Hidden Environmental Footprints of AI Training: Evidence from Google Gemini 1.0

Junho Choi

Columbia Sustainable Development

September 13, 2025

CU Environmental and Resource Economics Workshop

Motivation: Computation and Electricity Usage

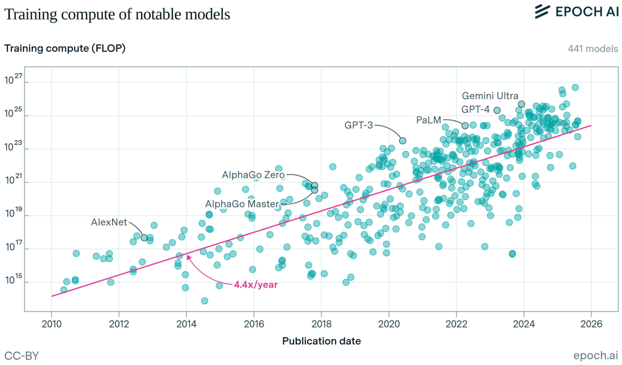

- Estimated floating-point operations for Gemini 1.0: 5×10^{25} FLOPs
- A modern laptop ≈1 trillion FLOPs per second → 1.6 million years
- BOTE calculation: "enough to power a small city for a year" (6-10K HH)

Why Gemini 1.0 and Iowa?

Gemini 1.0 training: one of the most computationally intensive

Data center at Council Bluffs, IA

- Well-connected to wind and solar
- **Marginal emissions**: such locations are still connected to fossil-fuel power plants
- Suspected deployment of **TPUs** (Patel, Nishball, & Ontiveros 2024)
- Other candidate: Mayes County, OK – evidence of TPU presence (Morgan 2022)

Figure: Council Bluffs, IA data center (Davis 2017)

"Hidden" cost: environment and climate

Main cost of AI training: **electricity**

- Data centers: power purchasing agreements (Verge 2014; EDF 2017)
  PPA
- PPAs as price hedges (Pombo-Romero et al. 2024)

"Hidden" cost: environment and climate

Main cost of AI training: **electricity**

- Data centers: power purchasing agreements (PPA) (Verge 2014; EDF 2017)
- PPAs as price hedges (Pombo-Romero et al. 2024)

What about the **social costs** of AI training?

- Price shocks to unprotected consumers & busi. (Penn and Weise 2025)
- Emissions: when computationally intensive, is energy demand at data centers **sufficiently** met by **renewables**?
- If relying on fossil fuels, is Google doing enough in offset efforts?
- Potential health and labor consequences? (Almond, Choi, and Papp *WP*)

More context & potential contribution

Fossil fuel plants: **electrically proximate** + **generative capacity**↑

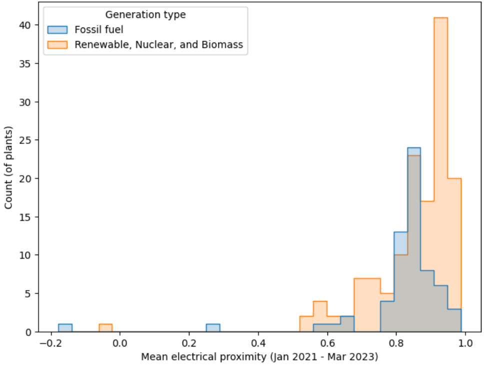

More context & potential contribution

Fossil fuel plants: **electrically proximate** + **generative capacity**↑

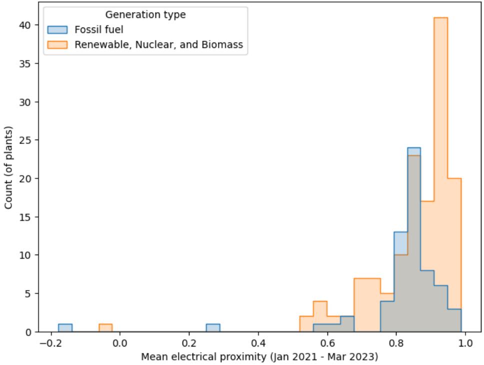

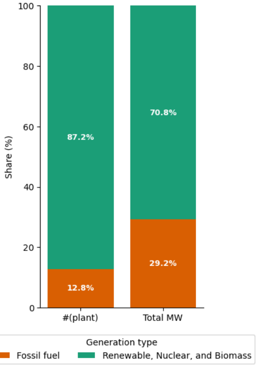

Research Question: Is there evidence that Google's Gemini 1.0 training led to significant rise in emissions (CO_2, NO_x, SO_2)?

