

# *Building and explaining data-driven energy demand models for Indian states*

Article

Published Version

Creative Commons: Attribution 4.0 (CC-BY)

Open Access

Hunt, K. M. R. ORCID: <https://orcid.org/0000-0003-1480-3755>  
and Bloomfield, H. C. ORCID: <https://orcid.org/0000-0002-5616-1503> (2025) Building and explaining data-driven energy demand models for Indian states. *Environmental Research: Energy*, 2 (2). 025003. ISSN 2753-3751 doi: 10.1088/2753-3751/adc7bc Available at <https://centaur.reading.ac.uk/122151/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: <http://dx.doi.org/10.1088/2753-3751/adc7bc>

Publisher: IOP Publishing

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

[www.reading.ac.uk/centaur](http://www.reading.ac.uk/centaur)

**CentAUR**

Central Archive at the University of Reading

Reading's research outputs online

PAPER • OPEN ACCESS

## Building and explaining data-driven energy demand models for Indian states

To cite this article: Kieran M R Hunt and Hannah C Bloomfield 2025 *Environ. Res.: Energy* **2** 025003

View the [article online](#) for updates and enhancements.

You may also like

- [India's National Green Hydrogen Mission: an analysis of the strategies, policies for net-zero emissions and sustainability](#)  
Balasubramanian Sambasivam and Rakesh Narayana Sarma
- [Quantifying the impact of energy system model resolution on siting, cost, reliability, and emissions for electricity generation](#)  
Anna F Jacobson, Denise L Mauzerall and Jesse D Jenkins
- [Household energy choices under fuel stacking scenarios: evidence for bundling welfare schemes for facilitating clean fuel use](#)  
M Manjula

# UNITED THROUGH SCIENCE & TECHNOLOGY



The Electrochemical Society  
Advancing solid state & electrochemical science & technology

## 248th ECS Meeting Chicago, IL October 12-16, 2025 *Hilton Chicago*



## Science + Technology + YOU!

Register by  
September 22  
to **save \$\$**

**REGISTER NOW**

# ENVIRONMENTAL RESEARCH ENERGY



## PAPER

### OPEN ACCESS

RECEIVED  
17 January 2025

REVISED  
25 March 2025

ACCEPTED FOR PUBLICATION  
1 April 2025

PUBLISHED  
9 April 2025

Original Content from  
this work may be used  
under the terms of the  
[Creative Commons  
Attribution 4.0 licence](#).

Any further distribution  
of this work must  
maintain attribution to  
the author(s) and the title  
of the work, journal  
citation and DOI.



## Building and explaining data-driven energy demand models for Indian states

Kieran M R Hunt<sup>1,2,\*</sup> and Hannah C Bloomfield<sup>3</sup>

<sup>1</sup> Department of Meteorology, University of Reading, Reading, United Kingdom

<sup>2</sup> National Centre for Atmospheric Sciences, University of Reading, Reading, United Kingdom

<sup>3</sup> School of Engineering, Newcastle University, Newcastle, United Kingdom

\* Author to whom any correspondence should be addressed.

E-mail: [k.m.r.hunt@reading.ac.uk](mailto:k.m.r.hunt@reading.ac.uk)

**Keywords:** India, renewable energy, energy demand, machine learning, meteorology

### Abstract

Accurate forecasts of energy demand are crucial for managing India's rapidly growing energy needs as it continues to decarbonise its grid. In this study, we develop state-level data driven models to predict weather-driven energy demand across India using the eXtreme Gradient Boosting framework. The models use as input population-weighted meteorological variables averaged over various timescales. The models are trained on daily energy demand data, scraped from reports issued by Grid-India, which we correct for trends in population and economic growth. The models demonstrate high skill, with half having  $r^2 > 0.8$ , significantly outperforming traditional multivariate linear regression models. We explain model behaviour through Shapley analysis and find a strong sensitivity to day of the week and public holidays, with reductions in energy demand on Sundays and varying impacts during holidays. While important variables vary by state and season, daily minimum temperature and 30 d mean temperature consistently emerge as key predictors, reflecting nighttime air conditioning use and seasonal heating or cooling needs. We also identify threshold behaviours, indicating large increases in energy demand once temperatures pass certain values. Using reanalysis, we extend our models to estimate all-India energy demand from 1979–2023, calibrated to 2023 conditions. We confirm a pronounced seasonal cycle, with greatest demand during the pre-monsoon and monsoon onset (May–June) and lowest demand in the winter (November–December). Combining these results with timeseries of renewable energy production, we find the largest energy deficit (demand minus renewable generation) occurs during or after monsoon withdrawal (September–October). Extreme deficit days, posing a risk to the national grid, are associated with early monsoon withdrawal or late monsoon breaks, leading to low wind speeds and persistently high dewpoint temperatures and cloud cover. The demand dataset created here can be used for energy grid management, siting of future renewable energy generation, and to aid with ensuring security of supply.

## 1. Introduction

### 1.1. The Indian energy system

India's energy system in 2025 is navigating a complex transition toward sustainability while meeting the growing demands of its population. As the third-largest global energy consumer, India is rapidly scaling up renewables to shift away from coal and natural gas (which still dominates the electricity mix). To address this, the government has targeted reaching net-zero emissions by 2070, with ambitious and rapid growth in existing installed renewable capacity to 500 GW of installed capacity by 2030, and 7400 GW needed by 2070 (Chaturvedi and Malyan 2021).

The country faces several key challenges in this transition. Although renewables are expanding, fossil fuels remain necessary to meet increasing energy demands (with peak demand expected to be 340 GW in about 2030) and maintain energy security. Efforts are also underway to boost electric vehicle adoption, but

progress is slow due to the limited charging infrastructure and high upfront costs. Additionally, energy security efforts include addressing dependencies on imported critical minerals required for renewable technology production, such as lithium and cobalt (Christensen *et al* 2021, Sun *et al* 2024).

A further key challenge to address will be the increased weather-dependence present in the Indian energy system with both the increasing level of electricity demand and renewable energy supply. Particularly challenging times will be those of high demand and low renewable generation (Bloomfield *et al* 2016), often called *energy shortfall events* (van der Wiel *et al* 2019). Due to the rapidly developing nature of the Indian energy system, historical records are a poor proxy for future and synthetic multi-decadal reconstructions of present and future energy demand are key to understand the expected frequency and severity of renewable energy droughts (Raynaud *et al* 2018, Kapica *et al* 2024). To do this statistical electricity demand models, such as the one developed in this study, are required.

### 1.2. The weather and climate of India

To understand the outputs of the demand models created in this study some background on the large-scale weather and climate of India is useful. The **pre-monsoon** occurs between March and May and is usually associated with very hot and relatively cloud free conditions across most of India. This is the season in which heatwaves predominantly occur (Chaudhury *et al* 2000, Ratnam *et al* 2016, Rohini *et al* 2016, Satyanarayana and Rao 2020) and so demand for air conditioning is typically the largest.

During June, the onset of the **summer monsoon** occurs, moving northwards until India is covered by early- or mid-July (Fasullo and Webster 2003, Wang *et al* 2009, Parker *et al* 2016). The summer monsoon peaks in July and August and is characterised by cloudy, humid, and very windy conditions, with lower dry bulb temperatures than the pre-monsoon. The monsoon is subject to strong intraseasonal variability, driven by the Boreal Summer Intraseasonal Oscillation (Kikuchi *et al* 2012, Lee *et al* 2013) and low-pressure systems (Hunt and Fletcher 2019, Thomas *et al* 2021). These lead to active and break periods in the monsoon, respectively characterised by cloudier, windier conditions and hotter, drier conditions across most of India (Annamalai and Slingo 2001, Pai *et al* 2016, Kulkarni *et al* 2018).

Following the withdrawal of the summer monsoon, the subcontinent experiences **post-monsoon** conditions, which persist through October and November until the arrival of the upper-level subtropical jet over northern India, which marks the start of winter. North Indian Ocean tropical cyclones are most common during this season (Mohapatra *et al* 2012). The other important feature of this season is the so-called 'northeast monsoon', where northeasterly winds bring heavy precipitation to southeast India (Dhar and Rakhecha 1983, Rajeevan *et al* 2012). This is followed by the **winter** season (December–March), which is characterised by dry conditions across southern India, while winter storms—known as western disturbances—impact north India, particularly along the Himalayas, bringing large-scale cloud cover and precipitation (Dimri *et al* 2015, Hunt *et al* 2025).

### 1.3. Modelling weather-dependent electricity demand

National electricity demand follows a number of regular and (to some extent) predictable patterns, such as seasonal cycles, weekly cycles and a diurnal cycle. It has been well documented how these behaviours relate to weather and climate.

The seasonal cycle of demand is predominantly dependent on temperature, with colder countries having a winter-peak demand (when they have their largest heating requirements), and warmer countries having a summer-peak demand (when demand for air-conditioning peaks). The statistical relationship between demand and temperature has a characteristic U-shaped response (Baker and Blundell 1985, Sailor and Pavlova 2003). Each arm can be characterised using either near-surface temperatures, heating degree days and cooling degree days (Bloomfield *et al* 2020, Staffell *et al* 2023) (where the thresholds used to define relative heating or cooling needs will vary by country) or composite weather variables which include lagged effects, effects of wind chill or lighting demands (Rowley 2016). Temperature data is often weighted by gridded population density to reflect the centres of demand in large countries (Sailor and Muñoz 1997). In warmer countries (e.g. tropical regions) the demand is impacted by local relative humidity (Parkpoom and Harrison 2008, Liu *et al* 2021). Precipitation rate can also be an important variable in countries experiencing a rainy season (Fan *et al* 2019).

As well as the weather-dependence, there are human-influences on demand such as the day of the week, national holidays, or long-term trends (e.g. changes in gross domestic product, improvements in energy efficiency, or more recently the example of the COVID-19 pandemic). The diurnal cycle of demand is predominantly controlled by human behaviour with many countries having a peak daily demand between 4 pm and 9 pm when households have returned from work and require energy for cooking, heating, lighting and entertainment (Deakin *et al* 2021).



Over the last decade, there has been a considerable rise in the number of demand studies using data-driven approaches. The vast majority of these have focused on modelling energy demand for single buildings or drawing on techniques from economics to model energy market behaviour (Seyedzadeh *et al* 2018, Antonopoulos *et al* 2020). Studies focusing on single buildings tend to focus on hour-ahead or day-ahead forecasts, and usually rely on autocorrelation or other timeseries techniques (Seyedzadeh *et al* 2018). Recent examples include using a long short-term memory network (LSTM) for forecasting single-household consumption in the UK (Yan *et al* 2019), combining convolutional layers with an LSTM to forecast energy consumption in a university building in India (Somu *et al* 2021), and using an ensemble of feedforward and eXtreme Gradient Boosting (XGBoost) models for day-ahead forecasts of energy use for an individual household in France (Harikrishnan *et al* 2025). Such models often use temperature and humidity data as predictors.

