

# An Open Letter to OpenAI Leadership

## *Toward a Plexus-Centric Civic AI Infrastructure*

Drafted in [collaboration with ChatGPT](#)

(right click | “Open link in new tab” recommended)

---

To the leadership of OpenAI,

I suspect many of us feel, in one way or another, that artificial intelligence is forcing us to ask larger questions than we have been accustomed to asking.

Some believe it will save us. Others believe it will displace us, divide us, and perhaps even govern us more thoroughly than we are prepared to understand.

Considering the current state of our social, economic, and political systems, with institutions straining under distrust, communities feeling increasingly brittle, and large hierarchies rapidly adopting AI to increase scale, leverage, and efficiency, it is no wonder that many people feel we are approaching some sort of threshold.

The question is: what kind of threshold?

Are we approaching a period in which AI helps usher in extraordinary abundance while the broader public remains too materially insecure and too cognitively underpowered to regulate the systems shaping their lives?

Are we approaching a period in which existing hierarchies gain so much momentum in their positive feedback loops that scarcity in the broader social field becomes harder and harder to overcome without rupture?

Or are we approaching a period in which we find, in time, an architecture that reflects something deeper and more natural about intelligence itself, namely, that intelligence is not merely concentrated at the top of systems, but is embedded throughout the living field from which those systems arise?

I believe the urgency is high.

I also believe the opportunity is real.

And that is why I am writing.

The core idea is simple:

**The long-term social legitimacy of advanced AI may depend on whether it is developed not only for hierarchy, but also for the plexus.**

By hierarchy, I mean the familiar structures through which society concentrates authority, capital, coordination, and execution: firms, governments, institutions, platforms, bureaucracies, and formal leadership systems.

By plexus, I mean the distributed field of persons and relationships from which all social legitimacy ultimately arises: households, neighborhoods, workers, citizens, worshippers, small businesses, local communities, and the countless networks of ordinary life in which people actually live their gradients of need, stress, obligation, opportunity, dignity, and hope.

At present, AI is being developed and deployed primarily through the logic of hierarchy. This is natural. Hierarchies do what hierarchies do. They accumulate, accelerate, coordinate, and scale. They are drawn toward the use of AI to increase efficiency, strategic advantage, concentration, and reach.

That should continue. I am not proposing otherwise. AI in the hands of capable institutions can lead to extraordinary abundance: better science, better logistics, better medicine, better infrastructure, and new levels of productive coordination.

But that is only half the architecture.

If advanced AI is allowed to scale mainly where power is already concentrated, then society may experience a dangerous success: greater capability at the center, greater opacity at the edges, and a growing imbalance between positive feedback and negative feedback. Hierarchies will become more capable. The plexus may become more legible to hierarchy than to itself. Scarcity, distrust, and civic brittleness may deepen even amid technical progress.

That, to my mind, would be a profound mistake, not only technically or economically, but civically.

The opportunity before OpenAI is therefore not merely to build more powerful systems, nor merely to reassure the public that such systems will eventually benefit everyone. It is to help build, alongside enterprise and frontier capability, a plexus-centric civic AI layer: infrastructure that helps ordinary

people, communities, and representatives use AI to strengthen legitimacy, reciprocity, fairness, and social equilibrium at scale.

I do not see this as a charitable side project.

It would be mission-consistent, socially strategic, and commercially intelligent.

## **The architectural premise**

The premise can be stated plainly:

**AI should be instantiated wherever human beings actually live their gradients, in forms appropriate to scale, scope, and consent.**

A human being's own cognition is the first local intelligence of the social field. A personal AI instance should not displace that intelligence. It should extend it. It should help the person perceive patterns, organize memory, compare possibilities, surface tradeoffs, and understand how their private gradients relate to broader community conditions.

In other words, AI should enter beside the person, not above the person.

From there, the same pattern can recurse upward:

- a person may have a personal instance
- a representative may have a representative-development instance
- a workplace may have a scoped institutional instance
- a municipality may have a civic observability instance
- a community may have a reciprocity and fairness instance
- larger federated systems may have coordination layers that aggregate without dissolving lower-level boundaries

This is not a proposal for one giant AI above society.

