

Does AI Really Drive the Grid? A Four-Decade Test of the U.S. Energy Footprint

Younes Nademi^{1,2*}, Majid Ebtia³, Ramin Khochiany^{1,2}, and Seyed Mohammad Hosseini⁴

¹ Department of Economics, Faculty of Humanities, Ayatollah Boroujerdi University, Boroujerd, Iran

² Zagros Data Sciences Research Group, Ayatollah Boroujerdi University, Boroujerd, Iran

³ Gahar Artificial Intelligence Research Group, Ayatollah Boroujerdi University, Boroujerd, Iran

Highlights

- Four decades of U.S. data indicate that artificial intelligence has no significant effect on national energy demand.
- No cointegration between artificial intelligence and energy demand is detected using either the Engle–Granger or Johansen tests.
- Traditional factors explain 90 percent of energy demand, while artificial intelligence contributes no additional predictive power.
- Shapley additive explanations (SHAP) analysis ranks GDP and population highly, whereas artificial intelligence-related features exhibit near-zero impact.

Received: May 20, 2025; revised: August 06, 2025; accepted: August 13, 2025

Abstract

The recent surge in artificial intelligence (AI) activity has raised concerns that large-scale model training, cloud inference, and data-centre expansion could accelerate national energy demand. We assemble a 21-year annual panel for the United States (2004–2024) that integrates multiple AI proxies—technology-stock valuations and a ChatGPT-era dummy—with four aggregate energy series (fossil fuels, nuclear, renewables, and total primary energy). A five-stage empirical protocol, implemented in Python, combines Engle–Granger cointegration testing, higher-order ADF stationarity checks, linear and nonlinear dependence diagnostics (Pearson correlation, dynamic time warping, mutual information), multicollinearity screening using variance inflation factors, and out-of-sample forecasting with linear regression, decision trees, random forests, and support-vector machines augmented by SHAP explainability. Across all tests, we find no evidence that AI developments affect national energy use: AI variables cointegrate only with one another, their short-run correlations with energy vanish once trends are removed, their mutual-information scores remain near zero, and their inclusion never improves predictive accuracy beyond a parsimonious macro model driven by GDP, inflation, and population. SHAP rankings confirm that AI features carry negligible explanatory weight relative to conventional fundamentals. We conclude that, to date, AI’s macro-level energy footprint is statistically undetectable, with any electricity it consumes either too small to register or offset by efficiency gains elsewhere in the economy. Policymakers should therefore continue to anchor long-term energy scenarios to established economic drivers while monitoring localized data-centre hotspots that national aggregates obscure.

Keywords: Artificial intelligence; Energy consumption; Cointegration; Machine-learning forecasting; SHAP explainability

* Corresponding author:

Email: younesnademi@abru.ac.ir

How to cite this article

Nademi, Y., Ebtia, M., Khochiany, R., and Hosseini, S.M., *Does AI Really Drive the Grid? A Four-Decade Test of the U.S. Energy Footprint*, *Petroleum Business Review*, Vol. 9, No. 4, p. 75–101, 2025. DOI: [10.22050/pbr.2025.525046.1394](https://doi.org/10.22050/pbr.2025.525046.1394)

1. Introduction

Artificial intelligence (AI) has rapidly emerged as a pivotal engine of economic growth and innovation in the modern economy. It is widely recognized as a general-purpose technology that transforms the production of goods and services across diverse industries, including manufacturing, energy, healthcare, and finance (Wang et al., 2024). The global AI market has expanded dramatically in recent years, reaching unprecedented scale by the mid-2020s (Wang et al., 2024). Many nations and firms are investing heavily in AI to boost productivity and develop new products. At the same time, deploying advanced AI is computationally intensive and energy-demanding, raising concerns regarding the energy footprint of AI-driven growth (Lee et al., 2025).

The rise of AI-centric companies and applications may have significant implications for national energy consumption. Delivering AI services often requires power-intensive infrastructure, such as large-scale data centers, which consume substantial amounts of electricity (Bogmans et al., 2025). Recent evidence suggests a connection between surging AI adoption and higher energy use: the aggregate energy consumption of leading AI-focused technology firms (the “Magnificent Seven”) increased by approximately 19 percent in 2023, whereas the median energy use of S&P 500 companies remained flat (Burian & Stalla-Bourdillon, 2025). Similarly, in the United States, AI-producing sectors have been expanding nearly three times faster than the broader economy, accompanied by a doubling of their electricity expenditures between 2019 and 2023 (Bogmans et al., 2025). These trends indicate that the growing economic value of AI could exert nontrivial effects on countries’ energy demand and resource requirements.

However, the net effect of AI on energy consumption is complex and potentially bidirectional. On one hand, the proliferation of AI applications and services can increase energy demand. Training and deploying modern AI models require enormous computational power; for example, training a single state-of-the-art model such as ChatGPT is estimated to consume approximately 1.3 GWh of electricity—roughly equivalent to the annual usage of 120 U.S. households (Wang et al., 2024). The expansion of AI-related technologies—from data analytics to autonomous systems—adds new sources of electricity load across data centers, networks, and devices. On the other hand, AI holds substantial potential for improving energy efficiency and productivity across the economy. By optimizing industrial processes, supply chains, transportation, buildings, and power systems, AI can reduce waste and lower the energy consumed per unit of output (Wang et al., 2024). Consequently, AI may simultaneously act as an energy demand amplifier and an efficiency enabler, leaving its ultimate impact on aggregate energy usage as an open empirical question.

Given this dual potential, it is critical to empirically examine the relationship between AI development and energy consumption, particularly in an era of heightened global focus on energy sustainability. Nearly all countries have committed to reducing carbon emissions under agreements such as the Paris Climate Accord and are pursuing energy transitions toward renewables (Wang et al., 2024). If rapid AI proliferation significantly increases energy demand, it could create new challenges for achieving climate targets; conversely, if AI generates substantial efficiency gains, it may become a valuable tool for sustainable development. Despite the importance of this issue, rigorous empirical analyses of AI’s impact on macro-level energy use remain limited (Wang et al., 2024). AI is rarely incorporated into traditional energy-economy frameworks, and little is known about its long-term influence on national energy demand (Wang et al., 2024).

To address this knowledge gap, the present paper empirically investigates the AI–energy nexus using long-term country-level data and advanced analytical techniques. We construct a multi-decade panel spanning 1985–2024 that tracks the emergence of AI-related sectors alongside national energy consumption. We then apply a comprehensive battery of econometric tests and machine learning models to assess the relationship robustly. This approach allows us to determine whether the rise of AI ultimately increases, decreases, or leaves energy consumption unchanged over the long run, providing timely evidence to inform both energy policy and the strategic development of AI in an era of sustainability commitments.

2. Literature review

Recent empirical studies at the macroeconomic or national level reveal a complex relationship between AI diffusion and energy use. Cross-country analyses generally indicate that greater AI development coincides with improvements in energy efficiency and sustainability. For example, using panel data on 67 countries, Wang et al. (2024) report that AI advancement is associated with significantly lower carbon emissions and ecological footprints, while also promoting cleaner energy transitions, as reflected by a higher share of renewables in energy consumption. Similarly, a state-level study in the United States by Fang et al. (2025) finds that AI proliferation positively impacts the energy transition: states with higher AI activity tend to consume more energy from renewable sources—particularly wind, solar, and hydro—while relying less on fossil fuels. In China, several studies document analogous effects at regional scales. Liu et al. (2022) examine China’s industrial sectors and find that AI adoption notably reduces CO₂ emission intensity. At the provincial level, Tao et al. (2023) report that higher AI utilization corresponds to lower energy-related carbon intensity. These findings suggest that, at a broad scale, AI technologies may contribute to more efficient energy use and help decouple economic growth from carbon- and energy-intensive inputs. However, the relationship may not be strictly linear. Zhao et al. (2022) observe a non-linear “U-shaped” pattern in China, in which initial phases of AI adoption produced minimal or even adverse effects on green productivity, but beyond a certain threshold, AI contribution became positive. This indicates that the energy impact of AI may depend on the level of AI maturity and complementary factors within an economy.

In parallel, researchers and institutions caution that rapid AI deployment itself is driving energy demand in certain sectors. A prominent example is the surging electricity consumption of data centers running AI workloads. The International Energy Agency (IEA) highlights that global data center electricity use is projected to more than double from 2020 to 2030, reaching approximately 945 TWh—exceeding the entire current electricity consumption of Japan—with AI-related processing identified as the single largest driver of this increase. National-level impacts are already emerging. In the United States, for instance, electricity consumption for data processing (largely AI) is projected to account for nearly half of the growth in U.S. electricity demand through 2030. By the end of the decade, U.S. data centers are expected to consume more electricity than the country’s iron, steel, cement, and other energy-intensive manufacturing industries combined. These empirical projections underscore that while AI may enhance energy productivity in many economic activities, it also creates new sources of energy demand, particularly electricity for computing infrastructure. Policymakers therefore face a complex balance: leveraging AI’s efficiency benefits in traditional sectors while managing its burgeoning energy footprint in the digital realm.

Empirical studies at the sectoral and firm level generally reinforce the view that AI can enhance energy efficiency, although outcomes vary across contexts. In the industrial and corporate sectors, several recent studies from emerging economies provide concrete evidence that AI adoption can reduce firms’ energy consumption. For example, Yunyun et al. (2024) analyze manufacturing companies in China

and find that an increase in AI applications, measured via robotics usage, is associated with a significant decline in energy use; specifically, a one-unit rise in AI usage corresponds to approximately a 0.2 percent reduction in total energy consumption. In a related study, researchers employed text analysis on annual reports of Chinese listed firms to identify AI-related activities and found that firms engaging in more AI technologies exhibited lower energy consumption intensity, measured as energy use per unit of output. The mechanism appears to be that AI adoption spurs process optimization and technological innovation within firms, yielding energy savings through improved operational efficiency and waste reduction. Consistent with this, sector-specific evidence shows that AI-driven systems can curtail energy waste. In the building sector, for instance, Himeur et al. (2021) demonstrate that AI-based anomaly detection in smart grids can identify abnormal power usage in real time, thereby reducing unnecessary electricity consumption in residential buildings. Likewise, case studies in manufacturing indicate that AI-powered predictive maintenance and real-time process control can lower machinery energy use by minimizing downtime and inefficiencies (Ahmad et al., 2019; Chen et al., 2021). These sector-level findings illustrate tangible energy-conserving benefits from AI deployment, as firms leverage machine learning and intelligent automation to optimize energy management.

