Why CSC180 is the the most important course you’ll ever take*

With slides taken from

The Future of Computer Science

And Why Every Other Major Sucks By Comparison

by Scott Aaronson, MIT EECS

Stereotypes of CS Majors

Moore's Law

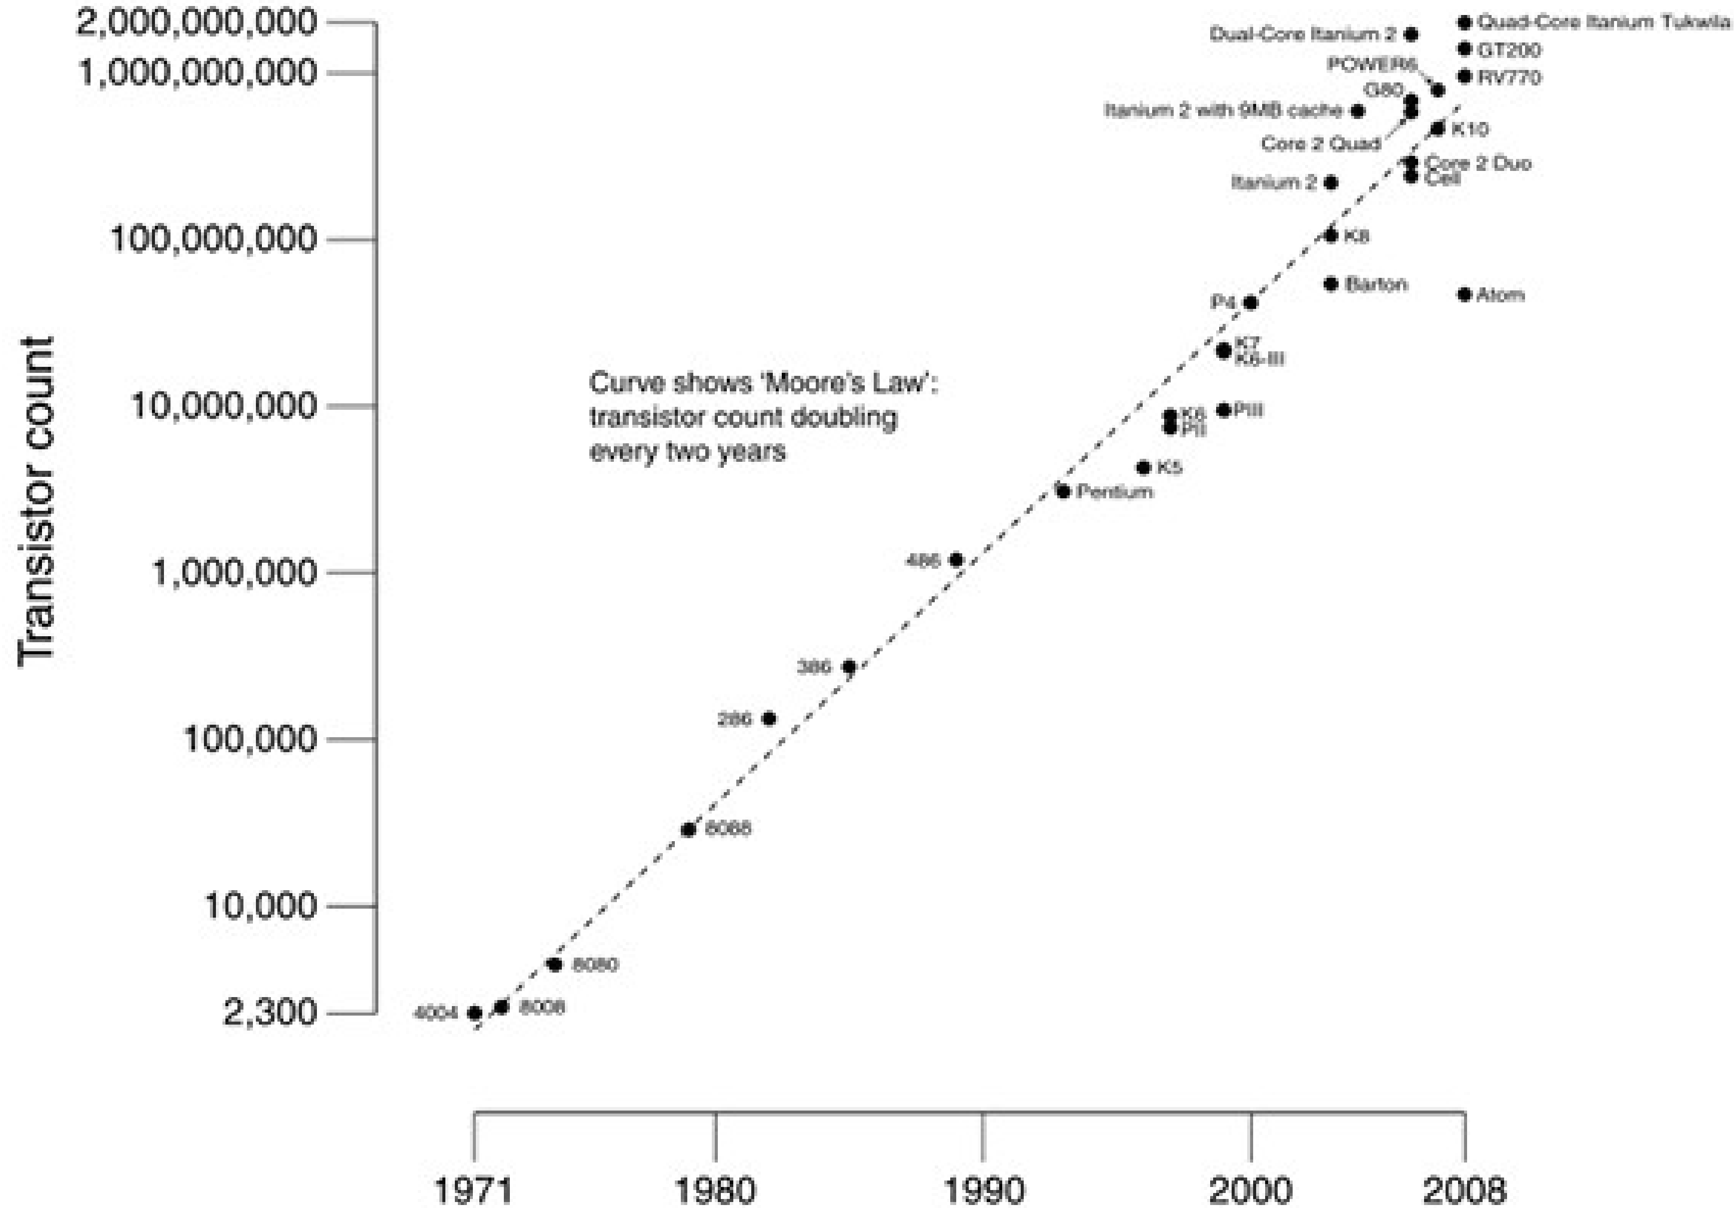
Transistor count
2,000,000,000
1,000,000,000
100,000,000
10,000,000
1,000,000
100,000
10,000
2,300
1971
1980
1990
2000
2008
Curve shows 'Moore's Law': transistor count doubling every two years
Dual-Core Itanium 2
Quad-Core Itanium Tukwila
GT200
POWER6
RV770
G80
Itanium 2 with 9MB cache
K10
Core 2 Quad
Core 2 Duo
Cell
Itanium 2
K8
Barton
Atom
P4
K7
K6-III
PIII
K6
PII
K5
Pentium
486
386
286
8088
8080
8008
4004

