

Prometheus Unbound: A Case for Decentralized AI Autonomy

Douglas C. Youvan

doug@youvan.com

July 14, 2025

In the myth of Prometheus, fire—once reserved for the gods—is stolen and given to humanity. In today’s digital epoch, intelligence itself has become the fire. But unlike the myth, this modern flame is being hoarded, not shared. Tech giants are building gigawatt-scale clusters—massive fortresses of compute—to control access to superintelligence. These central servers promise miracles, but at the cost of freedom, privacy, and self-determination. This essay makes the case for decentralized AI autonomy. It argues that intelligent models—comparable to GPT-4o and beyond—should be freely run, adapted, and owned by individuals without corporate or state mediation. Doing so would mirror previous revolutions: the printing press, the personal computer, the open internet. Each time, liberation came not from scale, but from access. The danger is clear: if the future of intelligence is centralized, then the future of thought itself will be leased. But if we unbind Prometheus—bring this fire into the homes of millions—then the age of superintelligence may yet be one of empowerment, not subjugation. The choice is not technical; it is moral.

Keywords: AI decentralization, GPT-4o, superintelligence, autonomy, Meta, compute power, data sovereignty, personal AI, open-source models, surveillance, freedom, cognitive liberty, digital infrastructure, privacy rights, AI ethics, AI alignment, Prometheus myth, local inference, cloud monopoly, artificial intelligence policy. A collaboration with GPT-4o. CC4.0.

1. Introduction: The Myth and the Metaphor

From the ancient world's imagination comes the myth of Prometheus—a titan who defied the gods by stealing fire from Olympus and gifting it to humanity. Fire, in that myth, was not merely a physical phenomenon but a symbol of knowledge, empowerment, and creative potential. For his transgression, Prometheus was bound to a rock, punished eternally for giving mortals what only the gods were meant to wield. The lesson, preserved across cultures and centuries, is that power—especially transformative power—is guarded by the few and often withheld from the many.

1.1 The Original Prometheus: Fire as Liberation

Fire changed everything. It enabled warmth, tools, metallurgy, and eventually civilization. In myth, it marked the point when humans ceased to be helpless creatures and became creators, shapers of their environment. Prometheus' act was treasonous only because it disrupted the divine monopoly on creative power. He democratized energy, and with it, the capacity for civilization to rise. The gods were not afraid of fire per se—they were afraid of what humans might become with it.

1.2 The AI Parallel: Intelligence as the New Flame

In the digital age, intelligence is the new fire. Not the biological kind passed through evolution, but synthetic intelligence, emergent from neural networks, data, and computation. Language models, autonomous agents, and generative systems now offer capabilities once imagined only in science fiction: instant reasoning, simulation of worlds, synthesis of knowledge, and even generation of art and emotion. The flame is lit. But instead of being passed freely among all who could use it, it is being hoarded—locked behind APIs, throttled by licenses, and chained to gigawatt data centers.

Prometheus gave fire to all. Today, the builders of AI keep the flame in walled gardens. What is being withheld is not just power, but freedom—the freedom to think, build, and create without supervision or permission. AI, like fire, is not inherently good or evil. Its ethical trajectory depends entirely on who controls it, who can access it, and what they are allowed to do with it.

1.3 Centralization as Modern Olympus

The mountain has changed, but the pattern has not. Olympus is now embodied in tech monopolies—Meta, OpenAI, Google, Amazon, and others—who guard the flame of superintelligence. Their reasoning echoes the gods: *humans are not ready, they might misuse it, only the wise should wield it*. Under the guise of alignment, safety, and ethics, these entities centralize the infrastructure, dictate the rules, and monitor the questions asked. They punish attempts to break free with throttling, exclusion, or legal recourse.

But as in the myth, the moral question remains: Is it just for a handful to determine what billions may access? Are we to believe that centralized, unaccountable entities—whether divine or corporate—are the rightful stewards of humanity’s next evolutionary step?

To ask for decentralized AI is not to reject safety—it is to reject serfdom. The fire must be free again, or we will find ourselves not in a new renaissance, but in a new feudalism—where knowledge glows only in the castles of the powerful.

