



Comments for the Record

Which Policies Will Foster the Growth of Artificial Intelligence?

WILL RINEHART | JULY 22, 2016

The techniques of artificial intelligence are quickly being adopted in medicine, life science, and for analysis of big data. Continuing progress will require the adoption of optimal legal and regulatory frameworks. Federal and local governments can foster these technologies by promoting policies that allow individuals to experiment. This can be achieved by considering the costs in applying liability rules, intervening only when there are demonstrable benefits, providing room to experiment, and allowing for trade in technology and ideas.

Defining AI

The term AI often conjures up an early 1990s image of Arnold Schwarzenegger and more recently, Samantha from *Her*. AI of this kind, often called strong AI, is far from our current technological capabilities and may never be achieved. While some fret over the risks from super intelligent agents with unclear objectives, task specific AI holds immediate promise. Narrow AI is a term for a collection of economic and computer models built using real world data to achieve specific objectives. These objectives might include [translating languages](#), better [predicting the weather](#), spotting tumors in [chest scans and mammograms](#), and helping [people identify caloric](#) information just from pictures of food.

Understanding AI Risks

In the course of searching for solutions, AI will encounter negative events. Risk management pioneer Aaron Wildavsky rightly defines risk this way, situating it as a byproduct of the search for welfare enhancing economic activity. Indeed, there is an important and deep relationship between wealth and risk, which should be little surprise considering that risk and return are correlated and the bedrock of finance theory.

Researchers in AI have identified three [concrete problems](#) that could contribute risk to AI:

1. The objective was incorrectly specified, leading to negative side effects or cheating by the AI;
2. The designer might have an objective in mind, but the costs in evaluating how well the AI is performing on these objectives make harmful behavior possible, or
3. The objective is clearly specified but some unwanted behavior occurs because the data is poorly curated or the model isn't sufficiently expressive of the environment it is interacting with.

Because of the variety of domains where AI is being applied, how these problems manifest will vary. AI applications within medicine will see different kinds of issues arise as compared to weather prediction models. How legal systems should be designed will and should vary considerably, depending on the already existing institutional environment, which includes the legal and moral 'rules of the game' that guide individuals' behavior, and the industry specific institutional arrangements, which generally define how companies are organized.

Because of the varying contexts, there is no workable one size fits all regulatory framework. All of this should give pause to any sort of broad regulatory effort to limit the application of AI, much as the [European Union is currently contemplating](#). Optimal levels of regulation must be discovered, so much of the work is likely to come from court cases.

Creating Rules for Liability

Since AI innovation is in its early stages, it is too soon to determine which liability rules, rules of evidence, and damage rules for various jurisdictions should apply.

Autonomous vehicles serve as a good example of the complexity of this question. While autonomous cars won't radically replace all cars in the next five years, they are likely to come into increasing contact with human drivers and cause accidents. Four [basic kinds of rules apply](#) for fault in car accidents, which states have adopted in varying degrees. Most jurisdictions have adopted comparative fault, so damages are apportioned based on their proportionate shares. Over time, courts will likely shift the allocation of burden to human drivers if driverless cars prove safe. However, one can also imagine cases where courts will assign some percentage of fault to the autonomous auto-manufacturer based on an adaptation of product liability law. How these cases play out will depend heavily on the degree of AI implemented and the control that manufacturers allow for drivers.

Product liability is a complex area of law and should be allowed to adapt to the challenges of AI. However, there should be more focus on the costs of the system. The [available empirical evidence suggests](#) that there isn't a measurable effect on the frequency of product accidents due to varying product liability regimes. In other words, the purported safety benefits of product liability might not exist when real world costs are considered. Moreover, given the current legal system, a [significant portion of the compensation](#) that is meant to pay damages to victims goes instead to transactions costs in the form of legal fees. In total, there are reasons to believe that enterprise is less likely to engage in those kinds of activities, which could be a deterrent to AI development. Thus, researchers should work towards better understanding how economically efficient these rules are in practice and jurisdictions should be careful on how they apply old rules onto new technologies, especially given the costs.

Openness to experimentation

Nevada was the first state to allow for the operation of autonomous vehicles in 2011 and has since been joined by five others, including California, Florida, Michigan, North Dakota, and Tennessee, as well as Washington D.C. While it is not guaranteed, these states will likely lead the way in developing autonomous vehicles since they are creating zones where the technology can start pushing down the risks. The long term benefits will come in the shape of investment and jobs. For policymakers, removing barriers to experimentation should be the utmost priority. This can broadly be achieved by adopting a mindset of permissionless innovation.

As we currently see in the computer security field, tools will be devised that search for these problems and correct them, similar to how algorithms can search for bad code and cybersecurity threats. Creating zones of experimentation where the three types of risk can be worked out will lead to a greater level of safety. The benefits may come in the form of laws passed in those five states and the District of Columbia, or perhaps via limited liability. Experimental spaces will ensure incentives are aligned to research and develop AI.

Given how promising these technologies are, prescriptive federal regulation is hardly justifiable at this time. In applying the old regulatory regime to these new spaces, regulators should be mindful of the [three-part test](#):

1. Prove the existence of market abuse or failure by documenting actual consumer harm;
2. Explain how current law or rules are inadequate, and show that no alternatives exist including market correctives, deregulatory efforts, or public/private partnerships to solve the market failure; and
3. Demonstrate how the benefits of regulation will outweigh the potential countervailing benefits, implementation costs, and other associated regulatory burdens.

Openness to trade

While the United States is at the forefront of AI development, there is no guarantee that advances will always be made here. Two basic principles flow from this. First, the US should maintain an openness to trade with other countries and ensure that there are not any trade related encumbrances, especially in data transfer. Second, we should be a leader in this space by encouraging our closest trading partners, including those in the EU, to abandon myopic views of AI and allow for more experimentation with the available tools. [Research and development has globalized](#) and only if we embrace that reality will the U.S. be able to reap the rewards.

Digital literacy

Digital literacy needs to be emphasized. As compared to media literacy and computer literacy, digital literacy focuses on imparting knowledge of complex network systems and big data, as well as critical thinking skills to understand how these systems relate to stand alone devices. For states and local government, this doesn't translate necessarily into a need for all students to be able to code, but to at least appreciate how technology works. While they are sure to involve educational institutions, strategies for digital literacy will likely better serve everyone if they originate from local communities and users of the technologies instead of the federal government via strict mandates.

Conclusion

Much like the beneficial uses for AI, the optimal legal and regulatory institutions for AI will have to be discovered. While many reflexively coil when hearing the term AI, the narrow version of AI might offer some real benefits. Federal and local governments can foster these technologies by being supportive but taking a hands off approach in helping to mitigate that risk and allowing the legal system to do its job. Progress in this space will depend on how comfortable we are with new machine-human partnerships. To accomplish this, we do not need more laws and institutions, but more trust in the ones that already exist.

The above comments were submitted to the White House's Office of Science And Technology Policy for their ongoing "Preparing for the Future of Artificial Intelligence" project.