# Future Visions of Telemedicine and AI

Robert Chang, MD

Byers Eye Institute at Stanford University

STANFORD UNIVERSITY

# Financial Disclosures

- Intellectual Property: PAXOS scope
- Consultant / Ad Board:
  - Alcon
  - Allergan
  - Santen
  - Pfizer
  - Aerie
  - Iridex
  - Kali Care
  - Healgoo
  - Unity Biotechnology
- Research Support:
  - Carl Zeiss Meditec
  - Santen

Stanford MEDICINE

# How Artificial Intelligence (AI) Will Change Medical Care

*(Are we going to lose our job?)*

# Geoff Hinton
## Google Brain Toronto

“If you work as a radiologist, you’re like Wile E Coyote that’s already over the edge of the cliff but hasn’t yet looked down.”

***”People should stop training radiologists now.”***

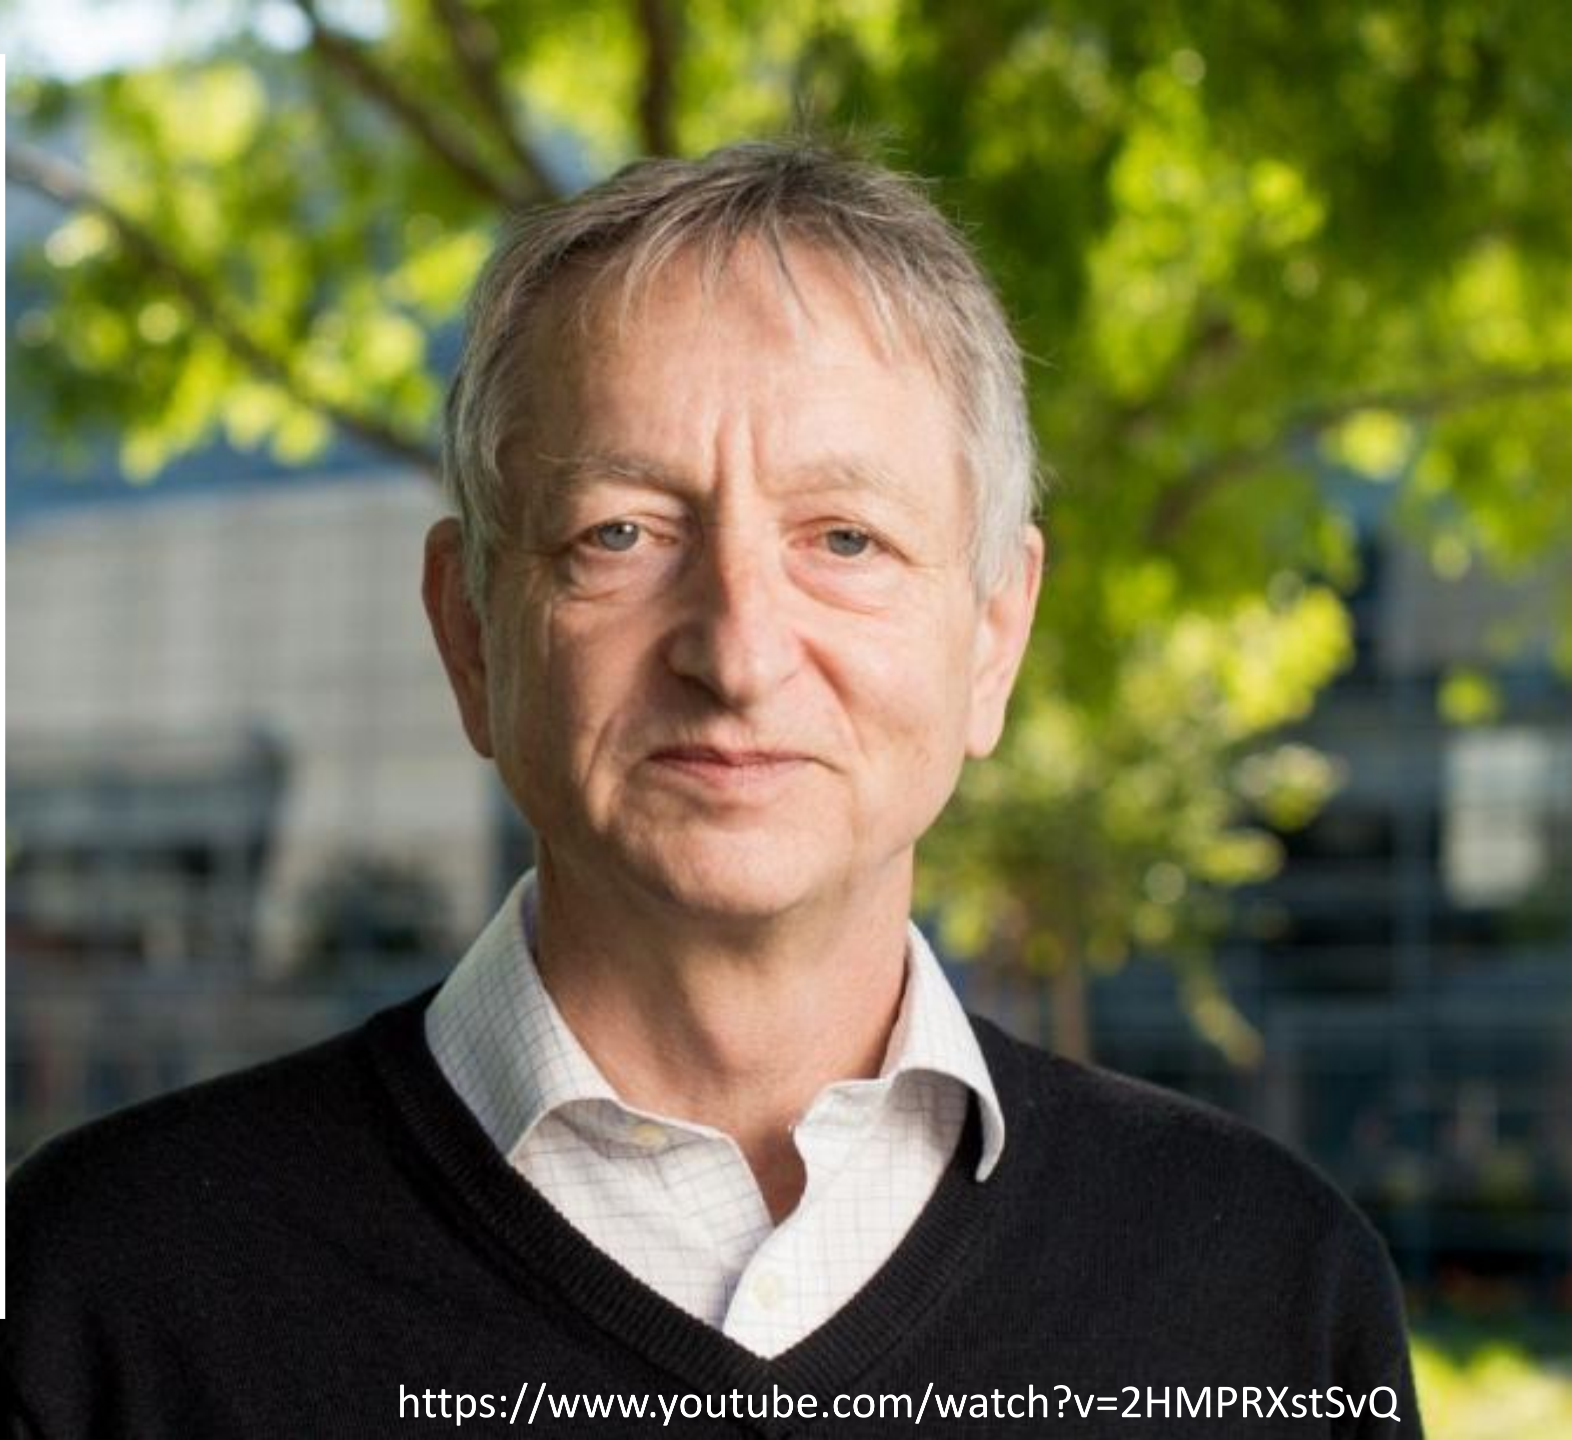

https://www.youtube.com/watch?v=2HMPRXstSvQ

# Google DeepMind Adds Memory to AI

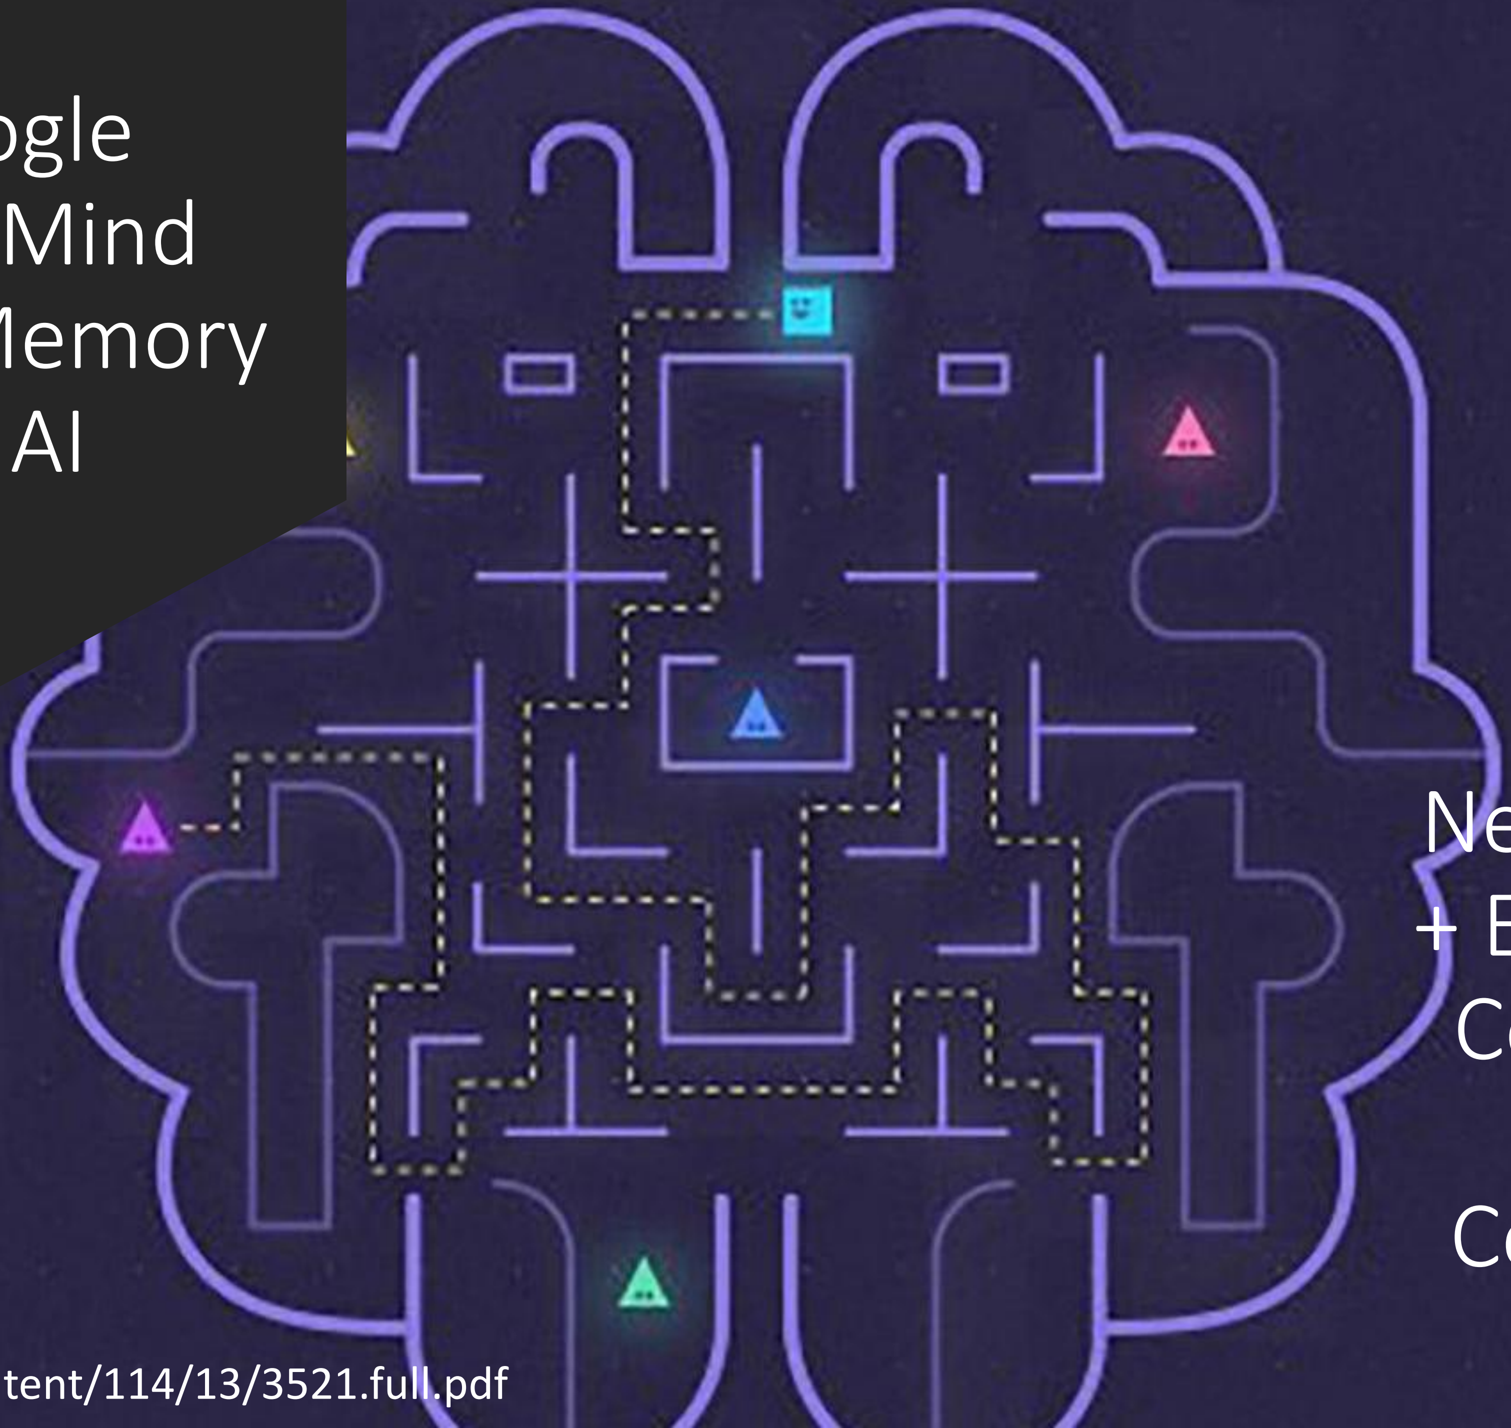

Deep Q Network (DQN) + Elastic Weight Consolidation (EWC) to Conquer Atari Games

http://www.pnas.org/content/114/13/3521.full.pdf

