

Frontier-AI Saliency and Electricity Generation in China: Open-Data Evidence for Sustainable Energy Monitoring

Onur Çelebi

Ankara University

Akın Usupbeyli

ausupbeyli@politics.ankara.edu.tr

Ankara University

Research Article

Keywords: artificial intelligence, electricity generation, sustainable energy monitoring, China, online attention, data centres

Posted Date: September 23rd, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-9873951/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License. [Read Full License](#)

Additional Declarations: No competing interests reported.

Abstract

This paper examines whether frontier-artificial-intelligence (AI) attention provides early signals of national electricity production. Taking China as a case study, it adds official monthly electricity-generation data to industrial activity, weather controls, and three DeepSeek-related public attention proxies: Wikipedia pageviews, Google Trends intensity and GDELT news volume. The sample period is March 2024-December 2025 and is considered a short, gap-aware monthly sample because January-February 2025 current-month electricity data are not available. It adopts a simple UECM-style lag-1 dynamic model for aggregate electricity output and ARDL-style lag-1 model for the share of electricity produced by thermal power stations. A small positive link between frontier-AI salience and total electricity output is found across attention measures. Although small in elasticity terms, this link is non-negligible in absolute volume given the scale of China's electricity system. Bootstrap inference on the electricity response residuals confirms the Wikipedia and GDELT specifications but is weaker for Google Trends. A placebo based on Black Myth: Wukong pageviews makes it less likely that the signal is simply a generic online-attention effect. In contrast, the evidence for a thermal-share effect is weak after adjusting for heating and cooling degree days. Thus, this article does not quantify the causal electricity demand of AI. Rather, it exemplifies how open data can help build an early sustainable-energy monitoring system, and the difficulty of finding national monthly data for assessing generation-mix effects related to SDG 7 and SDG 13.

1 Introduction

Artificial intelligence (AI) is increasingly discussed as a new source of electricity demand through data-centre expansion, model training, and rapidly scaling inference workloads. The issue has moved from a narrow technical concern to a broader sustainability and development question: how should energy systems plan for compute-intensive digital growth while maintaining affordability, reliability, and climate commitments? The International Energy Agency (2025) notes that data centres are becoming a more visible part of electricity systems. The same report also treats uncertainty as substantial, since AI deployment, efficiency gains and grid expansion do not proceed at the same pace.

Measurement is the main difficulty here. Official monthly electricity statistics are not usually built to separate AI-related load, data-centre demand or experimentation around frontier models. Direct measures of compute activity are also rarely available at the national monthly level. For sustainable-development planning, this gap matters because energy planners and climate-policy institutions need early indicators before dedicated categories capture emerging electricity pressures. In the present study, this measurement problem is examined through the DeepSeek attention episode, which provides a recent and policy-relevant case of AI-related public attention in China. The paper therefore uses DeepSeek-related Wikipedia, Google Trends, and GDELT indicators as public salience proxies, while interpreting them strictly as attention measures rather than direct measures of compute activity or electricity use. This article asks whether the broader measurement gap can be partially addressed using open data.

This framing connects the article to sustainable development research by treating compute-intensive AI growth as an energy-monitoring challenge rather than as a direct accounting exercise. If AI-related and data-centre loads become increasingly visible in electricity systems, policymakers will need low-cost indicators that can flag emerging demand-side pressures before they are separately recorded in official statistics. These open attention proxies are not direct measures of AI electricity use. With that caveat in mind, the proxies may still serve as early public signals for sustainability-oriented energy planning. Through them, the paper asks whether frontier-AI

salience coincides with movements in national electricity generation. The point, then, is monitoring rather than direct measurement.

China offers a demanding empirical setting. It has the world's largest electricity system, substantial industrial demand and an unusually visible role in debates over frontier AI models and platform competition. Policy context also matters here. Under the dual-carbon framework, China is pursuing carbon peaking and carbon neutrality (State Council of the People's Republic of China 2021), while the Eastern Data, Western Computing initiative seeks to redirect computing resources from advanced eastern regions to less developed, resource-rich western regions (NDRC 2022). Taken together, these features mean that China can be treated as a useful case for examining digital infrastructure, electricity demand, and sustainable energy planning.

The article does not treat online attention as a mechanical cause of electricity demand, nor use it as a direct consumption measure. Instead, the paper asks a more modest question. Is attention to a frontier AI model observably associated with monthly electricity-generation dynamics in China? To address this question, the empirical design brings together National Bureau of Statistics electricity data, industrial value-added growth, weather controls, and three salience proxies drawn from Wikipedia, Google Trends, and GDELT. The central contribution is a transparent monitoring framework rather than a causal estimate of AI electricity consumption.

Given the short sample, the empirical design is kept deliberately parsimonious. The observed sample contains 20 monthly rows, and the gap-aware dynamic equations use 18 usable months after lag construction. January and February 2025 require special treatment. Because current-period electricity data are missing for those months, the gap is represented explicitly instead of being silently bridged. In such a short sample, the estimates should be read as exploratory. The models are intended to generate interpretable early signals, not definitive structural parameters.

