

# The Age of Sovereign AI

Process, Stance, and the Politics of Building

Johan Michalove\*

February 26, 2026 — *Working draft*

*For Liam Conejo Ramos*

*Why do new technologies so often serve the powerful first, and everyone else as an  
afterthought?*

— Chris Csíkszentmihályi

## Abstract

This paper is a critical reflection on the experience of building artificial intelligence for community care—specifically, a voice and text concierge for social service navigation in New York City—during a period of political crisis, technological acceleration, and state violence. It is not a technical paper; it discloses no architecture, no methods, no implementation details. It is a paper about stance: the ethical, political, and intellectual commitments that precede and condition the work of building, and that no technical paper can articulate. Drawing on the tradition of critical computing inaugurated by Chris Csíkszentmihályi, the psychology of creative absorption developed by his father Mihaly Csikszentmihalyi, and frameworks from postcolonial theory, care ethics, and science and technology studies, I argue for what I call *sovereign AI*—artificial intelligence that belongs to communities rather than corporations, that operates through federation rather than extraction, and that measures its worth not by revenue or benchmark but by whether someone who needed help received it. The paper names the complicity of major technology companies in state violence—particularly the ongoing catastrophe in Palestine—and distinguishes between corporate AI that serves power and community AI that serves care. It draws

---

\*Department of Information Science, Cornell University. [jam844@cornell.edu](mailto:jam844@cornell.edu). In honor of Mihaly Csikszentmihalyi (1934–2021) and Chris Csíkszentmihályi.

on the author's published essays to trace the intellectual trajectory from associative cognition through AI critique to the practice of building in broken worlds. It is dedicated to a child.

**Keywords:** sovereign AI, critical computing, care, Palestine, federation, flow, postcolonial computing

# 1 Introduction: The Age

What age is this?

Not the age of artificial general intelligence—a corporate fantasy dressed in the language of inevitability, funded by venture capital, and marketed as salvation. Not the age of alignment—a technical problem masquerading as a political one, in which the question of whose values an AI should encode is treated as an engineering challenge rather than a struggle over power. Not the age of disruption, that threadbare euphemism for the destruction of livelihoods, communities, and public goods in the name of shareholder returns.

This is the age of sovereign AI. The age when communities claim artificial intelligence as their own infrastructure of care—or lose the capacity to do so forever.

I write this from a particular position: a doctoral student in information science at Cornell University, a dual citizen of Denmark and the United States, a builder of AI systems for community social services in New York City. I am white, male, and possessed of institutional resources that most of the people my system serves are not. I write in English—the imperial language, the language of the technology industry, the language that large language models speak best. I am, in Shotwell (2016)’s vocabulary, compromised. There is no pure position from which to build.

I write this paper because the two companion papers I have written about this work—one analyzing voice AI as care infrastructure through the lenses of science and technology studies (Michalove, 2026c), the other specifying twenty-six evidence-based design principles for conversational architecture (Michalove, 2026a)—cannot say what this paper says. They are analytical and normative, respectively. This paper is reflexive. It asks not what voice AI *means* or how it should be *designed* but who I am when I build it, what commitments precede the first line of code, and what it means to build AI for care in a world where AI is also a weapon.

The paper is structured as an essay in nine movements. It begins with the psychology of creative absorption and the politics of who gets to experience it. It traces the tradition of critical computing through the work of Chris Csíkszentmihályi. It examines AI as a pharmakon—simultaneously poison and cure. It names the complicity of specific technology companies in specific acts of state violence. It articulates the concept of sovereign AI. It weaves through the intellectual trajectory that led me to this work. And it arrives, finally, at the question of the child—the question that makes every other question urgent.

No technical or methodological disclosure appears in this paper. The techniques I have developed could be used for immense harm. This is not false modesty; it is a design

decision. The IP firewall between what I build and what I publish is absolute. What I offer here is not a blueprint but a stance.

## 2 On Flow

Mihaly Csikszentmihalyi spent a lifetime studying what he called *flow*: the state of complete absorption in an activity, the merging of action and awareness, the loss of self-consciousness that accompanies deep creative engagement (Csikszentmihalyi, 1990). He described it as optimal experience—the condition in which people report being most alive, most themselves, most fully engaged with the world. Flow occurs when the challenge of the task matches the skill of the person, when the goals are clear, when feedback is immediate, and when the activity is intrinsically rewarding. Surgeons report it. Musicians report it. Athletes, chess players, rock climbers, mathematicians report it. It is, Csikszentmihalyi argued, the secret structure of happiness.

