Flowers and Light: What to Expect from AI

Matthew Turk
Professor and Chair
Dept. of Computer Science
UC Santa Barbara

47th Annual SCALL Institute

WHAT IS A.I.?

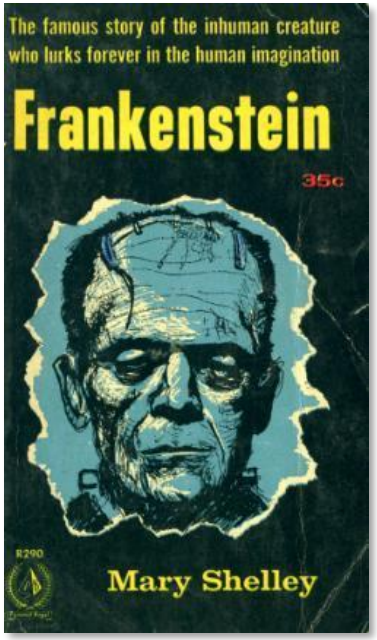

1818

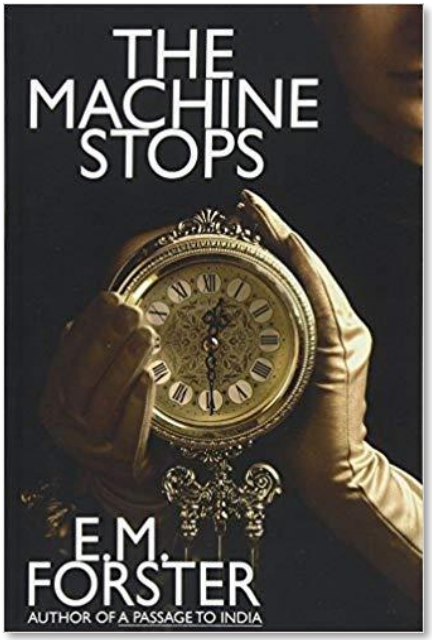

1909

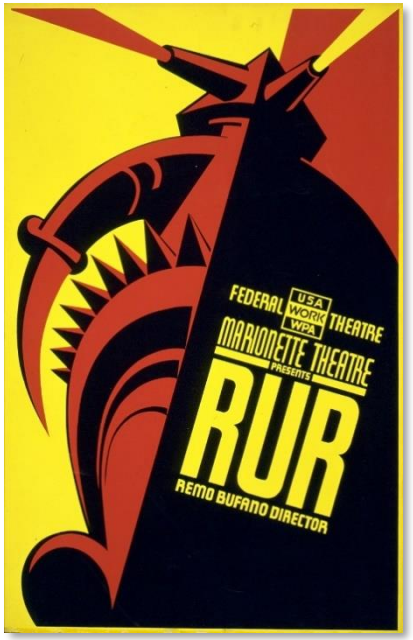

1920

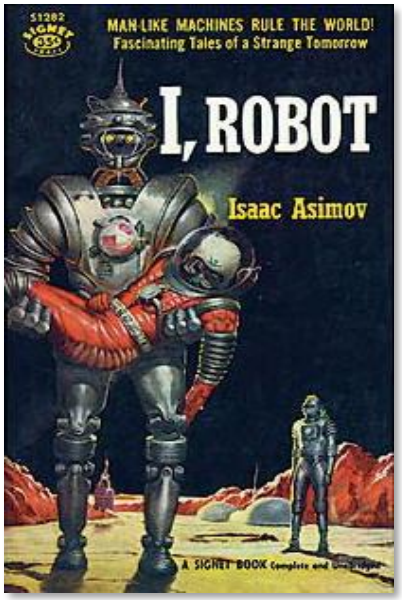

1942

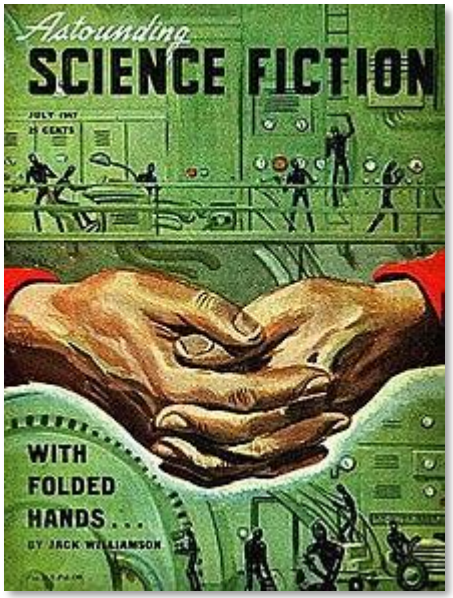

1947

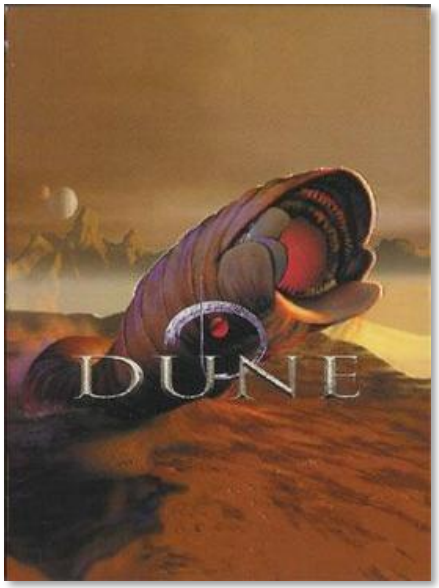

1965

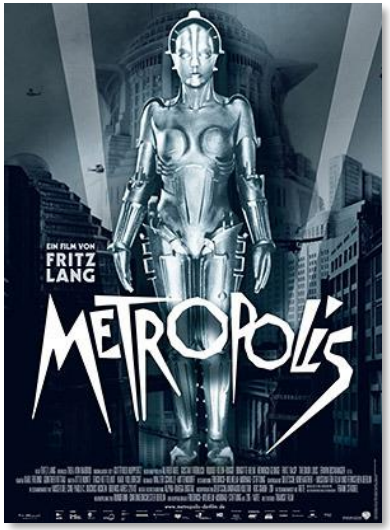
EIN FILM VON
FRITZ
LANG
METROPOLIS


DER HERR
DER WELT
REGIE
HARRY PIEL

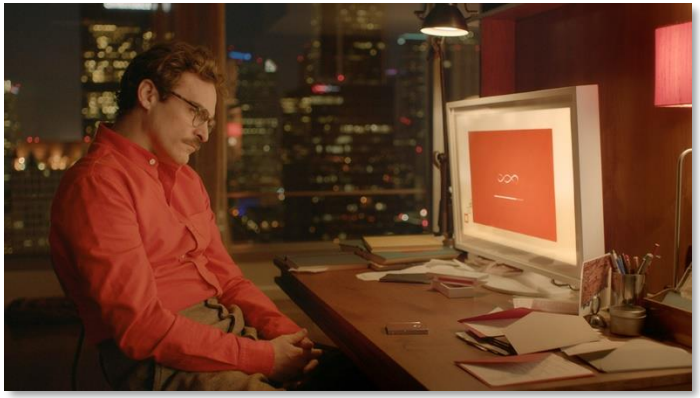

2001: A Space Odyssey (1968)

HAL was "born" on January 12, 1997 in Urbana, Illinois

- January 12, 1992 in the film

What does HAL do?

- Plays chess
- Talks, converses w/people
- Exhibits common sense
- Sees
- Appreciates art
- Reads lips!
- Has feelings, emotions
- Tries to manipulate people
- Exhibits fear of non-existence