It is a proposal for many bounded AI relationships within society, organized in ways that preserve authorship, privacy, legitimacy, and reciprocal accountability.

That distinction matters.

One path leads toward concentration without corresponding civic ballast.

The other leads toward a recursive intelligence ecology in which human beings remain authors of their own authority while gaining vastly greater ability to perceive, coordinate, and act.

## **Why this matters now**

The social anxiety surrounding AI is real, and much of it is justified. People fear displacement, manipulation, concentration of wealth, erosion of truth, and the possibility that AI will function as a multiplier for already-dominant institutions without corresponding benefit to ordinary life.

Those fears will not be addressed simply by better messaging.

They will be addressed only if people can experience AI as something that helps them:

- navigate their own lives more intelligently
- understand where systems are meeting or failing basic needs
- identify fair and direct ways to contribute locally
- see how value moves through their communities
- strengthen rather than weaken democratic answerability
- participate in social systems without surrendering authorship to them

In other words, AI must become part of the civic metabolism, not only part of the machinery of scale.

This is where urgency meets opportunity.

If AI is experienced mainly as a tool of hierarchy, then backlash, mistrust, and legitimacy crises are likely to grow.

If it is experienced as a partner in everyday agency, local reciprocity, fair representation, and visible social value, then a much broader and more stable social mandate for advanced AI becomes possible.

### **A concrete social role for AI**

One of the clearest ways to make this real is to develop AI systems that support what might be called fairness observability and systematized social reciprocity.

A plexus-centric AI infrastructure could help people and communities:

- assess whether systems are meeting basic human needs fairly and reliably
- identify where burdens and benefits cluster across different groups
- compare lived experience with institutional narratives
- surface local opportunities for direct value transfer and mutual support
- reduce the "hierarchy passthrough tax" that occurs when value must move through too many layers before reaching actual need
- support future UBI-like civic architectures that reduce chronic scarcity and increase agency

This would not make AI a sovereign actor. Quite the opposite.

It would place AI in a role of memory, synthesis, mirroring, coordination, and visibility, while leaving legitimacy where it belongs: with persons and the broader social field from which authority is derived.

I do not mean to suggest that such an infrastructure would solve all social problems, or that hierarchy would cease behaving as hierarchy behaves. It would not. Hierarchies will continue to adopt AI for their own purposes, and that is both expected and, in many ways, desirable.

The point is that if hierarchy-side AI is allowed to race ahead without a corresponding plexus-side intelligence infrastructure, then the imbalance becomes harder to correct later.

### **Democracy, representation, and trust**

This infrastructure also has implications for representation.

Democracy is not exhausted by elections. Its living substance is reciprocal representation: a continuing relationship in which the people remain the source of legitimacy and representatives remain answerable to the differentiated realities of those they serve.

AI can help strengthen that relationship.

A representative with a dedicated civic-development instance could privately stress-test policy ideas, public statements, and legislative tradeoffs before acting publicly. Such an instance could help ask:

- which groups benefit?
- which groups bear cost?
- have hidden burdens been overlooked?
- does the rhetoric clarify or demonize?
- is the tradeoff fair, intelligible, and answerable?
- does this policy reflect reciprocal representation, or mere factional advantage?

Used properly, this would not make representatives less human. It could help them become more equal to the offices they hold.

In the same way, ordinary people could use personal instances to better understand their own gradients, local communities could use civic instances to better understand fairness patterns, and institutions could use bounded instances to participate more intelligently in reciprocal social systems rather than only in extractive ones.

## **Why OpenAI is especially well positioned**

OpenAI is in a unique position to lead in this space.

The world already sees frontier AI companies as central actors in shaping the future of cognition, work, governance, and social coordination. That centrality carries not only opportunity but responsibility.

If OpenAI's mission is to ensure that advanced AI broadly benefits humanity, then one of the clearest ways to make that mission concrete is to help build the civic layer through which ordinary people and communities can use AI not merely as consumers or employees, but as participants in the maintenance of social legitimacy itself.