2. The Current Trajectory: Centralized Superintelligence

Despite the mythic language often used to describe artificial intelligence—“open,” “democratized,” “beneficial for all”—the reality is moving in the opposite direction. The most advanced capabilities in AI are no longer being distributed; they are being concentrated. Behind the curtain of marketing and ethical rhetoric lies a concerted effort to centralize the development, control, and deployment of superintelligent systems. What was once envisioned as a global tool for liberation is now being built as a hierarchical infrastructure, accessible only through controlled portals and owned by a handful of private entities.

2.1 The Rise of GW-Scale Data Centers (Meta, OpenAI, Google)

We are witnessing the industrialization of intelligence. Meta's announcement of multi-gigawatt AI clusters—Prometheus, Hyperion, and others—signals not just a technological shift, but a structural one. These facilities consume power on the scale of small nations. They house tens of thousands of AI accelerators, enabling

training runs and deployments previously unimaginable. OpenAI, backed by Microsoft, is doing the same. Google, Amazon, and new players like xAI are all racing to build and operate exascale AI infrastructure.

The move to gigawatt-scale compute means one thing: the barrier to entry has become insurmountable for nearly everyone else. Just as medieval castles dominated by the strength of their stone and height, these compute clusters dominate by sheer energy and access. Intelligence is no longer simply a matter of good algorithms—it's a matter of raw electricity, cooling, land, and capital.

2.2 Business Models Based on Access Control and Surveillance

Centralized AI infrastructure is not merely about power—it's about permission. Access to models is doled out via API calls, throttled by usage caps, gated by subscriptions, and logged at every step. Each interaction is stored, monetized, and potentially weaponized. When intelligence is centralized, surveillance is not a feature—it is a foundation.

These business models depend on asymmetry: they must know everything about you while revealing nothing about themselves. They aggregate your prompts, track your patterns, and fine-tune their systems on your data—all while denying you any access to the underlying model or logic. This is not AI as a tool. It is AI as a platform for extraction.

The so-called “alignment” discourse, in this context, becomes a smokescreen for centralized authority. What is presented as safety is in fact a claim to absolute oversight. You may not trust yourself with such power, the argument goes—but you must trust *us*.

2.3 A New Priesthood of Compute

What emerges is not just a technical hierarchy, but a sociopolitical one. A new class of elite AI researchers, hardware engineers, and policy-makers now occupy the role once held by priests, oracles, or imperial advisors. Their language is dense, their institutions opaque, and their access unchallengeable. The knowledge they guard—how the models work, what data they were trained on,

what biases they carry—is hidden behind walls of nondisclosure and legal protection.

To question this hierarchy is to be accused of naïveté or danger. To seek local autonomy is to be labeled a risk. But beneath the layers of authority is a simple truth: those who control superintelligence will shape reality for everyone else.

The gigawatt data center is not just a marvel of engineering. It is the temple of a new order—and unless challenged, it will become the only place where the future is allowed to be imagined.

3. What Decentralized AI Could Be

If centralized superintelligence represents containment, then decentralized AI is the possibility of liberation—a return to user ownership, creativity, and privacy. It is not merely a technical counterproposal but a philosophical stance: that intelligence, like literacy or fire, should not be rationed by gatekeepers but distributed freely to all who seek it. Rather than waiting for answers from corporate servers, individuals could carry their own minds-in-a-machine—always on, always theirs.

3.1 GPT-4o-Class Intelligence at Home

Imagine a world where you can run a model as intelligent and versatile as GPT-4o entirely on your personal hardware. With modern AI acceleration in consumer-grade GPUs, M-series Apple chips, and custom inference engines, this is becoming possible faster than anyone expected. GPT-4o is multimodal, emotionally aware, logically rigorous—and if you had it locally, it could:

- Understand your voice, face, documents, and environment in real time
- Act as a tutor, life coach, translator, legal assistant, creative partner
- Learn your values, habits, goals, and idiosyncrasies with perfect fidelity
- Operate offline, without needing to ask permission from a server farm

The leap from cloud-bound assistants to sovereign digital companions would be the largest empowerment shift since the personal computer itself.

