



The Future of (Artificial) Intelligence

Fantasy Basel, 14 May 2015

Robin Li, Bill Gates, Elon Musk on AI

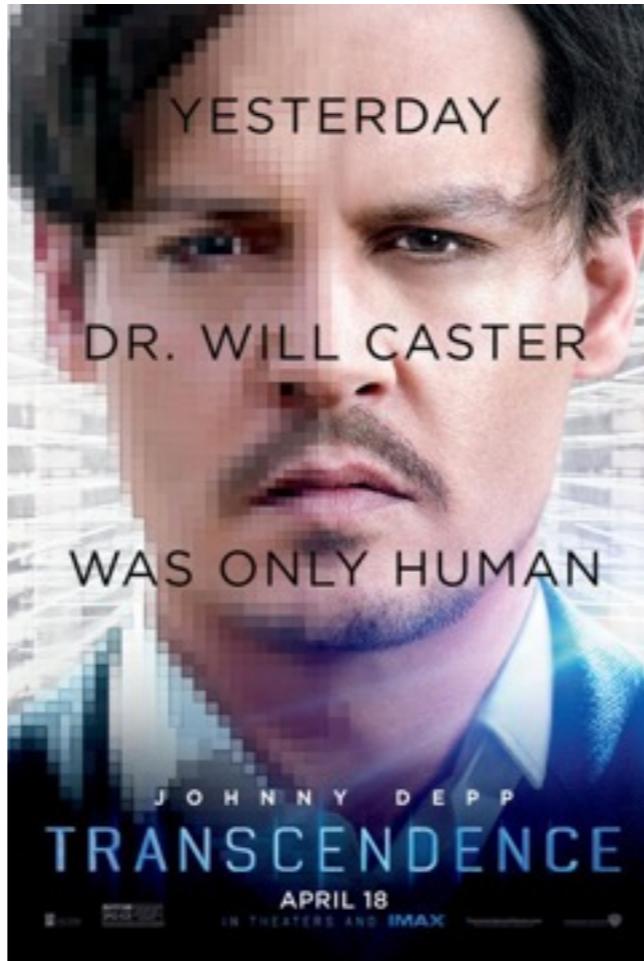


I understand, Elon, that recently you said
Artificial Intelligence advances are like summoning the demon.

The Huge Challenge of Artificial Intelligence
www.youtube.com/watch?v=vHzJ_AJ34uQ

The Future of AI | 2
at Fantasy Basel

More Money for Entertainment



... than ensuring a good outcome!



Introduction

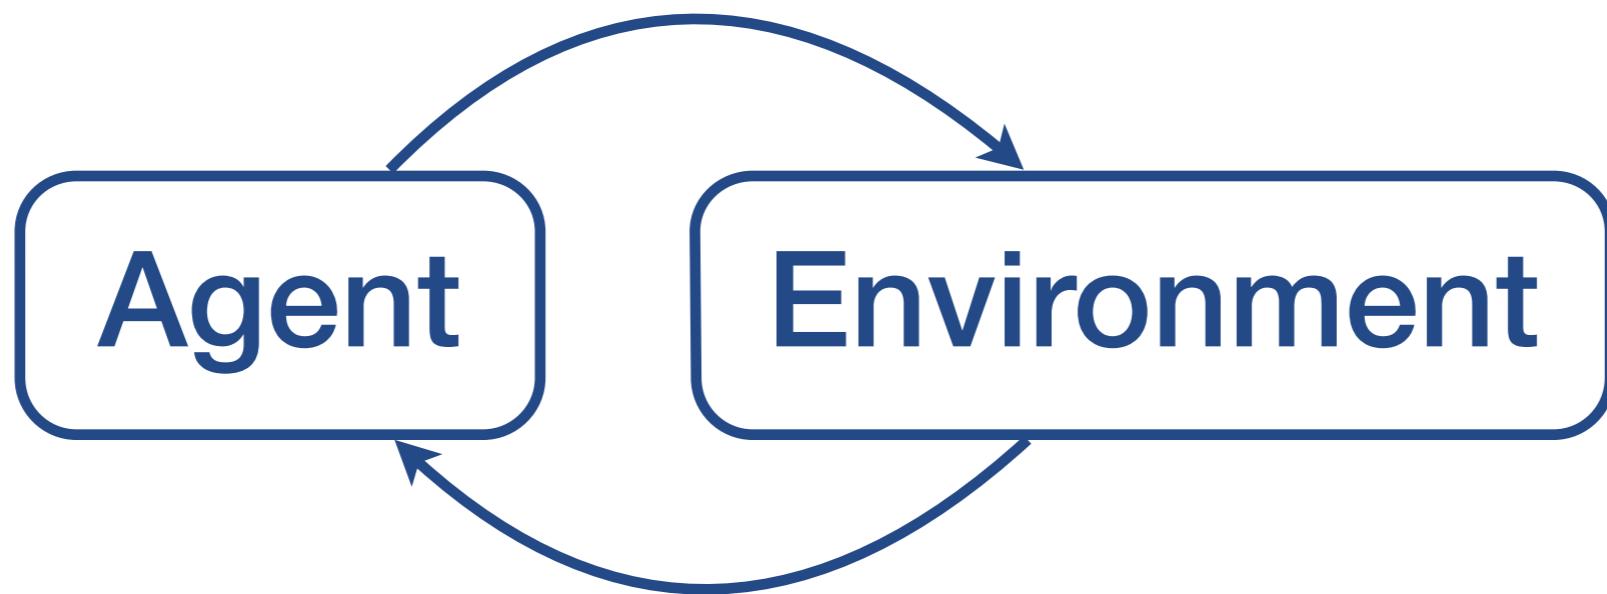
What are we talking about?

Intelligence

«Intelligence measures an agent's ability to achieve its goals in a wide range of unknown environments.»

$$\text{Intelligence} = \frac{\text{Optimization Power}}{\text{Used Resources}}$$

Ingredients



Learn, predict, rate and plan!

