

The Hidden Environmental Footprints of Large-Scale AI Training: Evidence from Google’s GLaM, PaLM, and Gemini 1.0 Models

Jun Ho Choi

PhD in Sustainable Development, Columbia University

January 2026*

This thesis is submitted in partial fulfillment of the requirements for the Master of Arts degree in Sustainable Development, which forms part of the PhD in Sustainable Development at Columbia University.

Abstract

This study provides new empirical evidence on the environmental impact of large-scale artificial intelligence (AI) model training. I examine two major training events: Google’s GLaM and PaLM models (trained in 2021–2022 at Mayes County, Oklahoma) and Gemini 1.0 (partially trained in 2023 at Council Bluffs, Iowa). Using hourly Continuous Emissions Monitoring Systems (CEMS) data, I identify fossil-fuel power plants that are electrically proximate to these data centers using a principal component analysis (PCA) of congestion prices, following [Godin and Ibrahim \(2021\)](#). I then employ a synthetic difference-in-differences (SDID) framework to quantify the causal impact of training-induced load shocks on local generation. The results reveal that AI training creates a sustained baseload demand that is primarily met by local coal-fired generation. Electrically proximate coal plants exhibited massive surges in output, with average CO₂ emissions rising by approximately 100 to 160 kilotons per plant per month during peak training periods relative to counterfactuals. In contrast, natural gas response was muted, primarily serving to flatten diurnal demand gaps. I estimate that these two events alone generated between 0.8 and 1.6 million tons of excess CO₂ emissions, corresponding to monetized social damages of \$150 to \$300 million (using the EPA 2023 social cost of carbon). These findings underscore that without coordination between compute scheduling and grid dynamics, AI training can significantly intensify fossil-fuel dependence even in regions with high renewable penetration.

*I am deeply grateful to my advisor, Professor Douglas Almond, for his guidance, thoughtful feedback, and encouragement throughout this project. I also thank Professor John Mutter, Director of Ph.D. in Sustainable Development, for his support and careful review. I am indebted to my fellow Sustainable Development students and to program coordinator Tamara Aldrich for their kind assistance and encouragement along the way. I am also thankful to Professor Michael Best, during whose course I first conceptualized this project, and to participants of the CU ERE Workshop, whose insightful comments helped strengthen this work. All remaining errors are my own.

1 Introduction

The rise of large-scale artificial intelligence (AI) systems has introduced a new class of electricity demand: high-intensity, episodic training workloads concentrated in specific data centers. These bursts of computation can impose short-run shocks on regional power grids, yet there is limited empirical evidence on their environmental consequences. Understanding how this digital innovation interacts with physical power infrastructure has become increasingly central to environmental and energy economics, as it links two major policy goals: technological progress and the mitigation of climate change.

Recent evidence underscores the scale of this challenge. In 2023, U.S. data centers consumed roughly 176 terawatt-hours (TWh) of electricity – about 4.4 percent of total national demand – up from 1.8 percent in 2014. Projections by researchers at Lawrence Berkeley National Laboratory and the U.S. Department of Energy suggest that, under current trends, data-center electricity use could double or triple by 2028, reaching 6.7 to 12 percent of U.S. demand (Shehabi et al., 2024). The rapid rise of compute-intensive AI workloads is a key driver of this trajectory.

The environmental footprint of AI has recently gained wide attention in both academic and policy circles. Studies and industry analyses have estimated that the energy intensity of large-scale model training may rival that of major industrial processes (Bashir et al., 2025; CarbonBrief, 2025). Related work has examined the energy and water implications of AI operations (E360, 2024), the carbon cost of inference workloads (Google Cloud, 2025b), and forecasts of data-center CO₂ emissions under alternative technological scenarios (Guidi et al., 2024). Most recently, Bonfiglioli et al. (2025) exploit a shift-share design to show that U.S. commuting zones with higher AI penetration experience faster emissions growth, driven by scale effects and data center proximity. However, while such aggregate studies establish a broad link between AI and carbon intensity, there remains a lack of empirical evidence tracing how specific AI training events alter real-world electricity generation in the short run. This paper fills that gap using high-frequency, plant-level data.

This study focuses on Google’s GLaM and PaLM models, suspected to be trained from November 2021 to February 2022, and Gemini 1.0, trained primarily from May to August 2023. These training runs likely occurred at the company’s Mayes County, Oklahoma and Council Bluffs, Iowa data centers, respectively. While Google reports that these facilities were matched with “carbon-free energy” on an annual basis, such accounting reflects contractual matching rather than physical grid flows. In wholesale markets, electricity from clean sources cannot be physically dedicated to a single consumer. When these facilities increase their power demands to train large models, the incremental load must be met by the marginal generating units available within the Southwest Power Pool (SPP) or Midcontinent Independent System Operator (MISO) networks. Given the limited

short-run substitutability of renewable sources, such incremental demand often necessitates ramping up fossil-fuel generation. This mechanism implies that AI training can increase system-wide carbon intensity even when a data center’s reported consumption is nominally carbon-free.

Empirically, I test this hypothesis using a synthetic difference-in-differences (SDID) framework. This approach constructs a weighted control group of power plants that matches the pre-treatment trends of plants affected by the data center’s load. I compare fossil-fuel power plants in the same electrical region that are electrically proximate to the data center with those that are electrically distant, before and during the training periods. Electrical proximity is measured as the pre-training correlation in the congestion components of locational marginal prices (LMPs) between each generation node and the data center’s load zone. This measure captures the extent to which plants co-move in response to shared grid constraints, identifying those most likely to respond to training-induced load shocks.

For the Mayes County data center, I examine the effect of training for PaLM and GLaM. The results indicate a sharp discontinuity in emissions patterns beginning in late 2021. During the significant training months (November 2021 through February 2022), monthly CO₂ emissions at electrically proximate coal power plants increased by approximately 113 to 151 kilotons per plant relative to the pre-treatment baseline. The estimated total excess CO₂ attributable to the statistically significant months of this training cycle is approximately 0.8 million metric tons.

For the Council Bluffs data center, I study the effect of the Gemini 1.0 training cycle. During June and July 2023, monthly CO₂ emissions at electrically proximate coal power plants increased by 105 to 160 kilotons per plant. In contrast, electrically proximate natural gas plants did not exhibit higher aggregate emissions, though within-day analysis reveals evidence of altered dispatch patterns, with gas units reducing their diurnal peakiness to support overnight load. Combining both events, I estimate a total excess carbon footprint of roughly 0.8 to 1.6 million metric tons, translating to social damages of \$150 to \$300 million.

Translating these physical externalities into monetary terms using a social cost of carbon of \$51 per ton ([Interagency Working Group on the Social Cost of Greenhouse Gases, 2021](#)) implies damages exceeding \$86 million, which is more than double the amount Google pledged for carbon removal in similar periods ([Spock, 2024](#)). These estimates likely represent a lower bound. First, similar increases are observed for local criteria pollutants, including SO₂ and NO_x. Second, unless Google massively expanded capacity undisclosed to the public, non-AI compute loads were likely displaced to other regions, leaking emissions elsewhere. Finally, the identification strategy relies on clean pre-treatment periods, yet the Mayes County facility likely saw continuous training activity, potentially attenuating the estimated treatment effect.

Taken together, the results highlight the need for dynamic coordination between AI compute scheduling and renewable generation availability to ensure that digital progress does not undermine

decarbonization goals. By linking large-scale AI training to observable changes in power-plant emissions, this study contributes to three strands of literature: (i) the economics of energy demand from digital infrastructure (Masanet et al., 2020; Shehabi et al., 2024); (ii) the environmental consequences of technological innovation (Acemoglu et al., 2012); and (iii) the empirical analysis of networked externalities in energy systems (de Sousa Costa and Ferreira, 2023; Davis et al., 2023). More broadly, this paper adds to emerging work on the environmental footprint of AI itself, offering the first plant-level evidence on how large-scale model training reshapes grid emissions and fossil-fuel generation patterns (Bashir et al., 2025; Han et al., 2024; Guidi et al., 2024).

Looking ahead, the environmental implications of these findings are likely to intensify. The training of increasingly complex models appears inevitable as firms compete to resolve current limitations, such as hallucinations, and expand multimodal capabilities (OpenAI, 2023). As model architectures scale, the energy required for training may grow non-linearly. Without a fundamental breakthrough in clean, firm power generation, or a shift away from the current paradigm of massive, centralized training, this escalation in demand will likely continue to be met by the existing fossil-heavy grid. The results of this study suggest that relying on the grid “as is” to fuel this competition is not merely inefficient, but actively deepens fossil-fuel dependence by increasing the utilization of legacy coal and gas infrastructure.

The remainder of the paper proceeds as follows. Section 2 describes the institutional background and data sources. Section 3 elaborates on the construction of electrical proximity measures based on locational marginal prices’ congestion components and principal component analysis. Section 4.1 outlines the empirical methodology, specifically the synthetic difference-in-differences strategy. Section 5 presents the main estimates regarding emissions and peakiness. Section 6 concludes and discusses implications for policy and future research.

2 Background and Data

2.1 Institutional Setting

Google operates one of the largest networks of hyperscale data centers worldwide (Datacenters.com, 2025). Table 1 presents the major U.S. facilities relevant to this study. Among these, the Council Bluffs, Iowa facility is one of the company’s oldest and most capacity-intensive sites. It is strategically located within the Midcontinent Independent System Operator (MISO) grid and recorded the highest single-site water use in the fleet, approximately 3.7 million m³ in 2023 (Google, 2023c). The Mayes County, Oklahoma data center, located within the Southwest Power Pool (SPP), follows closely with the second-largest water consumption of 3.1 million m³. This facility has been publicly associated with the deployment of Tensor Processing Units (TPUs) v4, Google’s proprietary AI-

Table 1: Major U.S. Google Data Centers and Indicators of AI Training Activity

City / County	State	WU	Primary Evidence	Notes / Relevance
Council Bluffs	IA	3.7	High reported site-level electricity and water demand; historical TPU deployments	Among the earliest and largest Google data centers; highest known water use; clear traceability to MISO generation nodes
Mayes County	OK	3.1	Public documentation of TPU (v4/v5) deployment and AI workloads	Known TPU cluster; likely secondary training or validation site for Gemini 1.0
Douglas County	GA	1.9	Primarily cloud-service operations	Large but not publicly linked to AI-training workloads
Berkeley County	SC	2.5	Mixed-use (Search, YouTube)	High water use but limited evidence of large-scale training activity
The Dalles	OR	1.6	Expansion tied to cloud storage and future AI workloads	Renewable-heavy grid; more relevant for inference or storage, not training
Lenoir	NC	1.0	Legacy infrastructure	Smaller, older site; no TPU or AI-specific disclosure

NOTE: Water consumption estimates (converted to millions of m³ from reported US gallons) for 2023 are derived from [Google \(2024\)](#); these values reflect water consumed rather than total withdrawal. Grid-region assignments are based on balancing-authority boundaries from FERC and ISO maps. TPU deployment information is drawn from public hardware announcements and press coverage ([Gupta and Ardito, 2022](#); [Patel et al., 2023](#); [Google Cloud, 2025a](#)). Council Bluffs, IA and Mayes County, OK are boldfaced as likely training grounds for Gemini 1.0 and GLaM/PaLM, respectively.

training hardware ([Morgan, 2022](#)). Together, these two facilities serve as the primary infrastructure for the large-scale model training operations analyzed in this paper.

This study utilizes both facilities to capture the environmental externalities of AI training across distinct grid topologies and model generations, with careful selection to ensure valid identification. The Mayes County site is analyzed as the training ground for the GLaM and PaLM models. While this facility is also suspected to have hosted training runs for Gemini 1.0, I restrict the analysis of the Gemini model to the Council Bluffs location. This exclusion is necessary because the continuous training activity at Mayes County following the deployment of PaLM creates a contaminated baseline, rendering the identification of a clean pre-treatment period for Gemini 1.0 infeasible in the SPP region. Consequently, the Council Bluffs facility provides the primary, uncontaminated setting for isolating the impact of the Gemini 1.0 training cycle.

The Mayes County, Oklahoma data center appears to have served as the primary training ground for Google’s GLaM and PaLM models between late 2021 and early 2022. Google released the GLaM model architecture in December 2021 ([Du et al., 2022](#)), followed by the PaLM model announcement in April 2022 ([Chowdhery et al., 2023](#)). In May 2022, shortly after the PaLM announcement, Google Cloud unveiled that its Oklahoma facility housed the “world’s largest publicly available ML cluster,”

comprising 9,000 TPU v4 chips (Gupta and Ardito, 2022). Given that PaLM was trained on 6,144 TPU v4 chips, a hardware scale available nowhere else at the time (Chowdhery et al., 2023), it is highly probable that the model’s training runs occurred at this location during the preceding months (November 2021 through February 2022). This timeline aligns with the facility’s observed water consumption spikes and serves as the treatment period for the Oklahoma analysis.

Evidence from company statements and contemporaneous media reporting suggests that Gemini 1.0 was trained between January and September 2023. By January 2023, Google leadership declared a “code red” following the emergence of OpenAI’s GPT models, and DeepMind CEO Demis Hassabis indicated that a new conversational model would be released later that year (Cuthbertson, 2023; Grant, 2023). In April 2023, Google and DeepMind announced their formal merger as “Google DeepMind,” unifying research and compute resources (Google DeepMind, 2023; Google, 2023a). In May 2023, Google confirmed that Gemini “was still in training” (Google, 2023e), and by August–September internal testing and commercialization planning (e.g., paywalls being set up, beta-testers being recruited) signaled that the model had entered its fine-tuning phase (UC Today, 2023; Victor, 2023).

Because detailed facility-level power data for Google are unavailable, this study turns to grid-level evidence from the MISO node MEC.MECB, which is the load-zone node managed by MidAmerican Energy Company (MEC) that includes the Council Bluffs area.¹ Load increases beginning in early to mid-2023 coincide closely with the documented Gemini 1.0 training period. Public sources and industry reporting describe Microsoft’s continued data-center construction in West Des Moines as a major future source of load growth for MEC, but most of these campuses remained under construction and only partially energized by late 2023 (Business Record, 2023; Weitz Company, 2023; Microsoft Local, 2023). They are therefore unlikely to explain the sharp rise in emissions observed in mid-2023. The timing instead points to intensified utilization of existing hyperscale capacity, most plausibly Google’s Council Bluffs facility, which has long operated within MISO and was reported to have exceptionally high water and energy use in 2023 (Google, 2023c). Other contributors (e.g., increases in cryptocurrency mining and industrial construction loads) were comparatively small or localized. Hence, the most consistent interpretation is that the mid-2023 load uptick in MEC.MECB reflects a temporary but power-intensive surge associated with Google’s AI-training operations rather than new physical capacity coming online. An overview of potential sources of load growth in the MEC territory in 2023 is provided in Table 9.