Of greater relevance to this study, however, are previous efforts to build demand models for significant populations, such as cities or countries. Regression models have proven popular over Europe given both the high population density and strong relationship between temperature and demand (Taylor and McSharry 2007, Bessec and Fouquau 2008, Bloomfield *et al* 2016). Staffell *et al* (2023) built a model for estimating energy demand across spatial scales (e.g. building to national) using as inputs heating- and cooling-degree days as well as gridded meteorological and population data. Their model comprised several modules to handle issues such as time lags and hot-weather correction, and performed well at the national level in the global south, including over India. Autoregression models, such as ARIMA, have also proven popular, and have been applied to short-term forecasts of energy demand for Qatar (Charfeddine *et al* 2023), India and Brazil (Khan and Osińska 2023), different sectors in Iran (Javanmard and Ghaderi 2023), and peak load in Korea (Lee and Cho 2022). Most of these studies did not use any weather data as model input.

Beyond these regression models, several studies have applied data-driven approaches to predict energy demand at these larger scales. These have included using support vector machines (SVMs) to predict day-ahead demand over Gujarat, India (Chandra *et al* 2024); SVMs and recurrent neural networks to predict daily electricity demand over Bali, Indonesia (Aisyah *et al* 2022); LSTMs to make short-term predictions of demand for Chandigarh, India (Bedi and Toshniwal 2019) and for Germany and five US states (Li *et al* 2022); feedforward networks to predict next-day demand over Jeddah, Saudi Arabia (El Desouky and El Kateb 2000); recurrent neural networks to forecast daily electricity demand in Poland (Zielińska-Sitkiewicz *et al* 2021); and XGBoost or other decision tree-based models to forecast daily energy demand in Chinese districts (Dong *et al* 2023), a region of France (Qinghe *et al* 2022), a north Indian state (Gulati *et al* 2021), and Australian states (Richardson *et al* 2024).

The most skilful of these models tend to include as inputs weather data, population data, and information about the day of week and public holidays. Skill has also been improved through de-aggregation, e.g. using clustering to separate out different socioeconomic responses.

The novel aims of this paper are:

1. To scrape, quality-control, and make available daily demand data at the state-level for India.
2. To use these data, alongside weather, population, and calendar data to build explainable, data-driven state-level demand models for all of India.
3. To use these models to build state-level and national-level demand profiles for India, understanding the varying responses to meteorological conditions.
4. To combine these profiles with renewable supply models in order to identify weather conditions that may lead to energy stress events (e.g. potential for blackouts) in the future.

The state-wise demand models presented here could be implemented for any country with available demand data. The pre-trained models could also be used with climate data to assess regional changes to energy demand in future climate scenarios. The methodology could also be used for highly weather-dependent timeseries such as mortality rates (Wu *et al* 2024) or tourism demand (Falk 2014).

The paper is structured as follows: section 2.1 describes the meteorological data required to train the energy demand model and choices made around climate predictors, followed by the electricity demand data used for training and verification and previously developed renewable energy data used to explore times of high system stress. Section 2.2 describes the machine learning methods used to create the demand model. In section 3, the model is first evaluated (section 3.1) then Shapley analysis is employed to understand the key explainable components (section 3.2). Extreme renewable supply and demand conditions are then considered (section 3.3), followed by a discussion (section 4) and conclusions (section 5).

## 2. Data and methods

### 2.1. Datasets overview

#### 2.1.1. ERA5

ERA5 is the fifth generation atmospheric reanalysis of global climate produced by the Copernicus Climate Change Service at the European Centre for Medium-Range Weather Forecasts (Hersbach *et al* 2020). Data from ERA5 (available from <https://cds.climate.copernicus.eu/cdsapp#!/home>) cover the entire globe at a horizontal resolution of  $\sim 25$  km and resolve the atmosphere on 137 levels from the ground up to 80 km in altitude. It covers the period from January 1940 until present at hourly frequency. We use single-level data<sup>4</sup> to investigate fields closely related to energy demand: 10 m winds, 2 m temperature, 2 m dewpoint temperature (the temperature to which an air parcel must be cooled in order to become saturated), downward shortwave radiation at the surface and cloud cover. ERA5 is commonly used as a tool for energy-meteorology analysis (van der Wiel *et al* 2019, Bloomfield *et al* 2020, Richardson *et al* 2024) due to the limited length of observational records from existing renewable facilities.

#### 2.1.2. Universal Thermal Climate Index (UTCI)

We use UTCI data derived from ERA5 variables, described in Di Napoli *et al* (2021) and downloaded from <https://cds.climate.copernicus.eu/cdsapp#!/dataset/derived-utci-historical?tab=overview>. UTCI is designed to measure the human physiological response to a wide range of meteorological factors, including temperature, humidity, wind speed, and radiation (Bröde *et al* 2012). The relationship between temperature, wind speed, relative humidity, and UTCI is shown in figure A2.

#### 2.1.3. Holidays

To identify public holidays, we use the python holidays package<sup>5</sup>. Then for each state model, we define a variable *is\_holiday*, which is 1 if the date is either a national or state holiday, and 0 otherwise.

#### 2.1.4. Indian electricity demand data

To train our models, we need reliable observed electricity demand data. We extract these from daily reports issued by the Grid Controller of India (Grid-India), known before November 2022 as the Power System Operation Corporation (POSOCO). The reports are in PDF format and contain the total energy supplied by the national grid to each state on each day.

To convert these data from reports to quality-controlled machine-readable demand data required a lengthy and technical scraping pipeline. This is described in detail in appendix A. We show examples from the scraped and quality-controlled data in figure 1. Before being used to train the model, we also remove trends in the mean and interannual variance of the demand data, as these large arise from non-weather (i.e. socio-economic) sources. This process is explained in appendix B.


#### 2.1.5. Renewable energy production

For historical wind and solar energy production data for India, we use the modelled dataset developed in Hunt and Bloomfield (2024). They built models for hourly wind and solar production using a combination of ERA5 reanalysis data and information from multiple sources on installed capacity and locations of wind turbines and solar panels. They reported high accuracy, with linear correlation coefficients of 0.922 and 0.951 officially between reported production and their modelled wind and solar production respectively. The historical timeseries then uses ERA5 data to compute what wind and solar production would have been were 2023 installations exposed to historical weather. The all-India installed capacity in 2023, which we assume in this study, was 42.8 GW of wind power and 67.1 GW of solar power (Hunt and Bloomfield 2024).

In this paper, we aggregate these data over all-India to give national production. We also aggregate from hourly to daily resolution, matching the resolution of our demand data. We use these datasets at the end of the paper to identify potential causes of energy deficit events. When compared with demand timeseries, both the demand and supply data are standardised by subtracting the mean and dividing by the standard deviation. The original modelled historical production data are available to download from <https://doi.org/10.5281/zenodo.7824872>.

<sup>4</sup> doi: <https://doi.org/10.24381/cds.adbb2d47>.

<sup>5</sup> <https://pypi.org/project/holidays/>.



## 2.2. Model framework

### 2.2.1. XGBoost overview

In our study, we will build models using XGBoost (Chen and Guestrin 2016), a highly efficient and scalable decision-tree based algorithm that works well with large tabular datasets. Decision-tree models, in their simplest form, are machine learning algorithms that make predictions based on a series of nested binary questions, ultimately creating a ‘tree’ of decisions. Each node of the tree represents a question, with the end points or leaves signifying the final predicted outcomes (Quinlan 1986, Kotsiantis 2013, Breiman *et al* 2017).

Starting with a simple decision tree model, XGBoost iteratively adds small new trees that aim to correct the errors of the existing ensemble. The new trees are fitted to the residual errors of the current ensemble rather than the original targets, a process known as ‘boosting’. By doing so, XGBoost gradually reduces the prediction error, leading to a robust model with extremely strong performance metrics (Vidhya 2016, Dataaspirant 2020, C-SharpCorner 2021).

A key strength of XGBoost is its interpretability. As the decision-making process can be visualised as a series of ‘if-then’ rules, the model’s workings are transparent and understandable. As we will see (section 2.4), there are various tools to measure the importance of each feature in the model, which provide insights into the factors driving the predictions.

In the following sections, we will detail the specifics of our approach, explaining how we build and interpret our XGBoost models.

### 2.2.2. Choice of input features

Data-driven models require careful selection of input variables (also known as features or predictors). Here, we use a short list of weather variables, all of which are associated with human thermal comfort, as evidenced by their inclusion—in some form—in the calculation for UTCI (section 2.1.2). Those variables are: 10 m  $u$  and  $v$  wind, 2 m temperature and dewpoint temperature, surface solar radiation, and total cloud cover.

We must account for the fact that heating and cooling demand are sensitive to weather across a range of timescales—e.g. some turn their heating on for a season, some only for a day at a time when needed. Therefore, alongside daily means, we include means for the last 7 d and last 30 d for temperature, dewpoint, solar radiation, and cloud cover. Subdaily variability, e.g. overnight minimum temperature, is also an important timescale. To match the daily resolution of our demand data, we therefore add daily maxima and minima of all weather predictors, excluding the solar radiation minimum (which is always zero) and minimum and maximum winds.



Each of these variables is then averaged for each state, giving us a total of 20 weather predictors per state per day. Initial tests showed slightly better model performance when using a population-weighted mean, rather than a simple area-weighted mean, and so we adopted this method throughout. We also include day of the week, day of the year, and a flag for state or national holidays (see section 2.1.3). This accounts for different heating and cooling use in domestic and industrial sectors.

We acknowledge that many of these variables are correlated with each other, either seasonally (e.g. during the monsoon, wind speed increases while temperature decreases) or through thermodynamic relationships (e.g. temperature and dewpoint; surface solar radiation and cloud cover). This has no impact on the eventual model skill, but is an important caveat when we come to explainability.

### 2.2.3. Model configuration and training

We train one XGBoost model per state, reflecting the spatial resolution of our demand data and the fact that populations used to different climates respond to weather in different ways. This approach has the additional advantage that it reflects the differing socioeconomic conditions between states—wealthier states such as Karnataka (south India) are more likely to have domestic air conditioning than poorer states such as Bihar (north central India). For each model, we randomly split our data into training samples (constituting 80%) and testing samples (the remaining 20%). Keeping a distinct test dataset, which we use to evaluate the model, ensures that we are not overfitting and thus overly confident in its performances.

As each state has a distinct model, they also have distinct hyperparameter tuning procedures, which we now briefly describe. Hyperparameters describe the model learning process, but are not directly learned from the data themselves. XGBoost has a large number of tunable hyperparameters, mostly different ways of preventing overfitting, of which we consider eight of the most important. They are: learning rate (the speed at which the model tries to optimise a solution), maximum tree depth (the complexity of each weak learner—deeper trees can capture greater complexity but are prone to overfitting), subsample ratio (the fraction of training data used for growing each tree), column subsample ratio (the fraction of predictors used for each tree), three regularisation parameters (alpha, gamma and lambda—each prevent overfitting in different ways by discouraging overly complex models), and minimum child weight (makes sure that splits only occur if both resulting nodes have a population above this threshold).