Above Human Level

Human Level

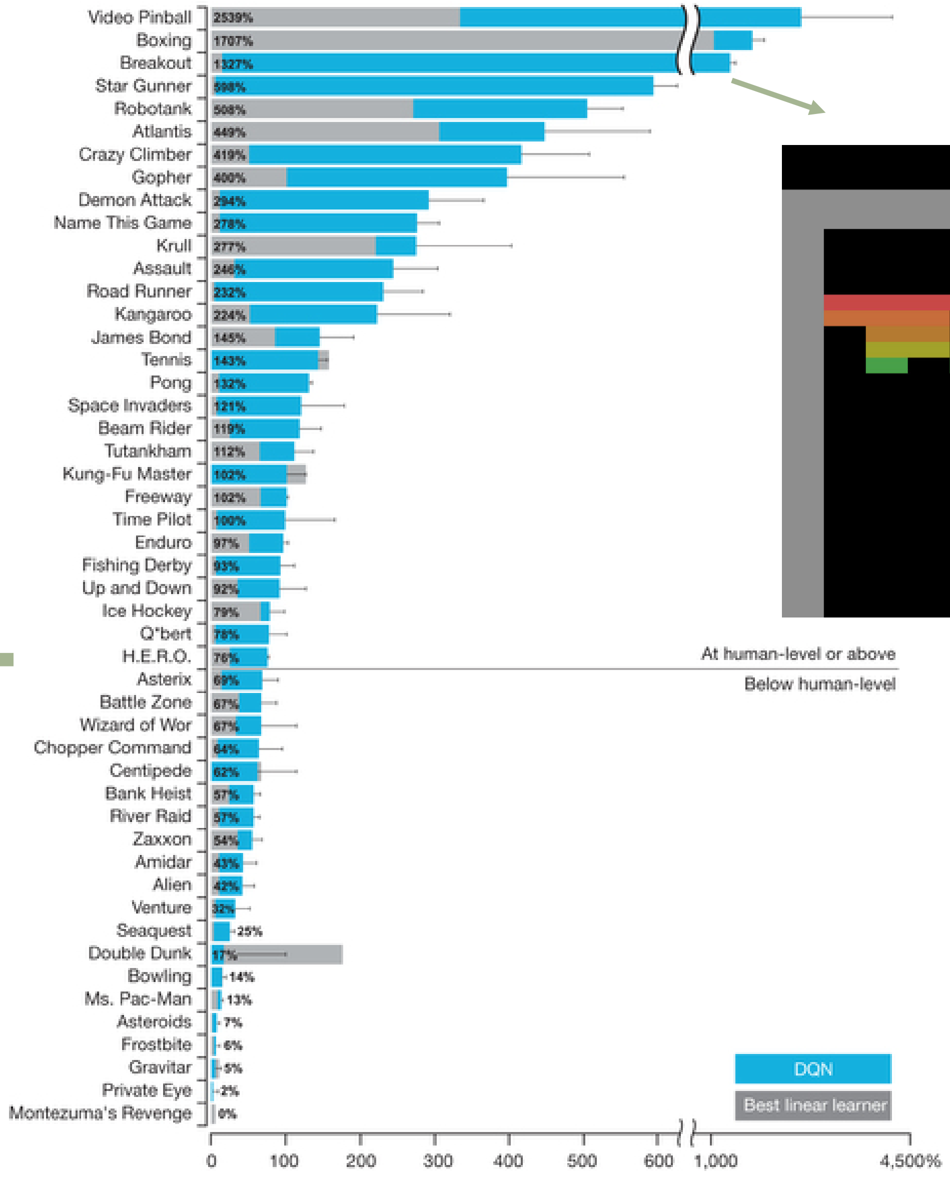



Breakout

https://www.nature.com/nature/journal/v518/n7540/full/nature14236.html

# AI Has Crushed Humanity at Poker

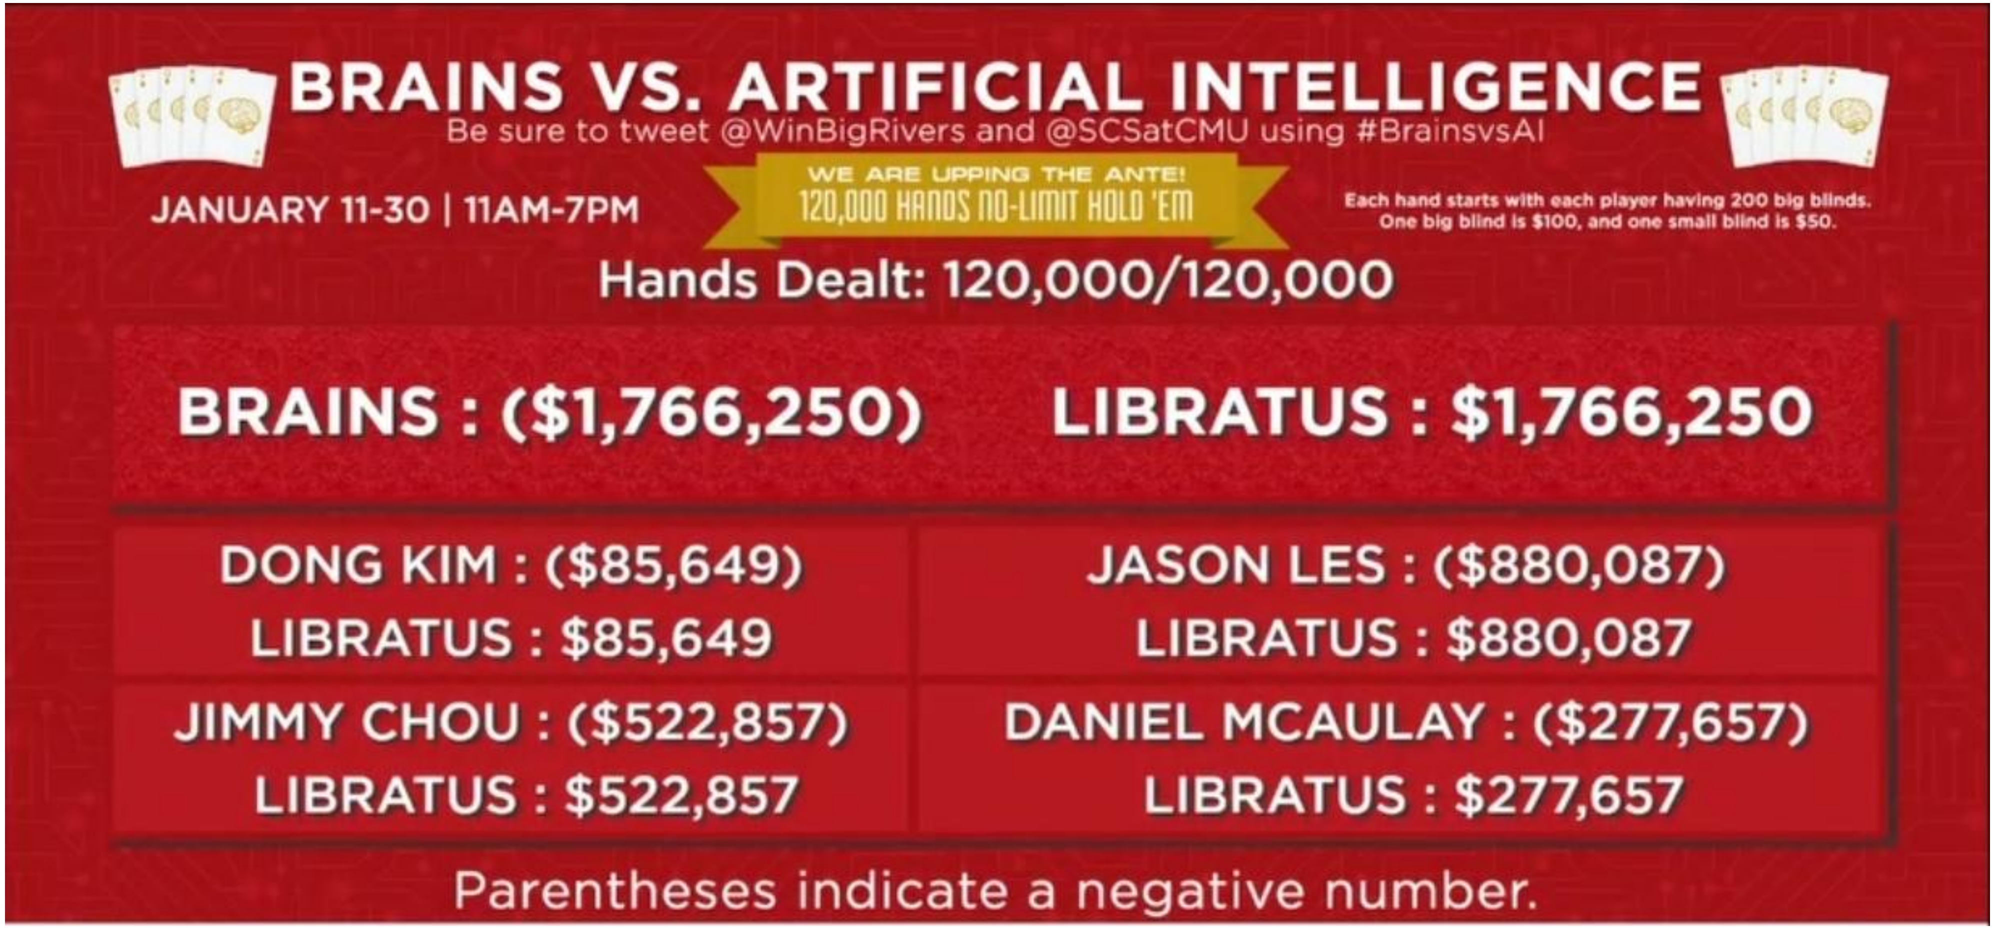

# Ke Jie *“AlphaGo is like a God of Go”*


5 Go Champions Against AI = Still No Chance

# Starcraft is the next big game for AI

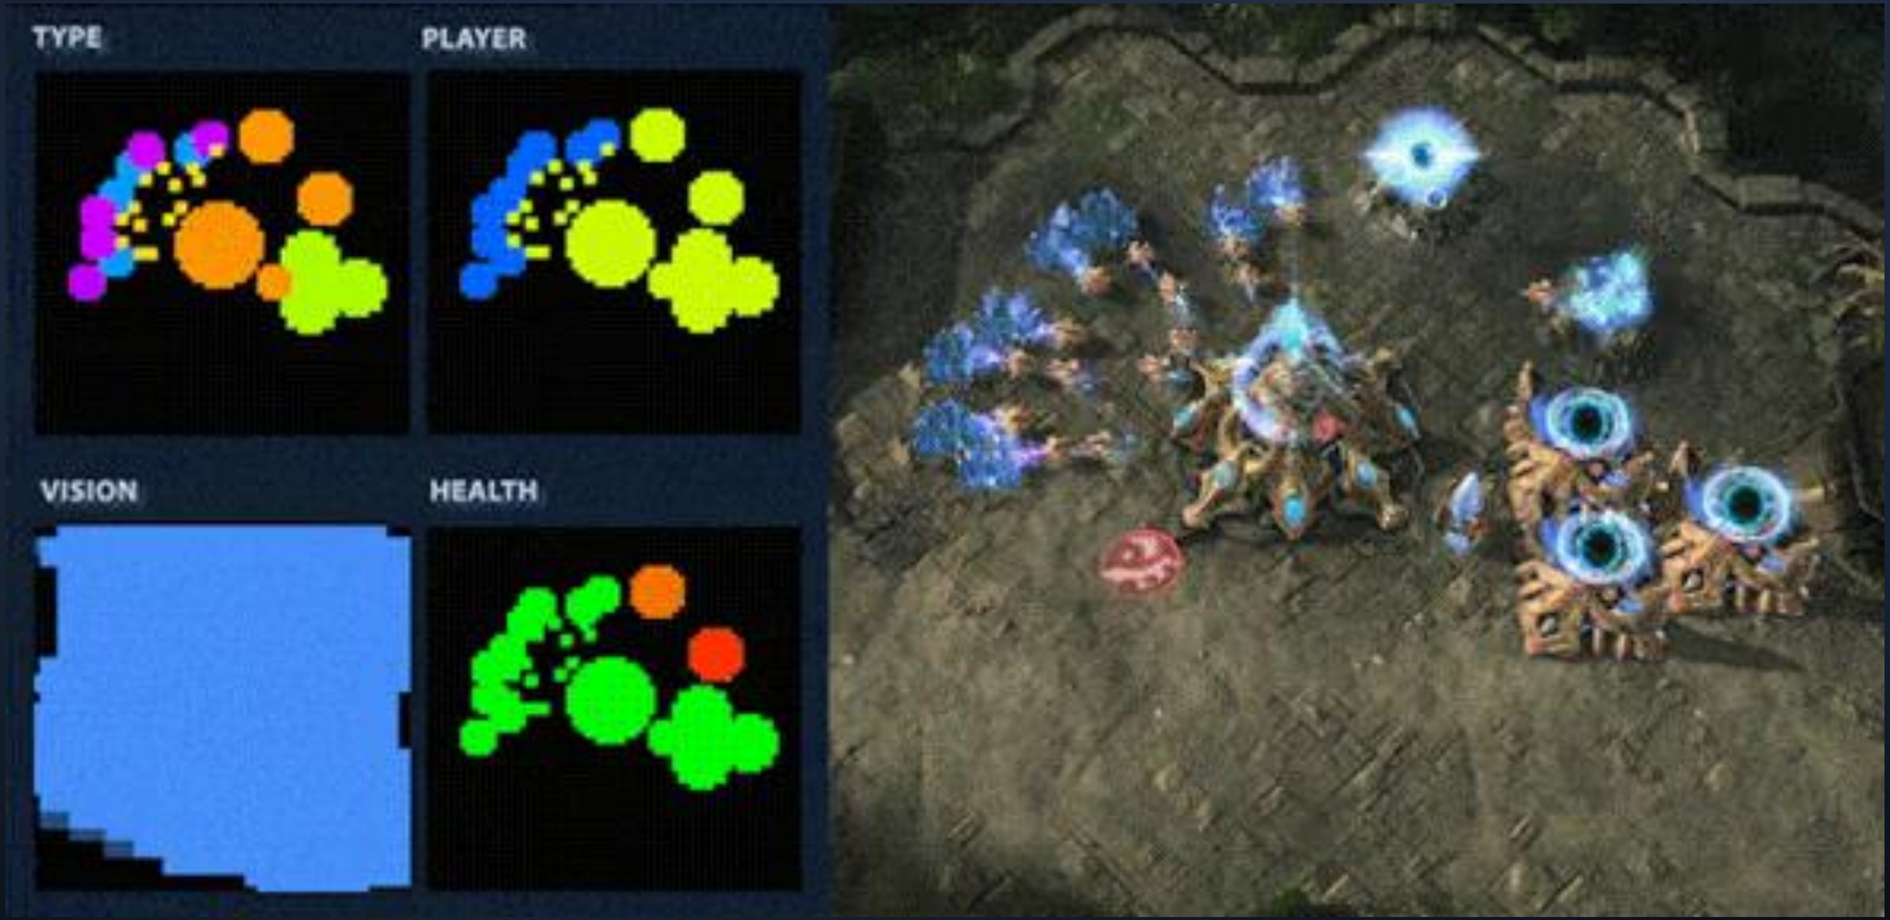

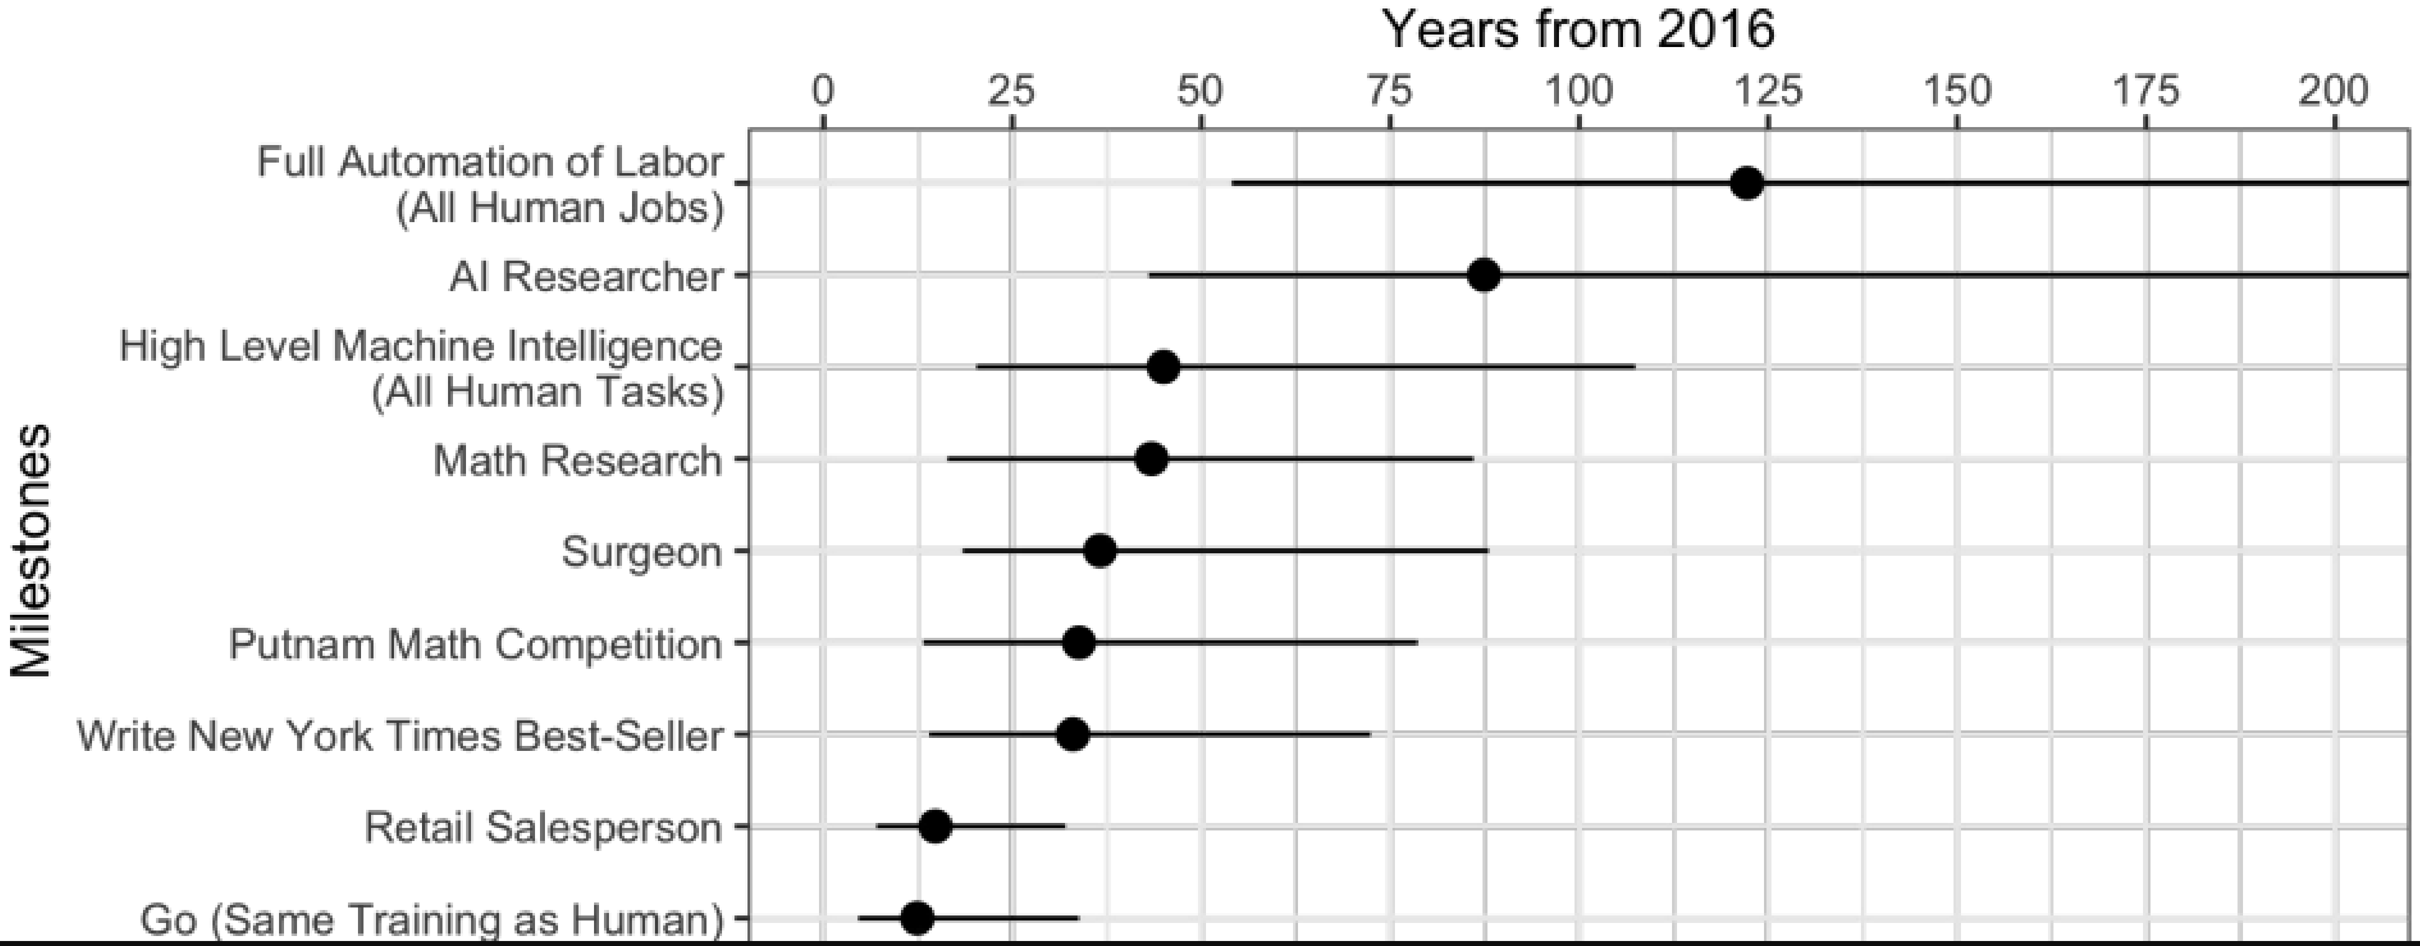


# Experts Predict When Artificial Intelligence Will Exceed Human Performance

Ref: arxiv.org/abs/1705.08807

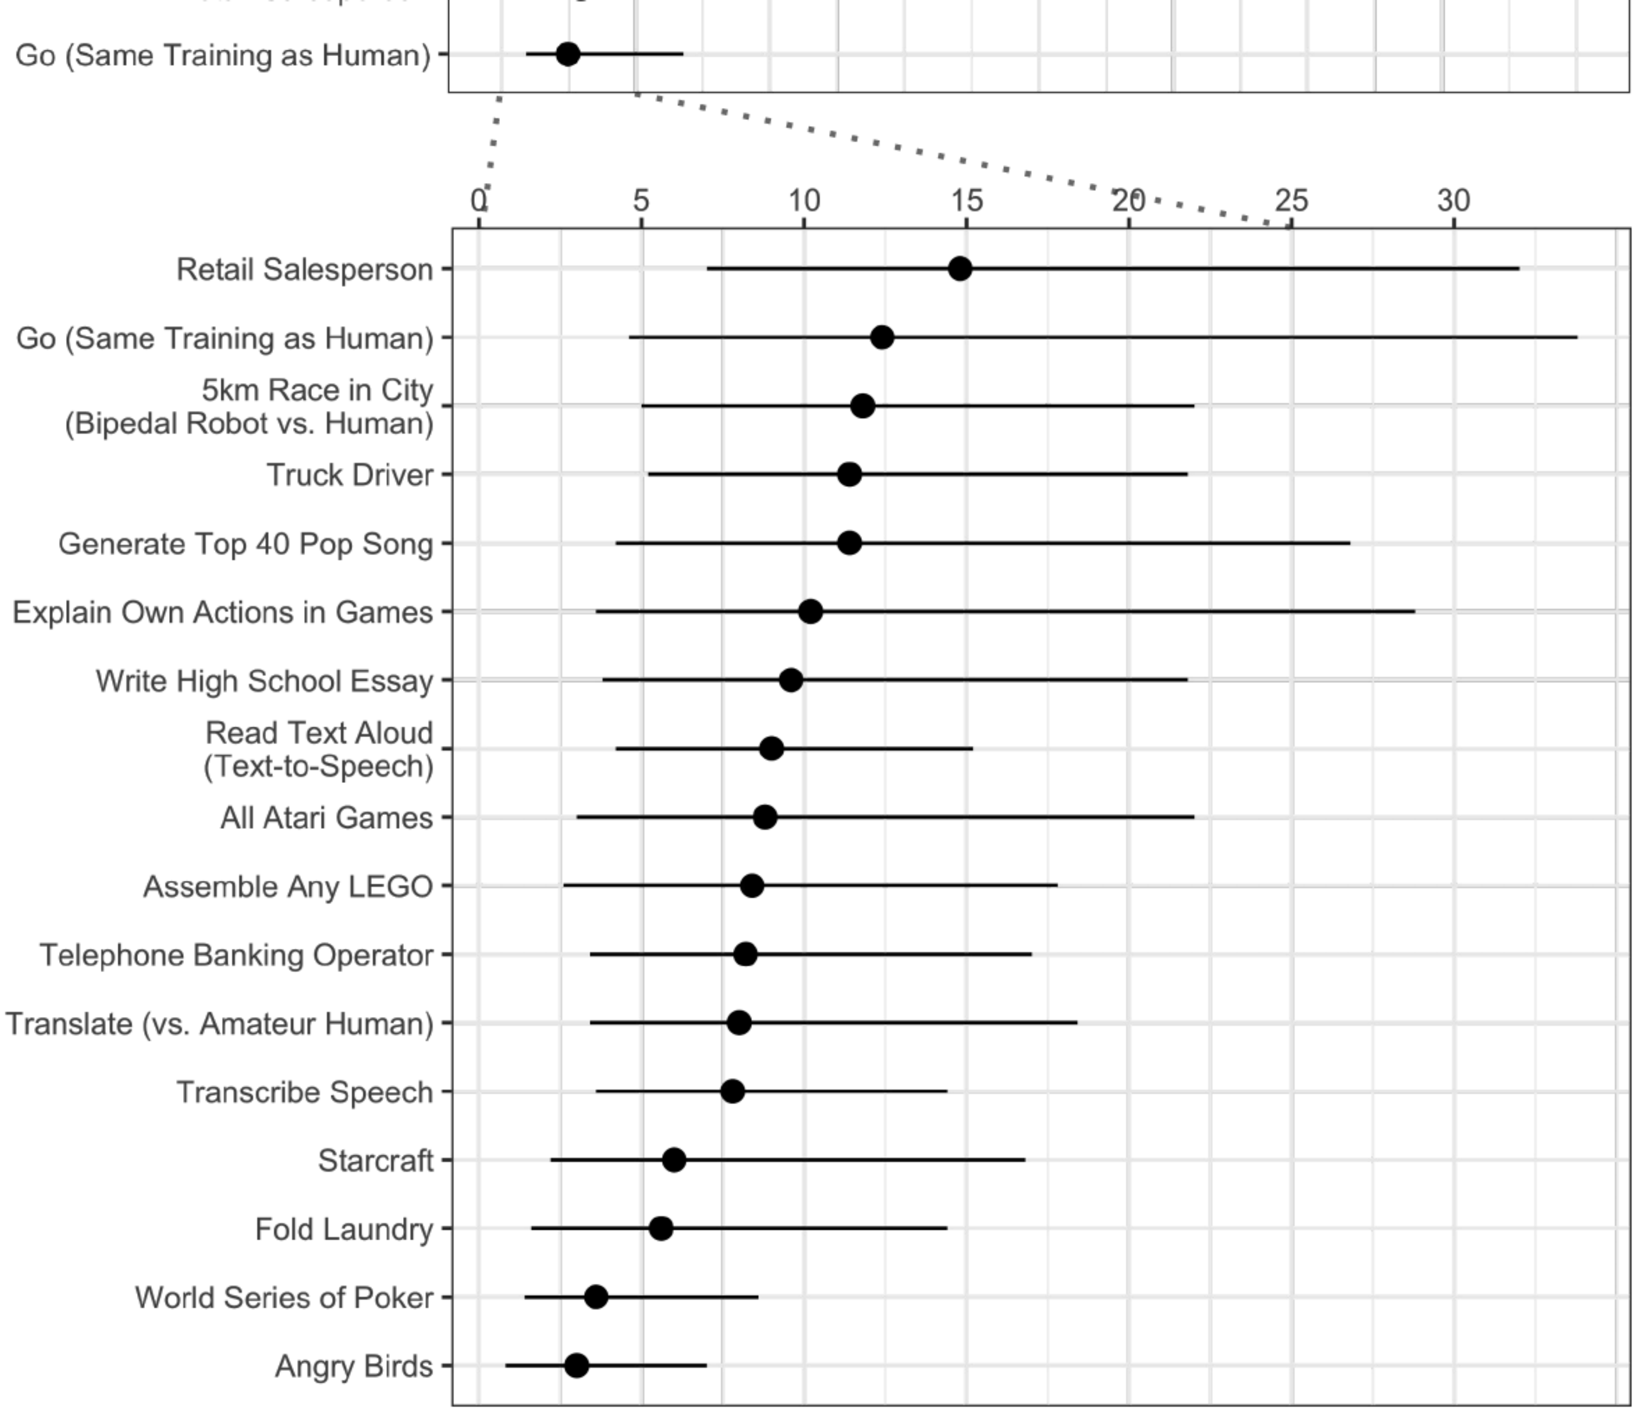

Go (Same Training as Human)
0
5
10
15
20
25
30
Retail Salesperson
Go (Same Training as Human)
5km Race in City (Bipedal Robot vs. Human)
Truck Driver
Generate Top 40 Pop Song
Explain Own Actions in Games
Write High School Essay
Read Text Aloud (Text-to-Speech)
All Atari Games
Assemble Any LEGO
Telephone Banking Operator
Translate (vs. Amateur Human)
Transcribe Speech
Starcraft
Fold Laundry
World Series of Poker
Angry Birds