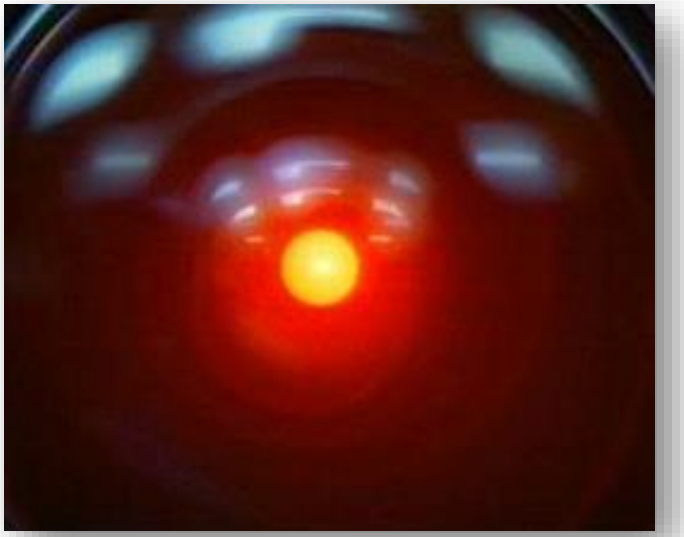

Meanwhile, in the real world…

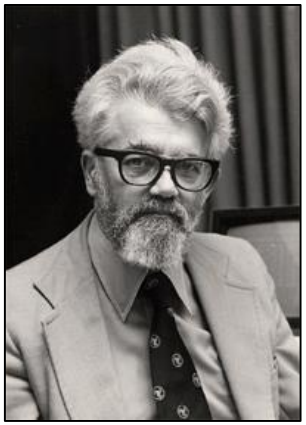

John McCarthy

Herb Simon

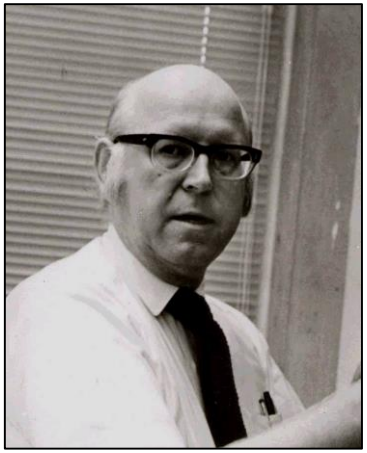

Allen Newell

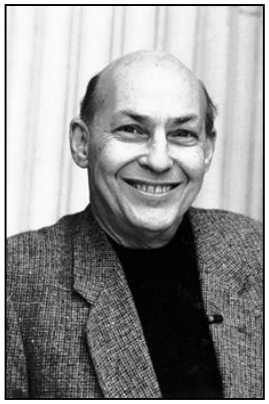

Marvin Minsky

Seminal event for Artificial Intelligence as a field, in 1956:
The **Dartmouth Summer Research Conference on Artificial Intelligence**

"We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer."

Foundations of AI

Philosophy

- Framed the ideas of AI
- Dualism/materialism, logical/rational/empirical, causality, consciousness, mind/body…

Mathematics

- Formalized computation, logic, probability
- Possibilities and limitations of computation

Psychology

- Experimental: the brain as an information processing device (Cognitive Science, Neuroscience, Psychophysics)

Computer Science and Engineering

- Algorithms
- Built real machines, Moore’s Law progress

So what is Artificial Intelligence?

"The science and engineering of making intelligent machines, especially intelligent computer programs."

"The study of how to make computers do things at which, at the moment, people are better."

"The design of flexible programs that respond productively in situations that were not specifically anticipated by the designer."

"The construction of computations that perceive, reason, and act effectively in uncertain environments."

"The branch of CS concerned with enabling computers to simulate such aspects of human intelligence as speech recognition, deduction, inference, creative response, the ability to learn from experience, and the ability to make inferences given incomplete information."

"Modeling aspects of human cognition on computers"

"What AI people do"

What AI people study

Logic
Knowledge representation
Search
Reasoning/inference
Non-monotonic reasoning
Planning
Probabilistic reasoning
Naïve physics
Machine learning

Speech recognition
Natural language processing
Computer vision
Pattern recognition
Intelligent agents
Robotics
Neural networks
Data mining
Expert systems

… and more…

IJCAI topics (keywords for submission)

Agent-based and Multi-agent Systems
Agent/AI Theories and Architectures
Agent-based Simulation and Emergent Behavior
Agent Communication
Argumentation
Auctions And Market-Based Systems
Coordination And Collaboration
Distributed AI
E-Commerce
Game Theory
Information/Mobile/Software Agents
Multiagent Learning
Multiagent Planning
Multiagent Systems (General/other)
Negotiation And Contract-Based Systems
Social Choice Theory
Constraints, Satisfiability, and Search
Applications
Constraint Optimization
Constraint Satisfaction (General/other)
Distributed Search/CSP/Optimization
Dynamic Programming
Search, SAT, CSP: Evaluation and Analysis
Global Constraints
Heuristic Search
Search, SAT, CSP: Meta-heuristics
Meta-Reasoning
Quantifier Formulations
Satisfiability (General/other)·
SAT and CSP: Modeling/Formulations
Search (General/other)
SAT and CSP: Solvers and Tools
Knowledge Representation, Reasoning and Logic
Action, Change and Causality
Automated Reasoning and Theorem Proving
Belief Change
Common-Sense Reasoning
Computational Complexity of Reasoning
Description Logics and Ontologies
Diagnosis and Abductive Reasoning
Geometric, Spatial, and Temporal Reasoning
Knowledge Representation Languages
Knowledge Representation (General/other)
Logic Programming
Many-Valued And Fuzzy Logics
Nonmonotonic Reasoning
Preferences
Qualitative Reasoning
Reasoning with Beliefs
Machine Learning
Active Learning
Case-based Reasoning
Classification
Cost-Sensitive Learning
Data Mining
Ensemble Methods
Evolutionary Computation
Feature Selection/Construction
Kernel Methods
Learning Graphical Models
Learning Preferences/Rankings
Learning Theory
Machine Learning (General/other)
Neural Networks
Online Learning
Reinforcement Learning
Relational Learning
Time-series/Data Streams
Transfer, Adaptation, Multi-task Learning
Semi-Supervised/Unsupervised Learning
Structured Learning
Multidisciplinary Topics and Applications
AI and Natural Sciences
AI and Social Sciences
Art And Music
Autonomic Computing
Cognitive Modeling
Computational Biology
Computer Games
Computer-Aided Education
Database Systems
Philosophical and Ethical Issues
Human-Computer Interaction
Intelligent User Interfaces
Interactive Entertainment
Personalization and User Modeling
Real-Time Systems
Security and Privacy
Validation and Verification
Natural-Language Processing
Dialogue
Discourse
Information Extraction
Information Retrieval
Machine Translation
Morphology and Phonology
Natural Language Generation
Natural Language Semantics
Natural Language Summarization
Natural Language Syntax
Natural Language Processing (General/other)
Psycholinguistics
Question Answering
Speech Recognition And Understanding
Text Classification
Planning and Scheduling
Activity and Plan Recognition
Hybrid Systems
Markov Decisions Processes
Model-Based Reasoning
POMDPs
Plan Execution And Monitoring
Plan/Workflow Analysis
Planning Algorithms
Planning under Uncertainty
Planning (General/other)
Scheduling
Theoretical Foundations of Planning
Robotics and Vision
Behavior And Control
Cognitive Robotics
Human Robot Interaction
Localization, Mapping, State Estimation
Manipulation
Motion and Path Planning
Multi-Robot Systems
Robotics
Sensor Networks
Vision and Perception
Uncertainty in AI
Approximate Probabilistic Inference
Bayesian Networks
Decision/Utility Theory
Exact Probabilistic Inference
Graphical Models
Preference Elicitation
Sequential Decision Making
Uncertainty Representations
Uncertainty in AI (General/other)
Web and Knowledge-based Information Systems
Information Extraction
Information Integration
Information Retrieval
Knowledge Acquisition
Knowledge Engineering
Knowledge-based Systems (General/other)
Ontologies
Recommender Systems
Semantic Web
Social Networks
Source Wrapping
Web Mining
Web Search
Web Technologies (General/other)

“Artificial”

+

“Intelligence”

“Artificial”

?

Weak AI

Artificial Narrow Intelligence (ANI)

Strong AI

Artificial General Intelligence (AGI)

“Intelligence”

Interact/communicate using language and speech

Appropriately apply vast amounts of knowledge

Adaptive, goal-oriented behavior

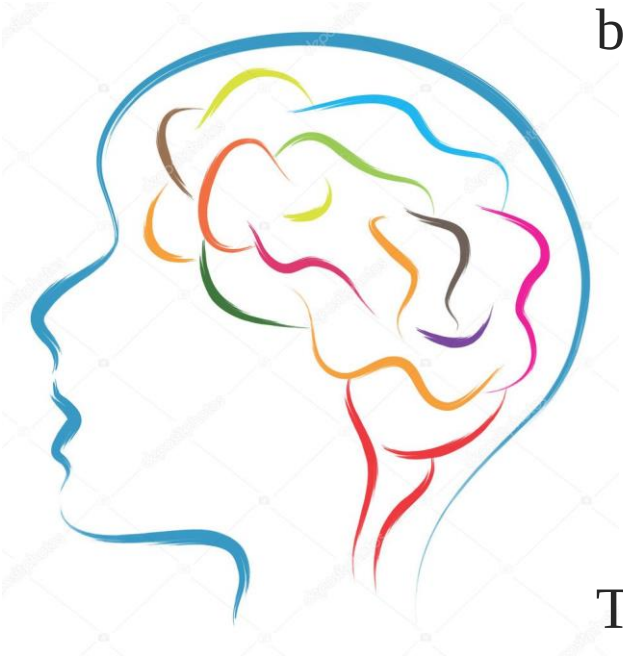

Self-awareness

Learn from experience

Tolerate error and ambiguity in the world; robustness

Timely response

Sound reasoning

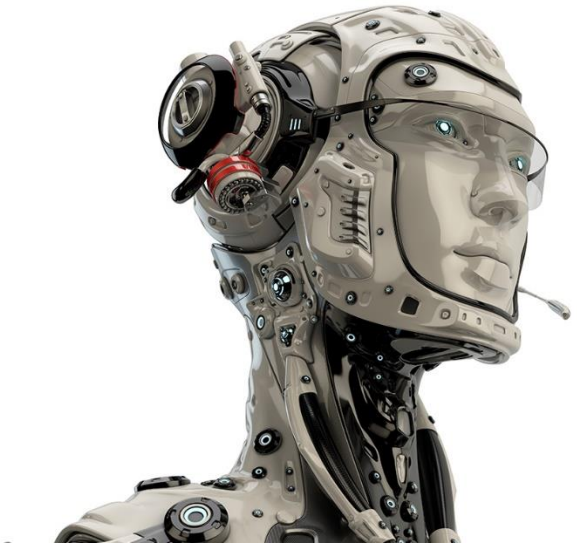

How could we know if a machine is intelligent?