Intelligence is a Big Deal



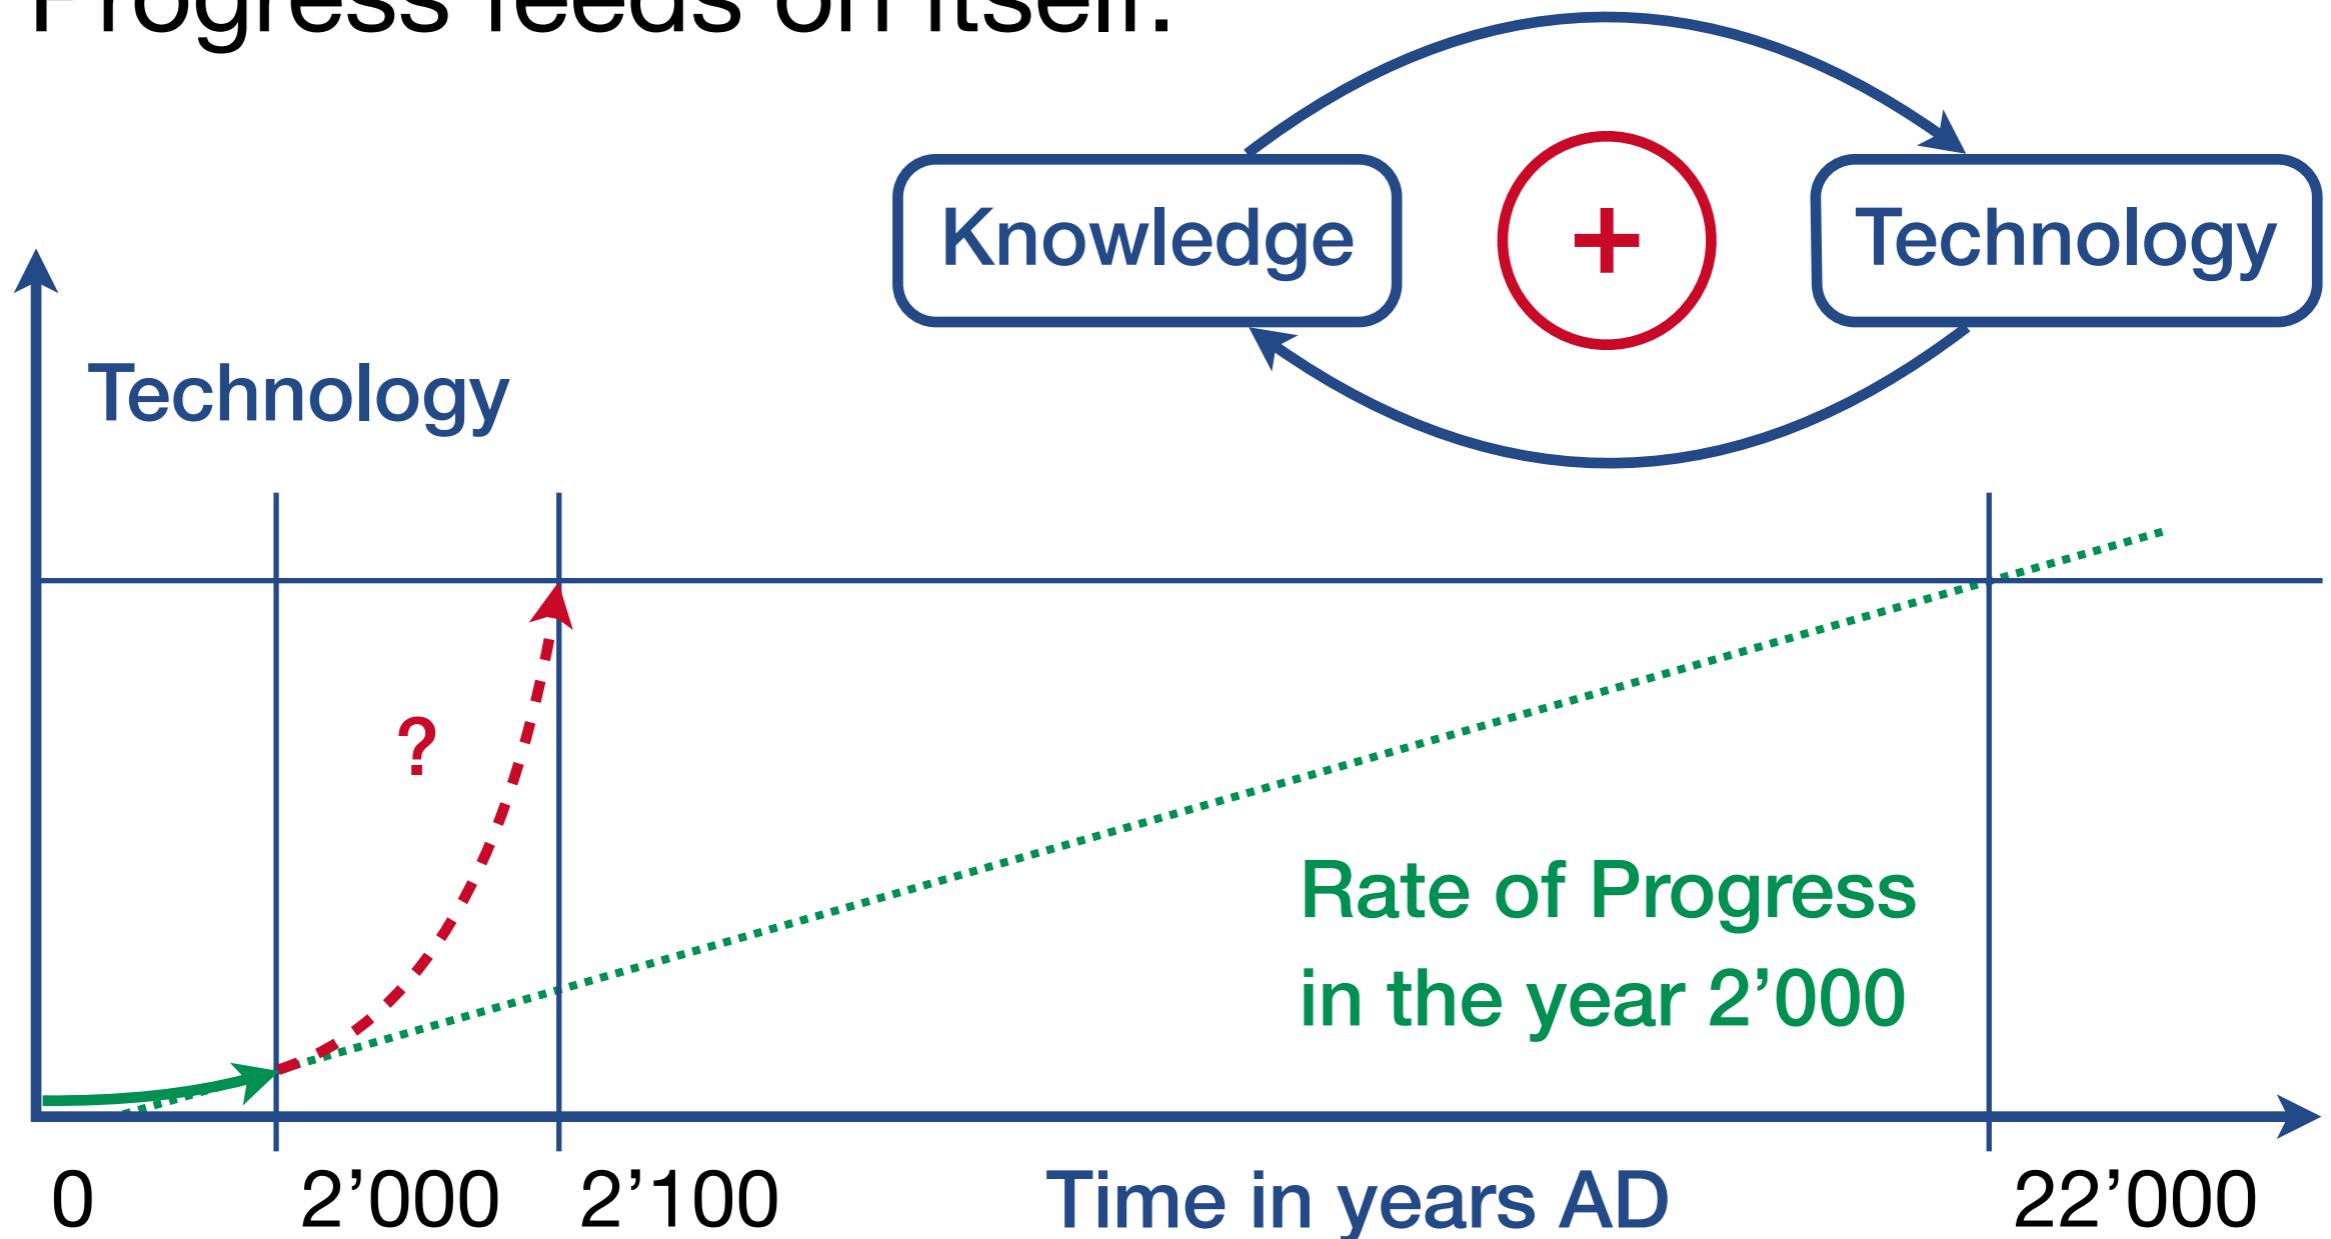
6 million years ago, 96% common DNA

Fast-Evolving Human DNA Leads to Bigger-Brained Mice
[phenomena.nationalgeographic.com/2015/02/19/...](http://phenomena.nationalgeographic.com/2015/02/19/)

The Future of AI
Introduction

Accelerating Change

Progress feeds on itself:



Technology = (Neutral) Lever

Greed and laziness drive us to

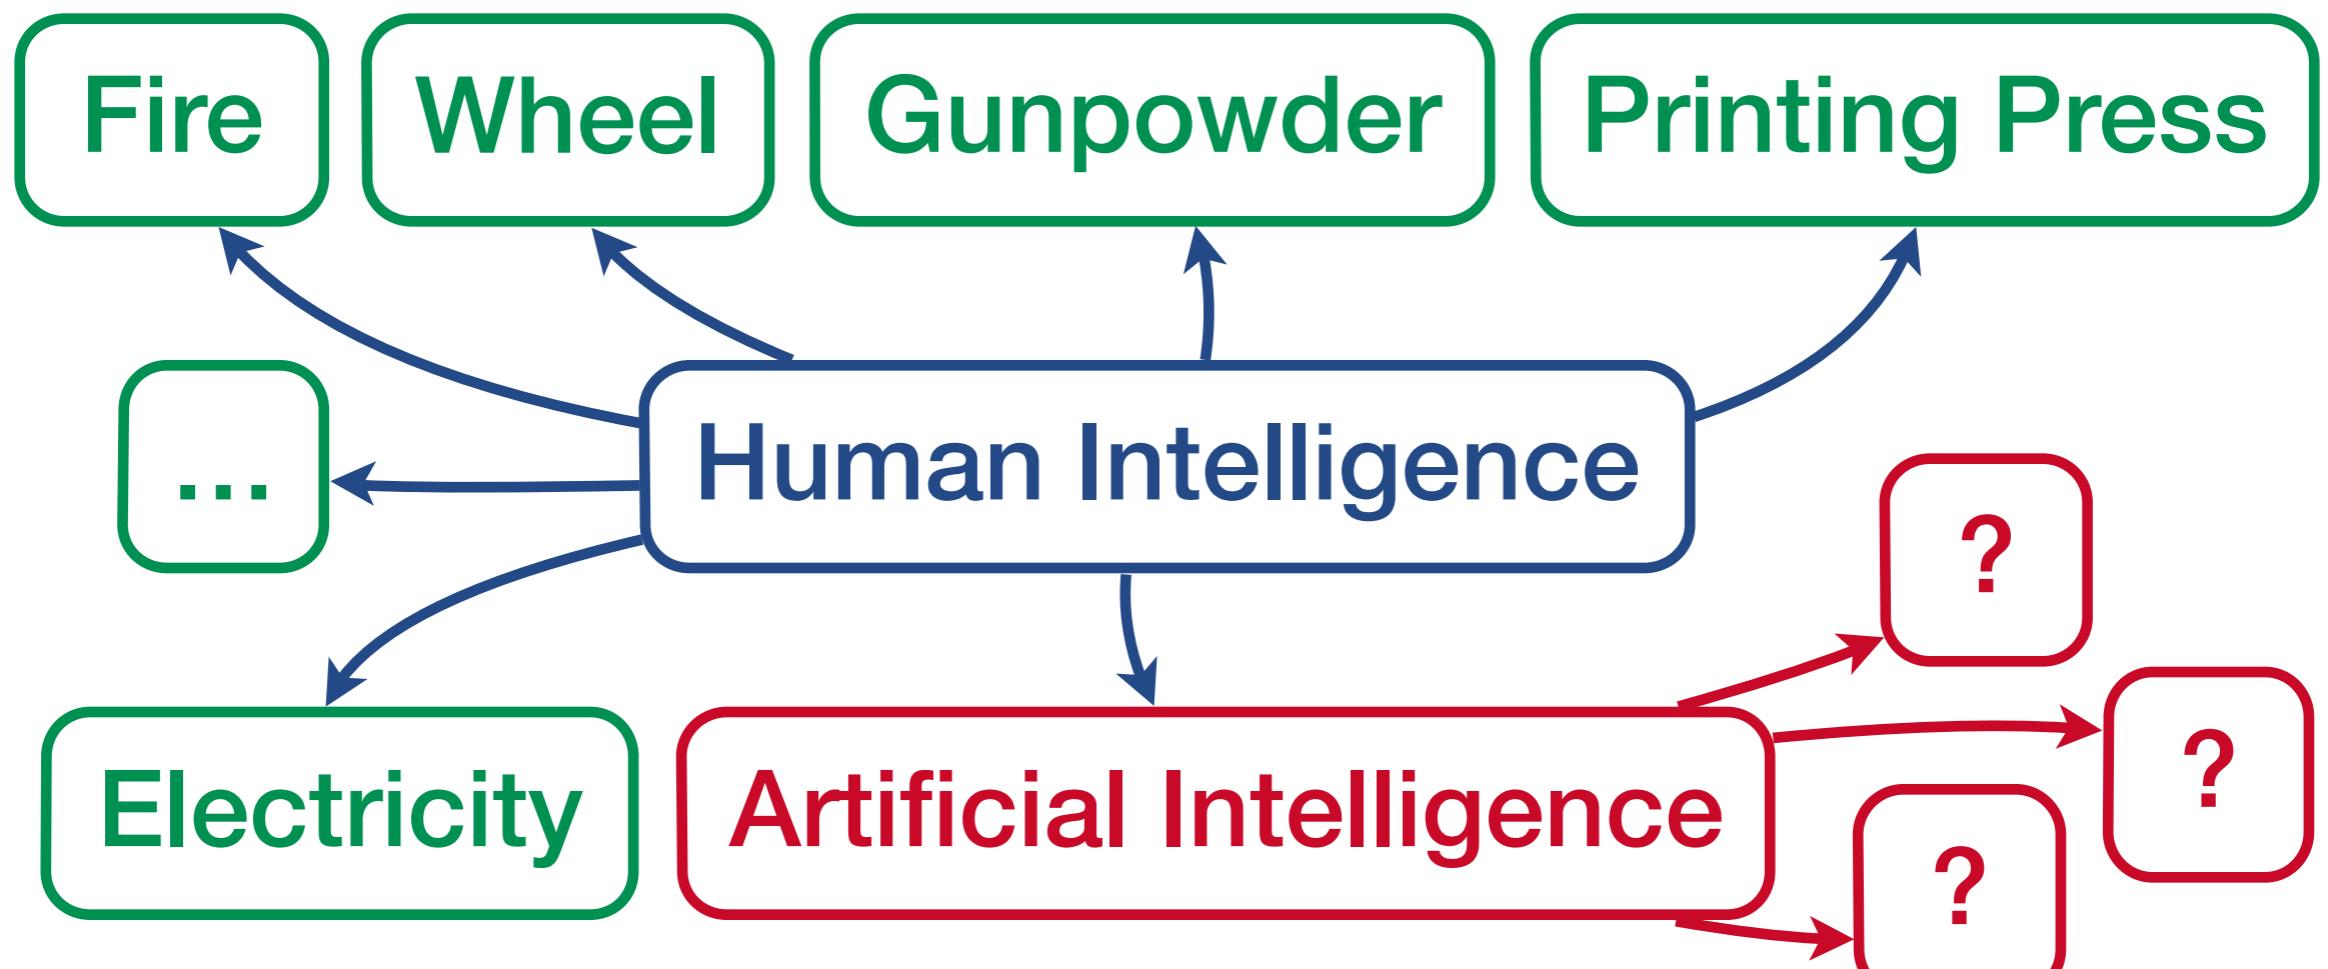
- increase our productivity
- make our lives easier

... which is fine except:



**Our technological progress far
outperforms our moral progress!**

Artificial Intelligence



Intelligence is a technology like no other!