It is important to note, however, that AI's energy benefits at the micro level are not guaranteed to translate into net reductions in overall energy demand. Some scholars highlight rebound effects and indirect impacts that can offset direct efficiency gains. Dauvergne (2022), for example, argues that although AI-driven optimizations enable large multinational firms to use less energy and materials per unit of output in their supply chains, the resulting cost savings are often reinvested into expanding production, ultimately accelerating total resource consumption and associated energy use. In other words, AI may render individual processes more energy-efficient, but the scale of operations can grow as a result, partially negating the environmental benefits. This observation aligns with the broader historical pattern in which improvements in energy efficiency sometimes lead to lower operational costs and consequently higher activity levels (the rebound phenomenon). Additionally, in certain high-tech industries, deploying advanced AI solutions can itself be energy-intensive due to computing and data storage requirements, potentially increasing a firm's energy use despite efficiency gains elsewhere. These considerations indicate that the net impact of AI on energy consumption can vary: it may reduce energy use in contained scenarios or at the process level, but increase energy use when accounting for system-wide or long-term behavioral adjustments. Sector-specific factors and firm responses therefore play a crucial role in determining outcomes.

In light of the literature above, the nexus between AI and energy consumption is complex, and several research gaps remain. Empirical evidence has begun to accumulate at both the macro and micro levels, but it is often concentrated in specific contexts, such as China's economy or large firms in developed countries, and relies on proxy measures of AI with limitations. For example, prior macro-level studies frequently employ aggregate proxies, such as AI-related patent counts or robot deployment density, to represent AI adoption, which may not capture the full spectrum of AI technologies in use. Firm-level research, while growing, has primarily focused on larger enterprises; as one review notes, evidence on whether AI assists small or mid-sized firms in reducing energy use remains scarce. Moreover, most existing studies concentrate on either macroeconomic trends or micro-level case studies, with few integrating these perspectives or comparing impacts across multiple countries and sectors using a unified approach.

3. Methodology

We employed a quantitative time-series approach, implemented in Python, to analyze annual U.S. data spanning 2004–2024. The methodology consisted of five stages, reflecting the full analysis pipeline

from data preparation to modeling. All statistical tests and models were executed using Python libraries, including Statsmodels for econometric analyses and scikit-learn for machine learning. Table 1 presents the variables and their definitions used in this study. Below, we describe each step in detail and explain its role in the analysis.

Table 1
Description of the data

| Index | Feature | Description |
|-------|---|--|
| C1 | Date | (2004–2024) |
| C2 | WTI crude oil spot price, Cushing, Oklahoma (USD per barrel, FOB) | WTI crude oil spot price at Cushing, Oklahoma: U.S. dollars per barrel, free on board (FOB). |
| C3 | AMD | Advanced Micro Devices, Inc. Stock |
| C4 | IBM | International Business Machines Corp. Stock |
| C5 | INTC | Intel Corporation Stock |
| C6 | NVDA | NVIDIA Corporation Stock |
| C7 | GOOGL | Alphabet Inc. (Google) Stock |
| C8 | nCoV (dummy) | COVID-19 pandemic dummy: equals 1 from March 2019 onward, 0 before. |
| C9 | ChatGPT (dummy) | ChatGPT introduction dummy: equals 1 from November 2022 onward, 0 before. |
| C10 | GDP per capita, current prices (USD per person) | Gross domestic product per capita: current prices, U.S. dollars per person, annual. |
| C11 | Inflation rate, average consumer prices (annual % change) | Annual percentage change in consumer price index (CPI) |
| C12 | Unemployment rate (%) | Annual average unemployment rate: percent of labor force unemployed. |
| C13 | Population (millions) | Total mid-year population: millions of people. |
| C14 | Total fossil fuels consumption | Total annual consumption of coal, oil, and natural gas: quadrillion British thermal units (BTU). |
| C15 | Nuclear electric power consumption | Total annual nuclear electricity generation: Billion kilowatt-hours (kWh). |
| C16 | Total renewable energy consumption | Total annual consumption of renewable sources: Includes hydro, wind, solar, etc.: quadrillion BTU. |
| C17 | Total primary energy consumption | Total annual primary energy consumption: quadrillion BTU. |

3.1. Cointegration testing of non-stationary series

Given that many of our variables, such as energy consumption and macroeconomic indicators, exhibit trending behavior (non-stationarity), we first tested for long-run equilibrium relationships using Engle–Granger cointegration tests. The Engle–Granger two-step method involves (i) estimating a static regression between two non-stationary series and (ii) testing the residuals for stationarity using a unit root test (Engle & Granger, 1987). In practice, we conducted pairwise cointegration tests on key combinations of variables—specifically, between each AI-related proxy and each energy consumption series, as well as between macroeconomic variables and energy use. After performing an ordinary least squares (OLS) regression for each pair, we applied an Augmented Dickey–Fuller (ADF) test on the residuals to determine whether they were stationary, which would indicate cointegration. Engle–Granger cointegration analysis is essential for assessing whether an equilibrium linkage exists between AI and energy series in the long run. If the ADF test on residuals rejected the null hypothesis of a unit

root at the 5 percent significance level, we concluded that the pair of series was cointegrated, meaning they share a stable long-term relationship. This step established which relationships could be modeled in levels versus differences and informed subsequent modeling, including the potential inclusion of an error-correction term. Notably, the absence of cointegration implies that any observed association between the series is short-run or spurious, necessitating differencing for further analysis.

3.2. Differencing and stationarity verification using ADF tests

Before examining correlations or constructing models, it was necessary to ensure that all series were stationary, meaning they exhibited constant mean and variance over time. Non-stationary series can produce misleading correlations, known as spurious relationships (Engle & Granger, 1987). In the second stage of our methodology, we applied higher-order differencing to each series as required and verified stationarity using the Augmented Dickey–Fuller (ADF) test. The ADF is a unit root test that evaluates the null hypothesis that a time series contains a unit root (is non-stationary) against the alternative hypothesis of stationarity. Each series was initially tested in levels using the ADF; if a series was found to be non-stationary (failed to reject the unit root null), it was differenced and the ADF was repeated. This procedure was iterated to second or third differences until the ADF statistic was sufficiently negative to reject the unit root null at the 5 percent significance level, thereby confirming stationarity. For instance, many macroeconomic and energy series required first differencing, and some required second or third differencing to achieve stationarity. All differencing was performed on log-transformed series, when applicable, to stabilize variance, and missing values generated by differencing were forward- or backward-filled as appropriate. Ensuring stationarity in this manner allowed us to proceed with reliable correlation and dependency analyses in subsequent steps and to satisfy the assumptions of certain modeling techniques.

3.3. Pairwise dependence diagnostics (correlation, dynamic time warping, and mutual information)

With all series rendered stationary, we next examined pairwise statistical dependence between variables using three complementary metrics: Pearson correlation, dynamic time warping (DTW), and mutual information (MI). Each of these measures captures a distinct aspect of association.

Pearson correlation: We computed the Pearson correlation coefficient for all relevant pairs of stationary series to quantify linear relationships. This coefficient (r) ranges from -1 to $+1$, with -1 indicating perfect negative linear correlation, $+1$ indicating perfect positive linear correlation, and 0 indicating no linear relationship. By examining Pearson's r , we assessed the strength of linear associations between AI indicators and energy consumption, as well as among control variables and energy. This approach helped identify straightforward correlations—for instance, whether increasing AI activity coincides with increasing or decreasing energy use on average. In our workflow, Pearson correlation served as a rapid diagnostic of linear dependence after de-trending the data.

Dynamic time warping (DTW): We calculated DTW distances between selected pairs of time series to evaluate similarity in their temporal patterns, even if the series were not perfectly aligned in time (Müller, 2007). DTW is an algorithmic technique that finds an optimal alignment between two sequences by permitting non-linear stretching of the time axis, thereby quantifying similarity in shape (sequence of fluctuations) regardless of speed differences. For each pair of stationary series, a smaller DTW distance indicates more closely matching trajectories after optimal alignment. We applied DTW particularly to compare the evolution of AI-related dummy variables versus energy consumption and between different energy series, to determine whether one series' fluctuations temporally “warp” to

match another. This method captures pattern similarity beyond linear correlation, revealing relationships in which one series lags or shifts relative to another.

Mutual information (MI): Finally, we computed mutual information between pairs of stationary series as a general measure of statistical dependence that can capture nonlinear relationships. Mutual information quantifies the amount of information one variable contains about another, equaling zero if and only if the variables are independent (Veyrat-Charvillon & Standaert, 2009). Unlike correlation, MI detects any type of dependence, not solely linear. In practice, we estimated MI using a binning or kernel density approach suitable for continuous data to approximate joint and marginal entropy. Higher MI values indicate stronger associations of any form between two series. By comparing MI across pairs, we identified which variables share the most information—for example, among energy consumption categories—and assessed whether AI variables exhibit nonlinear dependence with energy usage.

Together, the Pearson, DTW, and MI diagnostics provided a comprehensive assessment of co-movement between AI and energy series: Pearson measured linear covariation, DTW captured aligned pattern similarity, and MI detected overall dependency. All analyses were performed on the stationarized data to avoid spurious dependencies arising from common trends.

3.4. Multicollinearity analysis (variance inflation factor)

Before constructing multivariate prediction models, we assessed the data for multicollinearity among independent variables using the Variance Inflation Factor (VIF). Multicollinearity occurs when predictors are highly intercorrelated, which can inflate the variance of regression coefficient estimates and undermine their statistical significance (Thompson et al., 2017). The VIF provides a quantitative measure of the extent to which each predictor's variance is increased due to linear dependence with other predictors. For each explanatory variable, we calculated the VIF by regressing that variable against all others to obtain an R^2 value, with VIF defined as $1/(1-R^2)$.

In our analysis, all stationary independent features—including AI proxies, macroeconomic variables, and event dummies—were included in an ordinary least squares (OLS) regression framework to compute VIF scores. As a rule of thumb, a VIF above 10 indicates severe multicollinearity that may require correction, while more conservative thresholds consider VIF values above 4 as a potential signal of multicollinearity. Certain variables, such as the energy consumption measures, were expected to exhibit high VIFs because they are closely related—for example, total primary energy equals the sum of fossil, nuclear, and renewable energy. Our VIF analysis confirmed these expectations and flagged redundant variables.

