

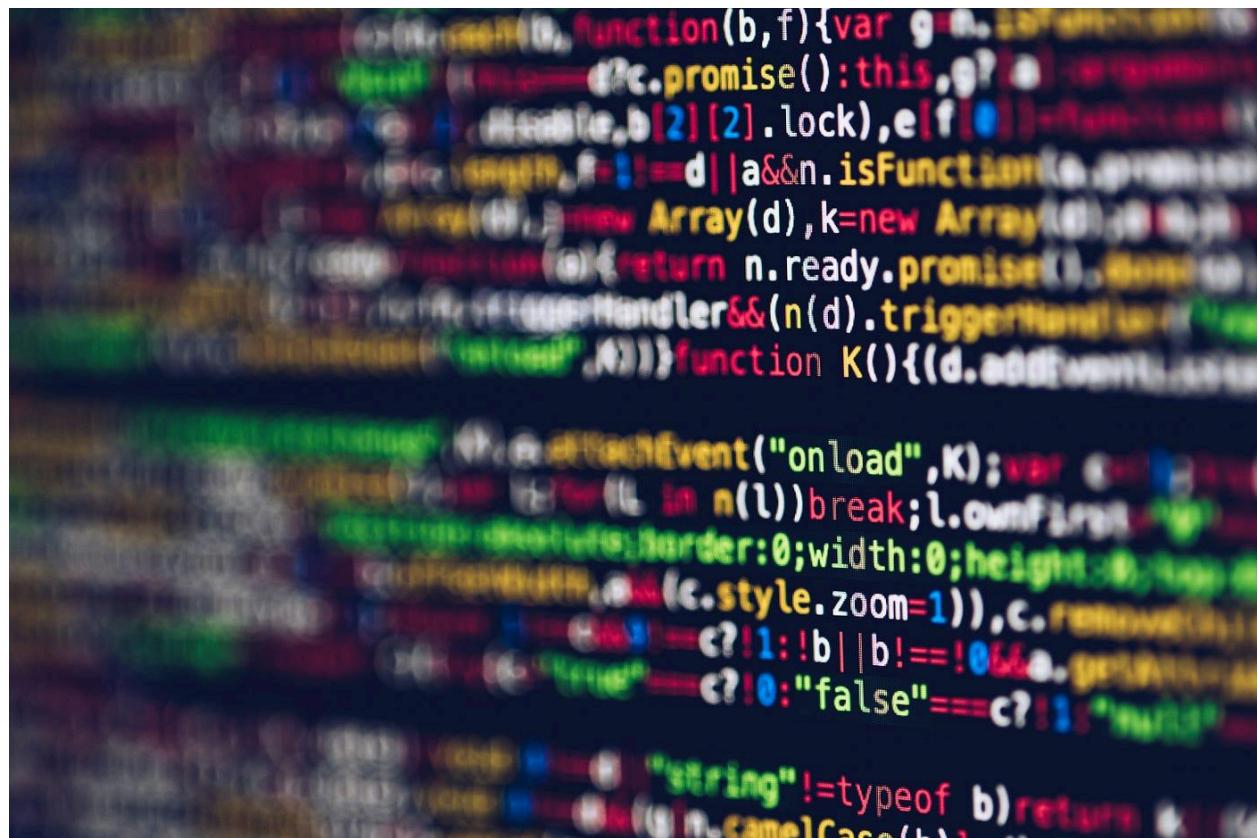
The below was originally published and is also available at:

<https://engineadvocacyfoundation.medium.com/startups-and-ai-policy-how-to-mitigate-risks-seize-opportunities-and-promote-innovation-ab3e66cea78f>



Startups and AI policy: how to mitigate risks, seize opportunities, and promote innovation

By Min Jun Jung and Nathan Lindfors, Engine Advocacy & Foundation



AI is dominating headlines and occupying the minds of [policymakers in Washington](#) concerned with how the technology will transform the economy. There are opportunities and challenges presented by AI, prompting several key policy debates, including about bias, competition, intellectual property, the workforce, and more. But the AI ecosystem is vast and diverse — including companies of all sizes who rely on different business models and touch many industries — and it's important policymakers consider the entire ecosystem as they debate changing legal and regulatory frameworks to preserve the ability of startups to innovate with AI, grow, and succeed.

