Energy 281 (2023) 128101

Contents lists available at ScienceDirect

Energy

journal homepage: www.elsevier.com/locate/energy

Extreme event counterfactual analysis of electricity consumption in Brazil: Historical impacts and future outlook under climate change

Gianlucca Zuin [a,c,*], Rob Buechler [b], Tao Sun [b], Chad Zanocco [b], Francisco Galuppo [a,c], Adriano Veloso [a], Ram Rajagopal [b]

[a] Universidade Federal de Minas Gerais, Belo Horizonte, Brazil
[b] Stanford University, Stanford, USA
[c] Kunumi, Belo Horizonte, Brazil

ARTICLE INFO

Dataset link: https://cds.climate.copernicus.eu, http://www.ons.org.br/, https://pangeo-data.github.io/pangeo-cmip6-cloud/

Keywords:
Counterfactual
Electric consumption
Forecast
Extreme events

ABSTRACT

Climate change and other disruptive events have a significant impact on the electrical grid, affecting both power supply and consumption. With the rise in frequency and severity of extreme events, like heatwaves and droughts, the stability and operations of the system are increasingly at risk. The COVID-19 pandemic, unrelated to climate, has also brought about dramatic shifts in global energy patterns. We apply machine learning to model electricity consumption counterfactuals for Brazil, one of the largest hydropower producers, to understand the effects of these events. By training our model on 23 years of data (1999–2021), we achieved a .848 R^2 and 2.6% MAPE. This enabled us to assess the impact of historical events on electricity consumption at both hourly and daily levels. Next, we use climate change scenarios to forecast electricity consumption and find that Brazil's capacity is unlikely to meet demand from 2070 on-wards. Our research provides much needed insight into the impact of extreme events on Brazil, with implications for understanding energy system responsiveness and resiliency. The counterfactual approach proposed is also transferable to other countries and contexts, with the potential for new application areas given interactions between extreme events, climate change, and transitioning energy systems.

1. Introduction

Energy systems are increasingly coupled with human, environment, and climate systems. For example, the deployment of renewable energy is critical for mitigating future climate change and improving population health [1,2], and hydropower, the leading renewable resource in the world [3], is dependent on water availability. Energy and climate experts are therefore raising concerns regarding the capacity of renewable generation sources (e.g., hydro, wind, and solar) to reliably meet electricity demand under future climate change and human development pressures. Increases in electricity consumption due to population and economic growth, worsening severity and frequency of extreme weather events, long term effects of climate change on surface temperatures and precipitation, and rapid land use change can present significant challenges to energy systems, particularly so for those dependent on renewable resources [4]. A crucial component of energy planning, therefore, requires an understanding of how these challenges translate to impacts on electricity consumption and supply.

Such issues are typified by Brazil, a country with an electricity system that is heavily dependent on hydropower. Brazil is currently undergoing substantial economic and social transitions with populations living in geographically and climatologically diverse regions, all contributing to changing patterns in electricity consumption. As of April 2023, Brazil's total generation capacity was 190.8 GW [5], the third-largest electricity sector in the Americas and only second to China in total hydroelectricity capacity. Brazil's electricity sector is comprised of 56.98% hydropower and 25.42% thermal [6] with an expected 2.7% yearly growth in grid capacity [7]. There are also indications that Brazil could be facing an energy crisis in the near future, spurred by low reservoir levels and drought, and exacerbated by increases in electricity consumption [8].

Historically, Brazil has relied almost completely on clean energy, but recently (see Fig. 1) there has been increased reliance on thermal generation sources. While in 2010 over 90% of electricity demand was satisfied by hydropower, a decade later hydropower generation

* Corresponding author at: Universidade Federal de Minas Gerais, Belo Horizonte, Brazil.
E-mail addresses: gzuin@dcc.ufmg.br, gianlucca@kunumi.com (G. Zuin), rbuec@stanford.edu (R. Buechler), luke18@stanford.edu (T. Sun), czanocco@stanford.edu (C. Zanocco), franciscogaluppo@dcc.ufmg.br, francisco@kunumi.com (F. Galuppo), adrianov@dcc.ufmg.br (A. Veloso), ram.rajagopal@stanford.edu (R. Rajagopal).

https://doi.org/10.1016/j.energy.2023.128101
Received 12 December 2022; Received in revised form 19 May 2023; Accepted 11 June 2023
Available online 14 June 2023

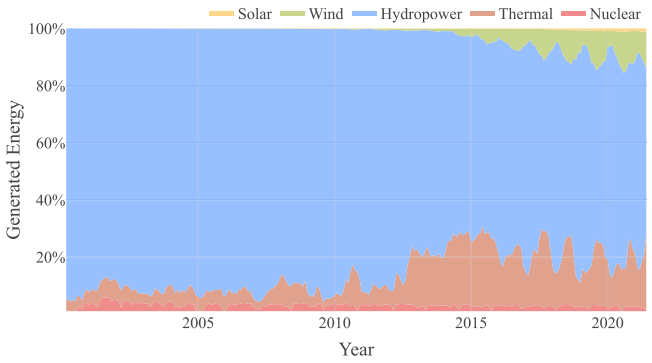

Fig. 1. Electricity generation profile of Brazil. Although hydropower constitutes over 50% of electricity generation, Brazil is experiencing a growth in the thermal generation share driven mostly by lower reservoir levels and increases in electricity demand.

fell to 50%–60%, with thermal constituting approximately 20% of total generation. The wider adoption of wind-based alternatives has offset some of this reliance on thermal, however, the main concern still stands: renewables may not be able to supply the majority of Brazilian electricity demand due to steady increases in population consumption, decreasing water availability and hydro-related crises, and unexpected shocks from extreme or disruptive events. For these reasons, developing a deeper understanding of the Brazilian energy system is both timely and of critical societal importance [9]. In this research, we narrow our study's focus to one aspect of the energy system: determinants of electricity consumption.

Given the tight coupling of weather to the Brazilian electrical system at a national and regional level, the impact of disruptive events, and the need for future planning under climate change, we propose the following research questions:

RQ1 What are the primary drivers of electricity consumption in Brazil and how do these drivers vary across regions?
RQ2 What are the impacts of disruptive or atypical events on consumption?
RQ3 What is the projected demand in the future and what capacity is needed to meet this future consumption?

To address these research questions, we test and deploy a suite of machine learning approaches on multiple data sources including electrical system operations, weather observations, historical events, and future climate projections. Building upon previous work [10], we expand this research by leveraging data covering over 70% of the Brazilian population to model hourly electricity consumption, investigating historical disruptive/extreme events in detail, and then applying future climate scenarios to quantify projected grid capacity, hydropower capacity, and consumption forecasts at regional and national-levels. As a result, our research provides a much needed deep exploration into historical, current, and future determinants of electricity consumption in Brazil and proposes a generalized counterfactual framework for measuring impacts of extreme or disruptive events that is transferable to other energy systems with comparable levels of data availability.

2. Related work

While many statistical methods for electricity demand prediction have been proposed and are demonstrated to be performant, one previously identified limitation has been an overreliance on linear or quasilinear relationships [11]. For example, [12] compares classical statistical methods utilizing linear relationships against more complex statistical approaches and found that gradient-boosting methods have higher performance when predicting energy consumption in urban areas. However, this is expected as tree-based approaches are known to struggle with out-of-distribution data [13] and if applied to demand forecasting this could be problematic as impacts from extreme events and climate change can exceed historical records. To address the challenge of out-of-distribution observations, [14,15] proposes augmenting datasets to improve extrapolation using the Data-Augmented Regression for Extrapolation or Data Augmented Regression (DAR) algorithm which consists of repeat sampling on the desired domain and using a background model to provide weighted responses. In our work, we propose applying the DAR algorithm to a tree-based electricity consumption model so that it can better accommodate out-of-distribution data, which we hypothesize will enable state-of-the-art performance in the modeling of extreme and disruptive events.

Disruptive events can impact electricity consumption, with a notable recent event being the COVID-19 pandemic, which has contributed to substantial uncertainty in consumption forecasting. While the pandemic is an example of one type of disruptive event [16], other disruptive events such as economic recessions can have dramatic impacts on the power sector [17]. In addition to such disruptions, long-term social and economic trends, such as population growth and urban development, will likely contribute to load growth. There is also the potential that future climate change – which will result in warmer yearly temperatures – will in turn result in increasing consumption. Warmer temperatures could also change the timing and intensity of daily electricity consumption, with peak electricity demand occurring at different hours of the day and at higher levels due to new deployment and utilization of building cooling technologies [18]. Moreover, these challenges and impacts may not be ubiquitous, and in the Brazilian context differ substantially due to climate, geographic, and economic variations across regions [19].

While some disruptive events may have impacts that persist across longer time horizons (e.g., greater than one year), such as the COVID-19 pandemic, economic recessions, and climatological changes, there are also shorter-duration extreme events (e.g., ranging from days to months) such as heat/cold waves, hurricanes, and flooding that can generate dramatic impacts on energy supply and demand [20], such as a 2021 Texas winter storm that left millions without power [21]. Research has identified the importance of generating better seasonal, subseasonal, and even week ahead forecasts to account for extreme events and their potential impacts on everything from renewable energy generation to electricity demand [22]. However, forecasting these extreme events, and even generating a better understanding of their direct impacts on the electricity system, remains an ongoing challenge [23–25]. In acknowledgment of these challenges, industry and regulatory groups are now calling for the introduction of reliability and resource adequacy metrics that better account for extreme weather events [26].

Counterfactual modeling has emerged as a technique to understand the impacts of disruptive events on the electricity system in a more generalized manner [27], allowing for comparisons across the world, including North America, South America, Europe, and Asia. For example, Buechler et al. (2022) used a counterfactual modeling approach to understand the impact of COVID-19-related restrictions on electricity consumption across 58 different countries/regions. In this work, counterfactual models of electricity consumption were developed for each country/region to simulate what consumption would have been in the absence of the pandemic and then compared against actual demand during this same period.

Concerning the Brazilian energy system, [28] evaluated 20 years of electricity load for the Brazilian grid and investigated scenarios for the balance of electricity supply and demand through 2030 considering projected renewable penetration. In related work, [29] generated a large number of realistic weather pathways and produced a probabilistic electricity demand forecast for Brazil for 2016–2100. In addition to these research efforts, a spatial econometrics approach was applied in [30] to forecast regional electricity consumption in Brazil. Similarly, [31] quantified the effects of changing Gross Domestic Product (GCP)

on demand for electricity in Brazil, allowing for a scenario-based analysis of possible future conditions. Focusing on the residential sector in Brazil, [32] proposed spatiotemporal models to estimate electricity demand. Recent research forecasting Brazilian electricity demand has compared more traditional time series analysis with machine learning methods, such as artificial neural networks, and found machine learning approaches to be a powerful tool when sufficiently large data is available [33]. On the energy generation side, [34] explored the uncertainties in future hydropower generation capacity in Brazil using insights from climate models.

We also conducted a review of the extant literature pertaining to Brazil and South America in the context of energy and related topics using the Scopus database [35]. Our search focused on articles related to weather, extreme events, demand, and consumption of energy, published between 2019 and 2023. In the entire Scopus database, we identified 41,634 relevant articles, for which studies encompassing South America accounted for 6976 (16.8%), and Brazil appearing in 4087 (9.8%) of them. This is comparatively lower than other regions such as North America with 22,444 (53.9%), with the United States accounting for 20,613 (49.5%), or Europe with 30,297 (72.7%). However, when filtering articles that incorporate analysis for either North America or Europe, only 249 (0.6%) studies mentioned Brazil without also including North America or Europe. This imbalance in the representation of Brazil and South America for energy-related topics suggests a discrepancy in the literature, a pattern also observed for other topics and for other countries in the global south more generally [36].