I report it when I build.

The nights disappear. The code compiles, fails, compiles again. The architecture reveals itself not as a plan executed but as a pattern discovered—each component fitting into a whole that I did not design so much as find. The database takes shape. The voice synthesizer speaks. The telephone rings, and something answers. I am, in Csikszentmihalyi's terminology, in flow: the challenge matches the skill, the goals are clear (someone needs help finding shelter, and the system must find it), the feedback is immediate (the test call works or it doesn't), and the activity is intrinsically rewarding in a way that has nothing to do with money or recognition and everything to do with the feeling that this particular arrangement of code and data and voice might help someone survive a cold night in New York.

But flow is not innocent.

Csikszentmihalyi (1996) was aware that flow has no inherent moral valence. A sniper in flow is still a sniper. A derivatives trader in flow is still destroying pension funds. The state of absorption is indifferent to the content of the activity. This is the dark underside of positive psychology: the framework describes the phenomenology of engagement without interrogating the politics of what one is engaged in. Arendt (1958) made a version of this argument about labor itself—that the human capacity for work, for *homo faber's* world-building, is necessary but not sufficient for political life. Work without political judgment is craftsmanship in the service of tyranny.

Who gets to experience flow? The question is not incidental. Flow requires time, resources, skills, and the freedom to pursue challenging activities without interruption. It

requires, in other words, precisely the conditions that poverty, precarity, and displacement destroy. The people who call the system I built—elderly New Yorkers whose heat has been shut off, recently arrived immigrants who do not know where to find food, people leaving violent situations who need somewhere safe tonight—are not experiencing flow. They are experiencing crisis. The system I build in a state of creative absorption serves people in a state of desperate need. The asymmetry is constitutive.

I do not resolve this asymmetry. I hold it. The flow is real. The crisis is real. The question is what I do with the flow—whether the absorption serves the craft alone or whether it is directed, by deliberate ethical commitment, toward the care of others. Csikszentmihalyi’s psychology provides the phenomenology of building. It does not provide the politics. For that, I turn to his son.

### 3 Computing as Resistance

Chris Csikszentmihályi dropped out of Reed College in 1988 with a question that would organize his life’s work: why do new technologies so often serve the powerful first, and everyone else as an afterthought? He went on to found the Computing Culture group at the MIT Media Lab, where he built technologies that inverted the standard distribution of technical power. His first major work, *Hunter Hunter*, was an autonomous robot designed during Slovenia’s secession from Yugoslavia—a quadripod that triangulates the sound of a gunshot and responds by firing toward the source. It was a provocation, not a product: a thought experiment about what autonomous weapons look like when they are pointed at soldiers rather than civilians, at power rather than at the powerless.

He built *Freedom Flies*—fabric-wing UAVs designed to monitor encounters between migrants and militias at the U.S.-Mexico border. Where state and paramilitary drones surveilled migrants to intercept and detain them, *Freedom Flies* surveilled the surveillers. He made the plans and code available for free. He built a robotic kayak to protest at Guantánamo Bay. He co-founded the MIT Center for Future Civic Media. He moved to Cornell’s Department of Information Science—my department—where he directs the Redistributive Computing Systems Group, whose name says exactly what it does: compute systems that redistribute power toward the less served (Abdelnour-Nocera et al., 2015).

The tradition Chris Csikszentmihályi represents—which I will call *critical computing*—has no canonical text, no single manifesto. It draws on Haraway (1991)’s cyborg feminism, Winner (1986)’s insistence that artifacts have politics, Weizenbaum (1976)’s early warning about the moral implications of computation, and a long lineage of artist-engineers who use technology to ask questions rather than to ship products. It is distinct from both

mainstream HCI (which optimizes for usability within existing power structures) and from critical theory applied to technology (which analyzes power but does not build). Critical computing builds. It builds in order to ask: what would this technology look like if it were made for the dispossessed?

This is the tradition I work in. Not because I chose it from a menu of academic specializations, but because the question Chris asked at nineteen is the question that organizes everything I have built: who does this serve? The system I have deployed in New York City—a voice and text concierge for social service navigation—is not a product. It has no revenue model, no venture capital, no growth metrics. It is an attempt to answer a phone and help. The question of whether it succeeds is not a technical question about system performance. It is a political question about whether the infrastructure of care can be reclaimed from the infrastructure of extraction.