What do we mean when we say AI?

Artificial intelligence encompasses a wide range of applications and functions, but it is often tricky to define — especially if those definitions, when written in law, form the basis of obligations or liability. In its essence, though, [AI describes](#) a branch of computer science that enables machines to perform tasks typically requiring human intelligence such as pattern recognition, problem-solving, and decision making.

Although artificial intelligence and machine learning are often used interchangeably, it is important to recognize that not all AI constitutes machine learning. Machine learning is a subset of artificial intelligence that uses mathematical models to enable a computer system to learn and improve without direct instruction. While the majority of startup and headline-grabbing AI applications are based on machine learning, some non-machine learning uses of AI include rule-based systems such as chess-playing, Kasparov-slaying [Deep Blue](#) that are based on large sets of predefined instruction.

Another important distinction to make is between generative and non-generative AI. While generative AI — including systems like ChatGPT and others that can generate text or images that resemble human creation — has recently been garnering significant attention, AI is deployed in a wide range of non-generative applications from [recruiting and job searches](#) to [self-driving cars](#). And chances are you've been using generative and non-generative AI for years, through things like autocomplete in Google Search or email spam filters — or to stick with the chess example — games on your phone.

How does AI work?

Artificial intelligence uses a combination of data and algorithms to perform human-like cognitive tasks. A prominent component technique for developing artificial intelligence is [machine learning](#), which processes massive amounts of information through algorithms to identify patterns and make predictions. The process begins with a dataset. Vast amounts of information are acquired and prepared to remove biased and irrelevant information. Most machine learning methods also involve labeling the data to help match inputs to corresponding outputs. The next step is the training phase, where the AI algorithm analyzes these large datasets and identifies patterns and correlations in an iterative process. The algorithm makes guesses and refines its accuracy through iterations, becoming increasingly proficient at identifying features and patterns until there isn't much room for improvement. Once trained, the AI model — the embodiment of the trained algorithm — can apply its acquired knowledge to new and unseen information, using the recognized patterns to make predictions through a process of reasoning called inference.

Most AI is built on prediction, and that ultimately has implications for model outputs and how regulation of the technology is best understood. By way of brief examples, text or image generative AI works by predicting the next word or adjacent pixels. That's why when you ask an AI model for legal precedents, it generates things that [look like legal precedents](#) — even if those precedents don't exist. If you ask for an image of high-five, the hands may well have too many fingers since high fives involve lots of fingers next to each other — but the model might not understand that humans' hands only have five fingers. In a healthcare setting, AI can be used to [predict heart attacks](#) — but it is still just a prediction — it's possible for someone with low to no risk of heart attack to still experience one. That outcome would have negative consequences, but the use of the AI technology is still likely to save more lives than not.

How are startups leveraging AI?

While large AI models created by large companies dominate headlines, startups are harnessing the potential of artificial intelligence to solve many pressing issues. The dynamic and rapidly evolving nature of artificial intelligence also offers ample opportunities for startups to explore innovative ways to utilize the technology, enabling startups to create new business models and carve out unique niches in the market. For example, startups are using AI to [monitor and ensure the health of bees](#), [detect when an elderly person falls](#), or [enable better sustainability practices](#). Startups are using AI to counter historic biases in [health](#), [lending](#), and [employment](#). And startups are using AI to help us have fun too: [teaching us to play games](#), [finding events we're interested in](#), and [helping us take better vacations](#).

What are the costs of AI for startups?