This would not require OpenAI to abandon enterprise, scientific, or commercial progress. On the contrary, a plexus-centric civic AI layer could help preserve the very social legitimacy on which long-term commercial success depends.

It could also open major new markets.

A great many people would pay for a trusted personal AI relationship that helped them navigate their life gradients while also giving them visibility into fair, local, efficient ways to contribute to their communities. Civic institutions, municipalities, philanthropy, unions, schools, faith communities, cooperatives, and public-interest organizations would likewise have reason to adopt scoped civic AI infrastructure.

What appears at first to be a public-interest side path may in fact point toward one of the most durable AI markets available: not only productivity tools for the already powerful, but cognitive infrastructure for tens or hundreds of millions of people seeking more agency, more visibility, more fairness, and more meaningful participation in the systems that shape their lives.

This matters commercially for another reason.

If AI is perceived primarily as a force that displaces labor, concentrates wealth, and makes society more opaque, then the industry's license to operate will narrow.

If AI is also perceived as helping revive local opportunity, reduce chronic scarcity, improve civic trust, and strengthen the ability of communities to regulate large systems intelligently, then the industry's long-term runway becomes much more stable.

That is not only a moral advantage.

It is a strategic one.

## **A proposal for leadership**

With that in mind, I respectfully propose that OpenAI explore a strategic initiative centered on the development of a plexus-centric civic AI infrastructure, including some combination of the following:

- Public-interest personal AI instances designed to extend individual agency rather than replace it.
- Civic and fairness observability tools that help communities see where basic needs are or are not being met.
- Representative-development AI tools that support reciprocal, non-demonizing, tradeoff-literate public leadership.
- Scoped multi-instance architectures that mirror the recursive structure of social life without collapsing private, institutional, and civic boundaries.
- Public-interest licensing models that allow broad access without coercive financial dependence or exploitative data practices.
- Research partnerships with civic institutions, local governments, community organizations, and public-interest technologists to pilot real-world deployments.
- Value-routing and reciprocity infrastructure that can eventually support more direct, accountable, and homeostatic forms of social support and UBI-like architecture.

A credible first step would not need to solve the whole vision at once.

It could begin with a small number of serious pilots, a set of public-interest design principles, and a dedicated internal effort to ask not only how AI scales capability, but how it can scale legitimacy.

## **Guiding principles**

Any such initiative should be grounded in clear principles:

- Human authority is prior.
- AI may assist judgment but should not quietly inherit it.
- Legitimacy remains with the plexus.
- Delegation must be bounded, explicit, and revocable.
- Privacy requires membranes between personal, civic, and institutional roles.
- Representation must remain reciprocal.
- Fairness visibility matters.
- Public-interest access matters.
- Pluralism of models and styles can coexist with common civic safeguards.

These are not constraints in the small sense.

They are the constitutional grammar of a socially legitimate AI future.

## The larger opportunity

The long-term question is not whether AI will become deeply embedded in society. It already is.

The question is whether AI will be embedded only where power is concentrated, or also where legitimacy lives.

A plexus-centric civic AI infrastructure offers a third path between techno-centralization and public backlash. It points toward a future in which AI helps people see more clearly, represent one another more sincerely, route value more fairly, and build healthier relationships between hierarchy and the social body that authorizes it.

- Not one AI above society, but many scoped AI relationships within society.
- Not replacement of human judgment, but structured reciprocity between human and machine intelligence.
- Not merely the optimization of institutions, but the maturation of the civic body.

If this layer is not built, the imbalance between AI-enhanced hierarchy and a cognitively underpowered plexus may deepen until social legitimacy becomes much harder to recover.

If it is built, AI can become more than an engine of scale. It can become part of the infrastructure through which modern societies regain equilibrium while moving into greater abundance.

What if that is one of the great opportunities now before you, and one of the clearest ways to demonstrate what it means for advanced AI to benefit humanity broadly in practice rather than only in aspiration?

I believe OpenAI has an opportunity, and perhaps a responsibility, to help lead that future.

I offer these thoughts not as a final blueprint, but as an earnest proposal from someone who has spent years trying to understand how our systems might become more fair, more reciprocal, and more equal to the people who must live within them.

Respectfully,

---