



HAL
open science

Investigation of the Impact of Economic and Social Factors on Energy Demand through Natural Language Processing

Yun Bai, Simon Camal, Andrea Michiorri

► **To cite this version:**

Yun Bai, Simon Camal, Andrea Michiorri. Investigation of the Impact of Economic and Social Factors on Energy Demand through Natural Language Processing. 2024. hal-04608605

HAL Id: hal-04608605

<https://hal.science/hal-04608605v1>

Preprint submitted on 11 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Investigation of the Impact of Economic and Social Factors on Energy Demand through Natural Language Processing

Yun BAI^{1*}, Simon CAMAL^{1†} and Andrea MICHIORRI^{1†}

^{1*}Centre for Processes, Renewable Energy and Energy Systems (PERSEE), Mines Paris - PSL, 1, Rue Claude Daunesse, Sophia Antipolis, 06904, France.

*Corresponding author(s). E-mail(s): yun.bai@minesparis.psl.eu;
Contributing authors: simon.camal@minesparis.psl.eu;
andrea.michiorri@minesparis.psl.eu;

[†]These authors contributed equally to this work.

Abstract

The relationship between energy demand and variables such as economic activity and weather is well established. However, this paper aims to explore the connection between energy demand and other social aspects, which receives little attention. Through the use of natural language processing on a large news corpus, we shed light on this important link. This study was carried out in five regions of the UK and Ireland and considers multiple horizons from 1 to 30 days. It also considers economic variables such as GDP, unemployment and inflation. We found that: 1) News about military conflicts, transportation, the global pandemic, regional economics, and the international energy market are related to electricity demand. 2) Economic indicators are more important in the East Midlands and Northern Ireland, while social indicators are more useful in the West Midlands and the South West of England. 3) The use of these indices improved forecasting performance by up to 9%.

Keywords: Power systems, regional electricity demand, economy and society, multi-horizon forecasting, probabilistic forecasting

Context

Electricity demand modelling is a fundamental process in power system planning, operation, and energy trading [1]. In order to avoid additional carbon emissions from excess electricity generation and the high costs of electricity storage, electricity demand and supply should be matched over time [2]. Demand forecasting has become a means of enabling power dispatch, planning generation schedules, and integrating renewable energy sources [3].

Electricity demand forecasting is linked to various factors, including weather, economic activity, and major events. These factors have been extensively studied across different spatial and temporal scales [4]. First, weather conditions significantly impact electricity demand by changing human activities, such as the use of heating, cooling, and lighting devices. This is due to temperature variations, humidity, solar radiation, and large meteorological events [5, 6]. Second, the intensity and development level of economic activity are key elements influencing electricity demand. For example, business hours and weekdays are characterised by higher electricity demand [7]. The authors of [8] proposed that electricity consumption in Brazil is spatially correlated with regional economic development. Macroeconomic indicators such as gross domestic product (GDP) per capita have been found to influence national electricity demand significantly [9]. Third, large events lead to changes in consumption patterns both temporally and spatially. For example, the lockdown policies during COVID-19 reduced electricity usage [10, 11].

In the early 2000s, researchers pointed out that social sciences should be integrated into energy demand studies in order to gain a comprehensive understanding of the formation and evolution of energy demand, which is influenced by market prices, consumer awareness, and social norms [12]. However, social factors pose challenges due to their ambiguous definition and measurement. In social sciences, scientists depict the complex relationship between human social characters and energy demand through household interviews, surveys, and relational sociology. These studies offer insights into human behavior and the correlation with peak demand, flexibility volumes, and energy usage patterns [13–15]. More recently, in order to better model electricity demand, studies have been carried out with large datasets related to social activities, including mobile phone network data [16], online search queries shown by Google Trends [17], and high-resolution satellite images [18].

Public news contains extensive but overlooked social knowledge that is potentially valuable for electricity demand modelling. Usually, this knowledge is noisy, unstructured, and sparse, making its distribution difficult to understand [19]. The emergence of Natural Language Processing (NLP) technologies allows us to discover knowledge from massive text. Attempts have already been made to forecast energy by forecasting incorporating news text, such as crude oil prices forecasting [20–22]. The study by [23] not only forecast crude oil prices, but also extended the application scope to gold and natural gas futures using the same news corpus. In electricity demand forecasting, the study by [24] used key word frequency extracted from social media to improve forecasting and explain the structural change in the load profiles. This research is an inspiring attempt at text-based forecasting in electricity demand, but the incorporation of richer text features remains an issue. Other studies have worked

on latent features processed by deep networks, for example Convolutional Neural Network (CNN) and Long Short-Term Memories (LSTM) [25, 26]. Although they provided improved forecasts compared to benchmarks, these studies lacked explanation of how the textual features worked, which is more compelling in this research. In [27], the authors discovered several social variables (geopolitical tensions, transportation, global pandemics, etc.) in unstructured news that could cause changes in electricity demand. However, the use cases only concern UK and Ireland national day-ahead load forecasting. Regional scenarios and longer forecasting terms should be explored to verify the time-varying impacts of social factors on electricity demand.

This study aims to explore more subtle relationships between economic-social activities and electricity demand. Building on [27], this paper extends the approach in four key ways: **1)** improving the methodology with a better explanation of causation, **2)** increasing the spatial resolution from national load to multi-regional loads, **3)** testing forecasts on longer horizons from 1 to 30 days ahead, and **4)** studying the effect on extreme cases. This study provides a more nuanced understanding of the complex interplay between economic-social activities and energy consumption, which will undoubtedly improve demand forecasting models.

Gaps, contributions, and research questions

In this study, we work on regional electricity demand forecasting incorporating both economic and textual social drivers. Text-based prediction has been applied in many areas, but there is a clear need to explain the improved performance and avoid the effects of spurious correlations. The difficulty of explaining the improved performance is due to the complicated and unintuitive steps of text preprocessing and feature extraction, and the black-box nature of the models used. Recent studies of demand forecasting with text have focused on single time series or forecast horizons, overlooking the potential for generalising the results and preventing the opportunity to determine the most appropriate application. Finally, the predominance of deterministic forecasting prevents a comprehensive analysis of the ability of the method to deal with forecasting uncertainty.

To bridge the gaps, this study contributes on the following points:

1. We propose a methodology to analyse the relationship between news content and demand from multiple perspectives. This makes it possible to highlight the relative importance of economic and social indicators.
2. We evaluate the method on a robust case study consisting of five regions in the UK and Ireland, and forecasting horizons from 1 to 30 days ahead. This allows us to understand where and when the social indicators studied are more relevant.

This study answers the following five research questions of relevance to power system practitioners, with more detailed discussion provided in [Results 5](#):

1. Is there any relationship between national news and regional electricity demand?
2. Can textual news be used for practical applications to improve regional demand forecasting?
3. Which topics have a higher impact on electricity demand?
4. Are the results consistent for different regions?
5. For which forecast horizons is the impact higher?

In order to facilitate the application to the wider research community, the codes and datasets are made available at <https://github.com/YunBAI-PSL/Regional-electricity-demand-forecasting-with-news>.

Results

This section first provides an analysis of the causality between economic indicators, social indicators, and regional electricity demand, with a causality map and a route as examples. Considering both deterministic and probabilistic forecasting, this section compares the forecasting errors of the models with and without economic or social factors in different regions. To understand the mechanism of how economic and social factors impact forecasting, the fourth part of this section analyses feature importance both overall and by horizon. The dependencies of key factors are also included. Finally, this section answers the five research questions in [Gaps, contributions, and research questions](#). Note that this section only shows results for the East Midlands region; the results of the other regions can be found in the [Supplemental materials](#).

Causality analysis




Fig. 1: (a) Granger causality map of economic and social factors and electricity demand in the East Midlands. Single and double rows indicate unidirectional and bidirectional Granger causality, respectively. (b) Granger causality route from inflation rate to electricity demand in the East Midlands.

To present the relations of economy, society, and electricity demand, we take the East Midlands as an example. In Figure 1, we plot the causality map of direct relations and a route of second order relations according to the results of the Granger test. In Figure 1a, we discover that there is no direct link between economic factors and electricity demand. The causation flows from economic factors to demand through

social factors. In particular, electricity demand is influenced by news about transport, pandemic control, regional economics, energy markets, strikes, and military conflicts. For news about pandemic control and regional economics, a bidirectional causality with electricity demand is also observed. Figure 1b presents a more complex causality chain observed in the East Midlands region. This can be interpreted as: inflation influences family life, which in turn influences travel & leisure. At this point, news about travel & leisure shows bidirectional causality with strikes and pandemic control, which, as seen before, directly influences electricity demand.

Deterministic forecast

To compare how economic and social factors impact electricity demands separately and jointly, we designed an experiment consisting of the comparison of four different prediction models trained with different sets of data. The prediction models are introduced in Methods 3 and built with Gradient Boosting Machines (GBM). The four models are: GBM, a benchmark forecast model trained only with historical load, calendar and weather information; GBM-E, trained with the inputs of GBM and three economic indicators (GDP, inflation, and unemployment); GBM-S, to which text-extracted social features are added; and finally, GBM-SE, trained with all the variables available for the previous models. Table 1 shows the Root Mean Squared Errors (RMSE) and Mean Absolute Percentage Error (MAPE) performance across all the regions on the average of 30 horizons.

Table 1: Horizon-averaged model performance of RMSE and MAPE with and without economic and social factors. S and E indicate social and economic factors. N (No) and Y (Yes) are indicators to show if the factor (S or E) participates in the model.

Regions	Metrics	RMSE (MW)		MAPE (%)	
	E S	N	Y	N	Y
East Midlands	N	135.76	125.78	3.19	2.89
	Y	133.96	123.31	3.19	2.84
West Midlands	N	115.31	109.53	3.53	3.34
	Y	109.05	107.91	3.32	3.27
South Wales	N	103.11	109.18	8.46	9.00
	Y	101.82	108.61	8.35	8.93
South West	N	124.23	117.25	9.20	8.83
	Y	116.69	116.15	8.78	8.77
Northern Ireland	N	53.66	49.42	5.64	5.15
	Y	50.03	48.68	5.28	5.07
Ireland	N	154.08	165.90	3.52	3.88
	Y	149.54	160.16	3.46	3.76

The analysis in Table 1 suggests that in each region considered, the use of social features improves forecasting performance. In almost all the regions, the use of economic indicators also contributes to improving forecasting performance. The overlap

effect of economic and social factors can be observed in the East Midlands, West Midlands, South West, and Northern Ireland. For example, in the East Midlands, economic and social factors can improve the benchmark model by 1.3% and 7.4% respectively. The improvement can reach up to 9.2% if both factors are combined. According to the performance of the factors, we summarised the economic-sensitive regions as the East Midlands and Northern Ireland, and the rest are social-sensitive regions.