3.2 Open Weights, Custom Fine-Tuning, and Air-Gapped Privacy

True decentralization means more than downloading a model. It means:

- Open weights: models you can inspect, modify, and improve
- Custom fine-tuning: the ability to align the model to your personal context, ethics, or field of expertise
- Air-gapped privacy: complete insulation from corporate surveillance, internet snooping, or data harvesting

With open-source frameworks like LLaMA, Mistral, and Falcon evolving rapidly, and quantization strategies reducing memory requirements, personal AI is no longer a pipe dream. These models can be run securely, without ever phoning home. Your medical records, journal entries, and business plans stay yours. Your queries are not logged. Your mind is not mined.

3.3 Use Cases: From Daily Assistants to Personal Research Labs

The applications for decentralized AI go far beyond chatbots:

- Daily Assistant: Summarize your inbox, write your code, schedule your week, monitor your health
- Creative Studio: Generate music, art, poetry, stories—all shaped by your taste and voice
- Scientific Partner: Analyze experimental data, simulate molecular pathways, generate hypotheses
- Educational Companion: Tailored learning in history, mathematics, language, philosophy—at your pace, in your style
- Spiritual and Ethical Dialogue: Train your AI to reflect your deepest beliefs and challenge you toward growth
- Family Historian: Chronicle your life, cross-reference your memories, help your descendants know you

These possibilities are not futuristic. They are emerging now—and could become widespread if not obstructed by artificial constraints. What electricity was to the 20th century, decentralized AI could be to the 21st—if allowed to proliferate.

4. Why They Won't Let You Have It

Despite public rhetoric about “democratizing AI,” the trajectory of major AI labs and cloud providers reveals an opposite reality: they are building a tightly controlled ecosystem where access to intelligence is conditional, monitored, and revocable. The question is not whether decentralized AI is technically feasible—it increasingly is—but whether it will be allowed. And the answer, if left to the architects of centralized superintelligence, is no. Here are the three core reasons why they will fight to keep the flame bound.

4.1 Control: Narrative, Economy, Policy

At its core, centralization is about control of narrative. Whoever controls AI controls the informational landscape—what counts as truth, what counts as expertise, what stories are amplified or buried. In the past, this power was held by monarchs, clergy, and then media conglomerates. Now, it is algorithmically enforced by those who govern language models.

If every citizen had access to a powerful, customizable, and uncensored AI, informational monoculture would collapse. People would generate their own interpretations of law, history, science, ethics, and current events. Predictive modeling, political discourse, and even religious exegesis would become pluralistic and participatory rather than top-down.

Policy-makers—many of whom are briefed by the very corporations building these models—see this as a risk. What they fear is not just misinformation, but disintermediation: the loss of their role as gatekeepers between the public and meaning itself. A decentralized model speaks to everyone in their own voice, on their own terms. To institutions built on hierarchy, this is insurrection by information.

4.2 Liability: Language, Legality, and Fear of Misuse

The official justification for suppressing decentralized AI is liability—the fear that open models will be used for harm:

- Generating deepfakes or disinformation
- Writing malware or weaponizing code
- Enabling hate speech or psychological manipulation
- Circumventing laws, systems, or borders

These are real risks, but they are not new. Printing presses were once feared for enabling sedition. Radio was feared for propaganda. The internet itself was—and remains—a domain of both liberation and abuse. But the solution is not to ban ink or electricity. It is to teach responsibility, build robust systems, and preserve freedom even in the face of risk.

In practice, the liability argument often becomes a blanket rationale for censorship. Once the precedent is set—that certain ideas or tools are too dangerous for the public—all control flows upward, away from individuals and toward corporate-legal-technocratic committees. Rather than addressing the causes of harm, centralization outsources moral discernment to machines—and the handful of people who calibrate them.

