

History and Definitions

CPSC 433: Artificial Intelligence
Fall 2024

Jonathan Hudson, Ph.D.
Assistant Professor (Teaching)
Department of Computer Science
University of Calgary

August 8, 2024

Copyright © 2024



Outline

- History of field of AI
- Connectionist vs. Symbolic
- The state of the art

Standing on the shoulders of ?


(not AI quite yet but maybe soon)

Origin of AI

- The concept of a “**thinking machine**” can be dated back to ancient times, when Greek myths and legends portrayed mechanical beings endowed with intelligence.
- 17th century philosopher **René Descartes** compared the human brain with a machine
 - debated that mathematics can explain how the human brain operates.

Pre-history

Philosophy	logic, methods of reasoning mind as physical system foundations of learning, language, rationality
Mathematics	formal representation and proof algorithms, computation, (un)decidability, (in)tractability Probability
Psychology	adaptation phenomena of perception and motor control experimental techniques (psychophysics, etc.)
Economics	formal theory of rational decisions
Linguistics	knowledge representation Grammar
Neuroscience	plastic physical substrate for mental activity
Control theory	homeostatic systems, stability simple optimal agent designs




“The Analytical Engine has no pretensions whatever to originate anything. It can do whatever we know how to order it to perform. It can follow analysis; but it has no power of anticipating any analytical relations or truths.”

Ada Lovelace [Mid 1800s]

(First computing AI philosopher?)

Ride the wave

- 1940-1950: Early days
 - 1943: McCulloch & Pitts: Boolean circuit model of brain
 - **1950: Turing's "Computing Machinery and Intelligence"** (To a degree a rebuttal of or reframing of disagreement with Lovelace's statement)
https://books.google.com/books?id=CEMYUU_HFMAC&pg=PA67
- 1950—74: Excitement: Look, Ma, no hands!
 - 1950s: Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
 - **1956: Dartmouth Conference: "Artificial Intelligence" adopted**
 - 1965: Robinson's complete algorithm for logical reasoning
 - 1969-79: Early development of knowledge-based systems