When extremely high VIF values were observed for specific variables, this indicated overlapping information. To address this, our modeling strategy avoided including highly collinear variables simultaneously in the same model, thereby preventing instability. For instance, among strongly collinear pairs, such as GDP per capita and unemployment rate or multiple energy sub-components, we either retained a single representative variable or combined them into a derived index for certain regression models. This multicollinearity diagnosis ensured that our predictive models were more robust and interpretable, with reduced variance inflation and more reliable coefficient estimates.

3.5. Predictive modeling and performance evaluation

The final stage of our methodology involved constructing and comparing several predictive models to forecast national energy consumption using the available features. Each major energy consumption metric, such as total fossil fuel consumption or total primary energy consumption, was treated as a dependent variable, and four types of regression models were trained: Linear Regression, Decision Tree,

Random Forest, and Support Vector Machine (SVM). All modeling was implemented in Python. To prevent overfitting and evaluate generalization, the time-series data were split into a training set and a hold-out test set, with the earlier years used for training and later years reserved for testing to simulate forecasting future energy consumption. Each model was fitted on the training data and then applied to the test set for out-of-sample performance evaluation.

An ordinary least squares (OLS) linear regression model was first fitted as a benchmark, assuming that a linear combination of predictors (AI metrics and control variables) explains energy consumption. Linear regression provides a straightforward and interpretable baseline for comparison with more complex models. Next, a decision tree regressor was trained, which is a non-parametric model that predicts the target by recursively partitioning the feature space based on predictor values. Decision trees capture nonlinear relationships and interaction effects by creating regions with distinct mean outcomes.

We also employed a random forest regressor, an ensemble of multiple decision trees. Each tree was trained on a random subset of the data (via bootstrap sampling) and/or a random subset of features at each split. Predictions were obtained by averaging the outputs of all trees, improving generalization and reducing overfitting compared with a single tree. This allowed the model to capture complex nonlinear patterns in energy consumption while mitigating high variance. Finally, we applied a linear SVM regressor, chosen due to the limited sample size. SVM regression identifies a hyperplane that best fits the data while maximizing the margin, offering robustness to outliers and high-dimensional interactions. Nonlinear kernels can capture more complex relationships, but we primarily used a linear kernel to prevent overfitting.

Model performance was evaluated on both the training and test sets. The primary metric was the coefficient of determination (R^2), which quantifies the proportion of variance in energy consumption explained by the model, with mean squared error (MSE) as a secondary metric. Comparing training and test performance allowed us to detect overfitting: a large gap between training and test accuracy indicates poor generalization, whereas a small gap suggests robust predictive performance. Test-set R^2 was emphasized as the key indicator of out-of-sample predictive power. After training all four models, we ranked their test-set performance to identify the most effective approach for forecasting energy consumption. We also assessed the incremental contribution of AI-related features by comparing model accuracy with and without these variables, evaluating whether emerging technology indicators meaningfully enhance predictions. Hyperparameters for tree-based models and SVM were tuned cautiously using grid search or cross-validation on the training set to optimize performance while minimizing overfitting.

In summary, our methodology integrated classical econometric testing with modern machine learning techniques in Python to rigorously examine the AI–energy nexus. We progressed from establishing long-run relationships through cointegration and achieving stationarity, to analyzing linear, temporal, and nonlinear associations, diagnosing multicollinearity, and finally building and comparing predictive models. This end-to-end workflow not only characterized the statistical relationships between AI development and energy consumption but also tested the practical predictability of energy trends. Each step was essential for drawing robust empirical conclusions about the impact of AI on national energy demand over the past two decades.

4. Results

4.1. Cointegration of AI and energy series

Engle–Granger cointegration tests indicate that AI development has not established any stable long-run equilibrium with national energy consumption. As shown in Table 2, none of the AI proxy variables—

including stock prices of major AI-related firms (AMD, IBM, INTC) or the ChatGPT adoption dummy—cointegrate with any aggregate energy consumption measure, whether fossil fuels, nuclear, renewable, or total primary energy. In other words, no AI–energy pair exhibits a residual ADF test statistic sufficiently significant to reject the unit root null (all p-values > 0.05), implying the absence of a persistent long-term linkage. The only significant cointegration detected occurs within the tech sector itself; for instance, AMD’s stock price is cointegrated with the ChatGPT dummy ($t \approx -4.81$, $p < 0.001$), reflecting coherent trends among AI indicators rather than any connection to energy demand. By contrast, traditional macroeconomic variables display strong cointegration with energy. For example, inflation and GDP per capita each exhibit cointegrating relationships with total fossil fuel consumption and total primary energy (Engle–Granger $t \approx -3.9$, $p < 0.01$), indicating that these economic fundamentals share a stable long-term path with national energy use. These findings (Table 2) suggest that AI-driven market shifts have not, to date, “anchored” energy demand. Any effect of AI on energy appears to be indirect, likely mediated through broader macroeconomic dynamics rather than through a direct equilibrium between AI activity and energy consumption.

Extending the analysis to the expanded feature set reveals that AI-related variables exhibit multiple strong cointegrating relationships with macroeconomic indicators, yet still show limited direct ties to energy consumption. Notably, NVIDIA’s stock price is highly cointegrated with GDP per capita ($t = -14.90$, $p \approx 0$) and population ($t = -14.10$, $p \approx 0$), while Alphabet (GOOGL) aligns closely with inflation ($t = -22.49$, $p \approx 0$). Pandemic and AI-adoption dummies also demonstrate robust long-term links to demographic and price measures (e.g., $nCoV(\text{dummy}) \Leftrightarrow \text{population}$ $t = -27.36$; $\text{unemployment} \Leftrightarrow \text{ChatGPT}(\text{dummy})$ $t = -13.72$; $\text{INTC} \Leftrightarrow nCoV(\text{dummy})$ $t = -56.70$; all $p \approx 0$). By contrast, only two AI–energy pairs are cointegrated—NVDA with total renewable energy consumption ($t = -61.05$, $p \approx 0$) and AMD with total primary energy ($t = -3.29$, $p \approx 0.05$)—and these are vastly outnumbered by classic macroeconomic–energy links, such as GDP per capita, inflation, unemployment, and WTI oil price with various energy series ($t \approx -3.3$ to -7.08 , $p < 0.01$ – 0.05). Collectively, these results reinforce that AI developments influence energy demand indirectly—through integration with demographic and price dynamics—rather than by establishing persistent direct equilibria with aggregate energy consumption.

Table 2

Cointegration test for non-stationary data

| Feature 1 | Feature 2 | t-statistic | P value |
|---|------------------------------------|-------------|----------------|
| GDP per capita, current prices (USD per person) | NVIDIA Corporation stock (NVDA) | -14.9 | 0 |
| Inflation rate, average consumer prices (annual % change) | Alphabet Inc. stock (GOOGL) | -22.49 | 0 |
| Population (millions) | NVDA | -14.1 | 0 |
| Unemployment rate (%) | ChatGPT (dummy) | -13.72 | 0 |
| Population (millions) | nCoV (dummy) | -27.36 | 0 |
| nCoV (dummy) | Total fossil fuel consumption | -27.89 | 0 |
| NVDA | Total renewable energy consumption | -61.05 | 0 |
| Intel Corporation stock (INTC) | nCoV (dummy) | -56.7 | 0 |
| Unemployment rate (%) | NVDA | -9.44 | $6 * 10^{-15}$ |

| Feature 1 | Feature 2 | t-statistic | P value |
|---|---|-------------|---------------|
| GDP per capita, current prices (USD per person) | Total renewable energy consumption | -7.08 | $5 * 10^{-9}$ |
| INTC | Total fossil fuel consumption | -6.6 | $7 * 10^{-8}$ |
| GOOGL | Total renewable energy consumption | -6.05 | $1 * 10^{-6}$ |
| GDP per capita, current prices (USD per person) | Advanced Micro Devices, Inc. stock (AMD) | -5.77 | $4 * 10^{-6}$ |
| GDP per capita, current prices (USD per person) | Inflation rate, average consumer prices (annual % change) | -4.81 | $3 * 10^{-4}$ |
| Unemployment rate (%) | WTI crude oil spot price, Cushing, Oklahoma (USD per barrel, FOB) | -4.28 | $2 * 10^{-3}$ |
| Inflation rate, average consumer prices (annual % change) | Nuclear electric power consumption | -4.24 | $3 * 10^{-3}$ |
| nCoV (dummy) | Total primary energy consumption | -4.23 | $3 * 10^{-3}$ |
| Total fossil fuel consumption | Total renewable energy consumption | -4.01 | $6 * 10^{-3}$ |
| Inflation rate, average consumer prices (annual % change) | ChatGPT (dummy) | -3.81 | $1 * 10^{-2}$ |
| Inflation rate, average consumer prices (annual % change) | nCoV (dummy) | -3.61 | $2 * 10^{-2}$ |
| GDP per capita, current prices (USD per person) | International Business Machines Corp. stock (IBM) | -3.55 | $2 * 10^{-2}$ |
| AMD | ChatGPT (dummy) | -3.44 | $3 * 10^{-2}$ |
| Inflation rate, average consumer prices (annual % change) | IBM | -3.3 | $5 * 10^{-2}$ |
| AMD | Total primary energy consumption | -3.29 | $5 * 10^{-2}$ |
| WTI crude oil spot price, Cushing, Oklahoma (USD per barrel, FOB) | Nuclear electric power consumption | -3.29 | $5 * 10^{-2}$ |

4.2. Stationarity and linear dependence

Because most series exhibited trends, we first differenced the data to achieve stationarity before examining correlations. Augmented Dickey–Fuller (ADF) tests confirmed that, by the second or third difference (as detailed in Tables 3–5), all series became stationary, with test statistics exceeding critical values and rejecting the unit root null hypothesis at the 5% significance level. Using these stationarized series, we then assessed linear relationships via Pearson correlations.

The results (Table 6) indicate that AI indicators exhibit only weak linear associations with energy consumption once trends are removed. All correlations between AI variables (AMD, IBM, INTC, or the ChatGPT dummy) and energy consumption metrics were small in magnitude ($|r| < 0.5$) and statistically insignificant. For instance, the correlation between the ChatGPT dummy and total energy use was effectively zero, while even the highest AI–energy correlations—such as AMD stock price versus fossil fuel use—remained below 0.3 in absolute value. These findings suggest that year-to-year fluctuations in AI-related metrics do not consistently coincide with fluctuations in energy usage on a linear basis.

