalysis Version 5 Land; TIGGE = THORPEX Interactive Grand Global Ensemble; CHIRPS = Climate Hazards Group InfraRed Precipitation with Station data.

Source: Kruitwagen et al. 2022a.

4.2 Customizable model

While our data service provides an open-source forecast for all to use, we acknowledge that others may find it useful to expand or customize our underlying model. Therefore, to make our process entirely open source and transparent, we have turned it into a data science package on GitHub: <https://github.com/H2Oxford> (Kruitwagen et al. 2022a). We not only share our project code in this repository, but provide a way for users to replicate and even experiment with our method in Python. For example, we give instructions on how the user can easily edit the configuration file to train a model for just one reservoir, or to remove the graph convolutional layer. Stakeholders may also use this service to train models outside of our extent, add new data like relative humidity and sedimentation rates, or even insert new data sources such as local weather records.

5 LIMITATIONS

The most obvious limitation to our model is our inability to incorporate reservoir management decisions as an input, such as operating limit, interstate water transfers, recent damming projects, sedimentation rates, dredging, or authorized changes to water release practices. While we attempt to capture interbasin dynamics with our adjacency matrix, our model is unable to predict sharp changes in water levels due to human interventions. Therefore, our forecast should be used to flag when drought-like conditions threaten water supply; it should not be used to monitor human interventions, or as a tool to inform reservoir management operations.

Our pursuit of near-real-time production has limited the types of data evaluated in the model. Globally available Earth observation data were best suited for our data service infrastructure, which requires that data be delivered in a consistent, timely, and efficient format via an API. For example, we were able to process four datasets—historic and forecasted precipitation and historic and forecasted temperature—from one data format delivered by ECMWF. Global data also offer more flexibility for scaling the tool in the future. There may be other sources—such as local weather gauge data—that could produce a higher-quality forecast and improve local acceptance of the results but were less-suited for near-real-time infrastructure.

We are also limited by the length of our training data, which start in 2010 (when many of the CWC water-level time series start). One goal of our model is to forecast drought conditions that threaten energy production. While there were noticeable droughts that took place within our time range, such as the drought conditions in southern India over 2016–18 (Mishra et al. 2021b), ideally we could have trained our model on the most significant droughts over the last century to better capture extreme lows.

Finally, the timeline of this research prevented us from exhausting additional machine learning experiments. For example, we did not examine how selecting different years for our validation and test subsets impacted our results.

6 NEXT STEPS

Though this paper focused only on the Water4Power pilot study, our data service is producing forecasts of water volumes for 66 reservoirs in India (see Kruitwagen et al. 2022b). We hope to work with other stakeholders, such as city officials, to implement similar dashboards on their websites. These expansions need not be limited to the power sector: Our data can deliver insights into a variety of other sectors, including urban water supply, food security and agriculture, and long-term resilience building through risk reduction. To the best of our knowledge through conversations with our expert technical advisory group panel, there are no other comparable forecasts being used by the government, despite the need for actionable water information.

We also believe this methodology has the potential to be scaled into a global product, given we identify a suitable global reservoir dataset, likely in the form of a satellite-derived product (the only data in our catalog that are not global come from our current reservoir water-level dataset from the CWC). New datasets, such as the reservoir surface water extents from Global Water Watch, could be used in lieu of the CWC reservoir data to scale our forecasts to anywhere in the world (Deltares et al. 2022). Such a dataset could be used to provide early warnings for disaster risk reduction, to prioritize nature-based solutions around reservoirs, or even by cities to inform resilience and adaptation policies, among other ideas.

APPENDIX A. WATER4POWER DASHBOARD

The mission of the Water4Power beta dashboard (<https://vasudhapower.in/analytics/generation/water>) is to alert stakeholders of potential water shortages that threaten power production, and to advocate for more water-prudent renewable power production to prevent future water scarcity-driven electricity outages. The current design of the dashboard was borne out of a yearlong engagement with a technical advisory group (TAG). The TAG—a group of external academics, practitioners, and public servants who are experts in the fields of water and energy—was asked to review and provide feedback on our process from both technical and user perspectives (see Appendix C for more details). The beta dashboard is hosted on the Vasudha Power Portal.

The data featured on the dashboard (summarized in Table A1) are organized into four main sections: discovery map, reservoir time series, water volume, and the water-energy nexus.