GPS



no not that one

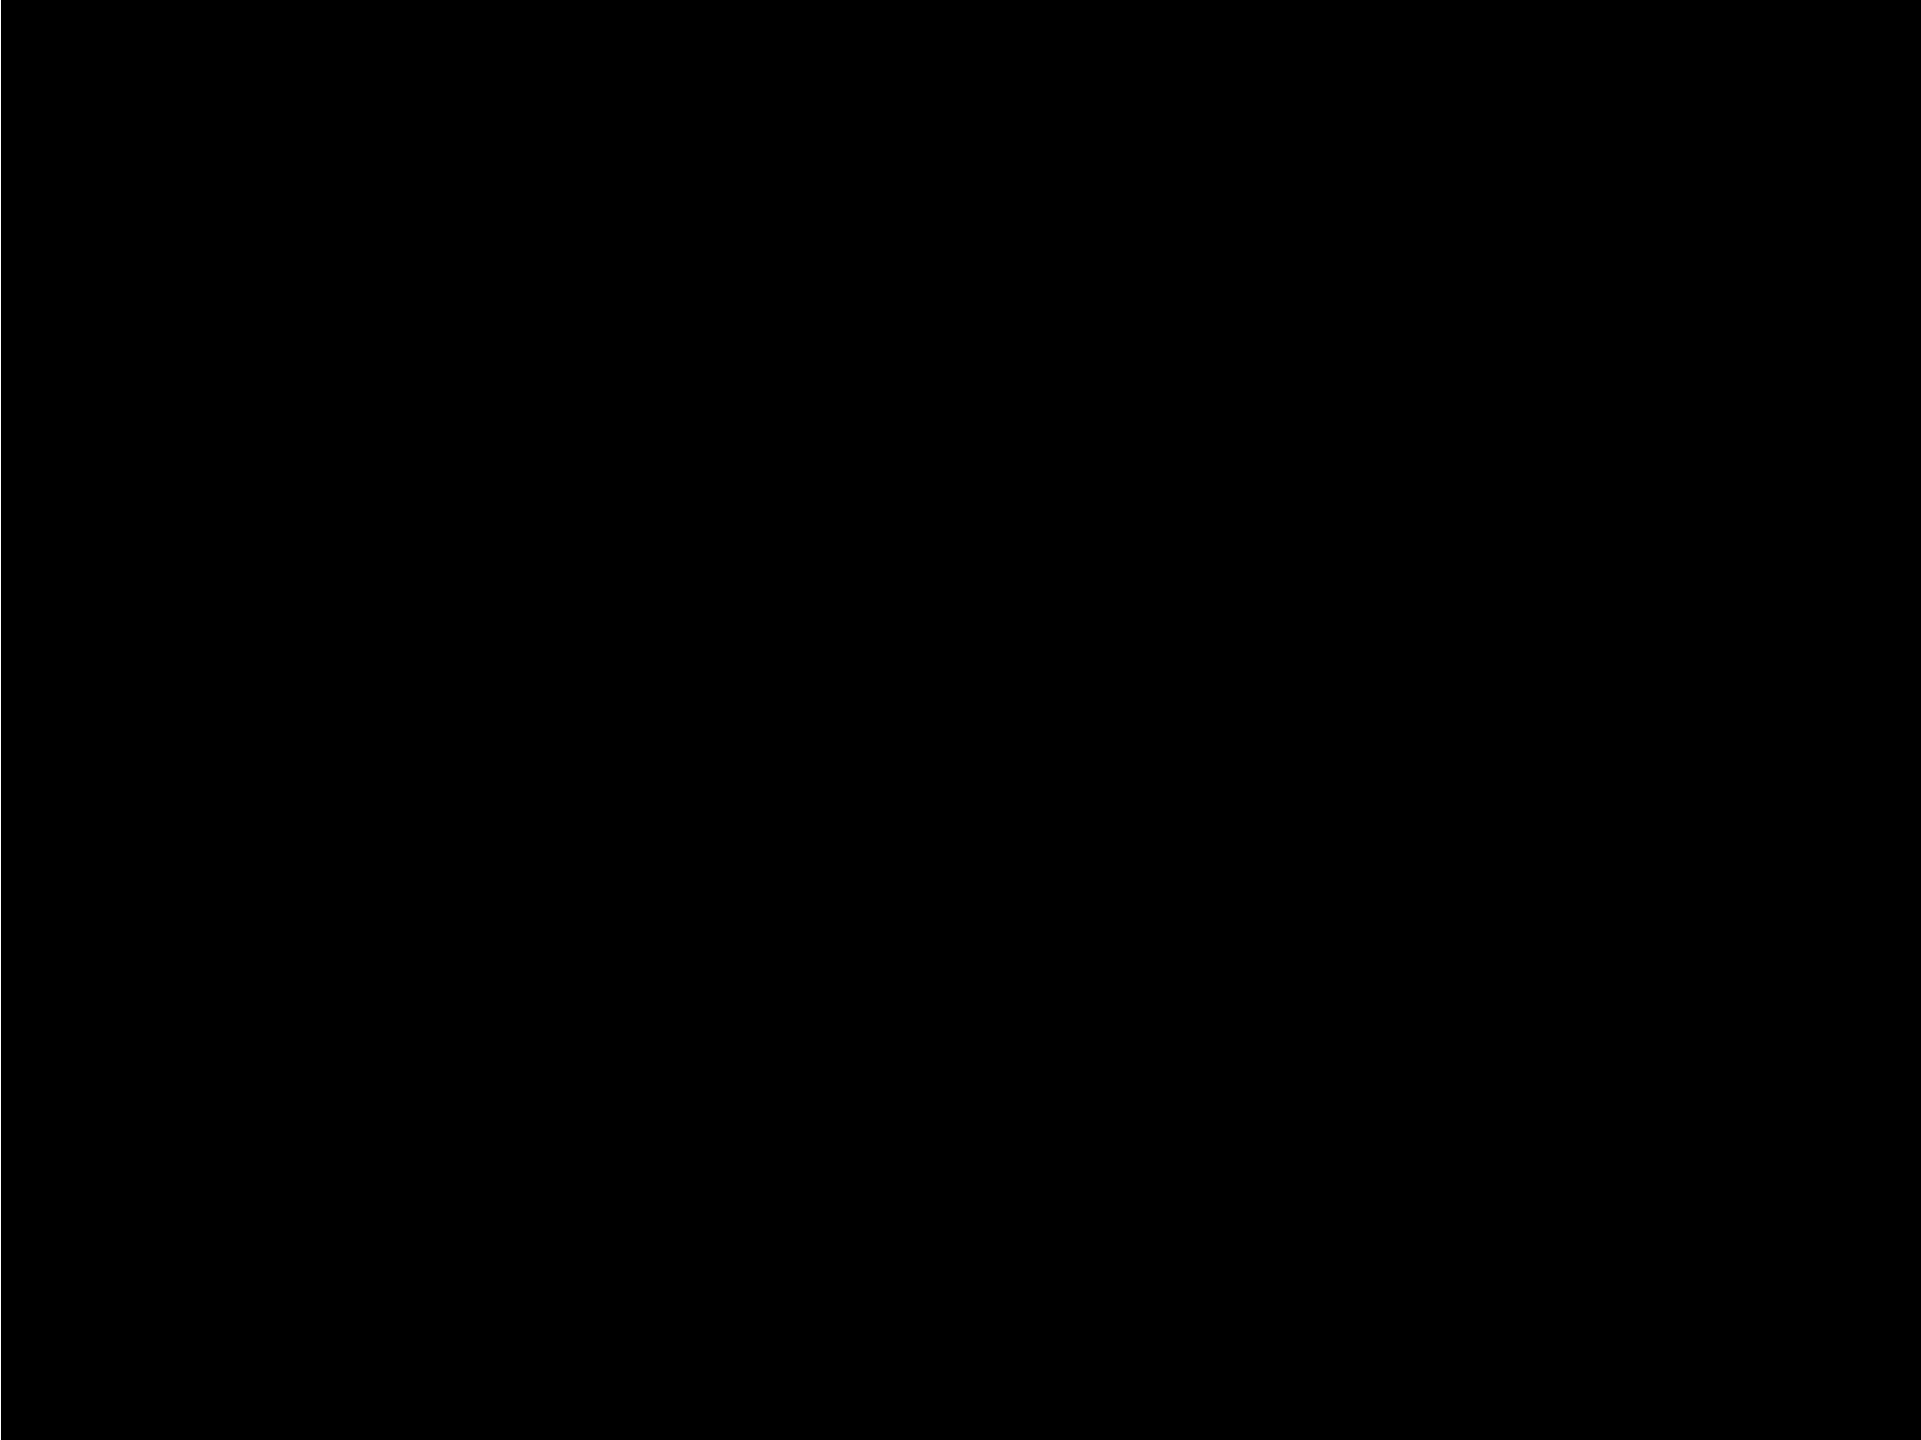
General Purpose Problem Solvers

Goals: Develop computer programs that show an **intelligent behavior** and that can serve as **companion** to human beings.

Example: the **GPS** program (**G**eneral **P**roblem **S**olver) “describe problem in logic and it will solve it using logical inferences”

Result: **Complete failure** and enormous **negative reaction** in media (and elsewhere)

Problems: One **knowledge representation** scheme with one **knowledge processing** mechanism not enough + **search spaces** are enormous.




AI Winter



We're number one!

Failure and New Hopes

- 1974—80: AI Winter I
- 1970—90: Knowledge-based approaches
 - 1969—79: Early development of knowledge-based systems
 - 1980—88: **Expert systems industry booms**



Bubbling

Goals: Select a **field** in AI and an application area and “make it work”

Examples:

- Natural language understanding and dictating
- Expert systems for configuration and diagnosis in small application areas
- Playing a game
- Scoring high in intelligence tests

Expert Systems: We did it mom

Result: It (**expert systems**) worked **within the expectations**

Expert systems – develop one knowledge representation and solution for one specific problem at a time

Problems:


