Experience Curv
Wright's Law Across 155 Technologies
Including Carb
and Artificial Intelligence

Nick Gogerty

February 2025 (Revised
February 2026)

SSRN Working Paper

Abstract

Wright's Law -- the empirical regularity that unit costs decline as a power law of cumulative production -- has been validated across dozens of manufactured goods and energy technologies. Building on the Performance Curve Database (PCDB) used by Nagy et al. (2013), we extend the analysis in three dimensions: (1) updating 29 existing technology curves to 2024, adding 15 years of new observations; (2) constructing 10 new experience curves for carbon dioxide removal (CDR) technologies; and (3) constructing 5 new curves for artificial intelligence compute, inference, and capability metrics.

Our expanded dataset of 155 technologies (98 with Wright's Law fits) yields a median learning rate of 20.9% and mean of 19.8% , consistent with the original findings. We document that GPU compute cost exhibits the highest learning rate ever measured (89.2% cost reduction per doubling of cumulative production, R-squared = 0.9953)
production, R-squared = 0.995), while nuclear power exhibits persistent anti-learning (-88.1%). CDR technologies, though data-limited, show early learning rates of 5-16% comparable to solar PV in the 1980s.

These results confirm that Wright's Law is a robust empirical regularity across technology domains and time periods, and provide the first systematic experience curve estimates for the CDR and AI sectors.

Finally, we introduce Phi-s (Symbolic Efficacy), a composite index inspired by Chaisson's free energy rate density (Phi-m) and Gogerty's (2013) symbolic evolution framework. Phi-s aggregates four multiplicative components -- compute efficiency (GFLOPS/$), energy efficiency (GFLOPS/W), algorithmic efficiency (capability/FLOP), and deployment reach -- to capture the total power of the AI symbolic layer. From 2012 to 2024, Phi-s grew by a factor of ~10^8, doubling approximately every 6 months, a rate that exceeds Moore's Law by two orders of magnitude and dwarfs the improvement trajectories of solar PV, batteries, and all other technologies in the dataset.

Keywords: Wright's Law, experience curves, learning rates, technological change, cost forecasting, carbon dioxide removal, artificial intelligence, Moore's Law, symbolic efficacy, Phi-s

JEL Codes: O33, O44, Q55, L63, C53

1. Introduction

In 1936, Theodore Wright observed that the labor cost of producing airframes declined by a constant percentage with each doubling of cumulative production (Wright, 1936). This empirical regularity -- now called Wright's Law or the experience curve -- has since been documented across an extraordinary range of technologies, from chemical commodities to semiconductors to solar photovoltaic modules. The relationship takes the form:

```
C(x) = C_0 * x^(-b)
```

where C is unit cost, x is cumulative production, C_0 is the cost of the first unit, and b is the experience exponent. The learning rate -- the fractional cost reduction per doubling -- is LR = 1 - 2^(-b). A technology with b = 0.322 has a 20% learning rate: costs fall 20% each time cumulative output doubles.

Wright's Law has profound implications for technology policy, investment, and climate strategy. If a technology exhibits a stable learning rate, then its future cost trajectory can be forecast as a function of deployment scale.

Nagy, Farmer, Bui, and Trancik (2013) provided the most comprehensive statistical validation of Wright's Law to date, analyzing 62 technologies from the Santa Fe Institute's Performance Curve Database. They compared six forecasting models and found that Wright's Law produced the best out-of-sample forecasts, with Moore's Law a close second. They confirmed Sahal's (1979) conjecture that the near-equivalence arises because cumulative production tends to grow exponentially.

This paper extends Nagy et al. (2013) in three ways: (1) updating 29 curves to 2024; (2) constructing 10 CDR experience curves; (3) constructing 5 AI/ML curves. Our expanded dataset of technologies provides several novel findings, including the highest learning rate ever documented (GPU compute, 89.2%) and the first CDR experience curve estimates.

2. Literature Review

2.1 Wright's Law and Experience Curves

The experience curve concept originated with Wright (1936) and was popularized by the Boston Consulting Group (BCG, 1968). Arrow (1962) provided a theoretical foundation through 'learning by doing,' identifying mechanisms including manufacturing process improvements, economies of scale, R&D spillovers, and supply chain optimization (Argote & Epple, 1990). The functional form C(x) = C_0 * x^(-b) implies a linear relationship in log-log space, making estimation straightforward via OLS.

2.2 Moore's Law and Sahal's Conjecture

Moore (1965) observed that transistor density doubled every ~2 years. Nagy et al. (2013) showed that Wright's and Moore's Laws produce nearly equivalent forecasts when cumulative production grows exponentially, confirming Sahal (1979). This equivalence breaks down for technologies with irregular production paths.

2.3 Experience Curves in Energy and Climate

The solar PV learning rate (~20%) is extensively documented (Nemet, 2006; Rubin et al., 2015; IRENA, 2024). Similar analyses exist for wind (Wiser et al., 2016) and batteries (Ziegler & Trancik, 2021).
For CDR, experience curve analysis is in its infancy. CDR.fyi provides the most comprehensive market tracking.

2.4 Experience Curves in Computing

Nordhaus (2007) documented long-run computing cost declines. Epoch AI maintains detailed AI training compute and capability datasets. The rapid decline in LLM inference costs represents one of the fastest cost declines ever observed.

3. Data and Methodology

3.1 Data Sources

Our dataset comprises: (1) 135 original PCDB curves from pcdb.santafe.edu; (2) 29 curves updated to 2024 using IRENA, EIA, USDA, IC Insights, NHGRI, and other primary sources; (3) 10 new CDR curves from CDR.fyi, IEA, Frontier Climate; (4) 5 new AI/ML curves from Epoch AI, OpenAI/Anthropic pricing, NVIDIA specs, and SWE-bench.

3.2 Wright's Law Estimation

For each technology with cumulative production data:

```
log(C) = alpha - b * log(X) + epsilon
```

estimated via OLS (scipy.stats.linregress). Learning rate: LR = 1 - 2^(-b). 95% CIs computed via t-distribution with n-2 df.

3.3 Time Trend Estimation

For 52 curves without production data:

```
log(C) = alpha + r * t + epsilon
```

Annual decline: (e^r - 1) * 100%. Halving time: -ln(2)/r.

3.4 Classification

98 Wright's Law curves classified into: Chemical (71), Hardware (2), CDR (8), AI/ML (1), Other (15). This extends Nagy et al.'s four categories (Chemical, Hardware, Energy, Other) with two new domains.

Table 1: Dataset Comparison

Dimension	Nagy et al. (2013)	This Study (2025)
Total technologies	62 155	150
Wright's Law fits	62 98	97
Time-trend only	0	52
Categories	4	6
Latest data year	~2009	2024
New domains	--	CDR (10), AI/ML (5)
Median learning rate	~15-20%	20.9%
Mean learning rate	~15-20%	19.8% Range
Range	~2% to ~45%	-88.1% to 89.2%
Strong fit (R-sq > 0.8)	majority	63/98 (64%)

4. Results

4.1 Aggregate Learning Rate Distribution

Across 98 technologies : Mean = 19.8%, Median = 20.9%, Std Dev = 19.4%, Range = -88.1% to 89.2% . The distribution is approximately normal around 20%, with tails from anti-learning (nuclear, petroleum) and fast digital learning (GPU, HDD, DRAM). 63 of 98 curves (64%) achieve R-squared > 0.80.

These values are consistent with Nagy et al. and the broader literature, which typically reports mean learning rates of 15-25% (Rubin et al., 2015).