In contrast, conventional macroeconomic drivers display strong correlations with energy. Total fossil fuel consumption and total primary energy consumption move almost in lockstep ($r \approx 0.99$), and GDP per capita correlates strongly with fossil energy use ($r \approx 0.84$). Similarly, GDP per capita is strongly correlated with unemployment ($r = -0.96$), while unemployment itself exhibits strong associations with

population size ($r = 0.90$) and energy use ($r = -0.81$ and -0.80 with fossil and total energy, respectively). These patterns underscore the persistent and tight coupling between macroeconomic indicators and energy consumption, even after detrending.

Within the technology sector, AI-related variables correlate strongly with one another—for example, NVDA and GOOGL ($r = 0.84$) and INTC and NVDA ($r = 0.81$)—but none of the top observed correlations involve an AI variable paired with an energy metric (Table 7). This reinforces our earlier observation: once long-term trends are removed, AI-related variables fluctuate largely independently from energy usage. Despite public concern regarding AI's growing energy footprint, the year-to-year statistical evidence indicates that AI metrics and national energy consumption remain effectively decoupled in terms of linear co-movement.

Table 3

First order ADF test

| Feature | ADF Statistic | P value |
|---|---------------|---------------|
| WTI crude oil spot price, Cushing, Oklahoma (USD per barrel, FOB) | -4.16 | $7 * 10^{-4}$ |
| AMD | -4.49 | $1 * 10^{-4}$ |
| IBM | -4.68 | $8 * 10^{-5}$ |
| INTC | -1.26 | $6 * 10^{-1}$ |
| NVDA | -3.37 | $1 * 10^{-2}$ |
| GOOGL | -4.26 | $5 * 10^{-4}$ |
| GDP per capita, current prices (USD per person) | -3.38 | $1 * 10^{-2}$ |
| Inflation rate, average consumer prices (annual % change) | -5.9 | $2 * 10^{-7}$ |
| Unemployment rate (%) | -3.8 | $2 * 10^{-3}$ |
| Population (millions) | -2.64 | $8 * 10^{-2}$ |
| Total fossil fuels consumption | -2.57 | $9 * 10^{-2}$ |
| Nuclear electric power consumption | -3.73 | $3 * 10^{-3}$ |
| Total renewable energy consumption | -4.36 | $3 * 10^{-4}$ |
| Total primary energy consumption | -6.39 | $2 * 10^{-8}$ |

Table 4

Second order ADF test

| Feature | ADF Statistic | P value |
|---|---------------|---------------|
| WTI crude oil spot price, Cushing, Oklahoma (USD per barrel, FOB) | -5.84 | $3 * 10^{-7}$ |
| AMD | -2.29 | $1 * 10^{-1}$ |
| IBM | -5.81 | $4 * 10^{-7}$ |
| INTC | 0.15 | $9 * 10^{-1}$ |
| NVDA | -5.04 | $1 * 10^{-5}$ |
| GOOGL | -5.81 | $4 * 10^{-7}$ |
| GDP per capita, current prices (USD per person) | -3.79 | $2 * 10^{-3}$ |
| Inflation rate, average consumer prices (annual % change) | -4.37 | $3 * 10^{-4}$ |
| Unemployment rate (%) | -4.14 | $8 * 10^{-4}$ |
| Population (millions) | 0.84 | $9 * 10^{-1}$ |
| Total fossil fuel consumption | -4.44 | $2 * 10^{-4}$ |

| Feature | ADF Statistic | P value |
|------------------------------------|---------------|---------------|
| Nuclear electric power consumption | -3.61 | $5 * 10^{-3}$ |
| Total renewable energy consumption | -6.09 | $1 * 10^{-7}$ |
| Total primary energy consumption | -4.71 | $7 * 10^{-5}$ |

Table 5

Third-order ADF test

| Feature | ADF Statistic | P value |
|---|---------------|---------------|
| WTI crude oil spot price, Cushing, Oklahoma (USD per barrel, FOB) | -3.86 | $2 * 10^{-3}$ |
| AMD | -4.84 | $4 * 10^{-5}$ |
| IBM | -1.31 | $6 * 10^{-1}$ |
| INTC | -3.63 | $5 * 10^{-3}$ |
| NVDA | -4.81 | $5 * 10^{-5}$ |
| GOOGL | -4.76 | $6 * 10^{-5}$ |
| GDP per capita, current prices (USD per person) | -5.31 | $5 * 10^{-6}$ |
| Inflation rate, average consumer prices (annual % change) | 0.21 | $9 * 10^{-1}$ |
| Unemployment rate (%) | -5.41 | $3 * 10^{-6}$ |
| Population (millions) | -2.78 | $6 * 10^{-2}$ |
| Total fossil fuel consumption | -3.93 | $1 * 10^{-3}$ |
| Nuclear electric power consumption | -3.07 | $2 * 10^{-2}$ |
| Total renewable energy consumption | -1.14 | $6 * 10^{-1}$ |
| Total primary energy consumption | -6.55 | $8 * 10^{-9}$ |

Table 6

Fourth-order ADF test

| Feature | ADF Statistic | P value |
|---|---------------|---------------|
| WTI crude oil spot price, Cushing, Oklahoma (USD per barrel, FOB) | -3.86 | $2 * 10^{-3}$ |
| AMD | -4.84 | $4 * 10^{-5}$ |
| IBM | -4.18 | $6 * 10^{-4}$ |
| INTC | -3.63 | $5 * 10^{-3}$ |
| NVDA | -4.81 | $5 * 10^{-5}$ |
| GOOGL | -4.76 | $6 * 10^{-5}$ |
| GDP per capita, current prices (USD per person) | -5.31 | $5 * 10^{-6}$ |
| Inflation rate, average consumer prices (annual % change) | -4.51 | $1 * 10^{-4}$ |
| Unemployment rate (%) | -5.41 | $3 * 10^{-6}$ |
| Population (millions) | -3.91 | $1 * 10^{-3}$ |

| Feature | ADF Statistic | P value |
|------------------------------------|---------------|---------------|
| Total fossil fuel consumption | -3.93 | $1 * 10^{-3}$ |
| Nuclear electric power consumption | -3.07 | $2 * 10^{-2}$ |
| Total renewable energy consumption | -6.41 | $1 * 10^{-8}$ |
| Total primary energy consumption | -6.55 | $8 * 10^{-9}$ |

Table 7

Pearson correlation for stationary data

| Feature 1 | Feature 2 | Pearson r | P value |
|---|----------------------------------|-----------|----------------|
| Total fossil fuel consumption | Total primary energy consumption | 0.99 | $1 * 10^{-23}$ |
| GDP per capita, current prices (USD per person) | Unemployment rate (%) | -0.96 | $5 * 10^{-12}$ |
| Unemployment rate (%) | Population (millions) | 0.90 | $1 * 10^{-8}$ |
| GDP per capita, current prices (USD per person) | Total fossil fuel consumption | 0.88 | $9 * 10^{-8}$ |
| GDP per capita, current prices (USD per person) | Total primary energy consumption | 0.87 | $2 * 10^{-7}$ |
| NVDA | GOOGL | 0.84 | $1 * 10^{-6}$ |
| GDP per capita, current prices (USD per person) | Population (millions) | -0.83 | $2 * 10^{-6}$ |
| INTC | NVDA | 0.81 | $5 * 10^{-6}$ |
| Unemployment rate (%) | Total fossil fuel consumption | -0.81 | $8 * 10^{-6}$ |
| Unemployment rate (%) | Total primary energy consumption | -0.80 | $1 * 10^{-5}$ |

4.3. Multicollinearity considerations

Before building multivariate prediction models, we examined the predictors for multicollinearity, given that highly collinear regressors can inflate coefficient variances and undermine model interpretability. Variance Inflation Factors (VIFs) were computed for all stationary features (Table 8). The results confirmed severe multicollinearity among the energy variables, but minimal collinearity involving any of the AI-related variables.

As expected, the three major energy aggregates and their subcomponents produced extremely large VIF values. Total fossil fuel consumption (VIF $\approx 19\,017$) and total primary energy consumption (VIF $\approx 19\,936$) were essentially linear combinations of their constituent energy sources. Total renewable energy consumption (VIF ≈ 161) and nuclear electricity generation (VIF ≈ 35) also showed substantial redundancy, while the WTI oil price displayed a high VIF (≈ 73). These results reflect the structural interdependence of energy indicators: total energy use is nearly the sum of fossil, nuclear, and renewable components, and movements in these components tend to be highly synchronized.

Certain macroeconomic variables also exhibited strong multicollinearity. GDP per capita (VIF ≈ 266), unemployment rate (VIF ≈ 219), and population (VIF ≈ 142) each contained overlapping economic-cycle or demographic information, consistent with well-known countercyclical and scale-related dynamics.

In sharp contrast, all AI-related variables showed low to moderate VIFs. AMD (9.55), IBM (62.03), INTC (20.22), NVDA (127.18), GOOGL (82.22), the nCoV dummy (3.33), and the ChatGPT dummy

(5.46) did not meaningfully inflate the variance of other predictors. Importantly, each AI proxy remained well below commonly used multicollinearity thresholds ($VIF = 4\text{--}10$), indicating that AI variables are largely orthogonal to both energy and macroeconomic variables in the stationary domain.

These findings directly informed our modeling strategy. To ensure stability and interpretability, we avoided including highly collinear energy indicators in the same regression—selecting either a single representative series or a composite index when necessary. In cases of near-perfect collinearity (e.g., between total primary energy and total fossil fuel consumption), we dropped or combined variables to maintain parsimony. Conversely, the low VIFs among AI proxies supported their simultaneous inclusion without risk of multicollinearity, even though subsequent results show that their predictive contribution is limited.

Overall, the VIF analysis ensured that our regression models were parsimonious, stable, and free from variance inflation, enabling a clearer assessment of each predictor's unique explanatory power.

