Center for AI and
Digital Policy

Artificial Intelligence and Democratic Values:
The Role of Data Protection

Marc Rotenberg, President and Founder
Center for AI and Digital Policy.

Global Privacy Assembly
19 October 2021
Mexico City, Mexico

Dear Friends, It is an honor to be with you today. Thank you, Commissioner, and the National Institute of Mexico for organizing this conference. Thank you to Elizabeth Denham and the Global Privacy Assembly for your ongoing work. The BBC recently wrote that Elizabeth Denham leaves “a substantial legacy” as UK Commissioner. We all owe Liz a debt of gratitude for her good work.

* * *

In the book *Computer Power and Human Reason*, the MIT computer scientist Joseph Weizenbaum wrote that we should never allow computers to make important decisions because computers will always lack human qualities, such as compassion and wisdom.

Yet we find ourselves today surrounded by machines that make life-altering decisions about all of us. They decide if we get an interview for a job, if we may cross a national border, if we receive public benefits, if our children are granted admission to university, even if we are to go to jail or to maintain our freedom.

The outcomes of these decisions are not simply the concerns of the experts or the academics. There are the day-to-day lived experiences of our friends, our neighbors, and our family members in this data driven, digital economy.

Many are resigned to this world. Others are not. The students who rose up this past year against the assignment of their grades by an algorithm that favored the privileged over the disadvantaged, *even when the achievements of the* *students in the less prosperous schools were greater,* recognized that automated decisions speak directly to basic concepts of fairness, justice, and transparency. And so they said, “Fuck the Algorithm,” “put an end to this.” And we saw not only a successful campaign against automated decision-making, we witnessed a campaign against unjust outcomes.

This is precisely the warning Safiya Noble presents in *Algorithms of Oppression* when she writes of, “the power of algorithms in the age of neoliberalism and the ways those digital decisions reinforce oppressive social relationships and enact new modes of racial profiling."

Still, Weizenbaum would have smiled knowing that more than forty years after he wrote, students could see beyond the scientism that often characterizes opaque and unaccountable decisions, organize themselves, and obtain a political goal.

But Weizenbaum was hardly alone in the early days of computing to recognize the risks of Artificial Intelligence. In fact, by the time he published *Computer Power and Human Reason*, many democratic governments had established comprehensive legal frameworks to regulate the processing of personal data.

Here in the United States, we enacted the Privacy Act of 1974, a foundational law that set out the essential framework of modern data protection.

The central goals of “fairness, accountability, and transparency” were well understood and established in law, as was the need for effective remedies and ongoing oversight. The framers of the Privacy Act also made clear that statistical data should be widely available, as long as it does not pose a risk to the rights of an identifiable individual.

And here I must make two further points about modern privacy law.

First, data protection is, at its core, about fairness, it is about justice, it is about how we treat each other regardless of what a computer decides. The “notice and choice” paradigm is an artificial construct, an effort to turn a fundamental right into a market commodity, and to provide liability immunity to those who obtain personal data. “Notice and choice” purposefully ignores the fact that the vast majority of data collection occurs without participation, consensual or otherwise, of the individual.

My second point is that the real “paradox of privacy” is that privacy requires transparency. That was understood by the drafters of the Privacy Act almost fifty years ago. That is the reason that individuals have the right to know all of the information about them, and the logic and factors, that contribute to decisions about them.

And modern privacy law has always recognized transparency as the most effective technique for oversight and accountability.

Weizenbaum wrote his famous book about the limits of Artificial Intelligence in part because he was surprised that people who interacted with his simple computer program ELIZA placed such confidence in outcomes that were hidden in a few lines of compute code. ELIZA merely restated the comments of the user as questions, appearing to produce the actual concern of a human therapist. It was almost too easy to pass that Turing Test.

We need to see the code to understand the outcome. We need to interrogate the decisions that an AI system proposes. As Karl Popper, the author of the *Logic of Scientific Discovery*, wrote “In so far as a scientific statement speaks about reality, it must be falsifiable: and in so far as it is not falsifiable, it does not speak about reality.”

* * *

Weizenbaum was one of the first people to encourage others to think critically about artificial intelligence. He came of age in a world where dramatic advances in science quickly turned into the tools of destruction. He understood that the wrong values could send societies off in terrible directions.