Similarly, facility-level data for the Mayes County data center are seemingly confidential, requiring reliance on regional load references within the Southwest Power Pool (SPP), specifically the Grand River Dam Authority settlement location GRDA_HUB and PNode (price node) GRDA.H. This node is connected to the MidAmerica Industrial Park, where Google is the largest tenant,

¹For more information on other MISO load zones under MEC, see Appendix Section A.1 and Table 8.

and the load patterns observed during late 2021 and early 2022 correspond with the suspected training windows for GLaM and PaLM. While the industrial park attracted other energy-intensive occupants around this time, their operational timelines do not align with the observed winter 2021 surge. Most notably, Northern Data, a large-scale cryptocurrency mining firm, did not announce its partnership with GRDA until March 2022, with initial 10 MW operations commencing only in the second quarter of that year ([Oklahoma Department of Commerce, 2022](#); [Northern Data AG, 2022](#)). Likewise, the electric vehicle manufacturer Canoo was in the early planning phases, with major manufacturing announcements and battery facility construction occurring in late 2022 ([Canoo Inc., 2022](#)). Other legacy industrial tenants, such as chemical and aerospace manufacturing facilities, exhibit stable demand profiles inconsistent with the rapid, episodic spikes associated with AI training. Thus, the most plausible driver of the load anomalies in the GRDA nodes during this period is the pre-announcement energization and testing of Google’s TPU v4 cluster, which was publicly unveiled in May 2022 ([Gupta and Ardito, 2022](#)). An overview of potential load sources in the GRDA territory is provided in Table 10.

2.2 Data

The empirical analysis combines multiple high-frequency and administrative datasets to capture plant-level emissions, market conditions, and local environmental drivers. Core data sources include the U.S. EPA’s Continuous Emissions Monitoring System (CEMS) and MISO’s hourly locational marginal price (LMP) reports. These are supplemented by ERA5 meteorological reanalysis (accessed via Open-Meteo), county-level socioeconomic indicators from the Bureau of Labor Statistics (BLS) and Bureau of Economic Analysis (BEA), and facility-level disclosures from Google and federal filings.

2.2.1 Continuous Emissions Monitoring System (CEMS)

CEMS reports hourly stack emissions and net generation for all large fossil-fuel power plants in the United States. I use data for all facilities within the Midcontinent Independent System Operator (MISO) and Southwest Power Pool (SPP) balancing authority between 2020 and 2024, extracting hourly CO₂, SO₂, and NO_x emissions, hourly heat input, gross generation (i.e., gross load), and operational status. Observations with missing generation values or flagged quality codes are excluded. The original datasets are published by the U.S. Environmental Protection Agency, but I access them through the Public Utility Data Liberation (PUDL) project, which standardizes and harmonizes the EPA’s Continuous Emissions Monitoring System (CEMS) and other datasets like the U.S. Energy Information Administration’s Form 860 and Form 923 ([U.S. Environmental Protection Agency, 2024](#); [Catalyst Cooperative, 2024](#)). Emissions are aggregated to the plant-day level

for the main analysis, with hourly and weekly aggregations used in robustness checks.

2.2.2 MISO and SPP Market Data on Locational Marginal Prices

The analysis utilizes hourly real-time and day-ahead locational marginal prices (LMPs) published by the Midcontinent Independent System Operator (MISO) and the Southwest Power Pool (SPP). For the Iowa analysis, the focal point is the MEC.MECB node, which represents the aggregated load zone for the MidAmerican Energy Company and encompasses the Council Bluffs data center. Analogously, for the Oklahoma analysis, I identify the relevant load point using the Settlement Location-PNode tuple (GRDA_HUB, GRDA_H). This specific identifier captures the pricing dynamics for the Grand River Dam Authority, which serves the Mayes County industrial complex. Data from these nodes provide the congestion price signals necessary to construct the electrical proximity measures described in Section 3.

Economically, LMP represents the optimal spot price of electricity at a specific node. At the market equilibrium, this price equals to the marginal social cost of providing an additional unit (MW) of electricity, adjusted for location-specific transmission losses and congestion. In the canonical optimal power flow formulation of [Bohn et al. \(1984\)](#), the system operator maximizes total consumer value net of generation costs, subject to constraints on energy balance, unit capacities, and transmission limits. Solving the Lagrangian yields the following decomposition of the optimal LMP at node i :

$$\begin{aligned} \text{optimal LMP at node } i = & [\text{social cost of additional demand at } i] \\ & \times [1 + \text{incremental losses caused by } i] \\ & + [\text{transmission constraint terms, over all transmission lines}] \end{aligned}$$

where the term

$$[\text{social cost of additional demand}] \times [\text{incremental losses caused by } i]$$

is referred to as marginal loss component (MLC), capturing incremental losses associated with delivering power to node i , and

$$[\text{transmission constraint terms, over all transmission lines}]$$

is referred to as marginal congestion components (MCC), which reflects the cost of transmission constraints binding across the network.

MISO and SPP both report the marginal congestion component (MCC), marginal loss compo-

ment (MLC), and overall LMP at an hourly frequency. For this study, the MCC series preceding the respective AI training events are of particular importance, as they serve as the basis for computing the electrical proximity measures described in Section 3. Specifically, I utilize the MCC data prior to the Gemini 1.0 training period for the Council Bluffs analysis and the data prior to the GLaM and PaLM training cycles for the Oklahoma analysis.

2.2.3 Weather and Economic Controls

Daily meteorological variables, including surface temperature, wind speed, and precipitation, are obtained from the ERA5 reanalysis product (Hersbach et al., 2020) via the Open-Meteo historical weather API (Open-Meteo, 2024). Each variable is spatially matched to the ERA5 grid cell nearest to the geographic coordinates of each generating unit. In the empirical analysis, I apply B-spline transformations to surface temperature and precipitation to account for the non-linear relationship between weather conditions and electricity generation. Additionally, I construct cooling degree-day (CDD) and heating degree-day (HDD) indices aggregated at the load-zone level. These metrics are included to control for temperature-driven shocks to regional electricity demand, as well as to account for the thermodynamic sensitivity of data center cooling systems, whose power consumption increases during extreme heat events due to degrading power usage effectiveness (PUE).

County-level economic controls are constructed from the Bureau of Labor Statistics’ Quarterly Census of Employment and Wages (QCEW) and the Bureau of Economic Analysis’ Local Area Personal Income (LAPI) datasets, merged to plants by ZIP code. These variables capture local labor-market and income conditions that may drive regional differences in electricity demand or emissions. In robustness checks, I also incorporate county-level unemployment rates from the Federal Reserve Bank of St. Louis’ Federal Reserve Economic Data (FRED) (Federal Reserve Bank of St. Louis, 2024), which serve as an additional proxy for short-term economic activity.

2.2.4 Administrative and Supporting Data

Supplementary data on data-center characteristics and resource use are drawn from Google’s annual Sustainability Reports and state environmental filings. A physics-based measure of electrical proximity would be possible if one were to use transmission and Power Transfer Distribution Factor (PTDF) information from FERC Form 715. However, my previous request for access to these materials under FERC’s Critical Energy/Electric Infrastructure Information (CEII) program was denied. As a result, future work aims to incorporate open grid-topology and flowgate data from the GridStatus.io dataset as a potential alternative. These efforts will allow for the construction of an engineering-grounded proximity metric to complement the principal components analysis- and congestion-driven approach employed in this paper.

3 Electrical Proximity from LMP Congestion Components

3.1 Background on Electrical Proximity

To quantify how individual power plants (or generator nodes, hereafter “gennodes”) respond to demand fluctuations in the load zone of interest, specifically the MEC.MECB node, I follow the approach of [Godin and Ibrahim \(2021\)](#) to apply principal component analysis (PCA) to the marginal congestion components (MCC) of locational marginal prices (LMPs). This method identifies latent dimensions of congestion behavior across the transmission network, and the cosine similarity between loading vectors of these latent factors serves as an indicator of electrical proximity. In essence, this approach infers the extent to which how two nodes experience similar transmission constraints, thereby capturing how closely they are connected in terms of congestion dynamics.

Below, I explain why, although this method is not the ideal way to quantify electrical proximity, it remains the most practical under realistic data limitations when only MCCs are available. The PCA-based procedure itself is detailed in Subsection [3.2](#).

In principle, the ideal measure of how gennodes react to demand fluctuations at a given node would rely on power transfer distribution factors (PTDFs), which describe the physical sensitivities of transmission-line flows to injections and withdrawals at specific nodes or buses. Following [Wood et al. \(2014\)](#), the PTDF along transmission line ℓ between injector node i and recipient node j is defined as

$$\text{PTDF}_{i,j,\ell} = \frac{\Delta f_\ell}{\Delta P_{i \rightarrow j}}$$

where $\Delta P_{i \rightarrow j}$ denotes the change in power transferred from node i to node j , and Δf_ℓ represents the change in flow on line ℓ resulting from that transfer. In words, $\text{PTDF}_{i,j,\ell}$ represents the sensitivity of the flow on line ℓ to a shift of power from node i to node j ([Wood et al., 2014](#)). Aggregating these sensitivities across all lines connecting i and j yields the overall responsiveness of power flow between them. This overall sensitivity between two nodes can be viewed as the “first-best” measure of electrical proximity.

In practice, however, it is rarely feasible to obtain detailed PTDF data, as these depend on confidential physical parameters of the transmission grid such as resistance and impedance. Although access may be requested through Critical Electricity/Energy Infrastructure Information (CEII) channels, such requests are often denied citing security reasons, as was the case for my own application. In such cases, a “second-best” alternative relies on injection shift factors (ISFs), where

ISF of a node i along line ℓ is defined and empirically derived as:

$$\text{ISF}_{i,\ell} = \frac{\partial P_\ell}{\partial P_i} \approx \frac{\Delta P_\ell^i}{\Delta P_i}$$

where P_i is the injection (or retrieval) of power at node i , P_ℓ is the transmission of power along line ℓ , and ΔP_ℓ^i is the change in active power along line ℓ due to ΔP_i . In words, ISF measures how much active power on line ℓ changes in response to an incremental injection at node i (Chen et al., 2014). Active power here refers to the portion of electricity that performs useful work, to be distinguished from reactive power that oscillates or dissipates in transmission.

ISFs and PTDFs differ in two main ways. First, PTDFs are defined over a pair of nodes or buses (an injector and a recipient), whereas ISFs are defined over just a single node with respect to a transmission line. In addition, while PTDFs apply to transmission lines that are directly connecting those nodes, ISFs can be calculated for any transmission line, even for those that are not directly adjacent to the node. Despite these distinctions, the two measures are conceptually similar. For instance, in a simple two-node system where line ℓ connects nodes i and j directly, $\text{PTDF}_{i,j,\ell}$, $\text{ISF}_{i,\ell}$, and $\text{ISF}_{j,\ell}$ would all be equivalent. ISFs can therefore be regarded as generalizations of PTDFs that remain valid under complexities of transmission grids.

ISFs are also closely related to the construction of MCCs. Following Bohn et al. (1984), the MCC for node j at time t can be expressed as:

$$\text{MCC}_j(t) = [\text{transmission constraint terms, over all transmission lines}] \equiv \sum_\ell \frac{\partial Z_\ell(t)}{\partial D_j(t)} \eta_\ell(t)$$

where $\eta_\ell(t)$ is the shadow value (or shadow price) of additional transmission capacity of line ℓ at time t , $D_j(t)$ is the electricity demanded by node j at time t , and $Z_\ell(t)$ is the transmitted electricity along line ℓ . The terms $\frac{\partial Z_\ell(t)}{\partial D_j(t)}$ are directly analogous to ISFs, implying that each node's MCC reflects a weighted sum of ISFs, with the weights determined by the corresponding shadow prices. A similar expression applies for supplier nodes i :

$$\text{MCC}_i(t) \equiv \sum_\ell \frac{\partial Z_\ell(t)}{\partial S_i(t)} \eta_\ell(t)$$

where $S_i(t)$ is the electricity supplied by node i at time t .²

Given ISFs, one may calculate the electrical proximity between nodes i and j as the cosine similarity between their respective ISF vectors. That is, denoting the vector of ISFs for node k as

²While the signs of $\frac{\partial Z_\ell(t)}{\partial S_i(t)}$ and $\frac{\partial Z_\ell(t)}{\partial D_j(t)}$ could be considered opposite (former pertaining to supply and latter to demand), I follow MISO conventions to have them be of same signs if power is moving in the same direction in net (e.g., from node i to j instead of the opposite).

\mathbf{ISF}_k , for a recipient node j and injector node i ,

$$\mathbf{ISF}_j = \left[\frac{\partial Z_\ell(t)}{\partial D_j(t)} \right]_{\ell \in L} \quad \text{and} \quad \mathbf{ISF}_i = \left[\frac{\partial Z_\ell(t)}{\partial S_i(t)} \right]_{\ell \in L}$$

and

$$\text{Second-best electrical proximity}_{i,j} = \frac{\mathbf{ISF}_i \cdot \mathbf{ISF}_j}{\|\mathbf{ISF}_i\| \|\mathbf{ISF}_j\|}$$

where the term “second-best” is referring to the idea that this method is the next best one to calculate electrical proximity, when PTDFs are not available. This metric captures the extent to which two nodes affect the same transmission lines in similar ways.

In the absence of ISFs or their associated shadow prices of transmission capacity (due to data limitations), MCCs themselves may provide indirect information about electrical proximity.³ Because MCCs encapsulate congestion dynamics across transmission lines, their covariance structure across nodes can be analyzed to uncover common congestion patterns. [Godin and Ibrahim \(2021\)](#) propose identifying dominant latent congestion patterns via PCA and grouping nodes that exhibit similar loading structures through correlation analyses and clustering.⁴ I extend their framework by computing cosine similarities between node loading vectors, to provide a continuous, quantitative measure of electrical proximity based on the inferred congestion dynamics; this method is described further in Subsection 3.2.

3.2 Method

PCA and parallels to ISF-shadow price decomposition of MCC.—My method differs from that of [Godin and Ibrahim \(2021\)](#) in two key aspects. First, because I focus on measuring the electrical proximity of individual gennodes to the load-zone node of interest (MEC.MECB for the Council Bluffs data center and (GRDA_HUB, GRDA_H) for the Mayes County data center), I do not aggregate MCC values at the balancing authority (BA) or ISO level. Instead, I work directly with node-level MCCs. Second, I conduct the analysis on real-time rather than day-ahead MCCs, in order to capture realized (as opposed to forecasted) load patterns, unexpected demand spikes, and real-time telemetry effects.

Aside from these differences, the procedure for implementing PCA follows [Godin and Ibrahim](#)

³While CEII requests do not specify the precise contents of the corresponding FERC-715 filings, it is likely that ISFs are included among the sensitive information covered by these confidentiality restrictions.

⁴[Godin and Ibrahim \(2021\)](#) group nodes by region and perform PCA on day-ahead, zonal-level congestion components rather than real-time, node-level data. They argue that day-ahead MCCs provide a cleaner signal with less noise than real-time prices, and they identify congestion patterns through correlations and clustering of zonal loadings rather than explicit similarity metrics.