Suspected timeline

- April 2023: Google DeepMind is AI effort is announced (Deepmind 2023)
- May 2023: Official "Gemini still in training" (Google 2023)
- Aug-Sep 2023: Paywalls being setup, and recruitment of testers (Thompson 2025)
- Nov 2023: Still in fine-tuning but Deepmind CEO "happy" CNBC 2023
- Dec 2023: Official launch of Gemini 1.0

Suspected timeline

- April 2023: Google DeepMind is AI effort is announced (Deepmind 2023)
- May 2023: Official “Gemini still in training” (Google 2023)
- Aug-Sep 2023: Paywalls being setup, and recruitment of testers (Thompson 2025)
- Nov 2023: Still in fine-tuning but Deepmind CEO “happy” CNBC 2023
- Dec 2023: Official launch of Gemini 1.0

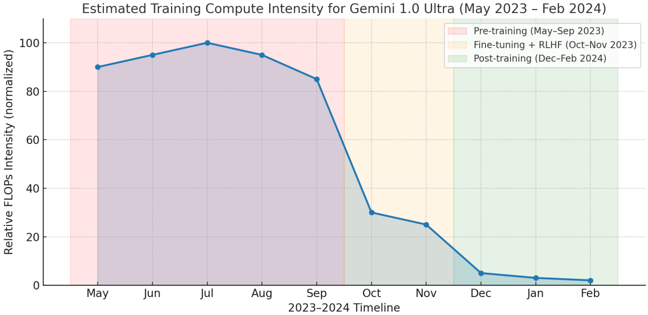

(just a sketch!)

Assembling the dataset

Emissions data – Continuous Emissions Monitoring System

- Power plant-level hourly data on NO_x, SO_2, CO_2 emissions (in mass) and heat input (in BTU); also available from **PUDL**
- CO_2 focus, but also SO_2: potentially **health channel**
  - EPA and WHO: acute respiratory attacks and forms PM
- Missing: *exact mix* of fuels plants are using
- TBC: transport modeling (AERMOD; suitable for short-range)

Assembling the dataset

Emissions data – Continuous Emissions Monitoring System

- Power plant-level hourly data on NO_x, SO_2, CO_2 emissions (in mass) and heat input (in BTU); also available from **PUDL**
- CO_2 focus, but also SO_2: potentially **health channel**
  - EPA and WHO: acute respiratory attacks and forms PM
- Missing: *exact mix* of fuels plants are using
- TBC: transport modeling (AERMOD; suitable for short-range)

Locational marginal prices by node – MISO:

- Midcontinent Independent System Operator (MISO)
- Focus on North (e.g., IA, MT, MN) and Central regions (e.g., WI, MO), excluding Manitoba and South region
- **Congestion components** of LMP at the hourly level

Assembling the dataset

Emissions data – Continuous Emissions Monitoring System

- Power plant-level hourly data on NO_x, SO_2, CO_2 emissions (in mass) and heat input (in BTU); also available from **PUDL**
- CO_2 focus, but also SO_2: potentially **health channel**
  - EPA and WHO: acute respiratory attacks and forms PM
- Missing: *exact mix* of fuels plants are using
- TBC: transport modeling (AERMOD; suitable for short-range)

Locational marginal prices by node – MISO:

- Midcontinent Independent System Operator (MISO)
- Focus on North (e.g., IA, MT, MN) and Central regions (e.g., WI, MO), excluding Manitoba and South region
- **Congestion components** of LMP at the hourly level