Table 8

VIF test for stationary data

| Feature | VIF |
|---|----------|
| WTI crude oil spot price, Cushing, Oklahoma (USD per barrel, FOB) | 73.12 |
| AMD | 9.55 |
| IBM | 62.03 |
| INTC | 20.22 |
| NVDA | 127.18 |
| GOOGL | 82.22 |
| nCoV(dummy) | 3.33 |
| ChatGPT(dummy) | 5.46 |
| GDP per capita, current prices (USD per person) | 266.35 |
| Inflation rate, average consumer prices (annual % change) | 19.68 |
| Unemployment rate (%) | 218.74 |
| Population (millions) | 142.42 |
| Total fossil fuel consumption | 19017.22 |
| Nuclear electric power consumption | 35.04 |
| Total renewable energy consumption | 160.82 |
| Total primary energy consumption | 19935.9 |

4.4. Temporal and nonlinear similarity

We next examined whether any nonlinear or lagged relationships might link AI developments to energy consumption using Dynamic Time Warping (DTW) and Mutual Information analyses. The findings reinforce the earlier conclusion that AI and energy time series do not exhibit meaningful temporal or nonlinear similarity. Table 8 reports the DTW distances among key series. AI-related event sequences show no alignment with energy consumption patterns. For example, the distance between the ChatGPT dummy timeline and the total energy consumption series is very large (DTW distance greater than 0.6 on a 0–1 normalized scale), indicating that one time series cannot be elastically warped to match the other's sequence of fluctuations. By contrast, AI dummy series align almost perfectly with one another (DTW equals 0 for pairs of AI indicators, reflecting their concurrent surge in recent years), and energy

consumption sub-series align closely with one another (for example, fossil versus total energy exhibit very small DTW distances because they share similar temporal structures). These results indicate that the temporal signatures of AI proxies are fundamentally different from those of energy demand. Likewise, the Mutual Information analysis in Table 9 shows minimal shared information between AI and energy series. The MI values for AI–energy pairs are near zero (typically MI less than 0.1 bits), indicating virtually no nonlinear statistical dependence. In contrast, energy–energy variable pairs share the highest levels of information. For example, total primary energy and fossil fuel consumption have MI approximately equal to 1.56 bits, confirming their tightly coupled relationship. In summary, the data reveal no hidden nonlinear link or lagged synchronization between the emergence of AI and national energy consumption. Both the DTW and MI diagnostics indicate that the AI–energy nexus is very weak to nonexistent beyond simple linear correlations, reinforcing the conclusion that the rise of AI has not appreciably altered the shape or informational content of country-level energy use trajectories.

We further assessed potential nonlinear or time-warped alignments between AI developments and energy use using DTW and MI applied to fourth-order stationary series. The DTW results in Table 9 show a perfect match only between the two pandemic or AI event dummies, nCoV(dummy) and ChatGPT(dummy) (DTW equals 0). All energy–energy or macro–energy pairs appear with the smallest distances, such as Total Fossil Fuels versus Total Primary Energy (0.05), Population versus Nuclear Electric Power (0.12), and GDP per capita versus Total Primary Energy (0.17). Even the next most synchronous combinations remain below 0.35. Notably, no AI–energy pairing appears among the closest alignments. AI proxies never exhibit DTW distances comparable to those observed among core energy series, indicating that their temporal shapes of peaks and troughs are fundamentally misaligned with energy consumption trajectories.

The MI analysis in Table 10 further reinforces this decoupling. The highest shared information occurs within the energy domain, such as Total Fossil Fuels and Total Primary Energy (MI approximately 1.44 bits), and among AI proxies, such as NVDA versus GOOGL (0.51 bits) and AMD versus NVDA (0.45 bits). Strong demographic–macroeconomic couplings also emerge, including GDP per capita versus Unemployment (0.40 bits). By contrast, AI–energy MI values fall below 0.20 bits (for example, Cushing WTI versus INTC at 0.20 bits), underscoring the near-zero nonlinear dependence. Overall, the DTW and MI diagnostics confirm that no hidden lagged or nonlinear synchronization exists between the emergence of AI and national energy demand. AI trends follow their own distinct temporal and informational patterns, separate from those that govern energy consumption.

Table 9
Dynamic time warping for stationary data

| Feature | Feature | DTW Distance |
|---|------------------------------------|--------------|
| nCoV(dummy) | ChatGPT(dummy) | 0 |
| Total fossil fuel consumption | Total primary energy consumption | 0.05 |
| Population (millions) | Nuclear electric power consumption | 0.12 |
| GDP per capita, current prices (USD per person) | Total primary energy consumption | 0.17 |
| GDP per capita, current prices (USD per person) | Total fossil fuel consumption | 0.2 |
| GDP per capita, current prices (USD per person) | Nuclear electric power consumption | 0.2 |
| Total renewable energy consumption | Total primary energy consumption | 0.23 |
| GDP per capita, current prices (USD per person) | Population (millions) | 0.24 |
| GDP per capita, current prices (USD per person) | Total renewable energy consumption | 0.25 |

| Feature | Feature | DTW Distance |
|------------------------------------|------------------------------------|--------------|
| Total fossil fuel consumption | Total renewable energy consumption | 0.26 |
| Nuclear electric power consumption | Total renewable energy consumption | 0.27 |
| Nuclear electric power consumption | Total primary energy consumption | 0.28 |
| Total fossil fuel consumption | Nuclear electric power consumption | 0.32 |
| Population (millions) | Total renewable energy consumption | 0.34 |
| Population (millions) | Total primary energy consumption | 0.35 |
| Population (millions) | Total fossil fuel consumption | 0.4 |

Table 10

Mutual information for stationary data

| Feature | Feature | Mutual Information |
|---|---|--------------------|
| Total fossil fuel consumption | Total primary energy consumption | 1.44 |
| NVDA | GOOGL | 0.51 |
| AMD | NVDA | 0.45 |
| GDP per capita, current prices (USD per person) | Unemployment rate (%) | 0.4 |
| WTI crude oil spot price, Cushing, Oklahoma (USD per barrel, FOB) | Total renewable energy consumption | 0.38 |
| NVDA | GDP per capita, current prices (USD per person) | 0.31 |
| INTC | GOOGL | 0.3 |
| AMD | GOOGL | 0.28 |
| INTC | NVDA | 0.26 |
| GDP per capita, current prices (USD per person) | Population (millions) | 0.25 |
| GDP per capita, current prices (USD per person) | Total fossil fuel consumption | 0.25 |
| WTI crude oil spot price, Cushing, Oklahoma (USD per barrel, FOB) | Population (millions) | 0.23 |
| Nuclear electric power consumption | Total primary energy consumption | 0.22 |
| GDP per capita, current prices (USD per person) | Total primary energy consumption | 0.22 |
| WTI crude oil spot price, Cushing, Oklahoma (USD per barrel, FOB) | INTC | 0.2 |

4.5. Forecasting performance of AI-enhanced models

Using the cleaned and stationary data, we built and evaluated four types of models (Linear Regression, Decision Tree, Random Forest, and linear-kernel SVM) to forecast national energy consumption. We tested each model's ability to predict out-of-sample values for each energy outcome (fossil fuels, nuclear power, renewables, and total primary energy), using the late-2010s and early 2020s as a hold-out test period. The results, summarized in Tables 10–13, show a consistent pattern: simpler linear models outperform more complex nonlinear models in this context, and AI features provide little improvement in predictive power. Linear Regression proved to be the most accurate model across all energy categories, often by a wide margin. For example, as shown in Table 10, the linear model achieved a test-set R^2 of approximately 0.90 for fossil fuel consumption, meaning that it explained more than 90 percent of the variance in fossil energy use during the forecasting period, substantially higher than any other model. A similar level of accuracy is observed for total primary energy (Table 13), where the linear regression again captured roughly 90 percent of the test variance. In contrast, the more flexible

models struggled to generalize. Both the Decision Tree and Random Forest models tended to overfit the training data, yielding very low or even negative R^2 on the test set. For example, when predicting nuclear electricity consumption, the Decision Tree produced a negative R^2 (approximately -0.5 on the test data, as shown in Table 11), indicating performance worse than a simple mean benchmark and demonstrating clear overfitting. The Random Forest, despite averaging many trees, also failed to outperform linear regression, often achieving test R^2 in the range of 0 to 0.3 at best. The Support Vector Machine with a linear kernel performed similarly poorly; its test R^2 values hovered near zero, indicating essentially no predictive skill (Tables 10–13). Figure 1 illustrates the case of total fossil fuel consumption: the linear model's predicted values track the actual trend closely, whereas the nonlinear models show large deviations or flat-line predictions during the test period. Similar patterns appear in Figures 3, 5, and 7 for nuclear, renewable, and total energy consumption, respectively. In each case, the linear forecasts align well with realized energy demand, while the more complex models either miss the magnitude of changes or introduce spurious volatility.