(2021). At each time step, I first demean the nodal MCCs to account for changes in reference nodes. Each demeaned time series is then standardized across time to make it suitable for PCA. PCA decomposes the standardized MCC matrix into a weighted sum of orthonormal vectors (or loading vectors), where the weights are given by scores. Formally, the standardized MCC of node i at time t , $\widehat{\text{MCC}}_{i,t}$, can be written as:

$$\widehat{\text{MCC}}_{i,t} = \sum_{d=1}^D \phi_{i,d} s_{d,t}$$

where $\phi_{i,d}$ is the loading of node i on the d -th principal component, reflecting how strongly that node’s congestion dynamics align with the d -th latent congestion factor, and $s_{d,t}$ is the corresponding score at time t , capturing the time-varying intensity of the said congestion factor across the system.

Notice that this PCA-based decomposition has both mathematical and economic parallels to the physical decomposition of MCCs into injection shift factors (ISFs) and transmission shadow prices, expressed as:

$$\text{MCC}_i(t) = \sum_{\ell} \text{ISF}_{i,\ell}(t) \eta_{\ell}(t)$$

where $\eta_{\ell}(t)$ and $\text{ISF}_{i,\ell}(t)$ are transmission shadow values and ISF, respectively, as mentioned previously. To reiterate the economic interpretation once more, this expression states that a node’s congestion cost reflects the sum of its marginal contributions to each binding transmission constraint, weighted by the associated scarcity value of line capacity.

The PCA decomposition can be viewed as an empirical counterpart to this theoretical identity. The loadings $\phi_{i,d}$ are analogous to ISFs in that they describe how strongly each node’s MCC co-moves with a given latent congestion factor. The said latent factor can be interpreted as an aggregated congestion pattern encompassing multiple correlated transmission lines. The scores $s_{d,t}$ act similar to the shadow prices $\eta_{\ell}(t)$, as they encapsulate the time-varying intensity or activation of these latent congestion patterns. In this sense, the first few principal components summarize the dominant congestion dynamics that guide observed MCC variation, much similar to how the most binding transmission lines dominate the ISF-shadow price representation of the physical grid.

Cosine similarities of nodal loading vectors as electrical proximity.—After having obtained the principal loadings $\phi_{i,d}$ for each node i and component d , I next quantify the electrical proximity between each gennode and the load-zone node of interest by comparing their respective loading profiles. Intuitively, if two nodes exhibit similar loadings across principal congestion factors, their MCC values will tend to co-move in response to the same set of transmission constraints, demand responses, network bottlenecks, and so forth. Such nodes are therefore considered “electrically close” under the empirical, congestion-based notion of proximity. nsidered “electrically close” under

the empirical, congestion-based notion of proximity.

More formally, let $\phi_i = (\phi_{i,1}, \phi_{i,2}, \dots, \phi_{i,\bar{D}})$ denote the vector of loadings for node i , where I retain the first \bar{D} components that together explain a large share of total MCC variation (about 80%) and where each component $d \in \{1, 2, \dots, \bar{D}\}$ explains at least 1% of the total variance, following Jolliffe (2002). The electrical proximity between node i and the reference node r (i.e., MEC.MECB), denoted $EP_{i,r}$, is then defined as the cosine similarity between their loading vectors:

$$EP_{i,r} = EP_i = \frac{\phi_i \cdot \phi_r}{\|\phi_i\| \|\phi_r\|}$$

For brevity, I write EP_i since the reference node r remains fixed throughout the analysis. This measure ranges from -1 to 1 , where values close to 1 indicate that node i and the load-zone node r share highly aligned congestion dynamics (i.e., respond similarly to system-wide congestion shocks), values near 0 imply little or no shared response, and values near -1 reflect opposite responses to the dominant congestion characteristics.

For empirical implementation, I compute this cosine similarity for all gennodes within the MISO footprint relative to MEC.MECB and for all gennodes within the SPP footprint relative to (GRDA_HUB, GRDA_H). The resulting proximity values serve as continuous measures of electrical connectedness, which I later use to define treatment intensity in the difference-in-differences analysis and to identify subsets of power plants that are “electrically close” versus “electrically distant” from the load zone most affected by AI-related demand shocks. The analysis retains up to ten principal components, which together explain roughly 80% of the total variation, with each individual component contributing at least one percent of the variance.⁵

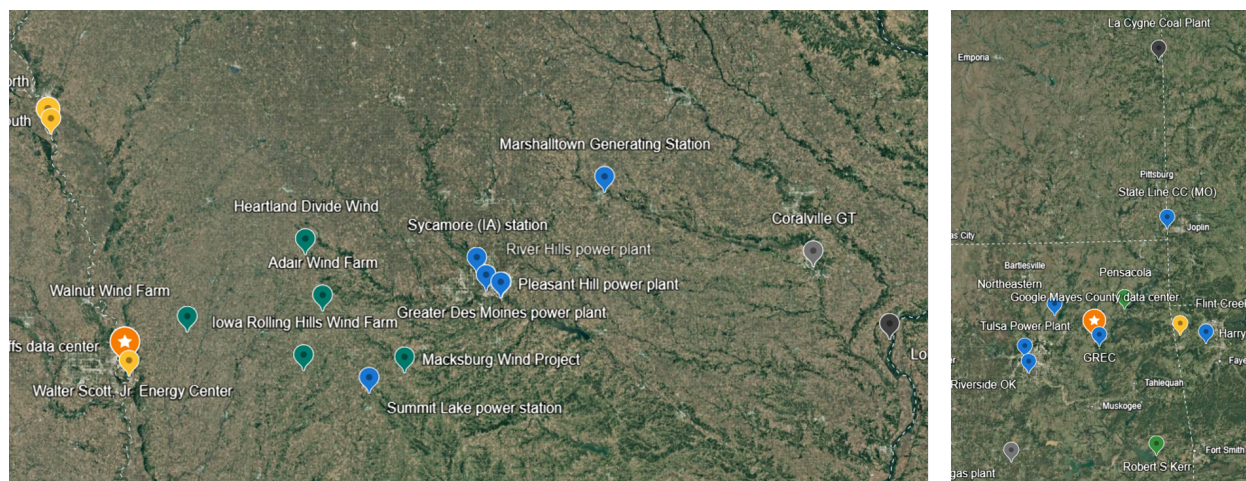
3.3 Descriptive Statistics of Electrical Proximity

I use real-time MCCs of MISO commercial pricing nodes (CPNodes) and SPP pricing nodes (PNodes) to compute the electrical proximity of each generator node relative to the load zones most likely affected by AI training. For the Council Bluffs analysis, the reference node is MEC.MECB, and I utilize data from January 2021 to November 2022. For the Mayes County analysis, I identify the relevant load point using the Settlement Location-PNode tuple (GRDA_HUB, GRDA_H), constructing electrical proximity over the period from January 2020 to September 2021. These 23-month and 21-month horizons, respectively, provide a sufficiently long pre-event window to capture stable comovement patterns in congestion dynamics prior to the suspected onset of AI training (January 2023 for Gemini 1.0 and November 2021 for GLaM/PaLM). In both regions, I retain up to 10 principal components, with each individual component contributing at least 1% of the explained variance;

⁵The cumulative variance explained by successive principal components (up to twenty) is illustrated in Figure 17.

together, these components explain at least 80% of the total variation in congestion dynamics.⁶

Figure 1: Sample of electrically proximate and distant plants to the data centers



(a) Near Council Bluffs, IA

(b) Near Mayes County, OK

NOTE: Orange indicates the Google Council Bluffs data center (panel a) and the Mayes County data center (panel b). Yellow represents electrically proximate coal power plants; green represents electrically proximate wind farms in the left panel and hydroelectric plants on the right panel; and blue represents electrically proximate natural gas power plants. Light gray and dark gray denote selected electrically distant natural gas and coal power plants, respectively.

While not directly testable, heuristic checks of the spatial distribution of electrically proximate power plants relative to Council Bluffs, IA, and Mayes County, OK, provide corroborating evidence for the validity of the electrical proximity measure. Despite geographic distance not being equivalent to electrical proximity, basic transmission physics suggest that electrically proximate generators should not be geographically remote, considering resistance and impedance. This pattern is evident in the PCA-derived electrical proximity values for both regions. Figure 1 displays the electrically proximate fossil-fuel power plants with $EP_i \geq 0.67$ (rationale for the threshold discussed below). Accounting for the fact that there are no natural gas power plants closer to Council Bluffs in the MISO territory, these coal- and natural gas-fired plants are geographically closer to the Council Bluffs and Mayes County data centers, respectively, than their electrically distant counterparts ($EP_i \in (-0.1, 0.1)$, with the rationale for thresholds elaborated below).

Figure 2 presents histograms of power plants by electrical proximity, distinguishing fossil-fuel (coal, natural gas, and petroleum) from non-fossil (wind, solar, nuclear, hydroelectric, and biomass) units. The results for the MISO region indicate that the MEC.MECB load-zone node is electrically proximate to a host of non-fossil fuel plants, particularly wind farms. However, as shown in Figure 3, focusing solely on the count of electrically proximate plants obscures the potential scale of

⁶Graphs of explained variance for each individual principal component (up to 20) for MISO is provided in Appendix Figure 17.

Figure 2: Histogram of plants by electrical proximity

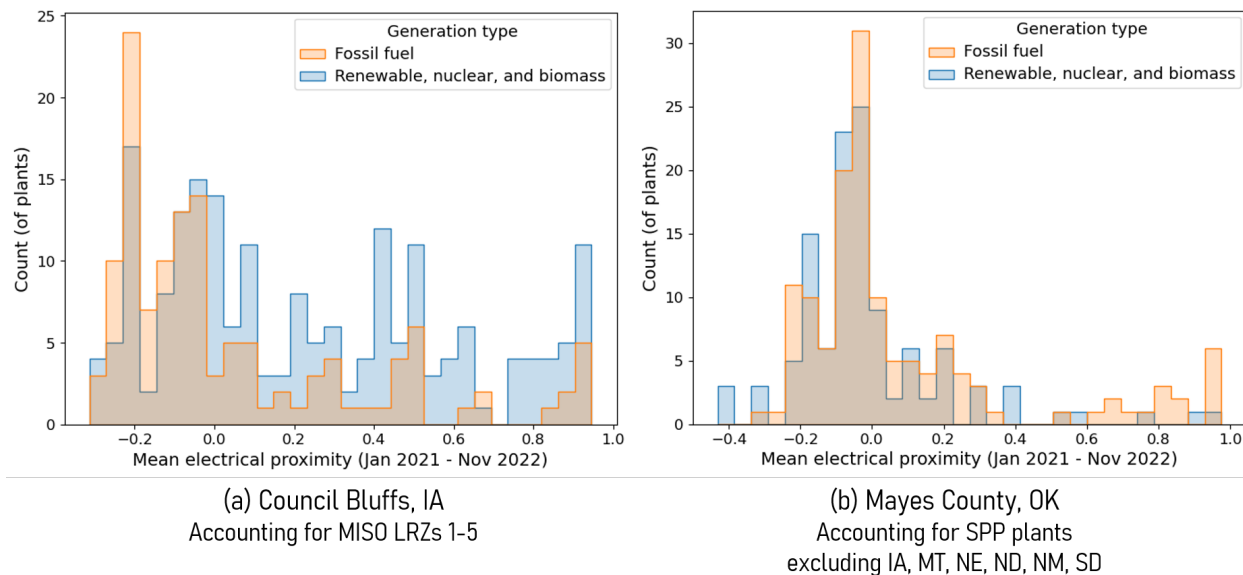
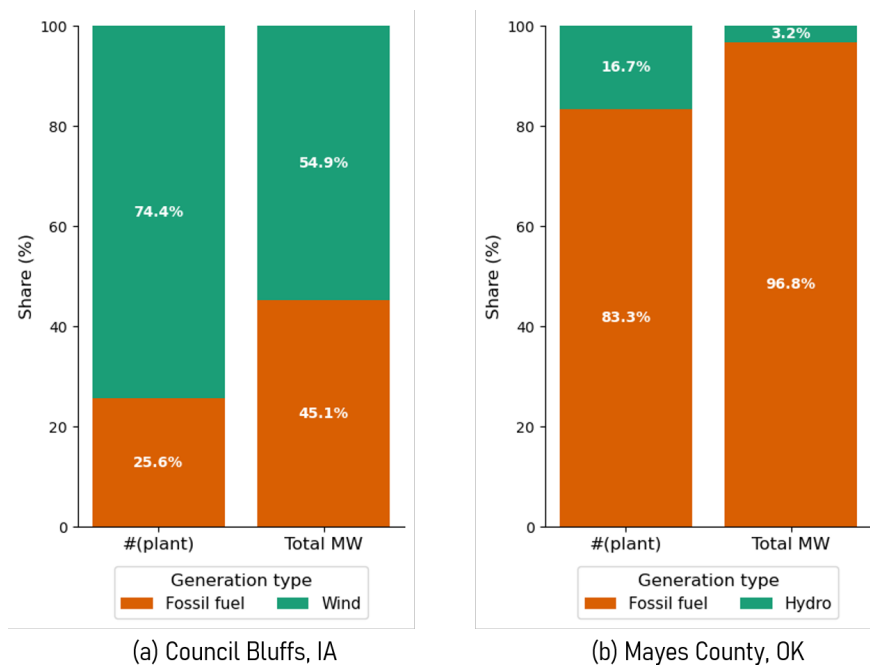


Figure 3: Shares of count and generative capacity for electrically proximate plants to data centers



accessible generation; the total generative capacity (in MW) of proximate fossil-fuel plants is comparable to that of non-fossil plants, despite their smaller numbers. The profile near the (GRDA_HUB, GRDA_H) node in the SPP region is markedly more fossil-heavy. Here, fossil-fuel plants dominate the immediate grid topology, accounting for over 80% of the total number of proximate units and more than 90% of the available generative capacity.

I limit the analysis to power plants located within MISO Local Resource Zone (LRZ) 3 (covering Iowa and southern Minnesota) and the LRZs that are geographically adjacent to it (1, 2, 4, and 5). This restriction excludes plants that are too geographically distant from Council Bluffs, such as those in eastern Texas that are still within the MISO territory.⁷ Similarly, for the SPP analysis, I restrict the sample to plants located in Oklahoma and its neighboring states (Kansas, Missouri, Arkansas, Texas, New Mexico, and Louisiana), as SPP lacks formal Local Resource Zone designations comparable to MISO. In both regions, electrically proximate plants are defined as those with $EP_i \geq 0.67$ (approximately the top 10th percentile among all identified gennodes in each region), while electrically distant plants have $EP_i \in (-0.1, 0.1)$ (roughly the 20th–60th percentile). Further interpretation of these categories, including the economic meaning of negative values below -0.1 , is provided in Subsection 4.1. I am able to match roughly 82% of generator nodes in the MISO North and Central regions (covering LRZs 1–5) and over 90% of generator nodes in the filtered SPP region to identifiable power plants; the remainder lack sufficient administrative metadata for reliable matching.