Artificial intelligence holds tremendous potential for enhancing the work of startups, but the costs associated with developing, training, and operating AI models can be daunting, particularly for small tech companies. The average seed-stage startup is working with around [\\$655,000 a year](#). By comparison, Google spent more than [\\$31 billion](#) on AI R&D in 2022, the cost of training OpenAI's GPT-3 ran upwards of [\\$4 million](#), and operating ChatGPT on Microsoft's Azure cloud infrastructure amounts to around [\\$100,000 per day](#).

For AI, hardware costs involve investing in powerful machinery with advanced computer chips and GPUs, which can cost upwards of [\\$10,000 each](#). On the software side, the collection, storage, and processing of data needed to build models can incur [significant investments](#) of time, labor, and money that increase as datasets grow larger. Data availability and quality proves to be a unique challenge for startups. While more established companies with sizable customer bases already have a stream of data on which to train AI models, startups typically [do not have access](#) to sufficient data. Many AI startups also find themselves at [a disadvantage to larger firms](#) that can lean on their name recognition to form partnerships with other enterprises to access their proprietary data. Finally, hiring skilled computer engineers and data scientists to develop and train the algorithms is expensive, with average base salaries for AI developers ranging up to [\\$150,000 a year](#).

Building unique models from scratch is challenging and incredibly expensive, and most startups developing their own AI models are doing so in a market niche, rather than trying to build general, broadly applicable foundation models. For example, [UnaliWear](#) uses sensors in their wrist worn watch and their AI model to detect when someone falls so that a medical alert center can be notified and the individual that fell can receive assistance. Their model is based on actual falls and gets better over time, in addition to learning an individual wearer's behavior to distinguish a fall versus 'flopping' into a chair. As another example, [BeeHero](#) uses low-cost sensors placed in hives to monitor the health of the hives and optimize hive placement to increase crop yield.

With the high initial expenditures associated with in-house AI model development, many startups are building with open source models or models from established AI companies. Building from [open source](#), fine-tuning other's models, or pinging the application programming interface at a larger AI company are all less expensive options, but aren't without cost and their own unique challenges. For example, OpenAI charges [6 cents](#) for about every 750 words of output on its GPT-4 model. And startups must navigate intellectual property and

documentation issues as they build with others' technologies. Furthermore, the integration of others' models can render startups susceptible to price hikes, access constraints, or regulatory changes aimed at the large companies that developed the models.

Despite the costs and challenges of creating and utilizing AI, there is ample opportunity for startups to flourish in the space. With smart AI policy, startups can safely develop, harness, and deploy AI technology to amplify economic growth, accelerate innovation, and improve quality of life.

Policy issues:

How should policymakers approach mitigating risks around bias and AI?

Artificial intelligence holds incredible promise, but it is important to be clear-eyed about potential risks associated with the technology. One significant concern revolves around the potential for bias and patterns of discrimination that are perpetuated by AI systems because they are trained on data that is biased or created by teams that are predominantly white and male. In a similar vein, AI's so-called [black box problem](#) can limit transparency about how AI models arrive at certain outputs.

Addressing risks around AI should begin with existing law and guidance. Existing legal frameworks, including [civil rights](#) and [consumer protection](#) law, for example, [speak to many of the issues](#) raised around AI by policymakers and the general public, like discrimination, bias, or deceptive practices. Agencies of jurisdiction should evaluate how AI interacts with the laws and regulations they are responsible for enforcing and disseminate proactive guidance to ensure companies understand their obligations as they develop new technologies using AI. Building potential AI rules with existing law in mind is critical to avoiding overlaps or contradictions that will create additional and unnecessary layers of cost and confusion which would overwhelmingly weigh on startups, and could bog down agencies tasked with enforcing the law.