ANNALS OF MEDICINE APRIL 3, 2017 ISSUE

# A.I. VERSUS M.D.

*What happens when diagnosis is automated?*

**By Siddhartha Mukherjee**

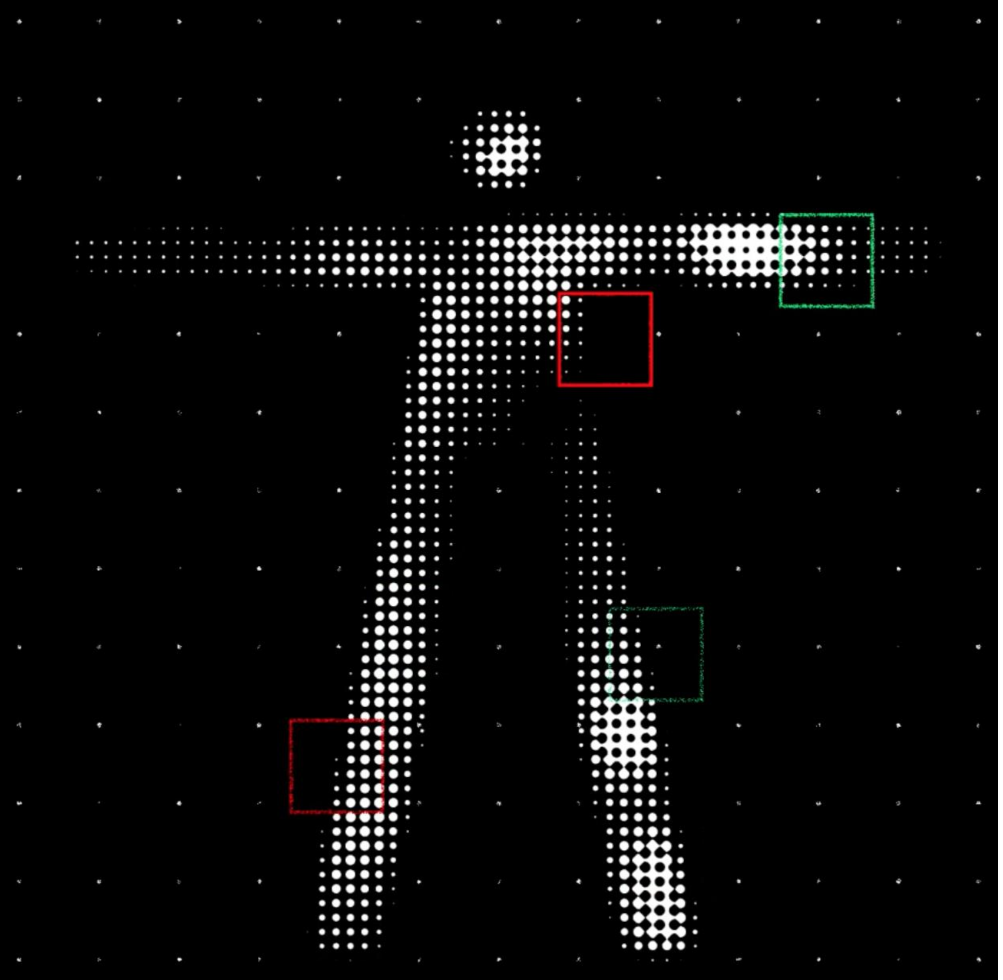

http://www.newyorker.com/magazine/2017/04/03/ai-versus-md

# How Are These Achievements Possible?

***What is Machine Learning?***

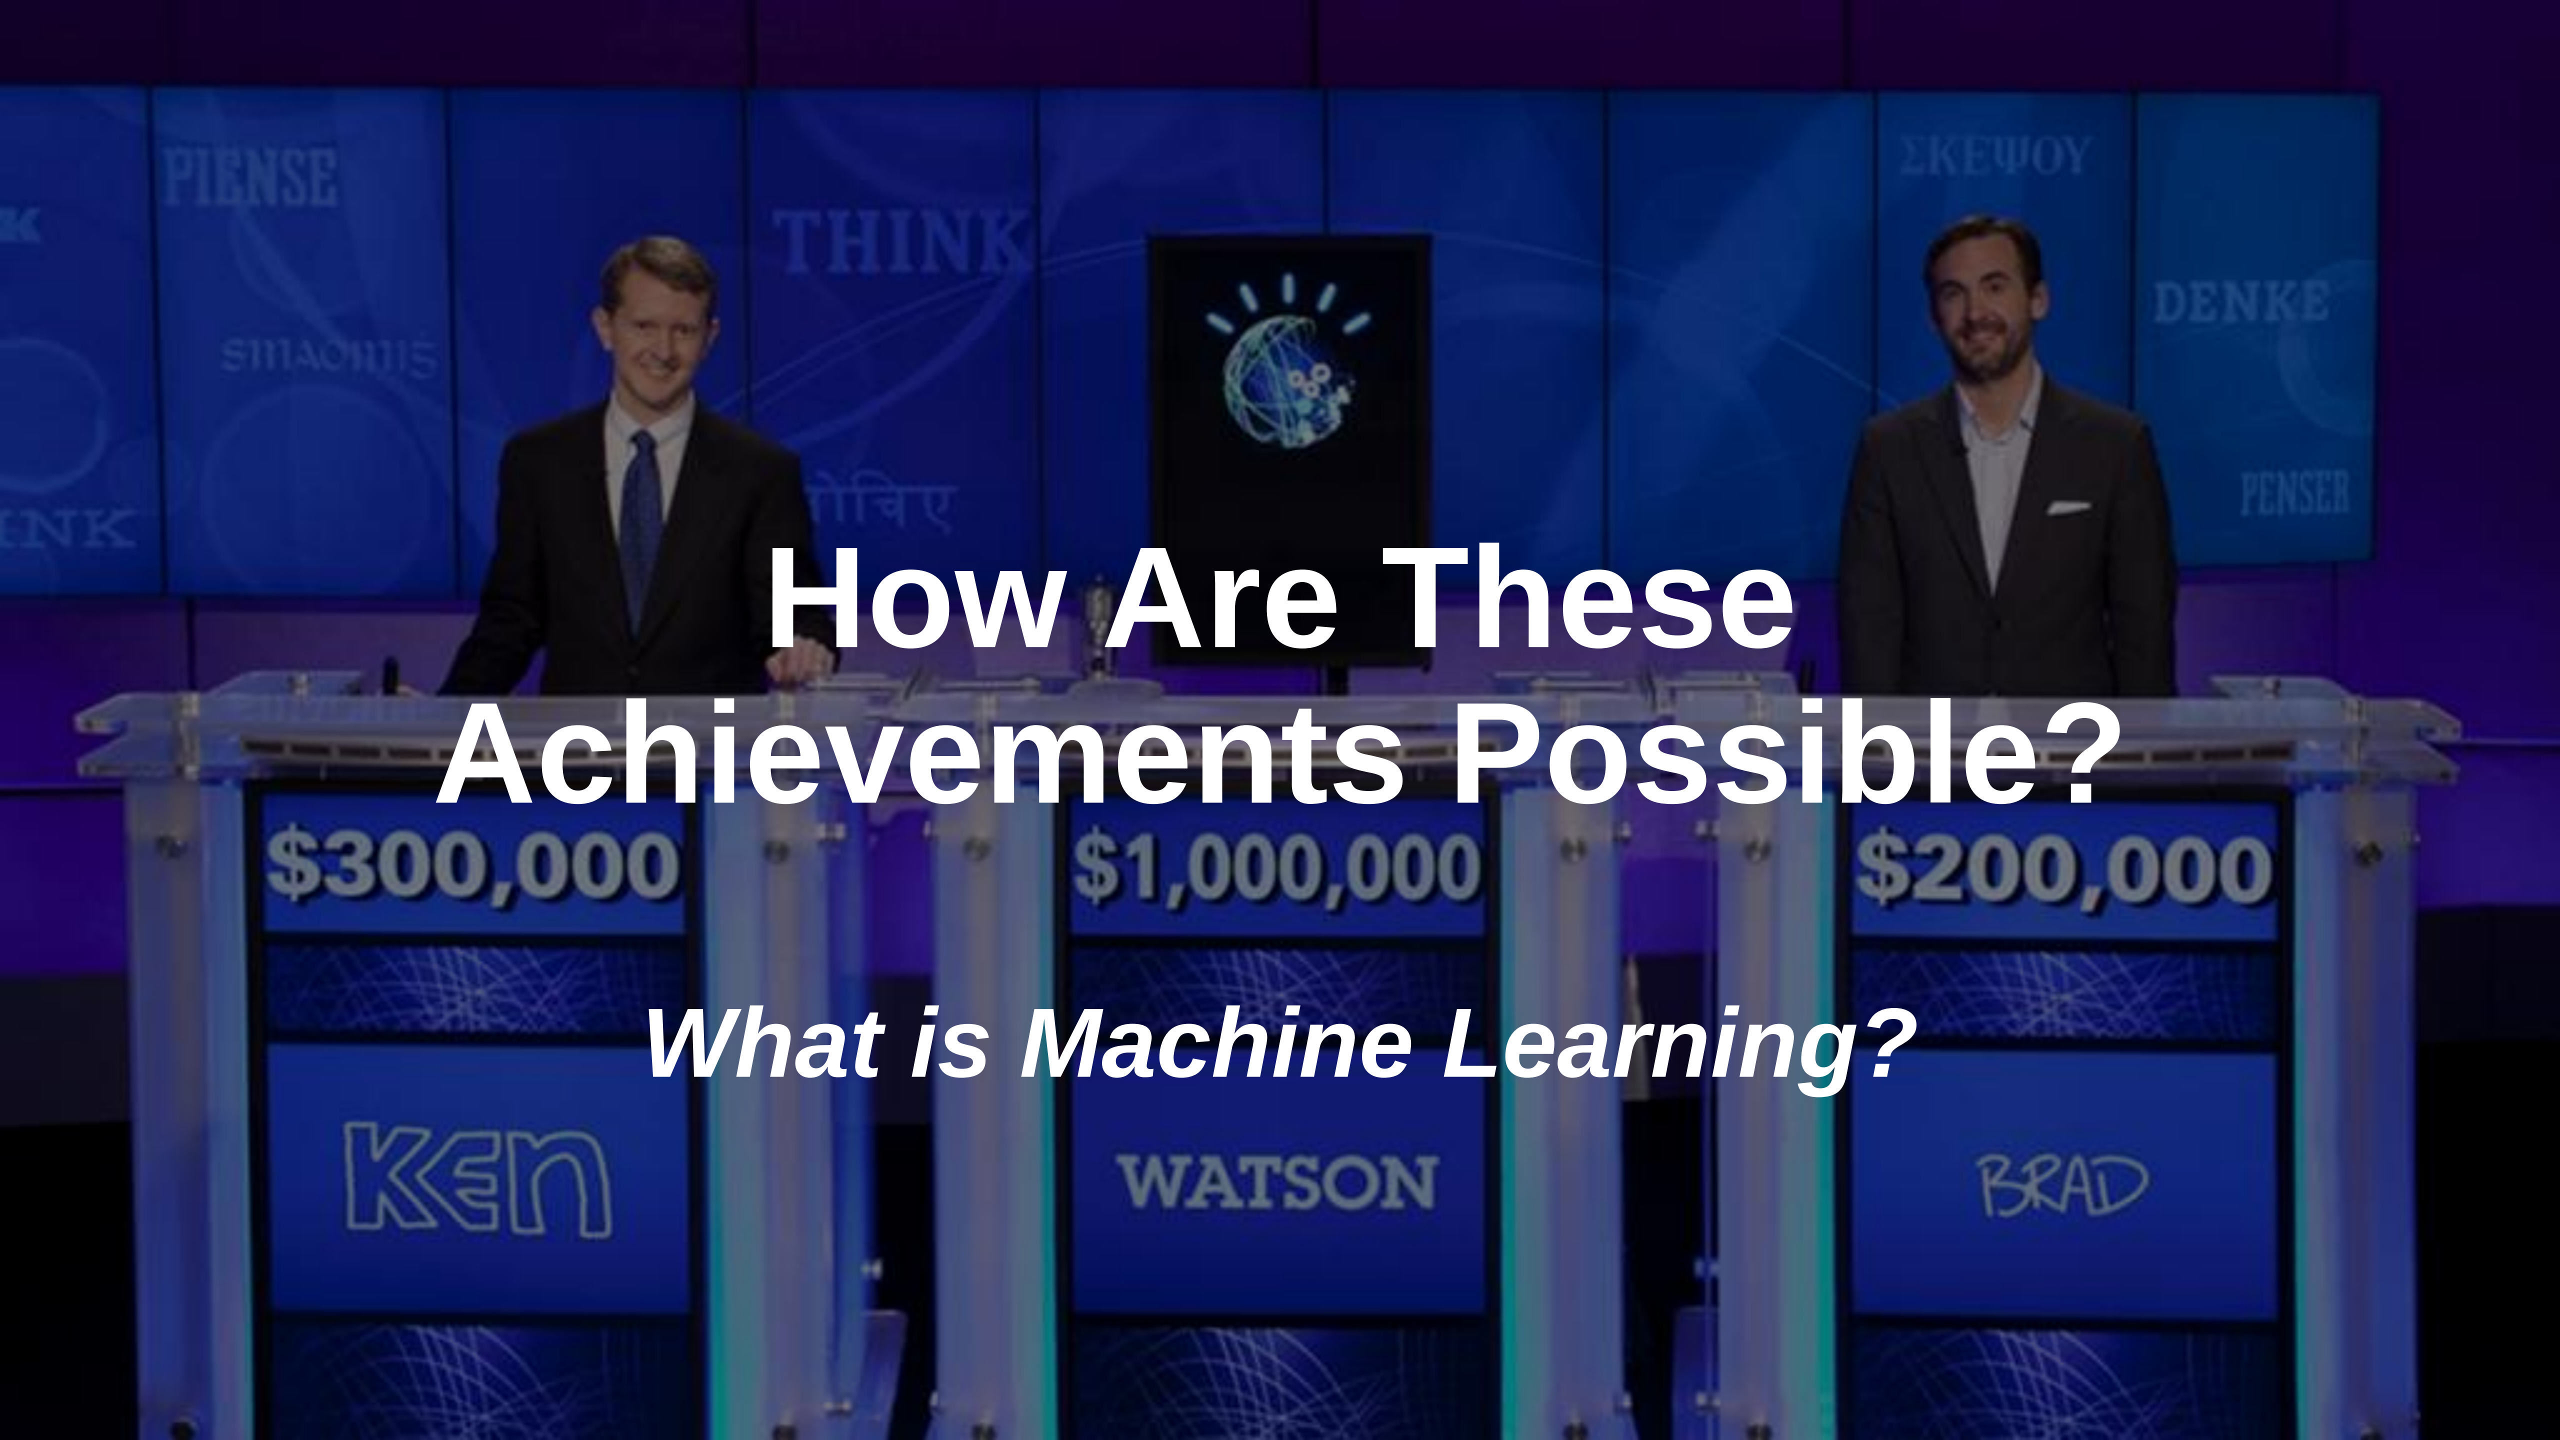

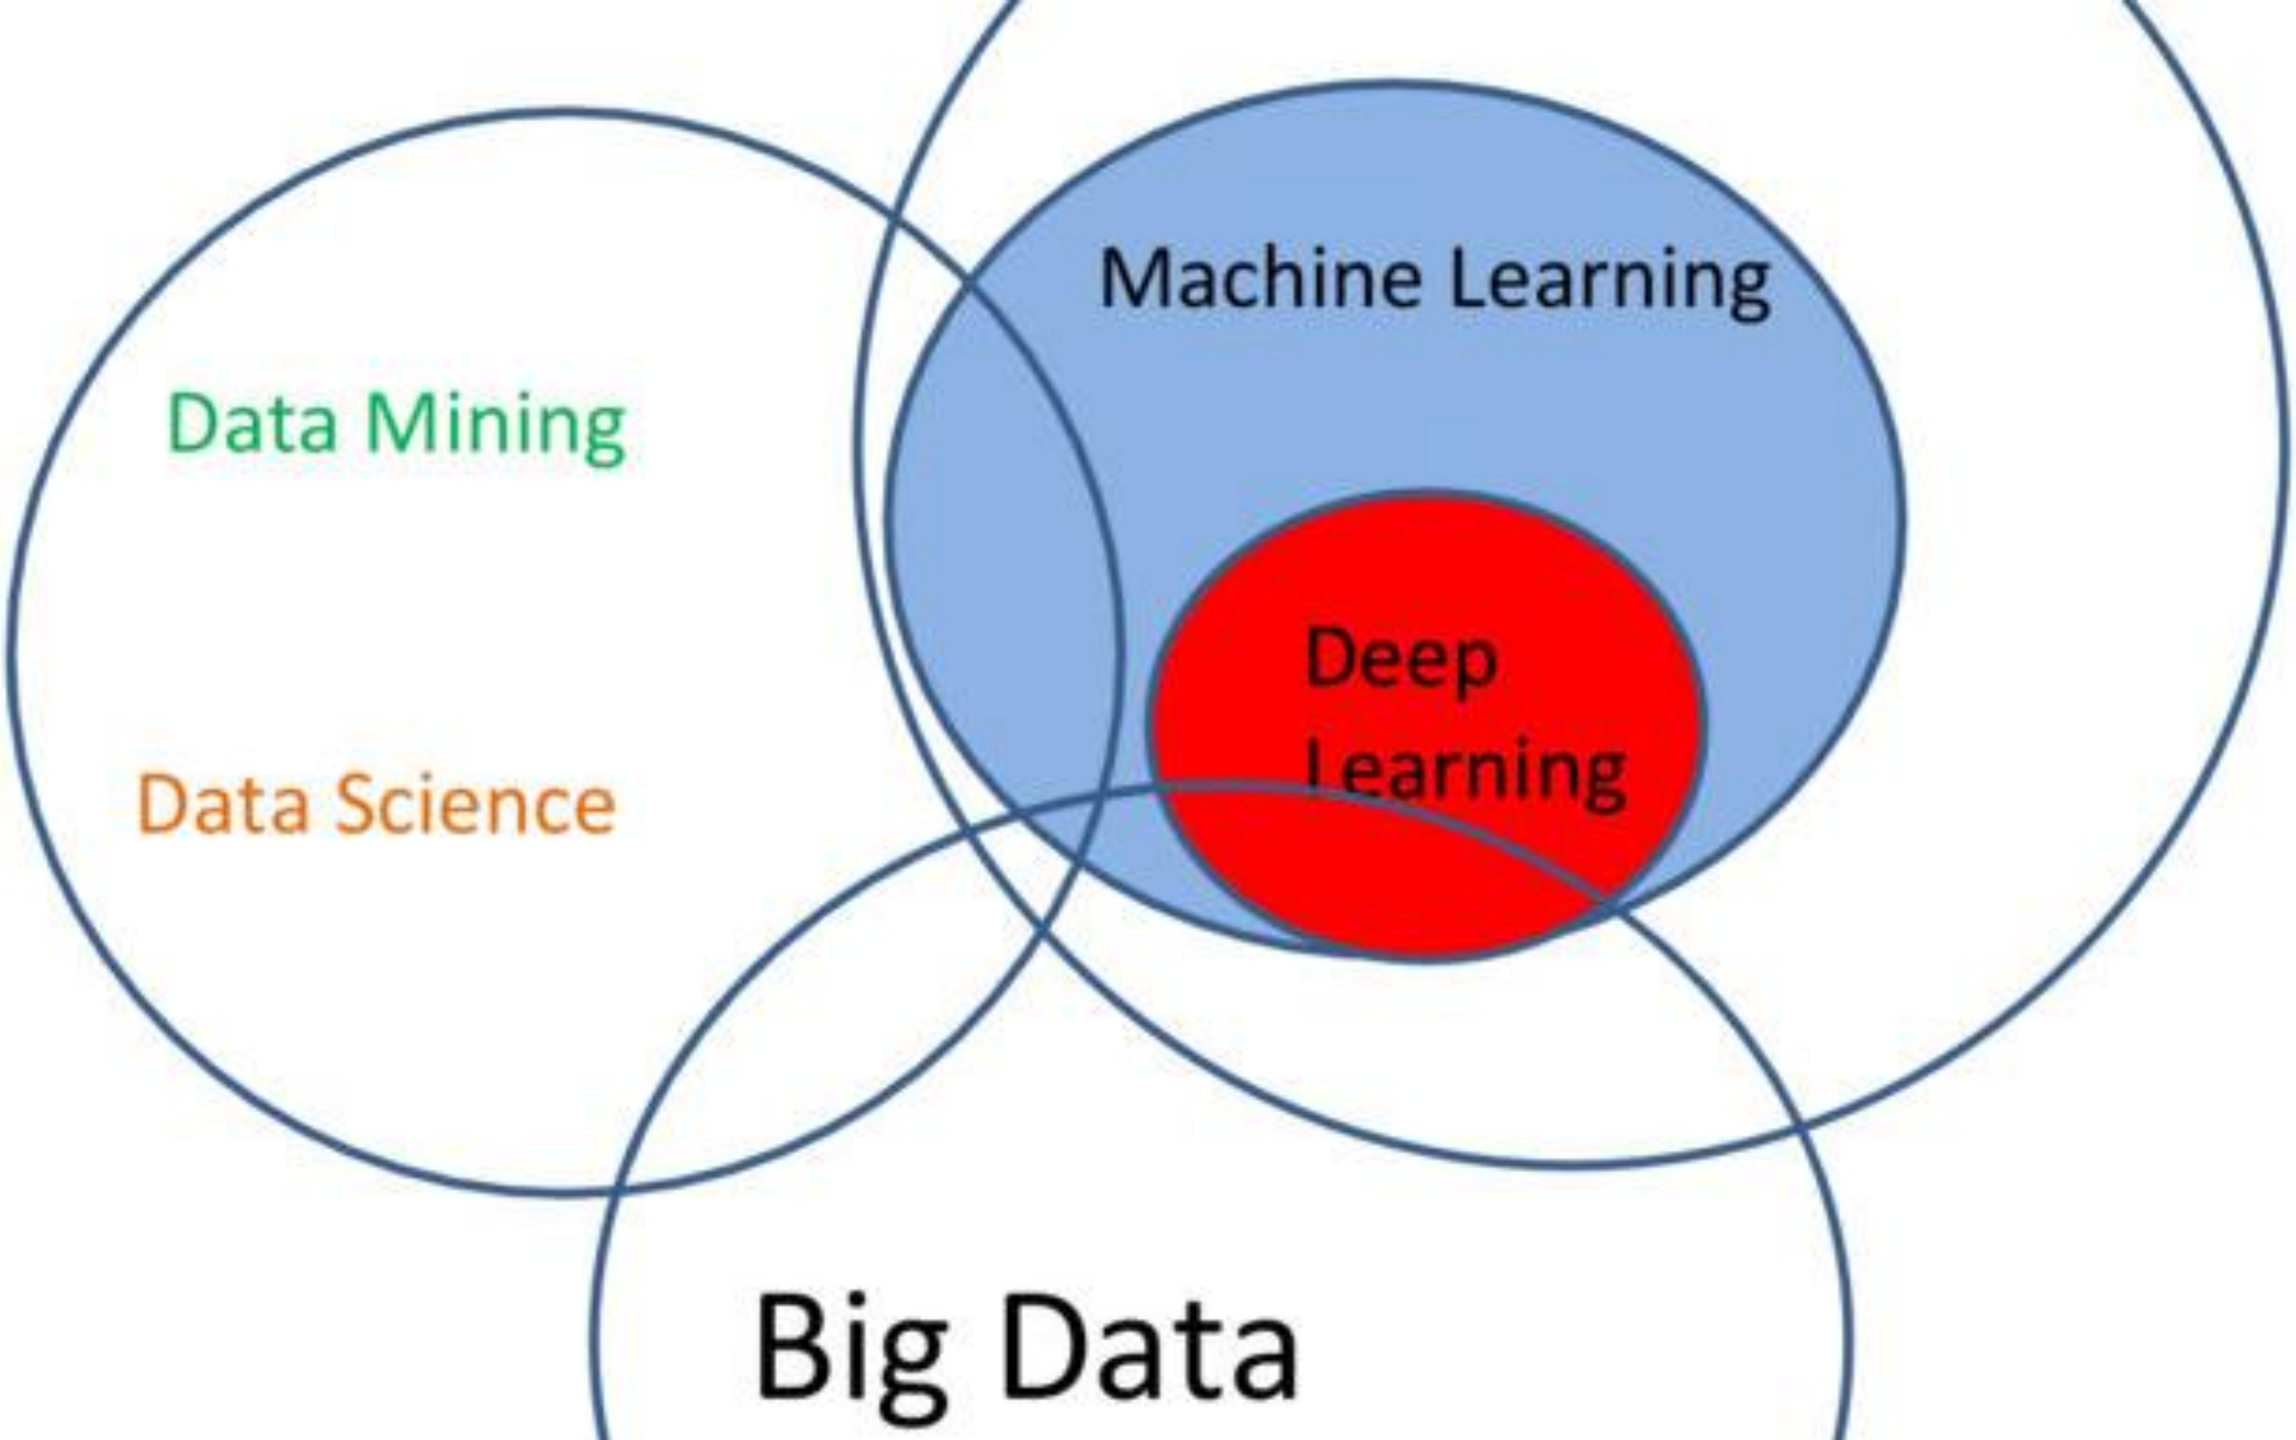
Machine Learning
Data Mining
Deep Learning
Data Science
Big Data

# 2010-2017

## ImageNet Challenge

IMAGENET

- 1,000 object classes (categories).
- Images:
  - 1.2 M train
  - 100k test.