Using these definitions, I compare the pre-treatment characteristics of electrically proximate (“treated”) and distant (“control”) power plants, separated by generation type, in Table 2. This comparison serves to verify that the candidate units are broadly similar prior to the treatment. With the exception of proximity-related variables (by construction) and geographic distance to the data center, the groups are balanced across most observables. The only notable deviation appears among non-peaker natural gas plants in the Mayes County analysis, where electrically proximate units exhibit larger generation capacities. While the groups are largely comparable in levels, the validity of the synthetic difference-in-differences (SDID) design rests on the ability to match pre-treatment dynamics; consequently, assessing pre-trends remains essential to confirm the suitability of the synthetic control construction.

One potential concern is that entities located outside an ISO’s formal territory (such as those near Omaha, NE, or elsewhere in Nebraska for the MISO region) might have supplied electricity to the data centers. However, from the standpoint of electrical proximity, these entities—captured in the data as “interface” nodes due to cross-ISO exchanges—consistently appear as electrically distant or even negatively correlated. For instance, in the MISO network, these nodes exhibit proximity values ranging approximately from -0.18 to 0.22 , well outside the treatment range. This pattern

⁷For a geographic depiction of these LRZ boundaries, see Figure 16, reproduced from [Midcontinent Independent System Operator \(2025\)](#).

indicates that such entities are unlikely to have responded directly to the load shocks originating from the Council Bluffs data center. A similar pattern holds for the Mayes County data center, where external interface nodes exhibit similarly low proximity scores, ruling out significant leakage to external grids.

Table 2: Summary Statistics: Treated vs. Control Power Plants, Pre-treatment

Variable	Coal Plants			Non-peaker Gas Plants		
	Treated	Control	<i>p</i> -value	Treated	Control	<i>p</i> -value
<i>A. Council Bluffs, IA data center</i>						
Total capacity (MW)	889.60	794.43	(0.797)	582.85	487.61	(0.618)
Distance to data center (km)	87.31	603.41	(<0.001)	231.62	656.73	(<0.001)
Electrical proximity	0.762	0.015	(0.007)	0.893	-0.041	(0.023)
CO ₂ emissions (metric kilotons)	7.56	11.11	(0.573)	3.60	8.63	(0.434)
SO ₂ emissions (metric tons)	10.21	9.18	(0.890)	0.02	0.04	(0.405)
NO _x emissions (metric tons)	6.076	6.925	(0.842)	0.23	0.88	(0.199)
Gross load (GW)	8.314	11.890	(0.603)	8.838	23.526	(0.412)
Number of plants	3	13	—	2	6	—
<i>B. Mayes County, OK data center</i>						
Total capacity (MW)	733.60	698.72	(0.910)	681.76	249.59	(0.046)
Distance to data center (km)	39.55	632.84	(<0.001)	78.85	579.00	(<0.001)
Electrical proximity	0.974	-0.016	(<0.001)	0.848	-0.044	(<0.001)
CO ₂ emissions (metric kilotons)	4.30	7.82	(0.322)	1.58	0.24	(0.850)
SO ₂ emissions (metric tons)	1.66	7.80	(0.015)	0.01	0.01	(0.918)
NO _x emissions (metric tons)	3.70	5.38	(0.569)	0.78	0.19	(0.030)
Gross load (GW)	4.315	8.320	(0.321)	3.583	0.497	(0.023)
Number of plants	2	17	—	7	15	—

NOTE: This table compares mean characteristics of treated and control power plants by fuel type. *p*-values correspond to two-sided mean-difference *t*-tests. Emission quantities and gross loads are annual averages. Pre-treatment refers to January 2022 to April 2023 for Panel A, and January 2020 to October 2021 for Panel B.

4 Methodology

4.1 Setup

I employ a synthetic difference-in-differences (SDID) event-study design to estimate the dynamic environmental impacts of Google’s Gemini 1.0 training. This approach extends the standard SDID framework of Arkhangelsky et al. (2021) to an event-study setting, as formalized by Ciccia (2024) and implemented in the `sdid_event` package (Ciccia et al., 2024). This method is preferred to standard difference-in-differences in this setting because it does not rely on the parallel trends assumption holding for an unweighted control group. Instead, it constructs a “synthetic control” from a donor pool of electrically distant plants, re-weighting them to match the pre-treatment

trends of the proximate plants exactly.

More specifically, I compare emissions and generation outcomes across two distinct quasi-experimental settings. The first analyzes the Gemini 1.0 training cycle at the Council Bluffs, Iowa data center. For this specification, I define the pre-treatment period as January 2022 through April 2023, and the treatment period as May 2023 through August 2023. The second setting examines the GLaM and PaLM training cycles at the Mayes County, Oklahoma data center. Here, the pre-treatment period spans January 2020 through October 2021, with the treatment period defined as November 2021 through February 2022.

To capture the heterogeneity of grid impacts, I stratify power plants based on their electrical proximity (EP) to the respective data center load nodes. I define the treatment group as plants with EP values between 0.67 and 1.0, representing the units most physically coupled to the data center’s demand. While this is the main treatment group specification, I also use an alternative setting in which the treatment EP values are between 0.3 and 1.0 to test for broader regional effects. The donor pool for the synthetic control is limited to plants with EP values between -0.1 and 0.1. These plants are considered electrically distant enough to be unaffected by the load shock but not “electrically opposite” (i.e., below -0.1), ensuring they provide a neutral counterfactual. Robustness checks exploring a wider donor pool (e.g., -0.3 to 0.3) are reserved for future validation. From this restricted donor pool, the SDID procedure constructs a synthetic counterfactual by identifying a weighted combination of control units that best reproduces the pre-treatment trajectory of the treated group.

4.2 Estimation

The estimation strategy disaggregates the static SDID estimator into dynamic treatment effects. Following [Ciccia \(2024\)](#), the dynamic estimator $\hat{\tau}_\ell^{sdid}$ for each event-time lag or lead ℓ (relative to the event onset at $\ell = 0$) is computed by comparing the outcome trajectory of the treated unit to that of the synthetic control, adjusted by their time-weighted pre-treatment differences.

The estimator for a specific lag ℓ is given by:

$$\hat{\tau}_\ell^{sdid} = \left(Y_{tr,\ell} - \sum_{i \in \text{Donor}} \hat{\omega}_i Y_{i,\ell} \right) - \sum_{t \in \text{Pre}} \hat{\lambda}_t \left(Y_{tr,t} - \sum_{i \in \text{Donor}} \hat{\omega}_i Y_{i,t} \right)$$

where $Y_{tr,\ell}$ is the average outcome of the treated units at event-time ℓ , and $\hat{\omega}_i$ and $\hat{\lambda}_t$ are the synthetic control unit and time weights estimated via the ridge-regularized optimization procedure of [Arkhangelsky et al. \(2021\)](#). The summation sets are defined as Donor = $\{1, \dots, N_{co}\}$, representing the never-treated units, and Pre = $\{1, \dots, a - 1\}$, representing the periods prior to the adoption date a . The second term captures the weighted pre-treatment gap, ensuring the estimator accounts

for pre-treatment trend differentials rather than simple level differences.

This specification allows me to trace the month-by-month evolution of emissions, testing for pre-trends (for $\ell < 0$) and quantifying treatment dynamics (for $\ell \geq 0$). Standard errors are calculated using the bootstrap inference procedure (Algorithm 2) proposed by [Clarke et al. \(2023\)](#), which is specifically adapted for synthetic difference-in-differences event studies.

The dependent variables Y_{it} comprise the natural logarithm of emissions (specifically $\ln(x + 1)$) for CO₂, SO₂, and NO_x. To capture shifts in diurnal dispatch patterns, I also construct a peakiness ratio, $\text{Peaky}_{i,t}$, defined as the normalized difference between average generation during on-peak (8 AM–10 PM) and off-peak (12 AM–6 AM) hours:

$$\text{Peaky}_{i,t} = \frac{\text{On-peak}_{i,t} - \text{Off-peak}_{i,t}}{\text{On-peak}_{i,t} + \text{Off-peak}_{i,t}}$$

The vector of controls includes B-splines of site-level temperature and precipitation, load-zone cooling and heating degree days (CDD and HDD), county-level socioeconomic indicators, and the historical monthly average fuel mix. Standard errors are clustered at the plant level.

4.3 Identification

The identification strategy exploits variation in the extent to which individual power plants are electrically proximate to the load-zone nodes most affected by Google’s AI training activity. I apply this framework to two distinct settings: the Council Bluffs facility (connected to the MEC.MECB node) during the Gemini 1.0 training cycle, and the Mayes County facility (connected to GRDA settlement point, PNode GRDA_H) during the PaLM and GLaM training cycles. The central hypothesis is that the sharp increase in electricity demand associated with these large-scale training events produced localized shocks to their respective regional grids. These shocks propagated through the transmission network, inducing a stronger operational response among power plants that are more electrically proximate relative to those that are electrically distant. The SDID design tests whether such differential responses can be detected in emissions and generation patterns by comparing treated plants not to the raw control group, but to a data-driven synthetic counterfactual constructed to optimally match their pre-treatment trajectory.

Identification in the SDID framework relies on a distinct set of assumptions that relax the strict parallel trends requirement of conventional difference-in-differences ([Arkhangelsky et al., 2021](#)). First, rather than requiring that treated and control units follow parallel trends unconditionally, SDID assumes that a stable relationship exists between the treated unit and a weighted combination of control units. Specifically, it assumes that the average outcome of the treated unit would have evolved in parallel to the weighted average of the synthetic control group in the absence of treatment.

By explicitly re-weighting control units (via unit weights $\hat{\omega}_i$) and pre-treatment time periods (via time weights $\hat{\lambda}_t$), SDID adjusts for both time-invariant unobserved heterogeneity and shared time-varying shocks more flexibly than standard fixed-effects models.

Second, the method maintains the no-anticipation assumption, requiring that power plants did not adjust their operations in expectation of the relevant training cycles (Gemini 1.0 or GLaM/PaLM) prior to their onset. This remains reasonable given the proprietary nature of Google’s AI training schedules, which are unlikely to influence day-to-day generation decisions *ex ante*.

Third, the approach assumes that the donor pool (electrically distant plants) is sufficiently rich to reconstruct the pre-treatment dynamics of the treated plants. SDID relies on the convex hull of the donor units covering the treated unit’s trajectory; if the treated plants exhibit behavior that is fundamentally unique and cannot be approximated by any linear combination of distant plants, the synthetic control may be biased. I assess the validity of this construction by examining the event-study coefficients prior to the shock.

Fourth, treatment exogeneity implies that the assignment of electrical proximity is orthogonal to transient error terms affecting emissions. While the physical transmission structure determines proximity, SDID effectively safeguards against selection bias driven by permanent unit-specific characteristics or aggregate time trends. The double-robustness property of the estimator ensures consistency even if the true data-generating process aligns more closely with either a pure unit-fixed-effects model or a pure synthetic control model.

Finally, the timing assumption posits that the onset of the respective training runs marks a discrete structural break in demand. As in the standard framework, while exact training logs are private, the event windows are anchored to publicly verifiable milestones: the capacity expansions and press reports surrounding the Gemini program in Iowa, and the TPU v4 cluster announcements matching the GLaM and PaLM timeline in Oklahoma.

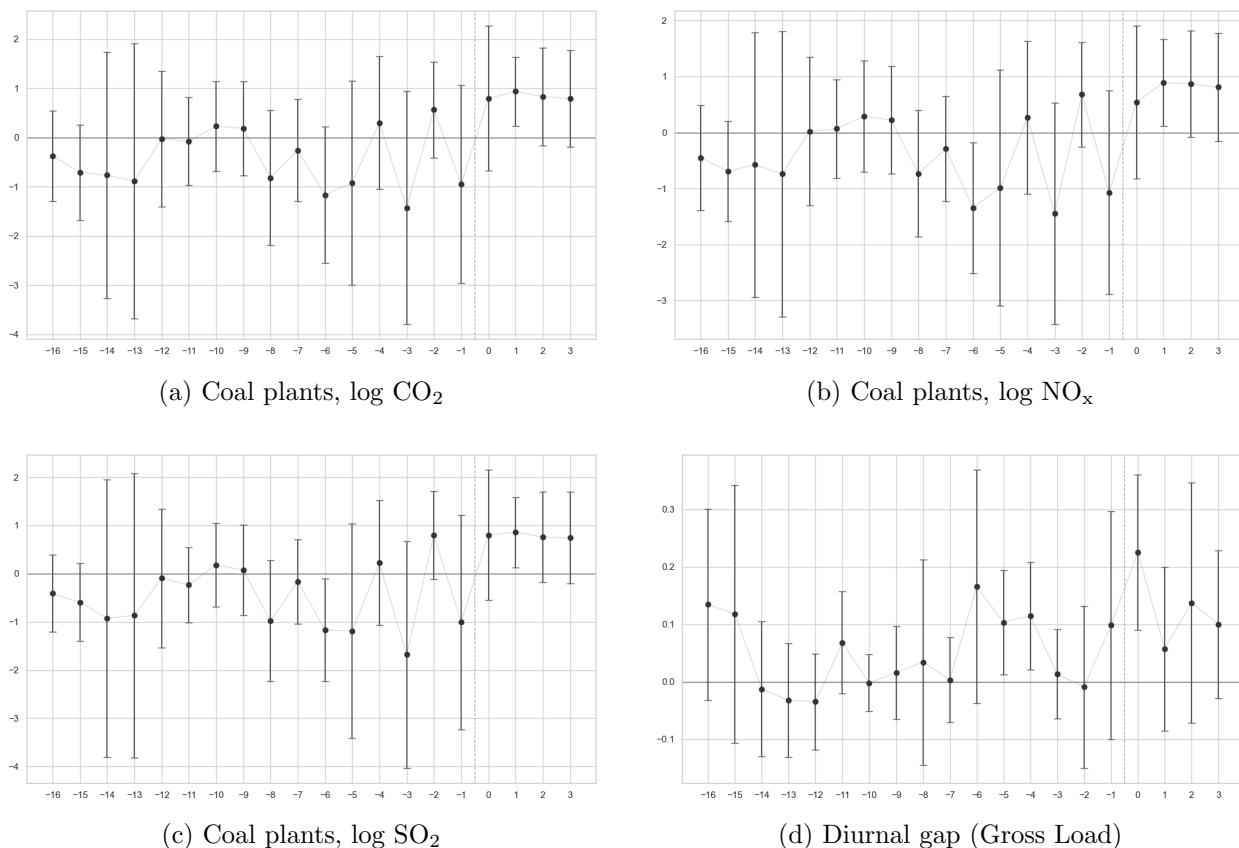
5 Results

This section presents the empirical results linking AI training-induced electricity demand to power plant behaviors, focusing on emissions and diurnal patterns, at an attempt to test the hypothesis that such AI training activity led to sharp localized shocks in the electricity grid, as described in Subsection 4.3. I first present the main results with DiD event studies, with a robustness check of the control groups being defined as those with $EP_i \in (-0.1, 0.1)$. Then, I present the additional results with the more traditional DiD design, which groups different months together based on training profiles. All results are reported separately for coal and natural gas power plants.