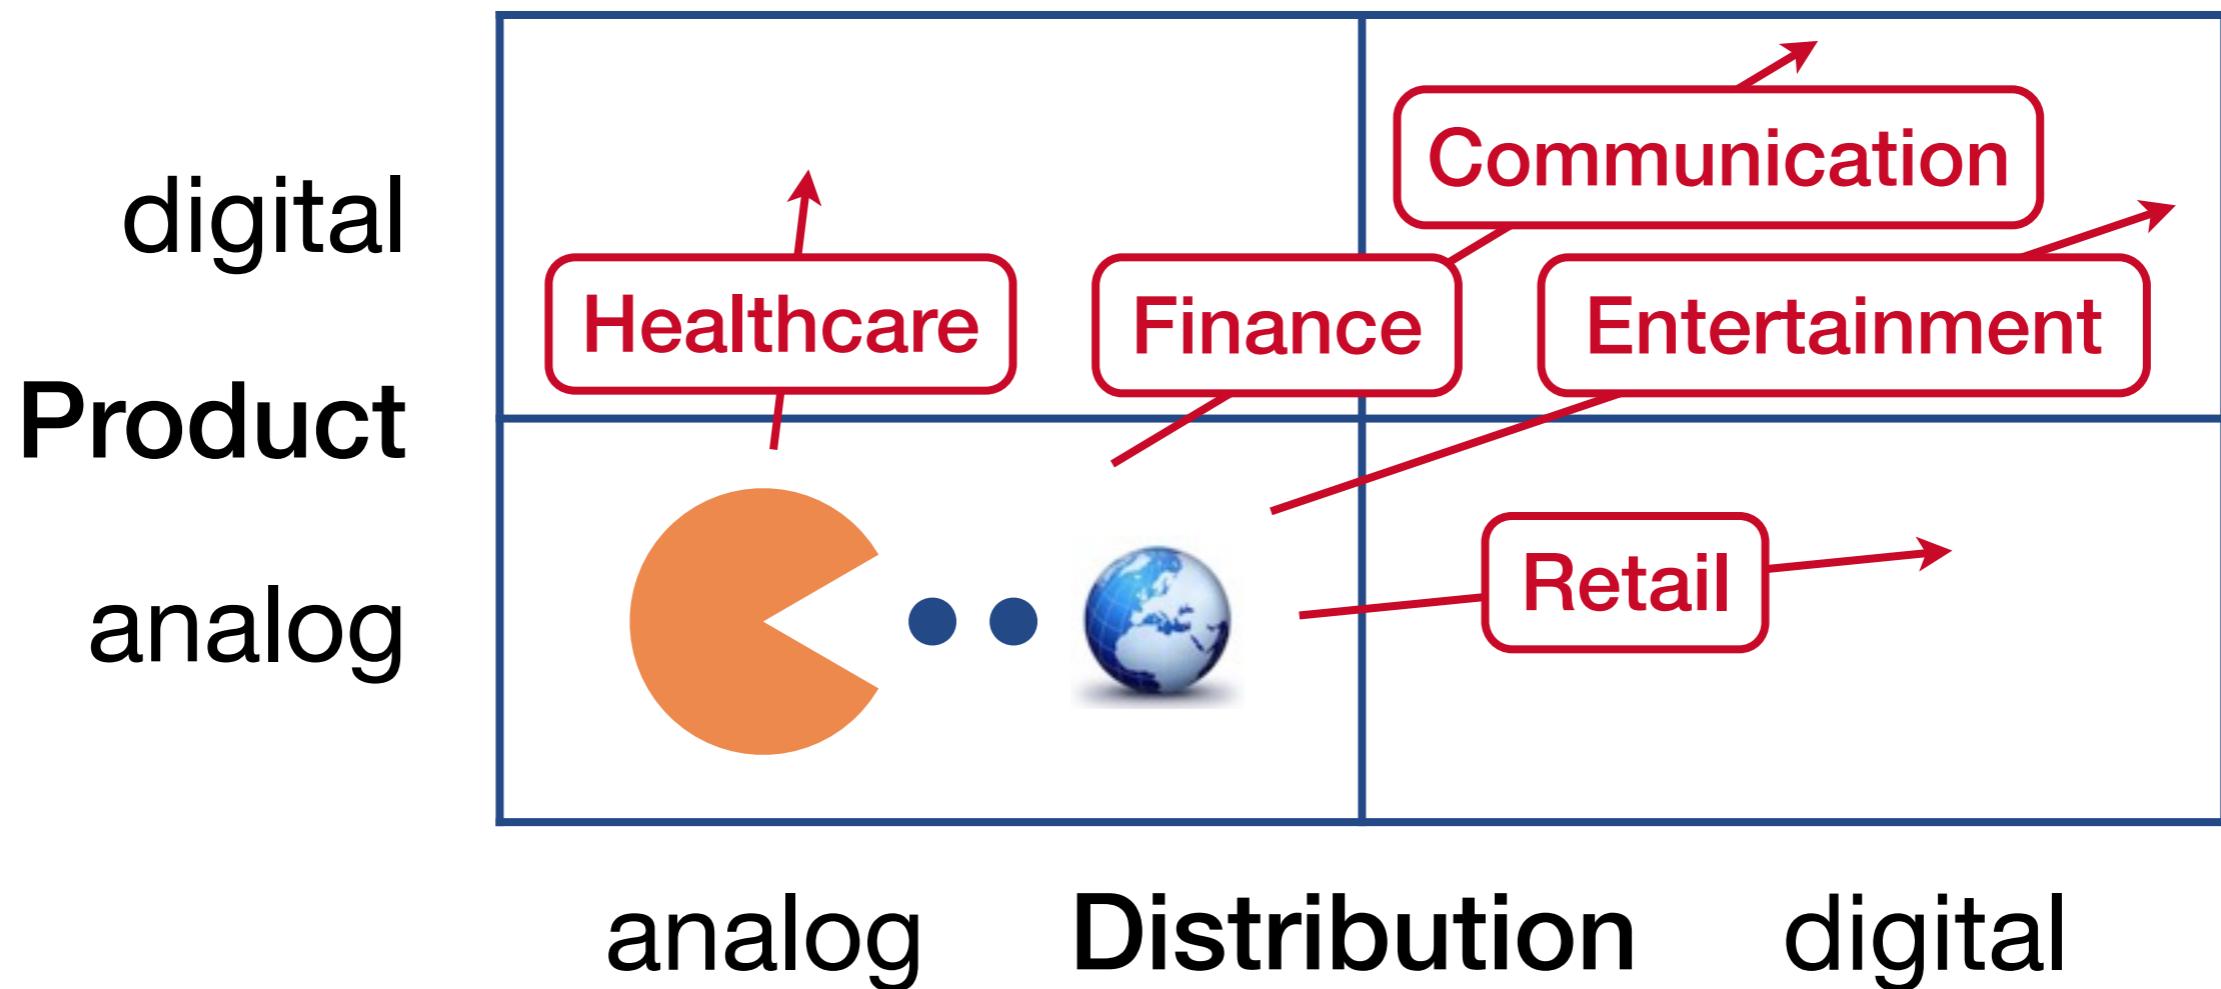


Current Trends

Where are we heading to?

Software is eating the world

... because it is more productive!



Describes without errors



A person riding a motorcycle on a dirt road.



A group of young people playing a game of frisbee.



A herd of elephants walking across a dry grass field.

Describes with minor errors



Two dogs play in the grass.



Two hockey players are fighting over the puck.



A close up of a cat laying on a couch.

Somewhat related to the image



A skateboarder does a trick on a ramp.



A little girl in a pink hat is blowing bubbles.



A red motorcycle parked on the side of the road.

Unrelated to the image



A dog is jumping to catch a frisbee.



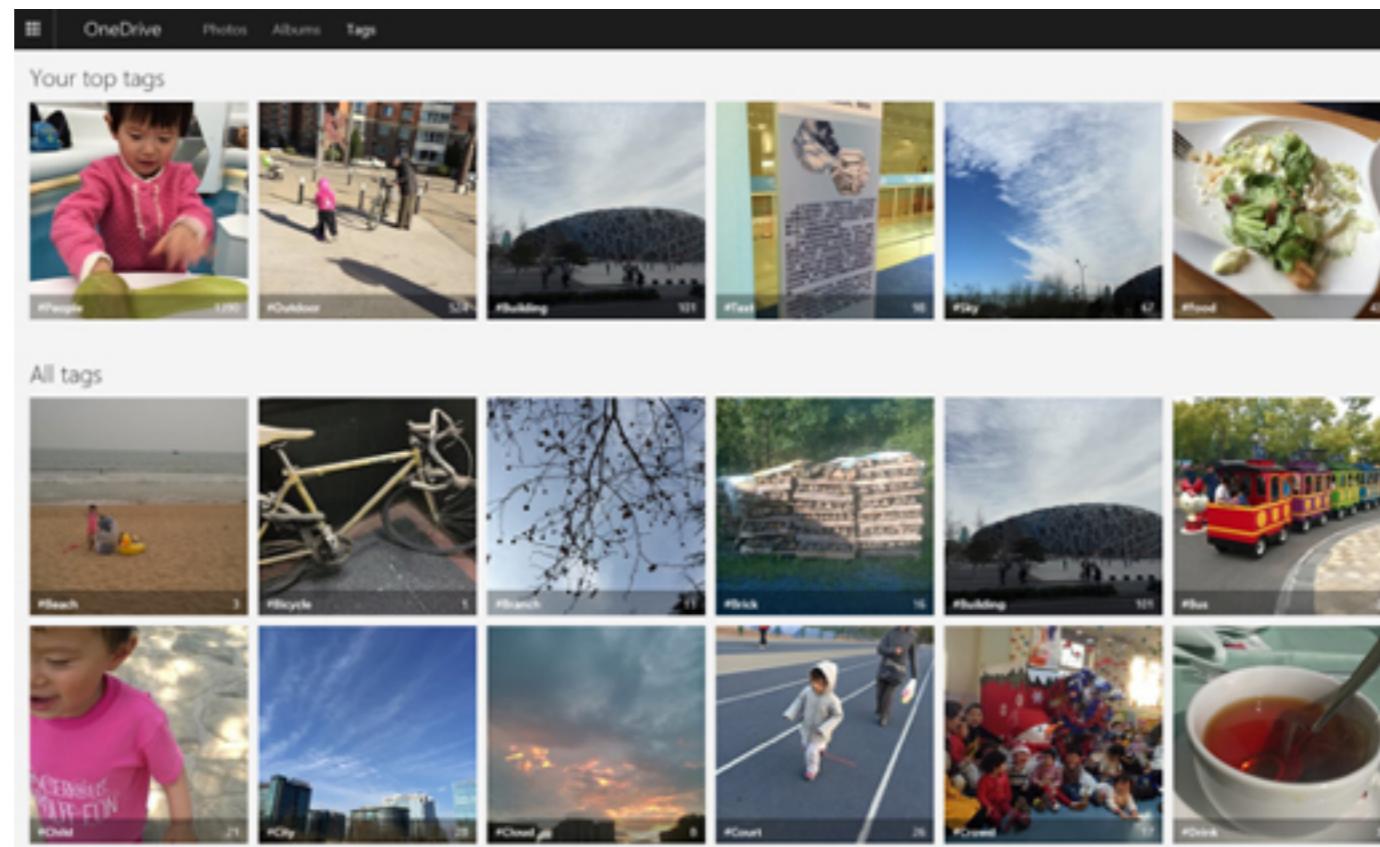
A refrigerator filled with lots of food and drinks.



