

Robotics, AI, and Humanity: Science, Ethics, and Policy

Vatican City, 16-17 May 2019

“Regulating AI: considerations that apply across domains”

By Angela Kane

While we think of AI as a phenomenon that has rapidly arisen over the last few years, we should remember that it was already eighty years ago that Alan Turing laid down the mathematical basis of computation. ARPANET began in 1969, the Internet Protocol in 1974, and the World Wide Web thirty years ago, in 1989.

Do any of you remember when we first started to access the internet, with modems? The distinctive whirring burpy sound they made when connecting – ever so slowly – to the web? This now seems very quaint, as the improvements in speed and performance, as well as the cost reductions in memory and information technology, have made possible the enormous expansion of data that now fuels the engine of global growth.

Harnessing AI has challenges and opportunities in many areas and domains: technical, ethical, political, social, cultural. These are accompanied by the need for accountability, algorithmic explainability and even legal liability. If we don't understand how a system works, then it blurs lines of who can or who should be responsible for the outcome or the process of the decision. Should that be the innovator? The regulator? The operator? And how can both policy-makers and the public trust technology when it is not properly understood?

These are vexing questions that have been further compounded by the rise in disclosures of data and privacy leaks, of hacking into sites containing sensitive personal information, of spoofing, of selling consumer data without consent and, to make matters worse, of concealing or delaying disclosure of such egregious violations of privacy.

The debate about these issues has become louder and more polarized; it is pitting powerful companies against governments and consumers. Scientists are weighing in - as do employees of technology companies, as we have seen with Google. Until 2015, Google's motto was “Don't be evil” but it was then changed to “Do the right thing” within its corporate code of conduct. Swarms of bots, dark posts and fake news websites inundate the web, ricochet around chatrooms, overwhelm the legitimate media outlets.

Let us remember just a few recent events: in the US presidential elections in 2016, Russia supported one candidate (who subsequently won) by waging a campaign with paid advertisements and fake social media accounts that contained polarizing content.

Concerns also abound in China about millions of cameras deployed with face recognition software which record streams of data about citizens. In India, it was reported that the “fake news problem plagues several popular social networks”¹ by spreading misinformation, doctored photos and videos which resulted in several cases of killing and even lynching.

More and more thoughtful questions about social platforms are being asked that do not lend themselves to easy answers. Technology companies are coming under increasing scrutiny, as they are seen to be operating without accountability. Facebook CEO Mark Zuckerberg testified last year in US Congress on efforts to address privacy issues and data sharing, but subsequent Facebook data leaks showed that his assurances to prevent a recurrence were hollow. Talking about regulation, he said: “My position is not that there should be no regulation. I think the real question is, as the internet becomes more important in people’s lives, is what is the right regulation, not whether there should be or not.”²

I will take stock of some of the efforts to address the attempts to regulate AI and technology, fully aware that the paper will outdate very quickly, as new initiatives and considerations are coming up quickly.

Curbing lethal autonomous weapons systems: an early effort at regulating AI

Following the publication of a report by the United Nations Special Rapporteur on extrajudicial, summary or arbitrary execution, Christof Heyns, in 2013, on the use of lethal force through armed drones from the perspective of protection of the right to life, sixteen countries put the questions related to emerging technologies on the agenda of the Convention on Certain Conventional Weapons (CCW) in Geneva. The first meetings on these issues took place in 2014 and they showed that few countries had developed any policy on the matter. Thematic sessions, with significant input from AI scientists, academics and activists, dealt with legal aspects, ethical and sociological aspects, meaningful human control over targeting and attack decisions, as well as operational and military aspects.