Such national studies are critical as each electricity grid system can include unique and place-specific features, with Brazil having among the highest hydropower penetration in world, so insights from other scholarship from other countries may not necessarily be generalizable to the Brazilian context. To help address this research gap, we apply Brazilian electricity consumption data in both counterfactual modeling and to build future climate-energy scenarios, using a wider time range with higher spatial and temporal resolution than extant scholarship, described in detail in the following section.

3. Materials and methods

3.1. Data and input variables

We applied three primary datasets in our study: Brazil's National Energy System Operator (ONS) historical reports [37], Copernicus ERA5 reanalysis data [38], and CMIP6 projections data [39], all available through public sources (see Data Availability Statement). Energy data was retrieved directly from the ONS website, which includes a diverse range of information about the Brazilian power system, such as load, maximum consumption, mean and total megawatts (MW), in hourly, monthly, and yearly intervals. Complete data from the ONS website is available from 1999 onward.

It is important to note that the Brazilian electrical grid is split into four subsystems associated with its main geopolitical regions. Although there is some exchange of energy between adjacent subsystems, they are mostly independent in supplying local consumption and ONS provides information regarding each subsystem separately. Supplementary Figure 1 illustrates the load profiles for each of Brazil's macro-regions (1999–2022). Notably, there was rapid load growth in each macro-region, with an overall increase of approximately 60% in consumption across these 20 years. Additionally, there is a close relationship between urban development and load, with Brazil's technological hubs (such as São Paulo in the Southeast region) having a higher energy consumption than less developed areas (such as the North region). Finally, there was a sharp drop in consumption at the end of 2002 related to a curtailment intervention initiated by the Brazilian government. This phenomenon will be further analyzed in Section 4.

Weather data was collected from the Copernicus Data Store, an European observation program that provides free and open access to environmental and climate data. The Copernicus Data Store is a repository of satellite data, in-situ measurements, and model outputs, covering a wide range of environmental and climate-related parameters, all accessible via an API. We applied the ERA5 reanalysis dataset, which consists of global hourly estimates from 1950–2022 for atmospheric variables, with a spatial resolution of 0.25 degrees (approximately a 30 × 30 km grid cell). To associate this weather data to Brazilian subsystems we extracted the mean values from grid cells using coordinates from each city with at least 100,000 inhabitants, weighting the values by the population so that larger cities will have a greater impact on mean values compared to smaller ones. Fig. 2 displays all included cities, encompassing over 70% of the total Brazilian population. Our final weather-related feature subset consists of daily temperature minimums, means and maximums, humidity, wind speed, precipitation, heating degree days (HDD), cooling degree days (CDD), heat index, wind chill index, apparent temperature, HDD derived from the wind chill index and CDD derived from the heat index. Supplementary Tables I–VIII contains descriptions and summary statistics for all of the variables applied in our research.

Degree days are a measure of heating or cooling, which has extensively been used for estimating energy consumption required for building occupants to perceive a comfortable indoor temperature and can be computed as an integral of a function of time over temperature. However, since temperature measurements in our data are not continuous, but rather recorded at discrete intervals, we can approximate heating degree days and cooling degree days for a single day as:

$$CDD = \frac{\sum_{i=1}^{T}(\theta_i - \theta_b)_{((\theta_i-\theta_b)>0)}}{T} \tag{1}$$

$$HDD = \frac{\sum_{i=1}^{T}(\theta_b - \theta_i)_{((\theta_i-\theta_b)<0)}}{T} \tag{2}$$

in which θ_i represents the temperature at instant i, θ_b is some constant representing the base comfort temperature and T is the number of equally spaced temperature measurement intervals. HDD is related to the amount of time that the temperature remained below some established threshold and the size of this difference, implying that household heating was needed to reach the comfort temperature. Analogously, CDD is related to temperatures rising above this same threshold, leading to the need for cooling. We consider the base comfort temperature in Brazil as 65 °F or 18.3 °C.

However, as noted by [40], the perceived temperature often differs from ambient dry bulb temperature. Thus, [40] defines the apparent temperature as the temperature equivalent perceived by humans caused by the combined effects of air temperature, relative humidity, and wind speed. The NOAA National Digital Forecast Database (NDFD) states that when the temperature falls below 50 °F, wind chill is a suitable measure of apparent temperature, and when the temperature rises above 80 °F, heat index should be employed. Between these two thresholds, humans experience the combined effect of both wind speed and humidity, and the ambient dry bulb air temperature is a reasonable measurement of apparent temperature. Heat index (HI) and wind chill index (WCI) can be computed as:

$$HI = c_1 + c_2\theta + c_3\phi + c_4\theta\phi + c_5\theta^2 + c_6\phi^2 + c_7\theta^2\phi + c_8\theta\phi^2 \tag{3}$$

$$WCI = 35.74 + 0.6215\theta - 35.75v^{0.16} + 0.4275\theta v^{-0.16} \tag{4}$$

in which θ is the dry bulb air temperature, ϕ is the relative humidity, v is the air speed and c_i are constants used to approximate [40] original heat index table [41,42], valued as $c_1 = 0.363445, c_2 = 0.988622, c_3 = 4.777114, c_4 = -0.114037, c5 = -8.50208e^{-4}, c6 = -2.071619e^{-2}, c7 = 6.87678e^{-4}, c8 = 2.74954e^{-4}$. Since both HI and WCI are measured in °F, we can replace the temperature from Eqs. (1) and (2) to obtain both

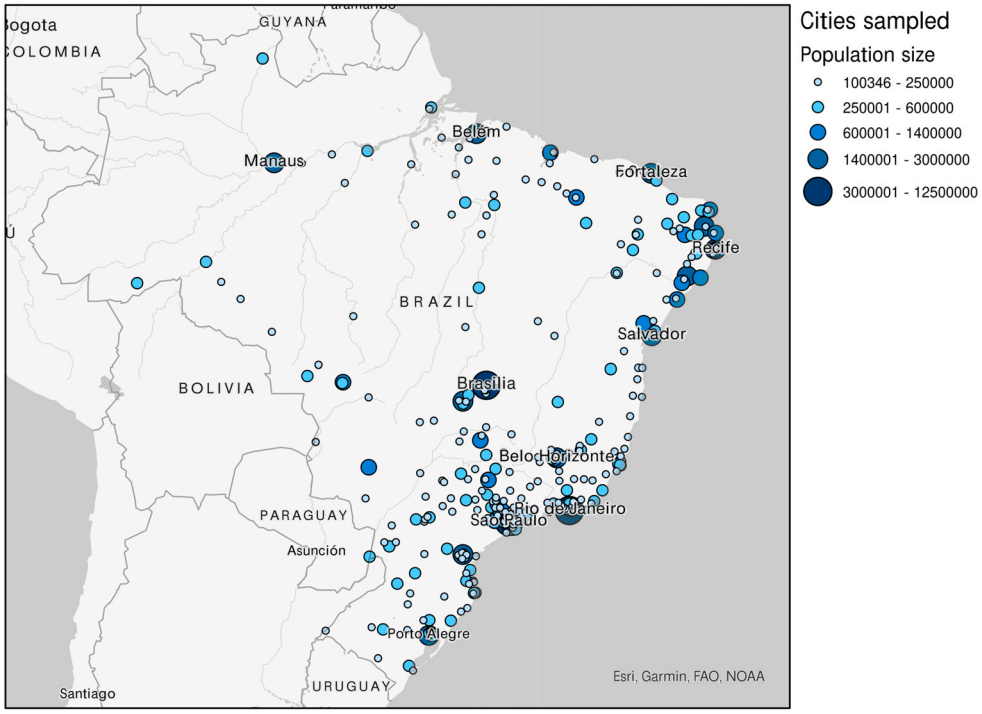

Fig. 2. Brazilian cities applied in our analysis with at least 100 000 inhabitants.

CDD and HDD derived from either heat index, wind chill or apparent temperature.

Brazil has comparatively high climate diversity. This is in part due to the size of the territory, the extent of its coastal regions, the variation in elevation, and the presence of different air masses that influence the temperature and humidity of each region. Moreover, according to [43], Brazil possesses six distinct climate zones: equatorial, tropical, semi-arid, coastal, subtropical, and tropical high-altitude. This climate diversity, coupled with demographic heterogeneity, would make building a unified model for the entirety of Brazil a challenging task. By taking advantage of the political macro-regions and examining each of them individually, a more manageable and interpretable analysis can be made. This approach results in each region of interest encompassing at most two climate domains. Fig. 3 illustrates a two-dimensional reduction of the evaluated daily weather variables employing the TSNE algorithm [44], in which each point represents a day in our dataset and point color corresponds to a region and its mean temperature.

We observed that each region's weather appears to be well-defined, with a clear separation between most regions with little overlap. This validates our analysis of distinct regions and the proposed city-to-region aggregation approach for climate variables. Given that the main focus of this work is on the impact of temperature on energy consumption, we also seek to understand how temperature impacts this separation, and, after visualizing these measures, the relationship between geography and temperature is evident. Furthermore, we can also draw some conclusions regarding the intersection between regions. From the temperature plot in Fig. 3, we observed that colder days in the Southeast are similar to the hotter days in the South. An analogous pattern can be identified between the Southeast and the North, while Northeast temperature-days appear to be the most distinct. Perhaps related to this, in our experiments we found that predicting consumption for the Northeast region presents the most challenges.

Finally, we also included data from various Coupled Model Intercomparison Project Phase 6 (CMIP6) models to forecast demand under future climate change, ranging from SSP1-2.6 to SSP5-8.5 [39], available via the Google storage API and managed by Pangeo and the Earth System Grid Federation. SSP-RCP (shared socioeconomic pathway-representative concentration pathway) narratives were designed to capture possible future scenarios for energy demand due to temperature and the shared socioeconomic pathways represent varying assumptions regarding future global development. Each of the five SSPs corresponds to a projection of greenhouse gas emissions, namely: SSP5 (fossil-fueled development), SSP4 (inequality), SSP3 (regional rivalry), SSP2 (middle-of-the-road development), and SSP1 (sustainable development). By design, SSPs were devised to work in conjunction with the representative concentration pathways. The various pathways relate to the levels of radiative forcing by the year 2100 and range from 1.9 to 8.5 W/m^2, in which higher values depict greater levels of anthropogenic influence. Although various combinations of SSP and RCP are possible, many do not lead to feasible scenarios of interest. We employed four Tier 1 SSP-RCP scenarios from the Scenario Model Intercomparison Project (ScenarioMIP) within CMIP6, SSP1-2.5, SSP2-4.5, SSP3-7.0 and SSP5-8.5 [45,46].

However, individual models can introduce distinct biases driven by their methodologies and the represented physical phenomena. As such, evaluation from a single model might lead to contradicting results. Applying the means from a group of models can allow for a better evaluation of the overall trend of the forecasted RCP-SSP [47,48]. Specifically, we considered an ensemble for each pathway compromised by the models ACCESS-CM2 [49], AWI-CM-1-1-MR [50], CESM2-WACCM [51], CMCC-CM2-SR5 [52], CMCC-ESM2 [53], CanESM5 [54], EC-Earth3 [55], GFDL-ESM4 [56], IITM-ESM [57], INM-CM4-8 [58], INM-CM5-0 [59], IPSL-CM6A-LR [60], MPI-ESM1-2-HR [61], MRI-ESM2-0 [62], and NorESM2-MM [63], sampling maximum and minimum near-surface air temperature (tasmax and tasmin), precipitation flux (pr), near-surface eastward and northward component of wind (uas and vas), and near-surface relative humidity (hurs). Like in the ERA5, we sample data using the coordinates from each Brazilian city with at least 100,000 inhabitants and compute the weighted values for each variable, enabling for comparison with past historical values. As depicted in Fig. 4, climate change is expected to occur at near-constant relative humidity at a global scale [64], although some small regional trends may be present. Fig. 4 shows that perceived temperature will increase throughout the century, largely driven by near-surface air temperatures. No clear pattern was identified concerning wind speed and precipitation.