Computing machinery and intelligence (Turing 1950)

I propose to consider the question, "Can machines think?" This should begin with definitions of the meaning of the terms "machine" and "think."

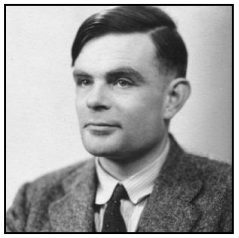

Instead of attempting such a definition I shall replace the question by another, which is closely related to it and is expressed in relatively unambiguous words.

The new form of the problem can be described in terms of a game which we call the "imitation game."

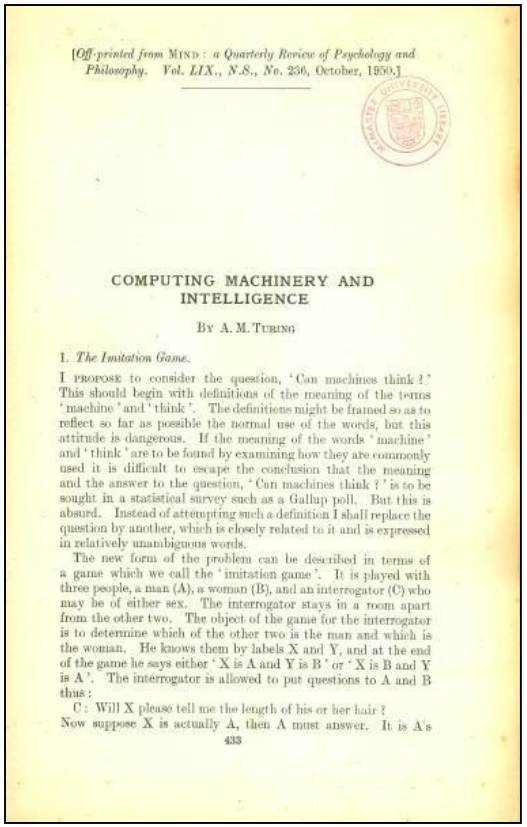

[*Off-printed from* MIND : *a Quarterly Review of Psychology and Philosophy. Vol. LIX., N.S., No. 236,* October, 1950.]

COMPUTING MACHINERY AND INTELLIGENCE

BY A. M. TURING

1. *The Imitation Game.*

I PROPOSE to consider the question, 'Can machines think?' This should begin with definitions of the meaning of the terms 'machine' and 'think'. The definitions might be framed so as to reflect so far as possible the normal use of the words, but this attitude is dangerous. If the meaning of the words 'machine' and 'think' are to be found by examining how they are commonly used it is difficult to escape the conclusion that the meaning and the answer to the question, 'Can machines think?' is to be sought in a statistical survey such as a Gallup poll. But this is absurd. Instead of attempting such a definition I shall replace the question by another, which is closely related to it and is expressed in relatively unambiguous words.

The new form of the problem can be described in terms of a game which we call the 'imitation game'. It is played with three people, a man (A), a woman (B), and an interrogator (C) who may be of either sex. The interrogator stays in a room apart from the other two. The object of the game for the interrogator is to determine which of the other two is the man and which is the woman. He knows them by labels X and Y, and at the end of the game he says either 'X is A and Y is B' or 'X is B and Y is A'. The interrogator is allowed to put questions to A and B thus:

C: Will X please tell me the length of his or her hair?

Now suppose X is actually A, then A must answer. It is A's

The Imitation Game

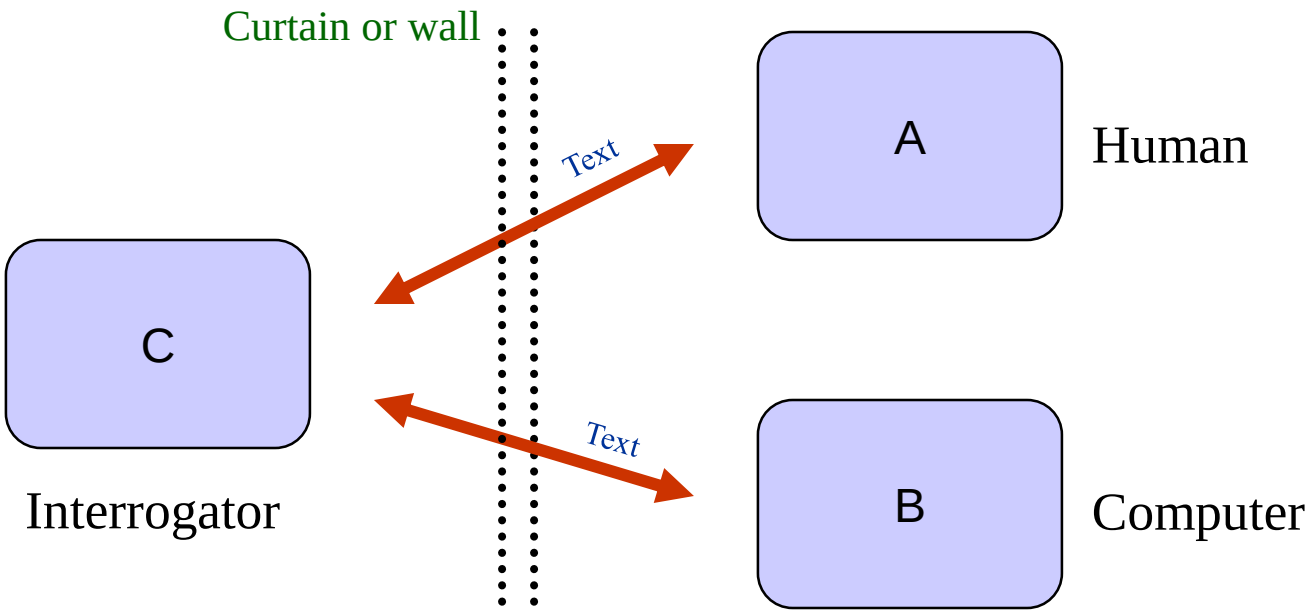

Q: Please write me a sonnet on the subject of the Forth Bridge.

A : Count me out on this one. I never could write poetry.

Q: Add 34957 to 70764.

A: (Pause about 30 seconds and then give as answer) 105621.

Q: Do you play chess?

A: Yes.

Q: I have K at my K1, and no other pieces. You have only K at K6 and R at R1. It is your move. What do you play?

A: (After a pause of 15 seconds) R-R8 mate.

Computing machinery and intelligence (Turing 1950)

I believe that in about fifty years' time it will be possible, to programme computers, with a storage capacity of about 10^9, to make them play the imitation game so well that an average interrogator will not have more than 70 per cent chance of making the right identification after five minutes of questioning.

The original question, "Can machines think?" I believe to be too meaningless to deserve discussion. Nevertheless I believe that at the end of the century the use of words and general educated opinion will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted.