- Very different methods used in different fields
- Application areas very **specialized**
- Systems resemble “**idiot savante**” (person without any interest outside his/her field)
- Hard problems still a problem

AI Winter²

We're number two!


Failure Mark II and but an awakening

- AI Winter II
 - 1987—93: Expert systems industry busts: “AI Winter”
- 1990—: Statistical approaches
 - Resurgence of probability, focus on uncertainty
 - **Connectionist ‘Machine Learning’ ascendance**
 - General increase in technical depth
 - **Agents and learning systems... “AI Spring”**
 - Willingness to use AI in marketing again (usually machine learning or deep learning instead)



Artificial Intelligence (AI) in Computer Science

AI HAS A LONG HISTORY OF BEING “THE NEXT BIG THING”...



Timeline of AI Development

- **1950s-1960s:** First AI boom - the age of reasoning, prototype AI developed
- **1970s:** AI winter I
- **1980s-1990s:** Second AI boom: the age of Knowledge representation (appearance of expert systems capable of reproducing human decision-making)
- **1990s:** AI winter II
- **1997:** Deep Blue beats Gary Kasparov
- **2006:** University of Toronto develops Deep Learning
- **2011:** IBM's Watson won Jeopardy
- **2016:** Go software based on Deep Learning beats world's champions

Connectionist vs Symbolic


Connecting the symbols?

