

Unlocking Public Value:

A Proposal for AI Opportunity Zones

Kevin Frazier





Author

Kevin Frazier

Kevin T. Frazier is the AI Innovation and Law Fellow at the University of Texas School of Law. His research focuses on artificial intelligence, innovation policy, and regulatory design. He has published with leading law journals such as the *Tennessee Law Review*, the *Rutgers Law Review*, the *Chapman Law Review*, and the *Fordham Intellectual Property, Media & Entertainment Law Journal*. He graduated from the Roberts D. Clark Honors College at the University of Oregon, earned his J.D. at UC Berkeley School of Law, and completed his M.P.A. at the Harvard Kennedy School.



Executive Summary

Every major technology has faced the same early verdict: too risky, too disruptive, too complicated to trust. The steam engine was accused of killing jobs; the internet of spreading lies; now artificial intelligence stands trial for everything from plagiarism to apocalypse. Yet, beneath the noise, AI is already helping farmers conserve water, drones fight wildfires, and teachers tailor lessons to each student. The real danger is not that we move too fast—it is that we move too slow and miss these gains entirely. Today, governments write rules around AI's costs but rarely measure its benefits, which leaves innovators to retreat and the public to wait. The answer is not to pause progress but to prove it: to create AI Opportunity Zones—new partnerships where governments get open access to advanced AI systems in exchange for testing them, transparently and rigorously, on society's toughest problems.



Introduction

The prevailing narrative surrounding artificial intelligence (AI) is one of caution; it is dominated by a focus on potential costs, risks, and liabilities. This discourse is prudent, but has become dangerously imbalanced, and is creating a policy environment that inadvertently stifles the very innovation necessary to unlock AI's profound societal benefits. Lawmakers, driven by measurable and often anecdotal, politically salient harms, are crafting regulations that do not account for the vast, though harder to quantify, public good that widespread AI adoption could generate. This has created a chilling effect, particularly for private actors wary of litigation and negative press. Rather than boldly pursuing projects with the potential to promote human flourishing, innovators retreat from sensitive domains. The result is a major societal loss. The fact that governments, the actors best positioned to manage these risks and deploy AI in high- impact public systems, are themselves hampered by the prohibitive costs of cutting-edge models, a lack of technical support, and public unease, compounds this loss. Instead, AI labs push ahead with products of questionable social utility that sustain public skepticism about AI's value on its face.

AI now faces a reliability Catch-22: we will not trust it with the domains that matter most until it performs flawlessly there—but it cannot reach that level of performance until it is allowed to learn in those very domains. This essay proposes a new framework to break this impasse through the creation of “AI Opportunity Zones.” This initiative calls for leading AI laboratories to establish competitive grant programs offering unlimited, no-cost access to their most advanced models for public sector entities. They would award these grants to governments—from municipal departments and school districts to state agencies and public universities—committed to fundamentally redesigning a public service or system with AI at its core. This is not a call for merely tinkering at the margins through short-term pilot initiatives. Instead, adoption of this framework would result in systemic, AI-forward transformation.

In exchange for access, training, and technical support, participating governments would commit to long-term projects, rigorous, transparent assessment, and the implementation of control groups to empirically demonstrate the value AI adds. This private-public partnership creates a powerful win-win: governments can fulfill their mandate to serve the public more effectively and efficiently, while AI labs can generate robust, real-world data on their

technology's societal benefits, combat the narrative that their sole goal is profit, and accelerate a virtuous cycle of innovation and adoption.



A Proposal for AI Opportunity Zones

The Peril of an Unbalanced Narrative

Technological progress has always been a jagged journey. From the Industrial Revolution to the dawn of the internet, innovation has introduced significant societal shifts. It has created new efficiencies and opportunities while simultaneously displacing established industries and imposing short-term, concentrated costs—often on already vulnerable communities. History's lesson is not that we should halt progress to avoid these costs, but that we must anticipate them, measure them, and implement responsive policies, particularly for the most affected communities. The challenge with AI is that our current policy and public discourse has become fixated on the costs, including job displacement, algorithmic bias, and privacy concerns, while largely ignoring the benefits.