Five years after the first discussions, there is now a Group of Governmental Experts to advance the issue – operating on a consensus basis, which essentially gives a veto right to any participant. Twenty-eight States now openly call for a ban on these weapons. Austria, Brazil and Chile have recently proposed a mandate to “negotiate a legally-binding instrument to ensure meaningful human control over the critical functions of weapons systems”, but the prospects for such a move are slim. So far, no legally binding or political actions have been adopted by the Group due to the objections of several States: Australia, Israel, Republic of Korea, Russian Federation and

¹ <https://www.thenextweb.com/in/2019/01/29/its-not-just-whatsapp-indias-fake-news-problem-plagues-several-popular-social-networks/> of 29 January 2019

² <https://theguardian.com/technology/2018/apr/11/mark-zuckerbergs-testimony-to-congress-the-key-moments>
11 April 2018

the United States. These States argue that concrete action on LAWS is “premature” and that the Group could instead explore “potential benefits” of developing and using LAWS.

The opposing positions do not augur well for any legislative progress in the issue of LAWS. Yet the voices in favour of a total ban are getting louder and louder. Already in 2015, at one of the world’s leading AI conferences, the International Joint Conference on Artificial Intelligence (IJCAI 15), an Open Letter from AI & Robotics Researchers - signed by nearly 4,000 of the preeminent scientists such as Stuart Russell, Yann LeCun, Demis Hassabis, Noel Sharkey and many many others – and over 22,000 endorsers including Stephen Hawking, Elon Musk, Jaan Tallinn, to name just a few – who warned against AI weapons development and posited that “most AI researchers have no interest in building AI weapons, and do not want others to tarnish their field by doing so”³. The decision by Google to end cooperation with the US Department of Defense on Project Maven – a minor contract in financial terms – was ended in 2018 due to strong opposition by Google employees who believed that Google should not be in the business of war. UN Secretary-General Guterres, former High Commissioner for Human Rights Zeid Ra’ad Al Hussein, and Pope Francis have weighed in, calling autonomous weapons “morally repugnant” and calling for a ban.

There are also parliamentary initiatives in capitals. In April last year, for example, the Lord’s Select Committee on AI challenged the UK’s futuristic definitions of autonomous weapons systems as “clearly out of step” with those of the rest of the world and demanded that the UK’s position be changed to align these within a few months.

Yet the Government’s response was limited to one paragraph which stated that the Ministry of Defense “has no plans to change the definition of an autonomous system” and notes that the UK will actively participate in future GGE meetings in Geneva, “trying to reach agreement (on the definition and characteristics of possible LAWS) at the earliest possible stage”⁴.

Interest in other European parliaments is also high, as awareness of the issue has grown exponentially. It’s the hot topic of the day.

The European Commission issued a communication⁵ in April 2018 with a blueprint for “Artificial Intelligence for Europe”. While this does not specifically refer to LAWS, it demands an appropriate ethical and legal framework based on the EU’s values and in line with the Charter of Fundamental Rights of the Union.

In July 2018, the European Parliament adopted a resolution⁶ that calls for the urgent negotiation of “*an international ban on weapon systems that lack human control over the use of force*”. The

³ <https://futureoflife.org/open-letter-autonomous-weapons/>

⁴ See Government response to House of Lords Artificial Intelligence Select Committee’s Report on AI in the UK: Ready, Willing and Able?, (recommendations 60-61), June 2018, <https://www.parliament.uk/lords-comittees/Artificial-intelligence/AI-Government-Response2.pdf>

⁵ https://ec.europa.eu/newsroom/dae/document.cfm?doc_id=51625

⁶ European Parliament resolution of 12 September 2018 (2018/2752(RSP))

resolution calls on the European Council to work towards such a ban and “urgently develop and adopt a common position on autonomous weapon systems”. In September 2018, EU High Representative Federica Mogherini told the EU Parliament that “the use of force must always abide by international law, including international humanitarian law and human rights laws. (...) How governments should manage the rise of AI to ensure we harness the opportunities while also addressing the threats of the digital era is one of the major strands of open debate the EU has initiated together with tech leaders”⁷

The issue of lethal autonomous weapons has clearly raised the profile of legislating AI. Advocacy by civil society, especially the Campaign to Stop Killer Robots, a coalition of NGOs seeking to preemptively ban lethal autonomous weapons, has been instrumental in keeping the issue prominent in the media, but this single-issue focus is not easily replicable in other AI-driven technologies.

