



Lecture 10-1

Artificial Intelligence

What is Artificial Intelligence?

AI is the study, design, and development of computational processes that solve problems that **previously required human intelligence**

Not all of computer science is focused on AI.

But, it's hard to think of a field in CS that isn't **contributing fundamentally** to it

The “**AI Paradox**”: once we become familiar with a technology, we stop considering it AI

OpenAI's Dota 2 AI steamrolls world champion e-sports team with back-to-back victories

The Int

By Nick Sta

HEALTH + BEHAVIOR

Artificial intelligence performs as well as experienced



New AI Model Exceeds Human Performance at Question Answering

Microsoft's new AI translates Chinese-to-English as well as a human translator

By Michael Hicks March 14, 2018 World Of Tech

It's all in how you teach the AI



Denise Hea

FINDINGS

UCLA resea
diagnose pro
evaluating m
experienced
10 years of e

BACKGR

Radiologists
practicing on
accurately ex
highly specia

METHOD

is primarily diagnosed visually,

initial clinical screening and fol

dermoscopic analysis, a biopsy

examination. Automated classifi

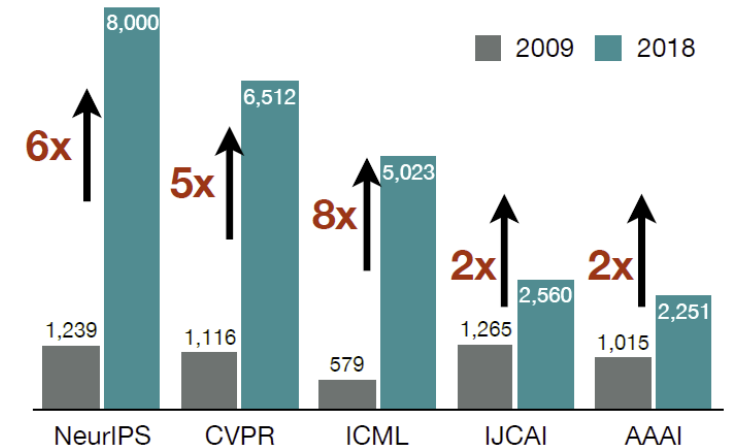
using images is a challenging ta

grained variability in the appea

Deep convolutional neural netw

potential for general and high

3 Rising attendance at AI conferences



Source: The AI Index 2018 Annual Report

IBM computer Watson wins clash

DeepFace: Choosing the Crown

Microsoft, Google Image Recognition

Deep learning algorithms compete

By R. Colin Johnson, 02/18/15

PORTLAND, Ore. -- First computers hurdle is image recognition -- surely too. Now Microsoft has programmed

The competition is fierce, with the In the 2015 championship on December they have one-upped humans too. For benchmark of 5.1% errors with a 4.94 Microsoft by 0.04%.

In modern face recognition consists of four stages. We revisit both the classification step by employing apply a piecewise affine representation from a deep network involve using several locally defining, rather than the we trained it on the large facial database unconstrained environment. Our method reaches a Faces in the Wild (LFW) current state of the art human-level performance.

The top row is a representative of the categories found in the database and the bottom row shows the algorithm's output. (Source: Microsoft)

1. Introduction

Face recognition in front of the algorithm and cultural implications are far reaching, yet the main barrier between machine as a buffer from having

ImageNet, with hundreds of object categories since 2010 with about 50,000 images. All the contestants algorithms, which are all derived from various versions of a human brain works to varying degrees. Most of the contestants algorithms in great detail. In the spirit of open source without

the most experts out

Deep convolutional neural network

potential for general and high

Deep Blue

On May 11, 1997, IBM's Deep Blue beat the world chess champion Garry Kasparov.

Fast supercomputers

The supercomputer, named Deep Blue, was the classic of the contest, because it was the first to use a human-like strategy.

The supercomputer, named Deep Blue, was the classic of the contest, because it was the first to use a human-like strategy.

