IDENTIFYING THE PROMISES AND CHALLENGES OF ARTIFICIAL INTELLIGENCE:

Shaping the future of California public policy to prepare and protect the state's population, including its workforce and consumers, while simultaneously driving innovation and productivity as a leader in the global marketplace

March 6, 2018
1:30 p.m.
State Capitol, Room 126

BACKGROUND PAPER

_____

I.   Background

Though first coined by a Dartmouth professor, John McCarthy, in the 1950s, there still does not appear to be any singular, consistent definition of "artificial intelligence" (AI) in use today, over 60 years later. McCarthy described the process as "that of making a machine behave in ways that would be called intelligent if a human were so behaving."[1] Today, there are many proposed definitions of AI, but most appear to still be roughly aligned around that same concept of creating computer programs or machines capable of "intelligent" behavior if exhibited by humans.[2] A recent public hearing in the State Capitol provided a useful working definition of AI along these lines: "AI is the theory and development of computer systems able to perform tasks that normally require human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages."[3]

Definitional issues notwithstanding[4], examples of what are perceived to be AI applications are ubiquitous. As those examples multiply, the conversation over the opportunities that AI

---

[1] Kaplan, Artificial Intelligence: What Everyone Needs to Know (2016) p. 1, citing J. McCarthy, M.L. Minsky, N. Rochester, and C. E. Shannon, "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence," 1955, http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth/html.

[2] *Id*. Defining, let alone measuring, human "intelligence" can be very difficult given how abstract and subjective it is.

[3] On January 25, 2018, the Little Hoover Commission (LHC) held a public hearing entitled "Artificial Intelligence: Applications and Implications." The first of a series of hearings and roundtables that the commission announced it would be holding on the topic, the LHC is intending to ultimately produce a report and policy recommendations about what the State of California can do in and around AI. The intent of that hearing, much like this joint informational hearing, was to examine to promises and challenges of AI.

[4] AI can also refer to "artificial general intelligence," "deep learning" or "neural networks." Artificial general intelligence refers to a not-yet existent software that would display a humanlike ability to adapt to different environments and tasks, and transfer knowledge between them. Deep learning refers to a machine learning technique in which data is filtered through self-adjusting networks of math loosely inspired by neurons in the brain, known as artificial neural networks. *See* Simonite, *The Wired Guide to Artificial Intelligence*, Wired (Feb. 1, 2018) <https://www.wired.com/story/guide-artificial-intelligence/> [as of Mar. 4, 2018].

presents, as well as its associated challenges, is becoming more pronounced and undoubtedly will present itself in policy form for the Legislature to consider in years to come. As described in a recent Wired article entitled *The Wired Guide to Artificial Intelligence*:[5]

> For most of us, the most obvious results of the improved powers of AI are neat new gadgets and experiences such as smart speakers, or being able to unlock your iPhone with your face. But AI is also poised to reinvent other areas of life. One is health care. Hospitals in India are testing software that checks images of a person's retina for signs of diabetic retinopathy, a condition frequently diagnosed too late to prevent vision loss. Machine learning is vital to projects in autonomous driving, where it allows a vehicle to make sense of its surroundings. There's evidence that AI can make us happier and healthier. But there's also reason for caution. Incidents in which algorithms picked up or amplified societal biases around race or gender show that an AI-enhanced future won't automatically be a better one.

Indeed, in this way, AI could very well exacerbate problems, as much as it could solve them. On the one hand, there is an incredible ability for AI to create a global paradigm shift that may propel society into an automation age and propose unique solutions to some of the world's greatest problems. On the other hand, AI comes with its own risks, including bias or loss of privacy, among other things. Thus far, academics and technology companies have led the discussion on the future of AI. Given the potential challenges of AI to society, as well as the benefits, government arguably needs to become more actively involved in understanding AI, and the future of AI, both in terms of technology and public policy.

On Tuesday, March 6, 2018, the Assembly Committees on Privacy & Consumer Protection and Emerging Technologies & Innovation will begin the process of bringing the Legislature more actively to the table, alongside academics and technology professionals, and public policy experts, to explore these questions: what exactly is encompassed by the term "AI;" what are the opportunities and challenges it provides to California's economy, workforce, consumers, and general population alike; and what might the appropriate balance between regulations and uninhibited innovation look like in preparing for those challenges and opportunities?

