

MUSTAFA SULEYMAN, DEEPMIND

INTRO

Thank you Mr. President and your excellencies for the invitation to speak to you about what artificial intelligence can do to help achieve the Sustainable Development Goals.

As a global community, we've made significant strides over the last few decades, taking on some of the world's greatest tragedies like child mortality. Every single day 17,000 new lives get to be lived, by children who would have died just 25 years ago. Peaceful cooperation and relentless innovation has been the driving force of this truly spectacular progress.

And, yet, in many ways inequality hasn't improved—it's become worse. Malnutrition and preventable disease continue to kill millions, impacting both developed and developing world healthcare systems, and the devastating threat of climate change looms, with the effects set to hit the poorest the hardest.

If we are to decrease this level of suffering, then we'll have to make sure humanity is doing everything possible to come up with bold, new solutions. But to have any chance of truly solving these problems - one of the great goals of human civilisation **and I would argue core to the founding purpose of this great institution** - then what's possible won't be enough. We must take on what is currently **impossible**, and do all we can to push back the limits of what humanity can achieve.

These limits are real, and they cap all of our aspirations for change.

Our efforts to tackle disease are capped by a desperate shortage of trained nurses and doctors — in the rich West as well as in developing nations. Our efforts to reduce energy consumption are held back by the insatiable demand for new products and services, and what seem like hard limits on how efficient our energy infrastructure can become.

This is why science and technology has been so critical to the history of civilization. Technological progress expands the possibilities of human achievement, increasing our collective capability to solve problems that were once considered unsolvable.

As a technologist as well as an activist, my message today is simple: first, that the progress of science and technology is about to go through **the greatest acceleration of all time**.

And, second, that **this is the greatest opportunity we have had for generations** to advance the causes of social justice, equality and the reduction of human suffering.

AI'S POTENTIAL TO HELP ACHIEVE SDGS

My confidence is based on what I see every day at places like DeepMind, the company I co-founded, working at the cutting edge of artificial intelligence research and application.

AI means different things to different people: what I mean is technology that people can use in complex domains to discover new knowledge, ideas and strategies well beyond our reach today.

This AI-enabled future has already begun.

Take energy, where we've already seen AI make dramatic impact. By applying our algorithms to Google's data centers we saved 40% of the energy used for cooling, stunning experts who worked in the field for decades and thought that scale of improvement was impossible. Industrial systems make up one-third of the world's total energy consumption, and there is widespread potential for these AI techniques in the fight against climate change.

Or take health, where costs keep rising unsustainably and millions suffer preventable harm every single day. Each year 285 million people worldwide are affected by sight loss. This number is expected to triple by 2050, with the majority of cases being entirely curable if caught early enough. The eye-scanning technology needed to detect early sight loss is widely available. But what's lacking is enough qualified doctors who can keep pace with analyzing these scans to figure out which patients are most at risk.

We have partnered with world leading Moorfields Eye Hospital in London to develop systems to identify multiple sight-threatening conditions, with absolutely **ground-breaking** results that will be published soon, which will **revolutionise** treatment of eye disease, transforming hundreds of millions of lives, saving billions of dollars for economies, and allowing doctors to focus more of their time and attention on patient care.

Similar advances are happening across multiple medical disciplines. Deep learning algorithms have identified more malignant melanomas and misidentified fewer false positives, than human analysts. Another machine learning study of breast cancer diagnosis showed the potential to reduce unnecessary surgeries by [30%](#).

And this is still only the beginning. AI could dramatically cut the time needed to test the new molecules that could be the foundation of clean energy, from [10-20 years](#) down to 1 or 2. It could [help make farmers more productive](#) by identifying where fertilisers are needed. It could help craft more reliable and low-cost materials for housing, infrastructure and transport. It could help allocate scarce resources—clean water and food—more equitably and help identify new sources of growth.

WHAT'S HOLDING US BACK

So what's holding us back? While it's clear how AI can help, we must be honest about the challenges too.

First, the technology industry remains misaligned with broader society in a number of ways. Our brightest minds aren't focused on solving the most serious problems facing humanity—including the SDGs. Despite the hype, Silicon Valley often gravitates toward the safest and most commercially short term ideas, creating personalised soda drinks when half a billion people don't have access to clean water, or new ways to order food when more than 800m people are malnourished.

We need new incentive structures that encourage technologists to take on society's gravest challenges, and to do so with ethics at the heart.

Second, if we want technologists to be more ambitious in solving big problems, then we also need societies to be more ambitious about how these technologies are governed, directing them toward intended benefits not unintended harms, and ensuring there are new, real-time forms of representative oversight over such powerful systems that affect us all in such profound ways.

We can already see the early consequences of getting this wrong. There is increasing public concern about elements of the technology industry, with growing discomfort about how personal data is used and the impact that algorithms can have on our daily lives. From the spread of facial recognition in drones to biased predictive policing to the use of AI in weaponry, the risk is that individual rights are left by the wayside in the race for technological advantage. This needs to change.

WHAT GOVERNMENTS AND STATE ACTORS CAN DO TO HELP

How can this be achieved?

As you all know better than anyone, a fairer world won't emerge by accident. We need our institutions to guarantee ethical outcomes from those in power, and to preserve human dignity as societies and technologies evolve.

One of the UN's greatest achievements is the adoption of the Universal Declaration of Human Rights. Presenting to the General Assembly, one of its great champions, Eleanor Roosevelt, said: "We stand today at the threshold of a great event both in the life of the United Nations and the life of all mankind." To this day it remains foundational to our sense of a just society, a good life and what it means to be human.

Now, as technology transforms every aspect of our lives, it's time to go one step further. We should reimagine the Universal Declaration of Human Rights to include digital rights—which urgently need affirmation. This would address the many fears around privacy, welfare and fairness, and could go even further with binding principles to ban unethical uses like lethal autonomous weapons, which I believe is an urgent priority. If AI is to serve society, it must be held to the highest ethical standards, and to incorporate the intrinsic rights that have proven to be bedrocks of fair and just societies.

Alongside getting the principles right, we need to get the practicalities right too. I said earlier that the majority of tech investment has flowed into areas that are tangential to social progress. The side-effect is that some of our most societally important infrastructure—the areas that would benefit most from the careful and ethical use of AI—is far from ready for change. Much of the world's health data is still stored on paper, and public information is often kept in inconsistent formats, limiting the ability of AI to help provide solutions. Supporting a new era of open, verifiable data and investing in digitization will lay the necessary groundwork for the breakthroughs we desperately need.

What's at stake is something truly world-changing. And we need it to be. Together, we have the opportunity to put AI - the next phase of the technological revolution and one of the most important of all time - at the service of our most pressing societal needs. If we can create the right structures, ethics and incentives, then we can look forward to incredible

scientific and social progress. And if we do that, we won't meet the SDGs by 2030 - we will far exceed them. Thank you.