Artificial Intelligence

Artificial Intelligence Systems:

Artificial intelligence: Something old, something new We have been here before. In fact, society has repeatedly developed, deployed and adjusted to advanced technologies throughout the course of history, from the ancient discoveries of fire, the wheel and simple tools, to the inventions of the steam engine, the printing press and the Internet itself. Each has led to significant changes in the way we live and work. Each has supported major improvements in quality of life. And each has been accompanied by unexpected, and in some cases unwanted, consequences that must be managed. Most recently, it was the Industrial Revolution that forever changed the course of human history. Powerful machines allowed us to make and move things that were previously inconceivable. It ushered in an unprecedented era of production and progress that transformed every life and every industry in every country.

But it also brought with it major upheaval in labor markets and unacceptable environmental impacts, both of which we continue to address in various ways, from economic and education policies that support the development of new jobs, to global coalitions on climate change.

Myth:

In the interests of grounding our discussion of the ethics of artificial intelligence in facts, we'd like to briefly address a few of the most common myths that are coloring the public discourse. First, while artificial intelligence will almost certainly redefine work in many industries, it will also lead to net new industries, companies and jobs, many of which are difficult to even conceive at this early stage. In fact, study after study, from the most respected economic scholars and research organizations in the world, indicates that technological advances like AI lead to net job growth. Perhaps the Organization for Economic Cooperation and Development (OECD) said it most unambiguously: "Historically, the income generating effects of new technologies have proved more powerful than the labor-displacing effects: technological progress has been accompanied not only by higher output and productivity, but also by higher overall employment." Second, when it comes to the protection of personal information, many of the same concerns that exist in today's computer systems also apply to AI. It is true that AI systems will be more capable of uncovering net new information from personal information, and that new insight will need to be protected with the same level of rigor as

before. But we think that AI will actually help solve this problem, and be better at protecting privacy through advanced techniques like de-identification and privacy-preserving deep learning. And third, the notion of an artificial general intelligence (AGI) AI system with all human abilities including consciousness is an extremely ambitious goal for which our scientific understanding is in a supremely early phase. We believe that much progress and benefit will come from the practical approach of specialized AI – i.e., systems that support tasks in well-defined domains – before AGI can even be contemplated. In the meantime, we are working with our clients, business partners as well as competitors to put in place best practices for safe deployment of a range of AI systems.

Embedded Values:

AI systems should function according to values that are aligned to those of humans, so that they are accepted by our societies and by the environment in which they are intended to function. This is essential not just in autonomous systems, but also in systems based on human-machine collaboration, since value misalignment could preclude or impede effective teamwork. It is not yet clear what values machines should use, and how to embed these values into them. Several ethical theories, defined for humans, are being considered (deontic, consequentialist, virtue, etc.) as well as the implications of their use within a machine, in order to find the best way to define and adapt values from humans to machines. But we do have an idea of how to embed ethical values into AI systems. And there are two main approaches to

this. First, the so-called “top-down” approach recommends coding values in a rigid set.

The work of understanding our responsibilities in developing and deploying safe and ethical AI systems is ongoing. And the development of trust will come through use over time, just as trust was built with all technologies that preceded AI, and all that will follow it. As the technology develops and matures, we encourage other technology companies, as well as experts of many other scientific disciplines, to join us in the study and development of robust, dependable and trustworthy AI applications. Artificial intelligence is not without its risks. But we believe the risks are manageable. And that the far greater risk would be to stifle or otherwise inhibit the development of a technology with the potential to greatly improve the quality of life around the world.