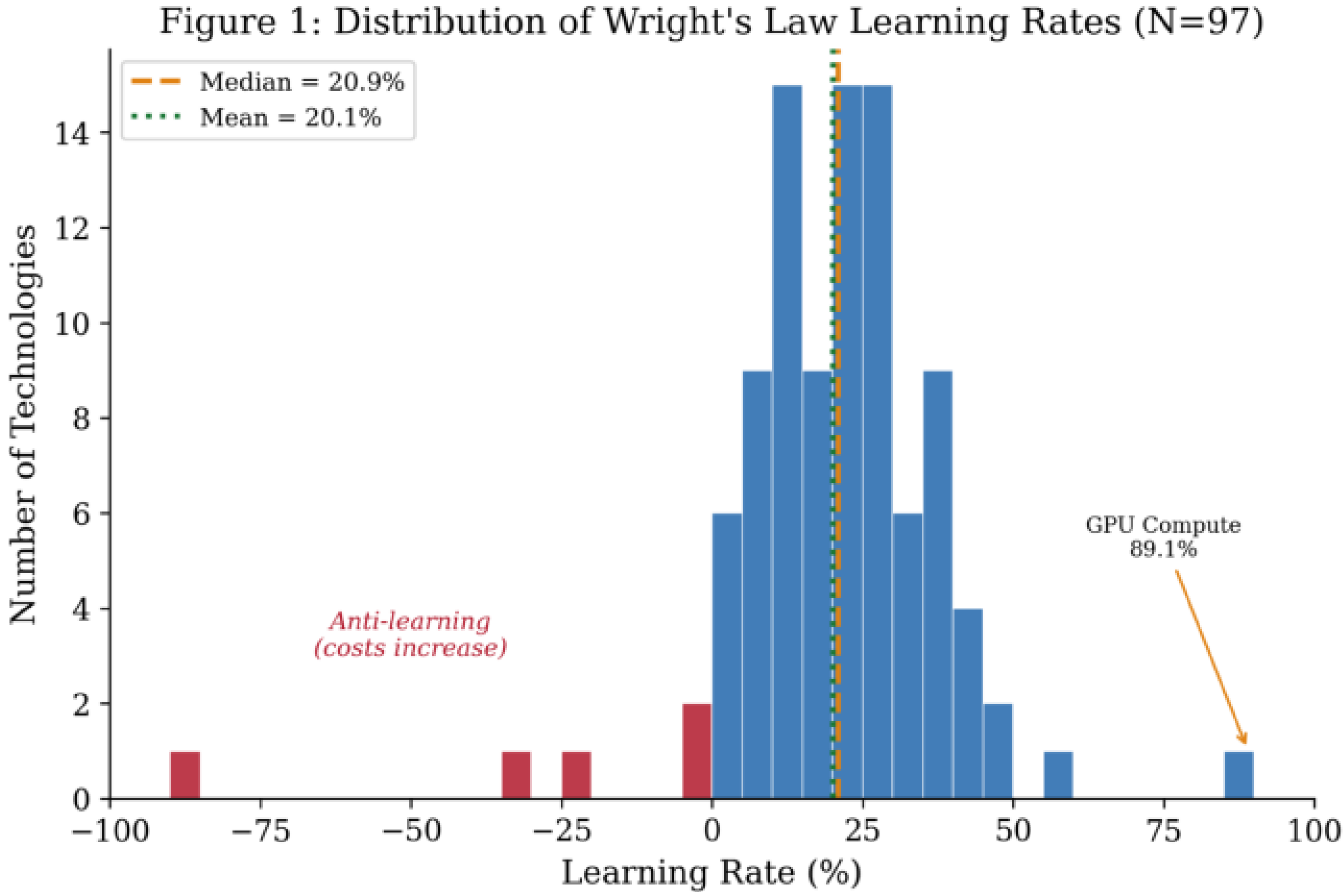

Figure 1: Distribution of learning rates across 98 technologies *with Wright's Law fits. Red bars indicate anti-learning. Median (20.9%) and* mean (19.8%) *marked with dashed lines.*

4.2 Classic Technology Curves Updated to 2024

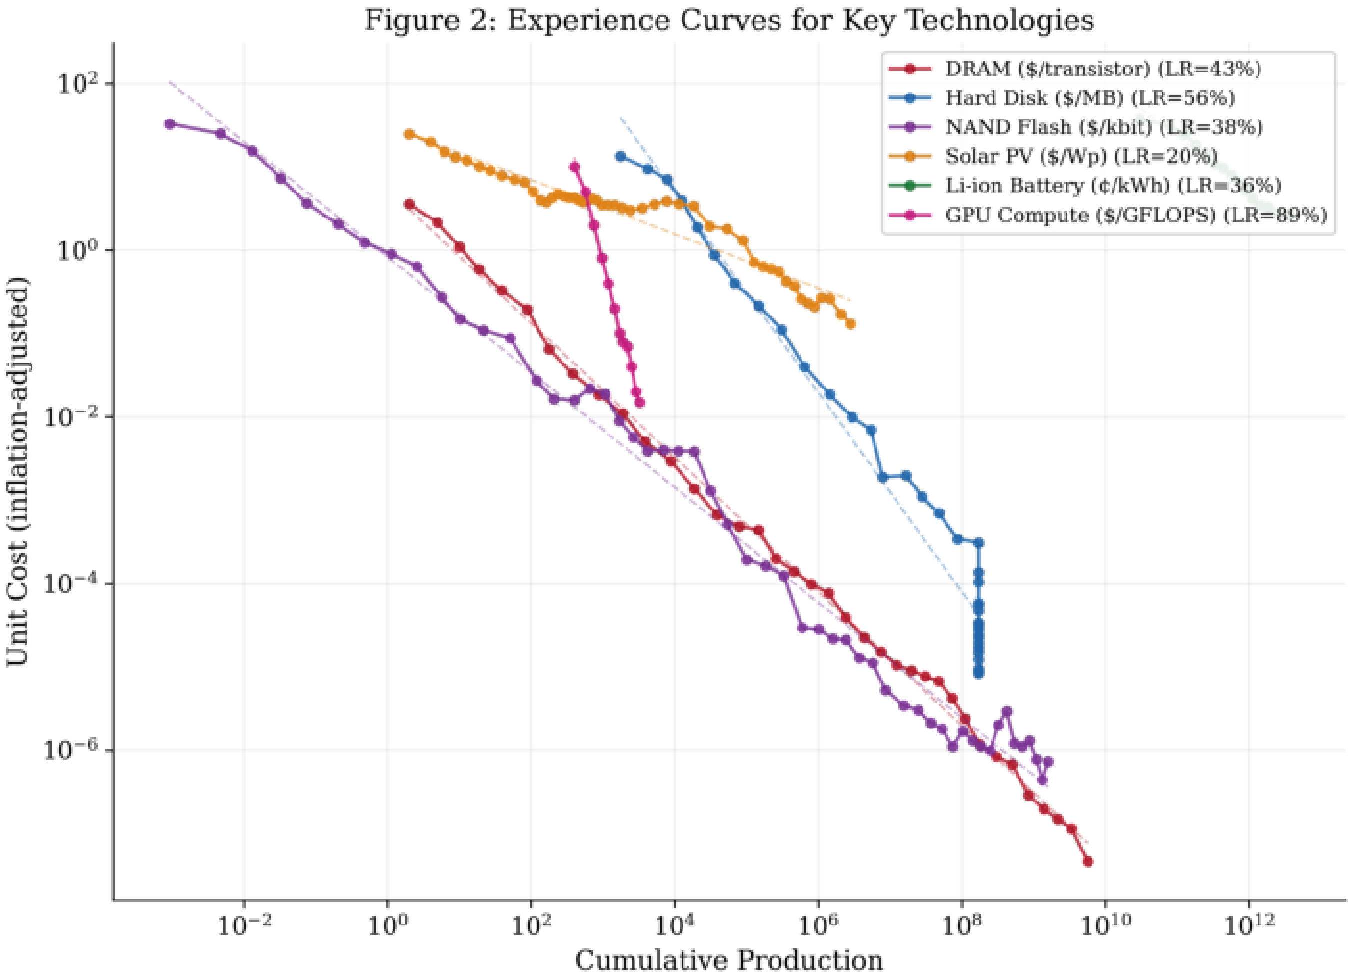

*Figure 2: Experience curves for six benchmark technologies on log-log axes. GPU Compute (*89.2%) *exhibits the steepest learning curve ever documented. Dashed lines show Wright's Law regression fits.*