Symbolic AI

- Top-down method of intelligence formulation
- Look at how human solves problem at the top-level
- try to form symbols and symbolic structures
- to break down the problem until you can write program a computer to do it


Connectionist?

- Inspired by the Human Brain.
- The human brain has about 86 Billion neurons and requires 20% of your body's energy to function.
- These neurons are connected to between 100 Trillion to 1 Quadrillion synapses!
- Deep learning neural networks really popular right now
- Connectionist method
 - Make network, train, hope result is useful

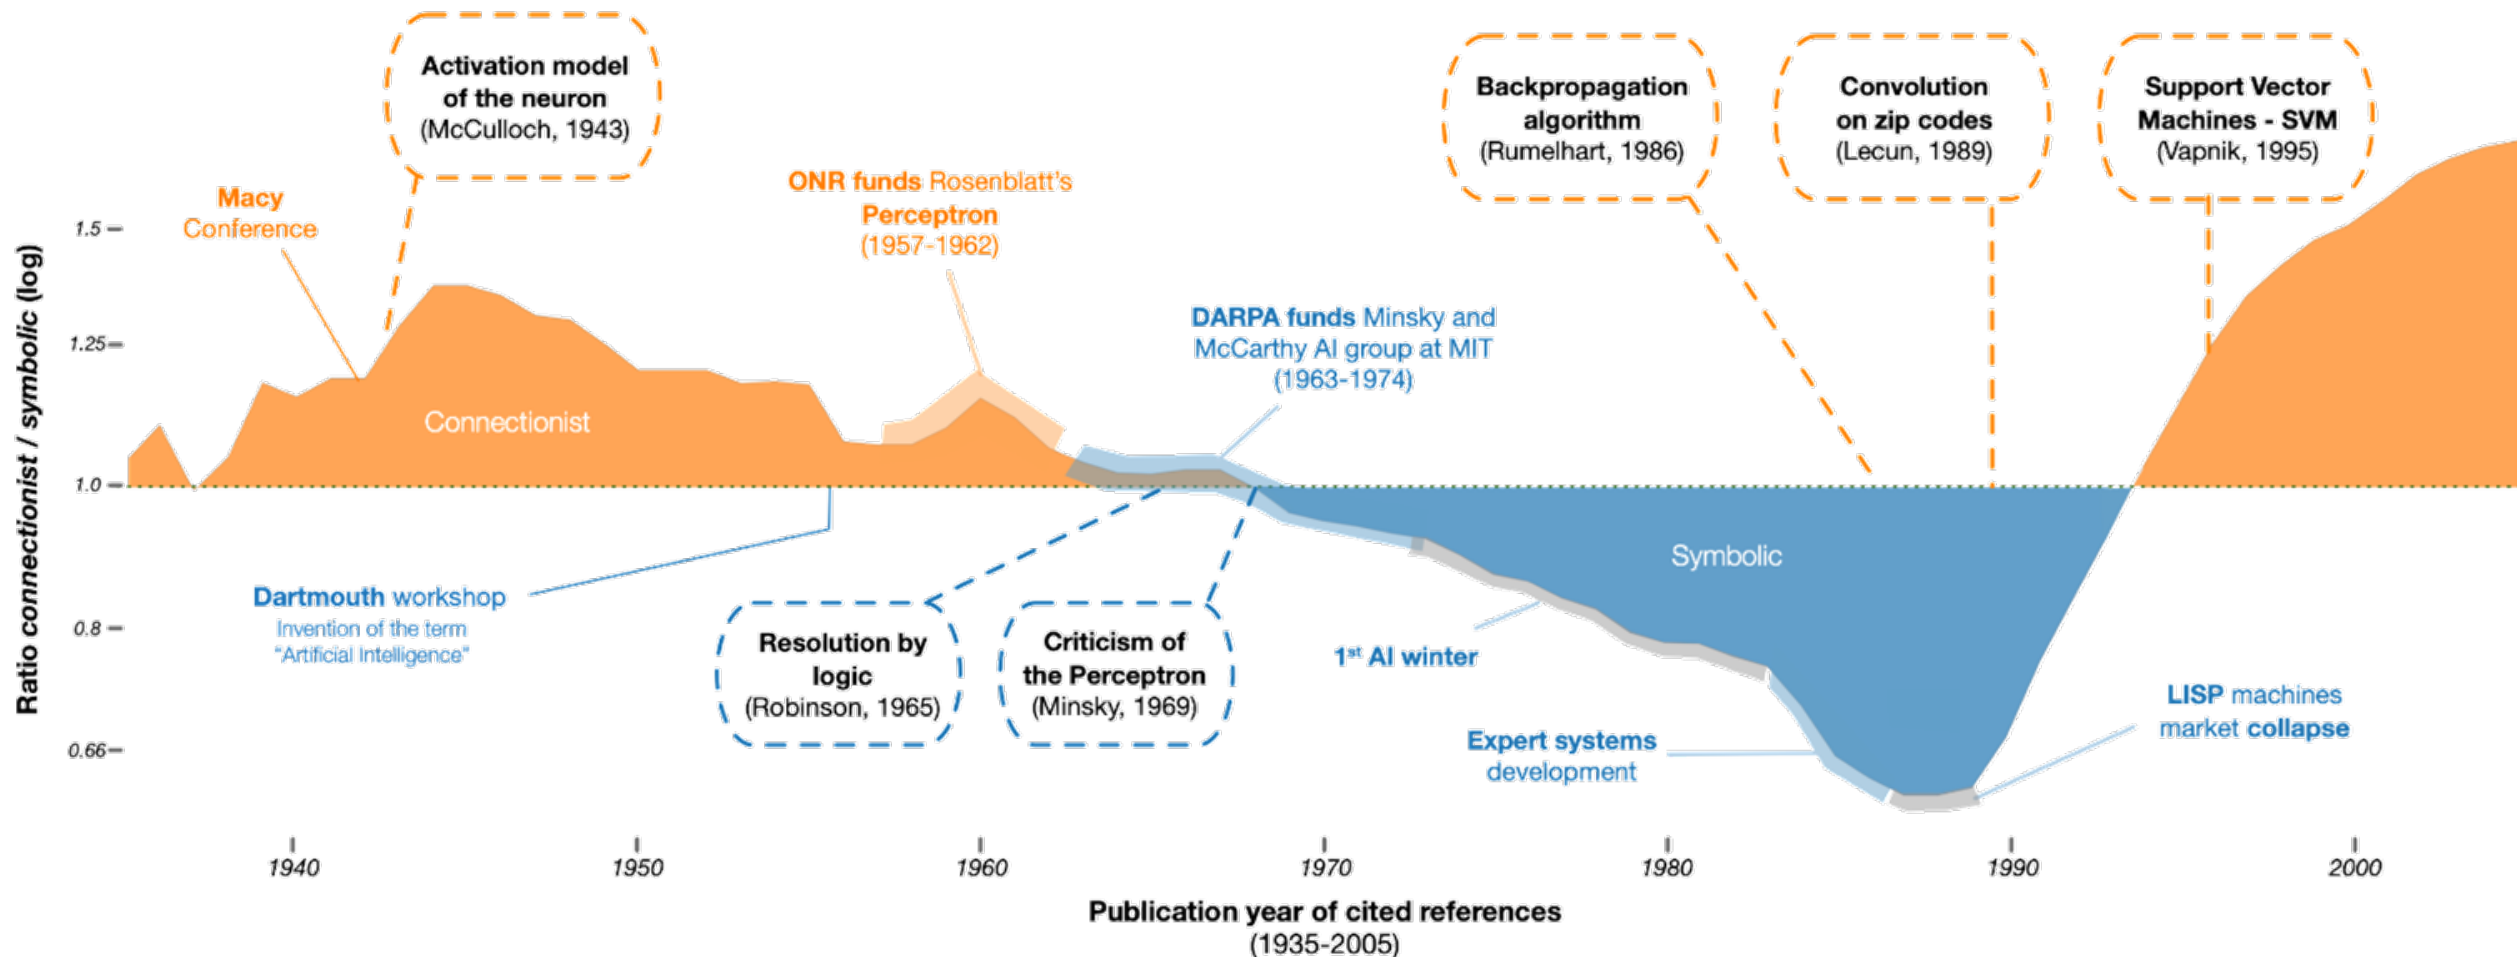


Neuron Model of Connections

- Developed to mimic the human neural system (in the brain) and its processing capabilities
- Decentralized knowledge representation and processing
👉 hopefully very efficient
- Simple components, the intelligence is in the **connections**



Connectionist vs Symbolic AI




Cardon, D., Cointet, J.-P. & Mazieres, A. (2018). Neurons spike back. The invention of inductive machines and the artificial intelligence controversy. *Réseaux*, 36(211), 173-220.