Can we ever hope to regulate and govern AI?

As in the case of lethal autonomous weapons, I believe it is easier to address a specific AI application than general AI that is broad, adaptive, and advanced as a human being across a range of cognitive tasks.

In the first case, automated decision systems are currently being used by public agencies, in criminal justice systems, in predictive policing, in college admissions, in hiring decisions and many more. Are these automated decision systems appropriate? Should they be used in particularly sensitive domains? How can we fully assess the impact of these systems? Whose interests do they serve? Are they sufficiently nuanced to take into account complex social and historical contexts? Do they cause unintended consequences?

The difficulty in finding answers to these questions is the lack of transparency and information. Many of these systems operate in a black box and thus outside the scope of understanding, scrutiny and accountability. Yet algorithms are endowed with a specific structuring function, as designed by individuals. In the US, this has already led to several lawsuits which proved that decision-making formulas were corrupt due to data entry errors and biased historical data. While this showed the limits of AI use in public policy, it is clear that lawsuits set precedent in law but cannot establish regulations and the rule of law.

But if the litigation shows us anything, it is that AI-driven technology has become an important issue for people and for governments. In response, we are seeing two distinct trends:

⁷ https://eeas.europa.eu/topics/economic-relations-connectivity-innovation/50465/autonomous-weapons-must-remain-under-human-control-mogherini-says-european-parliament_en

- The AI and tech industry have become a hub for ethics advisory boards and related efforts to buff their credentials in what I would call “responsible AI”;
- Private organizations have been established like Partnership for AI (mission: to benefit people and society), or Open AI (mission: to ensure that artificial general intelligence benefits all of humanity);
- Academic institutions – such as New York University - have set up institutes like AI Now, a research institute examining the social implications of AI; the Massachusetts Institute of Technology (MIT) conducts a project on AI Ethics and Governance to support people and institutions who are working to steer AI in ethically conscious directions;
- Workshops and conferences with a range of tech and non-tech stakeholders are being organized to debate the scope of the challenges as well as exploring solutions.

Governments are stepping up

The second trend is the increasing focus by Governments on the disruption by artificial intelligence and the search for shaping the ethics of AI. Let me mention some statements by leaders.

When Russian President Putin said last year to a group of school children that “whoever controls AI, will become the ruler of the world”, it made headlines. China’s blueprint – issued in 2017 and called the “New Generation Artificial Intelligence Development Plan” – outlined China’s strategy to become the world player in AI by 2030. The Plan barely mentions information on laws, regulations and ethical norms, since China is not hampered by values and fundamental rights as well as ethical principles such as accountability and transparency. In the two years since its publication, China is already starting to overtake the US as the leader in AI.

In Europe, French President Macron last year called the technological revolution that comes with AI “in fact a political revolution”, and said that in shaping how AI would affect us, you have to be involved at the design stage, *and set the rules* (italics added).

A French data protection agency (Commission Nationale de l’Informatique et des Libertés, CNIL) issued a 75-page report in December 2017 about the results of a public debate about AI, algorithms, ethics and how to regulate it. The report set out six areas which predominate the ethical dilemmas:

1. Autonomous machines taking decisions;
2. Tendencies, discrimination and exclusion which are programmed, intentionally or unintentionally;
3. Algorithmic profiling of people;
4. Preventing data collection for machine learning;
5. Challenges in selecting data of quality, quantity and relevance;
6. Human identity in the age of artificial intelligence.

Recommendations made in the report primarily focus on the individual by urging enhanced information and education but also request private industry to focus on ethics by establishing ethics committees and an ethics code of conduct or an ethics charter.

In the UK, the House of Lords Select Committee on Artificial Intelligence issued a report in April 2018 with the catchy title “AI in the UK: ready, willing and able?”⁸ The report was based on extensive consultations and contains an assessment of the current state of affairs as well as numerous recommendations on living with AI, and on shaping AI.