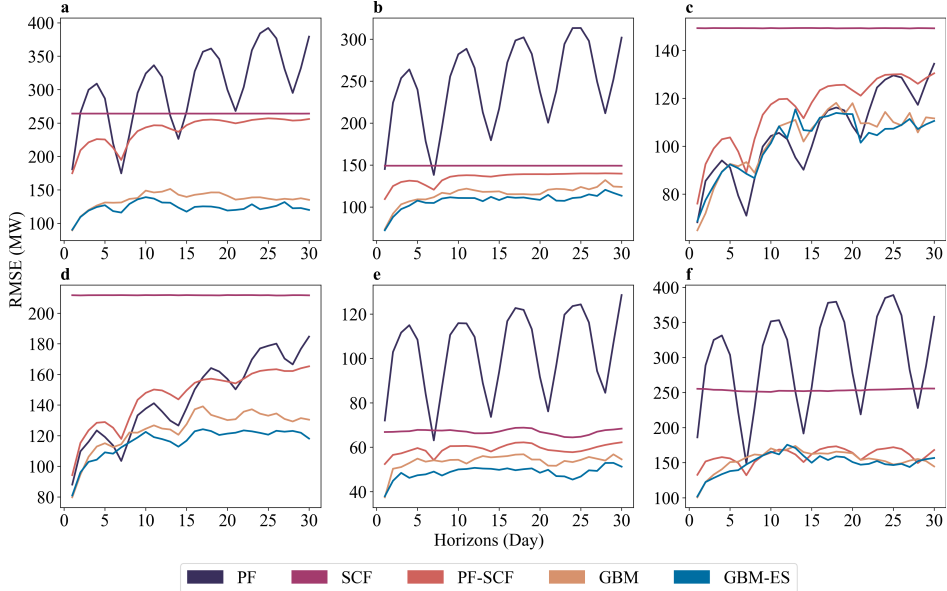


Fig. 2: RMSE error plots for model PF, SCF, PF-SCF, GBM, and GBM-ES in all the regions. a-f, GBM-ES is the best-performing GBM model with economic, social or both factors. (a): East Midlands, (b): West Midlands, (c): South Wales, (d): South West, (e) Northern Ireland, (f) Ireland.

The evolution of the forecasting error along the horizons from 1 to 30 days is presented in Figure 2. Here, the models GBM and GBM-ES are compared to naive forecasts: Persistence Forecasts (PF), Smart Climatology Forecasts (SCF), and the linear combination of PF and SCF (PF-SCF), described in [Methods 4](#). The charts show how, in almost all cases, the proposed model outperforms naive approaches, and the use of social and economic factors additionally reduces the errors.

Probabilistic forecast

In this case, the experiments presented above are repeated, but the evaluation is carried out against probabilistic forecasts metrics. In particular, the Continuous Ranked Probability Score (CRPS) is used as the metric to evaluate the performance of the probabilistic forecast. [Table 2](#) lists the CRPS values for all the regions on the average

of 30 horizons. Its analysis shows that CRPS values can be decreased by 5% on average by adding social factors in the West Midlands and South Wales regions. In the South West and Northern Ireland, economic factors improve probabilistic forecasts by 7.3% on average. Although both factors cannot reduce CRPS in the East Midlands, their combination can still improve by 1%.

Figure 3 shows the improvements by percentage on all the horizons. It demonstrates that in the East Midlands, economic factors improve forecasting on the future third week with an average of 6.6%, while social factors are merely beneficial on days 2, 5, and 6 with an improvement of 2.6%. The performance in the West Midlands and the South West are similar. Economic factors contribute to the second half of the horizons, reflecting delayed impact, but positive improvements in social factors cover almost all the horizons. Considering the second half of the horizons, in the West Midlands, the average improvement of economic factors is 3.8%, 10% for social factors, and 9.5% for economic-social factors. In South Wales, social factors can reduce the CRPS by 1.1% and 1% in the first and second halves of month. In Northern Ireland, economic and social factors can improve probabilistic forecasting, and higher improvement appears in the second half of the month. The improvement percentages for ‘E’, ‘S’, and ‘ES’ are 9.4%, 4.9%, and 9.2%, which shows that economic factors are more beneficial for reducing forecasting uncertainty. Social factors from the UK news can slightly improve the probabilistic forecasting in Ireland by less than 1%, and only in the first week.

As a short conclusion, combining economic and social factors is more beneficial in deterministic forecasting. A single social or economic factor performs better for probabilistic forecasting. In the West Midlands and South Wales, social factors can improve both forecasting tasks, while Northern Ireland prefers economic factors.

Table 2: Horizon-averaged model performance of CRPS with and without economic and social factors. S and E indicate social and economic factors. N (No) and Y (Yes) are indicators to show if the factor (S or E) participates in the model.

Regions	E		N	Y
	S			
East Midlands	N		0.2841	0.2849
	Y		0.3037	0.2822
West Midlands	N		0.2759	0.2735
	Y		0.2517	0.2578
South Wales	N		0.7329	0.8133
	Y		0.7233	0.7956
South West	N		0.4629	0.4387
	Y		0.4471	0.4436
Northern Ireland	N		0.3815	0.3454
	Y		0.3627	0.3463
Ireland	N		0.3306	0.3910
	Y		0.3500	0.3968




Fig. 3: Probabilistic forecasting improvements of models with economic and social factors. a-f, improvement percentages for 30 horizons in the different participation of economic-social factors. ‘E’ for economic, ‘S’ for social, ‘ES’ for both economic and social. (a): East Midlands, (b): West Midlands, (c): South Wales, (d): South West, (e) Northern Ireland, (f) Ireland.

A final analysis for probabilistic forecasts was carried out on a single special day. In the UK, 2023-05-08 (Monday) was a bank holiday to celebrate the coronation of King Charles III. However, due to its exceptional nature, this date was not present in the calendar used to identify non-working days. The government actually announced that 2023-05-08 would be a bank holiday on 2022-11-06, according to the news: *Extra bank holiday approved to mark King’s coronation* [28]. The python package `holidays` 0.40 used in the benchmark model did not include 2023-05-08 as a holiday [29]. Thus, the benchmark model treated the date 2023-05-08 as a normal Monday and overestimated the demand on this day. Nevertheless, the model incorporating social factors captured the relevant textual information that was mentioned more frequently as May 8th approached. The model featuring social factors outperformed the benchmark on different horizons as shown in Figure 4.

Figure 4 shows the comparison of actual values and forecasts of electricity demand in the East Midlands on date 2023-05-08. Both the benchmark and the model with




Fig. 4: Actual values and forecasts of electricity demand in the East Midlands on date 2023-05-08. The forecasts are generated by benchmark and the model with social factors. The horizons are 1, 10, and 30 days before 2023-05-08, which means that the datasets used in the three subplots end on 2023-05-07, 2023-04-28, and 2024-04-08 respectively. The RMSE improvements and 90% confidence intervals are given in the figure.

social factors are PGBM models in probabilistic forecasting. It is obvious that the demand is overestimated, because the models recognised this day as a regular Monday instead of holiday. However, the model with social factors can decrease forecasts especially in daytime, corresponding to the higher level of human activity, and brings 19.89%, 18.72%, and 10.96% of RMSE improvements for horizons 1, 10, and 30 days with respect to the benchmark. This case clearly illustrates well the impact of sudden events on the benchmark model. The adjustments for holidays prevented the benchmark from responding in time, leading to forecasting biases. This demonstrates that our model with social factors has the ability to handle sudden events, especially when traces of these events can be found in historical news texts, which the benchmark fails to detect.

Feature importance analysis

In this subsection, we applied SHAP (SHapley Additive exPlanations) methods to explain how economic and social factors contribute to the forecasting process [30]. We analysed the overall feature importance ranking and how the importance of top features varies with forecasting horizons. We also investigated the dependency of top factors with the most correlated feature. Figure 5 shows the results of the analysis for the East Midlands.

In Figure 5a, GDP, energy markets, and regional economics are the top three factors impacting electricity demand in the East Midlands. As expected, higher GDP impacts demand positively, as with the text-based social factors of energy markets and regional economics. In Figure 5b, GDP is always more important than the other two factors, and shows weekly patterns. The SHAP values of energy markets and regional




Fig. 5: Feature importance analysis plots by SHAP values in the East Midlands. **(a)** and **(b)** are SHAP values overall and on all horizons. In **(b)**, G is for GDP, E is for energy markets, R is for regional economics, **(c)** is the dependency plot of GDP and unemployment, and **(d)** is the dependency plot of energy markets and strikes.

economics are relatively stationary. The top features can be different in the other regions. In the West Midlands and Ireland, inflation ranks first in all the economic and social factors. Unemployment ranks in the top three in South Wales, the South West, and Northern Ireland.

Figure 5c and Figure 5d demonstrate the top two features and related interaction with the features detected by SHAP values on the whole dataset. In Figure 5c, the SHAP values and the actual GDP are positively correlated, and the SHAP values are higher under low unemployment rates. In Figure 5d, the SHAP values of energy markets follows an inverted U-shape. At the lower level of strikes, the SHAP values are positive correlated with the actual energy markets. The SHAP values start decreasing with the actual energy markets under a higher level of strikes.

Discussion

From the analysis of the results above, this study aims at answering the following research questions, and is thus beneficial for a better understanding of electricity demand, its potential application in forecasting, and the advantages and limitations of using NLP.

1. Is there any relationship between national news and regional electricity demand?

This question was initially answered in the work of [27] at the national level. This second study shows that there is still a stable relationship between national news and regional electricity demand. Results 1 identifies that social factors from national news, such as energy markets and regional economics, can result in a Granger-cause change in regional demand. National economic features can impact regional demand by Granger causing social factors. According to Results 4, news information plays a rather important role in forecasting models and we can also detect the complex interaction between different economic or social factors.