Discovery map

The user can view the pilot reservoirs on an interactive map. They can explore the hydrologic connections of these reservoirs and understand their water stress conditions. Here, the user must select a reservoir of interest to trigger the rest of the dashboard. Once a reservoir is selected, the map will classify the current volume using a traffic light alert.

Reservoir time series and forecast

First, the user can see an extended time series of the reservoir's volume. This graphic shows the past year of reservoir volumes followed by the 90-day forecast. It also shows the past year of precipitation and the dead zone line for additional analysis.

Precipitation time series

The user can view the Annual Cumulative Precipitation graphic (accumulation restarts January 1 of every year). They can compare the current precipitation conditions to a 10-year average baseline, which has been calculated by averaging daily precipitation from 2010 to 2020 to see if there has been a recent meteorological drought affecting the reservoir.

Water volume

Next, the user can view how full the reservoir currently is in the Bathtub graphic. They can compare the current storage—expressed as current volume relative to full (gross) volume—against a 10-year average baseline. The baseline represents the average daily volume from 2010 to 2020 relative to the full volume. In addition, the user can compare the current storage to the dead storage—the volume in which no more water can be released from the reservoir.

Power plant water use

The final section focuses on the water-energy nexus. We estimate the approximate water required to meet the generation in every thermal power plant dependent on the reservoir. For each plant, we list the generation (in million units, MU) and the water efficiency rate (volume of water in cubic meters to produce one megawatt-hour of energy). These two values are multiplied (after the units are converted) to produce the approximate water used in million cubic meters (MCM). Cumulative power generation as collated by Vasudha Power Info Hub is used here, thus water use is estimated in a cumulative manner for the calendar year.

APPENDIX B. DASHBOARD SOURCES

Table B1 | **Dashboard data providers**

DATA POINT	RESERVOIRS	PROVIDER OR SOURCE
Baseline water stress	All	WRI Aqueduct (Hofste et al. 2019b)
Historic, current, and baseline accumulated precipitation	All	H2Ox API
Historic, current, and baseline reservoir volume	All	H2Ox API
Reservoir volume forecast	All	H2Ox API
Daily power outages related to water shortages	All	Vasudha API
Dependent power plant	All	Vasudha API
Potential generation (MU)	All	Vasudha API
Compliance status	All	CSE (Yadav and Arora 2021)
OTC-CT	All	CSE (Yadav and Arora 2021)
Capacity in MW	All	CSE (Yadav and Arora 2021)
Water efficiency (MCM/MW)	All	CSE (Yadav and Arora 2021)
Gross storage volume, dead zone volume	Bargi	CWC n.d.
	Hirakud	DOWR Odisha n.d.
	Indira Sagar	CWC n.d.
	Maithon	CWC n.d.; DVC n.d.
	Mettur	CWC n.d.
	Panchet	CWC n.d.; DVC n.d.
	Pench	CWC n.d.
	Rihand	UPJVN n.d.
	Sriram Sagar	CWC n.d.
	Tenughat	CWC n.d.
	Tilaiya	CWC n.d.; DVC n.d.

Note: API = application programming interface; MU = million units; CSE = Centre for Science and Environment; OTC-CT = once-through cooling versus cooling tower; MCM = million cubic meters; MW = megawatt; CWC = India's Central Water Commission; H2Ox = winning hackathon team.

Source: Authors.

APPENDIX C. TECHNICAL ADVISORY GROUP

A technical advisory group—made of up external academics, practitioners, and public servants who are experts in the fields of water and energy—was convened to guide the research process and ensure our work had real-world application. The TAG met twice to review the research and data and provide guidance on next steps. Membership included the following individuals:

- Mr. Shashi Shekhar, Former Secretary, Ministry of Water Resources, Government of India
- Mr. Srinivas Krishnaswamy, CEO, Vasudha Foundation
- Mr. Ashish Fernandes, CEO, Climate Risk Horizons
- Dr. Casey Brown, Professor, University of Massachusetts Amherst
- Mr. Shripad Dharmadhikary, Policy Researcher, Manthan Adhyayan
- Dr. Sukanya Randhawa, Auroville Consulting