Computers in 1950

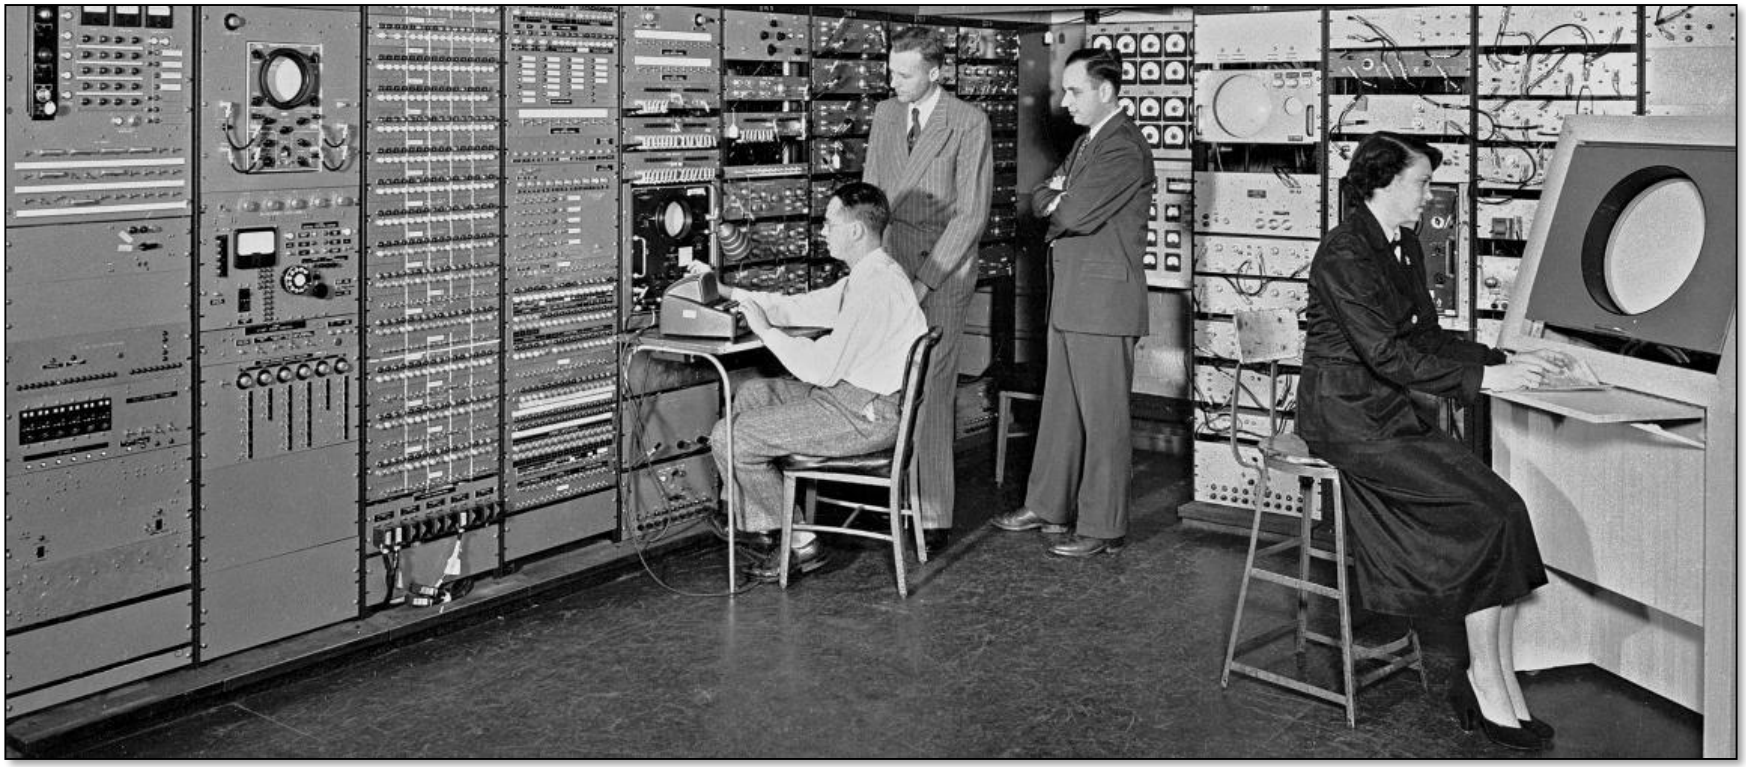

…and before that

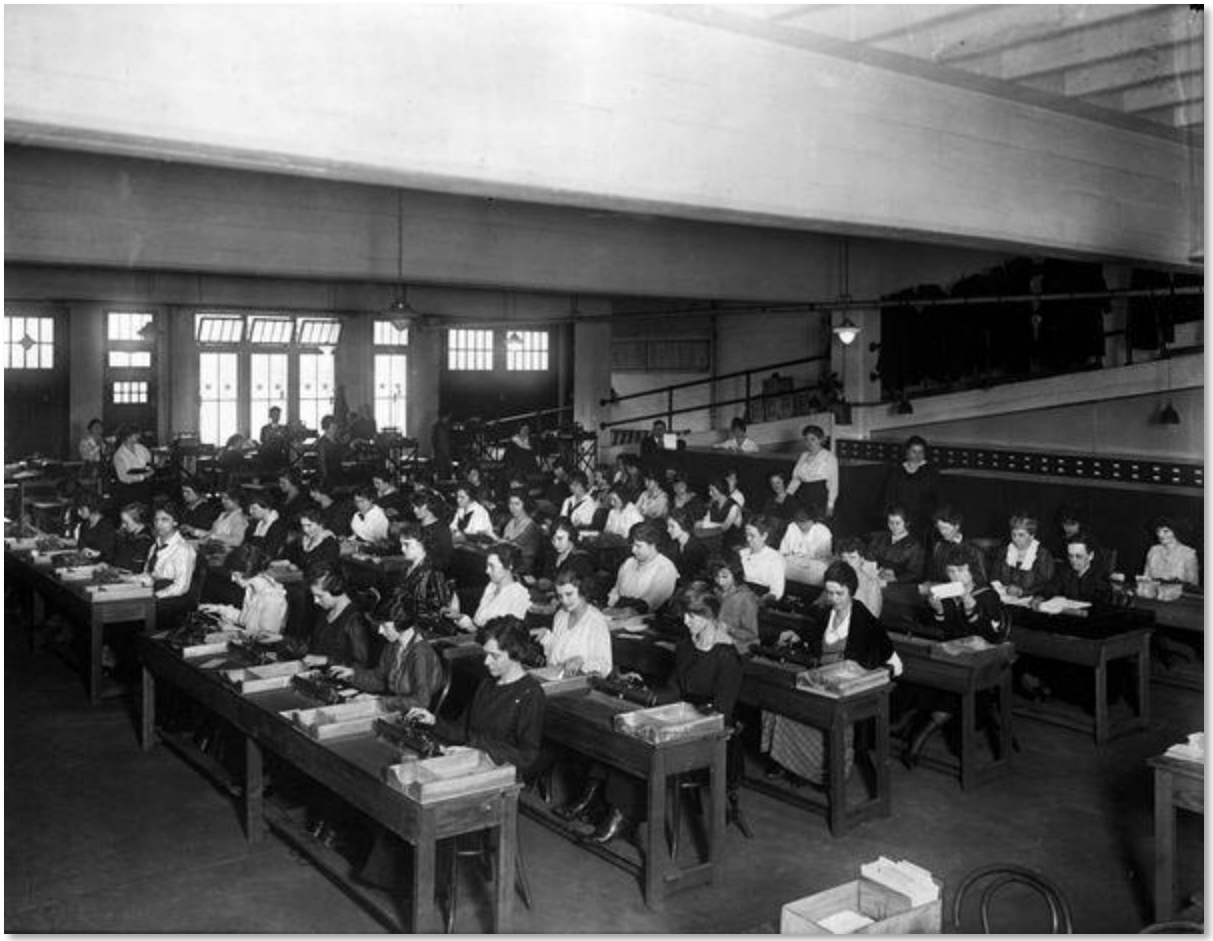

Turing paper (cont.)

The Turing Test

"Are there imaginable digital computers which would do well in the imitation game?"

I.e., Can a computer fool an interrogator into thinking it is a person?

Properties of the Turing Test

Operational/functional/behavioral definition of intelligence

Distinguishes between physical and intellectual capacities

Question and answer method – language comprehension and generation

Might there be other kinds of Turing Tests?

Emotional, physical, visual…

Different perspectives on AI

	Human	Ideal
Thought processes and reasoning	Systems that think like humans *Cognitive modeling*	Systems that think rationally *Laws of thought, logic*
Behavior	Systems that act like humans *Operational intelligence*	Systems that act rationally *Optimized performance*

“Strong AI” vs. “Weak AI”

“Strong AI” (Artificial General Intelligence)

- Goal: To understand the principles and mechanisms that account for intelligent action
- Makes the bold claim that computers can be made to think on a level (at least) equal to humans. [Intelligence = information processing]

“Weak AI” (Artificial Narrow Intelligence)

- Goal: To design intelligent systems that can survive and operate in the real world and solve problems of considerable scientific difficulty at high levels of competence
- Some “thinking-like” features can be added to computers to make them more useful tools
- Examples: expert systems, speech recognition, natural language understanding, recommender systems, etc.

“Strong AI” vs. “Weak AI” (cont.)

Principles of Strong AI

- Intelligent behavior is explicable in scientific terms; a rigorous understanding of intelligence is possible
- Intelligence can take place outside the human skull
- Intelligence = symbol manipulation (perhaps grounded in perception and action)
- The computer is the best laboratory instrument for exploring these propositions

Principle of Weak AI

- Who cares??

AI and human intelligence

AI (even Strong AI) doesn't *necessarily* seek to replicate human intelligence

- Sometimes more, sometimes less…

Creating the "essence of X" vs. creating "X"

Examples

- Physical vs. electronic newspaper
- Physical vs. virtual shopping
- Birds vs. planes



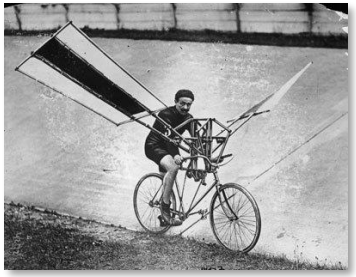


What is an AI Program?