The supercomputer, named Deep Blue, was the classic of the contest, because it was the first to use a human-like strategy.

The supercomputer, named Deep Blue, was the classic of the contest, because it was the first to use a human-like strategy.

The supercomputer, named Deep Blue, was the classic of the contest, because it was the first to use a human-like strategy.

1995 - 1998

2011

2014

2015

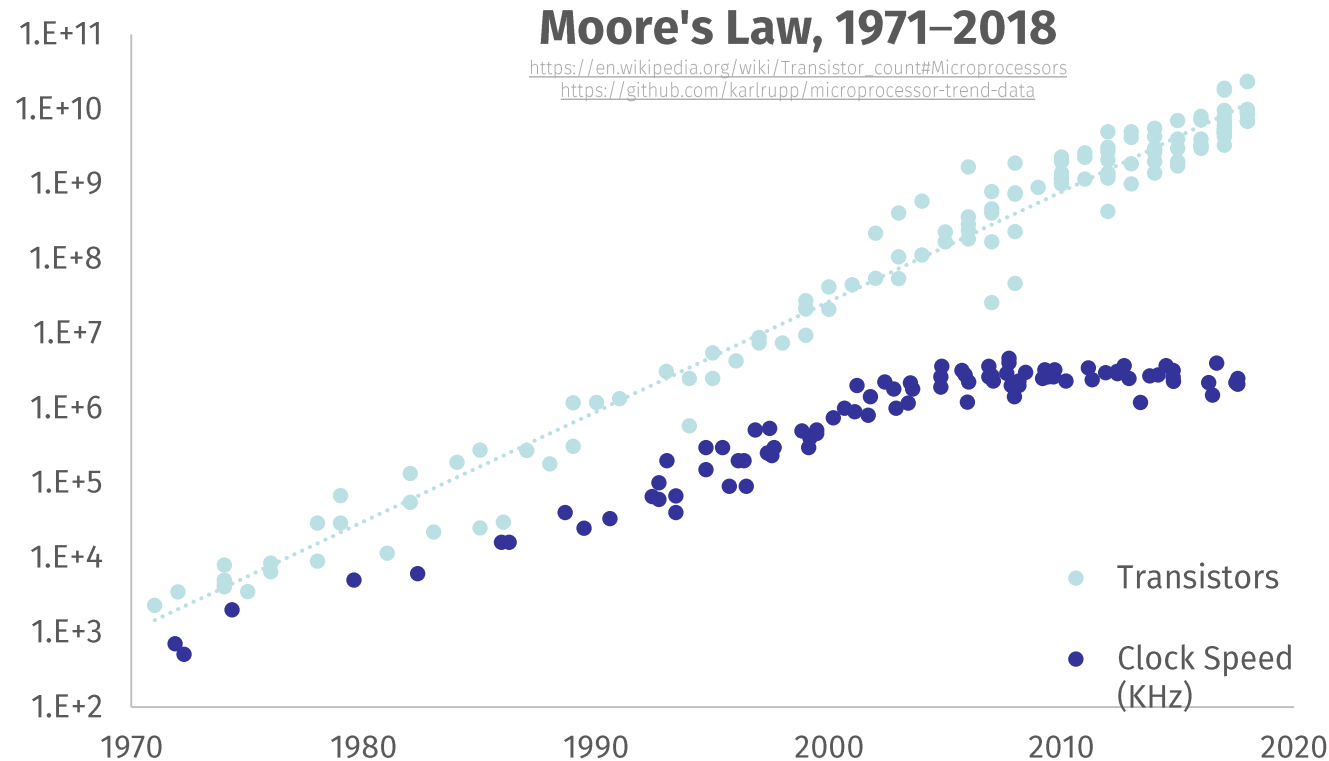
2016

2017

2018+

Why is it happening?

