

The Ethics of Artificial Intelligence

ARTIFICIAL INTELLIGENCE OPINION

BY MATTHEW MEE & NEIL JACOBSTEIN 05 NOV 2015

In a recent interview, Matthew Mee, MediaCom's Global Chief Strategy Officer, and Neil Jacobstein, Co-Chair of the Artificial Intelligence and Robotics Track at Singularity University (SU), discussed the societal implications of AI. Below are excerpts

Matt Mee: What is Singularity University, Neil, and what are you working on right now?

Neil Jacobstein: [Singularity University](#) is a benefit corporation that provides educational programs, corporate partnerships, and a startup accelerator for individuals, businesses, institutions, investors, NGOs, and governments. We enable our students and partners to utilize exponential technologies that double in price performance every 18 months. These

technologies, such as AI, robotics, synthetic biology and nanotechnology can help address the challenges of education, energy, environment, food, global health, prosperity, security, water, disaster response, and space. We call these the “Global Grand Challenges,” because we believe they represent the biggest opportunities to positively impact the lives of billions of people. Today, I am writing from our 2015 Graduate Summer Program, which has 80 students from 44 countries. The students live on the [NASA Research Park](#) campus in Silicon Valley, and participate in educational activities in AI, robotics, synthetic biology, nanotechnology, 3D printing, neuroscience, and computing infrastructure. They form startup businesses or non-profits centered on the use of exponential technology to improve the lives of at least a billion people within 10 years. We have an incubator program that helps these new organizations grow, and can bring in a number of SU partner companies to provide real-world feedback and support.

MM: For decades, there has been a huge, shall we say, difference in opinion about the future of AI: the optimistic see it as utopian and the anxious, dystopian. SU clearly leans toward the former. Why do you think AI sparks such polarization among super-smart people, and should I worry?

NJ: The polarization is real, but not helpful. We encourage people to recognize both the extraordinary potential of exponential technology to transform our lives for the better, and the potential to be misused in ways that produce powerful negative consequences. People who work in artificial intelligence tend to focus on constructive applications. However, many technical people are appropriately concerned about the potential risks of exponential technologies like AI. The fact is that all powerful and rapidly-growing technologies have significant downside risks. I think that our best course of action is to move forward while working to: (a) reduce the probability that such risks come to fruition, and (b) increase our ability to respond in very fast machine time if we or our monitoring AIs detect activity considered risky or malicious. This is not a trivial endeavor, and many people in the AI community signed a letter this year to increase R&D funding for: 1) verification – making sure AI programs conform to formal specifications; 2) validation – ensuring that such specifications produce the behavior we want without side effects; 3) security – ensuring that AI systems can’t be tampered with from the inside or outside; and 4) control – building redundant pathways that can be accessed to re-establish control if the systems exhibit

undesired behavior. Worrying about AI doesn't help. Supporting research on improving AI safety and control methods could help. Ultimately, AI will have many positive environmental and health-promoting consequences.

MM: Using AI to manage spiraling complexity is already a reality in our business, and this has been accelerated by the digital transformation of nearly all available marketing communication channels. What's interesting is that there is still a great deal of tension between our industry's legacy view of "creativity" and the increasingly important role of technology and automation. What's your view on AI/exponential technology and creativity?

NJ: AI technology has been used productively in marketing and communications for decades. The advent of powerful machine learning, fueled by improved algorithms, has contributed to the processing and interpretation of massive amounts of consumer data, and the development of advertising placement systems that help move billions of dollars in media every day. Creativity is not substrate dependent. We see creativity in both biological brains and in human thinking amplified by computer tools such as CAD, simulation and machine learning. We also see very encouraging results from computer software utilizing genetic programming algorithms that design new products by simulating evolution. One can get twisted up worrying about the operational definition of creativity, but – if you want a truly new product that meets a large set of design criteria and constraints – you can use genetic programming and many other techniques to get human competitive design results. Machine learning is also routinely used to infer customer sentiment and make marketing recommendations based on massive amounts of data.

MM: AI has already transformed the nature of work in many industries, but this is probably only the tip of the iceberg. What's your view on the shape of business in the future?

NJ: Most people assume that AI and machine learning will transform what is considered blue collar work, but it will also perform tasks associated with white collar work. Augmenting human capability with machine learning will be the most productive strategy for white collar work in the future. Combining human sensibilities, plus machine learning, plus good business practices is a winning combination in any business area, from marketing to sales

and manufacturing.

MM: The scale and pace of technological change and AI capability are things that are quite difficult to grasp using one's usual frame of reference. Gasoline-powered cars, for example, have been around for more than 100 years and have only recently evolved beyond a very basic template. Can you help us understand what's ahead?

NJ: Because we are accustomed to thinking linearly, it's important to understand what exponential growth actually is. For example, taking 30 one-meter steps and counting 1, 2, 3, 4, 5, etc., gets us 30 meters away. But exponential growth requires us to double with each step: 1, 2, 4, 8, 16, 32, etc. Note that 30 doubling steps takes us over a billion meters away, or more than 26 times around the planet. Understanding this non-intuitive difference is vital. As for what to expect from technology, check out the current science and exponential technology capability reports from all reliable sources. Based on the rapid growth in price performance of exponential technology, we can expect to see dramatic technology improvements over the next ten years, with hyperfast adoption of new products and services.

MM: You have talked about the need for all of us to become more "literate" in the mathematical, ecological and ethical issues surrounding AI, but this is hard to do when most people only see bad interpretations of AI in movies. How can people get started?

NJ: We have discussed the need to become mathematically literate; understanding the difference between linear and exponential growth is a minimum requirement for appreciating technological consequences. Decision making benefits from a quantitative approach to data. Society is still largely ignorant of how ecological systems work, although a great deal of information on the topic is available. Good places to start include [Barry Commoner's thoughts](#) on the "Four Laws of Ecology," Garrett Hardin's article, "[Tragedy of the Commons](#)," and recent literature on massive species extinction, including Elizabeth Kolbert's book, [The Sixth Extinction](#).

Finally, ethical literacy requires an understanding of the functional and secular wisdom inherent in Judeo-Christian values. There are also useful insights from Buddhism and Islam.

All religions carry the burden of some superstition and noise, and that should not be confused with ethics. Some technical people think that ethics are arbitrary, but many ethical principles are not. There is nothing arbitrary about considering the destruction of our life support system unethical. Managing technology ethically requires us to deal with long-range consequences systematically and proactively.

MM: Tell us a joke about AI.

NJ: I can't – I'm still waiting for my AI to develop a sense of humor!

ARTIFICIAL INTELLIGENCE OPINION