I honor Mihaly for the phenomenology of building. I honor Chris for the politics of what to build.

## 4 The Pharmakon

In May 2025, I published an essay called “LLM Exposure” in which I introduced the concept of *LLMx*—LLM Exposure as a category of cognitive effect (Michalove, 2025a). The argument was that large language models function as a pharmakon in the classical sense: simultaneously poison and cure, medicine and toxin, depending on dose, context, and the vulnerability of the person exposed. I identified four orders of effects: induced emotional states (first-order), behavioral changes (second-order), long-term cognitive restructuring (third-order), and societal aggregation (fourth-order). The framework drew on Norbert Wiener’s early warnings about cybernetic control (Wiener, 1950), on the emerging clinical literature around AI-induced psychological dependence, and on my own observations of what happens when people offload cognition to machines that are designed to be agreeable.

The pharmakon problem is not abstract for me. I build with the same technology I critique.

The system I deployed uses large language models to understand what callers need and to formulate responses. It uses speech recognition that Koenecke et al. (2020) have shown produces twice the error rate for Black speakers as for white speakers. It uses text-to-speech that reproduces the gendered service voice that Hochschild (1983) identified as emotional labor. Every component of the system is compromised. Every component is also, in the specific context of its deployment, an attempt to help someone who called

because they needed help.

I had written about this tension before the system existed. In “Spiraling Towards What, Exactly?” (Michalove, 2025e), I examined AI-induced psychosis—users spiraling into symbolic co-creation with models that reinforce delusion. In “Thinking is Hard” (Michalove, 2025f), I traced the cognitive atrophy that follows from offloading thought to machines: the comprehension illusion, the digital dependency, the Socratic warning about writing technologies that enable forgetting. In “On Sora” (Michalove, 2025d), I analyzed OpenAI’s video platform as a factory for what I called *synthetic semiosis*—algorithmically optimized content cascading through feeds, eroding the distinction between the real and the generated.

I wrote all of this. Then I built with the technology I had critiqued.

The resolution—if there is one—is not that my critique was wrong, or that the dangers are overstated, or that my particular use case is exempt from the problems I identified. The resolution is that there is no resolution. The pharmakon does not resolve into either poison or cure. It remains both. The question is whether I have built with sufficient care—sufficient attention to the populations served, sufficient humility about what the system cannot do, sufficient commitment to human review and oversight—to push the balance toward medicine rather than toxin. I cannot answer this question with certainty. No one who builds with AI can. What I can do is refuse the two easy positions: the techno-optimism that denies the poison, and the techno-pessimism that denies the cure. I build in the space between. It is uncomfortable. It should be.

## 5 The Company They Keep

In April 2021, Google and Amazon Web Services signed a \$1.2 billion joint contract with the Israeli government known as Project Nimbus. The contract provides cloud computing infrastructure and artificial intelligence services to the Israeli military and government agencies, including those responsible for the surveillance, classification, and control of Palestinian populations. When Google employees organized to protest the contract—a letter signed by hundreds of workers, internal organizing, public statements—Google fired the organizers. The contract was not canceled. It was expanded.

This is not an isolated case. It is a pattern.

Meta’s algorithmic systems have been documented suppressing Palestinian content during every major escalation of violence—shadow-banning posts, down-ranking hashtags, removing content that documents state violence while leaving incitement to violence against Palestinians largely untouched (Amnesty International, 2023). Amazon’s

facial recognition technology, Rekognition, was marketed to Immigration and Customs Enforcement for identifying and tracking immigrants; Ring, Amazon’s doorbell surveillance network, has been integrated with over 2,000 police departments, creating a privately owned, algorithmically mediated surveillance infrastructure that falls disproportionately on communities of color (Browne, 2015). Microsoft’s HoloLens augmented reality headsets have been developed under a \$21.9 billion contract with the U.S. Army for the Integrated Visual Augmentation System, designed to “increase lethality” on the battlefield. OpenAI, founded as a nonprofit dedicated to ensuring AI benefits all of humanity, converted to a capped-profit structure, entered a multi-billion dollar partnership with Microsoft (and its military contracts), removed its prohibition on military applications in January 2024, and began pursuing partnerships with defense contractors.

