The Senate’s hearing on AI regulation was dangerously friendly



The most unusual thing about this week’s Senate hearing on AI was how *affable* it was. Industry reps — primarily OpenAI CEO Sam Altman — merrily agreed on the need to regulate new AI technologies, while politicians seemed happy to hand over responsibility for drafting rules to the companies themselves. As Senator Dick Durbin (D-IL) put it in his opening remarks: “I can’t recall when we’ve had people representing large corporations or private sector entities come before us and plead with us to regulate them.”

This sort of chumminess makes people nervous. A number of experts and industry figures say the hearing suggests we may be headed into an era of industry capture in AI. If tech giants are allowed to write the rules governing this technology, they say, it could have a number of harms, from stifling smaller firms to introducing weak regulations.

Industry capture could harm smaller firms and lead to weak

regulations

Experts at the hearing included IBM’s Christina Montgomery and noted AI critic Gary Marcus, who also raised the specter of regulatory capture. (The peril, said Marcus, is that “we make it appear as if we are doing something, but it’s more like greenwashing and nothing really happens, we just keep out the little players.”) And although no one from Microsoft or Google was present, the unofficial spokesperson for the tech industry was Altman.

Although Altman’s OpenAI is still called a “startup” by some, it’s arguably the most influential AI company in the world. Its launch of image and text generation tools like ChatGPT and deals with Microsoft to remake Bing have sent shockwaves through the entire tech industry. Altman himself is well positioned: able to appeal to both the imaginations of the VC class and hardcore AI boosters with grand promises to build superintelligent AI and, maybe one day, in his own words, “capture the light cone of all future value in the universe.”

At the hearing this week, he was not so grandiose. Altman, too, mentioned the problem of regulatory capture but was less clear about his thoughts on licensing smaller entities. “We don’t wanna slow down smaller startups. We don’t wanna slow down open source efforts,” he said, adding, “We still need them to comply with things.”

Sarah Myers West, managing director of the AI Now institute, tells *The Verge* she was suspicious of the licensing system proposed by many speakers. “I think the harm will be that we end up with some sort of superficial checkbox exercise, where companies say ‘yep, we’re licensed, we know what the harms are and can proceed with business as usual,’ but don’t face any real liability when these systems go wrong,” she said.

“Requiring a license to train models would … further concentrate power in the hands of a few”

Other critics — particularly those running their own AI companies — stressed the potential threat to competition. “Regulation invariably favours incumbents and can stifle innovation,” Emad Mostaque, founder and CEO of Stability AI, told *The Verge*. Clem Delangue, CEO of AI startup Hugging Face, tweeted a similar reaction: “Requiring a license to train models would be like requiring a license to write code. IMO, it would further concentrate power in the hands of a few & drastically slow down progress, fairness & transparency.”

But some experts say some form of licensing could be effective. Margaret Mitchell, who was forced out of Google alongside Timnit Gebru after authoring a research paper on the potential harms of AI language models, describes herself as “a proponent of some amount of self-regulation, paired with top-down regulation.” She told *The Verge* that she could see the appeal of certification but perhaps for individuals rather than companies.

“You could imagine that to train a model (above some thresholds) a developer would need a ‘commercial ML developer license,’” said Mitchell, who is now chief ethics scientist at Hugging Face. “This would be a straightforward way to bring ‘responsible AI’ into a legal structure.”

Mitchell added that good regulation depends on setting standards that firms can’t easily bend to their advantage and that this requires a nuanced understanding of the technology being assessed. She gives the example of facial recognition firm Clearview AI, which sold itself to police forces by claiming its algorithms are “100 percent” accurate. This sounds reassuring, but experts say the company used skewed tests to produce these figures. Mitchell added that she generally does not trust Big Tech to act in the public interest. “Tech companies [have] demonstrated again and again that they do not see respecting people as a part of running a company,” she said.

Even if licensing is introduced, it may not have an immediate effect. At the hearing, industry representatives often drew attention to hypothetical future harms and, in the process, gave scant attention to known problems AI already enables.

For example, researchers like Joy Buolamwini have repeatedly identified problems with bias in facial recognition, which remains inaccurate at identifying Black faces and has produced many cases of wrongful arrest in the US. Despite this, AI-driven surveillance was not mentioned at all during the hearing, while facial recognition and its flaws were only alluded to once in passing.

Industry figures often stress future harms of AI to avoid talking about current problems

AI Now's West says this focus on future harms has become a common rhetorical sleight of hand among AI industry figures. These individuals "position accountability right out into the future," she said, generally by talking about artificial *general* intelligence, or AGI: a hypothetical AI system smarter than humans across a range of tasks. Some experts suggest we're getting closer to creating such systems, but this conclusion is strongly contested.

This rhetorical feint was obvious at the hearing. Discussing government licensing, OpenAI's Altman quietly suggested that any licenses need only apply to *future* systems. "Where I think the licensing scheme comes in is not for what these models are capable of today," he said. "But as we head towards artificial general intelligence … that's where I personally think we need such a scheme."

Experts compared Congress' (and Altman's) proposals unfavorably to the EU's forthcoming AI Act. The current draft of this legislation does not include mechanisms comparable to licensing, but it does classify AI systems based on their level of risk and imposes varying requirements for safeguards

and data protection. More notable, though, is its clear prohibitions of known and current harmful AI uses cases, like predictive policing algorithms and mass surveillance, which have attracted praise from digital rights experts.

As West says, “That’s where the conversation needs to be headed if we’re going for any type of meaningful accountability in this industry.”

https://www.theverge.com/2023/5/19/23728174/ai-regulation-senate-hearings-regulatory-capture-laws