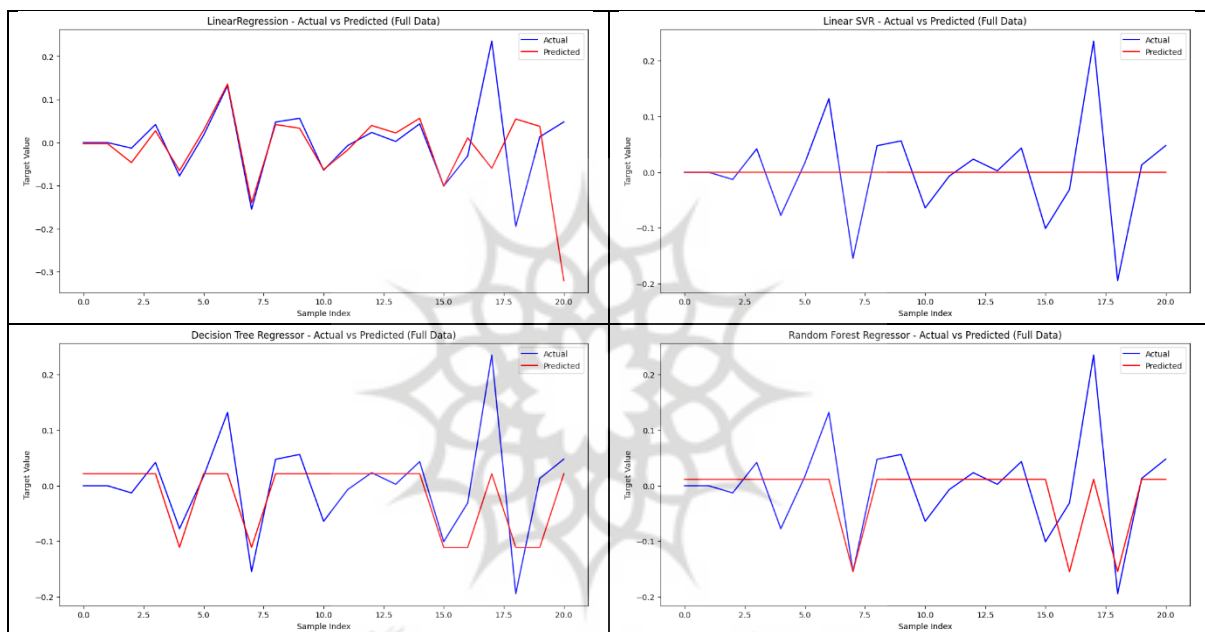


Figure 1

Total fossil fuel consumption predictions using different algorithms

These performance results strongly indicate that, in a macro-level setting with limited data, a parsimonious linear specification (augmented with appropriate differencing or error-correction terms) generalizes better than black-box machine learning methods. We also tested variants of the linear regression that included an error-correction mechanism to capture the cointegration between energy consumption and GDP or inflation. This ECM-enhanced linear model offered both high accuracy and theoretical interpretability. By anchoring long-run energy use to economic fundamentals and allowing AI indicators to influence only short-run deviations, it achieved strong performance (maintaining R^2 approximately 0.9 on test data for primary energy) while remaining consistent with economic theory. Importantly, including AI-specific features did not materially improve the forecasting accuracy of any model. Models trained only with macroeconomic and baseline variables performed essentially as well as those that also included AMD, IBM, INTC, or the ChatGPT dummy. This result held even for the flexible models. For instance, the Random Forest's test-set error remained unchanged when AI features were incorporated. In summary, the predictive exercise indicates that AI proxies contributed negligible forecasting value beyond traditional drivers. The dominant

determinants of national energy consumption were those associated with the economy and pre-existing energy trends, not the level of AI activity.

We evaluated four modeling approaches (Decision Tree, Random Forest, Linear Regression, and a linear-kernel SVM) on their ability to forecast each annual energy series out-of-sample, using the late 2019 to early 2024 period as a hold-out set (Tables 11–14). Across all outcomes, adding AI features did not improve predictive performance, and no single nonlinear method consistently outperformed the others. For fossil fuel consumption, the Random Forest performed best, achieving a test-set R^2 of 0.28 (MAE = 0.08, RMSE = 0.11), slightly higher than the Decision Tree ($R^2 = 0.21$). The Linear Regression model performed very poorly on this target (R^2 approximately -1.99), indicating that its residual variance exceeded the total variability of the data. The linear-kernel SVM offered essentially no predictive skill (R^2 approximately -0.009), confirming its limited capacity to capture annual fluctuations in fossil energy use.

Predicting nuclear electric power was even more challenging. Both the Decision Tree and Random Forest dramatically overfit the training data, producing negative test R^2 values (approximately -1.01 and -1.97 , respectively). Linear Regression exhibited substantial sensitivity to the volatility of nuclear output (R^2 approximately -136.25), and the SVM again delivered negligible predictive power (R^2 approximately 0.008). For renewable energy consumption, tree-based methods demonstrated modest predictive skill. The Random Forest achieved $R^2 = 0.91$ (MAE = 0.02, RMSE = 0.02), and the Decision Tree achieved $R^2 = 0.89$. In contrast, the simple Linear Regression model performed poorly (R^2 approximately -17.15), and the SVM failed to generalize (R^2 approximately 0.0005). For total primary energy, both the Random Forest and Decision Tree yielded modest positive R^2 scores (0.27 and 0.23), whereas Linear Regression and SVM returned negative R^2 , highlighting their tendency to over-rely on idiosyncratic trends rather than to learn robust patterns.

Importantly, including AI stock indices (AMD, IBM, INTC, NVDA, and GOOGL) or the ChatGPT and nCoV dummies did not alter these conclusions. Model performance remained dominated by baseline economic and historical energy features. In summary, in a macro-level, annual forecasting context with limited data, parsimonious tree-based ensembles offer the best, although still modest, predictive skill, while AI proxies contribute negligible additional value beyond traditional drivers of national energy consumption.

Table 11

Total fossil fuel consumption results

| Algorithm | Train | | | | Test | | | |
|--------------|------------|-------------|------------|----------|------------|-------------|------------|----------|
| | <i>mse</i> | <i>rmse</i> | <i>mae</i> | r^2 | <i>mse</i> | <i>rmse</i> | <i>mae</i> | r^2 |
| DT | 0.001 | 0.04 | 0.03 | 0.59 | 0.01 | 0.12 | 0.1 | 0.21 |
| RF | 0.002 | 0.05 | 0.03 | 0.33 | 0.01 | 0.11 | 0.08 | 0.28 |
| LR | 0.0002 | 0.01 | 0.01 | 0.95 | 0.05 | 0.23 | 0.19 | -1.99 |
| SVM (Linear) | 0.004 | 0.06 | 0.04 | -0.001 | 0.01 | 0.13 | 0.1 | -0.009 |

Table 12

Nuclear electric power consumption result

| Algorithm | Train | | | | Test | | | |
|--------------|------------|-------------|------------|---------|------------|-------------|------------|-----------|
| | <i>mse</i> | <i>rmse</i> | <i>mae</i> | r^2 | <i>mse</i> | <i>rmse</i> | <i>mae</i> | r^2 |
| DT | 0.0006 | 0.02 | 0.01 | 0.31 | 0.0008 | 0.028 | 0.025 | -1.009 |
| RF | 0.001 | 0.03 | 0.02 | -0.24 | 0.001 | 0.03 | 0.03 | -1.97 |
| LR | 0.0001 | 0.01 | 0.01 | 0.82 | 0.05 | 0.23 | 0.22 | -136.25 |
| SVM (Linear) | 0.0009 | 0.03 | 0.02 | 0.005 | 0.0003 | 0.01 | 0.01 | 0.008 |

Table 13

Total renewable energy consumption result

| Algorithm | Train | | | | Test | | | |
|--------------|------------|-------------|------------|-------|------------|-------------|------------|--------|
| | <i>mse</i> | <i>rmse</i> | <i>mae</i> | r^2 | <i>mse</i> | <i>rmse</i> | <i>mae</i> | r^2 |
| DT | 0.001 | 0.04 | 0.03 | 0.66 | 0.0006 | 0.02 | 0.01 | 0.89 |
| RF | 0.002 | 0.04 | 0.04 | 0.62 | 0.0005 | 0.02 | 0.02 | 0.91 |
| LR | 0.0004 | 0.02 | 0.01 | 0.92 | 0.11 | 0.33 | 0.28 | -17.15 |
| SVM (Linear) | 0.005 | 0.07 | 0.05 | 0.007 | 0.006 | 0.07 | 0.07 | 0.0005 |

Table 14

Total primary energy consumption result

| Algorithm | Train | | | | Test | | | |
|--------------|------------|-------------|------------|--------|------------|-------------|------------|--------|
| | <i>mse</i> | <i>rmse</i> | <i>mae</i> | r^2 | <i>mse</i> | <i>rmse</i> | <i>mae</i> | r^2 |
| DT | 0.001 | 0.04 | 0.03 | 0.55 | 0.01 | 0.1 | 0.09 | 0.23 |
| RF | 0.002 | 0.04 | 0.03 | 0.33 | 0.01 | 0.1 | 0.07 | 0.27 |
| LR | 0.0002 | 0.01 | 0.01 | 0.93 | 0.04 | 0.2 | 0.17 | -1.86 |
| SVM (Linear) | 0.003 | 0.06 | 0.04 | -0.002 | 0.01 | 0.12 | 0.09 | -0.007 |

To further investigate which features the most effective models utilized, we applied SHAP (Shapley Additive Explanations) analysis to the fitted models, enabling quantification of each feature's contribution to individual predictions. The SHAP explainability results, presented in Figures 2, 4, 6, and 8 for fossil fuels, nuclear energy, renewables, and total primary energy, respectively, reinforce previous conclusions regarding AI's limited role. In all cases, the most influential predictors of energy consumption are traditional socio-economic variables, while AI-related features rank near the bottom in importance.

For example, Figure 2 presents the SHAP summary for the fossil fuel consumption model. GDP per capita exhibits the largest positive SHAP values, indicating that increases in GDP per capita strongly drive predictions toward higher fossil energy use, consistent with the relationship between economic activity and energy demand. Similarly, variables such as population size and inflation—potentially proxying for energy prices or broader macroeconomic conditions—also demonstrate substantial SHAP contributions, underscoring their significance in shaping the model's output. In contrast, the ChatGPT dummy variable and technology stock indices (AMD, IBM, INTC) exhibit minimal SHAP impact, with their contributions on the SHAP importance plot either very small or nearly zero. This suggests that variations in these AI proxies do not meaningfully influence the predicted energy consumption.

In some instances, an AI feature's SHAP value fluctuates between slight positive and negative contributions across different years, showing no consistent effect and providing further evidence that the model identifies no clear signal from AI metrics. The same pattern is observed for other energy categories: as illustrated in Figure 4 (nuclear energy model), Figure 6 (renewables model), and Figure 8 (total primary energy model), AI-related features never emerge as primary drivers. Instead, the top SHAP-ranked features consistently correspond to macroeconomic or structural variables, such as GDP, population, or historical energy consumption levels. For instance, in the renewable energy consumption model, SHAP attributes most predictive power to factors like the renewable energy policy index and GDP, while an AI stock price contributes a nearly negligible SHAP value in any given year, effectively lost within the noise.

This consistent result from the SHAP analysis aligns closely with earlier statistical findings: the influence of AI proxies on energy consumption is marginal at best. Even when the models are capable of capturing complex nonlinear interactions, the data-driven interpretation is that AI features do not meaningfully affect national energy use compared with established determinants. The SHAP analysis thus provides an interpretable confirmation that the models appropriately focus on economic fundamentals, while the inclusion of AI variables has little to no effect on predicted outcomes.

Taken together, the empirical evidence presents a clear and consistent picture. AI, as measured by our proxies, does not exert a significant direct impact on national energy consumption during the analyzed period. No long-run equilibrium relationship was identified linking AI development to energy demand, and even in the short run, there is no strong linear or nonlinear association between the two. Predictive modeling further demonstrates that incorporating AI-related variables does not enhance forecast accuracy for aggregate energy consumption, and interpretability tools confirm that AI features are not among the primary drivers of energy use.