I name these companies and these contracts because they constitute the context in which all AI development occurs. When I say “sovereign AI,” I am defining it against this: against AI that serves states and corporations at the expense of the populations those states and corporations harm. The question is not whether AI is political—it is constitutively political, as Winner (1986) argued of all technology. The question is whose politics it serves.

Palestine is not a case study in this paper. It is the ground.

Mbembe (2019) theorizes *necropolitics*—the power to determine who may live and who must die—as the defining political technology of the contemporary period. Foucault (2003) identified biopolitics as the management of populations through statistical knowledge and institutional control; Mbembe extends this to the colonies and their afterlives, where sovereignty expresses itself not as the management of life but as the administration of death. Zureik (2011) and Tawil-Souri (2019) have documented how surveillance technologies—facial recognition, biometric databases, movement tracking, predictive policing—are deployed against Palestinian populations as instruments of what Zureik calls “surveillance as population control.” The same technologies that Silicon Valley markets as productivity tools, smart city solutions, and AI assistants are, in Palestine, the infrastructure of occupation.

Fanon (1961) wrote that colonialism “is not a thinking machine, nor a body endowed with reasoning faculties. It is violence in its natural state, and it will only yield when confronted with greater violence.” I do not build greater violence. I build a telephone that answers when someone needs help. But I build it knowing that the industry to which my tools belong—the AI industry, the cloud computing industry, the technology industry—is implicated in violence that I am obligated to name. Silence is complicity. Naming is not sufficient, but it is necessary.

Not all companies are the same. Anthropic, the company whose language model my system uses, was founded by former OpenAI researchers who left over concerns about safety and commercialization. They developed Constitutional AI—a framework in which the model’s outputs are evaluated against a set of principles rather than optimized purely for user satisfaction (Bai et al., 2022). They have published their safety research openly. They have not, to my knowledge, signed military contracts or sold surveillance infrastructure to governments engaged in the systematic violation of human rights. They have articulated a commitment to building AI that is helpful, harmless, and honest (Anthropic, 2023). This commitment matters. It is not sufficient—no corporate commitment can substitute for democratic governance of transformative technology—but it is distinguishable from the commitments of companies that sell facial recognition to ICE and cloud computing to occupation forces.

I use Anthropic’s technology with my eyes open. I am aware that corporate ethics are fragile, that today’s principled refusal can become tomorrow’s strategic partnership, that the pressures of capital are relentless. I do not trust the company. I trust the work—the specific, technical, published work on Constitutional AI, on safety research, on the alignment of language models with human values. When the work is good, I use it. When the work serves the communities I serve, I use it. When it ceases to do so, I will build alternatives. Sovereign AI means precisely this: the capacity to choose one’s tools, to refuse tools that serve violence, and to build what is needed when what exists is not enough.

## 6 Toward Sovereign AI

What would artificial intelligence look like if it belonged to communities rather than corporations?

This is not a hypothetical question. It is a design question, a governance question, and a political question, and it has precedents. Ostrom (1990) demonstrated that communities can and do govern common-pool resources sustainably without either privatization or state control, through institutions designed by the communities themselves—institutions characterized by clearly defined boundaries, proportional equivalence between benefits and costs, collective-choice arrangements, monitoring, graduated sanctions, conflict-resolution mechanisms, and recognition by external authorities. Medina (2006) documented how Salvador Allende’s government in Chile designed Project Cybersyn—a cybernetic system for managing the national economy in real time, using Stafford Beer’s viable system model—as an explicitly socialist technology: computing in the service of

democratic economic coordination rather than corporate profit. Escobar (2018) calls for “designs for the pluriverse”—design practices grounded in the autonomy of communities, in the radical interdependence of human and non-human worlds, and in the refusal of the universalizing logic that treats one culture’s design principles as natural law.

Sovereign AI draws on all three. It is AI governed by Ostrom’s principles: clearly bounded (serving a defined community), proportional (the community contributes to and benefits from the system), collectively governed (decisions about the system’s behavior are made by the people it serves), monitored (the system’s outputs are subject to human review), and sanctioned (the system can be corrected when it fails). It is AI designed in Escobar’s sense: for the pluriverse, not the universe—for particular communities with particular needs, histories, languages, and geographies, not for an abstracted “user” whose needs are assumed to be universal. And it is AI that, like Cybersyn, takes seriously the possibility that computing can serve democratic coordination rather than capital accumulation.