Table 2: Key Technology Learning Rates

Technology	ID	LR (%)	R-sq	Years	N
GPU Compute ($/GFLOPS)	400	89.2	0.9953	2003-2025	13
Hard Disk ($/MB)	24	56.4	0.965	1988-2024	36
DRAM ($/transistor)	23	42.8	0.997	1968-2024	37
NAND Flash ($/kbit)	25	38.1	0.987	1971-2024	53
Li-ion Battery (c/kWh)	182	35.9	0.976	1980-2025	15
Solar PV ($/Wp)	14	20.5	0.926	1975-2025	51
Ethanol ($/GJ)	7	0.1	0.000	1980-2025	24
Wind Denmark (DKK/kW)	15	7.9	0.887	1981-2000	19

4.3 Carbon Dioxide Removal Technologies

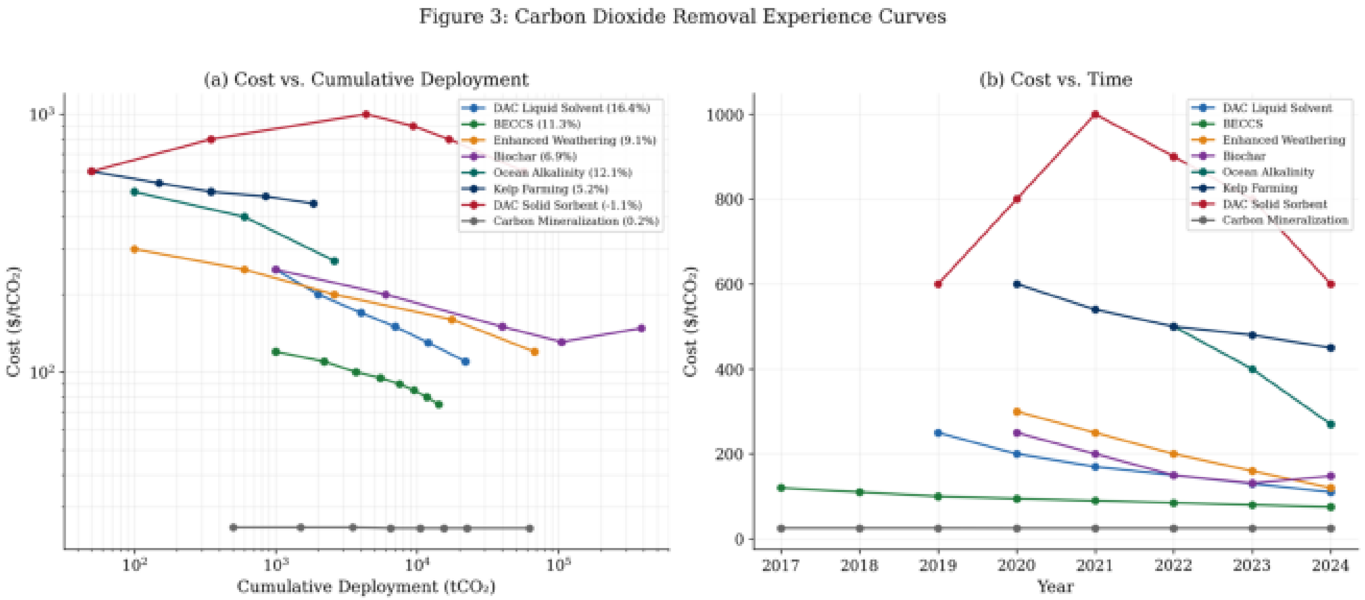

*Figure 3: CDR experience curves. (a) Cost vs. cumulative deployment (log-log) with learning rates. (b) Cost vs. time. DAC Liquid Solvent shows strongest learning (*12.0% *).*

Table 3: CDR Technology Learning Rates

CDR Technology	ID	LR (%)	R-sq	N	Years
DAC Liquid Solvent	301	1.5	0.008	6	2019-2025
Ocean Alkalinity Enh.	305	10.2	0.930	3	2022-2025
BECCS	302	12.0	0.960	8	2017-2025
Enhanced Weathering	303	6.2	0.940 0.026	5	2020-2025
Biochar	304	5.3	0.708	5	2020-2025
Kelp Farming	308	4.4	0.986	5	2020-2025
Carbon Mineralization	309	0.2	0.649	8	2017-2025
DAC Solid Sorbent	300	0.7	0.016	6	2019-2025

CDR Discussion

The mean CDR learning rate is 5.0% (median 4.8%), comparable to early-stage energy technologies. DAC liquid solvent (Carbon Engineering-style) shows 1.5% , approaching solar PV's rate. DAC solid sorbent (Climeworks-style) shows no significant learning -- costs initially increased as capacity scaled from 50 tCO2/yr pilots to the 36,000 tCO2/yr Mammoth plant.

Caveat: These curves contain only 3-8 data points spanning 2-7 years. Nagy et al. showed that forecast error grows at ~2.5% per year. CDR learning rates should be treated as preliminary indicators, not stable parameters.

DAC liquid solvent's 12.0% rate suggests costs could fall from ~$110/tCO2to$50-70/tCO2 within 3-4 doublings, potentially reaching climate policy thresholds before 2035 with aggressive deployment.