This imbalance stems from a fundamental measurement problem that predictable patterns of human cognition have magnified. The costs of AI, whether real or projected, are often specific, identifiable, and easily sensationalized. A factory worker displaced by AI is a tangible story with a human face. An instance of algorithmic bias in a hiring tool can be documented, litigated, and broadcast, triggering our innate sense of injustice. These are critical issues that demand attention and redress.

AI's benefits, however, are frequently diffuse. They are distributed across a large population and accrue over a longer time horizon, making them far more difficult to measure, label, and discuss. It is not the case that AI lacks potential to do good; we are not ignoring its benefits because they are insignificant, but because they are systemic. Consider the work of researchers at Texas A&M University on Soma Tech, an AI-powered irrigation system. By analyzing soil, crop, and weather data, it precisely delivers the amount of water needed and reduces water waste in drought-prone regions. Its benefit is not a single, dramatic event, but a fractional improvement in water efficiency across thousands of acres, leading to greater food security, lower water prices for consumers, and enhanced environmental resilience. Similarly, the [WATERWISE project in Greece](#) uses AI to analyze climate scenarios and tourism data to

forecast water demand and ensure the stability of a vital resource for an entire nation. These are not easily captured in a headline, yet their cumulative effect is immense.

The same dynamic applies in security and defense. [Florida International University's SHIELD](#) system uses machine learning to allow drones to detect and recover from cyberattacks mid-flight, which is a critical capability for autonomous systems. Meanwhile, the U.S. Army is rolling out [TurbineOne's AI](#), which allows soldiers to detect threats on the battlefield using devices that work without a cloud connection. The value of a successful cyberattack being thwarted or a threat being neutralized on the front lines is nearly impossible to quantify, yet the security it provides underpins society's basic functioning. A failure to properly identify and disclose these and other AI benefits is a problem that transcends the realm of communications to become a profound policy concern. Each of these AI use cases demonstrates the kind of domain where Opportunity Zones would convert incremental progress into validated systemic redesign.

Because policymakers and the public are not adequately weighing these distributed, preventative, and systemic benefits, the entire innovation ecosystem is being skewed toward excessive caution. This environment creates a powerful incentive for private actors to avoid risk. Why would a company deploy a novel AI system in a sensitive field like education or healthcare when the potential downside includes class action lawsuits and reputational ruin, while the upside—broad societal progress—is a public good from which they cannot fully capture the value? This reticence is a major loss for the AI labs themselves, which depend on the wide adoption of their technologies to fuel further research and development and to earn and sustain their users' trust and that of the public generally. More importantly, it is a staggering loss for society, which is being denied the fruits of innovation.



The Government's Role and the Current Bottleneck

Governments are the actors best suited to break this logjam. Public entities oversee the very systems—healthcare, education, transportation, public safety—where AI holds the most transformative potential. Unlike private corporations, governments can run large-scale experiments with a clear public mandate and explicit legal authority. They possess unique, population-scale datasets and a mission of public welfare rather than profit drives them, allowing them to pursue projects with long-term, societal returns. They can operate with a degree of transparency that fosters public trust and can structure AI initiatives in a way that is clearly permissible under existing law.

Recent pilot programs, such as the [State of Pennsylvania is engagement with generative AI](#) to assist state employees, are important and laudable first steps. They demonstrate a willingness among public officials to explore new technologies and build institutional capacity. These efforts, however, are distinct from what I am calling for here. Such pilots typically focus on augmenting existing workflows—using AI to draft memos more quickly, summarize regulations, or answer common citizen inquiries. This approach is valuable for building familiarity and achieving incremental efficiencies, but it does not unlock the technology's transformative power. It grafts AI onto legacy systems rather than using AI to design new, fundamentally better systems.

We will realize AI's full potential through incremental adjustments to already antiquated agency systems and processes. The path to harnessing AI instead requires enabling it to do entirely new things or to solve old problems in completely new ways. Yet, for governments to take this larger leap, they face a significant barrier: cost. The leading AI models are computationally expensive to build and operate, and labs must recoup these investments. If public universities, city governments, and federal research entities are expected to pay enterprise-level prices for access, they will lack the budgets for the ambitious, long-term experiments in systemic redesign that they desperately need. This is not a failure of will, but a failure of resources.

