# Artificial Intelligence (AI)

Mark Duewiger

March 21, 2019

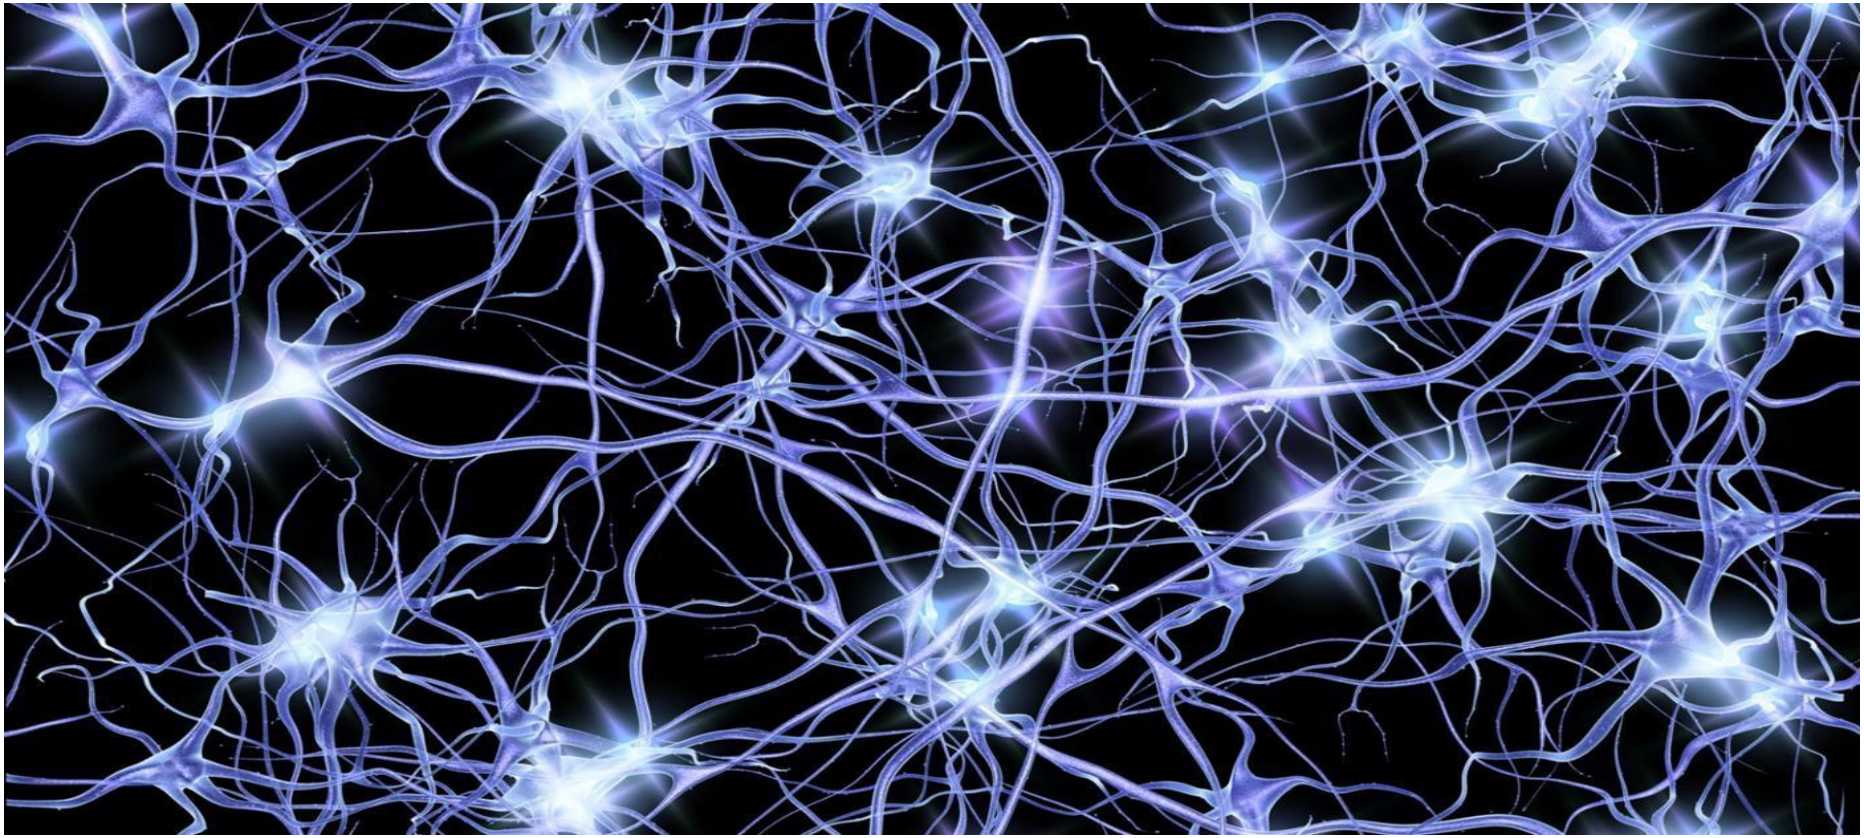

# AI Will be Covered in Three Evenings


- ARTIFICAL INTELLIGENCE:
  - Concepts
    - Rule Based AI
    - Fuzz Logic
    - Agents
    - Expert Systems
    - Neural Networks
  - Definitions
  - Programmed vs. Learning
  - AI History
  - Extra information to review at your leisure
  - thoughts on how we can augment our work with AI
- Presented by Mark