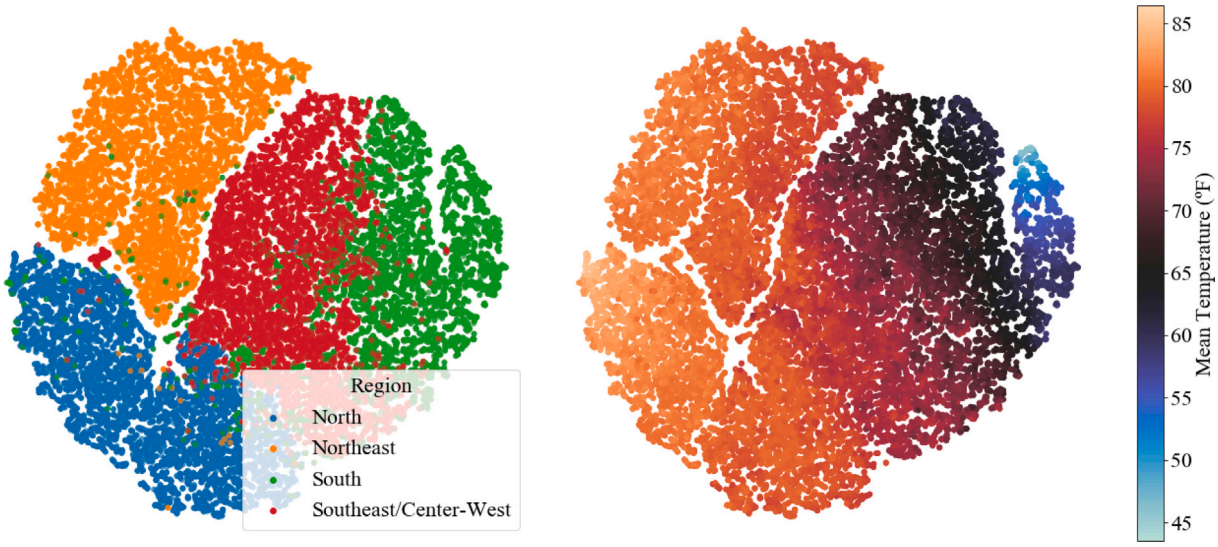

Fig. 3. TSNE representation of Brazilian weather. Each point represents a day in one of Brazil's sub-regions. There appears to be a relationship between temperature and geographical location that produces well-defined partitions in the plot on the left without explicitly performing clustering.

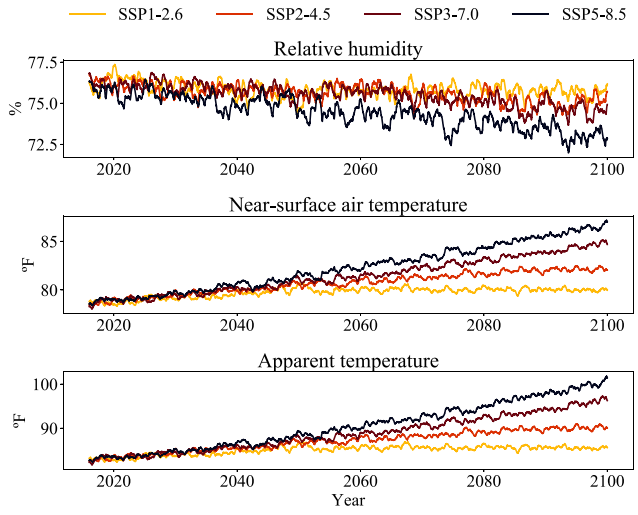

Fig. 4. CMIP6 projections for temperature and humidity, and the resultant increase in apparent temperature.

3.2. Counterfactual modeling and forecasting approach

Our main objective is to predict electricity consumption in the absence of any abnormal events, such as public health emergencies, government interventions, and extreme weather events. By excluding these abnormal events from our prediction model, we can estimate their contribution to changes in consumption. We formulate this problem as a regression. Given a set of $w \in W$ weather descriptors and a set of $t \in T$ time descriptors, we apply a function $f(w;t;\sigma)$ parameterized by σ that maps to a period of electricity consumption. To exclude disrupting factors, we search for optimal $W' \subset W$ and $T' \subset T$.

Counterfactual Construction: Guided by existing literature [65], we have identified three major drivers of electricity consumption: load growth (e.g., increase in population, GDP, industrialization, etc.) [66,67], historical events (e.g., pandemics, large-scale changes in government policies, noteworthy events, etc.) [27], and weather events (e.g., heatwaves, droughts, etc.) [68].

For instance, we observe that yearly electricity consumption follows a logistic growth trend as exemplified by Fig. 5(a). When plotting the daily energy curves in Fig. 5(b), a similar pattern can be seen as the shape of the curves remain relatively similar but there exists a step between subsequent years — the exception being the year 2020 which diverges from this trend. To develop a counterfactual model that specifically examines the influence of weather and temporal factors on electricity consumption, we employ a normalization technique by adjusting the daily consumption data with a load growth function. This function is obtained through yearly load interpolation. By normalizing the data in this manner, we mitigate the impact of factors such as population growth, consumer purchasing power, GDP, and industrialization on load growth. Furthermore, to enhance the precision of our analysis, we exclude from the training data any periods associated with notable historical events, such as the COVID-19 pandemic in 2020. This approach allows us to isolate the contribution of weather and temporal factors on electricity consumption, while minimizing the confounding effects of the two other major driver groups. Fig. 5 illustrates the consumption for the Southeast/Center-West region after normalization. Consumption for the remaining regions can be found in Supplementary Figures 2–5.

Several machine learning methods were employed in this study which are often used in literature concerning energy prediction [12], including LightGBM, Multivariate Adaptive Regression Splines (MARS), Linear Regression, and Support Vector Machines (SVM). LightGBM [69], short for Light Gradient Boosting Machine, is a tree-based ensemble learning algorithm that uses gradient boosting to optimize the training process, making it efficient for large high-dimensional datasets such as in our proposed problem. MARS [70] is a non-linear regression technique that uses piecewise linear functions to model the relationships between the predictor and response variables. Linear Regression [71] is a widely used method that models the relationship between the predictor and response variables as a linear equation. SVM [72] works by finding the optimal hyperplane that best separates the data and predicts the target.

In addition, both first and higher order features were applied in this study. First-order features refer to the original predictor variables used in the analysis, while higher-order features are constructed by combining pairs or interactions of the original features. These higher-order features capture more complex relationships and interactions, allowing for a more nuanced analysis. Due to the exponential computational constraints of higher-order feature sets, we limit ourselves to only second-order features. Since some of the aforementioned algorithms are sensitive to domain range, we normalize data prior to training the Linear Regression, MARS and SVM in a one-year-walk-forward

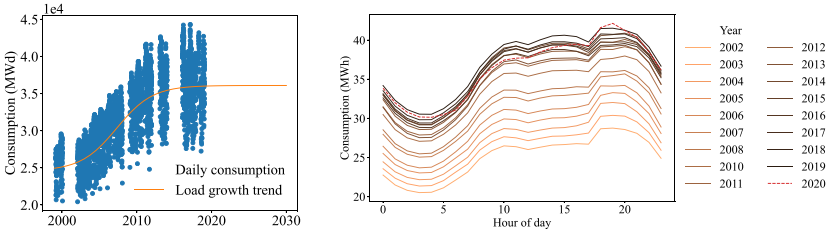

(a) Daily consumption. (b) Mean consumption curves.

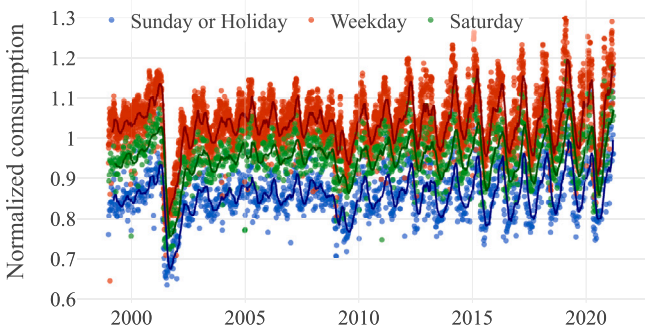

(c) Daily consumption normalized by load growth.

Fig. 5. Consumption and load for the Southeast/Center-west subsystem. Data for the remaining regions can be found in the Supplementary Figures 2–5.

validation. That is, we train the data on one year while using it to establish the normalization ranges, and validate our results on the following years, as to avoid data leakage. Since there is no missing data, imputation was not needed. However, a key consideration is the selection of optimal features.

Feature Selection: As not all variables are relevant for forecasting, we represent the model space as a directed acyclic graph (DAG) and each node depicts a model built from a distinct feature subset. Let A and B be two nodes representing two distinct models. Vertex $A \rightarrow B$ exists if B can be reached by simple feature addition from A, thus constituting a transitive reduction of the complete model space. This approach presents two desirable properties: (i) any vertex is reachable from the [$\emptyset$] model, and (ii) there exists a topological hierarchy with an ordering of all vertices into a sequence such that for every edge their is a beginning vertex that occurs earlier in the sequence than the ending vertex of the edge for any feature set path. We propose that there exists a set of optimal feature expansions that lead to performant models [73]. These properties imply a partial ordering starting from the root node, enabling the search of the $N!$ combinatorial space. In our experiments, we apply the greedy beam search algorithm [74], employing as a heuristic the R^2 of the model represented by a vertex and a beamwidth $\beta = 5$.

The beam search algorithm is an optimization of the best-first search parameterized by a beamwidth β. At each iteration, the best β nodes are considered expansion candidates, and all of their children are evaluated. If $\beta = 1$, then beam search behaves as best-first, in which only the best child in each iteration is expanded. If $\beta = \inf$, then beam search becomes a breadth-first search in which all possible expansion paths are considered. This allows us to search the $N!$ combinatorial space of feature subsets to select the best-performing specialized models. In our experiments, we employ a beamwidth β of 5.

Extrapolation: Although extrapolation beyond the range of available data can produce unreliable results [75], it can be useful for gaining insights into future scenarios. In our study, we employed a method called Data-Augmented Regression (DAR) [14,15] to extend the temperature range beyond the limits of our training data. We augmented our original training data by 10% through 1872 Monte Carlo simulations uniformly sampling temperatures between 50 °F and 120 °F (10 °C and 65 °C), and used a linear regression as a background model. DAR is a generalized regression algorithm that can be used to approximate the penalty function in situations where it is not possible to evaluate it analytically. In this method, an augmenting set of random data points is generated and added to the observed data, which is then used to train the model.

Mathematically, this approach involves generating additional N data points that are used to augment the original training set. Given a set of input values x and a corresponding output value y, new data points $\tilde{x}$ and $\tilde{y}$ are generated using a Monte Carlo approximation. Specifically, $\tilde{x}$ is drawn randomly according to some underlying probability distribution $\mu(x)$ (i.e., uniform), and $\tilde{y}$ is generated by drawing from the conditional distribution $p(y|\tilde{x})$ (i.e., a background model). These new data points are added to the original training set. The resulting data-augmented regression approach can be used to estimate a function f that extrapolates beyond the range of the original training data. The method involves minimizing a loss function on the augmented training set, which includes both the original and augmented data. In summary, we used the DAR method to expand our training data and cope with temperatures outside the range of our original data. By extending the domain, we were able to gain insights into future scenarios, albeit with some limitations.

