Exploring 6 AI Myths

Google

Myth: AI, machine learning, and deep learning are all the same thing

While artificial intelligence (AI) is a convenient and commonplace term, it has no widely agreed-upon technical definition. One helpful way to think about AI is as the science of making things smart. Much of the recent progress we've seen in AI is based on machine learning (ML), a subfield of AI where computers learn and recognize patterns from examples, rather than being programmed with specific rules. There are many different ML techniques, but deep learning is a particularly popular one right now. Deep learning is based on neural network technology, an algorithm whose architecture is inspired by the human brain and can learn to recognize pretty complex patterns, such as what "hugs" are or what a "party" looks like.

Myth: All AI systems are "black boxes," far less explainable than non-AI techniques

As with human-based processes or traditional software, some AI systems are quite simple and easy to explain, while others are far more complex. Explainability is a rich area of research that is producing new methods and tools that give us insight into why an AI system behaves in a certain way (i.e., which parts of an image are triggering a medical diagnosis tool to identify a disease). And we shouldn't miss out on the opportunity to use AI to improve transparency of decision-making, even if it's not fully explainable. Explanations for human decisions, for example, may not accurately reflect their influencing factors or unconscious biases. In fact, even if every individual decision made by some AI systems cannot be fully explained, we may be able to understand how they make decisions in general better than we understand how humans make similar decisions.

Myth: AI systems are only as good as the data they train on

There are four ingredients needed for AI innovation: data, algorithms, hardware, and human talent. While data is an important part of training a model, no real-world dataset will be perfect. It is possible to address shortcomings in training data—such as data scarcity, low quality data, and unbalanced data—through techniques like careful problem formulation, targeted sampling, synthetic data, or building constraints into models.

Myth: AI systems are inherently unfair

Unfair bias in AI is the result of human decisions about how an AI application is designed, tested, and deployed. There are many instances where human decision-making—from employment decisions to credit allocation—results in unfair outcomes for vulnerable groups, and if AI is trained to mimic the behavior of those human decision-makers it can also reflect those biases. Designing systems to address these biases is challenging, and requires careful consideration not just of the technology, but of the societal context in which it will be deployed. But well-designed, thoroughly vetted AI systems can limit unfair bias, and may even help us to identify and combat bias in human decision-making.

Myth: AI will make human labor obsolete

Technology breakthroughs—from the cotton gin to the personal computer—have long been met by the fear of mass unemployment. However, in the long run new technologies have increased productivity, created new jobs and new industries, and raised standards of living. Currently, AI systems excel at narrow tasks, while occupations consist of many interrelated tasks. AI will undoubtedly cause jobs to shift, as transformative technologies always have, but AI is also enhancing worker productivity and creating new types of jobs. The greater risk is that shifting jobs may increase income inequality and create challenges for workers whose jobs are displaced or require new skills. This challenge is larger than any one organization can solve, so we need to work together on policies and programs that both equip people for new jobs and allow for stable careers in a shifting landscape.

Myth: AI is approaching human intelligence

While AI systems are nearing or outperforming human beings at increasingly complex tasks like generating musical melodies or playing the game of Go, they remain narrow and brittle, and lack true agency or creativity. It's not that AI has chosen to create a melody, or understands that sounds make music. Rather, researchers have built tools to recognize patterns in melodies and use them to project similar patterns based on their guidance. AI systems that generate melodies cannot currently be used to generate realistic speech, much less paint a picture or play chess. Techniques like transfer learning are bringing us closer to AI systems that can apply their learning to multiple problems, but machines with human intelligence remain a long way off.