# AI Will be Covered in Three Evenings


- Jobs in the world of
  - Robots
  - Artificial Intelligence

- Presented by Phill

# AI Will be Covered in Three Evenings


MACHINE LEARNING

- what is it?
- how far have we come?
- how do machines learn; what are they learning?
- how good are they?
- where is it already being used?
- where are we heading?
- what problems might be involved in this development?

- Presented by Torben

# What is AI?

AI is an evolution in computing as, or more, important than the shifts to mobile or cloud computing. "It's hard to overstate," Amazon CEO Jeff Bezos wrote, "how big of an impact AI is going to have on society over the next 20 years."

# Rule-Based AI

- A typical **rule-based system** has four **basic** components:
  - A list of **rules** or **rule base**, which **is** a specific **type** of knowledge **base**.
  - An inference **engine** or semantic reasoner, which infers information or takes action **based** on the interaction of input and the **rule base**.
  - Memory
  - Communications to other devices.

- Rule Based Systems = Expert Systems = Knowledge Systems

- https://www.youtube.com/watch?v=kmRLeI1P9zw

# Fuzzy Logic

- **Fuzzy logic** is a form of many-valued logic in which the truth values of variables may be any real number between 0 and 1 inclusive.
  - It is employed to handle the concept of partial truth, where the truth value may range between completely true and completely false.
  - By contrast, in Boolean logic, the truth values of variables may only be the integer values 0 or 1.
- The term *fuzzy logic* was introduced with the 1965 proposal of fuzzy set theory by Lotfi Zadeh.
  - Fuzzy logic had however been studied since the 1920s, as infinite-valued logic—notably by Łukasiewiczand Tarski.
- It is based on the observation that people make decisions based on imprecise and non-numerical information, fuzzy models or sets are mathematical means of representing vagueness and imprecise information, hence the term fuzzy.
- These models have the capability of recognizing, representing, manipulating, interpreting, and utilizing data and information that are vague and lack certainty.
- Fuzzy logic has been applied to many fields, from control theory to artificial intelligence.

- https://www.youtube.com/watch?v=rln_kZbYaWc

# Agents Agents Everywhere

- **Intelligent agent** (**IA**) is an autonomous entity which acts, directing its activity towards achieving goals (i.e. it is an agent), upon an environment using observation through sensors and consequent actuators (i.e. it is intelligent). Intelligent agents may also learn or use knowledge to achieve their goals. They may be very simple or very complex. A reflex machine, such as a thermostat, is considered an example of an intelligent agent