5.1 Gemini 1.0 analysis: Council Bluffs data center

5.1.1 Coal-fired power plants: increases in baseload emissions

Figure 4: SDID event study with MISO coal plants, log emissions and diurnal gap



Note: The panels display point estimates from the synthetic difference-in-differences (SDID) event study. The y-axis represents the event-study coefficients (approximate percentage change for log-transformed outcomes), and the x-axis indicates event time in months relative to the suspected start of Gemini 1.0 training (0 = May 2023). Confidence intervals are constructed using a bootstrap estimator with clustering at the plant level.

Table 3 presents the SDID event-study estimates for electrically proximate coal plants during the Gemini 1.0 training window. Figure 4 complements these estimates by plotting the dynamic effects, confirming the stability of pre-treatment trends. Together, the results reveal a potential two-phase response to the AI training load.

In the initial month of May 2023, we observe a statistically significant increase in the diurnal peakiness ratio of approximately 0.23 ($p < 0.05$). This positive coefficient indicates that coal plants were dispatched more intensively during on-peak hours relative to off-peak hours compared to the control group. This suggests that in the early stages of the training cycle, the additional

Table 3: SDID estimates for coal plants during Gemini 1.0 training (May–August 2023, Council Bluffs data center)

Event Month	Log Emissions			Diurnal Gap
	CO ₂	NO _x	SO ₂	Gross Load
May 2023	0.791 (0.751)	0.540 (0.695)	0.798 (0.689)	0.225** (0.069)
June 2023	0.935** (0.358)	0.894** (0.396)	0.859** (0.372)	0.057 (0.073)
July 2023	0.829 (0.507)	0.872 (0.485)	0.758 (0.479)	0.137 (0.107)
August 2023	0.789 (0.502)	0.809 (0.494)	0.746 (0.486)	0.100 (0.066)
Treated Plants	3	3	3	3

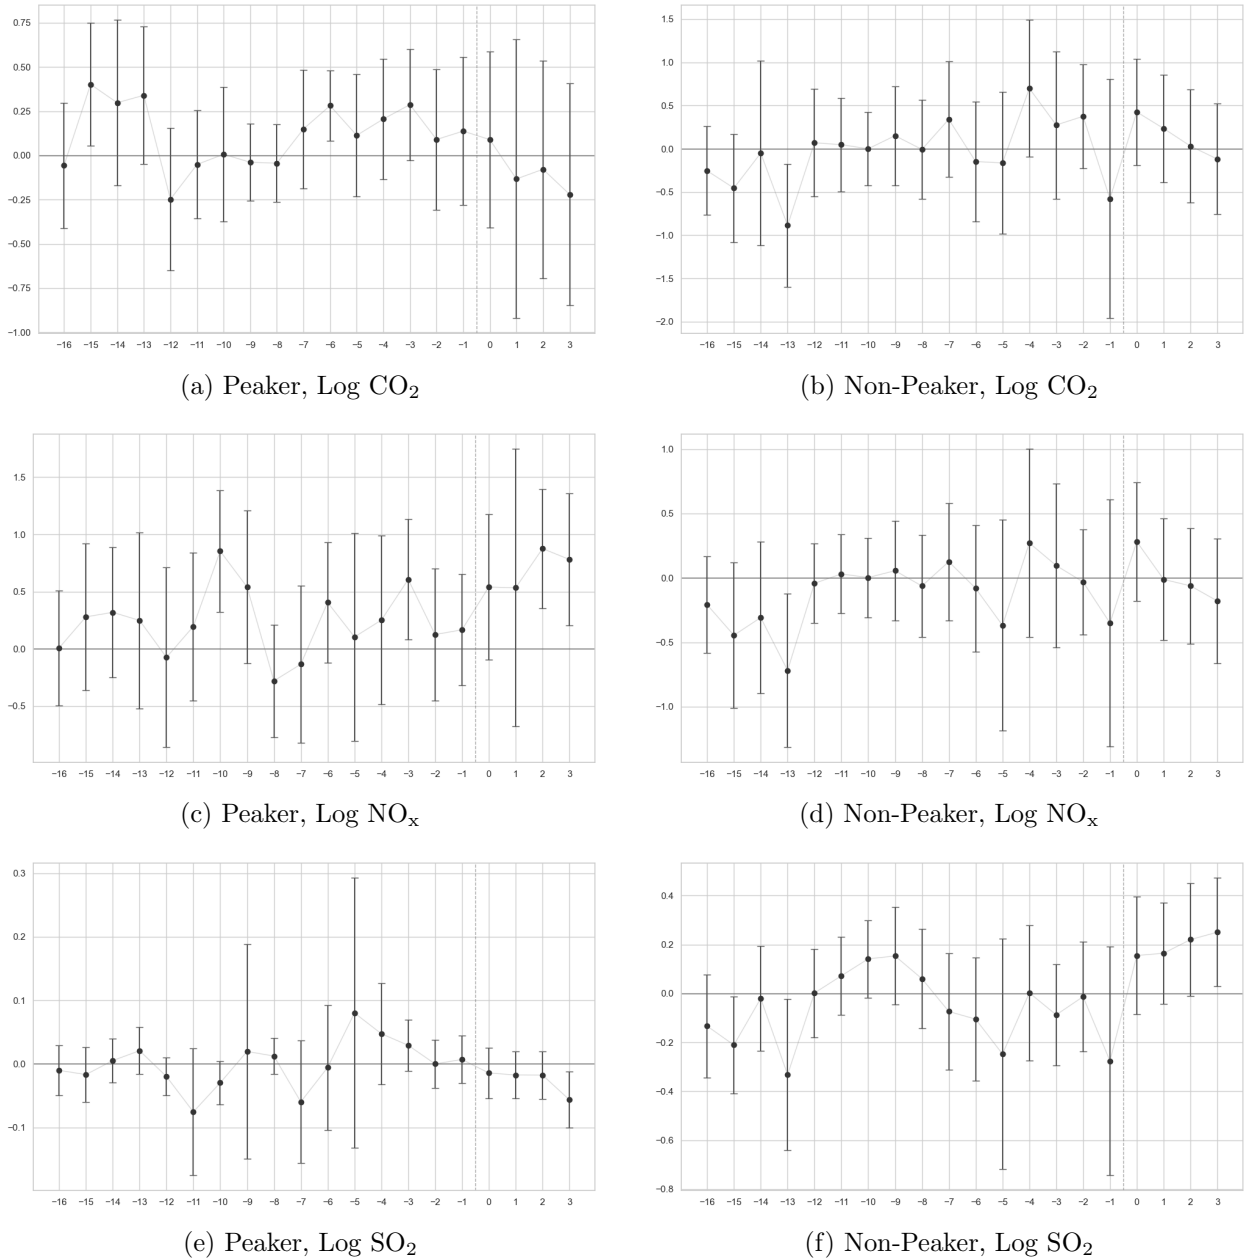
NOTES: The table reports point estimates and standard errors (in parentheses, clustered at the plant-level) from the SDID event study for electrically proximate coal power plants to the Council Bluffs data center. The dependent variables in the first three columns are log transformations of emissions; the fourth column is the diurnal peakiness ratio. ** : $p < 0.05$.

computational load may have coincided with daytime grid constraints, forcing coal units to ramp up specifically during the day rather than running as a constant baseload.

By June 2023, the operational pattern shifts. The diurnal gap coefficient falls to 0.057 and loses statistical significance, implying that the treated plants moved toward a flatter, more continuous generation profile. Simultaneously, log emissions exhibit a sharp, statistically significant surge. CO₂ emissions show a log-point increase of 0.935, with comparable increases in NO_x (0.894) and SO₂ (0.859). While recognizing that coefficients from log-transformed regressions are only approximately interpretable as percentage changes (Chen and Roth, 2024), these estimates imply a substantial ramping of fossil generation, corresponding to an approximate 90% increase in emissions.

While the emission coefficients for July and August remain large and positive (ranging from roughly 0.79 to 0.83), they lose statistical significance due to wider confidence intervals. This likely reflects seasonal confounding; as general cooling demand peaks in mid-to-late summer, the baseline variance in coal generation increases, making it harder to statistically isolate the specific contribution of the AI workload. Thus, the results suggest that the “cleanest” signal of AI-driven fossil dependence appears in June, before peak seasonal noise obscures the treatment effect.

Figure 6: SDID event study estimates for SPP natural gas plants by operational type



NOTE: The figure displays point estimates from the SDID event study for electrically proximate natural gas plants, stratified by operational type. The left column displays results for peaker plants, while the right column displays results for non-peaker plants. “Peaker” refers to single-cycle units that have capacity factors of 0.15 or less; “Non-Peaker” refers to all other natural gas units (e.g., combined-cycle or high-utilization plants). The x-axis indicates event time in months relative to the start of GLaM/PaLM 1.0 training (0 = November 2021). Confidence intervals are constructed using a bootstrap estimator with clustering at the plant level.

Table 4: SDID estimates for natural gas plants by operational type (Gemini 1.0 training)

Event month	Peaker (Log Emissions)			Non-Peaker (Log Emissions)			Diurnal Gap
	CO ₂	NO _x	SO ₂	CO ₂	NO _x	SO ₂	Gross Load
May 2023	0.090 (0.254)	0.540 (0.325)	-0.014 (0.020)	0.426 (0.314)	0.282 (0.235)	0.155 (0.123)	-0.166 (0.140)
June 2023	-0.133 (0.402)	0.533 (0.619)	-0.017 (0.019)	0.235 (0.317)	-0.012 (0.241)	0.165 (0.105)	-0.255** (0.107)
July 2023	-0.078 (0.314)	0.876** (0.266)	-0.018 (0.019)	0.031 (0.333)	-0.062 (0.230)	0.220 (0.117)	-0.246** (0.116)
August 2023	-0.220 (0.320)	0.781** (0.295)	-0.057** (0.023)	-0.119 (0.326)	-0.179 (0.248)	0.251** (0.113)	-0.282** (0.125)
Treated N	3	3	3	2	2	2	5

NOTES: The table reports point estimates and standard errors (in parentheses) from the SDID event study for electrically proximate natural gas plants. “Peaker” refers to single-cycle units that have capacity factors of 0.15 or less; “Non-Peaker” refers to all other natural gas units (e.g., combined-cycle or high-utilization plants). For diurnal gap, all gas plants are analyzed together. Significance levels are inferred from 95% confidence intervals: ** : $p < 0.05$.


5.1.2 Natural gas plants: limited response and diurnal substitution

To investigate the heterogeneity in gas plant responses, I classify natural gas units into two operational categories: “peakers,” defined as single-cycle combustion turbines with a historic capacity factor of 0.15 or less, and “non-peakers,” which include combined-cycle units and baseload gas generators. Table 4 presents the results stratified by these categories.

Unlike the coal analysis, where emissions rose uniformly across all pollutants, the results for natural gas plants lack physical consistency. For instance, while peaker units exhibit a statistically significant spike in NO_x emissions in August 2023 (coefficient 0.781, $p < 0.05$), there is no corresponding statistically significant increase in CO₂ or SO₂. Similarly, non-peaker units show an isolated increase in SO₂ in August without parallel movements in carbon emissions. This disjointed pattern suggests that while gas units may have experienced operational stress (potentially leading to inefficient combustion spikes in NO_x), they did not provide the massive net increase in energy generation observed in the coal fleet.

However, a robust structural shift is evident in the diurnal dynamics. The last column of Table 4 and Figure 8 show that the diurnal peakiness ratio for gas plants decreased significantly from June through August (coefficients ranging from -0.246 to -0.282). Since the peakiness ratio is defined as the normalized difference between on-peak and off-peak generation, a negative coefficient indicates a narrowing of this gap. This implies that gas plants, which historically operate primarily

Figure 8: SDID event study estimates for MISO natural gas plants, diurnal peakiness



NOTE: The figure displays point estimates from the SDID event study for electrically proximate natural gas plants. The x-axis indicates event time in months relative to the start of Gemini 1.0 training (0 = May 2023). Confidence intervals are constructed using a bootstrap estimator with clustering at the plant level.

during peak daytime hours, increased their relative output during off-peak (nighttime) hours. This behavior is consistent with the demands of AI training workloads: because data centers require a constant, flat stream of power 24/7, gas plants appear to have shifted their operational profile to fill overnight supply gaps, effectively acting less like intermittent peakers and more like intermediate baseload providers during the training cycle.


5.2 GLaM/PaLM analysis: Mayes County data center

5.2.1 Coal-fired power plants: baseload surge and shrinking diurnal gaps


The analysis of the Mayes County data center in the SPP region yields results that strikingly mirror the structural breaks observed in Iowa. Table 5 presents the SDID event-study estimates, accompanied by the plot presented in Figure 9.

Pre-treatment trends, as seen from Figure 9 remain stable, confirming the validity of the synthetic control. Evidence of load recruitment begins to appear in November 2021, coinciding with the start of the GLaM training cycle. While the coefficients for this month are large (approximately 3.0 log points) and statistically significant at the 10% level, the effect becomes fully pronounced in December 2021. In this month, where the training cycles of GLaM and PaLM likely overlapped, we observe a massive, statistically significant surge in emissions. Log CO₂ emissions rose by 3.16 log points, with similarly large increases in NO_x (2.94) and SO₂ (2.49). While coefficients of this


Figure 9: SDID event study with SPP coal plants, log emissions and diurnal gap




(a) Coal plants, log CO₂



(b) Coal plants, log NO_x



(c) Coal plants, log SO₂



(d) Diurnal gap (Gross Load)

Note: The panels display point estimates from the synthetic difference-in-differences (SDID) event study. The y-axis represents the event-study coefficients (approximate percentage change for log-transformed outcomes), and the x-axis indicates event time in months relative to the suspected start of GLaM and PaLM training (0 = November 2021). Confidence intervals are constructed using a bootstrap estimator with clustering at the plant level.

Table 5: SDID estimates for coal plants during GLaM/PaLM training (Nov 2021–Feb 2022, Mayes County data center)

Event Month	Log Emissions			Diurnal Gap
	CO ₂	NO _x	SO ₂	Gross Load
November 2021	3.033* (1.678)	2.811* (1.568)	2.247* (1.257)	-0.013 (0.025)
December 2021	3.164** (1.245)	2.936** (1.143)	2.487** (0.966)	0.009 (0.031)
January 2022	1.526 (1.224)	1.354 (1.054)	0.635 (0.999)	-0.024 (0.017)
February 2022	1.911 (1.336)	1.710 (1.211)	1.124 (1.044)	-0.041** (0.018)
Treated Plants	2	2	2	2

NOTES: The table reports point estimates and standard errors (in parentheses, clustered at the plant-level) from the SDID event study for electrically proximate coal power plants to the Mayes County data center. The dependent variables in the first three columns are log transformations of emissions; the fourth column is the diurnal peakiness ratio. * : $p < 0.1$, ** : $p < 0.05$.

magnitude in log-transformed models indicate multiplicative rather than percentage increases (suggesting these plants were ramped up from very low utilization baselines), the direction and scale of the effect are unmistakable: the local coal fleet was heavily recruited to power the peak of the training cycle.