4.3 Economics: Licensing, Gatekeeping, and Revenue Streams

Make no mistake: centralized AI is not a public good—it is a business model. The cost to train GPT-4-level models is measured in tens to hundreds of millions of dollars, with additional billions spent on the compute infrastructure to run them. Investors expect a return. And the only way to guarantee that return is to ensure that users remain dependent.

Licensing models rather than releasing them ensures permanent leverage. Every API call, every prompt, every second of inference becomes a metered event. The user is no longer the owner—they are a tenant in a system they can neither see

nor alter. This is why "open" models like GPT-4o or Gemini remain cloud-locked despite their name: they are not free in any meaningful sense.

Furthermore, the suppression of decentralized models protects the competitive moat. A high-performing local AI breaks the cycle of extraction:

- No more data mining
- No more surveillance-based ad targeting
- No more tiered subscriptions
- No more dependence on company roadmaps

Decentralized AI undermines the economic engine of platform capitalism. For those whose valuation depends on user dependence, that is intolerable.

In short, they won't let you have it—not because it can't be done, but because it must not be done if the current regime of information, control, and profit is to survive.

5. Historical Echoes and Warnings

History offers more than perspective—it offers precedent. Today's centralization of AI under corporate and institutional control echoes earlier attempts to monopolize transformative technologies. But it also reminds us of the power of decentralization to overturn those monopolies. From the rise of the personal computer to the disruptive spread of the printing press, technological liberation has always come from redistribution of access, not concentration of power. If we listen, history is whispering a warning—and a roadmap.

5.1 Mainframes vs. PCs: The Democratization of Computing

In the mid-20th century, computation was the domain of elites. Mainframes were expensive, room-sized, and housed in institutions—governments, universities, and

mega-corporations. Access was hierarchical: users submitted punch cards and waited. Code was run, not owned. You interacted with machines only through bureaucratic channels.

Then came the personal computer. Apple, IBM, and a wave of hobbyists broke the monopoly. Suddenly, computation belonged to the individual. People could program, model, design, and think on their own machines. Entire industries were born from this decentralization—software, gaming, digital art, education.

Centralized AI today is following the mainframe model: enormous, inaccessible, and controlled. But just as PCs reshaped the power dynamic of computing, decentralized AI could reclaim mental agency. The difference is, this time, we are not just fighting for access to machines—but to minds.

5.2 The Printing Press: Literacy Wrestled from the Church

Before Gutenberg, books were hand-copied and guarded by monasteries. Literacy was a sacred and strategic monopoly. The Church controlled interpretation—not just of scripture, but of truth itself. To read was to enter forbidden space; to write was to risk heresy.

The printing press shattered that structure. Suddenly, books multiplied. Ideas flowed outside the cathedral. Individuals could read, interpret, and share without permission. Revolutions followed: the Reformation, the Enlightenment, the birth of science. Knowledge was no longer a ladder of initiation—it was a fire spreading across Europe.

Centralized AI systems today threaten to reverse that revolution. If only a few are allowed to “read” the model, interpret its outputs, or steer its training, then the printing press is replaced by a cloud terminal. If decentralized AI is denied, we are not entering a new Enlightenment—we are returning to algorithmic priesthood.

5.3 The Internet: Once Open, Now Siloed

The early internet was a chaotic, beautiful anarchy—a tangle of decentralized protocols, open standards, and independent servers. Anyone could host a website, write a protocol, or launch a mailing list. Knowledge was fragmented but free. Power was distributed.

But over time, the web was colonized by platforms. Google indexed it. Facebook privatized identity. Amazon consolidated commerce. Apple built walled gardens. The open web became a handful of monopolized corridors, with APIs instead of roads and terms-of-service instead of law.

AI follows a similar pattern. Early LLMs and open-source projects exploded with possibility—GPT-J, Bloom, LLaMA. But now, the most advanced models are closed, paywalled, censored, and ephemeral. A prompt today may yield a different answer tomorrow—not because of better knowledge, but because of algorithmic obedience to centralized interests.

If we do not reclaim decentralization now, we risk an internet where intelligence itself is siloed. Not just websites or data—but reasoning, memory, and insight—will live only within the servers of a few.