# Biological Neuron

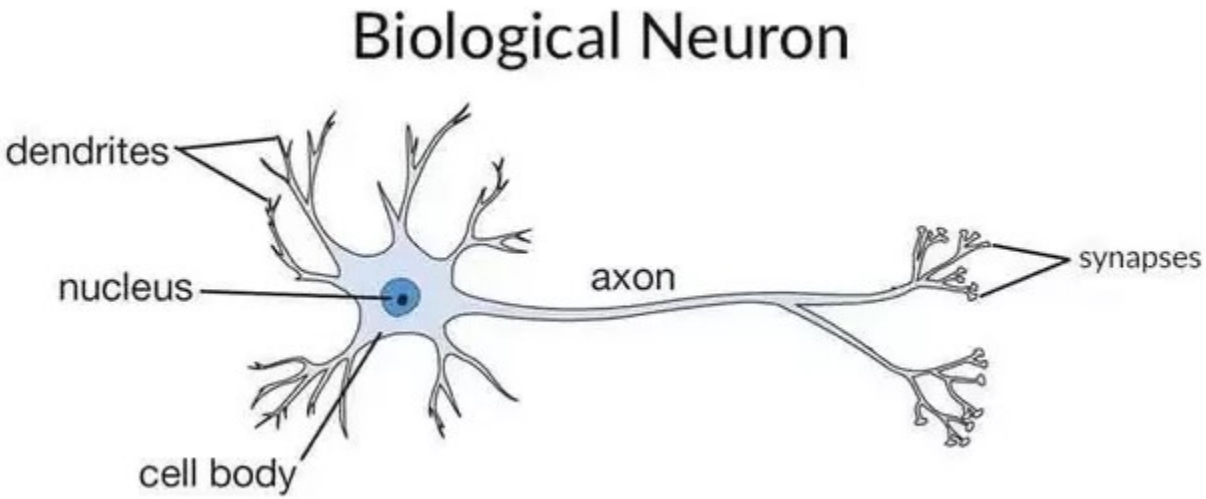

# Artificial Neural Network

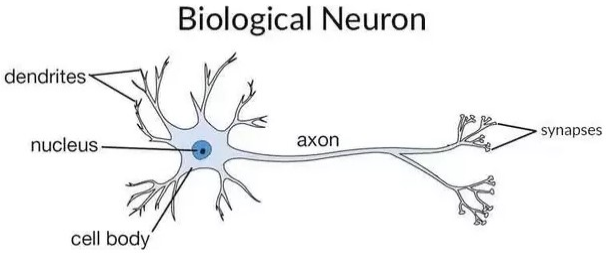


- Assign weights to interconnects and bias to nodes.
- Run known data through system and modify weights and bias
- Run not known date through system and modify weights and Bias

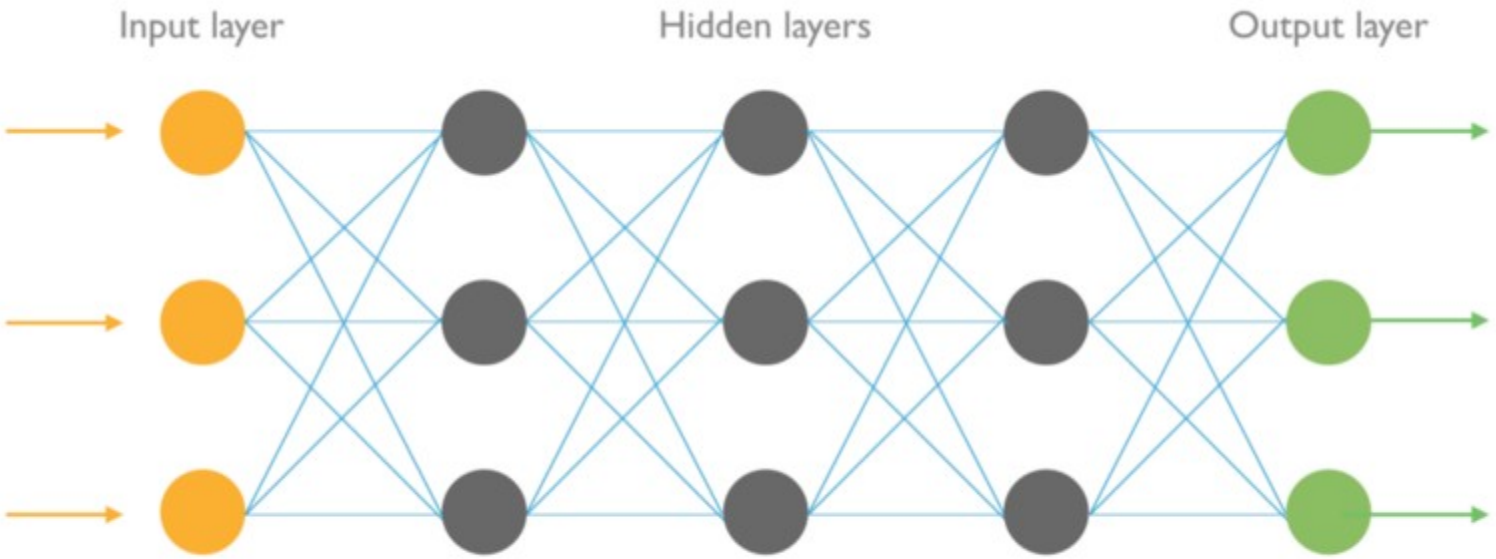

# Difference Between Computer and Human Brain

- **Difference Between Computer and Human Brain**

|  |  |
|---|---|
| Biological neurons or nerve cells | Silicon transistors |
| 200 billion neurons, 32 trillion interconnections. | 1 billion bytes RAM, trillion of bytes on disk. |
| Neuron size: 10-6 m. | Single transistor size: 10-9m. |
| Energy consumption: 6-10 joules per operation per sec. | Energy consumption: 10-16 joules per operation per second. |
| Learning capability | Programming capability |
| Operates in u seconds * | Operates in n seconds * |

* Needs to be verified

# Neural Network

- ADDITIONAL INFORMATION

- A *neural network* (NN), in the case of artificial neurons called *artificial neural network* (ANN) or *simulated neural network* (SNN), is an interconnected group of natural or artificial neurons that uses a mathematical or computational model for information processing based on a connectionistic approach to computation. In most cases an ANN is an adaptive system that changes its structure based on external or internal information that flows through the network.
- In more practical terms neural networks are non-linear statistical data modeling or decision making tools. They can be used to model complex relationships between inputs and outputs or to find patterns in data.
- *neural network* (NN), in the case of artificial neurons called *artificial neural network* (ANN) or *simulated neural network* (SNN), is an interconnected group of natural or artificial neurons that uses a mathematical or computational model for information processing based on a connectionistic approach to computation. In most cases an ANN is an adaptive system that changes its structure based on external or internal information that flows through the network.