The 183-page report has only two paragraphs on “regulation and regulators” which state that “Blanket AI-specific regulation, at this stage, would be inappropriate. We believe that existing sector-specific regulators are best placed to consider the impact on their sectors of any subsequent regulation *which may be needed*” (emphasis added). It also urges the Government Office for AI to “ensure that the existing regulators’ expertise is utilized in informing any potential regulation that may be required in the future” and foresees that “the additional burden this could place on existing regulators could be substantial”, recommending adequate and sustainable funding⁹.

In its final paragraphs, the report refers to the preparation of ethical codes of conduct for the use of AI by “many organizations” and recommends that a cross-sectoral ethical code of conduct – suitable for implementation across public and private sector organizations be drawn up (...) with a sense of urgency. “In time, the AI code could provide the basis for statutory regulation, *if and when this is determined to be necessary*”¹⁰ (emphasis added).

In June 2018, the Government issued a 42-page response to the House of Lords’ report. As to paragraph 386 (no blanket AI-specific regulation needed), the Government agreed with the recommendation. It stated its commitment to work with businesses to “develop an agile approach to regulation that promotes innovation and the growth of new sectors, while protecting citizens and the environment”¹¹. It further promises horizon-scanning and identifying the areas where regulation needs to adapt to support emerging technologies such as AI and the establishment of a Centre for Data Ethics and Innovation that “will help strengthen the existing governance landscape”¹². Yet the Centre – established late last year - has only an advisory function, promoting best practices and advising how Government should address potential gaps in the regulatory landscape.

Others European countries also addressed AI. Sweden published a report in May 2018 on its National Approach (a digestible 12 pages) which highlights the Government’s goals to develop

⁸ <https://publications.parliament.uk/pa/ld201719/ldselect/ldai/100/100.pdf>

⁹ *Ibid.*, paras. 386-387.

¹⁰ *Ibid.*, para. 420.

¹¹ <https://www.parliament.uk/business/committees/committees-a-z/lords-select/ai-committee/news-parliament-2017/government-response-to-report/>

¹² *Ibid.*, para. 108

standards and principles for ethical, sustainable and safe AI, and to improve digital infrastructure to leverage opportunities in AI. Finland was a bit ahead of the curve, issuing its first report on “Finland’s Age of Artificial Intelligence” already in December 2017, but none of its eight proposals deal with rules and regulations.

Germany issued a 12-point strategy (“AI Made in Germany - a seal of excellence”) which focuses on making vast troves of data available to German researchers and developers, improve conditions for entrepreneurs, stop the brain drain of AI experts and loosen or adapt regulation in certain areas, but it also heavily emphasizes the rights and advantages of AI for the citizens and underlines the ethical and legal anchoring of AI in Europe.

The European Union: “placing the power of AI at the service of human progress”

Finally, let me focus on the European Union which in April 2018 issued “AI for Europe: Embracing Change”¹³. This was the launch of a European Initiative on AI with the following aims:

1. Boost the EU’s technological and industrial capacity and AI uptake across the economy
2. Prepare for socio-economic change
3. Ensure an appropriate ethical and legal framework

Under these three headings, ambitious plans were laid out, both in financial terms (stepping up investments) and in deliverables, with time lines until the end of 2020.

Let us not forget that the General Data Protection Regulation (GDPR) came into force the same year. While this regulation imposes a uniform data security law on all EU members, it is important to note that any company that markets good and services to EU residents, regardless of its location, is subject to the regulation. This means that GDPR is not limited to EU member states, but that it will have a global effect.

One of the deliverables was the setting up of an Independent High-Level Expert Group on Artificial Intelligence¹⁴ which was asked to draft AI ethics guidelines and, through an online framework called the European AI Alliance reached out to stakeholders and experts to contribute to this effort.