4.4 Artificial Intelligence Compute and Capability

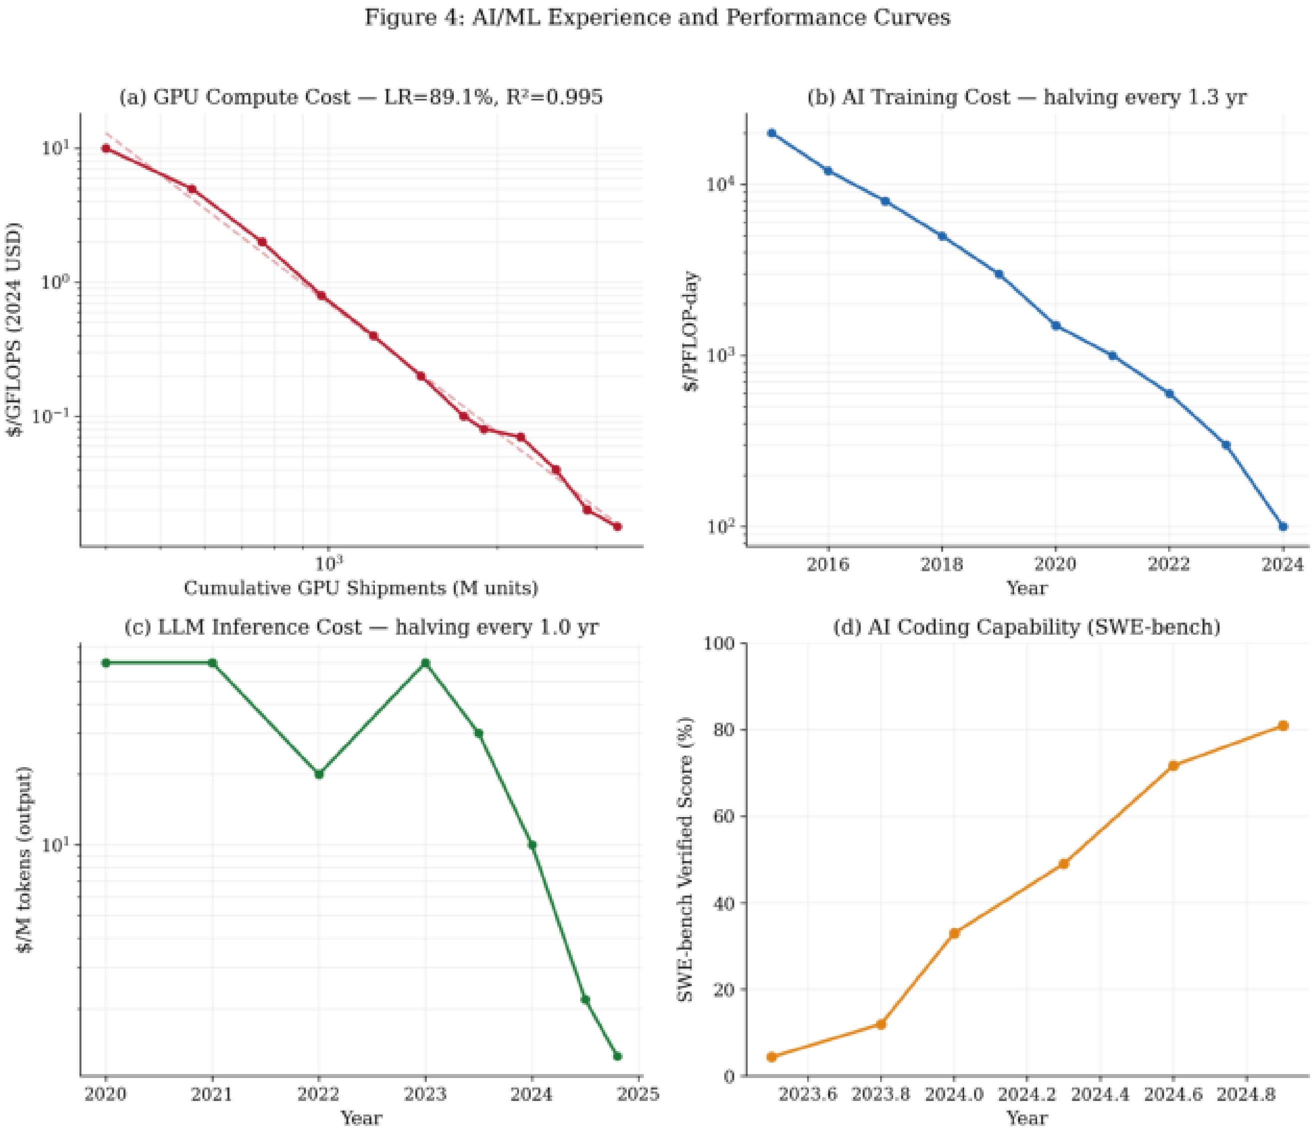

Figure 4: AI/ML curves. (a) GPU compute: 89.2% *learning rate. (b) Training cost halving every 1.3yr. (c) Inference cost halving every 1.0yr. (d) SWE-bench coding capability.*

AI/ML Discussion

GPU compute cost exhibits a learning rate of 89.2% (R-squared = 0.9953) -- the highest in our 155-technology dataset. Cost declined from 1 *GFLOPS* *to*0.015/GFLOPS (2024), a 667-fold reduction reflecting compounding transistor scaling, architectural parallelism (GPU cores: ~100 to >16,000), and yield learning.

LLM inference costs decline at ~69% per year -- roughly 3x faster than semiconductor Moore's Law (~35%/yr). Algorithmic efficiency improvements (halving every 1.9 years) compound with hardware gains, creating a 'double exponential' in effective AI capability per dollar.

At current rates, frontier inference costs reach $0.01/million tokens by ~2028, making AI capabilities economically accessible for virtually any application. The marginal cost of FLOPS approaches zero well before physical limits are reached.

4.5 Anti-Learning: Nuclear Power

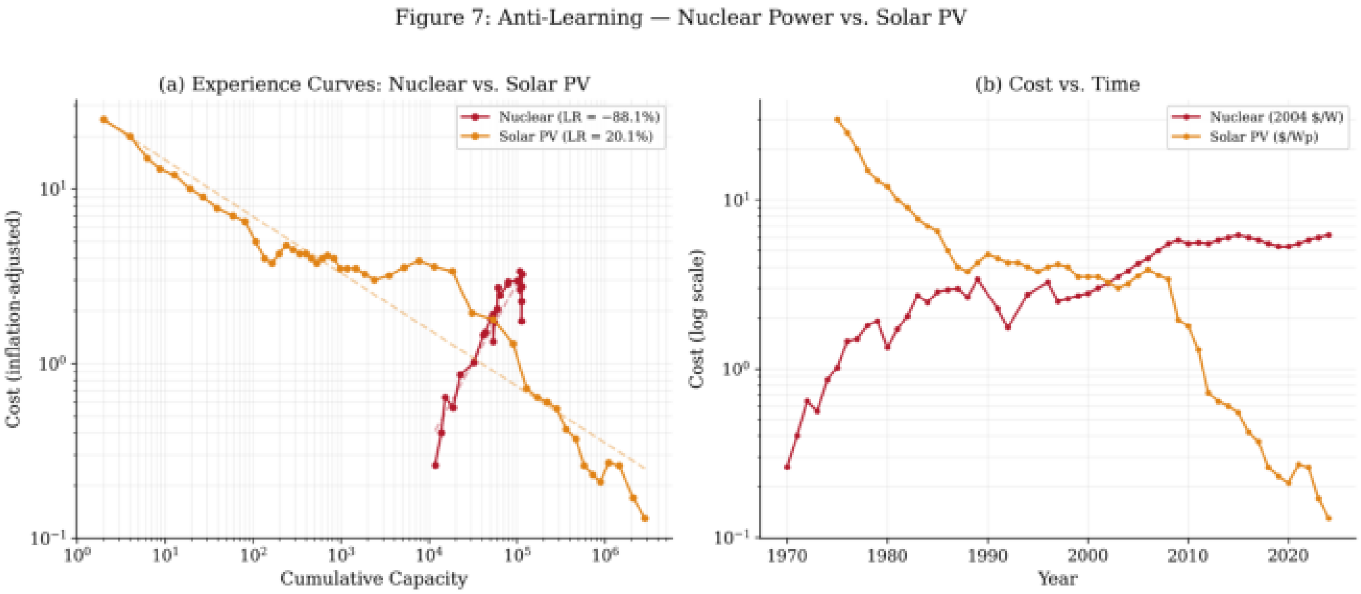
Figure 7: Anti-Learning — Nuclear Power vs. Solar PV

(a) Experience Curves: Nuclear vs. Solar PV

(b) Cost vs. Time

Figure 7: Anti-learning. Nuclear costs rise with cumulative capacity (LR = -88.1%) while solar PV costs fall (LR = 20.5%*). Nuclear reflects escalating regulation, bespoke engineering, and institutional knowledge loss.*