The findings should be read in the same cautious spirit. In the electricity-level equations, frontier-AI salience is associated with small positive implied dynamic responses in aggregate electricity output, although the evidence is not equally strong across proxies. Bootstrap support is clearer for the Wikipedia and GDELT specifications. The Google Trends specification remains directionally consistent but less decisive. The generation-mix results are weaker. Once heating and cooling degree days are retained, the thermal-share specifications provide little support for a stable salience-related response. This asymmetry is therefore important. Aggregate electricity data may contain a faint AI-salience signal before national monthly statistics can sustain stronger claims about carbon intensity or marginal dispatch.

2 Contribution and analytical propositions

The article makes its main contribution at the level of measurement. It shows how a question about frontier-AI salience can be examined with a transparent monthly workflow, even when direct data on compute activity or data-centre electricity use are not available. The empirical material is deliberately public like official Chinese electricity statistics, online-attention indicators, news-volume data, industrial activity and weather controls. This makes the framework relatively easy to audit and, as additional monthly observations become available, to update.

A second contribution concerns the outcome structure. Rather than treating electricity demand as a single sustainability issue, the analysis separates aggregate electricity output from the thermal generation share. That distinction matters. A positive association at the level of total generation need not imply a fossil-intensive generation response, especially in a system where weather, hydro conditions, renewable output, and dispatch

constraints may all affect the generation mix. The short sample cannot resolve all these mechanisms, but it can clarify which claims national monthly data are able to support at this stage.

The methodological choices follow from the same constraint. With a short monthly sample and a visible calendar discontinuity, the empirical design is kept deliberately parsimonious. The lag structure is fixed at one month, HDD18 and CDD18 are retained as core controls, HAC inference is supplemented with a residual-based bootstrap. In addition, a non-AI digital-attention placebo is used to probe whether the estimated pattern simply reflects generic online salience. Taken together, these choices define the contribution as an open-data monitoring design under tight data constraints, not as a claim of causal identification.

The empirical analysis is organized around four analytical propositions rather than causal hypotheses:

- AP1. Frontier-AI salience may be positively associated with aggregate electricity generation because public attention can co-move with experimentation, early adoption, and infrastructure-related activity.
- AP2. If the pattern reflects a broader salience signal rather than platform-specific noise, the electricity-level association should have a broadly similar sign across alternative attention proxies
- AP3. A positive electricity-level association does not necessarily imply a thermal-generation response, since the generation mix also depends on weather, hydro availability, dispatch constraints and system conditions.
- AP4. If the baseline result is specific to frontier-AI salience rather than generic digital attention, a non-AI attention spike should not reproduce the same positive electricity-level pattern.

3 Related literature and analytical motivation

3.1 AI, data centres, and electricity demand

Studies on data-centre energy use and AI workloads increasingly examine the energy implications of expanding digital infrastructure. Masanet et al. (2020) recalibrated global data-centre energy-use estimates and found that efficiency gains had historically offset demand growth. They also cautioned that this balance may not persist as workloads intensify. De Vries (2023) similarly highlighted the possibility that AI-specific electricity consumption could grow substantially if deployment scales rapidly.

The relationship between efficiency and induced demand has a broader economic background. Jevons (1865) first articulated the paradox that improvements in resource efficiency may expand rather than reduce total consumption by lowering effective costs. In the AI context, Patterson et al. (2022) argued that the carbon footprint of machine-learning training may plateau as hardware and algorithmic efficiency improve. The same study also noted, however, that inference workloads scale with adoption and can create a countervailing demand channel.

Strubell et al. (2019) provided one of the earliest empirical estimates of the energy costs of large natural language processing models. Subsequent reviews and institutional reports leave the energy-AI nexus difficult to quantify. They emphasise that training, inference, data-centre construction, and complementary infrastructure move on different timescales (International Energy Agency 2024, 2025; Kamiya and Coroamă 2025). Taken together, these studies suggest that AI-related electricity demand is plausible and policy-relevant, while direct national measurement remains difficult.

3.2 Online attention as an economic and behavioural signal

A second strand of research uses digital traces such as search queries, encyclopaedia pageviews, and news volume as high-frequency indicators of economic activity or public attention. Choi and Varian (2012) demonstrated that Google Trends data can improve short-term forecasts of automobile sales, unemployment claims, and travel demand. Ginsberg et al. (2009) showed that search query volume can track influenza epidemics faster than conventional surveillance systems.

In financial markets, Moat et al. (2013) found that changes in Wikipedia pageviews for financially relevant articles preceded stock-market movements, while Preis et al. (2013) showed that Google Trends could be used to construct trading signals. This article follows the same measurement logic but applies it to electricity monitoring. Wikipedia pageviews, Google Trends intensity, and GDELT news volume are not interpreted as compute use. They are treated as salience proxies that may move with discussion, experimentation, and early adoption around a frontier AI model.