ABBREVIATIONS

API	application programming interface	KL	Kullback-Leibler divergence term
BCM	billion cubic meters	LSTM	long short-term memory model
BLSTM	Bayesian long short-term memory model	MCM	million cubic meters
CHIRPS	Climate Hazards Group InfraRed Precipitation with Station data	MSE	mean squared error
CWC	Central Water Commission, Government of India	MW	megawatt
ECMWF	European Centre for Medium-Range Weather Forecasts	R²	coefficient of determination
ERA5-Land	ECMWF Re-Analysis Version 5 Land	seq2seq(2seq)	sequence-to-sequence(-to-sequence) model
H2Ox	Name of the winning hackathon team	TAG	technical advisory group
		TIGGE	THORPEX Interactive Grand Global Ensemble

ENDNOTES

1. Three percent is based on 23.8 billion cubic meters per year (BCM/year) in withdrawals in 2019 for coal power production (de Oliveira Bredariol et al. 2021) compared with 761 BCM/year in total withdrawals for India (FAO 2020).
2. Twenty percent is based on an estimated 40,520 megawatts (MW) of installed capacity for all power plants that source water from the 11 pilot reservoirs. The list of reservoirs as well as installed capacity came from Vasudha's online portal. This value was compared to the Central Electricity Authority's total installed capacity from coal found in Figure 1.
3. We only included adjacent reservoirs that were actively monitored by CWC in 2020.
4. Near-real-time access is coming with the release of ERA5-LandT (ECMWF 2020).
5. The reservoir adjacency matrix defined the connections among basin reservoirs: downstream and diagonal connections were assigned a positive value, and upstream a negative.
6. The graph convolutional layer was composed of three headers: two linear layers, each with bias and dropout, and graph convolution layer in the middle. Each reservoir was assigned its own header. The forecast and the predictor stages used the same headers so that data could be traversed between the two.

REFERENCES

- AP News. 2023. "UN Says India to Be World's Most Populous Nation by Mid-2023," April 19. Sec. Politics. <https://apnews.com/article/india-worlds-largest-population-china-un-47bcd360aba-6b7e3ecc4020702fa685f>.
- Blundell, C., J. Cornebise, K. Kavukcuoglu, and D. Wierstra. 2015. "Weight Uncertainty in Neural Networks." In *Proceedings of the 32nd International Conference on Machine Learning 37*: 1613–22. ICML'15. Lille, France: JMLR.org.
- CEA (Central Electricity Authority, Government of India). 2022. "All India Installed Capacity (in MW) of Power Stations (as of 30.09.2022)." In *Installed Capacity*, September. New Delhi: Government of India, Ministry of Power. <https://cea.nic.in/installed-capacity-report/?lang=en>.
- CWC (Central Water Commission, Government of India). n.d. "Live Storage Capacities of Reservoirs as per Data of: Large Dams/Reservoir Projects." New Delhi: Central Water Commission, Government of India. <http://cwc.gov.in/sites/default/files/admin/updation-lsc.pdf>. Accessed .January 2021.
- Deltares, WRI, and WWF. 2022. "Global Water Watch: Democratizing Earth's Data on Water Resources." Global Water Watch. <https://www.globalwaterwatch.io/>.
- de Oliveira Bredariol, T., M. Walton, and V. Chaturvedi. 2021. "Managing the Water-Energy Nexus Is Vital to India's Future—Analysis." International Energy Agency (blog), March 22. <https://www.iea.org/commentaries/managing-the-water-energy-nexus-is-vital-to-india-s-future>.
- DOWR Odisha (Department of Water Resources, Odisha, India). n.d. "Hirakud Dam Project." Odisha, India: Department of Water Resources. [www.dowrodisha.gov.in/Projects/MajMed/DetailsCompletedOngoing/Ongoing/CEUMB/Hirakud%20System%20\(Completed\)/SFHirakud.pdf](http://www.dowrodisha.gov.in/Projects/MajMed/DetailsCompletedOngoing/Ongoing/CEUMB/Hirakud%20System%20(Completed)/SFHirakud.pdf). Accessed July 2022.
- DVC (Damodar Valley Corporation). n.d. "Reservoir Overview & Salient Features." Kolkata, India: Damodar Valley Corporation. https://www.dvc.gov.in/dvcwebsite_new1/dams-barrages/. Accessed July 14, 2022.
- ECMWF (European Centre for Medium-Range Weather Forecasts). 2020. "ERA5-Land Timely Updates (ERA5-LandT) Release Date -." Copernicus User Support Forum, ECMWF Confluence Wiki. <https://confluence.ecmwf.int/pages/viewpage.action?pageId=185079782>.
- FAO (Food and Agriculture Organization of the United Nations). 2020. "AQUASTAT Core Database." Rome: Food and Agriculture Organization of the United Nations.