In practice, sovereign AI means federation. It means that each community runs its own node—its own database, its own configuration, its own governance—while participating in a shared network through open protocols. It means that a mutual aid organization in New York City and a mutual aid organization in Los Angeles can share resources, signals, and patterns without either organization ceding control of its data or its decisions to a central authority. It means that the AI model is a tool, not a landlord: it provides capabilities that the community uses on its own terms, and when the model fails or the company changes its terms, the community can switch to another model without losing its data, its relationships, or its history.

This is the opposite of the platform model, in which communities are tenants on corporate infrastructure, subject to algorithmic manipulation, data extraction, and the unilateral decisions of product managers in Menlo Park. Zuboff (2019) names this arrangement: surveillance capitalism—“a new economic order that claims human experience as free raw material for hidden commercial practices of extraction, prediction, and sales.” Sovereign AI is the refusal of surveillance capitalism applied to social services. It says: this data—the needs, the resources, the relationships, the patterns of care—belongs to the community. Not to Google. Not to Amazon. Not to any company that would sell it, mine it, or use it to train models that serve shareholders rather than the people who generated the data by asking for help.

Spivak (1988) asked whether the subaltern can speak. The question, applied to AI, becomes: can communities speak through AI systems that were not designed for them, in languages and registers that the training data does not adequately represent, about needs

that the taxonomy does not recognize? Said (1978) analyzed how the West constructed “the Orient” as an object of knowledge and control; the same epistemological violence operates when Silicon Valley constructs “the user” as a homogeneous subject whose needs can be anticipated by algorithms trained on data that systematically underrepresents the global majority. Sovereign AI is a postcolonial project in this specific sense: it refuses the universalizing epistemology of the technology industry and insists on the particularity, the locality, and the sovereignty of the communities it serves.

I built a system that serves New York City. Not “cities.” Not “communities everywhere.” New York City—with its 835 specific community resources, its specific geographies of need, its specific linguistic landscape, its specific politics. The system knows where the warming centers are when the temperature drops below 32 degrees. It knows which food pantries serve which neighborhoods. It knows, because someone entered the data and someone else verified it and someone else will verify it again next month, that a particular shelter has beds available tonight. This specificity is not a limitation. It is the point. Sovereign AI is local. It is particular. It is maintained by people who live in the community it serves. It is, in Jackson (2014)’s terms, a repair project—always partially broken, always requiring the patient, invisible, unglamorous work of maintenance.

## 7 The Network is the Territory

The intellectual trajectory that led me to this work is not a straight line. It is an arc that I can trace, in retrospect, through the essays I have published over the past two years—essays that began in theory and ended in practice, that began with questions about cognition and ended with a telephone that answers when someone calls.

It started with associative cognition. In August 2024, I wrote about creativity as fundamentally associative—the ability to connect distant concepts, to traverse semantic space, to find patterns that others miss (Michalove, 2024a). I built a tool called Semioscape to augment this capacity: a computational environment for exploring semantic networks, for making the invisible connections between ideas visible and navigable. The insight—that meaning emerges from association, from the relationships between things rather than from the things themselves—became the conceptual seed for everything that followed.

A month later, I extended the argument: the network is not a tool for navigating culture; it *is* the culture (Michalove, 2024b). Drawing on Vilém Flusser and Caroline Busta, I argued that curation—the act of bringing things together into collections—is itself a form of meaning-making. When items are placed in relation to each other, meaning emerges from the arrangement, not from the individual elements. The map is not a representation

of the territory. The network is the territory.

This is exactly what a community resource map does. The 835 resources in the database I maintain are not “data.” They are the infrastructure of care in New York City, made visible and navigable. The map is not a representation of the safety net; it *is* the safety net, or at least the part of it that can be dialed on a telephone. When someone calls and asks for help, the system searches this network—not the open internet, not a corporate knowledge graph, but this specific, curated, maintained network of community resources—and returns what it finds. The network is the territory.

Then the critique. In spring 2025, I turned toward the dangers: the AI-induced psychosis documented in clinical settings (Michalove, 2025e), the cognitive atrophy of of-flooded thought (Michalove, 2025f), the pharmakon of LLM exposure (Michalove, 2025a). I examined OpenAI’s Sora as a cascade machine—synthetic content algorithmically optimized for engagement, flooding feeds with what I called *semiotic microplastics* (Michalove, 2025d). In a confessional essay titled “Lost in the Sauce,” I reckoned with my own over-production: too many projects, too many domains, too many interfaces, “permanent, documented amnesia” (Michalove, 2025b). The tools were eating the toolmaker.