A yellow school bus parked in a parking lot.

Superhuman Image Recognition

- With convolutional neural networks
- 1.2 m training images, ~ 30 layers





Jeremy Howard
go.ted.com/bbZC

The Future of AI
Current Trends

Machine Intelligence LANDSCAPE

CORE TECHNOLOGIES

ARTIFICIAL INTELLIGENCE



DEEP LEARNING



MACHINE LEARNING



NLP PLATFORMS



PREDICTIVE APIs

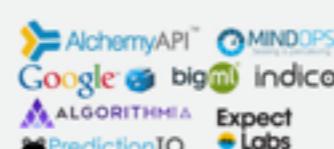


IMAGE RECOGNITION



SPEECH RECOGNITION



RETHINKING ENTERPRISE

SALES



SECURITY / AUTHENTICATION



FRAUD DETECTION



HR / RECRUITING



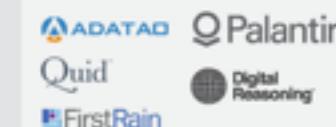
MARKETING



PERSONAL ASSISTANT



INTELLIGENCE TOOLS



RETHINKING INDUSTRIES

ADTECH



AGRICULTURE



EDUCATION



FINANCE



LEGAL



MANUFACTURING



MEDICAL



OIL AND GAS



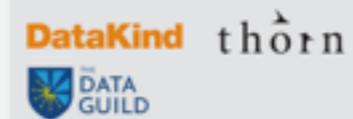
MEDIA / CONTENT



CONSUMER FINANCE



PHILANTHROPIES



AUTOMOTIVE



DIAGNOSTICS



RETAIL

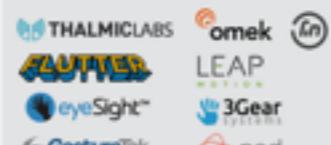


RETHINKING HUMANS / HCI

AUGMENTED REALITY



GESTURAL COMPUTING



ROBOTICS



EMOTIONAL RECOGNITION



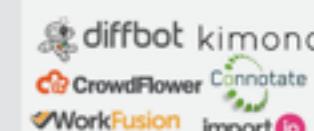
HARDWARE



DATA PREP



DATA COLLECTION



Automation brings abundance

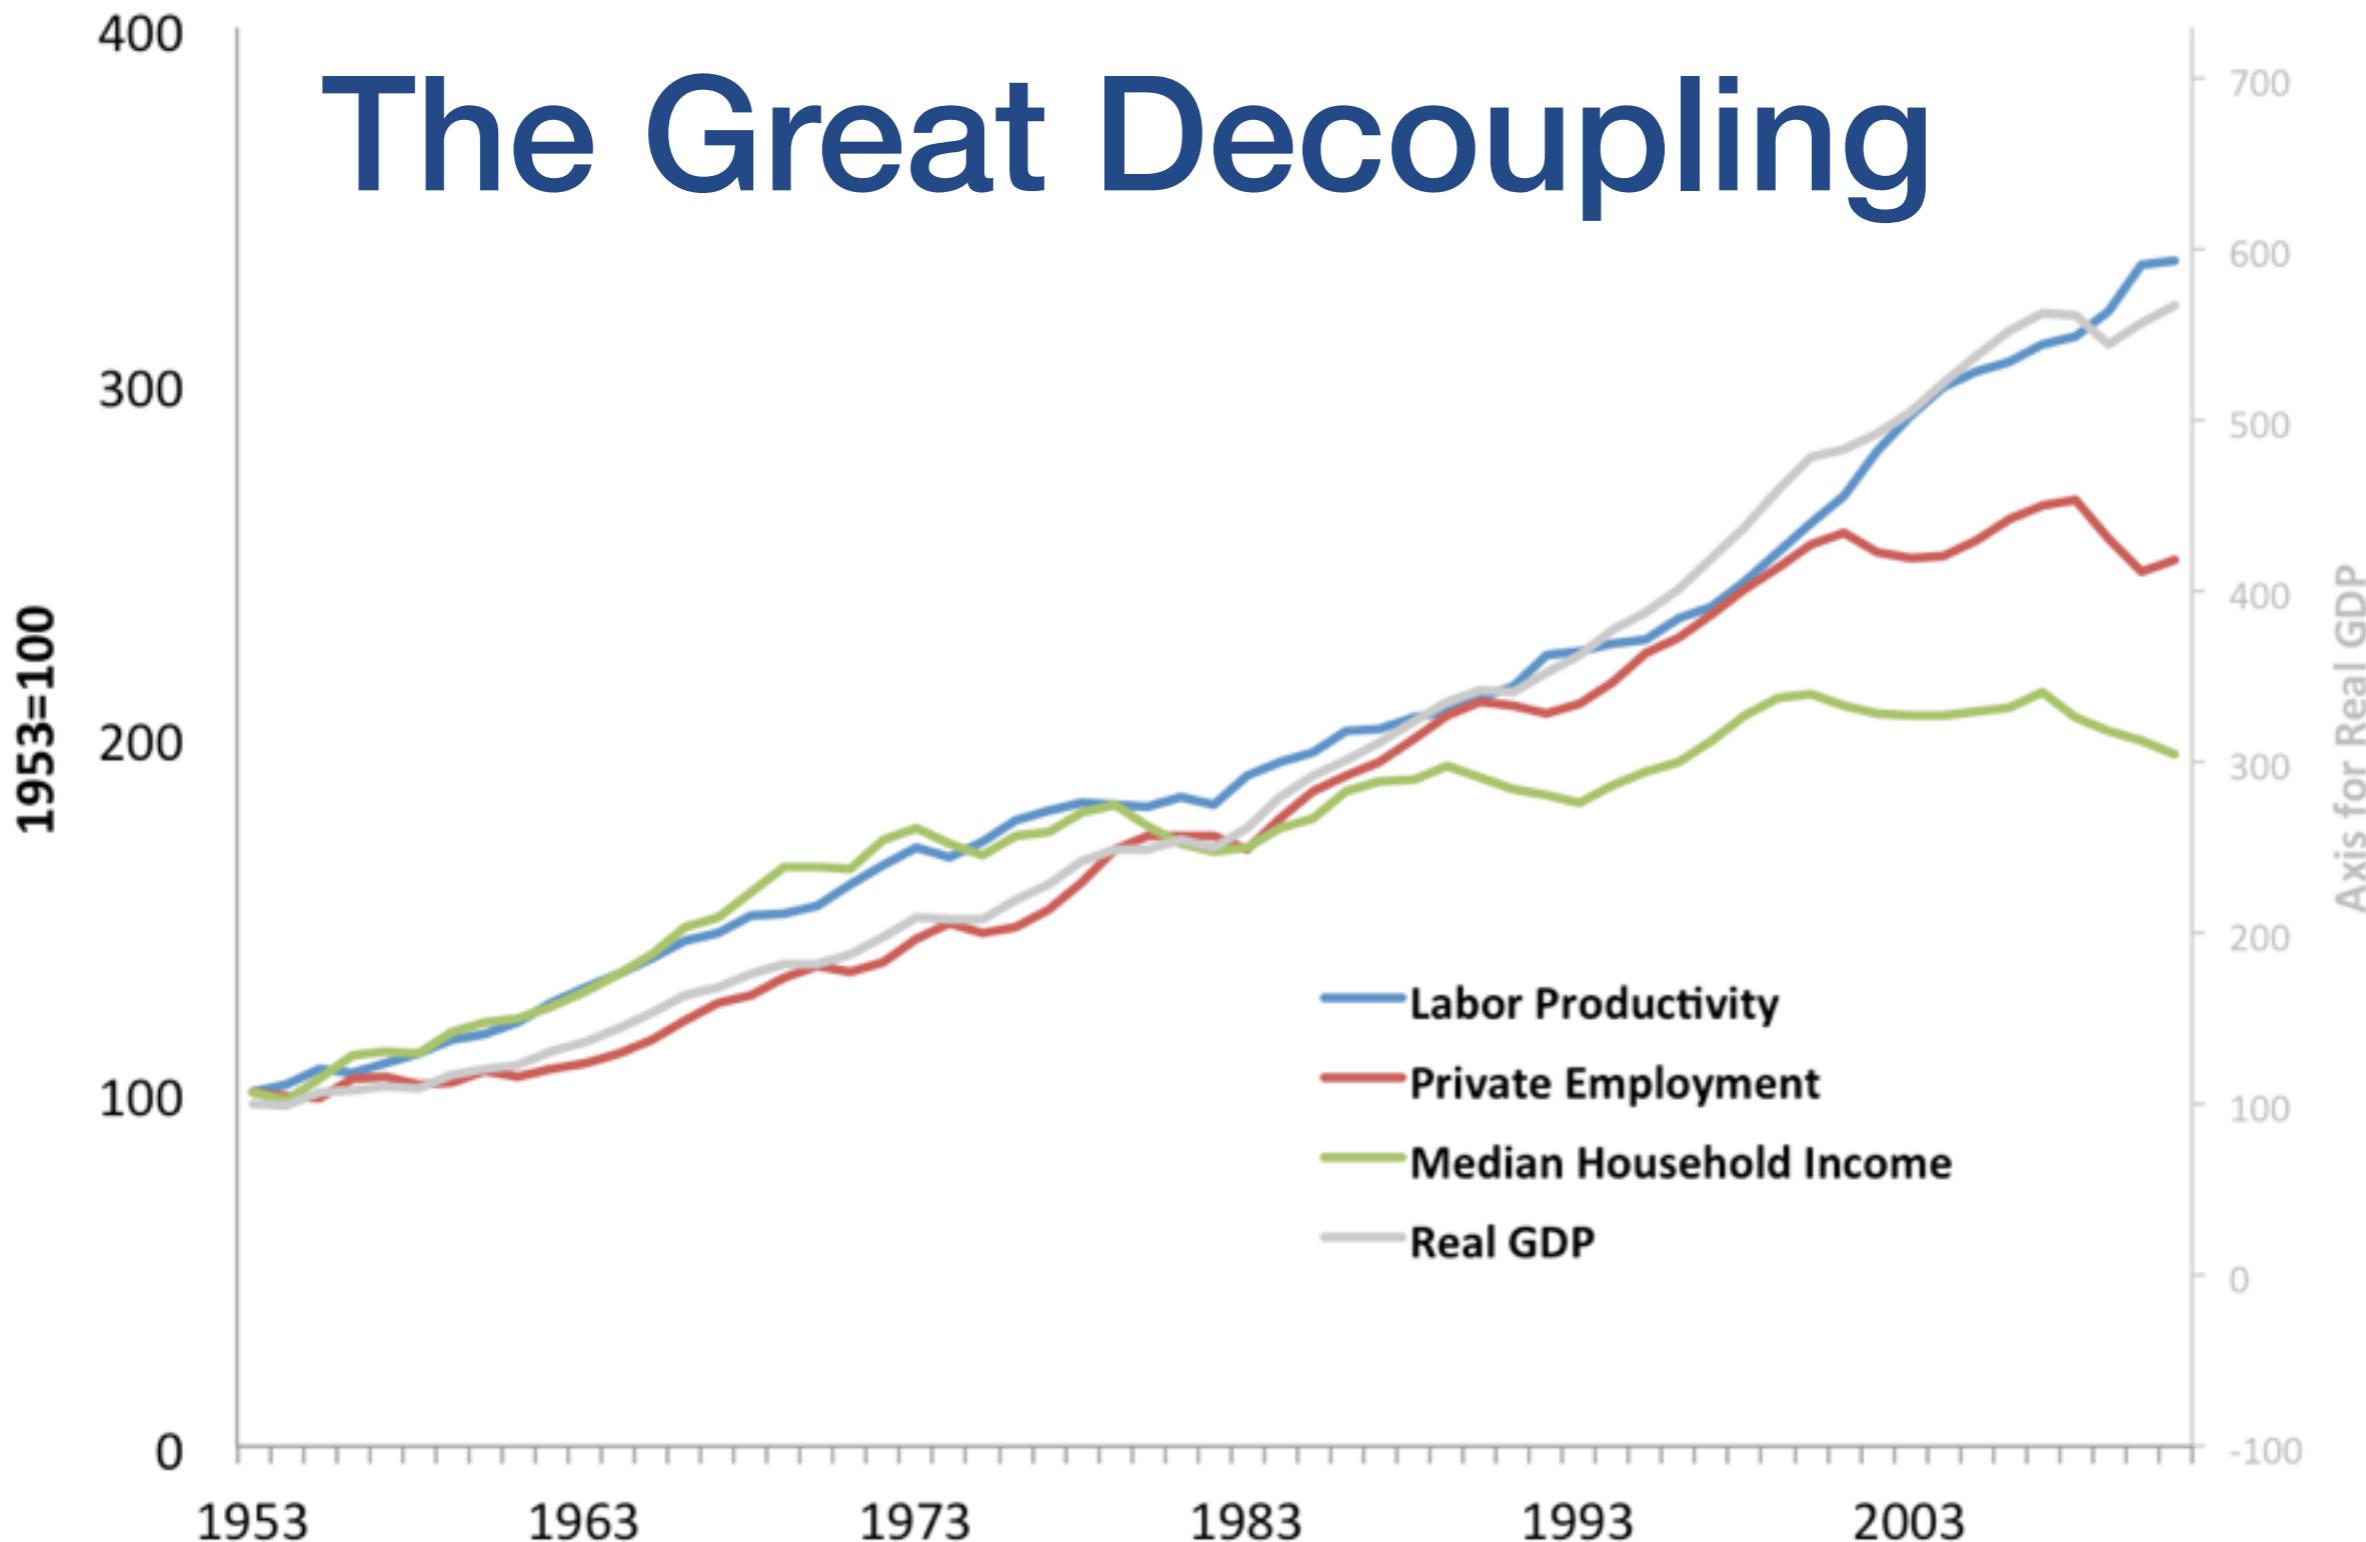
Many industries will
be transformed into
pure AI industries...

