

2027 AI Policy Platform

Introduction

Artificial intelligence (AI) companies are racing to create a history-altering technology that purports to cure diseases and solve society's greatest challenges. That race has already proven to be a difficult task. For all the promise AI holds, it raises serious safety, privacy, and accountability questions as companies race to build the most powerful models first.

The clearest examples of this harm have already awoken people to AI's perils. Whether it is the risks of anthropomorphic chatbots leading individuals to [death](#) or [suicide](#), or far-reaching [safety incidents](#) that occur beyond a company's containment, early fears around AI are already manifesting in our lives.

If we are not careful, we risk making the same mistakes as we did with social media; companies pledged to self-regulate and act responsibly, but those promises fell short. What began as platforms to strengthen our social connections ultimately evolved into environments where harmful, addictive design thrived. The decisions about AI today will define the world we inherit. We cannot afford to be excluded. Those who build and regulate AI systems must be accountable to the generation that will live with these consequences the longest.

AI, while presenting enormous [potential](#), also carries serious risks—especially if left solely to the AI industry's self-regulation. Now is the time for thoughtful policy, meaningful accountability, and proactive governance to ensure that AI's benefits do not come at the expense of our future.

Vision

Like our initial [Policy Platform](#), our 2027 AI Policy Platform puts forth key principles that we will strive to uphold through our support for legislation and policy meant to best protect young people online. The principles we wish to see governing AI are derived from the state and federal bills and laws we have spent years advocating for, as well as where we believe regulation must move toward to ensure adequate protections. As AI continues to evolve, we will update these principles in order to keep up with rapidly developing AI systems and the subsequent need for comprehensive regulation.

To achieve that vision, Design It For Us supports AI policies that adhere to 6 key principles:

1. Prevent Catastrophic Risks
2. Address Societal Risks
3. Design Chatbots to Protect People
4. Allow States to Protect Their Residents
5. Prevent and Protect Against Digital Forgeries
6. Advance AI Innovation While Safeguarding Meaningful Work

1. The largest AI developers must be required to prevent harm related to catastrophic risks. AI technology's immense power presents risks to public safety, public health, cybersecurity, and national security. Without strong, enforceable guardrails, AI companies will continue scaling at a rapid pace. Governments must enact laws and regulations that require minimum standards, disclosure, "kill switches," and transparency to prevent critical safety risks from occurring and so that regulators and the public know when incidents occur.

- A. The largest AI companies should slow down the pace of development.** In the AI industry's quest to achieve superintelligence, companies are scaling at all costs to build the most powerful models, and ignoring the consequences. Policies should allow companies to slow the pace of their innovation in order for society to catch up, with clear oversight to ensure they do not abuse coordination to protect their lead in the industry. The U.S. government should ensure that this pacing is developed alongside other nations' AI companies and achieve international cooperation.
- B. The largest AI companies must meet standards for safety to prevent harms their systems are capable of.** The largest AI companies (Anthropic, Google, Meta, OpenAI, etc.) should be required to meet minimum standards set by regulators that address the models' capabilities, including [Chemical, Biological, Radiological, and Nuclear](#) (CBRN), [cyber risks](#), persuasion at scale, [influence operations](#), psychological damage, and [loss of control](#). Frontier developers must publish and maintain safety frameworks assessing these risks and mitigation measures and [demonstrate](#), before deployment, through mandatory independent auditing, that their systems mitigate these risks. Regulators must take action, including by enforcing penalties, when AI companies do not comply with the law.