Classical AI programs can generally be thought of as comprising three *separated* parts

- Data / knowledge ("knowledge base")
- Operations / rules ("production rules")
- Control: Which rules to apply when

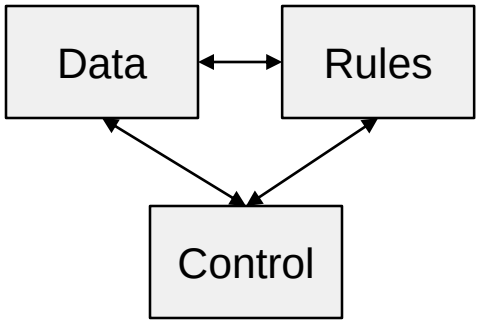

Data and *Operations* should be modular and easy to modify

On the other end of the spectrum, we have ***neural networks*** and specifically ***deep learning***

- Input/output black box!

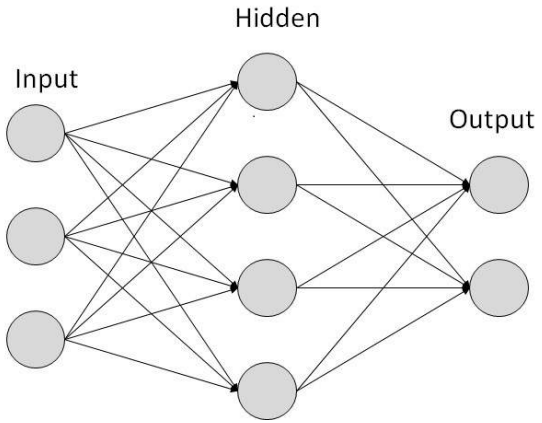

Classical vs. Deep Learning approaches

Classical AI

- Think about the world, interview domain experts
- Write down rules that encapsulate intelligent behavior
- Hope that these adequately cover the range of real-world situations

Deductive reasoning

NN / Deep Learning

- Collect massive amounts of data
- Provide an architecture for learning (primed with some goals)
- Hope that the data can be abstracted into meaningful general concepts

Inductive reasoning

How do people tend to reason?

Deep Learning

Deep Learning: Machine learning-based AI motivated by the structure and function of the brain, characterized by huge amounts of data, neural network architectures, advanced learning algorithms, and fast computers

- Learning hierarchies of features, factors, concepts
- Extracting multiple layers of representation from data