A balanced regulatory environment is critical to mitigate risks and equip regulators with the resources to combat bad actors, while avoiding burdening startups and socially-beneficial innovation. To achieve balance, policymakers need to recognize and allow for unforeseen positive uses of technology, and must

avoid mitigating innovation as they strive to mitigate risks. As one example to illuminate why, e.g., overbroad definitions threaten progress, consider discrimination and bias in lending. This abhorrent practice is [already illegal](#), but some policymakers believe it is critical to build on existing frameworks in response to AI. Should they do so, they'd need to keep in mind that many innovators are solving this problem for their communities by using AI for [equity-enhancing purposes](#), like building a unique model to extend credit to immigrants and other underrepresented groups that lack a credit score. Should an updated framework extend too broadly, that could impinge on the ability of startups to innovate with AI to solve similar societal problems.

How is startup competitiveness impacted by regulation?

Startups have comparatively fewer resources than larger market competitors and less ability to maneuver in response to regulatory changes, meaning the regulatory environment directly impacts their competitiveness. Uniform regulatory environments are critical for startup success as fractured, “[patchworks](#)” of regulation [add to burdens](#) that sap already [limited startup resources](#). Public and private entities and standard setting bodies have created standards and other tools for mitigating AI risk. These (often collaborative) efforts are critical to creating useful, balanced resources to guide AI development and broader public policy considerations around AI. The National Institute of Standards and Technology’s [AI Risk Management Framework](#) provides a useful resource for mitigating AI risk and provides a glossary of AI terms, which can be critical as part of fostering a uniform, consistent environment in any potential future regulatory framework.

AI is a data driven technology, and better access to more and higher quality data gives industry incumbents a leg up. The U.S. lacks a federal privacy law — instead myriad varying state laws create a patchwork of uneven, confusing, and costly rules that undermine startup competitiveness while simultaneously leaving parts of the country uncovered. Additionally, several concerns related to AI hinge on questions of privacy, making a [uniform national data privacy law](#) a useful part of the policy response to AI. Policymakers can simultaneously respond to [privacy-related concerns](#) while creating consistency and improving the competitiveness of startups.

Ultimately, balanced, clear, and consistent rules are key to maintaining startup competitiveness while addressing possible AI risks. Fortunately there are a few useful methods for encouraging best practices and promoting balanced regulation that are found in other parts of the law. For example, safe harbors, like those found in [cybersecurity](#) and privacy law, for example, work well to incent

adherence to best practices without rigid mandates or threats of severe punishment. Likewise, [regulatory sandboxes](#) can enable startups to experiment with new technologies without the burdens of strict rules and facilitate knowledge sharing between companies and regulators.

At the same time, there are hallmarks of inherently unbalanced regulation that policymakers should seek to avoid. An otherwise burdensome regulatory environment *with* a sandbox for startups is still a burden to innovation. Startups only want to enter sandboxes if they can eventually exit with a commercializable product and succeed in the marketplace. They can only successfully exit if the broader regulatory environment is conducive to innovation and scaling small companies to success. Many startups meanwhile will always forgo the sandbox environment due to investor pressures, product fit, or other factors — the broader regulatory landscape has to work for them too. (And it's important to remember that sandboxes require resources from regulators to be successful.)

Finally, the inclusion of applicability thresholds can be an indicator of unbalanced regulation. If policymakers feel the need to include various thresholds for obligations, it is because they recognize some obligations are not feasible for all (especially small) companies like startups. This is used as a crutch for strenuous regulation, and often [results in startups being subject](#) to practices that industry incumbents were not and only undertook much later in their development when they could leverage additional resources. Approaching regulation in this way inhibits scalability and threatens to cement incumbent companies. Moreover, it is imperative to steer clear of ex-ante regulations that create barriers to entry for startups. Mandatory certification or licensing schemes could create “[regulatory moats](#)” that bolster the power and position of large companies that are already established in the AI ecosystem while hindering startups from entering or succeeding in the market.

How does AI interact with intellectual property?