This large number of hyperparameters makes a grid search intractable, and so we use a Bayesian approach (Mockus 1994, Jones *et al* 1998, Brochu *et al* 2010, Wang *et al* 2020). Bayesian optimisation first reduces the problem to the minimisation of a single multivariate objective function,  $RMSE = F(\chi_1, \chi_2, \chi_3 \dots)$ , where  $\chi_i$  are the hyperparameters. It then selectively samples  $F$  by iteratively computing new sampling points using so-called ‘acquisition functions’ which compromise between regions where  $F$  is already known to be high and regions where there is low certainty in the predicted value of  $F$ . We use a five-fold cross-validation approach, i.e. we split the training dataset into five equally sized subsets (folds), training on four and validating on the fifth. This is repeated for each combination, and the RMSE used in the Bayesian objective function is the average from the five validations.

In this study, we use the Python implementation described at <http://krasserm.github.io/2018/03/21/bayesian-optimization/>. There are several more hyperparameters associated with XGBoost training but results are not typically sensitive to these except in edge cases (Wang and Chen 2019). For these hyperparameters, we retain the Python implementation defaults (<https://xgboost.readthedocs.io/en/stable/parameter.html><sup>6</sup>).

As each state model has its own set of hyperparameters, we do not state them explicitly here. However, we briefly note some common themes. All models use nearly all the available data and features when building each decision tree (subsample ratio and column subsample ratio respectively very close to 1). To avoid making decision based on too little information, the minimum child weight—the number of data points a decision node must have before it can be split further—is typically at least six. Alpha is typically quite high (10–20), meaning a strong  $L_1$  regularisation, encouraging the model to ignore unnecessary features. Other regularisation parameters (gamma and lambda) tend to be much lower, affording the models some flexibility. The learning rate is typically quite small (most models below 0.15), meaning the model makes smaller adjustments during each learning step, affording greater stability at the cost of longer training. The average maximum tree depth is quite high ( $\sim 10$ ), suggesting a preference for more complex base learners.

## 2.3. Benchmark models

To benchmark the performance of our XGBoost models, we compare them against two simpler models—multivariate linear regression and a dense feedforward neural network—using exactly the same training data and methodology for each state. Both models were chosen due to their common use in earlier

<sup>6</sup> Archived at <https://web.archive.org/web/20230124160343/https://xgboost.readthedocs.io/en/stable/parameter.html>.

demand modelling studies. Multivariate linear regression is used in many energy demand studies (especially in European contexts) because of its interpretability and simplicity (see, e.g. Bessec and Fouquau 2008, Bloomfield *et al* 2016). Simple feedforward networks have been used for short-horizon load forecasting, and of course capture nonlinearity better than multivariate linear regression models (see, e.g. El Desouky and El Kateb 2000, Zielińska-Sitkiewicz *et al* 2021, Javanmard and Ghaderi 2023). These models therefore show (a) how well a simple, commonly used linear model (especially under collinearity) and (b) a more flexible but still relatively basic neural network fare against our method. We do not consider sequential models, such as LSTMs or ARIMA, as there is no forecasting component to our models.

In our multivariate linear regression benchmark, we retain all meteorological variables despite potential collinearity (e.g. between daily mean temperature, daily minimum temperature, and 7 d mean temperature). We do so to ensure a fair comparison between all models. Collinearity can reduce the reliability of regression models through inflating the variance in its coefficients (Dormann *et al* 2013). In contrast, XGBoost is inherently robust to redundancy in predictors as it splits sequentially on input variables. We emphasise that the multivariate linear regression results serve only as a benchmark and are used to demonstrate that more advanced models are needed under these conditions.

The dense neural networks comprise a fixed architecture of three hidden layers with a decreasing number of neurons in each (128, 64, 32). Each layer has a ReLU activation function, and each model is trained for 100 epochs.

## 2.4. Explainability

To make our decision-tree ensemble models interpretable and hence explainable, we use Shapley value theory (Shapley 1953, Roth 1988, Lundberg and Lee 2017). Originating in cooperative game theory, Shapley values were originally devised as a way to assign payouts to players depending on their contribution towards the total payout by considering how different permutations of players perturb the outcome. In black- or grey-box models such as our decision-tree ensemble, Shapley values estimate the marginal contribution from a given predictor in forcing a prediction away from the distribution mean. They thus fairly attribute the model's prediction to each feature, ensuring consistency in how the contributions are allocated.

Formally, the sum of all the Shapley values for all predictors for a given prediction,  $\hat{Y}$ , therefore, is equal to the difference between the predicted value and the predictand mean, i.e.  $\hat{Y} - E(Y)$ . At the level of individual predictions, summing these values reproduces the deviation of each prediction from the mean. Aggregating them over all samples allows us to quantify the average impact of each feature on the model output.

In this study, we use the `shap` Python package (<https://github.com/slundberg/shap>), which contains TreeSHAP (Lundberg *et al* 2020), an efficient and exact algorithm for computing Shapley values in decision-tree ensembles. It thereby avoids the prohibitive combinatorial complexity of naive Shapley calculations.


## 3. Results

### 3.1. Model evaluation

We evaluate each of the 32 XGBoost state models over their test dates in figure 2(a). While only five states have  $r^2 > 0.9$ , half have  $r^2 > 0.8$  and all but two have  $r^2 > 0.4$  (i.e. a correlation coefficient of over 0.63). This is higher skill than we get from either a simple three-layer dense neural network (figure 2(b)) which does worse in every state than XGBoost (median skill  $r^2 = 0.63$ ), or a multivariate linear regression model (figure 2(c)) which does worse in every state than the neural network (median skill  $r^2 = 0.50$ ). The results for all models and states are given in table A1.

There is a large variability in model skill between the different state XGBoost models, ranging from  $r^2 = 0.23$  (Arunachal Pradesh) and  $r^2 = 0.32$  (Mizoram) to  $r^2 = 0.95$  (Chandigarh) and  $r^2 = 0.96$  (Delhi and Punjab). There are several potential causes for this. Firstly, in states where the average demand is very low, e.g. Arunachal Pradesh and Mizoram, the quality of the demand data appears to be poorer than for other states. This is evidenced by poor resolution (not shown), where small daily totals, especially in the northeast (grey line in figure 1) seem to lead to a high granularity in the reported figures. Secondly, model performance scales strongly with population ( $r = 0.44$ ,  $r^2 = 0.19$ ,  $p \sim 0.01$ ), and even more strongly with  $\log(\text{population})$  ( $r = 0.63$ ,  $r^2 = 0.40$ ,  $p < 0.01$ ). The relationship between model skill and state per capita GDP is not statistically significant ( $r = 0.14$ ,  $r^2 = 0.02$ ,  $p \sim 0.4$ ). Even so, the poorest and most rural states—Bihar, Jharkhand, and the northeast—have the weakest model skill and comprise almost all cases of  $r^2 < 0.6$ .

This likely happens not only because urban populations use more power, especially for heating and cooling (Osunmuyiwa *et al* 2020, Walia *et al* 2022), thus giving their energy usage a larger weather-dependent component, but also because rural populations in India are less likely to be connected to the grid and tend to use diesel generators as a primary power source (Richmond *et al* 2020, Ray and Pullabhotla 2023). Because of




the correlation between state model skill and state average energy demand, the state models with poorest skill contribute little to the all-India total, which, if computed by summing the output from the individual state models has  $r = 0.9$  (similar to the performance of the demand.ninja model used in Staffell *et al* (2023)), substantially higher than the model median  $r^2$  of 0.78.

### 3.2. Explainability

By applying Shapley analysis to the state models, and making them explainable, we can both check that the models are accurately representing presumed human behaviour in the context of heating and cooling demand, and learn more about the nuances of that behaviour.


First, we look at two variables that are unrelated to weather but that are important drivers of human behaviour, the day of the week and public holidays (figure 3). For each state model, we show the standardised mean Shapley value for each day of the week for both holidays and non-holidays. Despite the variables being



**Figure 3.** Impact of day of the week and holiday on energy demand. For each of the state demand models, we compute and standardise Shapley values for each day of the week. For each day, the mean value is denoted by the blue circle, with the error bars indicating the 5th and 95th percentiles of the population. The standardised mean Shapley values are also shown for state/national holidays, denoted by black crosses. Greyed out boxes indicate where the individual state model has low skill, i.e.  $r^2 < 0.5$ . All subplots have the same y-axis.

uncorrelated, the energy demand response is coupled and nonlinear—where holidays occur on a Sunday, they typically increase energy demand; where they occur on a weekday, they typically lower it. The effect of the weekly cycle, and of holidays, varies substantially between different states. Both are typically larger in the southern states—for example, in Kerala (KL), Goa (GA), and Tamil Nadu (TN), the standardised Shapley value for Sundays is about  $-0.5$ . In other words, about half of the total deviation from the overall mean energy demand on a given Sunday is due to the fact that it is a Sunday. All states have negative Shapley values on Sunday, then often slightly less negative on Saturday or Monday, depending on the state's working week. In some states, in particular Bihar, Uttar Pradesh, Odisha, and Chandigarh, the effect of both weekday and holiday on the predicted energy demand is negligibly small.

We use a similar approach to investigate the relationship between temperature and energy demand (figure 4). Here, we retrain the models, removing all variables correlated with temperature (e.g. dewpoint, cloud cover, and temperature itself) and replace them with UTCI. This reduces the effect of confounding variables when investigating the relationship between weather and energy demand, but in using UTCI, we retain in a single field the important qualities of the longer list and their effect on human physiology (e.g. cooling wind, harsh sunlight, high humidity). To incorporate the effect of seasonal inertia—a warm winter day and a cold summer day may have the same temperature but different human responses—we split UTCI into two orthogonal variables. These are a mean of UTCI over the previous thirty days, and the daily



anomaly to that mean. We then train each state model on four variables: 30 d mean UTCI, anomalous daily UTCI, weekday, and holiday. This reduced set of predictors has a moderate effect on model skill, weakening  $r^2$  by about 0.1–0.15 in most cases compared to the full model, and resulting in more models falling below our acceptable skill threshold ( $r^2 > 0.5$ ) than in the full version. The skills of each of the full models and UTCI-only models are given in table A1. However, this small subset, in removing any confounding variables, allows us to more accurately quantify both the effect of heat stress and the role of timescales in energy demand.