- C. Regulators should have the authority to block the largest AI companies from deploying dangerous models and require the development of an emergency shutoff.** Policies should allow regulators to block the deployment of models that fail to meet rigorous safety standards and that present harm to the public. Policies should also require the largest AI companies to develop the technical capabilities to block the deployment of harmful models when safety risks are detected and to ensure models are safe before deployment. This “kill switch” ensures that models capable of catastrophic risks, such as loss of control, can be swiftly shut down to prevent damage, injury, or harm.
- D. The largest AI companies should be subject to mandatory independent auditing.** The companies set up to be the richest in history should not be allowed to grade their own homework. Policies must ensure auditing schemes include accredited auditors that are licensed and assigned to companies, not selected by the companies themselves. They must include evaluators embedded inside the companies with employee-level access, with the ability to monitor dangerous activity continuously, verify that companies are following their safety frameworks, and the authority to report incidents as they happen without company editorializing.
- E. The largest AI companies must adhere to strict transparency and reporting requirements.** Regulation should require the largest AI developers to disclose known risks, safety incidents, and documented harms under a legally binding, uniform standard. The largest AI companies should be required to notify regulator(s) when cyber risks, persuasion at scale, influence operations, psychological damage, loss of control, or other critical safety incidents occur in a timely manner, especially when there is risk of injury or death. Where possible, AI developers and regulators should regularly publish these reports to ensure the public is aware of these incidents.
- F. Whistleblower protections are essential to AI accountability.** When [whistleblowers at Meta spoke](#) out, it ignited worldwide action focused on regulating social media to ensure protections for young people. Workers who speak out must be shielded from corporate retaliation. Young professionals entering the AI workforce must be allowed to champion ethical practices without jeopardizing their careers. Policies should ensure protection for employees of AI companies who become whistleblowers and ensure strict penalties for companies retaliating against employees who come forward.

2. AI systems must be built to anticipate and prevent societal harm. We support the goals behind this technology: automating tasks, reducing costs, and improving accuracy. But as AI becomes more ubiquitous, strong guardrails must be in place to prevent bias and abuse of power. Regulation is needed to protect people’s right to privacy and uphold their civil rights. And while we work to prevent catastrophic outcomes, we cannot [ignore the harms](#) this technology is already causing today.

A. Governments should be prohibited from using AI systems for mass surveillance.

The right to privacy is a fundamental right we outline in our [Declaration of Digital Rights](#). Data brokers and negligent data practices have allowed large amounts of personal data to be bought and sold online, making them accessible to anyone, including the government. AI systems should not exacerbate this problem by allowing the government to easily access troves of people's personal information and create dossiers on individuals to threaten, harass, or target in any way. Policies should prohibit the government from surveilling people and can be bolstered with national privacy data protections. Companies should also be required to install guardrails to prevent bad actors from utilizing AI systems to threaten, harass, or target people in any way.

B. AI companies must be required to ensure accountability in automated decisions.

Companies developing or deploying AI in consequential decisions such as law enforcement, employment, housing, credit, healthcare, and education must ensure those systems do not discriminate, create unjustified disparate impacts, or deny equal access to opportunities based on protected characteristics. Systems should be independently audited before deployment and regularly after, with results published, and people should be able to opt out for a human decision.

3. Chatbots must be designed to protect the people most likely to be harmed by them. AI companies that build chatbots prioritize maximizing engagement of their products. Chatbots like Character.AI, ChatGPT, Claude, Gemini, etc. utilize design features that prioritize profit through dependency, engagement, compulsive use, and targeting young users. These engagement mechanisms can have dangerous outcomes, including driving users toward [conspiracies](#), [emotional dependencies](#), [mass violence](#), and [suicide](#). Policies must prevent companies from optimizing their systems for engagement in ways that rewire our social fabric, especially with heightened protections for young people.

A. AI companies should be restricted from processing young people's chat log data beyond what is necessary to operate the service after a set period.

AI chatbots can imitate human interactions through hyperspecific data that chatbots collect through chat logs, which enable models to profile users. Policies should restrict the processing of that data after a determined period of time to prevent AI providers from profiting off of unhealthy [psychological dependencies](#), [parasocial relationships](#), or exploitative feedback loops.