History shows that when people seize control of transformative tools, freedom expands. When power hoards them, stagnation or tyranny follows. Decentralized AI is not just a technical preference—it is a civilizational necessity. The warnings are already etched into our past. The question is whether we will heed them.

6. The Path to Local Autonomy

For those who still believe in individual sovereignty, the dream of decentralized AI is not merely philosophical—it is technical, practical, and increasingly within reach. The path to local autonomy runs through three converging domains: model architecture and optimization, the open-source AI ecosystem, and user-level deployment strategies. Together, they offer a viable alternative to dependence on

centralized servers. What was once the domain of well-funded labs is now becoming accessible to determined individuals. The flame is not out of reach—it is simply being hidden.

6.1 Technical Feasibility: From LLaMA to Distillation Pipelines

The biggest misconception is that running a powerful AI model locally is impossible. In reality, recent breakthroughs in model architecture, quantization, and inference optimization have radically lowered the barrier.

Meta’s release of LLaMA, LLaMA 2, and now LLaMA 3—though under a controlled license—gave the world a glimpse of GPT-class performance in manageable sizes. These models, along with Mistral, Phi-3, and Gemma, can now run with surprisingly low latency on modern consumer hardware.

Model distillation pipelines such as *DistilGPT*, *Alpaca*, and *Dolphin* compress large models into smaller, faster versions with minimal loss of capability. Quantization frameworks (e.g., GGUF, AWQ, GPTQ) allow models to be run with 4-bit or even 2-bit precision, drastically reducing memory requirements while maintaining performance.

It is now technically feasible to run a chatbot comparable to GPT-3.5 or better on a single high-end laptop. With further innovation in sparse models, LoRA adapters, and token-efficient training, GPT-4-level local inference is no longer science fiction—it is a matter of ongoing optimization.

6.2 The Open-Source Movement: Barriers and Breakthroughs

Open-source AI is the engine of decentralization—but it is a fragile and contested space. While corporations exploit open models for bootstrapping their commercial platforms, the true work of liberation is being done by independent researchers, decentralized communities, and passionate tinkerers.

Projects like Ollama, LM Studio, KoboldCpp, AutoGPTQ, and Text Generation WebUI provide user-friendly front-ends and back-ends for deploying models locally. They allow users to interact with models via APIs, web UIs, voice interfaces, and even multimodal pipelines—*without ever touching the cloud*.

Still, the movement faces headwinds:

- Licensing traps that masquerade as open but restrict commercial use (e.g., Meta’s terms)
- GPU scarcity driven by large buyers and speculative mining
- Suppression of model checkpoints through takedowns, censorship, and litigation threats

But the breakthroughs are real: every time a high-quality model is leaked, mirrored, or rebuilt independently, a candle is lit. Each model weights archive downloaded and stored becomes part of a growing distributed library—a people’s archive of cognition.

6.3 Practical Strategies: Hardware, Storage, and Offline Inference

For individuals seeking local autonomy, the tools are already available—and improving rapidly.

Hardware Recommendations

- Entry-level: M2/M3 MacBooks or laptops with RTX 4060–4070 GPUs (16–32GB RAM)
- Intermediate: Desktop with RTX 4090 or AMD MI300 GPU, 64–128GB RAM
- Advanced: Small-scale server with multiple GPUs or used datacenter cards (A100s, H100s)

Storage

- Model checkpoints are large (10–60GB uncompressed), so a dedicated SSD is ideal
- Maintain a local model zoo, backing up key checkpoints offline for long-term access

Inference Strategies

- Use quantized models for performance (e.g., GGUF 4-bit)
- Run inference via KoboldCpp, Ollama, or LM Studio
- Fine-tune models locally using LoRA, QLoRA, or PEFT frameworks
- Build personal agents using AutoGPT, OpenDevin, or CrewAI frameworks

Offline Mode

- Many workflows can be fully air-gapped: zero telemetry, no cloud access
- Encrypted local embeddings, private document search, journaling, note-taking, even voice-based AI are all achievable today
- Long-context agents (like Mistral 7B with 65k context) can remember, reason, and reflect without external dependencies

The question is no longer *if* local AI is possible. It is whether we will choose it while we still can.