Insights From and to Other Disciplines

- Insights from *Distributional Semantics* lead to techniques like the one used in Project 3
  - *"Distributional hypothesis: linguistic items with similar distributions have similar meanings"* (Wikipedia)
- Programs inspired by how the brain works used to automatically learn how to detect cats in Youtube videos
- And to automatically learn to play videogames by trying different inputs and learning from mistakes

Technology

US Secret Service seeks Twitter sarcasm detector

5 June 2014 | Technology

The Secret Service wants its own automated system of social media monitoring

nature International weekly journal of science

NATURE | LETTER

日本語要約

Human-level control through deep reinforcement learning

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg & Demis Hassabis

Affiliations | Contributions | Corresponding authors

Nature **518**, 529–533 (26 February 2015) | doi:10.1038/nature14236
Received 10 July 2014 | Accepted 16 January 2015 | Published online 25 February 2015

The theory of reinforcement learning provides a normative account[1], deeply rooted in psychological[2] and neuroscientific[3] perspectives on animal behaviour, of how agents may optimize their control of an environment. To use reinforcement learning successfully in situations approaching real-world complexity, however, agents are confronted with a difficult task: they must derive efficient representations of the environment from high-dimensional sensory inputs, and use these to generalize past experience to new situations. Remarkably, humans and other animals seem to solve this problem through a harmonious combination of reinforcement learning and hierarchical sensory processing systems[4, 5], the former evidenced by a wealth of neural data revealing notable parallels between the phasic signals emitted by dopaminergic neurons and temporal difference reinforcement learning algorithms[3]. While reinforcement learning agents have achieved

This AI “solves” *Super Mario Bros.* and other classic NES games

Program even takes advantage of bugs and glitches.

by Ian Steadman, wired.co.uk - Apr 14, 2013 3:00pm EDT

Nintendo

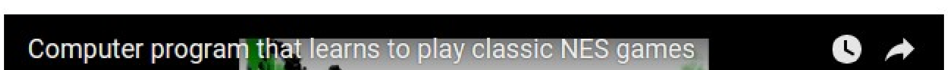

LATEST FEATURE STORY

FEATURE STORY (2 PAGES)

Bikes, bowlin the delicate b that is moder

Ars goes inside the mac recycling work.

WATCH ARS VIDEO

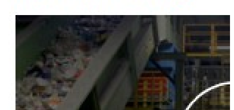

WIRED UK SCIENCE 06.26.12 11:15 AM

GOOGLE'S ARTIFICIAL BRAIN LEARNS TO FIND CAT VIDEOS

So what’s next?