In figure 4, we show 30 d mean UTCI plotted against its standardised Shapley values for each state model. The response to this seasonal cycle is typically more coherent in states where there is large annual variability in temperature—notably continental (as opposed to coastal) states towards the north: e.g. Punjab (PB), Delhi (DL), and Uttar Pradesh (UP). In these northern states, the U-shaped curve, which is common in many European countries (Bessec and Fouquau 2008) indicates the use of heating during conditions that feel cold and the use of air conditioning during conditions that feel hot. Note that that such conditions may not simply be caused by high temperatures alone, but also periods of intense solar radiation, low wind speed, or high humidity (see section 2.1.2 for the definition of UTCI). The relationship between temperature, wind speed, relative humidity, and UTCI is shown in figure A2.

Of particular interest in these states is that the inflection point of the U-shaped response varies considerably between states. It is lowest in states with colder winters (9 °C in Uttarakhand (UK), 16 °C in Chandigarh (CH)), and increases as we move into states with warmer climates (20 °C in Uttar Pradesh (UP), then 25 °C in Rajasthan (RJ) and Madhya Pradesh (MP)). This again demonstrates the value in having one



model per state, as populations used to different climates—of which India as many—respond in different ways to thermal stress. These inflection points are also of interest because they represent conditions with high demand predictability. In other words, as  $d(\text{demand})/d(\text{UTCI})$  tends to zero at these values, so does the sensitivity of demand to forecast errors. For comparison, the inflection points (also known as ‘dead zones’) are much wider, ranging from about 16 °C to 25 °C, in Europe (Bloomfield *et al* 2020).

As we move further south, the climate becomes warmer and more maritime, with a flatter annual cycle. As such, the response curves to 30 d mean UTCI represent only the right-hand side of the U-shape seen in the north, and become less coherent—instead showing greater sensitivity to daily changes in UTCI. Kerala (KL) and Tamil Nadu (TN) represent extreme cases of this, where the response curve effectively collapses onto a line.

There is also evidence that the two UTCI variables can combine nonlinearly. Considering the example of Delhi (DL), we see that below the 20 °C inflection point, warmer days reduce the seasonal heating burden, above the inflection point cooler days reduce the seasonal cooling burden; however, beyond a certain point in seasonal UTCI (about 32 °C), the daily UTCI anomaly becomes irrelevant as conditions are sufficiently warm to require constant AC use. This is indicated by the disappearance of the red–blue gradient from the bottom to the top of the swathe and is visible in many northern states, e.g. Delhi (DL), Uttar Pradesh (UP), Madhya Pradesh (MP), Punjab (PB), and Haryana (HR).

In summary, our simple UTCI models demonstrate important transitional behaviour with regional variability. The existence of a heating demand to cooling demand inflection point in many states implies a temperature-sensitive prediction skill for energy demand, whereas we also highlight the need to include both quick responses (e.g. to daily means) and slow responses (e.g. to seasonal means) in our models. Perhaps the most important point is that those states in north and central India with a coherent response to 30 d UTCI offer potentially very good predictability in forecasts of energy demand, even at relatively long lead times. From this model, it is not possible to ascertain the relative importance of monthly and daily UTCI for each state, as the latter is constructed to be orthogonal to the former. However, as we will see in the following analysis, these relationships can be obtained from the full models.

Using the same Shapley value explainability framework, now applied to the full models (including temperature, dewpoint temperature, wind speed, solar radiation, and so on) let us now look at some individual model predictions. We start with Kerala (figure 5), a wealthy, tropical state in the south. Here, we show the Shapley values for predictions of three days in the training dataset representing the highest, median, and lowest demand. The highest demand occurred in March 2015 (top of figure 5). The highest Shapley values, and thus the most important variables in predicting this extreme demand, are the 1 d and 7 d mean 2 m temperature. The 1 d mean of 28.5 °C represents about the 90th percentile for all March days (the second hottest month over Kerala) and about the 98th percentile for all days. The 7 d mean is almost the same, suggesting that these very high temperatures had persisted for at least a week prior. These unusually hot conditions were then compounded by high solar radiation and the fact this occurred on a weekday (Wednesday), presumably driving up office air-conditioning use. Crucially, this was a quick response, the 1 d mean temperature has about eight times more influence on the model prediction than the 30 d mean temperature.


The day of median demand (middle of figure 5) occurred in November 2018, during the winter. Here, the impact of predictors is much smaller than in the extreme days, and the signs of the impacts go in both directions. Of particular interest in this case is that this day was during Diwali, and so many offices were closed, reducing cooling demand despite it being unusually humid. So, although this was the median demand day, it was complex to model.

The day of lowest demand (bottom panel of figure 5) occurred just after the monsoon onset in 2015. Here, the model identifies reduced temperature, increased wind speed, and heavy cloud cover as import predictors that drive demand down, as well as it being a Sunday. These responses are all relatively short term—either daily or weekly means—indicating a quick response of energy usage once the monsoon onset has occurred and the summer heat has abated.

Together, these case studies indicate the rich tapestry of interactions that must be modelled to accurately predict energy demand. Over Kerala, these responses are typically fast (usually daily), and so accurate forecasts of energy demand will need accurate forecasts of these variables.

For comparison, we now apply the same analysis to a different state, Delhi (figure 6), which is in the north and has much larger annual variability in temperature and humidity than Kerala.

The day of highest demand over Delhi during our study period (top panel of figure 6) occurred just prior to the monsoon onset in 2021 (which was formally declared for Delhi on 12 July). The pre-monsoon is a very hot season, placing consistently high demands on air conditioning and fans, as evidenced by the importance of monthly and weekly average temperatures as predictors. The weekly average temperature (35.0 °C) was one of the highest in the study period. This case was further exacerbated by an extremely high overnight



**Figure 5.** Waterfall plots showing the relative contribution of each predictor to three predictions made by the full Kerala energy demand model. The three predictions correspond to the day with highest detrended demand over Kerala (top), the day of median detrended demand (middle), and the day with the lowest detrended demand (bottom). Predictors are sorted from top to bottom by the magnitude of their Shapley value, with red arrows indicating a positive contribution to the predicted demand and blue a negative contribution. Only the top nine variables are shown, with the effects of the remaining fourteen aggregated on the bottom row. Grey values to the left of the variable names indicate their values.  $f(x)$  is the predicted value (of detrended demand), and  $E[f(x)]$  is the mean value of observed detrended demand, against which the model predicts the deviation.

minimum temperature of 33.6 °C—the second-most important predictor—which likely resulted in many more households running air conditioning and fans overnight than usual. This overnight minimum was close to the record for our study period (34.5 °C). The daily mean and maximum were also about 1 °C less than the record (set on 5 June 2017).

The day of median demand (middle panel of figure 6) occurred in April 2017. The preceding thirty days had been relatively cool (22.3 °C), but the rising temperatures during the seasonal transition to pre-monsoon led to much higher temperatures over the preceding week (29.8 °C). These had contrasting effects on electricity demand, pushing it towards the centre of the distribution.

The day of lowest electricity demand in Delhi (bottom panel of figure 6) occurred in March 2015. The daily mean temperature (19.7 °C) is located at the inflection point of the demand-temperature curve that we saw in figure 4, i.e. between the low temperatures that require heating and the high temperatures that require air conditioning. Again, the average temperature for the preceding week is the most important predictor, with the average over the previous month, and the overnight minimum having a similar impact on the prediction. The low demand was reduced further by a national holiday—Holi was celebrated on 6 March 2015.

While such case studies are illuminating, they do not allow us to make general statements about the drivers of energy demand. To that end, we now aggregate the Shapley values for all model predictions across four states with different climates and socioeconomic conditions (figure 7). Kerala is ‘rich-hot’, Odisha is ‘poor-hot’, Delhi is ‘rich-hot/cold’, and Jammu and Kashmir is ‘poor-cold’.

For Kerala, aside from the day of the week, as we already discussed in figure 3, the 7 d and daily mean temperatures are the most important predictors. By extension, therefore, we claim that these are the two meteorological factors to which people respond most strongly when deciding whether to use energy for heating or cooling. Diurnal variations in temperature over Kerala are small due to its coastal location, and so there is no sensitivity to daily minima or maxima. There are lesser contributions from 30 d means of solar radiation, temperature, and dewpoint. These indicate a response to the monsoon onset, which lowers




Figure 6. As figure 5 but for Delhi.

temperature and solar radiation but increases dewpoint. This explains the counterintuitive response of increased energy demand with reduced dewpoint in this model.

For Odisha (top right panel of figure 7), temperature dominates the list of important predictors, and the relative importance scales with timeframe. Overnight minimum temperature is the most important, followed by daily mean temperature, then 7 d mean temperature, and then 30 d means of temperature and dewpoint provide important but relatively small adjustments. The response to the overnight minimum is strongly bimodal and suggests a threshold response, i.e. a minimum temperature beyond which people leave their air conditioning on overnight.

The model for Delhi (bottom left of figure 7) also exhibits a strong, but not bimodal, response to overnight minimum temperature, as already discussed in figure 6. Instead, the threshold response is to 30 d mean temperature, implying seasonal usage of heating and cooling units.


Unlike other states, Shapley values in the Jammu and Kashmir model (bottom right of figure 7) tend to decrease with increasing temperature. This is a cold state, due to its elevation, and so most energy variability here is weather-driven heating demand. As with other states, there is a threshold behaviour in the most important predictor (30 d mean dewpoint). Almost all the important predictors are 30 d or 7 d means, suggesting heating is turned on and off seasonally, rather than responding to specific daily weather.

In summary, the drivers of energy demand vary considerably by state—not only in the choice of weather variable, but its timeframe, from subdaily extremes to monthly means.

### 3.3. Extreme supply and demand

Having built the models and demonstrated they simulate plausible and robust relationships between weather and energy demand, we now use the models to understand the relationship between large-scale weather patterns and national energy demand. We thus use multi-decadal synthetic ERA5-driven data to extend the all-India energy demand record back to 1979 (figure 8), assuming the same population and economic conditions as 2023, summing predictions from the individual state models.

Lowest demand occurs in early winter, between mid-November and mid-December. Here, temperatures are typically close to their annual minima, but it is before the freezing temperatures in the northern states




necessitate heating. Further, the northeast monsoon provides cloud cover over the southern states, alleviating heat stress. Note that the diagonal streaks of reduced demand running through figure 8 represent Sundays.

Of particular interest to the grid are days where demand is extremely high—marked in white circles in figure 8. The majority of these high demand days occur just prior to the monsoon onset (mid-May–mid-June), with a few cases occurring during the monsoon itself, or later on during its withdrawal in September. In some cases, driven by synoptic-scale heatwaves, these can last for up to a week. Such events are huge draw on the grid as energy deficit rapidly accumulates, leading to a high risk of extended blackouts. Clearly, then, it is important to predict such weather-driven events in advance, such that appropriate action can be taken. As we will see, this is especially the case when periods of high demand co-occur with periods of low renewable energy production, leading to large energy deficits. Interestingly, these are very different to the days where low renewable production is considered in isolation, which commonly occur in December–February (see figure 1 of Dijkstra *et al* (2025) for further details). Exploring the weather patterns that drive these extreme events is the focus of the rest of this paper. We leave investigation of seasonal drivers, such as those leading to the high-demand monsoons of 1987 and 2015, for future work.