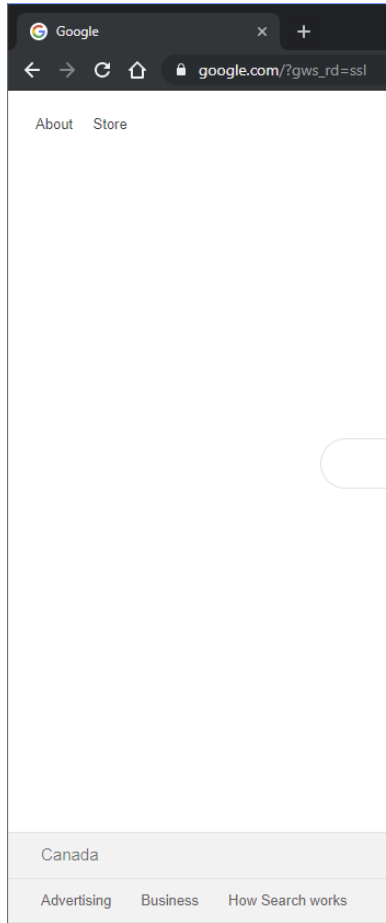
1. Scientific/mathematical **breakthroughs**, esp in machine learning
2. Growth in raw **computing power**



For comparison:

Species	# Neurons in Brain
Fruit Fly	10^5
Cat	10^9
Chimpanzee	10^{10}
Human	10^{11}

Today, most modern AI systems look like...



Current State of the Art

- **Predictive modeling** limited mostly by quality of training data
- **Image recognition, face recognition** becoming very reliable
- **Image generation** rapidly getting better
- **Language understanding & generation** maturing extremely quickly
 - try it out: <https://www.ai21.com/studio>
- **Robotics** remains a hard problem, but it's working well in some niche applications



Prompt: Here's a short rhyming poem by Dr. Seuss about Elon Musk's company, Neuralink which is building brain implants to achieve symbiosis with AI. Neuralink's mission statement is "If you Can't Beat'em, join'em". The poem has a structure typical of Dr. Seuss stories with simple alternate line rhymes. In the poem Elon explains why we need to merge with the AI to save humanity.

Neuralink, by Dr. Seuss

Response:

I have a friend, Elon Musk,
Who wants to join our brains to AI.
He thinks that humans can't compete
So he's gonna use AI to cheat.
They want to build a neural-link.
To link us to the AIs
And save humanity from its demise.
They want to put an AI in your brain.
They want to use your neurons as a frame.
They want to save humanity,
If you can't beat'em, join'em