The draft ethics guidelines were issued in December 2018 and received over 500 comments, according to the EU. What resulted were the “Ethics Guidelines for Trustworthy AI”¹⁵, issued in April 2019, which defines trustworthy AI as follows: “(It) has three components: (1) it should be lawful, ensuring compliance with all applicable laws and regulations (2) it should be ethical,

¹³ <https://ec.europa.eu/digital-single-market/en/news/communication-artificial-intelligence-europe> of 25 April 2018

¹⁴ Full disclosure: I was a reserve member of the High-Level Expert Group and participated in several of their meetings.

¹⁵ <https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai> issued on 8 April 2019

demonstrating respect for, and ensure adherence to, ethical principles and values and (3) it should be robust, both from a technical and social perspective since, even with good intentions, AI systems can cause unintentional harm. Trustworthy AI concerns not only the trustworthiness of the AI system itself but also comprises the trustworthiness of all processes and actors that are part of the system's life cycle."

The Guidelines then list seven essentials for achieving trustworthy AI:

1. Human agency and oversight
2. Robustness and safety
3. Privacy and data governance
4. Transparency
5. Diversity, non-discrimination and fairness
6. Societal and environmental well-being
7. accountability

Again, the Guidelines are currently in a pilot phase for more time to receive feedback and to ensure that they can be issued in the summer of 2019 and then implemented – which is expected in 2020.

At the same time, the EU Commission wants to bring their approach to AI ethics to the global stage: "because technologies, data and algorithms know no borders". Following the G7 summit in Canada in December 2018, where AI was prominently featured, the EU wants to strengthen cooperation with other "like-minded" countries like Canada, Japan and Singapore, but also with international organizations and initiatives like the G20 to advance the AI ethics agenda.

Before we break out the champagne in celebration of the ethics guidelines, let me mention one dissenting voice from the High-Level Group: Thomas Metzinger, Professor of Theoretical Philosophy in Germany, wrote a scathing article entitled "Ethics washing made in Europe"¹⁶ in which he called the Trustworthy AI story "a marketing narrative invented by industry, a bedtime story for tomorrow's customers". The narrative, he claimed is "in reality, about developing future markets and using ethics debates as elegant public decorations for a large-scale investment strategy".

Professor Metzinger considers that "industry organizes and cultivates ethical debates to buy time – to distract the public and to prevent or at least delay effective regulation and policy-making. And politicians like to set up ethics committees because it gives them a course of action when, given the complexities of the issues, they simply don't know what to do". Interestingly, he also mentions the use of lethal autonomous weapons systems as one of the "Red Lines", the non-negotiable ethical principles – which I outlined at the beginning of this paper.

¹⁶ <https://www.tagesspiegel.de/politik/eu-guidelines-ethics-washing-made-in-europe/24195496.html> published on 8 April 2019

Ethical AI – the new corporate buzz phrase

I agree that the jury on the EU Ethics Guidelines is still out, but the criticism of major tech companies and academic ethics boards, especially in the US, is very strong. Many tech companies have recently laid out ethical principles to guide their work on AI. Major companies like Microsoft, Facebook, and Axon (which makes stun guns and body cameras for police departments), all now have advisory boards on the issue. Amazon recently announced that it is helping fund research into “algorithmic fairness”, and Salesforce employs an “architect” for ethical AI practice, as well as a “chief ethical and human use” officer. More examples could be cited.

Yet are these actions are designed primarily to head off new government regulations? Is it a fig leaf or a positive step?

“Ethical codes may deflect criticism by acknowledging that problems exist, without ceding any power to regulate or transform the way technology is developed and applied,” wrote the AI Now Institute, a research group at New York University, in a 2018 report¹⁷. “We have not seen strong oversight and accountability to backstop these ethical commitments.”

The boards are also seen to mirror real-world inequality (mostly white men, very few women, few or no people of color or minorities)¹⁸ or to have members who do not represent ethical values. The establishment of an ethics board by Google (actually called Advanced Technology External Advisory Council, ATEAC) lasted barely a week before it was disbanded amid great controversy.