4. Results

We focused on answering RQ1 in our first set of experiments, with the main results summarized in Table 1. To understand the drivers behind the performant model, we applied the SHAP algorithm [76] to obtain an interpretation of the impact of each input variable. We assessed the statistical significance of our measurements through a pairwise t-test with p-value ≤ 0.05 and a one-year walk-forward validation. An exhaustive grid search was employed for hyperparameter tuning. The performant model was found to be a LightGBM [69] optimizing the L2 loss function, while the quantile regression optimizes the pinball loss function. The learning rate was set to $5e^{-}2$, with 64

Table 1
Baseline comparison of algorithms. Highlighted cells represents statistically superior results.

	Model	N		NE		S		SE/CW		Overall	
		R^2	MAPE	R^2	MAPE	R^2	MAPE	R^2	MAPE	R^2	MAPE
First order features	Linear regression	.610	2.57%	.738	2.40%	.875	3.23%	.826	2.59%	.788	2.76%
	GAM	**.621**	**2.52%**	.712	2.55%	.875	3.21%	.831	2.57%	.783	2.78%
	MARS	.614	2.58%	.723	2.52%	.875	3.16%	.824	2.53%	.787	2.78%
	SVM	.261	3.53%	.529	3.44%	.786	4.49%	.679	3.91%	.616	3.95%
Second order and logarithm transformations	Linear regression	**.621**	**2.52%**	.721	2.50%	.878	3.19%	.832	2.54%	.790	2.74%
	GAM	.618	**2.54%**	.713	2.55%	.877	3.17%	.831	2.58%	.784	2.79%
	MARS	**.624**	**2.51%**	.713	2.55%	.877	3.16%	.834	2.53%	.788	2.73%
	SVM	.313	3.39%	.566	3.33%	.821	4.11%	.710	3.61%	.649	3.72%
	LightGBM	.610	2.61%	**.773**	**2.33%**	**.887**	**3.05%**	**.850**	**2.50%**	**.854**	**2.64%**
	LightGBM + Linear extrapolation	.599	2.78%	**.773**	**2.33%**	**.885**	**3.07%**	.845	2.55%	.848	2.69%

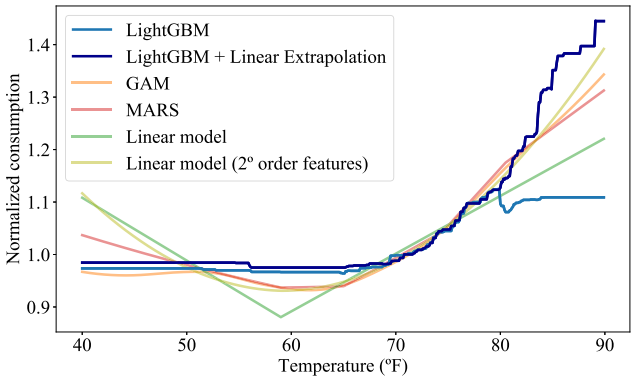

Fig. 6. Consumption simulation while varying temperature. Southeast/Center-west 2020-11-16, weekdays.

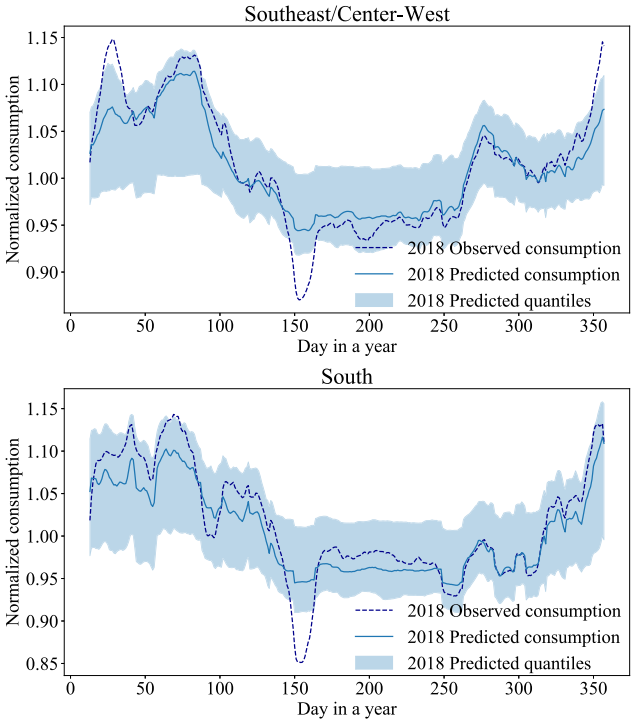

Fig. 7. In May 2018, the Southeast/Center-west and South subsystems experienced a significant decline in energy consumption due to a truck drivers' strike that disrupted distribution lines, resulting in temporary service interruptions and the population staying at home due to gas shortages between May 21 and June 1, 2018.

bins and training 100 trees with a maximum depth of 50 and 30 leaves. However, as mentioned previously, tree-based algorithms are known to not extrapolate well.

Fig. 6 presents a simulation in which we only altered temperature and observed the deviation in energy consumption. The LightGBM model trained on augmented data from a linear regression was able to extrapolate consumption past the 80 °F threshold, at which point the standard LightGBM produced an unrealistic flat consumption curve. However, we can expect a trade-off between extrapolating and predictive power. Table 1 presents the performance in terms of R^2 and MAPE for all evaluated methods in each Brazilian subregion and on the full dataset. The LightGBM approach outperforms all other methods in the overall scenario. Thus, introducing a dependence on the linear regression model, which presents a larger percent error – via the DAR algorithm – likewise reduces performance. Nevertheless, from Table 1, we verify that this loss in predictive power is minimal, with some regions presenting statistically similar performance to the standalone LightGBM, thus highlighting the ability to retain the robustness of gradient-boosting machines and the extrapolative capabilities of linear models. For more information, see Supplementary Figure 6 which depicts the SHAP importance plot and provides interpretable insight into the drivers of consumption for Brazil.

4.1. Atypical historical event impacts

To understand the impact of events on energy consumption, we compared the predicted consumption from our counterfactual model with the observed consumption. To estimate sensible ranges, we performed quantile regression through the pinball loss function to obtain the 5% and 95% quantile predictions. We have provided three relevant case study periods for analysis regarding RQ2. The first case study period is 2018 in which the Brazilian economy recovered after a long recession during a period with no severe heatwaves, droughts, or other extreme weather events of note, serving as a suitable year for baseline comparisons. Fig. 7 illustrates the consumption for the most developed region in Brazil for this time period. The only anomaly observed during 2018 refers to the truck drivers' strike that lasted from May 21st to May 30th, resulting in the disruption of entire supply chains across Brazil, empty gas stations, and people forced to work from home due to limited transportation options. The dramatic impacts of this strike on consumption are easily identifiable in Fig. 7.

The second case study period is the year 2020, in particular during the early stages of the COVID-19 pandemic which resulted in substantial restrictions on mobility and a reduction in economic activity that was also experienced on a global scale. In Fig. 8 we compared consumption with the Oxford Stringency Index, which measures the strictness of COVID-19-related government policies that restricted certain population behaviors [77]. During the early COVID pandemic, we observed a close relationship between the drop in consumption and the Oxford Stringency Index for Brazil. The exception is the South region, which overall was less impacted by the pandemic. One explanation for this is due to the region's economic dependency on tourism and lower

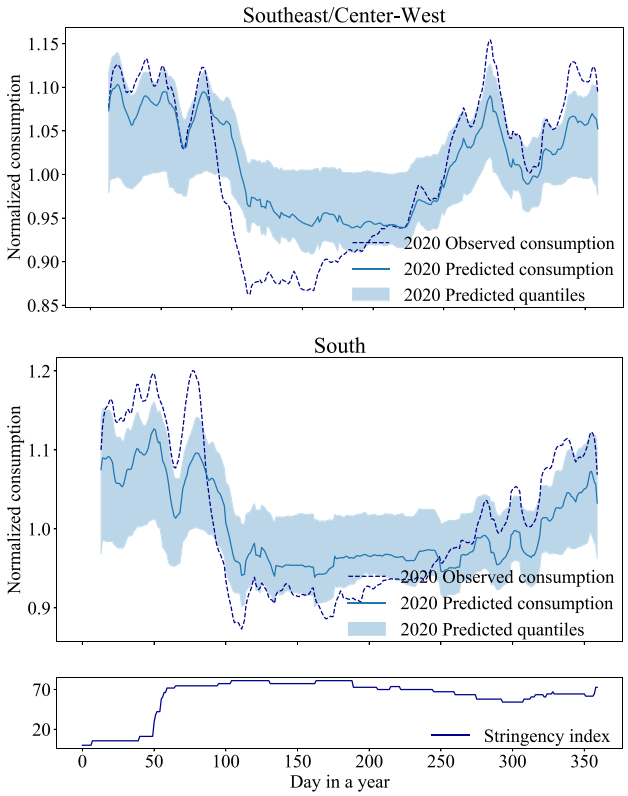

Fig. 8. As illustrated by the Oxford Stringency index [77,79], Brazil 2020 COVID-19 lockdown policies started around early March, which resulted in a drop in energy consumption. The population in the South had less adherence to lockdown policies and was not as impacted [79].

levels of population concern about the COVID-19 pandemic, with most restrictions lifted by April [78].

Our third case study spans from 2001 to 2002. After a long drought with low reservoir levels in early 2001, system operators were concerned that the electrical grid might collapse. To avoid a worst-case scenario, the Federal Government enacted a series of policies and programs aimed at reducing energy consumption by 20% to reduce grid strain [17]. These ranged from energy conservation awareness campaigns to time-based electricity rate pricing, with the intent of reducing electricity use during peaks periods and keeping all non-essential energy consumption to a minimum. An advantage of examining this period is the presence of an exact measure of expected impact, which is uncommon in counterfactual literature and allows for a direct evaluation of our approach. Fig. 9 highlights the impacts of these policies and programs across all regions except for the South which was excluded from these government initiatives.

Fig. 9 also displays residual plots that illustrate the difference between the actual energy consumption during an energy crisis in Brazil, the 2001 Apagão represented by solid lines, and the predicted energy consumption based on our counterfactual represented by dashed lines. The Apagão resulted in a notable reduction in energy consumption, as indicated by the negative residuals when compared to the expected counterfactual level. These negative residuals suggest that the rationing policies and planned outages implemented during the period had a substantial impact on energy consumption. Furthermore, the close match between the observed residuals and the expected 20% reduction in consumption confirms the effectiveness of our counterfactual model in predicting electricity consumption from weather variables and capturing the impact of exogenous extreme events.