# Artificial Intelligence: The science of Intelligent programs

- 'Artificial Intelligence' (AI) is a general term that refers to hardware or software that exhibits behavior which appears intelligent.
- It is "the science and engineering of making intelligent machines, especially intelligent computer programs."
- Forms of 'AI' have existed for decades,
  - rules-based programs that deliver rudimentary displays of 'intelligence' in specific contexts.
  - Expert systems
- Progress, however, has been limited — because algorithms to tackle many real-world problems are too complex for people to program by hand.

# Artificial intelligence: The science of intelligent programs

- These activities including making
  - medical diagnoses,
  - predicting when machines will fail
  - gauging the market value of certain assets
    - involve thousands of data sets and non-linear relationships between variables.
- In these cases, it's difficult to use the data we have to best effect — to '**optimise'**our predictions.
- In other cases, including recognising objects in images and translating languages, we can't even develop rules to describe the **features** we're looking for.
  - How can we write a set of rules, to work in all situations, that describe the appearance of a dog?
- What if we could transfer the difficulty of making complex predictions — the **data optimization** and **feature specification** — from the programmer to the program?
- This is the promise of modern artificial intelligence.

# Machine Learning

Machine learning lets us tackle problems that are too complex for humans to solve by shifting some of the burden to the algorithm.

As AI pioneer Arthur Samuel wrote in 1959, machine learning is the ‘field of study that gives computers the ability to learn without being explicitly programmed.’

The goal of most machine learning is  to develop a prediction engine for a particular use case. An algorithm will receive information about a domain (say, the films a person has watched in the past) and weigh the inputs to make a useful prediction (the probability of the person enjoying a different film in the future). By giving ‘computers the ability to learn’, we mean passing the task of optimization — of weighing the variables in the available data to make accurate predictions about the future — to the algorithm . Sometimes we can go further, offloading to the program the task of specifying the features to consider in the first place.

# Deep Machine Learning

- Machine learning algorithms learn through training. An algorithm initially receives examples whose outputs are known, notes the difference between its predictions and the correct outputs, and tunes the weightings of the inputs to improve the accuracy of its predictions until they are optimized. The defining characteristic of machine learning algorithms, therefore, is that **the quality of their predictions improve with experience**. The more data we provide (usually up to a point), the better the prediction engines we can create.

# History of Artificial Intelligence

- This section is important to understand -how long AI is taking to develop.
- AI is not easy as the History shows
  - AI has gone through 2 “AI winters” where research and money dried up
  - It is looking like a third “AI winter” could be averted.

# Thinking Rationally: “Laws of Thought"

- Aristotle: First to codify “right thinking”
- Several Greek schools developed various forms of logic:
  - Notation and rules of derivation for thoughts
- By 1965, programs existed that could, in principle, solve any solvable problem described in logical notation.

- Problems:
  - Not easy to state informal knowledge in logical notation
  - Big difference between solving a problem "in principle" and solving it “in practice”
    - Problems with just a few hundred facts can exhaust the computational resources of any computer