3.3 Electricity generation mix and system flexibility

A third literature asks how electricity systems respond to changes in demand composition and whether incremental load is met through thermal generation or low-carbon sources. Hawkes (2010) showed that marginal emission factors vary substantially across regions and time, implying that the environmental impact of new electricity demand depends on when and where it occurs. York and Bell (2019) argued that renewable energy additions do not always displace fossil generation one-for-one, because demand growth can absorb new clean supply.

Flexibility is central to this problem. Lund et al. (2015) reviewed flexibility measures that can support high shares of variable renewable electricity, including grid expansion, storage, demand-side management, and sector coupling. For the present paper, this literature helps justify the distinction between aggregate electricity output and the thermal generation share. A positive association at the electricity level would suggest demand co-movement with AI salience. It would not, by itself, establish a generation-mix effect. More direct evidence of such an effect would come from a thermal-share response. As shown below, the evidence is clearer for the electricity-level association than for the thermal-share channel.

4 Data, variables, and sample construction

4.1 Data sources and transformations

The analysis is based on the authors' monthly working dataset compiled from public sources. It covers March 2024 to December 2025 and contains 20 observed monthly rows. The file includes total electricity output and source components (thermal, hydro, wind, nuclear, and solar), industrial value-added growth, DeepSeek-oriented Google Trends and Wikipedia attention measures, a GDELT monthly news-volume proxy, and monthly weather controls based on heating and cooling degree days.

Electricity variables come from official Chinese monthly generation statistics. The main level outcome is `elec_total`, measured in 100 million kWh. The generation-mix outcome is `thermal_share`, defined as thermal generation divided by total electricity output. A clean share is available descriptively. However the main generation-mix specification focuses on `thermal_share` to avoid spreading the short sample across too many dependent variables.

Attention enters the dataset through three measures. The baseline proxy is Wikipedia pageviews, transformed as $\ln(1 + \text{wiki_total})$ to retain zero months and reduce skewness. Google Trends enters as $\ln(\text{gtrends_deepseek})$.

GDELT enters as $\ln(1 + \text{gdelt_vol_raw})$. Information-seeking, search intensity, and news coverage are therefore captured through different channels. The three measures should be read as tools for triangulation, not as direct measures of compute load.

Weather control remained in the main specifications. The file includes `china_hdd18` and `china_cdd18`, monthly heating and cooling degree days built on an 18°C base. Monthly electricity output and dispatch can be strongly weather-sensitive, so retaining both controls is part of the main design. In the revised design, weather variables are treated as core confounders rather than optional add-ons.

Table 1
Selected summary statistics (observed monthly sample, March 2024-December 2025)

Variable	N	Mean	Std. Dev.	Min	Max
elec_total	20	7,997.09	732.372	6,900.8	9,362.7
thermal_share	20	0.656	0.025	0.62	0.706
iva_growth	20	5.635	0.824	4.5	7.7
gtrends_deepseek	20	6.165	6.138	0.5	16.6
wiki_total	20	76,345.15	86,178.38	0	278,766
gdelt_vol_raw	20	590.7	813.167	1	2,957
wukong_views	20	176,281.45	341,654.773	20,621	1,536,320
china_hdd18	20	130.263	162.699	0	518.86
china_cdd18	20	129.901	115.329	0.92	323.15
ln_elec	20	8.983	0.09	8.839	9.144
ln_wiki	20	6.542	5.825	0	12.538
<i>Notes: Electricity output is measured in 100 million kWh. Wikipedia and GDELT are raw monthly totals before log transformation. HDD18 and CDD18 are retained weather controls. The observed sample contains 20 monthly rows, but the gap-aware dynamic equations use 18 usable months after lag construction.</i>					

Table 2
Core variables and transformations used in the revised specification

Variable	Definition	Units	Role
elec_total	Current-period monthly electricity output	100 million kWh	National total generation level
thermal_share	Thermal output / total output	Share	Dependent variable for generation-mix equation
iva_growth	Industrial value-added growth	Percent	Macro production control retained in all equations
ln_wiki	$\ln(1 + \text{wiki_total})$	Log points	Baseline online-attention proxy
ln_trends	$\ln(\text{Google Trends index})$	Log points	Alternative search-attention proxy
ln_gdelt	$\ln(1 + \text{GDELT volume})$	Log points	Alternative news-attention proxy
ln_wukong	$\ln(1 + \text{Wukong views})$	Log points	Placebo attention proxy
china_hdd18	Monthly heating degree days (18C base)	Degree days	Retained weather control
china_cdd18	Monthly cooling degree days (18C base)	Degree days	Retained weather control

Notes: Variable definitions correspond to the authors' working dataset and the transformations implemented in the Python analysis script. Attention measures are salience proxies and are not interpreted as direct measures of compute load.

4.2 Sample restrictions, missing months, and descriptive patterns

The short sample is informative but not frictionless. The observed working file contains March-December 2024 and March-December 2025. For dynamic estimation, the data are reindexed to the full monthly calendar. The missing January-February 2025 interval is therefore represented explicitly rather than silently skipped. This choice prevents the lag structure from treating December 2024 as if it were immediately followed by March 2025.

