



Artificial Intelligence and the future of work and living

Navigating between utopian dreams and dystopian fears.

Publication note

Originally published in [Cubed](#). Author's copy on [Substack](#).

Navigating between utopian dreams and dystopian fears

Note: A [previous version of this article](#) was published with [Cubed](#).

Today, most people can easily see that we are in the midst of a technological revolution driven by rapid advances in Artificial Intelligence (AI), or more specifically Machine Learning (ML). Unfortunately, most people, even including some AI/ML experts, don't seem to fully appreciate how much this revolution will change society. This rapid transformation will have a profound effect on society and it is occurring at an ever-accelerating pace and demands immediate action if we hope to steer toward a positive trajectory.

The world of AI/ML can be roughly categorized into specialized and general-purpose models. Specialized models excel in specific tasks, such as identifying mold on raspberries or detecting fraudulent credit-card transactions, and they often operate directly under human control. Conversely, general-purpose models attempt to emulate the versatility of a human being and they can tackle a broad range of activities, from identifying mold across different foods to creating recipes. Because general-purpose models can process and respond reasonably to natural language, it is increasingly the case that using one simply requires asking it to perform a task in the same way one might ask a person.

Despite the rapid bombardment of increasingly impressive results from these general-purpose models, many persist in believing that their own professions will somehow remain beyond the capabilities of general-purpose AI systems. Unfortunately, these beliefs appear to be supported by little more than wishful thinking. Appeals to human creativity and empathy are common, but they fall short both because most jobs don't require much creativity or empathy, and because general-purpose AI systems increasingly show the capacity to emulate human creativity and empathy.

My critical opinion of these unsupported beliefs about AI limitations is based predominantly on my 25 years of experience teaching Computer Science at UC Berkeley and more than 30 years of research experience with both traditional and AI/ML systems.

One specific anecdotal example occurred a little more than a year ago, during the Spring 2023 semester, when I tested GPT-4 — a general-purpose language model — on questions from the take-home final exam of my upper

division undergraduate course. Surprisingly, GPT-4 answered the questions correctly. I then tried to modify the questions, making them more complicated until GPT-4 was no longer able to answer them correctly.

Even after modifying the questions to challenge the model, it continued to perform well. I added complications and rewrote the problems in ways that typically challenge students, yet the model kept succeeding. By the point where I had made the questions convoluted enough that GPT-4 failed them, I found that my grad-student TAs also struggled with them and the contrived difficulty had made the questions themselves pedagogically questionable. GPT-4 has now been superseded by the next generation of AI models that are even more capable. Colleagues in UC Berkeley's math department have reported that GPT-5 is now able to solve the problems on their PhD qualifying exams. The next batch of AI models will likely be even more capable.

I find this situation very concerning because it suggests that students who are starting undergrad degrees today may find themselves obsolete upon graduation. Over the last year, a growing list of companies have been laying off tens of thousands of skilled workers, citing AI automation as the reason. Unlike previous technological revolutions, the new jobs being created are not more desirable high-level ones. They are lower-paid jobs of essentially babysitting the AIs, and there are fewer of them.

We are at a pivotal point in time. I believe that at some point in the next five-to-ten years we will reach a state where *every job* can be performed by an AI system. While inevitable, we can still decide what this future means to us humans.

One direction appears utopian: Individuals could enjoy the benefits of fully automated production, with work becoming optional and people free to pursue their passions rather than laboring for financial support.

The alternative presents a dystopian nightmare: Machines still do all the work, but only a relative few who own those machines reap massive rewards, while the majority struggle to find any work and survive at subsistence levels.

The direction we take depends on our willingness now to acknowledge the issues, plan for the impending future, and collaborate on implementing solutions. If we instead remain focused on political bickering, pointless animosity, and delusional denial, then most of us will inevitably face dire consequences.

The next few years will be a test of our societal ability to make thoughtful and intelligent decisions about our shared future. Will we pass this test?

About Me: James F. O'Brien is a Professor of Computer Science at the University of California, Berkeley. His research interests include computer graphics and animation, simulation of physical systems, human perception, rendering, artificial intelligence, computer vision, virtual reality, digital privacy and security, and image and video forensics.

If you would like to read my other articles, consider subscribing. You can also find me on LinkedIn, Instagram, and at UC Berkeley.

Disclaimer: Any opinions expressed in this article are only those of the author as a private individual. Nothing in this article should be interpreted as a statement made in relation to the author's professional position with any institution.

This article and all embedded images are Copyright 2024 by the author. This article was written by a human, and both an LLM and other humans may have been used for proofreading and editorial suggestions.