www.pipelinepub.com
Volume 22, Issue 6

The Case for Big Tech's Embrace of AI Regulation

By: Joshua Grossman

As others have stated, we have passed the event horizon for artificial intelligence. That is, the gravitational pull of AI deployments and technology have passed the moment whereby we can avoid its influence and "opt out" of its impact in our world. From the perspective of some AI operators, hardware and software providers, and members of the AI ecosystem, this can seem like a good and proper thing.

However, for many outside the world of AI, especially among the younger people that I speak with, there seems to be at best an ambivalence toward the changes AI is creating in the world and in many cases a deep antipathy. For example, let's consider "Cassie" (Cassandra) as a representative subject with whom I've discussed AI several times. Cassie's view is that AI is a tool of oppression built for our capitalist overlords; she sees a world where technology is used to bring about a complete surveillance state with unbridled authoritarian power and energy crises at best, and a world that looks a lot like Terminator: The Rise of The Machines at worst. Unfortunately, this is not an uncommon view. Even those at the frontier and creation of the best AI models admit that the possibilities range from a utopian "machines of loving grace world" to a dystopian authoritarian infrastructure for panopticon like control by malign actors. Unfortunately, Cassie is not alone in her dark view of AI. Recent polling comes to the same conclusion. Among voters 18-34 the net approval rating for AI is minus 44 per a recent NBC News poll!

If that wasn't reason enough for AI stakeholders to take the initiative in advocating for enlightened regulations, then the lack of operational efficiency in the way business is transacted should alarm them enough to prompt a re-evaluation in strategy and tactics with regard to creating a rules-based order for the uses of AI.

The recent conflict between The US government, open AI and Anthropic regarding the use of AI is a case in point. It is instructive as a lesson in the sort of business practices and sales motion that the industry almost certainly wants to avoid.

Very briefly summarizing the issues: the Pentagon in contracting with Anthropic determined that the guardrails and safeguards which Anthropic insisted upon, specifically, no use for domestic surveillance and no use of automated weapons by AI, were red lines for government acceptance of contract terms. Open AI, on the other hand, was more than willing to adapt its

requirements to suit the needs of the government. (Open AI argues they got the same restrictions but through modified contract language.)

Anthropic's position resulted in its listing as a supply chain risk with the US government which amounts to a black ball against any future contracts. For its part, OpenAI is now the target of boycotts and a plethora of bad press from both industry insiders as well as the public at large. The sort of race to the bottom between private companies and the government is just the sort of thing our industry should try to avoid. A much better approach would be to encourage lawmakers to create a set of standard federal guidelines and laws that regulate AI, as opposed to case-by-case negotiations pushing the ethical and moral frontiers of what sort of AI use cases are allowed. Moreover, Anthropic is challenging the government's designation in court with a lawsuit that claims the company suffered "public castigation" and a violation of its free speech and due process rights. Not a great outcome for the government or the companies.

So what does this mean for practitioners and those that see the benefits of AI in everything from medical breakthroughs to ecological wins and the preservation of endangered species? It means we need to think seriously about common sense regulation that will answer the valid concerns of those worried about the future impacts of AI balanced with the need to continue moving forward in a fraught geopolitical world. There are a number of existing frameworks for AI regulation with the European Union being most advanced thus far. Let's have a look at a couple of different regulatory schema that could be adopted to create real guardrails for safe and responsible AI.

The European Union system uses a four-tiered approach that is similar in some ways to the GDPR data privacy framework that the EU has adopted. That framework, in turn, has similarities to the US privacy enforcement regime as reflected in state privacy laws and federal rules like HIPPA. One benefit of these types of approaches is that they have already been deployed with relative success in both the EU and the US. In addition, there is a large cadre of privacy professionals familiar with this sort of approach who presumably could adopt a similar system for AI.The EU AI system is based on four different tiers of AI usage ranging from essentially safe uses to those that are considered dangerous and almost certainly prohibited. Here's a brief summary of the acts categories:

The Act establishes a four-tier hierarchy of risk, determining the level of regulatory intervention required for any given system:

Unacceptable Risk

- **Definition:** AI practices considered a clear threat to the safety, livelihoods, and rights of people.
- **Regulatory Posture:** Prohibited. These systems are banned from the EU market entirely.
- **Examples:** Social scoring by governments and manipulative AI that circumvents a person's free will.

High Risk

- **Definition:** AI systems that have a significant negative impact on people's safety or fundamental rights, specifically those used in critical infrastructures or essential private/public services.
- **Regulatory Posture:** Strictly Regulated. Must comply with mandatory requirements and undergo a conformity assessment before market entry.
- **Examples:** AI used in recruitment (CV sorting), credit scoring, and biometric identification.

Limited Risk

- **Definition:** AI systems with a specific risk of manipulation or deceit.
- **Regulatory Posture:** Transparency-only. Users must be made aware they are interacting with an AI system.
- **Examples:** Chatbots, deepfakes, and AI-generated text or images.

Minimal Risk

- **Definition:** AI applications that present negligible risk to citizens' right or safety.
- **Regulatory Posture:** No obligations. These systems may be used freely, though volunteary codes of conduct are encouraged.
- **Examples:** AI-enabled video games and spam filters.

Moreover, the EU AI Act prohibits specific AI practices deemed to be in contravention of Union values and fundamental rights. The rationale for these bans is the protection of human dignity and the prevention of mass surveillance and discriminatory social engineering. However, particularly in the US AI ecosystem, there are many influential voices opposed to federal AI regulation.

The counter argument to a regulation positive stance by industry players is often stated as somehow crippling the US as it strives to compete for AI superiority with the other main AI player, China. This argument overlooks the fact that with national regulatory legislation the government could have carve-outs for things like AI for defense-related use cases and thereby not hamstring development. As for AI uses in domestic surveillance, that use case is so antithetical to the constitution and our guaranteed rights that sacrificing the ability to surveil the domestic population of the US is a price worth paying to ensure a society in which we want to live.

US AI stakeholders would do well to implement a similar scheme backed through federal legislation, or they could take the approach that evolved over time with Privacy legislation where state privacy legislation frameworks like those in California became a blueprint for other states. Due to the scope of AI, a national system would be a much better alternative than a patchwork of state-based regulations, especially with regard to its uses in national defense. This is something that could be done through federal legislation along the lines of the legislative approaches that the Crypto industry is currently pursuing. That is, push for national laws and guidelines as opposed to state level enforcement. In sum, creating a national AI governance schema.

Ultimately, the benefits to a regulation forward approach by the AI business community will result in a better outcome for the residents of the US, a better outcome for the companies involved, and most importantly could be a huge step in establishing guardrails for the US AI industry to protect future generations. Not embracing regulation will result in a backlash due to destructive uses of AI like those prohibited by the EU schema and at worse could give rise to catastrophic consequences for humanity.

With a little luck, by adopting a “regulation forward” approach in our industry, perhaps we can act as the masters of our own fate. I believe regulation is coming whether the industry embraces it or not, better to be guiding that legislation than not. And importantly, maybe we change “Cassie’s” opinion and do some good while doing well.