Hunt and Bloomfield (2024) developed models for solar and wind energy production over India. We now combine those with our all-India aggregate demand model to investigate weather-driven and seasonal covariance of energy deficit over India.

We start with two randomly chosen sample years, 2002 and 2003 (figure 9). As already discussed, all-India energy demand tends to peak just before the monsoon onset, and this is reflected in the annual maximum of UTCI in both years in late May and early June. UTCI then falls steadily through the monsoon, with a smaller secondary peak around the withdrawal as wind speed drops and conditions become less overcast. Conveniently, renewable energy production follows a similar annual cycle. Wind and solar production are also relatively complementary, giving broader coverage across the annual cycle. Solar energy production in India peaks from March to June during the high solar zenith angle and cloudless conditions of the spring and pre-monsoon, whereas wind energy peaks during the monsoon from June to September, albeit with high day-to-day variability. In effect, the seasons that drive high energy demand—high temperature in the



pre-monsoon and high dewpoint in the monsoon—also bring weather suitable for renewable energy production—high solar radiation in the pre-monsoon and high wind speed in the monsoon.


Therefore, when we compare the standardised values of all-India demand and weather-driven renewable supply (figure 9(b)), we see that, to a large extent, they follow each other through the annual cycle each year. However, there is some misalignment. Renewable supply tends to lead demand in March and April, resulting in a surplus. In contrast, demand tends to lead supply in September and October, resulting in a deficit. Thus, as India continues to decarbonise its grid, it will need to account for the likelihood of high demand but low production as the monsoon withdraws—and as we see even from these two years, the size of that deficit (i.e. the total red area) can vary considerably.

We demonstrate the variability of these post-monsoon energy stress events in figure 10, showing the standardised difference in renewable supply and demand, averaged over all India. We note two key features. Firstly, that there is significant variability in the magnitude and duration of the standardised deficit in September and October—in some years, such as 1994 and 1996, it is almost completely absent; in others, such as 2007 and 2015, it last for several months, with periods of over a week where the standardised difference is below  $-2.5$ . Secondly, there is a significant trend in the frequency of the most extreme events, where the difference between standardised supply and demand is in the first percentile. The trend in these events,  $0.25 \text{ yr}^{-1}$ , is statistically significant, and driven both by a significant increasing trend in October UTCI, and significant decreasing trends in September and October solar supply and October wind supply. High-deficit days do still occur in the early period of the timeseries. However, unpacking the relative changes in return periods and/or likelihood of extreme events is left for future work.

To understand the variability in all-India energy demand and renewable supply and corresponding deficit/surplus events, we plot composite synoptic situations for lowest, median, and highest deciles of both supply and demand, where they overlap (figure 11).

The most important case is where the top decile of demand coincides with the bottom decile of renewable supply (bottom right of figure 11). Such extreme cases are quite rare—there have only been 77 occurrences since 1979, and we would expect about 150 if supply and demand were uncorrelated—and in agreement with figure 9 tend to occur around mid September. Here, the monsoon winds have retreated






southwards, reducing cooling, while there is still substantial cloud cover and high dewpoint (and hence high UTCI). Combined, these still winds and hot but cloudy conditions lead to both high demand and low production. The average date of 13 September is just before the climatological start of the monsoon withdrawal, and so these events either arise due to rapid withdrawals or monsoon breaks occurring near the end of the season. It is clear that such events will pose an increasing threat to the energy security of India as renewables are increasingly mixed into the national grid, and thus the associated weather patterns must be treated accordingly as natural hazards.

At the other end of the spectrum, overlapping cases of high supply and low demand (top left of figure 11) are even rarer. These typically occur during the monsoon onset (with a mean date of 13 June) where strong winds cover much of India—including the regions with high wind power capacity, the Western Ghats and Gujarat—and coincide with an anomalously cool pre-monsoon. This is in contrast to, e.g. days of high demand and high supply (top right of figure 11, mean date of 2 June) where the wind pattern is similar, though indicative of a faster onset, but conditions persisting from a much hotter pre-monsoon drive high demand.

#### 4. Discussion

Why should stakeholders use a demand model as complex as the one presented in this study? After all, an experienced grid operator could anticipate rising demand during heatwaves without any model at all. The largest single advantage is that it is qualitative. By feeding the model with weather forecast data, the operator can forecast precise values of electricity demand that allow for more robust operational planning than simple threshold-based approaches. Our model also learns regional variations, accounting for state differences in infrastructure, socioeconomic status, and behavioural responses. As we have shown, these vary substantially




between different states. Finally, this model framework is highly flexible. Using local explainability and case studies, we showed that the impact of different variables can change in context, and that non-weather variables (such as calendar effects) can also play a role. Our framework easily allows new variables, which we discuss shortly, to be included in future versions.

After building and testing the demand models, we explored two applications—explainability and historical variability—both of which raise important discussion points.

We noted in figure 2 that the multivariate linear regression models had relatively poor skill, with only about half of the state models having  $r^2 > 0.5$ . Such models perform much better over Europe (van der Wiel *et al* 2019, Bloomfield *et al* 2020, Staffell *et al* 2023) where they remain more-or-less state-of-the-art. We showed that there are at least two causes of this—a response timescale that varies by state, socioeconomic condition, and even season (figures 5 and 7); and a response that varies nonlinearly as a function of its predictors. The latter cause is further supported by the significant drop in model skill when only UTCI was considered as a predictor, denying the models the opportunity to learn the nuanced relationships between temperature, humidity, wind, cloud cover, and energy demand. In other words, weather-driven energy demand responds to more than just human heat stress.

We also note the emergence of threshold behaviour, evidenced by inflection points in figure 4 and bimodal distributions of Shapley values in figure 7. Formally, these are localised and very large maxima in the derivative of state energy demand with respect to some predictor (e.g. temperature), and represent, e.g. some temperature at which a significant fraction of the population turn heating off or turn air conditioning on. The thresholds vary between states, but their existence implies potential tipping points, whereby climate change causes a permanent transition beyond the threshold, resulting in a much larger draw on the national grid.

Since our Shapley value analysis revealed a range of behaviours including lagged and threshold responses (e.g. turning on air conditioning overnight once minimum temperatures reach a certain point) and strong weekend and holiday effects, these could inform targeted policies, such as energy conservation campaigns, demand-side management, and focused infrastructure investment. For example, there is an opportunity for



programs that encourage higher air conditioning set-points or opportunities for passive cooling in new building developments. The weekend and holiday-driven dips in demand indicate the potential for targeted demand shifting or tariffs that exploit reduced industrial load during these times. Infrastructure investment would vary significantly by state, but could include fuel-switching incentives for states with heating demand in winter (e.g. Jammu and Kashmir) or weatherisation to reduce cooling demand in coastal states with persistently high temperature and humidity (e.g. Kerala).

We also showed in section 3.1 that our models performed relatively poorly in states with either small or predominantly rural populations. This raises the question as to whether the methods presented here are able to support the energy security of all, as we have an affluence and data quality bias.

In our analysis of historical demand data, we showed that high demand events typically coincide with seasonally high renewable energy production (solar energy in the pre-monsoon; wind energy in the monsoon; figure 9, see Dijkstra *et al* (2025) for renewable-only results). Importantly, therefore, when future studies investigate the nature or predictability of extreme deficit days—those of greatest importance to the national grid—the must use a holistic approach that takes into account not only weather-driven energy demand but also weather-driven shortages in renewable energy.

One important shortcoming of our study pertains to the way in which we preprocessed our training data. In removing the trends of mean and variance from state-level demand data before training the model, we implicitly make several assumptions. Firstly, that human response to heating stress is consistent over time. While this is largely true (Thornton *et al* 2016), there is a growing movement in India towards conserving energy, e.g. by setting thermostats higher and thus reducing air conditioning use (Jain *et al* 2014). The reverse is also true, as India's middle class grows, more families can afford to install domestic air conditioning units. While these effects are mostly removed by our running variance filter, it is possible that they still result in reduced model skill.

Secondly, in detrending our training datasets, we remove the effects of climate change on energy demand. Except for figure 8, this has little impact in our study since the magnitude of interannual variability of temperature over India, forced, e.g. by large-scale modes of climate variability such as the El Niño Southern Oscillation, is much larger than its trend over the short length of our dataset. However, this is generally undesirable, especially in studies investigating the link between climate change and energy demand. One avenue may be to detrend the data before training the models and then add the trend back in before running them (whether for explainability or longer period analysis), but further work is needed to determine the most robust approach.

Thirdly, in our choice of predictor variables (section 2.2.2), we only consider those important for human health and comfort. These are appropriate to describe weather-driven variability in energy consumption in the domestic sector, and for a large part the industrial sector. However, as of 2023, these comprised only 26% and 42% of total energy consumption in India respectively (MOSPI 2024). The third important sector, agriculture, comprised an additional 17% of the national total and its energy consumption is also sensitive to weather conditions. In our study, we ignore it, assuming that some variance is explained by predictors already included, such as monthly-mean temperature and dewpoint, but it may explain poor model skill in states with large agriculture industries, such as Bihar. As such, future studies wishing to further increase model skill may consider more fully incorporating agricultural energy needs by including soil moisture in their model inputs.

Other extensions to our study could include: investigation of seasonal-scale drivers of high energy demand, such as the monsoons of 1987 and 2015, as we saw in figure 8; incorporating hourly or subdaily weather data into the predictors, or, if made available, the energy demand data used to train the models; applying the demand models to climate model output to quantify the effects of climate change on Indian energy demand; and identifying weather regimes leading to high energy demand or a large energy deficit, and then quantifying their predictability in subseasonal to seasonal forecasts.

## 5. Conclusions

In this study, we built data-driven models to predict weather-driven energy demand over India at the state level. The models were built using XGBoost, a decision tree-based model framework for tabular data that is highly efficient, interpretable, and has very competitive skill. We trained one model per state to take in a range of population-weighted meteorological variables as predictors—near-surface fields that relate to human heat stress such as 2 m temperature and dewpoint, 10 m winds, and cloud cover—and estimate the total daily energy demand. Before training, we aggregated predictors over a range of timescales—daily minima, means, and maxima as well as averages for the past 7 d and past 30 d—to reflect the variety of human responses to changing weather. Our training target dataset, state-by-state daily energy demand totals, was scraped from daily reports issued by the Indian national grid (Grid-India, formerly known as POSOCO). We then quality-controlled these data, detrended them, and corrected for increasing variance that arises due to the growing population and economy. The skill of the trained models, verified using a separate test dataset, was very high, with half of all states having  $r^2 > 0.8$ , substantially better than a multivariate linear regression model, which is the commonly used across Europe (Bloomfield *et al* 2020).