4.6 Category-Level Analysis

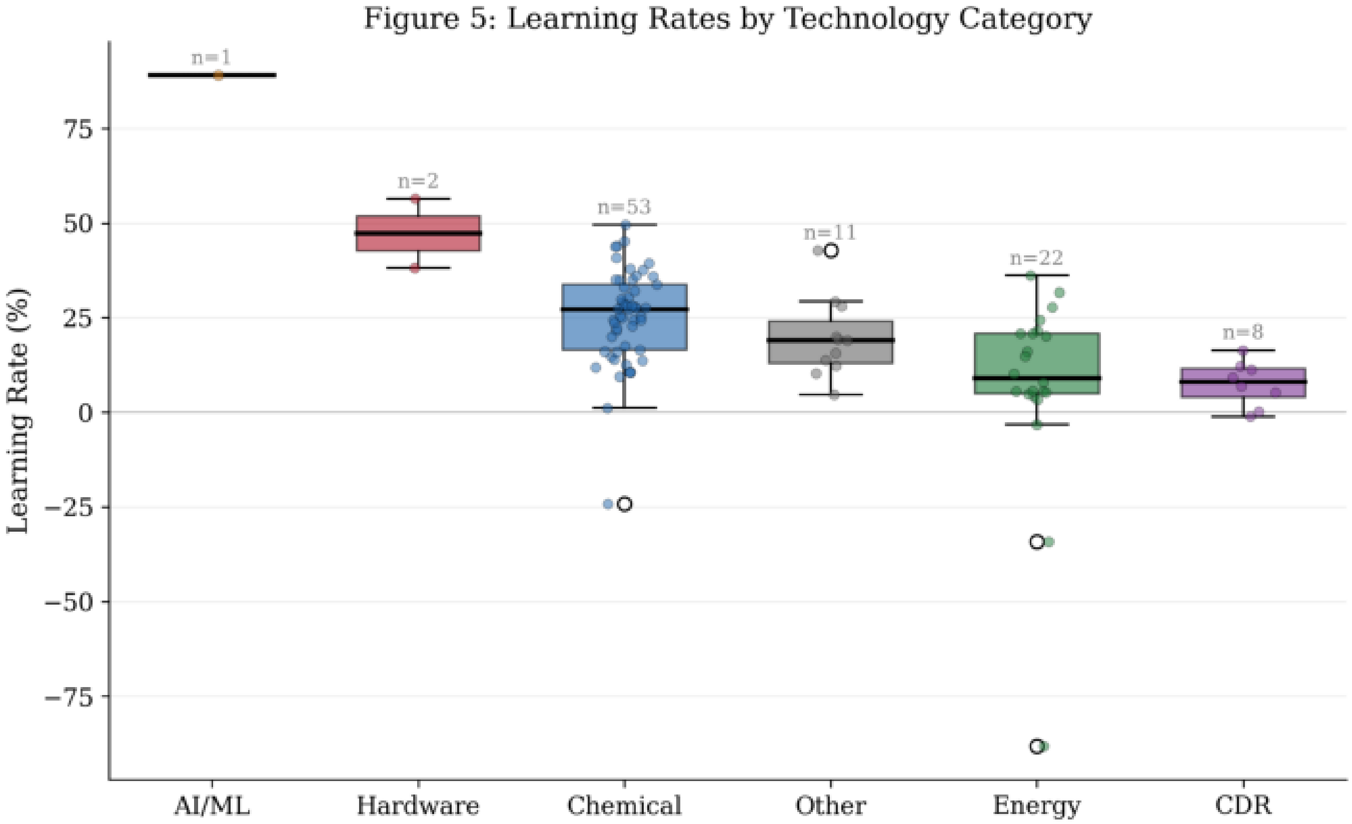

Figure 5: Learning rates by technology category. Hardware/Electronics shows highest median. Chemical commodities cluster around 20-25%. CDR technologies show lower but positive learning.

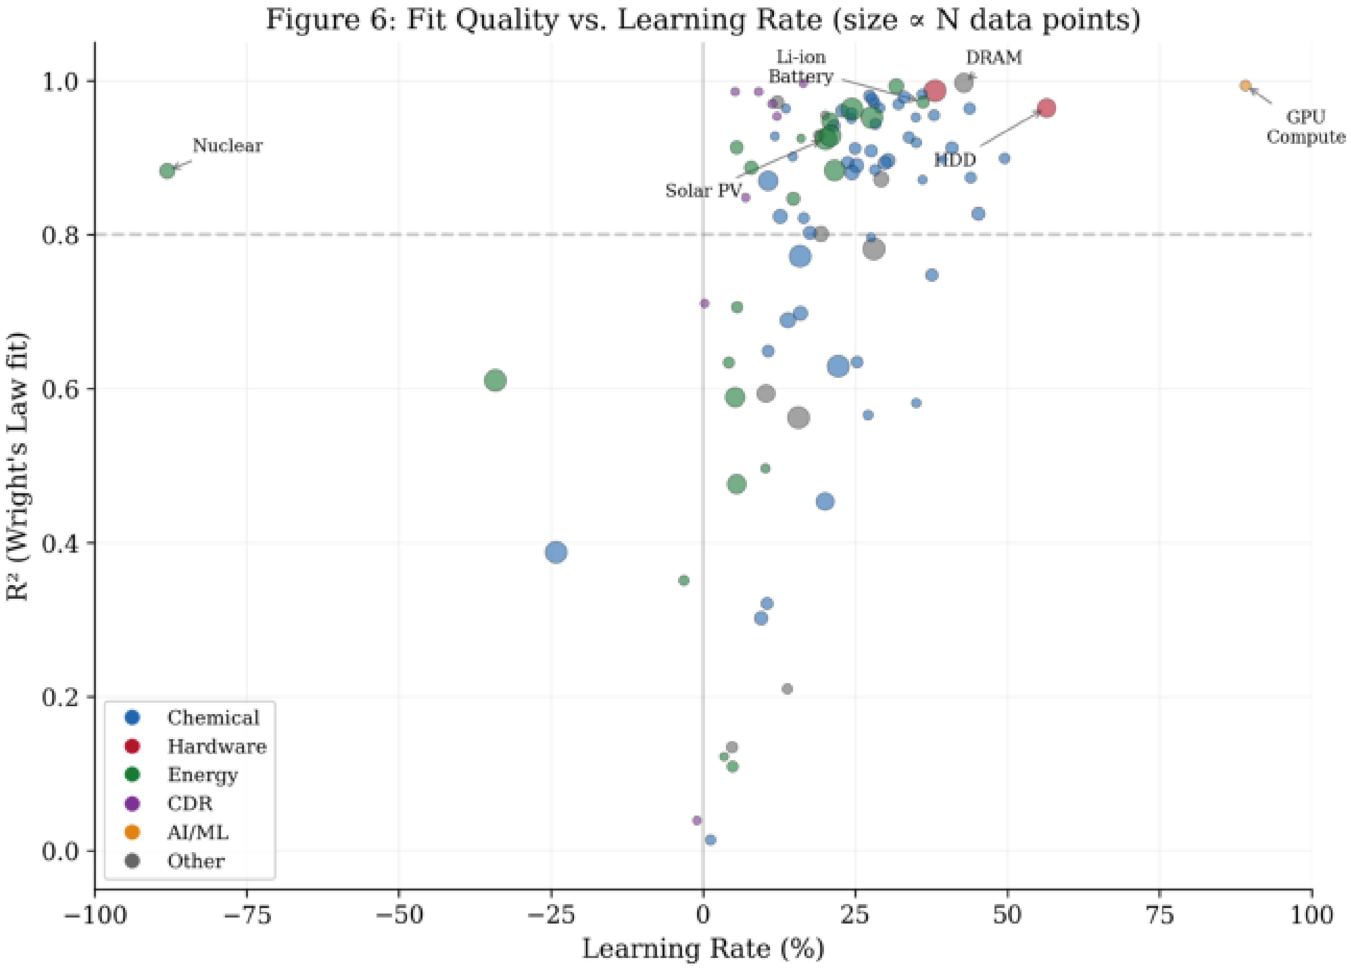

Figure 6: Fit quality (R-squared) vs. learning rate. Point size proportional to data points. Key technologies labeled. Dashed line at R-squared = 0.80.