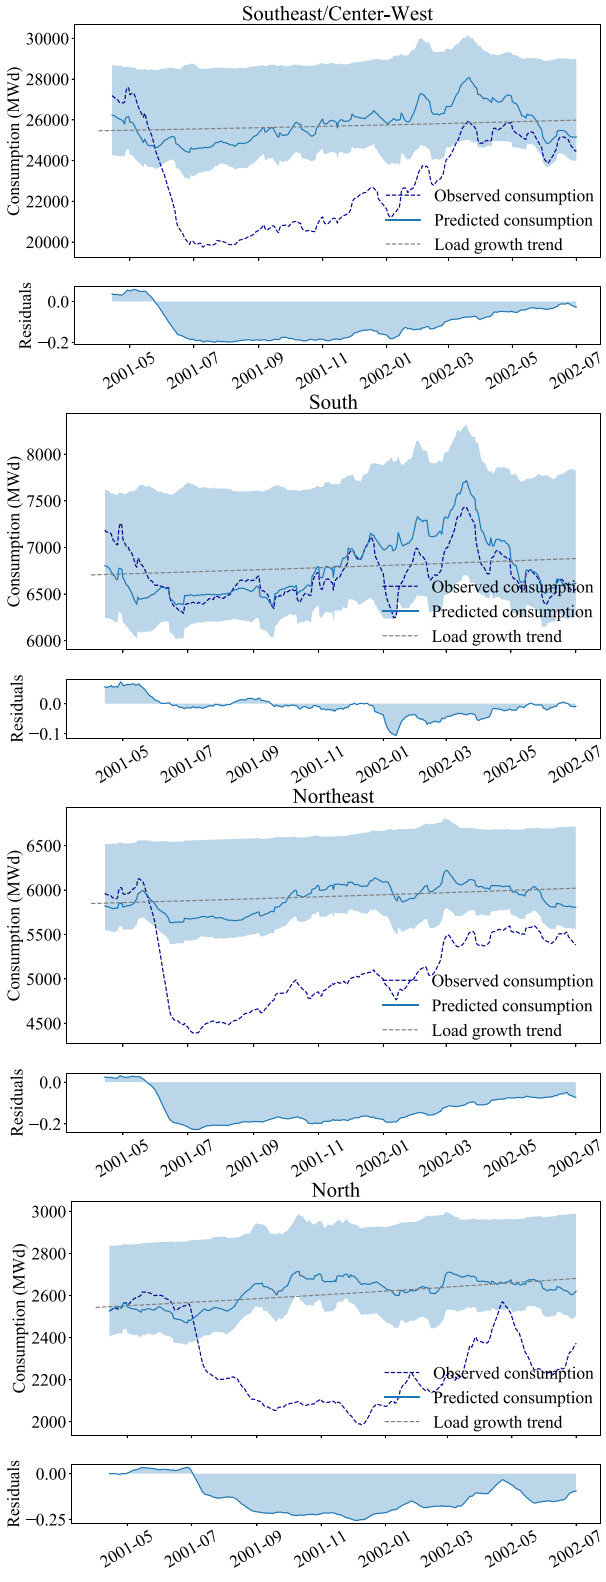

Fig. 9. In early 2001, Brazil experienced an energy crisis (*Apagão*, Blackout) due to reduced rainfall and inadequate infrastructure investments. In response, the government implemented nationwide measures, that only excluded the South, such as voluntary energy consumption reduction, electricity supply restrictions for industrial and commercial sectors, and efficiency programs. This led to a sharp drop in consumption until 2002 as pictured by the dashed line and the residuals.

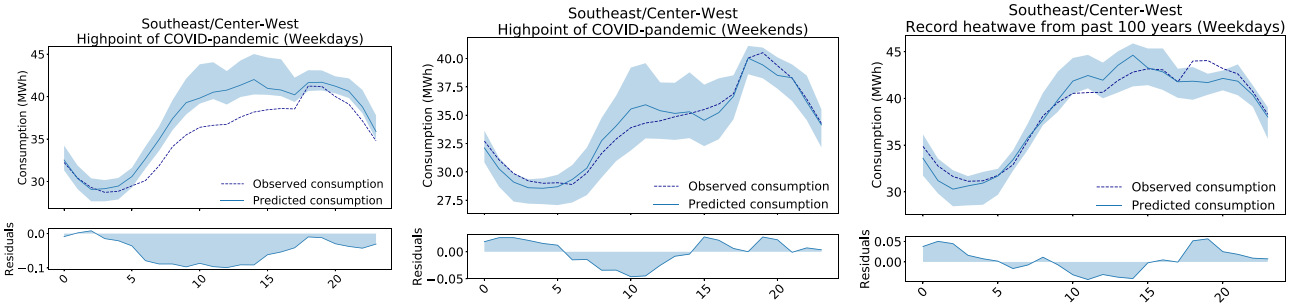

Fig. 10. Daily consumption shapes at an hourly level for weekdays and weekends in March and October 2020 for the Southeast/Center-west subsystem. During COVID-19 lockdowns, weekdays exhibited similar patterns to weekends, a possible effect of the lockdown policies. In October 2020, a significant heatwave caused an abnormal increase in energy consumption outside commercial hours, possibly due to increased housing cooling demand coupled with more relaxed lockdown policies, as the population began returning to workplaces that were usually cooled during the earlier parts of the day.

4.2. Hourly consumption shape impacts

To understand the change in consumption on a hourly level, we built an hourly counterfactual modeling approach following the same strategy as the previously presented daily consumption model. We validated our model on all regions using data from 2015 to 2019, verifying a mean percent error of <5%. Our experiments evaluated the change in consumption during the aforementioned atypical events. In particular, we focused on 2020 (see Fig. 10). The first half of 2020 was characterized by the initial onset of the COVID-19 pandemic and government restrictions intended to limit its spread. The second half of 2020 presents two overlapping events, the first being the relaxation of COVID-related government restrictions, and the second a record-breaking heatwave in October.

Our hour-by-hour energy consumption counterfactual analysis produced additional insights that augmented our daily experiments. During weekday commercial hours in early 2020, we observed a substantial drop in energy consumption. As shown in Fig. 10, a comparison of counterfactual predictions with actual consumption revealed negative residuals close to 10%. Furthermore, weekday consumption shapes matched those of weekends, with the most significant portion of the residuals being situated during commercial hours. This suggests that some of the primary drivers of this drop were the effects of COVID-19 lockdowns with closed and/or restricted commerce. We also found that the residuals were not as pronounced on weekends when businesses are typically closed, further supporting this interpretation.

Later in the year, Brazil experienced record-breaking heatwaves in October, which corresponded with positive residuals in our analysis. Both of these patterns, from early and late 2020, were already observed in Fig. 8, but the hour-by-hour counterfactual models provided a more in-depth understanding of the dynamics of each phenomenon. During October 2020, we observed small residuals during commercial hours, while in the later portion of the day, we observed positive residuals indicating an abnormal increase in consumption. One potential explanation is both a relaxation of lockdown policies and the need to cool homes in addition to already cooled workplaces. This allowed the hourly weekday shapes to regain their usual pattern with a comparative increase only in the later portion of the day.

We proceeded to analyze the relationship between errors and actuals. In doing so, we were not able to verify a clear pattern between temperature and predictions, thus suggesting that our model is robust to temperature changes, although we observed a larger error range near the middle portions of the distributions for the Southern and Centerwest regions. We did, however, identify a bias in predictions between midnight and 6 AM, in which the models appear to under-predict more often than in other hours of the day. One possible explanation could be related to sleep habits and the end of commercial hours in Brazil, which the model appears to be overemphasizing in late-night and early-morning predictions. This suggests that additional indicators of behavior may be needed to more accurately model consumption during this period.

4.3. Power outages and consumption impacts

Extreme events such as the COVID-19 pandemic may have longer-term impacts which, as demonstrated above, can be evaluated from daily energy consumption measures. However, very brief or short-duration events might not provide the same consumption signature. This is true for most outages, given that the recovery time can happen within a few hours. In this respect, using a daily consumption model might lead to an under-prediction of impact. For events such as these, we applied the hourly counterfactual model. Using the same methodology previously described, we can measure an outage impact by computing the residuals between the counterfactual model output and the observed consumption. Specifically, we computed the area between these curves and use it as a metric for the impact of the estimated loss of electricity consumption. Additionally, this allows us to measure grid recovery. Even if the utility indicates that electricity is completely restored, it might take some time for the grid to stabilize, and this approach accounts for this potential delay.

One of the largest outages in Brazil's history was the 2009 Brazil and Paraguay blackout. According to Brazil's Ministry of Mines and Energy, adverse "atmospheric factors" caused the failure of three transmission lines from the Itaipu Hydroelectric Power Plant, which resulted in its complete shutdown on November 10th, 2009 [80]. The massive blackout affected not only 18 of Brazil's 27 states, but also the entire population of Paraguay. As a result, four Brazilian states and 90% of Paraguay were left in complete darkness. It was reported that the outage started at 22:13 GMT-3 and that the Itaipu power plant returned to normal operations at 6:00 GMT-3 the following day.

However, there is some debate regarding recovery time from this outage event. For example, [81] claimed that power was restored at 2:45 GMT-3 as reported by national news media at the time. However, Globo, one of the largest news portals in Brazil, claimed that the southeast region began to recover after 3 h of blackout [82], therefore at 1:00 GMT-3. Fig. 11 illustrates the energy consumption for the Southeast and Center-west regions which contains all four states that experienced outages during this event. Indeed, at 1 AM the start of the recovery can be seen with consumption only returning to pre-outage levels between 6:00 and 7:00 GMT-3.

Another wide-scale outage occurred on March 21st, 2018, affecting all regions of the country, with greater intensity and duration in the North and Northeast. According to the Brazilian Electric System Operator, human failure was responsible for the outage, which resulted in an overload of the electric grid and eventually, its collapse [83]. This also raised concerns regarding the security of the power grid [84]. The outage began at 15:48 GMT-3 when the transmission line connected to the Belo Monte Power Plant failed after not being able to support an increase in load. Belo Monte lies in the Northern state of Pará close to the border of the Northeastern region. However, by 16:15 GMT-3 electricity was already restored in the Southeast and

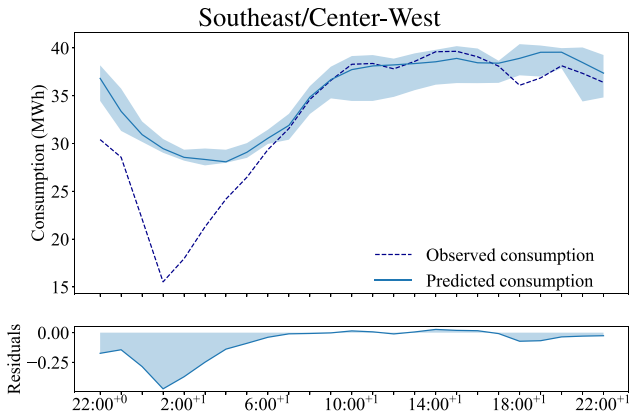

Fig. 11. The 2009 Brazil and Paraguay blackout on November 10th, 2009. Power was restored between 1:00 and 6:00 GMT-3, and affected an estimated 90 million people in Brazil.

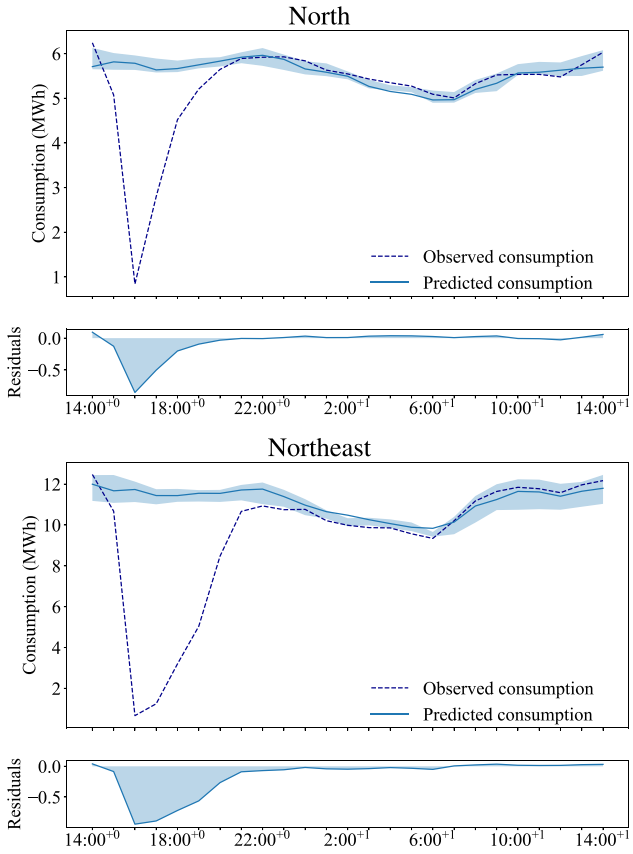

Fig. 12. The 2018 North and Northeast outage in Brazil, on March 21st, 2018. The energy was restored between 16:00 and 21:00 GMT-3 and affected at least 80 million people in Brazil. We estimate a loss of 12 MWh and 36 MWh respectively by computing the area between predicted and observed consumption.