Weather data – ECWMF and ERA5:

- Access through Openmeteo historical data API

Methodology

Difference-in-differences setup

$$Y_{i,t} = \beta\,(T_i \times C_t) + \boldsymbol{X}_{i,t}\gamma + \rho_i + \iota_{\mathrm{ym}(t)} + \epsilon_{i,t}$$

- $Y_{i,t}$: heat content (BTU) or emissions (CO_2, SO_2, NO_x)
- $T_i \in \{0, 1\}$: =1 if plant i is "electrically proximate"
- $C_t \in \{0, 1\}$: =1 if day t falls within training window (May-Sep 2023); trials with different timings as well
- ρ_i, $\iota_{\mathrm{ym}(t)}$: fixed effects
- $\boldsymbol{X}_{i,t}$: weather, (local) economic variables, etc.

Accompanied by a similar **event-study** (with control groups) setup

Electrical proximity

Currently using simpler, congestion-driven LMP differences to identify electrical proximity (Godin and Ibrahim 2021; also ideas from Brown, Zarnikau, & Woo 2020; Bushnell, Mansur, & Saravia 2008)
$\rightarrow$ strong co-movement of congestion price indicates injection/withdrawals are jointly limited by the same bottlenecks

1. Identify the "demand node" that Council Bluffs DC is connected to
2. Identify the power plants linked with each generation node ("gennode") within MISO North and Central
3. For the demand node and each gennode, calculate the correlation of **congestion component** of LMP. Specifically, fix the timeframe to a pre-training period
4. Rank the said correlations, defining $\geq$ 90pct as "treated" and $<$ 60pct as "control"

Alternative: resistance-based method, requiring specific transmission line-level information (FOIA + CEII form application for FERC Form 715 underway)

Electrical proximity (continued)

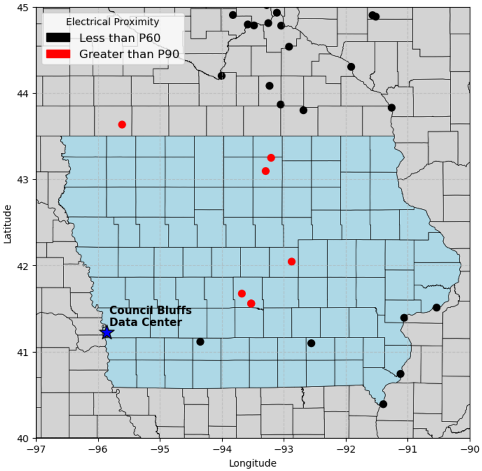

Electrical proximity (continued)

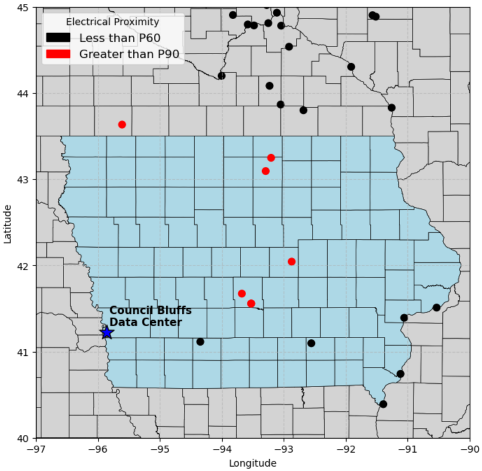

Variable	Electrical proximity Far	Close	t
Total MW	394.7	344.4	0.41
Natural gas	0.59	0.83	-1.29
Coal	0.34	0.00	3.84
N	29	6	