A New Framework: AI Opportunity Zones

To unlock the next chapter of public sector innovation, the country's leading AI laboratories should launch a competitive program—AI Opportunity Zones—that

gives governments open access to advanced AI systems in exchange for ambition and accountability. This set up is distinct from charity or a publicity stunt. If properly designed, it will instead take the form of a strategic investment: a partnership between the creators of transformative technology and the institutions most responsible for translating progress into public value.

The program would operate much like a national call for ideas. City agencies, state departments, school districts, and public universities would compete to redesign one essential public service from the ground up using AI. Each applicant would begin with a simple but demanding question: If you could rebuild this system today, with AI as a native component rather than a late add-on, how would it look?

An independent nonprofit consortium composed of AI labs, universities, and nonpartisan civic organizations would oversee the process to guarantee transparency and public accountability. This group would issue an open request for proposals, review submissions, and select the most promising projects. The competition would begin with short letters of intent describing the problem and the vision. From there, a smaller set of finalists would receive hands-on technical help to develop full proposals, which would ensure that even small or resource-constrained jurisdictions could compete on ideas rather than budgets.

Winning projects would receive far more than free model access. Each would be paired with a team of AI engineers and policy specialists—akin to the [forward deployed engineers](#) the private sector is using—who would work alongside public officials from day one. Government teams would attend an intensive onboarding “boot camp” designed to build technical fluency, ethical awareness, and experimental discipline. Embedded experts from the partner labs would spend much of the first year inside the agency, ensuring that the project is a true collaboration, not a handoff. This embedded model would allow governments to build lasting internal capacity rather than dependency on outside vendors.

Every project would be judged by the same demanding standards. Each must involve a systemic redesign, not a surface-level upgrade. Each must run long enough—at least a full year—to yield measurable results. And each must include a rigorous evaluation plan: a transparent method for comparing the redesigned system to a control group operating under the old model. These safeguards turn enthusiasm into



evidence. They ensure that success is provable, not asserted, and that failure yields lessons rather than scandal.

To maintain accountability and accelerate learning, participating governments would submit quarterly progress reports and join a national Community of Practice, a peer network where teams share data, troubleshoot challenges, and publish findings in public view. The same consortium that selects projects would also manage this community, building a national archive of what works, what does not, and why.

Finally, every proposal would have to meet the highest standards of privacy, data protection, and ethics. Citizen data would remain under public control. Models would be tested for fairness and security. And all outputs, successes and setbacks alike, would be published in plain language for the public to review.

AI Opportunity Zones would turn governments into laboratories of applied intelligence, places where public servants and technologists work side by side to prove, in measurable terms, that AI can make life better, fairer, and more efficient for everyone.

Addressing Counterarguments

Skepticism about a partnership of this depth is warranted. Critics will ask why the world's most profitable AI labs would give away their most advanced technology to governments that unable to pay for it. The reflexive assumption is that they would never agree to such terms. Yet, that view misses both the strategic and reputational logic at play.

For AI labs, the reputational dividend of this program could be enormous. [Right now, much of the public conversation](#) casts them as reckless, profit-driven enterprises racing ahead without regard for social consequence. An initiative like the AI Opportunity Zones offers a concrete rebuttal to that narrative. It would allow labs to demonstrate, in full public view, how AI can strengthen institutions rather than disrupt them, and how AI-forward organizations can operate safely, transparently, and in alignment with democratic values. The resulting stories would not be about another chatbot or social-media tool but about teachers, doctors, and city engineers doing their jobs better because of AI. In a policy environment defined by suspicion, that kind of proof is priceless.

Participation would also serve a clear business interest. These projects would show both public and

private prospective clients how to integrate AI into complex, regulated environments. A school district that successfully reimagines special education or a transportation agency that cuts congestion with AI-driven scheduling becomes a living advertisement for what responsible adoption looks like. By funding evidence of success, labs would seed new demand far beyond the public sector. The next hospital system, insurer, or logistics firm will not need to be convinced that AI can work, because they will have seen it working.

