

The Future of AI: What if We Succeed?

Stuart Russell

Thanks to AIJ for funding this panel

Why are we doing AI?

- ❖ To create intelligent systems

Why are we doing AI?

- ❖ To create intelligent systems
- ❖ The more intelligent, the better

Why are we doing AI?

- ❖ To create intelligent systems
- ❖ The more intelligent, the better
- ❖ We believe we can succeed
 - ❖ Limited only by ingenuity and physics
 - ❖ Progress is accelerating
 - ❖ Industrial and military AI race under way
 - ❖ We may soon have a roadmap

What if we do succeed?

- ❖ *“The first ultraintelligent machine is the last invention that man need ever make.”* I. J. Good, 1965
- ❖ Success would be the biggest event in human history

What if we do succeed?

- ❖ *“The first ultraintelligent machine is the last invention that man need ever make.”* I. J. Good, 1965
- ❖ Success would be the biggest event in human history...
- ❖ ... and perhaps the last event in human history

So, if that matters.....

- ❖ Along what paths will AI evolve?
- ❖ What is the (plausibly reachable) best case? Worst case?
- ❖ Can we affect the future of AI?
 - ❖ Technical or societal solutions?
 - ❖ What should we do now?

Panelists

- ❖ Henry Kautz, U. of Rochester
 - ❖ President of AAAI, 2010-12
- ❖ Joanna Bryson, U. of Bath
 - ❖ Co-author, EPSRC Principles of Robotics
- ❖ Anders Sandberg, U. of Oxford
 - ❖ Fellow, Future of Humanity Institute
- ❖ Sebastian Thrun, Udacity/Google/Stanford
 - ❖ Developer of Google driverless car, et al.