Once the gap is represented explicitly, the lag-1 electricity equation uses April-December 2024 and April-December 2025. March 2025 is not included in the dynamic estimation sample because its lag continuity is broken by the January-February gap. The same logic applies to the thermal-share equation. What remains is a small but internally consistent gap-aware estimation sample of 18 months.

In Fig. 1, log electricity and the three attention proxies are standardized around their sample means. The attention series remain low through much of 2024 and then rise sharply around early 2025. Electricity output, however, does not move mechanically with those spikes month by month. This is why the article avoids causal language and relies on a dynamic association model rather than simple contemporaneous comparisons.

Notes: The shaded band marks the January-February 2025 interval that is absent from the observed current-period electricity sequence. Standardisation is applied only for visual comparability; the econometric models use the original transformed variables.

Figure 2 compares the thermal share with the baseline Wikipedia attention series and the retained weather controls. The visible pattern is the seasonal structure of weather pressure. This is especially clear when HDD18 and CDD18 alternate in dominance. For the generation-mix equation, that pattern provides the main reason for retaining both weather controls.

Notes: HDD18 and CDD18 are plotted in standardised form for visual comparability. The thermal share and attention series are also standardised for the figure only.

Figure 3 focuses only on the weather series. Average temperature, HDD18, and CDD18 follow the expected seasonal pattern. This reinforcing the substantive reason for keeping them in the baseline equations. In a short sample, these controls help absorb variation that might otherwise be mistakenly attributed to technology salience.

Notes: Unlike the electricity variables, the weather controls are observed continuously across the paper window. Their seasonal structure supports retaining them in the main specifications.

5 Empirical strategy

The empirical strategy is designed as a gap-aware dynamic monitoring exercise. It is not a full causal-identification design and it is not used to deliver a formal cointegration verdict. The sample is too short for reliable unit-root pre-testing or asymptotic bounds-test inference. Instead, the paper uses parsimonious lag-1 dynamic specifications to separate lagged-level associations from short-run changes, and it treats bootstrap inference as the main small-sample safeguard.

5.1 Electricity-output specification

The main equation is a UECM-style lag-1 dynamic model. The dependent variable is the monthly change in log electricity output. The equation retains the lagged level of electricity, the lagged level of one attention proxy, lagged industrial value-added growth, first differences of attention and industrial activity, and contemporaneous weather controls in levels:

$$\Delta \ln(E_t) = \alpha + \phi \ln(E_{t-1}) + \theta A_{t-1} + \gamma IVA_{t-1} + \delta \Delta A_t + \kappa \Delta IVA_t + \lambda HDD_t + \mu CDD_t + \epsilon_t \quad (1)$$

$$\text{Implied dynamic association} = -\theta / \phi \quad (2)$$

Here E_t denotes total electricity output and A_t denotes one attention proxy at a time. The baseline proxy is $\ln_wiki$, but the same lag-1 structure is re-estimated with $\ln_trends$ and $\ln_gdelt$. The ratio $-\theta/\phi$ is reported as an implied dynamic association under the maintained specification, not as a confirmed structural elasticity or a formal cointegration coefficient.

5.2 Thermal-share specification

For the generation-mix outcome, the specification is an ARDL-style lag-1 model written in levels. The dependent variable is `thermal_share`. On the right-hand side, the model includes the lagged thermal share, current and lagged attention, current and lagged industrial value-added growth, and the contemporaneous HDD18 and CDD18 controls:

$$TS_t = \alpha + \rho TS_{t-1} + \beta_0 A_t + \beta_1 A_{t-1} + \psi_0 IVA_t + \psi_1 IVA_{t-1} + \lambda HDD_t + \mu CDD_t + \epsilon_t \quad (3)$$

$$\text{Implied dynamic multiplier} = (\beta_0 + \beta_1) / (1 - \rho) \quad (4)$$

The thermal-share multiplier is interpreted cautiously. In a national monthly sample of this length, dispatch conditions and other system-level influences cannot be observed fully enough to treat the thermal-share response as a structural generation-mix effect.

5.3 Small-sample inference and lag choice

The one-month lag structure is a deliberate restriction. In a short monthly sample with an explicit calendar gap, more heavily parameterized specifications would consume degrees of freedom quickly and could magnify instability without improving interpretation. With only 18 usable months in the dynamic equations and eight reported coefficients in the baseline tables, the effective degrees of freedom are necessarily limited. This limitation shapes the interpretation. The models are used for narrow monitoring, not for broad causal or structural claims.

The models are estimated by OLS and reported with Newey-West HAC standard errors using maxlag = 1 (Newey and West 1987). Since conventional asymptotic inference is unreliable in this setting, the paper supplements HAC standard errors with a 10,000-replication residual-based fixed-regressor bootstrap. The procedure resamples OLS residuals, keeps the regressor matrix fixed, re-estimates each model, and recomputes the implied dynamic association or multiplier. Percentile 95% confidence intervals and two-sided bootstrap p-values are reported alongside the HAC results.