- Funk, C., P. Peterson, M. Landsfeld, D. Pedreros, J. Verdin, S. Shukla, G. Husak, et al. 2015. "The Climate Hazards Infrared Precipitation with Stations—a New Environmental Record for Monitoring Extremes." *Scientific Data* 2 (1): 150066. <https://doi.org/10.1038/sdata.2015.66>.
- Ghosh, N., and J. Bandyopadhyay. 2009. "A Scarcity Value Based Explanation of Trans-boundary Water Disputes: The Case of the Cauvery Basin in India." *Water Policy* 11 (May): 141–67. <https://doi.org/10.2166/wp.2009.017>.
- Hochreiter, S., and J. Schmidhuber. 1997. "Long Short-Term Memory." *Neural Computation* 9 (8): 1735–80. <https://doi.org/10.1162/neco.1997.9.8.1735>.
- Hofste, R.W., P. Reig, and L. Schleifer. 2019a. "17 Countries, Home to One-Quarter of the World's Population, Face Extremely High Water Stress." *Insights* (blog), August. Washington, DC: World Resources Institute. <https://www.wri.org/insights/17-countries-home-one-quarter-worlds-population-face-extremely-high-water-stress>.
- Hofste, R.W., S. Kuzma, S. Walker, E.H. Sutanudjaja, M.F.P. Bierkens, M.J.M. Kuijper, M. Faneca Sanchez, et al. 2019b. "Aqueduct 3.0: Updated Decision-Relevant Global Water Risk Indicators." Technical Note. Washington, DC: World Resources Institute. <https://www.wri.org/research/aqueduct-30-updated-decision-relevant-global-water-risk-indicators>.
- India-WRIS (India Water Resources Information System). 2008. "India-WRIS: Reservoir Level." New Delhi: India Water Resources Information System. <https://indiawris.gov.in/wris/#/Reservoirs>.
- India-WRIS. 2021. "Welcome to Water Resources Information System—Wiki." India-WRIS Wiki. <https://indiawris.gov.in/wiki/doku.php>.
- Kruitwagen, L., C. Arderne, T. Lees, and L. Thalheimer. 2021. "H2Ox." H2Ox. <https://www.h2ox.org/>.
- Kruitwagen, L., C. Arderne, T. Lees, and L. Thalheimer. 2022a. "Wave2Web: H2Ox." GitHub. <https://github.com/H2Oxford>.
- Kruitwagen, L., C. Arderne, T. Lees, L. Thalheimer, S. Kuzma, and S. Basak. 2022b. "Wave2Web: Near-Real-Time Reservoir Availability Prediction for Water Security in India," June. Earth ArXiv. <https://eartharxiv.org/repository/view/3381/>.
- Lee, S., and J. Kim. 2021. "Predicting Inflow Rate of the Soyang River Dam Using Deep Learning Techniques." *Water* 13 (17): 2447. <https://doi.org/10.3390/w13172447>.
- Lehner, B., and G. Grill. 2014. "Global River Hydrography and Network Routing: Baseline Data and New Approaches to Study the World's Large River Systems." *Hydrological Processes* 27 (15): 2171–86.
- Lehner, B., C. Reidy Liermann, C. Revenga, C. Vörösmarty, B. Fekete, P. Crouzet, P. Döll, et al. 2011. "High-Resolution Mapping of the World's Reservoirs and Dams for Sustainable River-Flow Management." *Frontiers in Ecology and the Environment* 9 (9): 494–502. <https://doi.org/10.1890/100125>.
- Luo, T., D. Sriram Krishnan, and S. Sen. 2018. "Parched Power: Water Demands, Risks, and Opportunities for India's Power Sector." Working Paper. Washington, DC: World Resources Institute. <https://www.wri.org/parched-power-water-demands-risks-and-opportunities-indias-power-sector>.
- Mishra, V. 2020. "Long-Term (1870–2018) Drought Reconstruction in Context of Surface Water Security in India." *Journal of Hydrology* 580 (January): 124228. <https://doi.org/10.1016/j.jhydrol.2019.124228>.
- Mishra, V., S. Aadhar, and S. Shwarup Mahto. 2021a. "Anthropogenic Warming and Intraseasonal Summer Monsoon Variability Amplify the Risk of Future Flash Droughts in India." *NPJ Climate and Atmospheric Science* 4 (1): 1–10. <https://doi.org/10.1038/s41612-020-00158-3>.
- Mishra, V., K. Thirumalai, S. Jain, and S. Aadhar. 2021b. "Unprecedented Drought in South India and Recent Water Scarcity." *Environmental Research Letters* 16 (5): 054007. <https://doi.org/10.1088/1748-9326/abf289>.
- Muñoz Sabater, J. 2019. "ERA5-Land Hourly Data from 1981 to Present." Copernicus Climate Change Service, Climate Data Store. <https://doi.org/10.24381/cds.e2161bac>.
- PIB Delhi (Press Information Bureau, Government of India). 2022. "Cabinet Approves India's Updated Nationally Determined Contribution to Be Communicated to the United Nations Framework Convention on Climate Change." Press Release, August 3. New Delhi: Press Information Bureau, Government of India. <https://pib.gov.in/pib.gov.in/Pressreleaseshare.aspx?PRID=1847812>.
- Prakash, S. 2019. "Performance Assessment of CHIRPS, MSWEP, SM2RAIN-CCI, and TMPA Precipitation Products across India." *Journal of Hydrology* 571 (April): 50–59. <https://doi.org/10.1016/j.jhydrol.2019.01.036>.
- Saha, S., D. Roy, and A. Mazumdar. 2017. "Performance of a System of Reservoirs on Futuristic Front." *Applied Water Science* 7 (6): 2667–80. <https://doi.org/10.1007/s13201-016-0484-2>.
- Sen, S. 2021. "Integrated Basin Development of Damodar Valley: DVC and Lower Damodar." Conference Paper. ICOLD Symposium on Sustainable Water Development and River Basin and APG Symposium on Water and Dams. New Delhi. https://www.researchgate.net/publication/349683892_Integrated_Basln_development_of_damodar_valley_-dvc_and_lower_damodar.