→ Proceeding with natural gas power plants

Pre-trend comparisons

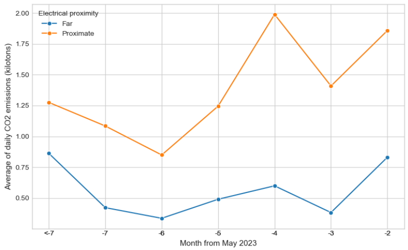

Figure: Comparison of CO_2 emissions

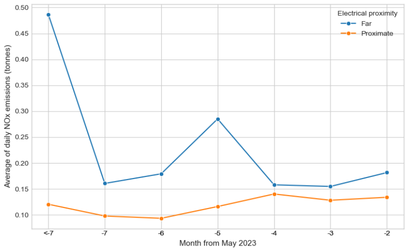

Figure: Comparison of NO_x emissions

→ suggestive that training may have started before May 2023

Diff-in-diff results, natural gas power plants

	DV: Daily CO_2 emissions (kilotons)			
	(1)	(2)	(3)	(4)
Gemini 1.0 training	0.888**	0.864***	0.841***	0.822***
	(0.067)	(0.072)	(0.068)	(0.822)
Gemini 1.5 training		0.081		0.101
		(0.050)		(0.052)
Gemini 1.5 Flash Fine-tuning		-0.603***		-0.552
		(0.040)		(0.043)
	DV: Daily heat input (billions of BTU)			
	(1)	(2)	(3)	(4)
Gemini 1.0 training	15.46***	14.93***	14.51***	14.08***
	(1.069)	(1.133)	(1.094)	(1.172)
Gemini 1.5 training		0.680		1.117
		(2.781)		(0.642)
Gemini 1.5 Flash Fine-tuning		-11.72***		-10.69***
		(0.353)		(0.422)
FE	— Year-month, power plant —			
Weather controls	N	Y	N	Y
Additional controls	— Local economic conditions —			
N	— 30,322 —			

Note: Reporting difference-in-differences results. Clustered (plant groups) SE in parentheses. *** : $p < 0.01$, ** : $p < 0.05$.

During the suspected training period, CO_2 (and heat input) significantly higher for electrically proximate plants

- $\approx$0.57SD$\uparrow$ in daily CO_2 emissions

Other training / fine-tuning

- G1.5 training: Oct-Dec 2023
- G1.5 fine-tuning: Mar-Apr 2023
- "Gemini 1.5 has been training in parallel": likely adding onto 1.0, not from scratch

Diff-in-diff results (continued)

	DV: Daily SO_2 emissions (tons)			
	(1)	(2)	(3)	(4)
Gemini 1.0 training	0.211	0.239	0.213	0.240
	(0.080)	(0.092)	(0.086)	(0.097)
Gemini 1.5 training		0.228		0.226
		(0.093)		(0.094)
Gemini 1.5 Flash Fine-tuning		0.222		0.218
		(0.094)		(0.090)
	DV: Daily NO_x emissions (tons)			
	(1)	(2)	(3)	(4)
Gemini 1.0 training	0.123	0.144	0.102	0.123
	(0.080)	(0.088)	(0.089)	(0.096)
Gemini 1.5 training		0.183		0.192*
		(0.067)		(0.063)
Gemini 1.5 Flash Fine-tuning		0.146		0.153
		(0.063)		(0.058)
FE	— Year, month, power plant —			
Weather controls	N	Y	N	Y
Additional controls	— Local economic conditions —			
N	— 30,322 —			

Note: Reporting difference-in-differences results. Clustered (plant groups) SE in parentheses. $^*: p < 0.1$.

NOx and SO2: not significant

- Likely due to natural gas
- Potentially differences in scrubber uses

Event study results

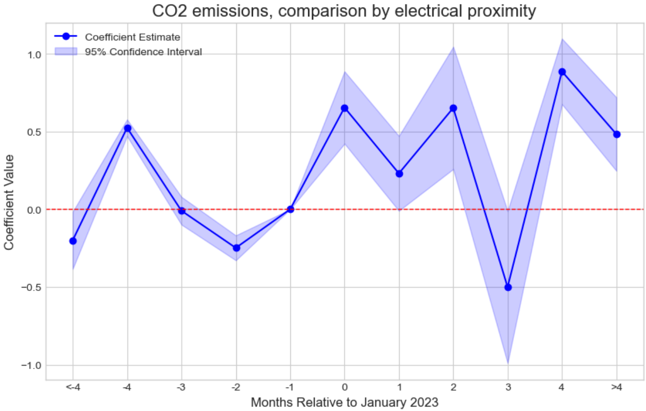

- Calculated relative to Jan 2023, due to training window uncertainty
- Pre-trends are not ideal; unsure what is occurring on April 2023