Or… neural networks with many layers

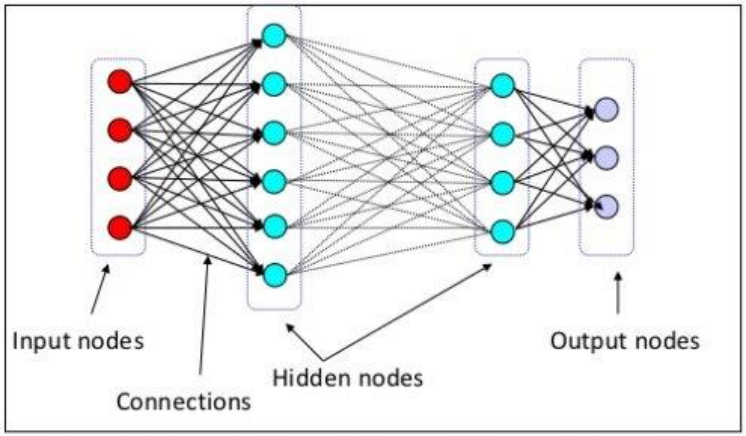

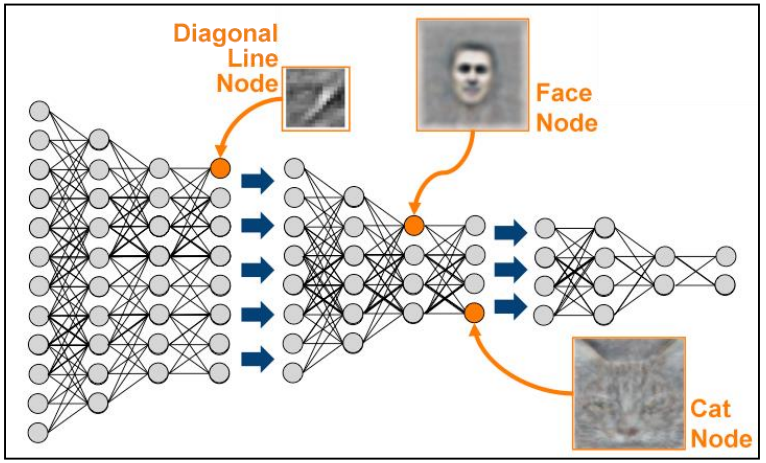

The hype cycle of AI

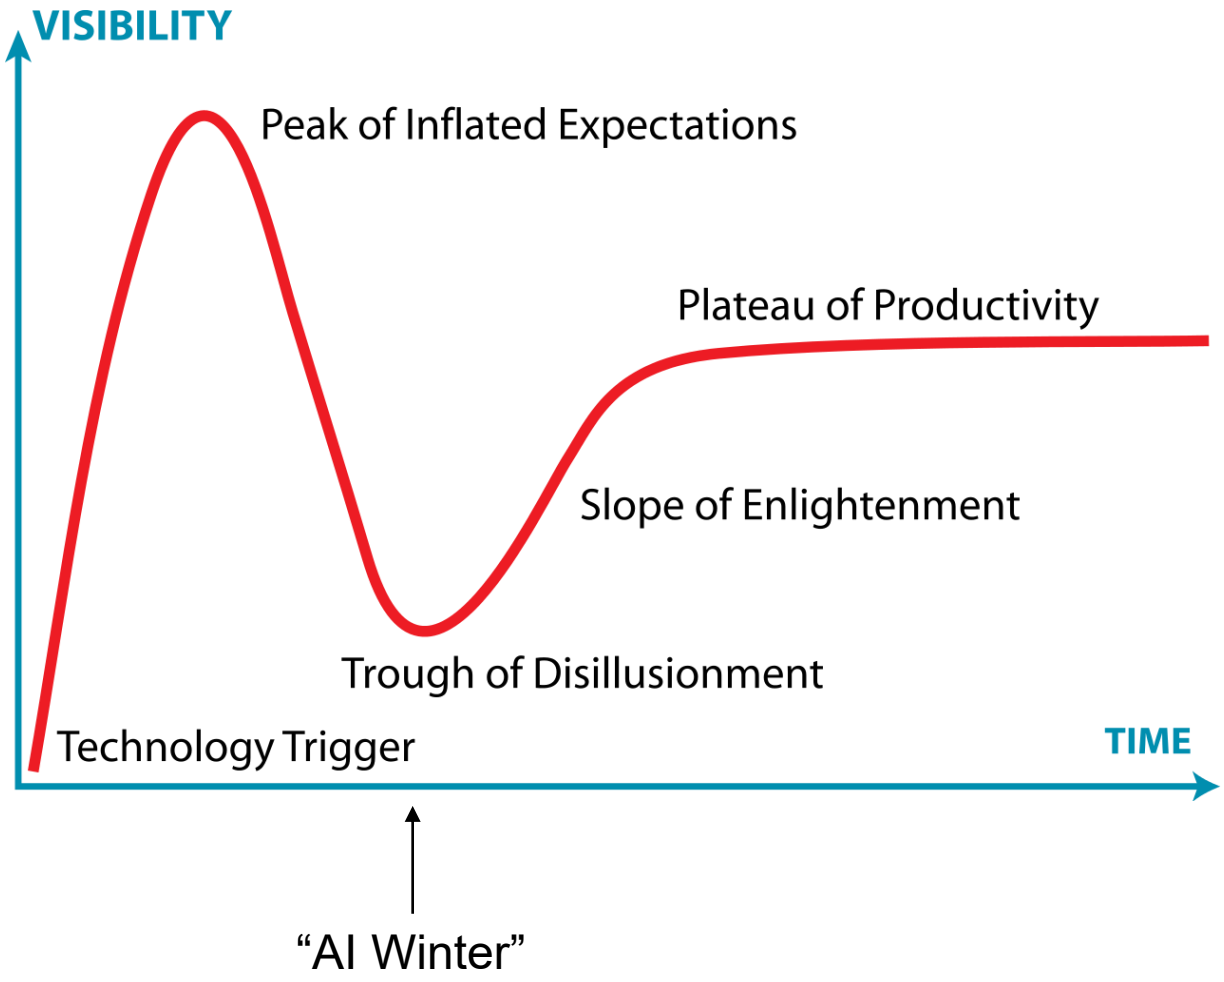

Some notable AI systems

IBM Deep Blue

- Beat world chess champion Gary Kasparov in 1997

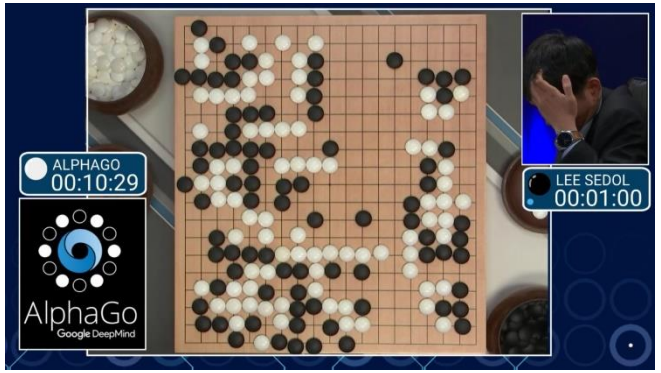

Google AlphaGo

- Defeated world Go champion in 2017

Expert systems

- Medical diagnosis
  - E.g., Leukemia diagnosis system

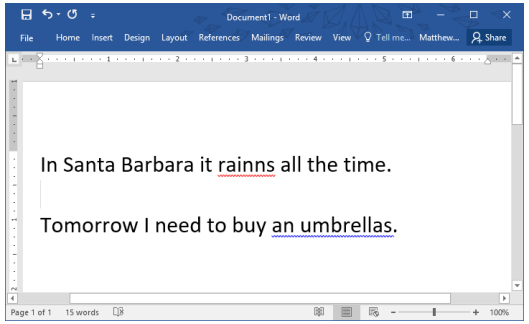

User interface

- Grammar and spelling checkers, automated help

Speech & language understanding & translation

- Web search (Google, Bing)
- Phone-based systems (e.g., airline reservations)
- Siri, Google Now, Amazon Alexa

Some notable AI systems (cont.)

IBM Watson

2011, an experiment in “Deep Q&A”

Recommender systems

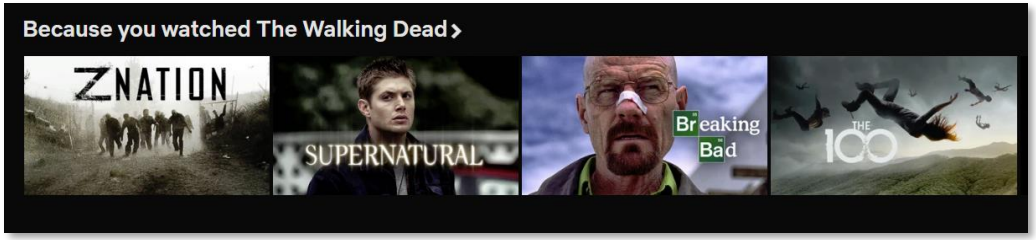

Netflix (Amazon, Pandora, etc.)

Skype Voice Translator

Some notable AI systems (cont.)

Data mining

- Fraud detection, credit scoring, customer profiles and preferences, genome analysis

Computer vision

- Face detection and recognition systems

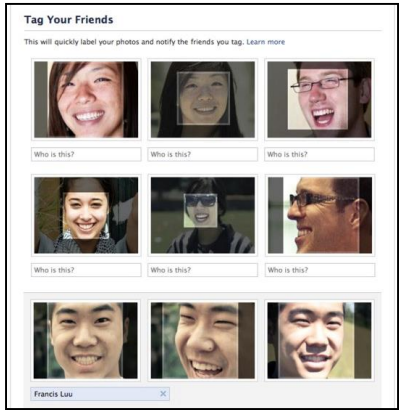

Robotics

- Mars Rover, robots for hazard environments, factory automation
- Sony, Honda, others: robot pets
- 1980s – DARPA ALV Program
- CMU Navlab drove across country (2797/2849 miles)
- E.g., “Hands Across America” 1995
- DARPA Grand Challenge, Urban Challenge
- Autonomous cars – Google, Tesla, Toyota, GM, …
- UAVs – unmanned aerial vehicles (drones)