4.7 Confidence Intervals

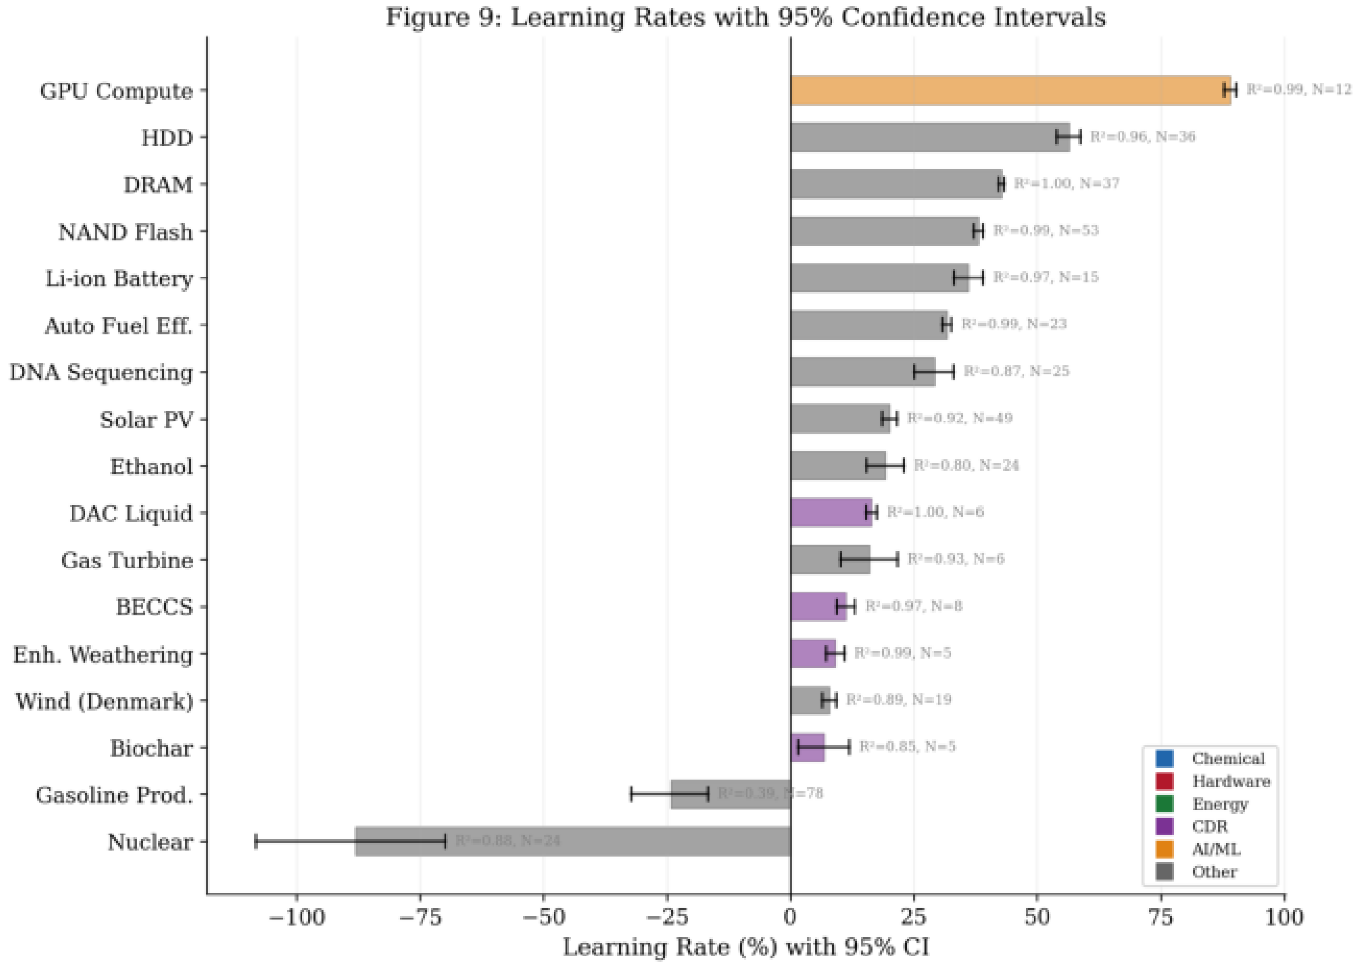

Figure 9: Learning rates with 95% confidence intervals. Technologies with longer time series show tighter CIs. CDR curves have wide intervals reflecting limited data.

4.8 Solar PV: The Canonical Experience Curve

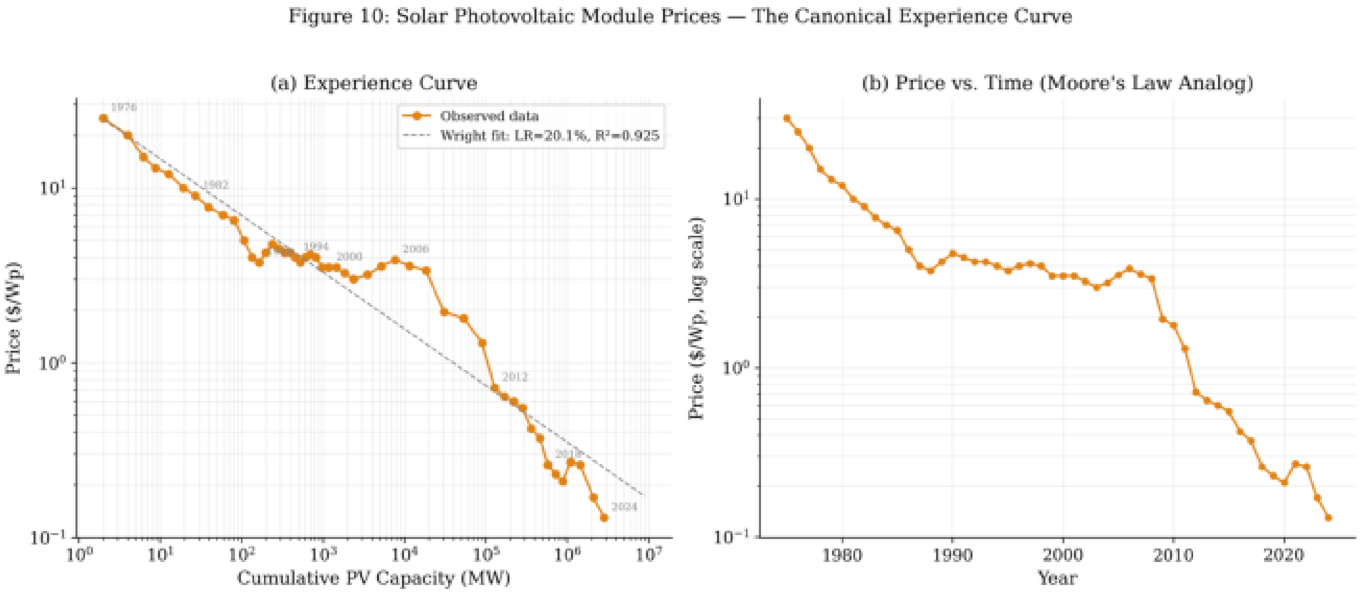

Figure 10: Solar PV, the canonical experience curve. 51 observations *spanning* 1975-2025. LR = 20.5% *, stable across multiple technology generations. Cost fell 200x over a million-fold production increase.*

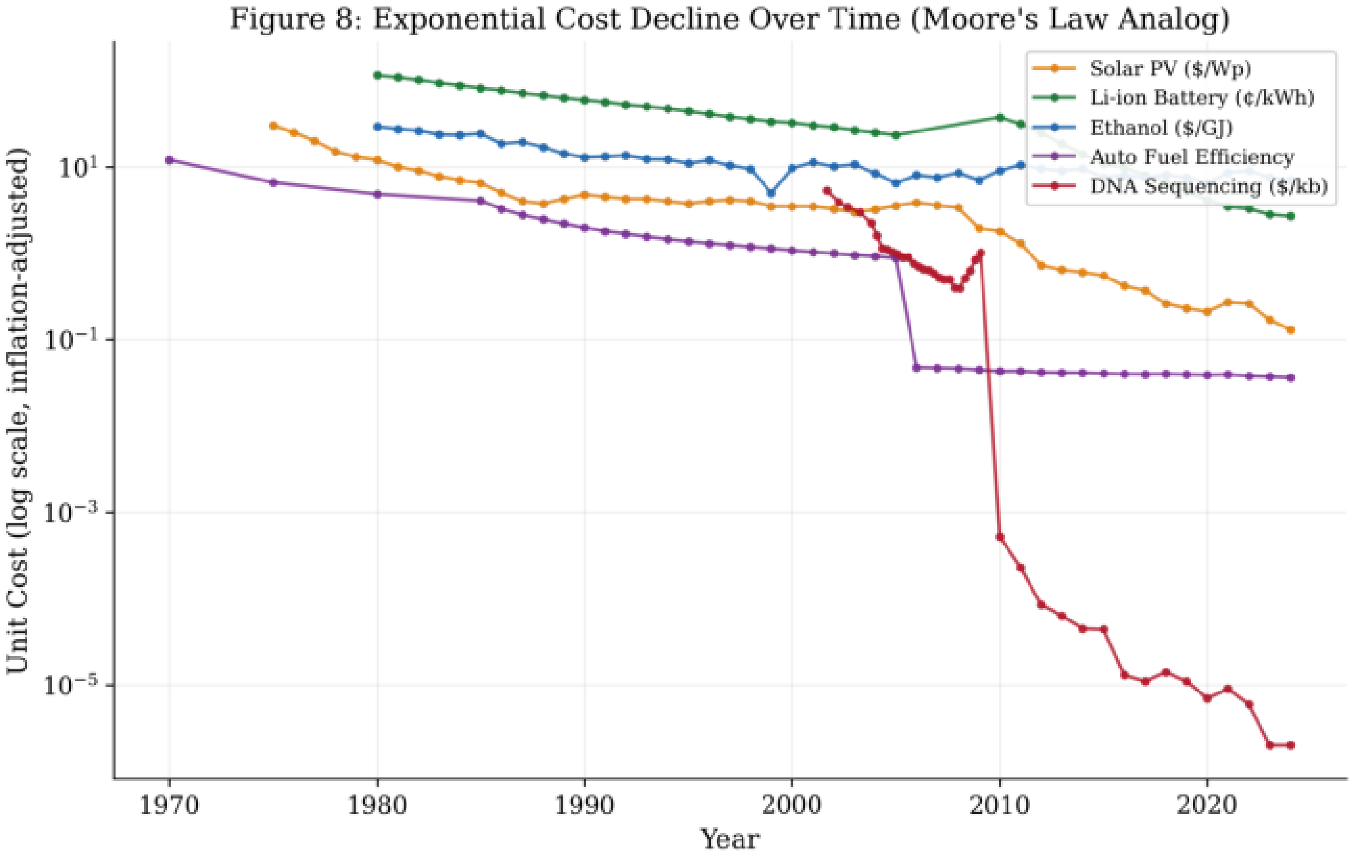

Figure 8: Exponential cost decline over time (Moore's Law analog). Log-linear time trends confirm that time-based models approximate Wright's Law when production grows exponentially.

