



Introducing OpenAI

by Greg Brockman, Ilya Sutskever, and the OpenAI team

December 11, 2015

OpenAI is a non-profit artificial intelligence research company. Our goal is to advance digital intelligence in the way that is most likely to benefit humanity as a whole, unconstrained by a need to generate financial return.

Since our research is free from financial obligations, we can better focus on a positive human impact. We believe AI should be an extension of individual human wills and, in the spirit of liberty, as broadly and evenly distributed as is possible safely.

The outcome of this venture is uncertain and the work is difficult, but we believe the goal and the structure are right. We hope this is what matters most to the best in the field.

Background

Artificial intelligence has always been a surprising field. In the early days, people thought that solving certain tasks (such as chess) would lead us to discover human-level intelligence algorithms. However, the solution to each task turned out to be much less general than people were hoping (such as doing a search over a huge number of moves).

The past few years have held another flavor of surprise. An AI technique explored for decades, deep learning, started achieving state-of-the-art results in a wide variety of problem domains. In deep learning, rather than hand-code a new algorithm for each problem, you design architectures that can twist themselves into a wide range of algorithms based on the data you feed them.

This approach has yielded outstanding results on pattern recognition problems, such as recognizing objects in images, machine translation, and speech recognition. But we've also started to see what it might be like for computers to be [creative](#), to [dream](#), and to [experience the world](#).

Looking forward

AI systems today have impressive but narrow capabilities. It seems that we'll keep whittling away at their constraints, and in the extreme case they will reach human performance on virtually every intellectual task. It's hard to fathom how much human-level AI could benefit society, and it's equally hard to imagine how much it could damage society if built or used incorrectly.

OpenAI

Because of AI's surprising history, it's hard to predict when human-level AI might come within reach. When it does, it'll be important to have a leading research institution which can prioritize a good outcome for all over its own self-interest.

We're hoping to grow OpenAI into such an institution. As a non-profit, our aim is to build value for everyone rather than shareholders. Researchers will be strongly encouraged to publish their work, whether as papers, blog posts, or code, and our patents (if any) will be shared with the world. We'll freely collaborate with others across many institutions and expect to work with companies to research and deploy new technologies.

OpenAI's research director is [Ilya Sutskever](#), one of the world experts in machine learning. Our CTO is [Greg Brockman](#), formerly the CTO of Stripe. The group's other founding members are world-class research engineers and scientists: [Trevor Blackwell](#), [Vicki Cheung](#), [Andrej Karpathy](#), [Durk Kingma](#), [John Schulman](#), [Pamela Vagata](#), and [Wojciech Zaremba](#). [Pieter Abbeel](#), [Yoshua Bengio](#), [Alan Kay](#), [Sergey Levine](#), and [Vishal Sikka](#) are advisors to the group. OpenAI's co-chairs are Sam Altman and Elon Musk.

Sam, Greg, Elon, Reid Hoffman, Jessica Livingston, Peter Thiel, Amazon Web Services (AWS), Infosys, and [YC Research](#) are donating to support OpenAI. In total, these funders have committed \$1 billion, although we expect to only spend a tiny fraction of this in the next few years.

You can follow us on Twitter at [@open_ai](#) or email us at info@openai.com.