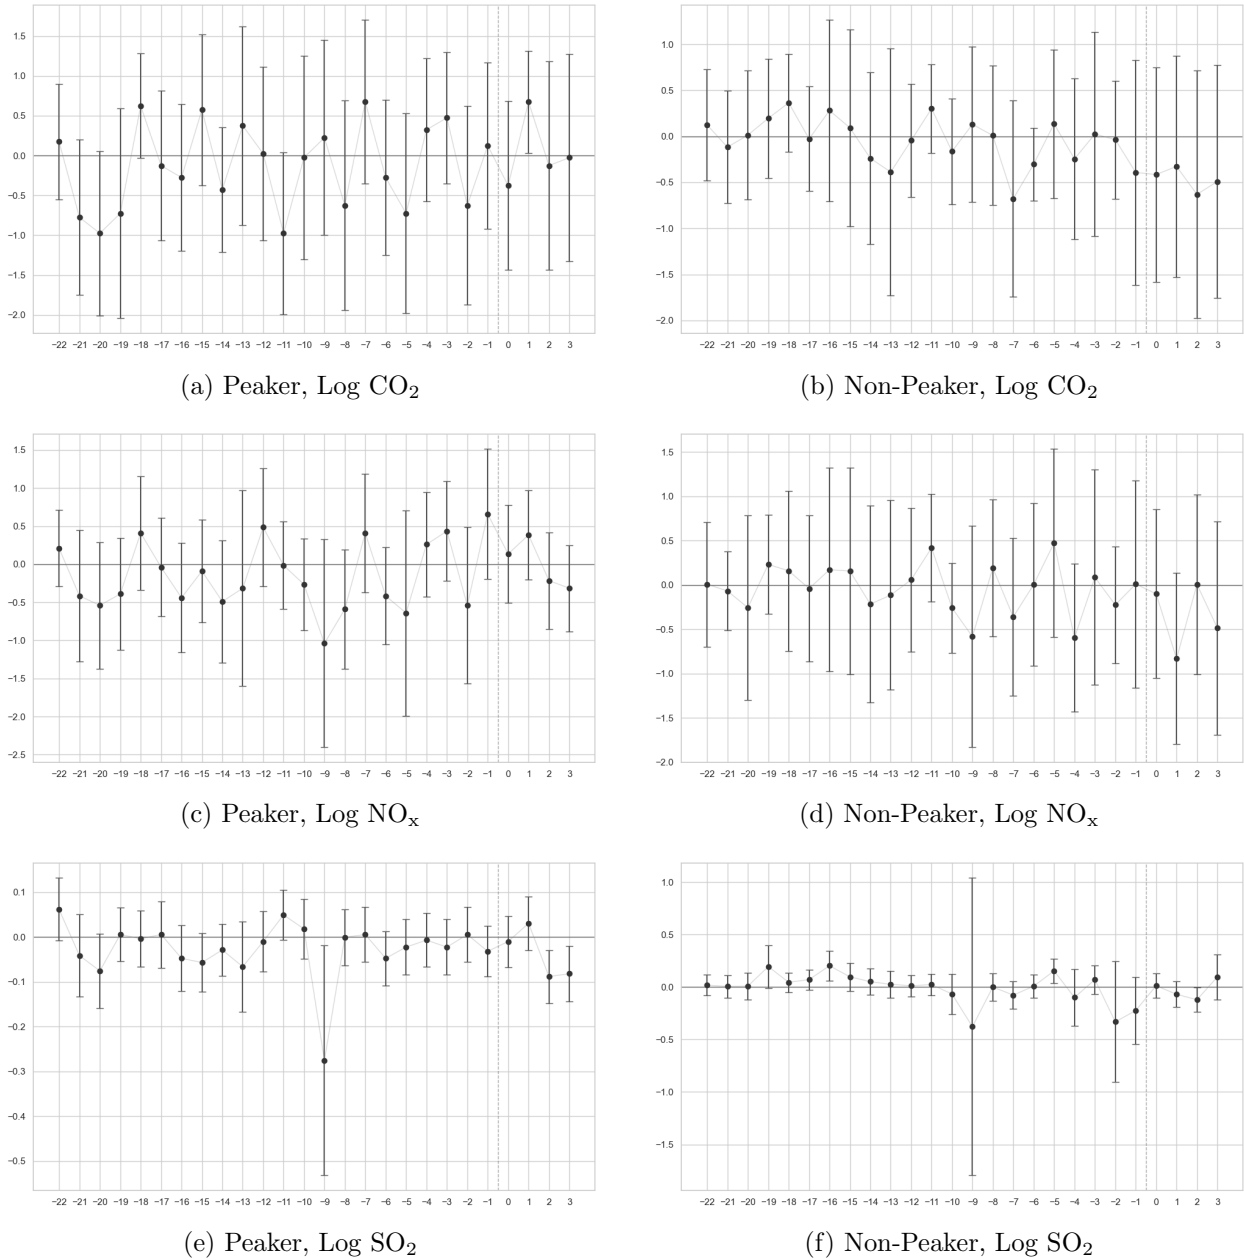
The diurnal dynamics tell a complementary story. While the peakiness ratio remained stable during the initial surge, we observe a statistically significant decrease in February 2022 (coefficient -0.041, $p < 0.05$). This narrowing of the diurnal gap implies that by the late stages of the training cycle, coal plants were filling in off-peak (overnight) generation valleys, effectively operating as constant baseload providers to satisfy the continuous computational demand.

5.2.2 Natural gas plants: marginal peaking support and muted aggregate response

Table 6 and Figure 11 presents the results for the natural gas fleet. In stark contrast to the massive, multi-pollutant surge observed in the coal sector, the aggregate response from natural gas units is largely muted, with one notable exception being peakers.

For peaker units, we observe a statistically significant increase in CO₂ emissions (approximately by 67.4%, $p < 0.05$) solely in December 2021. This timing is precise, as December represents the likely overlap of the GLaM and PaLM training cycles, creating the period of maximum aggregate

Figure 11: SDID event study estimates for SPP natural gas plants by operational type



NOTE: The figure displays point estimates from the SDID event study for electrically proximate natural gas plants, stratified by operational type. The left column displays results for peaker plants, while the right column displays results for non-peaker plants. “Peaker” refers to single-cycle units that have capacity factors of 0.15 or less; “Non-Peaker” refers to all other natural gas units (e.g., combined-cycle or high-utilization plants). The x-axis indicates event time in months relative to the start of GLaM/PaLM 1.0 training (0 = November 2021). Confidence intervals are constructed using a bootstrap estimator with clustering at the plant level.

Table 6: SDID estimates for natural gas plants by operational type (GLaM/PaLM training)

Event month	Peaker (Log Emissions)			Non-Peaker (Log Emissions)			Diurnal Gap
	CO ₂	NO _x	SO ₂	CO ₂	NO _x	SO ₂	Gross Load
November 2021	-0.376 (0.539)	0.135 (0.326)	-0.010 (0.029)	-0.415 (0.595)	-0.101 (0.485)	0.013 (0.060)	-0.100 (0.100)
December 2021	0.674** (0.327)	0.385 (0.300)	0.031 (0.030)	-0.327 (0.613)	-0.833 (0.493)	-0.069 (0.061)	-0.155 (0.198)
January 2022	-0.126 (0.667)	-0.215 (0.323)	-0.088** (0.030)	-0.630 (0.687)	0.005 (0.517)	-0.121** (0.060)	-0.007 (0.105)
February 2022	-0.026 (0.665)	-0.315 (0.289)	-0.082** (0.031)	-0.490 (0.646)	-0.487 (0.614)	0.097 (0.110)	0.062 (0.086)
Treated N	2	2	2	7	7	7	9


NOTES: The table reports point estimates and standard errors (in parentheses) from the SDID event study for electrically proximate natural gas plants to the Mayes County data center. “Peaker” refers to single-cycle units that have capacity factors of 0.15 or less; “Non-Peaker” refers to all other natural gas units. Significance levels are inferred from 95% confidence intervals: ** : $p < 0.05$.

load; an effect was seen from the electrically proximate coal plants reacting as well. The fact that peakers responded in this specific month suggests they were recruited only as a marginal resource to meet peak demand, rather than providing the sustained baseload support observed in the coal fleet.

While NO_x coefficients for peakers follow a similar positive trajectory in December (0.385), they lack statistical significance, likely due to the limited sample size ($N = 2$) and higher variance in emission rates during startup/shutdown cycling. Furthermore, while the table reports statistically significant decreases in SO₂ for January and February, these should be interpreted with caution. Since natural gas combustion emits negligible amounts of sulfur, these “significant” percentage changes likely reflect trivial absolute fluctuations near zero rather than meaningful changes in dispatch or pollution.

Finally, the diurnal gap analysis (also presented in shows no statistically significant movement. Unlike the MISO case, where gas plants shifted to fill overnight gaps, the SPP gas fleet appears to have adhered to its traditional role: providing occasional daytime peaking support during the single busiest month of the training cycle, without altering its fundamental diurnal generation profile.

Figure 13: SDID event study estimates for SPP natural gas plants, diurnal peakiness



NOTE: The figure displays point estimates from the SDID event study for electrically proximate natural gas plants. The x-axis indicates event time in months relative to the start of GLaM/ training (0 = Nov 2021). Confidence intervals are constructed using a bootstrap estimator with clustering at the plant level.

6 Discussion and Future Work

6.1 Summary of Findings and Limitations

The empirical results using the log emissions as reveal a statistically significant increase in CO₂ emissions from coal-fired power plants following the onset of the GLaM/PaLM and Gemini 1.0 training periods. In contrast, the natural gas fleet exhibited no comparable aggregate response, with the isolated exception of peaker units in the electrical vicinity of the Mayes County data center. This heterogeneity likely reflects the structural alignment between the demand profile of AI training and the dispatch characteristics of baseload generation. Unlike transient demand spikes that require flexible, fast-ramping gas units, AI model training imposes a sustained, constant load on the grid. Consequently, operators appear to have dispatched coal units to meet this continuous energy requirement, leaving natural gas plants largely on the margin.

To translate the statistical findings into aggregate physical emissions, I rely on the results from the levels specification rather than the log-transformed models. Figure 14 displays the graphical evidence for this specification. As detailed in Table 7, the cumulative excess CO₂ generated during the active training periods ranges from 0.788 megatons (restricting to $\alpha = 0.05$) to 1.582 megatons (including $\alpha = 0.10$).⁸ Applying the EPA’s 2023 social cost of carbon estimate of \$191 per ton,

⁸Estimates are derived from levels to avoid the potential distortions inherent in interpreting large percentage changes from log specifications with small baselines (Chen and Roth, 2024). This approach also ensures that the

Table 7: Estimated excess CO₂ emissions from levels regression and monetized social damages

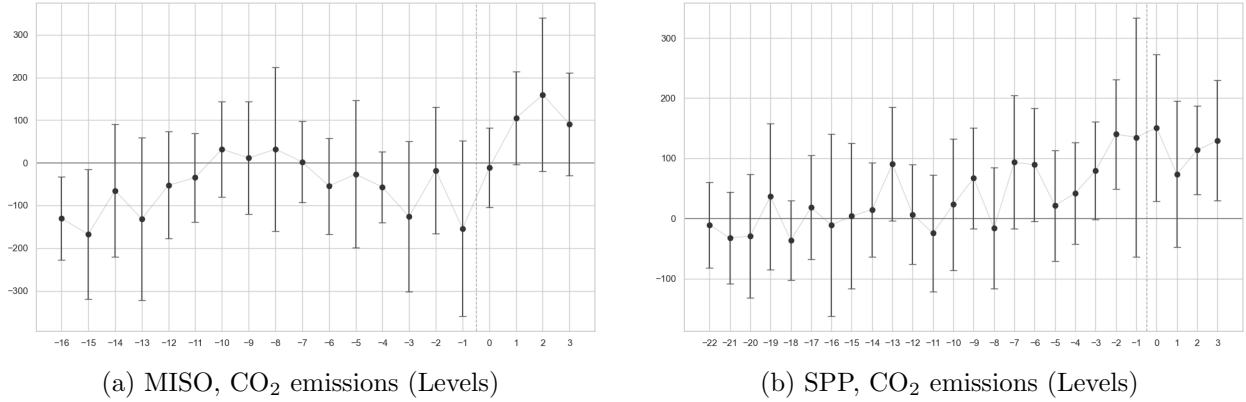
Event Source	Significant Window	α	N_{treat}	Estimated per-plant CO ₂ Increase (kt)	Excess CO ₂ (Mt)	Damages (EPA 2023 SCC)
<i>Panel A: Gemini 1.0 Training (Council Bluffs)</i>						
MISO Coal Plants	Jun 2023	0.10	3	+104.88	0.315	\$60.10 M
	Jul 2023	0.10	3	+159.73	0.479	\$91.53 M
<i>Panel B: GLaM/PaLM Training (Mayes County)</i>						
SPP Coal Plants	Nov 2021	0.05	2	+150.99	0.302	\$57.68 M
	Jan 2022	0.05	2	+113.60	0.227	\$43.39 M
	Feb 2022	0.05	2	+129.63	0.259	\$49.52 M
Total Estimated Impact ($\alpha = 0.05$)					0.788 Mt	\$150.59 M
Total Estimated Impact ($\alpha = 0.10$)					1.582 Mt	\$302.22 M

NOTE: Estimates are restricted to event months and plant types exhibiting statistically significant coefficients at the listed α significance levels in a levels-specification SDID analysis. “Estimated per-plant CO₂ Increase” is the point estimate (β) from the regression. “Excess CO₂” is calculated as $\beta \times N_{\text{treat}} \times 0.001$ where “ N_{treat} ” is the number of treated power plants. Damages are monetized using the EPA 2023 Social Cost of Carbon estimate of \$191 per metric ton (U.S. Environmental Protection Agency, 2023).

these emissions correspond to a total monetized social damage of 150.59 million USD to 302.22 million USD. While these estimates should be interpreted conservatively pending future analysis of granular dispatch data, the magnitude is economically significant. For context, Google pledged to contract for at least \$35 million worth of carbon credits in 2024 (Spock, 2024). Consequently, the estimated social damages from just these two AI training events are approximately 4 to 8 times larger than the company’s entire annual carbon removal commitment.