A short history of Neural Networks

- 1957: Perceptron (Frank Rosenblatt): one layer network neural network
- **First AI Winter (General problem solvers)**
- 1988: Backpropagation (Rumelhart, Hinton, Williams): learning a multi-layered network
- 1989: ALVINN: autonomous driving car using NN (CMU)
- 1989: (LeCun) Successful application to recognize handwritten ZIP codes on mail using a “deep” network
- **Second AI Winter (Expert systems)**
- 2015: Convolutional neural networks
- Late 2010s: Deep learning (advanced image recognition, AI masters Go/StarCraft 2, etc.)
- 2020s: LLM and Generative AI enabling (near-human capabilities for image recognition, speech recognition, and language translation)

Current

Current



After second AI Winter computer science made a lot of progress on taking the specialized system on the symbolic side and generalizing their ideas



On the other side our computing power reached the point that connectionist became powerful again

Post AI Winter II (1990s ->)

New developments and Applications :

- **Multi-Agent Systems:**
 - Cooperation concepts to bring together the many specialized systems to tackle harder tasks
- **The Internet**
 - Requires agents for “intelligent” routine tasks that are specialized and limited

Post AI Winter II (2000s ->)

- **Service-based computing:**
 - Offer functionality at your web site and allow others to use it (☞ describe service and how to access it, best in a form understandable by a program)
 - Decentralization of knowledge representation and processing
 - Tries to rescue the semantic web idea
- **Self-X Systems:**
 - Started as IBM's (and now others) code for learning/adaptive systems
 - Sees more and more areas within CS that want to include learning/adaptation into their concepts and systems
 - Self-organization, self-healing, ...
 - General goal: make systems easier to configure, maintain and adapt to usage

Post AI Winter II (2010s ->)

- **Machine Learning:**

- Around since Optical Character recognition research
- Prominence in 1990s with spam filters
- Good at seeing patterns in images, classifying things with many parts (text), filtering by pattern, chat systems, modeling complex things like, voice recognition, etc.
- Supervised Learning : Machine is given guidance (often based on pre-labeling training data)
 - Common examples are classifying something under labels, or finding parameters for a model that best matches data
- Unsupervised Learning: Machine is left unguided and has to decide most everything (often minor guidance like # of things to find)
 - Common example is clustering

- **Deep Learning:**

- Generally just means taking neural network machine learning to extreme with depth

Post AI Winter II (2010s ->)

- Areas in AI grow together:
 - Personal assistants like Alexa: restricted natural language understanding, machine learning, (intelligent) search techniques in the Internet and use of service-based computing
 - Watson: dedicated knowledge representation for an application area, machine learning to add to represented knowledge (often using natural language understanding) and restricted inference mechanisms
 - ...


Current

- **Popular areas:** Deep learning (neural networks), machine learning (AI for finding patterns), digital assistants (connect symbolic systems), service architectures, internet of things, self-automation, etc.
- **Things likely less close than you have been sold:** universal self-driving cars (visual identification has a lots of weaknesses), machine intelligence (?), true generic systems (we've generalized but there's usually a lot more specificity under the hood than you think)

What can AI do?

State of the art

- Which of the following can be done at present?
 - Play a decent game of table tennis
 - Drive safely along a curving mountain road
 - Drive safely on Deerfoot in the winter
 - Buy a week's worth of groceries on the web
 - Buy a week's worth of groceries from inside your local grocery store
 - Play a decent game of chess/poker
 - Discover and prove a new mathematical theorem
 - Design and execute a research program in molecular biology
 - Write an intentionally funny story
 - Give competent legal advice in a specialized area of law
 - Translate spoken English into spoken Swedish in real time
 - Converse successfully with another person for an hour (your miles may vary...)
 - Perform a complex surgical operation
 - Unload any dishwasher and put everything away



Future

AI - The Future ? (I)

- Specialized agents (**avatars**) helping and representing users (in the new **digital world**, i.e. moving Siri/Alexa/... forward)
- “Command-and-Control” systems for management and CEOs in large companies
 - ☞ Management by exception
- Cooperating intelligent robots in **disaster areas** and **hostile environments** represent humans acting out of safe environments
- Automation of trade and commerce
 - ☞ Management by Policies

AI - The Future ? (II)

In general AI enhances what a human being can do intellectually (assistants)

Still hard problems:

- Achieving necessary flexibility to interact with human beings in real world
- Find concepts to combine the methods of the different fields
- Common sense reasoning

Next... Agents

Jonathan Hudson, Ph.D.
jwhudson@ucalgary.ca
<https://cspages.ucalgary.ca/~jwhudson/>