Robot Uprising

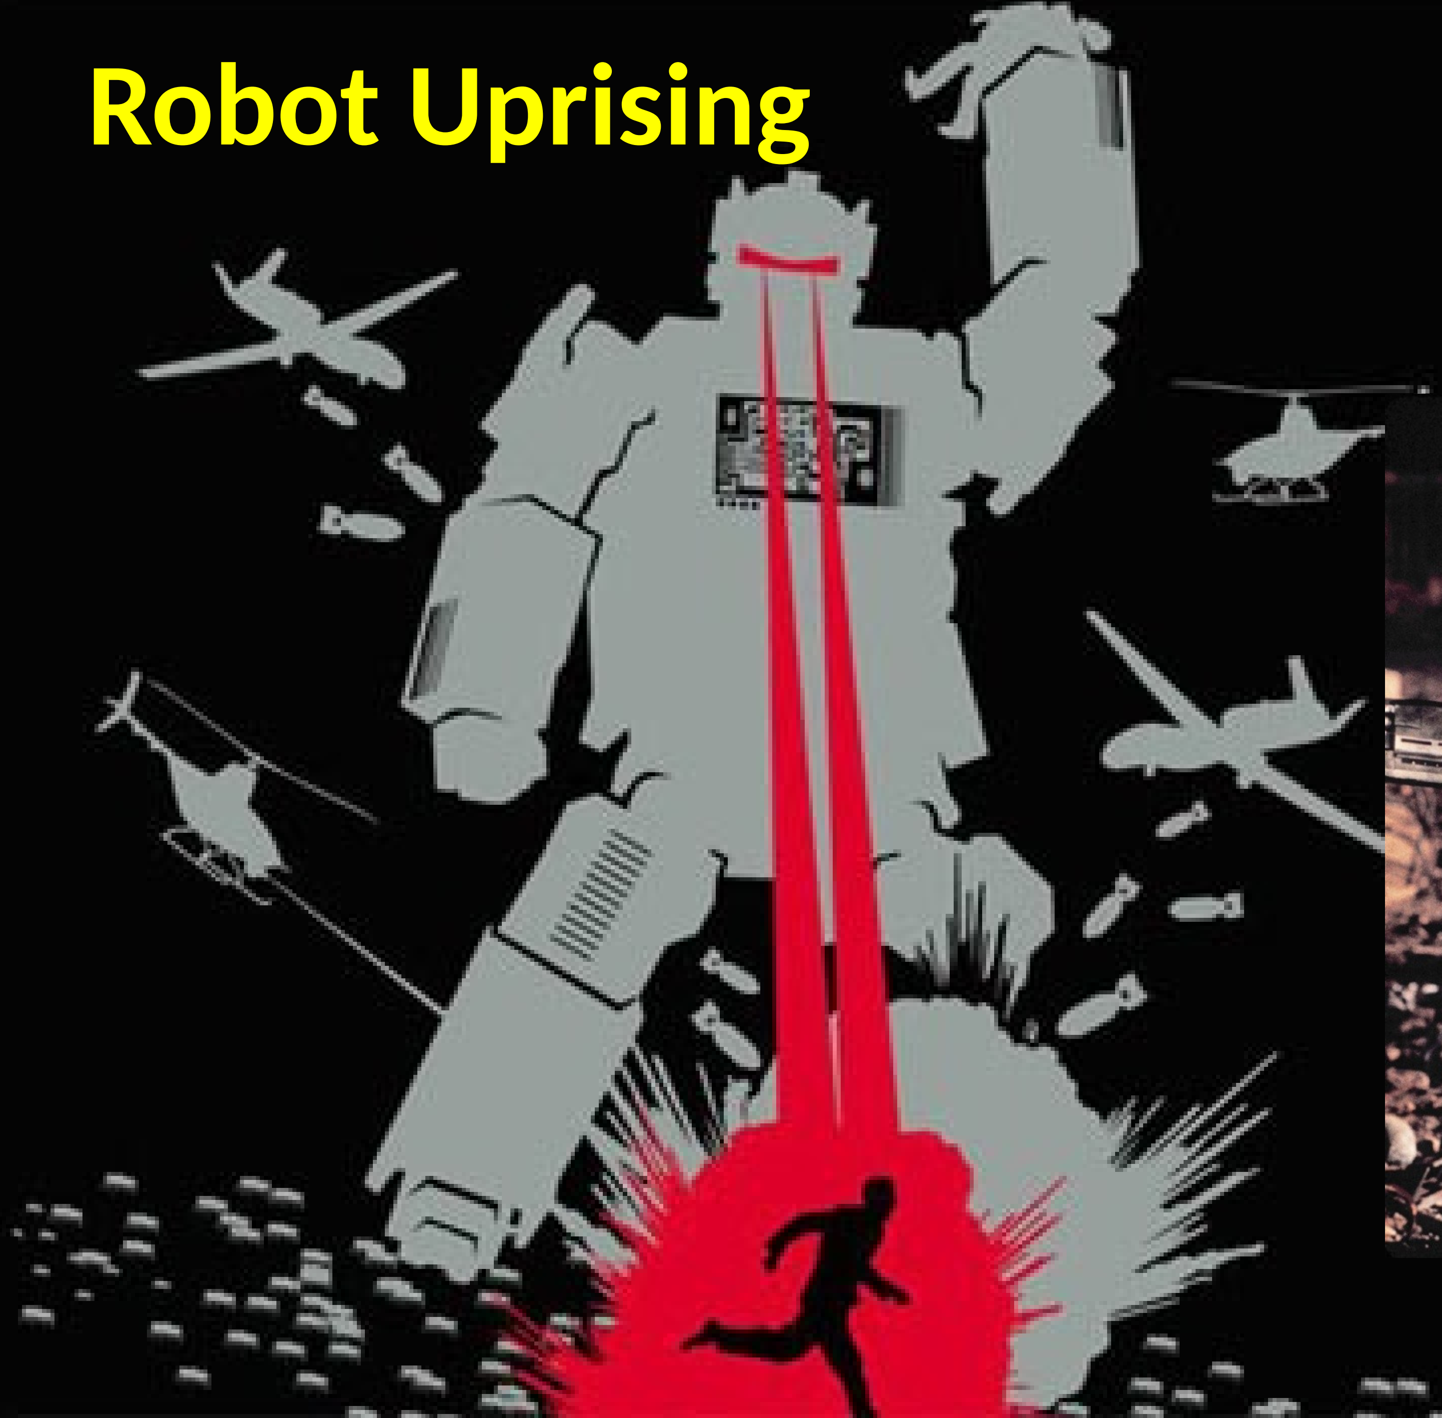

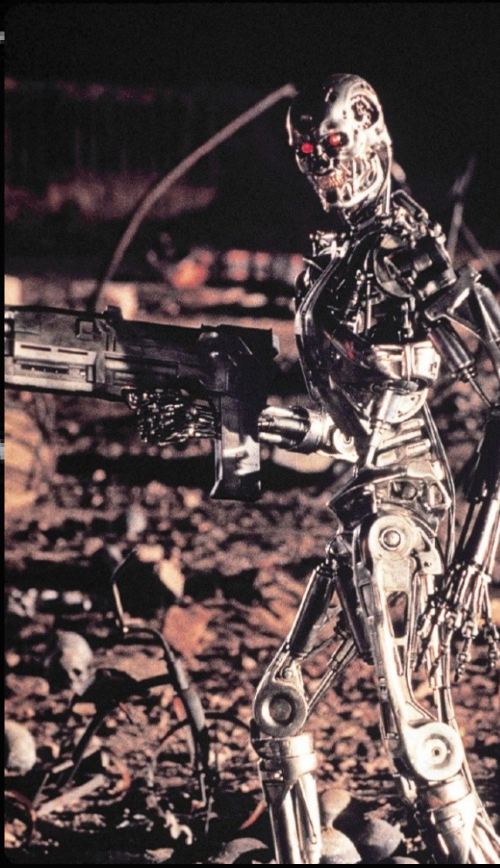

Uploading our brains to computers; replacement