Center-west region, thus these regions were not severely affected. The same cannot be said of the regions that depend on Belo Monte more heavily. Fig. 12 shows the consumption for March 21st, 2018, in the North and Northeast region. We verified a sharp drop in consumption at the 16:00 h mark, coinciding with the beginning of the outage. Recovery followed immediately after, with an estimated total grid recovery occurring shortly after 21:00, when consumption returned to pre-outage levels. However, our analysis hints at some instability later in the day in the Northeast region. This can be evaluated by comparing the solid and dashed lines in each plot. Computing the area between observed and counterfactual consumption, identified by the residuals, we obtained approximately a loss of 200% (12 MWh) and 300% (36 MWh) of energy for the North and Northeast regions respectively.

4.4. Forecasting consumption under a changing climate

Earth's mean temperature is expected to increase by the end of the century compared to pre-industrial levels, and how much of an increase will depend on future actions taken by human populations to mitigate carbon emissions. Climate change is expected to also have an effect on other climate variables like precipitation and drought. All these changes in climate are likely to impact energy usage, as there is a close relationship between weather and electricity consumption. However, not all regions will be subject to the same climatic changes and there exist many possible futures depending on different climate change projections. To account for this variability, we propose forecasting energy consumption under 106 distinct CMIP6 models encompassing SSP1-2.6, SSP2-4.5, SSP3-7.0, and SSP5-8.5 scenarios, building ensembles for each pathway that project future precipitation, temperature, humidity, and wind speed.

To address RQ3, we explored how each of Brazilian regions' energy consumption is expected to change under the different warming scenarios in Fig. 13. North and Northeast, the hotter regions, present a larger increase in energy usage than other regions. This conforms with the expected temperature increases for the continent, as the northern portion is expected to become significantly hotter than the southern portion. When comparing one subregion to another, we can further explain the disparity in electricity usage due to heat index. The Northern region encompasses the Amazon rainforest and has substantially higher humidity than the primarily semi-arid Northeast. As such, the effects of high temperatures are exacerbated in the form of larger increases in perceived temperatures.

We also considered the overall consumption increase for the whole country. Each pathway provides minimum and maximum temperature ranges, which we employed to obtain the consumption range in Fig. 14. Nearly all CMIP6 models provide similar minimum near-surface air temperatures (tasmin), regardless of the pathway. The same pattern, however, is not observed for maximum near-surface air temperatures (tasmax). As the scenarios become more severe in terms of projected climate change, the mean maximum temperature increases and there is a larger variance near the extremes of the distributions.

For the next five years (2021–2026), Brazilian's Energy Research Company (EPE), National Electric System Operator (ONS) and Electric Energy Trading Chamber (CCEE) expect a 3.6% and 2.7% yearly load and power grid expansion respectively [7]. Fig. 15 presents a simulation of daily consumption through the year 2070 assuming that Brazil retains its 60% hydro-based energy supply. We presented 360 rolling day averages to suppress seasonality and evaluate consumption tendencies. Although Brazil currently has 100 GW+ of total installed hydro capacity, it is unrealistic to assume that any hydropower will be able to constantly operate at 100% efficiency. Itaipu, the largest hydropower plant in Brazil and responsible for supplying over 10% of country-wide consumption, has historically operated at an average of 61.15%[1] power over the past 36 years. Thus, we considered a mean of 60% operating hydropower to estimate Brazil's total grid capacity. To the best of the authors' knowledge, official agencies only provide short-term predictions spanning five years and not longer-term ones. Nevertheless, extrapolating load growth using the SSP5-8.5 forcing scenario yielded precise short-term projections consistent with the official reports. However, as presented in Fig. 4, all studied atmospheric variables' projections are statistically similar until at least 2035 and all SSP scenarios led to similar short-term energy forecasts.

[1] Available in Portuguese at https://www.itaipu.gov.br/energia/geracao.

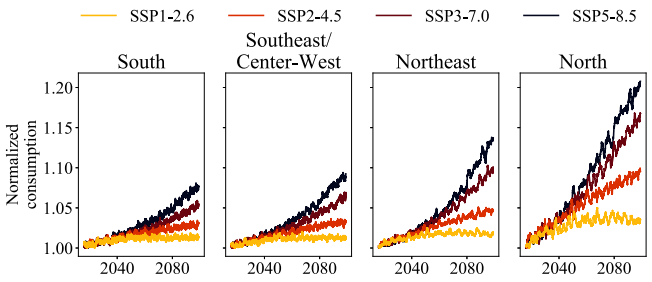

Fig. 13. Normalized consumption forecast for each Brazilian subsystem under all evaluated RCP-SSP scenarios, considering the means of the respective CMIP6 ensembles.

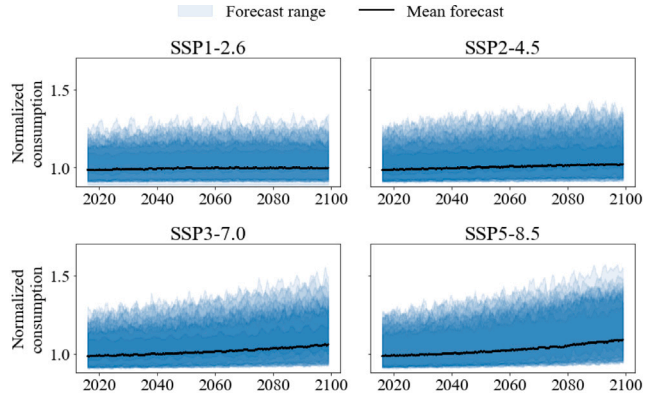

Fig. 14. Normalized consumption forecast for each evaluated SSP-RCP pathway on the entirety of Brazil. Each area illustrates the consumption forecast range from an individual CMIP6 model.

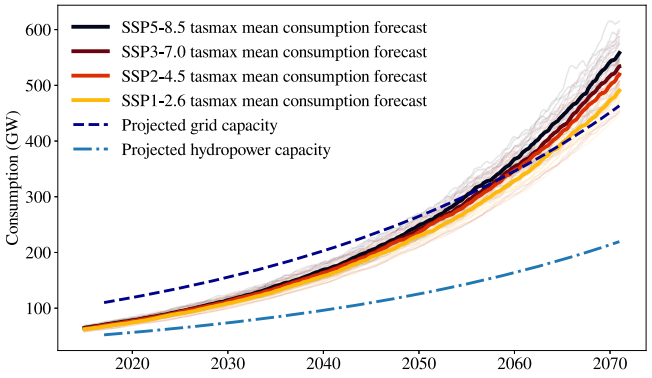

Fig. 15. Electricity consumption forecast for Brazil through 2070. Light-colored lines represent forecasts from individual CMIP6 models, while bold lines illustrate the ensemble mean.

5. Discussion

In our first set of experiments, our objective was to identify the best algorithm to build a counterfactual model of consumption. LightGBM, although performant, failed to extrapolate beyond training data. This meant it was unable to cope with unprecedented extreme weather events, such as heatwaves likely made more severe due to global warming. We then proposed augmenting the training data with the predictions from a linear regression model to generate temperatures outside the training scope. This enabled the LightGBM model to handle these previously unseen scenarios albeit at a small loss of performance.

In all experiments, the North subsystem remained the most challenging to predict. Before 2014, the North subsystem was not part of Brazil's integrated energy system, leading to a scarcity of data and inducing two distinct consumption distributions, as shown in Supplementary Figure 1. To counteract this, we filtered data pertaining to the North before 2014 leading to a scarcity of data. When exploring the drivers of consumption from SHAP, the most important feature identified was whether a given day was a weekend/holiday, with weekdays being associated with higher consumption. We observed that further stratifying by day of the week was helpful, as Saturdays tended to have higher consumption than Sundays, and Mondays lower than other weekdays, a characteristic also found for European electricity consumption patterns [85]. Since SHAP assumes feature independence, co-related features can share credit for importance. This was found to be particularly true for the temperature-derived variables, as they share a causal relationship. Nevertheless, we found that CDD remains one of the most impactful drivers, indicating a nearly proportional relationship that may be due to Brazil's warmer climate.

Some insights into disruptive event impacts emerged when comparing the counterfactual output with the observed consumption in 2001 when the Federal Government enacted a series of policies to reduce consumption by 20%. From our counterfactual model, the mean and median relative residuals for the Center-West region between June 2001 and January 2002 are $-18.1\%(\pm 1.6\%)$ and -18.6%. This matches the expected -20% reduction for the period in which the restriction policies were in place, from July 1st, 2001 to February 19th, 2002. A similar pattern was also observed in the North and Northeast regions with -19.1% and -18.9% relative residuals.

When evaluating the beginning of 2020 in Fig. 8, we observed that our counterfactual predictions closely match the observed consumption. In March 2020, state Governors started issuing recommendations about social distancing during the initial stages of the COVID-19 pandemic. During this initial period, electricity consumption decreased with a pattern that remained mostly constant from April to June 2020, with a mean $-8.87\% \pm 1.2\%$ residuals. However, by July 2020 we found that recovery to pre-pandemic levels coincided with the relaxation of COVID-19-related restrictions. Although more severe restrictions were imposed again in the latter part of the year, consumption was not as significantly impacted as it was during the initial stages of the pandemic.

We identified a clear distinction in behavior during weekdays and weekends and proceeded to evaluate them separately. For example, during the COVID-19 pandemic there was a larger margin of error between observed consumption and counterfactual output during weekdays. This finding aligns with other research assessing mobility and consumption change during the early pandemic, in which many commercial and industrial sectors either closed or adopted a home-office strategy. We proceeded to examine the daily consumption shapes and found that shapes during the weekdays tended to match weekends. Due to the higher levels of daily life restrictions indicated by the Oxford Stringency Index and the adoption of home office policies, we observed that weekday consumption patterns became more similar to weekends compared to pre-pandemic time periods. For most regions, there was also a change from a consumption peak near mid and late afternoon to a steady ramping of consumption into the evening hours. This phenomenon was particularly noticeable for the Southeastern-center western subsystem which is the most industrialized region.

Regarding heatwaves, we observed a compound phenomenon of the later-stage pandemic and higher temperature conditions. Many commercial establishments had reopened at this time but a portion of the population remained in home office. This likely contributed to commercial and industrial consumption returning to pre-pandemic levels, but residential consumption increasing. One potential explanation is an extreme heat event coupled with the record increase in energy-intensive cooling systems adoption (e.g., air conditioning), as well as less adherence of the population to COVID-19 restrictions [86].

Next, we examined the monthly consumption profiles for each region in Brazil. For all regions, the major sectors are commercial, industrial, and residential loads. We observed that the consumption

profile of the Southeast/Center-West region closely matches that of the South regions. This could explain why the model performed well on both of these subsystems. For the Southern region, we found that a significant part of the load was tied to agribusiness which presented a strong seasonality pattern. The Northern and Southeast regions likewise showed this same pattern, which was not seen in the North region. This was expected given the extractive focus and the large proportion of the Amazon rainforest in the region. Supplementary Figure 17 illustrates each region's profiles.