2. Can textual news be used for practical applications to improve regional demand forecasting?

NLP techniques are powerful for extracting numeric features, such as sentiments and topics, from textual news, and have been used to forecast crude oil prices [23], taxi demand [31], and so on. In this study, adding textual features to our forecasting model improved both deterministic and probabilistic performance, as evidenced by Results 2 and Results 3. In Table 3, the results have been averaged across all regions and horizons. They show better performance when textual features are added, with improvements of 3.9% and 2.3% for deterministic and probabilistic tasks, respectively.

Table 3: Horizon and regional averaged model performance of CRPS with and without economic and social factors. S and E indicate social and economic factors. N (No) and Y (Yes) are indicators to show if the factor (S or E) participates in the model.

	E	Deterministic		Probabilistic	
		N	Y	N	Y
S					
N		106.41	102.23	0.4275	0.4312
Y		102.31	100.93	0.4177	0.4251

3. Which topics have a higher impact on electricity demand?

SHAP values offer a method to quantify feature importance in forecasting models. As described in Results 4, the significance of features varies across regions. Key economic indicators such as GDP, inflation, and unemployment are influential on regional demand. Additionally, social factors such as energy markets, regional economics,

military conflicts, strikes, and transportation rank higher than the other factors.

4. Are the results consistent for different regions?

The study reveals a mixed picture of improvements across the six regions under investigation. Notable advancements were observed in four regions, demonstrating the importance of the textual features. However, in South Wales and Ireland, the results showed marginal improvements, indicating a potential need for further refinement or alternative approaches in these regions.

In Ireland, this phenomenon was in part expected, since this dataset was added for comparison with the UK regions, but the model was fed with textual data and economic indicators from the UK, not from Ireland. On the one side, Ireland is close to the UK, sharing a common border and a common language. Furthermore, news about politics in Northern Ireland was indicated as important for the prediction of UK national load in [27]. On the other hand, BBC news is clearly more related to the UK and not directly pertinent to Ireland, therefore, the lack of relevant impact observed is understandable. Regarding South Wales, its demand has been found to be the most challenging to predict, as shown in Figure 2 and Figure 3.

In the East Midlands and Northern Ireland, our results highlighted the significance of economic features as primary drivers of the observed outcomes, as seen in Table 1. This finding underscores the unique economic characteristics of the regions and their influence on the overall results. Conversely, in the remaining regions, social textual features emerged as more critical as seen in Table 1. This outcome emphasised the importance of incorporating textual features in future studies or strategies.

5. For which forecast horizons is the impact higher?

Table 5 to Table 10 and Table 11 in Supplemental materials show that economic or social factors are more powerful on the first ten horizons. The reasons are twofold. First, the predictive ability of the model decreases on longer horizons. Second, news is an external feature subject to timeliness, and current news is difficult to predict further.

Methods

Dataset

As an extended study of [27], this work still chooses the UK and Ireland as the research cases. Historical electric demand is obtained from [32] and the ENTSO-E transparency platform[33] for the Republic of Ireland and the regions Northern Ireland, East Midlands, West Midlands, South Wales, and South West.

Historical textual news sourced from the BBC news repository is maintained by [34]. The textual features such as word frequencies and topics are extracted by [27]. The thousands of features generated are clustered and aggregated into 'social factors' as described in Methods 2, summarised as: military conflicts, transportation, travel & leisure, sports events, pandemic control, regional economics, strikes, family life, election, and energy markets (see detailed description in Table 4).

In addition to the previous study, we also consider historical economic data. The objective is to verify a few key points. First, we want to ensure that any causality observed between social factors and electricity load is genuine. Second, we aim to confirm that improvements in forecast performance are not merely a reflection of easily obtainable economic data. Specifically, we want to avoid situations where social factors merely mirror economic developments that are readily available in tabular form. We apply the Gross Domestic Product (GDP) per capita, inflation rate, and unemployment rate from a national scale to incorporate the macroeconomic effects on electricity demands from [35].

This research involved building forecasting models based on the machine learning paradigm, where feature engineering is integral. In [Results 2](#), the benchmark model includes features of historical demands, and calendar indicators for holidays and weekends, as set in [27]. Regarding temperatures, 2-m temperatures for the major cities in each region were taken from ERA5, the global climate and weather dataset in the European Centre for Medium-Range Weather Forecasts (ECMWF) [36]. Concretely, we collected the temperatures in Leicester for the East Midlands, Birmingham for the West Midlands, Cardiff for South Wales, Bristol for the South West, Belfast for Northern Ireland, and Dublin for Ireland. Given that longer horizons are forecast in this study, the average hourly and monthly temperatures were employed instead of the instant observations to maintain the central trend and seasonality.

Feature Clustering

In the work [27], the authors extracted thousands of numerical features from vast news text with different NLP methods, covering statistical counts, word frequencies, public sentiments, topic distributions, and word embeddings. The authors stated that not all the features are useful for forecasting tasks, and features from different sources may have overlaps semantically. For example, the frequency of the word ‘war’ and war-related topic features refer to the same content, and the shapes of their feature series are similar.

Based on the above findings, we grouped all the features in [27] into ten clusters by hierarchical clustering and took the centroids from each cluster as social factors. Our method has three advantages: i) avoiding working with a large number of features by focusing on the ten centroids; ii) reducing the semantic overlaps between features by maximising the distances between clusters; iii) keeping the explanation while simplifying the analysis procedure. We summarised the content of each centroid as social factors (‘S’) by the features included as follows: S1 for military conflicts, S2 for transportation, S3 for travel & leisure, S4 for sports events, S5 for pandemic control, S6 for regional economics, S7 for strikes, S8 for family life, S9 for election, and S10 for energy markets.

Forecasting models and settings

A Light Gradient Boosting Machine (LGBM) is used as the deterministic forecasting model [37]. Based on decision trees, LGBM introduces Exclusive Feature Bundling (EFB) and Gradient-based One-Side Sampling (GOSS) to decrease computing time

and resources while maintaining accuracy. EFB can reduce the number of features by bundling them into one dimension, especially sparse dummy features. The GOSS method is effective in reducing the training samples. GOSS considers data samples with higher gradients more critical when computing information gain for splitting a node. By keeping the top a samples and randomly selecting from the rest of the samples with lower gradients, GOSS can speed up the process and lead to more accurate results.

In probabilistic forecasting, LGBM usually trains different models for each target quantile to depict the total forecasting distributions. If there are multiple quantiles, it will take more time to train and maintain the quantile regression models. Instead of training model $\hat{y}_q = f(\mathbf{x})$ for all quantiles q , Probabilistic Gradient Boosting Machines (PGBM) focus on learning the key parameters of an assumed target distribution [38]. In our case, we assume the data are normal distribution, that PGBM can learn $(\mu_{\hat{y}}, \sigma_{\hat{y}}^2) = f(\mathbf{x})$, and that the forecasting target \hat{y} can be sampled or inferred from $\hat{y} \sim N(\mu_{\hat{y}}, \sigma_{\hat{y}}^2)$.

Before model training, the dataset is split into a training set (2020.06.01 to 2022.11.30), a validation set (2022.12.01 to 2022.12.31), and a test set (2023.01.01-2023.05.31). All the data are standardised according to the training set to ensure the features are on the same scale and follow a normal distribution. For multi-horizon forecasting tasks in each region, we built 30 models for 30 horizons. In each horizon, we wrapped the hourly data into a multi-regression form, with 24 targets for 24 hours one day ahead. For sample i and horizon h , the formulas are $\hat{\mathbf{y}}_{i,h} = f_d(\mathbf{x}_i)$ for deterministic forecasting and $(\boldsymbol{\mu}_{\hat{\mathbf{y}}_{i,h}}, \boldsymbol{\sigma}_{\hat{\mathbf{y}}_{i,h}}^2) = f_p(\mathbf{x}_i)$ for probabilistic forecasting, where $\mathbf{x}_i = (x_{i,1}, x_{i,2}, \dots, x_{i,m})$ is the vector with m features, $f_d(*)$ and $f_p(*)$ are LGBM and PGBM models, $\hat{\mathbf{y}}_{i,h} = (\hat{y}_{1,i,h}, \hat{y}_{2,i,h}, \dots, \hat{y}_{24,i,h})$ is the target vector with 24 hours, $\boldsymbol{\mu}_{\hat{\mathbf{y}}_{i,h}} = (\mu_{\hat{y}_{1,i,h}}, \mu_{\hat{y}_{2,i,h}}, \dots, \mu_{\hat{y}_{24,i,h}})$ and $\boldsymbol{\sigma}_{\hat{\mathbf{y}}_{i,h}}^2 = (\sigma_{\hat{y}_{1,i,h}}^2, \sigma_{\hat{y}_{2,i,h}}^2, \dots, \sigma_{\hat{y}_{24,i,h}}^2)$ are mean and variance vectors with 24 hours. $f_d(*)$ and $f_p(*)$ both take Mean Squared Error (MSE) on the validation set as the loss function. For hyper-parameter searching, we employed Optuna, an automatic optimisation framework [39]. Optuna designs efficient searching and pruning strategies and allows users to build search space dynamically, thus achieving high performance with limited computing resources. For simplicity, we refer to both LGBM and PGBM as GBM in the main text.