Reasoning, learning, planning

Cognition

Computer vision

Vision

Hearing

Audio understanding

Speech & language

Natural language processing

Embodied Intelligence

Emotion

Affective computing

Manipulation

Touch

Haptics

Robotics

Locomotion

CMU Leg Lab 1983

Boston Dynamics 2018

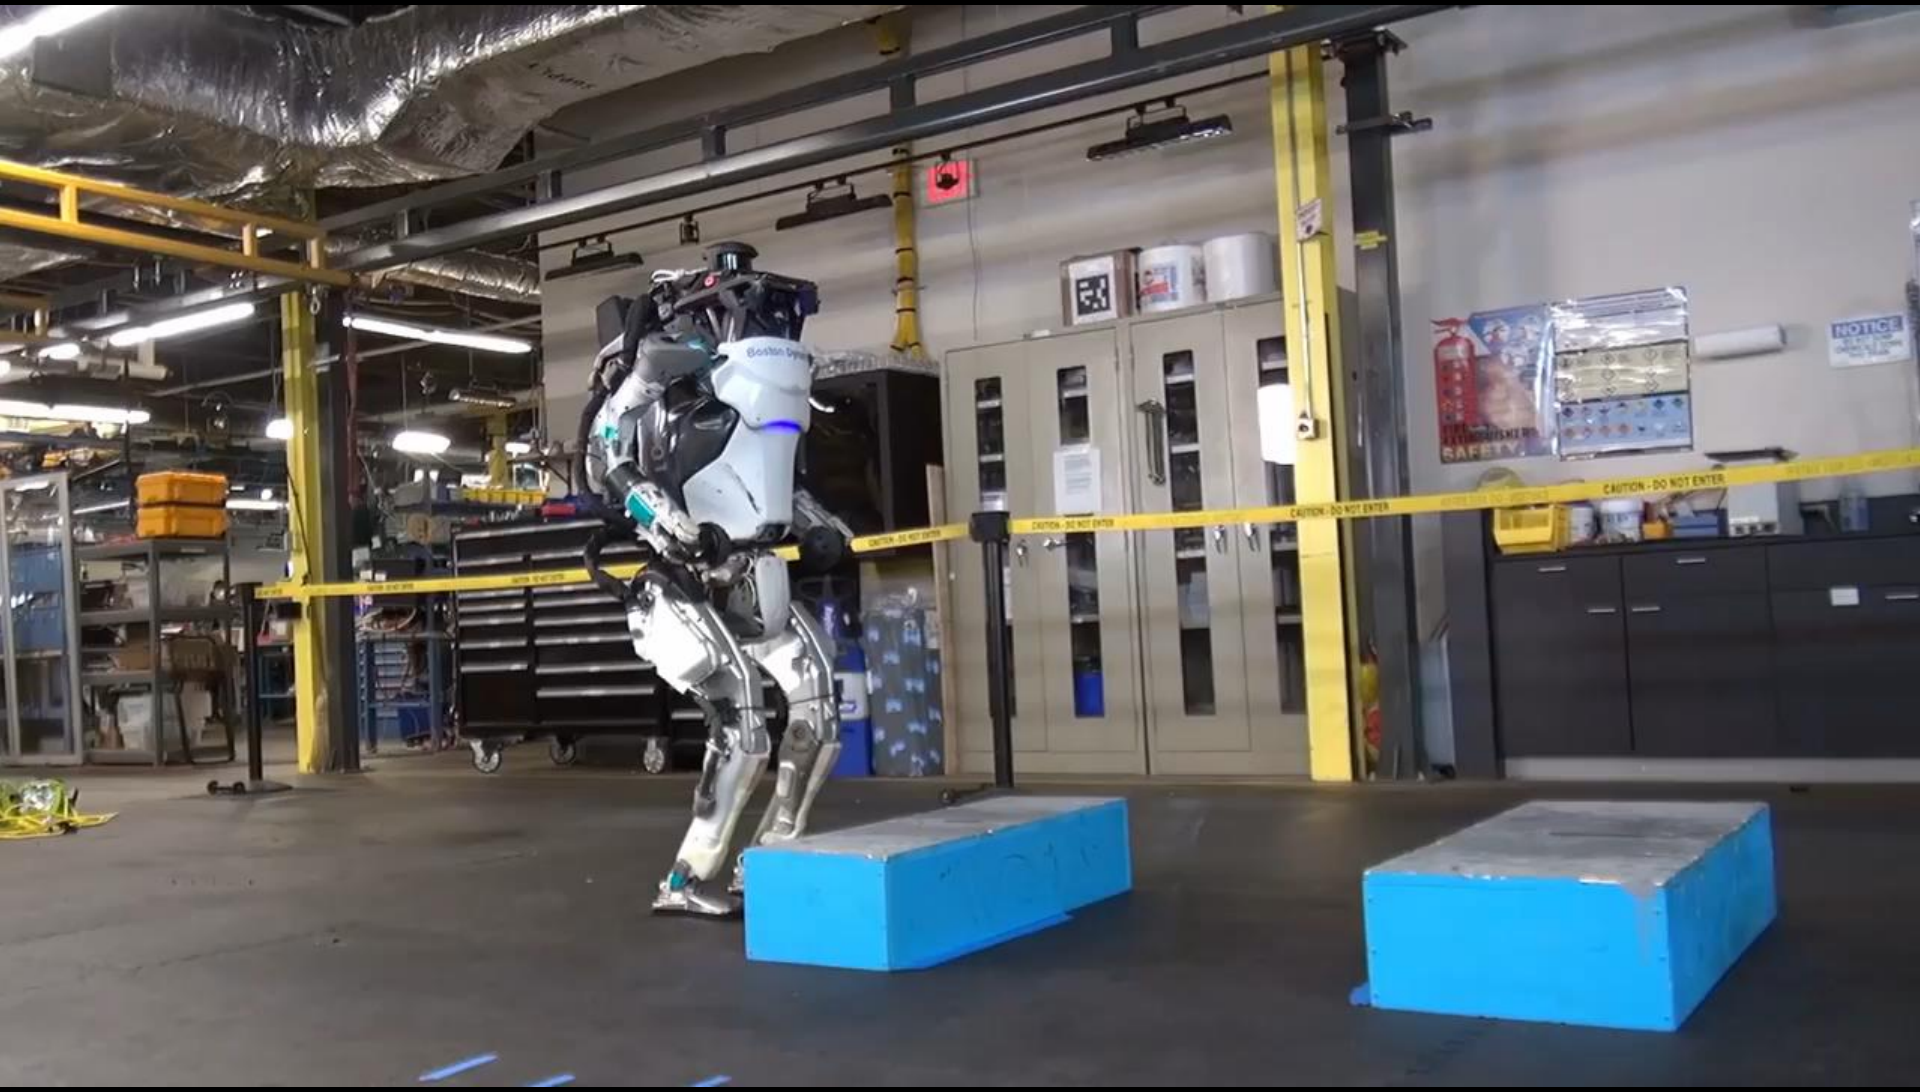
CAUTION - DO NOT ENTER
NOTICE

iRobot Roomba Pet Series

Vacuum Cleaning Robot

The smarter way to get it done.

So where’s it going?

The state of AI

After a long history of ups and downs, AI is now coming of age, with technical successes, huge investments, and practical benefits increasing daily.

- Data, algorithms, and computational power are enabling rapid advances in machine learning, which is being applied widely in every industry.
- Sensing and perception technologies are enabling safer driving and autonomous vehicles.
- AI systems are increasing productivity in the workplace.
- AI is improving medical diagnosis and healthcare.
- Automatic translation systems are enabling cross-cultural communication, large-scale document translation, and real-time counter-terrorism.
- Better web search, product recommendations, fraud detection, spam filtering, driving directions, computer security, drug discovery, genetics, etc. etc. etc.


© copyright Beyond Theory September 2016 - all rights reserved

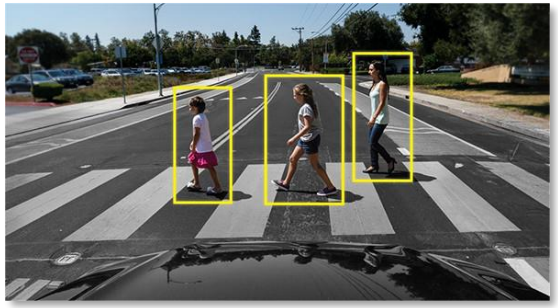


UBER
ADVANCED
TECHNOLOGIES
CENTER
UBERATC.COM/CAR

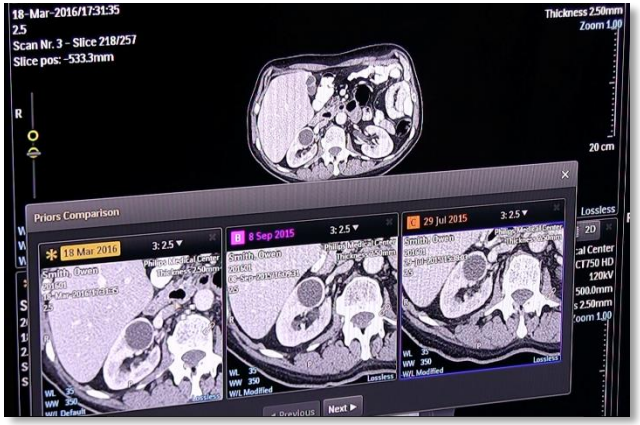
18-Mar-2016/17:31:35
2.5
Scan Nr. 3 - Slice 218/257
Slice pos: -533.3mm
Thickness 2.50mm
Zoom 1.00
20 cm
Priors Comparison
18 Mar 2016
8 Sep 2015
29 Jul 2015
Lossless
Next

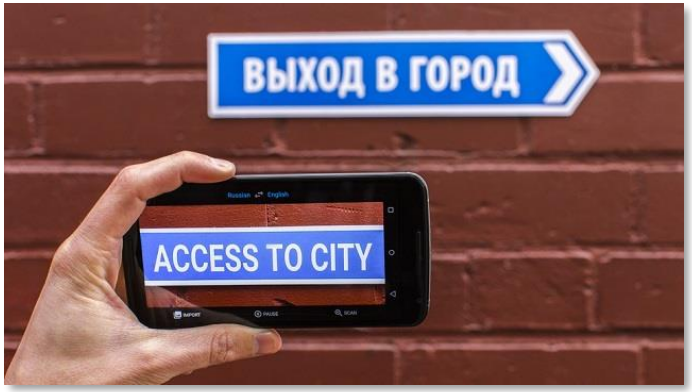
ВЫХОД В ГОРОД
ACCESS TO CITY

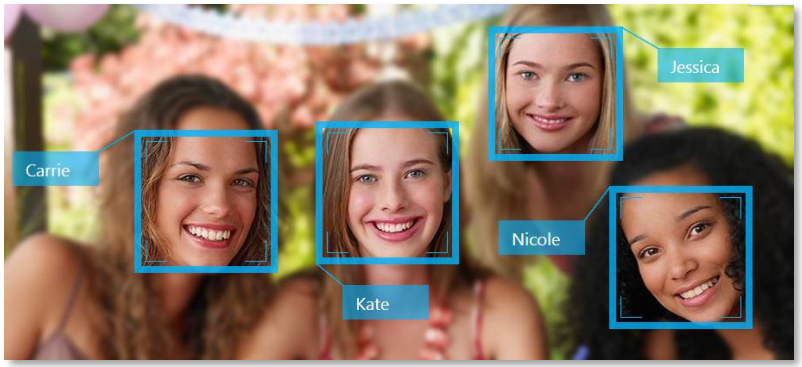
Jessica
Carrie
Nicole
Kate


amazon

The driver: Machine learning

Rapid advances in machine learning are largely responsible for the accelerated advances and optimism in AI.

Especially in deep learning and applications in visual detection and recognition, speech processing, and natural language processing

Why?

1. Better training algorithms
2. Better machines (especially GPUs)
3. More data

Coming soon…

Autonomous cars and trucks

More accurate radiology diagnosis

Better computer security – e.g., malware detection

Physical security and surveillance

Detection of fake news

Robot pets and companions

Search and rescue robots

AI-based tutoring and training

Artificial characters: entertainment and gaming

Document analysis and powerful information retrieval tools

Personalized everything

Continued improvements in productivity, efficiency, and accuracy across an organization

AI will transform…

Transportation

Healthcare

Finance

Energy

Agriculture

Manufacturing

Art and entertainment

Amara's Law: We tend to overestimate the effect of a technology in the short run and underestimate the effect in the long run.

Every industry is now investing heavily in AI!