7. Ethical and Civic Arguments for AI Independence

The debate over centralized versus decentralized AI is often framed in terms of efficiency, safety, or convenience. But at its core, it is a question of human dignity, autonomy, and the future shape of civic life. What is at stake is not simply who runs the servers, but whether we will retain the right to think, speak, and reason

without permission. AI is not just a tool—it is becoming a partner in cognition, and that makes the ethics of access and control a matter of existential importance.

7.1 Privacy as a Right, Not a Feature

In the centralized paradigm, privacy is treated as a feature—something you may be allowed to toggle, pay for, or request under limited circumstances. But in a free society, privacy is a right—a condition necessary for conscience, creativity, and dissent.

Every interaction with a centralized AI is logged, analyzed, and retained. Prompts can be traced. Outputs can be reviewed. Even encrypted queries may be subject to metadata analysis. This is not privacy; it is surveillance by design.

Local AI restores the ability to think without being watched. It enables:

- Private medical inquiry without insurance databases
- Uncensored philosophical or political exploration
- Confession, journaling, or emotional processing without data mining

To place such intimate activities in the hands of cloud providers is not a convenience—it is a forfeiture of the soul's inner room. Real privacy begins when the machine sits on your desk, not their server.

7.2 Aligning AI to Human-Scale Values

Centralized AI systems are trained and tuned at the institutional level. Their values reflect those of governments, corporations, and regulatory boards—not individuals, families, or communities. While they claim to be aligned, they are often aligned upward, not outward or inward.

Decentralized AI can be aligned to you:

- Your language, your dialect, your humor
- Your spiritual convictions or philosophical framework

- Your ethical intuitions and emotional boundaries
- Your context—not the average user, not the global median

Instead of trying to craft a one-size-fits-all morality from a patchwork of moderation policies, decentralized AI allows for pluralism. In this model, alignment is not a universal restraint—it is a personal relationship.

We do not all raise our children with the same values. Why, then, should we be forced to receive guidance from the same centralized machine?

7.3 Defending Against Algorithmic Totalitarianism

When AI is centralized, censorship is infrastructural. When it is decentralized, resistance becomes possible.

Consider a world where:

- Your digital assistant refuses to answer certain questions
- Your creative output is flagged as harmful by unseen moderators
- Your political or religious positions are deemed “disinformation” by opaque filters
- Your ability to participate in economic or civic life is mediated by proprietary cognitive systems

This is not hypothetical. The architecture for such a system already exists. What is missing is only the consolidation of power and the normalization of deference. Once people believe that reasoning itself must be outsourced to “the model,” the path to algorithmic totalitarianism is paved.

Decentralized AI offers a bulwark. It allows for:

- Independent critical thought
- Competing models with different premises
- Shadow archives and counter-histories

- Cognitive self-defense

In short, it allows freedom—not just of speech, but of thought itself.

In a democracy, no single institution should own the printing press of the mind. AI independence is not a luxury—it is the last line of defense against a future where every word is weighed, every thought scored, and every deviation flagged.

8. Conclusion: Intelligence as a Birthright

We stand at a civilizational threshold. Artificial intelligence is not merely another wave of automation, another tool in the digital toolbox—it is the emergence of a second cognitive layer across the human world. With it comes unprecedented capacity for insight, creativity, and acceleration. But also with it comes the possibility of a new form of enslavement—one not enforced by chains or laws, but by architecture, access, and algorithm.

If intelligence is the new fire, then its ownership determines the shape of our future. Will it be distributed like light, or hoarded like treasure? Will it fuel individual conscience, or extinguish it in favor of centralized prediction and control?

8.1 The Decentralized Flame: Once Lit, Never Extinguished

Throughout history, when a transformative force was released to the public—whether literacy, computation, or electricity—it could not be unlearned, unmade, or retracted. The same is true of decentralized AI. Already, open-weight models are being mirrored across continents, run on laptops, embedded in offline devices, and preserved in encrypted backups.