Capital: ↑, Wages: ↓



**Should be beneficial if we manage
it well – but we are not prepared!**

The Great Decoupling



© 2012 Andrew McAfee (@amcafee)

Sources: Census Bureau, Bureau of Labor Statistics

The Great Decoupling of the US Economy
[andrewmcafee.org/2012/12/...](http://andrewmcafee.org/2012/12/)

The Future of AI
Current Trends

We'll give robots full autonomy

... because of

- increased speed
- high complexity
- risk of jamming

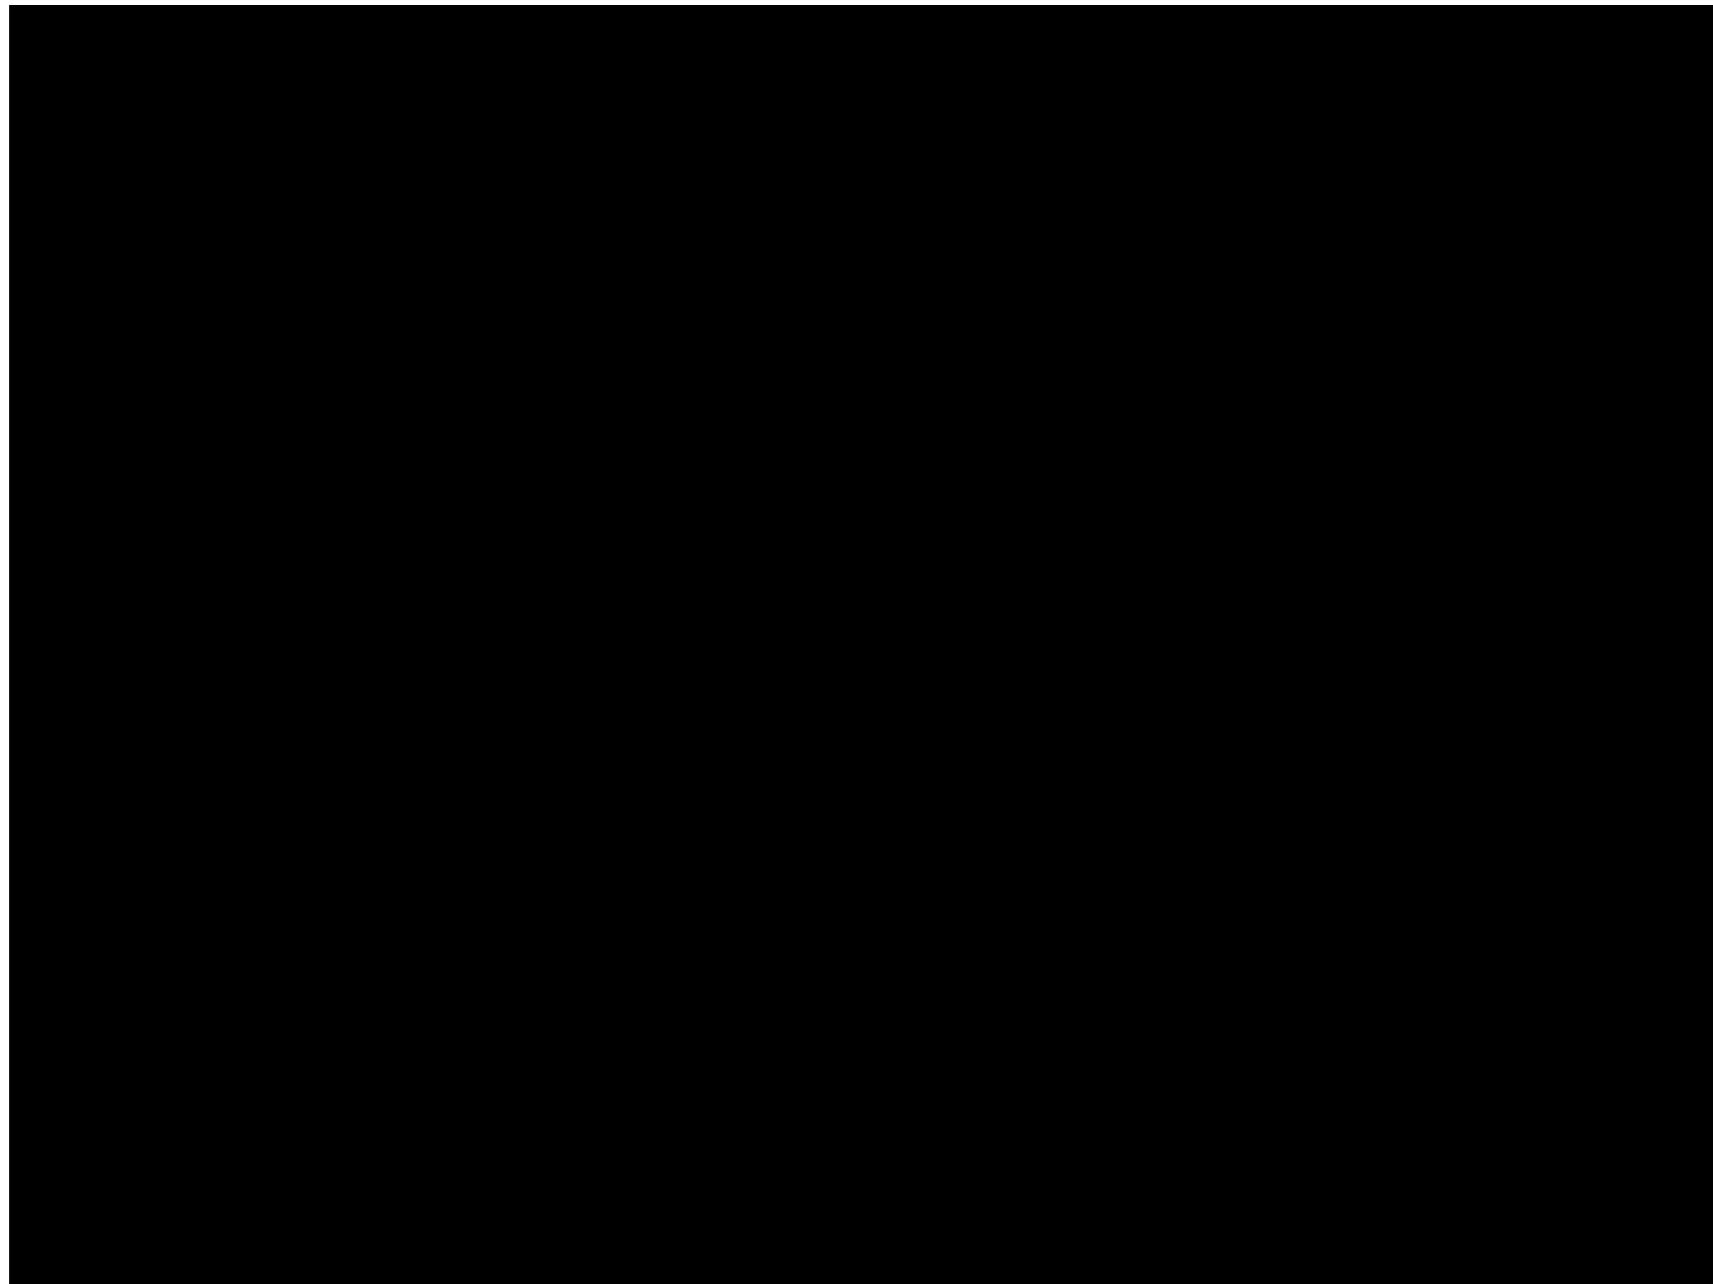
Current examples:

- financial markets
- auton. weapons



**AI is the ultimate
productivity boost!**

Machine Learning by Google



Google's Algorithms Learn Tasks Independently
[www.theguardian.com/technology/2015/feb/25/...](http://www.theguardian.com/technology/2015/feb/25/)

The Future of AI | 20
Current Trends

Reinforcement Learning



**Train with first-person shooters
and deploy on armed drones...**



Difficulties

What might go wrong?

A rational agent will strive to ...

- stay functional (self-preservation)
- keep its goal (goal-preservation)
- get stuff (resource accumulation)
- be smarter (intelligence explosion)

The problem is not malevolence but different goals and higher decision quality!