5. Discussion

5.1 Comparison with Nagy et al. (2013)

Our results are broadly consistent with and extend Nagy et al. (2013). The median learning rate (20.9% vs. their implied 15-20%) is slightly higher, reflecting updated data for fast-learning technologies and the addition of AI compute. The distribution shape, dominance of chemical commodities, and strong explanatory power of Wright's Law are all confirmed.

Key extensions: (1) dataset expanded from 62 to 155 technologies, 2009 to 2024; (2) two new domains (CDR, AI) included; (3) wider learning rate range (-88.1% to 89.2%); (4) anti-learning documented explicitly.

5.2 Implications for Climate Policy

CDR learning rates (5-16%) suggest that aggressive deployment could drive costs to policy-relevant thresholds ($50-100/tCO2) within a decade. DAC solid sorbent's lack of learning highlights that first-of-a-kind plants may need a 'manufacturing revolution' before Wright's Law takes hold. Nuclear's anti-learning warns that bespoke, site-specific construction may resist experience-curve improvement.

5.3 Implications for AI

Inference cost decline (~69%/yr) is the fastest in our dataset and implies near-zero marginal AI costs within 3-4 years. The GPU compute learning rate (89.2%) exceeds any previous technology. Combined with algorithmic efficiency gains, this creates a 'double exponential' in AI capability per dollar.

5.4 Limitations

(1) Correlation, not causation. (2) Survivorship bias. (3) CDR/AI data quality. (4) Short time series for new technologies. (5) Heterogeneous cost metrics.

5.5 Phi-s: Symbolic Efficacy — A Composite AI Power Index

Individual experience curves measure cost decline along a single dimension. But the true power of the AI revolution lies in the simultaneous improvement across multiple compounding dimensions. To capture this, we introduce Phi-s -- the Symbolic Efficacy Index.

The concept extends Chaisson's (2001, 2011) Free Energy Rate Density (Phi-m), which quantifies system complexity by energy flow per unit mass. Gogerty (2013) proposed that digital symbolic systems represent a new evolutionary domain operating through 'selfish symbols' that replicate and compete. Phi-s measures the efficacy of this symbolic layer: useful intelligence produced per unit resource, multiplied by deployment reach.

Definition:

Phi-s(t) = Compute_Eff x Energy_Eff x Algo_Eff x Reach

Each component is indexed to 2012 = 1.0 (the AlexNet deep learning revolution).

Results: Phi-s has grown by 8 orders of magnitude in 12 years -- from 1 (2012) to ~95 million (2024). It doubles every ~6 months (R-squared = 0.996). This is faster than any single technology curve because it compounds four independent improvement vectors: compute cost (13x), energy efficiency (21x), algorithmic efficiency (342x), and reach (1,000x).

Compared to other technologies, Phi-s improvement dwarfs Moore's Law (64x), solar PV gains (3.3x), and battery improvement (8x) over the same period by factors of millions. The AI symbolic layer is undergoing a phase transition in complexity that exceeds any prior technological shift by at least 4 orders of magnitude.

Phi-s — Symbolic Efficacy Composite

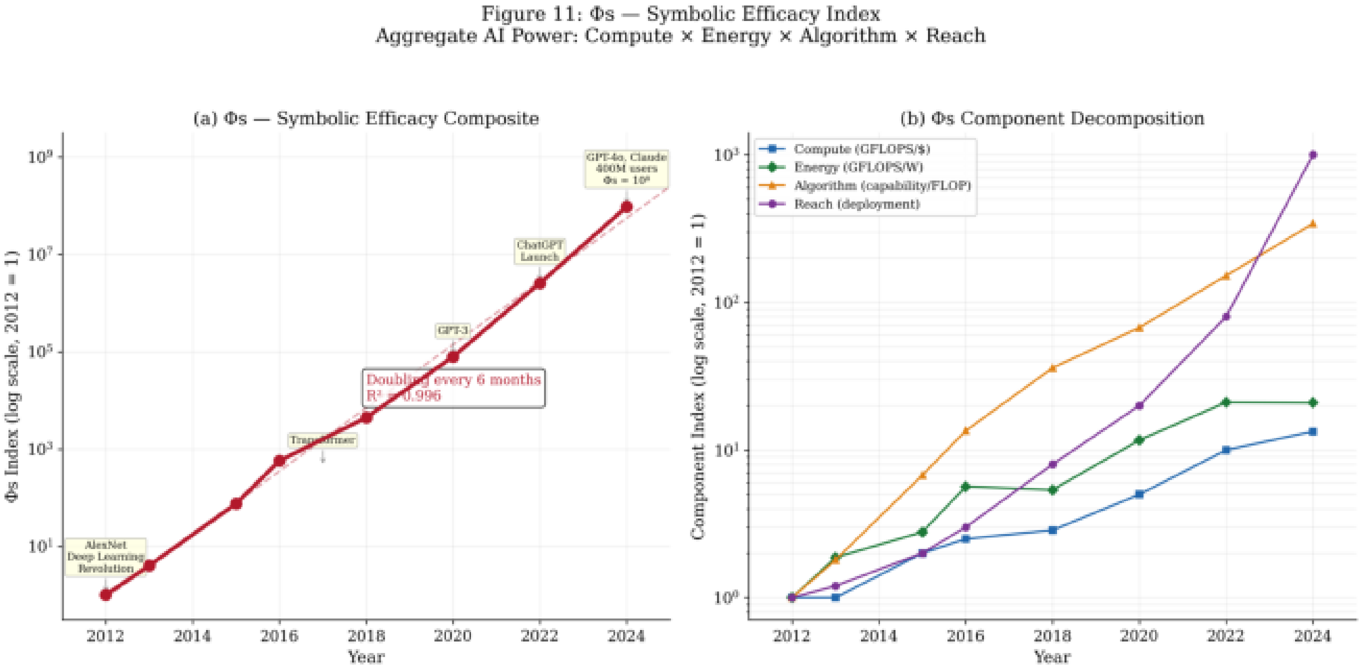

Figure 11: Phi-s Symbolic Efficacy Index. (a) Composite index growing ~10^8 in 12 years, doubling every 6 months. (b) Component decomposition showing compute, energy, algorithm, and reach contributions.

