Acceptance speech

19 June 2025

Michael I. Jordan, awardee in the Information and Communication Technologies category (17th edition)

I am deeply grateful to be recognized by the BBVA Foundation for my work in the area of Information and Communication Technologies.

It's a particular pleasure to receive an award in Europe, as I am as much European as American. I have an Italian passport, and I am living and working in France, as a researcher at Inria Paris. I speak fluent Italian and French and modestly fluent Spanish. I have a deep appreciation of the grand European traditions in mathematics, science, and literature, and in my move to Europe I hope to contribute in some modest way to that tradition.

I work in the field of machine learning, also referred to nowadays as AI. In my view, this field is best understood as an emerging engineering discipline, based on a set of tools that blends algorithmic thinking with inferential thinking. Machine learning uses data analysis at massive scale to augment classical methodology in science and technology. Machine learning is also noteworthy for being the first engineering discipline in which a major focus is data from and about humans, including human preferences, values, and decisions, and for aiming to provide mechanisms that can act in the human world in real time. It thus has vast, and currently poorly understood, implications for individuals, societies, and cultures.

I didn't plan to work in machine learning, and like many academics I arrived at the problems that interested me idiosyncratically. My original degree was in psychology, and my PhD was in cognitive science. My first faculty position was in neuroscience, at MIT, and I eventually arrived at a joint position in the two departments of Statistics and Computer Science, at UC Berkeley. During all of this wandering, I developed an interest in bringing algorithmic and statistical thinking into problem areas that hitherto were the realm of philosophers, including algorithmic theories of reasoning by individuals and groups, how individuals and groups cope with the vast uncertainty in which we live our lives, how order can arise out of disorder, and the algorithmic nature of learning. Problems such as these were not to be found solely in one field, and along the way I have engaged with ideas and people in a wide range of fields, including

statistical physics, control theory, operations research, linguistics, molecular biology, mathematics, and, most recently, economics.

Given my proximity to Silicon Valley, my thinking was also shaped by developments in information technology. I have always viewed information technology as principally a way to augment human intelligence and creativity. In the early days I felt that Silicon Valley was aligned with this perspective. I no longer believe that, and part of the reason for my coming to Europe is to flee the current Silicon Valley zeitgeist. A Frankenstein mindset has come to dominate Silicon Valley, where the goal is to imitate and somehow surpass human intelligence, with little thought to other goals that might be better for developing human-centered technology.

I believe that such goals can be usefully framed in terms of a collectivist perspective on AI that incorporates ideas from economics and other social sciences. Economics is about values, interactions, scarcity, tradeoffs, and the relationships between individuals and large-scale systems – all issues that arise when AI platforms are deployed. Also, issues such as privacy, fairness, bias and ownership that arise in such deployments are in part economic in nature, involving tradeoffs. Unfortunately they are often reduced to black-and-white legal and algorithmic concepts. Another concern is the need to allow creators to profit from their creative output when that output is aggregated into an AI system. This is in part a problem of market design, in markets in which data itself is a valued good, a novel concept for economics. Finally, the current dialog about AI often involves going straight to proposed regulatory mechanisms that aim at specific algorithmic behavior. We should instead be analyzing AI systems in their societal context, aiming to understand their equilibria, their out-of-equilibrium behavior, and their externalities. Regulation of AI systems will be needed, but effective regulation is best exerted at the level of equilibria, not at the level of algorithms.

Finally, the dialog about AI needs to involve the general public in a way that speaks to their daily lives. The current dialog, as conducted by technologists, journalists, and politicians, is dominated by hype, hysteria, and regulatory posturing. All humans engage with economic and social concepts on a daily basis, quite apart from academic theories, and the same can be said for cognitive concepts and reasoning under uncertainty. The general public will be able to contribute if AI is discussed in these terms. It will not be able to contribute if the dialog continues to focus on the purported imminent arrival of a superintelligence that will somehow change everything. That is science fiction, and science fiction is a childlike way to think about the future of humanity.