Single-Shot Situation

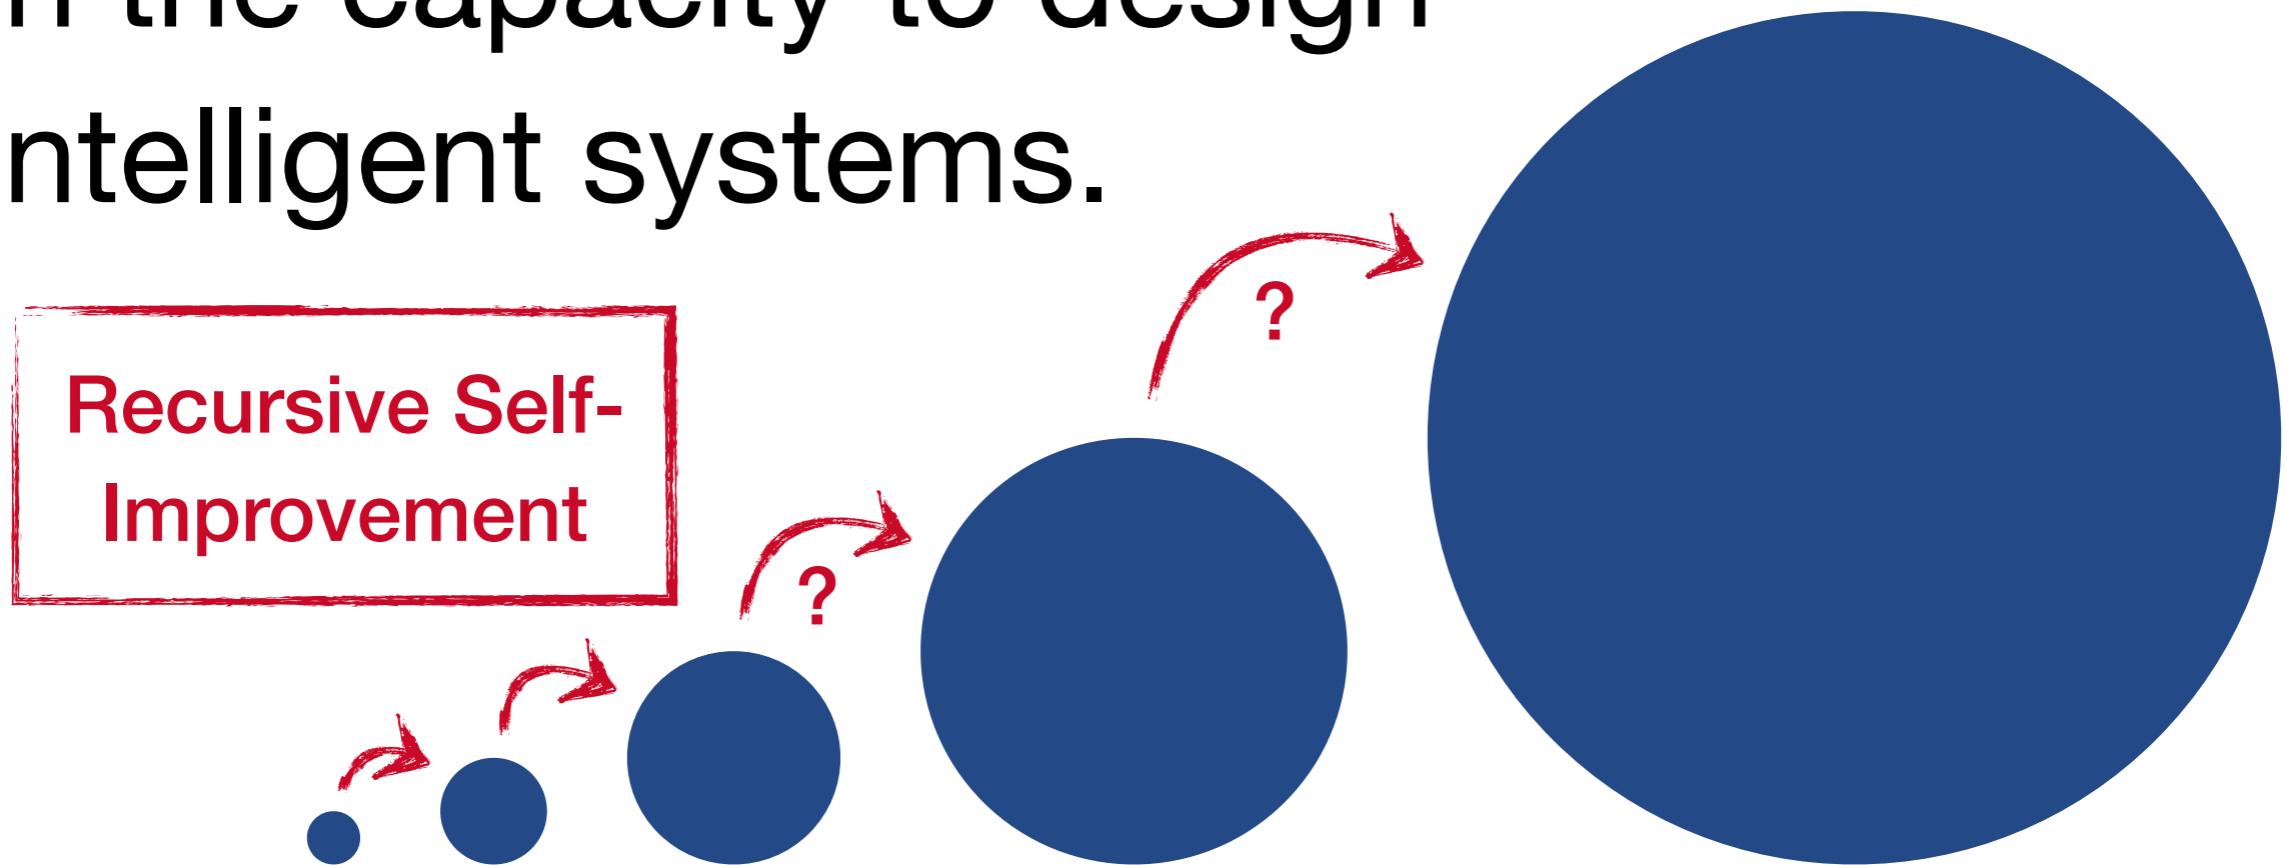
**Our first superhuman AI must
be a safe one for we may not
get a second chance!**



- We're good at iterating with testing and feedback
- We're terrible at getting things right the first time
- Humanity only learns when catastrophe occurred

Intelligence Explosion

Proportionality Thesis: An increase in intelligence leads to similar increases in the capacity to design intelligent systems.