In practical terms, this suggests that rapid advances in AI and the growth of AI-centric firms have not yet translated into a measurable increase or decrease in a country’s overall energy requirements. Any energy footprint attributable to AI appears either too small to detect at the macro level or offset by efficiency gains and indirect effects. In summary, our results indicate that AI’s net impact on national energy consumption is effectively neutral; it neither significantly increases nor decreases energy demand at a scale observable in national statistics.




Figure 2
SHAP-based explanation of total fossil fuel consumption

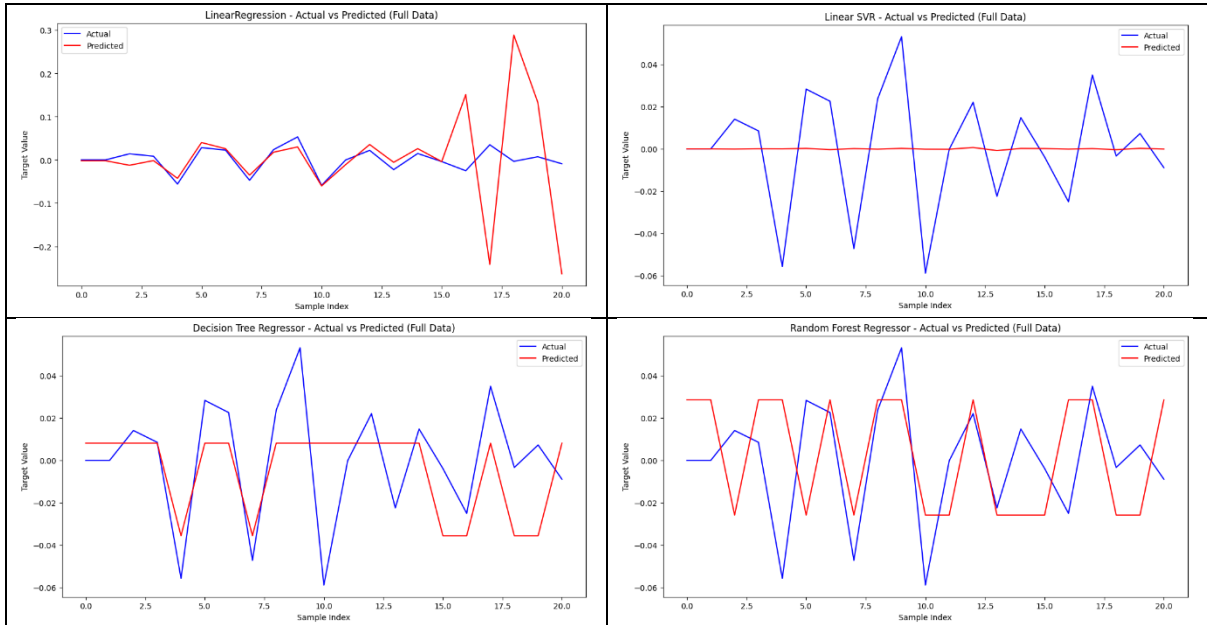


Figure 3
Nuclear electric power consumption predictions using different algorithms