In the first part of our results section, we used a combination of explainability techniques based on Shapley analysis to demonstrate the following interesting behaviours:

- Energy demand is sensitive to the day of the week and whether there is a public holiday. This effect is strongest in the central northern and southern states. In particular, energy demand is significantly reduced on Sundays (in Kerala by as much as half a standard deviation). Public holidays reduce consumption on weekdays but increase it at weekends.
- The most important predictors change considerably between states, and even seasons, but the two most commonly important variables across all the models are daily minimum temperature (reflecting nighttime usage of air conditioning) and 30 d mean temperature (reflecting a seasonal response to heating or cooling requirements).
- Several regional patterns emerge. For example, coastal states have low diurnal variability in temperature, and therefore energy demand there is more sensitive to 7 d means.
- Key predictor variables in many states have a strongly bimodal Shapley value distribution. This indicates threshold behaviour, where, e.g. air conditioning is turned on permanently once the temperature reaches a given value. This threshold behaviour varies between states and appears as a smoother inflection point in simpler models that only take into account human thermal comfort (i.e. UTCI).

In the second part of our results section, we used ERA5 reanalysis data to run the models to predict historical all-India energy demand since 1979, calibrated to 2023 population and economy.

- There is a pronounced seasonal cycle for all-India energy demand. It is lowest in winter (November and December) and highest in the pre-monsoon and monsoon onset (May and June). The onset of the monsoon leads to a fall in demand, although the values remain relatively high due to elevated dewpoint temperatures.
- However, by combining our results with earlier wind and solar production models, we showed that because the pre-monsoon has high solar energy production and the monsoon has high wind energy production, the largest energy deficits in fact occur during or after the monsoon withdrawal (September and October).
- Using a composite analysis, we found that the most extreme cases of energy deficit arise at the end of the monsoon season, caused either by a rapid and early monsoon withdrawal or a very late monsoon break. These lead to low wind speed across India, while the heavy cloud cover and high dewpoint conditions of the monsoon persist. Such events pose a substantial risk to the national grid and should be treated seriously in future studies.

India has been used as a case study for this method due to its strongly seasonally varying climate, and semi-openly available demand data for model training. The state-wise demand model presented here could be implemented for any country with available electricity demand data. The methodology could also be used for highly weather-dependent timeseries such as mortality rates (Wu *et al* 2024) or tourism demand (Falk 2014). Accurate and explainable models of this type are crucial if nations are to meet their net-zero targets while maintaining a reliable energy supply.

## Data availability statement

All the code used to build and train the models and create the figures in this paper will be openly available at <https://github.com/kieranmrhunt/india-renewable> following embargo. Code will be available from 1 May 2025. All data, namely the quality-controlled (but not detrended) daily demand data and population-weighted state averages of meteorological predictors, are openly available on Zenodo at <https://doi.org/10.5281/zenodo.14983361>.

## Acknowledgments

KMRH is supported by a NERC Independent Research Fellowship (MITRE; NE/W007924/1). HCB is funded by a Newcastle University Academic Track Fellowship.

## Appendix A. Scraping demand data

Prior to November 2022, POSOCO reports were stored at URLs of the form [https://posoco.in/download/21-03-17\\_nldc\\_psp/?wpdmdl=12345](https://posoco.in/download/21-03-17_nldc_psp/?wpdmdl=12345). The retrieved document depended only the final number with the choice of date being arbitrary, i.e. changing the date from 21 March 2017 to 30 October 2018 would still download the same file with ID 12345. These could therefore be scraped with a simple wget script that iterated through all values of XXXXX until error messages were returned. Since then, reports have been stored at URLs of the form [https://report.grid-india.in/ReportData/Daily%20Report/PSP%20Report/2024-2025/October%202024/06.10.24\\_NLDC\\_PSP.pdf](https://report.grid-india.in/ReportData/Daily%20Report/PSP%20Report/2024-2025/October%202024/06.10.24_NLDC_PSP.pdf). This more complicated format required a wget script to iterate through dates, such that the correct financial year (here, 2024–2025) and month name were used.

Running these download scripts leaves us with thousands of daily reports in PDF format. We provide a simplified version of our recipe<sup>7</sup> for extracting the data from these files and saving the output in a machine-readable tabular format in the appendix.

This process gives us a dataset for over ten years of daily energy demand for each Indian state, going from 1 January 2013 through to 28 April 2024. We retain the units of energy use from the reports—MU, a million units, equal to one gigawatt-hour—throughout our study.

A sample of these data is shown in figure 1. When we aggregate across the five main regions defined by the grid (top panel of figure 1), we see strong variability on both daily and seasonal timescales. In all regions, the seasonal variability has a stronger signal than the daily variability. The seasonal cycle varies considerably between regions, for example, demand in north India peaks later in the year (during the monsoon) than in

<sup>7</sup> The full code is available at <https://github.com/kieranmrhunt/india-renewable>.



south or west India (before the monsoon). Daily variability itself is also a function of season, with higher variability arising during seasons with higher demand.

This strong seasonality implies that much of the variance in energy demand can be explained through meteorology, as is the case for, e.g. Europe (van der Wiel *et al* 2019). This is readily demonstrated (bottom panel of figure 1) by comparing demand for a single state, here, Uttar Pradesh, with 2 m temperature, 2 m dewpoint, and 10 m wind speed. Most of the seasonal (i.e. intra-annual) variance is explained by variations in temperature; however, daily variability in demand appears to arise from a more complex interaction between all three variables. Modelling and understanding these interactions across different timescales is the focus of this paper.

### A.1. Scraping code


Here follows the recipe we developed for extracting the data from the POSOCO/Grid-India daily report files and saving the output in a machine-readable tabular format.

1. Define state name mappings. While most states use the same name in all documents (e.g. West Bengal, Maharashtra), the names of several vary between documents (e.g. Madhya Pradesh is sometimes rendered 'MP'; Jammu and Kashmir also appears as 'J&K', 'J&K(UT) & Ladakh(UT)', 'J&K(UT) and Ladakh(UT)'). We define a dictionary to make sure all such instances are assigned to the correct state. Note that in these reports, and thus in our models, energy demand for Ladakh and Jammu and Kashmir are combined.
2. Initialise data structures. We create empty lists to store dates and energy demand values for each state.
3. Load the report files in chronological order. We then use the `tika-python` package to extract the text into an unstructured dump.
4. Remove unwanted files. When we downloaded files from POSOCO using the greedy iterator, we also retrieved many irrelevant files (longer technical reports, error messages, memoranda, etc).
  - (a) We remove all files with an error message in the text dump.
  - (b) All daily demand reports are longer than one page. We thus check the file metadata and remove all files that only have one page.
  - (c) We also note that all daily demand reports contain the string 'NR WR SR ER' (denoting the regions at the top of the data table) and remove all files that do not contain this.
5. Identify the file date. For Grid-India files, this is trivial as it is prescribed during download and in the filename. For the bulk of our files, which were downloaded from POSOCO, we apply the regular expression `\d{1,2}-[a-zA-Z]{3}-\d{2,4}\n` to identify date-like strings in the text dump. These are then parsed and appended to the date list. If no date-like object is found, the file is removed.
6. Clean the output.
  - (a) We convert the dump to ASCII and remove all non-ASCII characters.
  - (b) We remove double spaces and any carriage returns within state names.
  - (c) In some files, garbled text appears in the table between certain pairs of states (e.g. Maharashtra and Goa). We identify and remove such instances using a regular expression search.
7. Extract the data block. We are fortunate that in every daily report, the first state in the table is Punjab and the last is Tripura. We thus use a regular expression to extract the whole table at once: `Punjab(?:[\s\S]*)Tripura.*(?:\r?n(?:!\r?n).*)*`.
8. Parse state data. We extract each line by splitting over `\n` and then find and extract the columnar data by using a regular expression to split over arbitrary-length whitespace: `\s(?:=[\d-])`.
9. Map state names and collect data. Find the state name using the dictionary defined earlier, then append the demand data to the appropriate list. As a final check, make sure all state demand data lists are of the same length. Finally, export to a CSV file.

We then apply a simple quality control procedure to the output file. For each state, we identify any days where the average difference between the reported energy demand and its neighbouring days is at least ten times larger than the average daily difference. Some of these arise from incorrect placement of a decimal point, and can be fixed automatically. In other cases, we replace the value through linear interpolation.

## Appendix B. Preprocessing the demand data

Even over the four years presented in figure 1, we see rising trends in both the mean and variance of energy demand. These trends become far more substantial when viewed over the length of our dataset (top row of figure A1). While some fraction of the trend in the mean is likely due to rising temperatures from climate



change, this is dwarfed by the contribution from population and economic growth. In 2013, the average all-India energy demand was 2650 MU. This increased to 3500 MU in 2018 (an increase of 32%), and to 4360 MU in 2023 (an increase of 65% from 2013). The same factors have also driven an increase in variance: the annual standard deviation in 2013 was 103 MU, rising to 196 MU in 2018 (an increase of 88%) and 355 MU in 2023 (an increase of 241% from 2013). We must also note the impact from the Covid-19 pandemic, which reduced energy demand across India for at least a year.

We must thus take steps to remove these three artefacts (rising trend, rising variance, and Covid). This requires multiple steps, which we lay out for all-India and two different states in figure A1. The first step (second row of figure A1) is to extract the long-term trend. We do this using LOESS smoothing (a form of local regression) because even though the trend is mostly linear, it has the additional advantage of capturing the effect of Covid.


Once this trend is removed (third row of figure A1), we can now clearly see the rising trend in variance in all three cases. To remove this, we divide the demand on each day by the standard deviation of demand over the previous 12 months, and then rescale by multiplying by the standard deviation for 2023. This leaves us (fourth row of figure A1) with a demand timeseries for each state that is corrected to reflect the socioeconomic conditions of 2023. We use these corrected demand timeseries as training targets for our models.

To reiterate, the variance removed by this preprocessing is specifically at the interannual scale. Early versions of our model trained without this variance adjustment overfitted on the *year* variable, reducing their ability to generalise and be used with past data (as we do later in the study) or future climate scenarios. By removing the interannual variance, which is primarily driven by economic and sociological factors rather than weather, we improved model skill by about 10%. This method does not affect weather-driven intra-annual variance, which is much larger (see figure 1) than its interannual counterpart and the clear

driver of variability in demand at daily, weekly, and seasonal scales. Consequently, and as we shall demonstrate in our model evaluation and explainability sections, our variance-adjustment approach preserves the core weather-demand relationship while mitigating spurious interannual variability that diminishes model robustness.

## Appendix C. The UTCI parameter space

Figure A2 shows the relationship between 2 m UTCI and 2 m temperature as a function of both wind speed and relative humidity, computed using the *thermofeel* library developed by Brimicombe *et al* (2022).