The Google debate shows that discussing these issues in the public eye also invites public scrutiny. While I consider it positive that private industry is studying the issues and inviting views on company ethics, it is ultimately the CEO who gets to decide which suggestions on AI ethics would be incorporated into what are essentially business decisions. A company is clearly more concerned with the financial bottom line rather than sacrificing profit for ethical positions taken by an external advisory board, as there is no legal obligation to follow what are well-intentioned recommendations.

So the issue revolves around accountability, and in my view, government regulation will be needed to enforce it.

Doteveryone, a UK organization (mission: Responsible Technology for a Fairer Future), recently issued a report¹⁹ entitled “Regulating for Responsible Technology” which calls for a new independent regulatory body with three responsibilities:

1. give regulators the capacity to hold technology to account;

¹⁷ https://ainowinstitute.org/AI_Now_2018_Report.pdf December 2018

¹⁸ As an example, see the following article of 30 March 2019, <https://www.theguardian.com/technology/2019/mar/28/big-tech-ai-ethics-boards-prejudice>

¹⁹ <https://doteveryone.org.uk/press-events/responsible-tech-2019/>, October 2018

2. inform the public and policymakers with robust evidence on the impacts of technology;
3. support people to seek redress from technology-driven harms.

In addition to outlining that we currently have a “system in need of a steward”, the organization also has a directory of regulation proposals in the UK to which it invites users to update²⁰. More surveys of such proposals might be very helpful in determining how best to go forward.

We should, however, also look at “soft law” which are substantive expectations that are not directly enforceable, as opposed to “hard law” which are legally enforceable requirements imposed by governments. As outlined by Wallach and Marchant²¹, soft law includes voluntary programs, standards, codes of conduct, best practices, certification programs, guidelines and statements of principles.

As an example of soft law being turned into hard law, they cite the Future of Life Institute Asilomar Principles adopted in 2017 as a soft law tool for AI governance, which have now been adopted by the State of California into its statutory law.

A paradigm shift is emerging

I believe one of the problems of the EU’s High-Level Expert Group on AI is that it tries to be all-comprehensive and therefore tends towards more general and lofty declarations rather than be prescriptive in application. As I noted at the beginning of this paper, it is easier to address regulation in one aspect of AI rather than the entire gamut of applications. Let me focus on one such aspect that has started to capture attention in a major way: facial recognition and the pervasive use of cameras.

The Turing Award has just been given to three preeminent computing scientists for their work on neural networks which has, inter alia, accelerated the development of face-recognition services. Yet they - together with some two dozen prominent AI researchers - have signed a letter to Amazon to stop selling its face-recognition technology (called “Rekognition”) to law enforcement agencies because it is biased against women and people of colour.

Facial recognition technology (FRT) has been used by government agencies, by retail industry, by Facebook with its millions of users posting photographs. In China, more than 176 million CCTV cameras are used for street monitoring and policing as well as in “cashless” stores and ATMs: where does consumer assistance start and surveillance begin?

Despite some positive aspects (reuniting missing children in India), there are major concerns about how to protect the privacy of those whose data is collected. With an industry quickly

²⁰ <https://docs.google.com/document/d/1b6xZtYNAL2O3DT7bDTHY2DdvTtNIOVGtTeRecFAwFI4/edit#>

²¹ See the article by Wendell Wallach and Gary Marchant, “Toward the Agile and Comprehensive International Governance of AI and Robotics”, Proceedings of the IEEE, March 2019, <https://ieeexplore.ieee.org/document/8662741>

mushrooming to an estimated more than \$10 billion in the next few years, alarms are beginning to sound about the lack of governmental oversight and the stealthy way it can be used to collect data on crowds of people – as we learned when it was revealed that the musician Taylor Swift had deployed FTR during her performances to root out stalkers. But is the technology only used for security?