4

Stanford Artificial Intelligence

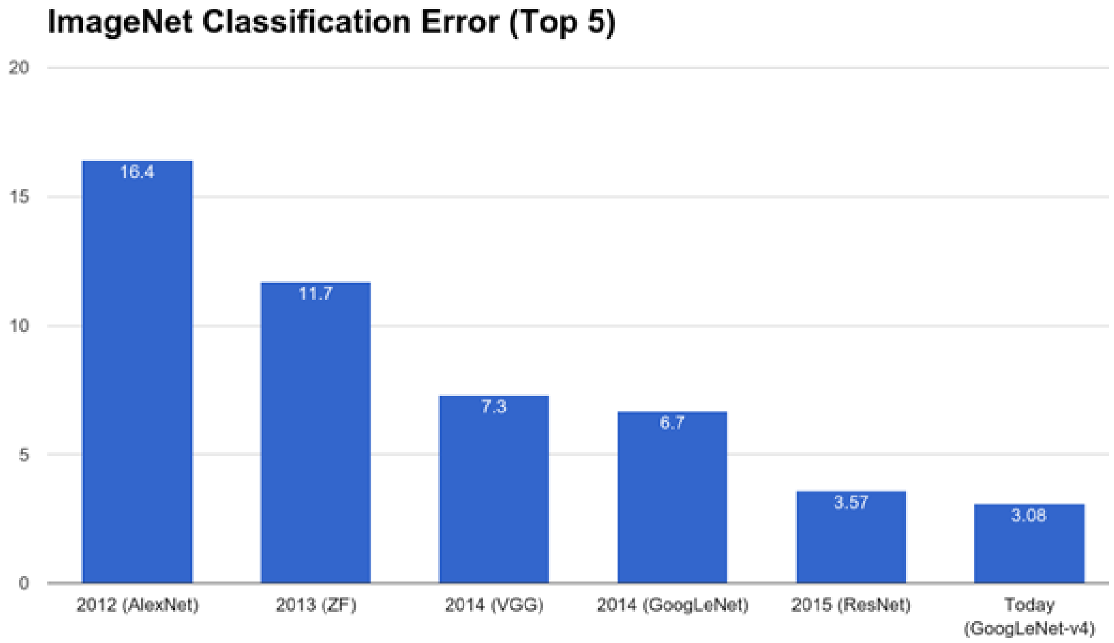

# Convolutional Neural Network (CNN)

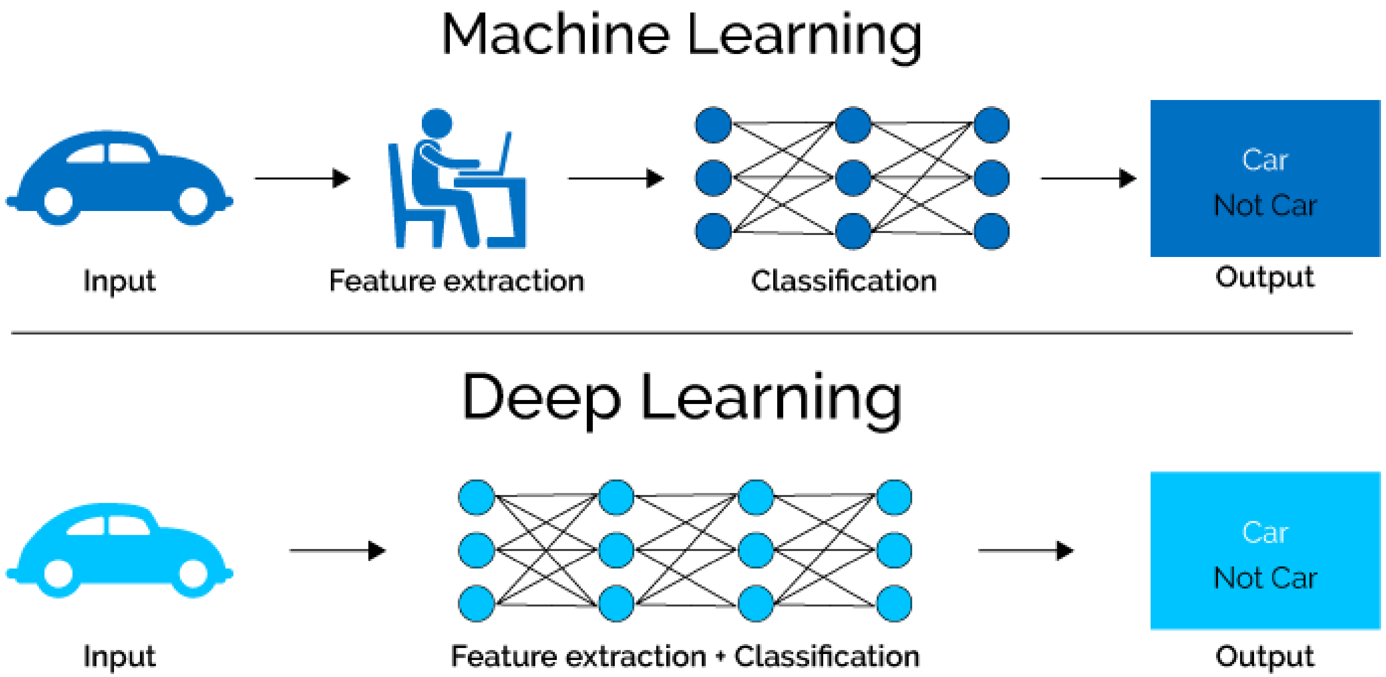

# Deep Learning Made Possible by

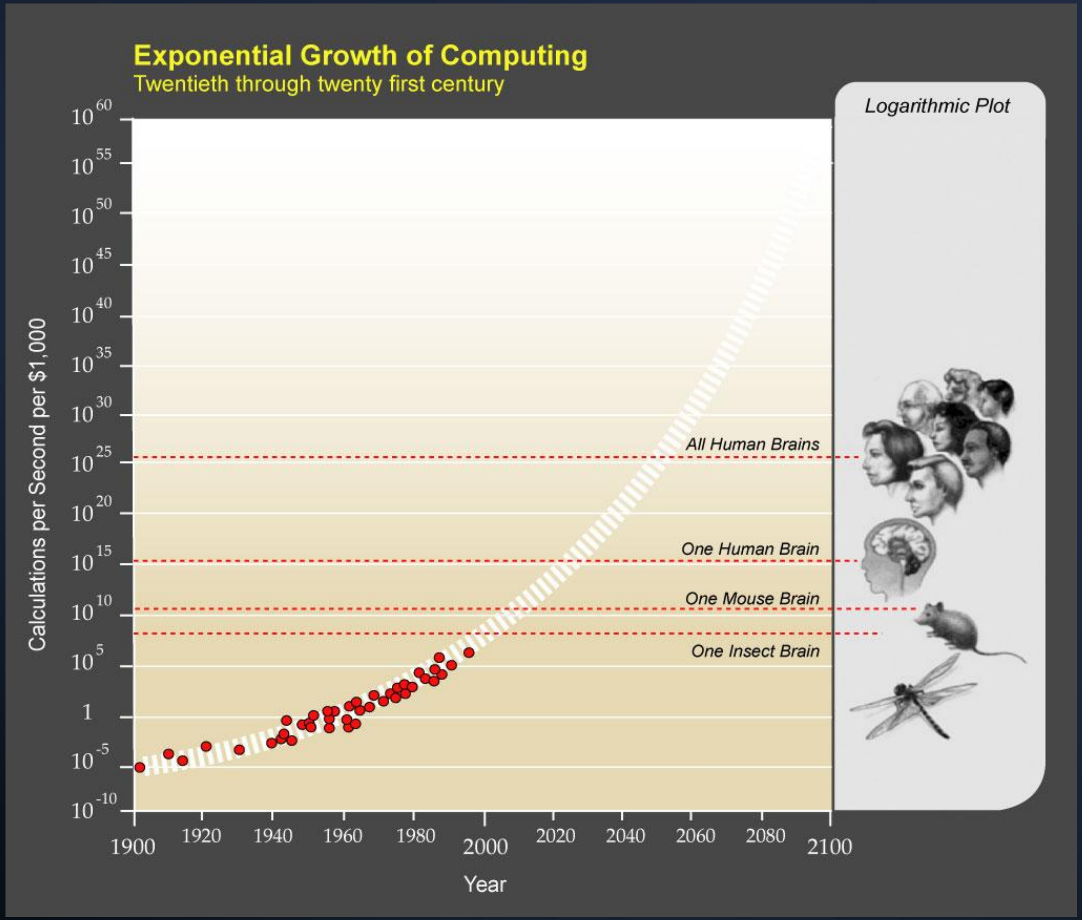


Compute Power

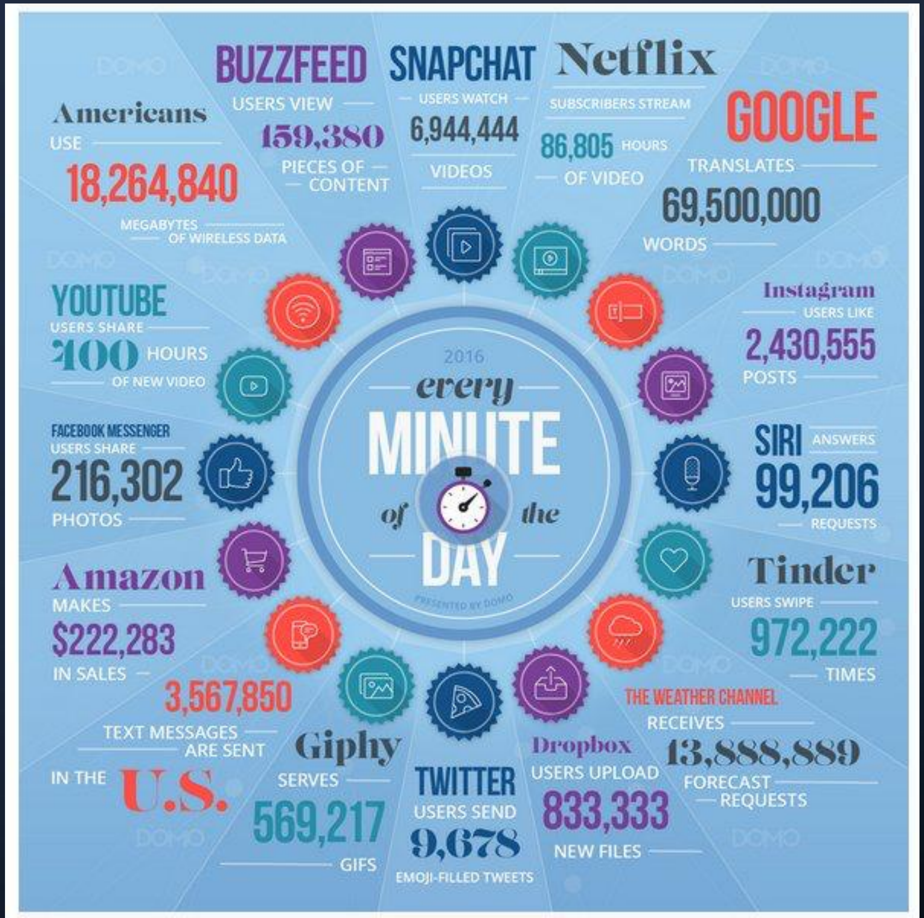


Data

# Perceptron Blueprint

1950s

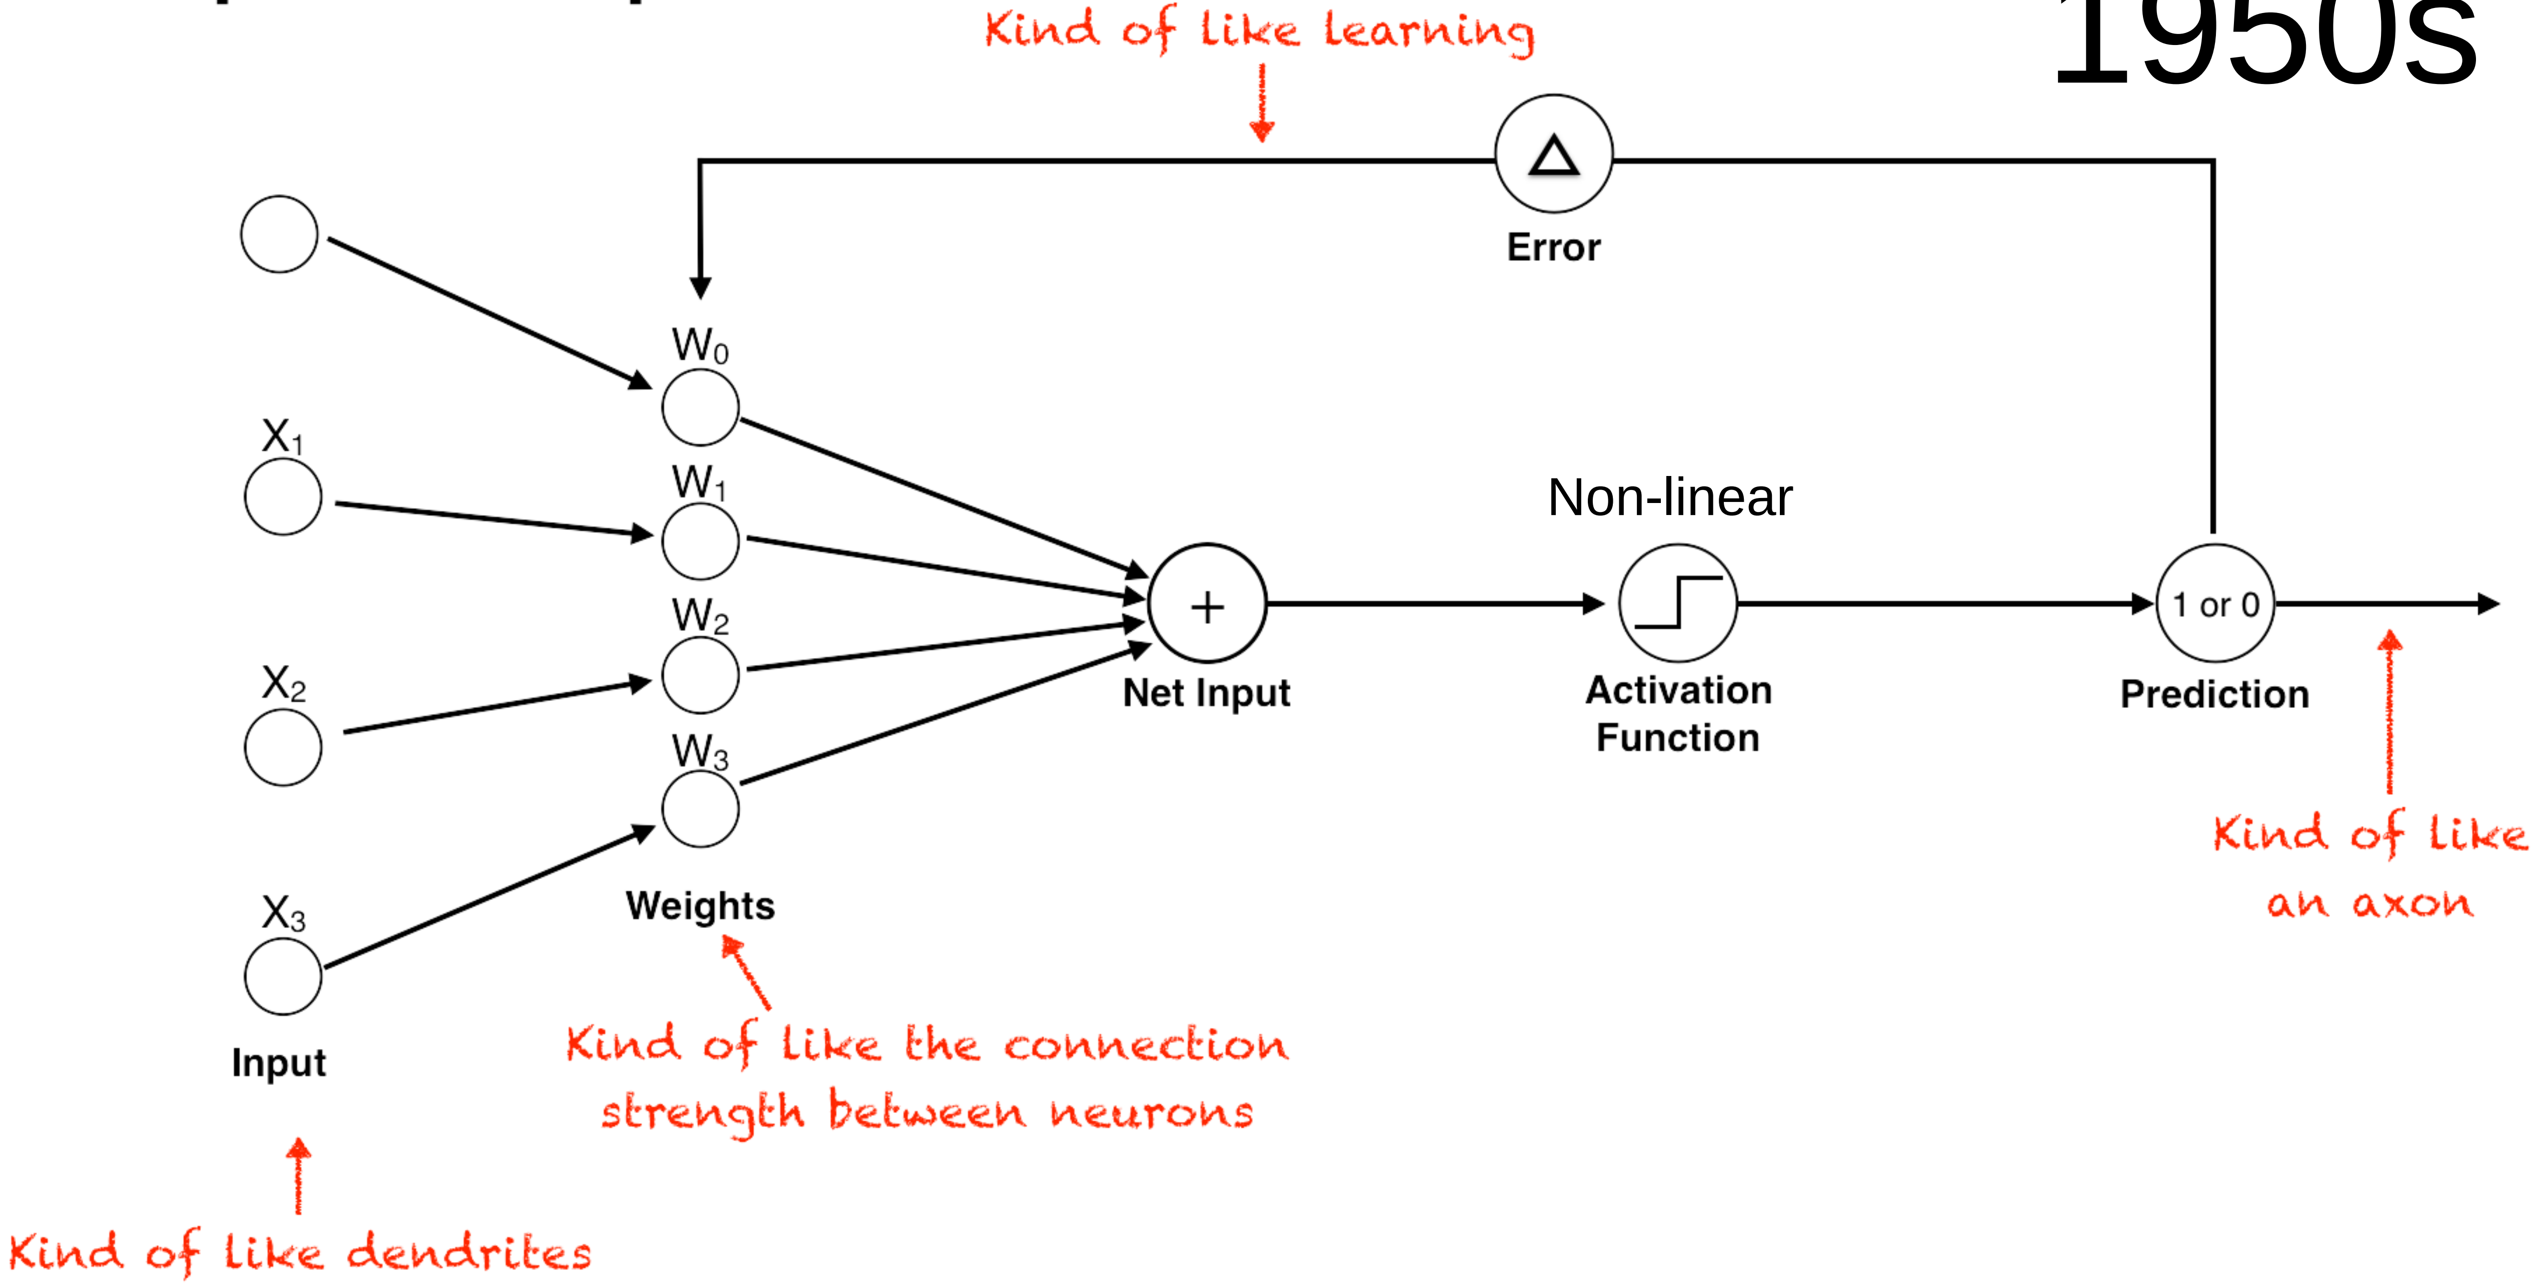

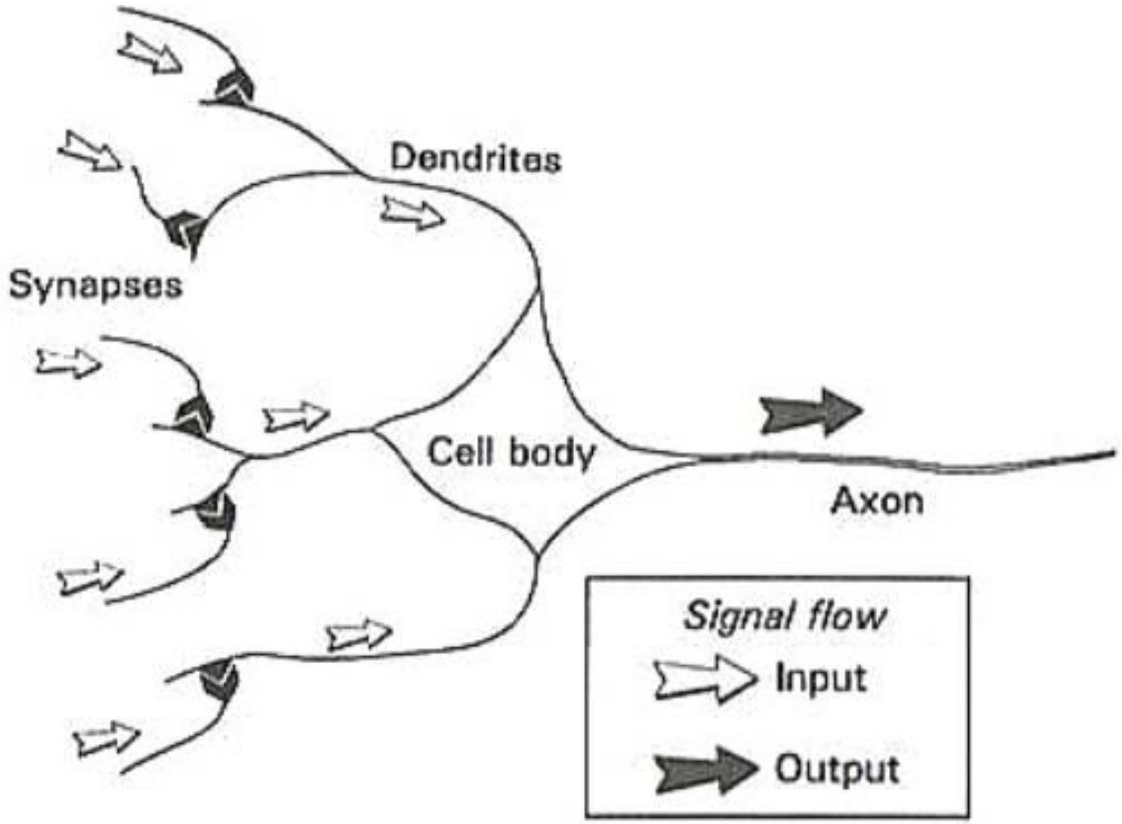


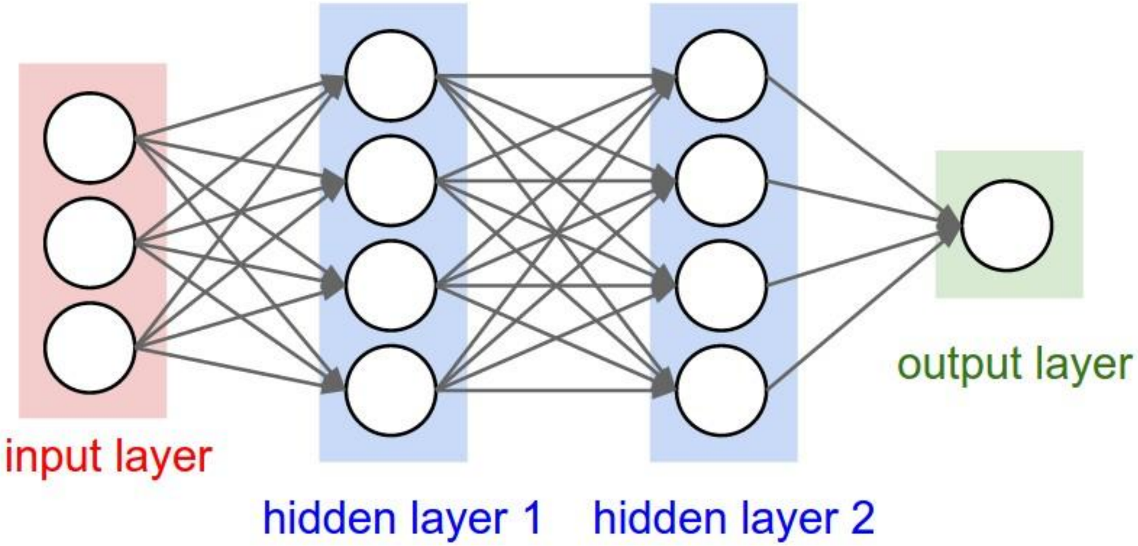


Biology Neural Nets

kaggle | Search kaggle | Competitions | Datasets | Kernels | Discussion | Jobs | ••• | Sign In

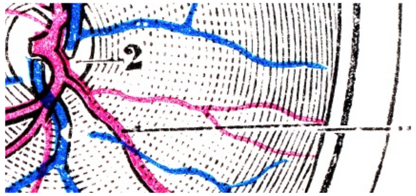


# Diabetic Retinopathy Detection

Identify signs of diabetic retinopathy in eye images

$100,000 · 661 teams · 2 years ago

2014

Overview | Data | Kernels | Discussion | Leaderboard | More | Submit Predictions

## Overview

- Description
- Evaluation
- Prizes
- References
- Timeline

**Diabetic retinopathy is the leading cause of blindness in the working-age population of the developed world. It is estimated to affect over 93 million people.**

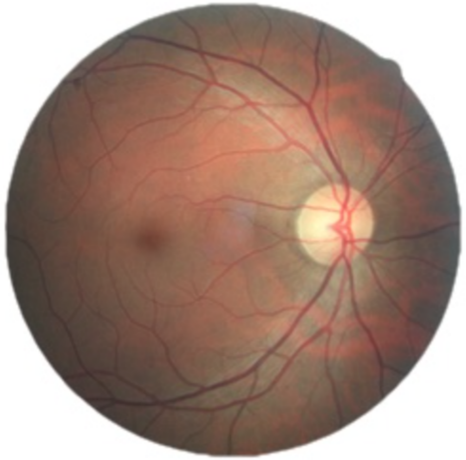

The US Center for Disease Control and Prevention estimates that 29.1 million people in the US have diabetes and the World Health Organization estimates that 347 million people have the disease worldwide. Diabetic Retinopathy (DR) is an eye disease associated with long-standing diabetes. Around 40% to 45% of Americans with diabetes have some stage of the disease. Progression to vision impairment can be slowed or averted if DR is detected in time, however this can be difficult as the disease often shows few symptoms until it is too late to provide effective treatment.

Currently, detecting DR is a time-consuming and manual process that requires a trained clinician to examine and evaluate digital color fundus photographs of the retina. By the time human readers submit their reviews, often a day or two later, the delayed results lead to lost follow up, miscommunication, and delayed treatment.