Benchmark models

Focusing on deterministic forecasting, this study considers several benchmark models: Persistence Forecasts (PF), Smart Climatology Forecasts (SCF), and a combination of PF and SCF (PF-SCF). PF and SCF are naive models with only history observations and are useful for short- and long-term forecasting [1]. PF takes the last observed value y_t as the forecast for all the horizons: $\hat{y}_{t+h} = y_t (h = 1, \dots, H)$, where H is the horizon of interest. As to SCF, it is a smarter version of Climatology Forecasts (CF). CF is based on the average of all the historical observations: $\hat{y}_{t+h} = \sum_{i=1}^N y_i (h = 1, \dots, H)$, where N is the number of observations. Considering the seasonal patterns in electricity demands, SCF uses the average history data at the same month and hour as the

forecasts:

$$\hat{y}_{t+h} = \frac{1}{N_h} \sum_{i=1}^N y_i \cdot \mathbf{1}_{(\text{month}(t+h)=\text{month}(t_i) \cap (\text{hour}(t+h)=\text{hour}(t_i)))}, \quad (1)$$

where N_h is the number of historical data with the same month and hour as timestamp $t+h$, $\mathbf{1}$ is the indicating function, $\text{month}(t+h)$ and $\text{hour}(t+h)$ are the month and hour of timestamp $t+h$, $\text{month}(t_i)$ and $\text{hour}(t_i)$ are the month and hour of timestamp t_i . According to [1], the combination of PF and SCF is better than the single model, with the forecast calculated as: $\hat{y}_{PF-SCF,t+h} = \alpha \hat{y}_{PF,t+h} + (1-\alpha) \hat{y}_{SCF,t+h}$, where $0 < \alpha < 1$ is the coefficient of the convex combination. In our case, we combine the forecasts from PF and SCF using the LASSO model for the PF-SCF model.

Evaluation metrics

For deterministic forecasting, Root Mean Square Error (RMSE) and Mean Absolute Percentage Error (MAPE) are used to measure the deviation between the truth and the forecasts. Continuously Ranked Probability Score (CRPS) measures the discrepancy between the cumulative distributions of the truth and forecasts. The three metrics can be calculated as follows:

$$\text{RMSE} = \frac{1}{30D} \sum_{H=1}^{30} \sum_{d=1}^D \sqrt{\frac{1}{24} \sum_{h=1}^{24} (y_{h,d,H} - \hat{y}_{h,d,H})^2}, \quad (2)$$

$$\text{MAPE} = \frac{100\%}{30 \times D \times 24} \sum_{H=1}^{30} \sum_{d=1}^D \sum_{h=1}^{24} \left| \frac{y_{h,d,H} - \hat{y}_{h,d,H}}{y_{h,d,H}} \right|, \quad (3)$$

$$\text{CRPS} = \frac{100\%}{30 \times D \times 24} \sum_{H=1}^{30} \sum_{d=1}^D \sum_{h=1}^{24} \int [F(\hat{y}_{h,d,H}) - \mathbf{1}_{\{\hat{y}_{h,d,H} \geq y_{h,d,H}\}}]^2 d\hat{y}, \quad (4)$$

where the RMSE and MAPE are first calculated on each day ahead (24 hours), then averaged on all of the days and all of the horizons. $y_{h,d,H}$ and $\hat{y}_{h,d,H}$ are the true electricity demand and forecasts. D is the number of days in the test set for each model. Hour h varies from 1 to 24, and horizon H varies from 1 to 30. In Formula 4, CRPS is calculated on each hour and averaged as RMSE and MAPE. $F(y_{h,d,H})$ is the Cumulative Distribution Function (CDF) of $\hat{y}_{h,d,H}$, and in our case it follows normal distribution. $\mathbf{1}$ is the indicator function.

Granger test

Based on the Vector Auto-Regression (VAR) model, the target of the Granger test is to detect whether a time series X can forecast another series Y [40]. If so, it is said that series X Granger causes Y . To conduct the test, the Granger method constructs two VAR models for Y :

$$Y_t = \alpha_0 + \alpha_1 Y_{t-1} + \dots + \alpha_p Y_{t-p} + \varepsilon_t, \quad (5)$$

$$Y_t = \beta_0 + \beta_1 Y_{t-1} + \cdots + \beta_p Y_{t-p} + \gamma_1 X_{t-1} + \cdots + \gamma_q X_{t-q} + \epsilon_t, \quad (6)$$

where p and q are lag orders, α_i , β_i , and γ_i are coefficients, ε_t and ϵ_t are error terms. As seen from the second VAR model, lag terms of X are added. If the second model is significantly better than the first, then X can forecast Y . F-statistic can be applied to test whether there is a significant difference between the two by comparing the Residual Sum of Squares (RSS) of both models. Given n samples of the data, $RSS = \sum_{i=1}^n (y_i - \hat{y}_i)^2$. Note R_0 and R_1 are RSS for the two VAR models, F-statistic is computed as follows:

$$F = \frac{(R_1 - R_0)/q}{R_0/(n - p - q)} \sim F(q, n - p - q). \quad (7)$$

The null hypothesis H_0 is that X cannot forecast Y . By comparing the F-statistic with F_δ under significant level δ , if $F > F_\delta$, then H_0 is rejected, indicating that X can Granger-cause Y .

Acknowledgement

The author Yun Bai was supported by the program of the China Scholarship Council (CSC Nos. 202106020064). Part of this work is carried out in the framework of the Fine4CAST project, funded by France 2030 (ANR reference: 22-PETA-0008).

Declaration of interests

- All affiliations are listed on the title page of the manuscript.
- All funding sources for this study are listed in the “acknowledgments” section of the manuscript.
- We, the authors and our immediate family members, have no financial interests to declare.
- We, the authors and our immediate family members, have no positions to declare and are not members of the journal’s advisory board.
- We, the authors and our immediate family members, have no related patent applications or registrations to declare.
- The authors declare no competing interests.

References

- [1] Metaxiotis, K., Kagiannas, A., Askounis, D., Psarras, J.: Artificial intelligence in short term electric load forecasting: a state-of-the-art survey for the researcher. *Energy conversion and Management* **44**(9), 1525–1534 (2003) [https://doi.org/10.1016/S0196-8904\(02\)00148-6](https://doi.org/10.1016/S0196-8904(02)00148-6)
- [2] Weron, R.: Electricity price forecasting: A review of the state-of-the-art with a look into the future. *International journal of forecasting* **30**(4), 1030–1081 (2014) <https://doi.org/10.1016/j.ijforecast.2014.08.008>

- [3] Balasubramanian, S., Balachandra, P.: Effectiveness of demand response in achieving supply-demand matching in a renewables dominated electricity system: A modelling approach. *Renewable and Sustainable Energy Reviews* **147**, 111245 (2021) <https://doi.org/10.1016/j.rser.2021.111245>
- [4] Gerossier, A., Girard, R., Bocquet, A., Kariniotakis, G.: Robust day-ahead forecasting of household electricity demand and operational challenges. *Energies* **11**(12), 3503 (2018) <https://doi.org/10.3390/en1123503>
- [5] MacMackin, N., Miller, L., Carriveau, R.: Modeling and disaggregating hourly effects of weather on sectoral electricity demand. *Energy* **188**, 115956 (2019) <https://doi.org/10.1016/j.energy.2019.115956>
- [6] De Felice, M., Alessandri, A., Catalano, F.: Seasonal climate forecasts for medium-term electricity demand forecasting. *Applied Energy* **137**, 435–444 (2015) <https://doi.org/10.1016/j.apenergy.2014.10.030>
- [7] Li, G., Li, Y., Roozitalab, F.: Midterm load forecasting: A multistep approach based on phase space reconstruction and support vector machine. *IEEE Systems Journal* **14**(4), 4967–4977 (2020) <https://doi.org/10.1109/JSYST.2019.2962971>
- [8] Assis Cabral, J., Legey, L.F.L., Freitas Cabral, M.V.: Electricity consumption forecasting in brazil: A spatial econometrics approach. *Energy* **126**, 124–131 (2017) <https://doi.org/10.1016/j.energy.2017.03.005>
- [9] Bianco, V., Manca, O., Nardini, S.: Electricity consumption forecasting in italy using linear regression models. *Energy* **34**(9), 1413–1421 (2009) <https://doi.org/10.1016/j.energy.2009.06.034>
- [10] Ruan, G., Wu, D., Zheng, X., Zhong, H., Kang, C., Dahleh, M.A., Sivaranjani, S., Xie, L.: A cross-domain approach to analyzing the short-run impact of covid-19 on the us electricity sector. *Joule* **4**(11), 2322–2337 (2020) <https://doi.org/10.1016/j.joule.2020.08.017>
- [11] Farrokhhabadi, M., Browell, J., Wang, Y., Makonin, S., Su, W., Zareipour, H.: Day-ahead electricity demand forecasting competition: Post-covid paradigm. *IEEE Open Access Journal of Power and Energy* **9**, 185–191 (2022) <https://doi.org/10.1109/OAJPE.2022.3161101>
- [12] Wilhite, H., Shove, E., Lutzenhiser, L., Kempton, W.: The legacy of twenty years of energy demand management: we know more about individual behaviour but next to nothing about demand. In: *Society, Behaviour, and Climate Change Mitigation*, pp. 109–126 (2000). https://doi.org/10.1007/0-306-48160-X_4
- [13] Powells, G., Bulkeley, H., Bell, S., Judson, E.: Peak electricity demand and the flexibility of everyday life. *Geoforum* **55**, 43–52 (2014) <https://doi.org/10.1016/j.geoforum.2014.04.014>