- 1940 First Commercial Computer

# Science Fiction Test

1968 2001: A Space Odyssey

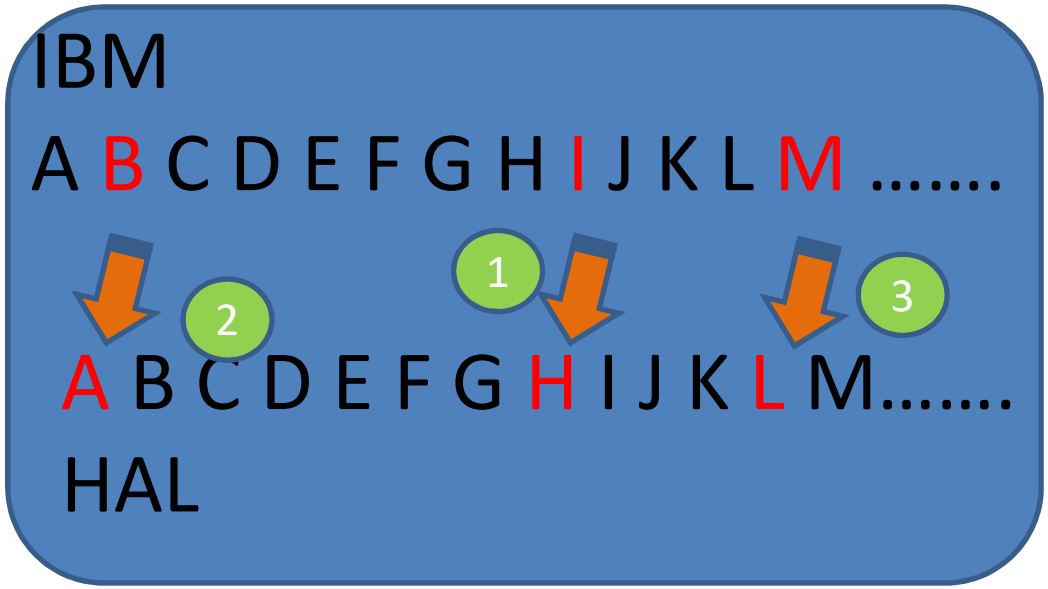


1950 Isaac Asimov published his

Three laws of robotics

1. A robot may not injure a human being or, through inaction, allow a human being to come to harm.
2. A robot must obey orders given it by human beings except except where such orders would conflict with the First Law.
3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.

# McCulloch & Pitts (1943)

- Proposed a model of artificial neurons
- Each neuron is characterized as being "on" or"off,"
- Switch to "on" occurring in response to stimulation by a sufficient number of neighboring neurons.
- The state of a neuron was conceived of as "factually equivalent to a proposition
- Any computable function could be computed by some network of connected neurons
- All the logical connectives (and, or, not, etc.) could be implemented by simple net structures.
- McCulloch and Pitts also suggested that suitably defined networks could learn.
- First Neural Network Computer (1950)

# Acting humanly: Turing Test

- Turing (1950) "Computing machinery and intelligence":
- "Can machines think?"  "Can machines behave intelligently?"
- Operational test for intelligent behavior: the Imitation Game

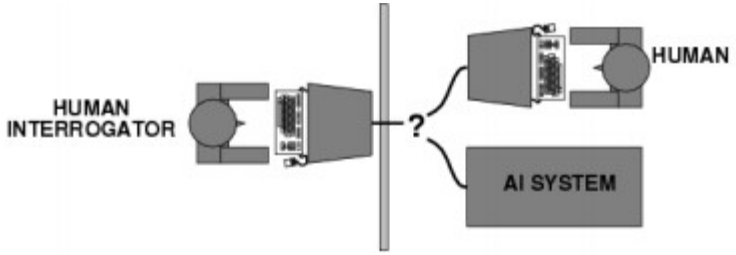


- Predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes
- Anticipated all major arguments against AI in following 50 years
- Suggested major components of AI: knowledge, reasoning, language understanding, learning
- ELIZA (Weizenbaum) simulated a psychotherapist

# Early Enthusiasm (1952 - 1969)

- GPS (thinking humanly) (General Problem S0lver)
- Herbert Gelemter (1959) constructed the Geometry Theorem Prover
- Arthur Samuel (1956) wrote a series of programs for checkers (draughts) that eventually learned to play at a strong amateur level
- LISP (1958) by John McCarthy

- Look  Ma No Hands

# Dartmouth (1956)

- Theme: – "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."
- 2 Month, 10 Man Study of AI
- Newell and Simon came up with a reasoning program, the Logic Theorist (LT)
- The program was able to prove most of the theorems in Chap 2, Principia Mathematica
- Participants
  - Marvin Minsky: converted from neural networks to symbol processing point of view
  - John McCarthy: coined term “AI”
  - Claude Shannon: invented switching circuits
  - Nathaniel Rochester: designed first popular IBM
  - Newell & Simon: wrote first AI program

# Start of DOD Funding

- MIT receives 2.2 million dollar grant from US
  - govt. to research Machine-Aided Cognition (AI)
- From Department of Defense
- US wants to stay ahead of Soviet Union
  - 1963

# Dose of Reality (1966 - 1973)

- In almost all cases, these early systems turned out to fail miserably when tried out on wider selections of problems and on more difficult problems.

- Failure to come to grips with the "combinatorial explosion" was one of the main criticisms of AI contained in the Lighthill report (Lighthill, 1973), which formed the basis for the decision by the British government to end support for AI research

# Prolog Language Revealed 1973

- General logic programming language
- Represented as facts and rules
- Developed in France