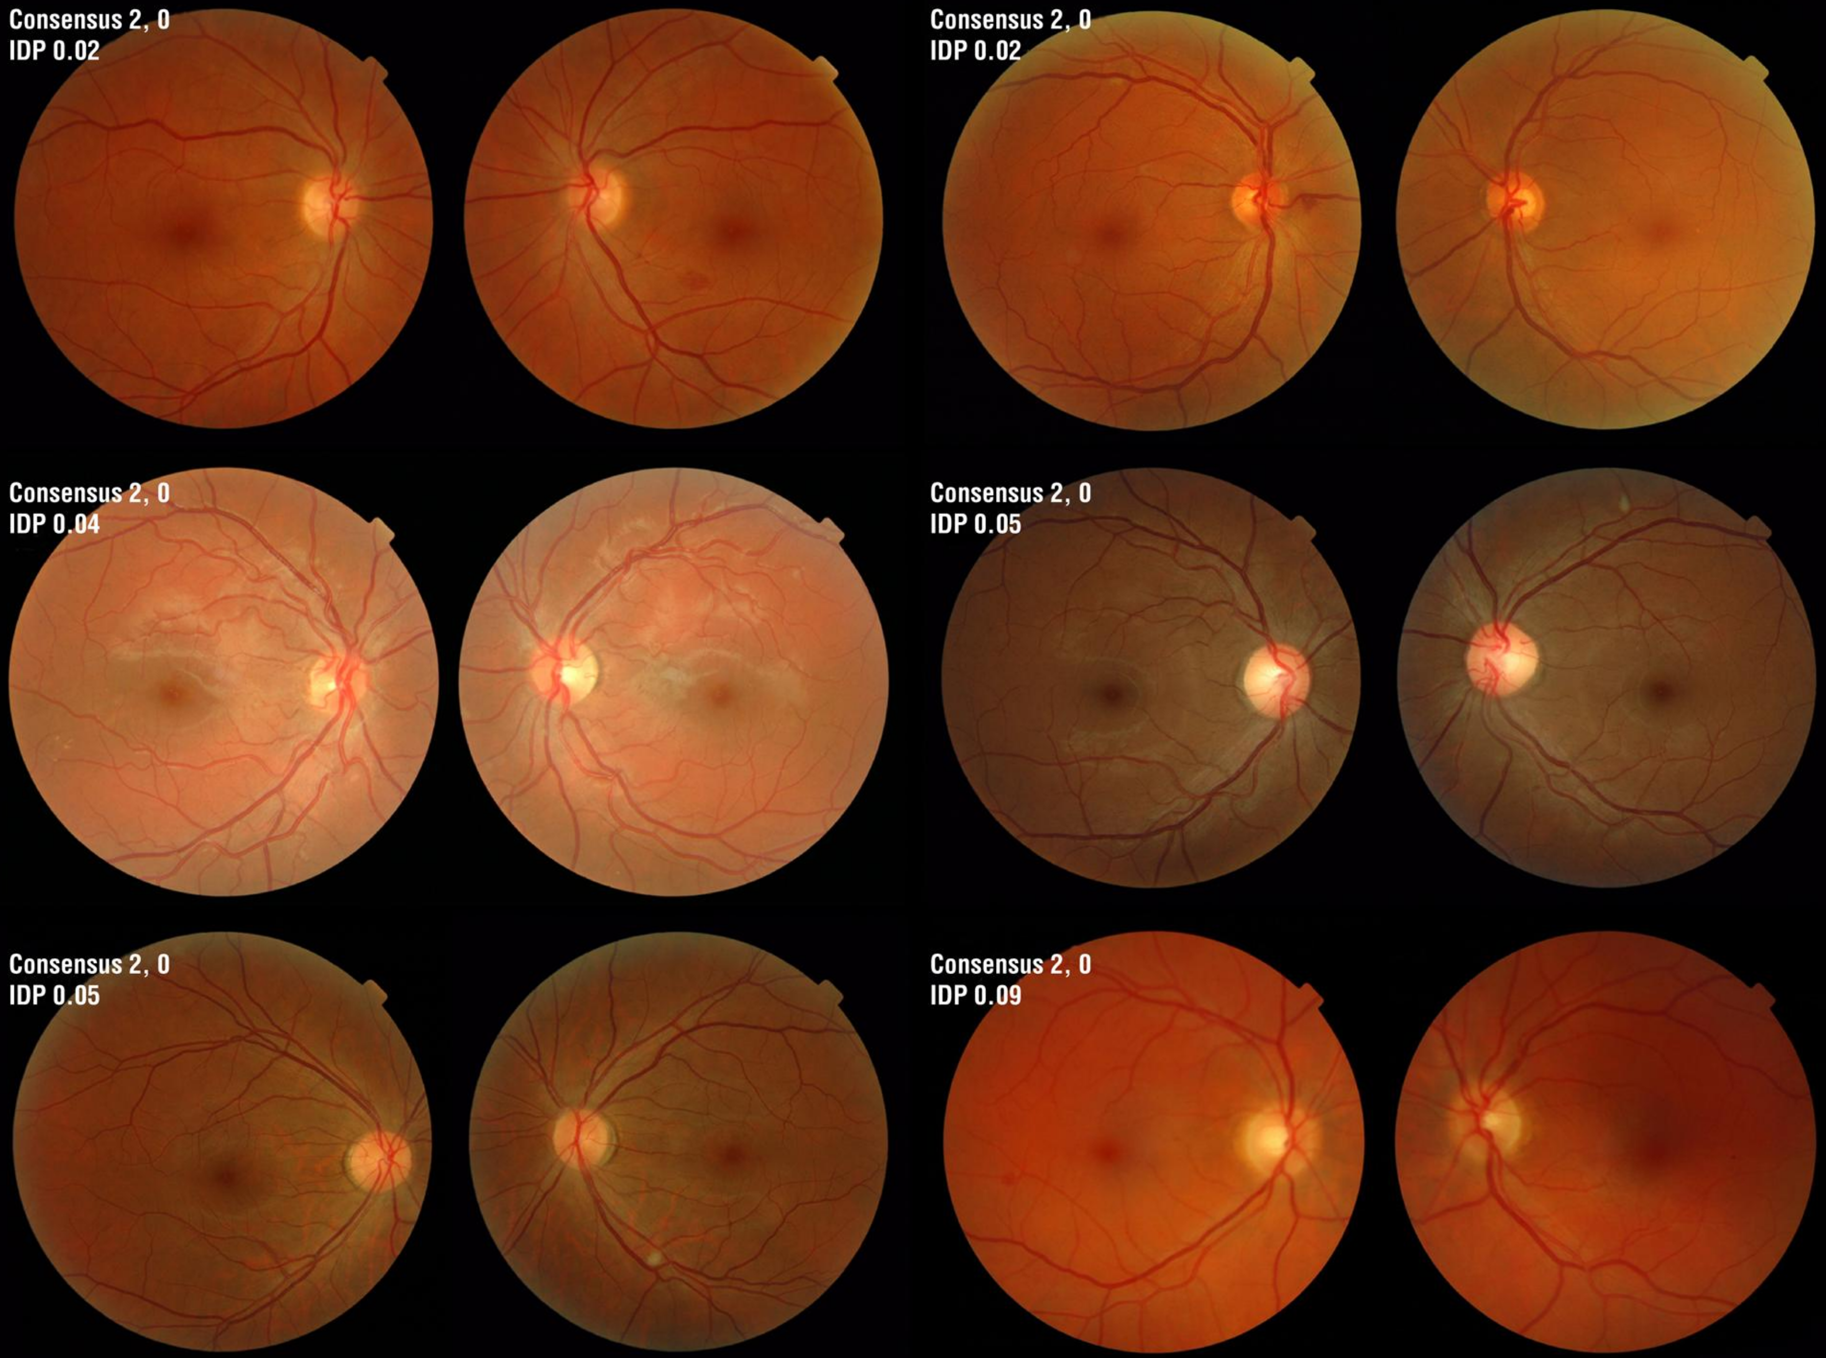
Consensus 2, 0
IDP 0.02
Consensus 2, 0
IDP 0.02
Consensus 2, 0
IDP 0.04
Consensus 2, 0
IDP 0.05
Consensus 2, 0
IDP 0.05
Consensus 2, 0
IDP 0.09

# Deep Learning = Learning Hierarchical Representations

Y LeCun

- **Traditional Pattern Recognition:** Fixed/Handcrafted Feature Extractor

Feature Extractor → Trainable Classifier →

- **Mainstream Modern Pattern Recognition:** Unsupervised mid-level features

Feature Extractor → Mid-Level Features → Trainable Classifier →

- **Deep Learning:** Representations are hierarchical and trained

Low-Level Features → Mid-Level Features → High-Level Features → Trainable Classifier →

