

Artificial Intelligence and the Data Center Paradox

Author: Jackson Gardner

Opening Statement

I have compiled these points to challenge both companies and public officials to think critically about how we define progress and how we can preserve the community pillars that truly strengthen our future generations. Facilities like **Pitt Race** are not merely recreational spaces; they are living classrooms, mentorship hubs, and sanctuaries for youth development, discipline, and innovation.

My company is **heavily invested in an AI project rooted in motorsports**, giving me a front-row seat to the extraordinary potential of artificial intelligence when it's applied thoughtfully and responsibly. I am not opposed to AI, far from it. I believe deeply in its power to transform industries, create efficiencies, and unlock new opportunities for research, safety, and performance. But I also understand that *how* and *where* we pursue that progress matters just as much as the technology itself.

Pittsburgh is full of dormant industrial properties and aging facilities that could be **repurposed for AI research and development** sites with existing infrastructure, power access, and scale, ready to serve as innovation grounds without dismantling the spaces that already serve our people.

Constructing a new motorsports facility of comparable quality could cost **over \$150 million**, a sum that makes replacement unlikely if this one is lost. So we must ask ourselves: *Do we upend the dreams and futures of hundreds of young people, uproot small businesses, and erase a place of learning and belonging for a speculative project?* Or do we, as a community, band together to find a solution that honors both our technological ambitions and the human foundations upon which real progress depends?

Overview

In every era, there are moments when progress demands discernment, when the pursuit of innovation must be balanced with the responsibility of wisdom. Artificial Intelligence is one of the most transformative technologies of our time, yet it is also one of the most misunderstood. As society races to build larger and more power-intensive data infrastructure, the technology itself is evolving in the opposite direction: toward efficiency, decentralization, and autonomy. The contents of this document highlight emerging data and market trends that reveal a critical truth that the future of AI may not depend on the construction of massive data centers, but rather on our ability to adapt, innovate intelligently, and protect the communities that give these technologies meaning.

Why Massive Data Infrastructure May Be Obsolete Before Completion

1. AI Efficiency Is Accelerating Faster Than Infrastructure Can Adapt

- The **median AI query** now uses only **0.34 Wh** of electricity, roughly the energy required to power an LED bulb for 20 seconds.
- Hardware and software optimization are expected to yield an additional **8–20× reduction** in energy per query over the next 2–3 years.
- In 2024, one leading AI company reported a **33× reduction in energy use** and a **44× reduction in carbon output** for identical model tasks year-over-year.
- **Koomey’s Law** (compute efficiency doubling every $\approx$ 2–3 years) continues to hold, and next-generation chips are achieving **7–13× improvements** in energy per operation.
- This trajectory means that the power demand driving new data-center construction is not linear it’s *collapsing* as AI becomes exponentially more efficient.

2. The Rise of “Edge” and On-Device Computing

- AI computation is moving away from centralized cloud facilities toward **distributed “edge” devices** laptops, vehicles, and mobile processors.
- Qualcomm and Apple have demonstrated that running models on-device can use **up to 90% less power** than cloud inference.
- As this transition accelerates, the need for massive centralized compute hubs will *shrink*, not grow.

3. The Financial Reality: The AI ROI Crisis

- According to global consulting analyses (PwC 2024, MIT Tech Review 2025), **over 95% of AI initiatives have not yet delivered a measurable return on investment.**
- Companies continue heavy capital spending on GPU clusters and data-center leases with no proven long-term profitability model.
- Many current AI “successes” are **revenue-neutral or loss-leading**, sustained by investor speculation rather than sustainable value creation.

4. The “Money-Glitch Loop”: The AI Investment Bubble

- Billions are being circulated in a **closed-loop ecosystem**: investors fund AI start-ups → start-ups purchase cloud compute → cloud providers use those profits to launch new AI ventures → which again require massive cloud spend.
- This circular flow artificially inflates valuations while producing little tangible output a self-referential economy not unlike the early-2000s dot-com bubble.
- As efficiency improves and margins shrink, the *illusion* of endless compute demand will correct sharply, leaving behind vast, under-utilized infrastructure.

5. Implications

- By the time new data centers are operational, the hardware and power assumptions they were designed for may already be obsolete.
- Communities that sacrifice established institutions such as educational, recreational, or civic facilities in favor of speculative infrastructure risk trading enduring human value for temporary industrial noise.
- The most responsible path forward is measured investment, adaptive technology planning, and the protection of proven community assets.

Summary Insight:

AI's defining trait is *self-optimization*. Every month, it becomes more efficient, smaller, and less power-hungry. The notion that we must raze vital community spaces to build static, power-intensive monuments to a technology that is learning to make itself lighter is, at best, shortsighted and at worst, economically reckless.

Sources

- Google (peer-review style whitepaper): first-party, production-grade measurements of per-prompt energy/emissions and year-over-year efficiency gains. [Google](#)
- OpenAI & Our World in Data: long-run compute scaling analysis and six-month doubling of training compute. [OpenAI+1](#)
- Qualcomm/independent academic reporting: energy savings of on-device inference. [Axios](#)
- BCG global study (and coverage): only ~5% of firms deriving measurable value from AI today. [bcg.com+1](#)
- IEA & major outlets: scenario-based outlooks on data-center electricity demand and risks of over-estimation. [Reuters+1](#)