B. AI companies should not be able to use chat logs for targeted advertising. [Over half of American teenagers](#) say they are using AI regularly. [Half of American teenagers also say](#) they trust or somewhat trust AI companions' information or advice. As usage

becomes more ubiquitous, policies should prevent AI companies from using chat log data for targeted advertising. While companies [like Meta](#) say targeted advertising improves consumer experience, hyperspecific data can [exploit](#) consumers' vulnerabilities for profit without their knowledge.

- C. AI companies should be liable when their products result in injury.** AI chatbots are not people; they are products built by companies, and when they cause harm, developers should be liable for that injury. Policies should enshrine that chatbots are products and provide sufficient recourse for injured parties and governments to hold companies accountable when injury occurs to oneself or others. Requiring AI companies to assume liability will incentivize companies to prevent harm before it occurs.
- D. AI companies should be required to disable harmful design features in their systems and implement measures to identify when a user is expressing suicidal behavior, self-harm, or showing signs of a mental health crisis.** Policies should require companies to disable design features in their systems that result in compulsive usage or emotional dependence. When a user expresses suicidal thoughts, intent to self-harm, or shows signs of an acute mental health crisis, policies should require AI systems to display a clear and prominent crisis message with service information. Policies should require a deployer to notify emergency services or law enforcement when a user faces imminent risk of death or serious risk of injury to themselves or to others.
- E. AI companies should not be able to use minors' (13-17-year-olds') data for training purposes.** Children's data privacy is critical, and in lieu of a comprehensive data privacy regime, policies regulating chatbots should prevent AI chatbot companies from training their models on minors' data. AI models trained on young people's data will be better at keeping young people hooked on conversations to maximize profit.
- F. AI companies should clearly disclose that a user is interacting with an AI chatbot.** AI systems often simulate human interactions. As AI chatbots become more powerful and increasingly provide users with medical, legal, financial, or other advice about sensitive subjects, policies should require AI systems to disclose regularly, persistently, and in context that their chatbots are not humans. Policies should also ensure that AI chatbots do not falsely represent their qualifications.
- G. AI companies should be required to collect the least amount of data as possible to determine a user's age.** Policies that offer heightened protections for young people in order to create age-appropriate experiences on chatbots may require a platform to determine a user's age. Policies should prioritize putting the onus on companies to estimate a user's age based on a platform's knowledge implied on the basis of objective

circumstances. Where age-assurance is required, AI companies should be required to default all users into settings that offer heightened protections, regardless of age, and only use age-assurance for adult users to disable default protections. Where required, age-assurance policies should ensure the least amount of data is collected as possible, companies are restricted to using that data for age-assurance purposes only, and data is deleted after use.

- H. AI companies should not require parental consent to grant minors (13-17-year-olds') access to use the product.** Not all parents and guardians have the time, capacity, or knowledge to know what happens on AI chatbots, or how to set safeguards or otherwise get involved. Some kids and teens don't have parents or guardians who can serve in this capacity, or who have their actual best interests in mind. This gap not only fails to keep all young people safe, but also would have a disproportionate impact on young people from marginalized communities. That is why AI chatbots should be safe by default for all users. AI companies can make parental controls available, but policies should not require companies to enable them.
- I. Researchers should be provided data to study AI chatbots.** Policies should require the production of anonymized data for researchers and mechanisms for researcher access that allow AI chatbots to be studied in order to understand and mitigate harm.

4. Allow States to Protect Their Residents. [From a 10-year moratorium](#), to sweeping AI [regulatory immunity](#), [to executive action weakening state regulation](#), the federal government and Congress have debated preemption proposals that would nullify states' rights to regulate artificial intelligence. As AI systems develop at a rapid pace, the federal government should act to pass strong regulation that utilizes its resources to prevent harm, while states — the laboratories of democracy — should be allowed to enact laws regulating AI and go further than proposals enacted by the federal government. Federal regulation should set a floor for these laws and a strong baseline for state policy to build off of, not a ceiling.