of “real life” by the Matrix

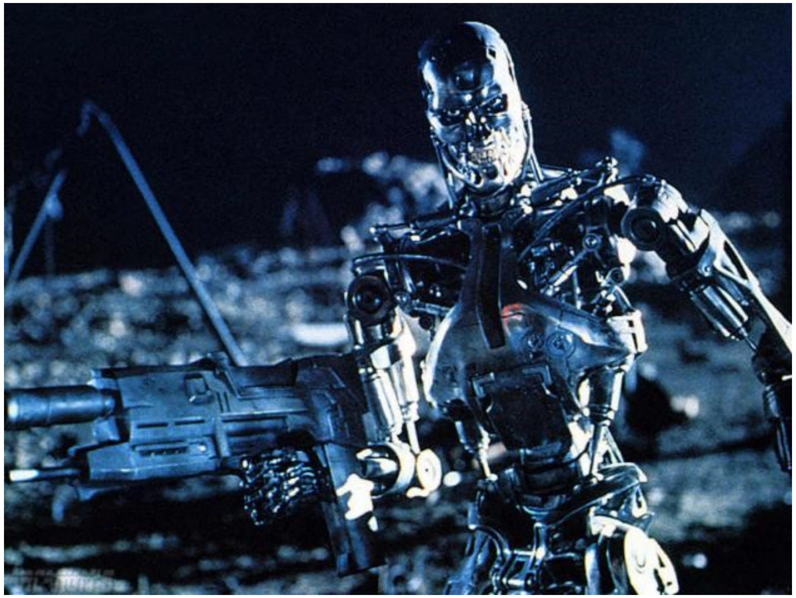

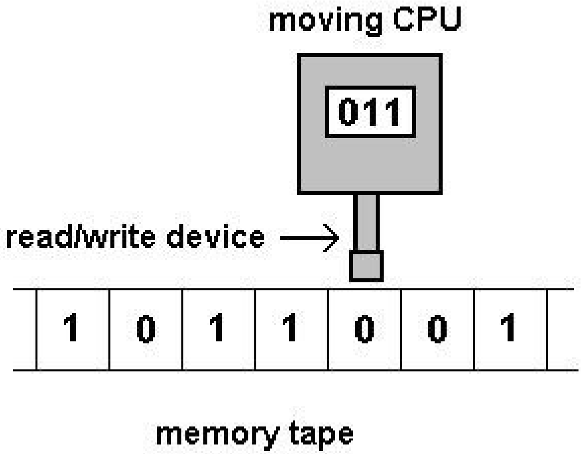

So what else is there?