- [14] Wagy, M.D., Bongard, J.C., Bagrow, J.P., Hines, P.D.: Crowdsourcing predictors of residential electric energy usage. *IEEE Systems Journal* **12**(4), 3151–3160 (2017) <https://doi.org/10.1109/JSYST.2017.2778144>
- [15] Hargreaves, T., Middlemiss, L.: The importance of social relations in shaping energy demand. *Nature Energy* **5**(3), 195–201 (2020) <https://doi.org/10.1038/s41560-020-0553-5>
- [16] Doumèche, N., Allieux, Y., Goude, Y., Rubrichi, S.: Human spatial dynamics for electricity demand forecasting: the case of france during the 2022 energy crisis. *arXiv preprint arXiv:2309.16238* (2023) <https://doi.org/10.48550/arXiv.2309.16238>
- [17] Fu, C., Miller, C.: Using google trends as a proxy for occupant behavior to predict building energy consumption. *Applied Energy* **310**, 118343 (2022) <https://doi.org/10.1016/j.apenergy.2021.118343>
- [18] Allen, M.R., Fernandez, S.J., Fu, J.S., Olama, M.M.: Impacts of climate change on sub-regional electricity demand and distribution in the southern united states. *Nature Energy* **1**(8), 1–9 (2016) <https://doi.org/10.1038/nenergy.2016.103>
- [19] Subramaniam, L.V., Roy, S., Faruquie, T.A., Negi, S.: A survey of types of text noise and techniques to handle noisy text. In: *Proceedings of The Third Workshop on Analytics for Noisy Unstructured Text Data*, pp. 115–122 (2009). <https://doi.org/10.1145/1568296.1568315>
- [20] Li, X., Shang, W., Wang, S.: Text-based crude oil price forecasting: A deep learning approach. *International Journal of Forecasting* **35**(4), 1548–1560 (2019) <https://doi.org/10.1016/j.ijforecast.2018.07.006>
- [21] Wu, B., Wang, L., Lv, S.-X., Zeng, Y.-R.: Effective crude oil price forecasting using new text-based and big-data-driven model. *Measurement* **168**, 108468 (2021) <https://doi.org/10.1016/j.measurement.2020.108468>
- [22] Liu, J., Huang, X.: Forecasting crude oil price using event extraction. *IEEE Access* **9**, 149067–149076 (2021) <https://doi.org/10.1109/ACCESS.2021.3124802>
- [23] Bai, Y., Li, X., Yu, H., Jia, S.: Crude oil price forecasting incorporating news text. *International Journal of Forecasting* **38**(1), 367–383 (2022) <https://doi.org/10.1016/j.ijforecast.2021.06.006>
- [24] Obst, D., De Vilmarest, J., Goude, Y.: Adaptive methods for short-term electricity load forecasting during covid-19 lockdown in france. *IEEE transactions on power systems* **36**(5), 4754–4763 (2021) <https://doi.org/10.1109/TPWRS.2021.3067551>
- [25] Bai, Y., Camal, S., Michiorri, A.: Electricity demand forecasting through natural

- language processing with long short-term memory networks. In: 2023 IEEE PES Innovative Smart Grid Technologies Europe (ISGT EUROPE), pp. 1–5 (2023). <https://doi.org/10.1109/ISGTEUROPE56780.2023.10408499>
- [26] Wang, D., Gan, J., Mao, J., Chen, F., Yu, L.: Forecasting power demand in china with a cnn-lstm model including multimodal information. *Energy* **263**, 126012 (2023) <https://doi.org/10.1016/j.energy.2022.126012>
- [27] Bai, Y., Camal, S., Michiorri, A.: News and load: A quantitative exploration of natural language processing applications for forecasting day-ahead electricity system demand. *IEEE Transactions on Power Systems* (2024) <https://doi.org/10.1109/TPWRS.2024.3361074>
- [28] Andersson, J.: Extra bank holiday approved to mark King’s coronation. <https://www.bbc.com/news/uk-63528830> (2022)
- [29] Montel, M., Yakovets, A., et al.: holidays documentation. pypi.org/project/holidays (2014)
- [30] Lundberg, S.M., Lee, S.-I.: A unified approach to interpreting model predictions. In: *Advances in Neural Information Processing Systems*, vol. 30 (2017). https://proceedings.neurips.cc/paper_files/paper/2017/file/8a20a8621978632d76c43dfd28b67767-Paper.pdf
- [31] Rodrigues, F., Markou, I., Pereira, F.C.: Combining time-series and textual data for taxi demand prediction in event areas: A deep learning approach. *Information Fusion* **49**, 120–129 (2019) <https://doi.org/10.1016/j.inffus.2018.07.007>
- [32] nged.data@nationalgrid.co.uk: National Grid. <https://www.nationalgrid.co.uk/our-network/live-data-feed-application-map>
- [33] ENTSO-E. Transparency Platform. <https://transparency.entsoe.eu>
- [34] Matthew, S.: BBC News Front Page Archive. <https://dracos.co.uk/made/bbc-news-archive/archive.php>
- [35] Office for National Statistics: Census Datasets. <https://www.ons.gov.uk/census>
- [36] Hersbach, H., Bell, B., Berrisford, P., Biavati, G., Horányi, A., Muñoz Sabater, J., Nicolas, J., Peubey, C., Radu, R., Rozum, I., Schepers, D., Simmons, A., Soci, C., Dee, D., Thépaut, J.-N.: ERA5 hourly data on single levels from 1940 to present. Copernicus Climate Change Service (C3S) Climate Data Store (CDS) (2023). <https://doi.org/10.24381/cds.adbb2d47>
- [37] 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. *Advances in neural information processing systems* **30** (2017)

- [38] Sprangers, O., Schelter, S., Rijke, M.: Probabilistic gradient boosting machines for large-scale probabilistic regression. In: Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, pp. 1510–1520 (2021). <https://doi.org/10.1145/3447548.3467278>
- [39] Akiba, T., Sano, S., Yanase, T., Ohta, T., Koyama, M.: Optuna: A next-generation hyperparameter optimization framework. In: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, pp. 2623–2631 (2019). <https://doi.org/10.48550/arXiv.1907.10902>
- [40] Granger, C.W.: Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: journal of the Econometric Society*, 424–438 (1969) <https://doi.org/10.2307/1912791>

Supplemental materials

Other benchmarks

Generalised Additive Models (GAM) is a generalised version of the linear model, a linear form of several nodes [2]. Each node f is a non-linear combination of features based on spline basis $B(x)$: $f(x) = \sum_{i=1}^m \beta_i B_i(x)$, where m is the dimension of $B(x)$. The GAM for forecasting electricity demand is:

$$\begin{aligned}
 D_t = & \sum_{i=1}^7 \sum_{j=0}^1 \alpha_{i,j} \mathbf{1}_{(dow_t=i) \cap (dls_t=j)} \\
 & + \sum_{i=1}^7 \beta_i D_{1d,t} \mathbf{1}_{dow_t=i} + \gamma D_{1w,t} \\
 & + f_1(t) + f_2(toy_t) + f_3(t, T_t) + f_4(T_{\delta=0.95,t}) \\
 & + f_5(T_{\delta=0.99,t}) + f_6(T_{Min99,t}, T_{Max99,t}) + \epsilon_t,
 \end{aligned} \tag{8}$$

where D_t is the electricity demand at timestamp t , $\alpha_{i,j}$, β_i , and γ are the coefficients that will be optimised, $\mathbf{1}$ is the indicating function, dow_t is the day of week variable, dls_t is a 0-1 variable to show if t is daylight saving time or standard time, $D_{1d,t}$ and $D_{1w,t}$ are the demand of last day and week at the same time t , toy_t is the time of year linearly growing from 0 to 1 during the year, T_t is temperature at time t , $T_{\delta=0.95,t}$ and $T_{\delta=0.99,t}$ are exponentially smoothed temperatures on $\delta = 0.95$ and $\delta = 0.99$: $T_{\delta,t} = \delta T_{\delta,t-1} + (1 - \delta)T_t$, $T_{Min99,t}$ and $T_{Max99,t}$ are the minimum and maximum of $T_{\delta=0.99}$ on a certain day before timestamp t .

Based on ordinary least square regression, Least Absolute Shrinkage and Selection Operator (LASSO) adds an L1 regularisation term to penalise model complexity and make the parameters sparse [3]. Supposing there are m samples and n features, the forecast \hat{y}_i is:

$$\hat{y}_i = \beta_0 + \beta_1 x_{i,1} + \beta_2 x_{i,2} + \dots + \beta_n x_{i,m}, \tag{9}$$

and the loss function of LASSO is:

$$L(\beta) = \frac{1}{2n} \sum_{i=1}^n (y_i - \hat{y}_i)^2 + \lambda \sum_{j=1}^m |\beta_j|, \tag{10}$$

where β_0 is an intercept, β_1 to β_n are coefficients, y_i is the truth value, and λ controls regularisation between 0 and 1.

Support Vector Regression (SVR) is a regression model based on Support Vector Machines (SVM) and captures the non-linear relationship between features and targets [4]. It uses kernel functions to map the inputs to high-dimensional space for a better data split. It also considers margin bounds and uses data outside bounds to fit the model with tolerance to specific errors, thus keeping the model smooth. The process

of solving an SVR model can be transferred to optimisation of a convex problem as:

$$\begin{aligned}
& \min_{w,b,\zeta,\zeta^*} \frac{1}{2} \|w\|^2 + C \sum_{i=1}^m (\zeta_i + \zeta_i^*) \\
& \text{subject to: } y_i - w^T \phi(x_i) - b \leq \epsilon + \zeta_i \\
& \quad w^T \phi(x_i) + b - y_i \leq \epsilon + \zeta_i^* \\
& \quad \zeta_i, \zeta_i^* \geq 0,
\end{aligned} \tag{11}$$

where w is the weight vector, b is the bias, $\phi(*)$ is the kernel function, ζ_i and ζ_i^* are slack variables, C is the penalty parameter, and ϵ is the error tolerance.

ExtraTrees Regressor (ETR) is an ensemble learning algorithm that integrates the results of multiple decision trees to improve the model generalisation [5]. Different from random forest, ETR randomly selects and splits a feature subset, and trains on the randomly selected samples. The random choice of ETR increases model diversity and reduces the risk of over-fitting. When splitting a node, given the randomly selected sample subset S and feature f_i , the maximum and minimum of f_i in S are noted as $f_{i,max}^S$ and $f_{i,min}^S$, a split can be defined as $s_i = [f_i < f_{i,c}]$, where $f_{i,c}$ is a cut-point. For the randomly selected N features $\mathbf{F} = [f_1, f_2, \dots, f_N]$, one can get N splits for all the features $\mathbf{S} = [s_1, s_2, \dots, s_N]$. The final optimal split s^* is computed by $s^* = \underset{s \in \mathbf{S}}{\operatorname{argmax}} \operatorname{score}(s, S)$. In the regression task, if S_l and S_r are two subsets of S split by s , and $S_l \cup S_r = S$, $\operatorname{score}(s, S)$ can be calculated as follows:

$$\operatorname{score}(s, S) = \frac{\operatorname{var}\{y|S\} - \frac{|S_l|}{|S|} \operatorname{var}\{y|S_l\} - \frac{|S_r|}{|S|} \operatorname{var}\{y|S_r\}}{\operatorname{var}\{y|S\}}, \tag{12}$$

where $\operatorname{var}\{y|S\}$ is the variance of output y in sample subset S , $|*|$ is the number of samples.