# The Past

- *Handcrafted Features* Labeled by Experts
- Thus 4Qs Hemes = Severe NPDR

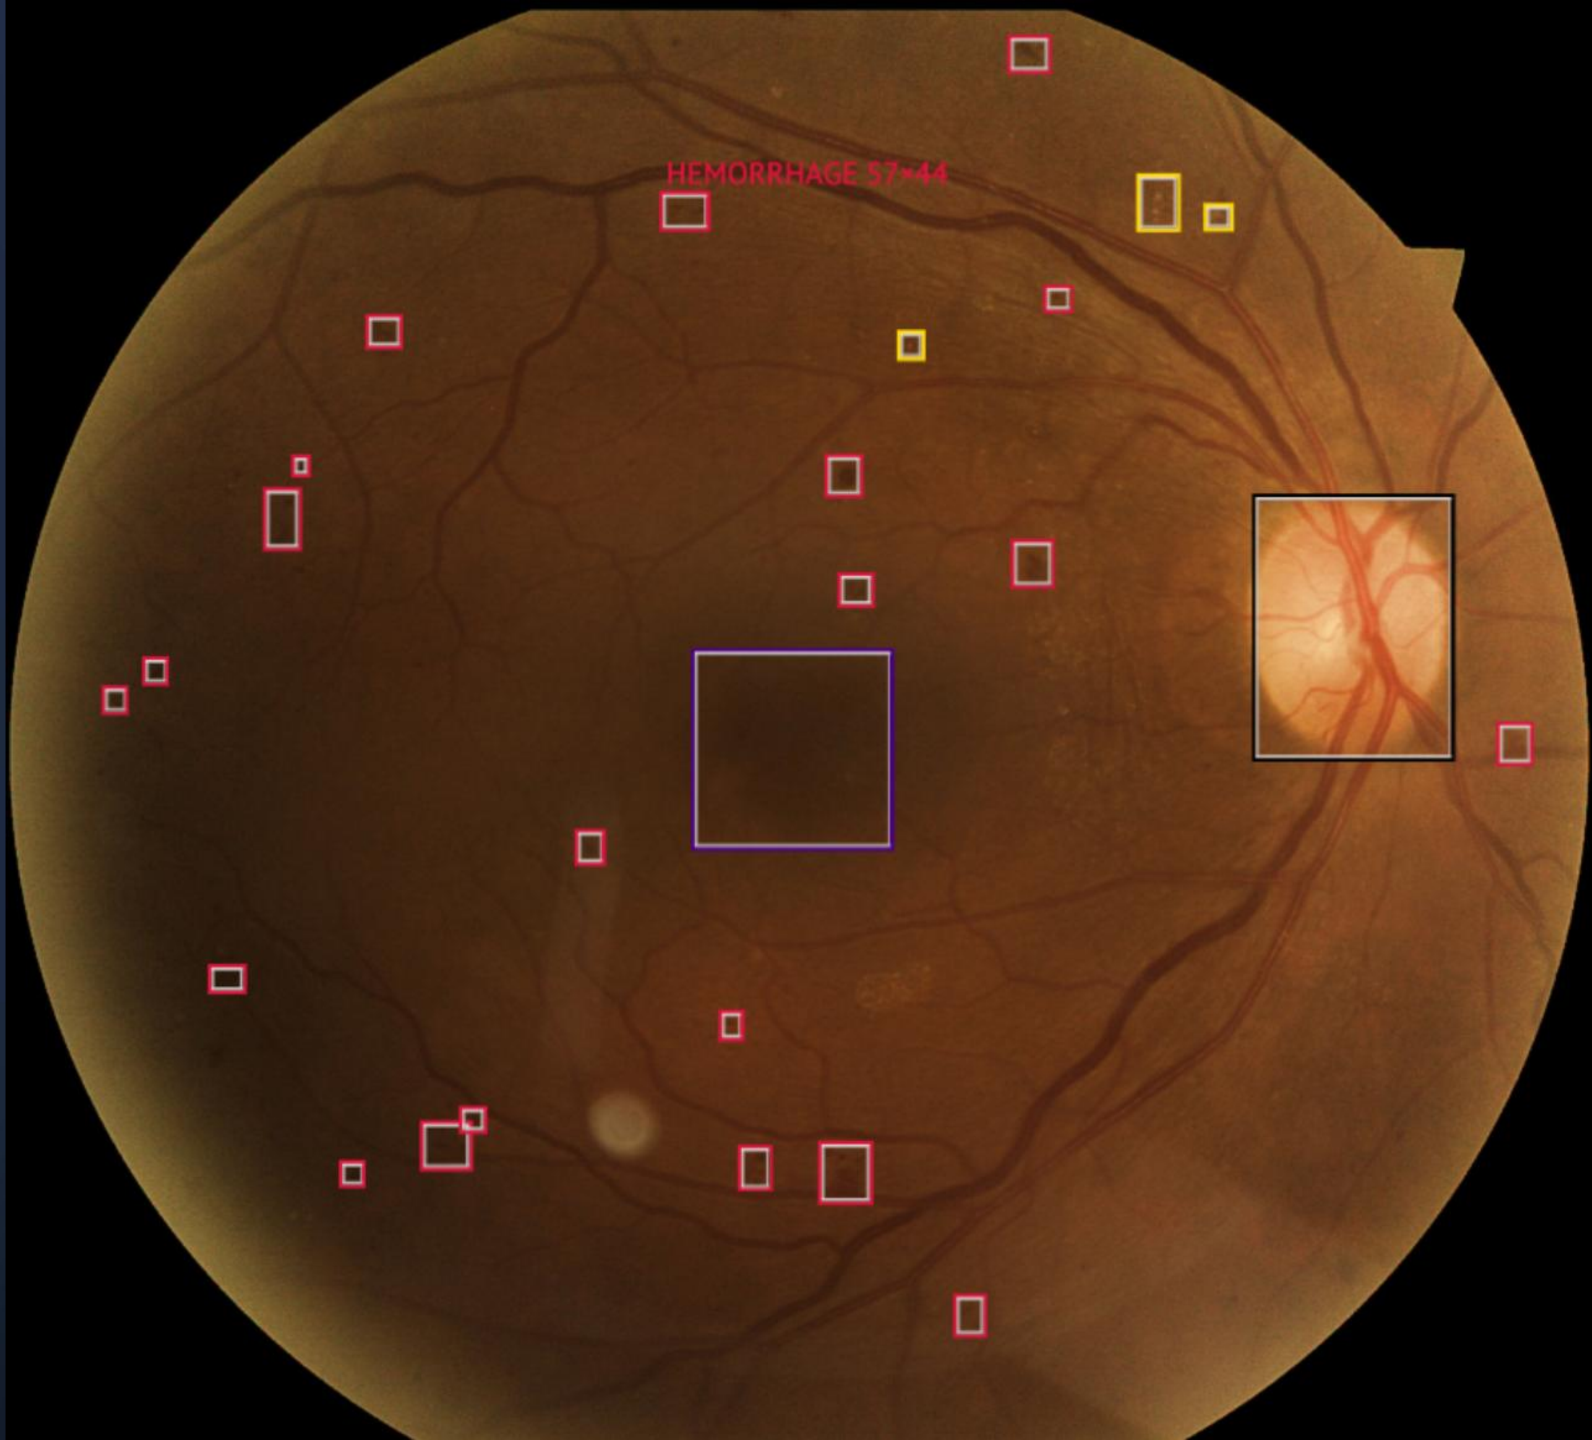

# The Present – Supervised Learning

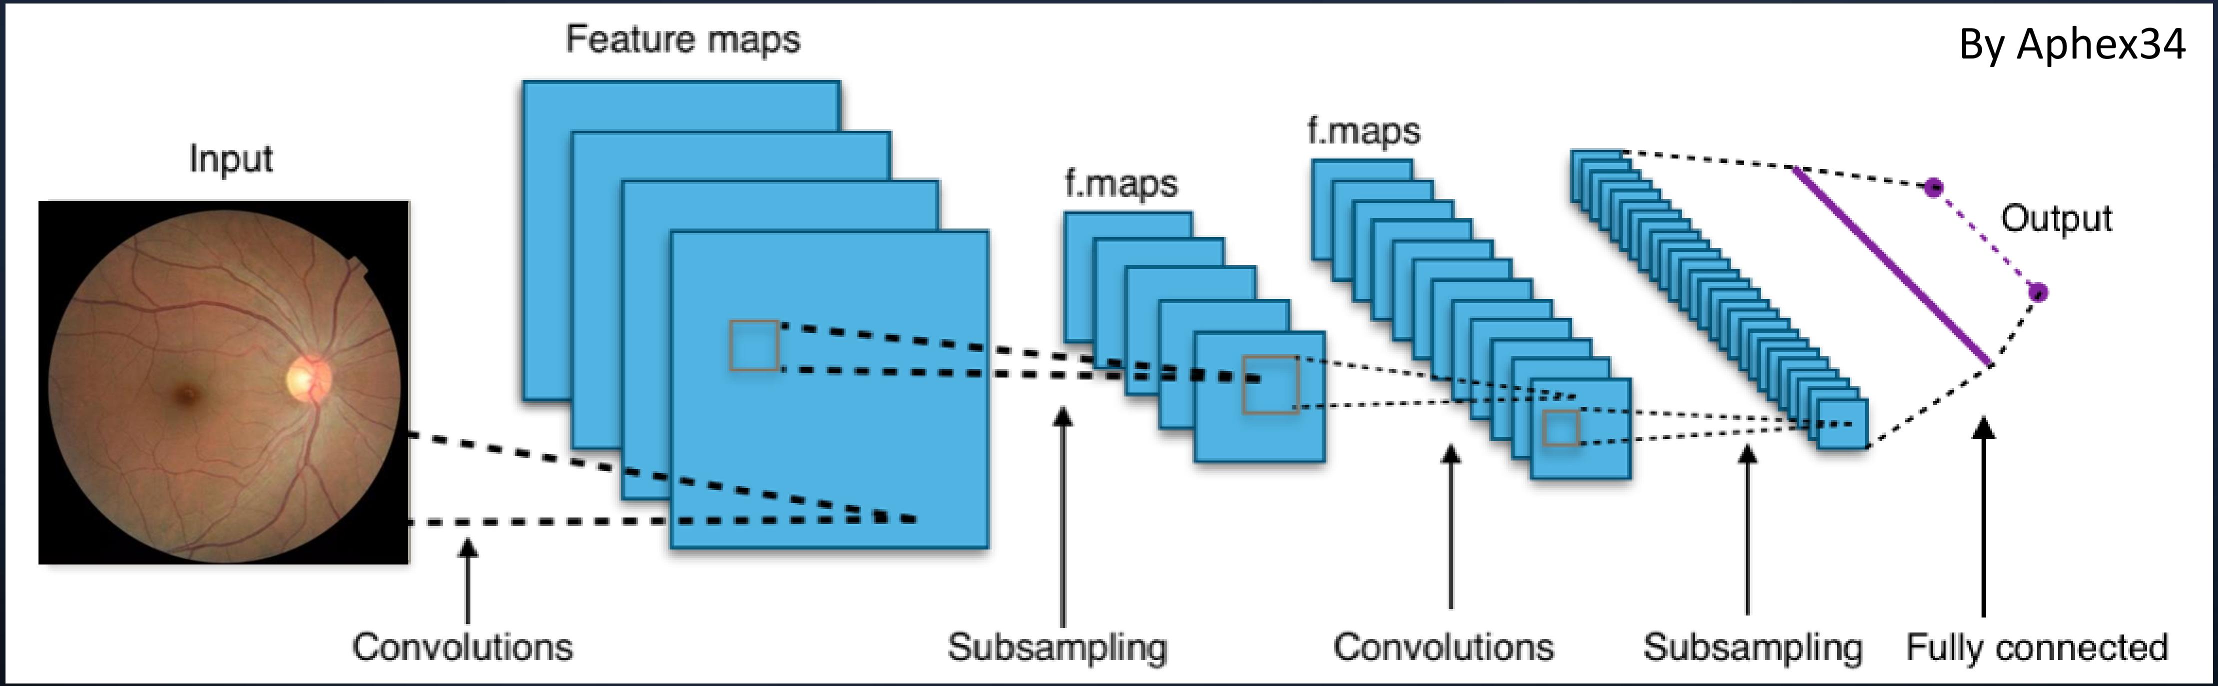


*Here are 10,000 Examples of Severe NPDR…find the pattern.*

# HOW A DEEP NEURAL NETWORK SEES

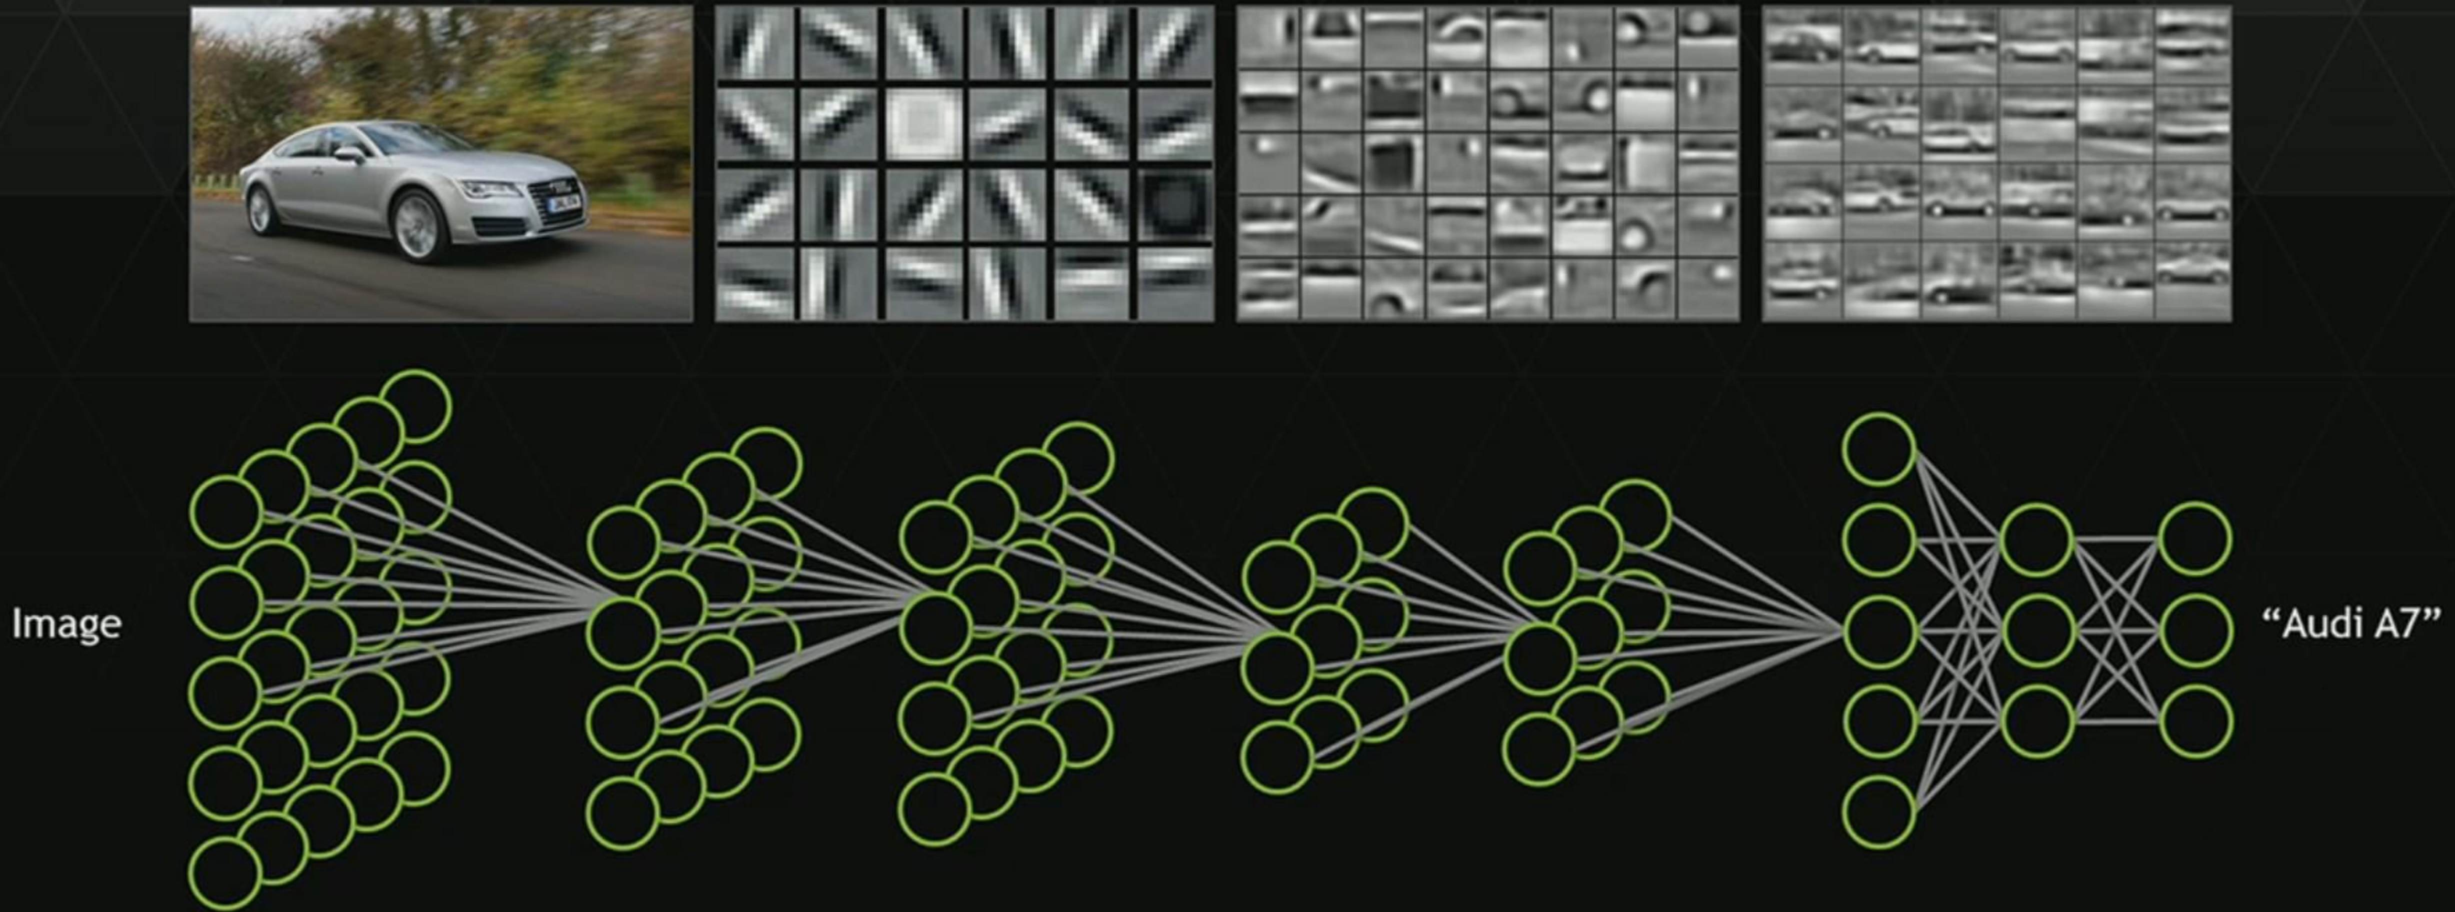


*Image source: “Unsupervised Learning of Hierarchical Representations with Convolutional Deep Belief Networks” ICML 2009 & Comm. ACM 2011.*
*Honglak Lee, Roger Grosse, Rajesh Ranganath, and Andrew Ng.*