## Appendix D. Model skill scores

Table A1 shows the  $r^2$  score for each of the four model architectures used in this paper, for each Indian state. **Table A1.** Values of  $r^2$  for each state model computed on test data only. Double asterisks (\*\*) indicate the best performing model for each state, single asterisks (\*) indicate the second-best performing model. 'XGBoost (full)' refers to XGBoost models trained using all variables; 'XGBoost (UTCI-only)' refers to XGBoost models trained using only 30 d mean UTCI, the daily anomaly to that mean, weekdays, and holidays; 'NN (full)' refers to a three-layer fully-connected neural network trained using all predictor variables; and 'MLR (full)' refers to a multivariate linear regression model trained using all predictor variables.

| State             | XGBoost (full) | XGBoost (UTCI only) | NN (full) | MLR (full) |
|-------------------|----------------|---------------------|-----------|------------|
| Andhra Pradesh    | 0.808**        | 0.703*              | 0.650     | 0.508      |
| Arunachal Pradesh | 0.228**        | 0.119               | 0.205*    | 0.095      |
| Assam             | 0.821**        | 0.706               | 0.740*    | 0.594      |
| Bihar             | 0.556**        | 0.454*              | 0.326     | 0.247      |
| Chandigarh        | 0.950**        | 0.822               | 0.933*    | 0.706      |
| Chhattisgarh      | 0.825**        | 0.685*              | 0.668     | 0.480      |
| Goa               | 0.533**        | 0.495*              | 0.417     | 0.334      |
| Gujarat           | 0.844**        | 0.596               | 0.627*    | 0.457      |
| Haryana           | 0.929**        | 0.818               | 0.893*    | 0.641      |
| Himachal Pradesh  | 0.426**        | 0.299               | 0.338*    | 0.132      |
| Jammu & Kashmir   | 0.652**        | 0.496               | 0.519*    | 0.508      |
| Jharkhand         | 0.416**        | 0.124               | 0.299*    | 0.108      |
| Karnataka         | 0.867**        | 0.806*              | 0.719     | 0.669      |
| Kerala            | 0.835**        | 0.724               | 0.751*    | 0.680      |
| Madhya Pradesh    | 0.936**        | 0.794               | 0.848*    | 0.691      |
| Maharashtra       | 0.845**        | 0.688               | 0.750*    | 0.617      |
| Manipur           | 0.407**        | 0.332*              | 0.310     | 0.305      |
| Meghalaya         | 0.587**        | 0.498               | 0.506*    | 0.454      |
| Mizoram           | 0.317**        | 0.301*              | 0.270     | 0.229      |
| NCT of Delhi      | 0.961**        | 0.892               | 0.936*    | 0.745      |
| Nagaland          | 0.342**        | 0.235*              | 0.199     | 0.124      |
| Odisha            | 0.758**        | 0.602*              | 0.559     | 0.496      |
| Puducherry        | 0.636**        | 0.572*              | 0.527     | 0.520      |
| Punjab            | 0.964**        | 0.897               | 0.931*    | 0.754      |
| Rajasthan         | 0.778**        | 0.577               | 0.637*    | 0.296      |
| Sikkim            | 0.494**        | 0.401               | 0.458*    | 0.410      |
| Tamil Nadu        | 0.858**        | 0.668               | 0.689*    | 0.651      |
| Telangana         | 0.865**        | 0.714*              | 0.697     | 0.396      |
| Tripura           | 0.412**        | 0.345               | 0.367*    | 0.349      |
| Uttar Pradesh     | 0.849**        | 0.721*              | 0.705     | 0.495      |
| Uttarakhand       | 0.717**        | 0.566               | 0.579*    | 0.308      |
| West Bengal       | 0.873**        | 0.744               | 0.812*    | 0.720      |

## ORCID iDs

Kieran M R Hunt  <https://orcid.org/0000-0003-1480-3755>

Hannah C Bloomfield  <https://orcid.org/0000-0002-5616-1503>

## References

- Aisyah S, Simaremare A A, Adytia D, Aditya I A and Alamsyah A 2022 Exploratory weather data analysis for electricity load forecasting using SVM and GRNN, case study in Bali, Indonesia *Energies* **15** 3566
- Annamalai H and Slingo J M 2001 Active/break cycles: diagnosis of the intraseasonal variability of the Asian summer monsoon *Clim. Dyn.* **18** 85–102
- Antonopoulos I, Robu V, Couraud B, Kirli D, Norbu S, Kiprakis A, Flynn D, Elizondo-Gonzalez S and Wattam S 2020 Artificial intelligence and machine learning approaches to energy demand-side response: a systematic review *Renew. Sustain. Energy Rev.* **130** 109899
- Baker P and Blundell R 1985 The microeconomic approach to modelling energy demand: some results for UK households *Oxf. Rev. Econ. Policy* **7** 54–76
- Bedi J and Toshniwal D 2019 Deep learning framework to forecast electricity demand *Appl. Energy* **238** 1312–26
- Bessec M and Fouquau J-M 2008 The non-linear link between electricity consumption and temperature in Europe: a threshold panel approach *Energy Econ.* **30** 2705–21
- Bloomfield H C, Brayshaw D J and Charlton-Perez A J 2020 Characterizing the winter meteorological drivers of the European electricity system using targeted circulation types *Meteorol. Appl.* **27** e1858
- Bloomfield H C, Brayshaw D J, Shaffrey L C, Coker P J and Thornton H 2016 Quantifying the increasing sensitivity of power systems to climate variability *Environ. Res. Lett.* **11** 124025

- Breiman L, Friedman J H, Olshen R A and Stone C J 2017 *Classification and Regression Trees* (Routledge)
- Brimicombe C, Di Napoli C, Quintino T, Pappenberger F, Cornforth R and Cloke H L 2022 Thermofeel: a Python thermal comfort indices library *SoftwareX* **18** 101005
- Brochu E, Cora V M and De Freitas N 2010 A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning (arXiv:1012.2599)
- Bröde P, Fiala D, Błażejczyk K, Holmér I, Jendritzky G, Kampmann B, Tinz B and Havenith G 2012 Deriving the operational procedure for the Universal Thermal Climate Index (UTCI) *Int. J. Biometeorol.* **56** 481–94
- C-SharpCorner 2021 XGBoost—the choice of most champions (available at: [www.c-sharpcorner.com/article/xgboost-the-choice-of-most-champions](http://www.c-sharpcorner.com/article/xgboost-the-choice-of-most-champions)) (Accessed 27 June 2023)
- Chandra P K, Bajaj D, Sharma H, Bareth R and Yadav A 2024 Electrical load demand forecast for Gujarat state of India using machine learning models 2024 *IEEE Int. Students' Conf. on Electrical, Electronics and Computer Science (SCEECS)* (IEEE) pp 1–6
- Charfeddine L, Zaidan E, Alban A Q, Bennasr H and Abulibdeh A 2023 Modeling and forecasting electricity consumption amid the COVID-19 pandemic: machine learning vs. nonlinear econometric time series models *Sustain. Cities Soc.* **98** 104860
- Chaturvedi V and Malyan A 2021 Implications of a net-zero target for India's sectoral energy transitions and climate policy *Technical Report* (Council on Energy, Environment and Water) (available at: [www.ceew.in/publications/implications-of-net-zero-target-for-indias-sectoral-energy-transitions-and-climate-policy](http://www.ceew.in/publications/implications-of-net-zero-target-for-indias-sectoral-energy-transitions-and-climate-policy))
- Chaudhury S K, Gore J M and Ray K C S 2000 Impact of heat waves over India *Curr. Sci.* **79** 153–5 (available at: [www.jstor.org/stable/24103439](http://www.jstor.org/stable/24103439))
- Chen T and Guestrin C 2016 XGBoost: a scalable tree boosting system *Proc. 22nd ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining* pp 785–94
- Christensen P A, Anderson P A, Harper G D J, Lambert S M, Mrozik W, Rajaeifar M A, Wise M S and Heidrich O 2021 Risk management over the life cycle of lithium-ion batteries in electric vehicles *Renew. Sustain. Energy Rev.* **148** 111240
- Dataaspirant 2020 How the kaggle winners algorithm XGBoost algorithm works (available at: <https://dataaspirant.com/xgboost-algorithm>) (Accessed 27 June 2023)
- Deakin M, Bloomfield H, Greenwood D, Sheehy S, Walker S and Taylor P C 2021 Impacts of heat decarbonization on system adequacy considering increased meteorological sensitivity *Appl. Energy* **298** 117261
- Dhar O N and Rakhecha P R 1983 Foreshadowing northeast monsoon rainfall over Tamil Nadu, India *Mon. Weather Rev.* **111** 109–12
- Di Napoli C, Barnard C, Prudhomme C, Cloke H L and Pappenberger F 2021 ERA5-HEAT: a global gridded historical dataset of human thermal comfort indices from climate reanalysis *Geosci. Data J.* **8** 2–10
- Dijkstra I, Bloomfield H C and Hunt K M R 2025 Identifying weather patterns responsible for renewable energy droughts over India *Adv. Geosci.* **65** 127–40
- Dimri A P, Niyogi D, Barros A P, Ridley J, Mohanty U C, Yasunari T and Sikka D R 2015 Western disturbances: a review *Rev. Geophys.* **53** 225–46
- Dong D, Wen F, Zhang Y and Qiu W 2023 Application of XGBoost in electricity consumption prediction 2023 *IEEE 3rd Int. Conf. on Electronic Technology, Communication and Information (ICETCI)* (IEEE) pp 1260–4
- Dormann C F *et al* 2013 Collinearity: a review of methods to deal with it and a simulation study evaluating their performance *Ecography* **36** 27–46
- El Desouky A A and El Kateb M M 2000 Hybrid adaptive techniques for electric-load forecast using ANN and ARIMA *IEE Proc.-Gener. Trans. Distrib.* **147** 213–7
- Falk M 2014 Impact of weather conditions on tourism demand in the peak summer season over the last 50 years *Tour. Manage. Perspect.* **9** 24–35
- Fan J-L, Hu J-W and Zhang X 2019 Impacts of climate change on electricity demand in China: an empirical estimation based on panel data *Energy* **170** 880–8
- Fasullo J and Webster P J 2003 A hydrological definition of Indian monsoon onset and withdrawal *J. Clim.* **16** 3200–11
- Gulati P, Kumar A and Bhardwaj R 2021 Impact of COVID19 on electricity load in Haryana (India) *Int. J. Energy Res.* **45** 3397–409
- Harikrishnan G R, Sreedharan S and Nikhil Binoy C 2025 Advanced short-term load forecasting for residential demand response: an XGBoost-ANN ensemble approach *Electr. Power Syst. Res.* **242** 111476
- Hersbach H *et al* 2020 The ERA5 global reanalysis *Q. J. R. Meteorol. Soc.* **146** 1999–2049
- Hunt K M R *et al* 2025 Western disturbances and climate variability: a review of recent developments *Weather Clim. Dyn.* **6** 43–112
- Hunt K M R and Bloomfield H C 2024 Quantifying renewable energy potential and realized capacity in India: opportunities and challenges *Meteorol. Appl.* **31** e2196
- Hunt K M R and Fletcher J K 2019 The relationship between Indian monsoon rainfall and low-pressure systems *Clim. Dyn.* **53** 1–13
- Jain M, Chhabra D, Mankoff J and Singh A 2014 Energy usage attitudes of urban India *ICT for Sustainability 2014 (ICT4S-14)* (Atlantis Press) pp 208–17
- Javanmard M E and Ghaderi S F 2023 Energy demand forecasting in seven sectors by an optimization model based on machine learning algorithms *Sustain. Cities Soc.* **95** 104623
- Jones D R, Schonlau M and Welch W J 1998 Efficient global optimization of expensive black-box functions *J. Glob. Optim.* **13** 455–92
- Kapica J, Jurasz J, Canales F A, Bloomfield H, Guezgouz M, De Felice M and Kobus Z 2024 The potential impact of climate change on European renewable energy droughts *Renew. Sustain. Energy Rev.* **189** 114011
- Khan A M and Osińska M 2023 Comparing forecasting accuracy of selected grey and time series models based on energy consumption in Brazil and India *Expert Syst. Appl.* **212** 118840
- Kikuchi K, Wang B and Kajikawa Y 2012 Bimodal representation of the tropical intraseasonal oscillation *Clim. Dyn.* **38** 1989–2000
- Kotsiantis S B 2013 Decision trees: a recent overview *Artif. Intell. Rev.* **39** 261–83
- Kulkarni S, Deo M and Ghosh S 2018 Impact of active and break wind spells on the demand–supply balance in wind energy in India *Meteorol. Atmos. Phys.* **130** 81–97
- Lee J-Y, Wang B, Wheeler M C, Fu X, Waliser D E and Kang I-S 2013 Real-time multivariate indices for the boreal summer intraseasonal oscillation over the Asian summer monsoon region *Clim. Dyn.* **40** 493–509
- Lee J and Cho Y 2022 National-scale electricity peak load forecasting: traditional, machine learning, or hybrid model? *Energy* **239** 122366
- Li Z, Ye H, Liao N, Wang R, Qiu Y and Wang Y 2022 Impact of COVID-19 on electricity energy consumption: a quantitative analysis on electricity *Int. J. Electr. Power Energy Syst.* **140** 108084
- Liu S, Zeng A, Lau K, Ren C, Chan P-W and Ng E 2021 Predicting long-term monthly electricity demand under future climatic and socioeconomic changes using data-driven methods: a case study of Hong Kong *Sustain. Cities Soc.* **70** 102936