Formal unit-root and bounds tests are not reported. At $N = 18$, ADF and Phillips-Perron tests would have very low power while bounds-test critical values are not reliable for the kind of inference required here. The paper therefore avoids a formal cointegration claim. In the electricity equation, the negative lagged-electricity coefficient is used more narrowly: it supports the use of the lag-1 dynamic specification in this setting. The main evidence rests on bootstrap-supported association estimates.

6 Empirical results

6.1 Baseline electricity equation

Table 3 reports the gap-aware electricity-output equation by using Wikipedia attention as the baseline proxy. The negative coefficient on lagged log electricity is estimated precisely, which is consistent with using a lag-1 dynamic specification in this setting. The lagged Wikipedia term is positive, while HDD18 and CDD18 enter with positive signs, in line with higher electricity pressure in colder and hotter months.

The coefficient on lagged log electricity is -0.6609, while the lagged Wikipedia term is 0.0023. Using $-\theta/\phi$ gives an implied association of 0.0035. This magnitude is small. It should therefore be read as evidence of a modest salience-related co-movement in aggregate electricity output, not as a direct estimate of AI electricity demand. Nevertheless, at China's monthly generation scale—averaging roughly 8,000 (100 million kWh), i.e. about 800 TWh per month—even an association of this size maps onto an economically non-trivial absolute volume of electricity, which is precisely why aggregate national output can register faint salience-related signals.

Table 3
Gap-aware lag-1 dynamic model for log electricity output (baseline attention proxy: Wikipedia)

Variable	Coefficient	HAC SE	p-value
const	5.7998	0.6145	< 0.001
L1_In_elec	-0.6609	0.0644	< 0.001
L1_In_wiki	0.0023	0.0009	0.0325
L1_iva_growth	-0.0134	0.013	0.3274
D_In_wiki	0.0032	0.003	0.3215
D_iva_growth	-0.0159	0.0057	0.0186
china_hdd18	0.0005	<0.001	< 0.001
china_cdd18	0.001	< 0.001	< 0.001

Notes: Dependent variable is Delta ln(electricity output). HAC standard errors use a Newey-West correction with maxlag = 1. HDD18 and CDD18 are retained in levels. The estimation sample contains 18 usable months after explicit representation of the January-February 2025 gap. Numerical values are reproduced from the original empirical output.

6.2 Generation-mix evidence

The thermal-share results are less supportive of a salience-related generation-mix response. Table 4 reports the lag-1 thermal-share model using Wikipedia attention. The lagged thermal share is positive, indicating persistence in the generation mix, but the current and lagged attention terms are imprecise. The implied multiplier is close to zero once HDD18 and CDD18 are retained. The electricity-output association should therefore not be read as evidence of a thermal-dispatch or carbon-intensity effect.

Weather-related variation is more visible in this equation. CDD18 is positive with $p = 0.0441$, while HDD18 is positive but weaker. Seasonal demand pressure therefore helps explain thermal-share movement, whereas attention adds little stable explanatory power in this small sample.

Table 4
Gap-aware ARDL-style lag-1 model for thermal generation share (baseline attention proxy: Wikipedia)

Variable	Coefficient	HAC SE	p-value
const	0.3667	0.1557	0.0403
L1_thermal_share	0.5155	0.1995	0.0272
ln_wiki	0.0054	0.0038	0.1824
L1_ln_wiki	-0.0056	0.0036	0.1524
iva_growth	-0.0116	0.0046	0.0296
L1_iva_growth	-0.0029	0.0045	0.5413
china_hdd18	0.0001	0.0001	0.0855
china_cdd18	0.0001	0.0001	0.0441
<i>Notes: Dependent variable is the thermal share of total electricity output. HAC standard errors use a Newey-West correction with maxlag = 1. Attention enters in current and lagged form, while HDD18 and CDD18 are retained contemporaneously. Numerical values are reproduced from the original empirical output.</i>			

7 Robustness checks

The robustness checks are restricted to results that are already supported by the original data and empirical output. No additional, unverified tests are introduced here. The aim is instead to make the existing robustness evidence more explicit.

7.1 Proxy triangulation and bootstrap inference

The electricity association is re-estimated with three AI-related attention proxies. In each case, the implied association is positive, but the strength of the bootstrap evidence differs across measures. Bootstrap inference supports the Wikipedia specification (bootstrap $p = 0.0300$, 95% CI [0.0001, 0.0064]) and the GDELT specification (bootstrap $p = 0.0233$, 95% CI [0.0007, 0.0128]). Google Trends points in the same direction, although its bootstrap confidence interval includes zero and its bootstrap p -value is 0.0889. For that reason, the Google Trends result is best read as supportive evidence, but not as decisive confirmation.