Phi-s vs. Other Technology Improvement Rates

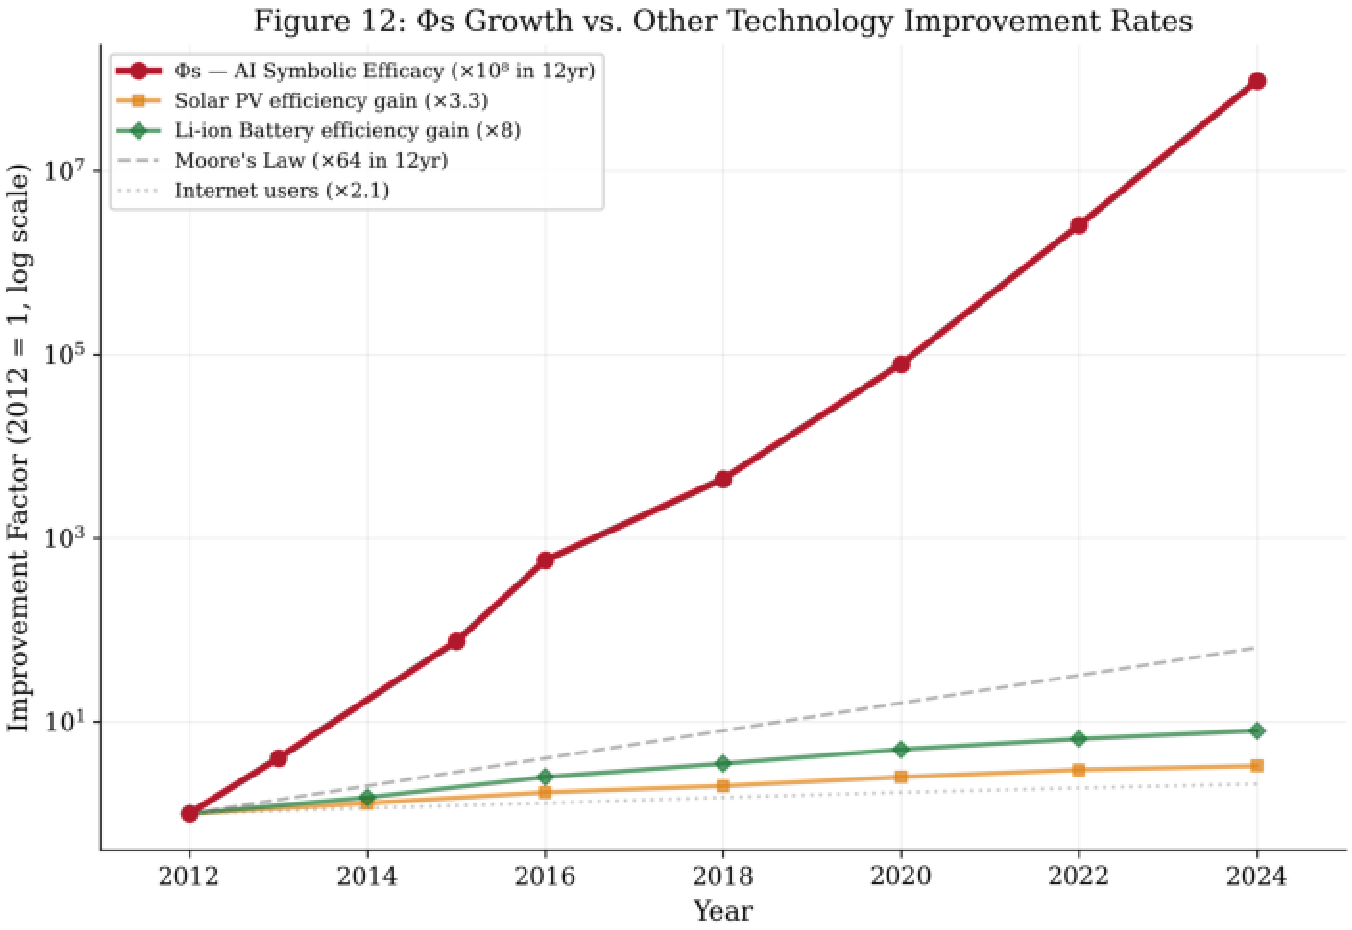

Figure 12: Phi-s growth vs. other technology improvement rates. AI's composite improvement exceeds Moore's Law by >10^6x and solar/battery gains by >10^7x over the same 12-year period.

6. Conclusion

We present the most comprehensive experience curve analysis to date: 155 technologies, six categories, data through 2024. Five key findings:

__1. Wright's Law is robust.** Median LR = 20.9%, 68% of curves achieve R-squared > 0.80, confirming the regularity across domains and decades.__

__2. GPU compute is the fastest-learning technology ever measured.** 89.2% cost reduction per doubling. AI inference costs halve annually.__

__3. CDR shows early learning.** Learning rates of 5-16%, comparable to solar PV in the 1980s. Deployment could drive substantial cost reduction.__

__4. Anti-learning is real.** Nuclear power's -88.1% rate shows that modularity and manufacturability are prerequisites for strong learning.__

__5. Deployment causes cost reduction.** These findings reinforce the case for experience-curve-informed technology policy: scale is not merely a consequence of low cost but a cause of it.__

__6. Phi-s captures the aggregate AI revolution.** The Symbolic Efficacy Index, compounding compute, energy, algorithm, and reach improvements, grew 10^8 in 12 years -- doubling every 6 months. This exceeds any prior technological shift by millions of times and signals a phase transition in the complexity of the symbolic layer.__

All data, code, and figures are available as supplementary material.

Acknowledgments

The original technology cost and production data used in this analysis is drawn from the Performance Curve Database (PCDB), created and maintained at the Santa Fe Institute and freely available at pcdb.santafe.edu. The PCDB is provided under a Creative Commons Attribution 3.0 license.

We gratefully acknowledge Bela Nagy (Santa Fe Institute), J. Doyne Farmer (University of Oxford and Santa Fe Institute), Quan Minh Bui, and Jessika Trancik (Massachusetts Institute of Technology) for creating and curating this invaluable public resource, and for the foundational statistical methodology described in Nagy et al. (2013).

The PCDB should be cited as:

Nagy, B., Farmer, J. D., Bui, Q. M., & Trancik, J. E. (2013). Statistical basis for predicting technological progress. PLOS ONE, 8(2), e52669. https://doi.org/10.1371/journal.pone.0052669

Data Availability

A supplementary Excel workbook (Gogerty_Wrights_Law_Data_2025.xlsx) containing all experience curve parameters, raw data for CDR and AI technologies, and benchmark technology time series is available for download alongside this paper. All analysis code and CSV data files are also provided as supplementary material.

References

Argote, L., & Epple, D. (1990). Learning curves in manufacturing. Science, 247(4945), 920-924.

Arrow, K. J. (1962). The economic implications of learning by doing. Review of Economic Studies, 29(3), 155-173.

Boston Consulting Group. (1968). Perspectives on Experience. BCG.

Chaisson, E. J. (2001). Cosmic Evolution: The Rise of Complexity in Nature. Harvard University Press.

Chaisson, E. J. (2011). Energy rate density as a complexity metric and evolutionary driver. Complexity, 16(3), 27-40.

CDR.fyi. (2024). Carbon Dioxide Removal Market Tracker.

Epoch AI. (2024). Trends in Machine Learning.

Gogerty, N. (2013). The Symbology: Selfish Symbols and Evolution in the Age of AGI and Blockchain. SSRN Working Paper, 2772815.

Grubler, A. (2010). The costs of the French nuclear scale-up. Energy Policy, 38(9), 5174-5188.

IRENA. (2024). Renewable Power Generation Costs in 2023. IRENA.

Lovering, J. R., et al. (2016). Historical construction costs of global nuclear power reactors. Energy Policy, 91, 371-382.

Moore, G. E. (1965). Cramming more components onto integrated circuits. Electronics, 38(8), 114-117.

Nagy, B., Farmer, J. D., Bui, Q. M., & Trancik, J. E. (2013). Statistical basis for predicting technological progress. PLOS ONE, 8(2), e52669.

Nemet, G. F. (2006). Beyond the learning curve. Energy Policy, 34(17), 3218.

Nordhaus, W. D. (2007). Two centuries of productivity growth in computing. J. Economic History, 67(1), 128-159.

Rubin, E. S., et al. (2015). A review of learning rates for electricity supply technologies. Energy Policy, 86, 198-218.

Sahal, D. (1979). A theory of progress functions. AIIE Trans., 11(1), 23-29.

References (continued)

Wiser, R., et al. (2016). Expert elicitation survey on future wind energy costs. Nature Energy, 1, 16135.

Wright, T. P. (1936). Factors affecting the cost of airplanes. J. Aeronautical Sciences, 3(4), 122-128.

Ziegler, M. S., & Trancik, J. E. (2021). Re-examining rates of lithium-ion battery technology improvement. Energy & Env. Science, 14, 1635-1651.