Insight

Primer: How to Understand and Approach AI Regulation

WILL RINEHART | JANUARY 10, 2019

Executive Summary

Countries and firms across the globe are racing to capitalize on artificial intelligence (AI). To ensure that the United States can capture the gains from this new technology, policymakers should recognize several things:

- Premature regulation is likely to be deleterious to innovation and progress in AI;
- Large firms have taken the lead on AI implementation and shouldn't be punished for doing so; and finally
- The United States shouldn't pursue a singular strategy of dominance, but rather a multiplicity of strategies that rely on the ingenuity of industries and individuals, working in concert, rather than conflict, with government agencies.

Calls for Regulation

Elon Musk, Bill Gates, Mark Cuban, and the late Stephen Hawking have been among the most vocal luminaries calling for the regulation of artificial intelligence (AI), but they are hardly alone. Countless papers, conferences, and talks dedicated to algorithms and artificial intelligence call for the same. Without detailing the harms, or explaining how the market has failed, many tend to focus on proposals to tax, regulate, and limit robots and artificial intelligence.

Embedded in these calls for new government power are countless uncertainties, as the track record of technology forecasts is far from stellar. One of the largest retrospective reviews of technology forecasts found that predictions beyond a decade were hardly better than a coin flip. In an analysis that focused specifically on AI predictions, the authors warned of "the general overconfidence of experts, the superiority of models over expert judgement, and the need for greater uncertainty in all types of predictions." Predictions that general AI is just around the corner have failed countless times across several decades.

This uncertainty indicates a fundamental reality about AI: It is a developing collection of technologies with a tremendous variety of applications. As a result, policymakers should embrace regulatory restraint, although there are opportunities for policy to strengthen AI development. The goal for policymakers should not be a singular AI policy or strategy, but a regulatory and policy approach that is sensitive to developments within society.

Terms and Origins

To understand artificial intelligence, it is helpful first to define terms, especially "narrow" AI and "general" AI. Narrow AI references models built using real-world data to achieve narrow, specific objectives such as translating languages, predicting the weather, spotting tumors in chest scans and mammograms, and helping people identify caloric information just from pictures of food.

Narrow AI can be contrasted with general AI, which refers to decision-making systems able to cope with any generalized task like a human. General AI is what most associate with the term AI, as it has found its way into popular culture through Arnold Schwarzenegger's early 1990's movies and, more recently, Samantha from the movie *Her*. While some fret over the risks posed by super-intelligent agents with unclear objectives, task-specific AI holds immediate promise while general AI is still far from full realization.

The diversity of what one thinks of as AI extends beyond these categories, too. Machine learning is another commonly referenced term, which denotes a process whereby a machine analyzes data and learns without supervision. Yet the barriers between AI and machine learning and more standard computer programming are blurry. In practice, there often isn't much difference between narrow AI and complex computer programming like machine learning. But there is an upside to this diversity, as Stanford's "One Hundred Year Study on Artificial Intelligence" noted: "[T]he lack of a precise, universally accepted definition of AI probably has helped the field to grow, blossom, and advance at an ever-accelerating pace."

This diversity stems from the technology's democratic origins. As the Obama White House noted in its "Preparing For the Future of Artificial Intelligence" report,

> The current wave of progress and enthusiasm for AI began around 2010, driven by three factors that built upon each other: the availability of big data from sources including e-commerce, businesses, social media, science, and government; which provided raw material for dramatically improved machine learning approaches and algorithms; which in turn relied on the capabilities of more powerful computers.

In January of 2010, the machine learning library, scikit-learn, was released to the public, democratizing the tools of AI and sparking the current rush. This program finds it genesis in Google's Summer of Code programs, and many different companies and entrepreneurs have applied these tools in manifold ways. As Representative Will Hurd said in June, "the United States boasts a creative, risk-taking culture that is inextricably linked to its free enterprise system."

Google, Facebook, Microsoft, and other large tech companies have played a large part in the development of AI. While it has been popular of late to criticize the largest tech companies, policymakers should be comfortable with large firms such as Google, Facebook, and Microsoft taking the lead on AI implementation. Even though these companies have been lambasted for their size, they shouldn't be penalized for adopting advanced technologies.

The democratic nature of AI development over the last decade means that there are a variety of experiments in the ecosystem, and shifting to AI-embedded processes will not be frictionless for firms or social institutions. As AAF has noted before, firms face significant practical hurdles in implementing AI-driven systems, as they aren't cheap and most automation schemes fail to achieve any positive results. The same kind of implementation problems exist in government institutions as well. In the most comprehensive study of its kind, George Mason University law professor Megan Stevenson tracked Kentucky's state-wide implementation of an algorithm meant to automate bail decisions by judges. While there were significant changes in bail-setting practices, over time these changes eroded as judges returned to their previous habits.

Regulation and AI

The shifting landscape and unclear implications of AI mean that policymakers should adopt three outlooks regarding narrow AI regulation.

First, AI is a general purpose technology, like electricity, the automobile, the steam engine, and the railroad, that will have a variety of regulatory impacts. AI isn't going to converge industry regulation but make it more variegated. Thus, calls to impose a singular regulatory framework on AI are misplaced. Some industries might need clarity, others might need a shift in liability rules, and yet others might need additional consumer safeguards. Still, we are a long way from those deep societal impacts. In the near term then, policymakers should be on alert to the potential barriers that could hobble growth in AI application, which might necessitate the liberalization of rules.