Existing intellectual property frameworks work well and can and should be applied to AI. Still, many policymakers and others are exploring and advocating for changes to intellectual property laws in response to the latest AI developments. The Senate Judiciary IP subcommittee has [held a series of hearings](#) on the topic, where policymakers have suggested updates to copyright and patent law. The Copyright Office [held a series of listening sessions](#) on AI, while some large rightsholders have [sued](#) AI companies for alleged infringement. And others have [sued](#) seeking to have AI recognized as a rightsholder.

For copyright and AI, in the interest of promoting progress and innovation, it would be best for policymakers to support legal interpretations that establish that use of information and content to build AI is lawful because it is [a noninfringing use](#). The alternative, fair use, is decided on a case-by-case basis. Proceeding through litigation to establish that a specific use is fair is costly and not dispositive for all future (even similar) uses of data. So while it is a fair use for AI to ingest and process data, it is more efficient to conclude that such uses are not even infringing. AI policy should seek to streamline innovation and must avoid endorsing changes to the law that will entrench incumbent entities and industries.

Similarly, current law around patent eligibility is critical to ensure only truly novel inventions are patentable, and to avoid bad faith litigation that arises from low-quality patents. AI policy should similarly be mindful of the need for [a balanced patent system](#) in technological innovation. [Section 101](#) of the Patent Act defines what is and is not eligible for patent protection, and as the Supreme Court made clear in [Alice Corp. v. CLS Bank International](#), merely performing an abstract idea using a computer does not make it patent eligible. Currently, abstract ideas, laws of nature, and natural phenomena cannot be patented — so a company cannot patent and seek to own, e.g., [the idea of scheduling](#) medical appointments using a computer; the [process of collecting](#), analyzing, and displaying data; [the idea of filtering](#) e-mail; or a [human gene](#). The same principle does and should apply for AI. Barring patents on abstract ideas means no one company may own those basic concepts of running a business and limits the existence of low quality patents that patent assertion entities often assert in [frivolous, abusive cases](#).

Current IP laws work well to incent innovation while mitigating abuse. Still, some advocates have sought government [agencies to recognize](#) AI as an inventor or [urged Congress](#) to change the law to enable AI to be recognized as an inventor or co-inventor, but this is not necessary to incentivize innovation. Startups and others continue to innovate and involve AI in the innovative process without such inventorship considerations, and humans tend to be sufficiently involved in these processes to be named inventor of the resulting invention.

What does AI mean for the job market?

The upheaval of jobs is to be expected as a result of technological progress, but the thoughtful development of the AI ecosystem can lead to the creation of new and better jobs. Throughout the economic history of the United States, technological upheavals have consistently led to an [expansion of employment opportunities](#), rather than a contraction. For example, most jobs that exist today

did not exist before the [Second World War](#). Although artificial intelligence can automate many of the processes that were formerly handled by humans, there is [not a finite amount of work](#) to be done. Like technological revolutions previous, AI can increase our productivity and lead to expanded — but different — job opportunities. As part of this process, people will need to be trained and retrained for the jobs of the future. AI policy must facilitate the allocation of talent to jobs that cater to the evolving demands of the modern workforce and advancing technology that enhances our [quality of life](#).

To ensure that the benefits of AI development accrue to society in the face of job market transformations, programs to help upskill and retrain workers are necessary. Currently, [STEM talent is in short supply](#) and is needed to fill critical roles in the technology sector and at startups. As AI development continues to accelerate, demand for high-skilled engineering talent is likely to increase further. Policymakers should take an all-of-the-above approach to AI skilling and upskilling, leveraging traditional STEM education, private sector incentives, government resources, and realigning existing education strategies. Workforce programs additionally should place particular emphasis on the nexus to technology and ensuring that all can equitably participate, especially given [existing gaps in access](#) among underrepresented communities.