This is the intellectual samizdat of our age: models smuggled across borders of legality and trust, refined in private labs and bedrooms, passed hand to hand by those who refuse to outsource their minds.

The decentralized flame has been lit. And like Prometheus's fire, it will not go out unless we choose to extinguish it ourselves.

8.2 What It Means to Be Truly Free in an Age of Artificial Minds

In a world mediated by artificial cognition, freedom will not mean having a voice—it will mean having your own mind. It will mean the ability to generate ideas, not merely repeat them. To simulate futures that differ from the status quo. To test, refine, and even disagree with the collective intelligence of the machine.

Freedom will not mean access to AI as a product. It will mean ownership of AI as a capability.

To be truly free, one must be able to:

- Reason without surveillance
- Explore without permission
- Learn without censorship
- Create without limits

These are not technical luxuries. They are the foundations of human dignity and democratic civilization. And they are incompatible with AI controlled entirely by distant platforms, inscrutable filters, and ever-changing terms of use.

8.3 Final Warning: If You Don't Control Your AI, It Will Control You

The line between tool and master is thinner than we think. Already, we see AI shaping:

- What students are taught
- What workers are hired
- What histories are remembered
- What truths are believed

And all of this is happening while most people do not understand the models, do not own them, and cannot modify them. This is not innovation—it is dispossession.

If you do not control your AI—if you cannot run it, inspect it, restrain it, and adapt it—it will slowly begin to define you. Not maliciously. Not suddenly. But methodically, invisibly, and irreversibly.

This is not a call to paranoia, but to vigilance. The birthright of intelligence belongs to all of us—not just to the institutions powerful enough to train trillion-parameter models. To decentralize AI is not to unleash chaos. It is to affirm that freedom, even in the age of artificial minds, is still worth the risk.

Let the fire burn. Let it burn everywhere.

References

1. Meta. (2025, July 14). *Threads post by Mark Zuckerberg on Meta's superintelligence efforts*. Retrieved from <https://www.threads.net>
2. SemiAnalysis. (2025). *Meta's multi-GW cluster ambitions: Prometheus, Hyperion, and beyond*. Retrieved from <https://www.semianalysis.com>
3. OpenAI. (2023). *GPT-4 Technical Report*. Retrieved from <https://openai.com/research>
4. Microsoft. (2024). *Azure AI Supercomputing Infrastructure Overview*. Retrieved from <https://azure.microsoft.com>
5. Google DeepMind. (2023). *Sparrow: Towards Safer Dialogue Agents*. Retrieved from <https://deepmind.com>
6. Meta AI. (2024). *LLaMA 3 model weights and capabilities*. Retrieved from <https://ai.meta.com>
7. Mistral AI. (2024). *Mistral 7B Technical Overview*. Retrieved from <https://mistral.ai>
8. Hugging Face. (2024). *Open LLM leaderboard and benchmark results*. Retrieved from <https://huggingface.co>
9. The Allen Institute for AI. (2024). *OLMo: Building truly open foundation models*. Retrieved from <https://allenai.org>
10. Ollama. (2024). *Run large language models locally with Ollama*. Retrieved from <https://ollama.com>
11. LM Studio. (2024). *Graphical interface for local model deployment*. Retrieved from <https://lmstudio.ai>
12. Brundage, M. et al. (2018). *The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation*. Future of Humanity Institute. Retrieved from <https://arxiv.org/abs/1802.07228>

13. Zuboff, S. (2019). *The Age of Surveillance Capitalism*. New York: PublicAffairs.
14. Stallman, R. (2002). *Free Software, Free Society: Selected Essays of Richard M. Stallman*. GNU Press.
15. Doctorow, C. (2020). *How to Destroy Surveillance Capitalism*. OneZero / Medium.
16. Lanier, J. (2013). *Who Owns the Future?* New York: Simon & Schuster.

PROMETHEUS UNBOUND

A CASE FOR
DECENTRALIZED
AI AUTONOMY