The study entails several limitations that outline avenues for future research. First, the definition of treatment relies on a binary classification of electrical proximity. This approach may obscure electrical proximity heterogeneity in dispatch behavior, particularly the distinction between plants closest to the load that are constrained by voltage support requirements and slightly more distant units that act as the true marginal suppliers. Future analyses could exploit granular data on spare capacity to better identify these spillover effects. Second, the analysis of the Mayes County facility is limited to the GLaM and PaLM training cycles. The rapid succession of subsequent models, including Imagen and PaLM 2, precluded the analysis of later events such as Gemini 1.0 in this region due to the absence of a stable pre-treatment baseline. Third, while the study monetizes the global damages from carbon emissions, it does not assess the local welfare impacts of increased criteria pollutants. The significant rise in sulfur dioxide and nitrogen oxides likely contributed to secondary particulate formation, the health and labor market consequences of which remain outside the scope of this paper. Finally, the analysis relies on approximate event windows derived from public announcements rather than precise computational load data. Access to high-frequency inter-estimated excess emissions remain within the physical generation capacity of the treated units.

Figure 14: SDID event study estimates for coal plants, in levels



NOTE: The figure displays point estimates from the SDID event study for electrically proximate coal-fired power plants using a levels specification (absolute CO₂ emissions in kilotons). The left panel displays results for the Council Bluffs data center (Gemini 1.0 training), while the right panel displays results for the Mayes County data center (GLaM/PaLM training). The x-axis indicates event time in months relative to the start of the respective training events: 0 = May 2023 for MISO and 0 = November 2021 for SPP. Confidence intervals are constructed using a bootstrap estimator with clustering at the plant level.

nal records would enable more rigorous identification strategies instrumental variables to quantify the marginal emissions intensity of specific AI training workloads.

6.2 Future Work and Conclusion

Looking ahead, future work will aim to address the identified limitations and extend the analysis in several directions. The first priority is to refine the definition of treatment by exploring heterogeneity in electrical proximity. The current binary classification based on congestion components may obscure distinctions between plants. Specifically, units immediately adjacent to a data center may be subject to voltage support constraints that render them inelastic to demand shocks, whereas slightly more distant “semi-close” plants may possess the spare capacity to act as the true marginal responders. Future specifications will aim to interact distance with pre-treatment capacity factors to test for these spatial spillover effects and correct for potential attenuation bias in the primary estimates.

Second, future research will leverage water consumption data to construct a more direct proxy for computational intensity. While current data availability is restricted to monthly utility-level aggregates that include non-Google users, this limitation can be addressed through an instrumental variable (IV) framework. By utilizing the specific timing of AI training cycles as an instrument for aggregate water usage, it is possible to isolate the variation in consumption driven by data center cooling from unrelated residential or commercial noise. This strategy relies on the exclusion

restriction that AI training schedules influence data center water use but do not affect neighboring users, thereby enabling the estimation of marginal emissions intensity per unit of computational load.

Third, while this study provided a comparative analysis of the Council Bluffs (MISO) and Mayes County (SPP) facilities, a broader geographical scope is necessary. The contrast between the coal-heavy response in SPP and the mixed operational dynamics in MISO highlights how local grid physics mediate environmental impacts. Future work should expand this sample to additional regions with distinct generation mixes, such as PJM, and analyze subsequent model generations beyond GLaM, PaLM and Gemini 1.0 to determine if the observed patterns persist as hardware efficiency improves and grid topologies evolve.

Finally, a key extension involves differentiating between emissions associated with AI training and those arising from inference activities. The current study benefits from the relative isolation of discrete training events, which allowed for clearer attribution of observed emissions changes. However, newer AI systems increasingly involve overlapping cycles of training and inference across distributed infrastructure. Future research should therefore seek to model both phases jointly, examining how ongoing inference operations contribute to sustained or incremental emissions over time. This distinction will become increasingly important as AI systems move from episodic model development toward continuous, large-scale deployment that interacts dynamically with the electricity grid.

To conclude, this study represents, to the best of my knowledge, the first empirical attempt to quantify the impact of AI model training on power-sector emissions using econometric techniques rather than purely model-based approaches. By combining insights from energy economics with synthetic difference-in-differences, this paper provides an evidence-based framework for examining how large-scale computational workloads interact with real electricity markets. The results indicate that the physical reality of grid dispatch often forces a reliance on local fossil generation, specifically coal, even when the corporate procurer is committed to renewable energy matching.

These findings bear critical implications for the trajectory of AI development. The computational requirements for training state-of-the-art models are scaling exponentially, and the current electricity grid, constrained by transmission bottlenecks and legacy infrastructure, appears ill-equipped to accommodate this surge in baseload demand sustainably. The contrast between the two case studies offers a suggestive mechanism: the emissions impact was notably more severe in the SPP region, where the electrically proximate fleet was dominated by coal-fired capacity. This heterogeneity suggests that expanding AI infrastructure in fossil-heavy regions creates a direct channel for increased carbon output, as legacy plants are ramped up to meet the unyielding energy requirements of model training.

Consequently, the path toward sustainable AI cannot rely on the re-animation of decrepit fossil

fuel facilities to provide firm power. While such plants offer the necessary ramping capabilities to support immediate training needs, relying on them entrenches a carbon-intensive energy mix. Instead, the results underscore the urgent necessity for deploying firm, low-carbon resources that can match the 24/7 profile of data center loads. Accelerating the integration of renewable energy storage and potentially firm baseload options such as nuclear power is essential to ensure that the technological breakthrough of artificial intelligence does not come at the cost of climate stability.

Ultimately, these findings should not be interpreted as a categorical criticism of AI development. Rather, they highlight the need to recognize and address the environmental externalities that accompany large-scale computing. This research is intended as an initial step toward formulating policies that mitigate the climate impacts of training cycles while enabling innovation to continue. Rather than constraining advancement, smarter policy approaches, such as scheduling large training runs during periods of low grid congestion ([Bernard et al., 2024](#)) or optimizing the spatial distribution of workloads, can help reduce the environmental cost of AI innovation and support a more sustainable integration of this technology into society.

References

- ACEMOGLU, D., P. AGHION, L. BURSZTYN, AND D. HEMOUS (2012): “The Environment and Directed Technical Change,” *American Economic Review*, 102, 131–166.
- ARKHANGELSKY, D., S. ATHEY, D. A. HIRSHBERG, G. W. IMBENS, AND S. WAGER (2021): “Synthetic Difference-in-Differences,” *American Economic Review*, 111, 4088–4118.
- BASHIR, N., E. A. OLIVETTI, ET AL. (2025): “The climate and sustainability implications of generative AI,” *MIT News (Institute-wide research brief)*, “Explained: Generative AI’s environmental impact”.
- BERNARD, L., A. HACKETT, R. D. METCALFE, AND A. SCHEIN (2024): “Decarbonizing Heat: The Impact of Heat Pumps and a Time-of-Use Heat Pump Tariff on Energy Demand,” Working Paper 33036, National Bureau of Economic Research.
- BOHN, R. E., M. C. CARAMANIS, AND F. C. SCHWEPPE (1984): “Optimal Pricing in Electrical Networks over Space and Time,” *The RAND Journal of Economics*, 15, 360–376.
- BONFIGLIOLI, A., R. CRINÒ, M. FILOMENA, AND G. GANCIA (2025): “Data, Power and Emissions: The Environmental Cost of AI,” Working Paper 12158, CESifo, available at SSRN: <https://ssrn.com/abstract=5553342>.
- BTC TIMES (2023): “Simple Mining Announces Major Expansion: Doubling Bitcoin Mining Capacity by Early 2024,” Accessed: 2025-10-30.
- BUSINESS RECORD (2023): “West Des Moines Approves Agreement with Microsoft for 6th Data Center,” Accessed: 2025-10-30.
- CANOO INC. (2022): “Canoo Announces EV Battery Module Manufacturing Facility in Oklahoma,” <https://www.prnewswire.com/news-releases/canoo-announces-ev-battery-module-manufacturing-facility-in-oklahoma-301665749.html>, accessed: 2026-01-15.
- CARBONBRIEF (2025): “AI: five charts that put data-centre energy use – and emissions – into context,” *CarbonBrief*.
- CATALYST COOPERATIVE (2024): “Public Utility Data Liberation (PUDL) Project,” Includes standardized data from the U.S. Environmental Protection Agency’s Continuous Emissions Monitoring System (CEMS) and the U.S. Energy Information Administration’s Forms 860 and 923. Accessed October 2025.
- CATTLEMEN’S HERITAGE BEEF COMPANY (2023): “Why Cattlemen’s Heritage — A New Approach to Beef Processing,” Accessed: 2025-10-30.
- CHEN, J. AND J. ROTH (2024): “Logs with Zeros? Some Problems and Solutions,” *The Quarterly Journal of Economics*, 139, 891–936.
- CHEN, Y. C., A. D. DOMINGUEZ-GARCIA, AND P. W. SAUER (2014): “Generalized injection shift factors and application to estimation of power flow transients,” in *2014 North American Power Symposium (NAPS)*, 1–6.

- CHOWDHERY, A. ET AL. (2023): “PaLM: Scaling Language Modeling with Pathways,” *Journal of Machine Learning Research*, 24, 1–113, first released as arXiv preprint Apr 2022.
- CICCIA, D. (2024): “A Short Note on Event-Study Synthetic Difference-in-Differences Estimators,” Tech. Rep. arXiv:2407.09565, arXiv preprint, arXiv:2407.09565 [econ.EM].
- CICCIA, D., D. CLARKE, AND D. PAILAÑIR (2024): “SDID_EVENT: Stata module providing Synthetic Difference-in-Differences (SDID) event-study estimators,” Tech. rep., Boston College Department of Economics, revised 24 Aug 2024.
- CLARKE, D., D. PAILAÑIR, S. ATHEY, AND G. IMBENS (2023): “Synthetic difference-in-differences estimation,” *arXiv preprint arXiv:2301.11859*.
- CORRIDOR BUSINESS JOURNAL (2023): “Data Centers Bringing Big Numbers to Iowa,” Accessed: 2025-10-30.
- CUTHBERTSON, A. (2023): “DeepMind’s AI chatbot can do things that ChatGPT cannot, CEO claims,” *The Independent*, accessed: 2025-10-30.
- DATACENTERS.COM (2025): “Search Data Centers: Colocation, Bare Metal, Cloud,” <https://www.datacenters.com/>, accessed: 2025-10-30.
- DAVIS, L. W., C. HAUSMAN, AND N. L. ROSE (2023): “Transmission Impossible? Prospects for Decarbonizing the US Grid,” *Journal of Economic Perspectives*, 37, 155–180.
- DE SOUSA COSTA, C. R. AND P. FERREIRA (2023): “A Review on the Internalization of Externalities in Electricity Generation Expansion Planning,” *Energies*, 16, 1840.
- DU, N. ET AL. (2022): “GLaM: Efficient Scaling of Language Models with Mixture-of-Experts,” in *Proceedings of the 39th International Conference on Machine Learning*, first released as arXiv preprint Dec 2021.
- E360, Y. (2024): “As use of A.I. soars, so does the energy and water it requires,” *Yale E360*.
- ELECTRIVE (2021): “Canoo moves up production launch to 2022,” Accessed: 2026-01-16.
- FEDERAL RESERVE BANK OF ST. LOUIS (2024): “Unemployment Rate, U.S. Bureau of Labor Statistics, retrieved from FRED, Federal Reserve Bank of St. Louis,” Accessed October 2025.
- GODIN, F. AND Z. IBRAHIM (2021): “An analysis of electricity congestion price patterns in North America,” *Energy Economics*, 102, 105506.
- GOOGLE (2023a): “April AI Update,” <https://blog.google/technology/ai/april-ai-update/>, accessed: 2025-10-30.
- (2023b): “Beyond the Forecast: A Deep Dive into the Drivers of Load Growth in MidAmerican Energy’s Territory, 2023,” Tech. rep., MidAmerican Energy Territory Analysis (unverified source), used for contextual reference; not an official MidAmerican filing.

- (2023c): “Google 2023 Environmental Report,” Tech. rep., Google, accessed: October 10, 2025.
- (2023d): “Google Sustainability Report 2023,” <https://sustainability.google/reports/>, accessed: 2025-10-30.
- (2023e): “PaLM 2: The Next-Generation Large Language Model,” <https://blog.google/technology/ai/google-palm-2-ai-large-language-model/>, accessed: 2025-10-30.
- (2024): “Google 2024 Environmental Report,” Tech. rep., Google, accessed: January 14, 2026.
- (2025): “Council Bluffs, Iowa — Data Center Location,” Accessed: 2025-10-30.
- GOOGLE CLOUD (2025a): “Gemini serving locations,” <https://docs.cloud.google.com/gemini/docs/locations>, accessed: 2026-01-15.
- (2025b): “Measuring the environmental impact of AI inference,” Tech. rep., Google LLC, blog + technical methodology paper.
- GOOGLE DEEPMIND (2023): “Announcing Google DeepMind,” <https://deepmind.google/discover/blog/announcing-google-deepmind/>, accessed: 2025-10-30.
- GRANT, N. (2023): “Google Calls In Help From Larry Page and Sergey Brin for A.I. Fight,” *The New York Times*, updated Feb 23, 2023. Accessed: 2025-10-30.
- GUIDI, G., F. DOMINICI, J. GILMOUR, ET AL. (2024): “Forecasting US data centre CO₂ emissions using AI models,” *Frontiers in Sustainability*, 5, 1507030.
- GUPTA, S. AND M. ARDITO (2022): “Google unveils world’s largest publicly available ML cluster,” <https://cloud.google.com/blog/products/compute/google-unveils-worlds-largest-publicly-available-ml-cluster>, accessed: 2026-01-15.
- HAN, Y., Z. WU, P. LI, A. WIERMAN, AND S. REN (2024): “The unpaid toll: quantifying the public health impact of AI,” *Preprint arXiv*, arXiv:2412.06288.
- HERSBACH, H., B. BELL, P. BERRISFORD, S. HIRAHARA, A. HORÁNYI, J. MUÑOZ-SABATER, J. NICOLAS, C. PEUBEY, R. RADU, D. SCHEPERS, A. SIMMONS, C. SOCI, S. ABDALLA, X. ABELLAN, G. BALSAMO, P. BECHTOLD, G. BIAVATI, J. BIDLOT, M. BONAVITA, G. DE CHIARA, P. DAHLGREN, D. DEE, M. DIAMANTAKIS, R. DRAGANI, J. FLEMMING, R. FORBES, M. FUENTES, A. GEER, L. HAIMBERGER, S. HEALY, R. J. HOGAN, E. HÓLM, M. JANISKOVÁ, S. KEELEY, P. LALOYAUX, P. LOPEZ, C. LUPU, G. RADNOTI, P. DE ROSNAY, I. ROZUM, F. VAMBORG, S. VILLAUME, AND J.-N. THÉPAUT (2020): “The ERA5 global reanalysis,” *Quarterly Journal of the Royal Meteorological Society*, 146, 1999–2049, european Union Copernicus Climate Change Service (C3S) funding.
- INTERAGENCY WORKING GROUP ON THE SOCIAL COST OF GREENHOUSE GASES (2021): “Technical Support Document: Social Cost of Carbon, Methane, and Nitrous Oxide: Interim Estimates under Executive Order 13990,” Tech. rep., U.S. Government, Washington, DC, executive Order 13990: Protecting Public Health and the Environment and Restoring Science to Tackle the Climate Crisis.