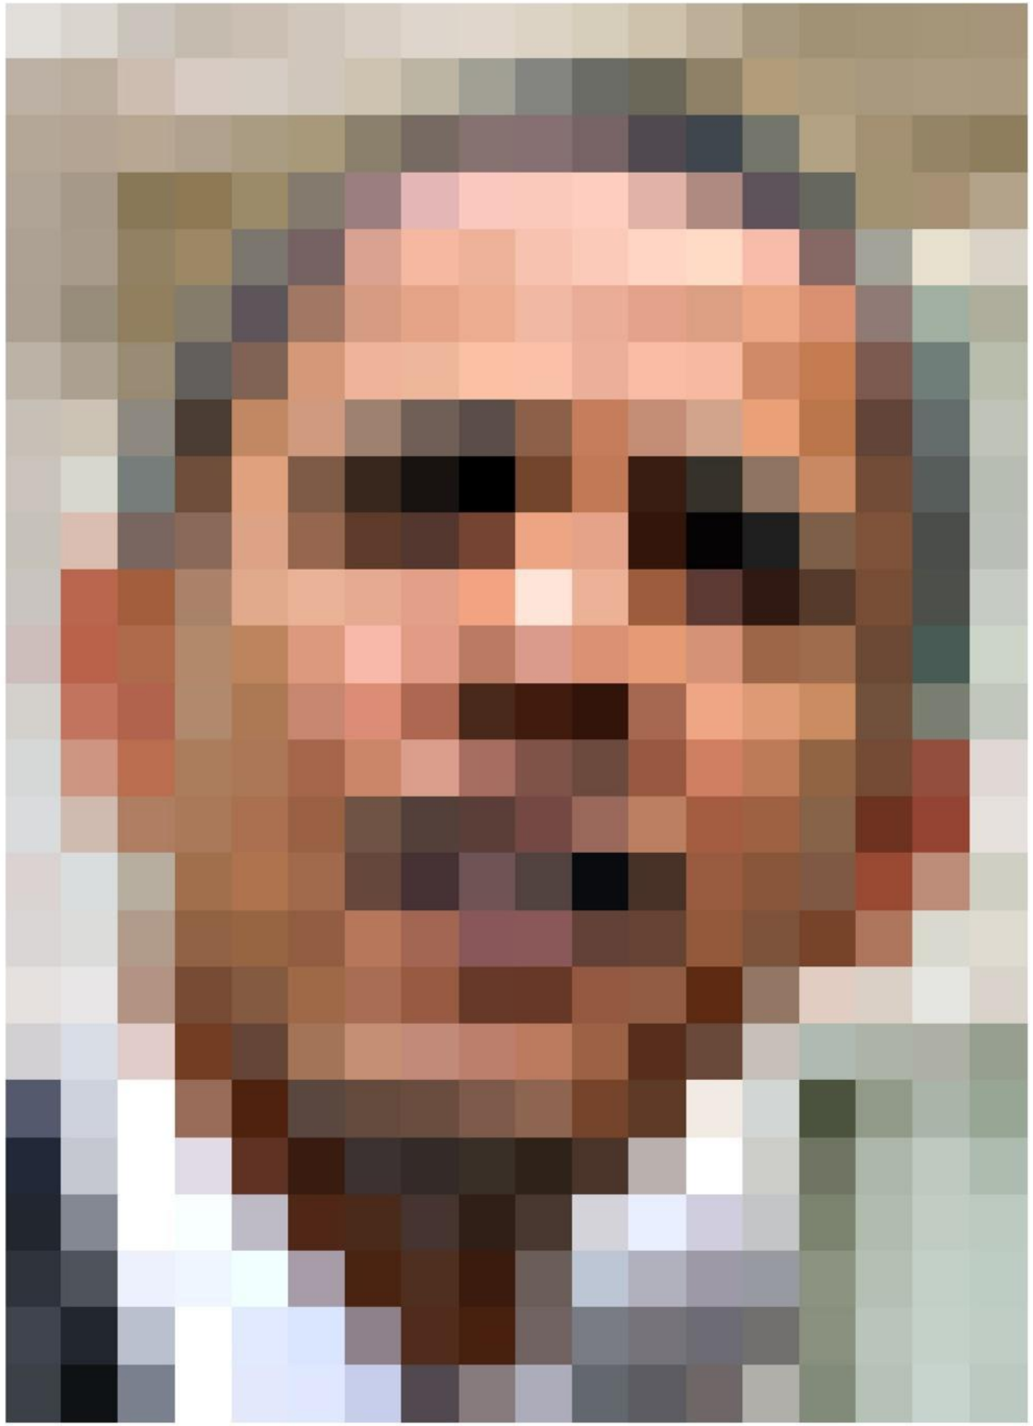

http://ihome.ust.hk/~yzhouas/

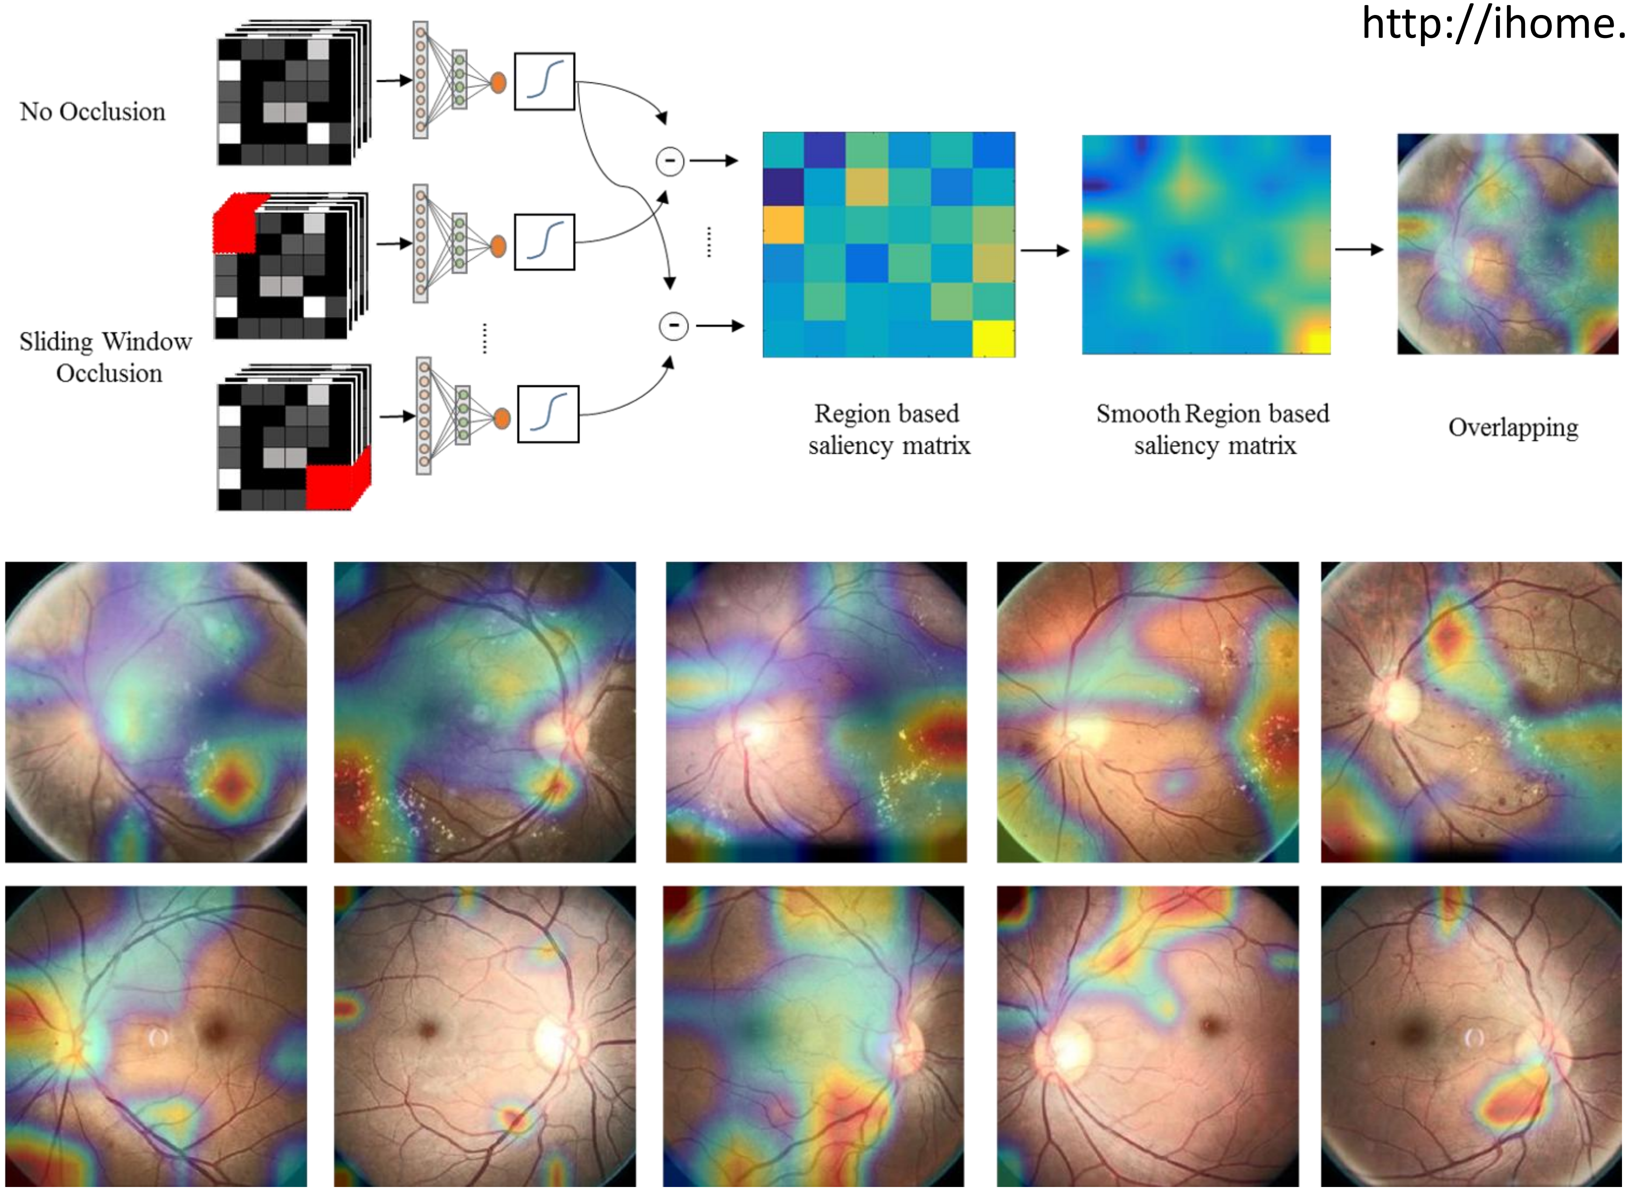

# Deep learning workflow

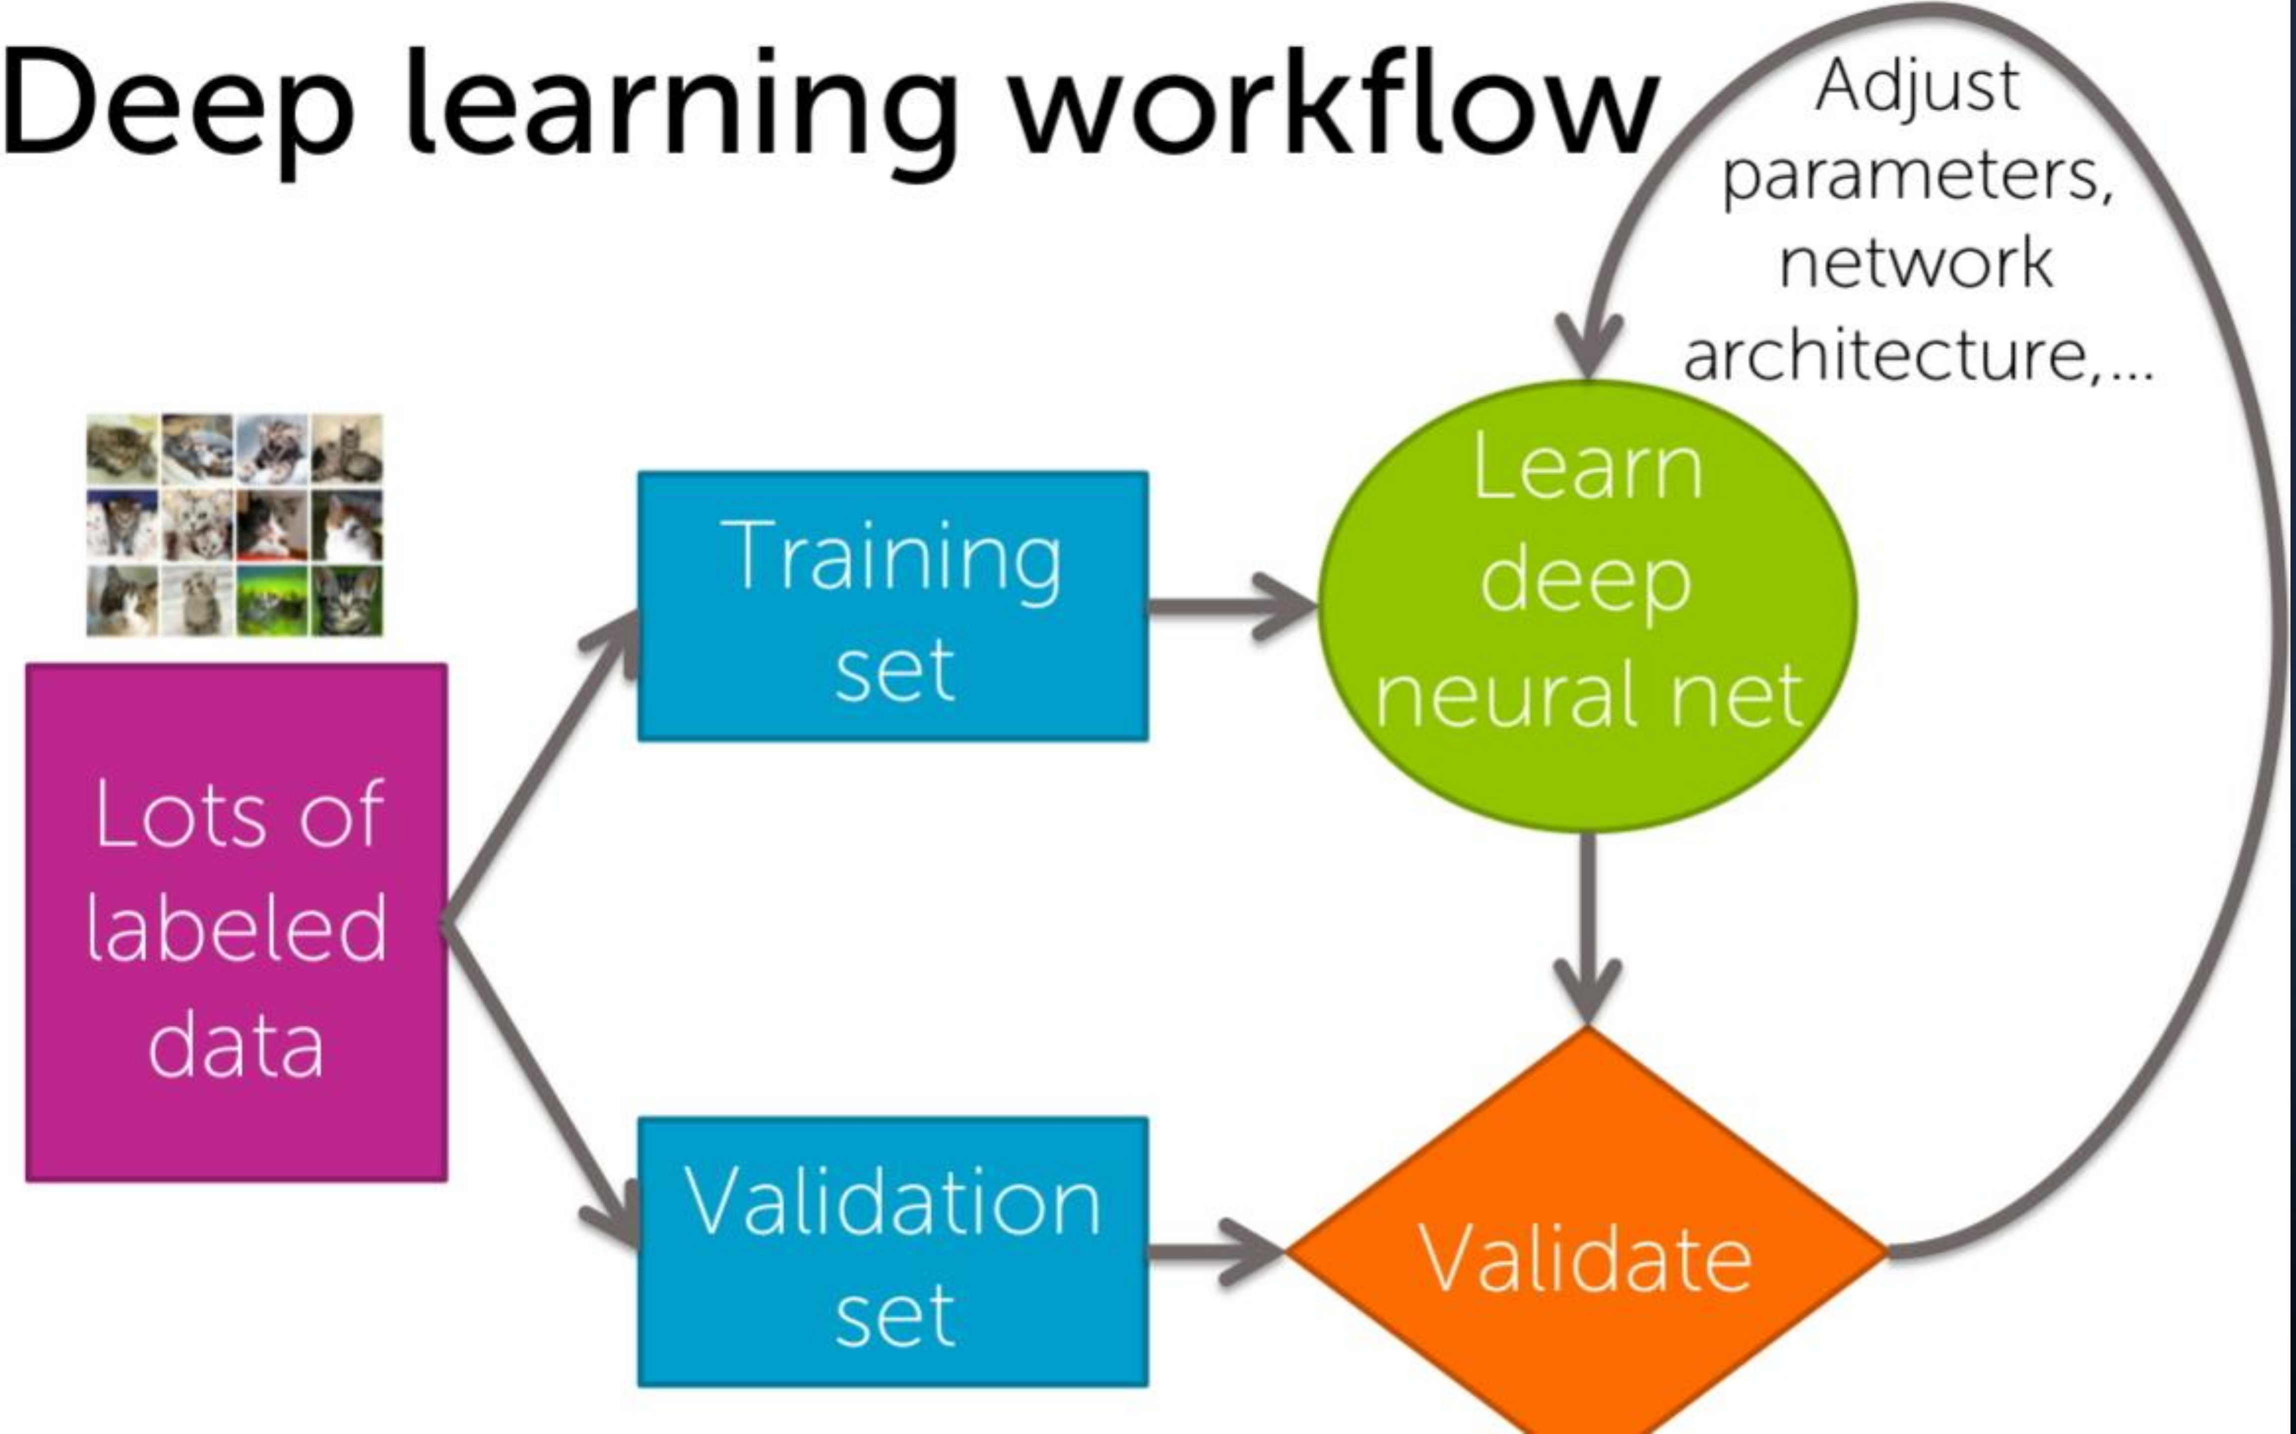

# For Pattern Recognition Tasks, Deep Learning Algorithms Will Outperform Humans Given Enough Objective Data

*Dermatology, Ophthalmology, Radiology, Pathology…*

Andrej Karpathy

# How Does AlphaGo Work?

# (Narrow AI)

# Under the Hood…

- Behavior cloning (supervised learning on human demonstration data)
  - 100,000 game examples, 30 million moves
- Reinforcement learning (REINFORCE)
  - Self Play
  - 1,000,000 games
- Value functions
  - Prevent Overfitting of Model
- Monte Carlo Tree Search (MCTS)
  - Determine Next Move

https://medium.com/@karpathy/alphago-in-context-c47718cb95a5

**<u>Trial and Error AI</u>**

Reinforcement learning

Counterfactual regret minimization = wider range of randomized bets

Pattern remover to avoid exploitation

---

# Libratus -- How Can It Bluff?

# Do These Algorithms Work in Medical Domains?

LETTER

doi:10.1038/nature21056

# Dermatologist-level classification of skin cancer with deep neural networks

Andre Esteva[1*], Brett Kuprel[1*], Roberto A. Novoa[2,3], Justin Ko[2], Susan M. Swetter[2,4], Helen M. Blau[5] & Sebastian Thrun[6]

**Skin cancer, the most common human malignancy[1–3], is primarily diagnosed visually, beginning with an initial clinical screening and followed potentially by dermoscopic analysis, a biopsy and histopathological examination. Automated classification of skin lesions using images is a challenging task owing to the fine-grained variability in the appearance of skin lesions. Deep convolutional neural networks (CNNs)[4,5] show potential for general and highly variable tasks across many fine-grained object categories[6–11]. Here we demonstrate classification of skin lesions using a single CNN, trained end-to-end from images directly, using only pixels and disease labels as inputs. We train a CNN using a dataset of 129,450 clinical images—two orders of magnitude larger than previous datasets[12]—consisting of 2,032 different diseases. We test its performance against 21 board-certified dermatologists on biopsy-proven clinical images with two critical binary classification use cases: keratinocyte carcinomas versus benign seborrheic keratoses; and malignant melanomas versus benign nevi. The first case represents the identification of the most common cancers, the second represents the identification of the deadliest skin cancer. The CNN achieves performance on par with all tested experts across both tasks, demonstrating an artificial intelligence capable of classifying skin cancer with a level of competence comparable to dermatologists. Outfitted with deep neural networks, mobile devices can potentially extend the reach of dermatologists outside of the clinic. It is projected that 6.3 billion smartphone subscriptions will exist by the year 2021 (ref. 13) and can therefore potentially provide low-cost universal access to vital diagnostic care.**

There are 5.4 million new cases of skin cancer in the United States[2] images (for example, smartphone images) exhibit variability in factors such as zoom, angle and lighting, making classification substantially more challenging[23,24]. We overcome this challenge by using a data-driven approach—1.41 million pre-training and training images make classification robust to photographic variability. Many previous techniques require extensive preprocessing, lesion segmentation and extraction of domain-specific visual features before classification. By contrast, our system requires no hand-crafted features; it is trained end-to-end directly from image labels and raw pixels, with a single network for both photographic and dermoscopic images. The existing body of work uses small datasets of typically less than a thousand images of skin lesions[16,18,19], which, as a result, do not generalize well to new images. We demonstrate generalizable classification with a new dermatologist-labelled dataset of 129,450 clinical images, including 3,374 dermoscopy images.