- Srinivasan, S., N. Kholod, V. Chaturvedi, P. Pratap Ghosh, R. Mathur, L. Clarke, M. Evans, et al. 2018. "Water for Electricity in India: A Multi-model Study of Future Challenges and Linkages to Climate Change Mitigation." *Applied Energy* 210 (January): 673–84. <https://doi.org/10.1016/j.apenergy.2017.04.079>.
- Sutskever, I., O. Vinyals, and Q.V. Le. 2014. "Sequence to Sequence Learning with Neural Networks." Ithaca, NY: arXiv, Cornell University. <https://doi.org/10.48550/arXiv.1409.3215>.
- Swinbank, R., M. Kyouda, P. Buchanan, L. Froude, T.M. Hamill, T.D. Hewson, J.H. Keller, et al. 2016. "The TIGGE Project and Its Achievements." *Bulletin of the American Meteorological Society* 97 (1): 49–67. <https://doi.org/10.1175/BAMS-D-13-00191.1>.
- UNDRR (United Nations Office for Disaster Risk Reduction). 2021. *GAR Special Report on Drought 2021*. Geneva: United Nations Office for Disaster Risk Reduction. <https://www.undrr.org/publication/gar-special-report-drought-2021>.
- UPJVN (Uttar Pradesh Jal Vidyut Nigam Limited Company). n.d. "Rihand." Uttar Pradesh Jal Vidyut Nigam Limited. <http://www.upjvn.org/Home/Rihand>. Accessed July 14, 2022.
- Viswanath, S., M. Saha, P. Mitra, and R.S. Nanjundiah. 2019. "Deep Learning Based LSTM and SeqtoSeq Models to Detect Monsoon Spells of India." In *Computational Science—ICCS 2019*, edited by J.M.F. Rodrigues, P.J.S. Cardoso, J. Monteiro, R. Lam, V.V. Krzhizhanovskaya, M.H. Lees, J.J. Dongarra, and P.M.A. Sloot. Lecture Notes in Computer Science. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-030-22741-8_15, 204–18.
- VPIH (Vasudha Power Info Hub). 2023. "Vasudha-Energy Mapping Initiative of Vasudha Foundation." Vasudha Power Info Hub. www.vasudha-emi.in.
- WRI India (World Resources Institute India). 2020. "Water4Power," November 21. World Resources Institute India. <https://water4power.wri-india.org/>.
- WRI India. 2021. "WAVE2WEB Hack," March 18. New Delhi: WRI India Ross Center for Sustainable Cities. <https://www.wricitiesindia.org/content/wave2web-hack>.
- Xiang, Z., J. Yan, and I. Demir. 2020. "A Rainfall-Runoff Model with LSTM-Based Sequence-to-Sequence Learning." *Water Resources Research* 56 (1): e2019WR025326. <https://doi.org/10.1029/2019WR025326>.
- Yadav, N.K., and S. Arora. 2021. "Water Inefficient Power: Implementing Water Norms and Zero Discharge in India's Coal-Power Fleet." New Delhi: Centre for Science and Environment. <https://www.cseindia.org/water-inefficient-power-10856>.