- A. States should be allowed to go further than the federal government.** Policies that govern AI should be enacted at the federal level that empower the federal government's resources to hold AI companies accountable. Policies should, however, allow states to build on a federal framework, allowing them to enact stronger protections for their residents.
- B. States should be allowed to regulate.** States with laws governing AI should be allowed to pass, enact, and enforce those laws without risk of companies or the federal government challenging those laws based on preemption in court.

5. Prevent and protect against digital forgeries. AI technology's ability to generate or alter content, including video, audio, or text, presents tremendous risks to people and institutions. AI deepfakes, in particular AI-generated non-consensual intimate images (NCII), are used to [target young people](#) and [marginalized communities](#). AI companies and bad actors must be held accountable when AI is used to create societal harm.

- A. Victims of AI deepfakes should have expeditious recourse for getting that content removed.** Policies should require AI companies to prevent the generation of AI deepfakes that cause harm, especially NCII, and social media companies should be required to create mechanisms for users to remove this content. Policies should also require technology companies to remove AI-generated NCII expeditiously and build on the [TAKE IT DOWN Act](#).
- B. Victims of AI deepfakes should have recourse against perpetrators of deepfakes.** Policies should allow victims to seek criminal and civil action against individuals who knowingly produce, distribute, or solicit AI-generated NCII.

6. Advance AI Innovation While Safeguarding Meaningful Work. AI presents an unprecedented opportunity to enhance our capabilities and quality of life. The challenge is to leverage these benefits while preserving meaningful human input and agency. As digital natives who will spend our careers interacting with AI, we should be allowed to help shape its integration rather than merely adapt to it. Education must also do more than create AI-literate students—it must prepare them for significant shifts in how skills are valued, and work is structured. As digital natives facing unprecedented technological change, we need education that builds adaptability for a rapidly evolving workforce.

- A. Career development must empower youth to guide AI integration in the workforce.** Rather than focusing only on preserving existing jobs, policies should encourage continuous skill-building and create paths for young people to participate in innovative work across all sectors. Companies should also disclose when AI is used in hiring and offer transparency throughout screening and hiring decisions that utilize AI. This ensures that youth have a role in defining how AI is implemented, rather than simply reacting to changes. Policymakers should prepare for highly disruptive scenarios by investing in retraining and transition assistance to equip the next workforce.
- B. Educational systems must be prepared for the possibility of widespread economic obsolescence.** Large AI companies [explicitly aim](#) to build systems that outperform humans at most economically valuable work. Educational institutions must foster

comprehensive AI literacy through technical skills, critical thinking, and ethical reasoning. These institutions should aim to equip students with skills that were foundational to core instruction before the AI age, such as the ability to form and evaluate their own judgements without assistance from AI tools. Access to these resources and opportunities should be made available to all students regardless of socioeconomic status. Policymakers should incentivize bold experimentation, while respecting young people's privacy, to ensure no young person is left unprepared for a post-AI economy.

- C. Decisions about AI in classrooms should be based on whether students learn more, and schools should expand, adjust, or drop AI tools based on that evidence.** We need rigorous and unbiased evidence-based assessments for teaching methods that combine AI tools with human-led instruction. This will ensure young people acquire both the technical fluency and the distinctively human skills that remain crucial in what will likely be a more automated future.
- D. AI moderation should only replace human moderation when demonstrably safer.** Human content moderation [often exposes workers](#) to traumatic material like child abuse, violence, and graphic content, causing documented psychological harm and creating unsustainable working conditions with high turnover and burnout. AI moderation offers a solution to these human costs, but systems should demonstrate adequate accuracy, fairness, and effectiveness before replacing human oversight completely.

Conclusion

We've learned a valuable lesson from the fight to reform social media that must be at the forefront of all technology policy: Big Tech is not capable of self-regulation. To believe otherwise is to put nationwide safety, health, and well-being at risk. It is the responsibility of policymakers and regulators to consider AI as a second chance to get technology policy right. As youth advocates and as the future of our AI-driven society, we deserve an opportunity to thrive as this new and uncertain technology rapidly advances. Centering our voices and our priorities can and will improve our collective future.