Results figures and tables




Fig. 6: Electricity demand plots for all the regions. a-f, (a): East Midlands, **(b):** West Midlands, **(c):** South Wales, **(d):** South West, **(e)** Northern Ireland, **(f)** Ireland.

References in supplemental materials

- [1] Murphy, A.: Climatology, persistence, and their linear combination as standards of reference in skill scores. *Weather and forecasting* **7**(4), 692-698 (1992). [https://doi.org/10.1175/1520-0434\(1992\)007<0692:CPATLC>2.0.CO;2](https://doi.org/10.1175/1520-0434(1992)007<0692:CPATLC>2.0.CO;2)
- [2] Obst, D., De Vilmarest, J., Goude, Y.: Adaptive methods for short-term electricity load forecasting during COVID-19 lockdown in France. *IEEE transactions on power systems* **36**(5), 4754-4763 (2021). <https://doi.org/10.1109/TPWRS.2021.3067551>
- [3] Santosa, F., Symes, W. W.: Linear inversion of band-limited reflection seismograms. *SIAM Journal on Scientific and Statistical Computing* **7**(4), 1307–1330 (1986). <https://doi.org/10.1137/0907087>
- [4] Smola, A.J., Schölkopf, B. A tutorial on support vector regression. *Statistics and computing* **14**, 199-222 (2004) <https://doi.org/10.1023/B:STCO.0000035301.49549.88>
- [5] Geurts, P., Ernst, D., Wehenkel, L.: Extremely randomized trees. *Machine learning* **63**, 3-42. (2006). <https://doi.org/10.1007/s10994-006-6226-1>

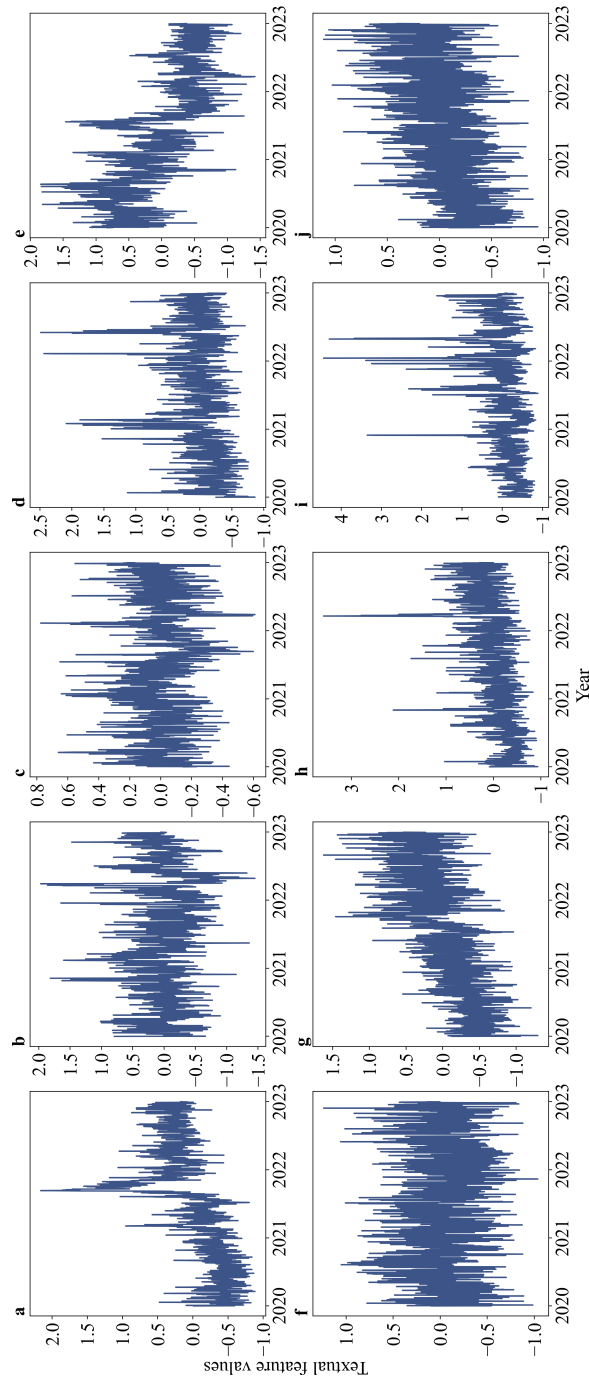


Fig. 7: Ten social factor plots. a-j, (a): military conflicts, **(b):** transportation, **(c):** travel & leisure, **(d):** sports events, **(e)** pandemic control, **(f)** regional economics, **(g)** strikes, **(h)** family life, **(i)** election, **(j)** energy markets. Note that each name of the centroids is summarised according to its content, for example the frequency of key words and news topics; see [27] for a detailed introduction.

Table 4: Selected key words in each clustering centroid.

Military conflicts	Transportation	Travel & leisure	Sports events	Pandemic control
international	road	places	sport	lockdown
control	park	summer	match	test
fight	car	travel	players	highest
injured	driving	university	final	restrictions
military	driver	france	events	positive
forces	bus	airport	win	health
defence	traffic	flight	points	cases
russian	bridge	airline	opening	virus
war	vehicle	holiday	football	deaths
Regional economics	Strikes	Family life	Election	Energy markets
million	times	couple	labour	global
jobs	forced	family	vote	power
scotland	strike	wife	prime	industry
business	pay	life	parliament	economy
increase	launched	mother	election	trade
wales	suffered	society	political	energy
council	staff	son	position	economic
cardiff	rail	daughter	leader	demand
glasgow	services	born	party	deal

Table 5: East Midlands: RMSE and MAPE errors on all the horizons. LASSO, SVR, ETR, and LGBM are models with prior grid parameter search, LGBM-FT is the fine-tuned LGBM with Optuna parameter searching. LGBM-FT-E features economic factors, LGBM-FT-S features factors, LGBM-FT-ES features both.

Methods	Horizons 1-10			Horizons 11-20			Horizons 21-30			Horizons 1-30		
	RMSE(MW)	MAPE(%)	RMSE(MW)	MAPE(%)	RMSE(MW)	MAPE(%)	RMSE(MW)	MAPE(%)	RMSE(MW)	MAPE(%)	RMSE(MW)	MAPE(%)
PF	258.73	8.03	310.01	10.04	342.03	11.50	303.59	9.86				
SCF	264.05	9.39	264.05	9.39	264.05	9.39	264.05	9.39				
PF-SCF	216.90	7.63	248.53	8.97	254.39	9.18	239.94	8.59				
GAM	200.42	5.84	246.11	7.65	239.60	7.34	228.71	6.94				
LASSO	166.07	5.61	226.10	8.23	231.46	8.36	207.88	7.40				
SVR	164.41	5.49	247.05	8.98	243.66	8.73	218.37	7.73				
ETR	225.40	7.96	262.44	9.37	257.39	9.22	248.41	8.85				
LGBM	209.10	7.21	253.24	9.11	245.15	8.76	235.83	8.36				
LGBM-FT	126.09	2.92	144.64	3.44	136.55	3.21	135.76	3.19				
LGBM-FT-E	123.56	2.85	128.79	2.95	124.99	2.88	125.78	2.89				
LGBM-FT-S	123.04	2.88	142.50	3.43	136.35	3.25	133.96	3.19				
LGBM-FT-ES	120.72	2.79	125.60	2.89	123.62	2.85	123.31	2.84				

Table 6: West Midlands: RMSE and MAPE errors on all the horizons. LASSO, SVR, ETR, and LGBM are models with prior grid parameter search, LGBM-FT is the fine-tuned LGBM with Optuna parameter searching. LGBM-FT-E features economic factors, LGBM-FT-S features social factors, LGBM-FT-ES features both.

Methods	Horizons 1-10			Horizons 11-20			Horizons 21-30			Horizons 1-30		
	RMSE(MW)	MAPE(%)	MAPE(%)	RMSE(MW)	MAPE(%)	MAPE(%)	RMSE(MW)	MAPE(%)	MAPE(%)	RMSE(MW)	MAPE(%)	MAPE(%)
PF	100.95	6.66	6.66	255.57	7.98	7.98	267.15	8.46	8.46	207.89	7.70	7.70
SCF	149.09	4.67	4.67	149.09	4.67	4.67	149.09	4.67	4.67	149.09	4.67	4.67
PF-SCF	127.67	3.95	3.95	137.95	4.31	4.31	139.62	4.36	4.36	135.08	4.21	4.21
GAM	155.36	4.14	4.14	169.55	4.94	4.94	148.98	4.53	4.53	157.96	4.54	4.54
LASSO	100.95	3.06	3.06	114.57	3.53	3.53	121.42	3.68	3.68	112.31	3.42	3.42
SVR	102.08	3.09	3.09	115.78	3.57	3.57	119.16	3.68	3.68	112.34	3.45	3.45
ETR	112.69	3.43	3.43	131.35	4.10	4.10	131.67	4.13	4.13	125.24	3.89	3.89
LGBM	107.06	3.23	3.23	121.42	3.75	3.75	125.63	3.90	3.90	118.04	3.63	3.63
LGBM-FT	105.68	3.19	3.19	117.14	3.60	3.60	123.13	3.81	3.81	115.32	3.53	3.53
LGBM-FT-E	102.31	3.09	3.09	111.06	3.39	3.39	115.21	3.53	3.53	109.53	3.34	3.34
LGBM-FT-S	101.65	3.06	3.06	110.42	3.37	3.37	115.08	3.52	3.52	109.05	3.32	3.32
LGBM-FT-ES	100.79	3.03	3.03	110.00	3.34	3.34	112.93	3.43	3.43	107.91	3.27	3.27

Table 7: South Wales: RMSE and MAPE errors on all the horizons. LASSO, SVR, ETR, and LGBM are models with prior grid parameter search, LGBM-FT is the fine-tuned LGBM with Optuna parameter searching. LGBM-FT-E features economic factors, LGBM-FT-S features social factors, LGBM-FT-ES features both.