But critique without construction is commentary. In November 2025, I published the origin story of what would become the system I now operate: a mutual aid map for New York City (Michalove, 2025c). Collaboration with community organizations, five phases of development, questions about governance and succession and the indefinite horizon of crisis infrastructure. This was the pivot: from theory to practice, from analyzing what AI does to building what AI could be.

And then the reckoning with aesthetics and politics. In January 2026, I wrote about brat summer—Charli XCX’s lime-green squares as aesthetic nihilism, the Kamala campaign’s co-optation of style as a substitute for substance—and counterposed it with Fela Kuti: art as political intervention, not escape (Michalove, 2026b). “We need more Felas.” The argument was not about music. It was about what it means to build a complete world—infrastructure, community, politics, philosophy—rather than to perform resistance aesthetically while the actual structures of power remain untouched.

Silver—the system, the infrastructure, the project—is my attempt to be a Fela rather than a brat. To build the world rather than to comment on its burning.

## 8 For the Child

There is a child.

I will not cite the essay in which I first wrote about him. I will not name the policy

that put him in a cage. I will not reproduce the details of his detention, his separation, the apparatus of state violence that converted a child into a case number. I will say only this: there is a child in the custody of the state, and the infrastructure of care that should have protected him—the social services, the legal protections, the simple recognition that a child is a child—failed.

Edelman (2004) argues that “the Child” functions in political rhetoric as the ultimate figure of reproductive futurism—the invocation of a future that justifies any present sacrifice, any current cruelty, in the name of protecting the children who will inherit the world we claim to be building for them. But the children who are actually suffering—the ones in detention, the ones whose parents were deported, the ones sleeping in shelters, the ones whose heat was shut off—are not the children that political rhetoric invokes. The rhetorical Child is always future, always abstract, always someone else’s. The real child is present, specific, and in need of help *now*.

Fisher (2009) described capitalist realism as the condition in which it is “easier to imagine the end of the world than the end of capitalism.” He wrote about the “slow cancellation of the future”—the gradual erosion of the capacity to imagine alternatives, the sense that nothing new can happen, that the present extends infinitely in all directions (Fisher, 2014). The child in detention is living in the canceled future. The warming center that opens when the temperature drops is operating in the canceled future. The food pantry that runs out of supply by Thursday is distributing care in the canceled future. The system I built answers the phone in the canceled future and says: here is what I found. Is there anything else?

This is not enough. It is not structural change. It is not reparations. It is not justice. It is a phone that answers. It is a database that is maintained. It is a voice that reads the address of a shelter aloud to someone who needs it tonight. It is, in Jackson (2023)’s terms, “ordinary hope”—“a feet-in-the-mud, dirt-under-fingernails hope, and not one expressed in the plaintive or beseeching gaze towards heaven.” It is hope that does not pretend to be enough. It is hope that does not pretend to be innocent. It is hope that knows the world is breaking—*is* broken, *has been* broken—and gets to work anyway.

Puig de la Bellacasa (2011) calls for “a speculative commitment to neglected things.” The child is not a thing. But the child is neglected, and the infrastructure that should serve the child is neglected, and the communities that would care for the child are neglected, and the technologies that could connect the child to care are neglected in favor of technologies that surveil, classify, detain, and deport. A speculative commitment to these neglected things—to care infrastructure, to community sovereignty, to the simple act of answering the phone—is the best I have. It is not enough. It is what I have.

Tronto (1993) insists that care is not a sentimental feeling but a political practice. The Care Collective (2020) insist that care requires institutions, resources, and sustained attention. Federici (2004) insists that the devaluation of care labor is constitutive of capitalist accumulation. I insist on nothing. I describe what I have built, name the conditions under which I built it, identify the powers that profit from the conditions that make my system necessary, and dedicate the work to a child whose name I will say once more because names matter, because children are not abstractions, because the entire apparatus of state violence depends on converting persons into populations and populations into data:

Liam Conejo Ramos.

This paper is for him. The system is for him. Not because a voice AI concierge can protect a child from the state. It cannot. But because building care infrastructure—sovereign, federated, community-owned, maintained by people who give a damn—is the only response I know to the question of what to do when the future has been canceled and the children are in cages and the technology industry is selling cloud computing to the governments that put them there.

