3rd Party LLC

TESTIMONY OF

Samir Goswami, Consultant, 3rd Party LLC

HEARING ON:

Artificial Intelligence: The Consequences for Human Rights

Before the Tom Lantos Human Rights Commission

May 22, 2018

Chairman Hultgren, Chairman McGovern and Members of the Commission, thank you for inviting me to testify today to discuss the implications of Artificial Intelligence, or AI, on human rights. AI is proving to be tremendously beneficial in the transportation, logistics, health care, defense, military intelligence and other sectors; and, while AI is showing great potential to uphold and promote human rights, conversely, it can also be used to suppress it. Today, I will provide examples and conclude with some initiatives already underway to help us better understand and guide AIs implications. However, an important point I want to leave you with is that while AI has tremendously improved our ability to process the world around us, we don't often act upon the insights we glean. That is, while machines may help us understand problems and human rights issues better—we, the humans, have to develop the political will to intervene; unfortunately, this is something we don't do enough of in the human rights space.

A. Introduction & Background:

Conceptually, AI is about building machines that are capable of thinking like, or at least mimicking the thinking processes of humans. Forbes Magazine states, "*AI can be thought of as simulating capacity for abstract, creative, deductive thought, and particularly the ability to learn*." If you have ever asked Siri a question, applied for a credit card online, or ordered a product through Alexa, you have most likely interacted with AI. In practical terms, AI is a collection of technologies that involve the processing of very large amounts of data. For example, banks can use AI to detect fraud patterns by analyzing millions of financial transactions; large grocery stores may use AI to accurately predict customer buying preferences resulting in better control of inventory and less waste. During flight, an airplane's engines may send a constant stream of performance data to a central server, which will analyze the information with other data, such as its age, routes it has flown, weather conditions it has encountered, and even records of the experiences of the pilot. It can process this information and provide an analysis to a human engineer containing a snapshot of system health, flagging anomalies or predicting which parts may need servicing.

Through a process referred to as "machine learning", a core component of AI, machines can be programmed to imitate certain ways our human brains process information; that is a machine can be taught to observe, identify and classify. It can even be taught to make and learn from its mistakes. This incredible progress is possible because of the massive amounts of data that we generate and the improvements in computing power that enables machines to quickly process all the information at its disposal—greatly expanding our ability to solve problems and understand the world around us.

B. **Opportunities, Responsibilities and Applications for Labor and Human Rights**

Amnesty International is piloting the use of machine learning and AI in their human rights investigations and response. When I was a managing director at Amnesty International USA we had access to over 30 years of meticulously recorded human rights data. Our goal was to test if we could use this historic data coupled with records of current occurrences to predict which of the 500 human rights incidents that we tracked every year needed the most attention from Amnesty's campaigners. Through a partnership with Purdue University and the nonprofit DataKind, we had volunteer computer scientists and programmers develop algorithms that sorted through 1.4 million lines and 11,000 files of data—all in a matter of hours. They were able to create a preliminary model that at least in our tests, correctly predicted a binary outcome with over eighty percent accuracy. This type of proactive analysis could enable human rights organizations to create heat-maps of urgency, warn human rights defenders of the severity of risks, or alert first responders to deploy interventions. Amnesty International continues groundbreaking research into AI applications for human rights and has also submitted a statement for the record to this Commission.

AI is also being used effectively to help counter human trafficking. The California based nonprofit organization Thorn is helping law enforcement identify human traffickers through machine learning and AI applications. Traffickers will often advertise the availability of their victims through classified advertisements, for example for escort or other adult services. The same trafficker or trafficking entity may use multiple phone numbers across hundreds of

different advertisements. To find connections and link victims to the same trafficker, law enforcement would have to manually scroll through thousands of these ads that are updated daily, which simply is too much data for a human to process. Thorn is using machine learning and AI to scan the web and the dark web to recognize common identifiers such as phone numbers, or similar styles of writing in such advertisements, and linking that data to the digital footprint that traffickers may leave.

AI is also being used to uncover human rights violations that many workers face around the world in the factories, farms and mines that they labor in. The U.S. Department of Labor finds that 139 goods from 75 countries may be made from child or forced labor. A U.S. company may have thousands of suppliers around the globe that provide raw materials, labor and other services to produce the goods that we consume. Most U.S. companies have codes of conduct that establish labor standards that they expect each one of their suppliers to abide by. However, each one of these suppliers operates in its own legal and social environments and has varying labor practices. Many U.S. companies will use on the ground audits to verify that the supplier is indeed complying with its standards. These audits generate a lot of data—close to 50 to 100 pages per factory, multiplied by thousands of factories, at least once a year.