Of course, guardrails are essential. The program's design prevents any single lab from monopolizing access or locking governments into proprietary systems. The competitive request-for-proposals process ensures multiple labs can participate on equal footing. Every grant agreement would include strict conditions on open standards and data portability, guaranteeing that governments retain control of their data and can migrate to other platforms once the grant period ends. The government—not the lab—remains the lead actor: setting the goals, owning the results, and defining what success means for the public.

Some will worry about the optics of failure. What if a high-profile project stumbles and becomes fodder for cynics eager to declare AI unfit for public use? The truth is that innovation without the possibility of failure is not innovation at all. The real cost lies in refusing to try. The Opportunity Zone model builds failure into the process as an engine of learning. Each project's requirement for a concurrent control group creates a built-in diagnostic tool: if the AI-enhanced system underperforms, the evidence will be clear, and the program can adjust in real time. Failure, when measured and published, becomes a public good.

Another legitimate concern is that grants might flow disproportionately to large or well-staffed governments. That risk can be mitigated by thoughtful design: the selection process should favor proposals from smaller or under-resourced jurisdictions and projects explicitly aimed at advancing equity. Technical assistance and embedded experts are not afterthoughts; they are the core mechanism for ensuring every community, regardless of its tax base or technical workforce, can compete on ideas rather than resources.

In short, what might look at first like an act of corporate altruism is, in fact, enlightened self-interest. By aligning their technology with the public's most trusted institutions, AI labs can earn



something no marketing budget can buy: legitimacy. Governments gain the capacity to lead rather than lag in technological change. And the public gains visible, verifiable proof that AI can deliver progress, not just promise.

A Win-Win for Labs and the Public

Ultimately, the AI Opportunity Zones framework is designed to create a partnership where public and private interests genuinely reinforce one another. For governments, it offers something that is otherwise out of reach: a supported and de-risked pathway to modernization. By sharing costs and embedding expertise, the program allows agencies to take a calculated, bounded leap toward innovation rather than a blind one. Governments gain access to frontier technology, engineering talent, and analytical capacity that would normally cost millions of dollars. They can then use those resources to redesign systems that have long outgrown their tools. In doing so, they fulfill their most basic democratic obligation: to deliver services that are not just efficient but adaptive, equitable, and resilient. The program would also help cultivate a new generation of technologically fluent public servants. These leaders will understand how to integrate emerging tools into mission-driven governance rather than being perpetually reactive to them.

For AI labs, the case is equally strong, if not stronger. Participation in Opportunity Zones would allow them to rewrite the story currently told about their industry. Instead of being seen as remote and self-interested, they would be recognized as active partners in solving society's most pressing problems. Each successful project would provide visible, verifiable evidence that AI can make public life better, from reducing wait times for social services to improving water efficiency in drought regions or detecting cyberattacks in real time. The data generated through these collaborations would be uniquely valuable: large-scale, real-world evidence of how AI performs under public constraints. This evidence can then feed directly into building safer, more capable systems.

And the benefits would not end when the grants do. By proving what AI can achieve in transparent, high-stakes settings, labs would open entirely new markets—not speculative, hype- driven ones, but mature markets built on trust. Hospitals, universities, and corporations will be more willing to adopt systems that have been tested in the open and validated by governments accountable to their citizens. In that sense, AI Opportunity Zones are not

a philanthropic gesture, but a strategic bet on the future of AI itself. The labs that help build that future will be the ones best positioned to lead it.

Conclusion

The current trajectory of AI policy and discourse is one of constraint. In many cases, safeguards around when and how AI is deployed are essential and necessary, but, a framework built solely on preventing harm will never allow us to realize the technology's full potential for good. We are meticulously pruning the weeds while failing to plant the seeds of a forest. The AI Opportunity Zones proposal offers a proactive, optimistic, and evidence-based path forward. It is a call for a strategic alliance between AI's creators and the stewards of our public institutions.

By empowering governments to become leaders in AI innovation, we can shift the narrative from one of fear to one of possibility, build a rich portfolio of evidence demonstrating the profound public value of this technology, and ensure that the next chapter of technological progress is one that benefits all of society.

Kevin Frazier is the AI Innovation and Law Fellow at the University of Texas School of Law and co-host of the Scaling Laws podcast.