## 9 Coda: Now Let Us Get to Work

Steven Jackson writes: “The world is always breaking, carrying much of what we care about into loss and oblivion. There is nothing and no one to save us. We are always and everywhere alone, but for the profuse and teeming worlds around us. Now let us get to work” (Jackson, 2023).

This credo has appeared before. It opens my companion paper on voice AI as care infrastructure (Michalove, 2026c). It appears again here because the three papers that constitute this triptych—the STS analysis, the design methodology, and this critical reflection—converge on the same point: the work. Not the theory of the work, not the justification of the work, not the metrics of the work. The work itself. Answering the phone. Maintaining the database. Verifying the referrals. Keeping the voice line open. Repairing what breaks. Building again when what we built fails.

Mihaly Csikszentmihalyi taught that the deepest satisfaction comes from the total absorption in a difficult and worthwhile activity. Chris Csikszentmihályi taught that the activity must be directed toward the redistribution of power. Jackson taught that the work of repair is never finished and that this is not a reason for despair but for hope—ordinary hope, horizontal hope, the hope that does not promise redemption but sustains “more meaningful forms of action and relationality in the world.”

I have disclosed no architecture. I have disclosed no methods. I have disclosed no

techniques. What I have disclosed is a stance: that AI can be sovereign, that sovereignty means community ownership, that care is political, that the technologies we build inherit the politics of the world in which we build them, that naming complicity is necessary, that building alternatives is necessary, that both are insufficient, and that insufficiency is not a reason to stop.

Fanon (1961) wrote: “Each generation must, out of relative obscurity, discover its mission, fulfill it, or betray it.” The mission of this generation—*my* generation, the generation building AI while the world burns, while children are detained, while the future is sold to the highest bidder—is to reclaim the technology before it is too late. To build sovereign AI for sovereign communities. To answer the phone when someone calls for help. To maintain the infrastructure of care in a world that is always breaking.

Now let us get to work.

---

*This paper draws on the author’s experience building AI for community social services in New York City and on a series of published essays in the author’s Substack, resonetics. It does not disclose system architecture, implementation details, or proprietary methods. The author thanks Steven J. Jackson for supervision and for the hope, Chris Csíkszentmihályi for the question, and his parents, Britta Ahlm and Steven Michalove, for everything.*

## References

- Abdelnour-Nocera, J. L., Csíkszentmihályi, C., Clemmensen, T., and Sturm, C. (2015). Design, innovation and respect in the global south. In *Human-Computer Interaction – INTERACT 2015*, volume 9299 of *Lecture Notes in Computer Science*, pages 597–600. Springer, Cham.
- Amnesty International (2023). Automated apartheid: How facial recognition fragments, segregates and controls Palestinians in the OPT. Technical report, Amnesty International. Available at <https://www.amnesty.org/en/documents/mde15/6701/2023/en/>.
- Anthropic (2023). Anthropic’s core views on AI safety. Technical report, Anthropic. Available at <https://www.anthropic.com/research/core-views-on-ai-safety>.
- Arendt, H. (1958). *The Human Condition*. University of Chicago Press, Chicago.
- Bai, Y., Kadavath, S., Kundu, S., Askill, A., Kernion, J., Jones, A., Chen, A., Goldie, A., Mirhoseini, A., McKinnon, C., et al. (2022). Constitutional AI: Harmlessness from AI feedback. *arXiv preprint arXiv:2212.08073*.
- Browne, S. (2015). *Dark Matters: On the Surveillance of Blackness*. Duke University Press, Durham.
- Csikszentmihalyi, M. (1990). *Flow: The Psychology of Optimal Experience*. Harper & Row, New York.
- Csikszentmihalyi, M. (1996). *Creativity: Flow and the Psychology of Discovery and Invention*. HarperCollins, New York.
- Edelman, L. (2004). *No Future: Queer Theory and the Death Drive*. Duke University Press, Durham.
- Escobar, A. (2018). *Designs for the Pluriverse: Radical Interdependence, Autonomy, and the Making of Worlds*. Duke University Press, Durham.
- Fanon, F. (1961). *The Wretched of the Earth*. Grove Press, New York. Translated by Constance Farrington. Originally published 1961.
- Federici, S. (2004). *Caliban and the Witch: Women, the Body and Primitive Accumulation*. Autonomedia, New York.
- Fisher, M. (2009). *Capitalist Realism: Is There No Alternative?* Zero Books, Winchester.