Our final experiments explored future periods by applying forecasting from CMIP6 projections. We found that given Brazil's current sources of generation, it will not be able to sustain its consumption, and it will face shortages sooner than initially thought even under optimistic assumptions. Under these scenarios, we can expect a considerable reliance on non-hydro energy sources by 2030 to meet demand. This suggests a need to promote renewable energy alternatives, as currently, the next most available energy resource in Brazil is carbon-intensive thermal power stations. Additionally, under the SSP5-8.5 forcing scenario, we projected that Brazil will be unable to meet its consumption by 2070.

We envision multiple pathways for extending our research in future work. A logical next step would be to collect data from other countries in South America and apply this counterfactual modeling approach, examining extreme or disruptive events that occurred across countries, such as the COVID-19 pandemic, as well as other events that were unique to that country. Comparing other countries with different mixes of hydro and thermal power generation, such as Peru and Argentina, could be particularly insightful. Additionally, we could continue our focus on the Brazilian power system, and collect more data such as the transmission network structure, electricity pricing, generation, etc. and apply it in a system-wide analysis. This could prove particularly interesting as the Brazilian power system is currently transitioning to accommodate new distributed energy resources to meet increased electrification and population demand.

Future research could also investigate emergent risks of the Brazilian power system to extreme climate events through a more sophisticated treatment of climate model output. While climate models are incomplete representations of global dynamics, they allow for a systematic exploration of plausible realizations of future climate states. The uncertainty in climate model output can be broadly sorted into three categories: internal variability uncertainty, model uncertainty (which can contain uncertainty from the downscaling method), and scenario uncertainty. While this work attempts to include heterogeneity within the model and scenario domains, inclusion of initial condition large ensembles and multiple downscaling methods would allow further research to consider a wider range of plausible climate futures and probe the susceptibility of the Brazilian region to unforeseen phenomenon, such as record-shattering heat extremes. Using established methods to partition these uncertainty sources may reveal which source is most important for evaluating impacts on power systems in Brazil at different timescales (i.e., short-term operation, planning, and long-term expansion). Finally, future research could consider the synthetic construction of extreme or disruptive events projected into future periods, perhaps by simulating past historical events at various impact magnitudes. Such scenario-based extreme event simulations could be useful for planning purposes, such as understanding the future risk to the power system posed by these event types, and is an exciting area of inquiry that we encourage other researchers to explore, both for Brazil and in other contexts.

6. Conclusion

In this work, we proposed a data-centric approach for counterfactual consumption modeling. We also present a method based on the DAR algorithm that allows tree-based models to extrapolate beyond the training distribution. After data augmentation, we achieved an R^2 of 0.848 with a MAPE below 2.7%. We employed our counterfactual model to both understand past historical scenarios and to provide forecasts for future warming. Furthermore, we found that our estimates generally matched the Brazilian government's projections when measuring the impact of 2001's federal power rationing policies. With the robustness of our model and the provided insights, we hope to raise awareness and enable future research regarding smart-grid planning. In future work, we plan to further refine our modeling approach and explore applications for smart-grid planning under climate uncertainty.

CRediT authorship contribution statement

Gianlucca Zuin: Methodology, Software, Validation, Investigation, Data curation, Writing – original draft, Visualization. **Rob Buechler:** Software, Validation, Data curation, Writing – original draft. **Tao Sun:** Validation, Writing – original draft. **Chad Zanocco:** Validation, Project administration, Writing – original draft, Visualization. **Francisco Galuppo:** Validation, Writing – review & editing. **Adriano Veloso:** Conceptualization, Supervision, Writing – review & editing. **Ram Rajagopal:** Conceptualization, Supervision, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The ERA5 weather data from the Copernicus Data store is available via API from https://cds.climate.copernicus.eu. Brazilian energy data owned by ONS is made directly available from http://www.ons.org.br/. The CMIP6 data is accessible via the Google storage API, managed by Pangeo and the Earth System Grid Federation, and is available at https://pangeo-data.github.io/pangeo-cmip6-cloud/.

Acknowledgments

The authors would like to thank Sarah Fletcher, Daniella Castro, and Jenny Skerter for their valuable feedback about this research. This work was funded in part by the National Science Foundation through a CAREER award (#1554178) and by a Stanford Precourt Pioneering Project award.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.energy.2023.128101.

References

[1] Sanyal S, Wuebbles DJ. The potential impact of a clean energy society on air quality. Earth's Future 2022;10(6):e2021EF002558.
[2] IEA. Hydropower has a crucial role in accelerating clean energy transitions to achieve countries' climate ambitions securely. 2021, [Online] https://www.iea.org/news/hydropower-has-a-crucial-role-in-accelerating-clean-energy-transitions-to-achieve-countries-climate-ambitions-securely.
[3] Moran EF, Lopez MC, Moore N, Müller N, Hyndman DW. Sustainable hydropower in the 21st century. Proc Natl Acad Sci 2018;115(47):11891–8.
[4] Moazami A, Nik VM, Carlucci S, Geving S. Impacts of future weather data typology on building energy performance–Investigating long-term patterns of climate change and extreme weather conditions. Appl Energy 2019;238:696–720.
[5] Agência Nacional de Energia Elétrica (ANEEL). Brasil ultrapassa os 190 GW em capacidade de geração de energia elétrica. Ministério de Minas e Energia; 2023, [Online]. Available: https://www.gov.br/aneel/pt-br/assuntos/noticias/2023/brasil-ultrapassa-os-190-gw-em-capacidade-de-geracao-de-energia-eletrica.
[6] ANEEL. Expansão da matriz elétrica brasileira - setembro/2021. Agência Nacional de Energia Elétrica; 2021.

[7] EPE, ONS, CCEE. 2ª Revisão quadrimestral das projeçes da demanda de energia elétrica do sistema interligado nacional 2021–2025. Estud Demanda 2021.

[8] Ferraz A. Economic impact of Brazil's energy crisis will last until 2023. 2021, [Online]. https://brazilian.report/business/2021/09/19/economic-energy-crisis-2023/.

[9] Hunt JD, Nascimento A, ten Caten CS, Tomé FMC, Schneider PS, Thomazoni ALR, de Castro NJ, Brandão R, de Freitas MAV, Martini JSC, et al. Energy crisis in Brazil: Impact of hydropower reservoir level on the river flow. Energy 2022;239:121927.

[10] Zuin G, Buechler R, Sun T, Zanocco C, Castro D, Veloso A, Rajagopal R. Revealing the impact of extreme events on electricity consumption in Brazil: A data-driven counterfactual approach. In: 2022 IEEE power and energy society general meeting (PESGM). IEEE; 2022, p. 1–5.

[11] Jiang P, Li R, Liu N, Gao Y. A novel composite electricity demand forecasting framework by data processing and optimized support vector machine. Appl Energy 2020;260:114243.

[12] Wang Z, Hong T, Li H, Piette MA. Predicting city-scale daily electricity consumption using data-driven models. Adv Appl Energy 2021;2:100025.

[13] Meyer H, Pebesma E. Predicting into unknown space? Estimating the area of applicability of spatial prediction models. Methods Ecol Evol 2021;12(9):1620–33.

[14] Hooker G. Diagnostics and extrapolation in machine learning. stanford university; 2004.

[15] Hooker G, Rosset S. Prediction-based regularization using data augmented regression. Stat Comput 2012;22:237–49.

[16] Ruan G, Wu D, Zheng X, Zhong H, Kang C, Dahleh MA, Sivaranjani S, Xie L. A cross-domain approach to analyzing the short-run impact of COVID-19 on the US electricity sector. Joule 2020;4(11):2322–37.

[17] Bardelin CEA. Os efeitos do racionamento de energia elétrica ocorrido no Brasil em 2001 e 2002 com ênfase no consumo de energia elétrica (Ph.D. thesis), Universidade de São Paulo; 2004.

[18] Burillo D, Chester MV, Pincetl S, Fournier ED, Reyna J. Forecasting peak electricity demand for los angeles considering higher air temperatures due to climate change. Appl Energy 2019;236:1–9.

[19] Swart J, Brinkmann L. Economic complexity and the environment: Evidence from Brazil. In: Leal Filho W, Tortato Ua, Frankenberger F, editors. Universities and sustainable communities: meeting the goals of the agenda 2030. Cham: Springer International Publishing; 2020, p. 3–45.

[20] Höltinger S, Mikovits C, Schmidt J, Baumgartner J, Arheimer B, Lindström G, Wetterlund E. The impact of climatic extreme events on the feasibility of fully renewable power systems: A case study for Sweden. Energy 2019;178:695–713. http://dx.doi.org/10.1016/j.energy.2019.04.128, [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0360544219307613.

[21] Kemabonta T. Grid resilience analysis and planning of electric power systems: The case of the 2021 texas electricity crises caused by winter storm Uri. Electr J 2021;34(10):107044.

[22] Orlov A, Sillmann J, Vigo I. Better seasonal forecasts for the renewable energy industry. Nat Energy 2020;5(2):108–10.

[23] Vitart F, Robertson AW. The sub-seasonal to seasonal prediction project (S2S) and the prediction of extreme events. Npj Clim Atmos Sci 2018;1(1):1–7.

[24] Kumar M, Gupta DK, Singh S. Extreme event forecasting using machine learning models. In: Advances in communication and computational technology. Springer; 2021, p. 1503–14.

[25] Lerch S, Thorarinsdottir TL, Ravazzolo F, Gneiting T. Forecaster's dilemma: Extreme events and forecast evaluation. Statist Sci 2017;106–27.

[26] Resource adequacy for a decarbonized future: A summary of existing and proposed resource adequacy metrics. Tech. rep. 000000003002023230, Electric Power Research Institute; 2022.

[27] Buechler E, Powell S, Sun T, Astier N, Zanocco C, Bolorinos J, Flora J, Boudet H, Rajagopal R. Global changes in electricity consumption during COVID-19. Iscience 2022;25(1):103568.

[28] Belançon MP. Brazil electricity needs in 2030: trends and challenges. Renew Energy Focus 2021;36:89–95.

[29] Trotter IM, Bolkesjø TF, Féres JG, Hollanda L. Climate change and electricity demand in Brazil: A stochastic approach. Energy 2016;102:596–604.

[30] de Assis Cabral J, Legey LFL, de Freitas Cabral MV. Electricity consumption forecasting in Brazil: A spatial econometrics approach. Energy 2017;126:124–31.

[31] Maluf de Lima L, Piedade Bacchi MR. Assessing the impact of Brazilian economic growth on demand for electricity. Energy 2019;172:861–73. http://dx.doi.org/10.1016/j.energy.2019.01.154, [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0360544219301707.

[32] de Assis Cabral J, de Freitas Cabral MV, Júnior AOP. Elasticity estimation and forecasting: An analysis of residential electricity demand in Brazil. Util Policy 2020;66:101108.

[33] Luzia R, Rubio L, Velasquez CE. Sensitivity analysis for forecasting Brazilian electricity demand using artificial neural networks and hybrid models based on autoregressive Integrated Moving Average. Energy 2023;274:127365. http://dx.doi.org/10.1016/j.energy.2023.127365, [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0360544223007594.

[34] Mendes CAB, Beluco A, Canales FA. Some important uncertainties related to climate change in projections for the Brazilian hydropower expansion in the Amazon. Energy 2017;141:123–38. http://dx.doi.org/10.1016/j.energy.2017.09.071, [Online]. Available: https://www.sciencedirect.com/science/article/pii/S036054421731592X.

[35] Ballew BS. Elsevier's Scopus® database. J Electron Resour Med Libr 2009;6(3):245–52.

[36] Liverpool L. Researchers from global south under-represented in development research. Nature 2021.

[37] ONS. Histórico da operação instalada, carga de energia. Oper Nacl Sist Elétr Rio Janeiro 2018.