When we launched the new Center for AI and Digital Policy, we took as our mission to understand the relationship between Artificial Intelligence and Democratic Values. Much as privacy advocates a generation ago explored the relationship between “privacy and human rights” or “cryptography and liberty ,” we aimed to bridge two large conceptual categories.

The focus on democratic values is purposeful. We can easily see two AI futures – one that favors pluralistic societies and respects human dignity, another that aggregates personal data and centralizes control.

In the first, AI offers insight into challenges such as the development of effective vaccines, the response to climate change, and the reduction of bias in criminal sentencing. AI helps fuel innovation and progress even as it subject to the regulations that builds trust and public confidence. AI remains accountable. This is the human-centric outcome.

But there is a different future, driven by the machine’s desire for personal data and the desire of governments to maintain power and suppress dissent. This future conceals decisions in layers of opacity that even those in charge may not fully understand. In such a world proponents will hold out impressive results – such as a reduction in crime -- to justify further deployment. But they will also reject independent evaluations, claiming perhaps that it is not even possible to replicate results. They will conflate correlation with causation. That is the AI-centric outcome

Of course, it is fine to prefer democratic governance over authoritarian rule but how precisely do we judge whether AI systems favor one outcome or the other?

To answer this question, we at the Center for AI and Digital Policy developed 12 metrics to assess AI. And in 2020 we published the first comprehensive review of the AI policies and practices in 30 countries. The report ***Artificial Intelligence and Democratic Values*** provides a basis to compare countries and also to follow particular countries over time to assess whether they are making progress toward democratic values.

We knew that our work would not be easy. It is simple to count patents issued, papers published, and national investments in AI research. But to assess subjective values is a more difficult task.

In constructing the Center's index, we looked closely at the history and structure of data protection law. Many of our metrics will be familiar to you. We asked, for example, whether the goals of "fairness, accuracy, and transparency" were included in a country's national AI strategy.

We asked whether the country had by law established a right to algorithmic transparency, as would be found in Article 22 of the GDPR or Article 9 of the Modernized Privacy Convention of the Council of Europe.

We also asked whether the country had an independent agency for oversight of AI practices. Data protection agencies, and also human rights commission and human rights institutes, can fill this responsibility.

And we looked specifically at the resolutions of the Global Privacy Assembly to evaluate a country's commitment to human-centric AI. Those countries that had sponsored the GPA resolutions on Artificial Intelligence received favorable scores.

As these examples illustrate, data protection law and data protection institutions provide the foundation for how we assess national AI policies and practices.

And consistent with the view that in democratic governments privacy and transparency are complimentary values, we also sought to assess the transparency of AI policy making. So, we asked whether countries had developed a process for *meaningful* public participation, and whether information about a country's AI policies were *readily* available online.

We also know from experience that words alone are not enough. There must also be investigation into *implementation* and *enforcement*.

We looked to the work of Civil Society and NGOs for independent evaluations, groups such as Access Now and BEUC that have repeatedly called for effective enforcement of the GDPR. We looked to Max Schrems and None of Your Business and their legal challenges to ensure that the regulators are doing what they are required to do. And we looked to Privacy International for its overall assessment of global enforcement and to AlgorithmWatch for its specific investigations of accountability for automated decision-making.

This is also why we ask both whether a country has *endorsed* the OECD AI Guidelines, a framework we support, and whether a country has *implemented* the OECD AI Guidelines.

This is why we ask both whether a country has endorsed the Universal Declaration of Human Rights, the foundation of modern human rights law, and whether a country has implemented the Universal Declaration.

Since publication of the report ***Artificial Intelligence and Democratic Values,*** we have engaged policymakers and urged the development of policies that safeguard democratic values.

- This week we asked the United States Office of Management and Budget to begin a public rulemaking on the use of AI by federal agencies. These regulations are required by law and are now overdue. As we explained, “Further delay by the OMB places at risk fundamental rights, public safety, and commitments that the United States has made to establish trustworthy AI.”

- Next week, we will urge the G20 countries to address the challenge of bias in AI. Earlier this year, the G7 leaders rightly called out algorithmic bias. They said they would “take bold action to build more transparency in technologies.” We would like to see the G20 leaders take similar steps to eradicate algorithmic bias at the Summit later this month in Rome.