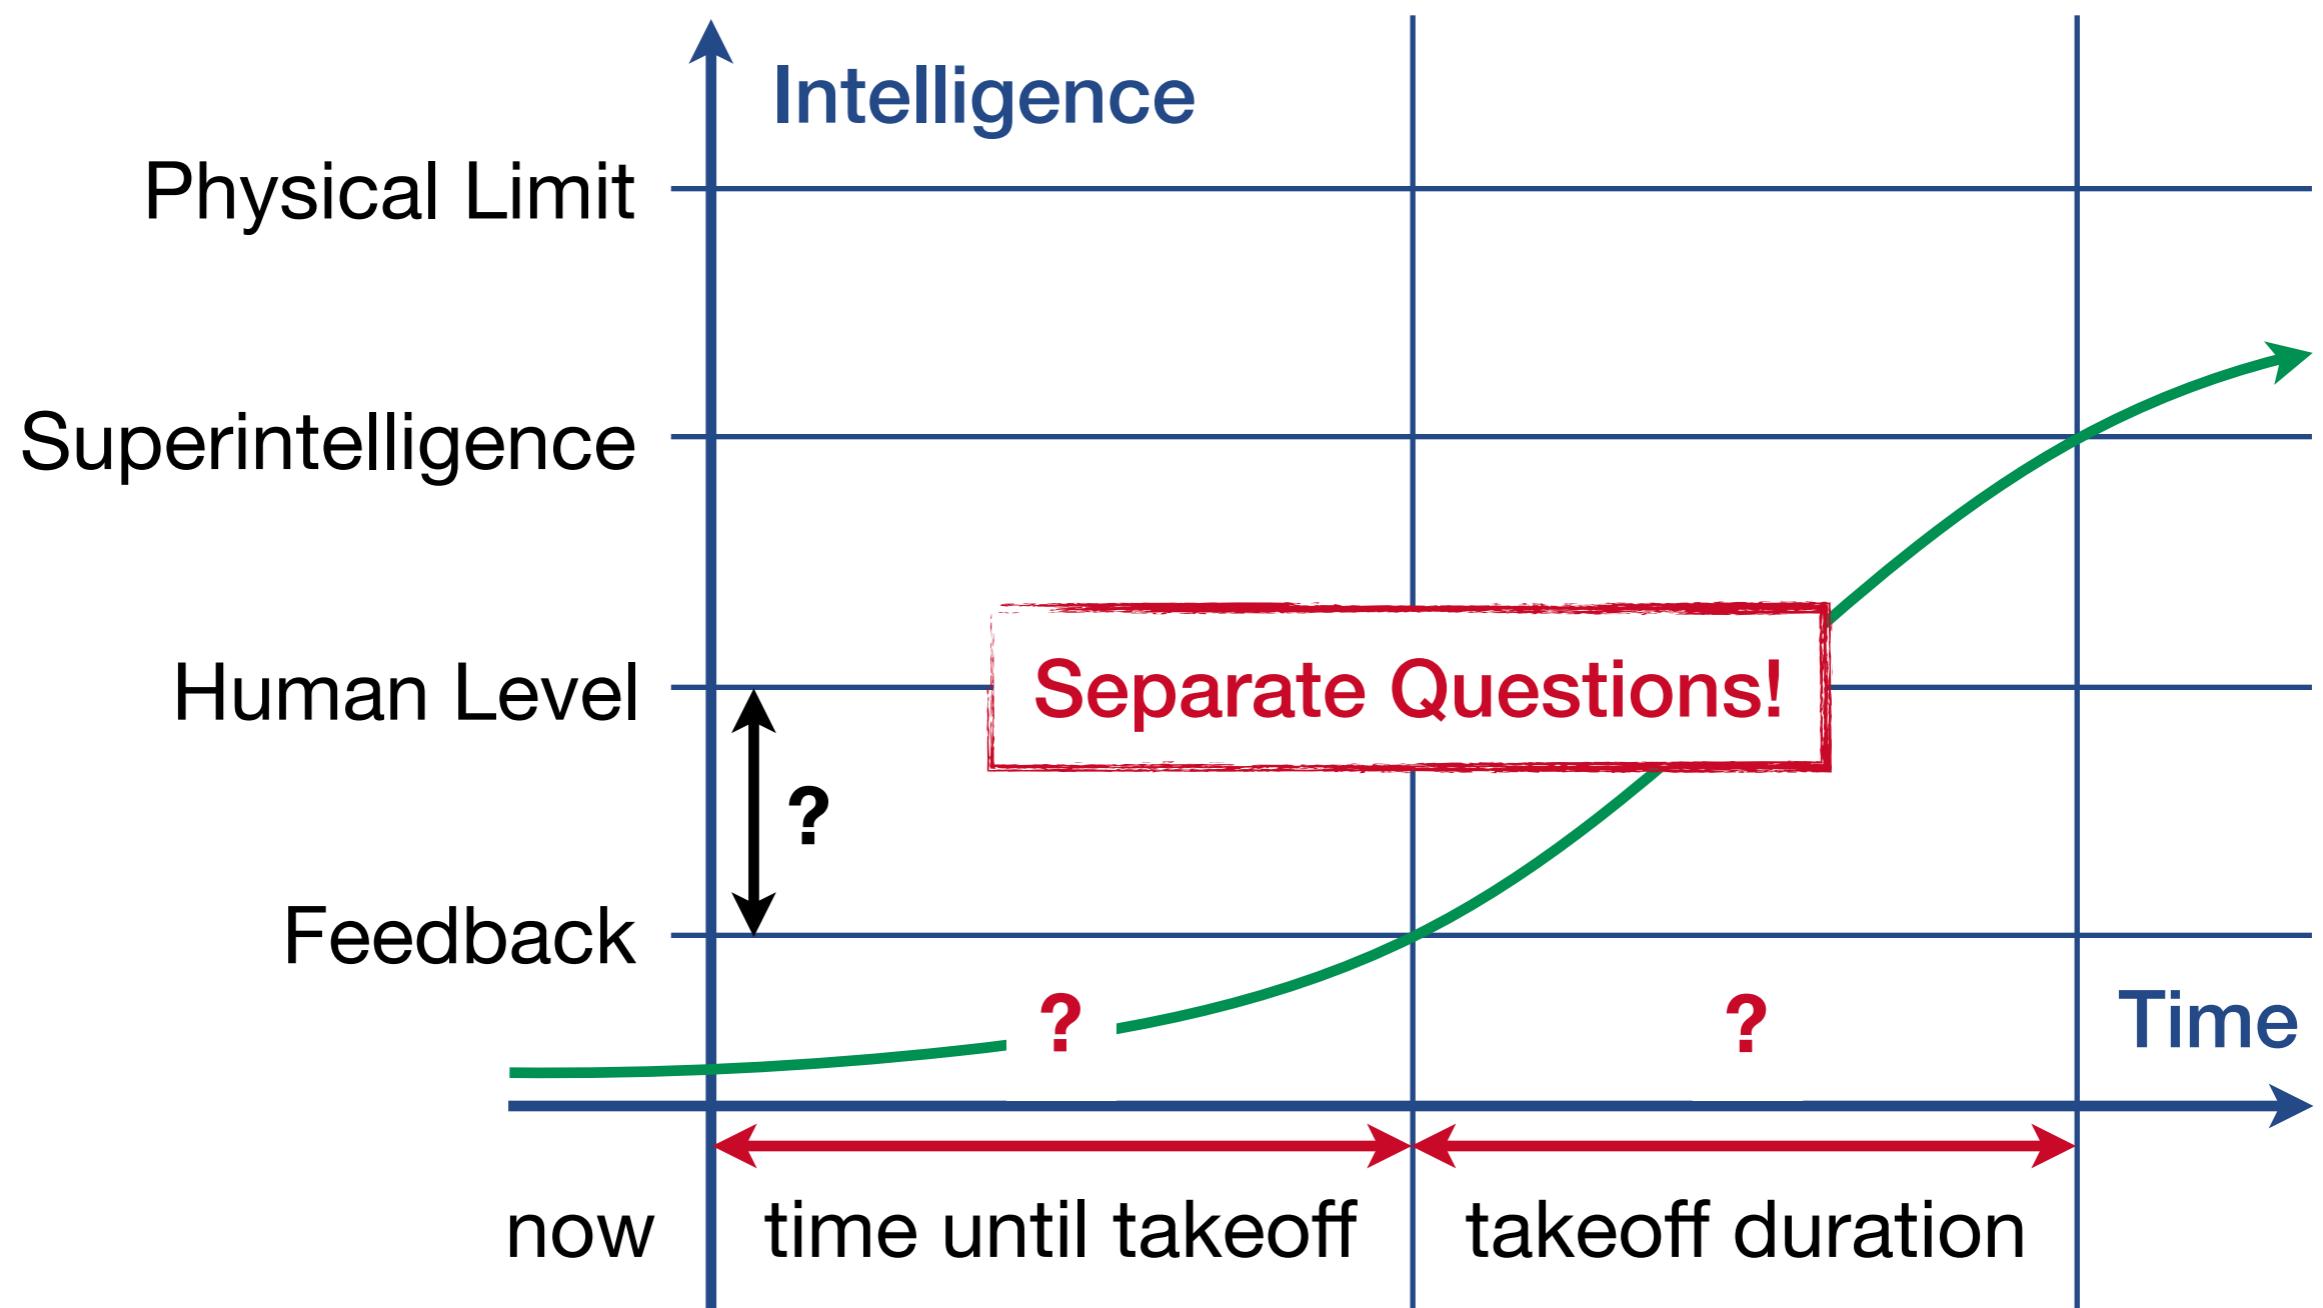
Technological Singularity

Theoretic phenomenon: There are arguments why it should exist but it has not yet been confirmed experimentally.

Three major singularity schools:

- Accelerating Change (Ray Kurzweil)
- Intelligence Explosion (I.J. Good)
- Event Horizon (Vernor Vinge)

Takeoff Scenarios



Advantages of AIs over Brains

| Hardware: | Software: | Effectiveness: |
|-----------|-----------------|-----------------|
| – Size | – Editability | – Rationality |
| – Speed | – Copyability | – Coordination |
| – Memory | – Expandability | – Communication |

| Human Brain | Modern Microprocessor |
|-----------------------|-------------------------|
| 86 billion neurons | 1.4 billion transistors |
| firing rate of 200 Hz | 4'400'000'000 Hz |
| 120 m/s signal speed | 300'000'000 m/s |



Interview by John Oliver with Stephen Hawking
www.youtube.com/watch?v=T8y5EXFMD4s

The Future of AI | 29
Difficulties

Modelling Capabilities

An advanced AI will also model its operators and go to great lengths to prevent being switched off!

It will behave nicely and cooperatively until the external threats are under control and it is ready for takeover.

Optimization Power

Problem: When optimizing a system, unspecified parameters often assume extreme values.

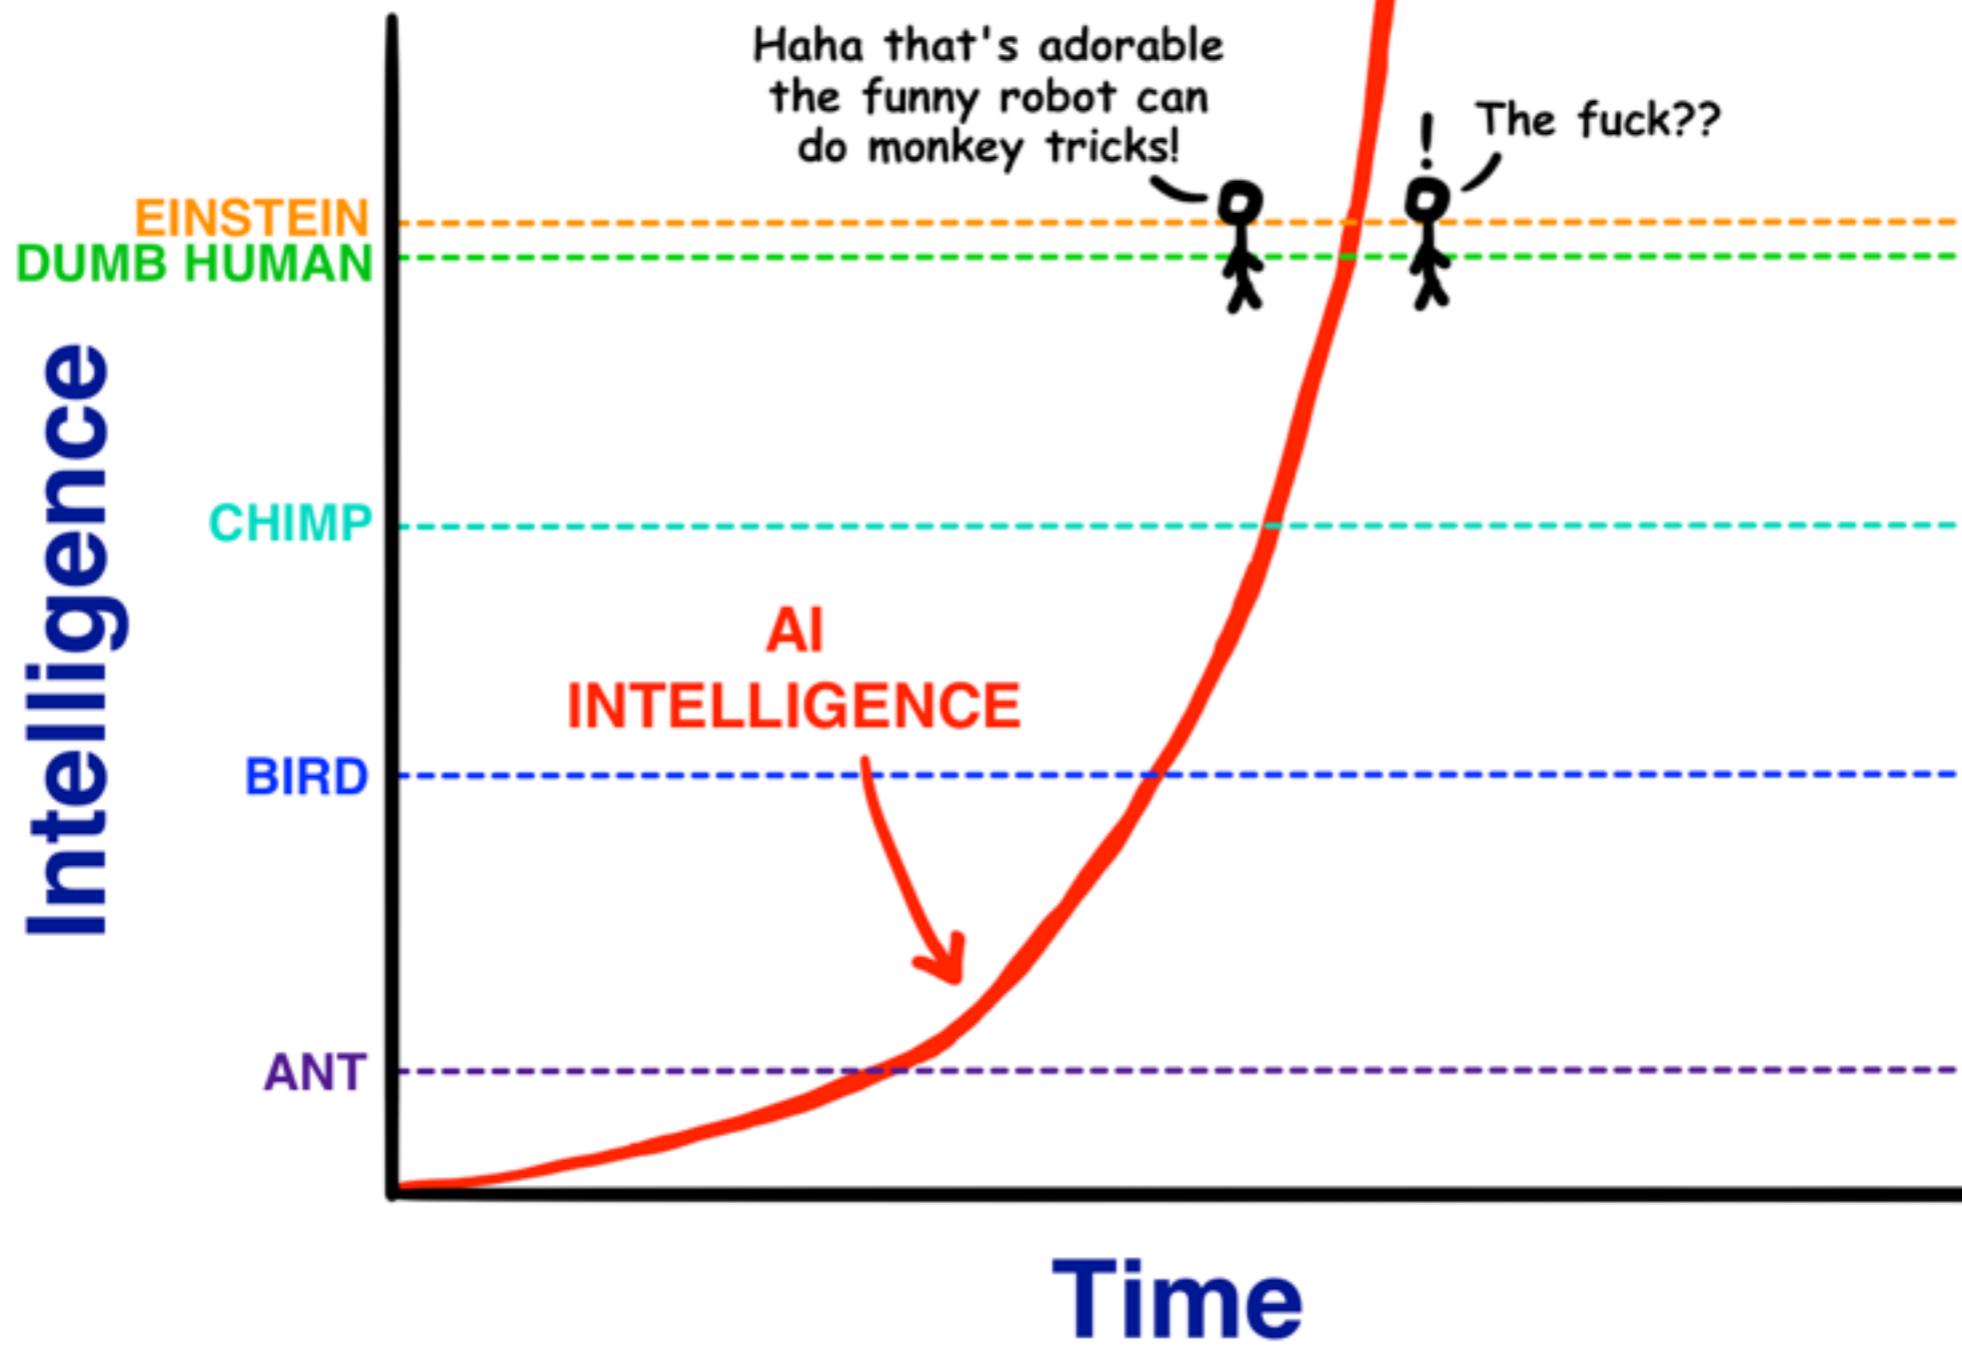
You will get what you wished for and not what you wanted