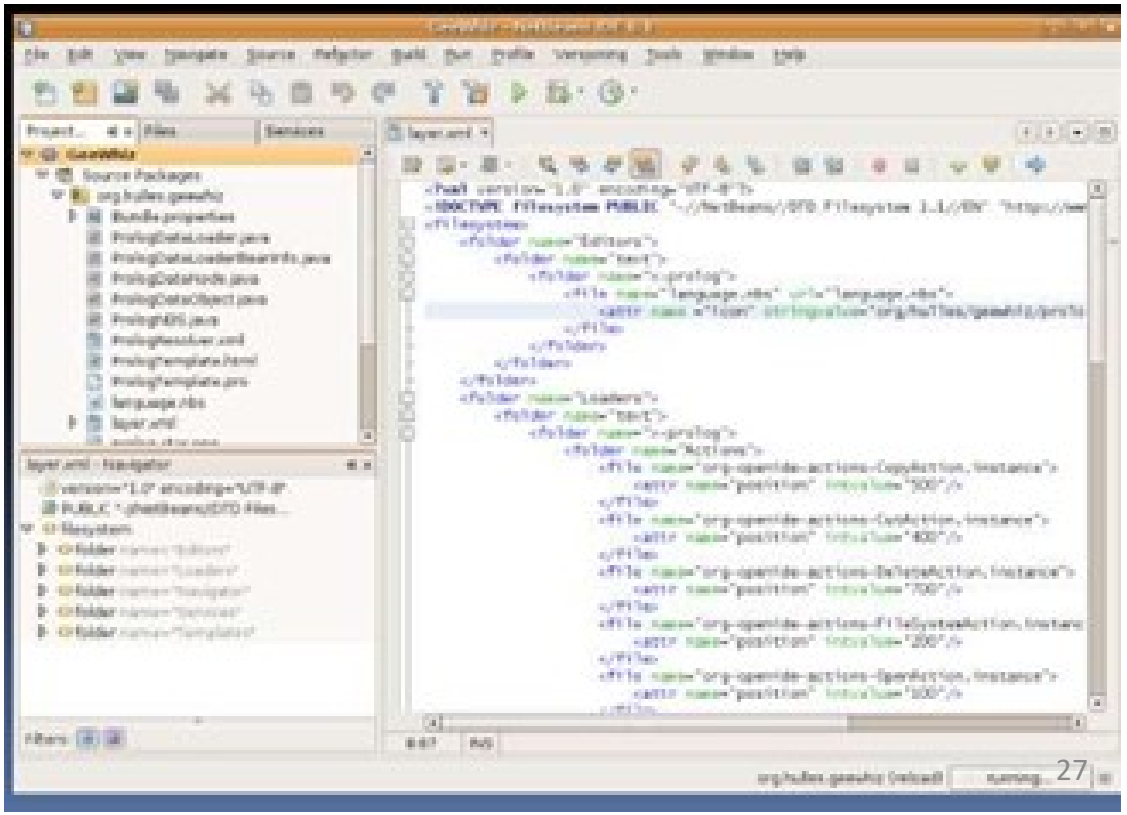

# Expert systems industry booms

- 1969 – 1979
- Early development of Knowledge based systems
- Expert systems for industry booms

# Expert systems industry busts

- 1988 – 1993
- Expert systems industry busts “AI Winter“
- Complexity

# Focus on Agents

- 1995
- Agents Agents Everywhere
- Neural nets return due to support learning

# Natural Language Processing

- Speech technologies
  - Automatic speech recognition (ASR)
  - Text-to-speech synthesis (TTS)
  - Dialog systems
- Language Processing Technologies
  - Machine Translation
  - Information Extraction
  - Information Retrieval
  - Text classification, Spam filtering.

# Others..

- Computer Vision:
  - Object and Character Recognition
  - Image Classification
  - Scenario Reconstruction etc.
- Game-Playing
  - Strategy/FPS games, Deep Blue etc.
- Logic-based programs
  - Proving theorems
  - Reasoning etc.

# AI Based Hardware/Software Sales $425M to companies 1986

- Expert systems in particular demand
- DuPont, General Motors, Boeing rely heavily on expert systems
- Companies develop that specialize in creating
  - Software that aids in producing expert
  - Systems

- Complexity and cost

# DESERT STORM 1986

Used Effectively in

- AI Military Systems
  - In missile systems
  - heads-up-displays
  - AI truly put the test

# Is Another AI Winter coming ??

- As of early November 2018 Many people though that there would be an AI Winter in 2019 due to:
  - Underpowered CPUs
  - Lack of AI specify memory structures
  - Obsolete algorithm's
- It looks now that an AI Winter has been over rated for now, do to:
  - A machine that does a billion billion operations per second was announced
  - New memory configurations are in development
  - A scalable ASIC designed to do AI type processing is near release

# Deep Learning

- Machine learning has been around for 30 years
- Deep Learning has been around for 10 years