Methods	Horizons 1-10		Horizons 11-20		Horizons 21-30		Horizons 1-30	
	RMSE(MW)	MAPE(%)	RMSE(MW)	MAPE(%)	RMSE(MW)	MAPE(%)	RMSE(MW)	MAPE(%)
PF	87.00	6.31	105.91	7.93	123.00	9.50	105.30	7.91
SCF	149.32	12.80	149.32	12.80	149.32	12.80	149.32	12.80
PF-SCF	99.47	8.22	120.91	10.31	127.77	10.94	116.05	9.82
GAM	177.77	10.91	217.14	13.95	189.13	13.09	194.68	12.65
LASSO	72.57	5.69	104.04	8.75	113.86	9.63	96.82	8.02
SVR	74.87	5.88	110.98	9.34	115.50	9.70	100.45	8.31
ETR	97.31	7.90	117.56	9.81	118.71	9.90	111.19	9.20
LGBM	91.50	7.40	115.22	9.79	115.99	9.73	107.57	8.97
LGBM-FT	87.49	6.99	111.45	9.26	110.39	9.13	103.11	8.46
LGBM-FT-E	97.27	7.83	121.49	10.12	108.78	9.05	109.18	9.00
LGBM-FT-S	87.54	7.03	110.57	9.18	107.35	8.85	101.82	8.35
LGBM-FT-ES	95.51	7.69	119.44	9.91	110.88	9.19	108.61	8.93

Table 8: South West: RMSE and MAPE errors on all the horizons. LASSO, SVR, ETR, and LGBM are models with prior grid parameter search, LGBM-FT is the fine-tuned LGBM with Optuna parameter searching. LGBM-FT-E features economic factors, LGBM-FT-S features social factors, LGBM-FT-ES features both.

Methods	Horizons 1-10			Horizons 11-20			Horizons 21-30			Horizons 1-30		
	RMSE(MW)	MAPE(%)	MAPE(%)	RMSE(MW)	MAPE(%)	MAPE(%)	RMSE(MW)	MAPE(%)	MAPE(%)	RMSE(MW)	MAPE(%)	MAPE(%)
PF	116.27	7.39	146.28	9.92	171.20	12.04	144.58	9.78				
SCF	211.72	16.24	211.72	16.24	211.72	16.24	211.72	16.24				
PF-SCF	125.62	9.24	151.98	11.68	161.42	12.43	146.34	11.12				
GAM	116.05	7.88	139.84	9.88	150.66	10.66	135.52	9.47				
LASSO	95.10	6.74	131.20	10.00	141.25	10.80	122.52	9.18				
SVR	91.76	6.43	121.52	9.20	131.41	9.96	114.90	8.53				
ETR	120.24	8.81	136.91	10.44	140.05	10.67	132.40	9.97				
LGBM	111.53	7.96	128.93	9.79	132.29	10.00	124.25	9.25				
LGBM-FT	110.39	7.84	129.54	9.77	132.76	9.98	124.23	9.20				
LGBM-FT-E	106.81	7.76	120.85	9.24	124.10	9.49	117.25	8.83				
LGBM-FT-S	105.14	7.63	120.60	9.20	124.32	9.51	116.69	8.78				
LGBM-FT-ES	107.08	7.86	119.43	9.15	121.96	9.29	116.16	8.77				

Table 9: Northern Ireland: RMSE and MAPE errors on all the horizons. LASSO, SVR, ETR, and LGBM are models with prior grid parameter search, LGBM-FT is the fine-tuned LGBM with Optuna parameter searching. LGBM-FT-E features economic factors, LGBM-FT-S features social factors, LGBM-FT-ES features both.

Methods	Horizons 1-10		Horizons 11-20		Horizons 21-30		Horizons 1-30	
	RMSE(MW)	MAPE(%)	RMSE(MW)	MAPE(%)	RMSE(MW)	MAPE(%)	RMSE(MW)	MAPE(%)
PF	97.06	9.94	104.80	10.70	107.10	11.00	102.99	10.55
SCF	67.37	7.06	67.37	7.06	67.37	7.06	67.37	7.06
PF-SCF	57.63	6.10	60.50	6.39	59.53	6.21	59.22	6.23
GAM	83.64	8.87	98.99	10.47	94.13	10.03	92.25	9.79
LASSO	48.85	5.11	53.86	5.66	54.02	5.60	52.24	5.46
SVR	49.79	5.21	55.67	5.88	54.26	5.63	53.24	5.57
ETR	54.65	5.74	60.05	6.39	57.99	6.06	57.56	6.06
LGBM	50.60	5.30	53.91	5.69	52.81	5.50	52.44	5.50
LGBM-FT	51.63	5.43	55.29	5.84	54.05	5.64	53.66	5.64
LGBM-FT-E	48.03	4.95	50.58	5.25	49.64	5.24	49.42	5.15
LGBM-FT-S	48.54	5.14	51.58	5.47	49.96	5.24	50.03	5.28
LGBM-FT-ES	46.78	4.86	50.06	5.21	49.20	5.15	48.68	5.07

Table 10: Ireland: RMSE and MAPE errors on all the horizons. LASSO, SVR, ETR, and LGBM are models with prior grid parameter search, LGBM-FT is the fine-tuned LGBM with Optuna parameter searching. LGBM-FT-E features economic factors, LGBM-FT-S features social factors, LGBM-FT-ES features both.

Methods	Horizons 1-10			Horizons 11-20			Horizons 21-30			Horizons 1-30		
	RMSE(MW)	MAPE(%)	RMSE(MW)	MAPE(%)	RMSE(MW)	MAPE(%)	RMSE(MW)	MAPE(%)	RMSE(MW)	MAPE(%)	RMSE(MW)	MAPE(%)
PF	270.09	6.15	310.58	7.18	315.32	7.37	298.66	6.90				
SCF	252.77	6.29	252.48	6.27	254.71	6.36	253.32	6.31				
PF-SCF	151.07	3.44	165.88	3.82	163.37	3.79	160.11	3.68				
GAM	187.42	4.09	232.09	5.19	227.39	5.10	215.63	4.79				
LASSO	111.85	2.52	129.08	3.00	130.92	3.05	123.95	2.86				
SVR	114.62	2.60	142.67	3.36	138.60	3.23	131.96	3.06				
ETR	149.09	3.56	169.52	4.09	160.26	3.84	159.62	3.83				
LGBM	129.15	3.01	148.20	3.54	141.01	3.33	139.45	3.29				
LGBM-FT	144.86	3.27	165.63	3.85	151.76	3.45	154.08	3.52				
LGBM-FT-E	157.85	3.63	178.82	4.27	161.02	3.74	165.90	3.88				
LGBM-FT-S	138.99	3.17	159.86	3.76	149.78	3.46	149.54	3.46				
LGBM-FT-ES	150.47	3.47	171.78	4.11	158.24	3.69	160.16	3.76				

Table 11: CRPS results on different models and horizons. PGBM-E is the model featuring economic factors, PGBM-S features social factors, PGBM-ES features both economic and social factors.

Regions	Methods	Horizons			
		1-10	11-20	21-30	1-30
East Midlands	PGBM	0.2571	0.3084	0.2867	0.2841
	PGBM-E	0.2691	0.2964	0.2892	0.2849
	PGBM-S	0.2672	0.3345	0.3092	0.3037
	PGBM-ES	0.2630	0.2991	0.2845	0.2822
West Midlands	PGBM	0.2436	0.2883	0.2957	0.2759
	PGBM-E	0.2423	0.2963	0.2818	0.2735
	PGBM-S	0.2301	0.2593	0.2655	0.2517
	PGBM-ES	0.2337	0.2754	0.2644	0.2578
South Wales	PGBM	0.5893	0.8206	0.7887	0.7329
	PGBM-E	0.6813	0.9394	0.8192	0.8133
	PGBM-S	0.5867	0.8075	0.7755	0.7233
	PGBM-ES	0.6573	0.9036	0.8258	0.7956
South West	PGBM	0.3804	0.4941	0.5141	0.4629
	PGBM-E	0.3815	0.4649	0.4696	0.4387
	PGBM-S	0.3820	0.4695	0.5141	0.4471
	PGBM-ES	0.3963	0.4615	0.4730	0.4436
Northern Ireland	PGBM	0.3656	0.3965	0.3825	0.3815
	PGBM-E	0.3457	0.3548	0.3356	0.3454
	PGBM-S	0.3531	0.3743	0.3608	0.3627
	PGBM-ES	0.3379	0.3560	0.3450	0.3463
Ireland	PGBM	0.3024	0.3640	0.3252	0.3306
	PGBM-E	0.3598	0.4461	0.3672	0.3910
	PGBM-S	0.3144	0.3880	0.3478	0.3500
	PGBM-ES	0.3611	0.4476	0.3820	0.3968

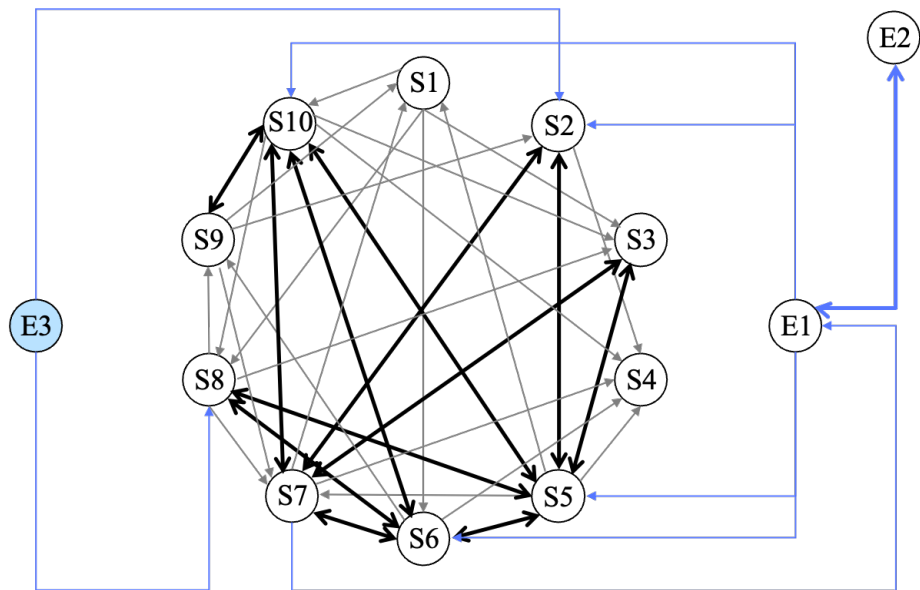


Fig. 8: Granger causality map of economic and social factors. Single and double rows indicate unidirectional and bidirectional Granger causality, respectively. E1 is for GDP, E2 is unemployment rate, E3 is inflation rate. S1 to S10 are social factors as listed in Section 1.