Table 5
Implied dynamic electricity associations across attention proxies

Attention proxy	N	Implied dynamic association	HAC SE	p-value	Boot 95% CI Lower	Boot 95% CI Upper	Boot p-value	AIC
Wikipedia pageviews (ln_wiki)	18	0.0035	0.0013	0.0064	0.0001	0.0064	0.03	-80.9142
Google Trends (ln_trends)	18	0.0101	0.005	0.0414	-0.003	0.0205	0.0889	-80.0745
GDELT news volume (ln_gdelt)	18	0.0071	0.0019	0.0001	0.0007	0.0128	0.0233	-80.7602
Black Myth: Wukong (Placebo)	18	-0.0286	0.0087	0.001	-0.0463	-0.0133	0.0013	-89.1733
Notes: Each row uses the same lag-1 electricity equation with retained IVA, HDD18, and CDD18 controls. The reported association is computed from the estimated lagged-level coefficients and should be interpreted as an elasticity-like association under the maintained specification, not as a structural demand elasticity. N = 18 in each specification.								

Table 6
Bootstrap inference for implied dynamic coefficients (10,000 replications)

Attention proxy	Model	Point estimate	Bootstrap 95% CI	Bootstrap p-value
Wikipedia pageviews	UECM – Electricity	0.0035	[0.0001, 0.0064]	0.0300
Google Trends	UECM – Electricity	0.0101	[-0.0030, 0.0205]	0.0889
GDELT news volume	UECM – Electricity	0.0071	[0.0007, 0.0128]	0.0233
Wukong (Placebo)	UECM – Electricity	-0.0286	[-0.0463, -0.0133]	0.0013
Wikipedia pageviews	ARDL – Thermal share	-0.0003	[-0.0032, 0.0047]	0.8265
Google Trends	ARDL – Thermal share	-0.0084	[-0.0609, 0.0328]	0.3022
GDELT news volume	ARDL – Thermal share	-0.0005	[-0.0087, 0.0199]	0.9026
Wukong (Placebo)	ARDL – Thermal share	0.0122	[-0.0138, 0.0404]	0.2045
Notes: Residual-based fixed-regressor bootstrap with B = 10,000 replications. Percentile confidence intervals are reported at the 95% level. Bootstrap p-values are two-sided and computed under the null by centering. Values are reproduced from the original empirical output.				

7.2 Placebo falsification

A remaining concern is that the positive electricity-attention association may reflect a generic response to any large digital salience shock rather than a pattern specific frontier-AI salience. To examine this possibility, the original analysis replaces the AI attention proxies with Wikipedia pageviews for the Chinese video game Black

Myth: Wukong. This case is useful as a placebo because the game produced a large and temporally concentrated online-attention spike after its August 2024 release, before the DeepSeek-related attention surge of early 2025.

The placebo electricity association is negative (-0.0286) and its bootstrap confidence interval lies below zero. This result makes it less likely that the baseline estimates simply capture any large online-attention spike. The placebo does not surely establish AI-specific causality, since the placebo differs in timing and context. It does, however, support a narrower interpretation: the positive baseline pattern does not appear to be a generic attention artefact.

7.3 Robustness of the null thermal-share result

The thermal-share result is also checked across attention proxies. Appendix Table A1 shows that the implied thermal-share multipliers are small and statistically weak across AI-related proxies once HDD18 and CDD18 are retained. The corresponding bootstrap p-values are 0.8265, 0.3022, and 0.9026. This leaves the main asymmetry intact. In this sample, the salience-related association is more visible in aggregate electricity output than in the generation mix.

8 Discussion

The results point to a cautious monitoring interpretation. Several public salience proxies move positively with aggregate electricity output in the gap-aware dynamic specification. The evidence, however, does not support a definitive causal AI-demand effect. The observed association is small and should be read as an early signal in national electricity data.

The level-vs-mix asymmetry is analytically useful. Aggregate electricity output may be broad enough to register diffuse demand-side pressures associated with experimentation, platform attention, or early adoption. The thermal share is shaped by dispatch conditions, seasonal weather, hydro availability, renewable output and other system constraints that are not fully captured in the national monthly sample. On this basis, weaker generation-mix inference is not surprising.

The industrial value-added control should also be interpreted cautiously. In the electricity equation the differenced IVA term is negative and significant, while the lagged level term is not. In the thermal-share equation the current IVA term is negative and significant. These signs may reflect short-sample timing, growth-rate measurement, or multicollinearity with other monthly dynamics. They should not be read as structural macroeconomic relationships.

9 Limitations and future research

Four limitations are especially important for interpreting the results.

First, the effective estimation sample is small. The observed file contains 20 monthly rows, and the gap-aware dynamic equations use 18 usable months. This is enough for an early monitoring exercise but not enough for stable structural estimation. The bootstrap is therefore used to guard against over-reliance on conventional asymptotic inference, and the reported coefficients are interpreted as associations rather than causal effects.

Second, the January-February 2025 discontinuity sits close to the period that makes the topic especially interesting. The paper handles this by explicitly representing the gap and preventing broken lag continuity. That is preferable to interpolation or silently chaining December 2024 to March 2025, but it also means the early-2025 shock period cannot be observed as fully as one would prefer.