Sensitivity checks (CO2) with control group definitions

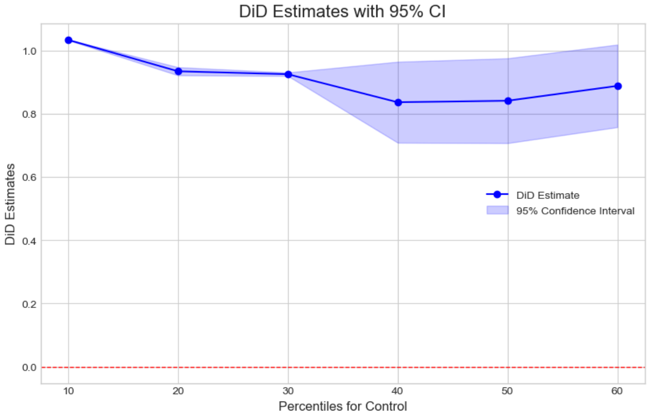

- Slight non-linearity
- In general, more stringent conditions to define a control group → greater DiD estimates

Continuous treatment

	DV: Daily CO_2 emissions (kilotons)			
	(1)	(2)	(3)	(4)
Gemini 1.0 training	0.908***	0.883***	0.863***	0.843***
	(0.057)	(0.060)	(0.060)	(0.822)
Gemini 1.5 training		0.071		0.091*
		(0.030)		(0.032)
Gemini 1.5 Flash		-0.603***		-0.554***
Fine-tuning		(0.027)		(0.028)
	DV: Daily NOx emissions (tons)			
	(1)	(2)	(3)	(4)
Gemini 1.0 training	0.100	0.116	0.078	0.094
	(0.049)	(0.054)	(0.055)	(0.060)
Gemini 1.5 training		0.139*		0.147*
		(0.051)		(0.051)
Gemini 1.5 Flash		0.109*		0.115*
Fine-tuning		(0.045)		(0.044)
FE	— Year-month, power plant —			
Weather controls	N	Y	N	Y
Additional controls	— Local economic conditions —			
N	— 39,822 —			

Note: Reporting difference-in-differences results. Clustered (plant groups) SE in parentheses. *** : $p < 0.01$, ** : $p < 0.05$.

Continuous treatment: coarsely generated from electrical proximity, normalized to be in $[0, 1]$

- Suggestions for a better continuous treatment?
- Aligned with discrete treatment results
- Slightly more NOx in the latter periods (barely stat. sig.)

Google's decarbonization pledge?

BIG
IF we are to believe the estimated effects..

Our pledge to support carbon removal solutions

Mar 14, 2024
1 min read

We plan to contract for at least $35 million worth of carbon credits in the next 12 months.

R Randy Spock
Carbon Credits and Removals Lead

Share

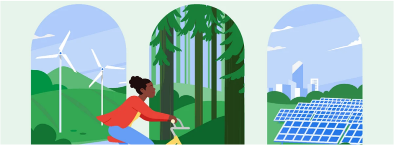

- Google pledged $35M in carbon removal credits
- Equivalent to $\sim$686k tCO_2 (using SCC $= \$51/t$) or $\sim$184k tCO_2 (using SCC $= \$190/t$)
- Extra emissions from 5 plants over 4 months: $\sim$480k tCO_2
- This equals $\sim$70% of pledge (SCC $51/t) or $\sim$260% of pledge (SCC $190/t)

Potential network or spillover effects?