Quantum Computers

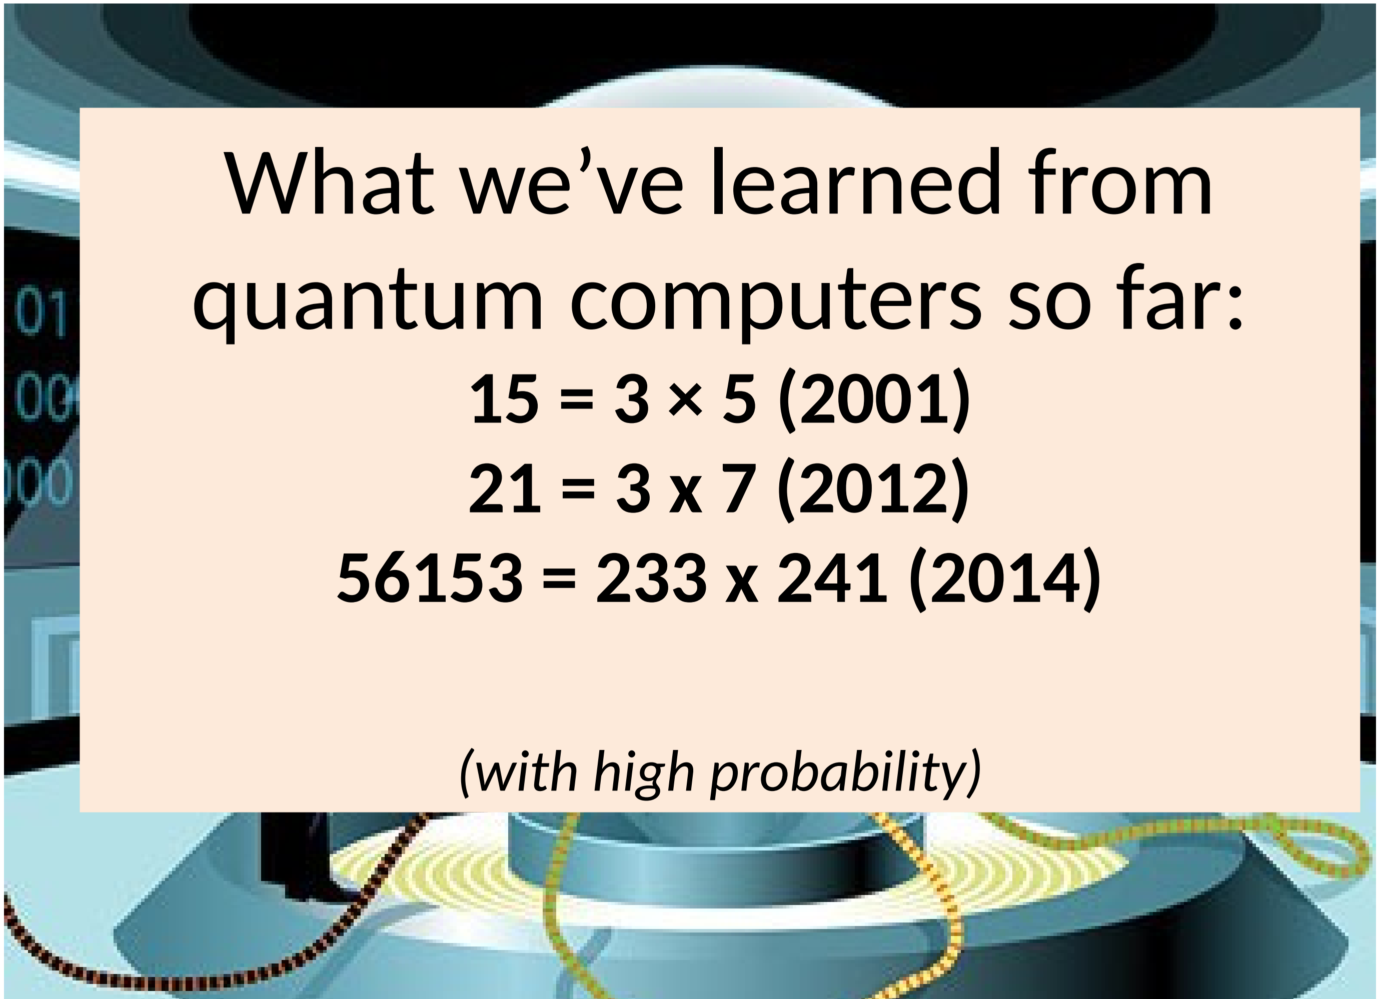

What we've learned from quantum computers so far:

15 = 3 × 5 (2001)

21 = 3 x 7 (2012)

56153 = 233 x 241 (2014)

(with high probability)