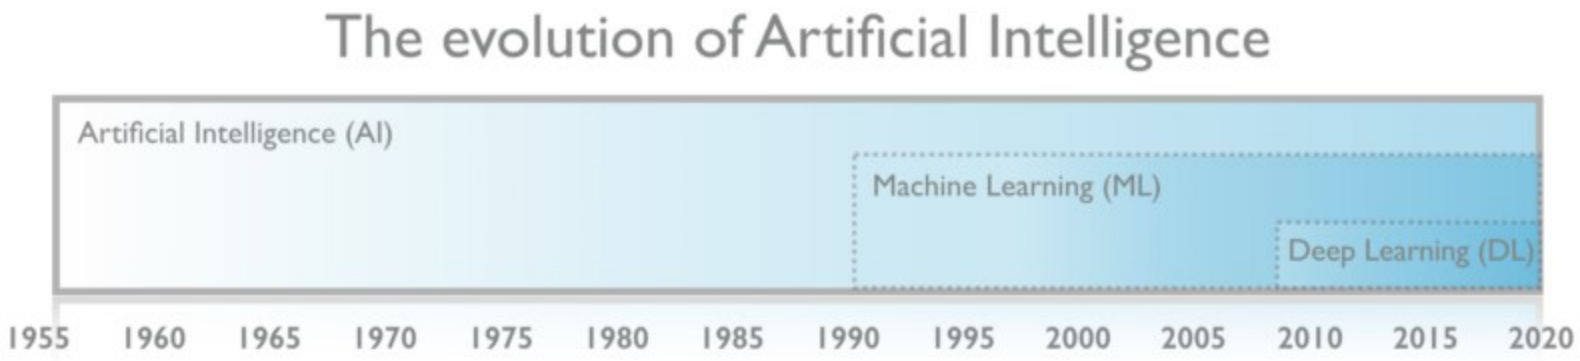


*Figure 4*

# Well Known Milestones

- Deep Blue defeated the reigning world chess champion Garry Kasparov in 1997
- No hands across America (driving autonomously 98% of the time from Pittsburgh to San Diego)
- During the 1991 Gulf War, US forces deployed an AI logistics planning and scheduling program that involved up to 50,000 vehicles, cargo, and people
- NASA's on-board autonomous planning program controlled the scheduling of operations for a spacecraft
- `Proverb` solves crossword puzzles better than most humans.

# Natural Language Processing

- Speech technologies
  - Automatic speech recognition (ASR)
  - Text-to-speech synthesis (TTS)
  - Dialog systems
- Language Processing Technologies
  - Machine Translation
  - Information Extraction
  - Informtation Retrieval
  - Text classification, Spam filtering.

# Others..

- Computer Vision:
  - Object and Character Recognition
  - Image Classification
  - Scenario Reconstruction etc.
- Game-Playing
  - Strategy/FPS games, Deep Blue etc.
- Logic-based programs
  - Proving theorems
  - Reasoning etc.

# What AI Can Do For You

APPLE EXPERT TALK

How AI can enhance our memory, work and social lives

Tom Gruber

https://youtu.be/DJMhz7JlPvA

# Appendix A

1) Slides not used
2) Elon Musk's thoughts]
3) Jobs

# Elon Musk is wrong. The AI singularity won't kill us all

It seems you can't open a newspaper without Elon Musk predicting that artificial intelligence (AI) needs regulating – before it starts World War III.

There are many technical reasons why the singularity might never happen. We might simply run into some fundamental limits. Every other field of science has fundamental limits. You can't, for example, accelerate past the speed of light. Perhaps there are some fundamental limits to how smart you can be?

Or perhaps we run into some engineering limits. Did you know that Moore's Law is officially dead? Intel is no longer looking to double transistor count every 18 months.

AI needs regulating because the big tech companies have got too big for their own good

Read for your self
https://www.wired.co.uk/article/elon-musk-artificial-intelligence-scaremongering

**List of jobs That No Longer Exist**
**Create new fields – add jobs**
**Loose both blue and white collar'**

# Will there be job los

**What should we do to mitigate job loss?**

Lector


http://www.ranker.com/list/jobs-that-no-longer-exist/coy-jandreau


**A few articles and videos on future Jobs.**
**Jobs that are likely to disappear and those that will flourish**

"The optimist's guide to the robot apocalypse" by Sarah
https://qz.com/904285/the-optimists-guide-to-the-robot-apocalypse/

# The Jobs We will Lose to Machines

https://youtu.be/gWmRkYsLzB4

Anthony  Goldbloom

And the ones we won't

Will lose jobs that have frequent/high volume tasks

Will not loose those with novel Situitions

4:30 min

# Don’t Fear Super-intelligent AI

https://youtu.be/z0HsPBKfhoI

Grady Booch

10 min

# The Jobs We will Lose to Machines

https://youtu.be/gWmRkYsLzB4

Anthony  Goldbloom

And the ones we won't

Will lose jobs that have frequent/high volume tasks

Will not loose those with novel Suitions

4:30 min

# The Jobs We will Lose to Machines

https://youtu.be/gWmRkYsLzB4

Anthony  Goldbloom

And the ones we won't

Will lose jobs that have frequent/high volume tasks

Will not loose those with novel Suitions

4:30 min

# The Jobs We will Lose to Machines

https://youtu.be/gWmRkYsLzB4

Anthony  Goldbloom

And the ones we won't

Will lose jobs that have frequent/high volume tasks

Will not loose those with novel Suitions

4:30 min

# Appendix B

1) An Impressive Script from Eliza

2) Script can go off topic if started differently

3) Have to be careful on the algorithm used is

4) Have to have specialized tools and skills to program AI

# An Impressive Script from Eliza

Men are all alike. (human)

IN WHAT WAY (machine)

They're always bugging us about something or other.