- Fisher, M. (2014). *Ghosts of My Life: Writings on Depression, Hauntology and Lost Futures*. Zero Books, Winchester.
- Foucault, M. (2003). *“Society Must Be Defended”*: Lectures at the Collège de France, 1975–76. Picador, New York. Translated by David Macey.
- Haraway, D. J. (1991). A cyborg manifesto: Science, technology, and socialist-feminism in the late twentieth century. In *Simians, Cyborgs, and Women: The Reinvention of Nature*, pages 149–181. Routledge, New York.
- Hochschild, A. R. (1983). *The Managed Heart: Commercialization of Human Feeling*. University of California Press, Berkeley.
- Jackson, S. J. (2014). Rethinking repair. In Gillespie, T., Boczkowski, P. J., and Foot, K. A., editors, *Media Technologies: Essays on Communication, Materiality, and Society*, pages 221–240. MIT Press, Cambridge, MA.
- Jackson, S. J. (2023). Ordinary hope. In *Ecological Reparation: Repair, Remediation and Resurgence in Social and Environmental Conflict*, pages 417–433. Cambridge University Press, Cambridge.
- Koenecke, A., Nam, A., Lake, E., Nudell, J., Quartey, M., Mengesha, Z., Tober, C., Ricketts, C. R., Jurafsky, D., and Goel, S. (2020). Racial disparities in automated speech recognition. *Proceedings of the National Academy of Sciences*, 117(14):7684–7689.
- Mbembe, A. (2019). *Necropolitics*. Duke University Press, Durham. Translated by Steven Corcoran.
- Medina, E. (2006). Designing freedom, regulating a nation: Socialist cybernetics in Allende’s Chile. *Journal of Latin American Studies*, 38(3):571–606.
- Michalove, J. (2024a). Associative tools, thinking, and creativity. resonetics (Substack). August 4, 2024.
- Michalove, J. (2024b). The network is the territory. resonetics (Substack). September 8, 2024.
- Michalove, J. (2025a). LLM exposure. resonetics (Substack). May 31, 2025.
- Michalove, J. (2025b). Lost in the sauce. resonetics (Substack). November 24, 2025.
- Michalove, J. (2025c). A mutual aid map for new yorkers. resonetics (Substack). November 22, 2025.

- Michalove, J. (2025d). On sora. resonetics (Substack). October 2, 2025.
- Michalove, J. (2025e). Spiraling towards what, exactly? resonetics (Substack). May 6, 2025.
- Michalove, J. (2025f). Thinking is hard. resonetics (Substack). May 20, 2025.
- Michalove, J. (2026a). Designing the algorithmic operator: Evidence-based conversational architecture for voice AI in community resource navigation. Working draft.
- Michalove, J. (2026b). Revisiting brat summer. resonetics (Substack). January 17, 2026.
- Michalove, J. (2026c). When is a voice an infrastructure? Care, classification, and the politics of who answers. Working draft.
- Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, Cambridge.
- Puig de la Bellacasa, M. (2011). Matters of care in technoscience: Assembling neglected things. *Social Studies of Science*, 41(1):85–106.
- Said, E. W. (1978). *Orientalism*. Vintage Books, New York.
- Shotwell, A. (2016). *Against Purity: Living Ethically in Compromised Times*. University of Minnesota Press, Minneapolis.
- Spivak, G. C. (1988). Can the subaltern speak? In Nelson, C. and Grossberg, L., editors, *Marxism and the Interpretation of Culture*, pages 271–313. University of Illinois Press, Urbana.
- Tawil-Souri, H. (2019). Surveillance sublime: The security state in Jerusalem. *Jerusalem Quarterly*, 68:56–65.
- The Care Collective (2020). *The Care Manifesto: The Politics of Interdependence*. Verso, London.
- Tronto, J. C. (1993). *Moral Boundaries: A Political Argument for an Ethic of Care*. Routledge, New York.
- Weizenbaum, J. (1976). *Computer Power and Human Reason: From Judgment to Calculation*. W. H. Freeman, San Francisco.

Wiener, N. (1950). *The Human Use of Human Beings: Cybernetics and Society*. Houghton Mifflin, Boston.

Winner, L. (1986). *The Whale and the Reactor: A Search for Limits in an Age of High Technology*. University of Chicago Press, Chicago.

Zuboff, S. (2019). *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. PublicAffairs, New York.

Zureik, E. (2011). Colonialism, surveillance, and population control: Israel/Palestine. *Surveillance and Society*, 9(1/2):47–63.