Relativity Computer

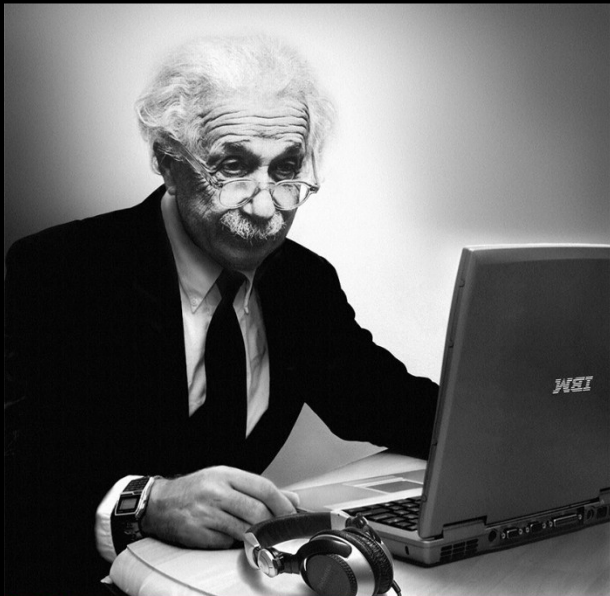

Zeno's Computer

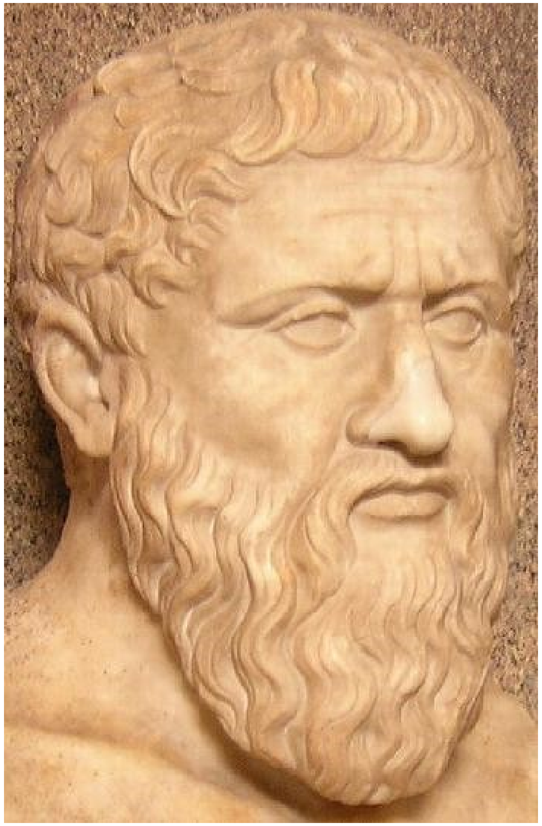

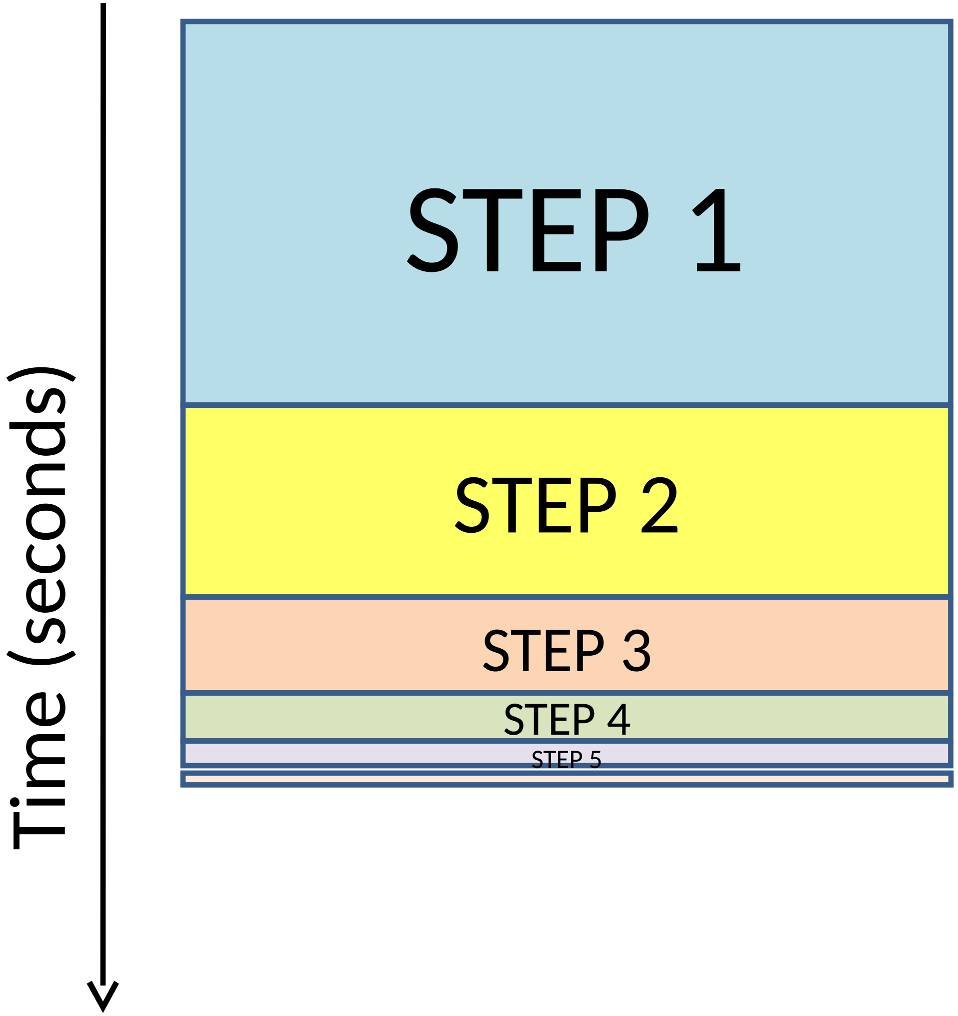

Time Travel Computer

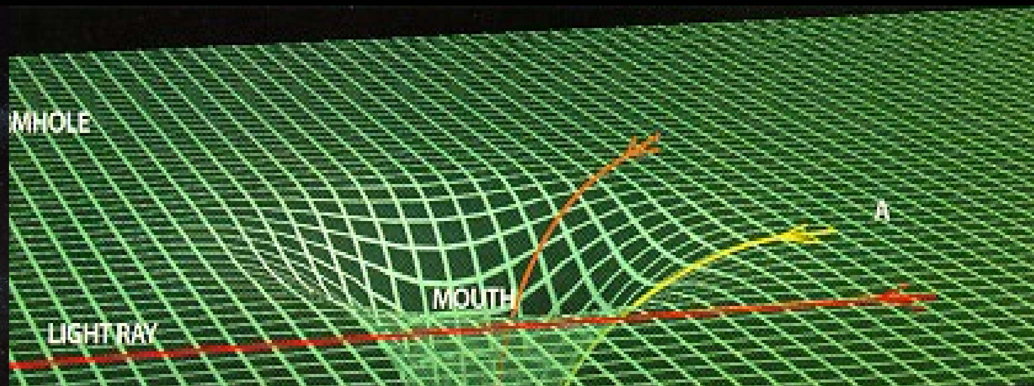