Supply chain managers can use machine learning and AI to process this vast amount of audit data to flag issues. However, audits can be forged or be susceptible to other influences—this erroneous data can contaminate an AI enabled analysis that may then not paint a complete and accurate picture of whether the supplier is acting ethically and workers' rights are being upheld. AI systems can be used to conduct outside validation to complement audit data through accessing and processing other information sources, such as news reports, court filings, public records, any materials that compromise the open source data footprint of a supplier and its business associates. Furthermore, workers also leave a data footprint—machines can scan social media, chat forums, message boards or public comment websites for any references about those suppliers made by workers. Mobile phones can be used to deploy surveys to workers directly, independent of supplier supervision, data from which can also be

incorporated into an analysis. All these various streams of data can and are being analyzed together by an AI enabled system to provide an independent human rights assessment of a supplier's labor practices.

It is important to note that all of this data and computing power are only useful if we actually act upon the AI informed insights we glean—which we don't often do, a point I will expand upon in the recommendations section of this testimony.

C. Consequences

I submit three broad categories of consequences for the Commission to consider. This is by no means an exhaustive list:

1. First, while AI can be a powerful tool in the labor rights space, the above example also leads to the obvious risks that AI can pose; that is, this level of analysis can also be flipped for illicit purposes. For example, machine learning and AI can be used to comb through worker social media posts to identify union organizers or those a State or factory owner may deem to be a "trouble maker". Facial recognition technology can be coupled with AI and machine learning to target migrant workers or human rights defenders who are challenging repressive labor regimes; and, predictive capabilities might flag workers and subject them to arbitrary detention or harassment based on the AI informed suspicion that they *might* challenge employment practices and poor working conditions in the future. While I may be able to point to numerous pilot projects that utilize AI to specifically advance human and labor rights—rogue states and actors can use technology enhanced by machine learning and AI to suppress civil and human rights at considerable scale.

2. Second, those developing AI systems have to be very aware of the human prejudices that the machine may inherit. AI is only as good as the data it learns from—if the data has biases they will be amplified by the machines. For example, if a computer is ingesting large quantities of employment data to inform an algorithm that selects candidates for a high

paying job, and that data, because of historic biases, mainly contains data sets of men, this may lead to programmed gender biases that can exacerbate discrimination against women.

3. Lastly, and importantly, AI can be used to generate wealth for a select few at the expense of others. Increased automation can lead to the decrease in employment, displace low wage workers and depress wages. A recent publication by the Council on Foreign Relations states, "*Accelerating technological change, including automation and advances in artificial intelligence that can perform complex cogitative tasks, will alter or replace many human jobs.*" A 2016 White House report on AI and the economy states that "*Because AI is not a single technology, but rather a collection of technologies that are applied to specific tasks, the effects of AI will be felt unevenly through the economy*." This White House report finds that anywhere from nine percent to forty seven percent of jobs over the next two decades could be disrupted by AI and automation. This also holds true for U.S. companies with overseas supply chains. For example, the International Corporate Accountability Roundtable estimates that two-thirds of all jobs in developing countries could face significant automation, mainly in the apparel, electronics and agricultural sectors. A 2016 report by the International Labour Organization identifies the risks that "automation, robots and artificial intelligence" will place on millions of workers in Asia. These impacts will be born on the shoulders of low-income women and migrants who are already some of the world's most vulnerable.

   However, these impacts have not been realized just yet, and can still be mitigated. The Council on Foreign Relations also states that, "*In the absence of mitigating policies, automation and artificial intelligence are likely to exacerbate inequality and leave more Americans behind."* We thus need to ensure that as technology is helping us make improvements in the marketplace to how we produce, distribute and consume things—we are not further exacerbating existing discriminatory tendencies, or making things worse for those already vulnerable—all grave human rights concerns.

D. Considerations & Recommendations:

Given these consequences, I submit the following considerations and recommendations:

1. First, the federal government should invest in "AI for good" and provide seed funding for such applications. However, we also need to act upon the insights that AI and machine learning deliver to us. Our investment criteria shouldn't just be that the technology was developed or deployed—we should measure the human outcomes that actually happened. That is, did the "good" that we envisioned in an "AI for good" technology application actually happen? A machine can help us better understand, but humans have to intervene. AI may help us find and understand a problem better and get down to the most accurate data point—but we still have to act upon it. For example, we can use AI to help pinpoint exactly which factory might utilize child labor—but that insight is wasted if we don't respond, and we deliberately ignore our ethical obligations to those children. Unfortunately, while some progress is being made, U.S. companies are simply not doing enough to act upon the technology enabled insights on labor and human rights abuses that they have access to.