TED Ideas worth spreading

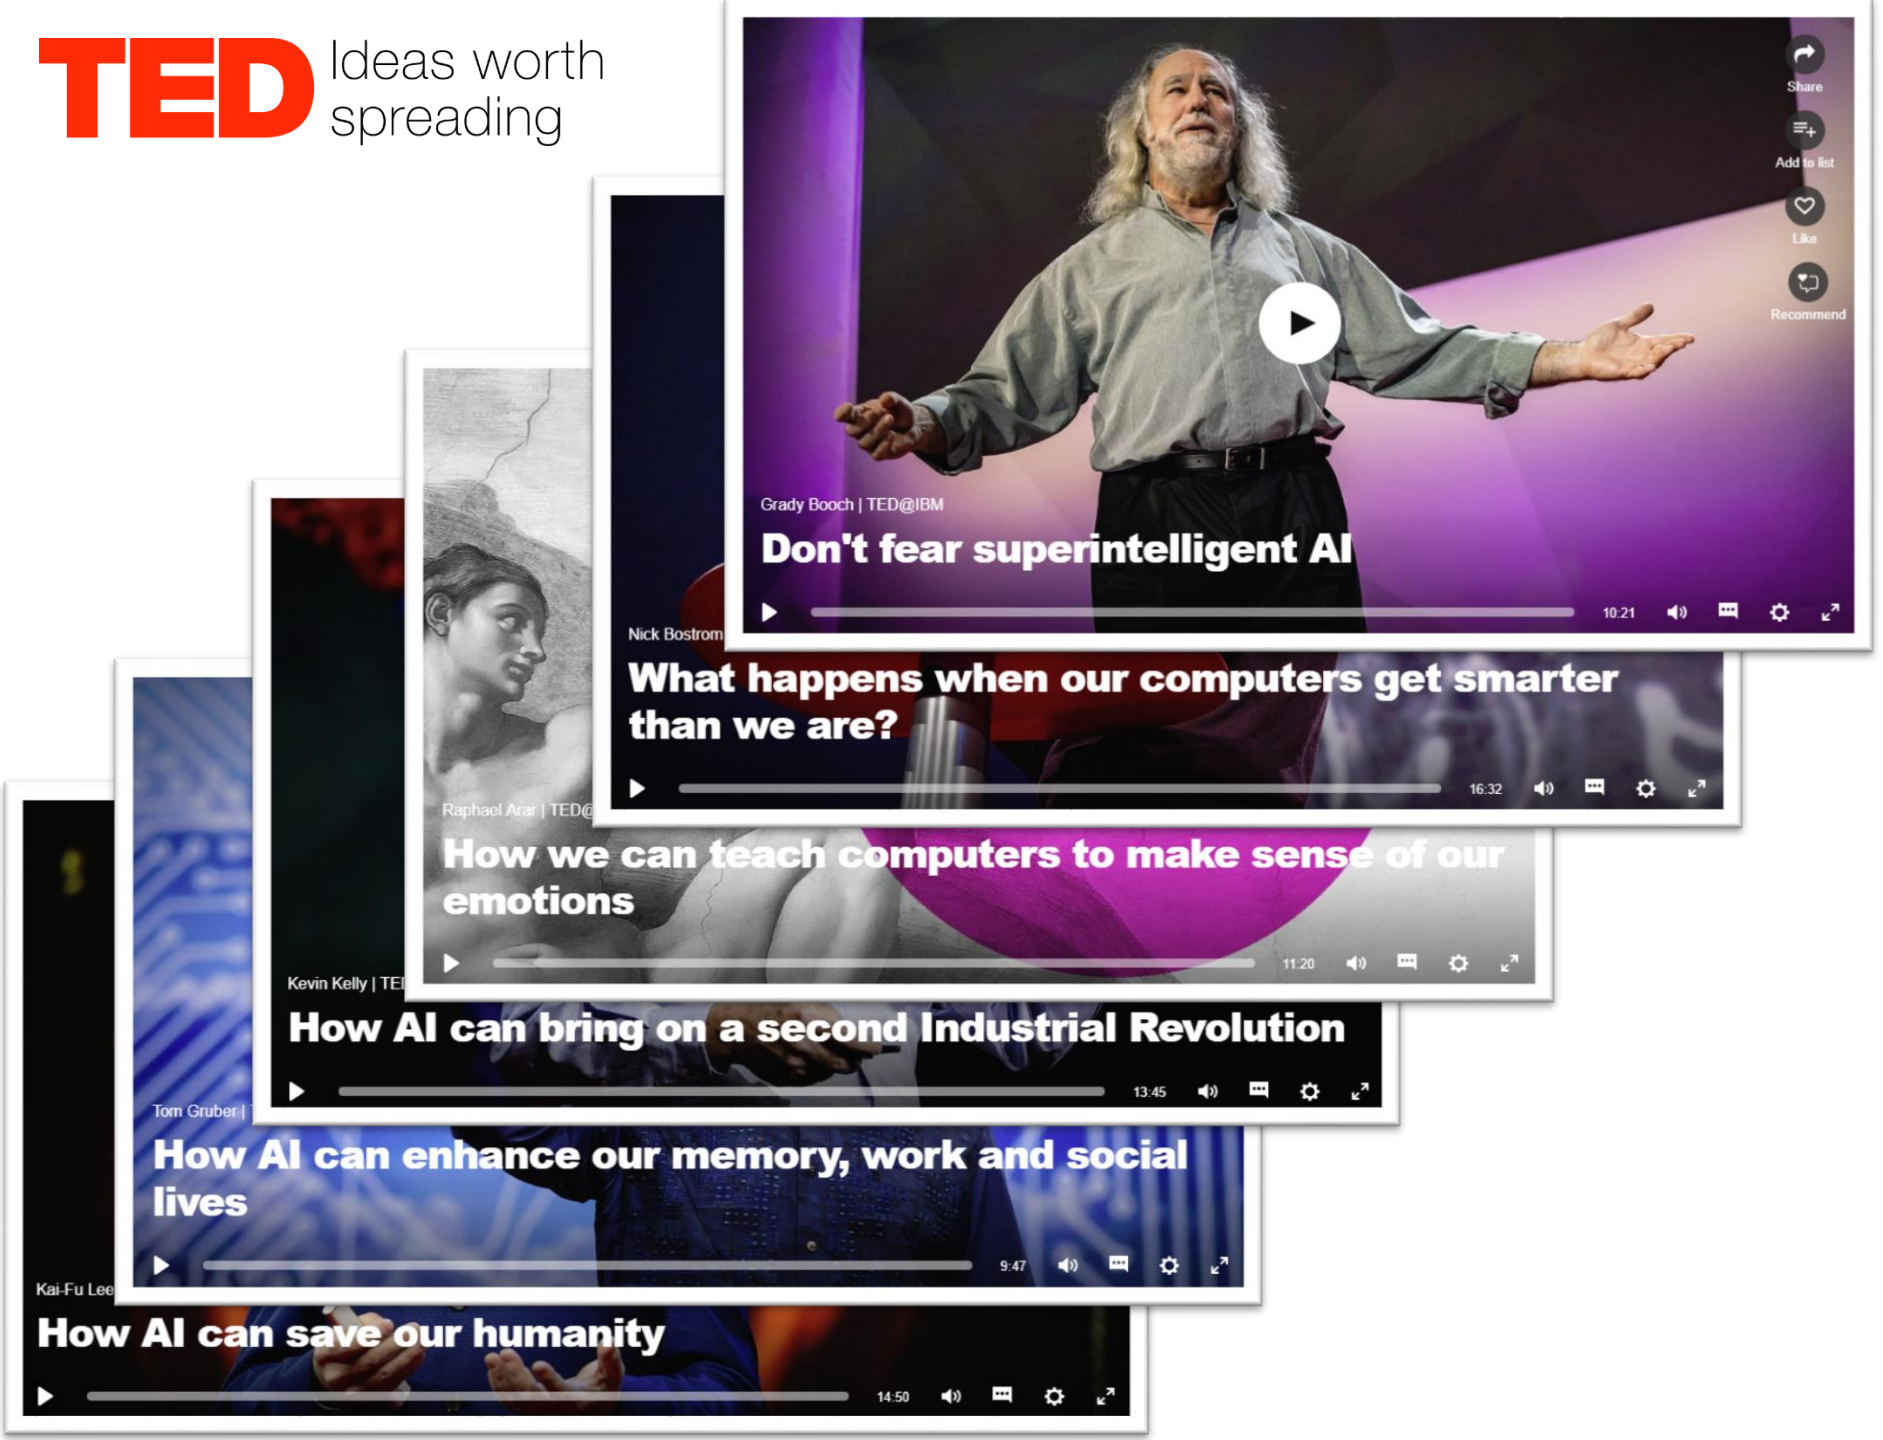

Newsweek

TECH & SCIENCE

STEPHEN HAWKING AI WARNING: ARTIFICIAL INTELLIGENCE COULD DESTROY CIVILIZATION

BY HANNAH OSBORNE ON 11/7/17 AT 4:43 AM

Stephen Hawking sits onstage during an announcement of the Breakthrough Starshot initiative with investor Yuri Milner in New York City, on April 12, 2016. Hawking, the English physicist, warns humanity needs to become a multiplanetary species to ensure its survival.

REUTERS/LUCAS JACKSON

From the Magazine

ELON MUSK'S BILLION-DOLLAR CRUSADE TO STOP THE A.I. APOCALYPSE

Elon Musk is famous for his futuristic gambles, but Silicon Valley's latest rush to embrace artificial intelligence scares him. And he thinks you should be frightened too. Inside his efforts to influence the rapidly advancing field and its proponents, and to save humanity from machine-learning overlords.

Perhaps we should learn from this old Twilight Zone episode?

Some issues to consider

Unemployment. What happens after the end of jobs?

Inequality. How do we distribute the wealth created by machines?

Access. How do we share access to powerful tools?

Humanity. How do machines affect our behavior and interaction?

Artificial stupidity. How can we guard against mistakes?

Racist robots. How do we eliminate AI bias?

Security. How do we keep AI safe from adversaries?

Explainability. How do we understand, evaluate, and trust AI decisions?

Evil genies. How do we protect against unintended consequences?

Robot rights. How do we define the humane treatment of AI?

The new Manhattan Project is not AI R&D – it's AI that is safe, responsible, accessible, and ethical.

http://www.cs.ucsb.edu/~mturk

http://www.cs.ucsb.edu