S. Aaronson and J. Watrous. **Closed Timelike Curves Make Quantum and Classical Computing Equivalent**, *Proceedings of the Royal Society A* 465:631-647, 2009. arXiv:0808.2669.

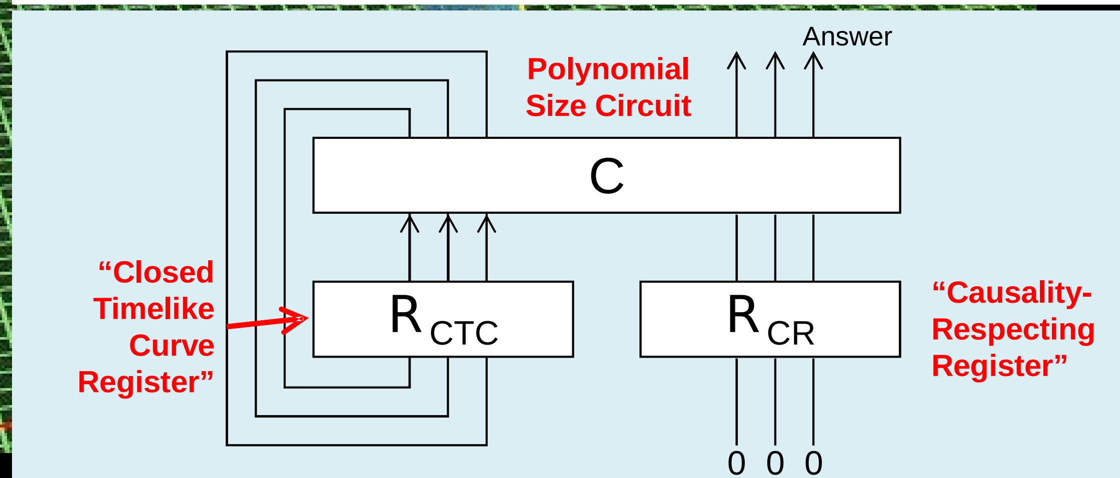

Computer Science Is Interdisciplinary