Containing FTR is easier in Europe, where strict privacy laws are being enforced with the GDPR, but in other countries (and continents) no regulations exist. Yet even here in Europe people are warning against the “surveillance state”. Looking at the increasing coverage and discussion of FTR, I am of the opinion that this will be one area of focus for regulation in the near future.

Could there be a role for international organizations or institutions?

UN Secretary-General Antonio Guterres weighed in on AI in July 2018, stating that “the scale, spread and speed of change made possible by digital technologies is unprecedented, but the current means and levels of international cooperation are unequal to the challenge”²². He set up a High-Level Panel on Digital Cooperation, with Melinda Gates and Jack Ma as Co-Chairs, and 18 additional members serving in their individual capacity. Their task is to submit a report by mid-2019 - contributing to the broader public debate - which identifies policy, research and information gaps, and makes proposals to strengthen international cooperation in the digital space.

The Panel has reached out and sought comments on their efforts from people all over the world, conducting a “global dialogue” to assist in reaching their final conclusions. What can be expected? First, it is important to bring this discussion to all member states, many of whom do not have the capacity to harness new technology and lack a sophisticated understanding of the matter. It is also important for the organization to embed this report in the universal UN values, and to consider practical ways to leverage digital technologies to achieve the Sustainable Development Goals.

The most important issue, in my opinion, is, however, to take stock of existing – and emerging – normative, regulatory and cooperative processes. I would not expect the UN to set rules and standards, but to have an inventory of the current state of affairs would be very valuable for national efforts to build on.

Past efforts by UN high-level panels have had mixed success. Despite the enormous work that goes into reports by high-ranking participants, their recommendations have at times been taken note of, politely debated – and then disappeared into a drawer without seeing implementation. Let us hope that the prominent co-chairs will contribute to a lively open debate and ensure that the recommendations in the final report will see further application.

²² SG/A/1817 of 12 July 2018

Summing Up: fifteen recommendations

Rapidly emerging technologies – AI and robotics in particular – present a singular challenge to regulation by governments. The technologies are owned by private industry, they advance in the blink of an eye, they are not easily understood due to their complexity and may be obsolete by the time a government has agreed to regulate them.

This means that traditional models of government regulation cannot be applied. So if not regulation, what can be done? Here are my proposals:

1. Expand AI expertise so that it is not confined to a small number of countries or a narrow segment of the population
2. Accept that the right decisions on AI technology will not be taken without strong input from the technologists themselves
3. Find therefore a common language for government officials, policy makers and technical experts
4. Begin dialogue so that (a) policies are informed by technical possibilities and (b) technologists/experts appreciate the requirements for policy accountability
5. Discuss how to build a social license for AI, including new incentive structures to encourage governments and private industry to align the development and deployment of AI technologies with the public interest
6. Focus on outcome not process: principles, privacy protection, digital policy convergence; differences in legal and regulatory systems and cultures between US, EU, China
7. Establish some “Red Lines” – no-go areas for AI technology, such as lethal autonomous weapons systems, AI-supported assessment of citizens by the government (“social scoring”)
8. Use the strategy of “soft law” to overcome limitations and challenges of traditional government regulation for AI and robotics
9. Discuss the challenges, costs, reliability and limitations of the current state of art
10. Develop strong working relationships, particularly in the defense sector, between public and private AI developers
11. Ensure that developers and regulators pay particular attention to the question of human-machine interface
12. Understand how different domains raise different challenges
13. Compile a list of guidelines that already exist and see where there are gaps that need to be filled to offer more guidance on transparency, accountability and fairness of AI tools
14. Learn from adjacent communities (cyber security, biotech, aviation) about efforts to improve safety and robustness
15. Governments, foundations and corporations should allocate funding to develop and deploy AI systems with humanitarian goals

I encourage others to add to the list. What is really important here is that we come to a common understanding of what needs to be done. How do we develop international protocols on how to develop and deploy AI systems? The more people ask that question, the more debate we have on it, the closer we will get to a common approach. I hope all of you will continue this discussion.

* * *