ACKNOWLEDGMENTS

We are pleased to acknowledge our institutional strategic partners that provide core funding to WRI: the Netherlands Ministry of Foreign Affairs, Royal Danish Ministry of Foreign Affairs, and Swedish International Development Cooperation Agency.

This publication was made possible thanks to the support of the Growald Family Fund, SED Foundation, Microsoft, and BlackRock.

The authors would like to thank the following people for providing invaluable insight and assistance in the development of this paper: Charles Iceland (WRI), Raj Bhagat (WRI India), Logan Byers (WRI), Daniel Gajardo (Stanford University), Bharath Jairaj (WRI India), Deepak Krishnan (WRI India), Renata Marson Teixeira de Andrade (WRI), Renee Pineda (WRI), and Gregory Taff (WRI). We're also grateful to our internal reviewers, John Feldmann, Elizabeth Saccoccia, Bina Shetty, and Thet Hein Tun, and external reviewers, Gennadii Donchyts (Google), Peter Kerins (Impact Observatory), and Hannah Kerner (University of Maryland).

The authors would also like to thank our partners at H2Ox and Vasudha Foundation as well as Vistar Media and Romain Warnault for their extensive guidance and feedback during the design and development of this study.

Lucas Kruitwagen (principal investigator), Chris Arderne, Thomas Lees, and Elizabeth Thalheimer of H2Ox designed and built the machine learning model, data service, and data science package. Samrat Basak (principal investigator), Samantha Kuzma, Sahana Goswami, and Anupriya Goyal of WRI India and WRI designed and managed the project, deliverables, and hackathon. They also designed the final dashboard in partnership with Vasudha.

Opinions or points of view expressed in this report are those of the authors and do not necessarily reflect the position of the reviewers, research partners, or the organizations they represent.

Maps are for illustrative purposes and do not imply the expression of any opinion on the part of WRI, concerning the legal status of any country or territory or concerning the delimitation of frontiers or boundaries.

ABOUT THE AUTHORS

Samantha Kuzma is the data lead and senior manager for Global Water Data Products at World Resources Institute.

Contact: samantha.kuzma@gmail.com

Dr. Lucas Kruitwagen is the co-founder of Oxford Earth Observation, a water risk intelligence start-up, and is a postdoctoral research associate at the University of Oxford.

Christopher Arderne is a consulting data scientist specializing in climate-related roles.

Sahana Goswami is senior program manager of urban water resilience at World Resources Institute India.

Anupriya Goyal is the program manager of urban water resilience at World Resources Institute India.

Dr. Thomas Lees completed his PhD at the University of Oxford looking at the intersection of hydrological science and deep learning. He now works as a renewable power analyst for a commodities team.

Dr. Lisa Thalheimer is a senior researcher at the Institute for Environment and Human Security, United Nations University. Contact: thalheimer@ehs.unu.edu

Samrat Basak is the program director of urban water resilience at World Resources Institute India.

ABOUT WRI

World Resources Institute is a global research organization that turns big ideas into action at the nexus of environment, economic opportunity, and human well-being.

OUR CHALLENGE

Natural resources are at the foundation of economic opportunity and human well-being. But today, we are depleting Earth's resources at rates that are not sustainable, endangering economies and people's lives. People depend on clean water, fertile land, healthy forests, and a stable climate. Livable cities and clean energy are essential for a sustainable planet. We must address these urgent, global challenges this decade.

OUR VISION

We envision an equitable and prosperous planet driven by the wise management of natural resources. We aspire to create a world where the actions of government, business, and communities combine to eliminate poverty and sustain the natural environment for all people.



Copyright 2023 World Resources Institute. This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of the license, visit <http://creativecommons.org/licenses/by/4.0/>