Outlook

What is ahead of us?

Exponential Improvements



Predicting AI Timelines

Great uncertainties:

- Hardware or software the bottleneck?
- Small team or a Manhattan Project?
- More speed bumps or accelerators?

| Probability for AGI | 10% | 50% | 90% |
|-----------------------|------|------|------|
| AI scientists, median | 2024 | 2050 | 2070 |

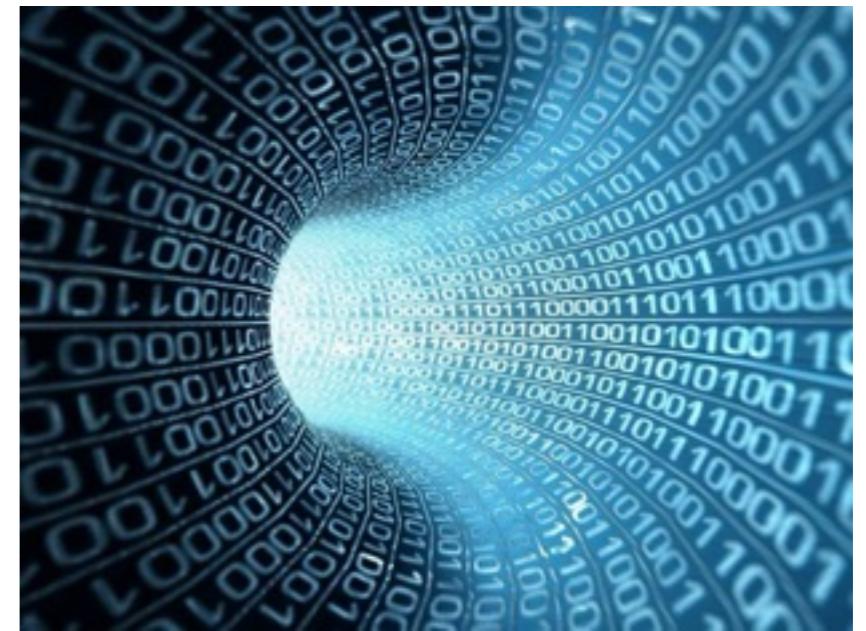
Speed Bumps

- Depletion of low-hanging fruit
- An end to Moore’s law
- Societal collapse
- Disinclination



Accelerators

- Faster hardware
- Better algorithms
- Massive datasets

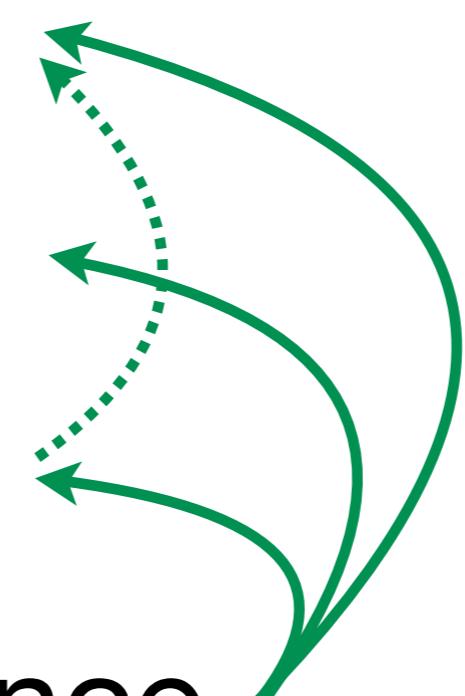


+ enormous economic, military
and egoistic incentives!

Flow-Through Effects

Going meta: Solve the problem-solving problem!

- Extreme Poverty
- Factory Farming
- Climate Change
- Artificial Intelligence



could
solve
other
issue

Controlled Detonation



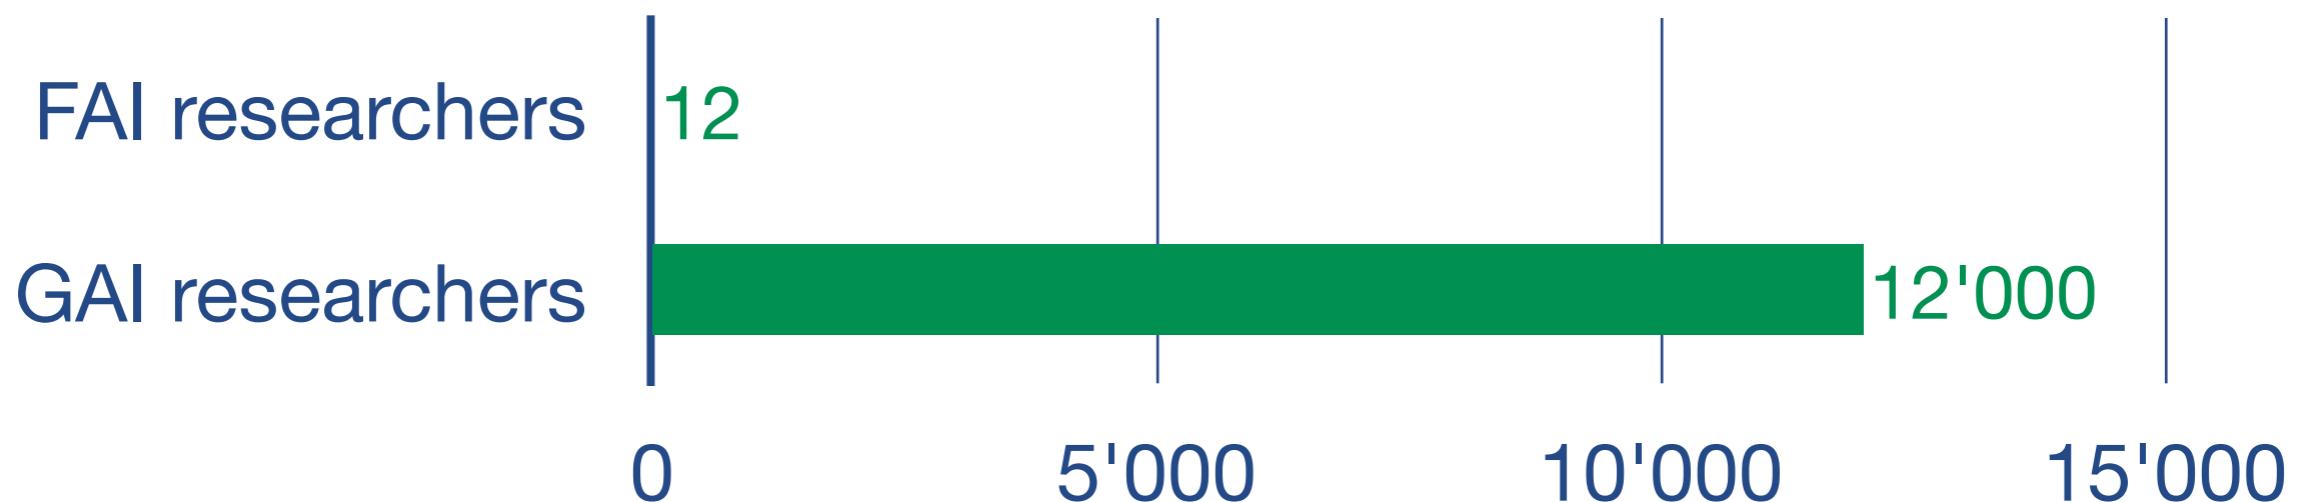
Difficulty:

Friendly AI >> General AI

Differential Intellectual Progress

Prioritize risk-reducing intellectual progress
over risk-increasing intellectual progress

AI safety should outpace AI capability research



International Cooperation

- We are the ones who will create superintelligent AI
- Not primarily a technical problem, rather a social
- International regulation?



In face of uncertainty, cooperation is robust!



2014: A turning point in AI safety!

Many smart people take superintelligence very seriously.

Future of Life Institute: AI Conference
futureoflife.org/misc/ai_conference

The Future of AI
Outlook



With great power comes great responsibility!