Second, premature action is likely to be deleterious to AI innovation and progress, as privacy regulation in Europe has shown. A rush to legislate AI applications, and thus constrain and narrow them, would signal to investors and innovators that their time, money, and talents should be put elsewhere. Such a shift would be a real loss, as the opportunities for AI applications are enormous. The United Kingdom's National Grid has turned to AI to reduce service outages. Facebook and MIT are using AI to give addresses to people throughout the world without them. And even *The New York Times* is getting into the game, by installing a recommendation feed for its users through AI.

Regulatory restraint does not mean consumers are exposed to harms. Consumers can be protected if policymakers choose the route of soft law. As Ryan Hagemann, Jennifer Huddleston Skees, and Adam Thierer explained, "soft law represents a set of informal norms, multi-stakeholder arrangements, and non-binding guidance standards that provide an adaptable alternative to more traditional regulations or legislation." These approaches have been successfully applied to autonomous vehicles, Internet of Things, advanced medical technologies, FinTech, and electric scooters. Relying on these strategies would be a smart strategy for AI regulation.

As a final matter, policymakers should temper concerns about the ethical implications of AI. The terminator scenario of AI might be well known, but it is not indicative of the current hurdles that AI researchers face. Instead, practitioners tend to be concerned with more concrete obstacles, such as avoiding side effects that reward hacking, ensuring that there is scalable supervision, and stopping undesirable behavior during the learning process.

Moreover, countless organizations are dedicated to these ethical problems, such as Data&Society, the Ethics and Governance of AI Initiative, and the AI Now Institute, just to name a few. Companies are beginning to hire researchers focused in AI ethics and are creating internal ethics boards for AI. Moreover, educational facilities

are beginning to implement ethics within curriculum. As Computer Science Professor Yevgeniy Vorobeychik explained in a filing, "the vast majority of AI researchers already set public good, broadly construed, as their aim." Policymakers should be optimistic about society's ability to consider and act on AI's ethical implications with both speed and nuance.

Policy Opportunities

Because of the United States' federal and divided government, implementing a comprehensive national strategy shouldn't be the goal. Rather, piecemeal changes through legislation, agency interpretations and workshops, local government initiatives, appropriations bills, and countless other actions will constitute our national strategy.

The Trump Administration, to its credit, has been prioritizing artificial intelligence within its own operations, and it should continue to do so. In the annual guidance to heads of executive departments and agencies, both the Office of Management and Budget and the Office of Science and Technology Policy directed agencies to focus on emerging autonomous technologies. The White House has also created a Select Committee on Artificial Intelligence, made up of the most senior R&D officials from across the federal government and has worked with European official to secure a digital compact on AI research.

While laudable, the administration might also consider something more direct for agencies. With Executive Order 13771, the Trump Administration committed to "a requirement that for every new federal regulation, two existing regulations need to be eliminated." The practical effect within agencies has been the creation of working groups to implement the order and improve regulations. While there might be some overlap with the Emerging Citizen Technology Office, a carefully crafted executive order could have a similar effect of reorienting agencies to be on the lookout for AI opportunities.

An enduring question is just how much the federal government should spend on research and development (R&D) for AI. As the Obama Administration argued, "there is an underinvestment in basic research…in part because it is difficult for a private firm to get a return from its investment in such research in a reasonable time frame." But there is a difference between underinvestment in basic research and underinvestment in AI-related technologies. The optimal amount of federal funding for research could be at the right level but not directed toward the right mix of technologies. Some research suggests that the optimal level of R&D spending by all actors in the economy is between 2.3 and 2.6 percent of gross domestic product, and the total amount of R&D spending today is near all-time highs. Policymakers should consider what place AI development has in the government's overall R&D programs.

Where government agencies will play a role in coming years will be in democratizing that ingenuity and helping lagging sectors update their production techniques. For example, the Federal Communication Commission convened a forum of experts to discuss how AI will affect the communications marketplace. The Federal Trade Commission has done something similar with its hearings on "Competition and Consumer Protection in the 21st Century," focusing on the implications of AI for competition law. Agencies should follow such efforts with a document or a set of resources for practitioners.

Congressional oversight will be key and, indeed, members have introduced various bills on the topic in the 116th Congress. At this point, regulation isn't on the table, but most of the bills would create independent working groups to study the potential of AI. For example, Representative John Delaney and Senator Maria Cantwell introduced the "Fundamentally Understanding The Usability and Realistic Evolution of Artificial Intelligence

Act of 2017," or the FUTURE of AI Act in both the House and the Senate. If passed, the bill would set up an advisory committee to study its impact and "to promote a climate of investment and innovation." The committee would only include 19 voting members, however, and would have to produce a report in a year and a half, which would be a tall order. A better approach, as noted above, would be for individual agencies to explore how AI will impact their stakeholders. While it is likely to create overlap, the resulting reports will be more industry-specific.

Research institutions will also be important players in ensuring change. Nearly 100 colleges and universities in the United States offer programs or are conducting research in artificial intelligence and data science. UC Berkeley plans to create a new Division of Data Science, one of its biggest reorganizations in decades, and this fall it began offering a major in data science since the fastest-growing class is its introduction to data science. And with $1 billion in funding, MIT will create a new college that combines AI, machine learning, and data science with other academic disciplines.

What the United States needs isn't a singular strategy of dominance, but rather a multiplicity of strategies that rely on the ingenuity of industries and individuals, working in concert, rather than conflict, with government agencies. To borrow a phrase, policymakers should be alert to the United States' AI strategy *on the ground*, not merely on the books.

For a list of specific actions that can be taken, see this AI agenda.