Third, the attention variables are salience proxies rather than measures of compute intensity. Wikipedia, Google Trends, and GDELT capture information seeking, search intensity, and news volume. They may co-move with experimentation and adoption, but they do not measure data-centre electricity use. Future work should add China-native attention measures such as Baidu Index or Chinese-language news counts if those data can be verified and replicated.

Fourth, causal direction remains unresolved. Attention and electricity may both respond to common shocks, and the present design lacks a credible monthly instrument for frontier-AI salience. Future research should use longer samples, regional grid data, direct data-centre measurements where available, event-study designs around AI releases, and provincial analysis linked to data-centre siting and the Eastern Data, Western Computing initiative.

10 Policy and sustainability implications

The findings carry three policy implications relevant to sustainable energy planning. First, AI-related electricity demand should be treated as a monitoring problem even before official statistics create dedicated categories for compute-intensive digital services. This is directly relevant to SDG 7, which includes reliable and sustainable energy access, and SDG 13, which concerns climate action and the integration of climate measures into policy planning (United Nations 2015). The framework shows that low-cost public indicators can be used to track emerging digital-infrastructure pressures without claiming to replace direct measurement.

Second, the divergence between electricity-output and thermal-share results cautions against translating AI-related electricity signals directly into carbon-intensity claims. A positive aggregate electricity association does not necessarily imply a fossil-intensive marginal response at national monthly frequency. This distinction matters for China's dual-carbon objectives, including carbon peaking before 2030 and carbon neutrality before 2060 (State Council of the People's Republic of China 2021). Policy interpretation should therefore separate demand monitoring from generation-mix attribution.

Third, the spatial dimension of Chinese computing infrastructure matters. The Eastern Data, Western Computing initiative was designed to rebalance computing resources across regions and to make use of resource-rich western areas (NDRC 2022). A national model such as the one presented here can provide a first monitoring layer, but it should be complemented by provincial electricity, data-centre, and renewable-generation indicators. The long-run policy value of the framework is therefore statistical: it highlights why official energy systems may need more explicit categories for data centres and other flexible digital loads.

11 Conclusion

This paper re-examines the link between frontier-AI salience and electricity generation in China with a "lean" design: lag-1 dynamics, retained HDD18 and CDD18 controls, explicit accounting for missing months, and publicly-available monthly data. The findings yield small positive implied dynamic relations between AI salience online and aggregate electricity generation, with the strongest and most consistent bootstrap results for Wikipedia and GDELT salience proxies.

The design simplification also alters the generation-mix results. With the weather variables retained, the evidence for the steady thermal-share response is weak. This is not a negative result. It defines the limits of what can be said about national monthly data: total electricity responses might be more apparent than generation-mix responses.

The main point thus is small but significant. Frontier-AI salience is positively related to total electricity generation in this brief Chinese monthly sample, but the current evidence does not warrant a definitive causal AI-demand effect or a steady thermal-generation response. The article is valuable for the monitoring framework as well as the coefficient estimates. As the sample grows and provincial data emerges, this framework can be expanded to better identification, province-level analysis, and direct use of data-centre electricity data.

Declarations

12 Conflict of interest

On behalf of all authors, the corresponding author states that there is no conflict of interest.

13. Funding Declaration

No funding was received for conducting this study or for the preparation of this manuscript.

Author Contribution

Both authors contributed to the conception and design of the study, data analysis, and interpretation. Onur Çelebi prepared the original draft and Akın Usupbeyli reviewed and edited the manuscript with Onur Çelebi. Both authors read and approved the final manuscript.

Data Availability

The processed monthly dataset used for the empirical analysis was compiled from the public sources identified in the manuscript, including official Chinese electricity statistics, public online-attention data, news-volume indicators, and weather data. The processed dataset is available from the corresponding author upon reasonable request, subject to the terms of use of the original data providers.

References

1. Choi, H., & Varian, H. (2012). Predicting the present with Google Trends. *Economic Record*, 88(s1), 2–9. <https://doi.org/10.1111/j.1475-4932.2012.00809.x>
2. De Vries, A. (2023). The growing energy footprint of artificial intelligence. *Joule*, 7(10), 2191–2194. <https://doi.org/10.1016/j.joule.2023.09.004>
3. Ginsberg, J., Mohebbi, M. H., Patel, R. S., Brammer, L., Smolinski, M. S., & Brilliant, L. (2009). Detecting influenza epidemics using search engine query data. *Nature*, 457(7232), 1012–1014. <https://doi.org/10.1038/nature07634>
4. Google (n.d.-a) FAQ about Google Trends data (Ed.). <https://support.google.com/trends/answer/4365533>. Accessed 10 Feb 2026.
5. Google (n.d.-b) Export, embed, and cite Trends data. <https://support.google.com/trends/answer/4365538>. Accessed 10 Feb 2026.