CAN YOU THINK OF A SPECIFIC EXAMPLE

Well, my boyfriend made me come here.

YOUR BOYFRIEND MADE YOU COME HERE

He says I'm depressed most of the time.

I AM SORRY TO HEAR YOU ARE DEPRESSED

It's true. I am unhappy.

DO YOU THINK COMING HERE WILL HELP YOU NOT TO BE UNHAPPY

# Eliza (cont'd)

I need some help, that much seems certain. WHAT WOULD IT MEAN TO YOU IF YOU GOT SOME HELP

Perhaps I could learn to get along with my mother.

TELL ME MORE ABOUT YOUR FAMILY

My mother takes care of me.

WHO ELSE IN YOUR FAMILY TAKES CARE OF YOU

My father.

YOUR FATHER

You are like my father in some ways.

WHAT RESEMBLANCE DO YOU SEE

# Eliza (cont'd)

You are not very aggressive but I think you don't want me to notice that.

WHAT MAKES YOU THINK I AM NOT VERY AGGRESSIVE

You don't argue with me.

WHY DO YOU THINK I DON'T ARGUE WITH YOU

You are afraid of me.

DOES IT PLEASE YOU TO BELIEVE I AM AFRAID OF YOU

My father is afraid of everybody.

WHAT ELSE COMES TO MIND WHEN YOU THINK OF YOUR FATHER

Bullies.

DOES THAT HAVE ANYTHING TO DO WITH THE FACT THAT YOUR BOYFRIEND MADE YOU COME HERE

# A Less Impressive Eliza Script

I don't understand your motives.
WHY DON'T YOU UNDERSTAND MY MOTIVES?
Shouldn't I be?
WHY DO YOU SAY THAT JUST NOW?
That's just the way it is.
SUPPOSE THAT'S JUST THE WAY IT WERE NOT.
So?
I AM NOT SURE I UNDERSTAND YOU FULLY.
You harp on the same questions too much.
WHAT DOES THAT SIMILARITY SUGGEST TO YOU?
I don't get you.
DO YOU WISH THAT YOU DID GET I?

# Eliza Script with Mistaken Identity

Do you think that I can use the computer this morning?

WHY DO YOU ASK?

I might be able to make some additional telecomp sales. WHY AREN'T YOU SURE?

My prospective customers have never seen a demonstration of this system.

DOES THAT MEAN VERY MUCH TO YOU?

Of course it does.

YOU SEEM VERY POSITIVE ABOUT THAT.

This is ridiculous.

WHAT DO YOU MEAN THIS IS RIDICULOUS?

Is there any reason I can't use the computer now?

TELL ME MORE ABOUT YOUR PROSPECTIVE CUSTOMERS.

# Appendix C

History sheets

# Abridged history of AI

- 1943 McCulloch & Pitts: Boolean circuit model of brain
- 1950 Turing's "Computing Machinery and Intelligence"
- 1956 Dartmouth: "Artificial Intelligence“ adopted
- 1952-69 Look, Ma, no hands!
- 1950s Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist,
- 1965 Robinson's algo for logical reasoning
- 1966-73 AI discovers computational complexity
  Neural network research almost disappears
- 1969-79 Early development of knowledge-based systems
- 1980-- AI becomes an industry
- 1986-- Neural networks return to popularity
- 1987-- AI becomes a science
- 1995-- The emergence of intelligent agents.

# 1.1 Brief history of AI

| | |
|---|---|
| *1943* | McCulloch & Pitts: Boolean circuit model of brain |
| *1950* | Turing's "Computing Machinery and Intelligence" |
| *1952–69* | Look, Ma, no hands! |
| *1950s* | Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine |
| *1956* | Dartmouth meeting: "Artificial Intelligence" adopted |
| *1965* | Robinson's complete algorithm for logical reasoning |
| *1966–74* | AI discovers computational complexity<br>Neural network research almost disappears |
| *1969–79* | Early development of knowledge-based systems |
| *1980–88* | Expert systems industry booms |
| *1988–93* | Expert systems industry busts: "AI Winter" |
| *1985–95* | Neural networks return to popularity |
| *1988–* | Resurgence of probability; general increase in technical depth<br>"Nouvelle AI": ALife, GAs, soft computing |
| *1995–* | Agents, agents, everywhere . . .<br>Machine learning comes to age, web intelligence, smart machines |
| *2003–* | Human-level AI back on the agenda |