Fig. 9: Granger causality map of economic-social factors and electricity demand in the West Midlands. Single and double rows indicate unidirectional and bidirectional Granger causality, respectively.




Fig. 10: Granger causality map of economic-social factors and electricity demand in South Wales. Single and double rows indicate unidirectional and bidirectional Granger causality, respectively.




Fig. 11: Granger causality map of economic-social factors and electricity demand in the South West. Single and double rows indicate unidirectional and bidirectional Granger causality, respectively.

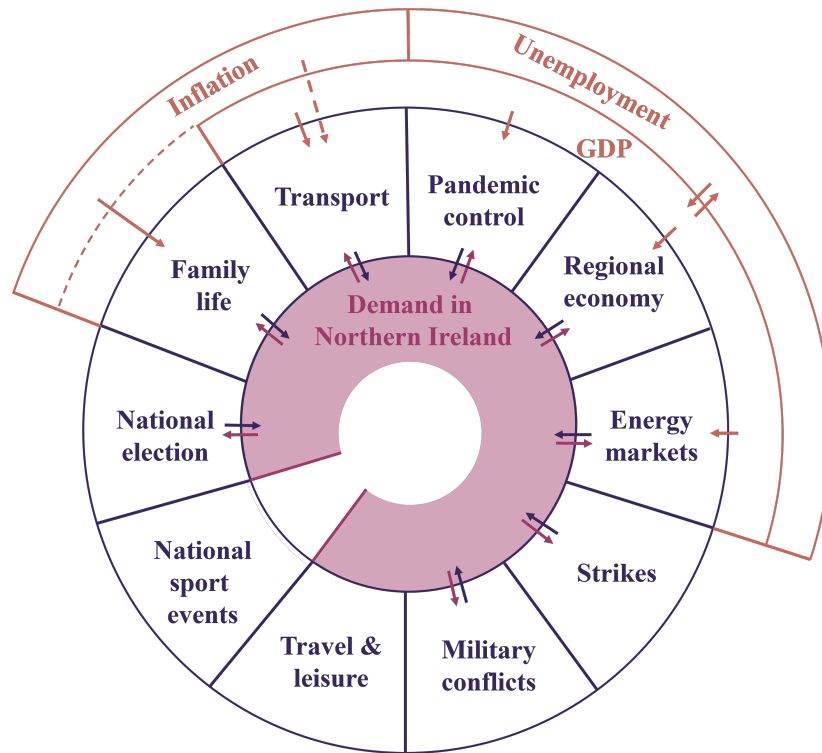


Fig. 12: Granger causality map of economic-social factors and electricity demand in Northern Ireland. Single and double rows indicate unidirectional and bidirectional Granger causality, respectively.




Fig. 13: Granger causality map of economic-social factors and electricity demand in Ireland. Single and double rows indicate unidirectional and bidirectional Granger causality, respectively.

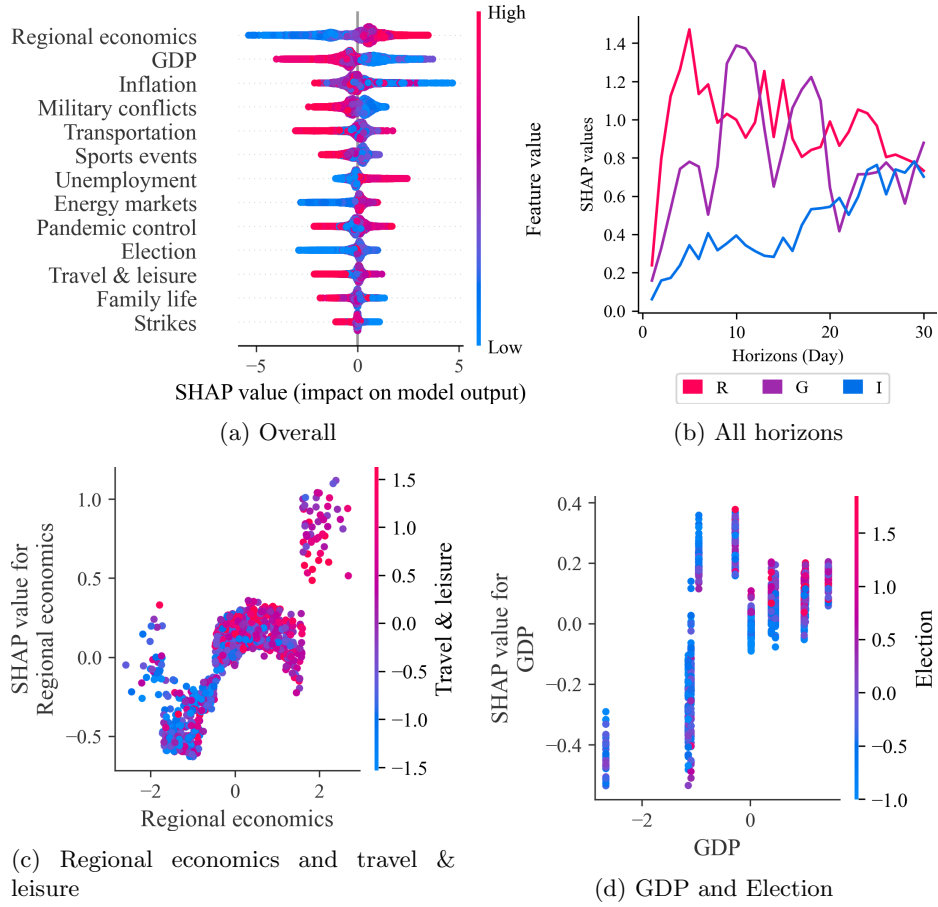


Fig. 14: Feature importance analysis plots by SHAP values in the West Midlands. **(a)** and **(b)** are SHAP values overall and on all horizons. In **(b)**, R is for regional economics, G is for GDP, I is for inflation. **(c)** is the dependency plot of regional economics and travel & leisure. **(d)** is the dependency plot of GDP and election.




Fig. 15: Feature importance analysis plots by SHAP values in South Wales. **(a)** and **(b)** are SHAP values overall and on all horizons. In **(b)**, U is for unemployment, G is for GDP, R is for regional economics, **(c)** is the dependency plot of unemployment and inflation, and **(d)** is the dependency plot of GDP and inflation.

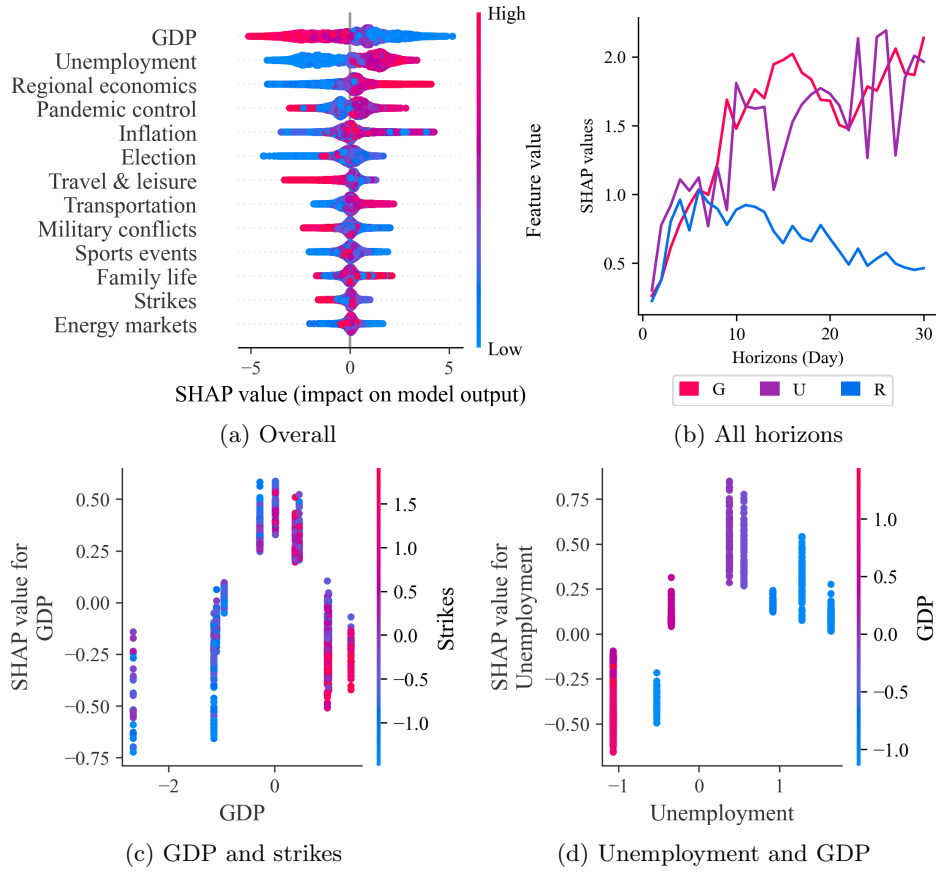


Fig. 16: Feature importance analysis plots by SHAP values in the South West. **(a)** and **(b)** are SHAP values overall and on all horizons. In **(b)**, G is for GDP, U is for unemployment, R is for regional economics, **(c)** is the dependency plot of GDP and strikes, and **(d)** is the dependency plot of unemployment and GDP.




Fig. 17: Feature importance analysis plots by SHAP values in Northern Ireland. **(a)** and **(b)** are SHAP values overall and on all horizons. In **(b)**, G is for GDP, R is for regional economics, U is for unemployment, **(c)** is the dependency plot of GDP and inflation, and **(d)** is the dependency plot of regional economics and sport events.




Fig. 18: Feature importance analysis plots by SHAP values in Ireland. **(a)** and **(b)** are SHAP values overall and on all horizons. In **(b)**, I is for inflation, G is for GDP, E is for energy markets, **(c)** is the dependency plot of inflation and GDP, and **(d)** is the dependency plot of GDP and unemployment.