Responsible Advanced AI Act (RAAIA)
Executive Summary

The Problem

Artificial intelligence is rapidly becoming more powerful: AI is already flying airplanes, designing new drugs, setting real estate prices, mass-producing fake videos, and getting top scores on graduate-level tests like the Bar and the Medical Licensing Exam.

Instead of making sure their products are safe, some AI companies are tempted to cut corners so they can launch their products first. Unsafe AI could be misused by terrorists or rogue states to create weapons of mass destruction or automatically hack into thousands of hospitals, banks and reservoirs.

Even without a criminal behind the wheel, AI is likely to misbehave in ways that pose severe safety risks. AI is fundamentally a "black box" technology; not even the people who design it understand how or why it works. As OpenAI CEO Sam Altman put it, "We do not know, and probably aren't even close to knowing, how to align a superintelligence."

The Solution

To protect the public against the growing risks of advanced AI, Congress needs to require that all AI developers adhere to minimum safety standards. The Center for AI Policy (CAIP) has released model legislation with a five-step plan to give the federal government the tools it needs to require safe AI.

1. Launch a 200-person office focused on the risk from advanced, general purpose AI, and use direct hiring authority to staff it with technical experts who can keep up with fast-paced tech.
2. Require the customers buying $30,000 advanced AI chips to fill out a one-page registration form, and use the forms to understand who's working on advanced AI, and where, and why.
3. Issue permits to the most advanced AI systems based on the quality of their safety testing and their cybersecurity, and force unsafe systems to fix their problems before they launch.
4. Ensure that AI labs that break the rules and cause catastrophic harm to the public are held liable for that harm, giving them the right incentives to make safer products.
5. Define a reasonable set of emergency powers for the government so that they can intervene and shut down an AI system that's in the process of going rogue.

How You Can Help

With Congress returning after the Easter recess, we're expecting intense interest from the new House AI Task Force and the Senate AI Caucus. We need support from organizations like yours to let Congress know that this model legislation lays out a good roadmap for a safe future. Please reach out to info@aipolicy.us to start the conversation – if you have concerns, we want to hear about them, and if you are ready to endorse the bill, then we want to get you involved.