   Additional policies and public and political pressure are needed to compel companies to actively monitor their supply chains for human rights abuses, and increased legal accountability is needed for those who don't. There is thus tremendous opportunity for the U.S. Department of Labor, USAID, DOJ, CBP and others to use AI and machine learning to verify how workers tied to U.S. public and private supply chains are being treated. These insights can be used to apply laws already on the books that prohibit forced labor and child labor made goods to enter the U.S., or to enforce trade agreements that have often-ignored labor protections in place. The U.S. government needs to act upon such technology-gleaned insights by compelling companies to drive supply chain improvements, enable law enforcement to prosecute those who abuse human rights, and press other governments to uphold workers' rights. In the supply chain management field, even without AI there is ample technology already available to determine if a supplier is treating its workers fairly; however, both governments and companies don't often act upon it.

2. Second, as the 2016 White House report on AI stated, "*Whether AI leads to unemployment and increases in inequality over the long-run depends not only on the technology itself but also on the institutions and policies that are in place*." Thus, at a macro level we need to ensure that we are properly preparing and training our workforce to avail of new technology jobs, and meaningfully assist in transitioning those who are at risk of being displaced. This includes workers in the territorial United States and those whose livelihoods are tied to the supply chains of U.S. companies overseas, with a particular focus on women and migrant workers who often conduct low wage work. Unfortunately, our collective track record is not great on these fronts. Employment loss due to technological advancements has happened on many occasions, while U.S. investments in labor market programs, such as job readiness and high-skills training, has decreased significantly over time, and is far less than those made by other industrialized countries.

3. Third, U.S. corporations developing and deploying AI need to incorporate a rights-based approach. AI systems need to be designed in ways that don't replicate human biases. Engineering teams designing AI systems need to be diverse, and the data that feeds into their systems also have to be corrected for biases. The Global Future Council on Human Rights recommends four central principles to combat bias and uphold human rights in machine learning: active inclusion, fairness, the right to understanding and access to remedy. For example, a consumer should be informed if AI was used to influence a decision about their lives (e.g. whether you get a mortgage), and should have access to a process for redress for erroneous or biased interpretations. The Center for Data Innovation recommends the concept of "algorithmic accountability", which they define as "*the principle that an algorithmic system should employ a variety of controls to ensure the operator can verify it acts in accordance with its intentions, as well as identify and rectify harmful outcomes*". In short, businesses need to ensure that the AI systems they create and utilize are not creating value for their shareholders and customers at the expense of human and civil rights.

E. Multi-stakeholder Initiatives to Understand AI and its Consequences on Human and Civil Rights

Overall, I believe that AI and other technologies have and can have a tremendous positive impact on human rights, and we need to prepare ourselves for the resulting implications in a collaborative manner. Already we are seeing progress that should be continued and supported. For example:

1. The Council on Foreign Relations, through their "The Work Ahead" project makes numerous recommendations including the need to strengthen the link between education and work through increased investments. The Council also calls on the nation's governors, Congress and the Administration to collectively establish a process to understand and address such technology implications.

2. The International Corporate Accountability Roundtable, a Washington DC based civil society organization through its "Robots and Rights" project is mapping sectors that rely heavily on low-skill human labor and how they will be impacted by automation and mechanization with a report to be released next month. The report will include "*policy and advocacy strategies and solutions from the perspective of both States and companies*".

3. Finally, the Partnership on AI, a San Francisco based non-profit association has been founded and funded by leading companies that are at the forefront of developing and applying AI technologies. The Partnership has brought together companies including Google, Amazon, Facebook, eBay and others with nonprofit civil and human rights organizations such as the ACLU, Amnesty International, the Center for Democracy and Technology, and the Electronic Frontier Foundation, to collectively identify solutions and safeguards to various AI influenced challenges. The Partnership aims to develop analysis and recommendations on how AI influences labor, the economy, and social good among other topics.

F. Conclusion:

Like any technology, AI has both positives and negative applications and effects. With the exponential surge in the availability of data and computing power comes an increase in our reliance on machines to help us do things we couldn't imagine a few years ago. Better data collection and analysis, finding patterns of human rights violations in large data sets, enabling quicker response to human rights incidents, all increase our ability to help one another. However, we have to choose to do so and not just limit AIs potential to commercial applications whose sole purpose is to increase wealth. And, AI will generate wealth—we need to ensure that that this prosperity will be shared broadly and not exacerbate existing economic disparities.

We need to proactively guide AIs myriad applications and prepare ourselves for the resulting implications in a collaborative manner. The development and deployment of AI technologies has to be within a policy framework that takes human rights principles into account, and the application of these technologies has to be matched with policies and programs that adequately prepare those who might stand to lose the most.

Most of all, we need to act upon the insights we glean: technology is just a tool to help us understand a problem better—it is not a replacement for the political will that is needed to drive change.

Thank you for your time and leadership and the opportunity to address this commission.