[38] Hersbach H, Bell B, Berrisford P, Hirahara S, Horányi A, Muñoz-Sabater J, Nicolas J, Peubey C, Radu R, Schepers D, et al. The ERA5 global reanalysis. Q J R Meteorol Soc 2020;146(730):1999–2049.

[39] Eyring V, Bony S, Meehl GA, Senior CA, Stevens B, Stouffer RJ, Taylor KE. Overview of the coupled model intercomparison project phase 6 (CMIP6) experimental design and organization. Geosci Model Dev 2016;9(5):1937–58.

[40] Steadman RG. A universal scale of apparent temperature. J Appl Meteorol Climatol 1984;23(12):1674–87.

[41] Steadman RG. The assessment of sultriness. Part I: A temperature-humidity index based on human physiology and clothing science. J Appl Meteorol Climatol 1979;18(7):861–73.

[42] Stull RB, Ahrens CD, et al. Meteorology for scientists and engineers. Brooks/Cole; 2000.

[43] Strahler A, Strahler A. Physical geography. John Wiley & Sons; 2007.

[44] van der Maaten L. Learning a parametric embedding by preserving local structure. In: Intl conf. on artificial intelligence and statistics. 2009, p. 384–91.

[45] Abram N, Gattuso J, Prakash A, Cheng L, Chidichimo M, Crate S, Enomoto H, Garschagen M, Gruber N, Harper S, Holland E, Kudela R, Rice J, Steffen K, von Schuckmann K. Framing and context of the report: Supplementary material. In: IPCC special report on the ocean and cryosphere in a changing climate. Switzerland: Intergovernmental Panel on Climate Change; 2019, p. 73–129.

[46] O'Neill BC, Tebaldi C, van Vuuren DP, Eyring V, Friedlingstein P, Hurtt G, Knutti R, Kriegler E, Lamarque J-F, Lowe J, Meehl GA, Moss R, Riahi K, Sanderson BM. The scenario model intercomparison project (ScenarioMIP) for CMIP6. Geosci Model Dev 2016;9(9):3461–82.

[47] Cheng L, Zhu J. Benefits of CMIP5 multimodel ensemble in reconstructing historical ocean subsurface temperature variations. J Clim 2016;29.

[48] Wang B, Zheng L, Liu DL, Ji F, Clark A, Yu Q. Using multi-model ensembles of CMIP5 global climate models to reproduce observed monthly rainfall and temperature with machine learning methods in Australia. Int J Climatol 2018;38(13):4891–902.

[49] Bi D, Dix M, Marsland S, O'farrell S, Sullivan A, Bodman R, Law R, Harman I, Srbinovsky J, Rashid HA, et al. Configuration and spin-up of ACCESS-CM2, the new generation Australian community climate and earth system simulator coupled model. J South Hemisph Earth Syst Sci 2020;70(1):225–51.

[50] Semmler T, Danilov S, Gierz P, Goessling HF, Hegewald J, Hinrichs C, Koldunov N, Khosravi N, Mu L, Rackow T, et al. Simulations for CMIP6 with the AWI climate model AWI-CM-1-1. J Adv Modelling Earth Syst 2020;12(9):e2019MS002009.

[51] Liu S-M, Chen Y-H, Rao J, Cao C, Li S-Y, Ma M-H, Wang Y-B. Parallel comparison of major sudden stratospheric warming events in CESM1-WACCM and CESM2-WACCM. Atmosphere 2019;10(11):679.

[52] Cherchi A, Fogli PG, Lovato T, Peano D, Iovino D, Gualdi S, Masina S, Scoccimarro E, Materia S, Bellucci A, et al. Global mean climate and main patterns of variability in the CMCC-CM2 coupled model. J Adv Modelling Earth Syst 2019;11(1):185–209.

[53] Lovato T, Peano D, Butenschön M, Materia S, Iovino D, Scoccimarro E, Fogli P, Cherchi A, Bellucci A, Gualdi S, et al. CMIP6 simulations with the CMCC earth system model (CMCC-ESM2). J Adv Modelling Earth Syst 2022;14(3):e2021MS002814.

[54] Swart NC, Cole JN, Kharin VV, Lazare M, Scinocca JF, Gillett NP, Anstey J, Arora V, Christian JR, Hanna S, et al. The Canadian earth system model version 5 (CanESM5. 0.3). Geosci Model Dev 2019;12(11):4823–73.

[55] Döscher R, Acosta M, Alessandri A, Anthoni P, Arneth A, Arsouze T, Bergmann T, Bernadello R, Bousetta S, Caron L-P, et al. The EC-earth3 Earth system model for the climate model intercomparison project 6. Geosci Model Dev Discuss 2021;1:2021.

[56] Dunne J, Horowitz L, Adcroft A, Ginoux P, Held I, John J, Krasting J, Malyshev S, Naik V, Paulot F, et al. The GFDL Earth System Model version 4.1 (GFDL-ESM 4.1): Overall coupled model description and simulation characteristics. J Adv Modelling Earth Syst 2020;12(11):e2019MS002015.

[57] Krishnan R, Swapna P, Vellore R, Narayanasetti S, Prajeesh A, Choudhury AD, Singh M, Sabin T, Sanjay J. The IITM earth system model (ESM): development and future roadmap. In: Current trends in the representation of physical processes in weather and climate models. Springer; 2019, p. 183–95.

[58] Volodin E. The mechanisms of cloudiness evolution responsible for equilibrium climate sensitivity in climate model INM-CM4-8. Geophys Res Lett 2021;48(24):e2021GL096204.

[59] Volodin E. The mechanism of 60-year and 15-year Arctic climate oscillations in climate model INM-CM5-0. In: EGU general assembly conference abstracts. 2020, p. 7265.
[60] Boucher O, Servonnat J, Albright AL, Aumont O, Balkanski Y, Bastrikov V, Bekki S, Bonnet R, Bony S, Bopp L, et al. Presentation and evaluation of the IPSL-CM6A-LR climate model. J Adv Modelling Earth Syst 2020;12(7):e2019MS002010.
[61] Müller WA, Jungclaus JH, Mauritsen T, Baehr J, Bittner M, Budich R, Bunzel F, Esch M, Ghosh R, Haak H, et al. A higher-resolution version of the max planck institute earth system model (MPI-ESM1. 2-HR). J Adv Modelling Earth Syst 2018;10(7):1383–413.
[62] Yukimoto S, Kawai H, Koshiro T, Oshima N, Yoshida K, Urakawa S, Tsujino H, Deushi M, Tanaka T, Hosaka M, et al. The meteorological research institute earth system model version 2.0, MRI-ESM2. 0: Description and basic evaluation of the physical component. J Meteorol Soc Jpn Ser II 2019.
[63] Tjiputra JF, Schwinger J, Bentsen M, Morée AL, Gao S, Bethke I, Heinze C, Goris N, Gupta A, He Y-C, et al. Ocean biogeochemistry in the Norwegian Earth System Model version 2 (NorESM2). Geosci Model Dev 2020;13(5):2393–431.
[64] Douville H, Qasmi S, Ribes A, Bock O. Global warming at near-constant tropospheric relative humidity is supported by observations. Commun Earth Environ 2022;3(1):237. http://dx.doi.org/10.1038/s43247-022-00561-z, [Online]. Available: https://www.nature.com/articles/s43247-022-00561-z.
[65] Giannakopoulos C, Psiloglou BE. Trends in energy load demand for Athens, Greece: weather and non-weather related factors. Clim Res 2006;31(1):97–108.
[66] Chen S-T, Kuo H-I, Chen C-C. The relationship between GDP and electricity consumption in 10 Asian countries. Energy Policy 2007;35(4):2611–21.
[67] Al-Bajjali SK, Shamayleh AY. Estimating the determinants of electricity consumption in Jordan. Energy 2018;147:1311–20.
[68] Alberini A, Prettico G, Shen C, Torriti J. Hot weather and residential hourly electricity demand in Italy. Energy 2019;177:44–56.
[69] Ke G, Meng Q, Finley T, Wang T, Chen W, Ma W, Ye Q, Liu T-Y. LightGBM: A highly efficient gradient boosting decision tree. In: Guyon I, Luxburg UV, Bengio S, Wallach H, Fergus R, Vishwanathan S, Garnett R, editors. Advances in neural information processing systems, Vol. 30. Curran Associates, Inc.; 2017, p. 3146–54.
[70] Friedman JH. Multivariate adaptive regression splines. Ann Statist 1991;19(1):1–67.
[71] Seber GA, Lee AJ. Linear regression analysis, Vol. 330. John Wiley & Sons; 2003.
[72] Zaki MJ, Meira W. Data mining and analysis: fundamental concepts and algorithms. Cambridge University Press; 2014.
[73] Zuin G, Araujo D, Ribeiro V, Seiler MG, Prieto WH, Pintão MC, dos Santos Lazari C, Granato CFH, Veloso A. Prediction of SARS-CoV-2-positivity from million-scale complete blood counts using machine learning. Commun Med 2022;2(1):72.
[74] Reddy DR, et al. Speech understanding systems: A summary of results of the five-year research effort, Vol. 17. Department of Computer Science. Camegie-Mell University, Pittsburgh, PA; 1977, p. 138.
[75] Hahn GJ. The hazards of extrapolation in regression analysis. J Qual Technol 1977;9(4):159–65.
[76] Lundberg SM, Lee S-I. A unified approach to interpreting model predictions. In: Guyon I, Luxburg UV, Bengio S, Wallach H, Fergus R, Vishwanathan S, Garnett R, editors. Advances in neural information processing systems, Vol. 30. 2017, p. 4765–74.
[77] Hale T, Petherick A, Phillips T, Webster S. Variation in government responses to COVID-19. Blavatnik school working paper, University of Oxford; 2020.
[78] Gullo MCR. A economia na pandemia Covid-19: algumas considerações. Rosa Ventos 2020;12(Esp. 3):1–8.
[79] Ritchie H, Mathieu E, Rodés-Guirao L, Appel C, Giattino C, Ortiz-Ospina E, Hasell J, Macdonald B, Beltekian D, Roser M. Coronavirus pandemic (COVID-19). Our World Data 2020.
[80] Brasil B. Apagão foi 'acidente' causado pelo mau tempo, diz ministro. 2009, [Online] https://www.bbc.com/portuguese/lg/noticias/2009/11/091111_apagao_rc.
[81] Hudedmani MG, Soppimath VM, Hubballi SV, Joshi D. Dawn after black out: A hope of light–A review. Int J Adv Sci Technol 2019;6(1):1264–71.
[82] Globo M. Apagão no Brasil em 2009. 2022-01-12, [Online] https://memoriaglobo.globo.com/jornalismo/jornalismo-e-telejornais/jornal-nacional/reportagens-e-entrevistas/noticia/apagao-no-brasil-em-2009.ghtml.
[83] Silveira D. Falha humana provocou apagão nas regies norte e Nordeste em março, diz ONS. 2018, [Online] https://g1.globo.com/economia/noticia/falha-humana-provocou-apagao-no-norte-e-nordeste-diz-ons.ghtml.
[84] Liu Y. Analysis of Brazilian blackout on march 21st, 2018 and revelations to security for Hunan Grid. In: 2019 4th international conference on intelligent green building and smart grid (IGBSG). IEEE; 2019, p. 1–5.
[85] Ziel F. Modeling public holidays in load forecasting: a German case study. J Mod Power Syst Clean Energy 2018;6(2):191–207.
[86] Freitas C. Preços de produtos típicos do verão sobem junto com a temperatura. Gazeta 2020-10-15. [Online]. Available: https://www.agazeta.com.br/es/economia/precos-de-produtos-tipicos-do-verao-sobem-junto-com-a-temperatura-1020.