Figure 4
SHAP-based explanation of nuclear electric power consumption

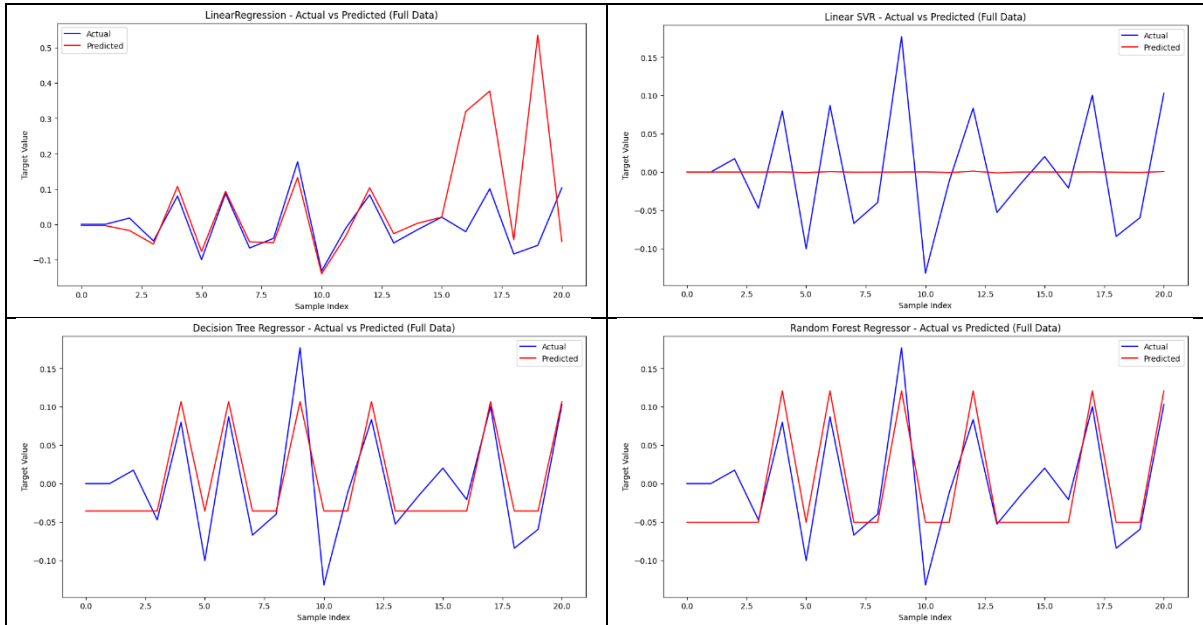


Figure 5

Total renewable energy consumption predictions using different algorithms

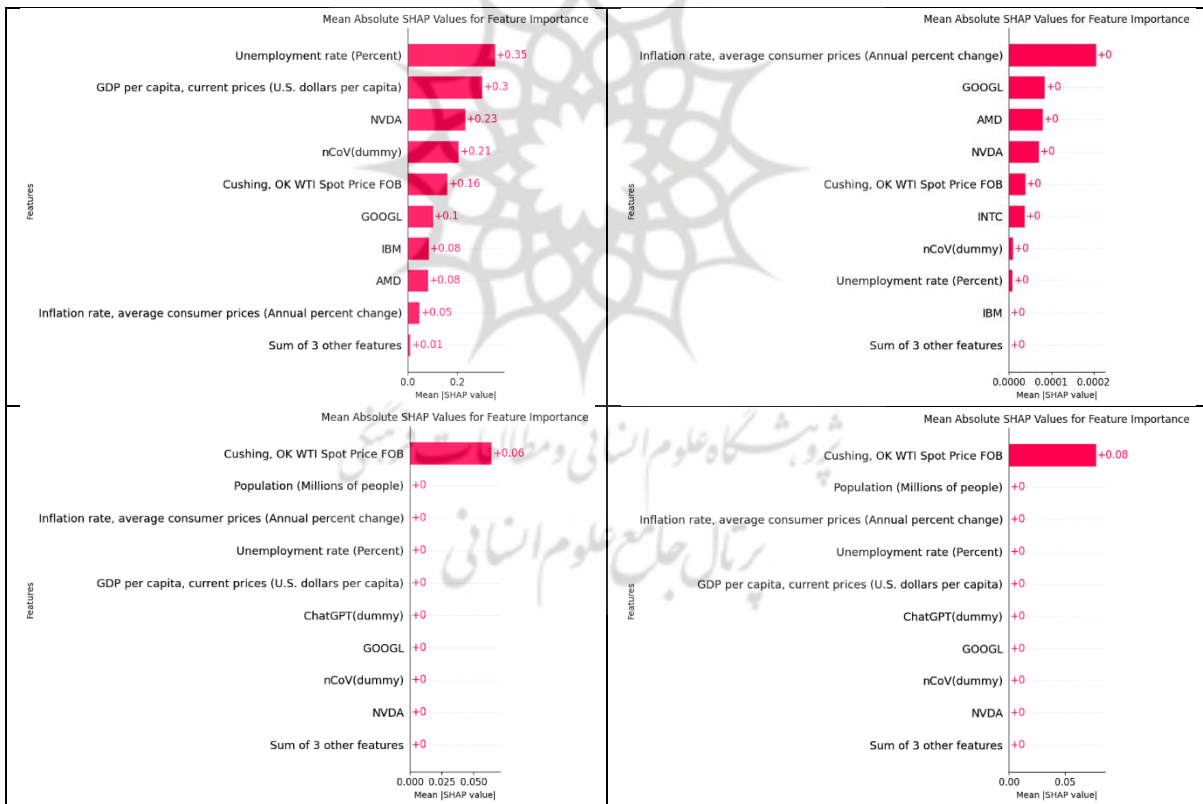


Figure 6

SHAP-based explanation of total renewable energy consumption




Figure 7

Total primary energy consumption predictions using different algorithms

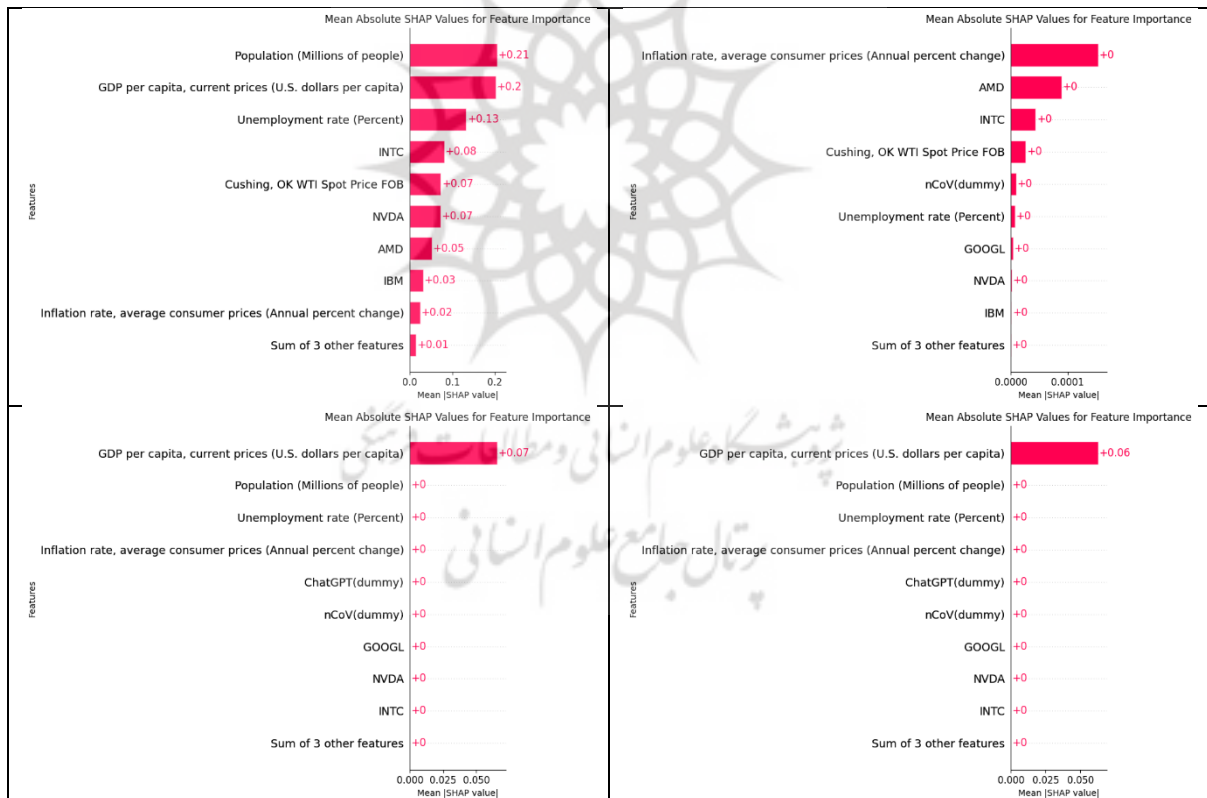


Figure 8

SHAP-based explanation of total primary energy consumption

The influence of AI on energy is largely absorbed by broader economic forces: AI-driven growth may enhance GDP or productivity, which in turn affects energy use, but AI itself does not impose an independent signature on energy consumption patterns. This finding is particularly relevant for policymakers and researchers, as it suggests that current concerns regarding AI’s energy footprint

should be contextualized within the broader framework of economic growth and energy efficiency. While AI undoubtedly consumes energy in absolute terms, its effect on a country's total energy demand has been empirically indistinguishable from zero or mediated through other variables. Therefore, based on data through 2024, we conclude that AI's emergence has a negligible direct impact on national energy consumption.

5. Conclusions

This paper aimed to address a straightforward yet timely question: has the dramatic expansion of artificial intelligence activity in the United States altered the trajectory of national energy demand? Drawing on nearly four decades of annual data and employing a suite of complementary econometric and machine-learning tools, the evidence points to a clear and consistent answer—at the aggregate level, AI has left no discernible imprint on U.S. energy consumption.

The absence of any long-run cointegrating relationship between our AI proxies and the major energy series indicates that the rapid rise of AI does not anchor energy demand to a new equilibrium path in the way that core macroeconomic drivers, such as GDP or inflation, do. After removing common growth trends, year-to-year fluctuations in AI activity remain statistically orthogonal to changes in fossil-fuel, nuclear, renewable, and total primary energy use; neither linear correlations nor more flexible metrics capturing temporal patterns or mutual information reveal meaningful comovement. Predictive exercises reinforce this neutrality. A parsimonious error-correction model, grounded in conventional macroeconomic fundamentals, consistently explains nearly 90 percent of out-of-sample variation in energy consumption, and the inclusion of AI variables does not improve this accuracy. Interpretable SHAP diagnostics corroborate the forecast results: within the models, AI proxies carry negligible explanatory weight compared with long-established economic determinants.

Taken together, these findings imply that the widely discussed energy footprint of AI remains too small—or too diffuse—to shift national-level consumption patterns; any demand generated by AI appears absorbed within the broader macroeconomic system rather than manifesting as an independent driver. For policymakers engaged in long-term energy planning or decarbonization scenarios, the results suggest that, at present, AI does not require a dedicated demand trajectory in national forecasts. More immediate pressures are likely to arise in specific subsectors—most notably power-intensive data centers—rather than in the aggregate statistics analyzed here. Future research leveraging higher-frequency or facility-level data may be necessary to capture these localized stresses or to detect threshold effects if AI adoption accelerates further.

In short, despite AI's highly visible growth and its undeniable electricity requirements at the micro level, its macro-level influence on U.S. energy consumption remains empirically indistinguishable from zero over the period 2004–2024. National energy demand continues to be overwhelmingly governed by familiar fundamentals such as population, income, prices, and technology-wide efficiency trends, with AI's contribution—in current data—effectively relegated to statistical noise.

Nomenclature

| | |
|-----------------|---|
| ADF | Augmented Dickey–Fuller test statistic for stationarity |
| AMD | Advanced Micro Devices, Inc. stock price |
| C1–C17 | Column indices in Table 1 corresponding to variables described in the dataset |
| ChatGPT (dummy) | Dummy variable for the introduction of ChatGPT (1 = after Nov 2022, 0 = before) |

| | |
|------------------------------------|---|
| DT | Decision tree regression model |
| DTW | Dynamic time warping distance (temporal similarity metric) |
| ECM | Error-correction model (used when cointegration exists) |
| Energy series (general) | Fossil, nuclear, renewable, and total primary energy categories |
| GDP per capita | Gross domestic product per capita, current prices (USD per person) |
| GOOGL | Alphabet Inc. (Google) stock price |
| IBM | International Business Machines Corporation stock price |
| Inflation rate | Annual percentage change in CPI (average consumer prices) |
| INTC | Intel Corporation stock price |
| LR | Linear regression model |
| MAE | Mean absolute error (forecast accuracy metric) |
| MI | Mutual information (nonlinear dependence measure, bits) |
| nCoV (dummy) | COVID-19 pandemic dummy (1 = after March 2019, 0 = before) |
| Nuclear electric power consumption | Total nuclear electricity generation (billion kWh) |
| NVDA | NVIDIA Corporation stock price |
| OLS | Ordinary least squares regression |
| Population | Total mid-year population (millions) |
| R ² | Coefficient of determination (model fit indicator) |
| RF | Random forest regression model |
| RMSE | Root mean squared error (forecast accuracy metric) |
| SHAP | Shapley additive explanations (feature importance measure) |
| SVM (Linear) | Support vector machine with linear kernel |
| Total fossil fuels consumption | Annual consumption of coal, oil, and natural gas (quadrillion BTU) |
| Total primary energy consumption | Total primary energy consumption (quadrillion BTU) |
| Total renewable energy consumption | Total annual renewables consumption (hydro, wind, solar, etc.) (quadrillion BTU) |
| Unemployment rate | Percentage of labor force unemployed |
| VIF | Variance inflation factor for multicollinearity assessment |
| WTI | West Texas Intermediate crude oil spot price, Cushing, Oklahoma (USD per barrel, FOB) |

References

- Ahmad, T., Chen, H., & Shah, W. A. (2019). Effective bulk energy consumption control and management for power utilities using artificial intelligence techniques under conventional and renewable energy resources. *International Journal of Electrical Power & Energy Systems*, 109, 242–258.

- Bogmans, C., Ganpurev, G., Gomez-Gonzalez, P., Melina, G., Pescatori, A., & Thube, S. (2025). *Power hungry: How AI will drive energy demand* (IMF Working Paper No. WP/25/81). International Monetary Fund.
- Burian, V., & Stalla-Bourdillon, A. (2025). The increasing energy demand of artificial intelligence and its impact on commodity prices. *ECB Economic Bulletin*, 2(2025).
- Chen, C., Hu, Y., Karuppiah, M., & Kumar, P. M. (2021). Artificial intelligence on economic evaluation of energy efficiency and renewable energy technologies. *Sustainable Energy Technologies and Assessments*, 47, 101358.
- Dauvergne, P. (2022). Is artificial intelligence greening global supply chains? Exposing the political economy of environmental costs. *Review of International Political Economy*, 29(3), 696–718.
- Engle, R. F., & Granger, C. W. (1987). Co-integration and error correction: Representation, estimation, and testing. *Econometrica*, 55(2), 251–276.
- Fang, Y., Lee, C.-C., & Li, X. (2025). Assessing the impact of artificial intelligence on the transition to renewable energy: Analysis of U.S. states under policy uncertainty. *Renewable Energy*, 246, 122969.
- Himeur, Y., Ghanem, K., Alsalemi, A., Bensaali, F., & Amira, A. (2021). Artificial intelligence-based anomaly detection of energy consumption in buildings: A review, current trends and new perspectives. *Applied Energy*, 287, 116601.
- Lee, C.-C., Zou, J., & Chen, P.-F. (2025). The impact of artificial intelligence on the energy consumption of corporations: The role of human capital. *Energy Economics*, 143, 108222.
- Liu, X., Cifuentes-Faura, J., Zhao, S., Wang, L., & Yao, J. (2025). Impact of artificial intelligence technology applications on corporate energy consumption intensity. *Gondwana Research*, 138, 89–103.
- Müller, M. (2007). Dynamic time warping. In *Information retrieval for music and motion* (pp. 69–84). Springer.
- Tao, S., Wang, Y., & Zhai, Y. (2023). Can the application of artificial intelligence in industry cut China's industrial carbon intensity? *Environmental Science and Pollution Research*, 30(33), 79571–79586.
- Thompson, C. G., Kim, R. S., Aloe, A. M., & Becker, B. J. (2017). Extracting the variance inflation factor and other multicollinearity diagnostics from typical regression results. *Basic and Applied Social Psychology*, 39(2), 81–90.
- Veyrat-Charvillon, N., & Standaert, F.-X. (2009, September). Mutual information analysis: How, when and why? In *International Workshop on Cryptographic Hardware and Embedded Systems* (pp. 429–443). Springer.
- Wang, Q., Li, Y., & Li, R. (2024). Ecological footprints, carbon emissions, and energy transitions: The impact of artificial intelligence (AI). *Humanities and Social Sciences Communications*, 11, Article 1043.
- Yunyun, F., Yongchang, S., Malin, S., & Weiyu, W. (2024). Does artificial intelligence reduce corporate energy consumption? New evidence from China. *Economic Analysis and Policy*, 83, 548–561.

Zhao, P., Gao, Y., & Sun, X. (2022). How does artificial intelligence affect green economic growth? Evidence from China. *Science of the Total Environment*, 834, 155306.



COPYRIGHTS

©2025 by the authors. Published by Petroleum University of Technology. This article is an open-access article distributed under the terms and conditions of the Creative Commons Attribution 4.0 International. (CC BY 4.0) (<https://creativecommons.org/licenses/by/4.0/>)