II. AI Applications

As indicated in the Wired article in Section I, above, AI is frequently associated with technologies linked to our smartphones, or new gadgets like virtual assistants or smart speakers like Alexa or Google Home. In cinema, it is often portrayed as "robot apocalypse." For the Legislature, contemplating AI applications of the "future" frequently includes autonomous vehicles and concerns displacement of workers with the automation of jobs. Beyond such examples, however, it is not as obvious what AI looks like five years down the line, let alone 10. In the commercial space, it is easy to foresee many applications of AI from handheld phones to

---

[5] *Id*.

speakers to food delivery options to transportation options potentially in a matter of years but not as obvious how those commercial applications might have broader impacts. Similarly, it is not as readily foreseeable as to what the AI applications are in the social and governmental spaces and how exactly AI may play a role in shifting and changing our society and our economy in the next decade.

For example, for many people, AI is not immediately associated with social justice. However, at the University of Southern California, the Center for Artificial Intelligence in Society (CAIS) has brought researchers together from around the world to focus on how computer science can be used to solve social problems. Indeed, from the CAIS' perspective, AI can be used to improve society and fight social injustice. Their current projects include: AI for Cybersecurity; HIV prevention among homeless youth; Wildlife Conservation with drones; AI for Wildlife Conservation in Africa; Predictive modeling of tobacco use and prevention among abused children; Predictive models of vulnerability and housing prioritization for youth and families; Gang violence prevention using game theory; Social network-based substance abuse prevention for homeless youth; Predictive modeling for early identification of suicidal thinking among active duty service members; Network-based suicide prevention for college students; AI for public safety and security using game theory; and others.

At the same time, while AI may present unique solutions to social problems or even governmental ones, as indicated above, it may very well exacerbate others if not done with adequate safeguards in place. For example, governmental entities may turn to AI for useful applications in everything from enhancing delivery of services to better addressing public safety concerns. Consider how some states' courts have sought to apply AI to conduct risk assessments (*i.e.* assessments of how likely a defendant is to commit future crimes) through the use of seemingly neutral algorithms. Already, concerns have been raised about how these algorithms may in fact reinforce or aggravate biases.[6]

At this joint informational hearing, the Committees will hear about both traditional and less traditional applications of AI that range from the commercial, social, and governmental spaces, as it prepares to grapple with the question of how it should view the opportunities and challenges of AI through the public policy lens.

III. Catching Policy up to Technology

As is often the case when framing public policy in relation to technology issues, the path forward is not a simple or straightforward one, and technology often outpaces the policy. Writing sound policy requires a thorough understanding of the universe being regulated, and emerging technologies rarely lend themselves to quick studies. As such, the goal of this joint

---

[6] Anwin, Larson, Mattu, and Kirchner, *Machine Bias*, ProPublica (May 23, 2016) <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing> [as of Mar. 4, 2018].

informational hearing is to start a series of discussions around the types of questions that will have to be answered in order to enable the State to formulate appropriate public policy around AI. The Committees will explore, among other things, the following questions:

- What technologies we are dealing with?
- What are the varying applications? What are they today and what might they look like five years from now? Ten years from now?
- What are their benefits?
- What are their challenges, unintended consequences, and potential pitfalls? Will those impacts be felt disproportionately by some job sectors, races, genders, classes, urban/rural areas?
- What is the appropriate role of government? What is the role of a government like that in California which is the 6th largest global economy? How does the State avoid hindering innovation, and at the same time adequately protect constituents against issues such as unemployment, bias, and loss of privacy?
- Are there any areas where the common law is sufficiently elastic to address legal issues that will be raised by AI technologies and their applications?
- How might the State begin to prepare the current workforce and employers? The future workforce and educational institutions? Businesses and consumers? Medical providers, patients, and caregivers? The general public?
- Is it enough to reframe educational curriculum or does society need to reframe how it approaches education and training in a person's lifetime more generally?
- Are there partnerships that the State should be looking to start now to address issues which may arise in five to 10 years' time?

The Committees recognize that the right solutions may not exist today, but are looking to identify the types of questions that the Legislature needs to consider in order to adequately prepare for the changes precipitated by AI applications. It may very well be that future hearings are necessary, perhaps in conjunction with other policy committees, to further explore topics raised at this initial informational hearing.