- IOWA DEPARTMENT OF TRANSPORTATION (2023): “Electric Vehicle (EV) Plan,” Accessed: 2025-10-30.
- JOLLIFFE, I. T. (2002): *Principal Component Analysis*, Springer Series in Statistics, New York: Springer-Verlag, 2nd ed.
- MASANET, E., S. SMITH, J. KOOMEY, ET AL. (2020): “Recalibrating Global Data Center Energy-Use Estimates,” *Science*, 367, 984–990.
- MICROSOFT LOCAL (2023): “Ginger East Datacenter Construction Update,” Accessed: 2025-10-30.
- MIDAMERICA INDUSTRIAL PARK (2022): “Mayes County-MidAmerica Industrial Development Project Plan,” Tech. rep., Mayes County and Oklahoma Ordnance Works Authority, review Committee Alternate 6 Draft. Accessed: 2026-01-16.
- MIDCONTINENT INDEPENDENT SYSTEM OPERATOR (2025): “Planning Year 2025-2026 Loss of Load Expectation Study Report,” Study Report PY 2025-2026, Midcontinent Independent System Operator (MISO), accessed: 2025-10-29.
- MININGSTORE (2023): “Iowa Bitcoin Mining Hosting Facility 3,” Accessed: 2025-10-30.
- MORGAN, T. P. (2022): “Google Stands Up Exascale TPuv4 Pods on the Cloud,” *The Next Platform*, accessed: 2025-10-30.
- NORTHERN DATA AG (2022): “Northern Data partners with State of Oklahoma and secures a home for up to 70,000 ASIC systems by end of 2024,” <https://northerndata.de/en/investor-relations/news/northern-data-partners-with-state-of-oklahoma-and-secures-a-home-for-up-to-70000-asic-systems-by-end-of-2024>, accessed: 2026-01-15.
- OKLAHOMA DEPARTMENT OF COMMERCE (2022): “Northern Data to Build North American Headquarters at MidAmerica Industrial Park,” <https://www.okcommerce.gov/northern-data-to-build-north-american-headquarters-at-midamerica-industrial-park/>, accessed: 2026-01-15.
- OPEN-METEO (2024): “Open-Meteo: Free weather forecast and historical data API,” <https://open-meteo.com/>, accessed via the Open-Meteo API.
- OPENAI (2023): “GPT-4 System Card,” Tech. rep., OpenAI, accessed: 2026-01-15.
- OPUS GROUP (2023): “First Phase of River Road Logistics Completed, Second Phase Starts,” Accessed: 2025-10-30.
- PATEL, D., D. NISHBALL, AND A. AHMAD (2023): “Multi-Datacenter Training: OpenAI’s Ambitious Plan,” <https://newsletter.semianalysis.com/p/multi-datacenter-training-openais>, accessed: 2026-01-15.
- SHEHABI, A., A. NEWKIRK, S. J. SMITH, A. HUBBARD, N. LEI, M. A. B. SIDDIK, B. HOLECEK, J. KOOMEY, E. MASANET, AND D. SARTOR (2024): “2024 United States Data Center Energy Usage Report,” Tech. Rep. LBNL-2001637, Lawrence Berkeley National Laboratory, Berkeley, CA, Energy Analysis and Environmental Impacts Division, Peer reviewed.

- SPOCK, R. (2024): “Our Pledge to Support Carbon Removal Solutions,” *The Keyword: Company News*, google Sustainability Blog post announcing \$35 million in carbon removal commitments.
- UC TODAY (2023): “Google Reportedly Nearing Release of GPT-4 Competitor Gemini,” *UC Today*, accessed: 2025-10-30.
- U.S. ENVIRONMENTAL PROTECTION AGENCY (2023): “Report on the Social Cost of Greenhouse Gases: Estimates Incorporating Recent Scientific Advances,” Tech. rep., U.S. Environmental Protection Agency, Washington, DC, updated Social Cost of Carbon based on the Greenhouse Gas Impact Value Estimation (GIVE) model.
- (2024): “Continuous Emissions Monitoring System (CEMS) Data,” Accessed through the Public Utility Data Liberation (PUDL) project, October 2025.
- VALVE WORLD (2023): “Clow Valve USD 75M Expansion Breaks Ground,” Accessed: 2025-10-30.
- VICTOR, J. (2023): “How Google is Planning to Beat OpenAI,” *The Information*, accessed: 2025-10-30.
- WEITZ COMPANY (2023): “Confidential Data Center Campus (Iowa),” Accessed: 2025-10-30.
- WOOD, A. J., B. F. WOLLENBERG, AND G. B. SHEBLÉ (2014): “Power System Security,” in *Power Generation, Operation, and Control*, Hoboken, NJ: John Wiley & Sons, chap. 7, 296–349, 3rd ed.

Appendix

A Notes on Load Zone Selections

A.1 MidAmerican Energy Load Zones for MISO

MISO’s network model defines multiple load zones under the MidAmerican Energy Company (MEC) balancing authority. All load-zone commercial pricing nodes in the MISO network, including those under MEC, follow the convention `<LSE>.<BA>`, where the prefix identifies the load-serving entity and the suffix identifies the balancing authority. Thus, `MEC.MECB` denotes MidAmerican’s own retail load aggregated under its balancing authority, whereas other prefixes (e.g., `ALTW`, `CFU`, `PELLa`) denote embedded municipal or cooperative utilities that settle through MEC’s balancing area. Because each municipal zone represents a distinct local utility with its own demand profile and service boundaries, none individually capture the Council Bluffs system load. In contrast, `MEC.MECB` aggregates the full MEC footprint, including the Council Bluffs area, making it the most appropriate publicly observable proxy for regional demand changes potentially driven by the Google data center. Table 8 lists all MEC-prefixed load zones and summarizes their relevance.

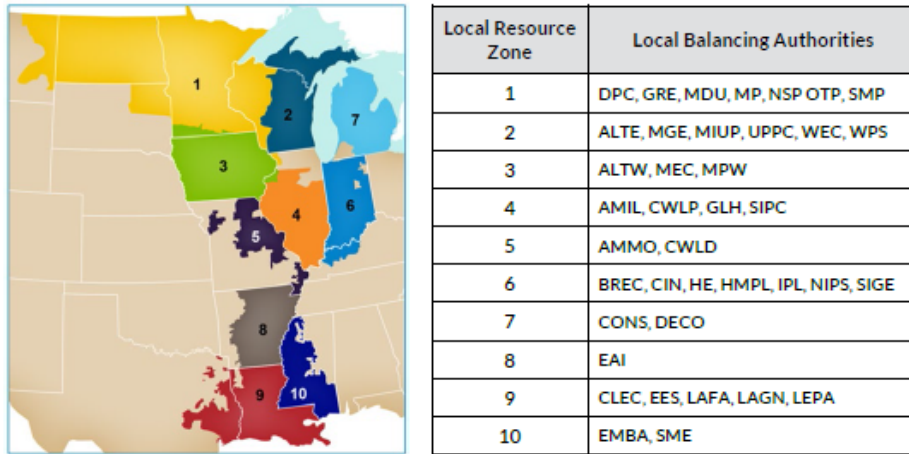
A.2 GRDA Load Zones for SPP

For the analysis of the Mayes County facility within the Southwest Power Pool (SPP), I select the Grand River Dam Authority Hub (`GRDA.HUB`, also listed with the PNode `GRDA.H`) as the reference node for identifying electrically proximate generators. This choice follows from the utility structure of the region, as the Grand River Dam Authority supplies power to the MidAmerica Industrial Park hosting the data center. The SPP commercial model defines several related settlement locations, including `GRDA.HUBSA` (PNode `GRDASA.H`) and `GRDAHUB2` (PNode `GRDAHUB2.H`), which correspond to distinct accounting or sub-regional aggregations within the same service territory. However, an analysis of the congestion components reveals that these nodes are electrically equivalent for the purpose of this study; they share near-identical sensitivity factors to the surrounding transmission constraints. Therefore, using `GRDA.HUB` provides a representative signal of the local grid conditions facing the data center, and the specific selection among these highly correlated GRDA nodes does not drive the identification of the treated generator set.

B Appendix Figures and Tables


B.1 Figures

Figure 16: Map of MISO LRZs (left) and local BA to LRZ designations



NOTE: This figure is a copy of Figure ES-2 and Table ES-2 from [Midcontinent Independent System Operator \(2025\)](#).

Figure 17: Explained Variance of Principal Components, from PCA with MISO real-time MCCs



NOTE: Up to 20 principal components are depicted, as the other, unreported principal components contribute less than 1% of explained variance each.

B.2 Tables

Table 8: MISO Load Zones under MidAmerican Energy Company (MEC)

Node	Entity / Area	Notes and Relevance
MEC.ALTW	Alta Municipal Utilities	Small municipal in northwest Iowa; limited load; not representative of Council Bluffs.
MEC.AMES	Ames Municipal Electric System	Central Iowa city with independent generation; geographically distant.
MEC.ATLN	Atlantic Municipal Utilities	Western Iowa municipal; small load; separate local service area.
MEC.BEPM.CPZD	Butler Electric Power (Central Power Zone)	Embedded municipal/cooperative; not near Council Bluffs.
MEC.CFU	Cedar Falls Utilities	Eastern Iowa municipal utility; large but distant from Council Bluffs.
MEC.CMMPA.ELD	CMMPA (Eldora)	Joint-action agency member in central Iowa.
MEC.CMMPA.MONT	CMMPA (Montezuma)	Joint-action agency member in east-central Iowa.
MEC.GENESEO	Geneseo Municipal Utility (IL)	Across the Iowa–Illinois border; not part of western Iowa load.
MEC.HCPD.ABRN	Harrison County Power District (Auburn area)	Western Iowa rural cooperative; small industrial load.
MEC.MEAN.LD	Municipal Energy Agency of Nebraska (Load)	Interconnected via MEC’s transmission ties; external to Iowa.
MEC.MECB	MidAmerican Energy Balancing Authority (Aggregate Load Zone)	Aggregates all MidAmerican-served load across Iowa and South Dakota, including Council Bluffs; most relevant proxy for data-center demand.
MEC.PELLA	Pella Municipal Utilities	Central Iowa municipal; not representative of Council Bluffs.
MEC.RLGY01	Regulated Load Group 01	Small aggregated MEC retail group; not geographically specific.

NOTE: Each node corresponds to a commercial pricing node (CPNode) identified as a load zone in MISO’s network model. While several municipalities operate under MEC’s balancing authority, only MEC.MECB encompasses the entire MidAmerican service territory and includes the Council Bluffs metropolitan area.

Table 9: Potential Sources of Load Growth in MidAmerican Energy (MEC) Territory, 2023

Project / Facility	Type	Status (2023)	Expected Load Characteristics or Temporal Relation to Observed Mid-2023 Uptick	Primary References
Google Council Bluffs Data Center	Hyperscale data center	Operational since 2009; active expansion; AI-training activity Jan-Sep 2023	Large, continuous 24/7 baseload (100–300 MW plausible). High water use (3.7 million m ³ in 2023) and timing coincides with observed mid-2023 zonal load increase, suggesting intensified utilization during Gemini 1.0 training.	Google (2023d, 2025); Grant (2023); Cuthbertson (2023)
Microsoft West Des Moines Data Centers (Campuses 4–6; Ginger East/West)	Hyperscale data center	Under construction; phased energization from late 2023; full operations extend to 2026	Confidential two-building complex (192 MW planned) under phased energization. Reported by multiple sources, but load primarily construction-related in early 2023. AI-generated “Deep Dive” report attributes regional load growth to Microsoft, though this remains unverified.	Business Record (2023); Weitz Company (2023); Microsoft Local (2023); Google (2023b)
Meta Altoona Campus	Hyperscale data center	Operational; expanding	Over 5 million sq. ft. campus with estimated 1.2 million MWh annual use (≈100,000 homes). Outside MEC but regionally affects power flow. Expansion continued during 2023; little direct bearing on Council Bluffs load.	Corridor Business Journal (2023)
River Road Logistics Park (Council Bluffs)	Warehousing / light industrial	Phase 1 completed (early 2023); Phase 2 construction ongoing	Typical daytime-peaked load (2–5 MW per building). Insufficient scale or temporal overlap to explain mid-2023 load surge.	Opus Group (2023)
Cattlemen’s Heritage Beef Processing Plant (Mills County)	Food processing / industrial	Land purchase finalized; groundbreaking 2023; operational 2024–25	Future moderate industrial load (5–10 MW expected). 2023 demand limited to construction. Included in “Deep Dive” report but not operational contributor to 2023 load.	Cattlemen’s Heritage Beef Company (2023); Google (2023b)
Clow Valve Foundry Expansion (Oskaloosa)	Heavy industry / manufacturing	Groundbreaking 2023	New 36,000 sq. ft. iron foundry expansion (≈\$80 million) with high electricity intensity. Modernization raises baseline load gradually; minor share of 2023 zonal change.	Valve World (2023)
Simple Mining and MiningStore Facilities (Cedar Falls, Grundy County, Muscatine)	Cryptocurrency mining	Active; expanded capacity through 2023	Combined operational capacity exceeding 75 MW by end of 2023. Continuous baseload, some within MISO. Contributes to rising regional consumption, but smaller magnitude than data-center demand near Council Bluffs.	BTC Times (2023); MiningStore (2023)
Rural Water and Municipal Electrification Projects	Utility infrastructure / transportation electrification	EV charging and municipal fleet initiatives announced 2023; early-stage build-out	Incremental, distributed, and low aggregate magnitude relative to hyperscale data-center load. Likely immaterial for observed mid-2023 increase.	Iowa Department of Transportation (2023)

Table 10: Potential Sources of Load Growth in MidAmerica Industrial Park, 2021-2022

Project / Facility	Type	Status (Winter 2021-22)	Expected Load Characteristics or Temporal Relation to Observed Uptick	Primary References
Google Mayes County Data Center	Hyperscale data center	Operational; massive expansion of TPU v4 cluster	Home to “world’s largest ML cluster” (9,000 TPU v4 chips) announced May 2022. Pre-announcement testing and training of PaLM (Nov 2021–Feb 2022) perfectly matches observed load surge.	Gupta and Ardito (2022); Chowdhery et al. (2023); Google (2023d)
Northern Data (Crypto Mining)	Cryptocurrency / HPC	Not yet operational	Partnership announced March 2022; initial 10 MW energization began Q2 2022. Load ramp-up occurred after the PaLM training window.	Oklahoma Department of Commerce (2022); Northern Data AG (2022)
Canoo (EV Manufacturing)	Automotive / Battery	Planning / Pre-construction	Major battery module facility announced Dec 2022. Production schedules moved up but significant load not present during winter 2021.	Canoo Inc. (2022); Electrive (2021)
MidAmerica Industrial Park Legacy Tenants (DuPont, Nordam)	Heavy Industry / Chemical	Steady state	Long-standing industrial loads with predictable diurnal patterns. No major expansion announcements coincide with the specific winter 2021-22 spike.	MidAmerica Industrial Park (2022)