6. International Energy Agency. (2024). *Electricity 2024: Analysis and forecast to 2026*. IEA. <https://www.iea.org/reports/electricity-2024>
7. International Energy Agency. (2025). *Energy and AI*. IEA. <https://www.iea.org/reports/energy-and-ai>
8. Jevons, W. S. (1865). *The coal question: An inquiry concerning the progress of the nation, and the probable exhaustion of our coal mines*. Macmillan.
9. Hawkes, A. D. (2010). Estimating marginal CO₂ emissions rates for national electricity systems. *Energy Policy*, 38(10), 5977–5987. <https://doi.org/10.1016/j.enpol.2010.05.053>
10. Kamiya, G., & Coroama, V. C. (2025). Data Centre Energy Use: Critical Review of Models and Results. EDNA - IEA 4E TCP. <https://www.iea-4e.org/edna/publications/data-centre-energy-use-critical-review-of-models-and-results/>
11. Lund, P. D., Lindgren, J., Mikkola, J., & Salpakari, J. (2015). Review of energy system flexibility measures to enable high levels of variable renewable electricity. *Renewable and Sustainable Energy Reviews*, 45, 785–807. <https://doi.org/10.1016/j.rser.2015.01.057>
12. Masanet, E., Shehabi, A., Lei, N., Smith, S., & Koomey, J. (2020). Recalibrating global data center energy-use estimates. *Science*, 367(6481), 984–986. <https://doi.org/10.1126/science.aba3758>
13. Moat, H. S., Curme, C., Avakian, A., Kenett, D. Y., Stanley, H. E., & Preis, T. (2013). Quantifying Wikipedia usage patterns before stock market moves. *Scientific Reports*, 3, 1801. <https://doi.org/10.1038/srep01801>
14. NASA POWER Project (nd.) Prediction Of Worldwide Energy Resources (POWER) API. <https://power.larc.nasa.gov/>. Accessed 13 Feb 2026.
15. National Bureau of Statistics of China (nd.) National Data (StatChina). <https://data.stats.gov.cn/english/>. Accessed 10 Feb 2026.
16. National Development and Reform Commission (2022). China approves mega project for greater computing power, digital future. https://en.ndrc.gov.cn/news/mediar/sources/202202/t20220218_1315947.html. Accessed 26 Apr 2026.
17. Newey, W. K., & West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3), 703–708. <https://doi.org/10.2307/1913610>
18. Patterson, D., Gonzalez, J., Hoelzle, U., Le, Q., Liang, C., Munguia, L. M., Rothchild, D., So, D. R., Texier, M., & Dean, J. (2022). The carbon footprint of machine learning training will plateau, then shrink. *Computer*, 55(7), 18–28. <https://doi.org/10.1109/MC.2022.3148714>
19. Pesaran, M. H., Shin, Y., & Smith, R. J. (2001). Bounds testing approaches to the analysis of level relationships. *Journal of Applied Econometrics*, 16(3), 289–326. <https://doi.org/10.1002/jae.616>
20. Preis, T., Moat, H. S., & Stanley, H. E. (2013). Quantifying trading behavior in financial markets using Google Trends. *Scientific Reports*, 3, 1684. <https://doi.org/10.1038/srep01684>
21. Seabold, S., & Perktold, J. (2010). Statsmodels: econometric and statistical modeling with Python. In: *Proceedings of the 9th Python in Science Conference*, pp 57–61. <https://doi.org/10.25080/Majora-92bf1922-011>
22. State Council of the People's Republic of China (2021). Action Plan for Carbon Dioxide Peaking Before 2030. https://english.www.gov.cn/policies/latestreleases/202110/27/content_WS6178a47ec6d0df57f98e3dfb.html. Accessed 26 Apr 2026.
23. Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp 3645–3650.

<https://doi.org/10.18653/v1/P19-1355>

24. The GDELT Project (nd.) GDELT 2.0/2.1 documentation and APIs. <https://www.gdeltproject.org/>. Accessed 13 Feb 2026.
25. United Nations (2015). Transforming our world: The 2030 Agenda for Sustainable Development. <https://sdgs.un.org/publications/transforming-our-world-2030-agenda-sustainable-development-17981>
26. Wikimedia Foundation (n.d.) Pageviews Analysis. <https://pageviews.wmcloud.org/>. Accessed 10 Feb 2026.
27. York, R., & Bell, S. E. (2019). Energy transitions or additions? Why a transition from fossil fuels requires more than the growth of renewable energy. *Energy Research & Social Science*, 51, 40–43. <https://doi.org/10.1016/j.erss.2019.01.008>

Figures




Figure 1

Electricity output and attention proxies (standardised)

Notes: The shaded band marks the January-February 2025 interval that is absent from the observed current-period electricity sequence. Standardisation is applied only for visual comparability; the econometric models use the original transformed variables.




Figure 2

Thermal share, weather pressure, and attention (standardised)

Notes: HDD18 and CDD18 are plotted in standardised form for visual comparability. The thermal share and attention series are also standardised for the figure only.




Figure 3

Observed monthly weather controls (standardised)

Notes: Unlike the electricity variables, the weather controls are observed continuously across the paper window. Their seasonal structure supports retaining them in the main specifications.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [SUPPORTINGMATERIAL.docx](#)
- [AppendixAAdditionalthermal.docx](#)