- And we will stand with our civil society friends in the EU and elsewhere on the need for a prohibition on remote biometric identification. *Our review of country AI practices found that the clearest distinction*

> *between AI systems in authoritarian countries and AI systems in democratic countries is the use of facial recognition for mass surveillance*. Such indiscriminate ongoing surveillance is intended precisely to coerce social behavior and limit freedom and dissent. This AI technique has been used against political protesters and religious minorities, and will almost certainly be more widely deployed unless a clear prohibition is established.

* * *

In presenting the case to you today that data protection is central to the evaluation of AI policies and practices, it is also important to emphasize that many AI applications fall outside the realm of data protection will benefit from the insights of data protection officials.

Professor Stuart Russel, one of the foremost experts on AI, has warned specifically about AI-enabled drone warfare urged a ban on lethal autonomous weapons. In his book, *Human Compatible and the Problem of Control*, Professor Russel writes how we choose to control AI is "possibly the most important question facing humanity"

Today we also confront challenges with the Internet of Things, that tie together devices, from thermostats and power grids, to door locks, electric vehicles, and hydro generators. The Internet of Things amplifies security vulnerabilities and makes it possible to take down power grids, communications networks, and municipal services. And AI-techniques could optimize the time, location, and target for an attack.

Creating autonomous and semi-autonomous devices -- whether network control mechanisms or aerial drones – leaves open fundamental questions about reliability and accountability.

Here Isaac Asimov's intuition that we must establish clear rules to ensure that these automated devices remain under human control is crucial.

However, we do not have Asimov’s Rules for Robots in force today. But we do have Algorithm Impact Assessments that provide the foundation for AI accountability. *Article 22 is now the cornerstone for oversight of AI techniques.* Without the ability to understand and assess the consequences of new AI systems, we will lose control, we will place the public at risk, and we will witness new forms of damage and destruction.

And if we find that we are unable to maintain control of these systems, then we should follow the recommendations of the **Universal Guidelines for AI** and enforce a Termination Obligation on those who have deployed these systems. As with data protection, those who design systems must remain responsible for the consequences of their actions.

*. * *

Perhaps it is too familiar in the data protection world to discuss a new technology, outline its risks to fundamental rights and public safety, and then assess its compliance with legal norms. But that is not my purpose today.

With AI we may indeed begin to find solutions to the pandemic, to global climate change, to the management of power grids, to the disparities in our societies that drive people apart. Instead of embedding racial bias in systems of automated decision-making, data analytic techniques should make it possible to unpack bad models, and end practices that replicate bias.

We might also bring AI to bear on the challenge of data protection enforcement. We need a lot more data from companies so that we can assess their business practices, learn how many complaints they have received, and how they have responded. Consumers sent 28,000 complaints to the US Federal Trade Commission before the agency finally took action against Facebook.

We could also introduce blockchain to verify changes in a company’s privacy practices, to modifications in its AI systems, to provide a verifiable record of changes that may have legal consequences. There is no necessary reason that these techniques cannot be used to advance fundamental rights and the rule of law.

Garry Kasparov, the former world chess champion and ever the optimist, wrote in *Deep Thinking,* “The more that people believe in a positive future for technology, the greater the chance there is of having one." And Garry wrote that book after he lost his world chess champion title to a computer program!

But I also disagree profoundly with the noted computer scientist Kai-Fu Lee who argues in his recent books that AI will necessarily lead to the betterment of humanity, that AI decisions are preferable to human decisions, and that we should remove barriers to data collection, such as privacy laws. In *AI 2041*, for example, he specifically objects to the GDPR goal of data minimization fearing it will “handicap the AI systems.”

This is technological determinism, a commitment to the advance of technology without regard to its human impact. It takes Weizenbaum’s insight about his simple ELIZA program and turns it on its head. Kai-Fu Lee elevates computer decisions over human choices.

That is the world of “Inverted accountability,” with machines making hidden decisions about people, unexplainable, unaccountable, and without human control. It is the world that modern data protection law -- *and the Alianza of data protection officials* -- has stood against for more than 50 years ago

If we are to have a human-centric AI, then humans must remain in control of AI, and we must remain in control of our personal data.

The UN High Commissioner for Human Rights Michelle Bachelet has stated the case directly. She has called data privacy “an essential prerequisite” for the protection of human rights in the context of AI. And she has urged a ban on AI applications that cannot be operated in compliance with international human rights law.

We have two futures ahead. Let us choose the one in which our technologies reflect our values.

Thank you for your attention.