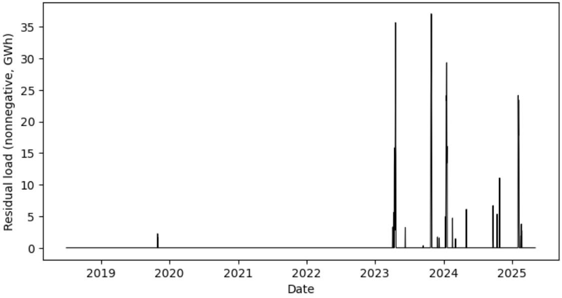

Figure: Energy imports into Bonneville Power Authority

- Potential evidence that Google's resource shifts are causing spillovers to other data centers, such as that in The Dalles, OR
- Fossil fuel energy (e.g., from ID and WY) likely to reach, say, Pacific NW due to projects like Boardman-to-Hemingway (B2H) **single-circuit line**

Hints of spillover effects

Table: Transmission Flow Effects by Plant and Training Period

	Plant 4162			Plant 8066		
Variable	CO_2	NO_x	SO_2	CO_2	NO_x	SO_2
PACE-to-BPAT, not training	0.138	0.098	0.104	0.249	0.183	0.202
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
PACE-to-BPAT, training	0.195	0.141	0.092	0.457	0.288	0.350
	(0.005)	(0.005)	(0.006)	(0.004)	(0.005)	(0.004)
Dependent variable SD	0.075	0.083	0.063	0.152	0.106	0.137
DOW FE	— Yes —					
Hour-of-day FE	— Yes —					
Year FE	— Yes —					
Month FE	— Yes —					
Number of observations	124,917			166,566		

NOTE: All variables are scaled by their respective **standard deviations**. Robust standard errors (SE) in parentheses. All estimates statistically significant at $\alpha = 0.01$. "Training" period refers to the suspected training period for Gemini 1.0, which is from June 2023 to November 2023. PACE-to-BPAT transmission is measured in MWh, CO_2 is measured in 1,000s of metric tonnes, and SO_2 and NO_x are measured in metric tonnes. All observations are at the hourly level.

Extension to Mayes County, OK

Compute

Google Cloud unveils world's largest publicly available ML hub with Cloud TPU v4, 90% carbon-free energy

May 11, 2022

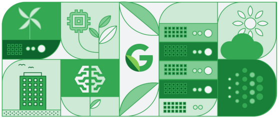

HOME › NEWS › THE ENERGY & SUSTAINABILITY CHANNEL

Google signs PPA with Leeward to power Oklahoma data center operations

The energy will be supplied from a 724MW portfolio of solar assets

January 17, 2025 By: Zachary Skidmore Have your say

Google has signed a series of Power Purchase Agreements (PPA) with Leeward Energy for energy from 724MW of solar projects under development in Oklahoma, US.

The solar portfolio comprises three large-scale installations strategically located adjacent to Google's data center in Pryor, Oklahoma.

- Location of Google's TPU v4 Pod cluster (world's largest public ML cluster, 2022)
- High energy efficiency: PUE $\sim$1.10 and $\sim$90% carbon-free grid mix
- Strong grid capacity in OK, with cheaper electricity relative to other regions
- Central U.S. geography: lower latency and good connectivity to both coasts
- Purpose-built for AI training workloads (supporting JAX, PyTorch, TensorFlow)
- Similar to IA center, Mayes County OK: a focal point for large-scale AI training

Next steps

1. Construct engineering-based electrical proximity relying on transmission + resistance for further robustness
   - Requires information from FERC Form 715
2. Incorporate Mayes County, Oklahoma data
   - Arguably the largest evidence on TPUv4/v5e deployment
   - More computational efficiency → more energy use?
3. Potential shift-share IV design for labor or health outcomes
   - Homebase data: worker-level data (hats off to Hannah Farkas!)
   - Impact through elec. price channel (e.g., Deschênes 2011)?
   - Shares as (normalized) elec. prox.; shift as training periods
4. Possible expansions to AI usage?
   - Fradkin WP: comparisons of AI usage through OpenRouter
5. TBA: **Your suggestions**!

THANK YOU!

Questions, comments, criticisms, & potential collaboration at:

jc5341@columbia.edu