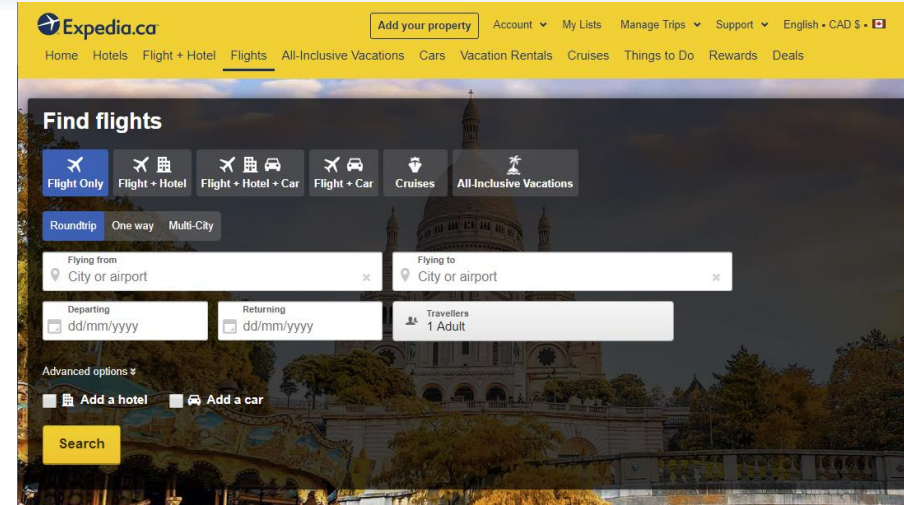
Growing Military Use



source: <https://www.youtube.com/watch?v=qsjPTtafilc>

Economic Impact

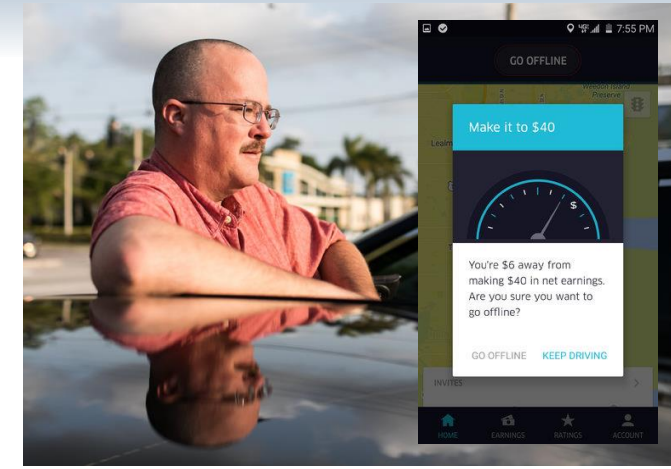
- Likely impacts:
 - Increased mechanization of **routine labor**
 - Automation of lower-end **knowledge work**
- Automation is nothing new
 - like other waves of innovation, AI will almost certainly make society as a whole **much richer**
 - but, it may also exacerbate **income inequality**
- What will this mean for the economy of the future?
 - **Post-scarcity** economy?
 - Human labor as a **luxury good**?



Ethics of AI

Will a new technology:

- disempower **individuals vs corporations?**
 - ⇒ user modeling; data mining; fostering addictive behaviors; developmental effects on children
- disempower **individuals vs governments?**
 - ⇒ facilitate disinformation (deep fakes; bots masquerading as people; filter bubbles); enable qualitatively new military or security tactics
- take **autonomous actions** in a way that obscures responsibility
 - ⇒ autonomous weapons; self-driving cars; loan approval systems
- disproportionately affect **vulnerable/marginalized groups**
 - ⇒ automated decision making tools trained in ways that may encode existing biases



Social Impact

- How will AI technologies **transform society**?
- Will there be a **social backlash** against AI?
 - If so, what will be considered AI?
- This **generation of children** will grow up taking for granted many technologies that strike us as magical
- How will **human relationships** change in the presence of always-available social agents?
- As we are increasingly **augmented by AI**, what are our inherent cognitive/emotional/motivational limitations, beyond which augmentation won't help?



Technology that will be in widespread use

- Tailored solutions for **specific tasks**, not general intelligence
- Prototypes **that work today** in labs & narrow deployments
- Some examples:
 - **Non-text input modalities** (vision; speech)
 - **Consumer modeling** (recommendation; marketing)
 - **Cloud services** (translation; question answering; AI-mediated outsourcing)
 - **Transportation** (automated trucking; some self-driving cars)
 - **Industrial robotics** (factories; some drone applications)
 - **AI knowledge work** (logistics planning; radiology; legal research; call centers)
 - **Policing & security** (electronic fraud; cameras; predictive policing)



Technologies that won't take off as quickly

- Overall, areas in which
 - major entrenched **regulatory regimes** need to be navigated
 - there exist **social/cultural barriers** to adoption
 - the **human touch** is crucial
 - substantial **new hardware** would need to be developed
- Some **examples**:
 - childcare, healthcare, eldercare
 - education
 - coaching, counselling
 - consumer robots beyond niche applications
 - semantically rich language understanding



Ethical consideration of advances in AI

- *Is it wrong to create machines capable of making human labor obsolete? Will humans become demoralized by the presence of vastly more intelligent robots?*
- *How can we ensure that intelligent robots will not be put to an evil purpose by a malevolent human? How can we ensure they do not adopt malevolent purposes themselves?*
- *Is it morally acceptable to create “personal” (self-conscious) AI?*