- Lundberg S M, Erion G, Chen H, DeGrave A, Prutkin J M, Nair B, Katz R, Himmelfarb J, Bansal N and Lee S-I 2020 From local explanations to global understanding with explainable AI for trees *Nat. Mach. Intell.* **2** 56–67
- Lundberg S M and Lee S-I 2017 A unified approach to interpreting model predictions *Advances in Neural Information Processing Systems* vol 30, ed I Guyon, U V Luxburg, S Bengio, H Wallach, R Fergus, S Vishwanathan and R Garnett (Curran Associates, Inc.) (available at: <https://proceedings.neurips.cc/paper/2017/file/8a20a8621978632d76c43dfd28b67767-paper.pdf>)
- Mockus J 1994 Application of Bayesian approach to numerical methods of global and stochastic optimization *J. Glob. Optim.* **4** 347–65
- Mohapatra M, Bandyopadhyay B K and Tyagi A 2012 Best track parameters of tropical cyclones over the North Indian Ocean: a review *Nat. Hazards* **63** 1285–317
- MOSPI 2024 Energy statistics India 2024 *Technical Report* (Ministry of Statistics and Programme Implementation, Government of India) (available at: [https://mospi.gov.in/sites/default/files/publication\\_reports/EnergyStatistics\\_India\\_publication\\_2024N\\_0.pdf](https://mospi.gov.in/sites/default/files/publication_reports/EnergyStatistics_India_publication_2024N_0.pdf))
- Osunmuyiwa O O, Payne S R, Ilavarasan P V, Peacock A D and Jenkins D P 2020 I cannot live without air conditioning! The role of identity, values and situational factors on cooling consumption patterns in India *Energy Res. Soc. Sci.* **69** 101634
- Pai D S, Sridhar L and Kumar M R R 2016 Active and break events of Indian summer monsoon during 1901–2014 *Clim. Dyn.* **46** 3921–39
- Parker D J, Willetts P, Birch C, Turner A G, Marsham J H, Taylor C M, Kolusu S and Martin G M 2016 The interaction of moist convection and mid-level dry air in the advance of the onset of the Indian monsoon *Q. J. R. Meteorol. Soc.* **142** 2256–72
- Parkpoom S and Harrison G P 2008 Analyzing the impact of climate change on future electricity demand in Thailand *IEEE Trans. Power Syst.* **23** 1441–8
- Qinghe Z, Wen X, Boyan H, Jong W and Junlong F 2022 Optimised extreme gradient boosting model for short term electric load demand forecasting of regional grid system *Sci. Rep.* **12** 19282
- Quinlan J R 1986 Induction of decision trees *Mach. Learn.* **1** 81–106
- Rajeevan M, Unnikrishnan C K, Bhatte J, Kumar K N and Sreekala P P 2012 Northeast monsoon over India: variability and prediction *Meteorol. Appl.* **19** 226–36
- Ratnam J V, Behera S K, Ratna S B, Rajeevan M and Yamagata T 2016 Anatomy of Indian heatwaves *Sci. Rep.* **6** 24395
- Ray S and Pullabhotla H K 2023 The changing impact of rural electrification on Indian agriculture *Nat. Commun.* **14** 6780
- Raynaud D, Hingray B, François B and Creutin J D 2018 Energy droughts from variable renewable energy sources in European climates *Renew. Energy* **125** 578–89
- Richardson D, Hobeichi S, Sweet L-B, Rey-Costa E, Abramowitz G and Pitman A J 2024 Predicting Australian energy demand variability using weather data and machine learning *Environ. Res. Lett.* **20** 014028
- Richmond J, Agrawal S and Urpelainen J 2020 Drivers of household appliance usage: evidence from rural India *Energy Sustain. Dev.* **57** 69–80
- Rohini P, Rajeevan M and Srivastava A K 2016 On the variability and increasing trends of heat waves over India *Sci. Rep.* **6** 1–9
- Roth A E 1988 *The Shapley Value: Essays in Honor of Lloyd S. Shapley* (Cambridge University Press)
- Rowley N 2016 Gas demand forecasting methodology *Technical Report* (National Grid Gas) (available at: [www.nationalgrid.com/sites/default/files/documents/8589937808-Gas%20Demand%20Forecasting%20Methodology.pdf](http://www.nationalgrid.com/sites/default/files/documents/8589937808-Gas%20Demand%20Forecasting%20Methodology.pdf))
- Sailor D J and Muñoz J R 1997 Sensitivity of electricity and natural gas consumption to climate in the USA—methodology and results for eight states *Energy* **22** 987–98
- Sailor D J and Pavlova A A 2003 Air conditioning market saturation and long-term response of residential cooling energy demand to climate change *Energy* **28** 941–51
- Satyanarayana G C and Rao D V B 2020 Phenology of heat waves over India *Atmos. Res.* **245** 105078
- Seyedzadeh S, Rahimian F P, Glesk I and Roper M 2018 Machine learning for estimation of building energy consumption and performance: a review *Vis. Eng.* **6** 1–20
- Shapley L 1953 A value for n-person games *Contributions to the Theory of Games II* (Princeton University Press)
- Somu N, Gauthama Raman M R and Ramamritham K 2021 A deep learning framework for building energy consumption forecast *Renew. Sustain. Energy Rev.* **137** 110591
- Staffell I, Pfenniger S and Johnson N 2023 A global model of hourly space heating and cooling demand at multiple spatial scales *Nat. Energy* **8** 1328–44
- Sun X *et al* 2024 Reducing supply risk of critical materials for clean energy via foreign direct investment *Nat. Sustain.* **7** 1–10
- Taylor J W and McSharry P E 2007 Short-term load forecasting methods: an evaluation based on European data *IEEE Trans. Power Syst.* **22** 2213–9
- Thomas T M, Bala G and Srinivas V V 2021 Characteristics of the monsoon low pressure systems in the Indian subcontinent and the associated extreme precipitation events *Clim. Dyn.* **56** 1859–78
- Thornton H, Hoskins B J and Scaife A 2016 The role of temperature in the variability and extremes of electricity and gas demand in Great Britain *Environ. Res. Lett.* **11** 114015
- van der Wiel K, Bloomfield H C, Lee R W, Stoop L P, Blackport R, Screen J A and Selten F M 2019 The influence of weather regimes on European renewable energy production and demand *Environ. Res. Lett.* **14** 094010
- Vidhya A 2016 Winning solutions of DYD competition—R and XGBoost ruled (available at: [www.analyticsvidhya.com/blog/2016/03/complete-guide-parameter-tuning-xgboost-with-codes-python/](http://www.analyticsvidhya.com/blog/2016/03/complete-guide-parameter-tuning-xgboost-with-codes-python/)) (Accessed 27 June 2023)
- Walia A, Tathagat T, Dhingra N and Varma P 2022 A residential end use energy consumption and appliance ownership patterns in India *Energy Efficiency in Domestic Appliances and Lighting: Proc. 10th Int. Conf. (EEDAL'19)* (Springer) pp 67–79
- Wang B, Ding Q and Joseph P V 2009 Objective definition of the Indian summer monsoon onset *J. Clim.* **22** 3303–16
- Wang J, Clark S C, Liu E and Frazier P I 2020 Parallel Bayesian global optimization of expensive functions *Oper. Res.* **68** 1850–65
- Wang Y and Chen Y 2019 Significant climate impact of highly hygroscopic atmospheric aerosols in Delhi, India *Geophys. Res. Lett.* **46** 5535–45
- World Bank 2023 India overview: development news, research, data (available at: [www.worldbank.org/en/country/india/overview](http://www.worldbank.org/en/country/india/overview))
- Wu Y *et al* 2024 Temperature frequency and mortality: assessing adaptation to local temperature *Environ. Int.* **187** 108691
- Yan K, Li W, Ji Z, Qi M and Du Y 2019 A hybrid LSTM neural network for energy consumption forecasting of individual households *IEEE Access* **7** 157633–42
- Zielińska-Sitkiewicz M, Chrzanowska M, Furmańczyk K and Paczutowski K 2021 Analysis of electricity consumption in Poland using prediction models and neural networks *Energies* **14** 6619