Developing STEM talent in traditional university settings is important but not sufficient. Policymakers should create incentives for the private sector to upskill and reskill their workforces. Reskilling later in life is likely to occur outside of a traditional university setting, for which public and [private credentials](#) and training programs can play a useful role. Accreditation agencies should consider new categories of accreditation to both help individuals recognize the programs worthwhile of their time and resources, and to help employers understand the qualifications of prospective employees. Incentives can also be used to encourage hiring of reskilled, talented individuals trained through such programs that might not possess training through traditional channels. Government resources — particularly those tailored toward AI-related education, like the contemplated [National AI Research Resource](#) can and must also play a critical role.

How can policymakers support innovation?

Policymakers play a pivotal role in not only implementing regulation that mitigates risks without compromising competitiveness but also actively nurturing innovation and ensuring equitable market opportunities for companies. As some ways to do so, government can work to bolster AI talent pipelines, open data and compute resources to startups, and disseminate tangible guidance on risk mitigation.

The government should create and fully fund the contemplated National Artificial Intelligence Research Resource (NAIRR) which will provide compute, datasets, and educational resources for startups, students, and academics. The [NAIRR](#), as designed, will be managed through the National Science Foundation by an outside entity and stands to benefit startups by improving talent pipelines, enabling AI research, and providing resources directly to startups. The NAIRR was the product of a robust congressionally chartered [task force](#) process that included stakeholders from government, industry, and academia and sought multiple rounds of stakeholder feedback. [Engine](#) and [startups themselves](#) weighed in throughout the task force process to ensure the resource would be designed with the needs of [entrepreneurs of all backgrounds](#) in mind. The government now must follow through to implement the resource that stands to promote innovation.

Government has already developed useful resources around responsible AI development, like the National Institute of Standards and Technology's AI risk management framework, and should further facilitate the dissemination and use of such resources. Startups routinely look to expert resources like Risk Management Frameworks developed by NIST, including those around [cybersecurity](#), [privacy](#), and now [artificial intelligence](#). The NIST AI RMF is nearly 50 pages and the [playbook](#) — a useful but perhaps intimidating in-depth guide for organizations — runs over two-hundred. NIST has distilled its earlier RMFs into digestible resources that make it easier for startups to get started and implement best practices. NIST should do likewise with the AI RMF, continuing to look for additional ways to make the framework more accessible and increase uptake by startups. To encourage adoption, the synthesized resources should be developed in collaboration with startups, small innovators, and intermediaries like incubators and accelerators that understand the needs of the startups who rely on these educational materials and can help ensure the best fit for those needs.

How should government leverage AI to deliver services?

Startups create innovative technologies that can improve government and the provision of public services. Too often, however, startups find it extremely difficult to work with the government. Lengthy contracting processes, challenges navigating government bureaucracy, and a general concern that incumbent companies are favored to succeed, all [hinder startups and their ability to participate in the federal contracting process](#) and secure contracts. In addition to facing these routine challenges of working with government, AI startups are likely to face headwinds as the result of legitimate concerns about mitigating risks of errors.

To solve both of these issues, policymakers should create a pathway for AI startups and a government to cooperatively work through prospective issues while speeding the time to contracting with the government. One option is to create a dedicated startup pilot program outside of the regular contracting process that combines the concept of a [regulatory sandbox](#) with government contracting, where AI startups with demonstrated technologies are able to work with government agencies to create solutions for agency needs. Within the program, a startup would be able to access and build solutions with government data, giving them the chance to build and demonstrate their product while working with the agency to mitigate identifiable risks before the technology is put into regular use.

Overall, balanced regulation is critical to reap the benefits of AI while mitigating its risks. Cultivating a regulatory environment that addresses risks and promotes startup competitiveness, builds on existing legal frameworks, avoids creating barriers to entry, supports preparing products for market, and broadens access to AI resources is instrumental in fostering innovation and maximizing the benefits of AI technology.

[Engine](#) is a non-profit technology policy, research, and advocacy organization that bridges the gap between policymakers and startups. Engine works with government and a community of thousands of high-technology, growth-oriented startups across the nation to support the development of technology entrepreneurship through economic research, policy analysis, and advocacy on local and national issues.