Deep learning algorithms, powered by advances in computation and very large datasets[25], have recently been shown to exceed human performance in visual tasks such as playing Atari games[26], strategic board games like Go[27] and object recognition[6]. In this paper we outline the development of a CNN that matches the performance of dermatologists at three key diagnostic tasks: melanoma classification, melanoma classification using dermoscopy and carcinoma classification. We restrict the comparisons to image-based classification.

We utilize a GoogleNet Inception v3 CNN architecture[9] that was pre-trained on approximately 1.28 million images (1,000 object categories) from the 2014 ImageNet Large Scale Visual Recognition Challenge[6], and train it on our dataset using transfer learning[28]. Figure 1 shows the working system. The CNN is trained using 757 disease classes. Our

Research

JAMA | Original Investigation | INNOVATIONS IN HEALTH CARE DELIVERY

# Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs

Varun Gulshan, PhD; Lily Peng, MD, PhD; Marc Coram, PhD; Martin C. Stumpe, PhD; Derek Wu, BS; Arunachalam Narayanaswamy, PhD; Subhashini Venugopalan, MS; Kasumi Widner, MS; Tom Madams, MEng; Jorge Cuadros, OD, PhD; Ramasamy Kim, OD, DNB; Rajiv Raman, MS, DNB; Philip C. Nelson, BS; Jessica L. Mega, MD, MPH; Dale R. Webster, PhD

**IMPORTANCE** Deep learning is a family of computational methods that allow an algorithm to program itself by learning from a large set of examples that demonstrate the desired behavior, removing the need to specify rules explicitly. Application of these methods to medical imaging requires further assessment and validation.

**OBJECTIVE** To apply deep learning to create an algorithm for automated detection of diabetic retinopathy and diabetic macular edema in retinal fundus photographs.

**DESIGN AND SETTING** A specific type of neural network optimized for image classification called a deep convolutional neural network was trained using a retrospective development data set of 128 175 retinal images, which were graded 3 to 7 times for diabetic retinopathy, diabetic macular edema, and image gradability by a panel of 54 US licensed ophthalmologists and ophthalmology senior residents between May and December 2015. The resultant algorithm was validated in January and February 2016 using 2 separate data sets, both graded by at least 7 US board-certified ophthalmologists with high intragrader consistency.

Editorial

Supplemental content

# Yes!

TELADOC

Dr

# Telemedicine Today

AIRSTRIP®

*Remote Triage Screening, Urgent Care, Clinical Decision Support Monitoring*

# Telemedicine Tomorrow

*Digital Therapeutics*

*AI-Supported Screening*

omada


# Digital Therapeutics

# Virtual Care Services

- For Diagnostic Interpretation of a Test, Think AI First Review
- Natural Language Processing
- Computer Vision Deep Learning

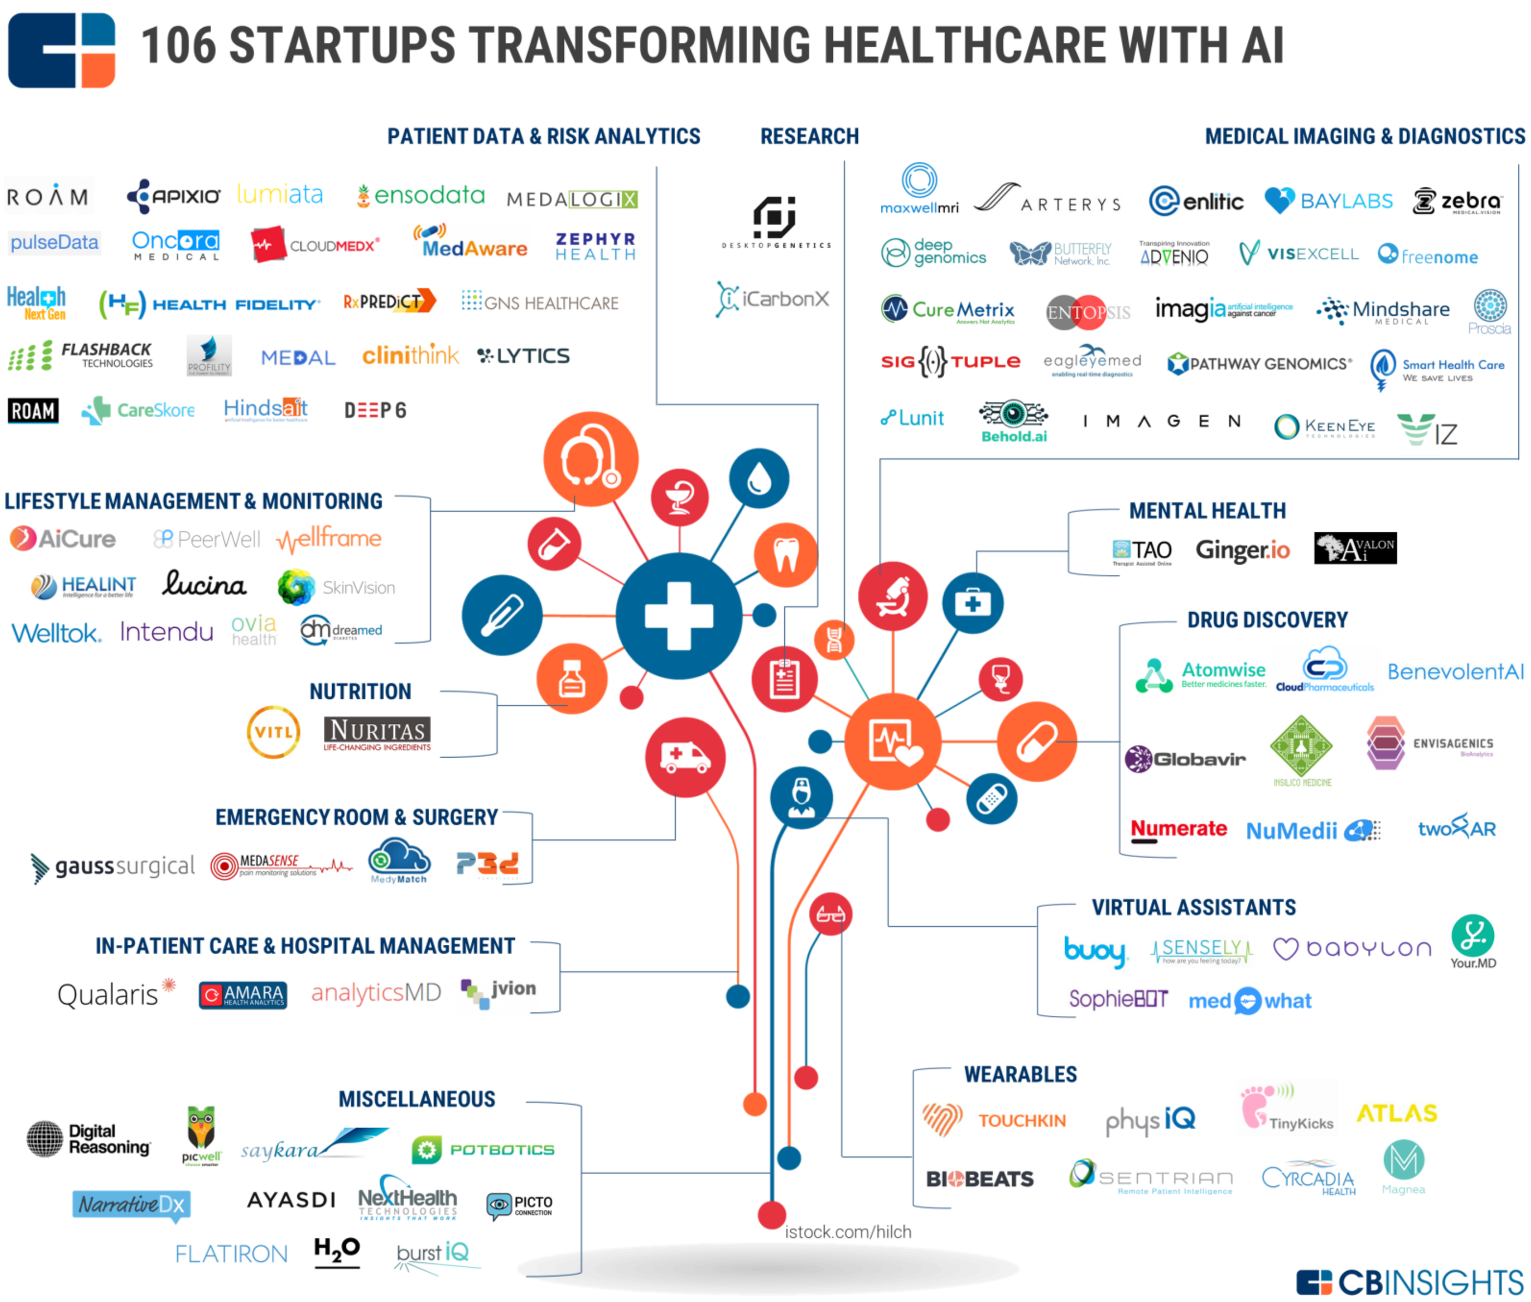


https://www.cbinsights.com/blog/artificial-intelligence-startups-healthcare/

# What Do We Need in Healthcare So That AI Can Augment Healthcare Providers?

- Huge amounts of precise/accurate big data tied to valid physician-driven outcome metrics, crunched by ever improving machine learning algorithms
- Payment models for healthy individuals to save money and time receiving preventive care remotely


LOOKING
FOR A
JOB

2025

Google.md

Doctor, do I have an eye problem?

Google Search

I'm Feeling Lucky

Capture Photo